


IBM

@server™

Management Central





Management Central

Contents

Chapter 1. Management Central	1
Chapter 2. What's new for V5R1?	3
Chapter 3. Getting started with Management Central	5
Installing and accessing Management Central	5
Setting up your central system	5
Adding endpoint systems with Management Central	6
Creating system groups with Management Central	7
Removing endpoint systems from groups with Management Central	7
Deleting system groups with Management Central	7
Sharing with other users in Management Central	8
What can I do with sharing and Management Central?	8
Chapter 4. Working with Management Central	11
Management Central hardware inventory	12
Management Central software inventory	12
Management Central users and groups inventory	13
Searching a Management Central users and groups inventory	13
Management Central system values inventory	14
Management Central fixes inventory	14
Managing software fixes	15
Obtaining fixes	16
Downloading fixes by using electronic customer support	16
Copying fixes from media with Management Central	16
Installing fixes	17
Comparing and updating fixes with Management Central	18
Management Central status descriptions for fixes	18
Managing software products	20
Packaging products with Management Central	20
Creating a product definition	20
Generating fixes	21
Using Management Central to add support for a product	21
Packaging and sending objects with Management Central	21
Creating a package definition with Management Central	22
Sending a package of files with Management Central	22
Monitoring system performance	23
Monitoring real-time system performance with system monitors	24
Creating a Management Central system monitor	24
Selecting Management Central metrics	25
Changing metric information for Management Central system monitors	26
Setting threshold commands for system monitors	27
Setting threshold actions for Management Central monitors	27
Selecting systems and groups for Management Central system monitors	28
Starting and stopping a Management Central monitor	28
Working with Management Central monitor graphs	29
Tips - Management Central monitor graphs	30
Example: Management Central's CPU Utilization performance metrics	31
The CPU Utilization (Interactive Feature) metric	32
The CPU Utilization (Database Capability) metric	34
The CPU Utilization (Secondary Workloads) metric	35
Uses for Management Central's CPU Utilization metrics	37
Monitoring jobs and servers with Management Central	38

Creating a new job monitor	38
Selecting metrics for job monitors	39
Specifying threshold values for a job monitor	40
Specifying the collection interval for a job monitor	43
Running commands for job monitors	43
Logging events for a job monitor	43
Applying thresholds and actions for a job monitor.	44
Viewing job monitor results	44
Reset triggered threshold for a job monitor	45
Monitoring message queues with Management Central	46
Creating a new message monitor	46
Defining message sets for a message monitor.	46
Running commands for message monitors	47
Logging events for a message monitor	48
Applying thresholds and actions for a message monitor	48
Viewing message monitor results.	49
Reset triggered threshold for a message monitor	49
Event Log	50
Collecting performance data with Collection Services	50
How Collection Services works	51
How to start Collection Services	51
Collection Services tasks.	51
Collection Services and performance database files	52
Collecting trace data	52
Dumping trace data.	52
Customizing data collections with Management Central	53
Time zone considerations for Collection Services	54
Creating database files to use with Management Central's Collection Services	54
Creating database files from an existing collection object	55
Managing collection objects with Management Central	56
Managing system values.	56
Manage users and groups across multiple systems with Management Central	57
Management Central user definitions	58
Creating users across multiple systems	59
Editing users across multiple systems	60
Deleting users across multiple systems	61
Scanning for owned objects with Management Central.	61
Collecting an inventory of users and groups.	61
View an inventory of users and groups	62
Searching an inventory of users and groups	63
Exporting an inventory.	63
Sending users and groups to multiple systems.	64
Synchronizing unique identifiers of users and groups (UID and GID).	65
Warning: Temporary Level 4 Header	65
Synchronizing unique identifiers when creating users or groups	65
Synchronizing unique identifiers when editing users or groups	65
Synchronizing unique identifiers when sending users or groups	65
Running commands with Management Central.	66
Creating a command definition with Management Central.	66
Running a command with Management Central on systems or groups	67
Scheduling tasks or jobs with Management Central scheduler	67
Management Central Scheduler	68
Scheduling jobs with the Advanced Job Scheduler	69
Install and customize the Advanced Job Scheduler	69
Installing Advanced Job Scheduler	69
Customizing the Advanced Job Scheduler	70

Assigning the general properties for the Advanced Job Scheduler.	71
Creating and working with applications for the Advanced Job Scheduler	71
Setting up a calendar for the Advanced Job Scheduler.	72
Setting up a holiday calendar for the Advanced Job Scheduler.	72
Working with library lists for the Advanced Job Scheduler.	73
Working with command variables for the Advanced Job Scheduler	73
Working with job controls for the Advanced Job Scheduler	74
Schedule a job	74
Schedule a job group	74
Job dependencies	74
Monitoring job activity for the Advanced Job Scheduler.	75
Chapter 5. Redbooks.	77

Chapter 1. Management Central

Are you interested in making your system administration tasks simpler, easier, less time-consuming, and much less repetitive? Are you looking to lower your overall total cost of server ownership? Then pay attention to Management Central! Management Central makes managing many servers as easy as managing a single server.

Management Central is a suite of easy-to-use systems management functions that come to you as part of your base operating system. You can use Management Central to manage multiple servers through a single central system. Simply select a server to use as your central system, then add endpoint systems to your Management Central network. You can create groups of similar or related endpoint systems to make managing and monitoring your servers even easier. Your central system will handle the communications for you. You can even take advantage of such options as scheduling and unattended operations. You'll find that Management Central is flexible and easily manipulated to suit your needs.

To make the most of Management Central's flexibility and function, see:

Getting started with Management Central

Start using Management Central! To get the most out of Management Central, set up your central system and endpoint systems in a way that makes sense for your business environment.

Working with Management Central

Find out about all the ways Management Central can help you streamline your server administration tasks, such as managing fixes, monitoring performance, and administering users and groups across multiple servers in your network.

Chapter 2. What's new for V5R1?

The Management Central topic contains information specific to the Management Central functions of Operations Navigator — information you need to effectively manage multiple iSeries and AS/400 servers.

In V5R1, Management Central offers you new and improved ways to manage the servers in your network.

Extreme Support

When you use Extreme Support, AS/400 delivers secure, personalized service and electronic support that is designed to help you keep your business running at peak performance. Through automated support, online tracking of service, and proactive maintenance, AS/400 offers support that is customized to your unique system environment.

Management Central - Pervasive



Administrators now have more flexibility to access Management Central information and monitor the systems they support. Management Central - Pervasive lets you remotely monitor system performance and status using a web phone, a personal digital assistant (PDA) with a wireless modem, or a traditional web browser.

Manage your products

Create and install a product that you have defined. A wizard is available to help you create and install your product. You can distribute and install your own applications. You can also create your own fixes to manage your products.

Work with system values

You can collect a system values inventory on one or more endpoint systems or system groups. You can schedule the collection of this information on a regular basis and store it on a designated central system.

Compare and update system values across multiple systems or a group of systems. You can manage system values quickly and easily across a network of systems running OS/400.

You can also work directly with the individual system values on any system to which you have a direct connection.

Manage users and groups

You can now manage your users and groups across multiple systems using Management Central. You can send, create, edit, or delete users and groups across multiple endpoint systems or system groups.

You can also work with an inventory of the users and groups on one or more endpoint systems. You can even create a user definition and then create multiple users across multiple systems based on the definition.

Monitor jobs and servers

Use Management Central job monitors to stay on top of job activity by monitoring a job or a list of jobs based on job name, job user, job type, subsystem, or server type.

Monitor message queues

Use Management Central message monitors to monitor your message queues for the information you need to manage your servers. For example, you could monitor a message queue to determine whether an application completed successfully. Or you could monitor the system operator message queue for a specific message that indicates when a critical storage condition exists. When you create a monitor, you can specify commands to run when the message is detected.

Enhanced monitor functions

Since V4R3, you have been able to view real-time performance metrics by using the system monitor support. Now in V5R1, you can display a graphical history of performance data. You can view collected data at your convenience instead of only viewing real-time data.

Other enhancements to the monitor support include:

- Changes to the system monitor window and its behavior to enhance usability.
- Ability to access the Performance Management/400 web site from the File menu.
- Ability to start and end Performance Management/400.
- Ability to hold, release, and end jobs from all Management Central monitors.

Enhancements to run command support

You no longer need to dig out your old CL Reference book or call your buddy down the hall to see if he remembers the parameters for the command you want to use. Just type the first few letters of an OS/400 command any place where you enter a command in Operations Navigator, click the **Prompt** button, and a complete command prompter is available to help you complete the command.

Enhancements to Collection Services

You can now collect TCP/IP-related data (TCP/IP Base and TCP/IP Interface). Data will be collected for system-wide performance information for TCP/IP and information for each active TCP/IP interface.

You can collect internal performance explorer data. You can collect data to enhance capacity planning capabilities or for other purposes (PEX Data - Processor Efficiency category).

Chapter 3. Getting started with Management Central

Start using Management Central! Follow a few simple steps to set up your Management Central network:

1. **Install and access Management Central**

Some of the Management Central functions that you will want to use are optionally installable components of Operations Navigator. Be sure you choose to install them when you install Operations Navigator. Find and open Management Central in your Operations Navigator window.

2. **Set up your central system**

You choose your central system when you first start Management Central. You can also change your central system easily at any time.

3. **Add endpoint systems**

Endpoint systems are the other systems in your network that you manage with your single central system. To start Management Central, add some endpoint systems that you want to manage.

4. **Create system groups**

Make the most of Management Central's ability to manage groups of systems. Create groups of systems to make managing them easier and more efficient.

5. **Share Management Central resources**

Sharing resources with other users will help you manage your systems more efficiently. You can share job monitors, message monitors, system groups, definitions, and system administration tasks with other administrators and operators.

When you have finished this preliminary work with Management Central, you're ready to see what you can do with Management Central.

Installing and accessing Management Central

Some functions of Management Central are optionally installable components of Operations Navigator, the graphical user interface (GUI) for iSeries 400. When you install Operations Navigator, be sure to choose to install Operations Navigator Base Support (which includes some of the Management Central functions), plus Configuration and Service, Users and Groups, Commands, Packages and Products, and Monitors.

If you did not install all the components you need when you installed Operations Navigator, do the following:

1. Go to your **Client Access** directory on your PC and double-click **Selective Setup**.
2. Use the Selective Setup wizard to install the additional components that you need for Management Central functions. To get all the Management Central functions, select Configuration and Service, Users and Groups, Commands, Packages and Products, and Monitors.

When you use the Selective Setup wizard, the components you select will be installed. Any components you deselect during the selective setup will be uninstalled. Be careful not to uninstall anything while you use the Selective Setup wizard.

Management Central appears automatically in the tree hierarchy in your Operations Navigator window. All you need to do is expand **Management Central** to access its functions. After you have added endpoint systems and created system groups, those endpoint systems and system groups will appear under Management Central as well.

Setting up your central system

Management Central allows you to manage multiple servers from a single system in a TCP/IP network environment. To do this, you need to have a central system. The other systems in your network are called **endpoint systems**. Once you add endpoint systems to your network, you only need to do your system administration tasks once. Your central system will initiate your tasks and store Management Central data.

Setting up your central system for the first time

You choose and set up your central system the first time you start Management Central. To set up a particular system as your central system, you must do the following:

1. Ensure that you have access to Management Central in Operations Navigator. Access is controlled with the Application Administration function in Operations Navigator.
2. Ensure that the system you want to use as your central system is:
 - Connected through Client Access.
 - Running the OS/400 operating system Version 4, Release 4 (V4R4) or later. For the latest Management Central and Operations Navigator functions, your central system should be running OS/400 Version 5, Release 1 (V5R1).

Changing your central system

Though you select and set up your central system the first time you use Management Central, changing central systems is quick and easy. To change your central system, do this:

1. Right-click Management Central and select **Change Central System**.
2. Use the **Change Central System** dialog 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 Operations Navigator network, right-click **My Connections** and select **Add connection**. When the new system is connected, you may change your central system to the new system.

Once you have set up your central system, you are ready to do the other tasks necessary for setting up Management Central.

Adding endpoint systems with Management Central

An endpoint system is any system in your TCP/IP network that you choose to manage through your central system.

To add endpoint systems for a large network, do the following:, do the following:

1. Right-click **Endpoint Systems** and select **Discover Systems**.
2. Specify the TCP/IP subnets that you want to search.
3. When you click **OK**, any connected systems that are found are added to your network as endpoint systems and the IP addresses of all your endpoint systems are updated.

To manually add one or more endpoint systems, do the following:

1. Right-click **Endpoint Systems** and select **New Endpoint System**.
2. Ensure that the system you want to add as an endpoint system is:
 - Connected to the central system, which means that the central system can access the endpoint system by using TCP/IP
 - Running the OS/400 operating system
3. Enter the name of the system and click **OK**.

That is all there is to it. The endpoint systems that you added appear automatically under **Endpoint Systems** in your Operations Navigator window. Next, you can create system groups to help you manage different sets of endpoint systems. The new system groups will appear in Operations Navigator as well.

Creating system groups with Management Central

A system group is a collection of endpoint systems that you define. Endpoint systems can belong to several system groups at the same time. Once 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 quick steps:

1. Open **Management Central** from your **Operations Navigator** window.
2. Right-click **System Groups** and select **New System Group**.
3. On the **New System Group** dialog, specify a unique name for the new system group. You can also enter a brief description that will help you later 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 will not be able to view or change this system group.
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 go back and add more endpoint systems or remove endpoint systems from your system group. You can even delete system groups from Management Central.

Removing endpoint systems from groups with Management Central

You can remove any endpoint system from a system group you created. (You can also remove endpoint systems from other system groups if the owner of the group has specified full sharing.)

To remove an endpoint system from a system group, do the following:

1. Select the system group from which you want to remove an endpoint system.
2. Right-click the endpoint system that you want to remove from the group.
3. Select **Remove**.

When you remove an endpoint system from a system group, you do not delete the endpoint system from Management Central or from the list of endpoint systems. However, when you delete an endpoint system under **Endpoint Systems**, you remove the endpoint system from all system groups.

You can always go back and work with your system groups, adding endpoint systems, creating new system groups, and even changing your central system if necessary.

Deleting system groups with Management Central

A system group is a collection of endpoint systems that you define. Once you have created a system group, you can manage the entire group from your central system as if it were a single system.

If you decide that the group is no longer useful for your environment, you can always go back and delete system groups from Management Central.

To delete a system group, follow these quick steps:

1. Open **Management Central** from your **Operations Navigator** window.
2. Select **System Groups**. Notice that the system groups you have defined appear in the right pane of your **Operations Navigator** window.
3. Right-click the name of the system group you want to delete and select **Delete**.

4. On the **Confirm Delete** dialog, verify that the system group shown is the one you want to delete.
5. Click the **Delete** button. That's all you have to do!

As soon as you click **Delete**, the system group you selected will be deleted from Management Central. Should you ever need it again, you can always go back and create a new system group using the same combination of endpoint systems.

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. Sharing allows users to use (or share) the same items — job monitors, message monitors, job and message monitor events, system groups, definitions, and system administration tasks. You can even set your user preferences to share **all** the new tasks you create.

To enable sharing, you must be the owner of the particular item you want to share. If you are the owner and want to use sharing, do the following:

1. Right-click the item you want to share and select **Properties**.
2. Select **Sharing**.
3. Set the level of sharing from the **Sharing** page. Choose one of the following levels of sharing as it is applicable to the item you are sharing:

None	Only the user who created the item can see or use it.
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 change this item or delete it. If the item (a task, job monitor, or message monitor) is running, other users cannot stop it.
Controlled	Other users can start and stop this item. Only the owner can 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.
Full	Other users can change and delete this item. Other users can also view this item and use it to create a new definition or system group.

Once you have set your level of sharing, what you can do with sharing is up to you.

What can I do with sharing and Management Central?

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

- **You can share job monitors and message monitors.**
When you share job and message monitors, others can use the monitors that you set up to measure the activity of jobs and messages on the systems in your network. If you choose **Read-Only** sharing, others can only open the monitor and its event log. 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
- **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 could 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 may 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 would 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 would take to create their own definitions.

- **You can share tasks.**

Tasks are long-running actions in Management Central. 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 could start the task and then go home. Let the second shift operator see the status on her PC while you relax at home!

- **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 dialog by right-clicking on Management Central. When you specify a value other than None, the sharing value applies to all future tasks. 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 could see what you did and work with the tasks you started — even when you are not there.

When you use sharing and Management Central, managing your systems becomes a quick, efficient process. See [What can I do with Management Central?](#) to learn more about getting the most out of Management Central.

Chapter 4. Working with Management Central

Use Management Central to streamline your system administration tasks. To make the most of Management Central, first plan your Management Central network. Then you can efficiently complete the tasks required to manage your system. To learn about building a Management Central network, see the getting started topic. To learn how to manage multiple servers easily and efficiently, keep reading.

You can use Management Central's powerful suite of functions to handle all your system administration tasks.

Collect inventory	<p>You can collect and manage inventory for hardware, software, users and groups, system values, and fixes. You can search these inventories for criteria that you specify. You can export the results of the search or the entire inventory to a PC file to perform other queries.</p> <p>You may have other applications installed that allow you to collect lists of other types of resources.</p>
Collect performance data	<p>Use Collection Services to collect system performance data that you can use for later analysis (for example, using Performance Tools for iSeries). See a graphical view of the metrics that you collected for an extended period of time with the Graph History window. You can use the Graph History function as long as you collect data with Collection Services; you do not need to have a system monitor running.</p>
Manage fixes	<p>Keep your fixes (or program temporary fixes, PTFs) current across multiple systems. Use Management Central to efficiently manage fixes. You can send, install, and compare and update fixes.</p>
Manage software products	<p>You can use Management Central to package and send software products to the systems in your network. You can also generate fixes for non-IBM software products.</p>
Monitor system performance and resources	<p>Use Management Central monitors to track what your systems are doing. You can monitor jobs, monitor messages, and monitor system performance. Using the event log, you can also track events created by these monitors. You can also monitor system performance and status remotely with Management Central - Pervasive.</p>
Manage system values	<p>Management Central lets you view, compare, and update system values—everything required to efficiently manage system values and maintain consistency across multiple systems in your network.</p>
Manage users and groups	<p>Keep track of the users, groups, and their privileges on multiple systems using Management Central's user administration function. You can also create, send, edit, and delete users across multiple systems.</p>
Package and send objects	<p>You can package and send objects to systems in your network using Management Central. You can create snapshots of your data to preserve more than one version of the data.</p>
Run commands	<p>You can use Management Central to run commands on multiple systems. For commands that you want to run on a regular basis, simply create a command definition and schedule it to run on any systems in your Management Central network. For assistance anytime you are entering or selecting a command, just click the Prompt button to see a complete list of the parameters and values for any OS/400 command.</p>

Schedule unattended tasks or jobs

Use Management Central's integrated scheduler to automate recurring tasks. You can choose to run a task immediately, or you can use the scheduler to select a later time. You can schedule a task to run just once, or you can schedule it to run daily, weekly, or monthly at a convenient time.

Finally, Management Central makes system administration even easier by allowing you to share Management Central resources with other users. Do not forget to use the online help available to you in Management Central. The online help offers tips and techniques for making the most of Management Central, including "What's This?" help, how-to information, and extended examples.

Management Central hardware inventory

To display the resource, status, and description of all hardware on the endpoint system, select Hardware Inventory (**Management Central** → **Endpoint Systems** → *any endpoint system* → **Configuration and Service** → **Hardware Inventory**). 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). It is recommended that you schedule collection of your hardware inventory on a recurring basis to keep your central system's inventory current.

You can right-click on 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.

Other types of inventory are software inventory, users and groups inventory, system values inventory, and fixes inventory.

Management Central software inventory

To view installed or supported products, expand Software Inventory (**Management Central** → **Endpoint Systems** → *any endpoint system* → **Configuration and Service** → **Software Inventory**).

Installed Products displays a list of the software products that are currently installed on the selected system. You can right-click on any software listed and select **Properties** to view additional information. You can send these products to one or more endpoint systems or system groups and install them on those systems. You can download fixes for an installed product regardless of whether or not its status is "Installed and supported".

Supported Products displays a list of the software products that the selected system currently supports for the other systems that it manages in the network. For example, this list can contain products that are not installed on this system. A system that provides support typically orders the fixes and sends them to systems where the product is installed.

If a product is installed, you can send this product to one or more endpoint systems or system groups and install it on those systems. You can also upgrade a software product that is installed and supported and still have the fixes from the previous release available in a save file. You would be concerned about this on the system you are using for your source system when distributing fixes. This would be necessary if you had to support several different releases within your network.

You can add support for a product whether or not it is installed on your system. When you add a product to the Supported Products list, you can copy save files to the source system for fixes to that product, even though the product is not installed. You can then send (or send and install) these fixes to other systems in your network.

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.

Other types of inventory are hardware inventory, users and groups inventory, system values inventory, and fixes inventory.

Management Central users and groups inventory

You can collect an inventory of users and groups on endpoint systems or system groups. Just like the other Management Central inventories, you can search these inventories and export them to your PC in your choice of formats.

Expand Users and Groups under any endpoint system (**Management Central** —> **Endpoint Systems** —> *any endpoint system* —> **Users and Groups**) and select **User Inventory** to see a list of the users on that system. Or select **Group Inventory** to see a list of the groups on that system. 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 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 under Management Central are automatically updated in the central system inventory.

Other types of inventory are hardware inventory, software inventory, system values inventory, and fixes inventory.

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. 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 were searching for all users on this endpoint system or system group with security officer authority, you could narrow the search to users in your Accounting department with security officer authority by clicking **And** and selecting **Department** and **Accounting**.

From the Search Results window, you can perform many of the actions that you can perform on a user or group elsewhere within Management Central. 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 when both the central system and the endpoint systems are running OS/400 V5R1 or later.

Management Central system values inventory

You can collect an inventory of the system values on any endpoint system that is running OS/400 V5R1 or later. Once you have collected these inventories, use Management Central to compare the system values on a model system to those on selected target systems. You can even choose to update the system values on the target systems to match those on the model system. To compare and update system values, right-click an endpoint system or system group, select **System Values**, and then select **Compare and Update**.

You will want to make sure that your system values inventories are current before doing a compare and update of system values on your systems. The **Compare and Update** window shows the date and time that the system values inventory was last collected on the target systems. You need a current inventory because the inventory data for the endpoints is used to do the compare and update. To collect inventory on a system or group, just right-click the endpoint system or system group, select **Inventory**, and then select **Collect**.

You can also export your system values inventory to a PC file. These PC files provide a history of the inventory and allow you to work with the data in a spreadsheet program or other application. To export a system values inventory, right-click the endpoint system or system group, select **System Values**, and then select **Export**. You can also click the **Export** button from the Compare and Update window.

Other types of inventory are hardware inventory, software inventory, users and groups inventory, and fixes inventory.

Management Central fixes inventory

Fixes inventory lets you periodically correct problems or potential problems found within a particular licensed program. Fixes, which are also known as program temporary fixes (PTFs), can correct problems that appear to be hardware failures or software failures, or they can provide new functions. Fixes are designed to replace one or more objects in the licensed program. You can manage your fixes inventory with the available graphical wizards. For example, use the Compare and Update wizard to automatically compare a group of systems to a model system, find the missing fixes and extra fixes, and send the missing fixes to each system and install them. You can launch the Compare and Update wizard from an endpoint system, a system group, or from a system in your list of connections.

When you select fixes inventory from the **Collect Inventory** dialog, software inventory will be automatically selected as well. You cannot select fixes inventory without including software inventory. The fixes inventory list (**Management Central** → **Endpoint Systems** → *any endpoint system* → **Configuration and Service** → **Fixes Inventory**) shows all products installed and the fixes contained within them. For each fix, you can view the status of the fix and other information such as the ID, associated product, release, or type.

From the fixes inventory list, you can do any of the following:

- Run the wizards to install fixes, send and install fixes, permanently install fixes, or uninstall fixes
- Clean up save files and cover letters
- Perform other advanced functions, such as canceling fix actions
- Schedule when to perform these actions
- Copy fixes from media

Because a collected inventory is used for Management Central tasks, it is important that you have an inventory that is current; therefore, you should collect the fixes inventory on a regular basis. You should also be aware that any changes made from the fixes inventory list are not automatically reflected in the inventory.

Other types of inventory are hardware inventory, software inventory, users and groups inventory, and system values inventory.

Managing software fixes

Keep your software current with fixes! A fix (or program temporary fix, PTF) contains new or changed objects that are used to correct current or potential problems in your software. They can also provide new functions. Fixes replace one or more objects in the licensed program.

Why should I manage fixes with Management Central?

One of the key benefits of Management Central is that it makes managing multiple systems as easy as managing a single system. Management Central simplifies the process of managing fixes with a variety of easy to use tools and descriptions for each fix status. These tools include several wizards to guide you through these tasks:

- Installing fixes
- Sending and installing fixes
- Permanently installing fixes
- Uninstalling fixes
- Comparing and updating fixes

For example, to install multiple fixes, select the fixes you want from the Management Central fixes inventory list and start the Install wizard. Similarly, you will find the Compare and Update wizard to be very helpful. The wizard compares the fix levels of a single system or multiple systems to a model system. You can send the save files of the missing fixes from a source system and then install the fixes to ensure the systems have the same level of fixes. You can launch the Compare and Update wizard from an endpoint system, a system group, or from a system in your list of connections.

How can I effectively manage fixes for software and licensed programs with Management Central?

Use the tasks that follow to help you manage your fixes:

- **Obtain fixes**
Find out how to download fixes from the internet using IBM Electronic Customer Support (ECS) or copy fixes from media. You can download or copy your fixes onto your source system. To use the Compare and Update wizard, you need to have on your source system the save files for all of the systems which are installed on your model system.
- **Install fixes**
Install and permanently install fixes with Management Central. You can also send them from a system that has the save files for the fixes and install the fixes on another system in your Management Central network. Refer to descriptions for each fix status to see a list of the Management Central statuses.
- **Compare and update fixes**
Use Management Central to compare fixes between a model system and any system in your Management Central network. You can find extra fixes or missing fixes when you run the compare. You can then update the systems in your network so that they have the same level of fixes found in your model system.
- **Clean up fixes**
Over time, save files and cover letters for fixes tend to accumulate, particularly on your source system. Management Central provides you with a way to delete save files and cover letters for fixes that are no longer needed.
- **Display cover letters**
Identify which problems are fixed or special instructions for installation.

To learn more about Management Central's capabilities, see the Management Central topic in the Information Center.

Obtaining fixes

Fixes are created in response to problems that occur with the OS/400 operating system, to other IBM licensed programs, and to your own products that you create with Management Central. IBM creates the fixes, but you have the responsibility to obtain and apply them to your systems. You can obtain fix save files for your source system in the following ways:

- Downloading the fix by using electronic customer support (ECS).
- Copying fixes from media.
- Ordering fixes from the Internet. In fact, you can select fixes, order fixes, and download fixes from the

AS/400 Technical Support  website (located at as400service.ibm.com).

You should use these methods only with individual fixes and not with cumulative fix package distributions. The cumulative package installation process does not create save files and is not designed for use over a network.

After you download your fixes onto your source system, you can use the Compare and Update wizard. To use the Compare and Update wizard, you need to have on your source system the save files for all of the systems which are installed on your model system.

See ordering and applying program temporary fixes for information as it relates to what you need to know about fixes before you begin the installation process.

To learn more about Management Central's capabilities, see the Management Central topic in the Information Center.

Downloading fixes by using electronic customer support

Management Central does not directly support electronic customer support. Instead you should use the Send PTF Order (SNDPTFORD) command. You may want to use the SNDPTFORD command from the system that you will use as your source system when you compare and update fix levels. The Send PTF Order (SNDPTFORD) command allows you to prepare an order for:

- A corrective or preventive PTF package
- Summary information for available PTFs
- Preventive Service Planning (PSP) information

Use Management Central to do more than manage your fixes. You can do many of the tasks required to manage your systems quickly and efficiently with this powerful tool.

Copying fixes from media with Management Central

Management Central provides the ability to copy fixes from media, which facilitates the loading of fix save files into service on your source system. To copy fixes from media, follow these steps:

1. Expand **Management Central**.
2. Expand the system onto which you want to copy the fix save files for distribution. This is the system you have chosen to be your model system. It can also be your model system.
3. Expand **Configuration and Services**, and then expand **Fixes Inventory**.
4. Right-click **All products** or the product for which you want to copy fixes.
5. Select **Copy from media**.
6. Complete the fields by using the online help.
7. Click **OK**. You see the Copying from Media window as the fixes are copied into save files.

After you have copied the fixes, you should collect your fixes inventory again. Just right-click your source system, select **Inventory**, and then select **Collect**. Once you have collected the inventory, you can install the fixes or distribute them to other systems.

Remember, you can use Management Central to do more than manage your fixes. You can do many of the tasks required to manage your systems quickly and efficiently with this powerful tool.

Installing fixes

You can use Management Central to install an individual fix on your system if:

- It is present on your system in the form of a save file (and the corresponding licensed program is installed on your system)
- It was loaded on your system using the Load PTF (LODPTF) command
- It was temporarily removed (uninstalled) from your system.
- You sent the fix without installing it.

In all these cases, before the fix is installed, Management Central reports the status of the fix as available.

If the corresponding licensed program is not on the system, the fix can be present only as a save file and cannot be installed on the system. The fix has a status then of supported only, which means that you can only send it to and install it on other iSeries servers that do have the product installed. To send fixes to an endpoint system, the save files containing the fixes must be present on the source system. These fixes will have a status of either available or supported only.

Management Central provides more than one path to install fixes. See the online help for a complete list of all those paths. One way to install a fix would be to do the following:

1. In Operations Navigator, expand **Management Central**.
2. Expand **Endpoint systems**.
3. Expand the system or group where the fixes you want to install are stored.
4. Expand **Configuration and Service**.
5. Expand **Fixes Inventory**.
6. Select the fixes that you want to install, right-click, and select **Install**. Be sure the fixes you select are **Available** (as shown in the Status column).
7. Follow the instructions of the wizard.

Management Central has its own descriptions for each fix status. You probably are not as familiar with these statuses as you are with the fix statuses that are available from Display PTF (DSPPTF) command.

You will find the **Compare and Update** wizard to be very helpful. This wizard compares the fix level of a single system or multiple systems to a model system and then allows you to install the fixes on the selected target systems so that they have the same level of installed fixes as the model system. When you use the collected inventory, you can have the system compare the fix levels for a group of systems to a model system, send the save files of the missing fixes from a source system, and then install the fixes.

Before you install or remove fixes, you should have the following available:

- A current backup of your user data
- A current backup of your operating system and licensed programs, or at least a backup taken since the last time you applied or removed fixes.

To learn more about Management Central's capabilities, see the Management Central topic in the Information Center.

Comparing and updating fixes with Management Central

Management Central provides tools and wizards to help you manage your fixes. You can take advantage of these tools and wizards by comparing and updating the levels of your fixes inventory.

To compare and update the levels of your fixes inventory, you need to define a model system and a source system. Your source system will have the save files on it. You will use your model system to compare against other systems in your network to ensure that your other systems have the same level of fixes like the model system. The Compare and Update wizard finds missing fixes and extra fixes on the target systems.

1. Set up your model system

Set up a model system that has the appropriate fixes installed for the particular products, all fixes for all products, or fixes for particular releases. In some instances, your model system might be your central system. You should define a model system that works best in your environment. To set up your model system, follow these steps:

- a. Determine which fixes you want installed on the model system.
- b. Install those fixes.
- c. Verify that the save file exists for the fixes on the model system.

2. Set up your source system

Set up a source system that has the save files on it. In some instances, your source system might be your model system. Get the save files to the source system by using the copy from media function.

3. Refresh your inventory

The comparison is done based on the information in the inventory, and now that you set up your model system, you may want to refresh the inventory at this time. If you do not refresh your inventory now, the Compare and Update wizard gives you the opportunity to refresh the inventory.

You can have the wizard perform a comparison, and optionally, send missing fixes, or send and install missing fixes after the compare has completed. As a general reminder, because a collected inventory is used to perform this task, it is important that you have an inventory that is current. You should collect your fixes inventory on all systems before you perform the compare and update task.

To **compare and update fixes on your target systems**, follow these steps:

1. In Operations Navigator, expand **Management Central**.
2. Expand **Endpoint Systems** or **System Groups**.
3. Right-click a system or a group and select **Fixes**, and then **Compare and Update**.
4. Use the Compare and Update wizard to determine what fixes are missing from the target system when compared to your model system. When you have finished, Management Central can send or send and install any missing fixes on the target system. When the missing fixes are installed, the target systems then have the same level of installed fixes as the model system.

Note: Only those fixes identified as missing can be sent and installed. You cannot uninstall extra fixes. You can only display them.

To learn more about Management Central's capabilities, see the Management Central topic in the Information Center.

Management Central status descriptions for fixes

The following table provides a detailed description of the fix statuses that are available in Management Central. Keep in mind that for those statuses that end with the words "action pending," the status represents either the ACN or PND suffixes that you find in the corresponding status description for the same fix when you use the Display PTF (DSPPTF) command.

Status	Description
--------	-------------

Available	The fix is ready to install on the local system. It either exists as a save file, or was loaded (by using the LODPTF command) but not applied (by using the APYPTF command), or has been temporarily removed. You cannot distribute the fix to other systems if it does not exist as a save file.
Cover letter	The cover letter for the fix is on the system, but the fix is not on the system.
Damaged	The fix is damaged. If you have the save file, you can uninstall or install the fix again. If you do not have the save file, you must get the save file and install or uninstall the fix again.
Install at next restart	The fix will be installed the next time the system is restarted.
Install permanently at next restart	The fix will be installed permanently the next time the system is restarted.
Installed	The fix is installed. You can either uninstall it from the system or install it permanently. The fix is not a permanent part of the system.
Installed permanently.	The fix is installed permanently. You cannot uninstall it. Permanently installing a fix means that you can no longer revert to the old objects. The fix is now a permanent part of the system until you supersede it. The only way to revert to the old objects is to reinstall the operating system from a system save that was made prior to permanently installing the fix. This action will be time consuming, so make sure that you are satisfied with any fix before you install it permanently.
Installed permanently - action pending	The fix is installed permanently, but you need to perform an action before the fix is active. Look at the cover letter to determine the required actions. If you have done the actions necessary to make the fix active, you do not have to restart the system at this time. The action pending status will be updated the next time that the system is started.
Installed - action pending	Indicates that the fix is installed, but that you need to perform an action before the fix is active. Look at the cover letter to determine the required actions. If you have done the actions necessary to make the fix active, you do not have to restart the system at this time. The action pending status will be updated the next time that the system is started.
On order	The fix has been ordered but has not yet arrived on the system.
Superseded	The fix has been replaced by a later fix.
Supported only	The fix exists on the system as a save file but cannot be installed on the system. A supported only fix can only be distributed to and installed on other systems.
Uninstall at next restart	The fix will be uninstalled the next time the system is restarted.
Uninstall permanently at next restart	The fix is installed or available and will be uninstalled permanently the next time the system is restarted.
Uninstalled permanently - action pending	The fix is uninstalled permanently, but you need to perform an action before the fix is no longer active. Look at the cover letter to determine the required actions. If you have done the actions necessary, you do not have to restart the system at this time. The action pending status will be updated the next time that the system is started.
Uninstalled - action pending	Indicates that the fix is uninstalled, but that you need to perform an action before the fix is no longer active. Look at the cover letter to determine the required actions. If you have done the actions necessary, you do not have to restart the system at this time. The action pending status will be updated the next time that the system is started.

Use Management Central to do more than manage your fixes. You can do many of the tasks required to manage your systems quickly and efficiently with this powerful tool.

Managing software products

Managing a product is easy with the Management Central product function or with the InstallShield Java Edition. A product is an application program that was assembled by using either the Management Central packaging function or the System Manager licensed program (SM1). The iSeries server provides management functions for software that is identified as a product. To use the management functions for your own software, the software must be identified to the server as a product.

- **Management Central product function**

Use the Management Central wizard to create and install a product that you define. Distribute and install your own applications. You can also create fixes to manage these products.

- **InstallShield Java Edition** 

With InstallShield Java Edition installed on either an iSeries server or a client workstation under a Java Virtual Machine (JVM), you can build Java install packages that can be installed to any operating system that supports Java, including the iSeries servers. Visit the InstallShield Java Edition website to find out how to enable your iSeries server to take advantage of this cross-platform install tool.

To learn more about Management Central's capabilities, see the Management Central topic in the Information Center.

Packaging products with Management Central

If you have your own applications that you would like to package, distribute, and collect inventory for, then the product function that is available in Management Central can help you. The advantage to packaging your application with Management Central is that you can distribute your applications and fixes in the same way that you distribute iSeries licensed programs. You can track your fixes the same way as you track IBM fixes. A Management Central product definition contains all the information that you need to send and install a product across multiple systems.

You need to create a Management Central product definition before you can package, or convert, your application into a product that your server recognizes as a product. The source system on which you create the Management Central product definition is used to manage the product. To package and distribute a product, do the following:

1. Create a product definition.
2. Install the product on the source system.
3. Send and install the product on other systems.

After you have installed the product on the source system, you can generate fixes to that product.

To learn more about Management Central's capabilities, see the Management Central topic in the Information Center.

Creating a product definition

To create a new product definition, follow these steps:

1. In Operations Navigator, expand **Management Central**.
2. Expand **Definitions**.
3. Right-click **Product** and select **New Definition**.
4. Follow the wizard's instructions for creating a new product definition.

The source system on which you create the Management Central product definition is used to manage the product. Once you create a Management Central product definition and install it on the source system, you can send and install the product on other systems and generate fixes on the source system.

To learn more about Management Central's capabilities, see the Management Central topic in the Information Center.

Generating fixes

Once you have installed a product, you can encounter situations when you need to provide fixes to correct problems or potential problems found within your installed product.

To generate fixes that can be used to change your installed product, follow these steps:

1. In Operations Navigator, expand **Management Central**.
2. Expand **Definitions**.
3. Select **Product**.
4. Right-click the product definition that you want to generate fixes for and select **Generate Fix**. The product definition must have a status of **Installed** or **Managed**.
5. Specify the appropriate information for the **Generate Fix** dialog. Click **OK**. A fix is created on the source system, but the fix is not yet installed.
6. Use Management Central to install the fixes.

For information about what criteria an object must meet to be part of a fix, see the Create Program Temporary Fix API available from the Software product exit programs.

To learn more about Management Central's capabilities, see the Management Central topic in the Information Center.

Using Management Central to add support for a product

Using Management Central, you can add support for a product that may or may not be installed on a system. To add support for a product, do the following:

1. Expand **Endpoint Systems**.
2. Expand the endpoint system where you want to add support for a product.
3. Expand **Configuration and Service**.
4. Expand **Software Inventory**.
5. Right-click **Supported Products** and select **Add Support**.
6. Click **Browse** to select from a list of all products in the central system inventory. When you select products from the list, the rest of the information is filled in for you.
7. When you have completed the appropriate fields, click **OK**.

Packaging and sending objects with Management Central

Use Management Central to package and send files and programs! A package definition allows you to group together a set of OS/400 objects or Integrated File System (IFS) 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.

Why should I use Management Central to package and send objects?

When you create a package definition, it is saved and can be reused at any time to send the defined set of files and folders to multiple endpoint systems or system groups. If you choose to 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 package definitions 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.

Use Management Central to do more than package and send objects. You can do many of the tasks required to manage your systems quickly and efficiently with this powerful tool.

Creating a package definition with Management Central

Management Central allows you to package and send objects. You first need to create a package definition before you can send your package. To create a package definition, follow these steps:

1. In Operations Navigator, expand **Management Central**.
2. Expand **Definitions**.
3. Right-click **Package** and select **New Definition**.
4. Specify a name and a brief description for the package definition.
5. Select the **Source system**.
6. Click **Add** and select the files and objects you want in the package definition.
7. If you want to save this version of the objects, check **Create snapshot** to create a snapshot of the selected files.
8. Select the **Sharing** tab to specify whether you want to share this package definition with other users.
9. Select the **Options** tab to specify additional options for this package definition. For example, you can specify whether to keep existing files or replace them when a file already exists on the target system.
10. Select the **Actions** tab to specify a command to be run on the target system when this package has been successfully sent. You can click **Prompt** to get assistance in entering or selecting a command.
11. Click **OK**.

Sending a package of files with Management Central

Management Central allows you to package and send objects. After you have created your package definition, you can send your package. To send your package, follow these steps:

1. Right-click the package definition that you want to send and select **Send**.
2. Select the endpoint systems or system groups to which you want to send the files or folders.
3. Click **OK** to start the send task immediately or click **Schedule** to specify when you want the task to start.


You can also send files and folders without creating a definition by following these steps:

1. Expand **My Connections**.
2. Expand the system from which you want to choose the files and folders to send.
3. Expand **File Systems**.
4. Expand **Integrated File System**.
5. Expand **Root** (or whichever file system you want).
6. Click the folder you want to select from.
7. Select the files or folders from the right pane.
8. Right-click any of the selected files or folders and select **Send**.
9. Select the endpoint systems or system groups to which you want to send the files or folders.
10. If you want to specify any options for this package, click the **Options** tab. You can specify whether or not to include subfolders in the package. You can also specify whether to keep or replace any file that already exists on the target system.
11. Click **OK** to start the send task immediately or click **Schedule** to specify when you want the task to start.

Monitoring system performance

IBM offers a variety of tools to help you monitor your system performance. For example, Management Central gives you the ability to graphically view the real-time performance on selected endpoint systems. The Performance Tools reports provide a way for you to effectively research areas of the system that are causing performance problems. After you have collected performance data over time, different reports offer you ways to see how and where system resources are used.

Take a moment to study the following table to see what tools are available to help you monitor performance and system activity:

Tools:	Capabilities:
Management Central system monitors	See real-time system performance data. Collect and display performance data as it happens or up to 1 hour. Detailed graphs help you visualize what is going on with your servers as it happens. Choose from a variety of metrics (performance measurements) to pinpoint specific aspects of system performance. Drill down into the graphs to get detailed real-time performance data. Drill down and work directly with your jobs.
Management Central job monitors	See real-time job performance data. Monitor a job or a list of jobs based on job name, job user, job type, subsystem, or server type. Choose from a variety of metrics to monitor the performance, status, or error messages for a job. Drill down and work directly with your jobs.
Management Central message monitors	Monitor your message queues. Find out whether your application completes successfully or monitor for specific messages that are critical to your business needs. From the Monitor window, you can see the details of a message, reply to a message, send a message, and delete a message.
Management Central Graph History	See near real-time data or historical performance data. See a graphical view of days, weeks, months, or years for the metrics that you collected. You do not need to have a system monitor running; as long as you use Collection Services to collect data, you can view the Graph History window.
Collection Services	Collect detailed performance data for later analysis. You can use Collection Services to collect multiple sets of performance data for comparison.
Performance Management/400	Receive graphs and reports when you transmit data. PM/400 automatically collects performance data by starting Collection Services. PM/400 collects the data and transmits the data to IBM. Graphs are created from the data and are viewable from the Web. PM/400 also enables you to view performance data for a longer period of time with the Graph History capabilities in Management Central.
Management Central Pervasive 	Performance monitoring goes wireless! Management Central - Pervasive lets you remotely monitor system performance and status using a web phone, a personal digital assistant (PDA) with a wireless modem, or a traditional web browser.
WRKSYSSTS command	View the current status of the system. It displays such things as the number of jobs currently in the system, the total capacity of the system auxiliary storage pool (ASP), the percentage of the system ASP currently in use, and the amount of auxiliary storage currently in use.
WRKSYSACT command	Collect and view performance data in a real-time manner. You view data for any jobs or tasks that were active during the last sample interval. Active means that the job or task consumed CPU. The data consists of CPU utilizations, synchronous and asynchronous I/O counts, storage amounts, and more.

See the performance overview topic if you would like more information about what iSeries performance encompasses.

Monitoring real-time system performance with system monitors

System monitors gather and present real-time performance data for your systems. You can use monitors to see your system performance as it happens across multiple systems and groups of systems. In contrast, you can use Collection Services to collect performance data for later analysis. Monitor graphs show your immediate system performance data, up to 1 hour, and you can use the Graph History function to see graphs of previous performance. Collection Services allows you to analyze multiple sets of true performance data for a longer period of your system performance history. You can see a graphical view of the metrics that you collected for an extended period of time with the Graph History window. You can use the Graph History function as long as you collect data with Collection Services; you do not need to have a system monitor running.

The system monitor graphs present system performance data in an easy-to-use graphical interface that you can directly manipulate to get different or more detailed data. Monitors allow you to collect performance data simultaneously for a wide variety of system metrics, for any system or system group, and for any length of time. Once you start a monitor, you are free to do other tasks on your server, in Operations Navigator, or on your PC. In fact, you could turn your PC off! It will continue to monitor your systems and perform any threshold commands or actions you specified. Your monitor will run until you decide to stop it. To effectively monitor real-time system performance, create a system monitor.

To get started working with system monitors, choose one of the following topics:

- **Creating a new system monitor**
Get step-by-step help through the process of creating a new monitor. Find information about performance metrics and how to set thresholds, threshold actions, and threshold commands.
- **Working with system monitor graphs**
View your system performance data in an easy-to-use graphical interface. To get the most out of your monitor graph, learn to find more detailed data, customize your display, and use the graph itself to change monitor properties.
- **Management Central's CPU Utilization metrics**
You can use single metrics or many metrics from the suite of performance metrics to target different types of performance data. This section introduces the variety of CPU Utilization metrics and ways you can use them to meet your information needs.

Finally, find some monitor tips to make monitoring system performance even easier. Create monitors and customize monitor graphs to reflect the way you manage your systems.

Creating a Management Central system monitor

Management Central system monitors are highly interactive tools you can use to gather and display real-time performance data for your endpoint systems. Creating a new monitor is a quick and easy process that begins at the **New Monitor** window. In Management Central, select **Monitors**, right-click **System**, and then select **New Monitor**. Then follow these steps:

1. **Specify a monitor name**
From the **New Monitor-General** page specify a name for your monitor. Provide a brief description so you can find the monitor in a list of monitors.
2. **Select metrics**
Use the **New Monitor-Metrics** page to select your metrics. You can monitor any number of metrics on any number of endpoint systems or system groups.
3. **View and change your metric information**
Use the **New Monitor-Metrics** page to edit the properties for each metric. You can edit the collection interval, maximum graphing value, and display time for each metric you select.

4. **Set threshold commands**
Use the **Thresholds** tab on the **Metrics** page to enable thresholds and specify commands to run on the endpoint system whenever thresholds are triggered or reset.
5. **Set threshold actions**
Use the **New Monitor-Actions** page to specify the actions you want to occur when a metric threshold is triggered or reset.
6. **Select your systems and groups**
Use the **New Monitor-Systems and Groups** page to select on which endpoint systems or system groups you want to start a monitor.
7. **Run the monitor**
Once you create your monitor, starting and stopping it are easy tasks to do.

After you have created your monitor, you are ready to begin working with monitor graphs. You may also want to explore further information that will help you understand performance metrics by taking a closer look at the Management Central CPU utilization metrics.

Selecting Management Central metrics: To effectively monitor system performance, you must decide which aspects of system performance you want to monitor. Management Central offers a variety of performance measurements, known as **metrics**, to help you pinpoint different aspects of system performance.

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, select **Monitors**, right-click **System**, and then select **New Monitor**. Fill in the required fields, and then click the **Metrics** tab.

When you create a monitor, you can use any metric, a group of metrics, or all the metrics from the list to be included in your monitor. Metric types you can use in your monitor include the following:

<p>Metric groups: CPU Utilization</p>	<p>Metric description: The percentage of available processing unit time consumed by jobs on your system. Choose from the following types of CPU Utilization metrics for use in your monitors:</p> <ul style="list-style-type: none"> • CPU Utilization (Average) • CPU Utilization (Interactive Jobs) • CPU Utilization (Interactive Feature) • CPU Utilization (Database Capability) • CPU Utilization (Secondary Workloads) • CPU Utilization Basic (Average) <p>To learn more about these metrics and how to use them, see Management Central's CPU Utilization metrics. You can also consult the online help available on the General tab of the New Monitor window or the Monitor Properties window in Management Central.</p>
<p>Interactive Response Time (Average and Maximum)</p>	<p>The response time that interactive jobs experience on your system.</p>
<p>Transaction Rate (Average)</p>	<p>The number of transactions per second completed by all jobs on your system.</p>

**Transaction Rate
(Interactive)**

The number of transactions per second completed on your system by the following types of jobs:

- Interactive
- Multiple requester terminal (MRT)
- System 36 environment interactive
- Pass-through

Batch Logical Database I/O

The average number of logical database input/output (I/O) operations currently performed by batch jobs on the system.

**Disk Arm Utilization
(Average and Maximum)**

The percentage of disk arm capacity currently used on your system during the time you collect the data.

**Disk Storage
(Average and Maximum)**

The percentage of disk arm storage that is full on your system during the time you collect the data.

**Disk IOP Utilization
(Average and Maximum)**

How busy the disk input/output processors (IOPs) are on your system during the time you collect the data.

**Communications IOP Utilization
(Maximum and Average)**

How busy the communications input/output processors (IOPs) are on your system during the time you collect the data.

**Communications Line Utilization
(Average and Maximum)**

The amount of data that was actually sent and received on all your system communication lines.

**LAN Utilization
(Maximum and Average)**

The amount of data that was actually sent and received on all your local area network (LAN) communication lines.

Machine Pool Faults

The number of faults per second occurring in the machine pool on the system.

**User Pool Faults
(Maximum and Average)**

The number of faults per second occurring in all of the user pools on the system.

If you need more help, click the **Help** button on the **New Monitor-Metrics** window. Once you become familiar with the Management Central metrics, which metrics you select will depend on the information needs of your computing environment. After you have selected metrics that target the information you are trying to see, you are ready to view and change detailed metric information for each metric you selected for your monitor.

Changing metric information for Management Central system monitors: You can customize your Management Central system monitors by changing information about the performance metrics you choose to use in your monitors. The **New Monitor-Metrics** page allows you to view and change detailed information for each metric. To access this page, select **Monitors**, right-click **System**, and then select **New Monitor**. Fill in the required information, and click the **Metrics** tab.

When you are done working with the metric information on the **New Monitor-Metrics** page, click the **Thresholds** tab, which allows you to enable your thresholds and to set your threshold commands to run on the endpoint system when thresholds are triggered or reset. You can also use the **New Monitor-Actions** page to set threshold actions to occur automatically when thresholds are triggered or reset.

To edit metric information, simply select the metric you want to edit from the list on the **New Monitor-Metrics** page and change any of the following properties:

- **Collection interval**
This information specifies how often metric information is collected. Increase or decrease this number as needed.
- **Maximum graphing value**
This information specifies the highest value that appears on the vertical axis of the monitor graph for this metric.
- **Display time**
This information specifies the number of minutes that appears on the horizontal axis of the monitor graph for this metric.

Setting threshold commands for system monitors: When you create a new monitor, you can choose to run commands on endpoint systems when thresholds are triggered or reset. 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 (except logging an event, which occurs on your central system), while threshold commands run on your endpoint systems.

What can I do with threshold commands?

Use threshold settings to automate any command you want to run when thresholds are triggered or reset. For example, you can set a command that stops any new job from starting when CPU utilization reaches 90%. You can then set another command that allows new jobs to start when CPU utilization falls back to 70%.

In another situation, you may have a monitor that is collecting data on average CPU utilization for a particular system. You can set thresholds and specify commands to keep the average CPU utilization between 20% and 90% or any boundaries you choose. 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 Threshold** option. You can then use this window to enter any commands you want to run when the threshold trigger and reset values are 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 click the **Prompt** button to assist you in entering or selecting a command. You can also type a number of special parameters to pass information to the command such as the time and actual value of the metric.

Setting threshold actions for Management Central monitors: When you create a new monitor, you can specify actions you want to occur on your PC when a threshold is triggered or reset. A **threshold** is a setting for a metric that is being collected by a monitor. **Threshold actions** happen on your PC to notify you when threshold events occur. For example, you can choose to open monitor graphs automatically when thresholds are triggered.

Threshold actions are different from any threshold commands you may have set. Threshold commands run on your endpoint systems, while threshold actions occur on your PC (except logging an event, which occurs on your central system).

What can I do with threshold actions?

Threshold actions allow you to determine what you want to happen when your metric reaches a certain value (called the **trigger** value). You can also specify what you want to happen when your metric reaches a second value (called the **reset** value). You can choose to add events to the **Event Log** on the central system when your metrics reach trigger or reset values. You can set alarms on your PC. You can even set monitor graphs to open automatically on your PC when your metrics reach trigger or reset threshold values.

How do I set threshold actions?

In Management Central, select **Monitors**, right-click **System**, and then select **New Monitor**. Use the **New Monitor-Actions** page to set threshold trigger actions and threshold reset actions. You can automate any of the threshold trigger actions and threshold reset actions shown below.

Trigger action and reset action:	Result:
Log event	Adds an entry to the Event Log on the central system, which indicates that the threshold was triggered.
Open event log	Displays the Event Log on your PC when a threshold trigger occurs. Much like the Open Monitor function, this action opens the Event Log only when you really need it.
Open monitor	Displays the monitor graph when system performance reaches a threshold trigger for a particular metric. This allows you to see a graphical view of your system performance data as it is being collected. You do not have to keep the monitor graph open on your PC all the time. It will open automatically if you select this action, and you can keep the monitor graph open even if you close Operations Navigator.
Sound alarm	Sounds an alarm on your PC when system performance reaches a trigger value.

When you finish setting threshold actions, be sure to turn on your thresholds. To turn on your thresholds, go to the **New Monitor-Metrics** window and click the **Threshold** tab. Use the **Enable threshold** option to enable your thresholds.

Finally, click the **General** tab on the **New Monitor-Metrics** page to view and change metric information. Use this window to specify properties for each metric you selected.

Selecting systems and groups for Management Central system monitors: You can specify on which systems and groups you want to run your system monitors. You can also add endpoint systems or system groups to a monitor that is already started. The **New Monitor-Systems and Groups** page allows you to select the systems or groups of systems on which you want to start your monitors. To access this page, right-click **Monitors**, select **System**, and then select **New Monitor**. Click the **Systems and Groups** tab and click **Browse** to select the systems or groups that you want to add to the monitor from the **Available systems and groups** list.

When you create a monitor, you can use any metric to be included in your monitor.

Starting and stopping a Management Central monitor: When you create a monitor by selecting metrics and setting threshold actions, the next step is to run the monitor on your endpoint systems. Starting and stopping monitors are easy tasks. In **Management Central**, select the monitor you want to start and use the start and stop icons in your toolbar. You can also follow these simple steps:

To start a monitor:

1. In **Management Central**, select **Monitors**, and then select **System**.
2. Right-click the monitor you want to start and select **Start**.
3. Select the systems or system groups on which you want to run the monitor and click **OK**.

To stop a monitor:

1. In **Management Central**, select **Monitors**, and then select **System**.
2. Right-click the monitor you want to stop and select **Stop**.

You do not have to stop the monitor to change what is monitored. You can, therefore, change the metrics, collection intervals, thresholds or even the systems without stopping it first.

Viewing status information

Once you start or stop a monitor, you can get information about the status of the monitor. When you select **Monitors**, and then select **System**, you will see monitor status information in the **Status** column of your monitors list. Status information is refreshed automatically. For more detailed status information, right-click a monitor and select **Status**.

You can also watch a monitor's icon in the list of monitors for quick status information. A green icon means the monitor is running successfully. A red icon means that the monitor failed to start. A yellow icon means that a threshold was triggered.

When your monitor is up and running, open the monitor to work with the monitor graphs to see real-time system performance data in an easy-to-use graphical interface.

Working with Management Central monitor graphs

After you create a new system monitor, double-click the monitor (or right-click and select **Open**) to open the monitor graph.

To interact with your system performance data, you can directly manipulate graph elements. You can also use the graphs' many visual cues to help you identify important system performance areas. You can use monitor graphs to find more detailed data, to change monitor properties, and to customize your graph display to target the information you want to see.

You can select a collection point on the graph to see the information that was summarized to calculate the collection point. You can also work directly with a job and perform tasks such as holding and releasing a job.

Finding more detailed data

You can use monitor graphs to find more information about system performance. For example, use your mouse to hover over a collection point (or data point) within a graph. You can see information about the exact time and the endpoint system for which the data was collected. If you click on any point in the graph, you will see detailed information for that data point in the upper-right **Details** pane of the monitor window. The lower-right **Properties** pane shows properties for the information you select in the **Details** pane. You can select either **Status** from the **File** menu or click the **Legend** button to get detailed information about the status of the monitor for each system that is being monitored.

Changing monitor properties using graphs

You can use monitor graphs to change your monitors. For example:

To change this:

Do this:

Threshold trigger and reset values	Select and drag threshold indicators on your monitor graph to new trigger or reset values. You can also right-click a monitor graph and select Thresholds .
Endpoint systems to a monitor	From the File menu, select Properties , and then select the Systems and Groups tab.
Monitor metrics	From the File menu, select Properties .

Customizing the graph's display

Use the built-in flexibility of Management Central system monitors to customize your graph display. For example, try the following:

- **Create different graph line colors for each endpoint system**
Your monitor graphs can become busy when you are monitoring metrics for many systems or groups. It can be difficult to determine which line belongs to which system. You can change the colors of graph lines and pick specific colors for specific systems. In Operations Navigator, select **User Preferences** from the **Options** menu. Use this window to set line colors for each of your endpoint systems.
- **Sort the bars on your monitor graphs**
You can sort the bars that appear in the monitor graphs **Details** pane in a way that makes sense to you. From the **View** menu, select **Sort Details** to arrange the properties in the **Details** pane by **Name** or by **Value**. You can also select **Normalize Details** to sort the bars by the scale the monitor graph uses or by the largest value present.
- **Change the Details pane**
By default, the values in the Details pane are normalized to the largest value. For example, this means that the job with the highest CPU utilization during the selected interval is represented by a bar across the entire width of the pane. The other jobs then display with proportionately less CPU utilization. You can change these values so that the actual CPU percentage used by each job is shown. To make this change, right-click on the **Details** pane and select **Normalize Details**. Then select **To Graph Scale**.
- **Size or collapse panes in the Monitor window**
You can resize the panes in the **Monitor** window to show only the information you want to see. For example, if you want to see only the graphs (and not the **Details** or **Properties** panes), resize the graphs panes so they fill the entire **Monitors** window. You can select the **View** menu to choose **Layout** which allows you to arrange the graph windows in one, two, three, or four columns. The **Zoom** option changes the size of the graph in approximately 25% increments. You can enlarge the graphs or make them smaller. Select the **Normal** option to return to the original graph size.

You can also pick up some useful tips for monitor graphs to make using system monitors even more efficient. When you need help, right-click any element of a monitor graph to get more detailed information. The built-in flexible functions of monitor graphs allows you to gather and see system performance data in ways that reflect how you manage your systems.

Tips - Management Central monitor graphs: System monitors, monitor graphs, and Graph History contain flexible functions that help you get the most out of your system performance data. Once you create a new monitor and become comfortable working with real-time graphs, use the tips that follow to make monitoring your system performance even easier.

- **Create shortcuts to your monitors**
Use a shortcut to access your monitors at any time without first opening Operations Navigator. In **Management Central**, expand **Monitors**, and then select the type of monitor for which you want to create a shortcut. Select **Create Shortcut** from the **File** menu to put a shortcut on your desktop.
- **Use the event log**
The **event log** is a detailed record you can use to see when system performance reaches threshold trigger and reset values. If you select **Include** from the **Options** menu in the **Event Log** window, you

can see detailed event information for particular systems, metrics, and monitors. You can delete events from this list and open a monitor by double-clicking an individual event.

- **Change threshold values directly on your graph**

You can change threshold trigger or reset values directly on your monitor graph. To adjust threshold values, just point to the threshold indicator on the graph and drag it up or down to change the trigger or reset value. A ToolTip shows the changing values.

- **Sort the bars on your monitor graphs**

You can sort the bars that appear in the monitor graphs **Details** pane in a way that makes sense to you. From the **View** menu, select **Sort Details** to arrange the properties in the **Details** pane by **Name** or by **Value**. You can also select **Normalize Details** to sort the bars by the scale the monitor graph uses or by the largest value present.

- **View long-term behaviors with the Graph History window**

Use the Graph History window to view long-term behaviors. Set up Collection Services to collect data, and over time, you will be able to view graphs of system monitor metrics over long periods of time. You can also export that data to a PC file for further analysis. You do not have to start a system monitor to use the Graph History function. You can access the function from the system monitors or directly from a management collection object (*MGTCOL).

- **Create reports to view information collected by the system monitor**

If you created a system monitor to view your performance data, and now you would like to extract that information to the performance database files, you can. The system monitors use a subset of Collection Services data. The data collected for the system monitors is also available in the management collection object (*MGTCOL). Use the Create Performance Data (CRTPFRDTA) command to extract this data to performance database files, and then use the Performance Tools reports to create reports or create your own queries to use the information in the performance database files.

- **Use the system monitor to view data collected by Collection Services**

If you configured Collection Services to collect performance data automatically, you can view that historic data with the system monitors and graphs, even if the monitor was not created or was not active when the data was collected. In this case, you cannot generate events on the previously collected data. Events are only generated if a system monitor is started.

Online help is available for every component of a monitor or monitor graph. You can always click the **Help** button or right-click elements in monitor graphs to find What's this help. In many instances, the online help offers ideas to improve the way you manage your systems.

Example: Management Central's CPU Utilization performance metrics

As the iSeries 400 expands its role as an e-business server, the workloads that it supports are expanding as well. Traditional workloads now include workloads that require less database processing and more CPU-intensive processing resources. Newer generations of iSeries processors deliver increasing amounts of central processing unit (CPU) cycles that allow you to take advantage of these e-business server workloads.

To meet the needs of your evolving business environment, Management Central offers a variety of metrics to help you effectively track different aspects of system performance. By designing CPU utilization metrics that specifically target your type of workload, IBM helps you gain an accurate understanding of how your system is performing given its unique workloads and demands. Compare the following CPU utilization metrics:

- **CPU Utilization (Average)**

This metric shows you the percentage of available processing unit time that is being consumed by all jobs on your server. This metric includes all work done on your system, both interactive and non-interactive.

- **CPU Utilization (Interactive Jobs)**

Formerly known as CPU Utilization (Interactive), this metric shows you the percentage of available processing unit time that is being consumed on the system for all jobs of the type I. Type I jobs include:

- Twinaxial data link control (TDLC)

- 5250 remote workstation
- 3270 remote workstation
- SNA pass-through
- 5250 Telnet

CPU Utilization (Interactive Jobs) helps you manage your interactive users' work compared to total CPU utilization capacity. The resulting data is relative to achieving good interactive response time and the amount of CPU left for non-interactive jobs.

- **CPU Utilization Basic (Average)**

This metric shows the percentage of available processing unit time that is being consumed by all jobs on the system. Unlike the CPU Utilization (Average) metric mentioned above, this metric does not track detailed data. This helps you save system resource, and therefore, takes less overhead to run.

- **CPU Utilization (Interactive Feature)**

This metric is designed to help you monitor and manage your system's interactive use. It determines whether a particular job is doing interactive work and measures the system's overall interactive workload.

This metric complements the other Management Central metrics shown above such as CPU Utilization (Interactive Jobs) and CPU Utilization (Average). This metric shows you when your system is approaching its interactive limits, helping you to get optimal performance and throughput on your system. For more information, see CPU Utilization (Interactive Feature).

- **CPU Utilization (Database Capability)**

This metric is intended to help you monitor your system's database use. Using this metric, you can see how much of your system CPU is being consumed by database activities and which jobs are contributing the most to this use. For more information, see CPU Utilization (Database Capability).

- **CPU Utilization (Secondary Workloads)**

This metric is designed for use on dedicated servers. For instance, this metric can be used on the IBM iSeries Dedicated Server for Domino to see how much non-Domino work is being done on the system. Because the iSeries Dedicated Server for Domino is intended to be used as a server devoted to Domino work, this metric helps you to identify and manage those workloads not directly contributing to that primary system activity. As with interactive utilization, managing secondary workloads helps you to maintain optimal system performance. For more information, see CPU Utilization (Secondary Workloads).

You can use each of these metrics independently or in combination with any of the other monitoring metrics. Use will vary according to the needs of your business environment. To learn more, see:

- **Uses for the CPU Utilization metrics**

Use the CPU Utilization metrics with other metrics to target different types of performance data.

- **Monitoring performance with Management Central Monitors**

Find information about how to create a new system monitor, including a discussion of the performance metrics available in Management Central. Get details about how to make the most of your Management Central monitor graphs.

For a broader discussion of system performance, refer to the performance overview topic.

The CPU Utilization (Interactive Feature) metric:

Use the CPU Utilization (Interactive Feature) metric to monitor and manage interactive use by the endpoint systems in your network. This metric is different from the CPU Utilization (Interactive Jobs) metric in subtle, yet significant, ways. The latter metric was named CPU Utilization (Interactive) in releases prior to V4R5, but has been renamed to CPU Utilization (Interactive Jobs) to help distinguish it from the CPU Utilization (Interactive Feature) metric. The CPU Utilization (Interactive Feature) monitor and metric is available on models preceding 7xx servers. For a more detailed discussion of this topic, refer to the

Performance Capabilities  book.

This feature includes all jobs and threads doing 5250 workstation I/O operations:

- Signed-on 5250 workstation jobs
- Autostart jobs, prestart jobs, or jobs submitted to a batch job queue that do I/O operations to a 5250 workstation.

The CPU Utilization (Interactive Feature) metric tracks the CPU utilization for all jobs doing 5250 workstation I/O operations relative to the capacity of the system for interactive work. Depending on the system and associated features purchased, the interactive capacity is equal to or less than the total capacity of the system.

How CPU Utilization (Interactive Feature) is measured

The CPU Utilization (Interactive Jobs) metric's graph and detailed data are scaled in comparison to 100% of CPU capability. It is intended to help you manage the workload generated by your interactive users. CPU Utilization (Interactive Feature), on the other hand, tracks all jobs identified by your machine as doing interactive work. Tracking of this nature helps you manage your total interactive workload in the same way your system tracks it. You can then anticipate and respond to instances when your interactive workload will exceed the limit of your available interactive capability. Staying within this limit allows you to optimize the performance of your server. CPU Utilization (Interactive Feature) is designed to assist you with managing your interactive workload in comparison with your total interactive capability. As a result, graph scaling is not based on 100% of CPU, but rather 100% of available interactive capability.

Thresholds

When using this metric with the monitor support in Management Central, you see these default threshold values:

	Trigger value	Reset Value
Threshold 1	70%	50%
Threshold 2	90%	70%

While these are good starting values, you may need to adjust the default values. For example, you could adjust according to the combination of 5250 (interactive) work and non-5250 (batch) work in your environment and the particular server on which you are operating.

If you are running on a system that has an interactive capacity that is the same as the total CPU capacity, such as a Model 620, this metric is the same as the total CPU capacity. The amount of impact depends on the percent of Interactive Feature CPU Utilization and the specific server on which you are operating.

The threshold trigger values of 70% and 50% and the reset values of 90% and 70% enable you to better manage your interactive workload. If your interactive feature CPU utilization approaches 100% on these CPU features, you still have significant CPU left for noninteractive work. However, if you need additional CPU power to complete your noninteractive work or your 5250 work consistently exceeds 100% of your interactive feature, you need to take capacity planning actions to determine if additional CPU processor performance is needed.

Interactive capability

To scale interactive usage to 100% of available interactive capability means that 100% of available interactive capability is the highest level of CPU utilization at which you have optimal system performance. Simply put, you get the best system performance by staying below 100%. This means that, unlike earlier metrics, you can exceed 100%.

As a result, graphs, thresholds, and other values that might logically be limited to 100% on other metrics are allowed to go higher with this metric. This allows you to view your actual interactive workloads when they exceed 100%, and to manage your unique computing environment, whether you normally run at 40% or 100% of purchased interactive capability. Exceeding this point for extended periods of time will cause degradation in your system's performance and throughput. In AS/400 or iSeries systems that support Interactive Feature cards, the 100% limit represents the capability purchased with those cards. With prior AS/400e servers and other AS/400 systems that have different interactive and server capacities, 100% represents the capability purchased indirectly with that hardware model. For other, traditional models, 100% may represent 100% of total CPU capability. Some older AS/400 systems do not support this metric.

In addition, scaling to interactive capability means that 40% is not 40% of total CPU use, but rather 40% of interactive capability. The two will not be the same unless interactive capability is equal to total CPU capability. For example, suppose that your server supports 25% of its CPU for interactive work. In this case, if 10% of total CPU capability is being used in interactive work, then four times that, or 40%, of interactive capability is being used. Scaling the performance data in this manner helps you to quickly see how the endpoint systems in your network are doing. If some systems support 25% of CPU for interactive work and others support 10% and even others support 80% or 100%, all of these systems will be normalized to graph against 100% of that available interactive resource. You simply track your network usage against the 100% mark. If you need to see how much of the total CPU on a particular server is available for interactive work, that information is available on the title bar of the Details pane.

Finally, be aware of another difference in the way CPU Utilization (Interactive Feature) is measured. When your machine measures your interactive workload on systems where interactive capability is limited to less than 100% of total CPU, the machine allows some flexibility in the form of a single-task exemption. If only one job is doing interactive work during a particular interval, no interactive work is charged to the system's interactive feature capability. Single-task exemption is designed to allow for a limited amount of normal systems management and maintenance activity that would normally get charged as interactive work. An example of this kind of an operation is the reclaim storage operation. Once there is more than one task running, though, the system counts all interactive work in its interactive totals and the single-task exemption no longer applies.

Note: If you measure CPU Utilization (Interactive Feature) while only one interactive job is active, you will not get an accurate picture of interactive capability consumed on your system. To get a better measure of the interactive capability used on your system, be sure to have more than one interactive job running while you run your monitor.

Using CPU Utilization (Interactive Feature) with other metrics

You may want to use CPU Utilization (Interactive Feature) with other performance metrics. To find more information, see:

- **Uses for CPU Utilization metrics**
Explore Management Central performance metrics by taking a closer look at CPU Utilization. Learn how to use CPU Utilization metrics to target different types of performance data.
- **CPU Utilization (Database Capability)**
Monitor how much processing time is being consumed by database activity.
- **CPU Utilization (Secondary Workloads)**
Monitor how much processing time is being consumed by work other than your primary workload on your dedicated server (such as the iSeries Dedicated Server for Domino).

For a full discussion of all performance metrics, refer to your online help or to the [Selecting metrics](#) page. For a broader discussion about system performance, refer to the [performance overview](#) topic.

The CPU Utilization (Database Capability) metric:

The CPU Utilization (Database Capability) metric allows you to monitor DB2 Universal Database for iSeries activity on your systems. This metric applies to all systems running V4R5 or later and includes all

database activity, including all SQL and data I/O operations. Use this metric to see how much of your system CPU is being consumed by database function. By selecting a specific point on the system monitor graph, you can see which jobs are doing the most database activity. In addition, you can find detailed data for each job, including the number of milliseconds of CPU used by that job in database processing during the particular sample interval being graphed.

CPU Utilization (Database Capability) measures how much processor time is being used on database activity. More specifically, CPU Utilization (Database Capability) shows database activity relative to total database capability.

When using this metric with the monitor support in Management Central, you see that the default threshold values are the same as those for the Interactive Feature CPU utilization metric:

	Trigger value	Reset Value
Threshold 1	70%	50%
Threshold 2	90%	70%

While these are good starting values, the default threshold values could be very normal in a nightly batch workload or a data warehousing environment. In another environment where moderate database activity is anticipated, exceeding these default thresholds could indicate unanticipated database activity, such as CPU-intensive queries. In that type of an environment, you may need to do additional analysis at the job level or you may need to do capacity planning for increased CPU requirements.

Using CPU Utilization (Database Capability) with other metrics

You may want to use CPU Utilization (Database Capability) with other performance metrics. To find more information, see:

- **Uses for CPU Utilization metrics**
Explore Management Central performance metrics by taking a closer look at CPU Utilization. Learn how to use CPU Utilization metrics to target different types of performance data.
- **CPU Utilization (Interactive Feature)**
Monitor how much processing time is being consumed by interactive work on your system.
- **CPU Utilization (Secondary Workloads)**
Monitor how much processing time is being consumed by work other than your primary workload on your dedicated server (such as the iSeries Dedicated Server for Domino).

For a full discussion of all performance metrics, refer to your online help or to the [Selecting metrics](#) page. For a broader discussion about system performance, refer to the [performance overview](#) topic.

The CPU Utilization (Secondary Workloads) metric:

The CPU Utilization (Secondary Workloads) metric is designed for use on dedicated servers. It measures how much CPU is being used on the system for work other than the primary workload for which the system is designed, which can include database activity. For instance, this metric can be used on the iSeries Dedicated Server for Domino to see how much non-Domino work is being done on the system.

The amount of CPU used by secondary workloads is currently reported only on Dedicated Servers for Domino systems running V4R5 or later and is a portion of the total CPU utilization capacity. On other systems and servers, the value is ignored and appears as 0%.

This metric does not collect job-level data; therefore, no details are shown in the right-hand section of the window.

Because the iSeries Dedicated Server for Domino is intended to be used as a server devoted to Domino work, this metric has been provided to help you to identify and manage those workloads not directly contributing to that primary system activity. As with interactive utilization, managing secondary workloads helps you to maintain optimal system performance. If a machine that is intended for dedicated use (such as Domino) is used too much for other activity, performance is affected. Although it varies by model, generally speaking you should plan not to exceed 10% to 15% of CPU in secondary work on a dedicated system.

Dedicated servers




A dedicated server is a server designed for use with a specific application environment. CPU Utilization (Secondary Workloads) was designed to work for both AS/400e Dedicated Server for Domino and iSeries Dedicated Server for Domino, as well as for dedicated servers that may be developed in the future.

You can use the CPU Utilization (Secondary Workloads) metric on a Dedicated Server for Domino to compare it to total CPU utilization. The best Domino performance can be achieved when Domino work is at least three times any non-Domino work that is active at the same time; therefore, good starting threshold trigger and reset values are:

	Trigger value	Reset Value
Threshold 1	5%	10%
Threshold 2	10%	15%

The AS/400e Dedicated Server for Domino is the industry's first server designed specifically for Lotus Domino and is built to deliver maximum value for small- to medium-sized businesses looking to run in a Domino mission-critical environment. It provides reliability, ease of management, and competitive cost for users who need a server solely for the use of Domino software (such as e-mail, office, Domino business applications, and Domino Web serving).

More information on the AS/400e Dedicated Server for Domino is available from the following sources:

- The Lotus Domino for AS/400  Web site (www.as400.ibm.com/domino)
- The Domino white paper , "Evaluating Appropriate Workloads for the AS/400e Dedicated Server for Domino" (www.as400.ibm.com/whpapr/dsd.htm)
- The Performance Capabilities  book section on the AS/400e Dedicated Server for Domino

Using CPU Utilization (Secondary Workloads) with other metrics

You may want to use CPU Utilization (Secondary Workloads) with other performance metrics. To find more information, see:

- **Uses for CPU Utilization metrics**
Explore Management Central performance metrics by taking a closer look at CPU Utilization. Learn how to use CPU Utilization metrics to target different types of performance data.
- **CPU Utilization (Interactive Feature)**
Monitor how much processing time is being consumed by interactive work on your system.
- **CPU Utilization (Database Capability)**
Monitor how much processing time is being consumed by database activity.

For a full discussion of all performance metrics, refer to your online help or to the Selecting metrics page. For a broader discussion about system performance, refer to the performance overview topic.

Uses for Management Central's CPU Utilization metrics:

Use Management Central's CPU Utilization metrics (or any other performance metrics) alone or in concert to target specific types of performance information. The information you want to glean from your systems may be more complex than a single metric will provide. Running metrics in combination allows you to analyze system data from multiple perspectives. Keep reading to learn more about ways to use the CPU Utilization metrics in Management Central.

Uses for CPU Utilization (Interactive Feature)

One of the best ways to interpret this metric is to use it in combination with others. For example, you might:

- Use CPU Utilization (Interactive Feature) with CPU Utilization (Average).
This combination gives you data about total system CPU use as well as interactive feature CPU use. This can be helpful in interpreting overall system activity.
- Use CPU Utilization (Interactive Feature) with CPU Utilization (Interactive Jobs).
Jobs of type I are a subset of all jobs included in those analyzed by the CPU Utilization (Interactive Feature) metric. As a result, this combination provides an accurate indication of interactive activity taking place on your system. This also helps you relate your system interactive capability to total CPU capability both at the system and job levels.

Uses for CPU Utilization (Database Capability)

CPU Utilization (Database Capability) can help you identify and resolve the impact of complex queries that are performed while other performance-critical applications are active. You may get the most value out of this specialized metric when you use it in combination with others to understand system database activity relative to other aspects of system activity. For example, you might:

- Use CPU Utilization (Database Capability) with CPU Utilization (Average).
This combination gives you an accurate perspective of database activity relative to overall CPU activity on your system.
- Use CPU Utilization (Database Capability) in addition to Disk Arm Utilization (Average) or Disk Storage (Average).
This mix of performance metrics provides valuable information in an environment where the effects of database activity are important relative to disk utilization.
- Use CPU Utilization (Database Capability) together with Communications IOP Utilization (Average).
You may choose to use this combination of performance metrics when analyzing a server application (such as Web serving or ODBC) for time spent in communications versus data storage and retrieval.

Uses for CPU Utilization (Secondary Workloads)

Much like the CPU Utilization Basic (Average) metric, CPU Utilization (Secondary Workloads) does not provide job-level detail. You can use CPU Utilization (Secondary Workloads) with other metrics, however, to obtain the additional job-level data if you so require. For example, you can:

- Use CPU Utilization (Secondary Workloads) with CPU Utilization (Database Capability).
This combination of metrics helps you monitor those Domino applications running ODBC or @DB commands to access DB2 Universal Database for iSeries data.
- Use CPU Utilization (Secondary Workloads) with CPU Utilization (Average), CPU Utilization (Interactive Jobs), and with other similar metrics.
These metrics, when used together, provide a summary of the secondary work taking place on your server. CPU Utilization (Secondary Workloads) does not provide job-level detail; however, you can use this metric with other metrics to obtain detailed job-level information about activity on your system. You may discover that different metrics, when used with CPU Utilization (Secondary Workloads), will provide related job-level detail that best suits your performance data needs.

To find more information about these metrics, see:

- **CPU Utilization (Interactive Feature)**
Monitor how much processing time is being consumed by interactive work on your system.
- **CPU Utilization (Database Capability)**
Monitor how much processing time is being consumed by database activity.
- **CPU Utilization (Secondary Workloads)**
Monitor how much processing time is being consumed by work other than your primary workload on your dedicated server (such as the iSeries Dedicated Server for Domino).

For a full discussion of all performance metrics, refer to your online help or to the *Selecting metrics* topic. For a broader discussion about system performance, refer to the *Management Central monitors* topic or to the *Managing system performance* topic.

Monitoring jobs and servers with Management Central

Use Management Central job monitors to stay on top of job activity by monitoring a job or a list of jobs based on job name, job user, job type, subsystem, or server type.

You can start a job monitor, and then turn to other tasks on your server, in Operations Navigator, or on your PC. In fact, you could even turn your PC off! Management Central will continue to monitor your jobs and perform any threshold commands or actions you specified. Your monitor will run until you decide to stop it.

To get started working with Management Central monitors, choose one of the following topics:

- **Creating a new monitor**
Get step-by-step help through the process of creating a new monitor.
- **Selecting job metrics**
Select any number of metrics for any selected set of jobs.
- **Specifying threshold values**
Find out how to set thresholds and threshold actions.
- **Running commands**
Manage your jobs by running a command on the server when a threshold value is triggered or reset.
- **Logging events**
Learn how to log an event whenever a threshold value is triggered or reset.
- **Applying thresholds and actions**
Choose when to apply the thresholds and actions you have defined.
- **Viewing job monitor results**
Open a Job Monitor window to see the status of the monitor and select from a menu of actions that can be performed on the jobs, such as reset triggered thresholds, display job properties, hold, release, or end a job.

Creating a new job monitor

Management Central job monitors are powerful tools that you can use to stay on top of job activity on your endpoint systems. Creating a new monitor is a quick and easy process that begins at the New Monitor window. In Operations Navigator, expand Management Central, expand **Monitors**, right-click **Job**, and then select **New Monitor**.

Once you have given your new monitor a name, the next step is to specify the jobs 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 impact on your system. You can specify the jobs to monitor in two 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 wild card 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.

Use the online help to assist you in creating your new job monitor. Read about selecting metrics to find out about the different aspects of job activity that you can measure with Management Central job monitors.

Selecting metrics for job monitors

When you create a job monitor, you must decide which aspects of job activity you want to monitor. Management Central offers several measurements, known as **metrics**, to help you pinpoint different aspects of job activity.

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, select **Monitors**, right-click **Job**, and then select **New Monitor**. Fill in the required fields, and then click the **Metrics** tab.

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 your monitor include the following:

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.

Job log messages

Monitor for messages based on any combination of Message ID, Type, and Minimum severity.

Job numeric values

CPU utilization

The percentage of available processing unit time used by each job that is being monitored 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.

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	
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.

Use the online help to assist you in selecting your metrics. Don't forget to specify threshold values that will allow you to be notified and to specify actions to be taken when a certain value (called the trigger value) is reached.

Specifying threshold values for a job monitor

When you have selected the metrics for your job monitor, you should consider setting a threshold for each metric. 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. You can set up to two thresholds for each metric that the job 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 for duration helps to avoid unnecessary threshold activity due to frequent spiking of values.

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. Depending on the type of metric you have selected, you can set your threshold values in the following ways:

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** (or press F4) 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.

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 impact 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** (or press F4) 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 **General** tab, select the statuses that you want to monitor for. Click the **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** (or press F4) 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** (or press F4) 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** (or press F4) 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.

Use the online help to assist you in setting your threshold values. Next, you will want to find out about running commands when a threshold is triggered or reset.

Specifying the collection interval for a job monitor: 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.

Click the **Metrics** tab to specify the number of collection intervals for each threshold.

Running commands for job monitors

When you create a new job monitor, you can choose to run commands on endpoint systems when thresholds are triggered or reset. 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.

What can I do with threshold commands?

Use threshold settings to automate any command you want to run when thresholds are triggered or reset. For example, if a certain batch job is supposed to complete before the first shift begins and it is still running at 6:00 a.m., you could set up Threshold 1 to send a page command to a system operator to look at it. You could 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 may want to notify your operators with a page command when the wait time values for the FTP and HTTP servers reach a median level. If the FTP server jobs end, you could 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** 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** 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 the 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.

Next, you will want to find out about logging events when a threshold when a threshold is triggered or reset.

Logging events for a job monitor

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

can select are:

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.
Run OS/400 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 or not 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.

Applying thresholds and actions for a job 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.

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 wanted to apply your thresholds and actions overnight on Monday night, you could select 11:00 p.m. as the **From** time and 6:00 a.m. as the **To** time. You would check Monday. The actions you specified would occur whenever the specified thresholds were 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 job monitor. The online help also contains instructions on starting your monitor. Then you will be ready to view your job monitor results.

Viewing job monitor results

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

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

the target systems (Summary Area) in the upper pane shows the status of the monitor on each system, the triggered events for the monitor, the summary data of the specified metrics, and the date and time that the monitor data was last collected.

When you select a system in the Summary Area, the jobs that are being monitored on that system are shown in the Job Area (lower pane). The list of jobs show 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 dialog to see a description of each column.

From the list of jobs in the Job Area, you can select from a menu of actions that can be performed on the jobs, such as reset triggered events, display job properties, hold, release, or end a job.

Remember to use other Management Central functions to manage your multiple servers easily and efficiently!

Reset triggered threshold for a job 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.

Remember to use other Management Central functions to manage your multiple servers easily and efficiently!

Monitoring message queues with Management Central

Use Management Central message monitors to monitor your message queues for the information you need to manage your servers. For example, you could monitor a message queue to determine whether an application completed successfully. Or you could monitor the system operator message queue for a specific message that indicates when a critical storage condition exists. When you create a monitor, you can specify commands to run when the message is detected.

Like the other Management Central monitors, you can start a message monitor, and then turn to other tasks on your server, in Operations Navigator, or on your PC. In fact, you could even turn your PC off! Management Central will continue to monitor your message queues and perform any threshold commands or actions you specified. Your monitor will run until you decide to stop it.

To get started working with Management Central monitors, choose one of the following topics:

- **Creating a new monitor**
Get step-by-step help through the process of creating a new monitor and defining up to two sets of messages to be monitored.
- **Defining message sets for a monitor**
Find out how to define a set of messages to be monitored.
- **Running commands**
Manage your messages by running a command on the server when a message count threshold is triggered or reset.
- **Logging events**
Learn how to log an event whenever a threshold value is triggered or reset.
- **Applying thresholds and actions**
Choose when to apply the thresholds and actions you have defined.
- **Viewing message monitor results**
Open a Message Monitor window to see the status of the monitor and select from a menu of actions that can be performed on the messages, such as reset triggered thresholds, display message properties, reply to a message, or delete a message.

Creating a new message monitor

Management Central message monitors are powerful tools that you can use to simplify the task of managing your servers by automating responses to critical messages. Creating a new monitor is a quick and easy process that begins at the New Monitor window. In Operations Navigator, expand Management Central, expand **Monitors**, right-click **Message**, and then select **New Monitor**.

Once you have given your new monitor a name, you can click the **Messages** tab to define the set of messages you want to monitor. You must select one message queue for each monitor. You can define up to two sets of messages for each monitor. For each set of messages, you can specify the threshold (number of messages received in the queue) that triggers an event and a server command that runs when the event is triggered. You can also choose whether to log these events and specify PC actions to be taken when a threshold trigger value is reached. And finally, you can specify whether to always apply the thresholds and actions you have specified, or to apply them only on the days and times you choose.

Use the online help to assist you in creating your new message monitor. The first step is to define message sets that will allow you to be notified and to specify actions to be taken when a certain value (called the trigger value) is reached.

Defining message sets for a message monitor

When you create a message monitor, you select one message queue to be monitored. For each message queue, you can define two separate and independent message sets to monitor. For each message set,

you can specify the threshold (number of messages received in the queue) that will trigger an event and a server command that runs when the event is triggered.

When you define your message sets, you want to specify each of the following:

- **What to monitor**

You can add a predefined set of messages (for example, “Probable modem problem”) to a message set. You can also add messages based on message ID, message type, and message severity. When you choose to monitor for inquiry messages (or all messages, which would include inquiry messages), you can specify a reply to be sent to the inquiry messages (such as Cancel, Ignore, or Retry).

Note: If an inquiry message is included more than once in a message set, the reply that is specified first is the one that is used. If a message is included both in Message Set 1 and Message Set 2, the action specified in Message Set 1 is taken first. This may mean that the message is not found when Message Set 2 is processed.

For example, you might specify messages in the first message set that would indicate that a problem was developing on the system. For that message set, you could specify a command that would page the system operator. For the second message set, you could specify a different set of messages that would indicate a critical situation. For those messages, you would specify a command that would start ending certain jobs on the system.

- **What to do when the messages are found**

When you have defined a set of messages to be monitored, you can choose either to permanently remove the monitored messages from the message queue or to trigger the monitor at the specified message count. If you select to trigger the monitor, you can specify the trigger value (how many messages must be received before the monitor is triggered) and a server command to run when the trigger value is reached. Optionally, you can also specify a server command to run when you reset the monitor. When you reset the monitor, you always have the option to reset without running the specified command. A message monitor is never automatically reset.

- **How often to look for the messages**

On the **Collection Interval** page, you can specify how often you want the message monitor to search the message queue for new messages. You can choose to search as often as every 15 seconds, or you can choose to search only once every hour, or you can choose any of several levels in between!

Use the online help to assist you in defining your message sets. Next, you will want to find out about running commands when a threshold is triggered or reset.

Running commands for message monitors

If you selected to trigger the monitor at a specified message count when you defined your message sets, you can choose to run commands on endpoint systems when thresholds are triggered or reset. A **threshold** for a message monitor is the message count that you specify to trigger an event. **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.

What can I do with threshold commands?

Use threshold settings to automate any command you want to run when thresholds are triggered or reset. For example, suppose you have a payroll application that sends a message when information is missing on someone’s timecard. You could set up the monitor to send a page command to the human resource person so they can take care of the problem..

In another situation, you may be running an application that sends a message when it times out. Typically, you find that the application needs to be restarted if more than three time-outs occur. You could set up your monitor to look for three occurrences of the message and then send a command to restart the

application. 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-Messages** page, click **Trigger at the following message count** to specify the message count and threshold trigger and reset commands. You can enter a 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 &MSGCOUNT) to pass information to the command, such as the time and actual message count.

Management Central monitors allow you to specify any batch commands to run on the server when the threshold is triggered or reset. The command you specify is run whenever a threshold trigger value is reached during the time you specify to apply thresholds and actions. When you reset the monitor, you always have the option to reset without running the specified command. A message monitor is never automatically reset.

Next, you will want to find out about logging events when a threshold when a threshold is triggered or reset.

Logging events for a message monitor

When you have defined your message sets, you can click the **Actions** tab to specify event logging and PC actions to be taken when a threshold is triggered. The following actions are available when a threshold is triggered or reset:

- **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 on which the event occurred, the message being monitored for, the severity of the message, the message monitor that logged the event, and any reply that was sent.
- **Run AS/400 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 Messages page. Whenever you reset a threshold, you can select whether or not to run the specified reset command.

In addition you can select any of the following actions to be taken when a threshold is triggered:

- **Open event log**
Displays the event log when the threshold is triggered (available only if you selected to log events).
- **Open monitor**
Displays a list of systems that are being monitored for the specified messages and a list of the specified messages that have been received for each system.
- **Sound alarm**
Sounds an alarm on the PC when the threshold for the monitor is triggered.

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

Applying thresholds and actions for a message monitor: When you have defined your message sets 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. Replies are sent and messages are removed from the message queue only during the time you specify to apply thresholds and actions for the monitor.

If you select to apply thresholds and actions during specified times, you must select the starting time and the stopping time and the times must be different. The time is based upon the time zone of 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.

For example, if you wanted to apply your thresholds and actions overnight on Monday night, you could select 11:00 p.m. as the **From** time and 6:00 a.m. as the **To** time. You would check Monday. The actions you specified would occur whenever the specified thresholds were 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 message monitor. The online help also contains instructions on starting your monitor. Then you will be ready to view your message monitor results.

Viewing message monitor results

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

Double-click the monitor name to open the Message Monitor window. In the Message Monitor window, you can see the overall status of the message monitor and a list of the target systems that the monitor is running on. A list of the target systems (Summary Area) in the upper pane shows the status of the monitor on each system, the message count on that system, and the date and time that the monitor data was last collected.

When you select a system in the Summary Area, the messages that are being monitored on that system are shown in the Message Area (lower pane). The list of messages shows the message ID, message type, and the severity of the message. The Event column shows the reason the message is in the list (either Message found or Message count = n).

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

You can select **Monitor Size** from the **View** menu to specify the maximum number of messages (starting with the most recent) to be displayed in the Message Area.

From the list of messages shown in the Message Area, you can select from a menu of actions that can be performed on the messages, such as reset a triggered threshold, display message properties, reply to a message, or delete a message.

Remember to use other Management Central functions to manage your multiple servers easily and efficiently!

Reset triggered threshold for a message monitor: When you are viewing the message 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 message level, the system level, or the monitor level:

- **Message level**

Select one or more messages in the Message Area of the Message Monitor window. Select **File**, select **Reset with Command** or **Reset Only**, and then select **Messages**. The thresholds for the selected messages 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 Message 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 Message 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 Message Area are ignored.

Remember to use other Management Central functions to manage your multiple servers easily and efficiently!

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 Properties page for each monitor whether or not you want events added to the Event Log. To see the Properties page 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 on System.

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



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.

Collecting performance data with Collection Services

Use Collection Services to collect performance data for later analysis by the Performance Tools for iSeries licensed program or other performance report applications, Management Central monitors, and the Graph History function. (If you prefer viewing real-time performance data, Management Central system monitors and the Graph History function provide an easy-to-use graphical interface for monitoring system performance.) Collection Services collects data that identifies the relative amount of system resource used by different areas of your system. Use Collection Services to:

- Easily manage your collection objects

- Collect performance data continuously and automatically with minimal system overhead
- Control what data is collected and how the data is used
- Move performance data between releases without converting the data
- Create performance data files that are used by Performance Tools

To find out more about Collection Services and performance data, keep reading.

How Collection Services works


Collection Services replaces the OS/400 performance monitor, which was called by the Start Performance Monitor (STRPFRMON) command. When you used the OS/400 performance monitor, your data was collected into as many as 30 database files.

Note: The performance monitor (STRPFRMON command) is not available after V4R5.

Collection Services capabilities introduce a new process of collecting performance data. Collection Services stores your 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 means a lower system overhead when collecting performance data. Even if you elect to create the database files during collection, you still experience a performance advantage over the OS/400 performance monitor because Collection Services uses a lower priority (50) batch job to update these files. The reduction in collection overhead makes it practical to collect performance data in greater detail and at shorter intervals on a continuous basis. Collection Services enables you to establish a network-wide system policy for collecting and retaining performance data and to implement that policy automatically. For as long as you retain the management collection objects, if the need arises, you have the capability to look back and analyze performance-related events down to the level of detail that you collected.

How to start Collection Services

You can start Collection Services by using any of the following methods. However, the information in the Performance topic focuses on Operations Navigator methods.

Starting method	Description
Operations Navigator	The section that follows below shows you how to do a variety of Collection Services tasks using Operations Navigator.
Management Central APIs	You can use Management Central APIs to start, customize, end, and cycle collections.
Traditional menu options	Type GO PERFORM in the character-based interface and select option 2 (Collect performance data) from the Performance Tools main menu. For additional information, go to the Performance Tools book  .
Performance Management/400	You can activate PM/400 which automates the start of Collection Services and then creates the database files during collection.

Collection Services tasks

You can use Collection Services and Management Central to perform a variety of data collection tasks as shown below.

Task	Description
Start Collection Services in a variety of ways	Create a customized performance data collection on an individual system or groups of systems with specific performance metrics. You can also use a Start Collector API in your startup program to start performance data collections automatically. For more information about how to perform these tasks, refer to the online help.
Create database files	Use Collection Services to automate the creation of performance database files. Learn how to control what data gets collected as you create database files.

Customize data collections	Customize your data collections! Find information about controlling what performance data you collect, and how often that data gets collected. You can also find information about important time zone considerations.
Manage collection objects	Find the information you need to manage collection objects, including the contents of collection objects, how long collection objects are saved, and what you can do with collection objects.
Collect sample data	Collection Services collects sample data. However, it does not collect trace data. Find out how to collect trace data.

Collection Services and performance database files

When you use Collection Services to collect performance data, you can create database files automatically as data is collected. You can also create database files from the collection object, where the data is stored after it has been collected. In addition, you can activate Performance Management/400, which automates the start of Collection Services and then creates the database files during collection. You can use these database files with PM/400, Performance Tools licensed program, or you can create your own queries to run against these files.

To find more information about performance database files, see Performance data files to find out what database files are available to you, as well as the field-level data included in each file.

See the Performance overview topic if you would like more information about iSeries performance.

Collecting trace data

Collection Services provides for the collection of sample data. Sample data is summary data that is captured at regular time intervals. You collect sample data for trend analysis and performance analysis. The data relates to things such as storage pools and response times. However, Collection Services does not support the collection of trace data. Trace data is detailed data that you collect to gain additional information about specific jobs and transactions.

If you want to collect trace data, you should use the Start Performance Trace (STRPFRTRC) command, which is a simplified interface to the Trace Internal (TRCINT) command for collecting multiprogramming level (MPL) and transaction trace data. Issue the STRPFRTRC command to collect trace data:

```
STRPFRTRC SIZE(*CALC)
```

When you issue the previous command, it collects the same trace data previously collected by the STRPFRMON (TRACE(*ALL)) command. You can use the Transaction Report to process the data.

You can end the trace and optionally write the data to the performance database file QAPMDMPT by using the End Performance Trace (ENDPFRTRC) command. This is the same file that was previously supported by the STRPFRMON and ENDPFRMON trace function and the DMPTRC command. To find out when you should dump your trace data, see dumping trace data.

From a character-based interface, you can also start and end a performance trace. From the Performance Tools main menu, select option 5 (Performance utilities) to access these options.

Dumping trace data

Deciding when to dump trace data is a significant decision because the dump affects system performance. The Dump Trace (DMPTRC) command puts information from an internal trace table into a database file. It is not good to dump trace data during peak activity on a loaded system or within a high priority (interactive) job. You can delay a trace dump, but you want to dump the data before you forget that it exists. If the trace table becomes cleared for any reason, you lose the trace data. However, delaying the dump slightly and then using the DMPTRC command to dump the trace in a batch job can preserve performance for the users.

To dump trace data, issue the following command:

DMPTRC MBR(member-name) LIB(library-name)

You must specify a member name and a library name in which to store the data. You can collect sample-based data with Collection Services at the same time that you collect trace data. When you collect sample data and trace data together like this, you should place their data into consistently named members. In other words, the names that you provide in the CRTPFRTA TOMBR and TOLIB parameters should be the same as the names that you provide in the DMPTRC MBR and LIB parameters.

Customizing data collections with Management Central

When you use Collection Services to collect performance data, you control what data is collected and how often it is collected. You can select from the collection profiles that are provided. The Standard profile corresponds to the settings in the OS/400 performance monitor function (STRPFRMON command) for system data. The Standard plus protocol profile corresponds to the STRPFRMON command settings for all data. Or you can select **Custom** to create your own customized profile. For your customized profile, you can select from a list of available data categories, such as System CPU, Local Response Time, Disk Storage, and IOPs (input/output processors).

For each category of data that you collect, you can specify how often the data will be collected. For many categories, you will want to select the default collection interval, which you can set from predefined settings between 15 seconds and 60 minutes. (The recommended setting is 15 minutes.)

The collected data is stored in a management collection object (type *MGTCOL), called a collection. To prevent these management collection objects from becoming too large, the collection must be cycled at regular intervals. Cycling a collection means to create a new collection object and begin storing data in it at the same time data collection stops in the original collection object. You can specify any interval from one hour to 24 hours, depending on how you plan to use the data.

How do I customize my data collections?

To customize Collection Services on a system, follow these steps:

1. In Operations Navigator, select either an endpoint system under **Management Central** or a system to which you have a direct connection under **My Connections** (or your active environment).
2. Expand **Configuration and Service**.
3. Right-click **Collection Services** and select **Properties**.
4. On the **General** page, you may want to specify a retention period longer than the default of 1 day. Collection Services may delete management collection objects and the data they contain from the system at any time after the retention period has expired. The expiration date is associated with the management collection object. Even if you move the collection object to another library, Collection Services will delete the object after it expires. You can specify **Permanent** if you do not want Collection Services to delete your management collection objects for you.
To view the Graph History window, you must specify a Collection retention period of either Graph or Summary. When you specify these options, you can take advantage of the historical reporting capabilities, which would allow you to do metric comparisons for multiple systems over extended periods of time.
You can also specify the path of the location where you want to store your collections, how often you want to cycle collections, and the default collection interval. You can select to create database files automatically during collection.
5. Click the **Data to Collect** tab.
6. For **Collection profile to use**, select **Custom**. You can specify the collection interval for each category you select for your customized list.
7. Click **OK** to save your customized values.

Once you have customized Collection Services to the settings you prefer, you can right-click **Collection Services** again and select **Start Collection Services** to begin collecting performance data.

Time zone considerations for Collection Services

When you review and analyze performance data, the actual local time of the collection can be significant. For example, you may need to be sure which data was collected during the busiest period of the day so that it represents the heaviest workload experienced by the system under review. If some of the systems from which you collect performance data are located in different time zones, you should be aware of these considerations:

- When you start Collection Services for a system group, you start Collection Services at the same time on all systems in the group. Any differences in system time and date settings due to some systems being located in different time zones are not taken into account.
- If you start Collection Services with the Management Central scheduler, the time at which the scheduler starts the task is based on the system time and date of your central system in Management Central.
- The management collection objects for each endpoint system reflect start and end times based on the QTIME and QUTCOFFSET (coordinated universal time offset) system values of that endpoint system and your central system. If the endpoint system is in a different time zone from your central system, and these system values are correctly set on both systems, the start and end times reported for collection objects are the actual times on the endpoint system. In other words, the start and end times reflect the value of QTIME on the endpoint system as it was at the actual point in time when those events occurred.
- The scheduling of a performance collection can cross a boundary from standard time to daylight savings time or from daylight savings time to standard time. If so, this time difference should be taken into account when scheduling the start time. Otherwise, the actual start and end times can be an hour later or earlier than expected. In addition, the reported start and end times for management collection objects are affected by this difference unless the QUTCOFFSET system value is adjusted each time the change to and from daylight savings time takes effect.

For more information about using Management Central's tool to collect performance data, see [Collecting performance data with Collection Services](#).

Creating database files to use with Management Central's Collection Services

You have many options that allow you to create database files.

- When you use Collection Services to collect performance data, you can create database files automatically as data is collected.
- You can use the Create Performance Data (CRTPFRDTA) command to create a set of performance database files from performance information stored in a management collection (*MGTCOL) object.
- You can create database files from the management collection object, where the data is stored after it has been collected. You can use either the Operations Navigator interface or the CRTPFRDTA command.
- You can activate Performance Management/400, which automates the start of Collection Services and then creates the database files during collection.

The database files can be used with the Performance Tools for iSeries licensed program or other applications to produce performance reports.

Why should I store the data in management collection objects instead of in the database files I need to run my reports?

Because you can manage the management collection objects separately from the database files, you can collect your performance data in small collection intervals (such as 5-minute intervals) and then create your database files with a longer sampling interval (such as 15-minute intervals).

From a single management collection object, you can create many different sets of database files for different purposes, by specifying different data categories, different ranges of time, and different sampling intervals.

For example, you might collect performance data on the entire set of categories (all data, or the **Standard plus protocol** profile) in 5-minute collection intervals for 24 hours. From that one management collection object, you can create different sets of database files for different purposes. You could create one set of database files to run your normal daily performance reports. These files might contain data from all categories with a sampling interval of 15 minutes. Then, to analyze a particular performance problem, you could create another set of database files. These files might contain only data for a single category that you need to analyze, a specific time period within the 24 hours, and a more granular sampling interval of 5 minutes.


How do I create the database files?

Collection Services places the data you collected into management collection objects. To use this data, you must first place the data in a special set of database files. To create database files automatically as data is collected, simply select **Create database files** on the **Start Collection Services** dialog. You can also create the database files later when you want to export data to them from an existing management collection object.

To export performance data from a management collection object to database files, follow these steps:

1. In Operations Navigator, select either an endpoint system under **Management Central** or a system to which you have a direct connection under **My Connections** (or your active environment).
2. Expand **Configuration and Service**.
3. Click **Collection Services**.
4. Right-click the management collection object you want to export to database files and select **Create Database Files**.
5. On the Create Database Files dialog, select the categories from the collection object to include in the database files. You can also select a different time period and sampling interval, as long as the collection object contains data to support your selections.
6. Click **OK**.

You can use the database files you have created with the Performance Tools for iSeries licensed program or other applications to produce performance reports. You can collect the performance data on one system and then move the management collection object (*MGTCOL) to another system to generate the performance data files and run the Performance Tools reports. This action allows you to analyze the performance data on another system without affecting the performance of the source system. For more

information about Performance Tools, see the Performance Tools book .

Creating database files from an existing collection object

You can use Management Central to export performance data from an existing management collection object to database files. Follow these steps:

1. Expand **Configuration and Service** for the system from which performance data is being collected.
2. Select **Collection Services**.
3. Right-click the management collection object from which you want to export data to the database files.
4. You can first select **Properties** to display the characteristics of the data in the collection object. On the Data properties page, you can see the categories of data collected in this collection object as well as the intervals at which they were collected. You can use this information in selecting the data that you want to export. When you have reviewed this information, click **OK**.
5. Right-click the management collection object again and select **Create Database Files**. Complete the fields using the online help.
6. Click **OK**.

After you convert the data in the database files, you can use the Performance Tools for iSeries licensed program or other applications to produce performance reports.

Managing collection objects with Management Central

When you use Collection Services to collect performance data, each collection is stored in a single object.

How do I find out what data is in my management collection objects?

To see a summary of the data in any management collection object, follow these steps:

1. In Operations Navigator, select either an endpoint system under **Management Central** or a system to which you have a direct connection under **My Connections** (or your active environment).
2. Expand **Configuration and Service**.
3. Select **Collection Services**.
4. Right-click any management collection object in the list and select **Properties** to see general information about that collection and a summary of the data that it contains.

How do I know when Collection Services will delete my old management collection objects?

Collection Services deletes only **cycled** management collection objects. A status of **Cycled** means that Collection Services has stopped collecting data and storing it in the object. The status of each management collection object is shown in the list of collection objects when you expand **Configuration and Service** and select **Collection Services**.

Collection Services deletes the cycled collection objects that have reached their expired date and time the next time it starts or cycles a collection. The expiration date is associated with the management collection object. Even if you move the collection object to another library, Collection Services will delete the object after it expires.

The expiration date for each management collection object is shown in the Properties for that collection object. To keep the object on the system longer, you simply change the date on the Properties page. Right-click any management collection object in the list and select **Properties** to see the information about that collection. You can specify **Permanent** if you do not want Collection Services to delete your management collection objects for you.

What else can I do with my collection objects?

You can right-click any collection object and select **Create database files** to specify the data categories, range of time within the collection period, and sampling interval that you want to include in the database files. You can also create database files from existing objects.

You can delete a collection object from the system by right-clicking the object and selecting **Delete**. If you do not delete the objects manually, Collection Services will delete them automatically after the expiration date and time.

You can right-click any collection object and select **Graph History** to view the data graphically in the management collection object.

Managing system values

Use Management Central to compare and update your system values across multiple systems in your network. As an administrator, you can manage system values across multiple systems. You can compare the system values on a model system to one or more target systems and then update the target system values to match the values of the model system. If you prefer, you can generate a list that shows the differences in values between the model system and the target system rather than actually changing the values on the target system.

Be sure you have current system value inventories on your target systems. It is possible to have your model system be a target system if you have collected inventory for the model system. You can also

export any system values inventory to a PC file. These PC files provide a history of the inventory and allow you to work with the data in a spreadsheet program or other application.

To compare and update your system values, follow these steps:

1. In Operations Navigator, expand **Management Central**.
2. Expand **Endpoint Systems, System Groups, or My Connections**.
3. Right-click an endpoint system or a system group that you want to be your target system, select **System Values**, and then select **Compare and Update**.
4. Complete the fields on the **Compare and Update** dialog.
 - Select the name of the model system against which you want to compare the target system or systems.
 - Select the categories and values that you want to include in the compare. For each system value that you want to update on the target system, select that item from the **Update** column.
 - Verify the target system or systems that are selected.
5. Click **OK** to perform the task immediately or click **Schedule** to run the task at a later time.

Additionally, if you have questions about where to find the system values in Operations Navigator, use the system value finder. When you have finished this task with Management Central, see *What you can do with Management Central* to learn about other tasks that you can perform.

Manage users and groups across multiple systems with Management Central

Management Central can help you as a system administrator to keep track of the users, groups, and their level of privileges on multiple systems. The following list gives you an idea of the many ways in which Management Central can make your job easier.

Create 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 a command, you can click **Prompt** to select appropriate parameters and values.

Create, edit, and delete 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.

Scan 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.

Collect 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.

Send users and groups

You can send users and groups from one system to multiple endpoint systems or system groups. Unlike the Copy/Paste action, the Send function copies as many user properties as possible to the target systems, including the user name and password, security settings, private authorities, and mail options.

Synchronize unique identifiers

You can synchronize the unique identifiers of users and groups (UID and GID) 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 UID and GID numbers are another way of identifying a user or group to a program. For example, the UID and GID numbers are used by programming interfaces in the Integrated File Systems environment.

Note: All OS/400 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 Management Central. This includes security administration (*SECADM) privileges, all object (*ALLOBJ) privileges, and authority to the profiles with which you are working.

Management Central user definitions

When you use Management Central to manage users and groups, you'll find that user definitions provide an easier way to create new users on multiple endpoint systems or system groups. Simplify your life by creating user definitions for the various types of users on your system. Then, when a request comes in for a new user, all special authorities, auditing values, session startup settings, and other information common to that type of user is already there!

You just 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.

What can I do with user definitions?

When you create a new user from a definition, you can change properties of the new user without affecting the properties defined in the user definition. Or you can simply use the definition properties for each new user you create — all you do is specify a name and password for the user!

When you create a new user from a definition, you can take advantage of these options:

- **Create the user now or later**

You can create the new user immediately or you can schedule a later time when you want the user to

be created. For example, you can create a user definition named Accounting Users, which specifies all the special authorities and other properties that the users in your accounting department need. Then, at any time, you can create one or more new users based on that definition on any endpoint system or system group.

- **Run a command after the user is created**

In the user definition, you can specify a command or program to run on the target system immediately after a user is created successfully on the system. The command or program is run when a user is created from the definition. This can be any command that can be used in the OS/400 batch environment. You cannot run an interactive command. If you don't remember the exact parameters and values that you want to specify on the command, you can click **Prompt** for detailed assistance.

- **Use &USER for the name of the new user**

You can use the replacement variable **&USER** any place in the command where you want the command to substitute the name of the user that is being created. For example, you could specify the command **CRTLIB &USER** to create a library with the user name as the name of the library. This will create a library each time the definition is used to create a user.

Do you want to change the properties of a user across multiple systems? Or delete a user who has left your organization? You can edit and delete users and groups across multiple endpoint systems or system groups using Management Central.

And remember, Management Central can help you in many other ways with managing users and groups across multiple systems.

Creating users across multiple systems

When you use Management Central to work with users across multiple systems, you can create a new user on one or more endpoint systems or on all the systems in a system group. (You can also create a new user based on a stored user definition by expanding **Management Central** and **Definitions**, and then selecting **User**. A list of existing definitions is shown.)

Some of the settings you can specify for the new user include:

- Groups and first selected group options
- Personal information, such as the user's name, location, and mail options
- Security information such as the user's privileges, auditing options, password expiration, and unique identifiers

Follow these quick steps to create users or groups on multiple endpoint systems:

1. Expand **Management Central**.
2. Select **Endpoint Systems** or **System Groups**.
3. Select one or more endpoint systems or system groups.
4. Right-click any selected endpoint system or system group, and select **Users and Groups**.
5. Select **New User** or **New Group**.
6. Use the online help to select all the settings for the new user or group.
7. When you have finished specifying your choices, you can click **OK** to create the user immediately or click **Schedule** to specify when you want the create task to start.

And remember, Management Central can help you in many other ways with managing users and groups across multiple systems.

Editing users across multiple systems

System administrators who use Management Central to work with users across multiple systems can save significant time when changing user and group settings. For example, you could quickly and easily change the authority level for several users or groups across multiple systems. You can specify this editing task and then schedule it to run at a convenient time.

Management Central helps you perform this complex task in these easy steps:

Select a list of users to edit

1. Select one or more endpoint systems or system groups, and right-click any selected system.
2. Select **Users and Groups**, and then select **Edit Users**.
3. Click **Browse** to select from a list of users. Or if you prefer, just type in the user names you want to edit, separating the names with either a comma or a space.
Note: System-defined users cannot be edited; these user names typically start with a Q, like QSECOFR.
4. Select the users you want to edit.
5. Click **OK** to close the Browse window.

Select the settings to edit

1. Select one of the options in the **Category** field. For example, you might select Privileges.
2. In the list of settings, select one or more settings that you want to edit for that category. If you selected the Privileges category, you might select All object access, Job control, and Save/restore in the list of settings.
3. Click the Properties button to change the selected settings.
4. In the **Edit Users** window, only the settings you selected to change are available. (To change other settings, just click Cancel to go back and change your selection.)
5. After you have made your changes, click **OK** to close this window.
6. Continue by selecting other settings or even other categories that you want to edit for the selected users.

Review a summary of changes

1. Select **Summary** in the **Category** field.
2. In the list of settings, you can see all the settings that you requested to change, the new values that you specified, and the category of each setting.

Correct any settings

1. Select any setting that you want to delete or change from the list.
2. Click the Clear button to delete the setting from the list.
3. To specify a different value for this setting, select the appropriate category, select the setting again, and then click the Properties button to specify a new value.

Start the edit task

Just click **OK** to start the Edit Users task immediately or click **Schedule** to specify when you want the task to start.

And remember, Management Central can help you in many other ways with managing users and groups across multiple systems.

Deleting users across multiple systems

System administrators who use Management Central to work with users across multiple systems can save significant time when deleting users and groups across multiple systems. When you use Management Central 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.

Follow these quick steps to delete users from multiple endpoint systems:

1. Expand **Management Central**.
2. Select **Endpoint Systems** or **System Groups**.
3. Select one or more endpoint systems or system groups.
4. Right-click any selected endpoint system or system group, and select **Users and Groups**.
5. Select **Delete Users**.
6. Specify the users that you want to delete. Click **Browse** to select from a list of all users in the central system inventory for the selected endpoint systems or system groups.
7. Select the action you want to take if the user owns objects on any system. Click **Scan for Owned Objects** to see what objects the selected users own on the selected endpoint systems or across the selected system groups.
8. If you select to transfer ownership of the objects to another user, specify the name of that user. Click **Browse** to select from a list of available users.
9. Click **OK** to start the delete task immediately or click **Schedule** to specify when you want the task to start.

And remember, Management Central can help you in many other ways with managing users and groups across multiple systems.

Scanning for owned objects with Management Central

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

To find out what objects any user on an endpoint system owns, follow these steps:

1. Expand **Management Central**.
2. Select **Endpoint Systems** or **System Groups**.
3. Select the endpoint systems or system groups that you want to scan.
4. Right-click any selected system or group, select **Users and Groups**, and then select **Scan for Owned Objects**.
5. Select the name of one or more users, and specify a maximum number of objects to display.
6. Click **OK**.

Collecting an inventory of users and groups

Collecting an inventory can help you manage users and groups across multiple systems. When you collect an inventory of the users and groups on one or more endpoint systems, you can view, search, or export that inventory to a PC file.

Follow these quick steps to collect an inventory of users and groups:

1. Expand **Management Central**.
2. Select **Endpoint Systems** or **System Groups**.
3. Select one or more endpoint systems or system groups.
4. Right-click any selected endpoint system or system group, and select **Inventory**.
5. Select **Collect**.
6. Select **Users and Groups** for **Inventory to collect**.
7. If you want an action to run on the central system when collection completes, select the action from the list. The list contains all actions defined by applications that are currently installed on the system.
8. Click **OK** to start collecting inventory immediately or click **Schedule** to specify when to collect inventory.

View an inventory of users and groups

Viewing an inventory can help you manage users and groups across multiple systems.

How do I view the inventory?

Follow these quick steps to view an inventory of users and groups:

1. Expand **Management Central**.
2. Expand **Endpoint Systems** or **System Groups**.
3. Expand the endpoint system whose inventory you want to view.
4. Expand **Users and Groups**.
5. Select **User Inventory** or **Group Inventory**.

How do I customize the inventory list?

You can customize the inventory list by doing any of the following:

- Sort the inventory entries 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.
- Specify which users or groups should be displayed in the list by selecting **Options** from the menu bar, and then selecting **Include**.
- Specify which columns of information you want to display in the list by selecting **Options** from the menu bar, and then selecting **Columns**.

What actions can I perform from the inventory list?

You can select one or more users in the list and perform any of the following actions:

User objects

Select a type of object (Printer output, Jobs, Server jobs, or Messages) to open a new window where you can view and work with the objects of this type that belong to this user. Or select **Scan for Owned Objects** to find all the objects that the user owns on this system.

New Based On

Create a new user on the endpoint system based on the properties of this user.

Edit

Change the properties for all the selected users on the endpoint system.

Delete

Delete the selected users from the endpoint system.

Send

Send the selected users to other endpoint systems. When you send users to another system, as many user properties as possible are copied to the new system, including the user name and password, security settings, private authorities, and mail options that are not copied when you select to copy a user to another system.

Properties

View the properties for a single selected user.

You can also search the inventory or export it to a PC file.

Searching an inventory of users and groups

Searching an inventory can help you manage users and groups across multiple systems. When you search an inventory of users or groups, you can extend your search beyond the name and description fields to any of the properties of the user or group. For example, you can search the inventory to see who has security officer privileges, as well as to query other profile properties.

You can define a more complex search by clicking **And** or **Or** to search on additional fields. For example, if you were searching for all users on these endpoint systems with security officer privileges, you could narrow the search to users in your Accounting department with security officer privileges by clicking **And** and selecting **Department** and **Accounting**.

Follow these quick steps to begin your search:

1. Expand **Management Central**.
2. Select **Endpoint Systems** or **System Groups**.
3. Select one or more endpoint systems or system groups whose inventory you want to search.
4. Right-click any selected endpoint system or system group, select **Inventory**, and then select **Search**.
5. Select **Users and Groups** for **Inventory to Search**.
6. On the **Basic** tab, specify a string you want to search for in the name and description fields of the inventory. You may leave this field blank if you want to search all users for criteria that you specify on the Advanced tab.
7. If you want to specify additional fields to search in the users and groups inventory, click the **Advanced** tab. When you specify advanced search criteria, the search results include all items that meet both the basic criteria and the advanced criteria. Advanced search is available only when both the central system and the endpoint systems are running OS/400 V5R1 or later.
8. Click **Search**.

You can also view the inventory or export it to a PC file.

Exporting an inventory

Exporting an inventory can help you manage your systems quickly and efficiently.

When you export your inventory into a PC file, you save a history of your inventory, which you can use to work with the data in a spreadsheet program or other application. You can also search the PC files for additional fields that are stored in the inventory but are not available in the basic or advanced search.

Follow these quick steps to export an inventory:

1. Expand **Management Central**.
2. Expand **Endpoint Systems** or **System Groups**.
3. Right-click the endpoint system or system group to export from and select **Inventory**.
4. Select **Export**.
5. Select the type of inventory you want to export.

6. Click **Export**.
7. Select the PC folder where you want to save the inventory.
8. Specify the type of file in which you want to save the inventory. You can select any of the following formats for your inventory data: ASCII Tab Delimited Text (*.txt), Comma Separated Variable (*.csv), Web Page (*.html), or Lotus 123. Compatible (*.csv).
9. Specify the name of the file in which you want to save the inventory.
10. Click **Save**.

You can export any of the following types of inventory: hardware inventory, software inventory, users and groups inventory, and fixes inventory. Following a slightly different procedure, you can also export a system values inventory.

Sending users and groups to multiple systems

Sending users and groups from one system to multiple endpoint systems or system groups can help you manage users and groups across multiple systems. The Send function copies more user properties to the target systems than the Copy/Paste action. User properties that are sent include user name and password, security settings, private authorities, and mail options. (For example, a LAN server password is sent in addition to the AS/400 password, while Copy/Paste includes only the AS/400 password.) 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 (UID) of the user on the target system based on the UID of the user being sent.

If you want your users to receive all their mail on the same system, regardless of how many systems they can sign onto in the network, you would select to use the system specified in the properties of the user being sent. If you are moving users to a new system and plan to delete them from the original system, you would select to use the local target system as the mail server.

Follow these quick steps to send users from an endpoint system:

1. Expand **Management Central**.
2. Expand **Endpoint Systems** or **System Groups**.
3. Expand the endpoint system from which you want to send users, and then expand **Users and Groups**.
4. Select **User Inventory**. If the inventory has never been collected, no list can be displayed. For information on how to collect the inventory, see Collect an inventory of users and groups.
5. Select one or more users to send.
6. Right-click any selected user and select **Send**.
7. On the **Send Users - General** page, select the endpoint systems or system groups to which you want to send the selected users.
8. Click the **Options** tab to specify the action to be taken if any user that you are sending already exists on the target system. 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.
9. Click **Advanced** to specify advanced send options. The advanced send options include specifying the mail system for the user and synchronizing the unique identifier (UID) of the user on the target system based on the UID of the user being sent.
10. Click **OK** to start the send task immediately or click **Schedule** to specify when you want the task to start.

And remember, Management Central can help you in many other ways with managing users and groups across multiple systems.

Synchronizing unique identifiers of users and groups (UID and GID)

When you are managing users and groups across multiple systems, it is important to synchronize the UID and GID numbers to ensure that each of these unique identifiers points to the same user or group on every system. This is especially important when you are working with systems in a clustering environment or a system with logical partitions. The UID and GID numbers are another way of identifying a user or group to a program. For example, the UID and GID 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.

Warning: Temporary Level 4 Header

Synchronizing unique identifiers when creating users or groups: To synchronize the unique identifiers (UID and GID) when creating users or groups across multiple endpoint systems, follow these steps:

1. Expand **Management Central**.
2. Select **Endpoint Systems** or **System Groups**.
3. Select one or more endpoint systems or system groups where you want to create users or groups.
4. Right-click any selected endpoint system or system group and select **New Users** or **New Groups**.
5. If you are creating a new user, you must specify the name for the new user. Then click **Capabilities**, and click the **Unique Identifier** tab. Select **Find a unique number across all selected systems**. The central system will find a unique number based on the inventory across all selected systems.
6. If you are creating a new group, you must specify the name for the new group. Then click **Networks**, and click the **Unique Identifier** tab. Select **Find unique numbers across all selected systems**. The central system will find a unique set of numbers based on the inventory across all selected systems.
7. When you are finished specifying the settings for the new user or group, click **OK** to start the create task immediately or click **Schedule** to specify when you want the task to start.

Synchronizing unique identifiers when editing users or groups: To synchronize the unique identifiers (UID and GID) when editing users or groups across multiple endpoint systems, follow these steps:

1. Expand **Management Central**.
2. Select **Endpoint Systems** or **System Groups**.
3. Select one or more endpoint systems or system groups where you want to edit users or groups.
4. Right-click any selected endpoint system or system group and select **Edit Users** or **Edit Groups**.
5. If you are editing users, specify the names of the users. Then select **Unique Identifier** as the **Category** under **Settings to edit for each user**. In the list, select **UID number** and then click the Properties button. Select **Find a unique number across all selected systems**. The central system will find a unique number based on the inventory across all selected systems.
6. If you are editing groups, specify the names of the groups. Then select **Unique Identifiers** as the **Category** under **Settings to edit for each group**. In the list, select **UID and GID numbers** and then click the Properties button. Select **Find unique numbers across all selected systems**. The central system will find a unique set of numbers based on the inventory across all selected systems.
7. When you are finished specifying the settings you want to edit, click **OK** to start the edit task immediately or click **Schedule** to specify when you want the task to start.

Synchronizing unique identifiers when sending users or groups: To synchronize the unique identifiers (UID and GID) when sending users or groups to a group of endpoint systems, follow these steps:

1. Expand **Management Central**.
2. Expand **Endpoint Systems** or **System Groups**.
3. Expand the endpoint system from which you want to send users or groups.
4. Expand **Users and Groups**.
5. Select **User Inventory** or **Group Inventory**.
6. Select the users or groups to send.
7. Right-click any selected user or group, and click **Send**.
8. On the **General** page, select the endpoint systems or system groups to which you want to send the selected profiles.
9. Click the **Options** tab to specify the action to be taken if any user or group that you are sending already exists on the target system. You can select not to change the profile that already exists, or you can select to update the existing profile with the settings from the profile you are sending.
10. Click **Advanced** to specify advanced send options.
11. Select whether to synchronize the unique identifiers on the target system based on the unique identifiers of the profile being sent.
12. Click **OK** to start the send task immediately or click **Schedule** to specify when you want the task to start.

If you are simply moving users or groups from one system to another (not in a clustering or logical partitions environment), you may save time by choosing not to synchronize the unique identifiers of the users or groups being sent.

Running commands with Management Central

Management Central enables you to define an action or a task and then perform that action or task on multiple endpoint systems or system groups. You can select one or more endpoint systems or system groups and then specify a command to run on those systems. You can click **Prompt** for assistance in entering or selecting a command. You can choose to run the command immediately or schedule it to run at a later time.

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, a task is created.

Why should I perform commands with Management Central?

If your day-to-day operations require you to perform repetitive tasks, you can take advantage of the Management Central command definition. For example, you could use a command definition to do any of the following tasks:

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

In fact, any control language (CL) command that you can run in batch, you can now send to multiple systems at the same time. Just create the command definition, and then run the command on endpoint systems or system groups.

Use Management Central to do more than run commands. You can do many of the tasks required to manage your systems quickly and efficiently with this powerful tool.

Creating a command definition with Management Central

Management Central enables you to define a command and then run that command on multiple systems. To do this, you must first create the command definition, using a few quick steps:

1. In Operations Navigator, expand **Management Central**.
2. Expand **Definitions**.
3. Right-click **Command** and select **New Definition**.
4. Specify a name for the definition, a brief description, and the command to be run. You can click **Previous Commands** to select from a list of commands you have previously run, or you can click **Prompt** to get assistance in entering or selecting a command.
5. To specify options concerning the job log or inquiry messages, click the **Options** tab.
6. Click the **Sharing** tab to specify whether you want to share this command definition with other users.
7. Click **OK**.

After you create the command definition, you can run the command on endpoint systems or system groups.

Running a command with Management Central on systems or groups

Management Central enables you to define a command and then run that command on systems or groups. To do this, you must first create a command definition. After you create the command definition, you can run the command using a few quick steps:

1. Right-click a command definition and select **Run**.
2. Select the endpoint systems or system groups on which you want to run the command.
3. Click **OK** to start the command task immediately or click **Schedule** to specify how often you want to run this task and when you want the task to start.

You can also run a command without creating a definition by following these steps:

1. In Operations Navigator, expand **Management Central**.
2. Select **Endpoint Systems** or **System Groups**.
3. In the right pane, select the endpoint systems or system groups where you want to run the command.
4. Right-click the selected endpoint systems or system groups, and select **Run Command**.
5. Specify the command to be run. You can click **Previous Commands** to select from a list of commands you have previously run, or you can click **Prompt** to get assistance in entering or selecting a command.
6. To specify options concerning the job log or inquiry messages, click the **Options** tab.
7. Click **OK** to start the command task immediately or click **Schedule** to specify how often you want to run this task and when you want the task to start.

Use Management Central to do more than run commands. You can do many of the tasks required to manage your systems quickly and efficiently with this powerful tool.

Scheduling tasks or jobs with Management Central scheduler

Management Central and Operations Navigator provide two different tools you can use to schedule tasks or jobs. Operations Navigator provides an integrated scheduler (the **Management Central Scheduler**) and the **Advanced Job Scheduler**. To find out about these two tools, keep reading.

What is the Management Central Scheduler and why should I use it?

Operations Navigator provides an integrated scheduler, the Management Central Scheduler, to organize when you want your tasks to occur. You have the option of choosing to perform a task immediately or choosing a later time.

You can use the Management Central scheduler to schedule a variety of tasks. For example, you could automate the process of collecting an inventory (for example, of hardware, software, or fixes) on whichever

day fits your operating schedule. You might schedule such a collection to occur every Saturday night at 10 PM. You could 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 simply want to install a set of fixes once. Using the scheduler function gives you the flexibility to do your work when it's convenient for you to do it. 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:

- Create, delete, edit, and send users and groups across multiple endpoint systems
- 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
- Run commands on selected endpoint systems and system groups
- Delete the save files and cover letters for selected fixes on selected endpoint systems and system groups
- Send fixes or packages of files and folders to selected endpoint systems and system groups
- Start installing fixes, uninstalling fixes, or installing fixes permanently
- Start and stop Collection Services on selected endpoint systems and system groups

To get started using this tool, see the Management Central Scheduler topic or the Operations Navigator online help, where you can learn how to schedule various Management Central tasks.

What is the Advanced Job Scheduler and why should I use it?

The Advanced Job Scheduler is a separate licensed program (5722-JS1) 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, you simply click the Schedule button from any Operations Navigator dialog to schedule tasks and jobs. To find more information about installing and using this tool, see the Advanced Job Scheduler topic.

Don't forget to use Management Central to do more than schedule tasks. You can do many of the tasks required to manage your systems quickly and efficiently with this powerful tool.

Management Central Scheduler

Operations Navigator lets you choose which scheduler you want to use for scheduling your tasks. Scheduling a task with Management Central is as easy as clicking a button — the **Schedule** button! To schedule a later time to perform a task, click **Schedule** from the appropriate dialog.

You can schedule a task to run just once, in which case the task runs a single time beginning at the specified date and time. Or, you can select any of the following options:

- **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. If you want full calendar management, you should choose Advanced Job Scheduler.

Scheduling jobs with the Advanced Job Scheduler

The Advanced Job Scheduler licensed program (5722-JS1) is a robust scheduler that allows unattended job processing 24 hours a day, 7 days a week. Choose standard, fiscal, or user-defined calendars to quickly describe and set up any job scheduling scenario that you need in the Advanced Job Scheduler. You may also view job completion history and manage notification of a job's status.

It is not necessary to install the Advanced Job Scheduler licensed program on each endpoint system in your Management Central network. When you install the Advanced Job Scheduler on the central system, jobs or tasks that you define on an endpoint system will gather job information that is needed from the central system. However, you must set up all job definition information on the central system.

If systems in your network have the Advanced Job Scheduler installed locally, you may schedule tasks outside of the Management Central network. Under **My AS/400 Connections** in Operations Navigator, you have access to the Advanced Job Scheduler on that local system when you expand **Work Management**.

Install and customize the Advanced Job Scheduler

See the following information to install and customize the Advanced Job Scheduler.

Install the Advanced Job Scheduler

Follow these steps to install the Advanced Job Scheduler.

Customize the Advanced Job Scheduler

If you have installed the program and this is your first time using the Advanced Job Scheduler, customizing is your next step. See how to specify the general properties that are used by the Advanced Job Scheduler, according to your needs.

The following information will help you use the Advanced Job Scheduler for scheduling and working with jobs.

Schedule a job

Schedule a job and specify the commands that are associated with the job.

Schedule job groups

Set up and schedule a series of jobs that run consecutively in a specified order. Jobs within a job group require completion before the next job is submitted for processing.

Schedule job dependencies

Set up jobs or groups of jobs that are dependent on each other. You can select the type of dependency that reflects how jobs are processed in your environment.

Monitor job activity

View a job or a job group's history or status. You can also set up the activity retention, which is how long you want to retain the activity records for a job.

Installing Advanced Job Scheduler

To install the Advanced Job Scheduler you must have previously installed Client Access Express. Then, follow these steps to install the Advanced Job Scheduler:

1. From your **Operations Navigator** window, click **File** from the menu bar.
2. Select **Install Plug-Ins**.
3. Select the source system where the Advanced Job Scheduler is installed and click **OK**. Check with the system administrator if you are not sure what source system to use.
4. Enter your AS/400 **User ID** and **Password**, and click **OK**.
5. Select **Advanced Job Scheduler** from the Plug-in selection list.
6. Click **Next** and then click **Next** again.

7. Click **Finish** to complete and exit the setup.

You have now installed the Advanced Job Scheduler.

To locate the scheduler, follow these steps:

1. Expand **Management Central**.
2. Click **Scan Now** in response to the message that Operations Navigator has detected a new component. You may see this message again when you access systems from **My AS/400 Connections**.
3. Expand **My AS/400 Connections** → the AS/400 server that has the Advanced Job Scheduler licensed program installed → **Work Management** → **Advanced Job Scheduler**.

Be aware that the Advanced Job Scheduler uses the time on the PC for the scheduled time. It does not use the time on the central system or the endpoint system. If the time on the PC is earlier than the central system time, you may schedule a job that does not run until the next scheduled day. For example, if the current central system time is 11:30 am, and the current PC time is 11:25 am, a job that you schedule to run at 11:28 am will not run until the next scheduled day. The time has already passed on the central system.

When you have finished this preliminary work with the Advanced Job Scheduler, you are ready to begin customizing the Advanced Job Scheduler. See scheduling jobs with the Advanced Job Scheduler to choose another task.

Customizing the Advanced Job Scheduler

To customize the Advanced Job Scheduler, select from the following tasks:

- **Assign the general properties**
Specify how long to retain activity and log entries for the Advanced Job Scheduler, as well as the period that jobs will not be allowed to run. You may specify the working days that jobs will process, and whether an application is required for each scheduled job. If you have a paging product installed, you can also set up the command that will be used to send a page whenever a job completes or fails.
- **Create and work with applications**
Applications are jobs that are grouped for processing. They are broader than job groups and do not necessarily process sequentially. Jobs in applications can process simultaneously and one job does not have to wait for another to process. All jobs within the application can be worked with and can have their own set of job defaults.
- **Set up a calendar**
Set up a calendar of selected days for scheduling a job or job group. This calendar can specify the dates to be used for scheduling a job, or it can be used in conjunction with other schedules.
- **Set up a holiday calendar**
Set up a calendar for days that you do not want to allow processing for a job. Alternate days may be specified for each exception day, or processing can be skipped completely for that day.
- **Work with library lists**
Library lists are user defined lists of libraries that are used by the Advanced Job Scheduler when a job is processing.
- **Work with command variables**
A command variable (previously known as a parameter) is a variable you may store and use in jobs submitted through the Advanced Job Scheduler. Examples of command variables include the beginning of each month, a division number, a company number, and so on.
- **Work with job controls**
Job controls are the defaults assigned to a job as you add it to the job scheduler.

For more information about these and other properties, refer to the online help for the Advanced Job Scheduler in Operations Navigator. When you have finished this preliminary work with the Advanced Job Scheduler, you are ready to begin scheduling jobs. See scheduling jobs with the Advanced Job Scheduler to choose another task.

Assigning the general properties for the Advanced Job Scheduler: Assign the General properties used by Advanced Job Scheduler. You may specify how long to retain activity records for jobs, as well as the period that jobs will not be allowed to run. You may specify the working days that jobs are allowed to process, and whether an application is required for each submitted job. You may have a paging product installed which allows you to receive a page(message) when a job ends. You can define the paging command that will send a page whenever a job completes or fails.

To set up the General properties for the Advanced Job Scheduler, follow these steps:

1. Expand **Work Management** from your **Operations Navigator** window.
2. Right-click **Advanced Job Scheduler** and select **Properties**.
3. Specify the **Activity Retention**. The activity retention is how long you want to retain the activity records for jobs. The possible values are 1 to 999 days or occurrences. Click **Days** to specify if you want to keep activity for a certain number of days, or click **Occurrences per job** if you want to keep activity for a certain number of occurrences per job.
4. Specify the **Log retention**. The log retention specifies, in days, how long you want to retain Advanced Job Scheduler log entries.
5. You may specify a **Reserved period**. Jobs will not run during this time.
6. Specify the working days from the list.
If a day is selected, it is designated as a working day and may be referenced when scheduling jobs.
7. Click **Application required for scheduled job** to designate whether an application is required for each scheduled job.
Applications are jobs that have been grouped together for processing. This cannot be selected if existing jobs do not contain an application.
8. Click **Base periodic frequency on start time** to base the next run time on the start time for jobs that are scheduled to run periodically. For instance, a job is to run every 30 minutes, starting at 8:00 am. (For a job to run around the clock, you would also specify 7:59 am as the ending time.) The job runs for a total of 20 minutes. With this field checked, the job would run at 8:00 am, 8:30 am, 9:00 am, etc. If this field is not checked, the job would run at 8:00 am, 8:50 am, 9:40 am, 10:30 am, etc.
9. You may specify the values for the **paging command**. This step only pertains if you have a paging product installed. The paging command is specified by your paging software and is used to send a pager message to a recipient that you specify. The command you specify is used to send pager messages for normal and abnormal completions of job scheduled entries.

If you chose to have an application required for certain jobs, go to working with applications. If your job will not require applications, go to setting up a calendar, or see customizing the Advanced Job Scheduler to choose another task.

Creating and working with applications for the Advanced Job Scheduler: **Applications** are jobs that have been grouped together for processing. For example, you might have a series of jobs that you use for payroll that you want to group together for an accounting process.

You can display all the existing applications on your system. You can add a new application, add a new application based on an existing application, or remove an application. You can also select an application and display its properties to make changes.

To create a new application, follow these steps:

1. Expand **Work Management** from your **Operations Navigator** window.

2. Right-click **Advanced Job Scheduler** and select **Properties**.
3. Click the **Applications** tab.
4. Click **New** and type a name for the application.
5. You may type a description for the application.
6. Choose the contacts for the application.
Contacts are the names of users who are contacted if you have a problem with a job within the application. You may specify up to 5 contacts per application. You may also choose to add or remove contacts from the contact list.
7. You may type additional information to help you identify the application.
The information is associated with the new application. This information may be useful if any problems occur.

See customizing the Advanced Job Scheduler to choose another task.

Setting up a calendar for the Advanced Job Scheduler: A **scheduling calendar** is a calendar of selected days that you may use for scheduling a job or job group. You may display scheduling calendars, add a new scheduling calendar, add a new scheduling calendar based on an existing one, or remove an existing calendar, provided it is not in use by a currently scheduled job.

You may select a calendar and display its properties to make changes. When you select a calendar, the details of the calendar are displayed under Details.

To set up a scheduling calendar, follow these steps:

1. Open **Work Management** from your **Operations Navigator** window.
2. Right-click **Advanced Job Scheduler** and select **Properties**.
3. Click the **Scheduling Calendars** tab.
4. Click **New** and type a name for the calendar.
5. You may type a description for the calendar.
6. Choose a **Reference calendar** if applicable.
This is a calendar that was previously set up, and its properties will be applied to the new calendar as if you merged the two calendars. You will not have reference calendars if this is your first time using the Advanced Job Scheduler.
7. Select the dates that you want to include on your calendar.
You must specify whether each date you have selected is for the current year or for every year, before you can add another date to the calendar. Otherwise, any date you select will be unselected when you click a different date.
8. Specify if you want certain days of the week to be included on the calendar.

See customizing the Advanced Job Scheduler to choose another task.

Setting up a holiday calendar for the Advanced Job Scheduler: A **holiday calendar** is an exception calendar for days that you do not want to process an Advanced Job Scheduler job. Alternate days may be specified for each exception day that you specify in a holiday calendar. You may display holiday calendars, add a new holiday calendar, add a new holiday calendar based on an existing one, or remove an existing calendar, provided it is not in use by a currently scheduled job.

You may select a calendar and display its properties to make changes. When you select a calendar, the details of the calendar are displayed under Details.

To set up a holiday calendar, follow these steps:

1. Expand **Work Management** from your **Operations Navigator** window.
2. Right-click **Advanced Job Scheduler** and select **Properties**.

3. Click the **Holiday Calendars** tab.
4. Click **New** and type a name for the calendar.
5. You may type a description for the calendar.
6. Choose a **Reference calendar** if applicable.
This is a calendar that was previously set up, and its properties will be applied to the new calendar as if you merged the two calendars. You will not have reference calendars if this is your first time using the Advanced Job Scheduler.
7. Select the dates that you want to include on your calendar.
You must specify whether each date you have selected is for the current year or for every year, before you can add another date to the calendar. Otherwise, any date you select will be unselected when you click a different date.
8. Select an alternate day for the job to run. You may choose the previous working day, next working day, a specific date or not at all. To select a specific date, click **Specific alternate date**, and type the date.
9. Select specific days of the week to be included on the calendar.

See customizing the Advanced Job Scheduler to choose another task.

Working with library lists for the Advanced Job Scheduler: A **library list** is a user-defined list of libraries that is used by the Advanced Job Scheduler job to search for information it needs while processing. You may display library lists, add a new library list, add a new library list based on an existing one, or remove a library list, provided that it is not being used by a currently scheduled job.

You may select a list and display its properties to make changes. You may place up to 250 libraries on the library list.

To add a new library list, follow these steps:

1. Open **Work Management** from your **Operations Navigator** window.
2. Right-click **Advanced Job Scheduler** and select **Properties**.
3. Click the **Library Lists** tab.
4. Click **New** and type a name for the library list.
5. You may type a description for the library list.
6. Click **Browse** to see a list of existing libraries, and select a library.
7. Click **Add** to add the list of selected libraries.

See customizing the Advanced Job Scheduler to choose another task.

Working with command variables for the Advanced Job Scheduler: **Command variables** (previously known as parameters) are variables that you store in the Advanced Job Scheduler and use in jobs submitted through the Advanced Job Scheduler. Command variables contain information that will be replaced inside the command string of a scheduled job. Examples of command variables include the beginning of each month, a company division number, a company number and so on. You may display command variables, add a new command variable, add a new command variable based on an existing one, or remove a command variable, provided it is not currently in use by a scheduled job.

You may select an existing command variable and display its properties to make changes.

To add a new command variable, follow these steps:

1. Open **Work Management** from your **Operations Navigator** window.
2. Right-click **Advanced Job Scheduler** and select **Properties**.
3. Click the **Command Variables** tab.
4. Click **New** and type a name for the command variable.
5. You may type a description for the command variable.

6. Type the length of the command variable. The length can range from 1 to 90.
7. Choose how you want to supply the replacement value:
 - a. Specify the data to use for the command variable. You may use any character in this field. The number of characters in the data cannot be greater than the length specified in the Length field.
 - b. Type a formula to calculate the date. (For examples, see the online Help.)
 - c. Type the program name that you use to retrieve the replacement value.
 - d. Type the library that you use to retrieve the replacement value.
 - e. Choose whether you want the replacement value retrieved from the system operator at run time.

See customizing the Advanced Job Scheduler to choose another task.

Working with job controls for the Advanced Job Scheduler: Job controls are the defaults assigned to a job as you add it to the job scheduler. Job control defaults include such things as default application, calendar, holiday calendar, and so on. All the existing job controls on your system are displayed on this page. You may select a job control and display its properties to make changes.

To work with a job control, follow these steps:

1. Open **Work Management** from your **Operations Navigator** window.
2. Right-click **Advanced Job Scheduler** and select **Properties**.
3. Click the **Job Controls** tab.
4. Select the job control you want to work with and click **Properties**.

See customizing the Advanced Job Scheduler to choose another task.

Schedule a job

To schedule a new job, follow these steps:

1. Open **Work Management** from your **Operations Navigator** window.
2. Right-click **Advanced Job Scheduler**.
3. Right-click **Scheduled Jobs** and select **New Scheduled Job**.

Refer to the online help for more information as you fill in details for the new job. See scheduling jobs with the Advanced Job Scheduler to choose another task.

Schedule a job group

Job groups are jobs that are grouped together to run consecutively in the order specified. A normal completion is required for each job in the group before the next job in the group is submitted for processing. If any job in the group does not complete normally, the processing stops for that group. To schedule a new job group, follow these steps:

1. Open **Work Management** from your **Operations Navigator** window.
2. Right-click **Advanced Job Scheduler**.
3. Right-click **Job Groups** and select **New Job Group**.

Refer to the online help for more information as you fill in details for the new job group. See scheduling jobs with the Advanced Job Scheduler to choose another task.

Job dependencies

The Advanced Job Scheduler allows you to set up dependencies that reflect how jobs are processed in your environment. Dependencies determine when a job or group of jobs can run. You can select to have all dependencies met before a job can run, or you can have at least one dependency met before the job can run. Dependencies include the following:

Job dependencies

Job dependencies refer to predecessor and successor relationships for jobs. Predecessor jobs are

those that must run before the successor job will run. A successor job is a job that runs after all the predecessor jobs have been processed. There can be multiple successor jobs for a single predecessor job.

Active dependencies

Active dependencies are lists of jobs that cannot be active when the selected job is to be submitted. If any of the jobs are active, the Advanced Job Scheduler will not let the specified job run. The selected job will be delayed until all the jobs in the list are inactive.

Resource dependencies

Following are the five types of resource dependencies:

- File
The job is dependent upon the status of a file in order to be processed.
- Object
The job is dependent upon the status of an object in order to be processed.
- Hardware configuration
The job is dependent upon the status of a hardware configuration in order to be processed.
- Network file
The job is dependent upon the status of a network file in order to be processed.
- Subsystem
The job is dependent upon the status of a subsystem in order to be processed.

To work with job dependencies, follow these steps:

1. Open **Work Management** from your **Operations Navigator** window.
2. Expand **Advanced Job Scheduler**.
3. Double-click **Scheduled Jobs**.
4. Right-click the **Job Name** whose dependencies you want to work with.
5. Select one of the following: **Job Dependencies**, **Active Dependencies** or **Resource Dependencies**. Refer to the online help for more information.

See scheduling jobs with the Advanced Job Scheduler to choose another task.

Monitoring job activity for the Advanced Job Scheduler

The Advanced Job Scheduler allows you to view your jobs activity through the following:

Scheduled Job Activity

The scheduled job activity allows you to specify how long the Advanced Job Scheduler activity records are to be retained. The possible values are 1 to 999 days or occurrences. You can specify to keep activity for a certain number of days, or for a certain number of occurrences per job. The following details about a scheduled job are displayed:

- Name
The name of the scheduled job.
- Group
The name of the job group for the job.
- Sequence
The sequence number of the job within the group, if the job is in a job group.
- Completion Status
The estimated time until the job completes processing. This can range between 0 and 100 percent.
- Started
When the job started running.

- Ended
When the job ended.
- Elapsed Time

To specify the activity retention, follow these steps:

1. Open **Work Management** from your **Operations Navigator** window.
2. Expand **Advanced Job Scheduler**.
3. Right-click **Scheduled Job Activity** and select **Properties**.

To view the scheduled job activity details, follow these steps:

1. Open **Work Management** from your **Operations Navigator** window.
2. Expand **Advanced Job Scheduler**.
3. Double-click **Scheduled Job Activity**.



Activity Log

The activity log displays activity within the scheduler such as a job added, changed, or submitted. Security violations, sequences processed by a scheduled job, and any errors received are displayed. The dates and times for the previous activities are also displayed. To view detailed message information, double-click on a date and time.

See scheduling jobs with the Advanced Job Scheduler to choose another task.

Chapter 5. Redbooks

View the following Redbooks from the International Technical Support Organization (ITSO) to find out more information about Management Central.

- **Managing AS/400 V4R4 with Operations Navigator** 
Operations Navigator brings a Windows-like graphical interface to configuring, monitoring, and managing the OS/400 environment. This book gives you insight into the wide range of AS/400 functions available through the AS/400 Operations Navigator graphical interface that comes packaged with AS/400 Client Access Express for Windows V4R4M0. It provides you with a moderate level overview of the AS/400 Operations Navigator interface and functionality, correlates Operations Navigator functions with corresponding OS/400 command functions, and, in many cases, gives tips on how to use these functions. This publication is intended to help two sets of AS/400 users who have some level of management responsibilities for an AS/400 system: those familiar with the OS/400 command level interface to system facilities and those new to the OS/400, but who are familiar with Windows-like graphical interfaces to system facilities.
- **Management Central: A Smart Way to Manage AS/400 Systems** 
Discover the benefits of Management Central and more. Management Central is a key component of Operations Navigator that provides AS/400 administrators with the ability to manage multiple AS/400 systems that are interconnected across a TCP/IP network. It provides several constructs, which help to manage groups of systems and their associated resources. This redbook discusses the capabilities of each of the functions that are available in V4R4, and introduces you to the graphical user interface way of managing the systems in your network. Find out how easy it is to manage your network with the click of a mouse!

For information about installing and getting started with Operations Navigator, be sure to see the Operations Navigator topic in the Information Center.



Printed in U.S.A.