IBM

IBM Systems - iSeries

Systems Management
Work with monitors

Version 5 Release 4

IBM

IBM Systems - iSeries

Systems Management
Work with monitors

Version 5 Release 4

> **Note**
> Before using this information and the product it supports, read the information in "Notices," on page 21.

Contents

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? iSeries™ Navigator provides the technology you need to do systems management tasks across one or more servers simultaneously.

Click Management Central in iSeries Navigator to find easy-to-use systems management functions that come to you as part of your base operating system. Management Central in iSeries Navigator allows you to manage one or more systems through a single central system. 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 systems 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.

With iSeries Navigator for Wireless, administrators have even more flexibility in how they access and interact with Management Central. The iSeries Navigator for Wireless overview contains tips on which devices to use, how to install and configure the required elements, and a comprehensive overview of the functions.

> **Related information**
> iSeries Navigator for Wireless overview

Work with Management Central

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

iSeries Navigator monitors

Monitors display current information about the performance of your systems. Additionally, you can use them to carry out predefined actions when a specific event occurs. You can use the system, message, job, file, and B2B transaction monitors to display and monitor information about your systems. The system and job monitors use the performance data collected by Collection Services.

The monitors included in iSeries Navigator use Collection Services data to track the elements of system performance of specific interest to you. Moreover, they can take specified actions when certain events, such as the percentage of CPU utilization or the status of a job, occur. You can use monitors to see and manage system performance as it happens across multiple systems and groups of systems.

With the monitors, you can start a monitor, and then turn to other tasks on your server, in iSeries Navigator, or on your PC. In fact, you could even turn your PC off. iSeries Navigator continues monitoring and performing any threshold commands or actions you specified. Your monitor runs until you stop it. You can also use monitors to manage performance remotely by accessing them with iSeries Navigator for Wireless.

iSeries Navigator provides the following types of monitors:

System monitor
> 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. For example, if you are monitoring the average CPU utilization on your server, you can click any collection point

on the graph to see a details chart that shows the 20 jobs with the highest CPU utilization. Then, you can right-click any of these jobs to work directly with the job.

Job monitor

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. To work directly with a job, just right-click the job from the list that is shown in the Job Monitor window.

Message monitor

Find out whether your application completes successfully or monitor for specific messages that are critical to your business needs. From the Message Monitor window, you can see the details of a message, reply to a message, send a message, and delete a message.

B2B activity monitor

If you have an application like Connect for iSeries configured, you can use a B2B activity monitor to monitor your B2B transactions. You can view a graph of active transactions over time, and you can run commands automatically when thresholds are triggered. You can search for and display a specific transaction as well as view a bar graph of the detailed steps of that specific transaction.

File monitor

Monitor one or more selected files for a specified text string, for a specified size, or for any modification to the file.

To find out more about monitors, see the following topics:

Monitor concepts

Monitors can display real-time performance data. Additionally, they can continually monitor your system in order to run a selected command when a specified threshold is reached. Learn how monitors work, what they can monitor, and how they can respond to a given performance situation.

The system monitors display the data stored in the collection objects that are generated and maintained by Collection Services. The system monitors display data as it is collected, for up to one hour. To view longer periods of data, you should use Graph history. You can change the frequency of the data collection in the monitor properties, which overrides the settings in Collection Services.

You can use monitors to track and research many different elements of system performance and can have many different monitors running simultaneously. When used together, the monitors provide a sophisticated tool for observing and managing system performance. For example, when implementing a new interactive application, you might use a system monitor to prioritize a job's resource utilization, a job monitor to watch for and handle any problematic jobs, and a message monitor to alert you if a specified message occurs on any of your systems.

Setting thresholds and actions

When you create a new monitor, you can specify actions you want to occur when the system metric reaches a specified threshold level, or an event occurs. When threshold levels or events occur, you can choose to run an i5/OS™ command on the endpoint systems, such as sending a message or holding a job queue. Additionally, you may choose to have the monitor carry out several predefined actions such as updating the event log and alerting you by either sounding an alarm on your PC or launching the monitor. Finally, you can automatically reset the monitor by specifying a second threshold level which causes the monitor to resume normal activity when it is reached.

Job monitors and Collection Services

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

The metrics that are available for a job monitor are:

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

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

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

> **Related information**
>
> About Collection Services

The QYRMJOBSEL job:

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

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

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

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

For example, suppose the job monitor server starts a job monitor at 5 minute collection intervals. Then another monitor that is using Collection Services starts, but uses a smaller interval. As a result, the QYRMJOBSEL receives the data at the smaller or more frequent interval. If the smaller interval is 30 seconds, there will be a 10 time increase in the amount of data QYRMJOBSEL processes, thereby increasing the need for CPU resources.

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

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

QZRCSRVS jobs and their impact on performance:

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

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

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

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

Tuning QZRCSRVS jobs

For jobs that pass work to the QZRCSRVS jobs, the subsystem that is specified on the QWTPCPUT API determines where the QZRCSRVS jobs run. QWTPCPUT is called during the processing of the QYSMPUT API. This API retrieves the subsystem information from the QUSRSYS/QYSMSVRE *USRIDX object and uses it on the QWTPCPUT call. As shipped, QZRCSRVS jobs are prestart jobs that run in the QUSRWRK subsystem and this is where the connections are routed.

If you end the prestart jobs in QUSRWRK with the ENDPJ command, then the QZRCSRVS jobs start as batch-immediate jobs in the QSYSWRK subsystem whenever a connection is requested. No jobs start in advance of the connection.

You can configure your system so that prestart jobs can be run from any subsystem. You can also configure your system to prevent batch-immediate jobs from being used at all. If the Job Monitor server jobs are calling Java Toolbox functions to pass work to QZRCSRVS, then they are using the QYSMPUT API, and the work should run in whichever subsystem is stored in the user index.

QZRCSRVS cleanup

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

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

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

Create a new monitor

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

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

You can specify the jobs to monitor in these ways:

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

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

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

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

Select the metrics:

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

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

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

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

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

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

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

Job monitor metrics

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

Job Count
: Monitor for a specific number of jobs matching the job selection.

Job Status
: Monitor for jobs in any selected status, such as Completed, Disconnected, Ending, Held while running, or Initial thread held.
: **Remember:** Metrics for job status can affect performance. Limit the number of jobs that you are monitoring to 40.

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

Job numeric values

CPU Utilization
: The percentage of available processing unit time used by all jobs that are included by this monitor on this system.

Logical I/O Rate
: The number of logical I/O actions, per second, by each job that is being monitored on this system.

Disk I/O Rate
: The average number of I/O operations, per second, performed by each job that is being monitored on this system. The value in this column is the sum of the asynchronous and synchronous disk I/O operations.

Communications I/O Rate
: The number of communications I/O actions, per second, by each job that is being monitored on this system.

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.

Specify the threshold values:

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

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

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

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

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

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

Job Count	When you define a threshold, you can specify a command to run on the endpoint system when the threshold is triggered. For example, selecting > **25 jobs** will trigger the threshold whenever the monitor detects more than 25 jobs running during the number of collection intervals you specify for **Duration**.
	You can then specify a command to be run on the endpoint system when the monitor detects more than 25 jobs. Enter the command name and click **Prompt** for assistance in specifying the parameters for the command. For more detailed information and examples of specifying commands to be run when thresholds are triggered, see the performance scenarios topic.
	Enable reset is optional, and cannot be selected until a trigger is defined. You can also specify a command to be run on the endpoint system when the threshold is reset.

Job Log Message	You must select **Trigger when any of the following messages are sent to the job log** before you can specify the conditions to trigger a threshold. You can specify messages to monitor for based on any combination of Message ID, Type, and Minimum severity. Each row in the Job Log Message table shows a combination of criteria that must be met for a message to trigger a threshold. A threshold will be triggered if it meets the criteria in at least one row. Use the online help to specify the conditions to trigger a threshold.
	Be careful to monitor the smallest number of jobs that will give you the information you need. Monitoring a large number of jobs for job log messages may have a performance 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** for assistance in specifying the parameters for the command.
	Be sure to click the Collection Interval tab to specify how often you want the monitor to check for job log messages.
	A message trigger can only be manually reset. You can specify a command to be run on the endpoint system when the threshold is reset. When you reset the monitor, you always have the option to reset without running the specified command.
Job Status	On the **Metrics - General** tab, select the statuses that you want to monitor for. Click the **Metrics - Status Threshold** tab to specify the conditions to trigger a threshold. You must select **Trigger when job is in any selected status** before you can specify the conditions to trigger a threshold. The threshold is triggered whenever the monitor detects that the job is in any selected status for the number of collection intervals you specify for **Duration**.
	You can then specify a command to be run on the endpoint system when the threshold is triggered. Enter the command name and click **Prompt** for assistance in specifying the parameters for the command.
	Reset when job is not in selected statuses is optional, and cannot be selected until a trigger is defined. You can specify a command to be run on the endpoint system when the threshold is reset.
Job Numeric Values	When you define the threshold, you can specify a command to run on the endpoint system when the threshold is triggered. For example, selecting › **101 transactions per second** for the Transaction Rate metric will trigger the threshold whenever the monitor detects more than 101 transactions per second on any of the selected jobs during the number of collection intervals you specify for **Duration**.
	You can then specify a command to be run on the endpoint system when the monitor detects more than 101 transactions per second. Enter the command name and click **Prompt** for assistance in specifying the parameters for the command.
	Enable reset is optional, and cannot be selected until a trigger is defined. You can also specify a command to be run on the endpoint system when the threshold is reset.
Summary Numeric Values (total for all jobs)	When you define a threshold, you can specify a command to run on the endpoint system when the threshold is triggered. For example, selecting › **1001 transactions per second** for the Transaction Rate metric will trigger the threshold whenever the monitor detects more than 1001 transactions per second on all of the selected jobs during the number of collection intervals you specify for **Duration**.
	You can then specify a command to be run on the endpoint system when the monitor detects more than 1001 transactions per second. Enter the command name and click **Prompt** for assistance in specifying the parameters for the command.
	Enable reset is optional, and cannot be selected until a trigger is defined. You can also specify a command to be run on the endpoint system when the threshold is reset.

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

Specify the collection interval:

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

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

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

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

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

Specify threshold run commands:

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

Use of threshold commands

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

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

How do I set threshold commands?

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

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

as &TIME or &NUMCURRENT) to pass information to the command, such as the time and actual value of the metric.

Specify event logging and actions:

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

Some of the actions you can select are:

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

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

How to read the event log

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

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

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

          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.

Apply thresholds and actions for a monitor:

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

Note: Because system monitors run continuously, the following information does not apply.

If you select to apply thresholds and actions during specified times, you must select the starting time and the stopping time. If the central system is in a different time zone from the endpoint system, you should be aware that the thresholds and actions will be applied when the starting time is reached on the endpoint system that you are monitoring. You must also select at least one day that you want the thresholds and actions to apply. The thresholds and actions apply from the selected starting time on the selected day until the next occurrence of the stopping time on the endpoint system.

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

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

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

View monitor results:

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

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

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

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

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

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

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

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

Graph history concepts:

Contains a description of the available options for managing and displaying records of performance data.

Graph history displays data contained in the collection objects created by Collection Services. Therefore, the type and amount of data available is dependent on your Collection Services configuration.

The amount of data that is available to be graphed is determined by the settings that you selected from the Collection Services properties, specifically the collection retention period. Use iSeries Navigator to activate PM iSeries over multiple systems. When you activate PM iSeries, you can use the graph history function to see data that was collected days ago, weeks ago, or months ago. You go beyond the realtime monitor capabilities, and have access to summary or detailed data. Without PM iSeries enabled, the graph data field supports 1 to 7 days. With PM iSeries enabled, you define how long your management collection objects remain on the system:

- **Detailed data** (attribute type *PFR in QMPGDATA.LIB or QPFRDATA.LIB)

  The length of time that management collection objects remain in the file system before they are deleted. You can select a specific time period in hours or days, or you can select **Permanent**. If you select **Permanent**, the management collection objects will not be automatically deleted.

- **Graph data** (attribute type *PFRDTL in QMGTC2.LIB)

  The length of time that the details and properties data that is shown in the Graph History window remains in the system before it is deleted. If you do not start PM iSeries, you can specify one to seven days. If you do start PM iSeries, you can specify 1 to 30 days. The default is one hour.

- **Summary data** (attribute type *PFRHST in QMGTC2.LIB)

  The length of time that the data collection points of a graph can be displayed in the Graph History window or remain in the system before they are deleted. No details or properties data is available. You must start PM iSeries to enable the summary data fields. The default is one month. The summary data is summarized in one-hour intervals and does not support second- and third-level details.

- **Graph history status**

  The Graph History window now displays the graph history status. You also can re-create the graph history data if it is missing.

Use graph history:

This topic contains step-by-step instructions to view graph history through iSeries Navigator.

Graph history is included in iSeries Navigator. To view the graph history of the data that you are monitoring with Collection Services, do these steps:

1. Follow the iSeries Navigator online help for starting Collection Services on either a single system or on a system group.
2. From the **Start Collection Services - General** page, select **Start IBM Performance Management for eServer iSeries** if needed.
3. Make changes to the other values for the collection retention period.
4. Click **OK**.

5. You can view the graph history by right-clicking either a system monitor or a Collection Services object and selecting **Graph History**.
6. Click **Refresh** to see the graphical view.

Tip: If the graph history data is missing, you can re-create it. To re-create the graph history data, right-click on the object in iSeries Navigator and choose **Create Graph History Data**.

Once you have launched a graph history, a window displays a series of graphed collection points. These collection points on the graph line are identified by three different graphics that correspond to the three levels of data that are available:

- A square collection point represents data that includes both the detailed information and properties information.
- A triangular collection point represents summarized data that contains detailed information.
- A circular collection point represents data that contains no detailed information or properties information.

The system adds data from the active collection object (*PFR attribute) to the *PFRDTL and *PFRHST collection objects when the following occurs:

- If the collection object properties is set to add graph data and summary data when cycled, the collection is cycled.
- If the already cycled object is selected and the menu option to summarize the data is selected.
- If a system monitor is running, then data is added to the *PFRDTL object only, as the system monitor is running.

Reset triggered threshold for a monitor:

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

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

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

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

Scenarios: iSeries Navigator monitors:

Provides scenarios that show how you can use some of the different types of monitors to look at specific aspects of your system's performance.

The monitors included in iSeries Navigator provide a powerful set of tools for researching and managing system performance. For an overview of the types of monitors provided by iSeries Navigator, see iSeries Navigator monitors.

For detailed usage examples and sample configurations, see the following scenarios:

Scenario: System monitor:

See an example system monitor that alerts you if the CPU utilization gets too high and temporarily holds any lower priority jobs until more resources become available.

Situation

As a system administrator, you need to ensure that the iSeries system has enough resources to meet the current demands of your users and business requirements. For your system, CPU utilization is a particularly important concern. You would like the system to alert you if the CPU utilization gets too high and to temporarily hold any lower priority jobs until more resources become available.

To accomplish this, you can set up a system monitor that sends you a message if CPU utilization exceeds 80%. Moreover, it can also hold all the jobs in the QBATCH job queue until CPU utilization drops to 60%, at which point the jobs are released, and normal operations resume.

Configuration example

To set up a system monitor, you need to define what metrics you want to track and what you want the monitor to do when the metrics reach specified levels. To define a system monitor that accomplishes this goal, complete the following steps:

1. In iSeries Navigator, expand **Management Central → Monitors**, right-click **System Monitor**, and select **New Monitor...**
2. On the **General** page, enter a name and description for this monitor.
3. Click the **Metrics** tab, and enter the following values:
   a. Select the **CPU Utilization Basic (Average)**, from the list of Available Metrics, and click **Add**. CPU Utilization Basic (Average) is now listed under Metrics to monitor, and the bottom portion of the window displays the settings for this metric.
   b. For **Collection interval**, specify how often you would like to collect this data. This will override the Collection Services setting. For this example, specify **30 seconds**.
   c. To change the scale for the vertical axis of the monitor's graph for this metric, change the **Maximum graphing value**. To change the scale for the horizontal axis of the graph for this metric, change the value for **Display time**.
   d. Click the **Threshold 1** tab for the metrics settings, and enter the following values to send an inquiry message if the CPU Utilization is greater than or equal to 80%:
      1) Select **Enable threshold.**
      2) For the threshold trigger value, specify **>= 80** (greater than or equal to 80 percent busy).
      3) For **Duration**, specify **1** interval.
      4) For the **i5/OS command**, specify the following:
         ```
         SNDMSG MSG('Warning,CPU...') TOUSR(*SYSOPR) MSGTYPE(*INQ)
         ```
      5) For the threshold reset value, specify **< 60** (less than 60 percent busy). This will reset the monitor when CPU utilization falls below 60%.
   e. Click the **Threshold 2** tab, and enter the following values to hold all the jobs in the QBATCH job queue when CPU utilization stays above 80% for five collection intervals:

1) Select **Enable threshold**.
2) For the threshold trigger value, specify **>= 80** (greater than or equal to 80 percent busy).
3) For **Duration**, specify **5** intervals.
4) For the **i5/OS command**, specify the following:

   `HLDJOBQ JOBQ(QBATCH)`
5) For the threshold reset value, specify **< 60** (less than 60 percent busy). This will reset the monitor when CPU utilization falls below 60%.
6) For **Duration**, specify **5** intervals.
7) For the **i5/OS command**, specify the following:

   `RLSJOBQ JOBQ(QBATCH)`

   This command releases the QBATCH job queue when CPU utilization stays below 60% for 5 collection intervals.

4. Click the **Actions** tab, and select **Log event** in both the **Trigger** and **Reset** columns. This action creates an entry in the event log when the thresholds are triggered and reset.
5. Click the **Systems and groups** tab to specify the systems and groups you want to monitor.
6. Click **OK** to save the monitor.
7. From the list of system monitors, right-click the new monitor and select **Start**.

Results

The new monitor displays the CPU utilization, with new data points being added every 30 seconds, according to the specified collection interval. The monitor automatically carries out the specified threshold actions, even if your PC is turned off, whenever CPU utilization reaches 80%.

Note: This monitor tracks only CPU utilization. However, you can include any number of the available metrics in the same monitor, and each metric can have its own threshold values and actions. You can also have several system monitors that run at the same time.

Scenario: Job monitor for CPU utilization:

See an example job monitor that tracks the CPU utilization of a specified job and alerts the job's owner if CPU utilization gets too high.

Situation

You are currently running a new application on your iSeries server, and you are concerned that some of the new interactive jobs are consuming an unacceptable amount of resources. You would like the owners of the offending jobs to be notified if their jobs ever consume too much of the CPU capacity.

You can set up a job monitor to watch for the jobs from the new application and send a message if a job consumes more than 30% of the CPU capacity.

Configuration example

To set up a job monitor, you need to define which jobs to watch for, what job attributes to watch for, and what the monitor should do when the specified job attributes are detected. To set up a job monitor that accomplishes this goal, complete the following steps:

1. In iSeries Navigator, expand **Management Central → Monitors**, right-click **Job monitor**, and select **New Monitor...**
2. On the **General** page, enter the following values:
   a. Specify a name and description for this monitor.
   b. On the **Jobs to monitor** tab, enter the following values:

1) For the **Job name**, specify the name of the job you want to watch for (for example, MKWIDGET).

2) Click **Add**.

3. Click the **Metrics** tab, and enter the following information:

   a. In the **Available metrics** list, expand **Summary Numeric Values**, select **CPU Percent Utilization**, and click **Add**.

   b. On the **Threshold 1** tab for the metrics settings, enter the following values:

      1) Select **Enable trigger**.

      2) For the threshold trigger value, specify **>= 30** (greater than or equal to 30 percent busy).

      3) For **Duration**, specify **1** interval.

      4) For the **i5/OS trigger command**, specify the following:

         ```
         SNDMSG MSG('Your job is exceeding 30% CPU capacity')
         TOUSR(&OWNER)
         ```

      5) Click **Enable reset**.

      6) For the threshold reset value, specify **< 20** (less than 20 percent busy).

4. Click the **Collection Interval** tab, and select **15 seconds**. This will override the Collection Services setting.

5. Click the **Actions** tab, and select **Log event** in both the **Trigger** and **Reset** columns.

6. Click the **Servers and groups** tab, and select the servers and groups you want to monitor for this job.

7. Click **OK** to save the new monitor.

8. From the list of job monitors, right-click the new monitor and select **Start**.

Results

The new monitor checks the QINTER subsystem every 15 seconds, and if the job MKWIDGET is consuming more than 30 percent of the CPU, the monitor sends a message to the job's owner. The monitor resets when the job uses less than 20% CPU capacity.

Scenario: Job monitor with Advanced Job Scheduler notification:

See an example job monitor that sends an e-mail to an operator when the threshold limit of a job is exceeded.

Situation

You are currently running an application on your iSeries server, and you want to be notified if the CPU utilization reaches the specified threshold.

If the Advanced Job Scheduler is installed on the endpoint system, you can use the Send Distribution using JS (SNDDSTJS) command to notify someone by e-mail when the threshold is exceeded. For instance, you could specify that the notification escalate to the next person if the intended recipient does not respond by stopping the message. You could create on-call schedules and send the notification to only those people that are on-call. You can also send the notification to multiple e-mail addresses.

Job monitor configuration example

This example uses the SNDDSTJS command to send a message to a recipient named OPERATOR, which is a user-defined list of e-mail addresses. You can also specify an e-mail address instead of a recipient or both. To set up a job monitor that accomplishes this goal, complete the following steps:

Note: By using the code examples, you agree to the terms of the "Code license and disclaimer information" on page 19.

1. In iSeries Navigator, expand **Management Central → Monitors**, right-click **Job monitor**, and select **New Monitor...**.
2. On the **General** page, enter the following values:
   a. Specify a name and description for this monitor.
   b. On the **Jobs to monitor** tab, enter the following values:
      1) For the **Job name**, specify the name of the job you want to watch for (for example, MKWIDGET).
      2) Click **Add**.
3. Click the **Metrics** tab, and enter the following information:
   a. In the **Available metrics** list, expand **Summary Numeric Values**, select **CPU Percent Utilization**, and click **Add**.
   b. On the **Threshold 1** tab for the metrics settings, enter the following values:
      1) Select **Enable trigger**.
      2) For the threshold trigger value, specify **>= 30** (greater than or equal to 30 percent busy).
      3) For **Duration**, specify **1** interval.
      4) For the **i5/OS trigger command**, specify the following:
         ```
         SNDDSTJS RCP(OPERATOR) SUBJECT('Job monitor trigger') MSG('Job
         &JOBNAME is still running!')
         ```
      5) Click **Enable reset**.
      6) For the threshold reset value, specify **< 20** (less than 20 percent busy).
4. Click the **Collection Interval** tab, and select **15 seconds**. This will override the Collection Services setting.
5. Click the **Actions** tab, and select **Log event** in both the **Trigger** and **Reset** columns.
6. Click the **Servers and groups** tab, and select the servers and groups you want to monitor for this job.
7. Click **OK** to save the new monitor.
8. From the list of job monitors, right-click the new monitor and select **Start**.

Message monitor configuration example

If you use a message monitor, you can send the message text to the recipient. Here is an example of a CL program that retrieves the message text and sends an e-mail to all on-call recipients with the SNDDSTJS command.

Note: By using the code examples, you agree to the terms of the "Code license and disclaimer information" on page 19.

```
PGM PARM(&MSGKEY &TOMSGQ &TOLIB)

DCL &MSGKEY  *CHAR 4
DCL &TOMSGQ  *CHAR 10
DCL &TOLIB   *CHAR 10

DCL &MSGTXT  *CHAR 132

RCVMSG   MSGQ(&TOLIB/&TOMSGQ) MSGKEY(&MSGKEY)
         RMV(*NO)  MSG(&MSGTXT)
               MONMSG CPF0000 EXEC(RETURN)

SNDDSTJS RCP(*ONCALL) SUBJECT('Message queue trigger')
MSG(&MSGTXT)
     MONMSG MSGID(CPF0000 IJS0000)

ENDPGM
```

This is the command that would call the CL program:

```
CALL SNDMAIL PARM('&MSGKEY' '&TOMSG' '&TOLIB')
```

Results

The monitor checks the QINTER subsystem every 15 seconds, and if the job MKWIDGET is consuming more than 30 percent of the CPU, the monitor sends an e-mail to the operator. The monitor resets when the job uses less than 20% CPU capacity.

See Work with notification for more information on the Advanced Job Scheduler notification function.

Scenario: Message monitor:

See an example message monitor that displays any inquiry messages for your message queue that occur on any of your iSeries servers. The monitor opens and displays the message as soon as it is detected.

Situation

You company has several iSeries servers running, and it is time-consuming to check your message queue for each system. As a system administrator, you need to be aware of inquiry messages as they occur across your system.

You can set up a message monitor to display any inquiry messages for your message queue that occur on any of your iSeries systems. The monitor opens and displays the message as soon as it is detected.

Configuration example

To set up a message monitor, you need to define the types of messages you would like to watch for and what you would like the monitor to do when these messages occur. To set up a message monitor that accomplishes this goal, complete the following steps:

1. In iSeries Navigator, expand **Management Central → Monitors**, right-click **Message monitor**, and select **New Monitor...**.
2. On the **General** page, enter a name and description for this monitor.
3. Click the **Messages** tab, and enter the following values:
   a. For **Message queue to monitor**, specify **QSYSOPR**.
   b. On the **Message set 1** tab, select **Inquiry** for **Type**, and click **Add**.
   c. Select **Trigger at the following message count**, and specify **1** message.
4. Click the **Collection Interval** tab, and select **15 seconds**.
5. Click the **Actions** tab, and select **Open monitor**.
6. Click the **Systems and groups** tab, and select the systems and groups you would like to monitor for inquiry messages.
7. Click **OK** to save the new monitor.
8. From the list of message monitors, right-click the new monitor and select **Start**.

Results

The new message monitor displays any inquiry messages sent to QSYSOPR on any of the iSeries servers that are monitored.

Note: This monitor responds to only inquiry messages sent to QSYSOPR. However, you can include two different sets of messages in a single monitor, and you can have several message monitors that run at the same time. Message monitors can also carry out i5/OS commands when specified messages are received.

Code license and disclaimer information

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

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

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

1. LOSS OF, OR DAMAGE TO, DATA;

2. DIRECT, SPECIAL, INCIDENTAL, OR INDIRECT DAMAGES, OR FOR ANY ECONOMIC CONSEQUENTIAL DAMAGES; OR

3. LOST PROFITS, BUSINESS, REVENUE, GOODWILL, OR ANTICIPATED SAVINGS.

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

Appendix. Notices

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

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

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

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

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

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

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

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

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

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

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

IBM Corporation

21

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

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

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

| AIX
| AIX 5L
| e(logo)server
| eServer
| i5/OS
| IBM
| iSeries
| pSeries
| xSeries
| zSeries

| Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States,
| other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the
United States, other countries, or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other
countries, or both.

| Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

Terms and conditions

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

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

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

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

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

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

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

IBM ®

Printed in USA