IBM i

IBM i integration with BladeCenter and System x

Version 7.2

IBM i

IBM i integration with BladeCenter and System x

Version 7.2

Note

Before using this information and the product it supports, read the information in “Notices” on page 137.

This edition applies to IBM i 7.2 (product number 5770-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.

Contents

Chapter 1. What's new for IBM i 7.2 . . . 1

Chapter 2. Concepts for integrated servers 3

Integrated server overview 3
Integrated server capabilities 4
Attaching servers to IBM i using iSCSI 6
Typical iSCSI-attached server installation 7
Single-server environment 9
Multiple-server environment. 11
Initiator system and service processor connection 13
Booting over the iSCSI network 13
Server management for integrated servers 14
Integrated Windows servers 14
Integrated VMware ESX servers 18
Integrated server console 21
Software updates for integrated servers 22
Storage management for integrated servers 23
Virtual storage for integrated servers 23
IBM i storage management for integrated servers 24
Predefined virtual storage and naming for integrated servers 27
Virtual storage linking for integrated servers . . 28
IBM i tape and optical devices shared with integrated Windows servers 29
Networking concepts for integrated servers. . . . 29
Service processor connection for integrated servers 29
Service processor functions and support 30
Service processor connection methods 31
iSCSI network for integrated servers 32
Integrated DHCP server for integrated servers. . 34
Networking between IBM i and integrated servers 35
Virtual Ethernet networks for integrated Windows servers 36
Physical networks for integrated servers. . . . 41
Network security for integrated servers 41
Performance concepts for integrated servers . . . 43
Storage performance for integrated servers . . . 43
Virtual Ethernet performance for integrated Windows servers 45
MTU considerations for the iSCSI network . . . 45
IBM i configuration objects for integrated servers. . 46
High availability concepts for integrated servers . . 52
Multipath I/O for integrated servers 52
Hot spare support for integrated servers. . . . 55
IBM i clustering for integrated servers 55
User and group enrollment concepts for integrated Windows servers 56
Enrolled user account options for integrated Windows servers 59
User enrollment templates for integrated Windows servers 60
Password considerations for integrated Windows servers 61
QAS400NT user and integrated Windows servers 62
IBM i NetServer for integrated Windows servers . . 63
IBM i Access for Windows and integrated servers 63

Chapter 3. Integrated server installation road map (Deprecated). 65

Prerequisites for installing an integrated server (Deprecated) 65
Preparing for the hardware installation (Deprecated) 65
iSCSI Network Planning Guide (Deprecated) . . . 65
Installing the hardware (Deprecated) 65
Preparing IBM i for the integrated server installation (Deprecated) 65
Installing the integrated server (Deprecated) . . . 66

Chapter 4. Managing integrated server environments 67

Creating and deleting integrated servers. 67
Installing integrated servers 67
Cloning integrated servers 67
Uninstalling integrated servers 67
Managing integrated servers. 69
Starting and stopping integrated servers. . . . 70
Viewing or changing integrated server configuration information 73
Configuring security between IBM i and integrated servers 75
Configuring multipath I/O for integrated servers (Deprecated) 77
Using hot spare hardware (Deprecated) 77
Configuring high availability for integrated servers (Deprecated) 77
Viewing integrated server messages 77
Launching the Web console for an integrated server 78
Managing integrated Windows servers (Deprecated) 78
Updating the integration software running on Microsoft Windows (Deprecated) 78
Managing virtual Ethernet and external networks (Deprecated) 78
Sharing tape and optical devices between IBM i and integrated Windows servers (Deprecated) . . 79
Running integrated Windows server commands remotely (Deprecated) 79
Managing integrated VMware ESX servers (Deprecated) 79
Updating the integration software for VMware ESX (Deprecated) 79
Configuring the QVMWINT user for integrated VMware ESX server management (Deprecated) . 79
Managing connections for integrated VMware ESX servers (Deprecated) 80

Migrating VMware ESX servers to new IBM i management infrastructure (Deprecated). . . . 80
Managing storage for integrated servers 80
Displaying information about integrated server virtual storage 80
Adding virtual storage to integrated servers . . 80
Copying virtual storage 85
Expanding virtual storage 85
Expanding a system disk for an integrated Windows server (Deprecated) 86
Unlinking virtual storage 86
Deleting virtual storage 87
Accessing the IBM i integrated file system from an integrated server 87
Administering integrated Windows server users from IBM i 88
Enrolling IBM i users to integrated Windows servers 88
Enrolling IBM i groups to integrated Windows servers 88
Configuring the QAS400NT user for user enrollment on integrated Windows servers . . . 89
Creating user enrollment templates for integrated Windows servers 91
Specifying a home directory in a user template 92
Changing the local password management user profile attribute 93
Configuring Enterprise Identity Mapping for integrated Windows servers 93
Ending user enrollment to an integrated Windows server 95
Ending group enrollment to an integrated Windows server 96
Preventing enrollment to an integrated Windows server 96
Managing network server host adapters 98
Creating a network server host adapter 98
Creating a network server host adapter based on another one 100
Displaying network server host adapter properties 101
Changing network server host adapter properties 101
Starting a network server host adapter 102
Stopping a network server host adapter . . . 102
Deleting a network server host adapter. . . . 103
Managing remote system configurations 104
Creating a remote system configuration . . . 105
Creating a remote system configuration based on another one 106
Displaying remote system configuration properties 106
Changing remote system configuration properties 107
Displaying remote system status 107
Deleting a remote system configuration . . . 108
Launching the Web console for a remote system 108
Managing service processor configurations . . . 108
Creating a service processor configuration. . . 109
Creating a service processor configuration based on another one 110
Displaying service processor configuration properties 110
Changing service processor configuration properties. 111
Initializing a service processor 111
Deleting a service processor configuration . . . 112
Launching the Web console for a service processor 112
Converting pre-IBM i 7.1 service processor configurations to use an IP address or host name 113
Managing connection security configurations . . . 113
Creating a connection security configuration . . 113
Creating a connection security configuration based on another one. 114
Displaying connection security configuration properties 114
Changing connection security configuration properties 115
Deleting a connection security configuration . . 115

Chapter 5. Backing up and recovering integrated servers 117
Backing up the NWSD and other objects associated with integrated servers 117
What objects to save and their location on IBM i 118
Backing up the NWSD of an integrated server 120
Backing up virtual storage for integrated servers 120
Backing up virtual storage for active Windows servers 121
Backing up NWSH objects and associated LIND objects and interfaces. 122
Backing up NWSCFG objects and validation lists 123
Backing up individual integrated Windows server files and directories 123
File-level backup restrictions for integrated Windows servers 124
Installing and configuring IBM i NetServer . . 125
Configuring integrated Windows servers for file-level backup 126
Creating shares on integrated Windows servers 126
Adding members to the QAZLCSAVL file . . . 127
Verifying that IBM i NetServer and the integrated Windows server are in same domain . 127
Saving integrated Windows server files. . . . 128
Saving user enrollment information for integrated Windows servers 129
Restoring the NWSD and other objects associated with integrated servers 130
Restoring the NWSD and virtual storage for integrated servers 130
Restoring NWSH objects 133
Restoring NWSCFG objects and validation lists 133
Restoring individual integrated Windows server files and directories 134
Restoring user enrollment information for integrated Windows servers 135

Notices 137
Programming interface information 139
Trademarks 139
Terms and conditions. 139

Chapter 1. What's new for IBM i 7.2

Read about new or changed information for the IBM® i integration with BladeCenter and System x topic collection.

Changes to iSCSI-attached integrated servers

Virtual Disk

- » Virtual disk with 4K sector size is supported.«
- » Support for specifying a preference to create virtual disk on solid-state drives.«
- » A new parameter of the Install Integrated Server command and the CRTNWSSTG command to customize the priority of IBM i resources allocated to formatting the client storage space.«

No new Windows Server 2003 (&R2) servers

» No new integrated Windows Server 2003 (&R2) servers can be installed/cloned on IBM i 7.2. Integrated Windows Server 2003 (&R2) servers that were installed on prior IBM i releases and then migrated forward to IBM i 7.2 can continue to run as is, but without service support. The suggested migration path for these servers is to upgrade to Windows Server 2008 or newer version.«

No support of IXS & IXA Hardware

» Integrated xSeries® Server (IXS)s and System x® servers attached using Integrated xSeries Adapter (IXA)s are no longer supported.«

Changed IBM i CL command

New parameters are added in IBM i CL commands to accommodate the new features:

» Install Integrated Server (INSINTSVR)«

» Installs an integrated server.«

» Create NWS Storage Space (CRTNWSSTG)«

» Creates a storage space used by a network server.«

What's new as of 31 July 2013

Information related to cloning integrated servers has been added. See “Cloning integrated servers” on page 67 for more information.

The following documentation has been migrated to the *IBM i iSCSI Solution Guide* PDF on the IBM i iSCSI Solution Guide Web page and is no longer included in this Information Center topic:

- The **Integrated server installation road map** chapter.
- BladeCenter and System x hardware installation and configuration information.

- Microsoft Windows Server and VMware ESX Server installation and configuration information.

The *iSCSI network planning work sheets* have been migrated to the *IBM i iSCSI Solution Work Sheets* PDF on the IBM i iSCSI Solution Guide Web page and are no longer included in this Information Center topic.

Note: All articles that have been removed are marked with **(Deprecated)** in the article titles.

How to see what's new or changed

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

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

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

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

Chapter 2. Concepts for integrated servers

Understand concepts for iSCSI-attached servers for the IBM i integration with BladeCenter and System x solution.

Integrated server overview

An integrated server is a combination of integrated server hardware, network components, virtual storage (virtual disks), shared devices, and IBM i integrated server configuration objects.

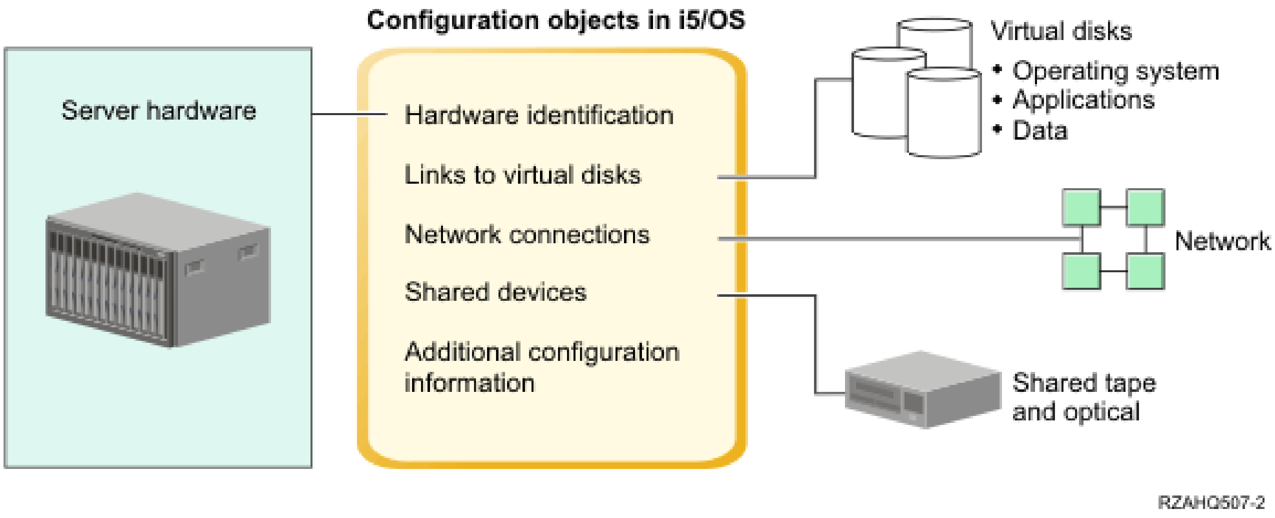

Figure 1. Integrated server overview

Server hardware

The server hardware is the physical hardware (such as the processor and memory) that the integrated server runs on. There are several types of server hardware that can be used for integrated servers, depending on your needs. The integrated server hardware is an external System x or BladeCenter product that is attached to IBM i with iSCSI initiator and target adapters. See “Attaching servers to IBM i using iSCSI” on page 6 for more information about the types of hardware that can be used for integrated servers.

iSCSI adapters

Both IBM i and the integrated server contain iSCSI adapters, which are connected over an Ethernet network. The integrated server uses its iSCSI adapter to connect to the iSCSI adapter in IBM i. Using this connection, the integrated server can access virtual storage, shared tape and optical devices, and virtual Ethernet resources on IBM i. See “Attaching servers to IBM i using iSCSI” on page 6 for more information.

Virtual storage

Each integrated server uses virtual storage that contain the integrated server operating system, applications, and data. This virtual storage is allocated from IBM i disk storage. The integrated server treats these drives as physical disk drives that are contained within the server. However, the integrated server does not actually have any physical disk drives of its own. See “Storage management for integrated servers” on page 23 for more information about virtual storage.

Shared tape and optical devices

An integrated Windows server can share supported tape and optical devices that are connected to the hosting IBM i partition. Shared IBM i devices are accessed as if they were local to the integrated Windows server. By default, IBM i tape and optical devices are automatically accessible by an integrated Windows server. You can choose to restrict which of these IBM i devices the integrated Windows server can access. A subset of IBM i tape devices are supported for use with various Windows versions. See the IBM i iSCSI Solution Guide for more information.

Note: IBM i devices cannot be shared with integrated VMware ESX servers.

Network

Each integrated server has one or more connections to a network. Physical network connections with a network adapter are supported for all types of integrated servers. Power server virtual Ethernet network connections are supported by integrated Windows servers. See “Networking concepts for integrated servers” on page 29 for more information about the types of network connections that can be used with integrated servers.

IBM i integrated server configuration objects

Configuration objects in IBM i describe each integrated server. The IBM i configuration objects identify the hardware that the integrated server runs on, the virtual storage that the integrated server uses, the iSCSI target and initiator adapters that the integrated server uses, the virtual Ethernet connections that an integrated Windows server uses, and many other attributes of the server. See “IBM i configuration objects for integrated servers” on page 46 for more information.

Integrated server capabilities

Integrated servers allow you to run supported versions of the Windows, or VMware ESX operating systems. With integrated servers, you can take advantage of IBM i capabilities such as storage management, high availability, and user propagation solutions.

There are fewer pieces of hardware to manage requiring less physical space. iSCSI-attached integrated servers can take advantage of BladeCenter hardware.

Greater accessibility and protection for your data

- Integrated servers use IBM i disk storage, which is more reliable than PC server hard disks.
- Integrated servers allow you to run AMD64 and Intel EM64T versions of Windows Server 2012,Windows Server 2008 and VMware ESX server.

- You have access to faster IBM i tape devices for integrated Windows server backups.
- You can back up the entire integrated server as part of your IBM i server backup. This allows you to recover a failed server much faster and easier than with typical file level recovery on the integrated server operating system.
- Integrated servers implicitly take advantage of superior data protection schemes which exist in IBM i such as RAID or drive mirroring.
- Typical integrated server configurations have storage space data spread across more IBM i disk drives than would be configured in standalone (non-integrated) server installations. This configuration can frequently provide better peak disk I/O capacity, since each server is not constrained to few dedicated drives.
- You can add additional disk storage to integrated servers without shutting down the server.
- It is possible to gain access to Db2® for i data through an enhanced Open Database Connectivity (ODBC) device driver using IBM i Access for Windows. This device driver enables server-to-server applications between integrated Windows servers and IBM i.
- Virtual networking for integrated Windows servers does not require additional LAN hardware and provides communications between IBM i logical partition and System x or BladeCenter blade servers attached using iSCSI adapters.

Simplified administration

- Your computer system is less complicated because of the integration of security, server management, and backup and recovery between the IBM i and integrated server operating systems. You can save your integrated server data on the same media as other IBM i data and restore individual Windows files as well as IBM i objects.
- For integrated Windows servers, user parameters, such as passwords, are easier to administer from IBM i using the user administration function. You can create users and groups and enroll them from IBM i to integrated Windows servers. The user administration function makes updating passwords and other user information from IBM i easy.

Remote management and problem analysis

- You can sign on to IBM i from a remote location and shut down or restart your integrated server.
- Since you can mirror integrated Windows server event log information to IBM i, you can remotely analyze Microsoft Windows errors.

Multiple servers

- Integrated Windows servers and logical partitions running on the same Power server have high-performance, secure virtual networking communications that do not require using LAN hardware.
- You can run multiple integrated servers on a single IBM i partition. Not only convenient and efficient, running multiple integrated servers also gives you the ability to easily switch to another hot spare server if the integrated server hardware fails.
- If you have multiple integrated servers installed on your IBM i partition, you can define their Windows domain roles in a way that simplifies user enrollment and access. For example, you might want to set up one of these servers as a domain controller. Then, you only have to enroll users to the domain controller, and users can log on from any Microsoft Windows machine on that domain.

Hot spare support

Server integration and storage virtualization provide innovative options that can enhance the reliability and recoverability of the integrated server environment. Hot spare hardware provides a way to quickly recover from certain types of hardware failures. This can reduce the server downtime from hours or days to minutes.

See “Hot spare support for integrated servers” on page 55 for more information.

Attaching servers to IBM i using iSCSI

A basic iSCSI network consists of an IBM i iSCSI target adapter and a System x or IBM BladeCenter blade iSCSI initiator adapter.

The target and initiator devices are connected over an Ethernet local area network (LAN). The iSCSI target for IBM i provides the storage devices for the iSCSI initiator. For an integrated Windows server, the iSCSI target also provides removable media devices and virtual Ethernet connections for the iSCSI initiator. Figure 2 illustrates a basic iSCSI network.

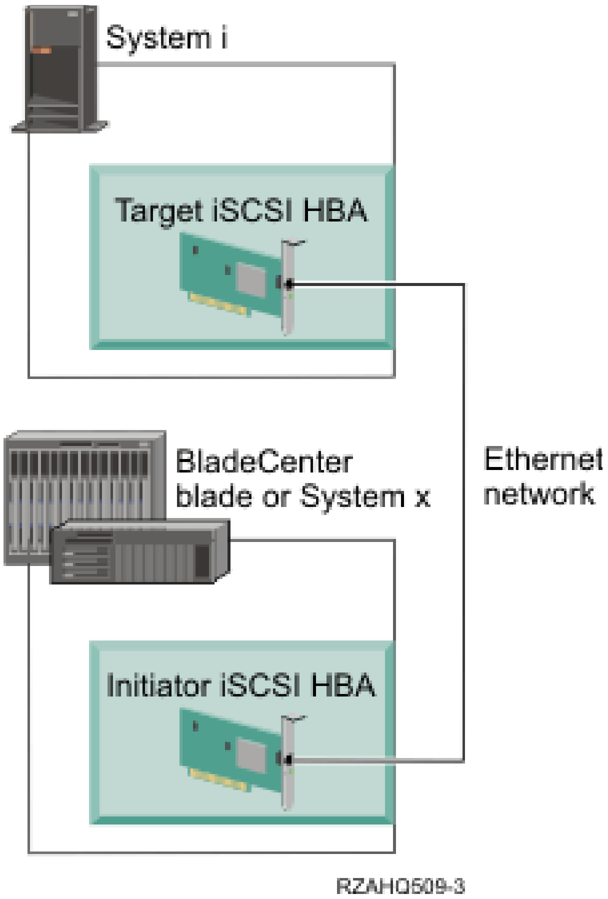

Figure 2. Basic iSCSI network

You need to configure both the iSCSI target and initiator adapters from IBM i. The iSCSI network is used for iSCSI traffic only.

There are two types of iSCSI target and initiator adapter implementations:

Software target or initiator (Ethernet NIC)

With a software target or initiator, the iSCSI protocol is implemented in the server operating system. Server resources (for example, CPU and memory) are used for the iSCSI protocol. The IBM i integrated server solution uses standard Ethernet network interface cards (Ethernet NICs) as software targets and initiators.

Hardware target or initiator (iSCSI HBA)

With a hardware target or initiator, the iSCSI protocol is implemented in firmware on the iSCSI adapter. The iSCSI protocol is offloaded from the server. The IBM i integrated server solution uses iSCSI host bus adapters (iSCSI HBAs) as hardware targets and initiators.

IBM i can use any combination of software-based or hardware-based iSCSI target adapters. The integrated server can use any combination of software-based or hardware-based iSCSI initiator adapters that are supported by the specific integrated server model and the operating system that is installed on the server. Either type of iSCSI initiator adapter can connect to either type of iSCSI target adapter.

Note: The level of support for software-based iSCSI initiator adapters (Ethernet NICs) depends on the integrated server operating system version.

See the IBM i iSCSI Solution Guide for information about supported iSCSI target and initiator adapters.

Typical iSCSI-attached server installation

iSCSI-attached integrated servers are standard System x or IBM BladeCenter models that have processors, memory, and expansion cards, but no physical disks. Integrated servers use virtual disks on IBM i that are managed by IBM i.

The installation procedure for an iSCSI-attached integrated server requires hardware to be installed and configured in both the IBM i and System x or BladeCenter products. You can use the System x expansion slots for additional options.

The following graphic illustrates a typical iSCSI-attached server installation:

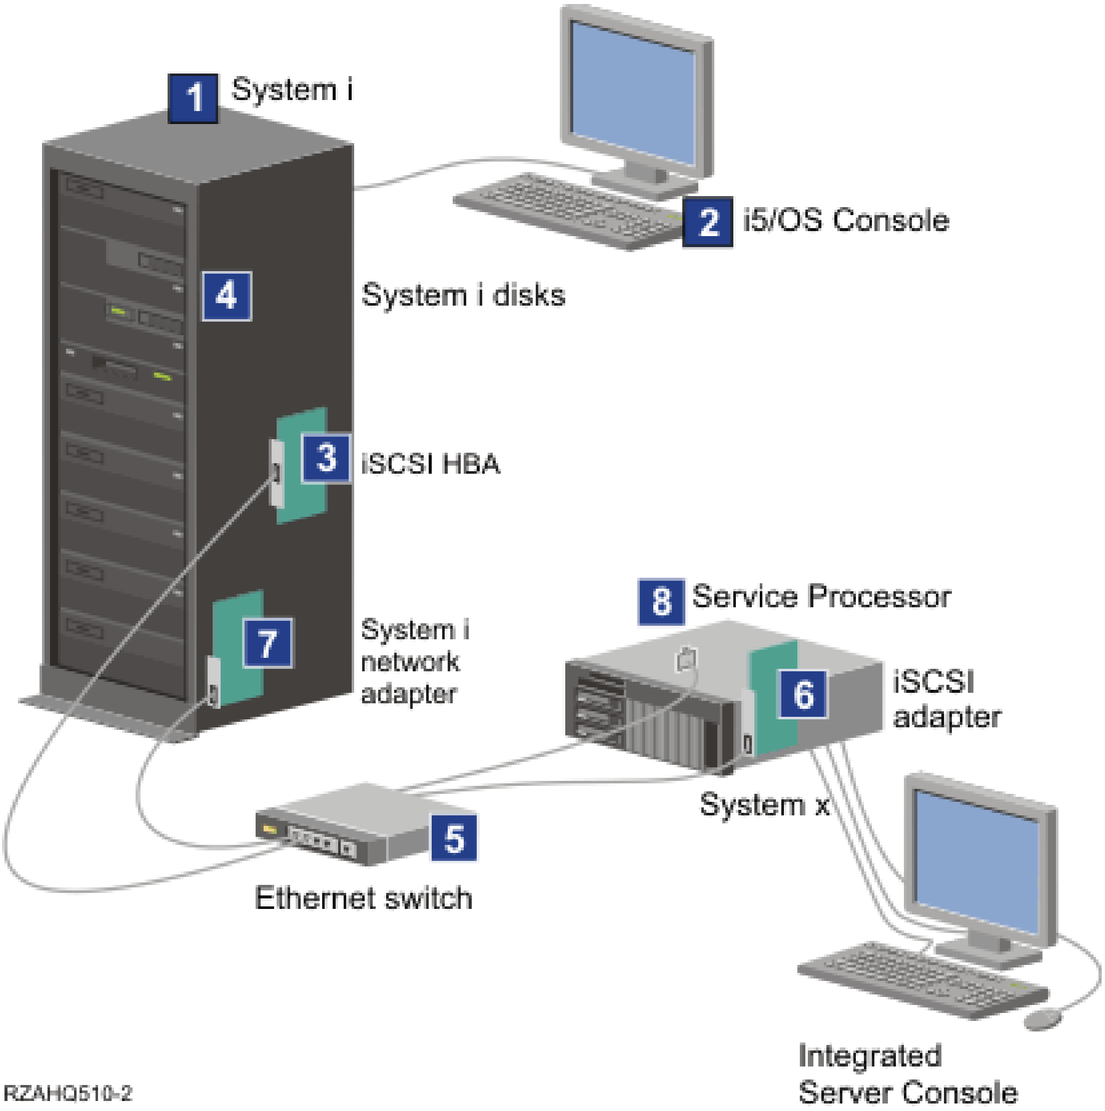

Figure 3. A typical iSCSI-attached integrated server installation

1. You need a compatible Power server model. See the IBM i iSCSI Solution Guide for information about supported Power server models.
2. The IBM i (IBM i) console, from which you connect to IBM i using *IBM Navigator for i* or the character-based interface, is shown to make clear the distinction between it and the integrated server console.
3. Depending on the type of the physical network, copper or fiber iSCSI adapters (Ethernet NICs or iSCSI HBAs) are available. This iSCSI adapter installed in IBM i is the target device and connects to an Ethernet network using standard Ethernet cables.
4. An integrated server does not have its own physical disk drive. IBM i emulates hard disk space for it to use from IBM i disks. These disks and other IBM i storage devices are accessed through the iSCSI target adapter.
5. The iSCSI adapter network cables are connected to a standard Gigabit Ethernet switch.
6. An additional iSCSI adapter is required in the System x or blade hardware. This adapter provides the connection to iSCSI target adapter in the Power

server. This adapter can be viewed from the System x or blade model as the storage adapter, where the disks are found across the network.

7. A typical Power server has a network card. An IBM i LAN connection is required to connect to and manage the System x or BladeCenter hardware.
8. A service processor allows IBM i to connect to and manage the system. The service processor is connected to IBM i over an Ethernet network.

For more information about hardware, see the IBM i iSCSI Solution Guide.

Single-server environment

A basic iSCSI-attached integrated server configuration requires iSCSI adapters and IBM i (IBM i) configuration objects.

The simplest form of the physical connection between an initiator system and IBM i is illustrated in Figure 4 on page 10.

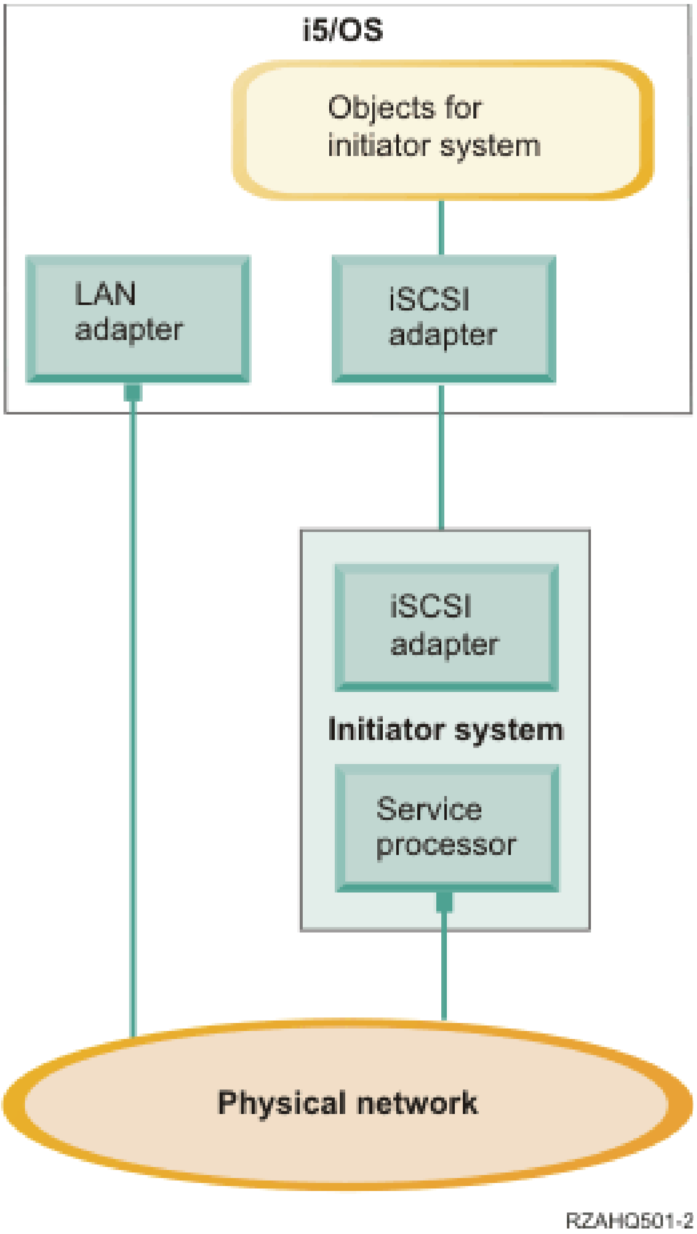

Figure 4. Single iSCSI-attached server

An iSCSI adapter is installed in each system. The Ethernet network between the iSCSI adapters is known as the iSCSI network. The initiator system (System x or BladeCenter system) uses this network to access storage through the IBM i iSCSI target adapter.

The initiator system has no physical disks and connects to virtual disks and virtual removable media devices in IBM i. The SCSI commands to access these devices are packaged in TCP/IP frames and travel over an Ethernet network from the initiator system to the IBM i iSCSI target adapter. This mode of communication is known as Internet SCSI or iSCSI.

The iSCSI-attached servers are configured in IBM i objects. For more information about these objects, see "IBM i configuration objects for integrated servers" on page 46.

IBM i can connect to and manage remote systems by sending commands to the service processor of the remote (initiator) system over an Ethernet network. For more information, see “Service processor functions and support” on page 30.

Two distinct networks are illustrated in “Single-server environment” on page 9. The iSCSI network uses an isolated switch or a direct connection. The service processor connection uses an external network (shared network).

Two distinct networks are not required. For example, the service processor connection can use the same isolated switch as the iSCSI network. This is one way to secure the service processor connection. However, the IBM i LAN adapter would not be available for other applications on the external network.

Both types of networks should be secured. For more information about security for iSCSI-attached servers, see “Network security for integrated servers” on page 41.

Multiple-server environment

You can use one IBM i iSCSI target adapter to host multiple initiator (System x or blade) systems.

This concept is illustrated in Figure 5 on page 12.

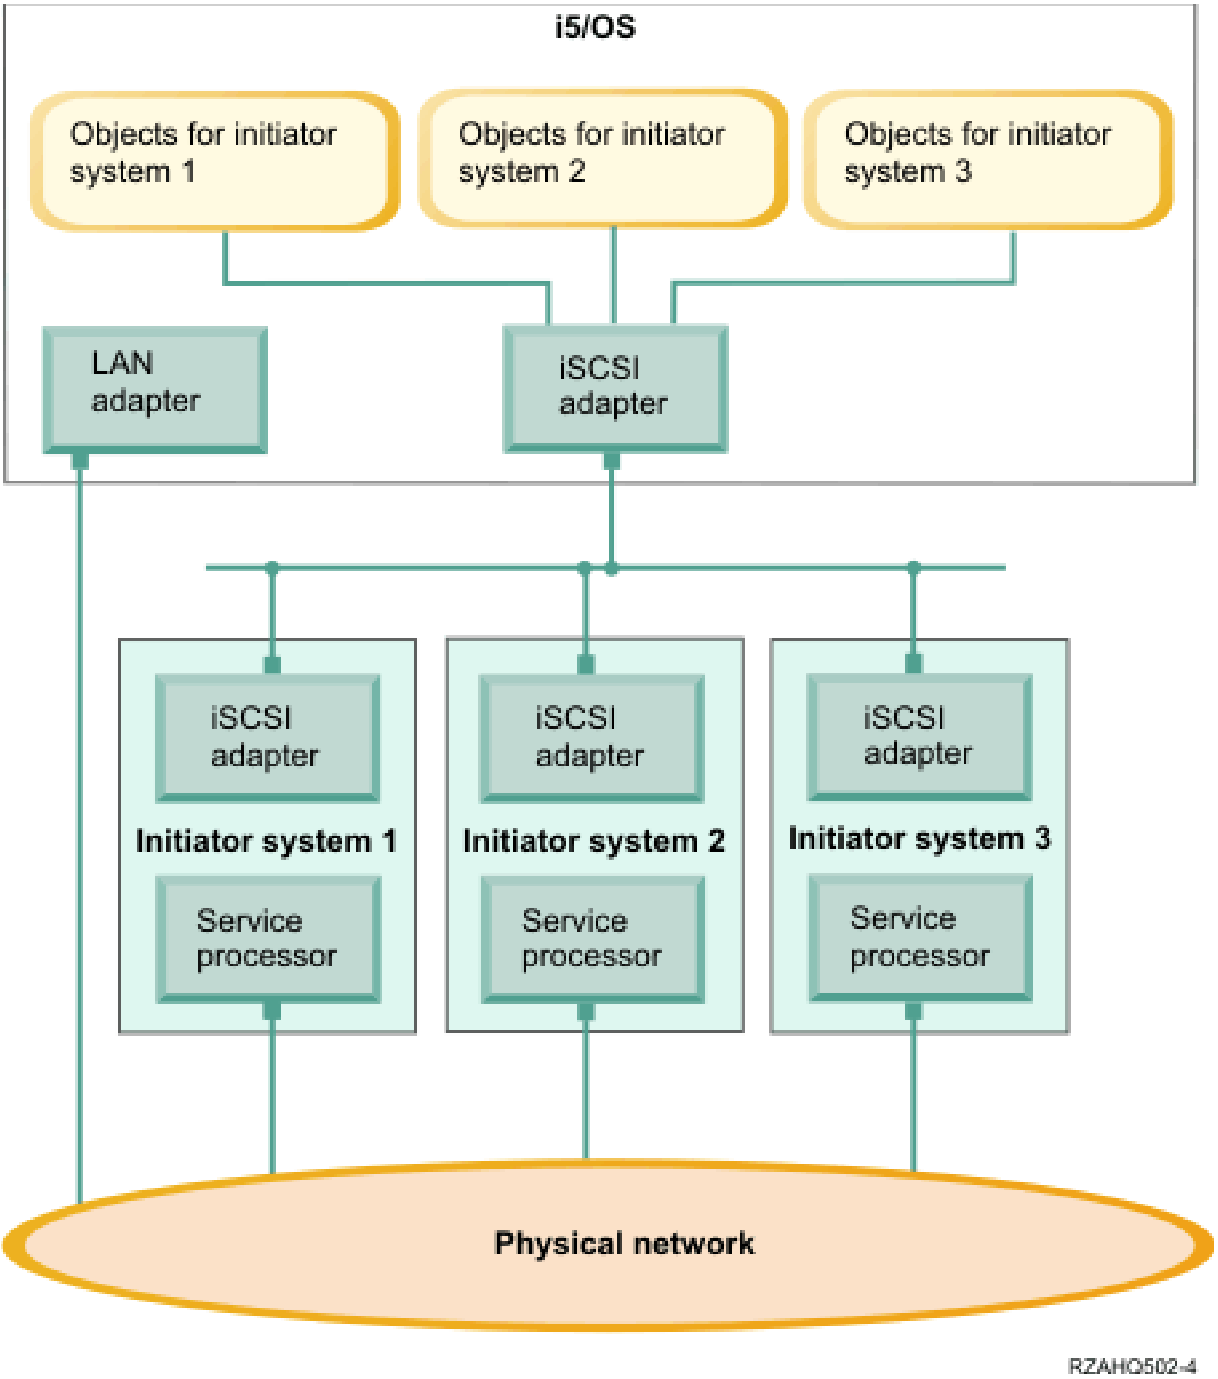

Figure 5. Multiple iSCSI-attached servers

The horizontal line in the diagram between the iSCSI adapters represents a switch. A switch is required when more than one iSCSI initiator adapter shares a single iSCSI target adapter.

You must install an iSCSI initiator adapter in each hosted System x or blade product. The iSCSI initiator adapters are connected by an Ethernet network. This network can be a physically secure or isolated network when a physically secure model is implemented. Each initiator system is represented by a set of IBM i (IBM i) objects. For more information, see "IBM i configuration objects for integrated servers" on page 46.

Each initiator system must have a service processor installed for remote connections and power management. Multiple service processors can be connected to a single IBM i LAN adapter over an external network.

Two distinct networks are not required. For example, the service processor connection can use the same isolated switch as the iSCSI network. This is one way to secure the service processor connection. However, the IBM i LAN adapter would not be available for other applications on the external network.

Initiator system and service processor connection

IBM i uses the Service Processor Manager function of IBM i Integrated Server Support to connect to System x or BladeCenter hardware on the network, to turn the initiator system hardware on and off, and to retrieve power status.

Initiator systems are identified by information stored in the remote system configuration and the service processor configuration in IBM i.

This is a different connection than the iSCSI network connection between the IBM i iSCSI target adapter and the iSCSI initiator adapter in the initiator system. The LAN adapter for the service processor of the remote server must be attached to a network that is reachable by an IBM i LAN adapter.

Both the IBM i objects and the service processor must be configured. You configure the connection options in the IBM i network server configuration objects.

Related tasks:

Configuring service processor connection (Deprecated)
Use the information from the IBM i remote system and service processor configurations to connect to the hardware of integrated System x and blade servers.

Booting over the iSCSI network

iSCSI-attached integrated server hardware is diskless. The boot device is a port configured on the iSCSI initiator adapter installed in the System x or blade hardware.

Both the IBM i remote system configuration and the iSCSI initiator adapter must be configured before you install or use a new integrated server. See “Remote system configuration” on page 49.

Boot modes and parameters

Boot parameters for an iSCSI initiator are configured with an iSCSI initiator configuration function. Boot parameter values must match the values in the IBM i remote system configuration. The parameters vary depending on the selected boot mode.

See the IBM i iSCSI Solution Guide for information about configuring the iSCSI initiator port as the iSCSI boot device. See “Changing remote system configuration properties” on page 107 for information about changing parameters for the remote system configuration object.

Enabling the hosted server boot device

The iSCSI initiator adapter installed in the System x or blade hardware acts as a boot device during the boot process, based on the configured parameters.

You must configure at least one port on the iSCSI initiator adapter as a boot device.

Server management for integrated servers

Concepts for managing servers that are integrated with IBM i.

Integrated Windows servers

When a Windows server is integrated with IBM i, there are some special things to consider when managing the server.

Installation

When you install an integrated server, parts of the installation process are performed on IBM i and parts of the installation process are performed on the integrated server console. For example, IBM i creates configuration objects and virtual storage for the server and starts the server. Then you install the integrated server operating system from the integrated server console. Unlike a stand-alone server, the integrated server installation process is initiated from IBM i, rather than at the server console.

You must also perform some other tasks both before and after the integrated server operating system is installed. See the installation road map in the IBM i iSCSI Solution Guide for the entire process.

Cloning

When you clone an integrated Windows server, parts of the cloning process are performed on IBM i and parts of the cloning process are performed on the integrated server console. For example, you prepare the server for cloning from the Windows console, then you use an IBM i GUI cloning wizard to duplicate the IBM i configuration objects and virtual storage for the server, then you perform some final setup on the clone server to get it ready for production use.

See the cloning road map in the IBM i iSCSI Solution Guide for the entire process.

Key IBM i integration features for integrated Windows servers

IBM i Integrated Server Support provides the following features for an integrated Windows server:

Startup and shut down from IBM i

The integrated server can be started or shut down from IBM i, allowing remote management of the server.

IBM i virtual storage

The integrated server uses virtual storage that is provided by IBM i. IBM i manages storage differently than a stand-alone server. Some techniques necessary to administer storage on a stand-alone server are unnecessary for integrated servers. See "Storage management for integrated servers" on page 23.

Dynamic virtual storage linking

IBM i virtual storage can be linked (added) to the integrated server while it is active. For example, if the server is running low on storage capacity, you can add virtual storage to the server without shutting it down.

Dynamic virtual storage unlinking

IBM i virtual storage can be unlinked (removed) from the integrated server while it is active.

Virtual storage backup and recovery from IBM i

You can back up entire IBM i virtual storage spaces that the integrated server uses along with your other IBM i data. This backup provides a snapshot of the storage that can be used for disaster recovery. The storage can even be backed up while the server is active.

You can restore entire IBM i virtual storage spaces that the integrated server uses. The integrated server must be shut down in order to restore a virtual storage space.

File level backup and recovery from IBM i

You can back up or recover individual Windows files within an IBM i virtual storage space from IBM i while the integrated Windows server is active.

File level backup and recovery from Windows

You can use supported IBM i tape devices from the integrated Windows server to back up or recover individual Windows files while the integrated Windows server is active.

Shared IBM i tape and optical devices

Supported IBM i tape and optical devices can be used by the integrated Windows server as if they were local devices on Windows. Note that a subset of IBM i tape devices are supported for use with various Windows versions. See the IBM i iSCSI Solution Guide for more information.

Enroll IBM i users and groups to Windows servers and domains

You can enroll IBM i users and groups to a Windows server or domain. User enrollment allows simplified administration of users that exist both on IBM i and in the Windows environment. For example, when the user changes their password on IBM i, their password is automatically changed in the Windows environment as well.

Submit remote commands from IBM i to Windows

The remote command feature enables IBM i to run Windows commands on the integrated Windows server, allowing remote management of the server.

Virtual Ethernet connections

The integrated Windows server can use virtual Ethernet connections. These connections allow TCP/IP communication with IBM i or other logical partitions on the Power server without requiring LAN adapters, cables, hubs, or switches.

Windows event log propagation

The integrated Windows server event log entries can be sent to an IBM i message queue or job log. IBM i administrators can view the Windows event log entries from IBM i.

IBM i management interfaces

You can manage IBM i integration features using the *IBM Navigator for i* Web GUI or using IBM i control language (CL) commands. Most integrated server management tasks in this topic are documented using the Web GUI, but links to the corresponding CL commands are also provided.

IBM i management infrastructure for integrated Windows servers

The infrastructure for managing integrated Windows servers from IBM i has the following key pieces:

IBM i software

Most of the software for installing and managing an integrated server runs on IBM i. This software consists of IBM i base operating system functions and various IBM i options such as **Integrated Server Support** (IBM i option 29). The IBM i software enables integrated server installation, IBM i configuration object management, virtual storage management, and much more.

IBM i configuration objects

On the IBM i side of server management, an integrated server is represented by a network server description (NWSD) and several other types of configuration objects. You can stop and restart the server from IBM i by varying the NWSD off and on. See "IBM i configuration objects for integrated servers" on page 46 for more information.

Integrated Windows server

The integrated Windows server that is integrated with IBM i.

Windows utilities

Windows utilities are used to access shared IBM i tape and optical devices. Windows utilities are also used to maintain the *Integrated Server Support* software that runs on the integrated Windows server. See "Windows utilities for IBM i integration with integrated Windows servers" on page 17 for more information.

Windows services

Windows services are used to perform many of the integration tasks for integrated Windows server. See "Windows services for IBM i integration with integrated Windows servers" on page 17 for more information.

QAS400NT user profile

This IBM i user profile is used when performing integrated Windows server administration tasks. See "QAS400NT user and integrated Windows servers" on page 62 for more information.

Windows utilities for IBM i integration with integrated Windows servers

Some of the utilities that provide the IBM i Integrated Server Support features are installed on the integrated Windows server. The following Windows utilities are provided in a Microsoft Management Console (MMC) plug-in named IBM i Integrated Server Support:

IBM i post-install utility for Windows Server 2008 and 2012(ibmsetup.exe)

This utility is used to install Integrated Server Support on an integrated Windows Server 2008 and 2012 server.

Software Level

View the level of IBM i Integrated Server Support software that is installed on IBM i and on the integrated Windows server. Optionally synchronize the Integrated Server Support software from IBM i to the integrated Windows server.

IBM i Devices

View the IBM i (IBM i) tape and optical devices that can be shared with the integrated Windows server. Lock (allocate) an IBM i device so that it can be used by the integrated Windows server. Unlock (deallocate) the IBM i device when it is no longer needed by the integrated Windows server.

Windows services for IBM i integration with integrated Windows servers

Some of the programs that provide the IBM i Integrated Server Support features are installed as Windows services that run on the integrated Windows server. The following Windows services are provided for integrated Windows server management:

IBM i Integration Manager

Manages integrated server startup and shutdown operations.

IBM i Shutdown Manager

Enables system shutdown from IBM i over the iSCSI network.

Important: If this service is stopped, the computer might not respond to a shutdown request from IBM i, which might result in data corruption.

IBM i Administration

Supports user enrollment, event log, disk, and statistics service requests from IBM i.

IBM i Remote Command

Enables processing Windows commands from IBM i.

IBM i Virtual Ethernet Manager

Manages the connection status (link state) for iSCSI-based virtual Ethernet network adapters.

Note: If this service is stopped, the computer does not respond to any virtual Ethernet link state changes.

Integrated VMware ESX servers

When a VMware ESX server is integrated with IBM i, there are some special things to consider when managing the server.

Installation

When you install an integrated server, parts of the installation process are performed on IBM i and parts of the installation process are performed on the integrated server console. For example, IBM i creates configuration objects and virtual storage for the server and starts the server. Then you install the integrated server operating system from the integrated server console. Unlike a stand-alone server, the integrated server installation process is initiated from IBM i, rather than at the server console.

You must also perform some other tasks both before and after the integrated server operating system is installed. See the installation road map in the IBM i iSCSI Solution Guide for the entire process.

Key IBM i integration features for VMware ESX servers

IBM i Integrated Server Support provides the following features for an integrated VMware ESX server:

Startup and shut down from IBM i

The integrated server can be started or shut down from IBM i, allowing remote management of the server.

IBM i virtual storage

The integrated server uses virtual storage that is provided by IBM i. IBM i manages storage differently than a stand-alone server. Some techniques necessary to administer storage on a stand-alone server are unnecessary for integrated servers. See “Storage management for integrated servers” on page 23.

Dynamic virtual storage linking

IBM i virtual storage can be linked (added) to the integrated server while it is active. For example, if the server is running low on storage capacity, you can add virtual storage to the server without shutting it down.

Virtual storage backup and recovery from IBM i

You can back up entire IBM i virtual storage spaces that the integrated server uses along with your other IBM i data. This backup provides a snapshot of the storage that can be used for disaster recovery. The storage can be backed up while the ESX server is active if the storage is linked to the ESX server with **exclusive update** access and all the virtual machines that use the storage are shut down. If storage is linked to the ESX server with **shared update** access, then the ESX server must be shut down before backing up the virtual storage space.

You can restore entire IBM i virtual storage spaces that the integrated server uses. The integrated server must be shut down in order to restore a virtual storage space.

IBM i management interfaces

You can manage IBM i integration features using the *IBM Navigator for i* Web GUI or using IBM i control language (CL) commands. Most integrated server management tasks in this topic are documented using the Web GUI, but links to the corresponding CL commands are also provided.

IBM i management infrastructure for integrated VMware ESX servers

The infrastructure for managing integrated VMware ESX servers from IBM i has the following key pieces:

IBM i software

Most of the software for installing and managing an integrated server runs on IBM i. This software consists of IBM i base operating system functions and various IBM i options such as **Integrated Server Support** (IBM i option 29). The IBM i software enables integrated server installation, IBM i configuration object management, virtual storage management, and much more.

IBM i configuration objects

On the IBM i side of server management, an integrated server is represented by a network server description (NWSD) and several other types of configuration objects. You can stop and restart the server from IBM i by varying the NWSD off and on. See “IBM i configuration objects for integrated servers” on page 46 for more information.

Integrated VMware ESX server

The VMware ESX server that is integrated with IBM i.

ESX platform manager (optional)

Software that manages one or more VMware ESX servers and their virtual servers. VMware vCenter Server is one example of an ESX platform manager.

Management server (integrated Windows server)

The IBM i Integrated Server Support software does not run directly on the VMware ESX server. Instead, an iSCSI attached integrated Windows server serves as a management server for the VMware ESX server. IBM i management tasks, such as shutdown and dynamic virtual storage linking, are sent to the management server over the point-to-point virtual Ethernet connection. Then the task is sent from the management server to the integrated VMware ESX server over a physical Ethernet connection. If an ESX platform manager (VMware vCenter) is configured, the task flows from the management server to the VMware vCenter server and then to the VMware ESX server.

The supported operating system versions on the iSCSI attached integrated Windows server are as follows:

- »Windows Server 2012«
- Windows Server 2008 R2
- Windows Server 2008
- Windows Server 2003

An integrated Windows server can serve as the management server for any number of integrated VMware ESX servers within the same IBM i logical partition. At least one integrated Windows server is required in each IBM i logical partition that hosts integrated VMware ESX servers.

Note: Only a small portion of the integrated Windows server capacity is needed to manage integrated VMware ESX servers. The integrated Windows server can be used for other workloads as well.

Windows utilities

A Windows utility is used to define and manage connection information so that IBM i can manage integrated VMware ESX servers. See "Windows utilities for IBM i integration with VMware ESX servers" for more information.

Windows services

A Windows service is used to perform requests initiated from IBM i to integrated VMware ESX servers that are managed from the integrated Windows server. See "Windows services for IBM i integration with VMware ESX servers" on page 21 for more information.

QVMWINT user profile

This IBM i user profile is used when performing integrated VMware ESX server administration tasks.

- QVMWINT is automatically created when the IBM i Integrated Server Support option is installed on IBM i. This profile is initially disabled.
- As part of the integrated VMware ESX server installation process, the QVMWINT profile must be enabled and then enrolled to the associated management server. The QVMWINT profile must also be created on the integrated VMware ESX server, the associated ESX platform manager (if one is used), or both. The QVMWINT user must have Administrator permissions on the management server and the VMware ESX server or ESX platform manager (vCenter) server.
- The QVMWINT password must match on IBM i, the integrated Windows server, and the integrated VMware ESX server or the associated ESX platform manager (if one is used). Note that at IBM i password level (QPWDLVL) 0 or 1, the QVMWINT password is converted to all lower case characters when it is set on the integrated Windows server.

Windows utilities for IBM i integration with VMware ESX servers

Some of the utilities that provide the IBM i Integrated Server Support features for VMware ESX servers are installed on an associated management server. The following Windows utilities are provided for integrated VMware ESX server management:

IBM i post-install utility for VMware ESX (ibmvmins.exe)

This utility runs on the integrated Windows server that serves as a management server for the integrated VMware ESX server. It is used to install Integrated Server Support for VMware ESX server.

IBM i connection utility for virtualization hosts (ibmvmcon.exe)

This utility runs on the integrated Windows server that serves as a management server for the VMware ESX server. It is used to define and manage connection information so that IBM i can manage integrated

VMware ESX servers. Connection information can be added, deleted, listed, verified, and managed from the connection utility.

Note: The connection information is used by the **IBM i Virtual Server Administration** service to establish connections between the integrated Windows server and the integrated VMware ESX servers that the integrated Windows server manages. See "Windows services for IBM i integration with VMware ESX servers" for more information.

Software Level

View the level of IBM i Integrated Server Support software that is installed on IBM i and on the integrated Windows server. Optionally synchronize the Integrated Server Support software from IBM i to the integrated Windows server.

Windows services for IBM i integration with VMware ESX servers

Some of the programs that provide the IBM i Integrated Server Support features for VMware ESX servers are installed on an associated management server. The following Windows services are provided for integrated VMware ESX server management:

IBM i Virtual Server Administration

Perform requests initiated from IBM i to integrated VMware ESX servers that are managed from the integrated Windows server.

Note: This service uses the connection information that is defined by the **IBM i connection utility for virtualization hosts**. See "Windows utilities for IBM i integration with VMware ESX servers" on page 20 for more information.

Services shared with integrated Windows servers

Several of the Windows services that are used for integrating Windows servers are also used when integrating VMware ESX servers. See "Windows services for IBM i integration with integrated Windows servers" on page 17 for more information.

Integrated server console

The integrated server console is a direct interface to the integrated server operating system.

Depending on your configuration of hardware and software, you can use a monitor, keyboard and mouse that are attached by one of the following methods:

Directly attached monitor, keyboard, and mouse

You can use a monitor, keyboard, and mouse that are directly connected to the System x or BladeCenter product. You interact with the integrated server through these devices exactly as you would with a regular stand-alone server.

A directly attached monitor, keyboard, and mouse are required for some iSCSI configuration tasks.

Remote GUI desktop application

You can use an application such as Microsoft Terminal Services, Remote Desktop, or another third party application to display the integrated server graphical user interface (GUI) desktop on a remote workstation. Most administration tasks that are normally performed on the server directly attached console can be performed on the remote desktop. See the Microsoft Terminal Services or other third party application documentation for information about how to configure and use a remote desktop for the server console.

IMM or RSA II graphical console redirection

For System x products equipped with an Integrated Management Module (IMM) or Remote Supervisor Adapter II (RSA II) service processor, the IMM or RSA II provides full hardware-based graphical console redirection. This redirection means that you can use a local desktop to access and control a remote server using a Web browser.

BladeCenter MM or AMM graphical console redirection

A BladeCenter enclosure (chassis) uses either a Management Module (MM) or an Advanced Management Module (AMM) which provides hardware-based graphical console redirection. This redirection means that you can use a local desktop to access and control a remote server using a Web browser.

Software updates for integrated servers

There are several types of software updates for integrated servers.

Updates to IBM i and firmware

You should update the following software and firmware for integrated servers.

Table 1. Methods for applying software updates for integrated servers

Component	Methods for applying software updates
IBM i, and related licensed products	Apply PTFs. See IBM i PTF group SF99369 for the latest PTFs.
IBM i Integrated Server Support software that runs on the integrated Windows server (for both Windows and VMware ESX integration)	Apply IBM i PTFs and then run a utility from the integrated Windows server. See the IBM i iSCSI Solution Guide .
iSCSI initiator BIOS and firmware	To update the iSCSI initiator firmware in a System x or blade server, see the IBM i iSCSI Solution Guide .
System x or BladeCenter updates	You might need to update the firmware for the System x, blade, or BladeCenter hardware. See the IBM i iSCSI Solution Guide .
Integrated server operating system	Apply updates at the integrated server console using the normal procedures for the operating system.

Updates for integrated Windows servers

The updates for the IBM i Integrated Server Support software that enables Microsoft Windows server to run on the integrated server are separate from the service packs for Windows itself, which you must get from Microsoft.

The process of installing Integrated Server Support software fixes on your integrated Windows server is called synchronization. When you synchronize an integrated Windows server, the integration software on the integrated server is updated to the same release level and service pack level that is on IBM i. The level of code on the Windows side is dependent on the level of code on the IBM i side.

When you synchronize an integrated server, there are four things that can happen:

1. If IBM i has been upgraded to a new release, for example, from i 6.1 to i 7.1, the integration software for the new release will replace that of the old release on the integrated Windows server.
2. If a new IBM i Integrated Server Support service pack has been installed on IBM i, it will be copied over to the integrated Windows server.
3. If an IBM i Integrated Server Support service pack has been removed from IBM i, it will be removed from the integrated Windows server as well, and replaced with the integration software currently existing in IBM i.
4. If the IBM i integration software and integrated Windows server integration software are at the same level, the synchronization operation can still be performed. This allows for recovery of a deleted or damaged integration software file on the integrated Windows server.

In all cases the integrated Windows server will be brought to the same level of integration software which exists in IBM i. See the IBM i iSCSI Solution Guide.

Updates for integrated VMware ESX servers

Install Integrated Server Support software updates to the integrated Windows server that serves as the management server for the integrated VMware ESX server. See the IBM i iSCSI Solution Guide.

Storage management for integrated servers

Integrated servers use virtual storage that is managed by IBM i.

Virtual storage for integrated servers

Integrated servers use virtual storage provided by IBM i instead of physical disk drives attached to the integrated server hardware.

Operating systems, such as Windows and VMware ESX, work with what they see as physical disk drives; there is little or no virtualization of storage at an operating system level. Because IBM i virtualizes all disk storage, you can use chunks of disk space from an IBM i storage pool to form virtual disk drives, which can then be allocated to the integrated server operating system. These virtual disks are also known as storage spaces or virtual storage. Integrated VMware ESX and Windows servers, as well as IBM i, AIX® and Linux running in Power server partitions, see these storage spaces as physical disk drives.

The IBM i object that is used to create virtual storage for an integrated server is called a Network Server Storage Space (NWSSTG), or storage space for short. These storage spaces are stored in the IBM i integrated file system (IFS) in a directory called /QFPNWSSTG. You can use the **File Systems** function of *IBM Navigator for i*, or the Work with Links (WRKLNK) command from an IBM i command line to view the contents of the /QFPNWSSTG directory. This storage space architecture is used by integrated Windows and VMware ESX servers and IBM i, Linux and AIX running in Power server logical partitions.

The amount of disk storage that you create for your servers is taken directly from the IBM i available storage, and each virtual disk is physically scattered across the physical disks in the IBM i storage pool. You can create virtual disks as large as 1 TB if there is available storage in the storage pool.

Storage spaces are different from other IBM i 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 servers need to be able to connect to and format a drive of a fixed size.

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 recover 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 integrated server at any time. These drives can be linked to the server while it is shut down or while it is started (dynamic linking). This means that you do not need to allocate large portions of your IBM i 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 integrated server virtual storage:

- Create new virtual storage (optionally copying preexisting virtual storage)
- Link virtual storage to an integrated server
- Unlink virtual storage from an integrated server
- Expand virtual storage
- Delete virtual storage

Virtual storage operations can be performed in these ways:

- Using *IBM Navigator for i* Web GUI.
- Using IBM i CL commands.

Related tasks:

"Adding virtual storage to integrated servers" on page 80
Use these tasks to add virtual storage to an integrated server.

IBM i storage management for integrated servers

Integrated servers use virtual storage (virtual disks) that are managed by IBM i.

This brief overview of IBM i storage management concepts is intended for administrators who are more familiar with how x86-based servers manage storage. Some techniques, such as defragmenting, are not necessary in an integrated server environment.

When the integrated server operating system is running, it uses a portion of the IBM i disk capacity. For this reason, the administration of integrated server storage has both an IBM i component and an integrated server operating system component. The IBM i component is used to create and link a chunk of storage to the integrated server. Many of the common disk administration tasks encountered in stand-alone servers (disk drivers, addressing, configuration and protection) are eliminated when you use an integrated server.

Disk storage administration tasks such as formatting and partitioning can be performed on integrated servers in exactly the same way as they are on stand-alone servers.

IBM i and disk drives

The key to understanding how disk storage is allocated to integrated servers is an understanding of how IBM i storage management works. The heart of storage management on IBM i is a technology called single-level storage. Single-level storage is a revolutionary storage management architecture that not only gives IBM i outstanding disk I/O performance, but greatly reduces the amount of administration required. IBM i does not directly manage disk drives. Beneath the operating system a level of software (called Licensed Internal Code) "hides" the disk drives and manages the storage of objects on those disk drives. A virtual address space is mapped over the existing disk space and used for addressing objects rather than disk drive IDs, cylinders, and sectors. Needed objects are copied ("paged in") from this address space on disk into the address space of main memory.

The major features of single-level storage are:

- Single storage pool

  The management of physical disk drives is implemented in the Licensed Internal Code, 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 IBM i and they automatically become part of the system ASP. Note that under some circumstances you might 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.

  IBM i 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 the licensed internal code. Therefore, IBM i does not require a Database Administrator. Licensed internal code also ensures that there is no disk fragmentation.
- Single address space

  Memory and disk on IBM i 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.

Because of the way IBM i manages disk data, you do not generally need to worry about partitioning high-growth databases, defragmenting disks, or disk striping on your integrated server. The integrated server uses device drivers to share the IBM i disk drives. These device drivers send and receive disk data to the IBM i storage management subsystem. IBM i storage management handles the hard disks, including spreading the integrated server disk drive images across multiple hard disk drives and applying RAID and file mirroring (if configured). Disk defragmentation software manages logical file fragmentation of the hard disk images. Because IBM i storage management handles these tasks, running a defragmentation program on the integrated server helps primarily in cases where "critical file system structures" can be defragmented.

Storage pools (ASPs)

In IBM i, physical hard disk drives are pooled together into one storage space called a disk pool, also called an auxiliary storage pool (ASP). If your file system runs out of space, you can add a new hard disk drive to the storage pool, and the new storage space will be available immediately. Every system has at least one storage pool, the system storage pool. The system storage pool is always ASP 1. You can configure additional *user* storage pools, numbered 2 - 255. You can use storage pools to distribute your IBM i data over different groups of disks. You can also use this concept to move less important applications or data to your older, slower disk drives. Support for independent ASPs (33-255) is provided through *IBM Navigator for i*. Both the Information Center and *IBM Navigator for i* refer to ASPs as storage pools or disk pools.

Disk protection

IBM i disks can be protected in these ways:

- **RAID-5:** The RAID-5 technique groups several disks together to form an array. Each disk holds checksum information of the other disks in the same array. If a disk fails, the RAID-5 disk controller can re-create the data of the failing disk with the help of the checksum information on the other disks. When you replace a failing disk with a new one, IBM i can rebuild the information from the failed disk on the new (and therefore empty) disk.
- **Mirroring:** Mirroring keeps two copies of data on two different disks. IBM i performs write operations on both disks at the same time, and can simultaneously perform two different read operations on the two disks of a mirrored pair. If one disk fails, IBM i uses information from the second disk. When you replace the failing disk, IBM i copies the data from the intact disk to the new disk.
- **Cross-site mirroring:** Cross-site mirroring, using the operating system geographic mirroring function for independent ASPs, mirrors data on disks at sites that can be separated by a significant distance.

To further increase the level of protection, you can attach the mirrored disks to two different disk controllers. Then if one controller fails, and with it one set of disks, the other controller can keep the system up. On larger Power server models, you can attach controllers to more than one bus. Attaching the two disk controllers that form a mirrored pair to two different buses increases availability even more.

You can define storage pools on IBM i to have different levels of protection or no protection at all. Then you can put applications and data into a storage pool with the right amount of protection, depending on how important their availability is. For more information about IBM i disk protection and availability options, see the Recovering your system topic collection.

Predefined virtual storage and naming for integrated servers

Predefined virtual storage (virtual disk drives) are automatically created when you install the integrated server operating system. The system uses this virtual storage for the integrated server support code and the operating system.

By default, IBM i creates these disks in the system storage pool (ASP), but you can choose a different location during the installation. IBM i also uses these disks to load and start the integrated server.

Predefined virtual storage and naming for integrated Windows servers

Integrated Windows servers have these predefined disks:

Boot and system drive (C)
This drive serves as the system drive. IBM i names this drive *server*1, where *server* is the name of the network server description (NWSD). This disk drive resides in the integrated file system and is automatically linked as the first drive.

The C drive size ranges from 2 GB to 1,000 GB.

Installation source drive (D)
IBM i names this drive *server*2, where *server* is the name of the NWSD. This disk drive resides in the integrated file system and is automatically linked as the second drive. IBM i formats the D drive as a file allocation table (FAT) disk.

Attention:

1. This drive must remain as a FAT drive. Do not make any changes to this drive. IBM i uses this drive to perform code updates, and changing the drive can make performing updates impossible.
2. Some third-party applications such as Citrix require that the drive letter for this drive be changed. This is supported as long as the drive remains linked to the server and has a FAT or FAT32 file system to allow configuration files to be written when the server is started.

Predefined virtual storage and naming for integrated VMware ESX servers

Embedded versions of VMware ESX server do not have any predefined virtual storage.

For installed versions of VMware ESX server, the installation process creates one predefined virtual disk. The disk corresponds to the first drive that the integrated server recognizes:

System disk (/dev/sda)
The VMware ESX operating system is installed on this disk.

You should allow at least 15 GB for this disk.

Do not configure virtual machines on the system drive. Create additional storage spaces and link them to the server for your virtual machines. For most environments, you can configure one virtual machine per storage space to simplify backup and other administration tasks.

Virtual storage linking for integrated servers

Integrated servers do not use physical disk drives. IBM i creates virtual storage (network server storage spaces) within its own file system and integrated servers use them as if they were normal physical disk drives.

To add virtual storage to an integrated server, you create the storage, link it to the server and then format it for the integrated server operating system.

iSCSI-attached integrated servers use only dynamic virtual storage links. The virtual storage link sequence position is assigned dynamically at the time that the virtual storage is linked to an active server. The virtual storage link sequence position can be specified, but it is not used until the server is restarted. The integrated server can either be shut down or active when adding a dynamic virtual storage link.

When dynamically linking virtual storage to an active server, the new virtual storage appears following all other linked virtual storage.

The following table shows the IBM i virtual storage features supported for network server descriptions (NWSDs) with NWSD type *ISCSI and various operating system (OS) types[1].

Table 2. Virtual storage features supported

Feature	**NWSD OS Types**		
	***WIN32 and *WIN64**	***ESX**	***ESXE**
Maximum number of storage spaces that can be linked to the server	64	64	64
Maximum capacity per storage space	1000 GB	1000 GB	1000 GB
Maximum total virtual storage capacity, assuming 1000 GB per storage space	61.5 TB	62.5 TB	62.5 TB
Can link virtual storage while the server is active?	Yes. Exceptions: storage linked at sequence 1-2	Yes. Exception: storage linked at sequence 1	Yes
Can unlink virtual storage while the server is active?	Yes. Exceptions: Storage linked at sequence 1-2 Storage cannot be part of a volume set Storage cannot be a volume mounted in a directory	No	No
Virtual storage format types allowed when linking	*NTFS, *FAT, *FAT32, *OPEN, » *REFS«	*NTFS, *FAT, *FAT32, *OPEN,» *REFS«	*NTFS, *FAT, *FAT32, *OPEN, » *REFS«
Virtual storage access types allowed when linking	Exclusive update, shared update[2]	Exclusive update, shared update[2]	Exclusive update, shared update[2]
Virtual storage requiring exclusive update access type	Storage linked at sequence 1-2	Storage linked at sequence 1	None
Number of shared access type links	62[2]	63[2]	64[2]

Note:

1. See the Create Network Server Desc (CRTNWSD) command documentation for a description of the NWSD types and the associated operating system (OS) types.
2. Shared storage spaces can be linked to multiple (up to 50) VMware ESX servers (OS types *ESX or *ESXE), but just one Windows server (OS types *WIN32 or *WIN64). Storage spaces cannot be linked to more than one Windows server at a time.

Network server storage spaces can reside in either the IBM i system storage pool (ASP 1) or a user storage pool. You can copy one storage space to another to move it to a different storage pool.

Network server storage spaces are one of the two types of IBM i storage that integrated servers use. Integrated servers can also access resources on IBM i that an administrator has shared with the network by using IBM i NetServer.

After you create virtual storage and link it to an integrated server, you must partition and format the storage using the standard utilities provided by the integrated server operating system.

Related tasks:

"Linking virtual storage to integrated servers" on page 83
Integrated servers can only access virtual storage that is linked to the Network Server Description (NWSD) for the server.

IBM i tape and optical devices shared with integrated Windows servers

Integrated Windows servers can use supported IBM i tape and optical devices.

Supported IBM i devices can be used by the integrated Windows server as if they were local devices on Windows for such tasks as installing applications and backing up data. A device can be used only by IBM i or one integrated Windows server at a time.

A subset of IBM i tape devices are supported for use with various Windows versions. For example, only IBM i **virtual** tape devices can be used with Windows Server 2008. See the IBM i iSCSI Solution Guide Web page for information about IBM i devices that have been tested with iSCSI-attached integrated Windows servers.

Note: IBM i devices cannot be used by iSCSI-attached VMware ESX servers.

Networking concepts for integrated servers

iSCSI-attached integrated servers use several types of network connections.

Service processor connection for integrated servers

This physical connection is required so that the hosting IBM i partition can communicate with the service processor of the initiator (System x or BladeCenter) system.

The connection can consist of a simple and switched network or a more complex and routed network. IBM i Integrated Server Support uses this connection to manage the state of the hosted system.

At one end of the connection is a LAN adapter or adapters that are controlled by IBM i. This LAN adapter can be available for other uses. The IP address and other attributes of this adapter are controlled using standard IBM i configuration methods. IBM i can automatically connect to the service processor using one or more IBM i TCP interfaces that are already configured.

At the other end of the connection is the service processor. The service processor has its own Ethernet port and TCP/IP stack. This TCP/IP stack is active whenever the system power cord is plugged into an energized alternating current (AC) outlet, even if the system is not in a powered on state.

Connection

There are multiple options that IBM i offers for connecting to the service processor. For more information, see "Service processor connection methods" on page 31.

Performance and maximum transmission unit (MTU)

There is not a requirement or advantage to having a high speed network or using a large MTU for the service processor connection.

Security

The security capabilities of your service processor hardware may affect your decision to use an isolated network or a shared network to provide the service processor connection. For more information, see "Configuring security between IBM i and integrated servers" on page 75.

Related concepts:

"Service processor functions and support"
Use the information from the IBM i remote system and service processor configurations to connect to and manage iSCSI-attached integrated servers.

"Service processor connection methods" on page 31
IBM i connects to the initiator blade or System x hardware on the network. Multiple connection methods are available.

Related tasks:

Configuring service processor connection (Deprecated)
Use the information from the IBM i remote system and service processor configurations to connect to the hardware of integrated System x and blade servers.

Service processor functions and support

Use the information from the IBM i remote system and service processor configurations to connect to and manage iSCSI-attached integrated servers.

Initiator systems are identified by information stored in the remote system configuration and the service processor configuration objects in IBM i.

This connection is different than the connection between the IBM i iSCSI target adapter and the iSCSI initiator adapter in the remote server. The LAN adapter for the service processor of the remote server must be attached to a network that is reachable by a LAN adapter that is installed and assigned to your IBM i partition.

Both the IBM i objects and the service processor must be configured. You can configure the connection method used in the IBM i network server configuration objects.

Static addressing for service processors

The service processor is configured with a specific IP address or host name.

Dynamic addressing for service processors

Using DHCP to obtain the service processor IP address is not supported for the IBM i integrated server solution. Use static addressing for the service processor.

Note: If a specific IP address or host name has not been set yet for the service processor, the factory default for most service processors is to use DHCP to obtain an IP address. The service processor initializes immediately when the server receives power and starts the DHCP process. This DHCP server is distinct from the DHCP server that is built into the IBM i side of the iSCSI network to assist with iSCSI boot of the integrated server operating system. If a service processor IP address cannot be obtained with DHCP, the service processor uses the default static IP address of 192.168.70.125. You can use the service processor Web interface to set a static address for the service processor, using the IP address obtained using DHCP, or the default IP address.

Supported functions by service processor type

The configuration options depend on the type of service processor. For information about the capabilities of each type of service processor, see the IBM i iSCSI Solution Guide.

Related concepts:

"Service processor connection for integrated servers" on page 29
This physical connection is required so that the hosting IBM i partition can communicate with the service processor of the initiator (System x or BladeCenter) system.

"Service processor connection methods"
IBM i connects to the initiator blade or System x hardware on the network. Multiple connection methods are available.

Service processor connection methods

IBM i connects to the initiator blade or System x hardware on the network. Multiple connection methods are available.

For information about the service processor connection methods, see the IBM i iSCSI Solution Guide.

Related concepts:

"Service processor connection for integrated servers" on page 29
This physical connection is required so that the hosting IBM i partition can communicate with the service processor of the initiator (System x or BladeCenter) system.

"Service processor functions and support" on page 30
Use the information from the IBM i remote system and service processor configurations to connect to and manage iSCSI-attached integrated servers.

iSCSI network for integrated servers

This physical network connects iSCSI target adapters in the hosting IBM i partition with iSCSI initiator adapters in the System x or BladeCenter system.

The iSCSI network is typically a simple, switched, Gigabit Ethernet network. The iSCSI target and initiator adapters can be connected directly to each other without a switch. Two kinds of traffic flow over this connection: storage (SCSI) and virtual Ethernet (LAN).

On one side of the network is an iSCSI target adapter or adapters controlled by IBM i. Each iSCSI target adapter port has up to two IP addresses: one for SCSI and one for LAN. For a hardware target (iSCSI HBA), separate IP addresses are used for the SCSI and LAN connections. For a software target (Ethernet NIC), the LAN connection uses the same IP address as the SCSI connection. You configure the IP addresses and other attributes of an adapter in an IBM i device description object known as the network server host adapter (NWSH). For more information, see "Network server host adapters" on page 48. Each iSCSI target adapter controlled by IBM i needs its own NWSH object. When you vary on an NWSH, an iSCSI target adapter controlled by IBM i uses the configured values. If you want different values to be used, you must vary off the NWSH, change the NWSH configuration, and vary on the NWSH again.

The iSCSI protocol is implemented differently, depending on the type of iSCSI target adapter:

Software target (Ethernet NIC)

The iSCSI protocol is implemented in IBM i, so IBM i resources (for example, CPU and memory) are used for the iSCSI protocol. The IBM i TCP/IP stack is aware of the IP address configured for the iSCSI target adapter.

Hardware target (iSCSI HBA)

The iSCSI protocol is implemented in firmware on the iSCSI adapter, so the iSCSI protocol is offloaded from IBM i. The TCP/IP stack is also implemented in hardware and is independent of the normal IBM i TCP/IP stack. The IBM i TCP/IP stack is unaware of the IP addresses configured for the iSCSI target adapter.

On the other side of the network is an iSCSI initiator adapter or adapters for the initiator system. You configure the IP addresses and other attributes of these adapters in an IBM i object known as the remote system configuration. For more information, see "Remote system configuration" on page 49. This configuration differs from the IBM i network server host adapter object in several ways:

- You can configure an iSCSI initiator adapter port with 1 or 2 IP addresses: SCSI, LAN, or both. There must be at least one SCSI and one LAN IP address among all the configured adapters.
- Whenever you configure an IP address for an iSCSI initiator adapter, you must also configure the corresponding adapter MAC address. Be careful to configure MAC addresses correctly.
- You configure all the iSCSI initiator adapters for an initiator system in the same IBM i remote system configuration. When the integrated server is later varied on, IBM i automatically ensures that iSCSI initiator adapters in the initiator system use values in the IBM i remote system configuration. If you want different values to be used, you must change the remote system configuration and vary on the server again.

- The iSCSI protocol is implemented differently, depending on the type of iSCSI initiator adapter:

  Software initiator (Ethernet NIC)
  The iSCSI protocol is implemented in the integrated server operating system, so integrated server resources (for example, CPU and memory) are used for the iSCSI protocol. The integrated server operating system TCP/IP stack is aware of the IP addresses configured for the iSCSI initiator adapter.

  Hardware initiator (iSCSI HBA)
  The iSCSI protocol is implemented in firmware on the iSCSI adapter, so the iSCSI protocol is offloaded from the integrated server. The SCSI traffic uses the iSCSI initiator adapter hardware TCP/IP stack, but LAN traffic uses the integrated server operating system TCP/IP stack. Consequently, the integrated server operating system TCP/IP stack is unaware of the iSCSI initiator adapter SCSI IP address, but is aware of the LAN IP address.

Note:

1. In IBM i configuration objects, network interface information is labeled as local or remote. These terms are relative to IBM i. Local interface information is for the IBM i side. Remote interface information is for the initiator system side.
2. The NWSH and the remote system configuration define IP address information for opposite sides of the iSCSI network. When connected by a simple, switched network, the following rules apply:
   - The SCSI IP addresses in these two objects that are connected by a switch must be in the same subnet. For example, with IP addresses of the form a.b.x.y and 255.255.255.0 subnet masks, a.b.x must be the same value for both objects.
   - The LAN IP addresses in these two objects that are connected by a switch must be in the same subnet.
   - In the NWSH, use the default value (none) for the gateway elements.
   - In the remote system configuration, use the default value (none) for the gateway elements.

Scaling the iSCSI Network

After you have installed an integrated server, you can scale the iSCSI network.

The basic installation process addresses integrated servers that use one IBM i iSCSI target and up to two System x or blade iSCSI initiators. After the server is installed, you can configure additional iSCSI targets or initiators if needed.

- Configure multipath I/O for the integrated server storage. See “Multipath I/O for integrated servers” on page 52
- Refer to the Scaling your iSCSI network chapter in the Implementing Integrated Windows Server through iSCSI to System i5® (www.redbooks.ibm.com/abstracts/sg247230.html) Redbooks® publication for more information.

Integrated DHCP server

There are several methods for delivering boot information to the initiator system. The default method of delivering IP and storage information to boot the integrated server uses an integrated Dynamic Host Configuration Protocol (DHCP) server on the IBM i side of the iSCSI network. For more information, see "Integrated DHCP server for integrated servers."

Even with DHCP, the IP address might be considered static because the DHCP server associates a single IP address with a MAC address.

Managing IBM i iSCSI target adapter function

Paths configured in the network server description control what storage and virtual Ethernet traffic, if any, can flow over an IBM i iSCSI target adapter. For more information, see the IBM i iSCSI Solution Guide.

Multiple initiator systems can use an IBM i iSCSI target adapter simultaneously if multiple network server descriptions use the same NWSH object.

Managing iSCSI initiator adapter function

You can configure an iSCSI initiator adapter with a SCSI IP address, a LAN IP address, or both. A SCSI IP address enables storage traffic, and a LAN IP address enables virtual Ethernet traffic.

Use of the iSCSI initiator adapter as a general-purpose external network connection is not supported. For more information about external network connections, see "Physical networks for integrated servers" on page 41.

For integrated Windows servers, each virtual Ethernet adapter is automatically assigned to an iSCSI initiator adapter. There is an option to select particular iSCSI initiator adapter on the advanced properties tab of each virtual Ethernet adapter. See the IBM i iSCSI Solution Guide.

Other considerations

The following items are additional considerations for iSCSI adapters.

- The iSCSI network only uses Internet Protocol version 4.
- The frame format is Ethernet version 2.
- The iSCSI network does not support Network Address Translation.

Security

For information about securing storage and virtual Ethernet traffic, see "Network security for integrated servers" on page 41.

Integrated DHCP server for integrated servers

The IBM i Integrated Server Support option provides an integrated DHCP server that is used for communication with iSCSI initiators in integrated servers.

This integrated DHCP server is not a general purpose DHCP server. It is intended to exclusively deploy boot parameters to the hosted server iSCSI initiator. This integrated DHCP server cannot be used for other types of networking. You should use the default configuration for most environments.

The integrated DHCP server is used to deploy boot parameters to the hosted-server iSCSI initiator when the **Dynamically delivered to the remote**

system via DHCP option is specified in the IBM i remote system configuration and when **AUTO** or **DHCP** mode is specified in the hosted-server iSCSI initiator. The following parameters are deployed to the hosted-server iSCSI initiator when an NWSD is varied on:

- IP addresses for the IBM i iSCSI target boot devices and the BladeCenter blade or System x iSCSI initiator devices.
- iSCSI Qualified Names (IQNs) that represent the target and initiator devices.

Both of these sets of IP addresses and IQNs are in the IBM i configuration objects used to define the hosted server. The target IP address is defined in the NWSH object. The initiator IP address and initiator IQN are defined in the remote system configuration. The target IQN is automatically configured and defined in the NWSD object. For more information about these objects refer to "IBM i configuration objects for integrated servers" on page 46.

Networking between IBM i and integrated servers

IBM i uses network connections to communicate with integrated servers for some administrative functions, such as linking storage and shutting down the server. Integrated Windows servers use a point-to-point virtual Ethernet network and integrated VMware ESX servers use multiple networks.

Point-to-point virtual Ethernet for integrated Windows servers

IBM i uses the point-to-point virtual network to communicate with integrated Windows servers. This type of virtual Ethernet network is specifically for integrated Windows servers and is different from the virtual Ethernet networks used for inter-partition communication on your Power server.

IBM i communicates with integrated Windows servers over a point-to-point virtual Ethernet network. When an integrated server is installed, a special virtual network is created between the integrated server and a controlling IBM i partition. This network is called point-to-point because it has only two end points (the integrated server, and the IBM i server). This point-to-point virtual Ethernet is emulated within IBM i and no additional physical network adapters or cables are used. In IBM i, it is configured as an Ethernet line description with Port Number value *VRTETHPTP.

When you install an integrated Windows server, IBM i automatically configures a point-to-point virtual Ethernet.

A point-to-point virtual Ethernet connection and a virtual Ethernet network are different in the following ways:

- A point-to-point virtual Ethernet is configured differently and can only have two end points (the IBM i system and an integrated Windows server).
- A point-to-point virtual Ethernet only supports the TCP/IP protocol, and by default uses restricted IP addresses in private domains, so the addresses are not passed through gateways or routers.

These addresses take the form of 192.168.xxx.yyy, where xxx ranges from 100 to 254 and results in a unique class C network. For example, the IBM i side of the point-to-point network is given the IP address 192.168.100.1, and the Windows operating system side has 192.168.100.2. As you create multiple integrated Windows servers, yyy is incremented so that the point-to-point network for each integrated Windows server is on a unique subnet.

You can automatically assign these IP addresses when you install an integrated Windows server, or you can manually configure them to prevent TCP/IP address collisions with other hosts on the system.

Multiple networks for integrated VMware ESX servers

IBM i does not communicate directly with a VMware ESX server to perform VMware ESX server management tasks (for example, to shut down the ESX server). Instead, IBM i uses an intermediate integrated Windows server that serves as a management server for the VMware ESX server. IBM i uses the point-to-point virtual network to communicate with the integrated Windows server. The integrated Windows server then uses a physical network between the two integrated servers to communicate with the VMware ESX server or an ESX platform manager to perform the task.

Virtual Ethernet networks for integrated Windows servers

Integrated servers can use a virtual Ethernet network that is configured on a Power server to communicate with the hosting IBM i partition, another partition, or other integrated servers.

Virtual Ethernet networks that do not include more than one logical partition

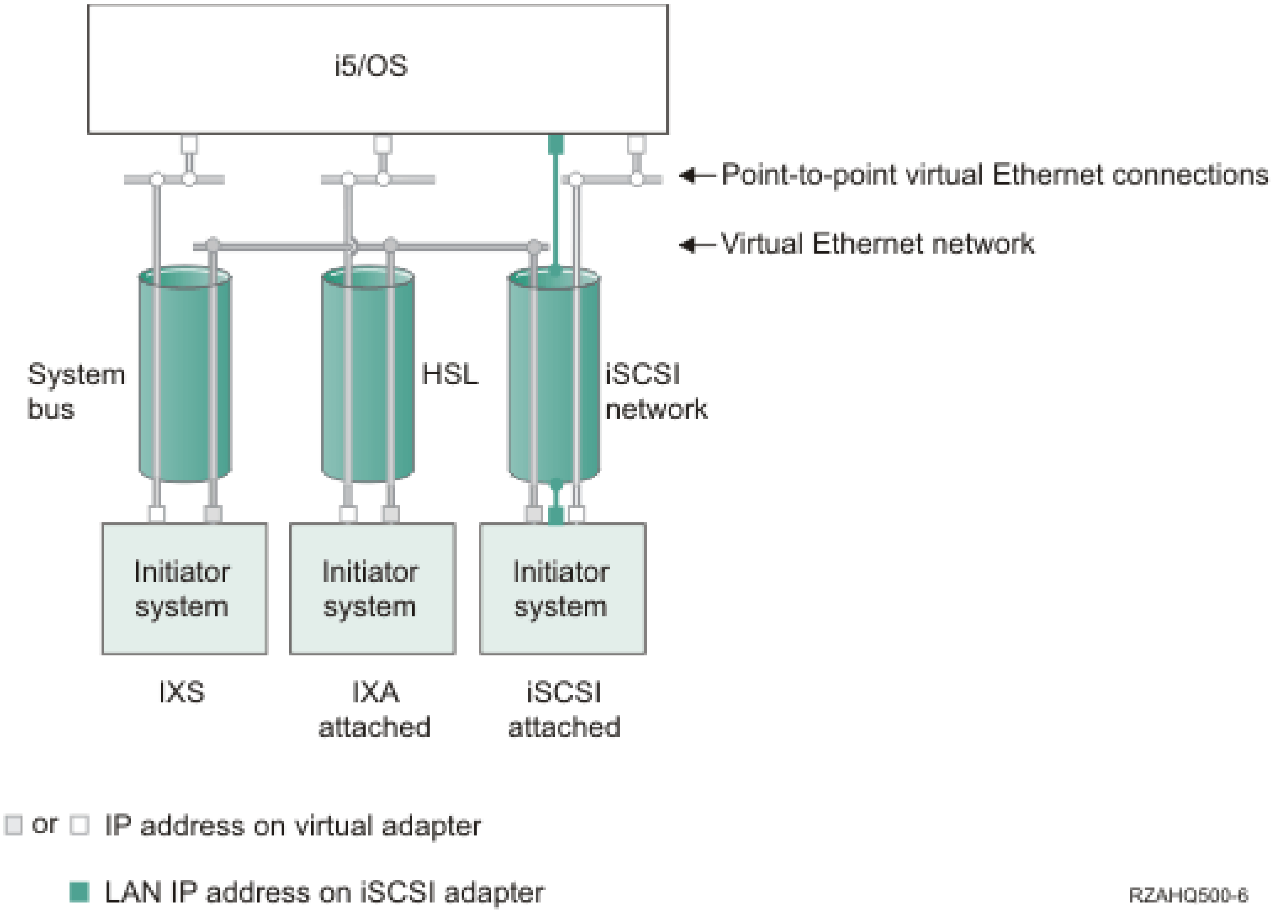

Figure 6. System bus, HSL, and iSCSI network tunnels

iSCSI-attached systems, Integrated xSeries Servers (IXSs), and systems attached using an Integrated xSeries Adapter (IXA) can all participate in virtual Ethernet networks and can communicate with each other.

- For iSCSI-attached servers, virtual Ethernet traffic is tunneled through a physical iSCSI network. Virtual Ethernet is needed when an iSCSI network is present for several reasons:
  - Virtual Ethernet can work with other virtual Ethernet support on your Power server.
  - Virtual Ethernet can provide multiple isolated virtual networks through each iSCSI target adapter even when switches in the iSCSI network do not support IEEE 802.1Q VLANs
  - Integrated servers can communicate with each other even if they are each attached by Ethernet switches that are not connected to each other.
- For IXSs, virtual Ethernet traffic flows over buses for Power servers.
- For IXA-attached servers, virtual Ethernet traffic flows through HSL cables.

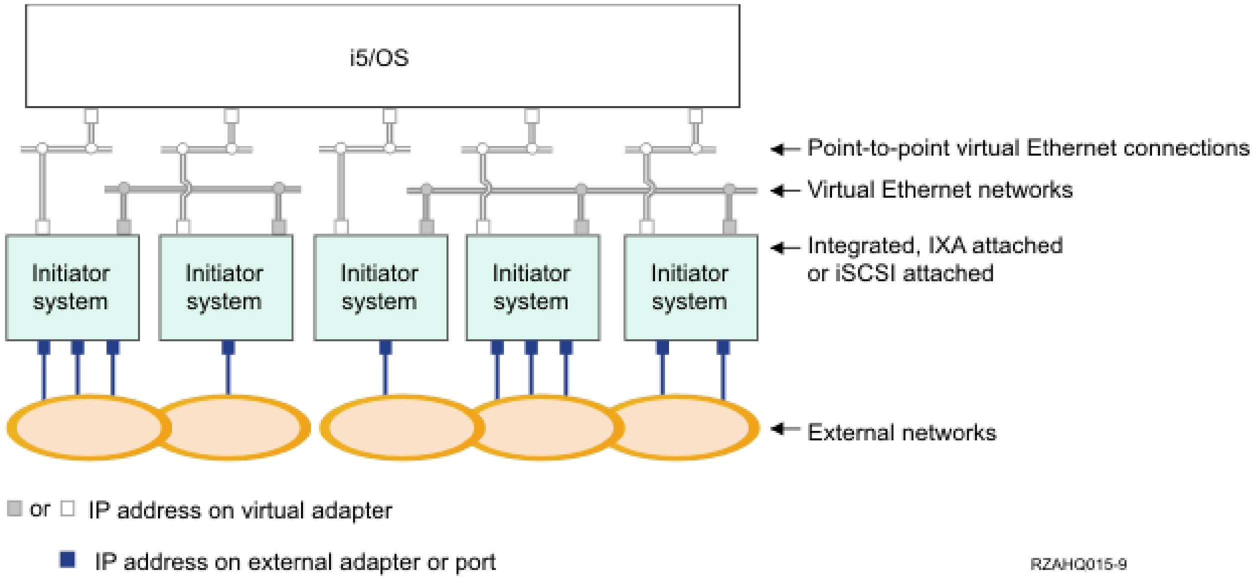

Figure 7. Two isolated groups of integrated Windows servers on the same Power server. Each group has its own virtual Ethernet network.

This figure illustrates how virtual networks work within the Power server. There are five separate integrated Windows servers. They are all connected to the single controlling IBM i partition with point-to-point virtual Ethernet networks. The boxes on the bottom of the integrated servers represent physical network adapter cards that allow the machines to make external network connections. The ovals to which they are connected represent external networks. Finally, there are two separate virtual Ethernet networks. Each integrated Windows server can participate in up to four virtual Ethernet networks simultaneously.

Like point-to-point virtual Ethernet, virtual Ethernet networks are configured through Ethernet line descriptions. An integrated server is connected to a virtual Ethernet network when its IBM i configuration (NWSD) is configured to have an Ethernet line description port number with a value of *VRTETH0 through *VRTETH9. Integrated servers that have NWSDs configured with the same port number values are connected to the same virtual Ethernet network. In the figure, the IBM i side of the line descriptions is not shown. Unlike when you use point-to-point virtual Ethernet, you do not configure a TCP/IP address on the IBM i side of a line description that is used in a virtual Ethernet network.

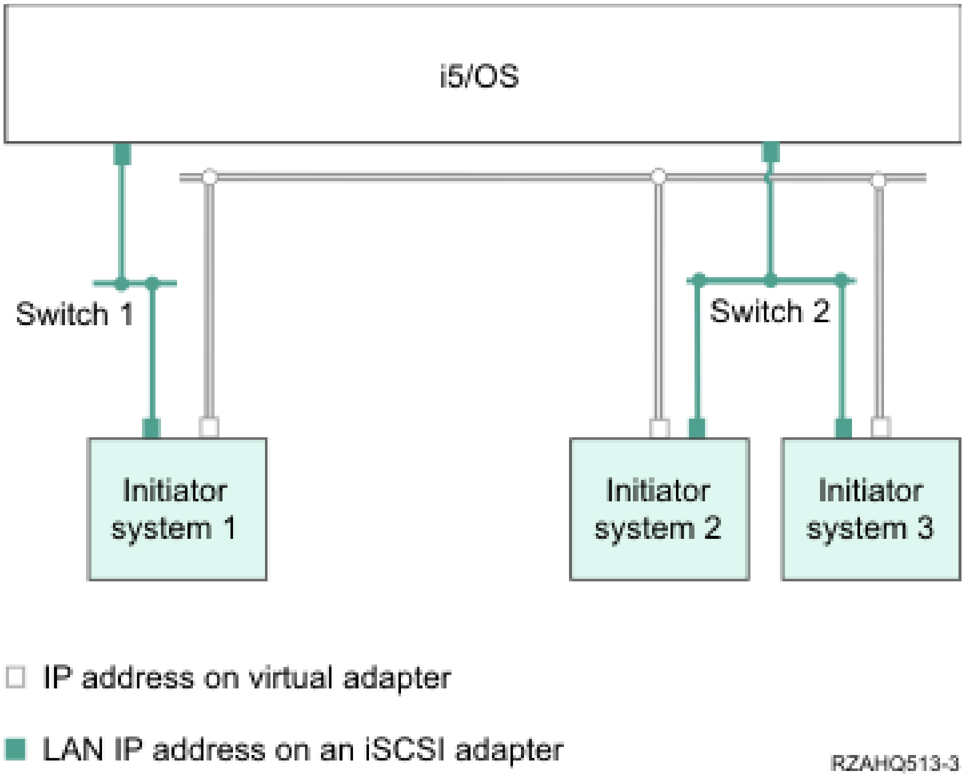

Figure 8. Virtual Ethernet tunneled through iSCSI networks

Virtual Ethernet tunneled through iSCSI networks has some special characteristics that are illustrated in Figure 8.

- Initiator system 1 can communicate with Initiator system 2 and with Initiator system 3, even though separate iSCSI networks (separate physical switches) are involved.
- Virtual Ethernet communication between Initiator system 2 and Initiator system 3 involves the Power server, even though both of these initiator systems are connected to the same physical switch.

Virtual Ethernet networks that include more than one logical partition

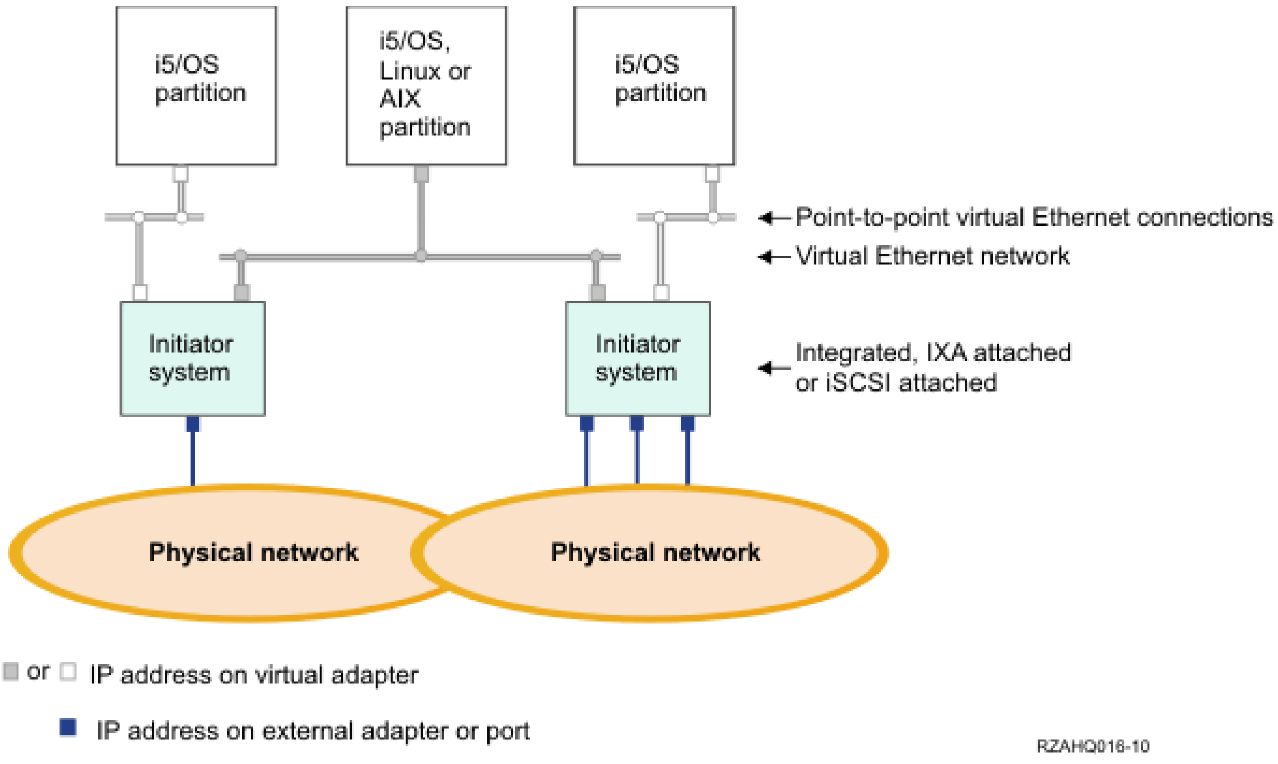
i5/OS partition
i5/OS, Linux or AIX partition
i5/OS partition
Point-to-point virtual Ethernet connections
Virtual Ethernet network
Initiator system
Initiator system
Integrated, IXA attached or iSCSI attached
Physical network
Physical network
▨ or □ IP address on virtual adapter
■ IP address on external adapter or port
RZAHQ016-10

Figure 9. A simple, inter-partition virtual Ethernet network

In Figure 9, the Power server has been partitioned and three separate virtual servers (logical partitions) have been created inside the Power server. Three virtual networks are represented in the figure; two point-to-point virtual Ethernet networks and one virtual Ethernet network. Each integrated server has a point-to-point virtual Ethernet network for communicating with its controlling partition. In this example, the virtual Ethernet network has three participants: two integrated servers, each controlled by a different IBM i partition, and a third partition running IBM i or another operating system. This is called an inter-partition virtual Ethernet network.

Inter-partition connections exist between partitions or integrated servers that are assigned the same virtual LAN ID. Participating integrated servers do not support virtual LAN IDs directly. Instead, each participating integrated server needs an Ethernet line description that associates a port value such as *VRTETH1 with a virtual adapter that has a virtual LAN ID. To create the virtual adapter, see Logical partitioning in the IBM Systems Hardware Information Center. Note that within the same partition, Windows servers can communicate with each other by using the same virtual Ethernet port number.

Related reference:

IBM i iSCSI Solution Guide

Physical networks for integrated servers

Integrated servers can use an integrated Ethernet controller, a network adapter installed in a PCI slot, or a BladeCenter I/O module to connect to an external network.

These are the normal networks which all integrated servers use, created by networking through physical adapters controlled by the integrated server operating system.

In an integrated server you can use any integrated network adapter or install a network adapter card as you would in a stand-alone server.

Network security for integrated servers

iSCSI-attached servers use two types of networks. You can add security to both the service processor connection and the iSCSI network.

Service processor connection security

Service processor security can involve one or more of the following mechanisms:

- Service processor password
- Network isolation and physical security

iSCSI network security

Consider the following types of iSCSI network traffic:

- Storage security can involve one or more of the following mechanisms:
  - Network isolation and physical security
  - Firewalls
  - Challenge Handshake Authentication Protocol (CHAP)
- Virtual Ethernet security can involve one or more of the following mechanisms:
  - Network isolation and physical security
  - Firewalls
  - Secure Sockets Layer (SSL) connection for sensitive data during user enrollment and remote command submission

Service processor password

This password is managed by IBM i and is used when IBM i starts a conversation with the service processor of the initiator system. The service processor checks the password to ensure that the IBM i configuration is authentic. New service processors have a default name and password. IBM i provides a way to change the password.

Network isolation and physical security

Network isolation minimizes the risk that data might be accessed by unauthorized devices or that data that could be modified when it traverses the network. You can create an isolated network by using a dedicated Ethernet switch or a dedicated virtual local area network (VLAN) on a physical VLAN switch or network. Note that IBM i iSCSI target adapters do not support VLAN tagging. When you configure a VLAN switch, do **not** configure it to add a special VLAN tag to the frames.

Physical security involves physical barriers that limit access to the network equipment and the network end points at some level (locked rack enclosures, locked rooms, locked buildings, and so on).

Firewalls

A firewall can be used between a shared network and IBM i to protect IBM i from unwanted network traffic. Similarly, a firewall can be used between a shared network and the initiator system to protect the initiator system from unwanted network traffic.

iSCSI-attached system traffic has the following attributes that should be helpful when configuring a firewall:

- iSCSI target and initiator adapters have static IP addresses (there is a DHCP boot mode, but the IP addresses involved are statically pre-configured)
- UDP and TCP ports that are deterministic and configurable. Each virtual Ethernet adapter on the hosted system uses a different UDP port to tunnel through the iSCSI network. Virtual Ethernet packets are encapsulated as follows, from outer header to inner header:
  - MAC and IP header for the iSCSI adapter using LAN (not SCSI) addresses.
  - UDP header. See “Configuring a firewall to allow integrated server connections” on page 76 for information about optionally controlling UDP port selection.
  - MAC and IP headers for the virtual Ethernet adapter.

Challenge Handshake Authentication Protocol (CHAP)

CHAP protects against the possibility of an unauthorized system using an authorized system's iSCSI name to access storage. CHAP does not encrypt network traffic, but rather limits which system can access an IBM i storage path.

CHAP involves configuring a secret that both IBM i and the hosted system must know. Short CHAP secrets may be exposed if the CHAP packet exchange is recorded with a LAN sniffer and analyzed offline. The CHAP secret should be random and long enough to make this method of attack impractical. IBM i can generate an appropriate secret. A hosted system uses the same CHAP secret to access all of its configured IBM i storage paths.

You can configure either target or bidirectional CHAP. Target CHAP authenticates the iSCSI initiator adapters that connect to the iSCSI target adapter in IBM i. Bidirectional CHAP involves both target CHAP and initiator CHAP. Initiator CHAP authenticates the iSCSI target adapters that connect to the iSCSI initiator adapter in the System x or blade hardware. Note that bidirectional CHAP is only supported for the iSCSI initiator that is used as the boot device.

Secure Sockets Layer (SSL) connection between IBM i and Windows

The IBM i Integrated Server Support option includes user enrollment and remote command submission functions, which may transfer sensitive data over the point-to-point virtual Ethernet. These applications automatically set up an SSL connection to encrypt their sensitive network traffic, and to ensure that each side of the conversation is authentic, based on automatically installed digital certificates. This security feature is provided by default and is not configurable. File data, command results, and traffic for other applications are not protected by this SSL connection.

Performance concepts for integrated servers

Integrated server performance is affected by the configuration of the virtual storage and network for the integrated server.

The iSCSI-attached systems have their own memory and one or more processors, but share the IBM i hard disk storage through virtual (simulated) disk drives (virtual storage). The disk drives are allocated to integrated servers by creating an IBM i virtual disk (network server storage space). The major difference between the integrated servers and stand-alone servers is that stand-alone servers tend to use dedicated disk drives and the integrated servers use IBM i storage spaces as virtual disks. Integrated Windows servers also include optional features to share IBM i tape, CD and DVD drives. Integrated Windows servers can use high-speed virtual Ethernet networks to communicate with other integrated servers or Power server logical partitions.

The use of IBM i storage spaces (virtual drives) provides performance benefits that are not typically available in stand-alone environments without significant storage fabric investment and maintenance costs. However, it also imposes some limitations. You should consider these limitations when planning and configuring integrated servers. The information below highlights some considerations affecting performance.

Related reference:

IBM i Performance Capabilities Reference

IBM i Performance Management

Storage performance for integrated servers

Storage performance depends on the configuration of the integrated server environment.

For performing processor or memory intensive work on an integrated server, the performance characteristics are equivalent to a stand alone server using dedicated disk drives. Since the integrated server disk drives (virtual storage) are allocated out of IBM i storage, the disk performance is dependent on IBM i.

Greater disk performance capacity with IBM i shared disks

On most standalone servers a few disks are dedicated to each server. For applications with a small average disk load, the performance is adequate. However, there can be periods of time where the server performance is limited by the capacity of those few dedicated disks.

When the same group of servers is integrated with IBM i, the virtual disks are spread across more IBM i hard disks. The total average disk load does not need to be any greater than for a group of servers with dedicated disks. But, when an individual server temporarily needs more disk performance capacity, it is available through the larger set of IBM i disks.

On servers with dedicated disks, the disk response times tend to be relatively steady.

On integrated Windows servers, you might take advantage of the predictable response time and configure the Windows Performance Monitor to produce alerts when disk response times exceed typical thresholds and indicate exceptional conditions which might need your attention.

On an integrated server, the IBM i storage, CPU and memory are shared between the integrated server and IBM i applications. It is normal for disk response to swing through a larger range. Short periods might occur where I/O operations from multiple integrated servers, or other IBM i operations contend for the same disk. Some disk intensive IBM i applications (like SAV and RST) can reduce the disk performance seen on the integrated server for a period of time. This can make it more difficult to choose a threshold value for short time periods.

Storage space balancing for integrated servers

The disks in the IBM i storage pool (ASP) may be configured to be unprotected, parity protected (RAID-5), or with mirrored protection. Unprotected disks provide no protection against disk failures. Parity protected disks maintain parity sets which allow recovery if a disk fails in a parity set (but at a performance cost). Mirroring provides protection against disk failures, but with much better performance than parity. The integrated server gains the benefits of the efficient IBM i storage architecture, regardless of how an ASP or independent ASP is configured.

IBM i has functions to help maintain the efficient spread of data across the disks. One example is the Start Disk Reorganization (STRDSKRGZ) operation, which balances disk storage utilization. Another is the “Add units to ASPs and balance data” available when hard disk resources are assigned to an ASP.

The location of the data associated with a storage space is usually automatically managed by IBM i. There is no need to configure striped volumes or software RAID of the disks within the integrated server operating system. Configuring these features in the integrated server operating system might actually slow the effective disk operations. For integrated Windows servers, continue to defragment the associated disk on Windows to maintain efficient file-system data structures.

You can monitor how well IBM i is fulfilling the integrated server's disk requirements by using the Work with Disk Status (WRKDSKSTS) and Work with NWS Storage Spaces (WRKNWSSTG) commands.

For integrated Windows servers, you can use the Microsoft Windows Performance Monitor as you would on any other server. See your Microsoft Windows documentation for information about using the Performance Monitor.

Consider the entire group of IBM i disks when you evaluate storage bottlenecks for integrated Windows servers

The IBM i storage space appears as one disk drive within Windows. When the Physical Disk average queue length (in Windows Performance Monitor) exceeds two, the server performance is not necessarily disk constrained. Assuming that memory paging issues have been ruled out, a queue length of two or a Windows disk utilization of 100% only points to a storage bottleneck if there is only one physical disk drive to perform the operations. There are usually multiple disks on IBM i in the storage space ASP operating in parallel. Typically, two times the number of disks in the ASP might point toward a disk bottleneck. You might also need to account for the average queue lengths of all the servers using the storage ASP.

Related reference:

IBM i Performance Capabilities Reference

Virtual Ethernet performance for integrated Windows servers

The Virtual Ethernet point-to-point connection is the default virtual network connection between the hosting IBM i partition and each integrated Windows server. The-point-to-point connection is used primarily for administrative operations which are part of the integration environment.

The IBM i and Windows CPU utilization cost of using the point-to-point connection is similar to the utilization cost of using a hardware network adapter. The connection is high speed, but total bandwidth is always shared with storage, tape and other integrated server operations. You can separate virtual Ethernet operations from storage operations by using another iSCSI target adapter.

A Virtual Ethernet connection between two or more integrated servers uses the IBM i CPU to switch the traffic between servers, even when IBM i is not an endpoint of the traffic. For most connections this utilization won't be significant. If you expect high sustained network loads across the virtual Ethernet connection between integrated Windows servers, you might want to balance the cost of using the Virtual Ethernet internal switch against external network adapters on the integrated servers.

Related reference:

IBM i Performance Capabilities Reference

MTU considerations for the iSCSI network

By default, iSCSI normally uses standard 1500 byte frames. You can configure the network to use other Ethernet frame sizes to adjust network performance.

High bandwidth and low latency is desirable for the iSCSI network. Storage and virtual Ethernet can take advantage of a maximum transmission unit (MTU) up to a 9000 byte ‘jumbo' frame if the iSCSI network supports the larger MTU. As a rule of thumb, a larger MTU typically decreases the amount of CPU utilization that is required on IBM i and the integrated server.

- Jumbo frames significantly improve performance for a software initiator to software target iSCSI network. Therefore, if your iSCSI network uses all software initiators and all software targets, and the network switches support jumbo frames, then use jumbo frames.
- However, if your iSCSI network uses any hardware initiators or hardware targets, or if the network switches do not support jumbo frames, then use standard frames.

Note: The frame sizes discussed here do not include the Ethernet 14 byte MAC header.

MTU considerations for each component of the iSCSI network:

iSCSI target

IBM i iSCSI target adapters automatically negotiate an MTU, up to 9000 bytes, that is compatible with initiators using the TCP/IP protocol. Therefore, you do not need to configure an MTU for the iSCSI target.

iSCSI initiator

iSCSI initiator adapters default to a frame size that can be transported in a standard 1500 byte Ethernet frame.

Hardware initiators (iSCSI HBAs) can be configured to use up to 9000 byte MTUs.

Some software initiators (Ethernet NICs) support larger MTUs and some do not. Check your Ethernet NIC documentation to determine if the Ethernet NIC can use a larger MTU.

If you want to use an MTU larger than 1500 bytes, you must configure it at each iSCSI initiator adapter. See *Changing the iSCSI initiator MTU* in the IBM i iSCSI Solution Guide .

Switch

Network switches typically use a default 1500 byte MTU.

Some switches support larger MTUs and some do not. Check your switch documentation to determine if the switch can use a larger MTU.

If you want to use an MTU larger than 1500 bytes, you must configure it on the switch. See your switch documentation for more information.

Tip: If you try to install an integrated server using jumbo frames and the installation fails, it could be a sign that your network hardware does not support jumbo frames.

IBM i configuration objects for integrated servers

IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

The following figure shows the objects that IBM i (IBM i) uses to configure iSCSI-attached integrated servers.

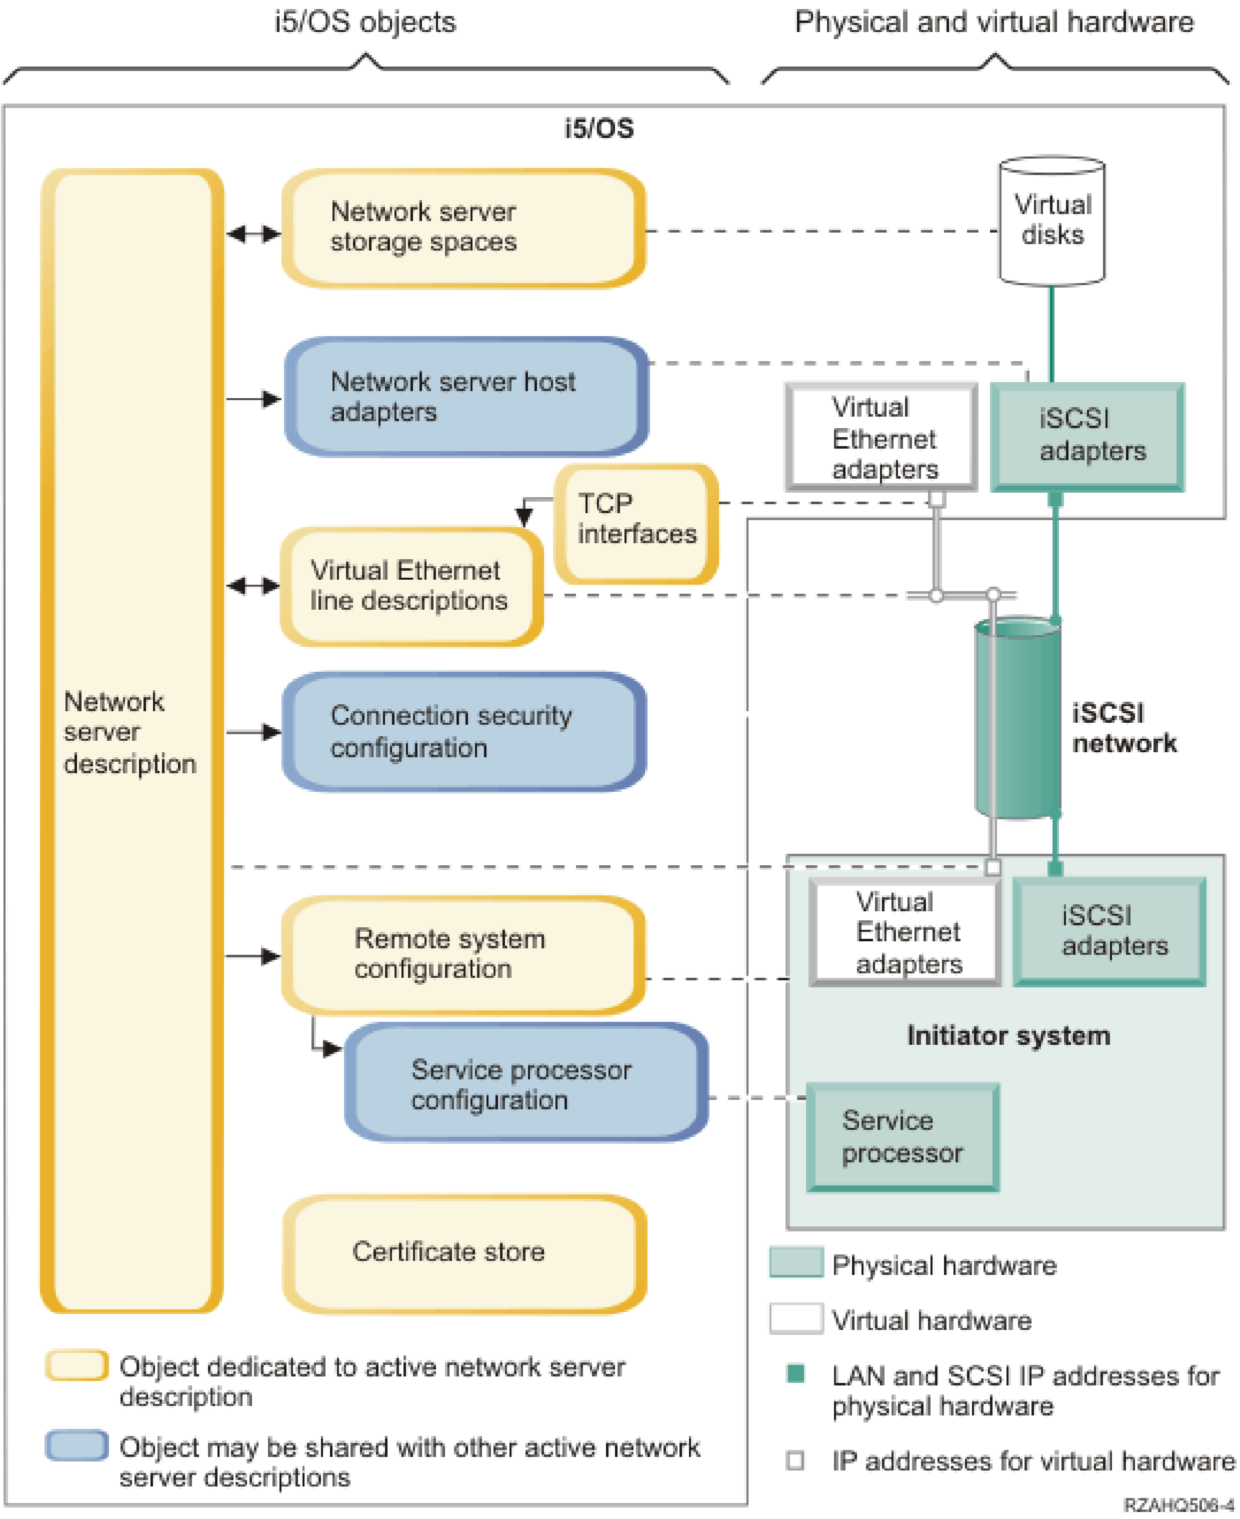
i5/OS objects
Physical and virtual hardware
i5/OS
Network server description
Network server storage spaces
Network server host adapters
TCP interfaces
Virtual Ethernet line descriptions
Connection security configuration
Remote system configuration
Service processor configuration
Certificate store
Virtual disks
Virtual Ethernet adapters
iSCSI adapters
iSCSI network
Virtual Ethernet adapters
iSCSI adapters
Initiator system
Service processor
Object dedicated to active network server description
Object may be shared with other active network server descriptions
Physical hardware
Virtual hardware
LAN and SCSI IP addresses for physical hardware
IP addresses for virtual hardware
RZAHQ506-4

Figure 10. iSCSI configuration objects in IBM i

Network server description

The network server description (NWSD) object is the main configuration object for an integrated server.

- It contains a reference to a remote system configuration.

- It contains references to the iSCSI and virtual Ethernet data paths for the integrated server.
  - You can define one or more storage paths. These storage paths reference the network server host adapter (NWSH) objects that are associated with the iSCSI target adapters that are used by the integrated server. You can choose which storage path is used for the SCSI data flows for each virtual disk drive. By associating your virtual disk drives with different storage paths, you can spread the overall server SCSI data flow workload across the storage path iSCSI target adapters for greater bandwidth. See "Multipath I/O for integrated servers" on page 52.
  - You can define one or more virtual Ethernet paths. These virtual Ethernet paths also reference the NWSH objects that are used by the integrated server. You can choose which NWSH is used for each virtual Ethernet port that the integrated server uses. By associating different virtual Ethernet ports with different NWSHs, you can spread the overall server virtual Ethernet data flow workload across the virtual Ethernet path iSCSI target adapters for greater bandwidth.
- The iSCSI-attached System x or BladeCenter hardware is controlled by IBM i.
  - An iSCSI-attached server is turned on and off by starting or stopping the NWSD for that server.
  - IBM i uses an Ethernet network to communicate with the service processor for the System x hardware or the BladeCenter management module for a BladeCenter server to perform the start and shut down tasks.

  Note: In case of a hardware failure, you can change the remote system configuration name that is specified in the NWSD and restart the server using spare hardware. See "Hot spare support for integrated servers" on page 55.

Network server host adapters

A network server host adapter (NWSH) device description object represents the iSCSI target adapter that is used by the IBM i side of the iSCSI connection.

- It identifies the iSCSI target adapter port.
  - For a hardware target (iSCSI HBA), it identifies the IBM i Network Server Host Port resource name (for example, CMNxx) for the iSCSI HBA port.
  - For a software target (Ethernet NIC), the NWSH uses a virtual port and also identifies the IBM i TCP/IP interface that is associated with the line description for the Ethernet NIC port.
- It defines how communications errors are logged and communications recovery information.
- It defines the IP addresses, ports, and so on. for the SCSI and LAN interfaces on the iSCSI target adapter.

IBM i can have multiple iSCSI target adapters. Each port on an iSCSI target adapter has an associated NWSH object.

- Each NWSH can be shared by multiple integrated servers. In configurations where bandwidth is not a concern, this results in a lower-cost solution.
- Each integrated server can use multiple NWSHs. Multiple NWSHs allow multiple SCSI and virtual Ethernet data paths between IBM i and the System x or blade system. Multiple NWSHs can provide greater bandwidth and connection redundancy.

Starting and stopping iSCSI target adapters.

- An iSCSI target adapter is started and stopped using the NWSH for that iSCSI target adapter.
- Alternatively, a software target (Ethernet NIC) can be started and stopped using the TCP/IP interface that is associated with the NWSH. The NWSH and the associated TCP/IP interface are started and stopped together.

  Notes:

  1. Do **not** use the same TCP/IP interface for multiple NWSHs. Only one NWSH that uses a particular TCP/IP interface can be active at a time.
  2. When starting a software target, the associated line description (LIND) is also started. However, when stopping a software target, the associated LIND remains active.

Remote system configuration

The remote system network server configuration (NWSCFG type RMTSYS) contains information that identifies the integrated server hardware to IBM i.

- It identifies the server hardware by serial number and type and model.
- It contains configuration information for the iSCSI initiator adapters that are used by the System x or blade hardware.
- It contains values required to boot the server.
- It contains a reference to the service processor NWSCFG object that is used to control the System x or blade hardware.
- It contains challenge handshake authentication protocol (CHAP) configuration values that are used to authenticate the remote system when it initially accesses storage.

The System x or blade server can have multiple iSCSI initiator adapters. Multiple iSCSI initiators allow multiple SCSI and virtual Ethernet data paths between IBM i and the System x or blade hardware. Multiple iSCSI initiators can provide greater bandwidth and connection redundancy.

The remote system configuration for an integrated server is referenced from the NWSD.

Service processor configuration

A service processor network server configuration (NWSCFG type SRVPRC) represents the System x service processor or the BladeCenter management module.

The service processor configuration contains the following information:

- It identifies the service processor or management module hardware by serial number and type and model.
- It defines how to find the service processor or management module on the Ethernet network using an IP address or host name.
- It contains a service processor user name and password that are used to sign on to the service processor.

Note: For a System x product, there is a one-to-one relationship between the service processor object and the remote system configuration. The service processor controls only one System x product. However, for BladeCenter systems, there can be a one-to-many relationship between the service processor object and the remote system configuration. Each management module can control any of the

BladeCenter systems that are contained within the BladeCenter chassis. Therefore, with iSCSI-attached BladeCenter systems it would be common for several remote system configurations to share (refer to) the same service processor object.

Connection security configuration

A connection security network server configuration (NWSCFG type CNNSEC) is used by the system. The integrated server installation process normally creates a default connection security configuration named `QCNNSEC` that is shared by all integrated servers on the IBM i system.

Certificate stores

Certificates are used to secure communications between IBM i and the initiator system for various functions. The certificates are kept in the following IBM i certificate store:

A certificate store that is associated with the network server description.
This certificate store is created and maintained automatically for you. It is used to store certificates that are generated and used internally by the IBM i Integrated Server Support. For example, certificates that are used when enrolling users to the hosted system. The certificates in this certificate store are used only when communicating with hosted systems that use the corresponding network server description.

Network server storage spaces (virtual storage)

A network server storage space (NWSSTG) represents a virtual disk drive (virtual storage) for an integrated server. Virtual storage can vary in size from 1 MB to 1000 GB each. Up to 64 virtual storage spaces can be linked to a server, depending on the server configuration. The storage capacity of an integrated server can range from several gigabytes to many terabytes. The virtual storage spaces are first created as stand-alone objects and then linked to the integrated server by identifying the NWSD of the integrated server that uses them.

Each server has up to two virtual disk drives that are automatically created by the server installation process. Each server can also have user-defined virtual disk drives.

- The system drive (typically the C: drive for Windows servers) contains the integrated server operating system (such as Windows Server or VMware ESX Server).
- For integrated Windows servers, the installation drive is used every time the server is started to pass configuration information from IBM i to the server. It also contains the IBM i Integrated Server Support (5770-SS1 option 29) code that runs on the Windows server. For Windows Server 2003 servers, the installation drive also contains a copy of the Windows server installation media.
- Additional user-defined drives are typically used for server applications and data.
- When linking the virtual disk drive to the NWSD, it is necessary to identify which of the NWSD storage paths to use for the SCSI data flows for that virtual disk drive. You can choose a specific storage path, the multipath group or let the default storage path be used.

The actual disk storage for the virtual disks is allocated from the IBM i integrated file system. The virtual disk drives can be allocated from the default system storage pool (also known as the system auxiliary storage pool, or system ASP), from a user-defined storage pool, or from an independent storage pool (independent ASP).

See "Storage management for integrated servers" on page 23 for more information about virtual storage.

Note:

1. Since virtual disks are objects in the IBM i integrated file system, an entire virtual disk drive image can be backed up and restored using the IBM i Save (SAV) and Restore (RST) commands. You can also do a file-level backup for the Windows operating system. For more information, see Chapter 5, "Backing up and recovering integrated servers," on page 117.
2. Even though storage spaces are allocated out of the integrated file system, storage operations are not performed by IFS while the integrated server is varied on. Therefore, operations like journaling are not enabled.

Related tasks:

"Viewing or changing integrated server configuration information" on page 73
Use either IBM Navigator for i or CL commands to view or change integrated server configuration information.

"Managing storage for integrated servers" on page 80
Use these tasks to manage storage for an integrated server.

"Managing network server host adapters" on page 98
Network server host adapter (NWSH) objects are used to configure the IBM i iSCSI target adapter. Use these tasks to manage NWSH objects.

"Managing remote system configurations" on page 104
Use these tasks to manage remote system configurations for iSCSI-attached integrated servers.

"Managing service processor configurations" on page 108
Use these tasks to manage service processor configurations for integrated servers.

"Managing connection security configurations" on page 113
Connection security network server configurations (NWSCFG subtype CNNSEC) are used by IBM i to connect to the integrated server hardware.

"Backing up the NWSD and other objects associated with integrated servers" on page 117
Do these tasks to back up the IBM i configuration objects and files related to integrated servers.

"Restoring the NWSD and other objects associated with integrated servers" on page 130
Do these tasks to restore the IBM i configuration objects and files related to integrated servers.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

Backing up your system

High availability concepts for integrated servers

Integrated servers can be made highly available through multipath storage connections, hot spare hardware, clustering, or by configuring the integrated server as a switchable device.

Related reference:

IBM i iSCSI Solution Guide

Multipath I/O for integrated servers

Multipath I/O enables multiple storage connections and provides automatic failover between connections to ensure that storage is accessible in case of a hardware failure.

A single IBM i iSCSI target port can support several servers or hosted systems.

iSCSI initiator port connection capabilities vary by operating system type. For details, see *Configuring multipath I/O for integrated servers* in the IBM i iSCSI Solution Guide .

You can configure the iSCSI environment to support multiple iSCSI targets, multiple iSCSI initiators, and multiple storage connections.

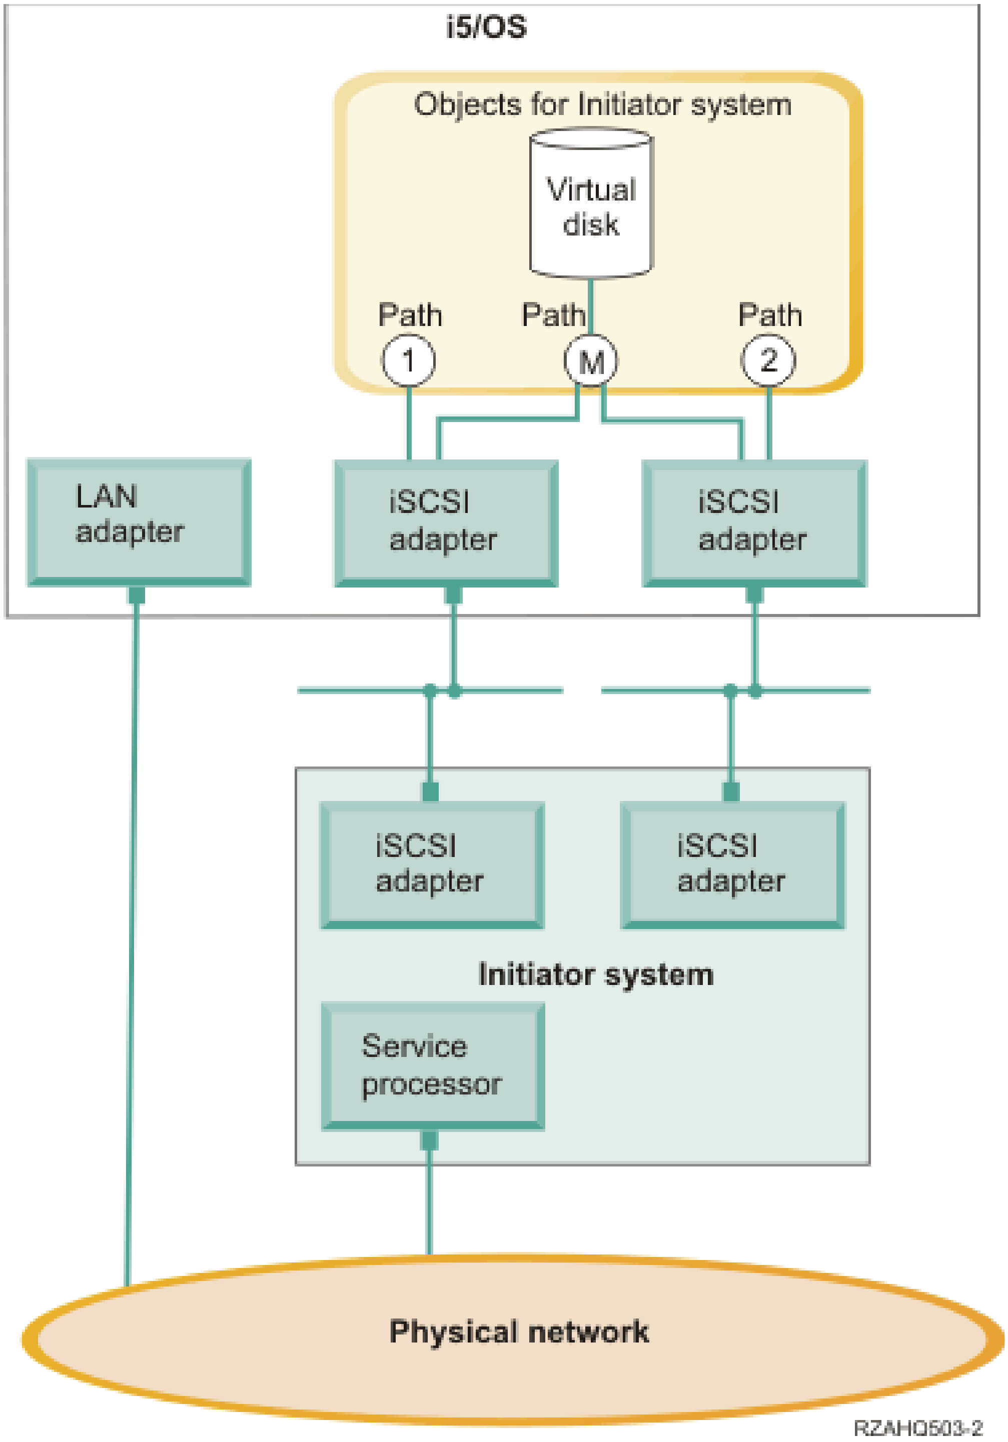

Figure 11. An environment with multiple iSCSI adapters installed in the target and initiator systems

Paths

Paths are connection points between virtual devices and iSCSI target adapters in IBM i. A virtual device being hosted by IBM i is said to be linked to a path. iSCSI initiator adapters access the virtual device through the path.

IBM i virtual storage or devices are linked to a network server host adapter (NWSH) object. For example, a configured virtual disk (such as Drive C:) hosted in IBM i is linked to the NWSH that represents the iSCSI target adapter.

There are several storage paths defined in Figure 11 on page 53. The paths labeled 1 and 2 each represent a single iSCSI target adapter. The path labeled M represents the multipath group, which is group of iSCSI target adapters.

You can configure storage for iSCSI-attached servers to use either a single path or a multipath group.

Removable media and virtual Ethernet connections use a single path. Connections for these devices cannot use the multipath group.

Multipath I/O and storage connection redundancy

A hosted system can use multiple iSCSI data paths to access virtual disks hosted by IBM i.

You can configure a multipath group of two or more iSCSI target adapters. Then specify that a virtual disk is accessed using the multipath group instead of a single iSCSI target adapter. With this configuration, the data on the virtual disk can be accessed using any of the iSCSI target adapters in the multipath group.

In Figure 11 on page 53, the multipath group is defined as path M. The virtual disks that are linked to the multipath group can be accessed by any of the iSCSI target adapters that are also linked to the multipath group. Only one multipath group can be defined per hosted system. This group can include up to four iSCSI target adapters.

For the most reliable storage network, do the following things:

- Configure multiple iSCSI targets in IBM i and define a multipath group that contains them.
- Configure multiple iSCSI initiators in the System x or blade product and configure them in the IBM i remote system configuration.
- Configure multiple switches to provide redundant network connections between the iSCSI targets and iSCSI initiators.
  - If you are using a BladeCenter system, configure multiple switch modules.
  - If you are using System x hardware, configure multiple switches in the iSCSI network.
- Link all storage to the multipath group.

Note: Removable media devices cannot use the multipath group.

The advantage of the multipath configuration is that, if there is a hardware failure, the hosted system can continue to access the disks that are configured to use the multipath group, using any of the iSCSI target adapters that are configured in the multipath group. This configuration can provide uninterrupted storage connections in case of a problem with an iSCSI target adapter, an iSCSI initiator adapter or a switch.

See *Configuring multipath I/O for integrated servers* in the IBM i iSCSI Solution Guide for more information about installing the required software components and linking storage to the multipath group.

Virtual Ethernet and initiator connection redundancy

Virtual Ethernet does not have the same multipath I/O concept that storage does. Virtual Ethernet supports iSCSI initiator redundancy, but not iSCSI target redundancy:

- If the integrated server has multiple iSCSI initiator adapters, the iSCSI initiator that is used for a particular virtual Ethernet adapter is automatically selected. If there are no failures, the virtual Ethernet adapter continues to use the selected iSCSI initiator. However, if the iSCSI initiator connection fails (for example, an initiator cable is pulled or the initiator card fails), a different iSCSI initiator adapter is automatically selected for the virtual Ethernet adapter and is used until another failure occurs.

  Note: In order for the automated selection process to work, the configured iSCSI target adapter must still be accessible by at least one iSCSI initiator adapter that is listed in the IBM i remote system configuration.
- There is no multipath group available for virtual Ethernet. A virtual Ethernet adapter is configured to use a specific iSCSI target and always uses that target. If the iSCSI target adapter fails or its cable is pulled, any virtual Ethernet adapters that are configured to use that iSCSI target adapter stop communicating. However, if the cable is plugged back in, communication automatically resumes.

For the most reliable virtual Ethernet network, do the following things:

- Configure multiple iSCSI initiators in the System x or blade product and configure them in the IBM i remote system configuration.
- Ensure that multiple iSCSI initiators can access the same IBM i iSCSI target.

Related reference:

IBM i iSCSI Solution Guide

Hot spare support for integrated servers

If your integrated server hardware fails, you can quickly configure your integrated server to use replacement hardware with your existing virtual storage.

Server integration and storage virtualization provide innovative options that can enhance the reliability and recoverability of the integrated server environment. Hot spare hardware provides a way to quickly recover from certain types of hardware failures. This can reduce the server downtime from hours or days to minutes.

- If the integrated server hardware fails, you can quickly and easily switch the server configuration to hot spare System x or BladeCenter hardware without restarting IBM i. Hot spare support also adds flexibility by enabling one spare server to be used to protect multiple production servers. This may reduce the overall number of servers needed to provide increased availability.
- If the IBM i iSCSI target adapters that the System x or blade system is using has a hardware failure, you can quickly switch the hosted system to use a spare iSCSI target adapter and restart the hosted system.

Related reference:

IBM i iSCSI Solution Guide

IBM i clustering for integrated servers

You can include the disks and configuration information for integrated servers in an IBM i cluster.

For more information, see the High availability topic collection.

User and group enrollment concepts for integrated Windows servers

Learn about how IBM i users and groups interact with integrated Windows servers.

One of the main advantages of using integrated Windows servers is the user administration function for IBM i and Windows user profiles. The user administration function allows administrators to enroll existing IBM i user and group profiles to Microsoft Windows.

Enrollment

Enrollment is the process by which an IBM i user or group profile is registered with the integration software.

The enrollment process happens automatically when triggered by an event such as using *IBM Navigator for i* or the Change NWS User Attributes (CHGNWSUSRA) command to enroll a user or group, an enrolled Windows user updating their IBM i user profile password or user attributes, or restarting the integrated server. If the integrated Windows server is active, the changes are made immediately. If the integrated server is varied off, the changes occur the next time the server is started.

Windows domains and local servers

Enrollment can be made to either a Windows domain or a local server. A Windows domain is a set of resources (applications, computers, printers) which are networked together. A user has one account across the domain and needs only to log onto the domain to gain access to all the resources. An integrated server can be a member server of a Windows domain and integrate IBM i user accounts into the Windows domain.

On the other hand, if you enroll IBM i users to an integrated server which is not part of a domain, it is called a **local server**, and user accounts will only be created on that integrated server.

Note: In Windows networking, groups of local servers can be loosely affiliated by using Windows workgroups. For example, if you open My Network Places and click Computers Near Me, you will see a list of the computers in the same workgroup as you. IBM i user accounts cannot be enrolled to Windows workgroups.

Special groups created on an integrated Windows server by IBM i

Two groups of users are created in Microsoft Windows as part of the installation to an integrated Windows server.

AS400_Users

Every IBM i user, when first enrolled to the Windows server, is placed in the AS400_Users group. You can remove a user from this group in the Windows server; however, the next time an update occurs from IBM i, the user will be replaced. This group is a useful place to check which IBM i user profiles are enrolled to the Windows server.

AS400_Permanent_Users

Users in this group cannot be removed from the Windows server by IBM i. It is provided as a way to prevent Windows users from being accidentally deleted by actions taken within IBM i. Even if the user profile is deleted from IBM i, the user will continue to exist in the Windows server. Membership in this group is controlled from the Windows server, unlike the AS400_Users group. If you delete a user from this group, it will not be replaced when an IBM i update is performed.

Using the IBM i user profile local password management (LCLPWDMGT) attribute

There are two ways to manage user profile passwords.

Traditional user (password managed by IBM i)

You may choose to have IBM i passwords and Windows passwords be the same. Enrolled Windows users manage their passwords in IBM i. This is configured by setting LCLPWDMGT(*YES) in the user profile.

Windows password-managed user

You may choose to manage enrolled Windows profile passwords in Windows without IBM i overwriting the password.This is configured by setting LCLPWDMGT(*NO) in the user profile.

For more information see "Enrolled user account options for integrated Windows servers" on page 59.

Using IBM i Enterprise Identity Mapping (EIM)

Defining EIM associations allows IBM i to support Windows single sign-on using an authentication method such as Kerberos. There are two ways to take advantage of the IBM i EIM support:

1. You can automatically create an EIM association using functions in the EIM Windows registry. Auto-creation and deletion of Windows EIM source associations are done when the IBM i Create, Change, or Delete user profile (CRTUSRPRF, CHGUSRPRF, or DLTUSRPRF) commands are used, specifying the EIMASSOC parameter values of *TARGET, *TGTSRC, or *ALL.
2. You may manually define EIM associations in the EIM Windows registry. When an EIM IBM i target association and Windows source association is defined for an IBM i user profile, the enrolled IBM i user profile may be defined as a different user profile name in Windows.

   Note: SBMNWSCMD, QNTC, and file level backup operations only work with EIM Kerberos associations. IBM i user profiles mapped to different Windows user names using an EIM Windows registry are not recognized. Those operations still attempt to use equivalent names.

For more information see "Enrolled user account options for integrated Windows servers" on page 59 and "Configuring Enterprise Identity Mapping for integrated Windows servers" on page 93.

Enrolling existing Windows user profiles

You can also enroll a user who already exists in the Windows server. The password for the user must be the same on IBM i as for the already existing Windows user or group.

User enrollment templates

You can customize the authorities and properties a user receives during enrollment through the use of user enrollment templates. See “User enrollment templates for integrated Windows servers” on page 60. If you do not use a template when you enroll users, they receive the following default settings:

- Users become members of the AS400_Users group and either the Users group in a local integrated Windows server or the Domain Users group on a Windows domain.
- IBM i keeps track of the user's IBM i password, password expiration date, description, and enabled or disabled status.

Enrolling IBM i groups

Up to this point, only the enrollment of individual IBM i user profiles to the Windows server has been discussed. You can also enroll entire IBM i groups. Then, when you add users to those IBM i groups that have been enrolled to the Windows server, you automatically create and enroll those users in the Windows server as well.

Enrolling to multiple domains

You may enroll users and groups to multiple domains, but typically this is unnecessary. In most Windows servers, multiple domains set up trust relationships with each other. In such cases, you only need to enroll the user in one domain because trust relationships automatically give the user access to other domains. See your Windows documentation for additional information about trust relationships.

Saving and Restoring enrollment information

Once you have defined your user and group enrollments, you need to save the enrollment definitions. You may save the enrollment information using options 21 or 23 on the GO SAVE menu, by using the Save Security Data (SAVSECDTA) command, or by using the Save Object List (QSRSAVO) API. Restoring the user profiles is done using the Restore User Profiles (RSTUSRPRF) command and specifying USRPRF(*ALL) or SECDTA(*PWDGRP) values.

Using the NWSD propagate domain user (PRPDMNUSR) parameter

If you have multiple servers which are members of the same domain, you may prevent duplicate domain enrollment from occurring on each member server. Use the Propagate Domain User (PRPDMNUSR) parameter in the Change Network Server Desc (CHGNWSD) command. See “Preventing enrollment to an integrated Windows server” on page 96 for more information.

Using the NWSD disable user profile (DSBUSRPRF) parameter

You can specify whether you want user profiles on integrated Windows servers to be disabled when the corresponding IBM i user profiles are disabled. Use the Disable User Profile parameter on the Change Network Server Description (CHGNWSD) command.

Related tasks:

“Enrolling IBM i users to integrated Windows servers” on page 88
To enroll IBM i users to integrated Windows servers, follow these steps.

“Enrolling IBM i groups to integrated Windows servers” on page 88
To enroll IBM i groups to integrated Windows servers, follow these steps.

"Changing the local password management user profile attribute" on page 93
Use these steps to change the local password management (LCLPWDMGT) user profile attribute.

"Configuring Enterprise Identity Mapping for integrated Windows servers" on page 93
Use this information to configure a user account to use EIM.

"Saving user enrollment information for integrated Windows servers" on page 129
Use CL commands and APIs to save user profiles and enrollment information for an integrated Windows server.

"Restoring user enrollment information for integrated Windows servers" on page 135
Use CL commands and APIs to save and restore user profiles and enrollment information for an integrated Windows server.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

Backing up your system

Enrolled user account options for integrated Windows servers

You can manage passwords for enrolled Windows users in either Windows or IBM i.

Traditional user (password managed on IBM i)

By default, enrolled users are set to this type. This user works in both Windows and IBM i. The IBM i password and Windows password are synchronized. Each time that the integrated Windows server is restarted, the user password is reset to the IBM i password. Password changes can only be made in IBM i. This user type is recommended for running File Level Backup and remote Windows commands. To set an IBM i user profile for a Windows user to this configuration, set the local password management (LCLPWDMGT) user profile attribute to **Manage this password locally through IBM i**. See "Changing the local password management user profile attribute" on page 93 for instructions.

Windows password-managed user

This person does all or most of their work in Windows and might never, or rarely, sign on to IBM i. If the user signs-on to IBM i, they must use an authentication method such as Kerberos to access IBM i.

When the user profile attribute LCLPWDMGT(*NO) is defined for an IBM i user, the IBM i user profile password is set to *NONE. The IBM i enrollment password is saved until Windows enrollment is successfully completed. After the IBM i user is enrolled to Windows, the Windows user can change and manage their password in Windows without IBM i overwriting their password. Using this method allows for a more secure environment because there are fewer passwords being managed. To set an IBM i user profile for a Windows user to this configuration, set the local password management (LCLPWDMGT) user profile attribute to **Manage this password remotely through some other platform**. See "Changing the local password management user profile attribute" on page 93 for instructions.

Windows user with Enterprise Identity Mapping (EIM) associations automatically configured

Specifying the user profile attribute of EIMASSOC to be *TGT, TGTSRC, or *ALL allows the integrated server to automatically define EIM Windows source associations. Using the automatic definitions of associations makes configuring EIM easier. To read how to create a user of this type, see "Configuring Enterprise Identity Mapping for integrated Windows servers" on page 93.

Windows user with Enterprise Identity Mapping (EIM) associations manually configured

The user can choose to manually define EIM Windows source associations. This method can be used to set the IBM i user profile to be enrolled to a different Windows user profile name. The user must manually define an IBM i target association for the IBM i user profile and also a Windows source association for the same EIM identifier. See the Enterprise Identity Mapping topic collection for details.

Table 3. Types of user configurations

User type	Function provided	User profile definition
Traditional user	• Both IBM i and Windows fully functional. • Easy to configure. • Password is changed from IBM i. • IBM i and Windows user ID and passwords are identical. • Recommended for system administrators, users who frequently use IBM i, or for systems which use IBM i for backup and restoration of user profiles.	LCLPWDMGT(*YES) and no EIM Windows source associations defined.
Windows password-managed user	• Password can be changed from Windows. • Simple configuration. • Windows password administration makes this configuration more secure because the IBM i password is *NONE. • IBM i sign-on requires an authentication method such as System i® Navigator provides with its support of IBM i sign-on using Kerberos.	LCLPWDMGT(*NO)
Windows user with Enterprise Identity Mapping (EIM) associations auto configured	Automatic creation of Windows source associations makes it easier to set up and configure to use Kerberos enabled applications.	For example: EIMASSOC(*CHG *TARGET *ADD *CRTEIMID)
Windows user with Enterprise Identity Mapping (EIM) associations manually configured	Allows the user to define EIM associations for enrolled IBM i user profiles to be different user profiles in Windows.	Use *IBM Navigator for i* to manually define EIM IBM i target associations and Windows source associations.

User enrollment templates for integrated Windows servers

You can use templates to simplify the enrollment of new users to an integrated Windows server.

Rather than manually configure many new users, each with identical settings, use a user enrollment template to automatically configure them. Each template is a Windows user profile that defines user privileges, such as group membership, directory paths, and organizational unit containers.

When you enroll users and groups from IBM i to the Windows environment, you can specify a user template on which to base the new Windows users. For example, you could create a user template and name it USRTEMP. USRTEMP could be a member of the Windows server groups NTG1 and NTG2. On IBM i, you could have a group called MGMT. You could decide to enroll the MGMT group and its members to a Windows server. During the enrollment process, you could specify USRTEMP as the user template. During enrollment, you automatically add all members of the MGMT group to the NTG1 and NTG2 groups.

User templates save you from having to set up group memberships individually for each user. They also keep the attributes of enrolled users consistent.

You can make a user template a member of any Windows group, whether you enrolled that group from IBM i or not. You can enroll users with a template that is a member of a group that was not enrolled from IBM i. If you do this, the users become members of that non-enrolled group as well. IBM i does not know about groups that were not enrolled from IBM i. This means that you can only remove users from the group by using the User Manager program on Windows.

If you use a template to define a new user enrollment, and the template has a folder or directory **Path** or **Connect To** defined, the newly-created Windows user will have the same definitions. The folder definitions allow the user administrator to take advantage of folder redirection and to manage terminal service sign-on.

If you use a template when you define a new user enrollment, and the template is a user object in a Windows Active Directory organizational unit container, the newly created Windows user object will be in the same organizational unit container. An organizational unit provides a method to grant users administrative control to resources.

You can change existing user templates. Such changes affect only users that you enroll after you change the template.

You use templates only when you create a newly enrolled user in the Windows environment. If you perform enrollment in order to synchronize an existing Windows user with an IBM i counterpart, Windows ignores the template.

Related reference:

IBM i iSCSI Solution Guide

Password considerations for integrated Windows servers

You can change IBM i system values and Windows Server policies to configure the rules for passwords and ensure that they work correctly for your environment.

1. Enrolled users must use IBM i passwords containing only characters and password lengths allowed in Windows passwords.
2. IBM i and an integrated Windows server must enforce consistent password rules. If the password rules on the two systems are not consistent, then a password for an enrolled IBM i user might be rejected by the integrated

Windows server. You can adjust the password rules either on IBM i or on the integrated Windows server to make them consistent:

- IBM i password rules can be adjusted using the IBM i system values listed in the next section.
- Refer to your Windows Server documentation for the methods to change Windows Server policies that control the rules for passwords.

3. When the IBM i passwords of enrolled users expire, their Windows passwords also expire. Users can change their passwords on Windows, but they must remember to also change their passwords on IBM i. Changing the IBM i password first automatically changes the Windows password.

IBM i system values affecting passwords

IBM i uses system values to control password rules and other security-related items.

1. Make sure that the IBM i QRETSVRSEC system value is set to 1. You can set QRETSVRSEC using the Work with System Value (WRKSYSVAL) command. If you do not set QRETSVRSEC to 1, you cannot enroll users on your integrated Windows server until they sign on to IBM i.

   Note: This system value is also required for other integrated server support functions, such as powering on an integrated server.
2. The IBM i password level can be set to allow user profile passwords of 1 - 10 characters or to allow user profile passwords of 1 - 128 characters. An IBM i password level change of the system value QPWDLVL requires an IPL.
3. If system value QPWDLVL is set to allow user profile passwords of 1 - 128 characters, then system value QPWDMAXLEN also needs to be changed to allow passwords to be 128 characters in length.
4. The IBM i password level of 0 or 1 supports passwords of 1 - 10 characters and limits the set of characters. At password level 0 or 1, IBM i converts passwords to all lowercase for Windows.
5. The IBM i password level of 2 or 3 supports passwords of 1 - 128 characters and allows more characters including uppercase and lowercase characters. At level 2 or 3, IBM i preserves password case sensitivity for Windows.
6. If the IBM i system value QSECURITY is 10, the Windows users that are created do not require passwords to sign on. All other IBM i QSECURITY levels require that a user profile has a password to sign on. You can find more information about security levels in the Security reference topic collection.
7. If you are using a language other than English, be aware that using anything but invariant characters in user profiles and passwords can cause unpredictable results. The IBM i globalization topic contains information about what characters are in the invariant character set. This statement is only true when QPWDLVL is 0 or 1. When QPWDLVL is 2 or 3, invariant characters can be used without causing any problems.

QAS400NT user and integrated Windows servers

IBM i uses the QAS400NT user to sign on to the integrated Windows server operating system.

The QAS400NT user is used to enroll IBM i users and groups to Windows domains and servers.

Related reference:

IBM i iSCSI Solution Guide

IBM i NetServer for integrated Windows servers

You must configure NetServer to enable updates to the IBM i Integrated Server Support software that runs on the Windows server and to enable communication for integrated VMware ESX server administration tasks. You can also configure print and file sharing.

NetServer enables Windows clients to connect to IBM i shared directory paths and shared output queues by way of TCP/IP.

Notes:

- To install Integrated Server Support service packs on an integrated Windows server, you must be signed on with a Windows account that corresponds to an IBM i user profile with the same password, or you must have a guest NetServer user profile configured.
- When you install a VMware ESX server, a NetServer file share is automatically created and used exclusively for integrated VMware ESX server administration. Access to this share requires that user QVMWINT exists on both IBM i and the integrated Windows server that manages the integrated VMware ESX server.

To set up NetServer for Integrated Server Support tasks, use the method found in the Getting started with IBM i NetServer topic.

Once you have set up NetServer, you need to set up a Windows user with access to NetServer, or you can set up a NetServer guest user profile.

IBM i Access for Windows and integrated servers

IBM i Access for Windows enables you to connect to IBM i.

It features a complete set of integrated functions that enable desktop users to use IBM i resources as easily as their local PC functions. With IBM i Access for Windows, users and application programmers can quickly process information, applications, and resources for their entire company.

IBM i Access for Windows also provides an Open Database Connectivity (ODBC) driver that can be used for server-to-server applications between integrated servers and IBM i. You can enable Open Database Connectivity (ODBC) to run as a Windows service by installing IBM i Access for Windows on your integrated Windows server. This enables you to write server applications that call the ODBC device driver to access Db2 for IBM i.

To enable ODBC to be started from a Windows service, run the `CWBCFG` command with the `/s` option after you install IBM i Access for Windows.

As a single user signed on to Windows, you have full support for all other IBM i Access for Windows features.

Additional information sources:

- See IBM i NetServer vs IBM i Access for Windows in the IBM i NetServer topic collection.

Chapter 3. Integrated server installation road map (Deprecated)

The road map contained in the IBM i iSCSI Solution Guide provides an outline of the tasks to install an iSCSI-attached integrated server.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide .

Note: All articles that have been removed are marked with **(Deprecated)** in the article titles.

Prerequisites for installing an integrated server (Deprecated)

You need to have software, hardware, and documentation before your begin installing an integrated server.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide .

Preparing for the hardware installation (Deprecated)

Use these tasks to prepare for the integrated server hardware installation.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide .

iSCSI Network Planning Guide (Deprecated)

Use this guide to plan for the iSCSI network and the IBM i configuration objects that are needed to complete an iSCSI-attached integrated server installation.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide .

Installing the hardware (Deprecated)

Install the hardware that is used by the integrated server.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide .

Preparing IBM i for the integrated server installation (Deprecated)

Create the IBM i objects that will be used by your integrated server.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide .

Installing the integrated server (Deprecated)

Install the integrated server from IBM i to create the necessary IBM i configuration objects and virtual storage (if needed) for the server.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide .

Chapter 4. Managing integrated server environments

Use these tasks to manage integrated server environments.

Creating and deleting integrated servers

There are several methods to create and delete iSCSI-attached servers that are integrated with IBM i.

Installing integrated servers

Installing integrated servers is accomplished by following the server installation road map in the IBM i iSCSI Solution Guide.

Follow the server installation road map in the IBM i iSCSI Solution Guide to install the server. The installation process involves setting up and configuring the server hardware, configuring IBM i, starting the server operating system installation from IBM i using the **Create Server** Web GUI task or an IBM i CL command, and finally doing some configuration tasks in the server operating system. The *IBM i iSCSI Solution Guide* contains detailed instructions for each step of the process.

Cloning integrated servers

Cloning integrated servers is accomplished by following the server cloning road map in the IBM i iSCSI Solution Guide.

Follow the server cloning road map in the IBM i iSCSI Solution Guide to clone the server. The cloning process involves preparing the server for cloning, duplicating the IBM i objects for the server using the **Create Server** Web GUI task with a base network server description (NWSD), and finally doing some more configuration tasks after the server is cloned. The *IBM i iSCSI Solution Guide* contains detailed instructions for each step of the process.

Uninstalling integrated servers

To uninstall (delete) an integrated server, follow these steps.

Before uninstalling an integrated server, shut down the integrated server from IBM i. See “Stopping integrated servers” on page 72.

To uninstall an integrated server, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Servers**.
3. Click the menu icon for the server you want to delete and select **Delete**.
4. Click **Delete** on the confirmation page.

The nonshared IBM i objects for the server are deleted. Typically, these objects are the ones that the **Create Server** GUI task or the install server command created when the server was originally installed. The process also deletes objects related to virtual Ethernet LANs that are associated with the server. The objects that are deleted include:

- Network server description (NWSD)
- Predefined virtual storage linked to the server
- Virtual Ethernet LAN line descriptions
- TCP/IP interfaces bound to virtual Ethernet LAN line descriptions
- TCP/IP device descriptions for virtual Ethernets
- TCP/IP controller descriptions for virtual Ethernets

Tip: If you want to use a CL command, see:
Delete Integrated Server (DLTINTSVR)

Note: The **Delete Server** GUI task (and the DLTINTSVR CL command) cannot be used if the NWSD object for the server no longer exists. If the NWSD object no longer exists, or if you prefer to delete the objects manually, see the procedures referenced in the following table.

Table 4. Nonshared objects to delete for an integrated server (manual method)

Objects to Delete	**How to Delete**
Network server description (NWSD). **Note:** Before deleting the NWSD, unlink any virtual storage that is linked to it. See "Unlinking virtual storage" on page 86.	See Delete Network Server Desc (DLTNWSD).
System drive virtual storage. **Note:** This storage is typically named `nwsdname1`, where `nwsdname` is the name of the NWSD. There is no system drive to delete for VMware ESXi embedded servers.	See "Deleting virtual storage" on page 87.
Installation drive virtual storage. **Note:** This storage is typically named `nwsdname2`, where `nwsdname` is the name of the NWSD. There is no installation drive to delete for VMware ESX servers.	See "Deleting virtual storage" on page 87.
Virtual Ethernet LAN line descriptions. **Note:** The line description names typically start with the NWSD name, followed by 00, 01, 02, PP, V0, V1, V2, V3, V4, V5, V6, V7, V8, or V9. For example, if the NWSD name is `MYSERVER`, then the point-to-point line description name would be `MYSERVERPP`.	See "Deleting a line description" on page 104.
TCP/IP interfaces bound to virtual Ethernet LAN line descriptions. **Note:** You can identify the TCP/IP interfaces that are associated with the NWSD by looking at the name of the attached line description. See the line description naming convention described previously.	See "Deleting a TCP/IP interface" on page 103.
TCP/IP device descriptions for virtual Ethernets. **Note:** The name of the TCP/IP device description starts with the first five characters of the NWSD name, followed by 'TCP' and an optional two-digit number. For example, if the NWSD name is `MYSERVER`, then the device name might be `MYSERTCP` or `MYSERTCP01`.	See Work with Device Descriptions (WRKDEVD). Type `WRKDEVD *CMN`. Then use option 4=Delete for each device associated with the server.
TCP/IP controller descriptions for virtual Ethernets. **Note:** The name of the controller description starts with the first five characters of the NWSD name, followed by 'NET' and an optional two-digit number.	See Work with Ctl Descriptions (WRKCTLD). Type `WRKCTLD *CMN`. Then use option 4=Delete for each controller associated with the server.

If you no longer need user-created virtual storage or other IBM i configuration objects that were used by the integrated server, you can delete them.

If you remove all of your integrated servers from IBM i and do not plan to install any more, you can delete IBM i Integrated Server Support. Deleting the product frees up the storage that the product uses.

Deleting shareable IBM i objects for a deleted server

When you uninstall an integrated server, the nonshared objects for the server are deleted. You might also want to delete the shareable objects, such as user-created virtual storage or other IBM i configuration objects that were used by the server.

Delete the shareable IBM i objects listed in the following table if they are no longer needed.

Important: The objects listed in the following table can be shared among multiple integrated servers. Make sure that the objects are not used by other integrated servers before deleting them. See the notes for additional considerations.

Table 5. Shareable objects to delete if no longer needed

Objects to Delete (if no longer needed)	How to Delete
User-created virtual storage. **Note:** Do not delete virtual storage if it is linked to other servers. For example, when deleting one of several VMware ESX servers that share the same virtual storage.	See "Deleting virtual storage" on page 87.
Remote system configuration. **Note:** Do not delete the remote system configuration if it is used by other servers. If you plan to install another integrated server on the same hardware, then you can typically reuse the existing remote system configuration for the new server.	See "Deleting a remote system configuration" on page 108.
Service processor configuration. **Note:** Do not delete the service processor configuration if it is used by other remote system configurations (for example, multiple blades in a BladeCenter). If you plan to install another integrated server on the same hardware, then you can typically reuse the existing service processor configuration for the new server.	See "Deleting a service processor configuration" on page 112.
Connection security configuration. **Note:** Do not delete the connection security configuration if it is used by other servers. For example, there is typically one connection security configuration named QCNNSEC that is shared by all integrated servers on the system.	See "Deleting a connection security configuration" on page 115.
Network server host adapters (NWSHs). **Note:** Do not delete NWSHs if they are used by other servers or if you plan to install another server that uses the same iSCSI target adapter.	See "Deleting a network server host adapter" on page 103.

Uninstalling IBM i Integrated Server Support

If you remove all integrated servers from your IBM i system and do not plan to install others, you can remove IBM i Integrated Server Support, Option 29. Removing the program frees the storage space it occupied on IBM i.

Managing integrated servers

Use these tasks to manage integrated servers.

Starting and stopping integrated servers

Use these tasks to stop and start integrated servers.

Starting integrated servers

You can start integrated servers from IBM i.

Important: Before starting an integrated VMware ESX server, ensure that the management server (integrated Windows server) is started. The management server is needed for administrative communication to flow from IBM i to the integrated VMware ESX server.

To start an integrated server, complete the following steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Servers**.
3. Click the menu icon for the server you want to start.
4. Select **Start**. A progress page is shown. After awhile, the status changes to **Started**.

   Tip: To process multiple objects, select the objects in the list. Then use **Select Action** > **Start**.

Note: For VMware ESX servers, if the management connection has not been set up yet, the GUI progress panel does **not** go to completion and the ESX server status does **not** change to **Started** (**ACTIVE** on WRKNWSSTS). To set up the ESX management connection, see the installation road map in the IBM i iSCSI Solution Guide .

Tip: If you want to use a CL command, see:
Work with Configuration Status (WRKCFGSTS) (use `WRKCFGSTS *NWS`)
Vary Configuration (VRYCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Starting an integrated server when IBM i TCP/IP starts

Configure integrated servers to start when IBM i TCP/IP starts.

The integrated server must have a point-to-point virtual Ethernet port and an associated TCP/IP interface before performing this task. These items are automatically configured when an integrated Windows server is installed. However, these items are **not** automatically configured for an integrated VMware ESX server. If you are performing this task for an integrated VMware ESX server that does not have a point-to-point virtual Ethernet port, see Configuring a point-to-point virtual Ethernet port for an integrated VMware ESX server.

Note: In order for the integrated server to automatically start, the iSCSI target adapters that the integrated server uses must also be configured to automatically start.

- To automatically start a hardware target (iSCSI HBA), configure the **`Online at IPL`** attribute in the network server host adapter (NWSH) object. See "Changing network server host adapter properties" on page 101.
- To automatically start a software target (Ethernet NIC), configure the **`Start interface when TCP/IP is started`** attribute in the IBM i TCP/IP interface that

the NWSH uses. Starting the TCP/IP interface also causes the associated line description and NWSH to start. See the procedure shown in the following task steps.

Attention: If multiple integrated servers use the same System x or BladeCenter blade server hardware, configure only one of them to autostart. Only one integrated server can use the server hardware at a time. Configuration of multiple TCP/IP interfaces to autostart for integrated servers that share the same server hardware can cause unpredictable results.

To have an integrated server automatically vary on when you start TCP/IP, follow these steps:

1. Select **Network** from *IBM Navigator for i*.
2. Click **Show All Network Tasks**.
3. Select **Network** > **TCP/IP Configuration** > **IPv4** > **Interfaces**.
4. Click the menu icon for the interface for the point-to-point virtual Ethernet LAN line description for the server and select **Properties**.

   Note: The point-to-point virtual Ethernet LAN line description has a name that consists of the network server description (NWSD) name followed by 'PP'. For example, if the NWSD name is MYSVR, then the point-to-point virtual Ethernet LAN line description is MYSVRPP.
5. On the **Advanced** tab, select the **`Start interface when TCP/IP is started`** check box and click **OK** to save the change.

The integrated server automatically varies on when you start TCP/IP. TCP/IP can be automatically started by the system at IPL by changing the system IPL attributes. Any TCP interfaces that have been enabled for autostart are started along with TCP/IP at IPL.

Tip: If you want to use a CL command, see:
Configure TCP/IP (CFGTCP) (use `CFGTCP`, then option 1)

Configuring a point-to-point virtual Ethernet port for an integrated VMware ESX server:

Configure a point-to-point virtual Ethernet port for an integrated VMware ESX server so that it can be automatically started when IBM i TCP/IP starts.

Restriction: The point-to-point virtual Ethernet port on an integrated VMware ESX server can only be used to automatically start the integrated VMware ESX server when IBM i TCP/IP starts. The point-to-point virtual Ethernet port does **not** provide a virtual Ethernet communication connection between the integrated VMware ESX server and any other systems.

Do these steps to configure a point-to-point virtual Ethernet port for an integrated VMware ESX server:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Servers**.
3. Click the menu icon for the integrated VMware ESX server and select **Properties**.
4. On the server properties panel, click the **Virtual Ethernet** tab.
5. Click the **Add...** button to add a new virtual Ethernet port.

6. On the virtual Ethernet properties panel, specify the values for the point-to-point virtual Ethernet port:
   a. Type the **Internet address** for the integrated server side of the point-to-point virtual Ethernet.

      Note: This IP address is not used by the integrated VMware ESX server.
   b. Type the **IBM i internet address** for the IBM i TCP/IP interface.

      Note: This IP address is not used by IBM i for communications. Its only purpose is to provide a mechanism to automatically start the integrated VMware ESX server when IBM i TCP/IP starts.
   c. Type the **Subnet mask** for the point-to-point virtual Ethernet network.
   d. Leave the default values for the remaining items.
   e. Click **OK** to add the new port to the **Virtual Ethernet** tab on the server properties panel.
7. On the server properties panel, click **OK** to save the changes. The NWSD is updated and a line description and IBM i TCP/IP interface for the new point-to-point virtual Ethernet port are created.

Tip: If you want to use a CL command, see:
Change Network Server Desc (CHGNWSD) (see the VRTETHPTH and TCPPORTCFG keywords)
Create Line Desc (Ethernet) (CRTLINETH)
Add TCP/IP Interface (ADDTCPIFC)

Stopping integrated servers

You can shut down integrated servers from IBM i.

CAUTION:
Take special care when shutting down a VMware ESX server from IBM i:

- **You must shut down the ESX server before shutting down the integrated Windows server that serves as the management server for the ESX server, and also before shutting down the ESX platform manager (vCenter) server, if configured. If the management server or ESX platform manger (if configured) is not available, then IBM i powers down the ESX server without notifying ESX, which could cause data corruption on the ESX server.**
- **The IBM i system does not attempt to shut down any virtual machines that the ESX server is hosting. You must manually shut down the virtual machines before shutting down the ESX server to ensure a clean shutdown of the virtual machines.**

To shut down an integrated server, follow these steps.

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Servers**.
3. Click the menu icon for the server you want to stop and select **Shut Down**.
4. Click **Shut Down** on the confirmation page.

The status changes to **Shutting down...**, **Partially shut down**, and eventually **Shut down**.

Tip: If you want to use a CL command, see:
Work with Configuration Status (WRKCFGSTS) (use `WRKCFGSTS *NWS`)
Vary Configuration (VRYCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Shutting down your IBM i system when integrated servers are present

Learn how to safely shut down your system when integrated servers are installed.

The easiest way to ensure your integrated servers will be shut down safely is to always manually shut them down before shutting down IBM i system. The CL command `PWRDWNSYS *CNTRLD` will attempt to power-down each of the integrated servers, giving each of them a period of time (the NWSD attribute SHUTDTIMO, by default 15 minutes) in which to shut down. Note that there is no guarantee that they will finish shutting down within this time period.

CAUTION:
The CL command `PWRDWNSYS *IMMED` is not recommended. This will power down the IBM i system immediately, without attempting to shut down any integrated servers.

Table 6. Methods for shutting down the IBM i system

Action	Result
Shut down the integrated server manually.	The integrated server is varied off properly, with no risk of data loss.
Issue the CL command `pwrdwnsys *cntrld`.	The integrated server is given the length of time specified in the shut down timeout attribute of its NWSD in which to shut down, then the IBM i system continues to power down.
Issue the CL command `pwrdwnsys *immed`.	The IBM i system powers down immediately and does not shut down any integrated servers. This could cause data corruption on the integrated servers.

If your IBM i system uses the Power On/Off Schedule, the Power-Off exit program (QEZPWROFFP) should be changed to vary off all NWSDs before calling the Power Down System (PWRDWNSYS) command. Careful consideration must be given to scheduling as the number and activity of each server will determine the amount of time necessary to completely vary off each server. Use the Submit multiple jobs (SBMMLTJOB) and Job description (JOBD) parameters of the Vary Configuration (VRYCFG) command to vary multiple servers at the same time in batch. The scheduled power on must not occur before the system has a chance to vary off all servers and issue the PWRDWNSYS. See the Schedule a system shutdown and restart topic.

Viewing or changing integrated server configuration information

Use either IBM Navigator for i or CL commands to view or change integrated server configuration information.

Note: Some IBM i configuration object properties cannot be changed while the object is active, or in use by an active server. See the appropriate configuration object command documentation for restrictions.

- Change Network Server Desc (CHGNWSD)

- Change NWS Storage Space (CHGNWSSTG)
- Change Device Desc (NWSH) (CHGDEVNWSH)
- Change NWS Configuration (CHGNWSCFG)

If you want to change a property that cannot be changed while the configuration object or associated server is active, perform one of the following tasks before changing the configuration object properties.

- “Stopping integrated servers” on page 72
- “Stopping a network server host adapter” on page 102

View and change integrated server configuration information as follows:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select one of the following links to show the corresponding list of objects.
   - Servers
   - All Virtual Storage
   - Network Server Host Adapters
   - Remote Systems
   - Service Processors
3. Click the menu icon for an object in the list and select **Properties**.
4. If you change any of the object properties, click **OK** to save the changes and close the object properties page.
   Otherwise, click **Cancel** to close the object properties page.

Using the character-based interface you can view and change all integrated server configuration information. The following table summarizes the relevant CL commands.

Table 7. CL commands for managing integrated server configuration information

Tasks	**CL Command**
Vary on and off integrated servers, check the status of the integrated server and objects that are associated with the network server description (NWSD).	`WRKCFGSTS CFGTYPE(*nws)`
Manage your integrated servers.	`WRKNWSD` `WRKNWSSTS`
Manage line descriptions that are created when you install the integrated server.	`WRKLIND`
Manage TCP/IP interfaces that are created during server installation.	`NETSTAT OPTION(*IFC)` `CFGTCP`, option 1
Monitor virtual storage (network server storage spaces).	`WRKNWSSTG`
Vary on and off iSCSI targets and check the status of the network server host adapter (NWSH).	`WRKCFGSTS CFGTYPE(*DEV) CFGD(*NWSH)`
Manage network server host adapters	`WRKDEVD DEVD(*NWSH)`
Manage remote system and service processor network server configurations	`WRKNWSCFG`

Related concepts:

“IBM i configuration objects for integrated servers” on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Configuring security between IBM i and integrated servers

Use these tasks to manage security for integrated servers.

Configuring CHAP for integrated servers

Use these tasks to configure the challenge handshake authentication protocol (CHAP) for the remote system configuration for an integrated server.

You must have security administrator (*SECADM) special authority to create, change, or display CHAP information.

Configuring target CHAP for integrated servers:

Do these steps to configure the initiator to authenticate the target.

1. Vary off the network server description (NWSD) for your integrated server.
2. Select **Integrated Server Administration** from *IBM Navigator for i*.
3. Select **Remote systems**.
4. Click the menu icon for the remote system configuration for the integrated server and select **Properties**.
5. On the **CHAP Authentication** tab, click **Enable Challenge Handshake Authentication Protocol (CHAP)** to enable CHAP.
6. Specify information for **Target CHAP Values**.
   a. Select an option for **CHAP name**.
   b. Select **Generate CHAP secret once** or select **Specific CHAP secret** and specify a CHAP secret.
7. Configure target CHAP on the iSCSI-attached server. See the IBM i iSCSI Solution Guide .

Configuring initiator CHAP for integrated servers:

If you have configured a target CHAP, you can also use these steps to configure initiator CHAP for your iSCSI-attached integrated server.

1. Vary off the NWSD for your integrated server.
2. Select **Integrated Server Administration** from *IBM Navigator for i*.
3. Select **Remote systems**.
4. Click the menu icon for the remote system configuration for the integrated server and select **Properties**.
5. On the **CHAP Authentication** tab, click **Enable bidirectional CHAP** to enable initiator CHAP.
6. Specify information for **Initiator CHAP Values**.
   a. Select an option for **CHAP name**.
   b. Select **Generate CHAP secret once** or select **Specific CHAP secret** and specify a CHAP secret.
7. Configure initiator CHAP on the iSCSI-attached server. See the IBM i iSCSI Solution Guide .

Changing a service processor password for an integrated server

Change the service processor password for an integrated server.

Note: The service processor password cannot be changed while the service processor configuration is in use by an active server. If an associated server is active, stop the server before performing this task. See "Stopping integrated servers" on page 72.

To change the service processor password, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Service Processors**.
3. Click the menu icon for a service processor configuration from the list available.
4. Select **Initialize**.
5. Select the **Change service processor user ID and password** option.
6. Specify the new **User**, **Password**, and **Confirm new password values**.
7. Click **Initialize** to perform the operation.

Configuring a firewall to allow integrated server connections

Use this information to configure a firewall to allow integrated server connections.

If there is a firewall between IBM i and the iSCSI network for the integrated server, then the firewall must be configured to allow incoming iSCSI and virtual Ethernet traffic to pass.

The values that affect firewall configuration are listed below:

For storage paths and virtual Ethernet connections protected by the firewall:

Remote IP address

Use the procedure described in "Displaying remote system configuration properties" on page 106 to display the properties of the remote system configuration for the server. Go to the **Remote Interfaces** tab and note the `SCSI Internet Address` and `LAN Internet Address` values.

- **Local IP address and TCP port:** Use the procedure described in "Displaying network server host adapter properties" on page 101 to display the properties of the network server host adapter (NWSH). Go to the **Local Interfaces** tab to see information that is used by the NWSH. Record the following values:
  - Local SCSI interface: Internet address
  - Local SCSI interface: TCP port
  - Local LAN interface: Internet address
  - Local LAN interface: Base virtual Ethernet port
  - Local LAN interface: Upper virtual Ethernet port

  Note: Virtual Ethernet traffic is encapsulated in UDP packets. Each virtual Ethernet adapter is automatically assigned a UDP port from a range that begins at the specified base virtual Ethernet port number and ends at the base virtual Ethernet port number plus the number of configured virtual Ethernet adapters. Each virtual Ethernet adapter is also has a UDP port assigned at the Windows server. UDP ports for virtual Ethernet are normally automatically allocated by Windows. If you want to override automatic allocation, you can manually allocate a UDP port by performing the following steps at the Windows console.

  1. Navigate to the **Network Connections** Window.
  2. Double-click the **IBM Virtual Ethernet x** adapter that you want to configure.

    3. Click **Properties**.
    4. Click **Configure**.
    5. Click **Advanced**.
    6. Click **Initiator LAN UDP Port**.
    7. Enter the UDP port that you want the virtual Ethernet adapter to use.
- **TCP ports associated with all Local IP addresses:**
  1. Select **Integrated Server Administration** from *IBM Navigator for i*.
  2. Select **Servers**.
  3. Click the menu icon for the server from the list available and select **Properties**.
  4. Go to the **System** tab and click the **Advanced** button.
  5. Note the following values:
     - `Virtual Ethernet control port`

Configuring multipath I/O for integrated servers (Deprecated)

Use these tasks to configure IBM i and your integrated server operating system for multipath I/O.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide.

Using hot spare hardware (Deprecated)

If there is a problem with your integrated server or iSCSI target adapter hardware, use these steps to change your IBM i configuration objects to point to new hardware.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide.

Configuring high availability for integrated servers (Deprecated)

Use these tasks to configure high availability iSCSI-attached integrated servers.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide.

Viewing integrated server messages

View IBM i message logs for integrated servers.

The **monitor job log** is a key source of information when troubleshooting integrated server problems. It contains messages that vary from normal processing events to detailed error messages. The monitor job always runs in the QSYSWRK subsystem with the same name as the integrated server.

To find the job log:

1. Select **Work Management** from *IBM Navigator for i*.
2. Select **Active Jobs**.

3. One of the jobs listed under the QSYSWRK section has the same name as your integrated server. Click the menu icon for it and select **Job Log**.
4. The integrated server monitor job log window opens. Click the menu icon for a message ID and select **Properties** to see details.

To find the job log in the character-based interface

1. At an IBM i command line, enter `WRKACTJOB SBS(QSYSWRK)`
2. One of the jobs listed has the same name as your integrated server. Select option `5` (Work with job).
3. Type `10` and press **Enter** to display the job log.
4. Press **F10** to see the detailed messages.

Launching the Web console for an integrated server

Do these steps to launch the integrated server service processor Web console that is associated with an IBM i network server description (NWSD).

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Servers**.
3. Click the menu icon for an integrated server from the list available and select **Launch Web console**.
4. Click the **Web Console** link on the **Launch Web Console** page.

The Web console for the integrated server service processor is shown in a separate Web browser window. For example, if the integrated server is a blade in an IBM BladeCenter that has an Advanced Management Module (AMM) service processor, then the AMM Web interface is shown.

Note: If the service processor Web console page does not appear, review the notes on the **Launch Web Console** page for possible reasons.

Managing integrated Windows servers (Deprecated)

Use these tasks to manage integrated servers running Windows Server.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide .

Updating the integration software running on Microsoft Windows (Deprecated)

IBM i Integrated Server Support includes software that runs on the Windows operating system. Whenever updates to this software are loaded on IBM i, you must synchronize the software from IBM i to Windows.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide .

Managing virtual Ethernet and external networks (Deprecated)

Use these tasks to configure and manage the Ethernet networks for integrated Windows servers.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide.

Sharing tape and optical devices between IBM i and integrated Windows servers (Deprecated)

Use these tasks to configure an integrated Windows server to use IBM i tape and optical devices.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide.

Running integrated Windows server commands remotely (Deprecated)

You can use IBM i to remotely submit integrated Windows server batch commands. Windows server commands that can run in batch mode without user interaction will work.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide.

Managing integrated VMware ESX servers (Deprecated)

Use these tasks to manage integrated servers running VMware ESX Server.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide.

Updating the integration software for VMware ESX (Deprecated)

The integrated server software for VMware ESX Server has some components that run on IBM i, and others that run on an integrated Windows server that serves as the management server for the integrated VMware ESX server.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide.

Configuring the QVMWINT user for integrated VMware ESX server management (Deprecated)

You need to set up the QVMWINT user in order to perform IBM i management tasks, such as shutting down the integrated VMware ESX server and linking storage to the ESX server while it is active.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide.

Managing connections for integrated VMware ESX servers (Deprecated)

The IBM i connection utility for virtualization hosts (`ibmvmcon.exe`) is used to define connection information so that IBM i can manage integrated VMware ESX servers.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide .

Migrating VMware ESX servers to new IBM i management infrastructure (Deprecated)

IBM i 7.1 uses a new management infrastructure for integrated VMware ESX servers. When upgrading IBM i to version 7.1 or later, you need to migrate integrated VMware ESX servers that were installed before 7.1.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide .

Managing storage for integrated servers

Use these tasks to manage storage for an integrated server.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Displaying information about integrated server virtual storage

Do these steps to display information about an integrated server disk drive (network server storage space) from IBM i.

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **All Virtual Storage**.
3. Select virtual storage from the list available
4. Click the menu icon for the virtual storage and select **Properties**.

Tip: If you want to use a CL command, see:
Work with NWS Storage Spaces (WRKNWSSTG)
Display NWS Storage Space (DSPNWSSTG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Adding virtual storage to integrated servers

Use these tasks to add virtual storage to an integrated server.

Related concepts:

"Virtual storage for integrated servers" on page 23
Integrated servers use virtual storage provided by IBM i instead of physical disk drives attached to the integrated server hardware.

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Creating virtual storage for integrated servers

Do these steps to create virtual storage for an integrated server.

Creating a storage space in an independent storage pool (ASP) requires that the storage pool device is available.

To create virtual storage for an integrated server, do these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **New Virtual Storage**.
3. Click **Continue** on the **Select Base Object** page.
4. Specify the virtual storage drive name and description.

   Note:

   - Consider using a naming scheme to allow easy identification of storage spaces and to allow using generics (*) on the save commands. Otherwise you might have trouble correlating storage space names that you see from IBM i with disk drives that you see from the integrated server. Correlating storage space names can be especially difficult if you have both storage that was linked to the server while it was shut down or while it was started (dynamic linking).
   - This name is also used for the storage space object created in the /QFPNWSSTG directory of the integrated file system.
5. If you want to copy data from an existing storage space, select **Initialize storage with data from existing storage**. Then select the source storage space to copy data from.
6. Specify the storage capacity in megabytes (MB) or gigabytes (GB).
7. Select a storage pool (disk pool) to contain the virtual storage.
8. Select the planned file system.
   - For integrated Windows Server 2003 servers, use **NTFS**.
   - For integrated Windows Server 2008 and Windows Server 2008 R2 servers, use **NTFS**. Also use the **Advanced Data Offset** button and select **Align the first logical storage sector**.
   - » For integrated Windows Server 2012, use NTFS or ReFS. Also use the Advanced Data Offset button and select Align the first logical storage sector.«
   - For integrated VMware ESX servers, use **Open source**. Also use the **Advanced Data Offset** button and select an alignment value based on how the storage is used. See the following table.

Table 8. Data offset values to use for storage linked to integrated VMware ESX servers

Storage Usage	Storage Data Offset Value
ESX system disk	**Align the first logical partition sector**

Table 8. Data offset values to use for storage linked to integrated VMware ESX servers (continued)

Storage Usage	**Storage Data Offset Value**
Storage for guest operating systems: • Windows Server 2012 • Windows Server 2008 • Windows Server 2008 R2 • Windows Vista • Windows 7	**Align the first logical storage sector**
Storage for guest operating systems: • Windows Server 2003 • Windows XP • Windows 2000 • Windows NT 4.0 • Linux • Any other guest operating systems not previously listed	**Align the first logical partition sector**

When a sector size of 4096 bytes per sector is specified, the data offset can't be specified. It will be equivalent to select Align the first logical storage sector regardless of which file system is used.

9. Optional:
   - If you want to immediately link the storage to a server after it is created, check **Link storage to server**and provide the linking attributes.
   - Specify the storage **Sector size in bytes**.
   - Specify the **Preferred storage unit** for this storage.
   - Select the relative **Resource allocation priority** to be used by the hosting system to initialize the disk when the client operating system formats this storage space. This task can impact the operating system performance. Choose to run it at a priority which suits your system. Higher priority will consume more resources and complete more quickly.
10. Click **OK**.

The process of creating a storage space can range from a few minutes to a few hours, depending on the size. When IBM i finishes creating the storage space it is listed with the other storage spaces.

Tip: If you want to use a CL command, see:
Work with NWS Storage Spaces (WRKNWSSTG)
Create NWS Storage Space (CRTNWSSTG)

After creating the storage, you must link it to the network server description of your integrated server. Then you must format the storage using the integrated server operating system disk management utilities. For Windows servers, partition and format it using Windows **Disk Management** or by using the `DISKPART` command-line utility.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Linking virtual storage to integrated servers

Integrated servers can only access virtual storage that is linked to the Network Server Description (NWSD) for the server.

You must create the virtual storage before you can link it. See "Creating virtual storage for integrated servers" on page 81. After you create and link the virtual storage, it appears as a new hard disk drive to the integrated server. Then you must format it before you can use it.

Note: Virtual storage can be dynamically linked to a server while the server is active.

To link virtual storage to an integrated server, follow these steps:

1. If you are not linking a disk drive dynamically, then shut down your integrated server. See "Stopping integrated servers" on page 72.
2. Select **Integrated Server Administration** from *IBM Navigator for i*.
3. Select **All Virtual Storage**.
4. Click the menu icon for an available virtual storage and select **Add Link**.
5. Select the server you want to link the disk to.
6. Select one of the available link types and the link sequence position.
7. Select one of the available storage paths.
8. Select one of the available data access types.
9. Click **OK**.
10. Start the integrated server. See "Starting integrated servers" on page 70.
11. When the server is started, format the virtual storage. You can use the utilities provided by the integrated server operating system. See "Formatting virtual storage" on page 84.

Tip: If you want to use a CL command, see:
Work with NWS Storage Spaces (WRKNWSSTG)
Add Server Storage Link (ADDNWSSTGL)

Related concepts:

"Virtual storage linking for integrated servers" on page 28
Integrated servers do not use physical disk drives. IBM i creates virtual storage (network server storage spaces) within its own file system and integrated servers use them as if they were normal physical disk drives.

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Managing disk drives for Windows when running out of drive letters:
The maximum number of virtual disk drives (virtual storage) that can be linked to an integrated server is greater than the number of drive letters that are available on Windows. Since not all drives will have a drive letter, other options must be used to utilize all storage linked to the server. Here are two options to utilize all virtual disk drives which are linked to a server.

1. A disk drive letter can be made up of multiple disk drives using a spanned volume set.

   Note: When you create a volume set, all of the existing data on the partitions that you use for the new volume set is erased. You should consider volume sets while you are setting up your server.

   a. From **Disk Management**, right-click each disk drive number and select **Upgrade to Dynamic Disk...** from pop-up menu.
   b. Right-click a disk drive partition and select **Create Volume...** from pop-up menu.
   c. Follow the create volume wizard to create a spanned volume, making sure to add the multiple disks.

      Note: This feature is nice because if the volume gets full, a disk can be dynamically added, and it will be immediately joined to the spanned volume without ever requiring to reboot the server.
2. A disk drive can be mounted over a subdirectory of an existing disk drive letter.
   a. Create a directory on a disk drive letter that is formatted with NTFS. For example, MD C:\MOUNT1.
   b. From **Disk Management**, click over disk drive partition you want to format and select **Format** from the pop-up menu.
   c. Once drive is formatted, right-click over disk drive partition again and select **Change Drive Letter and Path...** from pop-up menu.
   d. Select **Add**.
   e. Select radio button **Mount in this NTFS folder:**
   f. Use **Browse** button to find directory C:\MOUNT1 that was created in step 1.
   g. Click **OK** to make that directory a mount point for this disk drive.

Formatting virtual storage

In order to use integrated server virtual disks (network server storage spaces), you must format them.

Before you can format virtual disks, you must first create them (see "Creating virtual storage for integrated servers" on page 81) and link them (see "Linking virtual storage to integrated servers" on page 83). Then start the integrated server from IBM i (see "Starting integrated servers" on page 70).

Formatting storage for VMware ESX servers:

For an integrated server with VMware ESX, format virtual storage based on the operating system that will use the storage.

Refer to VMware ESX documentation for information on partitioning and formatting storage for ESX and the associated virtual machines. See White Paper Aligning storage partitions for VMware ESX Server on iSCSI attached integrated servers for additional considerations to help you partition your storage for improved performance.

Formatting storage for Windows servers:

Do these steps to format virtual storage for an integrated server with the Microsoft Windows operating system.

1. On the integrated Windows server console, select **Start** > **All Programs** > **Administrative Tools** > **Computer Management**.
2. Double-click **Storage.**
3. Double-click **Disk Management.**

4. To create a new partition, right-click the unallocated space on the basic disk where you want to create the partition, and then click **New Partition**.
5. Follow the prompts to format the new drive.
   a. Specify the storage space name for the volume label.
   b. Select the file system you specified when you created the virtual storage.
   c. Select the quick format for a storage space that has just been created. It has already been low level formatted by IBM i when it was allocated.

Copying virtual storage

Do these steps to create new virtual storage (virtual disk) for an integrated server with information from an existing disk.

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **All Virtual Storage**.
3. Select virtual storage from the list available.
4. Click the menu icon for the virtual storage and select **New Based On**.
5. Specify a name and description.
6. Specify the virtual storage capacity. See the online help for details on valid disk sizes associated with a particular file system format. If you want to increase the size of the disk while copying it, you can specify a larger size. The extended portion of the disk will be unpartitioned free space.

   Note: For integrated Windows servers, you can use the DISKPART command line utility to expand an existing partition in order to utilize any additional free space. Refer to the Microsoft Knowledge base articles for DISKPART for details and limitations.
7. Select a storage pool (auxiliary storage pool) to contain the disk.
8. Optional:
   - » Specify the **Preferred storage unit** for this storage.«
   - » Select the relative **Resource allocation priority** to be used by the hosting system to initialize the disk when the client operating system formats this storage space. This task can impact the operating system performance. Choose to run it at a priority which suits your system. Higher priority will consume more resources and complete more quickly.«
9. » Click **OK**.«

Tip: If you want to use a CL command, see:
Work with NWS Storage Spaces (WRKNWSSTG)
Create NWS Storage Space (CRTNWSSTG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Expanding virtual storage

Do these steps to expand integrated server virtual storage (disk drive).

For information about expanding a boot disk, see the IBM i iSCSI Solution Guide .

To expand virtual storage, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **All Virtual Storage**.
3. Select virtual storage from the list available.
4. Click the menu icon for the virtual storage and select **Properties**.
5. Click on the **Capacity** tab of the virtual storage property sheet.
6. Specify the increased virtual storage size in the **New capacity** field. See the online help for details on valid disk sizes associated with a particular file system format. The extended portion of the disk will be unpartitioned free space.
7. Click **OK**.
8. If the virtual storage is linked to an active server, a confirmation panel is shown to indicate that the virtual storage will be temporarily unavailable to the server while the virtual storage is being expanded. Click **Change** on the confirmation panel to confirm that this is acceptable, or click **Cancel** on the confirmation panel to cancel the virtual storage expansion operation.

Tip: If you want to use a CL command, see:
Change NWS Storage Space (CHGNWSSTG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Expanding a system disk for an integrated Windows server (Deprecated)

To expand an integrated Windows server system disk, unlink the disk from the integrated server, expand the disk, and then relink the disk to the server.

Note: The information in this section has been migrated to the IBM i iSCSI Solution Guide .

Unlinking virtual storage

Unlink virtual storage to make it inaccessible to the integrated server.

Restrictions:

1. For an integrated Windows server, see "Virtual storage linking for integrated servers" on page 28 for information about when virtual storage can be dynamically unlinked while the server is active.
2. You cannot dynamically unlink virtual storage from an active integrated VMware ESX server.

If you do not want to dynamically unlink the virtual storage for an integrated Windows server, or if the server is an integrated VMware ESX server, shut it down. See "Stopping integrated servers" on page 72.

To unlink virtual storage, complete the following steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **All Virtual Storage**.
3. Click the menu icon for the virtual storage you want to unlink.
4. Select **Remove link** to open the **Remove Link from Server** window.

5. Optional: If multiple servers are linked to the storage, select one or more servers to unlink from the storage.
6. Optional: To eliminate gaps in the sequence of the virtual storage spaces, click **Compress link sequence**.
7. Click **Remove** to unlink the virtual storage.

Tip: If you want to use a CL command, see:
Work with NWS Storage Spaces (WRKNWSSTG)
Remove Server Storage Link (RMVNWSSTGL)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Deleting virtual storage

Use this task to delete virtual storage for an integrated server.

Before you can delete virtual storage, you must do the following actions:

1. If the virtual storage is linked to an active integrated VMware ESX server, stop the server. See "Stopping integrated servers" on page 72.
2. If the virtual storage is linked to any integrated servers, unlink the virtual storage from all the integrated servers. See "Unlinking virtual storage" on page 86

Do the following steps to delete virtual storage for an integrated server:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **All Virtual Storage**.
3. Click the menu icon for the virtual storage that you want to delete and select **Delete**.

   Tip: To process multiple objects, select the objects in the list. Then use **Select Action** > **Delete**.
4. Click **Delete** on the confirmation panel.

Tip: If you want to use a CL command, see:
Work with NWS Storage Spaces (WRKNWSSTG)
Delete NWS Storage Space (DLTNWSSTG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Accessing the IBM i integrated file system from an integrated server

You can access the IBM i integrated file system from an integrated server through IBM i Support for Windows Network Neighborhood (IBM i NetServer). This allows you to easily work with file system resources on IBM i.

For information about using IBM i NetServer, see:

- Getting started with IBM i NetServer
- Creating IBM i NetServer file shares
- Configuring and connecting your PC client

- Accessing file shares from a Windows client

Administering integrated Windows server users from IBM i

Use these tasks to manage integrated Windows server users from IBM i

One of the main advantages of using integrated Windows server is synchronized, simplified user administration. Existing IBM i user profiles and groups of profiles can be enrolled to integrated Windows servers, meaning that those users can log onto Windows server with the same user ID and password pair that they use to sign on to IBM i. If they change their IBM i password, their Windows password changes as well.

Enrolling IBM i users to integrated Windows servers

To enroll IBM i users to integrated Windows servers, follow these steps.

Create an IBM i user profile for the user if one does not already exist. You can find information about creating IBM i user profiles in the Security topic collection.

To enroll a single user to the integrated Windows server, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Servers** or **Domains**.
3. Click the menu icon for an available Windows domain or server from the list and select **Enroll Users**.

   Note: Do not select a Windows workgroup. Enrollment to a workgroup is not supported.
4. Select to enter the user name or choose the user name from the list.
5. Optional: If you want to use a user template as a basis for user settings, specify a Windows user to use as a template when creating the user on Windows. Remember that if you change the user template after you enroll a user, the changes will not affect the user.
6. Click **Enroll**.

Tip: If you want to use a CL command, see:
 Work with NWS User Enrollment (WRKNWSENR)
 Change NWS User Attributes (CHGNWSUSRA)

Related concepts:

"User and group enrollment concepts for integrated Windows servers" on page 56
Learn about how IBM i users and groups interact with integrated Windows servers.

Enrolling IBM i groups to integrated Windows servers

To enroll IBM i groups to integrated Windows servers, follow these steps.

You can find information about creating IBM i user and group profiles in the Security topic collection.

To enroll an IBM i group and its members to the integrated Windows server, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Servers** or **Domains**.

3. Click the menu icon for the domain or server, then select **Enroll Groups**.

   Note: Do not select a Windows workgroup. Enrollment to a workgroup is not supported.
4. Enter a group name or select an unenrolled group from the list.
5. Optional: To use a template to create new users, specify a Windows user to use as a template when creating users in the group on Windows. If you change the user template after you enroll a user, the changes do not affect the user.
6. Select **Global** if the group is being enrolled in a domain and the group should be visible to the domain. Otherwise, select **Local** . Windows server local groups can contain users and Windows server global groups, while Windows server global groups can contain only users. See the Windows online help for more information about group types.
7. Click **Enroll**.

Tip: If you want to use a CL command, see:
Work with NWS User Enrollment (WRKNWSENR)
Change NWS User Attributes (CHGNWSUSRA)

Related concepts:

"User and group enrollment concepts for integrated Windows servers" on page 56
Learn about how IBM i users and groups interact with integrated Windows servers.

Configuring the QAS400NT user for user enrollment on integrated Windows servers

You need to set up the QAS400NT user in order to successfully enroll an IBM i user or group profile on a domain or local server in these situations.

- You are enrolling on a domain through a member server.
- You are enrolling on a local server using a template which specifies a home directory path
- You are enrolling on a domain through an IBM i partition which contains both domain controllers and member servers on the same domain.

You do not need to set up the QAS400NT user in order to successfully enroll an IBM i user or group profile on a domain or local server in the following cases:

- You are enrolling on a domain through an IBM i partition which contains a domain controller but no member servers on the same domain.
- You are enrolling on a local server (or locally on a member server) using a template which does not specify a home directory path.

If you need to set up the QAS400NT user, follow these steps:

1. Create the QAS400NT user profile on IBM i with User class *USER. Take note of the password because you need it in the next step. Make sure that the password complies with the rules for Windows passwords if you are enrolling on a domain.
2. Create the QAS400NT user account on the Windows console of the integrated Windows server you are enrolling through. Note that the IBM i user profile password and Windows user account password must be the same for the QAS400NT user.
   a. Setting up QAS400NT on a domain controller

      On the domain controller of the domain you are setting up enrollment for, create the QAS400NT user account as follows:

1) From the integrated server console

   a) Click **Start** > **All Programs** > **Administrative Tools** > **Computer Management** > **System Tools** > **Local Users and Groups**.

   b) Select **System Tools –> Local Users and Groups**.

2) Right-click the **Users** folder (or the folder that the user belongs to), and select **New** > **User**...

3) Enter the following settings:

```
Full name: qas400nt
User logon name: qas400nt
```

4) Click Next. Enter the following settings:

```
Password: (the same password as you used for QAS400NT on IBM i)
Deselect: User must change password at next logon
Select: User cannot change password
Select: Password never expires
```

5) Click Next, then Finish

6) Right click the QAS400NT user icon and select Properties.

7) Click the **Member Of** tab and then Add.

8) Enter `Domain Admins` in the box and click OK, then OK again. This gives the QAS400NT user account sufficient rights to create users.

b. Setting up QAS400NT on a local server

On the local server (or member server if you are enrolling locally) you are setting up enrollment for, create the QAS400NT user account as follows:

1) From the integrated server console

   - In Windows Server 2003 click **Start** > **All Programs** > **Administrative Tools** > **Computer Management** > **System Tools** > **Local Users and Groups**.

2) Right-click the **Users** folder, and select **New User...**

3) Enter the following settings:

```
User name: qas400nt
Full name: qas400nt
Password: (the same password as you used for QAS400NT on IBM i)
Deselect: User must change password at next logon
Select: User cannot change password
Select: Password never expires
```

4) Click Create, then Close.

5) Right click the QAS400NT user icon and select Properties.

6) Click the Member Of tab and then Add.

7) Enter Administrators in the box and click OK, then OK again. This gives the QAS400NT user account rights to the User Administration Service.

3. Enroll the IBM i QAS400NT user profile on the domain or local server using *IBM Navigator for i* or the Change NWS User Attributes (CHGNWSUSRA) command. Do not try to use a template when enrolling QAS400NT.

4. Use *IBM Navigator for i* or the Work with NWS User Enrollment (WRKNWSENR) command to confirm that QAS400NT has been successfully enrolled. You may now enroll IBM i user profiles through domain controllers or member servers on the domain.

Notes:

- You may change the QAS400NT password from IBM i since it is now an enrolled user.
- If there are multiple integrated servers that belong to different domains on a single IBM i partition, you must set up QAS400NT for each domain. All

QAS400NT user accounts must have the same password as the IBM i user profile. Alternatively, consider using Active Directory or trust relationships between domains, and enroll users on only a single domain.

- If you have multiple IBM i partitions and multiple integrated servers, QAS400NT passwords on different IBM i partitions can be different as long as each domain does not contain integrated servers on more than one IBM i partition. The rule is, all IBM i QAS400NT user profiles and corresponding Windows user accounts must have the same password for a single domain.
- Be sure not to delete the QAS400NT user profile on IBM i, or let the password expire. To minimize the risk of the QAS400NT password expiring on one of multiple IBM i partitions on the same Windows domain, it is recommended that you allow only one IBM i partition to propagate changes to the QAS400NT user profile.
- If you have multiple IBM i partitions, each with an integrated Windows server on the same domain, failing to keep the QAS400NT password synchronized across all IBM i partitions can cause enrollment problems. To minimize this problem, it is recommended that you limit propagation of changes to the QAS400NT password to just one IBM i partition, but still allow other partitions to keep sufficient authority to enroll users. Then, failure to change a password on one of the other partitions prevents user enrollment from that partition only.

Creating user enrollment templates for integrated Windows servers

Follow these steps to create user enrollment templates.

A user enrollment template is a tool to help you enroll users from IBM i to the Windows environment more efficiently. You do not have to manually configure many new users with identical settings.

You can make a user template a member of any Windows server group, whether you enrolled that group from IBM i or not. You can enroll users with a template that is a member of a group that was not enrolled from IBM i. If you do this you can only remove users from the group by using the User Manager program on Windows server.

If you are creating a template that will be used to enroll administrators, you may want to make the template a member of the Windows server group *Administrators*. Likewise, if you want to protect Windows users from accidental deletion from IBM i, enroll the template in the *AS400_Permanent_Users* (or OS400_Permanent_Users) group.

Follow these steps to create a Windows template.

Creating user enrollment templates on a Windows domain

Do these steps at the integrated server console.

1. Click **Start** > **All Programs** > **Administrative Tools** > **Active Directory Users and Computers**.
2. Click the domain name.
3. Right-click **Users**, then select **New** > **User**.
4. In the **Username** and **Logon name** fields, enter a distinctive name for the template, such as *stduser* or *admtemp*. Click **Next**.
5. It is recommended that you also deselect the **User must change password at next logon** check box and select the **User cannot change password**, **Password**

never expires, and **Account is disabled** checkboxes. This prevents anyone using the template account itself to access the integrated server.

6. Do not enter a password for a template account.
7. Click **Finish**.
8. To set up group memberships, double-click the template name in the list of domain users and groups that appear in the right pane. Click the **Member of** tab and then click **Add** to add the groups that you want.

Creating user enrollment templates on a Windows server

Do these steps at the integrated server console.

1. Click **Start** > **All Programs** > **Administrative Tools** > **Computer Management**.
2. Select **System Tools** > **Local Users and Groups**.
3. Right-click **Users** and select **New User**.
4. In the **User name** field, enter a distinctive name for the template, such as *stduser* or *admtemp*.
5. It is recommended that you also deselect the **User must change password at next logon** check box and select the **Password never expires**, **User cannot change password**, and **Account is disabled** checkboxes. This prevents anyone using the template account itself to access Windows server.
6. Click **Create**, then **Close**.
7. Click **Users** or refresh to show the new user template.
8. To set up group memberships, double-click the template name in the list of domain users and groups that appears in the right pane. Click the **Member of** tab and then click **Add** to add the groups that you want.

Specifying a home directory in a user template

Follow these steps to specify a home director in a user template.

To allow integrated Windows servers to manage users in the most portable way possible, a home directory can be set up for each user to store user-specific information generated by applications. To minimize the amount of work that must be done, specify home directories in the template accounts so that each new profile created by the enrollment process has a home directory created for it automatically. To provide scalability, it is important not to lock home directories to a particular disk drive. Use the Universal Naming Convention (UNC) names to give portability.

To customize your template profiles to include a home directory, follow these steps from the integrated Windows server console:

1. Create the home directory folder on the appropriate server, and share it.
2. In a domain, click **Start** > **All Programs** > **Administrative Tools** > **Active** > **Directory Users and Computers** from the Windows console. On a local server, click **Start** > **All Programs** > **Administrative Tools** > **Computer Management** > **Local Users and Groups**.
3. Double-click the template (model user) to display its properties.
4. Click the Profile tab.
5. In the Home folder segment, click **Connect**. Select a drive letter (such as Z:). Move to the **To:** dialog, and enter the directory path of the home directory using a UNC name, for example: `\\systemiWin\homedirs\%username%`. In this example, **systemiWin** is the name of the server where the home directory folder resides, and **homedirs** is the name of the home directory folder. If you use the variable *%username%*, instead of the logon or user name, Windows

server automatically substitutes the user's name in place of the variable name when each new Windows server account is created. It also creates a home directory for the user.

Changing the local password management user profile attribute

Use these steps to change the local password management (LCLPWDMGT) user profile attribute.

1. Select **Users and Groups** from *IBM Navigator for i.*
2. Select **Change User**.
3. Type the user profile name you want to change, then click **OK**.
4. Select **Capabilities**.
5. On the **Password** tab, select one of the following options and click **OK**.
   - **Manage this password locally through IBM i**
   - **Manage this password remotely through some other platform**
6. Click **OK** again on the user properties page to save the change.

Tip: If you want to use a CL command, see:
Work with User Profiles (WRKUSRPRF)
Change User Profile (CHGUSRPRF)

Related concepts:

"User and group enrollment concepts for integrated Windows servers" on page 56
Learn about how IBM i users and groups interact with integrated Windows servers.

Configuring Enterprise Identity Mapping for integrated Windows servers

Use this information to configure a user account to use EIM.

What is EIM?

Enterprise Identity Mapping (EIM) is a way to consolidate a user's various UserIDs and passwords together under a single account. Using it, a user can log on just once to a system, and then EIM will work together with other services behind the scenes to authenticate the user to all of his accounts.

This is called a single sign-on environment. Authentication still takes place whenever users attempt to access a new system; however, they will not be prompted for passwords. EIM reduces the need for users to keep track of and manage multiple user names and passwords to access other systems in the network. Once a user is authenticated to the network, the user can access services and applications across the enterprise without the need for multiple passwords to these different systems.

The EIMASSOC user profile attribute

EIMASSOC is a user profile attribute specifically designed to aid in configuring EIM. At the IBM i command prompt type `CHGUSRPRF` and the user profile name and then press F4 to prompt. Then page down to the very bottom and you will see a section labled `EIM association`. Here is a summary of what the fields mean:

- **Element 1: EIM identifier** This is the UserID that EIM uses to identify you. Think of it as your Master ID under which all your other user IDs will be stored. If you specify *USRPRF the system will use your IBM i user profile name as the EIM identifier. Alternatively, you can specify any valid character-string. If you enter *DLT in this field and press enter, you will be presented with a list of changed options for deleting EIM associations.

- **Element 2: Association type** This value specifies how the IBM i user profile that you are editing will be associated with the EIM identifier. The values of *TARGET, *TGTSRC, or *ALL will allow auto-creation or deletion of IBM i target and Windows source associations.
- **Element 3: Association action** The special values are:
  - *REPLACE The Windows source associations will be removed from all EIM identifiers that have an association for this user profile. For the enrolled user, a new Windows source association will be added to the specified EIM identifier.
  - *ADD For the enrolled user, a Windows source association will be added.
  - *REMOVE The Windows source association will be removed.
- **Element 4: Create EIM identifier** This value specifies whether the EIM identifier should be created if it does not already exist. The special values allowed are, *NOCRTEIMID, an EIM identifier will not be created, or, *CRTEIMID, an EIM identifier will be created if it does not exist.

Automatic and Manual EIM associations

In a typical EIM configured environment, which uses single sign-on, IBM i target associations and Windows source associations are typically defined. With integrated Windows server user administration, the system administrator might decide to define enrolled Windows users to have EIM associations automatically defined. For instance, if an enrolled Windows user has EIMASSOC(*USRPRF *TARGET *ADD *CRTEIMID) specified, IBM i will automatically create an IBM i target and a Windows source association. The EIMASSOC information is not stored in the user profile. Also, this information is not saved or restored with the user profile. And, if the IBM i system is not configured for EIM, then no association processing is done and the EIMASSOC information is ignored.

If IBM i is configured to use EIM and EIMASSOC processing is defined for the enrolled user, integrated Windows server user administration will auto create or delete Windows source associations for the user in the Windows EIM registry. For a user enrolled locally to the Windows environment, the Windows EIM registry name is the fully qualified, local Domain Name System (DNS) name. The Windows EIM registry type is defined to be Windows 2000. For users enrolled to a Windows domain, the Windows registry name is the fully qualified domain DNS name and the Windows registry type is defined to be Kerberos - case ignore. If EIMASSOC is defined for a user, and IBM i is configured to use EIM, and the Windows EIM registry doesn't exist, integrated Windows server user administration will create the Windows EIM registry.

Use EIM associations to allow different Windows user profile names

EIM provides a mechanism to associate user profiles in a directory system. EIM allows for an EIM identifier to have an IBM i user profile target association defined and a Windows user profile source association to be defined. It is possible for a user administrator to define a Windows source association using a different Windows user profile name than the IBM i target association user profile name. Integrated Windows user administration will use the defined EIM Windows source association Windows user profile, if it exists, for Windows user enrollment. The

IBM i target association needs to be defined. Using the EIM identifier, the Windows source association needs to be defined by the administrator. The Windows source association needs to be defined for the same EIM identifier in the correct Windows EIM registry name and type. For a user enrolled locally to Windows, the Windows EIM registry name is the fully qualified, local domain name server (DNS) name. The Windows EIM registry type is defined to be EIM_REGTYPE_WIN2K. For users enrolled to a Windows domain, the Windows registry name is the fully qualified domain DNS name and the Windows registry type is defined to be EIM_REGTYPE_KERBEROS_IG.

Related concepts:

"User and group enrollment concepts for integrated Windows servers" on page 56
Learn about how IBM i users and groups interact with integrated Windows servers.

Ending user enrollment to an integrated Windows server

To end the enrollment of a user to Windows domains and servers, do these steps at the Windows console.

To end the enrollment of a user to Windows domains and servers, follow these steps on the integrated Windows server console:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Servers** or **Domains**.
3. Click the menu icon for the domain or server that contains the user that you want to unenroll, then select **Open**.
4. Click the menu icon for **Enrolled Users**, then select **Open**.
5. Click the menu icon for the user that you want to unenroll, then select **Unenroll**.
6. Click **Unenroll** on the confirmation window.

Effects of ending user enrollment to the integrated Windows server

When you end user enrollment from the Windows environment, you also remove the user from the list of enrolled Windows server users, as well as from the Windows server group AS400_Users (or OS400_Users). Unless the user is a member of the Windows server group AS400_Permanent_Users (or OS400_Permanent_Users), you also delete the user from the Windows environment.

You cannot delete users who are members of the Windows server group AS400_Permanent_Users (or OS400_Permanent_Users) from Windows server by either ending enrollment or deleting them from IBM i. However, ending enrollment does remove the user from the list of enrolled Windows server users and from the Windows server group AS400_Users (OS400_Users).

You can keep users on the Windows environment after you have ended their enrollment on IBM i. However, this practice makes it possible to add these users to groups on IBM i and change passwords on IBM i without these updates ever appearing in the Windows environment. These discrepancies can make it difficult to keep track of users on either system.

You can end user enrollment in a number of ways. Actions that end user enrollment include the following:

- Intentionally ending enrollment for the user.

- Deleting the IBM i user profile.
- Ending enrollment for all IBM i groups to which the user belongs.
- Removing the user from an enrolled IBM i group when the user does not belong to any other enrolled groups.

Ending group enrollment to an integrated Windows server

To end the enrollment of a group to Windows environment domains and servers, follow these steps.

When you end enrollment of a group to the integrated Windows server, all users whose enrollment is limited to that group also have their enrollment ended. If the group has only members that were enrolled through it, the group is deleted from the integrated Windows server.

However, if the group has any members that were added from the Windows operating system rather than enrolled from IBM i, the group is not deleted. The only members that the group can still have are nonenrolled users.

To end the enrollment of a group to Windows domains and servers, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Servers** or **Domains**.
3. Click the menu icon for the domain or server that contains the group that you want to unenroll, then select **Open**.
4. Click the menu icon for **Enrolled Groups**, then select **Open**.
5. Click the menu icon for the group that you want to unenroll, then select **Unenroll**.
6. Click **Unenroll** in the confirmation window.

Preventing enrollment to an integrated Windows server

Use these tasks to prevent users from being enrolled to an integrated Windows server.

There are several reasons why you might want to prevent IBM i user profile enrollment to a particular integrated server:

- If there are multiple integrated servers that belong to the same domain, and they are all on the same IBM i partition, user profile enrollment will, by default, go through all of the integrated servers in that partition. To reduce network traffic you can turn off enrollment to all integrated servers on the domain except one. This single integrated server would normally be the domain controller, if it is in the partition.
- If there are multiple integrated servers that belong to the same domain, but they are all on different IBM i partitions, there is a risk of the QAS400NT passwords getting out of synchronization and causing problems with user profile enrollment. By preventing enrollment of the QAS400NT user profiles from all IBM i partitions except one, you can reduce the risk of enrollment problems. Notice that the other IBM i partitions keep sufficient authority to enroll users. Then, failure to change a password on one of the other partitions prevents user enrollment from that partition only.

There are two methods to prevent IBM i user profile enrollment to a particular integrated server:

- Use the Propagate Domain User (PRPDMNUSR) parameter of the network server description (NWSD). See below for a description of how to do this.
- Create data areas with the Create data area (CRTDTAARA) command. See below for a description of how to do this.

Notes:

- Do not turn enrollment off for all of the integrated servers on the domain. Otherwise all your users may go to update pending (*UPDPND) status, and no further enrollment takes place.
- You may want to leave two integrated servers enabled for user enrollment so that you can still make changes if one of the servers is down.

Using the PRPDMNUSR parameter to prevent enrollment to a domain through a specific integrated server

The Propagate domain user (PRPDMNUSR) parameter of the Change network server description (CHGNWSD) command can be used to prevent user enrollment to a domain through a specific integrated server.

This option may be useful in the case where there is a single IBM i partition which controls multiple integrated Windows servers that belong to the same domain, because it can turn off enrollment for all integrated servers except one.

To use the PRPDMNUSR parameter to prevent user enrollment, do these steps.

1. Choose one integrated server you wish to stop domain enrollment on. (You do not need to shut down the server.)
2. Run the following Change Network Server Desc (CHGNWSD) command:

   `CHGNWSD NWSD(`*`nwsdname`*`) PRPDMNUSR(*NO)`where *nwsdname* is the name of the network server description for the integrated server.

Using the CRTDTAARA command to prevent enrollment of QAS400NT to a specific integrated server

The Create Data Area (CRTDTAARA) command can be used to prevent enrollment of the QAS400NT user profile only, for the specified integrated server. The enrollment of other user profiles is not affected.

This option may be useful in the case where there are multiple integrated servers that belong to the same domain, but they are all on different IBM i partitions. You want to enroll user profiles from these different IBM i partitions, but not have multiple QAS400NT user profiles propagating passwords to the domain. Follow these steps:

1. Choose one IBM i partition that you wish to use for enrollment of QAS400NT on the domain. Ensure that QAS400NT is enrolled on this IBM i partition.
2. If QAS400NT is enrolled on other IBM i partitions follow these steps:
   a. On the domain controller, add the QAS400NT user account to the OS400_Permanent_Users group to ensure that it is not deleted.
   b. On the IBM i partitions where you want to prevent enrollment of QAS400NT, delete the QAS400NT user profile.
3. On the IBM i partitions where you want to prevent enrollment of QAS400NT, create a data area with the following Create Data Area (CRTDTAARA) command:

   `CRTDTAARA DTAARA(QUSRSYS/`*`nwsdname`*`AU) TYPE(*CHAR) LEN(10) VALUE(*NOPROP)`where *nwsdname* is the name of the network server description for the integrated server, and ***NOPROP** is the keyword that signals that QAS400NT user profile parameters (including the password) are not propagated from this IBM i partition.

4. Create and enroll the QAS400NT user profile on each of the IBM i partitions you created the data area on. Notice that you still need to keep the QAS400NT password current (not expired) on all these IBM i partitions for enrollment of user profiles (other than QAS400NT) to occur. Because the QAS400NT password is not propagated, it does not matter what the password is, as long as it is not expired.

Managing network server host adapters

Network server host adapter (NWSH) objects are used to configure the IBM i iSCSI target adapter. Use these tasks to manage NWSH objects.

An NWSH object must be started (varied on) in order for an integrated server to use the corresponding iSCSI target for storage or virtual Ethernet data flows. Stopping (varying off) an NWSH object will make the corresponding iSCSI target unavailable to any integrated servers that have storage or virtual Ethernet paths defined to use it.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Creating a network server host adapter

A network server host adapter (NWSH) object must be created for each IBM i iSCSI target port.

Notes:

1. If you are using the *IBM i iSCSI Solution Work Sheets* PDF on the IBM i iSCSI Solution Guide Web page, use the following work sheets to help you do these tasks:

   IBM i iSCSI target network server host adapter

   IBM i TCP/IP interface for iSCSI software target NWSH

   IBM i line description for iSCSI software target NWSH

2. The network server host adapter and the remote system configuration define IP address information for opposite sides of the iSCSI network. When connected by a simple, switched network, the following rules apply:
   - The SCSI IP addresses in these two objects that are connected by a switch must be in the same subnet. For example, with IP addresses of the form a.b.x.y and 255.255.255.0 subnet masks, a.b.x must be the same value for both objects.
   - The LAN IP addresses in these two objects that are connected by a switch must be in the same subnet.
   - If you use CL commands to create the network server host adapter or remote system configuration, set the gateway elements to `*NONE`.

To create an NWSH, follow these steps:

1. Determine the IBM i hardware resource name that was assigned to the iSCSI target adapter port. For more information, see "Determining the hardware resource name for an iSCSI target adapter" on page 100.

2. Select **Integrated Server Administration** from *IBM Navigator for i*.
3. Select **New Network Server Host Adapter**.
4. Click **Continue** on the **Select Base Object** page.
5. On the **General** tab:
   a. Enter the NWSH device **Name** and **Description**.
   b. Select the **Hardware resource**.
      - Select **Virtual** if your iSCSI target adapter is an Ethernet NIC.
      - Select the resource name that was determined in step 1 if your iSCSI target adapter is an iSCSI HBA.
   c. Optional: Select **Online at IPL** if your iSCSI target adapter is an iSCSI HBA and you want it to automatically start when IBM i starts.

      Note: If your iSCSI target adapter is an Ethernet NIC, the equivalent function is accomplished by setting the corresponding TCP/IP interface to automatically start when TCP/IP is started on IBM i.
   d. Optional: Select the **Object authority**. You can use the default value **Change**.
6. On the **Local (Target) Interface** tab:
   a. Select the cable connection type. If the hardware is physically connected to an Ethernet switch, you can use the default value **Network**.
   b. Specify the remaining values based on iSCSI target adapter type.
      - For an **iSCSI HBA**:
        1) Enter a **Subnet mask**.
        2) Enter a SCSI interface **Internet address**.
        3) Enter a LAN interface **Internet address**.
      - For an **Ethernet NIC**:

        Select an IBM i TCP/IP interface for the SCSI interface **Internet address**.

        Tip: If you did not previously create an IBM i TCP/IP interface and corresponding line description for your iSCSI target adapter, click the **New** button to create them now:
        1) For **TCP/IP interface**:
           a) Enter an **Internet address**, **Subnet mask** and **Description**.
           b) Select **Start this TCP/IP interface every time TCP/IP is started** if you want the new NWSH to start automatically.
        2) For **Line description to use for the TCP/IP interface**:
           – If the line description exists, select it from the list.
           – Otherwise, enter the remaining values to create a line description.
             a) Enter a **Name** and **Description**.
             b) Select the **Hardware resource** for your iSCSI target adapter port that was determined in step 1.
             c) Set the **Maximum frame size**.
        3) Click **Create**.
7. Click **OK**.

Tip: If you want to use a CL command, see:
Work with Device Descriptions (WRKDEVD) (use `WRKDEVD *NWSH`)
Create Device Desc (NWSH) (CRTDEVNWSH)
Work with TCP/IP Network Sts (NETSTAT) (use `NETSTAT *IFC`)

Configure TCP/IP (CFGTCP)
Add TCP/IP Interface (ADDTCPIFC)
Work with Line Descriptions (WRKLIND)
Create Line Desc (Ethernet) (CRTLINETH) (use parameters: `MAXFRAME(8996) CMNRCYLMT(1 0)`)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Determining the hardware resource name for an iSCSI target adapter

You must determine the IBM i iSCSI target adapter hardware resource name. The resource name is used when creating a network server host adapter (NWSH), or when creating a line description (LIND) that is used with an NWSH.

Determine the IBM i hardware resource name that was assigned to the iSCSI target adapter port. Find the iSCSI target adapter port resource with physical location values that match the location of the iSCSI target adapter.

1. From the IBM i command line, run the following command to display a list of the communications resources:

   ```
   WRKHDWRSC *CMN
   ```

2. Use option **7=Display resource detail** on each iSCSI target adapter port resource until the correct one is found.

   Note: The iSCSI target adapter port resource description is:

   - **Ethernet Port** if your iSCSI target adapter is an Ethernet NIC.
   - **Network Server Host Port** if your iSCSI target adapter is an iSCSI HBA.

3. On the **Display Resource Detail** panel for the iSCSI target adapter port, examine the **Location** value to determine the frame ID, card position, and port values. If the location value corresponds to the iSCSI target adapter port for the new NWSH, record the **Resource name** value so that it is available when creating the NWSH or LIND.

   For example, if you are using the *IBM i iSCSI Solution Work Sheets* PDF on the IBM i iSCSI Solution Guide Web page, then record the **Resource name** value in one of the following work sheets:

   IBM i line description for iSCSI software target NWSH for a software target (Ethernet NIC).

   IBM i iSCSI target network server host adapter for a hardware target (iSCSI HBA).

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Creating a network server host adapter based on another one

Create a new network server host adapter (NWSH) object based on an existing object.

This saves time when some of the new NWSH attributes are the same or similar to the attributes of an existing NWSH.

To create a network server host adapter based on an existing one, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Network Server Host Adapters**.
3. Click the menu icon for the network server host adapter to copy from the list available.
4. Select **New Based On**.
5. Enter the new NWSH device **Name**.
6. Specify any other attributes that should be different from the NWSH that is being copied.
7. Click **OK**.

Tip: If you want to use a CL command, see:
Work with Device Descriptions (WRKDEVD) (use `WRKDEVD *NWSH`)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Displaying network server host adapter properties

A network server host adapter (NWSH) object contains configuration information for an IBM i iSCSI target adapter.

To display the attributes of a network server host adapter, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Network Server Host Adapters**.
3. Click the menu icon for a network server host adapter from the list available.
4. Select **Properties**.
5. Click on the appropriate tabs for the properties you want to display.
6. Click **Cancel** to close the panel.

Tip: If you want to use a CL command, see:
Work with Device Descriptions (WRKDEVD) (use `WRKDEVD *NWSH`)
Display Device Description (DSPDEVD)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Changing network server host adapter properties

A network server host adapter (NWSH) object contains configuration information for an IBM i iSCSI target adapter.

Note: Some NWSH properties cannot be changed while the NWSH is active. See the Change Device Desc (NWSH) (CHGDEVNWSH) command documentation for restrictions. If you want to change a property that cannot be changed while the NWSH is active, stop the NWSH before performing this task. See "Stopping a network server host adapter" on page 102.

To change the attributes of a network server host adapter, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Network Server Host Adapters**.
3. Click the menu icon for a network server host adapter from the list available.
4. Select **Properties**.
5. Click the appropriate tabs for the properties you want to change and change the properties.
6. Click **OK** to save any changes.

Tip: If you want to use a CL command, see:
Work with Device Descriptions (WRKDEVD) (use `WRKDEVD *NWSH`)
Change Device Desc (NWSH) (CHGDEVNWSH)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Starting a network server host adapter

Start a network server host adapter (NWSH) object to make an iSCSI target port available to an integrated server.

To start a network server host adapter, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Network Server Host Adapters**.
3. Click the menu icon for a network server host adapter from the list available.
4. Select **Start**.

If the NWSH is for a software target (Ethernet NIC), then the associated TCP/IP interface and line description (LIND) are also started.

Tip: If you want to use a CL command, see:
Work with Configuration Status (WRKCFGSTS) (use `WRKCFGSTS *DEV *NWSH`)
Vary Configuration (VRYCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Stopping a network server host adapter

Stopping (varying off) a network server host adapter (NWSH) object makes the corresponding IBM i iSCSI target adapter unavailable to any integrated servers that use it.

Stopping an NWSH that is being used by active servers can cause the servers to fail if critical storage resources can no longer be accessed. Normally, you should shut down any integrated servers that are using the NWSH before stopping the NWSH. See "Stopping integrated servers" on page 72 for more information.

To stop a network server host adapter, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Network Server Host Adapters**.
3. Click the menu icon for a network server host adapter from the list available.

4. Select **Stop**.
5. Click **Stop** on the confirmation panel.
6. If active servers are currently using the NWSH, a warning message is shown. Click **Continue**.

If the NWSH is for a software target (Ethernet NIC), then the associated TCP/IP interface is also stopped.

Tip: If you want to use a CL command, see:
Work with Configuration Status (WRKCFGSTS) (use `WRKCFGSTS *DEV *NWSH`)
Vary Configuration (VRYCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Deleting a network server host adapter

To delete a network server host adapter (NWSH), follow these steps.

Before deleting an NWSH, stop the NWSH from IBM i. See "Stopping a network server host adapter" on page 102.

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Network Server Host Adapters**.
3. Click the menu icon for a network server host adapter from the list available and select **Delete**.

   Tip: To process multiple objects, select the objects in the list. Then use **Select Action** > **Delete**.
4. Click **Delete** on the confirmation panel.
5. If the deleted NWSH represented a software target (Ethernet NIC), then the NWSH had an associated TCP/IP interface and line description (LIND). If the interface and LIND are no longer needed, delete them. For more information, see:

   "Deleting a TCP/IP interface"

   "Deleting a line description" on page 104

Tip: If you want to use a CL command, see:
Work with Device Descriptions (WRKDEVD) (use `WRKDEVD *NWSH`)
Delete Device Description (DLTDEVD)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Deleting a TCP/IP interface

You can delete a TCP/IP interface when it is no longer needed.

Shut down any integrated servers or network server host adapters (NWSHs) that use the TCP/IP interface before deleting the interface. See "Stopping integrated servers" on page 72 or "Stopping a network server host adapter" on page 102 for more information.

To delete a TCP/IP interface, follow these steps:

1. Select **Network** from *IBM Navigator for i*.
2. Click **Show All Network Tasks**.
3. Select **Network** > **TCP/IP Configuration** > **IPv4** > **Interfaces**.
4. Locate the TCP/IP interface to delete in the list.
5. If the TCP/IP interface status is not **Inactive**, click the menu icon for the TCP/IP interface and select **Stop** to stop the TCP/IP interface.
6. Click the menu icon for the TCP/IP interface and select **Delete** to delete the TCP/IP interface.

Tip: If you want to use a CL command, see:
Configure TCP/IP (CFGTCP) (use `CFGTCP`, then option 1)
End TCP/IP Interface (ENDTCPIFC)
Remove TCP/IP Interface (RMVTCPIFC)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Deleting a line description

You can delete a line description (LIND) object when it is no longer needed.

Shut down any integrated servers or network server host adapters (NWSHs) that use the line description before deleting the line description. See "Stopping integrated servers" on page 72 or "Stopping a network server host adapter" on page 102 for more information.

To delete a line description, follow these steps:

1. Type `WRKCFGSTS *LIN` from an IBM i command line.
2. If the line description status is not **VARIED OFF**, select option **2=Vary off** to vary off the line description.
3. Type `WRKLIND` from an IBM i command line.
4. Select option **4=Delete** to delete the line description.

Tip: If you want to use a CL command, see:
Work with Configuration Status (WRKCFGSTS)
Vary Configuration (VRYCFG)
Work with Line Descriptions (WRKLIND)
Delete Line Description (DLTLIND)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Managing remote system configurations

Use these tasks to manage remote system configurations for iSCSI-attached integrated servers.

Remote system network server configurations (NWSCFG subtype RMTSYS) are used to configure attributes of an iSCSI-attached remote System x or BladeCenter blade server.

The remote system configuration is used to identify the specific System x or BladeCenter hardware that the integrated server uses. It also defines how the remote system boots and communicates with IBM i.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Creating a remote system configuration

A remote system network server configuration (NWSCFG subtype RMTSYS) must be created for each System x or blade system that is used for an integrated server.

Attention: If you need to define more than one remote interface (more than one iSCSI initiator port on the BladeCenter blade or System x model), then use the *IBM Navigator for i* interface to create the remote system configuration. See the CRTNWSCFG and CHGNWSCFG Prompting Problems When defining more than one remote interface troubleshooting topic for more information.

Notes:

1. If you are using the *IBM i iSCSI Solution Work Sheets* PDF on the IBM i iSCSI Solution Guide Web page, use the following work sheet to help you do this task:

   IBM i remote system configuration
2. The network server host adapter and the remote system configuration define IP address information for opposite sides of the iSCSI network. When connected by a simple, switched network, the following rules apply:
   - The SCSI IP addresses in these two objects that are connected by a switch must be in the same subnet. For example, with IP addresses of the form a.b.x.y and 255.255.255.0 subnet masks, a.b.x must be the same value for both objects.
   - The LAN IP addresses in these two objects that are connected by a switch must be in the same subnet.
   - If you use CL commands to create the network server host adapter or remote system configuration, set the gateway elements to `*NONE`.

To create a remote system configuration, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **New Remote System Configuration**.
3. Click **Continue** on the **Select Base Object** page.
4. On the **General** tab:
   a. Enter the **Name** and **Description**.
   b. Select the **Service processor configuration**.
   c. Specify the **Remote system identity**.
   d. Optional: Select the **Object authority**. You can use the default value **Change**.
5. On the **Remote Interfaces** tab, enter information to define the SCSI and LAN interface attributes for the remote system.
6. Optional: Specify values on the **Boot Parameters** and **CHAP Authentication** tabs if wanted.
7. Click **OK**.

Tip: If you want to use a CL command, see:
Work with NWS Configuration (WRKNWSCFG)
Create NWS Configuration (CRTNWSCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Creating a remote system configuration based on another one

Create a remote system configuration based on an existing object.

You can copy an existing remote system network server configuration (NWSCFG subtype RMTSYS) object when creating a new one. This saves time when some of the new remote system configuration attributes are the same or similar to the attributes of an existing remote system configuration.

To create a remote system configuration based on an existing one, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Remote Systems**.
3. Click the menu icon for the remote system configuration to copy from the list available.
4. Select **New Based On**.
5. Enter the new remote system configuration **Name**.
6. Specify any other attributes that should be different from the remote system configuration that is being copied.
7. Click **OK**.

Note: There is no equivalent CL command for this task.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Displaying remote system configuration properties

A remote system network server configuration (NWSCFG subtype RMTSYS) object contains configuration information for an System x or BladeCenter system that will be used to run an iSCSI-attached integrated server.

To display the attributes of a remote system configuration, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Remote Systems**.
3. Click the menu icon for a remote system configuration from the list available.
4. Select **Properties**.
5. Click on the appropriate tabs for the properties you want to display.
6. Click **OK** to close the panel.

Tip: If you want to use a CL command, see:
Work with NWS Configuration (WRKNWSCFG)
Display NWS Configuration (DSPNWSCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Changing remote system configuration properties

A remote system network server configuration (NWSCFG subtype RMTSYS) object contains configuration information for a System x or BladeCenter system that is used for an integrated server.

Note: Some remote system object properties cannot be changed while the remote system object is in use by an active server. See the Change NWS Configuration (CHGNWSCFG) command documentation for restrictions. If you want to change a property that cannot be changed while the associated server is active, stop the server before performing this task. See "Stopping integrated servers" on page 72.

To change the attributes of a remote system configuration, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Remote Systems**.
3. Select a remote system configuration from the list available.
4. Select **Properties**.
5. Click the appropriate tabs for the properties you want to change.
6. Click **OK** to save any changes.

Tip: If you want to use a CL command, see:
Work with NWS Configuration (WRKNWSCFG)
Change NWS Configuration (CHGNWSCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Displaying remote system status

Do these steps to display the status for the System x or BladeCenter hardware for iSCSI-attached integrated servers.

You can use the status to help you determine if hardware is available for use by an iSCSI-attached integrated server.

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Remote Systems**.
3. Click the menu icon for a remote system configuration from the list available.
4. Select **Status**.
5. The status of the remote system hardware is shown.
6. Click **Cancel** to close the panel.

Tip: If you want to use a CL command, see:
Work with NWS Configuration (WRKNWSCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Deleting a remote system configuration

Do these steps to delete remote system configurations for integrated servers.

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Remote Systems**.
3. Click the menu icon for a remote system configuration from the list available and select **Delete**.

   Tip: To process multiple objects, select the objects in the list. Then use **Select Action** > **Delete**.
4. Click **Delete** on the confirmation panel.

Tip: If you want to use a CL command, see:
 Work with NWS Configuration (WRKNWSCFG)
 Delete NWS Configuration (DLTNWSCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Launching the Web console for a remote system

Do these steps to launch the integrated server service processor Web console that is associated with an IBM i remote system configuration.

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Remote Systems**.
3. Click the menu icon for a remote system configuration from the list available and select **Launch Web console**.
4. Click the **Web Console** link on the **Launch Web Console** page.

The Web console for the integrated server service processor is shown in a separate Web browser window. For example, if the remote server configuration represents a blade in an IBM BladeCenter that has an Advanced Management Module (AMM) service processor, then the AMM Web interface is shown.

Note: If the service processor Web console page does not appear, review the notes on the **Launch Web Console** page for possible reasons.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Managing service processor configurations

Use these tasks to manage service processor configurations for integrated servers.

Service processor network server configurations (NWSCFG subtype SRVPRC) are used to configure attributes of the service processor or Management Module of each iSCSI-attached remote System x or BladeCenter.

The service processor configuration defines attributes that are used to connect to the service processor or Management Module on the network. Remote system configurations contain a reference to the corresponding service processor configuration that is used to control the remote system hardware.

Note: A service processor configuration is not needed for each IBM BladeCenter blade server in a BladeCenter chassis. Just one service processor configuration is needed for the IBM BladeCenter chassis.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Creating a service processor configuration

A service processor network server configuration (NWSCFG subtype SRVPRC) must be created for the service processor or Management Module of each System x or BladeCenter system that is used for an integrated server.

Notes:

1. If you are using the *IBM i iSCSI Solution Work Sheets* PDF on the IBM i iSCSI Solution Guide Web page, use the following work sheet to help you do this task:

   IBM i service processor configuration
2. A service processor configuration is not needed for each blade in an IBM BladeCenter chassis. Just one service processor configuration is needed for the BladeCenter chassis.

To create a service processor configuration, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **New Service Processor Configuration**.
3. Click **Continue** on the **Select Base Object** page.
4. Enter the **Name** and **Description**
5. Specify either a **Host name** or **Internet address** to identify the service processor on the network.
6. Optional: Specify the System x or BladeCenter **Serial number** and **Manufacturer type and model**. If you leave these values blank, they are automatically retrieved when the service processor configuration is initialized.
7. Optional: Select the **Object authority**. You can use the default value **Change**.
8. Click **OK**.

Tip: If you want to use a CL command, see:
Work with NWS Configuration (WRKNWSCFG)
Create NWS Configuration (CRTNWSCFG)

A service processor configuration must be initialized with a user name and password before it can be used with an integrated server. See "Initializing a service processor" on page 111.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Creating a service processor configuration based on another one

You can copy an existing service processor network server configuration (NWSCFG subtype SRVPRC) when creating a new one. This saves time when some of the new service processor configuration attributes are the same or similar to the attributes of an existing service processor configuration.

To create a service processor configuration based on an existing one, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Service Processors**.
3. Click the menu icon for the service processor configuration to copy from the list available and select **New Based On**.
4. Enter the new service processor configuration **Name**.
5. Specify any other attributes that should be different from the service processor configuration that is being copied.
6. Click **OK**.

Note: There is no equivalent CL command for this task.

A service processor configuration must be initialized with a user name and password before it can be used with an integrated server. See "Initializing a service processor" on page 111.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Displaying service processor configuration properties

A service processor network server configuration (NWSCFG subtype SRVPRC) object contains configuration information for a service processor or Management Module of an System x or BladeCenter system that is used to run an iSCSI-attached integrated server.

To change the attributes of a service processor configuration, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Service Processors**.
3. Click the menu icon for a service processor configuration from the list available.
4. Select **Properties**.
5. Click on the appropriate tabs for the properties you want to display.
6. Click **OK** to close the panel.

Tip: If you want to use a CL command, see:
Work with NWS Configuration (WRKNWSCFG)
Display NWS Configuration (DSPNWSCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Changing service processor configuration properties

A service processor network server configuration (NWSCFG subtype SRVPRC) object contains configuration information for a service processor of a System x or BladeCenter system that is used to run an integrated server.

Note: Some service processor object properties cannot be changed while the service processor object is in use by an active server. See the Change NWS Configuration (CHGNWSCFG) command documentation for restrictions. If you want to change a property that cannot be changed while the associated server is active, stop the server before performing this task. See "Stopping integrated servers" on page 72.

To change the attributes of a service processor configuration, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Service Processors**.
3. Click the menu icon for a service processor from the list available.
4. Select **Properties**.
5. Click the appropriate tabs for the properties you want to change.
6. Click **OK** to save any changes.

Tip: If you want to use a CL command, see:
Work with NWS Configuration (WRKNWSCFG)
Change NWS Configuration (CHGNWSCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Initializing a service processor

A service processor configuration (NWSCFG subtype SRVPRC) must be initialized with a user name and password before it can be used with an integrated server.

Note: The service processor configuration cannot be initialized while it is in use by an active server. If an associated server is active, stop the server before performing this task. See "Stopping integrated servers" on page 72.

A service processor configuration contains configuration information for the service processor of a System x or BladeCenter system that is used for an integrated server. The service processor needs to be initialized before it can be used with an integrated server. You might also want to synchronize the user and password that are used to secure the service processor connection or change the user or password that are used to connect to the service processor.

To initialize a service processor, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Service Processors**.
3. Click the menu icon for a service processor configuration from the list available.
4. Select **Initialize**.
5. Choose one of the following options:

- **Validate service processor user ID and password and store in {NWSCFG NAME}**

  Tip: If you are initializing a configuration object of a service processor for the first time, use this option.
- **Validate and set user ID and password in a new service processor**
- **Change service processor user ID and password**

6. Enter the **User** and **Password**, if needed.
7. Click **Initialize** to perform the selected option.

Tip: If you want to use a CL command, see:
Work with NWS Configuration (WRKNWSCFG)
Initialize NWS Configuration (INZNWSCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Deleting a service processor configuration

To delete a service processor configuration, follow these steps.

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Service Processors**.
3. Click the menu icon for a service processor configuration from the list available and select **Delete**.

   Tip: To process multiple objects, select the objects in the list. Then use **Select Action** > **Delete**.
4. Click **Delete** on the confirmation panel.

Tip: If you want to use a CL command, see:
Work with NWS Configuration (WRKNWSCFG)
Delete NWS Configuration (DLTNWSCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Launching the Web console for a service processor

Do these steps to launch the integrated server service processor Web console that is associated with an IBM i service processor configuration.

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Select **Service Processors**.
3. Click the menu icon for a service processor configuration from the list available and select **Launch Web console**.
4. Click the **Web Console** link on the **Launch Web Console** page.

The Web console for the integrated server service processor is shown in a separate Web browser window. For example, if the service processor configuration represents an IBM BladeCenter Advanced Management Module (AMM), then the AMM Web interface is shown.

Note: If the service processor Web console page does not appear, review the notes on the **Launch Web Console** page for possible reasons.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Converting pre-IBM i 7.1 service processor configurations to use an IP address or host name

Specify either the service processor IP address or host name in the service processor configuration properties.

Before IBM i 7.1, multicast discovery using Service Location Protocol (SLP) could be used to discover the System x or BladeCenter service processor. As of IBM i 7.1, multicast discovery is no longer supported. A specific IP address or host name must be used to connect to the service processor.

If you upgrade your system to IBM i 7.1 or later and you have a service processor configuration that does not use an IP address or host name, you must change the properties of the service processor configuration to specify either the service processor IP address or host name. See "Changing service processor configuration properties" on page 111 for details.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Managing connection security configurations

Connection security network server configurations (NWSCFG subtype CNNSEC) are used by IBM i to connect to the integrated server hardware.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Creating a connection security configuration

Do these steps to create a connection security configuration for an integrated server.

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Click the **Show All Integrated Server Administration Tasks** button.
3. Select **Integrated Server Administration** > **iSCSI Connections** > **Connection Security** > **New Connection Security Configuration**.
4. Click **Continue** on the **Select Base Object** page.
5. Enter the **Name** and **Description**.
6. Optional: Select the **Object authority**. You can use the default value **Change**.
7. Click **OK**.

Tip: If you want to use a CL command, see:
Work with NWS Configuration (WRKNWSCFG)
Create NWS Configuration (CRTNWSCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Creating a connection security configuration based on another one

You can copy an existing connection security network server configuration (NWSCFG subtype CNNSEC) when creating a new one. This saves time when some of the new connection security configuration attributes are the same or similar to the attributes of an existing connection security configuration.

To create a connection security configuration based on an existing one, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Click the **Show All Integrated Server Administration Tasks** button.
3. Select **Integrated Server Administration** > **iSCSI Connections** > **Connection Security** > **Connection Security**.
4. Click the menu icon for the connection security configuration to copy from the list available.
5. Select **New Based On**.
6. Enter the new connection security configuration **Name**.
7. Specify any other attributes that should be different from the connection security configuration that is being copied.
8. Click **OK**.

Note: There is no equivalent CL command for this task.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Displaying connection security configuration properties

Do these steps to display the properties of a connection security configuration for an iSCSI-attached integrated server.

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Click the **Show All Integrated Server Administration Tasks** button.
3. Select **Integrated Server Administration** > **iSCSI Connections** > **Connection Security** > **Connection Security**.
4. Click the menu icon for a connection security configuration from the list available.
5. Select **Properties**.
6. Click **OK** to close the panel.

Tip: If you want to use a CL command, see:
Work with NWS Configuration (WRKNWSCFG)
Display NWS Configuration (DSPNWSCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software,

and virtual storage.

Changing connection security configuration properties

Do these steps to change the properties of a connection security configuration for an integrated server.

Note: Some connection security object properties cannot be changed while the connection security object is in use by an active server. See the Change NWS Configuration (CHGNWSCFG) command documentation for restrictions. If you want to change a property that cannot be changed while the associated server is active, stop the server before performing this task. See "Stopping integrated servers" on page 72.

To change the attributes of a connection security configuration, follow these steps:

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Click the **Show All Integrated Server Administration Tasks** button.
3. Select **Integrated Server Administration** > **iSCSI Connections** > **Connection Security** > **Connection Security**.
4. Click the menu icon for a connection security configuration object from the list available.
5. Select **Properties**.
6. Click **OK** to save any changes.

Tip: If you want to use a CL command, see:
Work with NWS Configuration (WRKNWSCFG)
Change NWS Configuration (CHGNWSCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Deleting a connection security configuration

Do these steps to delete a connection security configuration for an integrated server.

1. Select **Integrated Server Administration** from *IBM Navigator for i*.
2. Click the **Show All Integrated Server Administration Tasks** button.
3. Select **Integrated Server Administration** > **iSCSI Connections** > **Connection Security** > **Connection Security**.
4. Click the menu icon for a connection security configuration from the list available and select **Delete**.
5. Click **Delete** on the confirmation panel.

Tip: If you want to use a CL command, see:
Work with NWS Configuration (WRKNWSCFG)
Delete NWS Configuration (DLTNWSCFG)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Chapter 5. Backing up and recovering integrated servers

You can back up and recover integrated server data from either IBM i or the integrated server operating system.

You can use either IBM i or native integrated server utilities or a combination of both to manage backups. When you are planning your backup strategy, refer to the Backup and recovery topic, as well as Microsoft or VMware documentation.

To back up an integrated server on IBM i, you have these basic options:

- Do a full system backup on your IBM i system. See the topic Backing up your system.
- Back up the network server description (NWSD), virtual storage, and other objects that are associated with the integrated server on IBM i. See "Backing up the NWSD and other objects associated with integrated servers."
- Back up individual integrated Windows server files by using the IBM i Save (SAV) command and IBM i NetServer. See "Backing up individual integrated Windows server files and directories" on page 123.
- Back up individual integrated server files by using native integrated server operating system utilities, such as the Windows Server 2003 Backup utility.

Your recovery options depend on how you backed up your system, as well as what you need to recover.

- If you need to recover your entire system, see the Backup and recovery topic collection.
- If you need to restore a network server description and its associated IBM i virtual storage, or other IBM i objects, refer to "Restoring the NWSD and other objects associated with integrated servers" on page 130.
- To restore integrated server data (files, directories, shares, and the Windows registry) that you backed up with the Save (SAV) command, see "Restoring individual integrated Windows server files and directories" on page 134.
- If you used a program such as the Windows Server 2003 Backup utility or tar to save your files, use that program to restore the files.

Use these tasks to back up and recover integrated servers.

Backing up the NWSD and other objects associated with integrated servers

Do these tasks to back up the IBM i configuration objects and files related to integrated servers.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

Backing up your system

What objects to save and their location on IBM i

Use this table to determine which objects need to be saved when you save your integrated server.

Many objects are created as a result of installing integrated servers. Some of these objects are system-related, others user-related. You need to save them all if you want to restore properly. You can save these objects by using options of the IBM i GO SAVE command. Option 21 saves the entire system. Option 22 saves system data. Option 23 saves all user data (which includes storage spaces in QFPNWSSTG).

If you want to save a particular object, use the following table to see the location of that object on IBM i and the command to use. The topic Manually saving parts of your system has more information about using the save commands. In addition to saving the entire drive (storage space), you can also save and restore individual files and directories.

Important: Ensure that the auxiliary storage pool (ASP) is available when you save the data.

Object content	Object name	Object location	Object type	Save command
Virtual disks	Various	/QFPNWSSTG	Network server storage space	`GO SAVE`, option 21 or 23 `SAV OBJ('/QFPNWSSTG/stgspc') DEV('/QSYS.LIB/`*`TAP01`*`.DEVD')`
Messages from the integrated server	Various	Various	Message queue	`GO SAVE`, option 21 or 23 `SAVOBJ OBJ(msgq) LIB(qlibrary) DEV(`*`TAP01`*`) OBJTYPE(*MSGQ)`
IBM i config objects for integrated servers	Various	QSYS	Device config objects	`GO SAVE`, option 21, 22, or 23 `SAVCFG DEV(`*`TAP01`*`)`
IBM i Integrated Server Support code	QNTAP, NTAP and subdirectories	QSYS and /QIBM/ProdData/NTAP	Library and Directory	`SAVLICPGM LICPGM(5770SS1) OPTION(29)`
Windows server file shares	QNTC and subdirectories	/QNTC/servername/sharename	Directory	`GO SAVE`, option 21 or 22 `SAV`
IBM i TCP interfaces	QATOCIFC	QUSRSYS	physical file **Note:** TCP/IP must be ended when you save the TCP interface physical files.	`GO SAVE`, option 21 or 23 `SAVOBJ OBJ(QATOCIFC) LIB(QUSRSYS) DEV(TAP01) OBJTYPE(*FILE)`

Object content	Object name	Object location	Object type	Save command
IBM i TCP interfaces	QATOCLIFC	QUSRSYS	logical file **Note:** TCP/IP must be ended when you save the TCP interface physical files.	`GO SAVE`, option 21 or 23 `SAVOBJ OBJ(QATOCLIFC) LIB(QUSRSYS) DEV(`*`TAP01`*`) OBJTYPE(*FILE)`
iSCSI NWSCFG and associated validation list	Various	QUSRSYS	Network Server Configuration and associated values	`SAVOBJ LIB(QUSRSYS) OBJTYPE(*NWSCFG *VLDL)`
iSCSI path certificate store	nwsdname.*	/QIBM/UserData/ NWSDCert	Certificate store file	`GO SAVE`, option 21 or 23 `SAV OBJ('/QIBM/UserData/NWSDCert/ nwsdname.*')`
User profile enrollment information	Various	QSYS	User profile	`GO SAVE`, option 8, 21, or 23 `SAVSECDTA DEV(`*`TAP01`*`)`

Tip: If you want to see the CL command documentation, see:
Save (SAV)
Save Object (SAVOBJ)
Save Configuration (SAVCFG)
Save Licensed Program (SAVLICPGM)
Save Security Data (SAVSECDTA)

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

"User and group enrollment concepts for integrated Windows servers" on page 56
Learn about how IBM i users and groups interact with integrated Windows servers.

Related tasks:

"Backing up the NWSD and other objects associated with integrated servers" on page 117
Do these tasks to back up the IBM i configuration objects and files related to integrated servers.

"Restoring the NWSD and other objects associated with integrated servers" on page 130
Do these tasks to restore the IBM i configuration objects and files related to integrated servers.

"Saving user enrollment information for integrated Windows servers" on page 129
Use CL commands and APIs to save user profiles and enrollment information for an integrated Windows server.

"Restoring user enrollment information for integrated Windows servers" on page 135
Use CL commands and APIs to save and restore user profiles and enrollment information for an integrated Windows server.

Backing up the NWSD of an integrated server

Do these steps to save an NWSD with the Save Configuration (SAVCFG) command.

Note: when you save the associated storage space objects, you also need to save the Network Server Description (NWSD). To save an NWSD, you use the Save Configuration (SAVCFG) command:

1. On the IBM i command line, type `SAVCFG`.
2. Press Enter to have IBM i save the NWSD configuration.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

Backing up virtual storage for integrated servers

Use the Save (SAV) command to back up virtual storage for your integrated server.

You need to save both predefined virtual storage that IBM i creates, and user-defined virtual storage.

The virtual storage spaces for an integrated server are in the integrated file system. To save the virtual storage from IBM i, you use the Save (SAV) command.

Note: Treat the network server description, predefined virtual storage, and any user-defined virtual storage linked to an integrated server as a unit. Save and restore them at the same time. Together they constitute a complete system, and should be treated as such. Otherwise, the integrated server might not start or run correctly.

Note: You can use the same steps for backing up predefined virtual storage (the system drive and the installation drive) and user defined virtual storage.

Do these steps to back up virtual storage:

1. Ensure that the auxiliary storage pool (ASP) that contains the virtual storage is varied on.
2. If you are saving to tape, ensure that you have mounted a tape that is formatted for IBM i.
3. Select one of the following options.

Option	Description
Save virtual storage for an active Windows server.	See "Backing up virtual storage for active Windows servers" on page 121.
Shut down the integrated server to prevent users from updating files during the backup.	See "Stopping integrated servers" on page 72.

4. On the IBM i command line, type `SAV` and press F4.

5. If you are saving the storage space to tape, specify the name of your tape device . For example, specify `/QSYS.LIB/TAP01.DEVD` in the *Device* field.
6. If you are saving the storage space to a save file instead of to tape, specify the path to the save file as the device.

   For example, to use a save file named MYSAVF in library WINBACKUP, you would specify `'/QSYS.LIB/WINBACKUP.LIB/MYSAVF.FILE'` for the device.
7. In the `Name` field under `Objects:`, specify `'/QFPNWSSTG/stgspc'`, where `stgspc` is the name of the network server storage space.

   For example, if the NWSD for the integrated server is named *testserver*, you can save the system and installation drives by saving these network server storage spaces:

   - /QFPNWSSTG/*testserver*1
   - /QFPNWSSTG/*testserver*2
8. If you are saving virtual storage for an active server, specify the following values:
   a. Specify `*YES` for the **Save active** parameter. This option allows the storage space to be saved while it is still being used the system.
   b. Specify `*NWSSTG` for the **Save active option** parameter. This option allows network server storage spaces in directory '/QFPNWSSTG' to be saved when they are active.
9. Specify values for any other parameters that you want and press `Enter` to save the storage space.
10. If you stopped the integrated server, restart it now. See "Starting integrated servers" on page 70.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

Backing up virtual storage for active Windows servers

Use the FREEZE.BAT and THAW.BAT scripts to configure backup for active Windows servers.

The disks that you create for integrated Windows servers are stored in the integrated file system. To save these storage spaces from IBM i, you use the Save (SAV) command.

IBM i saves the changes that are made to the storage space during a save operation. This information is stored in a temporary file that can be up to 25% of the total size of the storage space. This default setting should work for most configurations.

Use the freeze and thaw scripts if you receive a message that too much space is being used by the process that tracks changes. You can also use the scripts if you know that applications on the Windows server will make frequent read and write requests to the storage space during the backup.

- The FREEZE.BAT script runs when IBM i starts to back up a storage space. Use this script to stop applications that might fill the temporary storage space.
- The THAW.BAT script runs when IBM i finishes backing up a storage space. Use this script to start any applications that you stopped with the FREEZE.BAT script.

Do the following steps to customize storage space backup.

1. Run these scripts when you start and finish backing up the storage space. You can modify them for your environment.
   a. %SYSTEMROOT%\AS400WSV\ADMIN\FREEZE.BAT
   b. %SYSTEMROOT%\AS400WSV\ADMIN\THAW.BAT
2. Edit the scripts.
3. Use the Save (SAV) command to save the storage space. For more information, see "Backing up virtual storage for integrated servers" on page 120.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

Backing up NWSH objects and associated LIND objects and interfaces

Use the Save Configuration (SAVCFG) command to back up a network server host adapter (NWSH) object and associated line description (LIND) object. Use the Save Object (SAVOBJ) command to back up the associated TCP/IP interface.

For hardware targets, just the NWSH needs to be backed up. For software targets, the NWSH and the associated LIND and TCP/IP interface need to be backed up.

Use SAVCFG to back up a hardware or software target NWSH and the LIND for a software target NWSH:

1. On the IBM i command line, type `SAVCFG`.
2. Press **Enter** to have IBM i save the NWSH and LIND configurations (and all other configuration objects on the system).

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

Backing up the TCP/IP interface for a software target NWSH

Note that TCP/IP must be ended on IBM i before you save the TCP/IP interface files.

Use SAVOBJ to back up the TCP/IP interface for a software target NWSH:

1. On the IBM i command line, type `SAVOBJ LIB(QUSRSYS) OBJ(QATOCIFC QATOCLIFC) DEV(TAP01) OBJTYPE(*FILE)`. Replace `TAP01` with your tape device.
2. Press **Enter** to have IBM i save the TCP/IP interface used by the NWSH (and all other TCP/IP interfaces on the system).

Backing up NWSCFG objects and validation lists

Network server configuration objects are stored in the QUSRSYS library. These include the network server configuration objects (type *NWSCFG) and an associated validation list object (type *VLDL).

Note: The *NWSCFG and *VLDL objects will share the same name.

To save the network server configuration and validation list objects, use the Save Object (SAVOBJ) command:

1. If you are saving to tape, ensure that you have mounted a tape that is formatted for IBM i.
2. Shut down the Windows server to release any object locks.
3. On the IBM i command line, type `SAVOBJ` and press `F4`.
4. In the **Objects** field, specify the NWSCFG names.
5. In the **Library** field, specify `QUSRSYS`.
6. If you are saving the objects to tape, specify the name of your tape device in the **Device** field (for example, TAP01). If you want to use a save file instead of tape, specify `*SAVF` as the device and enable the data compression option.
7. For **Object type**, specify both *NWSCFG and *VLDL.
8. If you are using a save file, press `F10` to see additional parameters.
9. In the **Save file** field, specify the path to your save file (for example winbackup/nwscfg).
10. If you are using a save file, page down change the value for Data compression to `*YES`.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

Backing up individual integrated Windows server files and directories

The Integrated Server Support option allows you to save integrated server data (files, directories, shares, and the Windows registry) to tape, optical or disk (*SAVF) along with other IBM i data and restore the data on an individual basis.

IBM i Integrated Server Support allows you to save integrated server data (files, directories, shares, and the Windows registry) to tape, optical or disk (*SAVF) along with other IBM i data and restore the data on an individual basis. However, you should not use this approach as your primary backup procedure. You should still periodically save your entire system and the NWSD associated with your Windows server for disaster recovery. Then you can choose to do daily backups of only the integrated server files that have changed. See "Backing up the NWSD and other objects associated with integrated servers" on page 117.

You can also use a utility such as the Backup program that comes with Windows.

Related reference:

"Examples: Saving parts of integrated Windows servers" on page 128
These examples show how to use the save (SAV) or restore (RST) commands for specific parts of an integrated server.

Backing up your system

File-level backup restrictions for integrated Windows servers

File-level backup for integrated Windows servers has some limitations and requirements for the environment.

Limitations

- This support is not available to stand-alone Windows servers because the code comes packaged with IBM i Integrated Server Support.
- This method does not back up files that are part of the IBM i Integrated Server Support code.
- You cannot stop users from signing on and accessing data on the server while the Save (SAV) or Restore (RST) command is running. IBM i Integrated Server Support can save a file that is in use as long as it can read the file. Consequently, you should back up integrated server files when you expect few users to be accessing the system. A note telling users to avoid accessing the server would be a good precaution.
- Windows Server 2003 provides function with its Volume Shadow copy Service (VSS). This allows applications that are backup aware the ability to save files while they are still in use when using file-level backup
- The QSECOFR user profile should not be used to perform a file-level backup. Even if enrolled to the integrated server, QSECOFR will not be used to back up the files. The Windows Local System Account will be used instead. It may not have the necessary authority to back up all of the requested files.
- If the user profile *LCLPWDMGT value is *YES, then the system value, QRETSVRSEC, must be set to 1 and the user password must be changed or the user have signed-on after QRETSVRSEC was changed.
- If the user profile *LCLPWDMGT value is *NO, then network authentication (kerberos) is used. The user must access the IBM i operation through an EIM enabled application. See *SBMNWSCMD and file level backup support for Kerberos V5 and EIM* in the IBM i iSCSI Solution Guide for more information.

Requirements

- The integrated server must be active and have a working TCP/IP point-to-point virtual Ethernet connection with IBM i. You must back up your integrated server files either before putting the system into restricted state to back up the rest of the IBM i files or after completing restricted state operations.
- This procedure requires that you have the same user ID and password on the integrated server and IBM i.
- Your integrated server user account must be a member of the Administrators group.
- File-level backup uses the QNTC file system (NetClient) to build the list of files to be saved. QNTC uses IBM i NetServer to locate servers in the domain. You need to have the NetServer in the same domain (see "Verifying that IBM i NetServer and the integrated Windows server are in same domain" on page 127) as the integrated server from which you are going to save files.

- Be careful about trying to restore all files on all drives that you previously saved through the QNTC file system. Certain Windows system files (for example, files in the Recycle Bin) can cause unexpected results after you restore them.
- On Windows Server 2003, you need to give special consideration to System File Protection when you are backing up and recovering Windows system files. Refer to Microsoft documentation.

Installing and configuring IBM i NetServer

IBM i NetServer is used for file-level backup from IBM i. NetServer is also used for some administration tasks, such as updating the Integrated Server Support software that runs on the Windows operating system.

If NetServer is not set up on your system yet, you can set it up using the Getting started with NetServer topic.

To install updates to the IBM i integrated server support software on the Windows operating system, you must be signed on with a Windows account that has local administrator authority. That Windows account must correspond to an IBM i user profile with the same password. Alternatively, you must have a guest NetServer user profile configured.

Creating a Windows user with authorities to access IBM i NetServer

Before you can apply updates to the Integrated Server Support software that runs on the integrated Windows server, you must be signed on with a Windows account that has the authorities that are required to access NetServer.

You can set up the Windows account for the user who performs Integrated Server Support tasks in one of two ways:

- The user signs onto Windows with an account that has a corresponding IBM i user profile with the same password. This Windows account must also be a member of **Windows Administrators** group. You can enroll the user to Windows after the server has been installed. See "Enrolling IBM i users to integrated Windows servers" on page 88.
- The user signs onto Windows with any account and a **guest** user profile is configured for NetServer. If a **guest** user profile is not configured for NetServer, you must set up the **guest** user profile before performing Integrated Server Support tasks that might need to use it. See "Creating a guest user profile for NetServer" for more information.

Creating a guest user profile for NetServer:

If a **guest** user profile is not configured for NetServer, you must set up the **guest** user profile before performing Integrated Server Support tasks that might need to use it.

You must have IBM i *SECADM special authority to perform this task.

You can use the *IBM Navigator for i* graphical interface to set up a guest user profile for NetServer with no special authorities and no password.

Alternatively, follow these steps to set up a guest user profile for NetServer:

1. On IBM i, create a user profile with no special authorities and no password:
   ```
   CRTUSRPRF USRPRF(username) PASSWORD(*NONE) SPCAUT(*NONE)
   ```

Note: See the Security topic collection for information about user profiles.

2. Enter the following command, where *username* is the name of the user profile that you created:

   ```
   CALL QZLSCHSG PARM(username X'00000000')
   ```

3. To stop NetServer, enter the following command:

   ```
   ENDTCPSVR SERVER(*NETSVR)
   ```

4. To restart NetServer, enter the following command:

   ```
   STRTCPSVR SERVER(*NETSVR)
   ```

Configuring integrated Windows servers for file-level backup

Do these steps to configure file-level backup for integrated Windows servers.

1. Ensure that the person who is saving and restoring files has the same password on IBM i and the integrated server. The easiest method is found at "Enrolling IBM i users to integrated Windows servers" on page 88. Also ensure that the user is a member of the Administrators group. Refer to "Creating user enrollment templates for integrated Windows servers" on page 91.
2. Create shares for each drive or volume that you want to save when you request to save all the files on a Windows server. IBM i Integrated Server Support accesses the file system and translates these shares into path-names. See "Creating shares on integrated Windows servers."
3. Add members to the QAZLCSAVL file in QUSRSYS that lists the share names that you want to be able to save. See "Adding members to the QAZLCSAVL file" on page 127.
4. Ensure that IBM i NetServer is in the same domain as the integrated server for which you want to save files. See "Verifying that IBM i NetServer and the integrated Windows server are in same domain" on page 127.
5. Ensure that the person performing the saves or restores has *ALLOBJ authority which gives the user full access to the programs and devices required for the save or restore process. If *ALLOBJ authority cannot be provided, the user must have at least *USE authority on object QNTAP/QVNASBM so the backup or restore request can be communicated to the integrated Windows server.

Creating shares on integrated Windows servers

Create a file share for each file or directory that you want to save at the integrated server console. IBM i will use this share to back up the Windows files.

To create shares on integrated Windows servers, do this from the integrated server console:

1. Open the **My Computer** icon to open **Windows Explorer**.
2. Right-click the drive or volume that you want.
3. From the pop-up menu, select **Sharing**.
4. Click **Share this folder**. Provide a **Share Name** (characters in the share name must be in the more restrictive code page 500 character set). The default share name is the same as the last part of the directory name. Share names can be no longer than 12 characters and can include embedded blanks.
5. You can choose unlimited access or limit the number of users who can access the share at one time. You can also use the **Permissions** button to set up the level at which you want to share (No Access, Read, Change, or Full Control).
6. Click on **Apply** to create the share.

Adding members to the QAZLCSAVL file

To enable file-level backup and recovery from IBM i, add a member for each integrated Windows server to the QAZLCSAVL file in QUSRSYS.

For the member name, use the NWSD name of the server (*nwsdname*).

To add a member, do this:

1. On the IBM i command line, use the Add Physical File Member (ADDPFM) command to add a file member. Type

   ```
   ADDPFM FILE(QUSRSYS/QAZLCSAVL) MBR(nwsdname)
    TEXT('description')  EXPDATE(*NONE) SHARE(*NO) SRCTYPE(*NONE).
   ```

2. In the file member that you just created, list all the shares that you want to be able to save. List each share name that you defined for the server on a separate line. The maximum length that the Windows share name can be is 12 characters. Share names can have embedded blanks. For example, if you defined cshare, dshare, eshare, fshare, gshare, and my share as shares on WINSVR1, your member name WINSVR1 would look like this:

   ```
   QUSRSYS/QAZLCSAVL
   WINSVR1
   0001.00  cshare
   0002.00  dshare
   0003.00  eshare
   0004.00  fshare
   0005.00  gshare
   0006.00  my share
   ```

   Note: If you specify multiple share names that point to the same integrated server directory, IBM i saves the data multiple times for a "save all" request. To avoid duplicating data when you save it, do not include multiple shares that include the same directory or data.

Verifying that IBM i NetServer and the integrated Windows server are in same domain

To save integrated server files for file-level backup, IBM i NetServer must be configured in the same domain as the files you want to save.

1. Check the domain for your integrated server:
   a. Select **Integrated Server Administration** from *IBM Navigator for i*.
   b. Select **Servers**.
   c. Find your integrated server in the list; then look in the Domain column to find the domain for that server.
2. Check the domain for NetServer:
   a. Select **Network** from *IBM Navigator for i*.
   b. Select **TCP/IP Servers**.
   c. Click the menu icon for IBM i NetServer in the list of TCP/IP servers, then select **Properties**. The domain name for NetServer appears on the **General** tab.
3. If NetServer is not in the same domain as the integrated server, change the domain for NetServer:
   a. Click the **Next Start** button on the **General** tab.
   b. In the **Domain name** field, type the name of the Windows server domain.
   c. Click **OK** to save the new domain name.
   d. Click **OK** to close the properties page.

e. Stop and start NetServer (click the menu icon for **IBM i NetServer** and select **Stop**, then **Start**).

Saving integrated Windows server files

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

After you finish the necessary preliminaries (see "Configuring integrated Windows servers for file-level backup" on page 126), you are ready to back up integrated Windows server files on IBM i. To be able to restore a directory or file by share name, you must specify that file or share name specifically on the Save (SAV) command.

Note: To avoid duplicating data, be careful specifying what you want to save on the SAV command. If you specify multiple share names that point to the same directory on the integrated server, IBM i saves the data multiple times.

To specify what you want IBM i to save, do this:

1. Ensure that the integrated Windows server is active (described in "Starting integrated servers" on page 70). Also ensure that the QSYSWRK subsystem, QSERVER, and TCP/IP are active (you can do this by using the Work with Active Jobs (WRKACTJOB) command).
2. On the IBM i command line, type `SAV` and press F4.
3. In the `Device` field, specify the device on which you want IBM i to save the data. For example, `'QSYS.LIB/TAP01.DEVD'` saves the data to tape.
4. In the `Object` field, specify what you want IBM i to save in the form '/QNTC/*servername*/sharename' You can use wildcard characters. Refer to "Examples: Saving parts of integrated Windows servers" for how to specify particular parts of the integrated server.
5. Use the `Directory subtree` field to specify whether you want to save subtrees under a directory. The default is to save all directories.
6. To specify that you want to save changes since the last save, specify *LASTSAVE in the `Change period` field. You can also specify a specific range of dates and times.
7. Press `Enter` to save the shares that you specified.

Examples: Saving parts of integrated Windows servers

These examples show how to use the save (SAV) or restore (RST) commands for specific parts of an integrated server.

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

To save or restore this:	**Specify this:**
All integrated server objects.	`OBJ('/QNTC/*') SUBTREE(*ALL)`
All objects for *server1*.	`OBJ('/QNTC/server1/*') SUBTREE(*ALL)`
All objects for *server1* that changed since you last saved the files.	`OBJ('/QNTC/server1/*') SUBTREE(*ALL) CHGPERIOD(*LASTSAVE)`
All objects for *server1* that changed during a certain period (in this case between 10/19/99 and 10/25/99).	`OBJ('/QNTC/server1/*') SUBTREE(*ALL) CHGPERIOD('10/19/99' '00:00:00' '10/25/99' '23:59:59')`

To save or restore this:	Specify this:
All directories, files, and shares to which a particular share (for example, 'fshare') refers. IBM i does not save and restore the directory over which the share is built.	`OBJ('/QNTC/server1/fshare/*') SUBTREE(*ALL)`
Only files to which the specified share (for example, 'fshare') refers that match the specified pattern (pay*). IBM i does not save directories nor shares.	`OBJ('/QNTC/server1/fshare/pay*')`
Only directories and shares (no objects) for 'fshare' and its immediate children.	`OBJ('/QNTC/server1/fshare') SUBTREE(*DIR)`
Directories, shares, and files for 'terry' and its subtrees (not directory 'terry').	`OBJ('/QNTC/server1/fdrive/terry/*') SUBTREE(*ALL)`
Only the specific file 'myfile.exe'.	`OBJ('/QNTC/server1/gdrive/myfile.exe')`
The registry for an integrated Windows server	`OBJ('/QNTC/server1/$REGISTRY')`

Tip: If you want to use a CL command, see:
Save (SAV)
Restore (RST)

Related tasks:

"Backing up individual integrated Windows server files and directories" on page 123
The Integrated Server Support option allows you to save integrated server data (files, directories, shares, and the Windows registry) to tape, optical or disk (*SAVF) along with other IBM i data and restore the data on an individual basis.

"Restoring individual integrated Windows server files and directories" on page 134
Use the Restore (RST) command to restore individual files for your integrated Windows server.

Saving user enrollment information for integrated Windows servers

Use CL commands and APIs to save user profiles and enrollment information for an integrated Windows server.

More information may be found in the *Backup and recovery of security information* section in the Security Reference topic collection.

User profiles may be saved using the Save Security Data (SAVSECDTA) command or the Save Object List (QSRSAVO) API. The IBM i system value QRETSVRSEC must be set to 1 for integrated Windows server enrollment support.

User profiles need to be saved using the above methods for integrated Windows server enrollment. User profiles saved using other commands or APIs are not supported for integrated Windows server enrollment.

Related concepts:

"User and group enrollment concepts for integrated Windows servers" on page 56
Learn about how IBM i users and groups interact with integrated Windows servers.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

Backing up your system

Restoring the NWSD and other objects associated with integrated servers

Do these tasks to restore the IBM i configuration objects and files related to integrated servers.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

Backing up your system

Restoring the NWSD and virtual storage for integrated servers

One method of restoring your integrated server data is to restore the Network Server Description (NWSD) and virtual disk drives (virtual storage) that IBM i associates with that server. It is the fastest method for restoring large amounts of data.

If you used file-level backup, you can also restore specific integrated server files.

When you restore saved objects from IBM i, you need to be aware of these considerations:

1. Treat a network server description (NWSD), its predefined disk drives (see "Predefined virtual storage and naming for integrated servers" on page 27), and any user-defined disk drives that are linked to it as a unit. Restore them at the same time. Otherwise, the integrated server may not be able to re-establish items such as Windows server File System permissions.
2. To have IBM i automatically relink restored disk drives in the integrated file system to the appropriate NWSD, restore the NWSD after you restore the disk drives.
3. If you restore an NWSD before restoring the predefined and user-defined disk drives in the integrated file system, you might need to relink those disk drives. The system will attempt to relink the storage space to the NWSD that it was linked to when it was saved. To manually link the storage, see "Linking virtual storage to integrated servers" on page 83 for each disk drive that is associated with the NWSD.
4. When you are done restoring the NWSD and all its associated virtual storage spaces, vary on the integrated server.
5. For integrated Windows servers, when you restore a domain controller, ensure that the domain database held on the server is synchronized with the other domain controllers.

   Follow normal Windows procedures to do this and refer to documentation from Microsoft as necessary.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

Restoring virtual storage for integrated servers

Use the Restore (RST) command to restore virtual storage for your integrated server.

The predefined virtual storage for the integrated server system drive and installation drive, as well as any user-defined virtual storage are stored in the integrated file system.

To restore virtual storage for integrated servers, use the Restore (RST) command:

1. Ensure that the auxiliary storage pool (ASP) that you are restoring data to is varied on and available.
   - By default, a storage space that is being restored will be recreated into the ASP from which it was saved. If you want to restore the data to a different ASP than it was saved from, do the following steps.

     Use the Create NWS Storage Space (CRTNWSSTG) command to create a temporary storage space with the same name as the storage space you are restoring and specify the name of the ASP that you want the data to be restored to.

     Use the following steps to restore the data to the temporary storage space. The restore command replaces the data in the temporary storage space with the data that is being restored.
   - By default, a storage space that is being restored will be recreated on the same preferred storage media from which it was saved. If you want to restore the data to a different preferred storage media than it was saved from, do the following steps.

     Use the Create NWS Storage Space (CRTNWSSTG) command to create a temporary storage space with the same name as the storage space you are restoring and specify the preferred storage media that you want the data to be restored to.

     Use the following steps to restore the data to the temporary storage space. The restore command replaces the data in the temporary storage space with the data that is being restored.
2. If you are restoring from save media, ensure that you have mounted your media.
3. If no network server storage spaces currently exist on the system (none appear when you use the Work with NWS Storage Spaces (WRKNWSSTG) command), you must create the /QFPNWSSTG directory before you can restore network server storage spaces. To create the /QFPNWSSTG directory, complete these steps:
   a. On the IBM i command line, type `CRTNWSSTG` to create a network server storage space and press **F4**.
   b. Provide a name for the storage space.

c. Use the minimal size allowed and specify the appropriate storage pool (ASP).

d. Press **Enter** to create the storage space. IBM i creates the storage space in the /QFPNWSSTG directory.

4. To restore the storage spaces, type `RST` and press **F4**.
5. In the `Name` field under `Objects:`, specify `'/QFPNWSSTG/`*`stgspc`*`'`, where `stgspc` is the name of the network server storage space.

   For example, to restore the system drive, use `/QFPNWSSTG/nwsdname1`. To restore the installation drive, use `/QFPNWSSTG/nwsdname2`.
6. If you are restoring a storage space that resided in a user ASP or an independent ASP and was saved on i 5.4 or earlier releases, you must also specify the UDFS object. Starting with i 6.1, the UDFS file is not specified on the save or restore commands since it is automatically included with the storage space directory.

   Note: To restore the .UDFS object to an independent storage pool, specify `dev/`*`independent ASP name`*`/stgspc.UDFS` where *`independent ASP name`* is the name of the independent storage pool and *`stgspc`* is the name of the network server storage space.
7. Specify values for any other parameters that you want and press **Enter** to restore the storage space.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

Restoring integrated server NWSDs

Use the Restore Configuration (RSTCFG) command to restore a network server description (NWSD) object for an integrated server.

In a disaster recovery situation, you would use Restore Configuration (RSTCFG) to restore all the configuration objects, one of which is the integrated server network server description (NWSD). In some situations, for example when you migrate to new integrated server hardware, you need to specifically restore the NWSD. To have IBM i automatically relink virtual disk drives (virtual storage) within the integrated file system to the restored NWSD, restore those disk drives first.

1. On the IBM i command line, type `RSTCFG` and press F4.
2. In the `Objects` field, specify the name of the NWSD followed by an '*'. This will restore both objects (NWSD, LIND) that have used the standard naming convention in one pass and in the proper sequence.
3. In the `Device` field, specify the device name if you are restoring from media. If you are restoring from a save file, specify *SAVF and identify the name and library for the save file in the appropriate fields.
4. Press Enter to have IBM i restore the NWSD.
5. When you are done restoring the NWSD and all its associated storage spaces, start the integrated server. See "Starting integrated servers" on page 70.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

Restoring NWSH objects

Use the Restore Configuration (RSTCFG) command to restore the Network Server Host Adapter (NWSH) objects for integrated servers.

In a disaster recovery situation, you would use Restore Configuration (RSTCFG) to restore all the configuration objects, one of which is the network server host adapter (NWSH).

1. On the IBM i command line, type `RSTCFG` and press F4.
2. In the `Objects` field, specify the name and type of the NWSH.
3. In the `Device` field, specify the device name if you are restoring from media. If you are restoring from a save file, specify *SAVF and identify the name and library for the save file in the appropriate fields.
4. Press Enter to have IBM i restore the NWSH.

Note:

1. When you restore an NWSH, you must start the NWSH before you start the integrated server.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

Restoring NWSCFG objects and validation lists

Use the Restore Object (RSTOBJ) command to restore network server configuration (NWSCFG) objects.

For servers attached by iSCSI, the configuration objects need to be restored to the QUSRSYS library. These include the network server configuration objects (type *NWSCFG) and an associated validation list object (type *VLDL).

Note: The *NWSCFG and *VLDL objects will share the same name.

To restore network server configurations, you use the Restore Object (RSTOBJ) command:

1. On the IBM i command line, type `RSTOBJ` and press `F4`.
2. If you are restoring from save media, ensure that you have mounted your media.
3. In the **Objects** field, specify the name the network server configuration.
4. In the **Save Library** field, specify `QUSRSYS`.

5. In the **Device** field, specify either the name of the device that contains the save media or specify `*SAVF` if you are restoring from a save file.
6. In the **Object** types field, specify both `*NWSCFG` and `*VLDL`.
7. If you are restoring from a save file, specify the name and library for the save file.
8. Press `Enter` to restore the network server configuration and associated validation list.

Related concepts:

"IBM i configuration objects for integrated servers" on page 46
IBM i uses objects to represent and control integrated server hardware, software, and virtual storage.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

Restoring individual integrated Windows server files and directories

Use the Restore (RST) command to restore individual files for your integrated Windows server.

IBM i Integrated Server Support supports file-level backup and recovery of your integrated Windows server files. You can recover a particular file from your IBM i backup without restoring the entire disk drive (virtual storage). Before using this method, however, consider the amount of data you need to restore. For large amounts of data, restoring an entire disk drive object is much faster than restoring all the individual files in the disk drive. To restore a smaller amount of data, this method works great.

You should restore the directory first, then files, then the registry, then reboot for new registry entries to take effect. To restore files that you saved by this method, use the RST command:

1. Ensure that the integrated Windows server and TCP/IP are running.
2. On the IBM i command line, type `RST` and press F4.
3. In the `Device` field, specify the device on which the data is available. (For example, `'QSYS.LIB/TAP01.DEVD'` restores the data from tape.)
4. In the `Object` field, specify what you want IBM i to restore in the form '/QNTC/*servername*/sharename'

   You can use wildcard characters. Refer to "Examples: Saving parts of integrated Windows servers" on page 128 for how to specify particular parts of an integrated Windows server. Avoid restoring Windows system files by this method because the restored files may behave unpredictably.
5. In the `Name` field, specify the path name of the object to restore.
6. You can use the `Include or omit` field to include or omit objects with the pattern that you specify in the `Name` portion of the `Object` parameter.
7. In the `New object name` field, leave the object name the same or specify a new path name. The new path name must be referenced by a share name that exists on the integrated Windows server.

   Note: When you save a directory that has shares defined over it, IBM i saves the share information with the directory. If you specify a new object name when you restore the directory, IBM i does not re-create these shares.

8. Use the `Directory subtree` field to specify whether you want to restore subtrees under a directory. The default is to restore all directories.
9. To specify that you want to restore files that were saved during a particular period, specify starting and ending dates and times in the `Change period` field.
10. Provide any other information that you want IBM i to use to restore the files and press Enter.
11. When the files are restored, reboot the integrated server for new registry entries to take effect.

Related reference:

"Examples: Saving parts of integrated Windows servers" on page 128
These examples show how to use the save (SAV) or restore (RST) commands for specific parts of an integrated server.

➡ Backing up your system

Restoring user enrollment information for integrated Windows servers

Use CL commands and APIs to save and restore user profiles and enrollment information for an integrated Windows server.

More information may be found in the *Backup and recovery of security information* section in the Security Reference topic collection.

User profiles saved with the Save Security Data (SAVSECDTA) command or the Save Object List (QSRSAVO) API may be restored using:

- The Restore User Profiles (RSTUSRPRF) command and specifying the parameter USRPRF(*ALL). If the parameter USRPRF(*ALL) is not specified, then user profiles may be restored if the parameter and value SECDTA(*PWDGRP) is specified.
- The Restore Object List (QSRRSTO) API.

If you save user profiles using the Save Object List (QSRSAVO) API, and a previous target release value is used, the user profile enrollment definitions will not be restored. After restoring the user profiles, the enrollment needs to be defined. Use *IBM Navigator for i* or the Change NWS User Attributes (CHGNWSUSRA) command to define the enrollment.

User profiles need to be saved and restored using the above methods for integrated Windows server enrollment. User profiles saved using other commands or APIs are not supported for integrated Windows server enrollment.

Related concepts:

"User and group enrollment concepts for integrated Windows servers" on page 56
Learn about how IBM i users and groups interact with integrated Windows servers.

Related reference:

"What objects to save and their location on IBM i" on page 118
Use this table to determine which objects need to be saved when you save your integrated server.

➡ Backing up your system

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:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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

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

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

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement 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.

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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

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

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. _enter the year or years_.

Programming interface information

This IBM i integration with BladeCenter and System x publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM i.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

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

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 and logos are trademarks of Oracle, Inc. in the United States, other countries, or both.

Other product and service names might be trademarks of IBM or other companies.

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.

Product Number: 5770-SS1

Printed in USA