IBM i

Systems management
Management Central

7.1

IBM Confidential IBM Confidential IBM Confidential IBM Confidential

IBM i

Systems management
Management Central

7.1

Note

Before using this information and the product it supports, read the information in "Notices," on page 49.

This edition applies to IBM i 7.1 (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

Management Central

As a part of System i® Navigator, Management Central provides the technology that you need to do systems management tasks across one or more systems simultaneously.

With Management Central, you can perform many systems management functions as part of your base operating system. Management Central allows you to manage one or more systems through a single central system. Select a system to use as your central system, and then add the endpoint systems to your Management Central network. You can create groups of similar or related endpoint systems to manage and monitor your systems easier. Your central system can handle the communications for you. You can use such options as scheduling and unattended operations. Management Central is flexible and easily manipulated to suit your needs.

What's new for IBM i 7.1

Read about new or significantly changed information for the Management Central topic collection.

Miscellaneous technical updates were made.

How to see what's new or changed

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

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

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

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

PDF files for Management Central

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

To view or download the PDF version of the Management Central topic collection, select Management Central (about 321 KB).

You can also view or download a PDF version of specific sections of the Management Central topic collection:

- Getting started with Management Central (about 130 KB)
- Working with Management Central monitors (about 131 KB)
- Advanced job scheduler (about 172 KB)

You can view or download these related topics:

- Performance (about 1041 KB)
- Maintaining and managing i5/OS® and related software (about 452 KB)

Saving PDF files

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

1. Right-click the PDF link in your browser.

2. Click the option that saves the PDF locally.
3. Navigate to the directory in which you want to save the PDF.
4. Click **Save**.

Downloading Adobe Reader

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

Related reference:

"Related information for Management Central" on page 46
Web sites, and other information center topic collections contain information that relates to the Management Central topic collection. You can view or print any of the PDF files.

Getting started with Management Central

To work with Management Central more effectively, set up your central system and endpoint systems in a way that makes sense for your business environment. When these preliminary steps are completed, you are ready to start working with Management Central.

To view or download a PDF version of this topic, select Getting started with Management Central (about 290 KB).

Related information:

Installing System i Navigator

Before you begin

To make sure that you complete a successful installation and connection to Management Central, it is suggested that you follow these instructions before you start the installation process.

Configuring TCP prerequisite checklist

To ensure a smooth installation and setup of Management Central, make sure that the environment is properly prepared. Use the checklist in this topic to make sure that everything is ready before you begin installing Management Central.

Prerequisite checklist

1. Your System i product is current with the latest fixes, service packs for the client, and Java™ PTF group.
2. Read the Frequently Asked Questions at the Navigator service Web site.
3. Use the QTIMZON system value to set the Java time zone for any system that is OS/400® V5R2 or earlier. (This is because in any systems V5R3 or later the QTIMZON system value is used for the Java time zone.)
4. Load all clients with System i Navigator and the latest service packs. (The release of the client may be at a higher release than the central system.)
5. Determine the IP addresses of all of the clients that you are using. If the client has multiple IP addresses, it might be necessary to set the IP addresses to be used, so that the central system can connect back to the PC. In such a situation, setting the value for QYPS_HOSTNAME in the MgmtCtrl.properties file identifies the IP addresses to use. The following steps can help you decide which IP addresses work. To do this, use the IPCONFIG command from a command prompt. Write the addresses down for future reference.
   a. Confirm a valid connection from the PC to the central system. Use the ping command (ping *xx.xx.xx.xx*, where *xx.xx.xx.xx* represents the IP address of the central system) on the PC.
   b. Run IPCONFIG from the command prompt on the PC and record all of the IP addresses.

c. From the central system, ping each IP address.
d. For the first IP address that works, create a C:\MgmtCtrl.properties file and add this line: QYPS_HOSTNAME==<ip address on which you performed the ping>.

6. If you are upgrading System i Navigator from a previous release, close all open System i Navigator windows. Start System i Navigator and try to connect to the central system.

Related information:

Service & Support (System i Navigator)

Setting the time zone before upgrading to i5/OS V5R3, or later

Setting the Time zone (QTIMZON) system value

Management Central connection considerations

Understanding how Management Central establishes a connection is an important contributing factor toward a successful installation and setup. Whether your system configuration is complex or simple, there are many considerations that affect a successful connection.

How Management Central establishes a connection

When the Management Central Java server (QYPSJSVR) starts, it obtains the IP address for itself, by long name (system + domain name), from TCP/IP. Typically, the clients that appear under My Connections and the Management Central endpoints are defined by the system name or short name.

By default, the System i Navigator lookup frequency is set to Always. This setting causes a system that is listed under My Connections to use Domain Name System (DNS) or the TCP/IP host table (Configure TCP/IP (CFGTCP), option 10) to determine the IP address, so that it can connect to the central system. The Host Name Search Priority (Configure TCP/IP (CFGTCP), option 12) option controls how the DNS search is done. If it is *LOCAL, it searches the TCP/IP host table first. If it does not find it there, it uses DNS. If it is *REMOTE, then DNS is searched first, followed by the TCP/IP host table.

Connection timeout delay

When the Management Central systems on an endpoint are not running, a connection failure happens right away. However, if the system is down or if a bad IP address is being used, the connection cannot be made and there will be a several-minute timeout delay before the connection failure is posted.

Connection tests

Management Central uses the IP address of the system located under My Connection to connect to the central system. When Management central performs a connection test, it does a ping on the PC of the name that is being used for the central system (typically short name) and then it returns the same IP address as a Ping on the central system by the long name. If this is not successful, then the client cannot connect with the Java server. You can resolve this issue by overriding the IP address on the central system.

To override the IP address on the central system, use the following character-based command:

```
CALL PGM(QSYS/QYPSCONFIG) PARM(xxxx 'y.y.y.y')
```

Where *xxxx* is the setting QYPSHOSTNAME and *y.y.y.y* is the value of the IP address to be used.

Important: Edit the file using the character-based interface. Do not use a mapped drive, or other method.

Lookup frequency

The system environment variable QYPS_DNS sets the Management Central lookup frequency (values 0 = Never, 1 = Always). You can set the QYPS_DNS system variable by using one of these methods:

- Management Central properties window
- The Connection tab on the client
- The character-based interface, which is used to add a configuration property

  ```
  CALL PGM(QSYS/QYPSCONFIG) PARM(xxxx 'y')
  ```

  Where QYPS_DNS is the setting and y is the value 0 or 1.

It is suggested that you set the lookup frequency to Always. When the lookup frequency is set to Always, the IP address in the properties of the endpoint is ignored and a request for the IP address through DNS or the Host Table on the central system is made. As a result, if IP addresses are changed or if DNS or host table is changed, the new IP address is automatically picked up by Management Central.

When the lookup frequency is set to Never, the IP address contained in the properties of the endpoint object is used. As a result, a client might successfully connect to the central system which uses the IP address determined by Management Central, but then have a task run to the central system and have a connection failure. Such an event indicates that the Management Central lookup frequency is set to Never and that the IP address in the endpoint for the central system is not correct. To resolve this problem, edit the IP address of the endpoint on the endpoint properties window.

Note: The Management Central lookup frequency setting is different from the lookup frequency setting for a system under My Connections.

Connecting to a Java server

When a client connects to a Java server, the Java server uses an authentication procedure that connects back to the PC. Therefore, the central system must be able to ping the PC.

A common connection problem occurs when the PC's address is one that is reserved for private networks (such as when an individual uses VPN from home to gain access to their network behind their router). For example, assume the PC's address is 10.100.46.143 and the IP address of the central system is 164.143.28.82. A connection failure occurs because addresses that start with 10 are not forwarded by routers. In such a situation, you need to find out what the external IP address of the PC is and then set up a client C:\MgmtCtrl.properties file, and then add the line QYPS_HOSTNAME=xxx.xxx.xxx.xxx (where the xxx's are the external IP address of the PC). This causes the Java server use the IP address specified in the properties file to connect to the PC.

Management Central bulk data transfer considerations

A *bulk transfer* is a function that is used in Management Central to transfer data from a source system to a target system (such as sending a package, sending PTFs, and so on). For a successful transfer, the target system needs to be able to connect back to the source system. The IP address that is used on the target system is determined by the lookup frequency on the target system. If the lookup frequency is Never, then the IP address that is used is the one that is provided by the central system for the source system. If the lookup frequency on the target system is set to Always, then DNS or the host table is used to determine the IP address of the source system.

Running Management Central tasks from My Connections

Some of the System i Navigator functions use Management Central to obtain information. For example, you can view PTFs that are in Inventory by using **My Connections** > **Configuration and Service**. If Management Central cannot connect to the central system, then the function that you are trying to access experiences a several-minute timeout delay. This results in a connection failure message. A good practice to follow is to expand Management Central before you attempt to run any Management Central functions that are located under My Connections. By doing so, you ensure that you can connect to the central system.

To run a Management Central task on a system under My Connections, the system must be defined as an endpoint under Management Central. To define a system as an endpoint, expand Management Central, right-click Endpoint Systems, and select New Endpoint System.

Related information:

TCP/IP setup

TCP/IP troubleshooter

Experience Report: Configuring Management Central Connections for Firewall Environments

Installing Management Central

After you have completed all of the prerequisite tasks, you are ready to install Management Central. This topic series covers the installation steps as well as how the connection function works. If you fail to connect successfully after you have installed Management Central, see the information about troubleshooting Management Central connections.

Related tasks:

"Troubleshooting Management Central connections" on page 14
Several factors can prevent a connection to the Management Central server. You can take these steps to troubleshoot a failed connection.

Checking for the most current MC code

You must have the most current Management Central server code, Management Central client code, and Management Central dependencies before you can successfully use Management Central.

Checking the Management Central systems for the most current code

The IBM® Software technical document, Recommended PTFs for Management Central Supported Releases  , provides a summary of the recommended fixes by release.

To access this page from the IBM Web site  , follow this navigation path:

1. From the menu bar, click **Products**.
2. From the Products page, select **System i (iSeries)** under Systems & Servers.
3. Select **Support** from the navigation tree on the left.
4. Select **Support search** from the navigation tree on the left.
5. From the IBM System i5® Support search page, type the document number (360059564) in the **Search for** field and click **Search**.

Checking the Management Central client for the most current code

The System i Access  page provides up-to-date information about the service packs (fixes) for IBM i Access for Windows. To access this page from the IBM Web site  , follow this navigation path.

1. From the menu bar, click **Products**.
2. From the Products page, select **System i (iSeries)** under System & Servers.
3. Select **Software** from the navigation tree on the left.
4. Select **System i software from A to Z** from the System i software page.
5. Under A, click iSeries® Access.
6. On the iSeries Access page, select **Service Packs (Fixes)** from the navigation tree on the left.

Related tasks:

"Changing the central system setup" on page 13
You can select a different system as your central system at any time. The central system must be a system to which you are directly connected. For the latest System i Navigator functions, your central system should be running i5/OS Version 5, Release 4 or later.

Installing and accessing Management Central

Some of the systems management functions that you can use are optionally installable features of System i Navigator, the graphical user interface (GUI) for the System i product.

The following Management Central functions are installed if you only install the basic System i Navigator feature and none of the subfeatures:

- Tasks (inventory only)
- Endpoint systems
- System groups

If you did not install all of the features that you need when you installed System i Navigator, complete these steps:

1. Select **Start** > **Control Panel** > **Add or Remove Programs** > **IBM i Access for Windows** > **Change**.
2. Select the modify option to install the additional features that you need for systems management functions. To get all the systems management functions, select Configuration and Service, Users and Groups, Commands, Packages and Products, and Monitors.

When System i Navigator has been installed, double-click the desktop icon to start System i Navigator. You are now ready to set up your central system.

Related information:

Connecting to System i: System i Navigator

Installing System i Access for Windows on the PC

Verifying the connection function

The Verify Connection function that is located under Management Central is different from the function that is located under My Connection. This topic discusses the purpose of each function and how they differ from each other.

Verifying Connection from My Connection

Expand **My Connections**, right-click a system and select **Diagnostics** > **Verify Connection**.

This Verify Connection function pings the different host servers to see if they are working correctly and can be reached from the PC. Because it is restricted to single System i Navigator functions, it is one of the first things you should rule out when you are troubleshooting a Management Central connection failure. (Many Management Central functions build on the single system functions.) After you have confirmed that the connection to the endpoint systems is successful, then you can proceed to verify the connection from Management Central.

Verifying Connection from Management Central

Right-click **Management Central** and select **Verify Connection**.

The Verify Connection function from the Management Central container is a diagnostic tool that checks the most common factors that can cause a failed connection. It then displays the status of these tests. If it reports any failures, you can obtain specific information about the failure as well as recovery information by clicking **Details**. The following is a list of what Management Central verifies.

- The Java setup is correct on the central system. This includes verifying that certain .jar files are present, and that certain integrated file system file and folder authorities have not been changed.
- The required files that were included with the operating system are not deleted from the central system, are not damaged, and are being journaled.
- The TCP/IP configuration on the central system is valid. This includes verifying that the host names of both the central system and the PC are in the host tables or in DNS as appropriate.
- A simple Navigator connection can be made to the central system.
- The VRM, host name, the IP address of the central system, and the VRM of System i Navigator are correct.
- The ports that Management Central uses are not in use by another application on the central system.
- On the central system, the user profiles that are needed to run Management Central are not deleted, or disabled, and that they have valid, unexpired passwords.
- If SSL is being used on the central system and configured correctly, both the PC and the central system are using SSL.
- The central system is not marked as a secondary system in a Management Central High Availability environment. Secondary systems cannot be used as central systems.
- The Management Central servers are working correctly on the central system.
- What type of authentication are supported on the central system.

Note: System i Navigator uses the Java toolbox code on the client side (PC) to start the Management Central Verify Connection function. If the toolbox code is not working correctly, then the Verify Connection function will not start. If the Java Virtual Machine (JVM) or the toolbox code on the server side is not working correctly, the Verify Connection function will work until the last few checks. The JVM must be started before these last few checks can be performed.

Related information:

IBM Toolbox for Java

Setting up the central system

To manage multiple systems from a single system, you need to have a central system. After you have installed Management Central and connected successfully, you are ready to set up the central system.

The systems in your network are called *endpoint systems*. You select one of these endpoint systems as your central system. After you add endpoint systems to your network and select your central system, you only need to do your system administration tasks once. Your central system initiates your tasks and stores the necessary systems management data. You choose your central system when you first start System i Navigator. You can also easily change your central system at any time.

Important: The release of the central system must be the latest release in the network.

Setting up your central system for the first time

This information outlines the requirements for configuring the central system for the first time.

To start using System i Navigator, double-click the desktop icon and select a system to connect to and define a System i connection. The first system you specify is assigned as your central system. Management Central is shown automatically at the top of the list in the left pane of your System i Navigator window. The Management Central system is automatically started on the central system.

To access the distributed systems management functions of System i Navigator, expand **Management Central**.

Management Central databases are located in libraries QMGTC and QMGTC2. For systems running releases earlier than i5/OS V5R3, the Management Central databases are located in the QUSRSYS library.

To complete an initialization, the Management Central sever requires that QSECOFR is enabled and active. If you use a different profile name with the same kind of authorization as QSECOFR, you need to run the following command on the central system.

```
CALL PGM(QSYS/QYPSCONFIG) PARM(QYPSJ_SYSTEM_ID 'XXXXX')
```

(xxxxx is a user ID other than the default of QSECOFR)

In some cases, the central system might have multiple IP addresses by which it can be accessed (CFGTCP option 10). You can use a ping command on the central system to display the IP address that will be returned to Management Central. If this is not the IP address that the clients use to connect to the system, you can override the default IP address with the address that the ping command displayed. You can use the following command to override the default IP address.

```
CALL PGM(QSYS/QYPSCONFIG) PARM(QYPS_HOSTNAME 'w.x.y.z')
```

(w.x.y.z is the IP address that Management Central should use for connection purposes)

If your central system is running OS/400 V5R2 or later (or V5R1 with PTF SI06917), you can right-click **Management Central** and select **Verify Connection** to verify that the central system connection is configured properly. To see detailed information about any Failed message, select the message and click **Details** (or double-click the message).

Note: The Verify Connection function only confirms that Management Central is working properly on the central system. TCP/IP configuration and firewalls also might prevent the Management Central client from successfully connecting to the central system.

For more information about these and other Management Central tasks and topics, refer to the detailed task help that is available from the System i Navigator window. Click **Help** from the menu bar and select **System i Navigator overview** > **Management Central**.

Related information:

Experience Report: Configuring Management Central Connections for Firewall Environments

TCP/IP troubleshooter

TCP/IP setup

Scenarios

Management Central settings and options

If you are upgrading from a release that is earlier than V5R3, you should note that the system environment variables were moved. This topic explains where you can find the client and server environment variables for systems running i5/OS V5R3 or later.

/QIBM/UserData/OS400/Mgtc/Config/McCSConfig.properties

QYPS_EARLIEST_RELEASE
QYPS_MAXPTF_SIZE
QYPS_FTP_DISCOVERY
QYPS_DISCOVERY_TIMEOUT
QYPS_DISC_LCLSUBNET
QYPS_SNMP_DISCOVERY
QYPS_IP_DISCOVERY
QYPS_DISCOVERY_STARTUP
QYPS_MAX_SOCKETS
QYPS_MAX_CONTIMOUT
QYPS_RETRY_TIMEOUT
QYPS_RETRY_INTERVAL
QYPS_AUTORETRY

QYPS_SOCKETTIMEOUT
QYPS_COLLECTPTF_IFCHANGED
QYPS_DNS
QYIV_QUERY_MAX_SIZE
QYPSJ_SAVF_RECORDS
QYPSJ_TOOLBOX_TRACE
QYPS_LOCATION
QYPS_LOCATION2
QYPSJ_CONNECT_INTERVAL

/Qibm/UserData/OS400/Mgtc/Config/McCSSecure.properties

(SSL setup)

QYPS_AUTH_LEVEL
QYPS_SSL

/Qibm/UserData/OS400/Mgtc/Config/McEPConfig.properties

QYPS_TRACE
QYPSJ_TRACE
QYPSJ_SYSTEM_ID
QYPS_MAX_TRANSFERS
QYPS_HOSTNAME
QYPS_MINIMUM_PORT
QYPS_MAXIMUM_PORT

/Qibm/UserData/OS400/Mgtc/Config/McEPSecure.properties

QYPS_USER_PASSWORD
QYPS_BASIC_AUTH
QYPS_TRUST_LEVEL
QYPS_KERBEROS_PRINCIPAL
QYPS_KERBEROS_CONFIG
QYPSJ_SYSTEM_ID
QYPS_ID_MAPPING_ONLY
QYPS_USE_ID_MAPPING

Settings

System i Navigator allows you to manage multiple systems from a single system in a IP network environment. Some aspects of your TCP/IP environment may require changes to your Management Central system configuration. For example, if you are using a firewall or if you want to use SSL encryption for Management Central server communications, you might need to change some of your Management Central server settings.

Table 1. Management Central settings set via System i Navigator

Name	Description	Values	System i Navigator Field Name (Right-click Management Central > Properties > Connection tab)
QYPS_AUTORETRY	Specifies whether to automatically restart monitors on failed systems	0 = No, 1 = Yes	Automatically restart monitors on failed systems
QYPS_COLLECTPTF_IFCHANGED	Update fixes inventory only if changes have occurred	0 = NO, 1 = YES; 0 is the default	When collecting inventory, only update when changes have occurred
QYPS_DNS	IP address lookup frequency	0 = Never, 1 = Always,	IP address lookup frequency

Table 1. Management Central settings set via System i Navigator (continued)

Name	Description	Values	System i Navigator Field Name (Right-click Management Central > Properties > Connection tab)
QYPS_MAX_CONTIMOUT	Maximum time (in seconds) to wait for a connection to a system to be established	1 to 3600 (The default value is 180 seconds.)	While connected to endpoint systems
QYPS_MAX_SOCKETS	Maximum number of sockets that can be created on a system	200 (This is the default value.)	Maximum connections
QYPS_MAXPTF_SIZE	Maximum data transfer size	-1 = No maximum size	Maximum data transfer size (MB)
QYPS_RETRY_INTERVAL	Specifies how often (in minutes) to attempt a monitor restart	5 (This is the default value.)	How often to attempt restart
QYPS_RETRY_TIMEOUT	Specifies how long (in minutes) to attempt a monitor restart	180 (This is the default value.)	How long to attempt restart
QYPS_SOCKETTIMEOUT	Maximum time (in seconds) to wait on a socket to return from a request	30 seconds (This is the default value.)	When connecting to endpoint systems

Table 2. Management Central settings set via character-based interface

Name	Description	Values	Use the character-based interface
QYIV_QUERY_MAX_SIZE	Maximum number of records in the Inventory query	200	
QYPS_HOSTNAME	The host name or IP address that you want the endpoints and the PC to connect to when they need to make a new connection back to the system. **Note:** If you use a host name, then you are relying on the endpoint or the PC to resolve the host name through their host table or DNS.		
QYPS_LOCATION	Library name where the Management Central databases are found	QMGTC	
QYPS_LOCATION2	Second library name where the Management Central databases are found	QMGTC2	
QYPS_ID_MAPPING_ONLY	Indicates whether only the Enterprise Identity Mapping (EIM) should be used for authentication	0=No, 1=Yes	
QYPS_MAXIMUM_PORT	Used by BDT (Bulk Data Transfer) QYPSBDTSVR job. Maximum of range of port numbers to be used.		
QYPS_MINIMUM_PORT	Used by BDT (Bulk Data Transfer) QYPSBDTSVR job. Minimum of range of port numbers to be used.	Name of host server	
QYPS_TRACE	C++ server tracing	-1 to turn Off; or 0 to turn On	
QYPS_USE_ID_MAPPING	Java server tracing	-1 to turn Off; or 2 to turn On	
QYPSJ_CONNECT_INTERVAL	How often (in seconds) to do the heartbeat to check connections	60	
QYPSJ_PORT	Port on which the Java server is listening to for incoming client requests	5544 (This is the default value.)	
QYPSJ_SAVF_RECORDS	Maximum number of records in the Java save file	100	
QYPSJ_SYSTEM_ID	User profile with all object authority	User profile which the Java server runs for certain tasks. This profile must have *SECOFR class authority. QSECOFR is the default, or you can specify the user profile name.	
QYPSJ_TOOLBOX_TRACE	Indicates whether to turn Toolbox trace on	0=Off, 1=On	
QYPSSRV_PORT	Port on which the C++ server is listening for incoming client requests	5555. This is the default value.	

Table 2. Management Central settings set via character-based interface (continued)

Name	Description	Values	Use the character-based interface
QYPSJ_TRACE	Port on which the C__ server is listening for incoming client requests	5555. This is the default value.	

Table 3. Management Central settings set via System i Navigator

Name	Description	Values	System i Navigator Field Name (Management Central > Right-click Endpoint Systems > Properties)
QYPS_DISC_LCLSUBNET	Discover local subnet	0 = No, 1 = Yes	
QYPS_DISCOVERY_STARTUP	Search every time the Management Central server starts	0 = No, 1 = Yes	
QYPS_DISCOVERY_TIMEOUT	Discovery timeout (in seconds)	15 (This is the default value.)	Timeout (seconds)
QYPS_EARLIEST_RELEASE	Earliest operating system release to search for	V5R4M0, this is the default	Earliest operating system release to search for
QYPS_FTP_DISCOVERY	Run discovery using File Transfer Protocol	0 = No, 1 = Yes	How to verify systems, FTP check box
QYPS_IP_DISCOVERY	Run discovery using Internet Protocol	0 = No, 1 = Yes	
QYPS_SNMP_DISCOVERY	Run discovery using Simple Network Mail Protocol	0 = No, 1 = Yes	How to verify systems, SNMP check box

The following table contains Property file settings that you might need to change in order to accommodate your system's needs. Unless it is otherwise indicated, use the character-based interface to make these changes.

Table 4. Management Central property file parameters

Parameter	Description	Values	
QYPS_SSL	Turns the Secure Sockets Layer (SSL) on or off.	0 = Off, 1 = On	System i Navigator Field Name **(Right-click Management Central > Properties > Security tab**) Field name = Use Secure Sockets Layer (SSL)
QYPS_AUTH_LEVEL	SSL authentication level. This value works with the QYPS_SSL.	0 = off (This is the default. It can only connect to a server without SSL), 1 = Sever Authentication on (This means it can connect to server with or without SSL.)	System i Navigator **(Right-click Management Central > Properties > Security tab**) Field name = Authentication level
QYPS_USER_PASSWORD	Require password on endpoint systems	0 = No, 1 = Yes	System i Navigator **(Right-click Management Central > Properties > Security tab**) Field name = Use profile and password authentication
QYPSJ_SYSTEM_ID	The user profile with which the Java Server runs for certain tasks	QSECOFR is the default value. You can also specify a user profile name, however its profile must have *SECOFR class authority.	

Related tasks:

"Troubleshooting Management Central connections" on page 14
Several factors can prevent a connection to the Management Central server. You can take these steps to troubleshoot a failed connection.

Adding endpoint systems to your Management Central network

An endpoint system is any system or logical partition in your IP network that you choose to manage through your central system.

When you add a connection to a system from System i Navigator (by clicking **File** > **Connection to Systems** > **Add connection** while your current environment is selected in the left pane), the system is

added to the list under your current active environment (typically named My Connections). Alternatively, when you add a new endpoint system, the system name is added to the list of Endpoint Systems under Management Central.

When you perform an action on a system under My Connections, a direct connection from the client (your PC) to the system is required, and actions are performed on one system at a time. In contrast, Management Central allows systems management tasks to be performed on multiple systems (in the Endpoint Systems list) and only one client connection (to the central system) is required.

The central system handles the connections to the endpoint systems. The Management Central property setting for the Lookup Frequency controls how the IP address for an endpoint system is determined. If it is set to Never then the IP address that is stored in the endpoint object is used. If it is set to Always, then the TCP/IP, on the system provides the IP address for the system name that is specified.

Note: If you are adding endpoint systems that are running OS/400 V5R1, you must have the following fixes (also known as PTFs) installed on the V5R1 system: SI01375, SI01376, SI01377, SI01378, and SI01838. Without these fixes, you will not be able to use all the systems management functions on the endpoint system.

To add one or more endpoint systems, complete the following steps:

1. Right-click **Endpoint Systems** and select **New Endpoint System**.
2. Enter the name of the system and click **OK**.

The endpoint systems that you added appear automatically under **Endpoint Systems** in your System i Navigator window. After you have added an endpoint system, you can view its properties. You can also change the description or the IP address as needed.

Next, you can create system groups to help you manage different sets of endpoint systems. The new system groups appear under Management Central in System i Navigator.

For more information about these and other Management Central tasks and topics, refer to the detailed task help that is available from the System i Navigator window. Click **Help** from the menu bar and select **System i Navigator overview** > **Management Central**.

How to completely remove endpoints

To completely remove an endpoint that is also defined as a My Connection system, all users that have the system defined must remove the My connection system so it will not be automatically added.

When connecting to a target system, Management Central requires and uses endpoint objects. Additionally, many Management Central functions appear under systems that are listed under My Connections. Thus, whenever a user creates a system under My Connections, an endpoint object is saved in the database on the central system as well as the client PC.

If you delete the endpoint from Management Central, only the entry in the central system database is deleted. You must also delete the system from all clients that have that system listed under My Connections. Otherwise, the next time a user that still has that system listed under My Connections starts System i Navigator, the endpoint is automatically added again to Management Central.

Creating system groups in your Management Central network

A *system group* is a collection of endpoint systems that you define. If you are working with multiple systems or multiple logical partitions, creating a system group allows you to perform tasks on all the systems without selecting each endpoint system. Just select the system group you created and start your task.

Endpoint systems can belong to several system groups at the same time. After you have created a system group, you can manage the entire group from your central system as if it were a single system.

To create a system group, follow these steps:

1. Open **Management Central** from your **System i Navigator** window.
2. Right-click **System Groups** and select **New System Group**.
3. On the **New System Group** window, specify a unique name for the new system group. You can also enter a brief description that will help you identify this group in a list of system groups.
4. From the **Available systems** list, select the endpoint systems that you want to include in this new group. Click the **Add** button to add the systems to the **Selected systems** list.
5. If you want to give other users the ability to view or change this system group, use sharing. Click the **Sharing** tab and specify **Read-only** or **Full** sharing. If you specify **None**, other users cannot view or change this system group unless they have special authority, which is administered under Host Applications in Application Administration. Users with this special authority, called Management Central Administration Access, can view all tasks, definitions, monitors, and system groups under Management Central in the System i Navigator window.
6. Click **OK** to create the new system group.

The system group you create will include all the endpoint systems you entered. You may decide later that you want to edit that list of endpoint systems. You can always add more endpoint systems or remove endpoint systems from your system group.

You can delete system groups from Management Central. When you delete a system group or remove endpoint systems from a system group, only the system group is changed. The endpoint systems that were in the system group are still listed under **Endpoint Systems** in the System i Navigator window. If you delete an endpoint system from the **Endpoint Systems** list, that endpoint system is removed from all system groups.

For more information about these and other Management Central tasks and topics, refer to the detailed task help that is available from the System i Navigator window. Click **Help** from the menu bar and select **System i Navigator overview** > **Management Central**.

Related information:

Management Central and Application Administration

Changing the central system setup

You can select a different system as your central system at any time. The central system must be a system to which you are directly connected. For the latest System i Navigator functions, your central system should be running i5/OS Version 5, Release 4 or later.

If your PC is running V5R2 or V5R3 System i Navigator, and you want to select a central system that is running OS/400 V5R1, you must have the following fixes (also known as PTFs) installed on the V5R1 system: SI01375, SI01376, SI01377, SI01378, and SI01838. Without these fixes, you cannot connect to the V5R1 system as a central system.

To change your central system, follow these steps:

1. Right-click Management Central and select **Change Central System**.
2. Use the **Change Central System** window to choose a system from your list of connected systems.
3. If the system you want to use as your central system is not currently connected to your System i Navigator network, right-click your active environment (typically My Connections) and choose **Connection to Systems** > **Add connection**. When the new system is connected, you can change your central system to the new system.

After you have added endpoint systems and created system groups, those endpoint systems and system groups will appear under Management Central as well. Once you have set up your central system, you are ready to do the other tasks necessary for setting up Management Central.

Important: The central system that you use should be equal to or at a later release than the releases of the endpoints that are being used.

For more information about these and other Management Central tasks and topics, refer to the detailed task help that is available from the System i Navigator window. Click **Help** from the menu bar and select **System i Navigator overview** > **Management Central**.

Management Central plug-ins

A plug-in is a separately installable component of System i Navigator. A plug-in adds folders and objects to the hierarchy tree, choices to System i Navigator menus, and property pages to the property sheet for a folder or object. There are several Management Central plug-ins that you can use to manage your system.

Backup, Recovery, and Media Services (BRMS)
IBM Backup, Recovery, and Media Services (BRMS) helps you implement a disciplined approach to managing your backups and provides you with an orderly way to retrieve lost or damaged data.

i5/OS Clusters technology
As companies strive to compete in today's environment, high availability has become an essential key to many businesses. The i5/OS cluster technology can be used to achieve high availability in System i environments. Clusters technology provides mechanisms that enables critical resources to be automatically available on backup systems. Those resources could include data, application programs, devices, or environment attributes.

Working with systems with partitions
The Systems with Partitions container that is located under Management Central allows you manage the logical partitions of all of the servers on the system from the central system.

Advanced Job Scheduler
The IBM Advanced Job Scheduler for i5/OS (5761-JS1) licensed program is a powerful scheduler that allows unattended job processing 24 hours a day, 7 days a week. This scheduling tool provides more calendar features and offers greater control over scheduled events than the Management Central scheduler. You can also view job completion history and manage notification of a job's status.

Troubleshooting Management Central connections

Several factors can prevent a connection to the Management Central server. You can take these steps to troubleshoot a failed connection.

First and foremost, make sure that the central system is running on the highest operating system release in the network. Problems can occur because there are clients in the network that are running an operating system that is at a higher release than the central system.

Failed connection to the central system

1. From the PC, verify that you can ping your central system using the name or IP address listed in System i Navigator as your central system. If this is unsuccessful then there is something wrong with either your network, or your DNS, or host table. You must fix this before you can connect.
2. From the central system, make sure that you can ping your PC using the IP address of your PC. If this is unsuccessful, you will not be able to use some of the Management Central functions. For more information, see the Information Center experience report, "Configuring Management Central Connections for Firewall Environments".
3. Verify the central system connection. (From System i Navigator, expand **My Connections**. Right-click the system that is your system and select **Verify Connections**.) If this reports any errors, click **Details**. This opens a window that displays information about what happened.

4. Use the Verify Connection function that is located under Management Central to further troubleshoot the problem. (From System i Navigator, right-click **Management Central** and select **Verify Connection**.) If this reports any errors, click **Details**. This opens a window that displays information about what happened.

What to do if you still cannot connect

If you still cannot connect, use the following procedures to further troubleshoot the problem:

1. Verify that the Management Central server QYPSJSVR is running on the Central System.
   a. From System i Navigator, expand **My Connections** > **system (that you are using as the central system)** > **Network** > **Servers** > **TCP/IP**.
   b. Look at the Management Central item to see if the server is started. If necessary, right-click Management Central under TCP/IP, and click **Start**.
   c. If the server still fails to start, view the job logs for possible problems, or continue with the next items to check for some common problems that can cause the servers not to start.
2. Check the TCP/IP configuration on the central system.

   It is important that the central system is able to ping itself using both the fully qualified domain name and the short name. If pinging either of these names fails, you will need to add the name and IP address to either the system's host table or DNS. Make sure that the IP address used in these pings is one that the PC can contact.
3. If you are using SSL with Management Central, verify that it is set up correctly. Make sure to configure your central system, all your endpoint systems, as well as System i Navigator on your PC.
4. Check the QSECOFR profile.
   a. Management Central requires a profile with *ALLOBJ and *SECOFR authority enabled, and a valid password must be set so that it does not expire.

      Important: You must make this change via the character-based interface, otherwise the system might not be able to read the file.

      By default, Management Central uses the QSECOFR profile. Thus, if this default has not been changed, then you can enable QSECOFR and set the password to never expire. If you choose to set the password with an expiration date, then you must remember to keep the password active. You can keep the password active by always changing the current password before it expires.

      If you are using a customized profile other than QSECOFR, then enable it and set the password to never expire. To change from using QSECOFR to a customized profile, open the properties file `/QIBM/UserData/OS400/MGTC/config/McEPSecure.properties`. Change the parameter `QYPSJ_SYSTEM_ID = QSECOFR` to `QYPSJ_SYSTEM_ID = YOURPROFILE` (where `YOURPROFILE` is the profile name replacing QSECOFR).
   b. Or, you can use the following command to change from using the QSECOFR profile to a customized profile:

      ```
      CALL PGM(QSYS/QYPSCONFIG) PARM(xxxx 'yyyy')
      ```

      where *xxxx* is QYPSJ_SYSTEM_ID and *yyyy* is the name of the profile to be used.
5. If both of the Management Central servers on the central system are started successfully and you have done the above troubleshooting, but you still cannot connect from System i Navigator, then most likely the problem is either the TCP/IP configuration related to the firewall. In either case, use the Configuring Management Central Connections for Firewall Environments experience report to troubleshoot this problem. A few important notes are listed below:
   - The central system needs to be able to initiate a connection with System i Navigator on the PC, so it is important that the Central System can ping the IP address of the PC.
   - The PC needs to be able to initiate a connection with System i Navigator that is using the following IPs:
     – The name or IP being used as the central system name in System i Navigator (the name of the system under my connections).

– The IP address that the central system gets when it pings itself.

Note: The initial connection to the central system uses the name or IP specified in System i Navigator for the central system. However, during this initial connection, the central system discovers its own IP address and sends that IP to the PC. The PC uses that IP address for all further communications. The ports that Management Central uses need to be open in any firewalls that are being used.

Failed connection from PC to the central system

1. Right-click Management Central and run Verify Connection.
2. Make sure that the single socket layer (SSL) for the Management Central servers is turned on. Look in `/QIBM/USERDATA/OS400/MGTC/config/McCSSecure.properties` and confirm that `QYPS_SSL>1` or `QYPS_AUTH_LEVEL>1`. If you change these values, remember to restart the Management Central servers.
3. If you are running OS/400 V5R2, determine if the QYPSSRV job fails to start. If it failed to start, then the Digital Certificate Manager (DCM) configuration was not done correctly. Make sure that you have assigned your certificate the Management Central Application identification as well as the host server IDs.
4. Is there a padlock icon next to the central system? If not, then the client is not using SSL to connect. Under My Connections, right-click the central system, go to the Secure Sockets tab, and then choose to use SSL. Then click **OK**. You must close System i Navigator and restart it before this value takes affect.
5. On that same Secure Sockets tab as mentioned in step 3, there is a button to Download the CA to your PC. Make sure that you have done this, using the operating system that you CREATED the CA on (not necessarily the central system).
6. On the same Secure Sockets tab mentioned in the above bullet, there is a Verify SSL Connection. Run this and look at the results.
7. If you are running OS/400 V5R2 verify that the file QIBM\ProdData\OS400\Java400\jdk\lib\security\java.security has the following properties defined as these can cause a connection problem.
   - os400.jdk13.jst.factories=true
   - ssl.SocketFactory.provider=com.sun.net.ssl.internal.ssl.SSLSocketFactoryImpl
8. If you are running OS/400 V5R2 on the client, on your PC, look at c:\Documents and Settings\All Users\Documents\ibm\client access\classes\com\ibm\as400\access\KeyRing.class. Is it size 0? If so, delete the file and download the Certificate Authority.

Failed connection from central system to endpoint

In addition to following the steps for troubleshooting a failed connection from the PC to the central system, you should also view the job log on the central system. It should give a reason for why the connection was rejected. (For example: (CPFB918) Connection to system mysystem.mydomain.com rejected. Authentication level 0. Reason Code 99. This means that the SSL is not active for the endpoint. Instead, it is at authentication level 0.) You can find the meanings for negative reason codes in /QSYS.LIB/QSYSINC.LIB/H.FILE/SSL.MBR.

Note: Endpoint systems do not require a padlock.

Additional considerations

Firewall considerations

All communication is TCP initiated from the PC to the central system. You can specify the exact port to use by adding the following line to the C:\MgmtCtrl.properties file:

```
QYPSJ_LOCAL_PORT=xxxx
```

where *xxxx* is the port number. The port number should be greater than 1024 and less than 65535. Additionally, the port number must not be used by another application on the PC. The port must be open through the firewall. Should the firewall require it, all sockets must be open.

Related concepts:

"Management Central settings and options" on page 8
If you are upgrading from a release that is earlier than V5R3, you should note that the system environment variables were moved. This topic explains where you can find the client and server environment variables for systems running i5/OS V5R3 or later.

Related information:

Scenario: Secure all connections to your Management Central server with SSL

Experience Report: Configuring Management Central Connections for Firewall Environments

Digital Certificate Manager

Working with Management Central monitors

Management Central monitors can be used to check your system performance, your jobs and servers, your message queues, and changes to selected files.

You can specify thresholds for various metrics on these monitors, and then specify actions to be taken whenever a monitor detects that a threshold has been triggered. For example, you can run an i5/OS command or start a program when the threshold is triggered. For specific examples that describe how you can use these monitors, see the related concept Scenarios: Performance.

You can use a *system monitor* to see detailed graphs that monitor the real-time performance of multiple i5/OS operating system. In the Graph History window, you can see a graphical view of the metrics that have been collected for an extended period of time by Collection Services. You can contrast this data with the real-time data for the last hour shown in a System Monitor window.

You can monitor your jobs and servers with a *job monitor*. For example, you might want to monitor a job's CPU usage, job status, or job log messages. For each of those metrics, you can specify a threshold and actions to take when that threshold is triggered. For example, you could set up your monitor to send a message to the system operator whenever the CPU usage exceeds a certain threshold. In addition to the i5/OS commands, you can use the Advanced Job Scheduler Send Distribution using JS (SNDDSTJS) command. This command notifies someone by e-mail that the threshold is exceeded, if the Advanced Job Scheduler (5761-JS1) licensed program is installed on the endpoint system.

You can create a *message monitor* to take action on a list of messages that are important to you. For example, when the message monitor detects CPI0953 (threshold of a disk pool is exceeded), you could specify to run a command that deletes objects that you no longer need from the disk pool.

You can use a *file monitor* to monitor for a specified text string or for a specified size. Or, you can monitor for any modification to one or more selected files. You can select one or more files to be monitored, or you can select the History log option, which will monitor the i5/OS history log (QHST).

Note: Integrated file system treats QSYS physical files as directories, with the physical file members actually treated as files.

You can use a *B2B activity monitor* to view a graph of active transactions over time, and you can run commands automatically when thresholds are triggered. You can search for and display a specific transaction as well as view a bar graph of the detailed steps of that specific transaction.

In System i Navigator, or on your PC. You can choose to be informed by an audible or visible alarm on your PC when important thresholds are reached. The monitor continues to run and perform any

threshold commands or actions you specified. Your monitor runs until you decide to stop it. You can view all your monitors, as well as all your Management Central tasks, remotely with System i Access for Wireless.

In the Management Central properties, you can specify whether you want the central system to automatically attempt to restart your monitors on endpoint systems where they failed to start. If you select to have the system automatically attempt to restart your monitors, you may also specify how long you want the central system to keep trying to restart the monitors and how often you want the system to try during that time period. For example, if you want the system to try to restart monitors every five minutes for a period of 3 hours, you can select **Automatically restart monitors on failed systems,** and then specify 180 minutes for **How long to attempt restart** and 5 minutes for **How often to attempt restart.**

The steps to create and run a monitor are basically the same for whichever type of monitor you choose to run.

To view or download a PDF version of this topic, select Working with Management Central monitors (about 194 KB)

Related information:

Monitor concepts

Scenarios: System i Navigator monitors

Management collection objects

Collection Services stores data for each collection in a single collection object from which you can create as many different sets of database files as you need. This introductory topic explains the management collection object, when it is created, and how the available Collection Services settings affect these objects.

A *management collection object* (also known as *MGTCOL) serves as an efficient storage medium to hold large quantities of performance data. Once you have configured and started Collection Services, performance data is continuously collected and stored in these objects. Then, when you need to work with performance data you can use the data that is stored in these objects to populate performance database files.

Each *MGTCOL object has one has one of these attributes:

***PFR (detailed data)**

*MGTCOL objects that have the *PFR attribute can become quite large. Their size depends on the number of active jobs in the system, performance metrics being collected, and the collection interval. Data in this type of object support the IBM Performance Management for System i5 (PM for System i5) performance metrics and reflect all of the requested system performance data. The **Location to store collections** field that is located on the Collection Services Properties window displays the library in which the *PFR objects are located. The job QYPSPFRCOL collects and stores this data in this object.

The collection is cycled (a new *PFR object is created) at least once in a 24 hour period and the QYPSPFRCOL job writes the performance data into the new object. You can schedule this to happen more frequently.

When PM for System i5 is running, the *PFR objects are placed in the QMPGDATA library. If you are not using PM for System i5, then the *PFR objects are placed in the QPFRDATA library. These are default settings.

Note: If you use the Create Database Files Now option you can specify a different library, however this does not change the default setting. All subsequent files will be written to the QMPGDATA (or the QPFRDATA) library.

***PFRDTL (graph data)**

Graph history and system monitors use *MGTCOL objects that have the *PFRDTL attribute. These objects are stored in the QMGTC2 library. The *PFRDTL object supports second and third level detail for the top twenty uses of the metric and the data retains the same interval by which it was collected.

The collection is cycled (a new *PFRDTL object is created) at least once in a 24 hour period and the job QYMEPFRCVT writes the data to a new object. The naming convention for *PFRDTL objects is Q0yyddd00, where yy is the year and ddd is the Julian day of the year. For best results when using the graph history function, you should retain a minimum of seven days of *PFRDTL objects.

***PFRHST (summary data)**

Graph history also uses *MGTCOL objects that have the *PFRHST attribute. These objects are stored in the QMGTC2 library. When the collection is cycled, the QYMEARCPMA job adds the data to the existing *PFRHST object. No detail data or properties data is available. You must start PM for System i5 to enable the summary data fields. The default retention period is one month. The summary data is summarized in one-hour intervals and does not support second and third level details.

Setting the retention period

You can set the retention period for these objects from the Collection Services Properties window.

1. From the System i Navigator, expand **Management Central** > **Endpoint Systems** > ***your system*** > **Configuration and Service**.
2. Right-click **Collection Services** and select **Properties**.

Viewing collection objects

Viewing collection objects through **System i Navigator**

You can use System i Navigator to view *MGTCOL objects with the *PFR attribute.

From the System i Navigator, expand **Management Central** > **Endpoint Systems** > ***your system*** > **Configuration and Service** > **Collection Services**.

You can also use this method. From the System i Navigator, expand **My Connections** > ***your system*** > **Configuration and Service** > **Collection Services**.

Each object that is listed under the Collection Name is a different management collection object. You can right-click the object to see its status and data summary.

Viewing collection objects through **Character-based interface**

The following command can be used to view objects for the *PFRHST and the *PFRDTL type collection objects in the library QMGTC2:

```
WRKOBJPDM LIB(QMGTC2) OBJTYPE(*MGTCOL)
```

Related information:

Collection Services

Creating database files from Collection Services data

Managing collection objects

System i Navigator monitors

Job monitors and Collection Services

In order to avoid creating a negative performance affect on your system, you should understand how the different metrics in the job monitor uses Collection Services.

The metrics that are available for a job monitor are:

- Job count
- Job log message
- Job status
- Job numeric values
- Summary numeric values

The data for the job numeric and summary numeric values metrics come from Collection Services. The overhead for obtaining this data is minimal and is not affected by the number of specific jobs that are being monitored. It takes two intervals of Collection services data before the first point or data metric value can be calculated. For example, if the collection interval is 5 minutes it will take more than 5 minutes before the first metric value is known.

The overhead for the job log message and job status metrics is much more costly in terms of the CPU resources required to obtain the information. Additionally, the number of jobs that are being monitored as well as the collection interval, affect the amount of CPU overhead that is required. For example, a job Monitor with a 5 minute interval will have six times the amount of overhead process to complete versus if the collection interval was set to 30 minutes.

Related information:

Collection Services

The QYRMJOBSEL job

For every job monitor that runs, a QYRMJOBSEL job starts. This topic explains the purpose of the QYRMJOBSEL job and what causes it to end.

The QYRMJOBSEL uses the information that is specified in the General page of the Job Monitor definition (**Management Central** > **Monitors** > **Job** > **Right-click a monitor and click Proprieties**) with Collection Services data (QYPSPFRCOL) to determine what specific jobs need to be monitored. These jobs are then shown in the bottom half of the Job Monitor status window.

Even if only one job is running, QYRMJOBSEL still examines all of the active job data from Collection Services to determine how many jobs are running, if new instances have started or if instances that were running during the previous interval have ended. The QYRMJOBSEL job does this analysis at each interval. Thus, the amount of CPU resource that is needed for QYRMJOBSEL to complete this function is determined by how many active jobs are on the system. The more active jobs, the more jobs for QYRMJOBSEL to analyze.

Additionally, the QYRMJOBSEL job registers with Collection Services the needed probe data, but it cannot provide the notification interval. So it is always at the lowest interval at which Collection Services is running. Thus, a smaller collection interval means that this processing is performed more frequently.

For example, suppose the job monitor server starts a job monitor at 5 minute collection intervals. Then another monitor that is using Collection Services starts, but uses a smaller interval. As a result, the

QYRMJOBSEL receives the data at the smaller or more frequent interval. If the smaller interval is 30 seconds, there will be a 10 time increase in the amount of data QYRMJOBSEL processes, thereby increasing the need for CPU resources.

When the job monitor is stopped, its associated QYRMJOBSEL job receives an ENDJOB immediate and terminates with a CPC1125 Completion 50 severity. This is the normal way that the QYRMJOBSEL is removed from the system.

Note: For QYRMJOBSEL to work properly, the Java time zone must be correctly set. This is done by setting the QTIMZON system value.

QZRCSRVS jobs and their affect on performance

Job monitors connect to a QZRCSRVS job for each job that is being monitored for the Job Log Messages and the Job Status metrics. The more jobs that are being monitored for these metrics, the more QZRCSRVS jobs are used.

QZRCSRVS jobs are not Management Central jobs. They are i5/OS TCP Remote Command Server jobs that the Management Central Java server uses for calling commands and APIs. In order to process the API calls for the Job Log Messages and Job Status metrics in a timely fashion within the job monitor's interval length, the APIs are called for each job concurrently at interval time.

When both metrics are specified on the same monitor, two QZRCSRVS jobs are started for each job. For example, if 5 jobs are monitored for Job Log Messages, 5 QZRCSRVS jobs are started to support the monitor. If 5 jobs are monitored for Job Log Messages and Job Status, then 10 QZRCSRVS jobs are started.

Thus, it is recommended that for standard systems, when you are using the Job Log Message and Job Status metrics, you limit the number of jobs monitored on a small system to 40 jobs or less. (With larger systems more jobs may be monitored. However, you need to have a clear understanding of the resources that are used when monitoring more jobs and determine the affordable number to monitor.) Also, severely limit using these two metrics for monitoring subsystems, as doing so can cause a large number of QZRCSRVS jobs to run. (A job monitor that uses just the other metrics and does not use Job Status or Job Log Message, does not use QZRCSRVS jobs.)

Tuning QZRCSRVS jobs

As shipped, QZRCSRVS jobs are prestart jobs that run in the QUSRWRK subsystem, and connections to these jobs will be routed to this subsystem. It is possible to configure your system so that prestart jobs can run in any subsystem. If you end the prestart jobs in QUSRWRK with the ENDPJ command, then the QZRCSRVS jobs start as batch-immediate jobs in the QSYSWRK subsystem whenever a connection is requested. No jobs start in advance of the connection.

You can also configure your system to prevent batch-immediate jobs from being used at all. See the Use of prestart jobs topic for details on how to configure the QZRCSRVS jobs so that they can run in any subsystem.

QZRCSRVS cleanup

A cleanup thread runs once an hour to determine whether a QZRCSRVS job is still being used by a Job Monitor. It determines if the job was used at least twice within the maximum job monitor interval length. If the job is not used during the previous two hours, it is ended. Java time stamps are used for this comparison, so it is imperative that the time zone value used by Java is correct (system value QTIMZON).

QZRCSRVS jobs are automatically removed two hours after the job it supports ends. Likewise QZRCSRVS jobs will end if the Job Monitor that created them stops, or if Management Central ends.

Note: Since the Management Central Job Monitor monitors active jobs, you might see messages like "Internal job identifier no longer valid" in the QZRCSRVS job. This normally happens when a monitored job with Job Log Messages or the Job Status metric ends while the monitor is running.

Related information:

Use of prestart jobs

Special considerations

When working with Management Central monitors, you need to consider these special points.

Special considerations when working with job monitors

- The Job Count metric monitors the number of active jobs that match the job selection criteria during a collection interval.
- The Job Monitor window (**Management Central** > **Monitors** > **Job** > **Righ-click a job monitor** > **Open**) shows jobs that meet the criteria even if the jobs are no longer active at the end of the interval. Collection services provides information that determines the job count as well as the jobs to display in the window. This data contains information about all of the jobs that are active during that interval. Nevertheless, it is possible that if a job uses negligible CPU, then information about that job is not passed to the job monitor and so it does not appear in the count or the detail status display.
- For the metrics Job Status and Job Log Message if a job monitor triggers it continues to display those jobs that created the condition even if a job has ended and is not active during the interval. For this condition the job displays with a gray icon, and continues to be displayed until the trigger resets or the monitor restarts.

Special considerations when working with file monitors

- The Text metric monitors for a specific text string. When you use this metric, the File Monitor obtains a shared read lock on the files that it is monitoring. Programs which obtain a shared update lock can update files without interfering with the monitor. However, users, programs and commands (such as the Work with Objects using Programming Development Manager (WRKOBJPDM) command or the Start Source Entry Utility (STRSEU) command) that obtain an exclusive lock will interfere with the file monitor and might cause it to either fail or to not be able to monitor the criteria during each interval.
- A file monitor uses an integrated file system to access the information that it needs about the files that it is monitoring. Integrated file systems treat QSYS physical files as directories. Only the physical file members are actually treated as files. If you want to monitor the size of the entire contents of the QSYS physical file you must monitor all of the members that it contains (typically a single file member).

  For example, to monitor the size of the database file QAYIVDTA in the QMGTC library enter `/qsys.lib/qmgtc.lib/qayivdta.file/qayivdta.mbr` in the Files To Monitor field (**Management Central** > **Monitors** > **File** > **Right-click a monitor** > **Properties** > **General tab**). You can view the size of the database file from within the System i Navigator File System.
- The Text metric is the only valid metric when monitoring the QHST file.

Special considerations when working with system monitors

The V5R3 PTF SI18471 introduced the ability for the central system to try to restart a system monitor regardless of the reason. (Before this PTF, the central system would only restart a system monitor if the failure was due to a connection failure with the endpoint and if the monitor was still in a started status. This meant that only monitors with multiple endpoints that suffered connection failures were restarted.) To use this feature the following conditions must be met:

- The central system must be running release V5R4 or later. (This capability is also available on V5R3 central systems provided the PTF SI18471 is installed.)
- The keyword &RESTART is in the name of system monitor.
- The Management Central property **Automatically restart monitors on failed systems** is checked. (**Right-click Management Central** > **Properties** > **Connection tab**)

Creating a new monitor

Creating a new monitor is a process that begins at the New Monitor window. In System i Navigator, expand Management Central, expand **Monitors**, right-click the type of monitor you want to create (for example, **Job**), and then click **New Monitor**.

After you have given your new monitor a name, the next step is to specify what you want to monitor. If you are creating a job monitor, you will select which jobs you want to monitor. Be careful to monitor the smallest number of jobs that will give you the information you need. Monitoring a large number of jobs may have a performance affect on your system.

You can specify the jobs to monitor in these ways:

Jobs to monitor
: You can specify jobs by their job name, job user, job type and subsystem. When specifying job name, job user and subsystem, you can use an asterisk (*) as a wildcard to represent one or more characters.

Servers to monitor
: You can specify jobs by their server names. Select from the list of **Available servers** on the **Servers to monitor** tab. You can also specify a custom server by clicking the **Add custom server** button on the New Monitor or Monitor Properties - General page under the **Servers to monitor** tab. To create a custom server, use the Change Job (QWTCHGJB) API

When multiple job selection criteria are specified, all jobs matching any of the criteria are monitored.

Selecting the metrics

For each type of monitor, Management Central offers several measurements, known as *metrics*, to help you pinpoint different aspects of system activity. A metric is a measurement of a particular characteristic of a system resource or the performance of a program or a system.

For a *system monitor*, you can select from a wide range of available metrics, such as CPU utilization, interactive response time, transaction rate, disk arm utilization, disk storage, disk IOP utilization, and more.

For a *message monitor*, you can specify one or more message IDs, message types, severity levels. You can also select from a list of predefined sets of messages that are associated with a specific type of problem, such as a communications link problem, a cabling or hardware problem, or a modem problem.

For a *file monitor*, you can select to monitor files across multiple endpoint systems for a specified text string or for a specified size. Or, you can select to trigger an event whenever a specified file has been modified. You can select one or more files to be monitored, or you can select the **History log** option, which will monitor the i5/OS history log (QHST).

For a *job monitor*, available metrics include job count, job status, job log messages, CPU utilization, logical I/O rate, disk I/O rate, communications I/O rate, transaction rate, and more.

The Metrics page in the New Monitor window allows you to view and change the metrics that you want to monitor. To access this page, click **Monitors**, right-click the type of monitor you want to create (for example, **Job**), and then click **New Monitor**. Fill in the required fields, and then click the **Metrics** tab.

Use the online help to assist you in selecting your metrics. Remember to specify threshold values that allow you to be notified and to specify actions to be taken when a certain value (called the trigger value) is reached.

System monitor metrics

Metrics that you can use in a system monitor include the following:

Table 5. System monitor metric definitions

Name	Description
CPU Utilization (Average)	The percentage of available processing unit time that is being consumed by all jobs, threads of a job, and Licensed Internal Code tasks on the system. Click any collection point on the graph to see a Details chart that shows the 20 jobs or tasks with the highest CPU utilization.
CPU Utilization (Interactive Jobs)	The percentage of available processing unit time that is being consumed on the system for all jobs which include the following: • A 5250 workstation that includes a Twinax attached remote line and local area network (LAN) line • Systems Network Architecture (SNA) attached line that includes SNA display station pass-through • All Telnet sessions, for example, LAN, IBM Personal Communications, System i Access PC5250, and other SNA or Telnet emulators Click any collection point on the graph to see a Details chart that shows the 20 interactive jobs (5250 jobs) with the highest CPU utilization.
CPU Utilization (Interactive Feature)	The percentage of available interactive capability. The model number of your server (and for some models, the optional interactive feature card) determines the interactive capability of your system. It is possible to operate at greater than 100% of your available interactive capability. However, optimal system performance is achieved by maintaining an interactive workload that does not exceed the 100% level for extended periods. A recommended range should be approximately equal to or less than 70%. Click any collection point in the graph to see a Details chart that shows the 20 jobs with the highest CPU contributing to this workload.
CPU Utilization Basic (Average)	The percentage of available processing unit time that is being consumed by all jobs on the system. This metric includes the same work as CPU Utilization (Average) but does not include active job details. No additional data is available for this metric. You save system resource by not tracking the more detailed information.
CPU Utilization (Secondary Workloads)	The percentage of available processing unit time that is being consumed by secondary workloads running on your dedicated server. For example, if your system is a dedicated server for Domino®, Domino work is considered the primary workload. CPU Utilization (Secondary Workloads) shows the available processing unit time that is being consumed by any work other than Domino work on your server and can include WebSphere® Java and general Java servlets that run as Domino applications. No additional data is available for this metric.
CPU Utilization (Database Capability)	The percentage of available database capability that is being consumed by i5/OS database functions on your system, which includes file I/O, SQL, and general query functions. The model number and features of your system determine the amount of CPU available for database processing on your system. A recommended range should be approximately equal to or less than CPU Utilization (Average). Click any collection point in the graph to see a Details chart that shows the 20 jobs with the highest database CPU utilization.
Interactive Response Time (Average)	The average response time, in seconds, being experienced by 5250 interactive jobs on the system. Click any collection point on the graph to see a Details chart that shows the 20 jobs with the highest response time.
Interactive Response Time (Maximum)	The maximum response time, in seconds, that has been experienced by any 5250 interactive job on the system during the collection interval. Click any collection point on the graph to see a Details chart that shows the 20 jobs with the highest response time.

Table 5. System monitor metric definitions (continued)

Transaction Rate (Average)	The number of transactions that are being completed per second by all active jobs on the system. Click any collection point on the graph to see a Details chart that shows the 20 jobs with the highest transaction rate.
Transaction Rate (Interactive)	The number of transactions that are being completed per second on the system by active 5250 jobs, which include the following: • A 5250 workstation that includes a Twinax attached remote line and local area network (LAN) line • Systems Network Architecture (SNA) attached line that includes SNA display station pass-through • All Telnet sessions, for example, LAN, IBM Personal Communications, System i Access PC5250, and other SNA or Telnet emulators Click any collection point on the graph to see a Details chart that shows the 20 jobs with the highest transaction rate.
Batch Logical Database I/O	The average number of logical database input/output (I/O) operations being performed per second by all non-5250 batch jobs on the system. A logical I/O operation occurs when data is transferred between the system and application I/O buffers. This metric indicates how much work your batch jobs are performing during any given interval. Click any collection point on the graph to see a Details chart that shows the 20 batch jobs with the highest number of logical database I/O operations per second.
Disk Arm Utilization (Average)	The average percentage of all disk arm capacity that was utilized on the system during the collection interval. This metric shows how busy the disk arms on the system are during the current interval. Click any collection point on the graph to see a Details chart that shows the utilization of each disk arm.
Disk Arm Utilization (Maximum)	The maximum percentage of capacity that was utilized by any disk arm on the system during the collection interval. This metric shows how busy the disk arms on the system are during the current interval. Click any collection point on the graph to see a Details chart that shows the utilization of each disk arm.
Disk Storage (Average)	The average percentage of storage that was full on all disk arms during the collection interval. This metric shows how full the disk arms on the system are during the current interval. Click any collection point on the graph to see a Details chart that shows the percentage of storage that was full on each disk arm.
Disk Storage (Maximum)	The maximum percentage of storage that was full on any disk arm on the system during the collection interval. This metric shows how full the disk arms on the system are during the current interval. Click any collection point on the graph to see a Details chart that shows the percentage of storage that was full on each disk arm.
Disk IOP Utilization (Average)	The average utilization of all the disk input/output processors (IOPs) during the collection interval. This metric shows how busy the disk IOPs on the system are during the current interval. Multifunction IOPs can perform both Disk and Communication I/O work and can therefore be reported under either or both categories. If they performed work in both areas, the division of utilization is unknown and is reported fully under each category. Click any collection point on the graph to see a Details chart that shows the utilization of each input/output processor (IOP).
Disk IOP Utilization (Maximum)	The maximum utilization of any disk input/output processor (IOP) during the collection interval. This metric shows how busy the disk IOPs on the system are during the current interval. Multifunction IOPs can perform both Disk and Communication I/O work and can therefore be reported under either or both categories. If they performed work in both areas, the division of utilization is unknown and is reported fully under each category. Click any collection point on the graph to see a Details chart that shows the utilization of each input/output processor (IOP).

Table 5. System monitor metric definitions (continued)

Communications IOP Utilization (Average)	The average utilization of all the communications input/output processors (IOPs) during the collection interval. This metric shows how busy the communications IOPs on the system are during the current interval. Multifunction IOPs can perform both Disk and Communication I/O work and can therefore be reported under either or both categories. If they performed work in both areas, the division of utilization is unknown and is reported fully under each category. Click any collection point on the graph to see a Details chart that shows the utilization of each input/output processor (IOP).
Communications IOP Utilization (Maximum)	The maximum utilization of any communications input/output processor (IOP) during the collection interval. This metric shows how busy the communications IOPs on the system are during the current interval. Multifunction IOPs can perform both Disk and Communication I/O work and can therefore be reported under either or both categories. If they performed work in both areas, the division of utilization is unknown and is reported fully under each category. Click any collection point on the graph to see a Details chart that shows the utilization of each input/output processor (IOP).
Communications Line Utilization (Average)	The average amount of data that was actually sent and received for all non-LAN lines that are active during the time you collect data. Line utilization is an approximation of the actual amount of data transmitted compared with the theoretical limit of the lines based on the line speed settings in the line descriptions. The communication lines included on this monitor are one of the following line types: Bisync, Async, IDLC, X25, LAPD, SDLC, or PPP. This metric shows how actively the system is using its communication lines. If you have communications lines, such as fax lines, that are very busy much of the time, you may want to exclude these heavily utilized lines from the system monitor graph. Click any collection point on the graph to see a Details chart that shows the utilization of each line on the system.
Communications Line Utilization (Maximum)	The maximum amount of data that was actually sent and received for all non-LAN lines that are active during the time you collect data. Line utilization is an approximation of the actual amount of data transmitted compared with the theoretical limit of the line based on its line speed setting in the line description. The communication lines included on this monitor are one of the following line types: Bisync, Async, IDLC, X25, LAPD, SDLC, or PPP. This metric shows how actively the system is using its communication lines. If you have communications lines, such as fax lines, that are very busy much of the time, you may want to exclude these heavily utilized lines from the system monitor graph. Click any collection point on the graph to see a Details chart that shows the utilization of each line on the system.
LAN Utilization (Average)	The average amount of data that was actually sent and received on all local area network (LAN) lines in the system, compared with the theoretical limit of the lines based on the line speed settings in the line descriptions. The LAN lines included on this monitor are one of the following line types: token-ring or Ethernet. This metric shows how actively the system is using its LAN lines. Click any collection point on the graph to see a Details chart that shows the utilization of each line on the system.
LAN Utilization (Maximum)	The maximum amount of data that was actually sent and received on any local area network (LAN) line in the system, compared with the theoretical limit of the line based on its line speed setting in the line description. The LAN lines included on this monitor run one of the following line types: token-ring or Ethernet. This metric shows how actively the system is using its LAN lines. Click any collection point on the graph to see a Details chart that shows the utilization of each line on the system.

Table 5. System monitor metric definitions (continued)

Machine Pool Faults	The average number of faults per second that occur in the machine pool of the system during the time you collect the data. Only Licensed Internal Code runs in the machine pool. This metric shows the level of faulting activity in the system's machine pool. Click any collection point on the graph to see a Details chart that shows the number of faults per second in the system's machine pool.
User Pool Faults (Average)	The average number of faults per second occurring in all of the user pools on the system during the time you collect the data. This metric shows how much faulting activity is occurring in the system's user pools. Click any collection point on the graph to see a Details chart that shows the number of faults per second in each auxiliary storage pool.
User Pool Faults (Maximum)	The maximum number of faults per second occurring in all of the user pools on the system during the time you collect the data. This metric shows how much faulting activity is occurring in the system's user pools. Click any collection point on the graph to see a Details chart that shows the number of faults per second in each auxiliary storage pool.

Job monitor metrics

You can use any metric, a group of metrics, or all the metrics from the list to be included in your monitor. Metrics you can use in a job monitor include the following:

Table 6. Job monitor metric definitions

Name	Description
Job Count	Monitor for a specific number of jobs matching the job selection.
Job Status	Monitor for jobs in any selected status, such as Completed, Disconnected, Ending, Held while running, or Initial thread held. **Remember:** Metrics for job status can affect performance. Limit the number of jobs that you are monitoring to 40.
Job Log Messages	Monitor for messages based on any combination of Message ID, Type, and Minimum severity.

Job numeric values

Table 7. Job numeric values definition

Name	Description
CPU Utilization	The percentage of available processing unit time used by all jobs that are included by this monitor on this system.
Logical I/O Rate	The number of logical I/O actions, per second, by each job that is being monitored on this system.
Disk I/O Rate	The average number of I/O operations, per second, performed by each job that is being monitored on this system. The value in this column is the sum of the asynchronous and synchronous disk I/O operations.
Communications I/O Rate	The number of communications I/O actions, per second, by each job that is being monitored on this system.

Table 7. Job numeric values definition (continued)

Transaction Rate	The number of transactions per second by each job that is being monitored on this system.
Transaction Time	The total transaction time for each job that is being monitored on this system.
Thread Count	The number of active threads in each job that is being monitored on this system.
Page Fault Rate	The average number of times, per second, that an active program in each job that is being monitored on this system refers to an address that is not in main storage.

Summary numeric values

Table 8. Summary numeric values definition

Name	Description
CPU Utilization	The percentage of available processing unit time used by all jobs monitored on this system. For multiple-processor systems, this is the average percent busy for all processors.
Logical I/O Rate	The number of logical I/O actions, per second, by all jobs monitored on this system.
Disk I/O Rate	The average number of I/O operations, per second, performed by all jobs monitored on this system. The value in this column is the sum of the asynchronous and synchronous disk I/O operations.
Communications I/O Rate	The number of communications I/O actions, per second, by all jobs monitored on this system.
Transaction Rate	The number of transactions per second by all jobs monitored on this system.
Transaction Time	The total transaction time for all jobs monitored on this system.
Thread Count	The number of active threads for all jobs monitored on this system.
Page Fault Rate	The average number of times, per second, that active programs in all jobs monitored on this system refer to an address that is not in main storage.

Specifying the threshold values

Setting a threshold for a metric that is being collected by a monitor allows you to be notified and, optionally, to specify actions to be taken when a certain value (called the *trigger value*) is reached. You can also specify actions to be taken when a second value (called the *reset value*) is reached.

For example, when you create a system monitor, you can specify an i5/OS command that stops any new jobs from starting when CPU utilization reaches 90% and another i5/OS command that allows new jobs to start when CPU utilization falls to less than 70%.

For some metrics, it is appropriate to specify a reset value, which resets the threshold and allows it to be triggered again when the trigger value is reached. For those thresholds, you can specify a command to be run when the reset value is reached. For other metrics (such as the File Status metric and the Text metric on file monitors, and any message set on a message monitor), you can specify to automatically reset the threshold when the trigger command is run.

You can set up to two thresholds for each metric that the monitor is collecting. Thresholds are triggered and reset based on the value at the time the metric collection is made. Specifying a higher number of collection intervals in the Duration field helps to avoid unnecessary threshold activity due to frequent spiking of values.

You can also choose to add an event to the Event Log whenever the trigger value or the reset value is reached.

On the New Monitor - Metrics page, the threshold tabs provide a place for you to specify a threshold value for each metric that you have selected to monitor. For example, if you are creating a job monitor, you can set your threshold values in the following ways depending on the type of metric you have selected:

Job Count	When you define a threshold, you can specify a command to run on the endpoint system when the threshold is triggered. For example, selecting **25 jobs** will trigger the threshold whenever the monitor detects more than 25 jobs running during the number of collection intervals you specify for **Duration**. You can then specify a command to be run on the endpoint system when the monitor detects more than 25 jobs. Enter the command name and click **Prompt** for assistance in specifying the parameters for the command. For more detailed information and examples of specifying commands to be run when thresholds are triggered, see the performance scenarios topic. **Enable reset** is optional, and cannot be selected until a trigger is defined. You can also specify a command to be run on the endpoint system when the threshold is reset.
Job Log Message	You must select **Trigger when any of the following messages are sent to the job log** before you can specify the conditions to trigger a threshold. You can specify messages to monitor for based on any combination of Message ID, Type, and Minimum severity. Each row in the Job Log Message table shows a combination of criteria that must be met for a message to trigger a threshold. A threshold will be triggered if it meets the criteria in at least one row. Use the online help to specify the conditions to trigger a threshold. Be careful to monitor the smallest number of jobs that will give you the information you need. Monitoring a large number of jobs for job log messages may have a performance affect on your system. You can specify a command to be run on the endpoint system when the threshold is triggered. Enter the command name and click **Prompt** for assistance in specifying the parameters for the command. Be sure to click the Collection Interval tab to specify how often you want the monitor to check for job log messages. A message trigger can only be manually reset. You can specify a command to be run on the endpoint system when the threshold is reset. When you reset the monitor, you always have the option to reset without running the specified command.

Job Status	On the **Metrics - General** tab, select the statuses that you want to monitor for. Click the **Metrics - Status Threshold** tab to specify the conditions to trigger a threshold. You must select **Trigger when job is in any selected status** before you can specify the conditions to trigger a threshold. The threshold is triggered whenever the monitor detects that the job is in any selected status for the number of collection intervals you specify for **Duration**. You can then specify a command to be run on the endpoint system when the threshold is triggered. Enter the command name and click **Prompt** for assistance in specifying the parameters for the command. **Reset when job is not in selected statuses** is optional, and cannot be selected until a trigger is defined. You can specify a command to be run on the endpoint system when the threshold is reset.
Job Numeric Values	When you define the threshold, you can specify a command to run on the endpoint system when the threshold is triggered. For example, selecting **101 transactions per second** for the Transaction Rate metric will trigger the threshold whenever the monitor detects more than 101 transactions per second on any of the selected jobs during the number of collection intervals you specify for **Duration**. You can then specify a command to be run on the endpoint system when the monitor detects more than 101 transactions per second. Enter the command name and click **Prompt** for assistance in specifying the parameters for the command. **Enable reset** is optional, and cannot be selected until a trigger is defined. You can also specify a command to be run on the endpoint system when the threshold is reset.
Summary Numeric Values (total for all jobs)	When you define a threshold, you can specify a command to run on the endpoint system when the threshold is triggered. For example, selecting **1001 transactions per second** for the Transaction Rate metric will trigger the threshold whenever the monitor detects more than 1001 transactions per second on all of the selected jobs during the number of collection intervals you specify for **Duration**. You can then specify a command to be run on the endpoint system when the monitor detects more than 1001 transactions per second. Enter the command name and click **Prompt** for assistance in specifying the parameters for the command. **Enable reset** is optional, and cannot be selected until a trigger is defined. You can also specify a command to be run on the endpoint system when the threshold is reset.

For more information about these and other Management Central tasks and topics, refer to the detailed task help that is available from the System i Navigator window. Click **Help** from the menu bar and select **System i Navigator overview** > **Management Central**.

Specifying the collection interval

When you are setting thresholds for the metrics you have selected to monitor, you should consider how often you want the data to be collected.

Click the **Collection Interval** tab to select whether to use the same collection interval for all metrics, or to use different collection intervals for each metric type. For example, you may want to collect job count data every 30 seconds, but you may want to collect the job log message data every 5 minutes because job log message data typically takes longer to collect than job count data.

If you want to monitor numeric and status metrics for less than 5 minutes, you must select **Use different collection interval**.

Note: The job count, job numeric values, and summary numeric values metrics must have an equal or lesser collection interval than the collection interval for the job status metric.

To specify the number of collection intervals for each threshold, click the **Metrics** tab and indicate the number of intervals in the **Duration** field.

Specifying threshold run commands

A *threshold* is a setting for a metric that is being collected by a monitor. *Threshold commands* run automatically on your endpoint system when threshold events occur. Threshold commands are different from any threshold actions you may have set. Threshold actions happen on your PC or central system, while threshold commands run on your endpoint systems.

Using threshold commands

Threshold settings are used to automate any i5/OS command you want to run when thresholds are triggered or reset. For example, suppose you are running a job monitor and a certain batch job that is supposed to complete before the first shift begins is still running at 6:00 a.m. To accomplish this, you can set up Threshold 1 to send a page command to a system operator to look at it. You can also set up Threshold 2 to send a command to end the job if it is still running at 7:00 a.m.

In another situation, you might want to notify your operators with a page command when the job monitor detects that the wait time values for the FTP and HTTP servers have reached a median level. If the FTP server jobs end, you can restart the server with a start server command (such as STRTCPSVR *FTP). You can set thresholds and specify commands to automatically handle many different situations. In short, you can use threshold commands in any way that makes sense for your environment.

How do I set threshold commands?

On the New Monitor-Metrics page, click the **Thresholds** tab to enable your thresholds. Before you can set any threshold commands, you must turn your thresholds on by selecting the **Enable trigger** (or similarly named) option. You can then use this window to enter any commands you want to run when the threshold trigger value is reached. Select the **Enable reset** (or similarly named) option if you want to specify a command to run when the threshold reset value is reached.

Management Central monitors allow you to specify any batch commands to run on the server when the threshold is triggered or reset. You can enter an i5/OS command name and click **Prompt** (or press F4) for assistance in specifying the parameters for the command. You can even use replacement variables (such as &TIME or &NUMCURRENT) to pass information to the command, such as the time and actual value of the metric.

Specifying event logging and actions

When you have specified the threshold values for your monitor, you can click the **Actions** tab to select event logging and the PC actions to be taken when a threshold is triggered or reset.

Some of the actions you can select are:

Table 9. Actions that you can select

Action	Description
Log event	Adds an entry to the event log on the central system when the threshold is triggered or reset. The entry includes the date and time the event occurred, the endpoint system being monitored, the metric being collected, and the monitor that logged the event.
Open event log	Displays the event log when an event occurs.
Open monitor	Displays a list of systems that are being monitored for the specified metrics and a list of the values for the specified metrics as they are collected for each system.
Sound alarm	Sounds an alarm on the PC when the threshold for the monitor is triggered.

Table 9. Actions that you can select (continued)

Run i5/OS command	If you have specified a server command to run when the threshold for this monitor is triggered or reset, those commands run only during times that actions are applied. This option cannot be changed from the Actions page. If you do not want the command to run, you can remove the command from the Metrics page. Whenever you manually reset a threshold, you can select whether to run the specified reset command.

When you have specified the actions that you want to take when a threshold value is reached, you are ready to specify when to apply the thresholds and actions you have selected.

How to read the event log

The Event log window displays a list of threshold trigger and reset events for all of your monitors. You can specify on the Monitor Properties - Actions page for each monitor whether you want events added to the Event Log. To see the Properties pages for any monitor, select the monitor in the Monitors list and then select Properties from the File menu.

The list of events is arranged in order by date and time by default, but you can change the order by clicking on any column heading. For example, to sort the list by the endpoint system where the event occurred, click System.

An icon to the left of each event indicates the type of event:

Table 10. Icons and meanings they indicate

Icon	Description
	Indicates that this event is a trigger event for which you did not specify a server command to be run when the threshold was triggered.
	Indicates that this event is a trigger event for which you specified a server command to be run when the threshold was triggered.
	Indicates that this event is a threshold reset event.

You can customize the list of events to include only those that meet specific criteria by selecting **Options** from the menu bar and then selecting **Include**.

You can specify which columns of information you want to display in the list and the order in which you want the columns to be displayed by selecting **Options** from the menu bar and then selecting **Columns**.

You can view the properties of an event to get more information about what triggered the event log entry.

You can have more than one Event Log window open at the same time, and you can work with other windows while the Event Log windows are open. Event Log windows are updated continuously as events occur.

Applying thresholds and actions for a monitor

When you have specified your threshold values and chosen to log events, you can select whether to always apply these thresholds and actions, or to apply them only on the days and times you choose.

Note: Because system monitors run continuously, the following information does not apply.
If you select to apply thresholds and actions during specified times, you must select the starting time and the stopping time. If the central system is in a different time zone from the endpoint system, you should be aware that the thresholds and actions will be applied when the starting time is reached on the endpoint system that you are monitoring. You must also select at least one day that you want the thresholds and actions to apply. The thresholds and actions apply from the selected starting time on the selected day until the next occurrence of the stopping time on the endpoint system.

For example, if you want to apply your thresholds and actions overnight on Monday night, you can select 11:00 p.m. as the **From** time and 6:00 a.m. as the **To** time and check **Monday**. The actions that you specified occur whenever the specified thresholds are reached at any time between 11:00 p.m. on Monday and 6:00 a.m. on Tuesday.

Use the online help to finish creating your monitor. The online help also contains instructions on starting your monitor.

For more information about these and other Management Central tasks and topics, refer to the detailed task help that is available from the System i Navigator window. Click **Help** from the menu bar and select **System i Navigator overview** > **Management Central**.

Viewing monitor results

When you have specified when to apply the thresholds and actions you have defined for your monitor, you are ready to view your monitor results.

Double-click the monitor name to open the Monitor window. In the Monitor window, you can see the overall status of the monitor and a list of the target systems that the monitor is running on.

For job, message and file monitors, a list of the target systems (Summary Area) in the upper pane shows the status of the monitor on each system and the date and time that the monitor data was last collected. The Summary Area also shows additional information related to the specific metrics being collected.

After you select a system, detailed information about what is being monitored on that system is shown in the lower pane. For example, if you are viewing a Job Monitor window, the list of jobs in the lower pane shows the triggered events, the last event that occurred, and the actual values for the specified metrics.

You can select **Columns** from the Options menu to display additional columns of information. Click Help on the Columns window to see a description of each column.

From the list in the lower pane, you can right-click any item and select from a menu of actions that can be performed. For example, if you select a job, you can select reset triggered events, display job properties, hold, release, or end a job.

For system monitors, detailed information displays as graphs that you can save and print.

You can view all your monitors, as well as all your System i Navigator systems management tasks, remotely with System i Navigator for Wireless.

Related information:

System i Navigator graph history

Viewing graph history

Resetting triggered threshold for a monitor

When you are viewing the job monitor results, you can reset a triggered threshold.

You can choose to run the server command that was specified as the reset command for this threshold, or you can choose to reset the threshold without running the command.

You can also choose to reset thresholds at the job level, the summary level, the system level, or the monitor level:

Job level
: Select one or more jobs in the Job Area of the Job Monitor window. Select **File**, select **Reset with Command** or **Reset Only**, and then select **Jobs**. The thresholds for the selected jobs will be reset. Other thresholds that have been triggered for this monitor remain in the triggered state.

Summary level
: Select one or more systems in the Summary Area of the Job Monitor window. Select **File**, select **Reset with Command** or **Reset Only**, and then select **Summary**. The thresholds for job count, job numeric values metrics, and summary numeric values metrics will be reset. Other thresholds that have been triggered for this monitor remain in the triggered state.

System level
: Select one or more systems in the Summary Area of the Job Monitor window. Select **File**, select **Reset with Command** or **Reset Only**, and then select **System**. All thresholds for this monitor on the selected systems will be reset. Thresholds for this monitor that have been triggered on other systems remain in the triggered state. Any selections you have made in the Job Area are ignored.

Monitor level
: Select **File**, select **Reset with Command** or **Reset Only**, and then select **Monitor**. All thresholds for this monitor on all systems will be reset. Any selections you have made in the Summary Area or the Job Area are ignored.

Using other features of Management Central

After Management Central has been set up, you can use it to streamline your server administration tasks.

Working with inventory

The System i Navigator inventory functions can help you collect and manage various inventories on a regular basis and to store the data on the system that you selected as your central system.

For example, you can collect the inventory for users and groups, fixes, system values, hardware resources, software resources, service attributes, contact information, or network attributes. You may have other applications installed that allow you to collect lists of other types of resources.

You can either collect an inventory immediately or schedule it to be collected at a later time. You can schedule the inventory collection to occur daily, weekly, or monthly to keep your inventory current.

For more information about these and other Management Central tasks and topics, refer to the detailed task help that is available from the System i Navigator window. Click **Help** from the menu bar and select **System i Navigator overview** > **Management Central**.

Viewing an inventory

After you have collected the inventory, you can view the inventory list and right-click any item in the list to see the actions you can perform on the item.

For example, to display the inventory of all installed products on an endpoint system, select Software Inventory (**Management Central** > **Endpoint Systems** > **any endpoint system** > **Configuration and Service** > **Software Inventory** > **Installed Products**) This is a very easy way to see what software is installed on the endpoint system. The **Status** column reflects the current status of the software (Installed or Installed and supported) at the time of the last inventory collection (which is shown above the list).

It is recommended that you schedule the collection of all your system inventories on a recurring basis to keep your central system's inventory current.

How to use inventories

When you view an inventory on an endpoint system, you can right-click any item in the inventory list to see the actions you can perform on the item. Also by selecting the properties menu option for an inventory item (such as hardware) more information is displayed about that item.

For example, here are just a few of the ways that you can use inventories to manage your systems:

- After you have collected *fixes inventory*, you can compare fixes on one or more endpoint systems to the fixes on a model system. You can then send the missing fixes to the target endpoint systems and install them on those systems. You can also export the fixes inventory to a PC file, which you can use to work with the data in a spreadsheet program or other application.
- When you are viewing a *software inventory*, you can select any software product in the list, send it to one or more target endpoint systems, and install it on those systems. You can also export the software inventory to a PC file, which you can use to work with the data in a spreadsheet program or other application.
- Display a *hardware inventory* list to see the resource, status, and description of all hardware on the endpoint system. This is a very easy way to check the operational status of your hardware. The Status column reflects the operational status at the time of the last inventory collection (which is shown above the list). You can right-click any hardware listed and select **Properties**. You can review a great deal of information under the General, Physical location, and Logical address tabs. You can use this information for upgrades as well as problem analysis. You can also export the hardware inventory to a PC file, which you can use to work with the data in a spreadsheet program or other application.
- When you display the list for a *user inventory*, you can right-click one or more users and select any of the following actions: delete, edit, view the properties, or scan for objects owned by a user. You can do similar actions with groups by selecting Group Inventory for an endpoint system.

  You can search these inventories based on criteria that you specify. Additional search function is available when you search a users and groups inventory. You can export the results of the search or an entire inventory to a PC file to work with the data in a spreadsheet program or other application.

Running actions on an inventory

You might have applications installed that define actions that you can run against the collected inventory. If you have installed an application program that offers an action, you will see that action in the **Available actions** list in the Run Actions window.

To see the Run Actions window, right-click any system in the System i Navigator window, select **Inventory**, and then select **Run Actions**.

When you select an action from the **Available actions** list, a list of related inventories is shown under **Inventory for selected action**. You should select all the recommended inventories and then click **Add** to add this information to the **Selected actions to run** list.

For example, if you have installed the IBM Electronic Service Agent™ option of i5/OS, you can select **Send Electronic Service Agent inventory to IBM** from the **Available actions** list to receive your inventory data in a series of reports that show your system's growth and maintenance.

Searching a Management Central users and groups inventory

Searching on users and groups provides you with a lot of flexibility to query the user and group inventory for the information you want.

To access the **Search** window, right-click an endpoint system and select **Inventory** > **Search**.

The Basic search is for quick searches to find a particular user or group. The Advanced search page gives you the flexibility to search on additional profile properties. For example, you can search for all users on this endpoint system or system group with security officer authority by selecting Privilege class, and then selecting Security officer.

You can click **And** or **Or** to search on additional fields. For example, if you are searching for all users on this endpoint system or system group with security officer authority, you can narrow the search to users in your Accounting department with security officer authority by clicking **And** and selecting **Department** and entering the string **Accounting**.

From the Search Results window, you can perform many of the actions that you can perform on a user or group elsewhere within System i Navigator. For example, you can delete a user or group, edit the profile (for example, remove its Security Officer authority), view its properties, or scan for objects owned by a user or group. Also from the results window, you can export the search results into a spreadsheet, text file, or HTML (Web) page.

Advanced search is available only for user and group inventories, which require that both the central system and the endpoint systems are running OS/400 V5R1 or later.

Working with systems with partitions

The Systems with Partitions container that is located under Management Central allows you manage the logical partitions of all of the servers on the system from the central system.

With logical partitioning (LPAR), you can address multiple system requirements in a single system to achieve system consolidation, business unit consolidation, and mixed production or test environments. By itself, LPAR does not provide a significant availability increase. It can, however, be used to complement other availability strategies. Because each partition is treated as a separate system, you can run a single environment on a single system image. This can provide a more cost-efficient solution.

Authority requirements

Access to logical partition information in System i Navigator, dedicated service tools (DST), and system service tools (SST) requires either operations or administration authority to the logical partition function. In addition, you need remote panel authorization if you want to use the Operations Console remote panel for secondary partitions from your PC.

Logical partitions can be created using System i Navigator. To access logical partition functions, you must first configure the service tools server. Service tools are used to configure, manage, and service your models 8*xx* and earlier or logical partitions. If you want to manage logical partitions on servers other than model 8*xx*, you must use the Hardware Management Console (HMC). You need to use a service tools user ID with LPAR administrator authority.

Related information:

Configuring the service tools server

Partitioning with System i

Logical partition concepts

Planning for logical partitions

Creating logical partitions

Managing logical partitions by using System i Navigator, DST, and SST

Scheduling moving logical partition resources

Related information for Logical partitions

Running commands with Management Central

System i Navigator enables you to define an action or a task and then perform that action or task on multiple endpoint systems or system groups. These are the same commands that you normally run using the character-based interface.

For example, you can use a command definition to perform any of the following tasks:

- Set network attributes on multiple endpoint systems or system groups
- Set up your own help desk or operations "procedures book" to handle customer and system needs.

Any control language (CL) command that you can run in batch, you can send to multiple systems at the same time. Create the command definition, and then run the command on endpoint systems or system groups.

To run a command with Management Central, complete the following steps:

1. Expand **Management Central** > **Endpoint System.**
2. Right-click the endpoint system on which you want to run the command and click **Run Command**. For more information about this window, click **Help**.

You can click **Prompt** for assistance in entering or selecting an i5/OS command. You can choose to run the command immediately or schedule it to run at a later time.

Starting with V5R3, the command runs under the CCSID of the user profile that is submitting the command. If the profile is set to 65535 (or is set to *sysval, and the sysval is 65535), it uses the default CCSID 37.

Note: Be sure that the command you specify is supported by the release of i5/OS that is running on the target endpoint system. For example, starting with V5R3 any outputs other than job logs that are produced by a Run command are viewed by expanding the system under **My Connections** > **Basic Output** > **Printer Output**.

Related information:

Defining commands

Creating command definitions

You can create a command definition to save a command that you want to run over and over on multiple endpoint systems and system groups. Storing a command definition on the central system allows you to share commonly used or complex commands with other users. When a command is run from a definition, a task is created.

To create a command definition, complete the following steps:

1. Expand **Management Central** > **Definitions**.
2. Right-click **Command** and select **New Definition**.
3. The New Command Definition window opens.

Packaging and sending objects with Management Central

A bulk data transfer is the process of sending packages, fixes, PDFs and so on, from a source system to a target system in a single transfer. This topic discusses package definitions, what happens when a package is sent, and how to troubleshoot a failed transfer.

What you can do with package definitions

Sending files to another system or group of systems is a simple point-and-click operation in System i Navigator. If you expect to send the same files again at a later date, you can create a *package definition*, which can be saved and reused at any time to send the defined set of files and folders to multiple endpoint systems or system groups. If you create a snapshot of your files, you can keep more than one version of copies of the same set of files. Sending a snapshot ensures that no updates are made to the files during the distribution, so that the last target system receives the same objects as the first target system.

Another benefit of using System i Navigator to package and send objects is that you can run a command when the distribution of the package is complete. This means that you can:

- Distribute a batch input stream and run it.
- Distribute a set of programs and start your application.
- Distribute a set of data files and run a program that acts on that data.

You can specify whether to include subfolders in the package. You can also specify whether to keep or replace any file that already exists on the target system. You can start the send task immediately or click **Schedule** to specify when you want the task to start.

You can select and send files and folders without creating a package definition. However, a package definition allows you to group together a set of i5/OS objects or integrated file system files. The package definition also allows you to view this same group of files as a logical set, or as a physical set, by taking a snapshot of the files to preserve them for later distribution.

For more information about these and other Management Central tasks and topics, refer to the detailed task help that is available from the System i Navigator window. Click **Help** from the menu bar and select **System i Navigator overview** > **Management Central**.

Troubleshooting a failed transfer

- Look at the task's job log and try to determine a cause. From the Task Status window, right-click the failed endpoint and click **Task Output**.

  Sending packages uses the Save/Restore function. When a save or restore operation issues an error or warning message, the Management Central package send function marks the status as failed. This does not necessarily mean that the entire process failed. You need to check the job log and determine the cause of the failure. It is possible that there is a message that indicates that the restore function worked with limitation and thus generated a warning.
- Make sure that the target system can connect back to the source system.

  On the endpoint system ping itself by the long name. If this is successful, then on the source system, ping the endpoint system using its long name.

  To complete a successful transfer, the target system must connect back to the source system. The IP address that is used on the target system is determined by the lookup frequency on the target system. If the lookup frequency is Never then the IP address that is used is the one that is provided by the central system for the source system.

  It might be that target system cannot connect to the source system via this IP address, but can connect by using a different IP address, one that is defined in its host table. If the lookup frequency on the target is set to Always then it will use DNS, the host table, or both to determine the IP address of the source system and it will not use the IP address that is provided by the central system.

Related information:

Distributing fixes to multiple systems with System i Navigator

Packaging and distribution considerations

When working with the packaging function, you need to keep these considerations in mind.

- The packaging function that does not use a snapshot, stores temporary save files in the QRPLOBJ library. These files are prefixed with QYDS. The packaging function that uses a snapshot stores temporary save files in the QUSRSYS library. (A *snapshot* is a file that contains the data at a particular instant in time for all the files that were selected to be in a package. Creating a snapshot allows you to capture the contents of the selected files at a given time and then distribute that version of the files at a future time.)
- Typically, the QRPLOBJ library is cleaned up when an IPL is done. However, if between IPLs, the temporary storage that is used in QRPLOBJ is a concern, you can use the following commands to view and clear the objects that are in this library.

```
DSPLIB LIB(QRPLOBJ)
WRKOBJPDM LIB(QRPLOBJ) OBJ(*ALL)
WRKOBJPDM LIB(QRPLOBJ) OBJ(QYDS*) OBJTYPE(*FILE) OBJATR(*SAVF)
CLRLIB LIB(QRPLOBJ)
```

- The packaging function allows you to send and restore QSYS objects, QSYS libraries, integrated file system directories, and integrated file system files.
- Database files with referential constraints might not work properly because of sequence dependency. Additionally, database files with referential constraints behave differently depending on whether the database file that is being distributed is being replaced or is a new file. Thus the packaging function does not support sending database files when there is a dependency on the sequence in which the files are restored (such as logical database files).
- The packaging function does not support IASP distributions.
- You cannot use Management Central to distribute CUM tapes/packages.
- Packaging was not designed for very large distributions. A long duration of time maybe required to send very large save files to the target systems. If the size of the files (save file or snapshot size) is over 1 gigabyte, then you should run tests in your environment to determine if the time that is required to perform the distribution to the target systems is acceptable.

  As an alternative, you might want to send very large files between systems is to use FTP. This can be faster.
- You cannot distribute the latest i5/OS release, or migrate to a later release using Management Central. LPPs and Base i5/OS Options can be distributed and installed, but not Base i5/OS (QSYS and SLIC).
- You cannot mix QSYS and integrated file system files in a single package. Management Central uses the save/restore function, and is therefore bound by the restrictions that it imposes regarding mixing different file systems.

  You can create a package containing QSYS files and another one containing integrated file system files, and then send each package to an endpoint system. But, you cannot combine them into a single package.

  As a work around you can place the integrated file system objects into a save file. Then include the save file with your QSYS objects. Next, perform the restore of the save file to integrated file system objects. Or you can use the post command capability in the package definition to do the restore.
- You can refresh the snapshot by right-clicking the package definition and selecting **Update Snapshot** from the context menu. However, remember to resend the package to the systems that you want the updates on after you have updated the snapshot.
- Save and restore operations are performed under the user profile of the user that is signed on to System i Navigator. The post distribution command runs under the user profile of the person who started the distribution (the person that is signed on to System i Navigator). The job description that is used is QSYS/QYPSJOBD.
- If you are distributing a QSYS object that you created, then you will need *RWX authority to the QRPLOBJ library on both the source and the target systems. If someone else created the object, then you might need additional authorities. Authority to RSTOBJ is required when you are sending all of the objects from a library.

  If the package that you are distributing is an integrated file system file that you have created, then you do not need any additional authorities.
- In V5R2 and earlier, the package function runs under the C++ server QYPSSRV. In V5R3 and later, the package functions runs under the Java server QYPSJSVR. Thus, if your central system is V5R3 or later, you cannot create a snapshot on a V5R2 or earlier source system. In this special situation the QYPSJSVR server is not able to properly communicate with the V5R2 source QYPSSRV server. Nonetheless, you can still send a package from a V5R2 source system to a target system running V5R3 or later.

Managing users and groups with Management Central

System i Navigator can help you as a system administrator to keep track of the users, groups, and their level of privileges on one or more endpoint systems.

For more information about these and other Management Central tasks and topics, refer to the detailed task help that is available from the System i Navigator window. Click **Help** from the menu bar and select **System i Navigator overview** > **Management Central**.

The following list gives you an idea of the many ways in which System i Navigator can make your job easier.

Creating a user definition

You can create a user definition and then create multiple users across multiple systems based on the definition. First, create user definitions for the types of users on your systems. Then, when a request comes in for a new user, all special authorities, attributes, and other information common to that type of user are already stored in the user definition. You can even specify a command to be run after a user is created from a user definition! If you need assistance in entering or selecting an i5/OS command, you can click **Prompt** to select appropriate parameters and values.

When you create a new user from the user definition, you specify the name for the user, a brief description to help you identify this user in a list of users, and a new password for the user. All other properties of the new user are based on the properties stored in the user definition, unless you choose to change them. You may also select the groups the user should belong to and provide personal information about the user at the time the user is created.

Creating, editing, and deleting users and groups

You can create, edit, and delete users and groups across multiple endpoint systems or system groups--and even schedule these actions. For example, use the Edit Users function to change the properties for one or more users on the selected endpoint systems or system groups. If you need to change the authority level for several users on multiple systems, or if a user who has access to multiple systems changes his or her name, you can easily edit that information and apply the change to all systems.

When you use System i Navigator to delete users, you can select an action to be taken if any of the selected users owns objects on any system from which that user is being deleted. You can click **Scan for Owned Objects** to see what objects the selected users own on the selected endpoint systems or across the selected system groups.

Collecting an inventory

You can collect an inventory of the users and groups on one or more endpoint systems, and then view, search, or export that inventory to a PC file. Extensive advanced search capabilities are provided for easy searching. For example, you can search the inventory to see who has Security Officer privileges, as well as query other profile properties. Also, you can sort these inventory lists by clicking on any column heading. For example, you can group together all users in the inventory who have Security Officer privileges by clicking the Privilege Class heading.

You can perform various actions from the User Inventory list by right-clicking one or more users and selecting an action from the menu. For example, you can delete a user, edit a user, view its properties, or scan for objects owned by a user. You can do similar actions with groups by selecting Group Inventory for an endpoint system.

It is recommended that you schedule collection of users and groups inventory on a recurring basis to keep your central system's inventory current. Changes that you make to the user or group inventory on an endpoint system or system group under Management Central are automatically updated in the current central system's inventory.

Sending users and groups

You can send users and groups from one system to multiple endpoint systems or system groups. All the user properties you need are sent to the target systems, including the user name and

passwords (LAN server password as well as the i5/OS password), security settings, private authorities, Enterprise Identity Mapping (EIM) associations, and mail options. If the user has an entry in the system distribution directory on the source system, an entry is created (or updated) for that user on the target system.

You can also specify the action to be taken if any user in the list that you are sending already exists on the target system. When you are sending users, you can select not to change the user that already exists, or you can select to update the existing user with the settings from the user you are sending. When you are sending users, you can click Advanced to specify advanced send options. The advanced send options include specifying the mail system for the user and synchronizing the unique identifier of the user on the target system based on the user identifier of the user being sent.

To **send** users or groups from one system to another, you must also have save/restore (*SAVSYS) authority.

When you send a user from a system running IBM i 6.1 or a subsequent release to a system running a previous release, the number of device sessions that user can have might not be copied. You must reset the number of device sessions to an appropriate value after copying the profile.

Scanning for owned objects

You can scan for owned objects to find out what objects a user or group owns across multiple endpoint systems or system groups, and you can even scan for objects owned by multiple users simultaneously.

Synchronizing unique identifiers

You can synchronize the unique identifiers of users and groups across multiple endpoint systems to ensure that each of these numbers points to the same user on every system. This is especially important when you are working with systems in a clustering environment or a system with logical partitions. The user identification and group identification numbers are another way of identifying a user or group to a program. For example, the user identification and group identification numbers are used by programming interfaces in the integrated file systems environment.

You can choose to synchronize unique identifiers when you create new users or groups, when you edit users or groups, or when you send users or groups from one system to another. Be sure to keep your user and group inventories current if you are synchronizing unique identifiers when you create or edit users or groups.

Note: All i5/OS special authorities and other authorities that are needed when working with users and groups in the character-based interface are honored when managing users and groups with System i Navigator. This includes security administration (*SECADM) privileges, all object (*ALLOBJ) privileges, and authority to the profiles with which you are working. However, even a user with the most restricted set of system privileges (*USER) can view, search, or export a user or group inventory that has been collected by another user with the correct authorities. The user with *USER authority cannot create or delete users, edit existing users, or send users to another system.

Related concepts:

"Synchronizing functions" on page 44
You can synchronize the configuration of key functions, such as EIM and Kerberos, across a group of endpoint systems.

Related information:

Scenario: Configuring the Management Central servers for single sign-on

Propagating system settings from the model system (System A) to System B and System C

Sharing with other users in Management Central

Sharing saves you time, makes system administration easier, and reduces the number of redundant tasks you need to do. As of V5R4 and later releases, you can share system monitors and system events.

Sharing allows you to use (or share) the same items: monitors, monitor events, system groups, definitions, and system administration tasks. You can even set your user preferences to share all of the new tasks that you create. For example you might give a user special authority (administered under Host Applications in Application Administration) to view all tasks, definitions, job monitors, message monitors, file monitors, activity monitors, system monitors, system events, and system groups under Management Central in the System i Navigator window.

Only the owner of an item can change the level of sharing. The owner can specify any of the following levels of sharing:

None	Other users cannot view this item. Only the owner of the item or a user with special authority administered under Host Applications in Application Administration can view this item. Users with this special authority, called Management Central Administration Access, can view all tasks, definitions, job monitors, message monitors, system monitors, system events, and system groups under Management Central in the System i Navigator window.
Read-Only	Other users can view this item and use it. Other users can create a new item based on this one and make changes to the new one as needed. However, other users cannot delete or change this item in any way. If you are the owner of a monitor and have specified actions (such as opening the event log window or sounding an alarm on the PC), these actions occur for all users of the monitor whenever a threshold is triggered or reset. The other users cannot change these actions. If the item (a task or a monitor) is running, other users cannot stop it.
Controlled	Other users can start and stop this task or monitor. Only the owner can delete the item or change any properties of this item, including the level of sharing. Other users can also view this item and use it to create a new item based on this one. If you are the owner of a monitor and have specified actions (such as opening the event log window or sounding an alarm on the PC), these actions occur for all users of the monitor whenever a threshold is triggered or reset. The other users cannot change these actions. Any actions that are associated with running a monitor that was created by another user (the owner) runs under the authority of the owner. Therefore, as the owner, you might be sharing a monitor with someone who does not have the same level of authority as you.
Full	Other users can change and delete this definition or system group. Other users can also view this item and use it to create a new definition or system group.

Uses for sharing objects and tasks

What you can do with sharing depends on the needs of your work environment. Consider these examples:

- **You can share job monitors, message monitors, system monitors, and file monitors.**

  When you share monitors, others can use the monitors that you set up to measure the monitored activity on the systems in your network. If you choose **Read-Only** sharing, others can open the monitor and its event log, and they can view the properties of the monitor. If you choose **Controlled** sharing, others can also start or stop the monitor. The level of sharing that you specify when you create a monitor also applies to any events that are logged when a threshold is triggered or reset. You can change the level of sharing for events after they have been logged.

- **You can share system groups.**

  When you share system groups, other users can view the system groups and use them to perform authorized actions. Unless you specify **Full** sharing, you control the endpoint systems in the system group for all authorized users. This ensures that the system group is always up to date. Suppose you created a system group called "West Coast Systems." If you chose to share that group, all system operators can use that system group to work with the West Coast systems. If you specify **Full** sharing, other users may update the contents of that group.
- **You can share definitions.**

  Part of your job might include maintaining a "run book" of commonly used commands. You can share the command definitions in that run book to ensure that the commands your system operators run are accurate. If you need to make a change to one of those commands, you only need to do it once. Your users can share that one set of accurate commands.

  You can also share package definitions, product definitions, and user definitions. By sharing definitions, you save other users the time it takes to create their own definitions.
- **You can share tasks.**

  Tasks are long-running actions in System i Navigator. You can share any actions that have been created and allow users to see the status of tasks. For example, suppose you needed to install 50 fixes on a system group containing 50 systems. If you shared that task, you can start the task and then go home while letting the second shift operators see the status on their PC.
- **You can use global sharing to share all tasks.**

  Use global sharing to specify the level of sharing for all your system administration tasks -- None, Read-Only, or Controlled sharing. You access global sharing through the User Preferences window by right-clicking Management Central. When you specify a value other than None, the sharing value applies to all future tasks that are created with System i Navigator on this PC. Existing tasks are not affected. For example, suppose you are in an environment where you are part of a five-person team that works around the clock. If you chose to globally share your tasks at the Controlled level, your team can see what you did and work with the tasks you started -- even when you are not there.

Synchronizing date and time values

Management Central provides a convenient way for you to synchronize date and time values across your network.

To synchronize the date and time values across your network, select your endpoint systems or system groups whose date and time values you want to update from the **Endpoint Systems** list under Management Central in System i Navigator. Then, right-click any selected system and select **System Values** > **Synchronize Date and Time**. Specify a model system that has the most accurate date and time values.

The date and time system values that are updated on the target systems include system date (QDAYOFWEEK, QDATE, QDAY, QMONTH, QYEAR), time of day (QTIME, QHOUR, QMINUTE, QSECOND), and time zone (QTIMZON). To verify that a time adjustment is being made, select the endpoint system from the list under My Connections (or your active environment) in System i Navigator. Then, go to **Configuration and Service** > **Time Management** > **Time Adjustment** to view the current time adjustment.

The time used from the model system is the software clock time rather than the QTIME system value. The software clock time is the same as the QTIME system value except when the SNTP (Simple Network Time Protocol) client is started on the model system. When SNTP is running on the model system, the software clock is synchronized to the time server specified in the SNTP configuration. For more information about configuring SNTP, see Simple Network Time Protocol (SNTP).

You can choose to synchronize the time without changing the time zone, or synchronize both the time and the time zone with those on the model system.

When a system changes to or from Daylight Saving Time (DST), the GMT offset (QUTCOFFSET) system value is automatically updated from the GMT offset attribute of the time zone (QTIMZON) system value.

Related information:

Simple Network Time Protocol (SNTP)

Synchronizing functions

You can synchronize the configuration of key functions, such as EIM and Kerberos, across a group of endpoint systems.

You select a model endpoint system and a set of target endpoint systems, and then use the Synchronize Functions wizard to duplicate the model system's Kerberos or EIM configurations (or both) on the specified target systems. Synchronizing these functions from the model system saves you time by eliminating the task of individually configuring each function on each target system. Synchronizing your EIM configurations allows you to create EIM associations between user identities within your network. This in turn allows a user who has different profiles on different systems to work with distributed applications that use Kerberos authentication without having to sign on to each of these systems individually.

For example, John Smith may be JSMITH on system CHICAGO1, JOHNSMITH on system DETROIT1, and JRSMITH on system DENVER. If EIM and Kerberos are configured on all three systems, and all three profiles are associated with the same EIM identifier, John Smith can use Management Central to manage these V5R3 systems. For example, he can run commands on these systems, and monitor performance, jobs, and other resources on these systems. John Smith can also access other services and applications that use EIM and Kerberos authentication without the need for multiple passwords to these different systems across the enterprise.

Using Kerberos and EIM together in this way is referred to as *single signon* because it eliminates the need to provide multiple user names and passwords for distributed applications. Single signon benefits users, administrators, and application developers by enabling an easier password management system across multiple platforms without the need to change underlying security policies. See Single signon for details on how to enable single signon by using network authentication service and Enterprise Identity Mapping (EIM).

Note: If the SNTP box is checked then a TCP job QTOTNTP should be running on the endpoint. If it is not running then Management Central will use information from the model system. If SNTP is checked and the client QTOTNTP job is running then you should not run multiple Time Synchronization tasks within one polling interval of the SNTP client. You can view the SNTP polling interval at **My Connections** > **system** > **TCP/IP** > **Right-click SNTP** > **Properties** > **Client tab** .

Related concepts:

"Managing users and groups with Management Central" on page 40
System i Navigator can help you as a system administrator to keep track of the users, groups, and their level of privileges on one or more endpoint systems.

Related information:

Scenario: Configuring the Management Central servers for single sign-on

Propagating system settings from the model system (System A) to System B and System C

Scheduling tasks or jobs with Management Central scheduler

System i Navigator provides two different tools you can use to schedule tasks or jobs: an integrated scheduler (the Management Central scheduler) and the Advanced Job Scheduler.

Management Central scheduler

The Management Central scheduler helps you to organize when you want your tasks to occur. You can choose to perform a task immediately or at a later time.

You can use the Management Central scheduler to schedule a variety of tasks. For example, you can automate the process of collecting an inventory (such as hardware, software, or fixes) on whichever day that fits your operating schedule. You might schedule such a collection to occur every Saturday night at 10 p.m. You can also schedule to clean up the save files and cover letters of the fixes from your systems on the first of every month. Or you might want to install a set of fixes once.

To schedule a later time to perform a task, click **Schedule** from any window in which the button is displayed. Your scheduling information is stored on the central system and submitted there. No scheduling function is needed at the endpoint system. You can then view the scheduled job in one of the Scheduled Tasks containers. You can also view the job by using the Work with job schedule entries (WRKJOBSCDE) command on the character-based interface. Scheduled jobs have a job name of Q*xxxxxxxx* where *xxxxxxxx* can be a hex number such as FFFFFF08.

Important: Do not use the Work with Job Schedule Entries (WRKJOBSCDE) command to alter or delete a scheduled job if that job was scheduled using the Management Central Scheduler or the Advanced Job Scheduler. If the job is altered or deleted by using the WRKJOBSCDE command, Management Central is not notified of the changes. The task might not run as expected, and error messages might be shown in the Management Central server job logs.

The following scheduling options are available from the Management Central scheduler:

- **Daily**

  The task runs every day at the specified time beginning on the specified date.
- **Weekly**

  The task runs every week at the specified time beginning on the specified date. You may either accept the default (today's date) or specify the day of the week when you want the task to run.
- **Monthly**

  The task runs every month at the specified time beginning on the specified date. You may either accept the default (today's date) or specify a day of the month (1-31), first day, or last day.

You can schedule any task for which a **Schedule** button is available. For example, you can schedule a specific time to collect inventory. Tasks that run only once will be removed from the Scheduled Tasks view when they are run. They then appear in a Task Activity folder. (If you want full calendar management, you should use the Advanced Job Scheduler.)

For more information about these and other Management Central tasks and topics, refer to the detailed task help that is available from the System i Navigator window.

What you can do with Management Central scheduler

Using the scheduler function gives you the flexibility to do your work when it is convenient for you. In addition, you can use the Management Central scheduler to do almost any task in Management Central. For example, you can schedule when to do any of the following tasks:

- Run commands on selected endpoint systems and system groups.
- Collect inventory on selected endpoint systems and system groups.
- Collect system values inventory on selected endpoint systems and system groups; then compare and update system values to those on a model system.
- Create, delete, edit, and send users and groups across multiple endpoint systems.
- Send fixes or packages of files and folders to selected endpoint systems and system groups.

- Start installing fixes, uninstall fixes, or install fixes permanently.
- Delete the save files and cover letters for selected fixes on selected endpoint systems and system groups.
- Start and stop collection services on selected endpoint systems and system groups.

You can schedule a task to run once, in which case the task runs a single time beginning at the specified date and time. Tasks that run only once are removed from the Scheduled Tasks container when they run. They then appear in a Task Activity container.

Advanced Job Scheduler

IBM Advanced Job Scheduler for i (5761-JS1) is a separate licensed program that you can install and use to schedule tasks and jobs. This scheduling tool provides more calendar features and offers greater control over scheduled events. If you have Advanced Job Scheduler installed, click the **Schedule** button from any System i Navigator window to schedule tasks and jobs.

After you have installed the plug-in, an Advanced Job Scheduler container is displayed under Management Central. Tasks that are scheduled with the Advanced Job Scheduler are in this container.

You can also use the Work with Jobs using Job Scheduler (WRKJOBJS) command to display jobs that are scheduled with the Advanced Job Scheduler. However, do not delete scheduled Management Central tasks or change the owner from the WRKJOBJS display. If the job is altered or deleted by using the WRKJOBJS command, Management Central is not notified of the changes. The task might not run as expected, and error messages can appear in the Management Central server job logs.

Related information:

Managing job scheduling

Code license and disclaimer information

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

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

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

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

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

Related information for Management Central

Web sites, and other information center topic collections contain information that relates to the Management Central topic collection. You can view or print any of the PDF files.

Web sites

You can use a variety of Web sites to find more information about Management Central. These include:

- System i Navigator 

  System i Navigator provides a wide variety of tools to simplify i5/OS management. Go to the System i Navigator home page to find information about System i Navigator, including functional release overviews, news about technical conferences, and other hot topics. You will find links to a variety of information including release updates, functional overviews, FAQs, and more.

Other information

You can find links from various places in the Management Central topic to other information center topics that relate to Management Central.

- Experience Report: Configuring Management Central Connections for Firewall Environments

  This report details Management Central connections and the configurations required to enable Management Central to operate within a variety of firewall environments. As a distributed management application, Management Central requires numerous incoming and outgoing TCP/IP socket connections. In contrast, the basic premise of a firewall is to restrict/modify incoming and outgoing connections.
- Single sign-on

  If you have been looking for a way to simplify the task of managing user profiles on the System i product, single signon may be the answer for you. This information presents a single signon solution for your system, which uses the technology of Enterprise Identity Mapping (EIM), paired with your system's network authentication service. The single signon solution simplifies the task of managing user profiles, while reducing the number of signons that a user must perform to access multiple applications and servers.

  This topic includes a scenario that demonstrates how to configure an entire system group to participate in a single signon environment. After administrators complete the scenario for propagating a single signon configuration across multiple systems, they can do the necessary configuration so that the entire system group can participate in the single signon environment.

Related reference:

"PDF files for Management Central" on page 1
You can view and print a PDF file of this information.

Appendix. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

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

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
3-2-12, Roppongi, Minato-ku, Tokyo 106-8711

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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

IBM Corporation

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

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

Programming interface information

This Management Central 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 Sun Microsystems, Inc. in the United States, other countries, or both.

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

Terms and conditions

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

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

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

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

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

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

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

IBM Confidential IBM Confidential IBM Confidential IBM Confidential

Printed in USA