



System i  
Systems management  
Performance

*Version 6 Release 1*







System i  
Systems management  
Performance

*Version 6 Release 1*

**Note**

Before using this information and the product it supports, read the information in “Notices,” on page 195.

This edition applies to version 6, release 1, modification 0 of IBM i5/OS (product number 5761-SS1) and to all subsequent releases and modifications until otherwise indicated in new editions. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

© Copyright International Business Machines Corporation 1998, 2008. All rights reserved.

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

---

# Contents

<b>Performance</b>	<b>1</b>	
What's new in V6R1	1	
PDF file for Performance	2	
Managing system performance	3	
Selecting a performance management strategy	3	
Determining when and how to expand your system	5	
Comparing performance metrics before and after system changes	6	
Tracking performance	6	
Researching a performance problem	7	
Displaying performance data	14	
Tuning performance	15	
e-business performance	18	
Applications for performance management	22	
Applications for performance management overview	22	
Collection Services	24	
IBM i5/OS Job Watcher	56	
IBM i5/OS Disk Watcher	56	
Performance explorer	57	
IBM Systems Director Navigator for i5/OS Performance interface	66	
System i Navigator monitors	73	
System i Navigator graph history	86	
IBM Performance Management for System i5	88	
IBM Systems Workload Estimator	104	
Performance Tools	104	
Extended Adaptive Cache	165	
Intelligent Agents	167	
Scenarios: Performance	185	
Scenario: Improving system performance after an upgrade or migration	185	
Scenario: System monitor	186	
Scenario: Message monitor	187	
Scenario: Job monitor for CPU utilization	188	
Scenario: Job monitor with Advanced Job Scheduler notification	189	
Related information for Performance	191	
<b>Appendix. Notices</b>	<b>195</b>	
Programming Interface Information	196	
Trademarks	197	
Terms and conditions	197	



---

## Performance

- | Monitoring and managing your system's performance is critical to ensure you are keeping pace with the changing demands of your business.

To respond to business changes effectively, your system must change too. Managing your system, at first glance, might seem like just another time-consuming job. But the investment pays off soon because the system runs more efficiently, and this is reflected in your business. It is efficient because changes are planned and managed.

- | Managing performance of any system can be a complex task that requires a thorough understanding of that system's hardware and software. i5/OS® is an industry leader in the area of performance management and has many qualities that are not found in other systems, including unparalleled performance metrics, always-on collection services, and graphical viewing of performance data. While understanding all the different processes that affect system performance can be challenging and resolving performance problems requires the effective use of a large suite of tools, the functions offered by i5/OS are intended to make this job easier for users.

This topic will guide you through the tasks and tools associated with performance management.

**Note:** By using the following code examples, you agree to the terms of the "Code license and disclaimer information" on page 193.

### Related concepts

Work management

Work management supports the commands and internal functions necessary to control system operation and the daily workload on the system.

---

## What's new in V6R1

This topic describes what information is new or significantly changed in this release.

### IBM Systems Director Navigator for i5/OS Performance interface

New topics have been added to describe the IBM® Systems Director Navigator for i5/OS Performance interface. This interface allows you to display and manage Collection Services, IBM i5/OS Job Watcher, and IBM i5/OS Disk Watcher performance data in one place. See the "IBM Systems Director Navigator for i5/OS Performance interface" on page 66 topic for information on how to work with performance data. See the "Investigate Data" on page 66 topic for information about viewing performance data.

### Collection Services

Collection Services is now self starting on IPL. Collection Services data can be analyzed through the IBM Systems Director Navigator for i5/OS Performance interface. The expiration of performance database collections is now supported.

Several new data files have been added and many existing files have changes and additional fields. See the Collection Services data files topics for information.

### Performance Management Agent

The portion of the IBM Performance Management for System i5™ offering that runs on i5/OS has been named Performance Management Agent (PM Agent).

PM Agent no longer starts Collection Services. The database collection cleanup function has been removed and added to Collection Services and is configurable.

The PM Agent CL commands have new names:

Old CL command	New CL command	Description
CFGPM400	CFGPMAGT	Configure PM Agent
PMLINMON	CFGPMLIN	Configure PM Agent Line

## CL commands

Some functions, including CL commands that were previously part of the Performance Tools (5761-PT1) licensed program (such as WRKSYSACT and PRTPEXRPT) have been moved to i5/OS. Several commands are updated to support data from all collections (Collection Services, Performance Explorer, IBM i5/OS Job Watcher, and IBM i5/OS Disk Watcher). See the CL commands for performance topic for a listing of some common CL commands.

## IBM i5/OS Disk Watcher

The IBM i5/OS Disk Watcher data collection support is available with the base operating system. New in this release is the IBM i5/OS Disk Watcher interface, which provides the analysis and viewing of IBM i5/OS Disk Watcher data. It is now available in an optional feature of the Performance Tools licensed program (5761-PT1, Option 1.)

New topics have been added to describe IBM i5/OS Disk Watcher. IBM i5/OS Disk Watcher data can be analyzed through the IBM Systems Director Navigator for i5/OS Performance interface. See the “IBM i5/OS Disk Watcher” on page 56 topic for information.

## IBM i5/OS Job Watcher



The IBM i5/OS Job Watcher performance tool is now available. IBM i5/OS Job Watcher is a valuable tool for diagnosing i5/OS performance problems. The IBM i5/OS Job Watcher data collection support is now available with the base operating system. The IBM i5/OS Job Watcher Web based interface, which provides the analysis and viewing of IBM i5/OS Job Watcher data is now available in a new optional feature of the Performance Tools licensed program (5761-PT1, Option 3).

If you have used the iDoctor for System i5 tools in the past and have questions on them please see the iDoctor for System i5 Web site.

New topics have been added to describe IBM i5/OS Job Watcher. IBM i5/OS Job Watcher data can be analyzed through the IBM Systems Director Navigator for i5/OS Performance interface. See the “IBM i5/OS Job Watcher” on page 56 topic for information.

## How to see what's new or changed

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

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

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

---

## PDF file for Performance

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




- To view or download the PDF version of the performance topic, select Performance (about 2,428 KB). This PDF does not include the reference information for performance.
- To view or download the PDF version of the reference information for performance, select Reference information for Performance (about 3,557 KB).

## Saving PDF files

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

1. Right-click the PDF link in your browser.
2. Click the option that saves the PDF locally.
3. Navigate to the directory in which you would like to save the PDF.
4. Click **Save**.

## Downloading Adobe Reader

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

### Related reference

“Related information for Performance” on page 191

Listed here are the product manuals and IBM Redbooks™ (in PDF format), Web sites, and information center topics that relate to the Performance topic. You can view or print any of the PDFs.

---

## Managing system performance

Successfully managing performance ensures that your system is efficiently using resources and that your server provides the best possible services to your users and to your business needs. Moreover, effective performance management can help you quickly respond to changes in your system and can save you money by postponing costly upgrades and service fees.

Performance management is necessary to optimize utilization of your computer system by measuring current capabilities, recognizing trends, and making appropriate adjustments to satisfy end user and management requirements such as response time or job throughput. It is needed to maintain business efficiency and avoid prolonged suspension of normal business activities. Therefore, managing performance is part of your daily operations.

Understanding the factors that affect system performance helps you respond to problems and make better long-term plans. Effective planning can prevent potential performance problems from developing and ensures that you have the system capacity to handle your current and growing workloads.

### Related information



Three-In-One Benchmark

See the Three-In-One Benchmark Web site for information about the IBM Three-In-One Benchmark.

## Selecting a performance management strategy

Developing a good performance management strategy will help you manage your system's performance.

Your performance management strategy depends in a large part on the amount of time you can afford to spend managing performance. If you are working with a small company, you may be managing many different aspects of your business and cannot devote many hours to managing performance. Many large companies employ performance specialists to keep their systems tuned and running effectively.

- Different business needs require different performance management strategies. For determining a basic performance management strategy and for identifying which performance applications to use, classify

- | your company in one of three categories: small business, mid-sized business, and large business. The
- | business resources vary for each size, and your management strategy will vary accordingly.

## **Small business**

- | A small business most likely has fewer resources to devote to managing performance than a larger
- | business. For that reason, use as much automation as possible. You can use Performance Management for
- | System i5 (PM for System i5 ) to have your performance data sent directly to IBM where it will be
- | compiled and generated into reports for you. This not only saves you time, but IBM also makes
- | suggestions to you when your server needs an upgrade.

The following is a list of recommended performance applications for a small business:

- | • IBM Systems Director Navigator Performance interface: Display and manage performance data.
- Collection Services: Collect sample data at user-defined intervals for later analysis.
- Performance Management for System i5: Automate the collection, archival, and analysis of system performance data.
- Performance Tools: Gather, analyze, and maintain system performance information.
- System i<sup>™</sup> Navigator Monitors: Observe graphical representations of system performance, and automate responses to predefined events or conditions.

## **Mid-sized business**

The mid-sized business probably has more resources devoted to managing performance than the small business. You may still want to automate as much as possible and can also benefit from using PM for System i5.

The following is a list of recommended performance applications for a mid-sized business:

- | • IBM Systems Director Navigator Performance interface: Display and manage performance data.
- Collection Services: Collect sample data at user-defined intervals for later analysis.
- Performance Management for System i5: Automate the collection, archival, and analysis of system performance data.
- Performance Tools: Gather, analyze, and maintain system performance information.
- System i Navigator Monitors: Observe graphical representations of system performance, and automate responses to predefined events or conditions.

## **Large business**

The large business has resources devoted to managing performance.

The following is a list of recommended performance applications for a large business:

- | • IBM Systems Director Navigator Performance interface: Display and manage performance data.
- Collection Services: Collect sample data at user-defined intervals for later analysis.
- Performance Management for System i5: Automate the collection, archival, and analysis of system performance data.
- Performance Tools: Gather, analyze, and maintain system performance information.
- | • IBM i5/OS Job Watcher: Collect detailed information about a specific job or thread resource.
- | • IBM i5/OS Disk Watcher: Collect detailed information about disk performance data.
- Performance explorer: Collect detailed information about a specific application or system resource.
- Performance Trace Data Visualizer (PTDV): View trace data from a Java<sup>™</sup> application.

### **Related concepts**

“IBM Systems Director Navigator for i5/OS Performance interface” on page 66

The IBM Systems Director Navigator for i5/OS Performance interface allows you to display and manage performance data by bringing together a variety of performance information into one place.

“Collection Services” on page 24

Collection Services provides for the collection of system management data. It is the primary collector of system data.

“IBM i5/OS Job Watcher” on page 56

IBM i5/OS Job Watcher provides for the collection of job data for any or all jobs, threads, and tasks on the system. It provides call stacks, SQL statements, objects being waited on, Java JVM statistics, wait statistics and more which are used to diagnose job related performance problems.

“IBM i5/OS Disk Watcher” on page 56

IBM i5/OS Disk Watcher provides for the collection of disk performance data to diagnose disk related performance problems.

“System i Navigator graph history” on page 86

Graph history provides a graphical display of performance data collected by Collection Services over a specified period of time.

“IBM Performance Management for System i5” on page 88

The IBM Performance Management for System i5 (PM for System i5) offering automates the collection, archival, and analysis of system performance data and returns reports to help you manage system resources and capacity.

“Performance explorer” on page 57

Performance explorer collects more detailed information about a specific application, program or system resource, and provides detailed insight into a specific performance problem. This includes the capability both to perform several types and levels of traces and to run detailed reports.

“Performance Trace Data Visualizer” on page 66

Performance Trace Data Visualizer for iSeries™ is a tool for processing, analyzing, and viewing Performance Explorer collection data residing in PEX database files.

#### **Related reference**

“Performance Tools” on page 104

The Performance Tools licensed program includes many supplemental features that supplement or extend the capabilities of the basic performance tools that are available in the operating system.

“System i Navigator monitors” on page 73

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.

## **Determining when and how to expand your system**

As your business needs change, your system must also change. To prepare for any changes, you will want to model the current system and then see what would happen if the system, the configuration, or the workload were changed.

As your business needs evolve, so do your system needs. To plan for future system needs and growth, you will need to determine what would happen if the system, the configuration, or the workload were changed. This process is called trend analysis and should be done monthly. As your system approaches resource capacity guidelines, you may want to gather this data more frequently.

Trend analysis should be done separately for interactive and batch environments. If your company uses a certain application extensively, you may want to perform a trend analysis for the application. Another environment that may be important to track would be the end-of-month processing. It is important that you collect trend analysis data consistently. If your system’s peak workload hours are between 10:00 AM and 2:00 PM and you collect trend analysis data for this time period, do not compare this data to data collected from other time periods.

To do a proper job of capacity planning and performance analysis, you must collect, analyze, maintain, and archive performance data. IBM offers several tools that help you with your capacity planning, resource estimating, and sizing:

- IBM Performance Management for System i5
- IBM Systems Workload Estimator

#### **Related concepts**

“IBM Performance Management for System i5” on page 88

The IBM Performance Management for System i5 (PM for System i5) offering automates the collection, archival, and analysis of system performance data and returns reports to help you manage system resources and capacity.

“IBM Systems Workload Estimator” on page 104

The IBM Systems Workload Estimator is a Web-based sizing tool for System i, System p™, and System x™. You can use this tool to size a new system, to size an upgrade to an existing system, or to size a consolidation of several systems.

#### **Related reference**

“Selecting a performance management strategy” on page 3

Developing a good performance management strategy will help you manage your system’s performance.

## **Comparing performance metrics before and after system changes**

Comparing performance metrics before and after system changes provide important information for both troubleshooting and planning.

You should establish a set of system performance metrics before any major change in the system configuration, for example adding a new application or performing a system upgrade. Maintaining accurate system performance metrics can provide essential troubleshooting information. At a minimum, system performance metrics should include current collection objects from Collection Services.

#### **Related concepts**

“Collection Services” on page 24

Collection Services provides for the collection of system management data. It is the primary collector of system data.

## **Tracking performance**

Tracking your system performance over time allows you to plan for your system’s growth and ensures that you have data to help isolate and identify the cause of performance problems. Learn which applications to use and how to routinely collect performance data.

Tracking system performance helps you identify trends that can help you tune your system configuration and make the best choices about when and how to upgrade your system. Moreover, when problems occur, it is essential to have performance data from before and after the incident to narrow down the cause of the performance problem, and to find an appropriate resolution.

The system includes several applications for tracking performance trends and maintaining a historical record of performance data. Most of these applications use the data collected by Collection Services. You can use Collection Services to watch for trends in the following areas:

- Trends in system resource utilization. You can use this information to plan and specifically tailor system configuration changes and upgrades.
- Identification of stress on physical components of the configuration.
- Balance between the use of system resource by interactive jobs and batch jobs during peak and normal usage.
- Configuration changes. You can use Collection Services data to accurately predict the effect of changes like adding user groups, increased interactive jobs, and other changes.

- Identification of jobs that might be causing problems with other activity on the system
- Utilization level and trends for available communication lines.

The following tools will help you monitor your system performance over time:

- | • IBM Systems Director Navigator Performance interface
- | • Collection Services
- | • IBM Performance Management for System i5

#### **Related concepts**

“IBM Systems Director Navigator for i5/OS Performance interface” on page 66

The IBM Systems Director Navigator for i5/OS Performance interface allows you to display and manage performance data by bringing together a variety of performance information into one place.

“Collection Services” on page 24

Collection Services provides for the collection of system management data. It is the primary collector of system data.

“System i Navigator graph history” on page 86

Graph history provides a graphical display of performance data collected by Collection Services over a specified period of time.

“IBM Performance Management for System i5” on page 88

The IBM Performance Management for System i5 (PM for System i5) offering automates the collection, archival, and analysis of system performance data and returns reports to help you manage system resources and capacity.

#### **Related reference**

“Network performance” on page 19

The network design, hardware resources, and traffic pressure often have a significant effect on the performance of e-business applications. You can use this topic for information on how to optimize network performance, and tune server communication resources.

## **Researching a performance problem**

There are many options available to help you identify and resolve performance problems. Learn how to use the available tools and reports that can help you find the source of the performance problem.

- | Most of the tools that collect or analyze performance use either trace or sample data. Collection Services regularly collects sample data on a variety of system resources. Several tools analyze or report on this sample data, and you can use this to get a broader view of system resource utilization and to answer many common performance questions. IBM i5/OS Job Watcher and IBM i5/OS Disk Watcher also collect sample data. For more detailed performance information, several tools generate trace-level data. Often, trace-level data can provide detailed information about the behavior and resource consumption of jobs and applications on your system. Performance Explorer and the Start Performance Trace (STRPFTRC) command are two common tools for generating trace data.

- | For example, if your system is running slowly, you might use the System i Navigator monitors to look for problems. If you see that the CPU utilization is high, you could identify any jobs that seem to be using an unusually large amount of resources. Then, you may be able to correct the problem by making configuration changes. However, some problems will require additional information. To get detailed information about that job's performance you could start an IBM i5/OS Job Watcher collection for the desired job, gather detailed information about that job's behavior on the server, and potentially make changes to the originating program.

## **Identifying a performance problem**

Learn the common steps involved with identifying a performance problem.

When you try to identify a performance problem, it is important to assess whether the hardware configuration is adequate to support the workload. Is there enough CPU capacity? Is the main storage

sufficient for the different types of applications? Answering these questions first, perhaps through capacity modeling techniques, prevents needless effort later.

With an understanding of the symptoms of the problem and the objectives to be met, the analyst can formulate a hypothesis that may explain the cause of the problem. The analyst can use commands and tools available with i5/OS and the Performance Tools licensed program to collect and review data related to the system performance.

Reviewing the data helps you to further define the problem and helps you to validate or reject the hypothesis. Once the apparent cause or causes have been isolated, a solution can be proposed. When you handle one solution at a time, you can redesign and test programs. Again, the analyst's tools can, in many cases, measure the effectiveness of the solution and look for possible side effects.

To achieve optimum performance, you must recognize the interrelationship among the critical system resources and attempt to balance these resources, namely CPU, disk, main storage, and for communications, remote lines. Each of these resources can cause a performance degradation.

Improvements to system performance, whether to interactive throughput, interactive response time, batch throughput, or some combination of these, may take many forms, from simply adjusting activity level or pool size to changing the application code itself. In this instance, an activity level is a characteristic of a subsystem that specifies the maximum number of jobs that can compete at the same time for the processing unit.

## Identifying and resolving common performance problems

Many different performance problems often affect common areas of the system. Learn how to research and resolve problems in common areas, for example, backup and recovery.

When performance problems occur on the system, they often affect certain areas of the system first. Refer to the following table for some methods available for researching performance on these system areas.

Area	Description	Available tools
Processor load	Determine if there are too many jobs on the system or if some jobs are using a large percentage of processor time.	<ul style="list-style-type: none"><li>• Work with Active Jobs (WRKACTJOB) command.</li><li>• Work with System Activity (WRKSYSACT) command.</li><li>• The work management function in System i Navigator.</li><li>• CPU utilization metrics within the System i Navigator monitors.</li><li>• Work with Disk Status (WRKDSKSTS) command.</li></ul>
Main storage	Investigate faulting and the wait-to-ineligible transitions.	<ul style="list-style-type: none"><li>• The IBM Systems Director Navigator Performance interface.</li><li>• Disk storage metrics within the System i Navigator monitors.</li><li>• Work with System Status (WRKSYSSTS) command.</li><li>• The Memory Pools function under Work Management in System i Navigator.</li></ul>



Area	Description	Available tools
Disk	Determine if there are too few arms or if the arms are too slow.	<ul style="list-style-type: none"> <li>• The IBM Systems Director Navigator Performance interface.</li> <li>• Work with Disk Status (WRKDSKSTS) command.</li> <li>• Disk arm utilization metrics within the System i Navigator monitors.</li> <li>• Performance Tools System and Component report.</li> </ul>
Communications	Find slow lines, errors on the line, or too many users for the line.	<ul style="list-style-type: none"> <li>• Performance Tools Component Report.</li> <li>• LAN utilization metrics within the System i Navigator system monitor.</li> </ul>
IOPs	Determine if any IOPs are not balanced or if there are not enough IOPs.	<ul style="list-style-type: none"> <li>• The IBM Systems Director Navigator Performance interface.</li> <li>• Performance Tools Component Report.</li> <li>• IOP utilization metrics within the System i Navigator system monitor.</li> </ul>
Software	Investigate locks and mutual exclusions (mutexes).	<ul style="list-style-type: none"> <li>• The IBM Systems Director Navigator Performance interface.</li> <li>• Performance Tools Locks report.</li> <li>• Performance Tools Trace report.</li> <li>• Work with Object Locks (WRKOBJLCK) command.</li> <li>• View details of suspected jobs under Work Management in System i Navigator.</li> <li>• Work with System Activity (WRKSYSACT) command.</li> <li>• Display Performance Data (DSPPRDTA) command.</li> </ul>
Backup and recovery	Investigate areas that affect backup and recovery and save and restore operations.	<ul style="list-style-type: none"> <li>• System i Performance Capabilities Reference (Save/Restore Performance chapter).</li> </ul>

### Related concepts

“IBM Systems Director Navigator for i5/OS Performance interface” on page 66

The IBM Systems Director Navigator for i5/OS Performance interface allows you to display and manage performance data by bringing together a variety of performance information into one place.

Work management

See the Work management topic for more information about the Work management function in iSeries Navigator.

### Related reference

“Monitor metrics” on page 76

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.

Backup and recovery frequently asked questions

See the Backup and recovery frequently asked questions topic for answers to common backup and recovery questions.

### **Related information**

CL commands for performance



System i Performance Capabilities Reference (Save/Restore Performance chapter)

See the Save/Restore Performance chapter of the Performance Capabilities Reference for information about backup and recovery related performance.

## **Collecting system performance data**

Collecting data is an important step toward improving performance.

When you collect performance data, you gather information about your system that can be used to understand response times and throughput. It is a way to capture the performance status of the system, or set of systems, involved in getting your work done. The collection of data provides a context, or a starting point, for any comparisons and analysis that can be done later. When you use your first data collections, you have a benchmark for future improvements and a start on improving your performance today. You can use the performance data you collect to make adjustments, improve response times, and help your systems achieve peak performance. Performance problem analysis often begins with the simple question: "What changed?" Performance data helps you answer that question.

There are four collectors that have access to and can collect the data.

- Collection Services
- Job Watcher
- Disk Watcher
- Performance Explorer

### **Related concepts**

"Collection Services" on page 24

Collection Services provides for the collection of system management data. It is the primary collector of system data.

"IBM i5/OS Job Watcher" on page 56

IBM i5/OS Job Watcher provides for the collection of job data for any or all jobs, threads, and tasks on the system. It provides call stacks, SQL statements, objects being waited on, Java JVM statistics, wait statistics and more which are used to diagnose job related performance problems.

"IBM i5/OS Disk Watcher" on page 56

IBM i5/OS Disk Watcher provides for the collection of disk performance data to diagnose disk related performance problems.

"Performance explorer" on page 57

Performance explorer collects more detailed information about a specific application, program or system resource, and provides detailed insight into a specific performance problem. This includes the capability both to perform several types and levels of traces and to run detailed reports.

## **Collecting information about system resource utilization**

Several tools monitor how resources like central processing unit (CPU), disk space, interactive capacity, and many other elements, are being used. You can use these tools to start identifying problem areas.

Many tools are available to help you monitor and track the way the system and your applications are using the available resources. You can use this information as a starting point for problem analysis, and to identify trends that will help you with capacity planning and managing the growth of your system.

See the following topics to learn how and when to use these tools:

- IBM Systems Director Navigator Performance interface



- | • System i Navigator monitors
- | • CL commands for performance
- | • IBM Performance Management for System i5

#### **Related concepts**

“IBM Systems Director Navigator for i5/OS Performance interface” on page 66

The IBM Systems Director Navigator for i5/OS Performance interface allows you to display and manage performance data by bringing together a variety of performance information into one place.

“IBM Performance Management for System i5” on page 88

The IBM Performance Management for System i5 (PM for System i5) offering automates the collection, archival, and analysis of system performance data and returns reports to help you manage system resources and capacity.

#### **Related reference**

“System i Navigator monitors” on page 73

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.

#### **Related information**

CL commands for performance

## **Collecting information about an application’s performance**

An application might be performing slowly for a variety of reasons. You can use several of the tools included in i5/OS and other licensed programs to help you get more information.

Collecting information about an application’s performance is quite different from collecting information about system performance. Collecting application information can be done only with certain performance applications such as Performance Explorer, Performance Trace Data Visualizer, and Job Watcher. Alternately, you can get an overview of application performance by using the Job monitor to track individual server performance and Performance Tools to track and analyze server jobs.

**Note:** Collecting an application’s performance data can significantly affect the performance of your system. Before beginning the collection, make sure you have tried all other collection options.

- | The Start Performance Trace (STRPFRTTC) command collects multiprogramming and transaction data.
- | After running this command, you can export the data to a database file with the Dump Trace (DMPTRC) command.

#### **Related concepts**

“IBM i5/OS Job Watcher” on page 56

IBM i5/OS Job Watcher provides for the collection of job data for any or all jobs, threads, and tasks on the system. It provides call stacks, SQL statements, objects being waited on, Java JVM statistics, wait statistics and more which are used to diagnose job related performance problems.

“Performance explorer” on page 57

Performance explorer collects more detailed information about a specific application, program or system resource, and provides detailed insight into a specific performance problem. This includes the capability both to perform several types and levels of traces and to run detailed reports.

“Performance Trace Data Visualizer” on page 66

Performance Trace Data Visualizer for iSeries is a tool for processing, analyzing, and viewing Performance Explorer collection data residing in PEX database files.

#### **Related reference**

“System i Navigator monitors” on page 73

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.

“Performance Tools” on page 104

The Performance Tools licensed program includes many supplemental features that supplement or extend the capabilities of the basic performance tools that are available in the operating system.

Start Performance Trace (STRPFTRC) command

See the Start Performance Trace (STRPFTRC) command to collect Multiprogramming level (MPL) and Transaction trace data.

“Java performance in i5/OS” on page 19

i5/OS provides several configuration options and resources for optimizing the performance of Java applications or services on the system. Use this topic to learn about the Java environment and how to get the best possible performance from Java-based applications.

### Related information



Performance Trace Data Visualizer

See the Performance Trace Data Visualizer Web site for information about the Performance Trace Data Visualizer.

### Dumping trace data:

The Dump Trace (DMPTRC) command puts information from an internal trace table into a database file.

It is not a good practice 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 information. 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 CRTPFRTDA TOMBR and TOLIB parameters should be the same as the names that you provide in the DMPTRC MBR and LIB parameters.

### Related concepts

“Collection Services” on page 24

Collection Services provides for the collection of system management data. It is the primary collector of system data.

### Related reference

Dump Trace (DMPTRC) command

See the Dump Trace (DMPTRC) command to put information from an internal trace table into a database file.

### Dumping memory:

The Dump Main Memory Information (DMPMEMINF) command dumps information about pages of main memory to a file.

To dump memory data, issue the following command:

```
DMPMEMINF OUTFILE(MYLIBRARY/DMPMEMFILE)
```

The command to view the dump could be something like the following SQL:

```
SELECT count(*),POOL, OBJNAME, LIBNAME FROM mylibrary/dmpmemfile
group BY POOL, OBJNAME, LIBNAME
order by 1 desc
```

#### **Related reference**

Dump Main Memory Information (DMPMEMINF) command

See the Dump Main Memory Information (DMPMEMINF) command to dump information about pages of main memory to a file.

## **Scenario: Improving system performance after an upgrade or migration**

In this scenario, you have just upgraded or migrated your system and it now appears to be running slower than before. This scenario will help you identify and fix your performance problem.

### **Situation**

You recently upgraded your system to the newest release. After completing the upgrade and resuming normal operations, your system performance has decreased significantly. You would like to identify the cause of the performance problem and restore your system to normal performance levels.

### **Details**

Several problems may result in decreased performance after upgrading the operating system. You can use the performance management tools included in i5/OS and Performance Tools licensed program (5761-PT1) to get more information about the performance problem and to narrow down suspected problems to a likely cause.

1. Check CPU utilization. Occasionally, a job will be unable to access some of its required resources after an upgrade. This may result in a single job consuming an unacceptable amount of the CPU resources.
  - Use WRKSYSACT, WRKSYSSTS, WRKACTJOB, or System i Navigator system monitors to find the total CPU utilization.
  - If CPU utilization is high, for example, greater than 90%, check the amount of CPU utilized by active jobs. If a single job is consuming more than 30% of the CPU resources, it may be missing file calls or objects. You can then refer to the vendor, for vendor-supplied programs, or the job's owner or programmer for additional support.
2. Start a performance trace with the STRPFRTTC command, and then use the system and component reports to identify and correct the following possible problems:
  - If the page fault rate for the machine pool is higher than 10 faults/second, give the machine pool more memory until the fault rate falls below this level.
  - If the disk utilization is greater than 40%, look at the waiting and service time. If these values are acceptable, you may need to reduce workload to manage priorities.
  - If the IOP utilization is greater than 60%, add an additional IOP and assign some disk resource to it.
  - If the page faults in the user pool are unacceptably high, you might want to automatically tune performance.
3. Run the job summary report and refer to the Seize lock conflict report. If the number of seize or lock conflicts is high, ensure that the access path size is set to 1TB. If the seize or lock conflicts are on a user profile, and if the referenced user profile owns many objects, reduce the number of objects owned by that profile.

#### **Related concepts**

"Adjusting performance automatically" on page 17

Most users should set up the system to make performance adjustment automatically. When new systems are shipped, they are configured to adjust automatically.

#### **Related reference**

STRPFRTRC command

See the Start Performance Trace (STRPFRTRC) command to collect Multiprogramming level (MPL) and Transaction trace data.

#### **Related information**

Performance Tools reports

Performance Tools reports provide information on data that has been collected over time. Use these reports to get additional information about the performance and use of system resources.

## **Displaying performance data**

After you have collected performance data, learn how to display the data using the most appropriate tool for your purposes.

Displaying performance data helps you analyze your system's performance more accurately. Performance data can be displayed in many different ways; however, you may find a certain performance application more appropriate in some situations. Most applications display data collected with either Collection Services or from a performance trace. The best way to access that data depends on whether you are attempting to resolve a performance problem, are monitoring your system performance to plan for future growth, or are identifying trends.

## **Displaying near real-time performance data**

Use the following tools to display current or recent performance information:

- | • IBM Systems Director Navigator Performance interface
- | • CL commands for performance
- | • Performance Tools plug-in
- | • System i Navigator monitors

## **Displaying historical performance data**

Use the following tools to view data that is stored on your system:

- | • IBM Systems Director Navigator Performance interface
- | • IBM Performance Management for System i5
- | • Performance Tools plug-in
- | • System i Navigator graph history

#### **Related concepts**

"IBM Systems Director Navigator for i5/OS Performance interface" on page 66

The IBM Systems Director Navigator for i5/OS Performance interface allows you to display and manage performance data by bringing together a variety of performance information into one place.

"Collection Services" on page 24

Collection Services provides for the collection of system management data. It is the primary collector of system data.

"IBM Performance Management for System i5" on page 88

The IBM Performance Management for System i5 (PM for System i5) offering automates the collection, archival, and analysis of system performance data and returns reports to help you manage system resources and capacity.

"System i Navigator graph history" on page 86

Graph history provides a graphical display of performance data collected by Collection Services over a specified period of time.

#### **Related tasks**

“Performance Tools plug-in” on page 106

You can view system resource utilization data in System i Navigator. You can view the data, graph it, and summarize it into reports. Find information about how to access this function here.

#### **Related reference**

“System i Navigator monitors” on page 73

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.

#### **Related information**

CL commands for performance

## **Tuning performance**

When you have identified a performance problem, you will want to tune the system to fix it.

The primary aim of performance tuning is to make the most efficient use of the system resources. Performance tuning is a way to adjust the performance of the system either manually or automatically. Many options exist for tuning your system. Each system environment is unique in that it requires you to observe performance and make adjustments that are best for your environment; in other words, you are required to do routine performance monitoring.

IBM also offers a tool that allows you to improve both the I/O subsystem and system response times by reducing the number of physical I/O requests that are read from disk. You can improve your system performance with Extended Adaptive Cache.

In addition, you may also want to consider some tuning options that allow processes and threads to achieve improved affinity for memory and processor resources.

#### **Related concepts**

“Extended Adaptive Cache” on page 165

You can use Extended Adaptive Cache to improve system performance by collecting disk usage data, and then using those statistics to create a cache, effectively reducing the physical I/O requests for the disk.

#### **Related reference**

Thread affinity system value

See the thread affinity system value to specify if secondary threads will have affinity to the same group of processors and memory as the initial thread.

Processor multitasking system value

See the processor multitasking system value to specify if processor multi-tasking is on, off, or determined by the system.

## **Performing basic system tuning**

To tune your system’s performance, you need to set up your initial tuning values, observe the system performance, review the values, and determine what to tune.

To begin tuning performance, you must first set initial tuning values by determining your initial machine and user pool sizes. Then, you can begin to observe the system performance.

### **Set initial tuning values**

Setting initial tuning values includes the steps you take to initially configure the system pool sizes and activity levels to tune your system efficiently. The initial values are based on estimates; therefore, the estimates may require further tuning while the system is active. The following steps set the initial tuning values:

- Determine initial machine pool size  
Tune the machine pool to under 10 faults/second.

- Determine initial user pool sizes

Tune user pools so that the sum of faults for all user pools is less than the number of processors times the processors percent busy. For example, in a system with four processors running at 50 percent busy ( $4 * 50 = 200$ ), you would set the faults to less than 200 faults/seconds.

### Observe system performance

To observe the system performance, you can use the Work with System Status (WRKSYSSTS), Work with Disk Status (WRKDSKSTS), and Work with Active Jobs (WRKACTJOB) commands. With each observation period, you should examine and evaluate the measurements of system performance against your performance goals.

1. Remove any irregular system activity. Irregular activities that may cause severe performance degradation are, for example, interactive program compilations, communications error recovery procedures (ERP), open query file (OPNQRYF), application errors, and signoff activity.
2. Use the WRKSYSSTS, WRKDSKSTS, WRKACTJOB and WRKSYSACT CL commands to display performance data.
3. Allow the system to collect data for a minimum of 5 minutes.
4. Evaluate the measures of performance against your performance goals. Typical measurements include:
  - Interactive throughput and response time, available from the WRKACTJOB display.
  - Batch throughput. Observe the auxiliary input/output (AuxIO) and CPU percentage (CPU%) values for active batch jobs.
  - Spooled throughput. Observe the auxiliary input/output (AuxIO) and CPU percentage (CPU%) values for active writers.
5. If you observe performance data that does not meet your expectations, tune your system based on the new data. Be sure to:
  - Measure and compare all key performance measurements.
  - Make and evaluate adjustments one at a time.

### Review performance

Once you have set good tuning values, you should periodically review them to ensure your system continues to do well. Ongoing tuning consists of observing aspects of system performance and adjusting to recommended guidelines.

To gather meaningful statistics, you should observe system performance during typical levels of activity. For example, statistics gathered while no jobs are running on the system are of little value in assessing system performance. If performance is not satisfactory in spite of your best efforts, you should evaluate the capabilities of your configuration. To meet your objectives, consider the following:

- Processor upgrades
- Additional storage devices and controllers
- Additional main storage
- Application modification

By applying one or more of these approaches, you should achieve your objectives. If, after a reasonable effort, you are still unable to meet your objectives, you should determine whether your objectives are realistic for the type of work you are doing.

### Determine what to tune

If your system performance has degraded and needs tuning, you need to identify the source of the performance problem and make specific corrections.

#### Related reference



“Researching a performance problem” on page 7

There are many options available to help you identify and resolve performance problems. Learn how to use the available tools and reports that can help you find the source of the performance problem.

## Adjusting performance automatically

Most users should set up the system to make performance adjustment automatically. When new systems are shipped, they are configured to adjust automatically.

The system can set performance values automatically to provide efficient use of system resources. You can set up the system to tune system performance automatically by:

- Adjusting storage pool sizes and activity levels
- Adjusting storage pool paging

## Adjusting storage pool sizes and activity levels

Use the QPFRADJ system value to control automatic tuning of storage pools and activity levels. This value indicates whether the system should adjust values at system restart (IPL) or periodically after restart.

You can set up the system to adjust performance at IPL, dynamically, or both.

- To set up the system to tune only at system restart (IPL), select **Configuration and Service** → **System Values** → **Performance in System i Navigator**. Click the **Memory Pools** tab and select **At system restart** under **Automatically adjust memory pools and activity levels**. This is equivalent to setting the QPFRADJ system value to 1.
- To set up the system to make storage pool adjustments at system restart (IPL) and to make storage pool adjustments periodically after restart, select **Configuration and Service** → **System Values** → **Performance in System i Navigator**. Click the **Memory Pools** tab and select both **At system restart** and **Periodically after restart** under **Automatically adjust memory pools and activity levels**. This is equivalent to setting the QPFRADJ system value to 2.
- To set up the system to make storage pool adjustments periodically after restart and not at system restart (IPL), select **Configuration and Service** → **System Values** → **Performance in System i Navigator**. Click the **Memory Pools** tab and select **Periodically after restart** under **Automatically adjust memory pools and activity levels**. This is equivalent to setting the QPFRADJ system value to 3.

The storage pool values are not reset at system restart (IPL) to the initial values.

## Adjusting storage pool paging

The dynamic tuning support provided by the system automatically adjusts pool sizes and activity levels for shared pools to improve the performance of the system. This tuning works by moving storage from storage pools that have minimal use to pools that would benefit from more storage. This tuning also sets activity levels to balance the number of threads in the pool with the storage allocated for the pool. To adjust the system, the tuner uses a guideline that is calculated based on the number of threads.

When dynamic adjustment is in effect, the following performance values are changed automatically to the appropriate settings:

- Machine (\*MACHINE) memory pool size (QMCHPOOL system value)
- Base (\*BASE) memory pool activity level (QBASACTLVL system value)
- Pool size and activity level for the shared pool \*INTERACT
- Pool size and activity level for the shared pool \*SPOOL
- Pool sizes and activity levels for the shared pools \*SHRPOOL1-\*SHRPOOL60

When dynamic adjustment is in effect (the QPFRADJ system value is set to 2 or 3), the job QPFRADJ that runs under profile QSYS is seen as active on the system.

### Related information

Memory pools

See the Memory pools topic for information about memory pools.

## Determining when to use simultaneous multithreading

Simultaneous multithreading allows sharing of process facilities to run two applications or two threads of the same application at the same time.

Although an operating system gives the impression that it is concurrently executing a very large number of tasks, each processor in a symmetric multiprocessor (SMP) traditionally executes a single task's instruction stream at any moment in time. The QPRCMLTTSK system value controls whether to enable the individual SMP processors to concurrently execute multiple instruction streams. Each instruction stream belongs to separate tasks or threads. When enabled, each individual processor is concurrently executing multiple tasks at the same time. The effect of its use will likely increase the performance capacity of a system or improve the responsiveness of a multithreaded application. Running multiple instruction streams at the same time does not improve the performance of any given task. As is the case with any performance recommendations, results vary in different environments.

The way that multithreading is done depends on the hardware model, and therefore, the performance capacity gains vary. Some models support this approach through a concept called simultaneous multithreading (SMT). This approach, called hyperthreading on some Intel® processors, shares processor facilities to execute each task's instructions at the same time. Older processors use an approach called hardware multithreading (HMT). In the hardware multithreading approach, the hardware switches between the tasks on any long processing delay event, for example, a cache miss. Some models do not support any form of multithreading, which means the QPRCMLTTSK system value has no performance effect.

Because the QPRCMLTTSK system value enables the parallel use of shared processor resources, the performance gains depend highly on the application and the model. Refer to the *System i Performance Capabilities Reference* for guidelines about what performance gains might be expected through its use. In some cases, some applications are better served by disabling this system value.

### Related reference

QPRCMLTTSK system value

Specifies whether processor multi-tasking is on, off, or determined by the system.

### Related information



System i Performance Capabilities Reference PDF

See the System i Performance Capabilities Reference for guidelines about performance gains that might be expected through the use of the QPRCMLTTSK system value.

## e-business performance

Managing performance in an e-business environment introduces several new problems for the system administrator.

In addition to routine tuning on the server, administrators must also monitor and optimize the hardware and services supporting their e-business transactions.

### Related information



Domino for iSeries sizing and performance tuning

See the IBM Redbooks publication Domino for iSeries Sizing and Performance Tuning for Domino for iSeries performance information.

## Client performance

While the system administrator often has little control of the client-side of the e-business network, you can use these recommendations to ensure that client devices are optimized for an e-business environment.



Clients consisting of a PC with a Web browser often represent the e-business component that administrators have the least direct control over. However, these components still have a significant effect on the end-to-end response time for web applications.

To help ensure high-end performance, client PCs should:

- Have adequate memory. Interfaces that use complex forms and graphics and resource intensive applets may also place demands on the client's processor.
- Use a high-speed and optimized network connection. Many communication adapters on a client PC may function while they are not optimized for their network environment. For more information, refer to the documentation for your communication hardware.
- Use browsers that fully support the required technologies. Moreover, browser support and performance should be a major concern when designing the Web interface.

## Network performance

The network design, hardware resources, and traffic pressure often have a significant effect on the performance of e-business applications. You can use this topic for information on how to optimize network performance, and tune server communication resources.

The network often plays a major role in the response time for web applications. Moreover, the performance impact for network components is often complex and difficult to measure because network traffic and the available bandwidth may change frequently and are affected by influences the system administrator may not have direct control over. However, there are several resources available to help you monitor and tune the communication resources on your server.

Refer to the following topics for more information:

### Related concepts

"IBM Systems Director Navigator for i5/OS Performance interface" on page 66

The IBM Systems Director Navigator for i5/OS Performance interface allows you to display and manage performance data by bringing together a variety of performance information into one place.

"Tracking performance" on page 6

Tracking your system performance over time allows you to plan for your system's growth and ensures that you have data to help isolate and identify the cause of performance problems. Learn which applications to use and how to routinely collect performance data.

### Related reference

"System i Navigator monitors" on page 73

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.

### Related information



System i Performance Capabilities Reference PDF

The Performance Capabilities Reference provides detailed information, reports and examples that can help you configure or tune your server for optimal performance. In particular, see Chapter 5: Communications Performance to help you plan for and manage communication resources.

SystemiNetwork.com

This Web site hosts many resources for optimizing your network plan and resources. In particular, refer to the articles "Cultivate your AS/400 Networks" and "8 tools for better network performance."

## Java performance in i5/OS

i5/OS provides several configuration options and resources for optimizing the performance of Java applications or services on the system. Use this topic to learn about the Java environment and how to get the best possible performance from Java-based applications.

Java is often the language of choice for web-based applications. However, Java applications may require some optimization, both of the i5/OS environment and of the Java application, to get optimal performance.

Use the following resources to learn about the Java environment in i5/OS and the available tips and tools for analyzing and improving Java performance.

#### **Related concepts**

“IBM Systems Director Navigator for i5/OS Performance interface” on page 66

The IBM Systems Director Navigator for i5/OS Performance interface allows you to display and manage performance data by bringing together a variety of performance information into one place.

Tuning Java program performance

There are several important configuration choices and tools to help you get the best performance from Java-based applications.

#### **Related reference**

“Collecting information about an application’s performance” on page 11

An application might be performing slowly for a variety of reasons. You can use several of the tools included in i5/OS and other licensed programs to help you get more information.

#### **Related information**



System i Performance Capabilities Reference

The Performance Capabilities Reference provides detailed information, reports and examples that can help you configure or tune your server for optimal performance. In particular, see Chapter 7: Java Performance, to help you optimize the performance of Java applications, and learn performance tips for programming in Java.

Java and WebSphere Performance on IBM eServer iSeries Servers

Use this IBM Redbooks publication to learn how to plan for and configure your operating environment to maximize Java and WebSphere performance, and to help you collect and analyze performance data.

WebSphere J2EE Application Development for the IBM eServer iSeries Server

This IBM Redbooks publication provides an introduction to J2EE, and offers suggestions and examples to help you successfully implement J2EE applications on the server.

## **IBM HTTP Server performance**

The IBM HTTP Server is often an important part of e-business performance. IBM provides several options and configuration choices that allow you to get the most out of this server.

- | IBM HTTP Server for i5/OS can play an important role in the end-to-end performance of your Web-based applications, and several functions allow you to effectively monitor and improve Web server performance. In particular the Fast Response Caching Accelerator (FRCA) may allow you to significantly improve HTTP Server performance, particularly in predominantly static environments. The IBM HTTP Server for i5/OS also provides a Web Performance Monitor and Web Performance Advisor.

Refer to the following resources for information on how to maximize HTTP Server performance.

#### **Related concepts**

“IBM Systems Director Navigator for i5/OS Performance interface” on page 66

The IBM Systems Director Navigator for i5/OS Performance interface allows you to display and manage performance data by bringing together a variety of performance information into one place.

#### **Related information**

Collection Services data files: QAPMHTTPB

Collection Services data files: QAPMHTTPD

IBM HTTP Server for i5/OS

Refer to this topic for information on setting up, configuring, and managing an HTTP Server on

i5/OS. This topic also includes descriptions of the Web Performance Monitor, Web Performance Advisor, and Fast Response Caching Accelerator (FRCA).



#### System i Performance Capabilities Reference

The Performance Capabilities Reference provides detailed information, reports, and examples that can help you configure or tune your system for optimal performance. In particular, see Chapter 6: Web Server and Web Commerce, for HTTP server performance specifications, planning information, and performance tips.



**IBM HTTP Server (powered by Apache): An Integrated Solution for IBM eServer iSeries servers**  
Use this IBM Redbooks publication to get an in-depth description of HTTP Server (Powered by Apache) for i5/OS, including examples for configuring HTTP Server in common usage scenarios.



**AS/400 HTTP Server Performance and Capacity Planning**  
Use this IBM Redbooks publication to learn about HTTP server impacts on performance tuning and planning. This publication also includes suggestions for using performance management tools to collect, interpret, and respond to Web server performance data.

## WebSphere performance

WebSphere® Application Server is the e-business application deployment environment of choice. Use this topic to learn how to plan for and optimize performance in a WebSphere environment.

Managing system performance in a WebSphere environment presents several challenges to the administrator. Web-based transactions may consume more resources, and consume them differently than traditional communication workloads.

Refer to the following topics and resources to learn how to plan for optimal performance, and to adjust system resources in a WebSphere environment.

### Related information

Collection Services data files: QAPMWASAPP

Collection Services data files: QAPMWASCFG

Collection Services data files: QAPMWASEJB

Performance data files: QAPMWASRSC

Collection Services data files: QAPMWASSVR



#### WebSphere Application Server Performance Web site

This Web site provides resources for each version of WebSphere Application Server, including many useful performance tips and recommendations. This resource is particularly valuable for environments using servlets, Java Server Pages (JSPs) and Enterprise JavaBeans (EJBs).



#### DB2 UDB/WebSphere Performance Tuning Guide

This IBM Redbooks publication provides an introduction to both the WebSphere and DB2 environments, and offers suggestions, examples, and solutions to common performance problems that can help you optimize WebSphere and DB2 performance.



#### Java and WebSphere Performance on IBM eServer iSeries Servers

Use this IBM Redbooks publication to learn how to plan for and configure your operating environment to maximize Java and WebSphere performance, and to help you collect and analyze performance data.



#### WebSphere V3 Performance Tuning Guide

This IBM Redbooks publication offers detailed recommendations and examples for optimizing WebSphere V3 performance.



#### System i Performance Capabilities Reference

The Performance Capabilities Reference provides detailed information, reports and examples that can

help you configure or tune your server for optimal performance. In particular, see Chapter 6, "Web Server and Web Commerce", for performance tips specific to WebSphere Application Server.

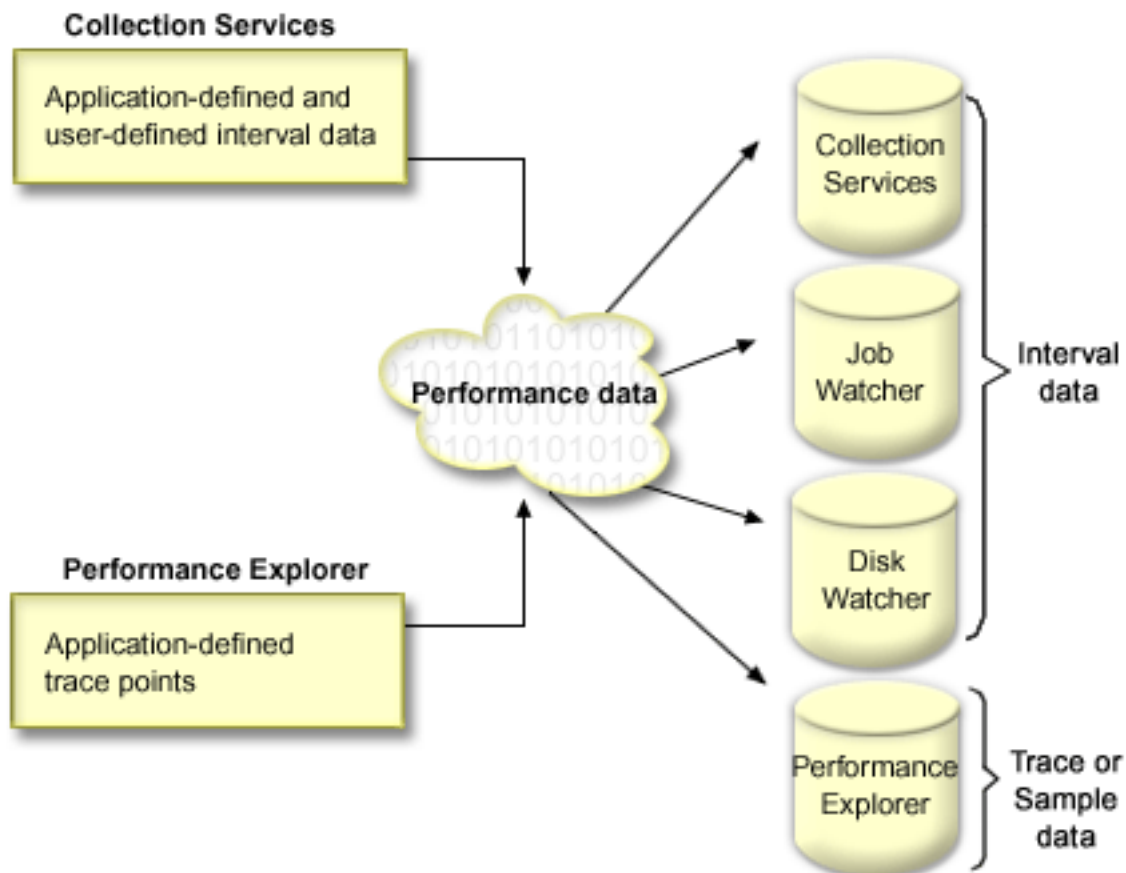
## Applications for performance management

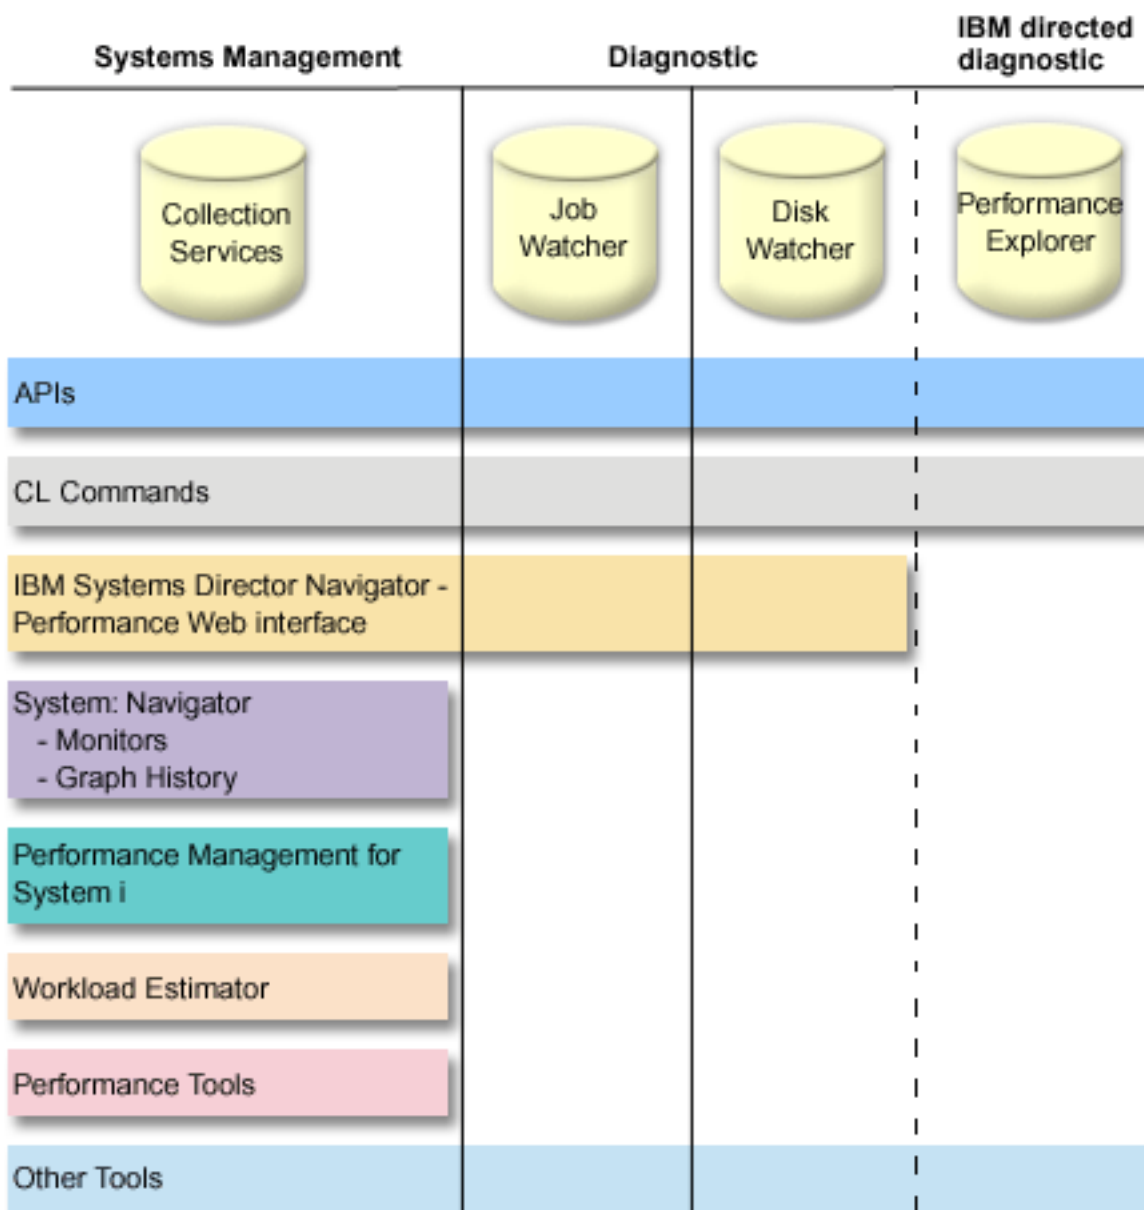
Many of the applications for performance management have several functions. The following topics provide detailed information about each of the performance management applications.

### Applications for performance management overview

Managing performance requires the use of a variety of specialized applications. Each of these applications offers a specific insight into system performance.

The following figures represent the main performance applications. The cloud shape represents all of the data that exists in the system that can be collected. There are four collectors that have access to and can collect the data. Ultimately the data that is collected by a collector is deposited into a set of database files.





Each of the collectors has unique characteristics.

#### Collection Services

Collection Services provides for the collection of system management data. It is the primary collector of system data. You can run this continuously to know what is happening with your system. Collection Services data is deposited into a management collection object and then converted and put into database files. The interval data that is collected is specified by either application defined or user defined interval data.

#### IBM i5/OS Job Watcher

IBM i5/OS Job Watcher provides for the collection of job data for any or all jobs, threads, tasks on the system. It provides call stacks, SQL statements, objects being waited on, Java JVM statistics, wait statistics and more which are used to diagnose job related performance problems.

### | **IBM i5/OS Disk Watcher**

| IBM i5/OS Disk Watcher provides for the collection of disk performance data to diagnose disk  
| related performance problems.

### | **Performance Explorer**

| Performance Explorer provides for the collection of detailed data at a program and application  
| level to diagnose problems. It also traces the flow of work in an application and can be used to  
| diagnose difficult performance problems. Application-defined performance explorer trace points,  
| such as with Domino<sup>®</sup>, NetServer<sup>™</sup>, or WebSphere servers specify the data that is collected. It is  
| intended to be used as directed by IBM. Performance Explorer data is deposited into a  
| management collection object and then converted and put into database files.

| The performance data contained in any of the database files can be accessed through APIs or CL  
| commands. The performance data contained in some of the database files can be investigated and  
| analyzed using one of more of a variety of tools that are described in the Applications for performance  
| topic.

## **Collection Services**

Collection Services provides for the collection of system management data. It is the primary collector of system data.

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
- Integrate your own programs to collect user-defined performance data into Collection Services.

### **Related concepts**

iSeries Navigator

See the System i Navigator topic for information about how to use System i Navigator to collect and manage performance data.

“Time zone considerations for Collection Services” on page 29

When you review and analyze performance data, the actual local time of the collection can be significant.

“Performance explorer concepts” on page 58

Performance explorer works by collecting detailed information about a specified system process or resource. This topic explains how performance explorer works, and how best to use it.

### **Related tasks**

“Activating PM Agent” on page 89

PM Agent is a part of the operating system and you must activate it to use its collecting capabilities.

### **Related reference**

Start Performance Collection (STRPFRCOL) command

See the Start Performance Collection (STRPFRCOL) command for information on how to start data collection.

Performance Management APIs

See the Performance Management APIs for information about how to use Performance Management APIs to collect and manage performance data.

### **Related information**

Performance data files

See the Performance data files topic for information about files that contain performance data.

## How Collection Services works

Collection Services stores data for each collection in a single collection object from which you can create as many different sets of database files as you need.

- | Storing the data in a single collection object results in lower system overhead when collecting
- | performance data. If you elect to create the database files during collection, Collection Services uses a
- | lower priority (50) batch job to update these files. This low collection overhead makes it practical to
- | collect detailed performance data at short 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.

The following figure provides an overview of the following Collection Services elements:

### User interfaces

- | Several methods exist that allow you to access the different elements of Collection Services. For
- | example, you can use CL commands, APIs, the System i Navigator interface, and the IBM
- | Systems Director Navigator Performance interface.

### General properties

General properties define how a collection should be accomplished, and they control automatic collection attributes.

### Data categories

Data categories identify the types of data to collect. You can configure categories independently to control what data is collected and how often the data is collected.

### Collection profiles

Collection profiles provide a means to save and activate a particular category configuration.

### Performance collector

The performance collector uses the general properties and category information to control the collection of performance data. You can start and stop the performance collector, or configure it to run automatically.

### Collection Object

The collection object, \*MGTCOL, serves as an efficient storage medium for large quantities of performance data.

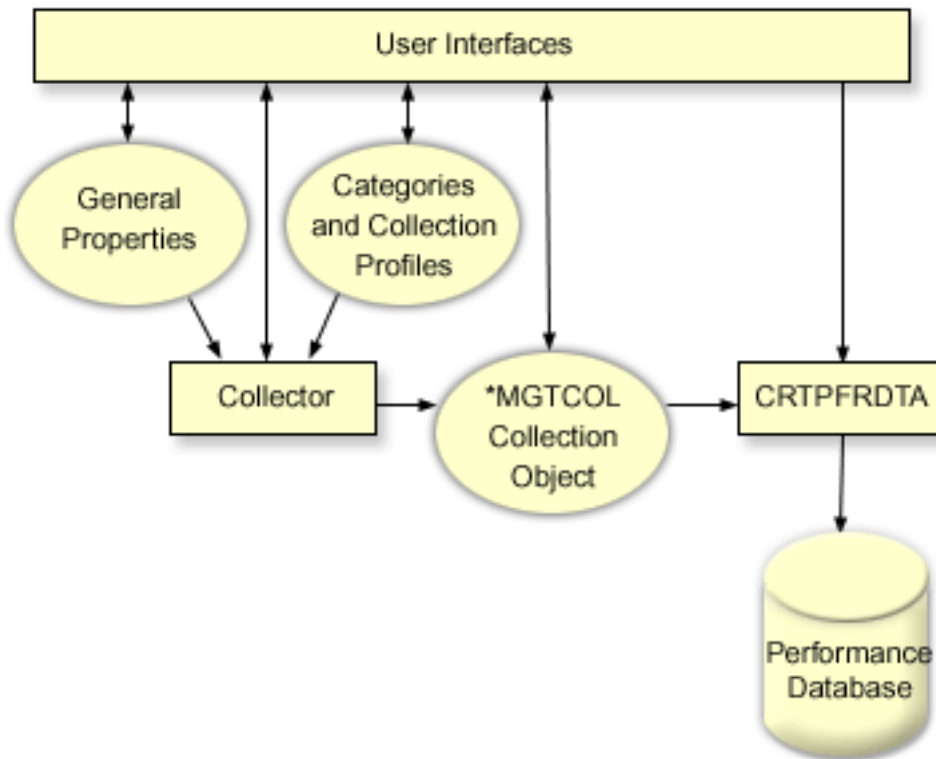
### Create Performance Data (CRTPFRTDA) command

The CRTPFRTDA command processes data that is stored in the management collection object and generates the performance database files.

### Performance database

The database files store the data that is processed by the CRTPFRTDA command. The files can be divided into these categories: Performance data files that contain time interval data, configuration data files, and trace data files.





## Creating database files from Collection Services data

Use this information to manually or automatically create database files from Collection Services data.

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.

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 create database files from the management collection object, where the data is stored after it has been 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 use either the System i Navigator interface or the CRTPFRDTA command.
  - You can activate PM Agent, which creates the database files during collection.
- | You can use the database files that you have created with the IBM Systems Director Navigator for i5/OS Performance interface 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 produce performance reports. This action allows you to analyze the performance data on another system without affecting the performance of the source system.

## Storing data in management collection objects instead of in database files

Why should you store the data in management collection objects instead of in the database files that you need to run your 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.

In addition, the single management collection object allows you to manage the data as a single object rather than as many files. The single collection object allows you to move the performance data between releases without converting the data. As long as you retain the collection objects, you can look back and analyze the performance-related events down to the level of detail that you collected.

#### Related tasks

“Creating database files” on page 70

To create database files, follow these steps.

#### Related reference

Create Performance Data (CRTPFRTDA) command

See the Create Performance Data (CRTPFRTDA) command for information about creating performance files.

## Customizing data collections

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 categories that are typically needed by Performance Tools, with the exception of communications data. The **Standard plus protocol** profile corresponds to all categories that are typically needed by Performance Tools, including communications data. Or you can select **Custom** to create your own customized profile. There are also several other profiles available. Refer to the online help for detailed descriptions. 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.)

- | **Note:** When the default value is set to any specified time, all categories use the specified time except  
| those categories with:
- | • Explicit time intervals as may be set up in the \*CUSTOM profile.
  - | • Categories with interval restriction such as, disk storage, input/output processors, and  
| communications-related categories which must collect at least every 5 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 cycle interval from one hour to 24 hours, depending on how you plan to use the data.

#### Related tasks

“Configuring Collection Services” on page 70

Configure Collection Services by doing the following.

#### | Collection Services collection profiles:

| Descriptions of the Collection Services collection profiles. The collection profiles define what is collected.

- | • **Minimum** - The most commonly used set of performance data. This is the minimum data collection recommended and includes the following categories:
  - | – System Bus: This category contains data on the operation of each system bus.
  - | – Storage Pool: This category contains storage pool configuration data and storage pool operation data.
  - | – Hardware Configuration: This category contains hardware resource information for the system. This category contains the same data that the Display Hardware Resources (DSPHDWRSC) command acquires. Only the first instance of this data will be reported in the database if more than one instance is encountered.
  - | – System CPU: This category contains data on system CPU usage for each processor.
  - | – System-Level Data: This category contains general system data that is used on a system-wide basis.
  - | – Job MI: This category contains information on every active task, job, and thread in the system. The data collected is provided by the machine interface (MI).
  - | – Job OS: This category contains information on every active job in the system. The data collected is provided by the operating system.
  - | – Disk Storage: This category contains system storage unit data. It includes base storage unit information and operational data for disk drives.
  - | – IOP: This category contains data on the system’s input/output processors (IOPs). It includes data on IOP bus use and IOP utilization by adapter resources.
- | • **Standard** - The data categories typically needed by the tools in Performance Tools for i5/OS, except for communications protocol data. The data categories in the Standard profile include all the categories in the Minimum profile plus the following:
  - | – Storage Pool Tuning: This category contains pool tuning configuration data for each system storage pool.
  - | – Subsystem: This category contains data on active subsystems and subsystem pools. Only the first instance of this data will be reported in the database if more than one instance is encountered.
  - | – SNADS Transaction: This category contains transaction boundary information specific to active SNADS jobs in the system.
  - | – Local Response Time: This category contains response time information for workstations connected to 5254 controllers. Response time data is reported for each workstation and is saved in a set of response time buckets.
  - | – APPN: This category contains data on the system’s APPN support. The data recorded contains both general information and data classified according to transaction type and work activity.
  - | – SNA: This category contains data on the system’s SNA support. Data is reported for each active T2 task consisting of controller, task, and session information.
  - | – TCP/IP Base: This category contains system-wide performance information for TCP/IP.
  - | – User Transactions: This category contains data for application-defined transactions rather than IBM-defined transactions. You can create your own user-defined transactions.
  - | – Domino: This category is included in this profile when the Domino for i5/OS licensed program is installed on the system.

- | – IBM HTTP Server (powered by Apache): This category is included in this profile when the IBM HTTP Server for i5/OS licensed program is installed on the system.
- | – WebSphere Application Server: This category is included in this profile when the IBM WebSphere Application Server licensed program is installed on the system.
- | – Java: This category contains IBM Technology for Java (J9) VM and thread data.
- | – Logical partition: This category contains performance data that is collected from eligible partitions if the IBM Director Server (xxxx-DR1) licensed program is installed on the partition that is running Collection Services. To collect data from other partitions, the IBM Director Agent (xxxx-DA1) licensed program must be installed on the other partitions and the server must be authorized to the other partitions.
- | • **Standard plus protocol** - The data categories typically needed by the tools in Performance Tools for i5/OS, including communications protocol data. The data categories in the Standard plus protocol profile include all the categories in the Standard profile with the addition of the following:
  - | – Network Server: This category contains information about network servers. For Integrated xSeries<sup>®</sup> Servers, data is reported for CPU utilization. For virtual I/O adapters on hosting partitions (partitions that provide the physical resources), data is provided about the I/O activity that occurs within this partition due to the virtual device support that it is providing on behalf of guest partitions.
  - | – Communications Base: This category contains base protocol information for each communication line that is available for use (varied on).
  - | – Communications Station: This category contains station information for certain communication lines. Data is reported for each station that is available for use (varied on). Protocols supporting this data are Token Ring, Ethernet, DDI, Frame Relay, and X.25.
  - | – Communications SAP: This category contains Service Access Point (SAP) information for certain communication lines. Data is reported for each configured SAP within lines available for use (varied on). Protocols supporting this data are Token Ring, Ethernet, DDI, and Frame Relay.
  - | – Data port services: This category contains performance data obtained from data port services. Data port services is Licensed Internal Code that supports the transfer of large volumes of data between a source system and one of any number of specified target systems in an iSeries cluster environment. Data port services is used by Licensed Internal Code clients, such as, remote independent auxiliary storage pool (ASP) mirroring.
  - | – TCP/IP Interface: This category contains information for each active TCP/IP interface.
- | • **Enhanced capacity planning** - The data categories in the Enhanced capacity planning profile include all the categories in the Standard plus protocol profile with the addition of the PEX Data - Processor Efficiency data category. The PEX Data - Processor Efficiency data category contains the cycles per instruction for performance explorer (PEX) data. Data may be collected to enhance capacity planning capabilities or for other purposes. Special considerations apply when using this category:
  - | – A performance explorer definition, QPMIPEXPEI, is created. If a performance explorer definition already exists, it is deleted and re-created.
  - | – This category requires Collection Services to start a performance explorer (PEX) collection (session-ID QPMINTPEXD). This collection can conflict with other performance explorer collections.
  - | – You should not end or start the QPMINTPEXD session manually because this will affect the validity of the data collected.
  - | – When collection of this category stops, it also ends the performance explorer collection for session QPMINTPEXD.
- | • **Custom** - A user-defined profile. You select the categories to collect.

#### 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 saving time or from daylight saving 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 saving time takes effect.

#### **Related concepts**

Date and time system values: Time of day

See the Date and time system values: Time of day topic for information about the QTIME system value.

Date and time system values: Offset from coordinated universal time (UTC)

See the Date and time system values: Offset from coordinated universal time (UTC) topic for information about the QUTCOFFSET system value.

## **Implementing user-defined categories in Collection Services**

The user-defined categories function in Collection Services enables applications to integrate performance data collection into Collection Services.

This allows you to gather data from an application by writing a data collection program, registering it, and integrating it with Collection Services. Collection Services will then call the data collection program at every collection interval, and will store the data in the collection object. You should use the Collection Object APIs listed below to access the data stored in the collection object. You may access the data in real-time, as it is being collected, or for as long as the collection object is retained.

To implement this function, you need to:

1. Develop a program to collect performance data for a new category in Collection Services.
2. Create a job description for your collection program. The job description QPMUSRCAT in QGPL provides an example, but does not represent default values or recommendations.
3. Register the new category and specify the data collection program.
  - Register: QypsRegCollectorDataCategory
  - De-register: QypsDeregCollectorDataCategory

After you register the category, Collection Services includes it in the list of available collection categories.

4. Add the category to your Collection Services profile, and then cycle Collection Services
5. Develop a program to query the collection object.

- Retrieve active management collection object name: QpmRtvActiveMgtcolName (Used only for querying the collection object in real-time)
- Retrieve management collection object attributes: QpmRtvMgtcolAttrs
- Open management collection object: QpmOpenMgtcol
- Close management collection object: QpmCloseMgtcol
- Open management collection object repository: QpmOpenMgtcolRepo
- Close management collection object repository: QpmCloseMgtcolRepo
- Read management collection object data: QpmReadMgtcolData

Your customized collection program now runs at each collection interval, and the collected data is archived in the collection objects.

You can also implement the Java versions of these APIs. The required Java classes are included in ColSrv.jar, in the integrated file system (IFS) directory QIBM/ProdData/OS400/CollectionServices/lib. Java applications should include this file in their classpath. For more information about the Java implementation, download the javadocs in a .zip file.

### Query the collection object in real-time

If your application needs to query the collection object in real-time, it will need to synchronize the queries with Collection Services. To do this, the application should create a data queue and register it with Collection Services. Once registered, the collector sends a notification for each collection interval and for the end of the collection cycle. The application should maintain the data queue, including removing the data queue when finished, and handling abnormal termination. To register and deregister the data queue, refer to the following APIs:

- Add collector notification: QypsAddCollectorNotification
- Remove collector notification: QypsRmvCollectorNotification

#### Related reference

##### QpmCloseMgtcol API

The Close Management Collection Object (QpmCloseMgtcol) API closes a management collection object that was previously opened by the Open Management Collection Object (QpmOpenMgtcol) API.

##### QpmCloseMgtcolRepo API

The Close Management Collection Object Repository (QpmCloseMgtcolRepo) API closes a repository of a management collection object that was previously opened by the Open Management Collection Object Repository (QpmOpenMgtcolRepo) API.

##### QpmOpenMgtcol API

The Open Management Collection Object (QpmOpenMgtcol) API opens a specified management collection object for processing and returns a handle to the open management collection object.

##### QpmOpenMgtcolRepo API

The Open Management Collection Object Repository (QpmOpenMgtcolRepo) API opens a specified repository of a management collection object for processing.

##### QpmReadMgtcolData API

The Read Management Collection Object Data (QpmReadMgtcolData) API returns information about a specific record in a repository of a management collection object.

##### QpmRtvActiveMgtcolName API

The Retrieve Active Management Collection Object Name (QpmRtvActiveMgtcolName) API returns the object name and library name of an active management collection object.

##### QpmRtvMgtcolAttrs API

The Retrieve Management Collection Object Attributes (QpmRtvMgtcolAttrs) API returns information about attributes of a management collection object and repositories of a management collection object.



#### QypsAddCollectorNotification API

The Add Collector Notification (QypsAddCollectorNotification) API registers with a collector to provide notifications to a specified data queue for a collection event.

#### QypsDeregCollectorDataCategory API

The Deregister Collector Data Category (QypsDeregCollectorDataCategory) API removes a user-defined data category from the Collection Services function of Management Central.

#### QypsRmvCollectorNotification API

The Remove Collector Notification (QypsRmvCollectorNotification) API removes a notification registration from a collector for a specified data queue and collection event.

#### QypsRegCollectorDataCategory API

The Register Collector Data Category (QypsRegCollectorDataCategory) API adds a user-defined data category to one or more collector definitions of the Collection Services function of Management Central.

### Collection program recommendations and requirements:

Collection Services calls the data collection program once during the start of a collection cycle, once for each collection interval, and again at the end of the collection cycle.

The data collection program must perform any data collection and return that data to a data buffer provided by Collection Services. In addition to providing a data buffer, Collection Services also provides a work area, which allows the data collection program to maintain some state information between collection intervals.

The data collection program should collect data as quickly as possible and should perform minimal formatting. The program should not perform any data processing or sorting. Although data from the user-defined category is not converted into database files, Collection Services may run the Create Performance Data (CRTPFRTDA) command automatically and add the data from the collection object to database files at the end of each collection interval. If the data collection program cannot complete its tasks within the collection interval, the CRTPFRTDA command does not run properly.

**Note:** By using the code examples, you agree to the terms of the “Code license and disclaimer information” on page 193.

You may create the data collection program in several environments:

- \*PGM for OPM languages. This environment may not be used to query the collection object and may result in poor performance. However, it is supported for older programming languages.
- \*SRVPGM, an entry point in a service program. This is for ILE languages.
- \*JVAPGM, the required Java classes are included in ColSrv.jar. This file is included in the IFS at QIBM/ProdData/OS400/CollectionServices/lib. Download the javadocs.zip file and open index.html for a description of the Java implementations of the APIs.

Collection Services sends the following requests to the data collection program:

Request	Description
Start collection	The data collection program should initialize any interfaces or resources used during data collection. Optionally, it may also initialize a work area, provided by Collection Services, that preserves state information between collection intervals. If you want to include a control record prior to the collected data, the data collection program may also write a small amount of data to the data buffer. Typically, this control record would be used during data processing to help interpret the data.

Request	Description
Collection interval	Collection Services sends an interval request for each collection interval. The data collection program should collect data and return it in the data buffer. Collection Services then writes that data to the interval record in the collection object. If the amount of data is too large for the data buffer, the data collection program should set a "More data" flag. This action causes Collection Services to send another interval request with a modifier indicating that it is a continuation. Collection Services resets the more data flag before each call. This process is repeated until all the data is moved into the collection object.
End of collection	When the collection for the category containing the data collection program ends, Collection Services sends this request. The data collection program should perform any cleanup and can optionally return a collection control record. The data collection program should also send a return code that indicates the result of the collection.
Clean up and terminate (Shutdown)	Collection Services sends this request if an abnormal termination is necessary. Operating system resources are freed automatically when the data collection program job ends, but any other shutdown operations should be performed by the data collection program. The data collection program can receive this request at any time.

For a detailed description of these parameters, the work area, data buffer, and return codes, refer to the header file QPMDCPRM, which is located in QSYSINC.

### Data storage in collection objects

Collection objects have a repository for each data collection category. This repository gets created by Collection Services when collections for that category are started. Each repository consists of the following records:

Record	Description
Control	This optional record can be the first or last record that results from the data collection program, and may occur in both positions. Typically, it should contain any information needed to interpret the record data.
Interval	Each collection interval creates an interval record, even if it is empty. The interval record contains the data written to the data buffer during the collection interval. It must not exceed 4 GB in size.
Stop	Collection Services automatically creates this record to indicate the end of a data collection session. If the collections for the user-defined category were restarted without ending or cycling Collection Services, you can optionally include a control record followed by additional interval records after the stop record.

### Example: Implementing user-defined categories:

Look here for sample programs that illustrate how you can use the provided APIs to integrate customized data collections into Collection Services.

*Example: Data collection program:*

This program example collects some test data and stores it in a data buffer, which Collection Services copies to the collection object.

**Note:** By using the code examples, you agree to the terms of the "Code license and disclaimer information" on page 193.

## C++ sample code

```
#include "string.h"           // memcpy(), memset(), strlen()
#include "stdio.h"            // printf()
#include "qpmcprm.h"          // data collection program interface
#include "time.h"

extern "C"
void DCPentry( Qpm_DC_Parm_t *request, char *dataBuffer,
               char *workArea, int *returnCode )
{
    static char testData[21] = "Just some test stuff";
    int i;

    /* Print contents of request structure */

    printf( "DCP called with parameters:\n" );
    printf( "  format name: \"%8.8s\"; category name: \"%10.10s\";\n",
            request->formatName, request->categoryName );
    printf( "  rsvd1: %4.4X; req type: %d; req mod: %d; buffer len: %d;\n",
            *(short *) (request->rsvd1), request->requestType,
            request->requestModifier, request->dataBufferLength );
    printf( "  prm offset: %d; prm len: %d; work len: %d; rsvd2: %8.8X;\n",
            request->parmOffset, request->parmLength, request->workAreaLength,
            *(int *) (request->rsvd2) );
    printf( "  rec key: \"%8.8s\"; timestamp: %8.8X %8.8X;\n",
            request->intervalKey,
            *(int *) (request->intervalTimestamp),
            *(int *) (request->intervalTimestamp + 4) );
    printf( "  return len: %d; more data: %d; rsvd3: %8.8X %8.8X;\n",
            request->bytesProvided, request->moreData,
            *(int *) (request->rsvd3),
            *(int *) (request->rsvd3 + 4) );

    switch ( request->requestType )
    {
        /* Write control record in the beginning of collection */
        case PM_DOBEGIN:
            printf( "doBegin(%d)\n", request->requestModifier );
            switch ( request->requestModifier )
            {
                case PM_CALL_NORMAL:
                    memcpy( dataBuffer, testData, 20 );
                    *(int *) workArea = 20;
                    request->moreData = PM_MORE_DATA;
                    request->bytesProvided = 20;
                    break;

                case PM_CALL_CONTINUE:
                    if ( *(int *) workArea < 200 )
                    {
                        memcpy( dataBuffer, testData, 20 );
                        *(int *) workArea += 20;
                        request->moreData = PM_MORE_DATA;
                        request->bytesProvided = 20;
                    }
                    else
                    {
                        *(int *) workArea = 0;
                        request->moreData = PM_NO_MORE_DATA;
                        request->bytesProvided = 0;
                    }
                    break;

                default:
                    *returnCode = -1;
            }
        }
    }
```



```

        return;
    }
    break;
/* Write control record in the end of collection */
case PM_DOEND:
    printf( "doEnd(%d)\n", request->requestModifier );
    switch ( request->requestModifier)
    {
        case PM_CALL_NORMAL:
            memcpy( dataBuffer, testData, 20 );
            *(int *)workArea = 20;
            request->moreData = PM_MORE_DATA;
            request->bytesProvided = 20;
            break;

        case PM_CALL_CONTINUE:
            if ( *(int *)workArea < 200 )
            {
                memcpy( dataBuffer, testData, 20 );
                *(int *)workArea += 20;
                request->moreData = PM_MORE_DATA;
                request->bytesProvided = 20;
            }
            else
            {
                *(int *)workArea = 0;
                request->moreData = PM_NO_MORE_DATA;
                request->bytesProvided = 0;
            }
            break;

        default:
            *returnCode = -1;
            return;
    }
    break;

/*Write interval record */
case PM_DOCOLLECT:
    printf( "doCollect(%d)\n", request->requestModifier );
    for ( i = 0; i < 10000; i++ )
        dataBuffer[i] = i % 256;
    request->bytesProvided = 10000;

    switch ( request->requestModifier)
    {
        case PM_CALL_NORMAL:
            *(time_t *)workArea + 4 = time( NULL );
            *(int *)workArea = 1;
            request->moreData = PM_MORE_DATA;
            break;

        case PM_CALL_CONTINUE:
            *(int *)workArea += 1;
            if ( *(int *)workArea < 20 )
                request->moreData = PM_MORE_DATA;
            else
            {
                *(time_t *)workArea + 8 = time( NULL );
                printf( "doCollect() complete in %d secs (%d bytes transferred)\n",
                    *(time_t *)workArea + 8 - *(time_t *)workArea + 4, 10000 * 20 );
                request->moreData = PM_NO_MORE_DATA;
            }
            break;

        default:
            *returnCode = -1;
    }

```

```

        return;
    }
    break;
/* Clean-up and terminate */
case PM_DOSHUTDOWN:
    printf( "doShutdown\n" );
    *returnCode = 0;
    return;
    break;

default:
    *returnCode = -1;
    return;
    break;
}

}/* DCPentry() */

```

### Related concepts

“Collection program recommendations and requirements” on page 32

Collection Services calls the data collection program once during the start of a collection cycle, once for each collection interval, and again at the end of the collection cycle.

*Example: Program to register the data collection program:*

This sample program registers the data collection program from the previous example with Collection Services. After running, Collection Services displays the data collection program in the list of data collection categories.

**Note:** By using the code examples, you agree to the terms of the “Code license and disclaimer information” on page 193.

### C++ sample code

```

| #include "stdlib.h"
| #include "stdio.h"
| #include "string.h"
| #include "qypscoll.h"
|
|
| int main( int argc, char *argv[] )
| {
|     int    CCSID = 0;
|     int    RC = 0;
|     Qyps_USER_CAT_PROGRAM_ATTR    *pgmAttr;
|     Qyps_USER_CAT_ATTR            catAttr;
|     char   collectorName[11] = "*PFR    ";
|     char   categoryName[11]  = "TESTCAT ";
|     char   collectorDefn[11] = "*CUSTOM "; /* Register to *CUSTOM profile only */
|
|     if ( argc > 2 )
|     {
|         int len = strlen( argv[2] );
|
|         if ( len > 10 ) len = 10;
|         memset( categoryName, ' ', 10 );
|         memcpy( categoryName, argv[2], len );
|     }
|
|     if ( argc < 2 || *argv[1] == 'R' )
|     {
|         pgmAttr = (Qyps_USER_CAT_PROGRAM_ATTR *)malloc( 4096 );
|         memset( pgmAttr, 0x00, sizeof(pgmAttr) );
|         pgmAttr->fixedPortionSize = sizeof( Qyps_USER_CAT_PROGRAM_ATTR );
|         memcpy( pgmAttr->programType,    "*SRVPGM    ", 10 );
|     }
|
| }

```

```

|         memcpy( pgmAttr->parameterFormat, "PMDC0100", 8 );
|         memcpy( pgmAttr->ownerUserId, "USERID ", 10 );
|         memcpy( pgmAttr->jobDescription, "QPMUSRCAT QGPL ", 20 );
|         memcpy( pgmAttr->qualPgmSrvpgmName, "DCPTST LIBRARY ", 20 );
|         pgmAttr->workAreaSize = 123;
|         pgmAttr->srvpgmEntrypointOffset = pgmAttr->fixedPortionSize;
|         pgmAttr->srvpgmEntrypointLength = 8;
|         pgmAttr->categoryParameterOffset = pgmAttr->srvpgmEntrypointOffset +
|                                           pgmAttr->srvpgmEntrypointLength;
|         pgmAttr->categoryParameterLength = 10;
|     /* Set entry point name */
|     memcpy( (char *) (pgmAttr) + pgmAttr->srvpgmEntrypointOffset,
|             "DCPentry", pgmAttr->srvpgmEntrypointLength ); /* Set parameter string */
|     memcpy( (char *) (pgmAttr) + pgmAttr->categoryParameterOffset,
|             "1234567890", pgmAttr->categoryParameterLength );
|
|     memset( &catAttr, 0x00, sizeof(catAttr) );
|     catAttr.structureSize = sizeof( Qyps_USER_CAT_ATTR );
|     catAttr.minCollectionInterval = 0;
|     catAttr.maxCollectionInterval = 0;
|     catAttr.defaultCollectionInterval = 30; /* Collect at 30 second interval */
|     memset( catAttr.qualifiedMsgId, ' ', sizeof(catAttr.qualifiedMsgId) );
|     memcpy( catAttr.categoryDesc,
|             "12345678901234567890123456789012345678901234567890", sizeof(catAttr.categoryDesc) );
|
|     QypsRegCollectorDataCategory( collectorName,
|                                   categoryName,
|                                   collectorDefn,
|                                   &CCSID,
|                                   (char*)pgmAttr,
|                                   (char*)&catAttr,
|                                   &RC
|                                   );
|
| }
| else
|     if( argc >= 2 && *argv[1] == 'D' )
|         QypsDeregCollectorDataCategory( collectorName, categoryName, &RC );
|     else
|         printf("Unrecognized option\n");
|
| } /* main() */

```

*Example: Program to query the collection object:*

This sample program illustrates how to query the data stored in the collection object using the Java classes shipped in the ColSrv.jar file in the QIBM/ProdData/OS400/CollectionServices/lib directory path.

**Note:** By using the code examples, you agree to the terms of the “Code license and disclaimer information” on page 193.

### Java sample code

```

import com.ibm.iseries.collectionservices.*;

class testmco2
{
    public static void main( String argv[] )
    {
        String    objectName = null;
        String    libraryName = null;
        String    repoName = null;
        MgtcolObj mco = null;
        int       repoHandle = 0;
        int       argc = argv.length;
        MgtcolObjAttributes
            attr = null;
    }
}

```

```

MgtcolObjRepositoryEntry
    repoE = null;
MgtcolObjCollectionEntry
    collE = null;
int      i,j;

if ( argc < 3 )
{
    System.out.println("testmco2  objectName libraryName repoName");
    System.exit(1);
}

objectName = argv[0];
libraryName = argv[1];
repoName   = argv[2];

if ( ! objectName.equals( "ACTIVE" ) )
    mco = new MgtcolObj( objectName, libraryName );
else
    try
    {
        mco = MgtcolObj.rtvActive();
    } catch ( Exception e)
    {
        System.out.println("rtvActive(): Exception " + e );
        System.exit(1);
    }
System.out.println("Object name = " + mco.getName() );
System.out.println("Library name = " + mco.getLibrary() );

try
{
    attr = mco.rtvAttributes( "MCOA0100" );
} catch ( Exception e)
{
    System.out.println("rtvAttributes(): MCOA0100: Exception " +
e );
    System.exit(1);
}

System.out.println("MCOA0100: Object " + mco.getLibrary() + "/" + mco.getName() );
System.out.println("    size = " + attr.size + " retention = " + attr.retentionPeriod +
    " interval = " + attr.dftInterval + " time created = " + attr.timeCreated +
    " time updated = " + attr.timeUpdated );
System.out.println("    serial = " + attr.logicalPSN + " active = " + attr.isActive +
    " repaired = " + attr.isRepaired + " summary = " + attr.sumStatus +
    " repo count = " + attr.repositoryCount );
if ( attr.repositoryInfo != null )
    for(i = 0; i < attr.repositoryCount; i++ )
    {
        repoE = attr.repositoryInfo[ i ];
        System.out.println("        name = " + repoE.name + " category = " + repoE.categoryName +
            " size = " + repoE.size );
        for( j = 0; j < repoE.collectionInfo.length; j++ )
        {
            collE = repoE.collectionInfo[ j ];
            System.out.println("            startTime = " + collE.startTime + " endTime = " + collE.endTime +
                " interval = " + collE.interval );
        }
    }

try
{
    attr = mco.rtvAttributes( "MCOA0200" );
} catch ( Exception e)
{
    System.out.println("rtvAttributes(): MCOA0200: Exception " + e );

```

```

        System.exit(1);
    }

    System.out.println("MCOA0200: Object " + mco.getLibrary() + "/" + mco.getName() );
    System.out.println("    size = " + attr.size + " retention = " + attr.retentionPeriod +
        " interval = " + attr.dftInterval + " time created = " + attr.timeCreated +
        " time updated = " + attr.timeUpdated );
    System.out.println("    serial = " + attr.logicalPSN + " active = " + attr.isActive +
        " repaired = " + attr.isRepaired + " summary = " + attr.sumStatus +
        " repo count = " + attr.repositoryCount );
    if ( attr.repositoryInfo != null )
        for(i = 0; i < attr.repositoryCount; i++ )
        {
            repoE = attr.repositoryInfo[ i ];
            System.out.println("        name = " + repoE.name + " category = " + repoE.categoryName +
                " size = " + repoE.size );
            for( j = 0; j < repoE.collectionInfo.length; j++ )
            {
                collE = repoE.collectionInfo[ j ];
                System.out.println("            startTime = " + collE.startTime + " endTime = " + collE.endTime +
                    " interval = " + collE.interval );
            }
        }

    if ( repoName.equals("NONE") )
        return;

    try
    {
        mco.open();
    } catch ( Exception e)
    {
        System.out.println("open(): Exception " + e );
        System.exit(1);
    }

    try
    {
        repoHandle = mco.openRepository( repoName, "MCO0100" );
    } catch ( Exception e)
    {
        System.out.println("openRepository(): Exception " + e );
        mco.close();
        System.exit(1);
    }
    System.out.println("repoHandle = " + repoHandle );

    MgtcolObjReadOptions  readOptions = new MgtcolObjReadOptions();
    MgtcolObjRecInfo  recInfo = new MgtcolObjRecInfo();

    readOptions.option = MgtcolObjReadOptions.READ_NEXT;
    readOptions.recKey = null;
    readOptions.offset = 0;
    readOptions.length = 0;

    while ( recInfo.recStatus == MgtcolObjRecInfo.RECORD_OK )
    {
        try
        {
            mco.readData( repoHandle, readOptions, recInfo, null );
        } catch ( Exception e)
        {
            System.out.println("readData(): Exception " + e );
            mco.close();
            System.exit(1);
        }
    }

```

```

        if( recInfo.recStatus == MgtcolObjRecInfo.RECORD_OK )
        {
            System.out.print("Type = " + recInfo.recType );
            System.out.print(" Key = " + recInfo.recKey );
            System.out.println(" Length = " + recInfo.recLength );
        }

    }/* while ... */

    mco.closeRepository( repoHandle );
    mco.close();

}/* main() */

}/* class testmco2 */

```

## Managing collection objects

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

You can delete a collection object from the system. If you do not delete the objects manually, Collection Services will delete them automatically after the expiration date and time.

- | Collection Services deletes only expired management collection objects that exist in the configured
- | collection library.

Collection Services deletes the cycled collection objects that have reached their expiration 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. You can specify **Permanent** if you do not want Collection Services to delete your management collection objects for you.

### Related tasks

“Managing collections” on page 66

View performance data and manage your performance collections using IBM Systems Director Navigator for i5/OS.

## User-defined transactions

Collection Services and performance explorer collect performance data that you define in your applications.

With the provided APIs, you can integrate transaction data into the regularly scheduled sample data collections using Collection Services, and get trace-level data about your transaction by running performance explorer.

For detailed descriptions and usage notes, refer to the following API descriptions:

- Start Transaction (QYPESTRT, qypeStartTransaction) API
- End transaction (QYPEENDT, qypeEndTransaction) API
- Log transaction (QYPELOGT, qypeLogTransaction) API (Used only by performance explorer)
- Add trace point (QYPEADDT, qypeAddTracePoint) API (Used only by performance explorer)

**Note:** You only need to instrument your application once. Collection Services and performance explorer use the same API calls to gather different types of performance data.

## Integrating user-defined transaction data into Collection Services

You can select user-defined transactions as a category for collection in the Collection Services configuration. Collection Services then collects the transaction data at every collection interval and stores that data in the collection object. The Create Performance Data (CRTPFRTDA) command exports this data to the user-defined transaction performance database file, QAPMUSRTNS. Collection Services organizes the data by transaction type. You can specify as many transaction types as you require; however, Collection Services will only report the first 15 transaction types. Data for additional transaction types is combined and stored as the \*OTHER transaction type. At every collection interval, Collection Services creates one record for each type of transaction for each unique job. For a detailed description, refer to the usage notes in the Start Transaction API.

Collection Services gathers general transaction data, such as the transaction response time. You can also include up to 16 optional application-defined counters that can track application-specific data like the number of SQL statements used for the transaction, or other incremental measurements. Your application should use the Start Transaction API to indicate the beginning of a new transaction, and should include a corresponding End Transaction API to deliver the transaction data to Collection Services.

## Collecting trace information for user-defined transactions with performance explorer

You can use the Start, End, and Log Transaction APIs during a performance explorer session to create a trace record. Performance Explorer stores system resource utilization, such as CPU utilization, I/O, and seize/lock activity, for the current thread in these trace records. Additionally, you may choose to include application-specific performance data, and then send it to performance explorer in each of these APIs. You can also use the Add Trace Point API to identify application-specific events for which performance explorer should collect trace data.

To start a performance explorer session for your transactions, specify \*USRTRNS on the (OSEVT) parameter of your Performance Explorer definition. After entering the ENDPEX command, performance explorer writes the data supplied by the application to the QMUDTA field in the QAYPEMIUSR performance explorer database file. System-supplied performance data for the start, end, and any log records is stored in the QAYPEMIUSR and QAYPETIDX database files.

### Related concepts

“Performance explorer” on page 57

Performance explorer collects more detailed information about a specific application, program or system resource, and provides detailed insight into a specific performance problem. This includes the capability both to perform several types and levels of traces and to run detailed reports.

### Related reference

QYPESTRT, qypeStartTransaction API

QYPEENDT, qypeEndTransaction API

QYPELOGT, qypeLogTransaction API

QYPEADDT, qypeAddTracePoint API

Create Performance Data (CRTPFRTDA) command

### Related information

Performance data files: QAPMUSRTNS

## C++ example: Integrating user-defined transactions into Collection Services:

This C++ example program shows how to use the Start Transaction and End Transaction APIs to integrate user-defined transaction performance data into Collection Services.

**Note:** By using the code examples, you agree to the terms of the “Code license and disclaimer information” on page 193.

```

//*****
// tnstst.C
//
// This example program illustrates the use
// of the Start/End Transaction APIs (qypeStartTransaction,
// qypeEndTransaction).
//
//
// This program can be invoked as follows:
// CALL lib/TNSTST PARM('threads' 'types' 'transactions' 'delay')
//     where
//     threads      = number of threads to create (10000 max)
//     types        = number of transaction types for each thread
//     transactions = number of transactions for each transaction
//                   type
//     delay        = delay time (millisecs) between starting and
//                   ending the transaction
//
// This program will create "threads" number of threads. Each thread
// will generate transactions in the same way. A thread will do
// "transactions" number of transactions for each transaction type,
// where a transaction is defined as a call to Start Transaction API,
// then a delay of "delay" millisecs, then a call to End Transaction
// API. Thus, each thread will do a total of "transactions" * "types"
// number of transactions. Each transaction type will be named
// "TRANSACTION_TYPE_nnn" where nnn ranges from 001 to "types". For
// transaction type n, there will be n-1 (16 max) user-provided
// counters reported, with counter m reporting m counts for each
// transaction.
//
// This program must be run in a job that allows multiple threads
// (interactive jobs typically do not allow multiple threads). One
// way to do this is to invoke the program using the SBMJOB command
// specifying ALWMLTTHD(*YES).
//
//*****

#define _MULTI_THREADED

// Includes
#include "pthread.h"
#include "stdio.h"
#include "stdlib.h"
#include "string.h"
#include "qusec.h"
#include "lbcpyrv.h"
#include "qypesvp.h"

// Constants
#define maxThreads 10000

// Transaction pgm parm structure
typedef struct
{
    int types;
    int trans;
    int delay;
} tnsPgmParm_t;

// Error code structure
typedef struct
{
    Qus_EC_t error;
    char      Exception_Data[100];
} error_code_t;

```



```

//*****
//
// Transaction program to run in each secondary thread
//
//*****

void *tnsPgm(void *parm)
{
    tnsPgmParm_t *p = (tnsPgmParm_t *)parm;

    char tnsTyp[] = "TRANSACTION_TYPE_XXX";
    char pexData[] = "PEX";
    unsigned int pexDataL = sizeof(pexData) - 1;
    unsigned long long colSrvData[16] = {1,2,3,4,5,6,7,8,
                                          9,10,11,12,13,14,15,16};

    unsigned int colSrvDataL;
    char tnsStrTim[8];

    struct timespec ts = {0, 0};

    error_code_t errCode;

    _DPA_Template_T target, source; // Used for LBCPYNV MI instr

    unsigned int typCnt;
    unsigned int tnsCnt;
    int rc;

    // Initialize error code
    memset(&errCode, 0, sizeof(errCode));
    errCode.error.Bytes_Provided = sizeof(errCode);

    // Initialize delay time
    ts.tv_sec = p->delay / 1000;
    ts.tv_nsec = (p->delay % 1000) * 1000000;

    // Loop doing transactions
    for (tnsCnt = 1; tnsCnt <= p->trans; tnsCnt++)
    {
        for (typCnt = 1; typCnt <= p->types; typCnt++)
        {
            // Set number field in transaction type
            source.Type = _I_UNSIGNED;
            source.Length = 4;
            source.reserved = 0;
            target.Type = _I_ZONED;
            target.Length = 3;
            target.reserved = 0;
            _LBCPYNV(tnsTyp + 17, &target, &typCnt, &source);

            // Set Col1 Svcs data length in bytes
            colSrvDataL = (typCnt <= 16) ? (typCnt - 1) : 16;
            colSrvDataL = colSrvDataL * 8;

            // Call Start Transaction API
            qypeStartTransaction(tnsTyp,
                                (unsigned int *)&tnsCnt,
                                pexData,
                                (unsigned int *)&pexDataL,
                                tnsStrTim,
                                &errCode);

            // Delay specified amount
            rc = pthread_delay_np(&ts);

            // Call End Transaction API

```

```

        typeEndTransaction(tnsTyp,
                           (unsigned int *)&tnsCnt,
                           pexData,
                           (unsigned int *)&pexDataL,
                           tnsStrTim,
                           (unsigned long long *)&colSrvData[0],
                           (unsigned int *)&colSrvDataL,
                           &errCode);
    }
}

return NULL;
}

//*****
//
// Main program to run in primary thread
//
//*****

void main(int argc, char *argv[])
{
    // Integer version of parms
    int threads; // # of threads
    int types;   // # of types
    int trans;   // # of transactions
    int delay;   // Delay in millisecs

    pthread_t threadHandle[maxThreads];
    tnsPgmParm_t tnsPgmParm;
    int rc;
    int i;

    // Verify 4 parms passed
    if (argc != 5)
    {
        printf("Did not pass 4 parms\n");
        return;
    }

    // Copy parms into integer variables
    threads = atoi(argv[1]);
    types   = atoi(argv[2]);
    trans   = atoi(argv[3]);
    delay   = atoi(argv[4]);

    // Verify parms
    if (threads > maxThreads)
    {
        printf("Too many threads requested\n");
        return;
    }

    // Initialize transaction pgm parms (do not modify
    // these while threads are running)
    tnsPgmParm.types = types;
    tnsPgmParm.trans = trans;
    tnsPgmParm.delay = delay;

    // Create threads that will run transaction pgm
    for (i=0; i < threads; i++)
    {
        // Clear thread handle
        memset(&threadHandle[i], 0, sizeof(pthread_t));
        // Create thread

```

```

    rc = pthread_create(&threadHandle[i],      // Thread handle
                       NULL,                  // Default attributes
                       tnsPgm,                // Start routine
                       (void *)&tnsPgmParm); // Start routine parms
    if (rc != 0)
        printf("pthread_create() failed, rc = %d\n", rc);
}

// Wait for each thread to terminate
for (i=0; i < threads; i++)
{
    rc=pthread_join(threadHandle[i], // Thread handle
                    NULL);           // No exit status
}

} /* end of Main */

```

### Java example: Integrating user-defined transactions into Collection Services:

This Java example program shows how to use the Start Transaction and End Transaction APIs to integrate user-defined transaction performance data into Collection Services.

**Note:** By using the code examples, you agree to the terms of the “Code license and disclaimer information” on page 193.

```
import com.ibm.iseries.collectionservices.PerformanceDataReporter;
```

```

// parameters:
//  number of TXs per thread
//  number of threads
//  log|nolog
//  enable|disable
//  transaction seconds

public class TestTXApi
{
    static TestTXApiThread[]  thread;

    static private String[] TxTypeString;
    static private byte[] [] TxTypeArray;

    static private String TxEventString;
    static private byte[] TxEventArray;

    static
    {
        int i;

        // initialize transaction type strings and byte arrays

        TxTypeString = new String[20];
        TxTypeString[ 0] = "Transaction type 00";
        TxTypeString[ 1] = "Transaction type 01";
        TxTypeString[ 2] = "Transaction type 02";
        TxTypeString[ 3] = "Transaction type 03";
        TxTypeString[ 4] = "Transaction type 04";
        TxTypeString[ 5] = "Transaction type 05";
        TxTypeString[ 6] = "Transaction type 06";
        TxTypeString[ 7] = "Transaction type 07";
        TxTypeString[ 8] = "Transaction type 08";
        TxTypeString[ 9] = "Transaction type 09";
        TxTypeString[10] = "Transaction type 10";
        TxTypeString[11] = "Transaction type 11";
        TxTypeString[12] = "Transaction type 12";
        TxTypeString[13] = "Transaction type 13";
    }
}

```

```

TxTypeString[14] = "Transaction type 14";
TxTypeString[15] = "Transaction type 15";
TxTypeString[16] = "Transaction type 16";
TxTypeString[17] = "Transaction type 17";
TxTypeString[18] = "Transaction type 18";
TxTypeString[19] = "Transaction type 19";

TxTypeArray = new byte[20][];
for ( i = 0; i < 20; i++ )
    try
    {
        TxTypeArray[i] = TxTypeString[i].getBytes("Cp037");
    } catch(Exception e)
    {
        System.out.println("Exception \"" + e + "\" when converting");
    }

}/* static */

public static void main( String[] args )
{
    int    numberOfTXPerThread;
    int    numberOfThreads;
    boolean log;
    boolean enable;
    int    secsToDelay;

    // process parameters
    if ( args.length >= 5 )
try
    {
        numberOfTXPerThread = Integer.parseInt( args[0] );
        numberOfThreads     = Integer.parseInt( args[1] );

        if ( args[2].equalsIgnoreCase( "log" ) )
log = true;
        else
        if ( args[2].equalsIgnoreCase( "nolog" ) )
            log = false;
        else
        {
            System.out.println( "Wrong value for 3rd parameter!" );
            System.out.println( "\tshould be log|nolog" );
            return;
        }

        if ( args[3].equalsIgnoreCase( "enable" ) )
enable = true;
        else
        if ( args[3].equalsIgnoreCase( "disable" ) )
            enable = false;
        else
        {
            System.out.println( "Wrong value for 4th parameter!" );
            System.out.println( "\tshould be enable|disable" );
            return;
        }

        secsToDelay = Integer.parseInt( args[4] );

    } catch (Exception e)
    {
        System.out.println( "Oops! Cannot process parameters!" );
        return;
    }
}

```

```

else
{
    System.out.println( "Incorrect Usage." );
    System.out.println( "The correct usage is:" );
    System.out.println( "java TestTXApi numberOfTXPerThread numberOfThreads
log|nolog enable|disable secsToDelay");
    System.out.println("\tlog will make the program cut 1 log transaction per start / end pair");
    System.out.println("\tdisable will disable performance collection to minimize overhead");
    System.out.print("\nExample: \"java TestTXApi 10000 100 log enable 3\" will call " );
    System.out.println("cause 10000 transactions for each of 100 threads");
    System.out.println("with 3 seconds between start and end of transaction");
    System.out.println("Plus it will place additional log call and will enable reporting." );
    return;
}

System.out.println( "Parameters are processed:" );
System.out.println( "\tnumberOfTxPerThread = " + numberOfTXPerThread );
System.out.println( "\tnumberOfThreads = " + numberOfThreads );
System.out.println( "\tlog = " + log );
System.out.println( "\tenable = " + enable );
System.out.println( "\tsecsToDelay = " + secsToDelay );

// cause initialization of a PerformanceDataReporter class
{
    PerformanceDataReporter pReporter = new PerformanceDataReporter();
    pReporter.enableReporting();
}

TestTXApi t = new TestTXApi( );

System.out.println( "\nAbout to start ..." );
t.prepareTests( numberOfTXPerThread, numberOfThreads, log, enable, secsToDelay );

long startTime = System.currentTimeMillis();

t.runTests( numberOfThreads );

// wait for threads to complete
for ( int i = 0; i < numberOfThreads; i++ )
    try
    {
        thread[i].join( );
    } catch(Exception e)
    {
        System.out.println( "***Exception \"" + e + "\" while joining thread " + i );
    }

long endTime = System.currentTimeMillis();

System.out.println( "\nTest runtime for " + ( numberOfTXPerThread * numberOfThreads ) +
    " TXs was " + ( endTime - startTime ) + " msec" );

}/* main() */

private void prepareTests( int numberOfTxPerThread,
    int numberOfThreads, boolean log,
boolean enable, int secsToDelay )
{
    System.out.println( "Creating " + numberOfThreads + " threads");
    thread = new TestTXApiThread[numberOfThreads];
    for ( int i = 0; i < numberOfThreads; i++ )
        thread[i] = new TestTXApiThread( i, numberOfTxPerThread,
            log, enable, secsToDelay );
}

}/* prepareTests() */

private void runTests( int numberOfThreads )

```

```

{
    for ( int i = 0; i < numberOfThreads; i++ )
        thread[i].start( );
}

/* runTests() */

private class TestTXApiThread extends Thread
{
    private int      ordinal;
    private int      numberOfTxPerThread;
    private boolean log;
    private boolean enable;
    private int      secsToDelay;

    private PerformanceDataReporter    pReporter;

    private long    timeStamp[];
    private long    userCounters[];

    public TestTXApiThread( int ordinal, int numberOfTxPerThread,
                           boolean log, boolean enable, int secsToDelay )
    {
        super();
        this.ordinal          = ordinal;
        this.numberOfTxPerThread = numberOfTxPerThread;
        this.log              = log;
        this.enable           = enable;
        this.secsToDelay      = secsToDelay;

        pReporter = new PerformanceDataReporter( false );
        if ( enable )
            pReporter.enableReporting();
        timeStamp = new long[1];
        userCounters = new long[16];
        for ( int i = 0; i < 16; i++ )
            userCounters[i] = i;
    }

    /* constructor */

    public void run()
    {
        int i;

        for ( i = 0; i < numberOfTxPerThread; i++ )
        {
            pReporter.startTransaction( TxTypeArray[i%20], i, TxTypeArray[i%20], 20, timeStamp );
            pReporter.startTransaction( TxTypeArray[i%20], i, TxTypeString[i%20], timeStamp );
            if ( log )
                pReporter.logTransaction( TxTypeArray[i%20], i, TxTypeArray[i%20], 20 );
            pReporter.logTransaction( TxTypeArray[i%20], i, TxTypeString[i%20] );
            if ( secsToDelay > 0 )
                try
                {
                    Thread.sleep(secsToDelay * 1000);
                } catch(Exception e) { }
            pReporter.endTransaction( TxTypeArray[i%20], i, TxTypeArray[i%20], 20, timeStamp,
                                     userCounters );
            pReporter.endTransaction( TxTypeArray[i%20], i, TxTypeString[i%20], timeStamp,
                                     userCounters );
        }
    }

    /* run() */

}

/* class TestTXApiThread */

/* class TestTXApi */

```

## Collecting performance data across partitions

Collection Services, along with IBM Director Multiplatform, can retrieve data about CPU usage and number of processors available from your server partitions regardless of the operating system running on them.

One of the analyses IBM Performance Management for System i5 (PM for System i5) provides is to plot the growth of the system to determine when an upgrade may be necessary. For a system that is not partitioned, this is a straightforward process. However, if your system has been partitioned into multiple i5/OS partitions, the data arrives at IBM from each partition separately, which makes forming a reliable view of the entire system performance more difficult. If the partitions are running AIX® or Linux®, or if any of the i5/OS partitions have PM Agent turned off, then no data is sent, which makes forming a view of the entire system nearly impossible.

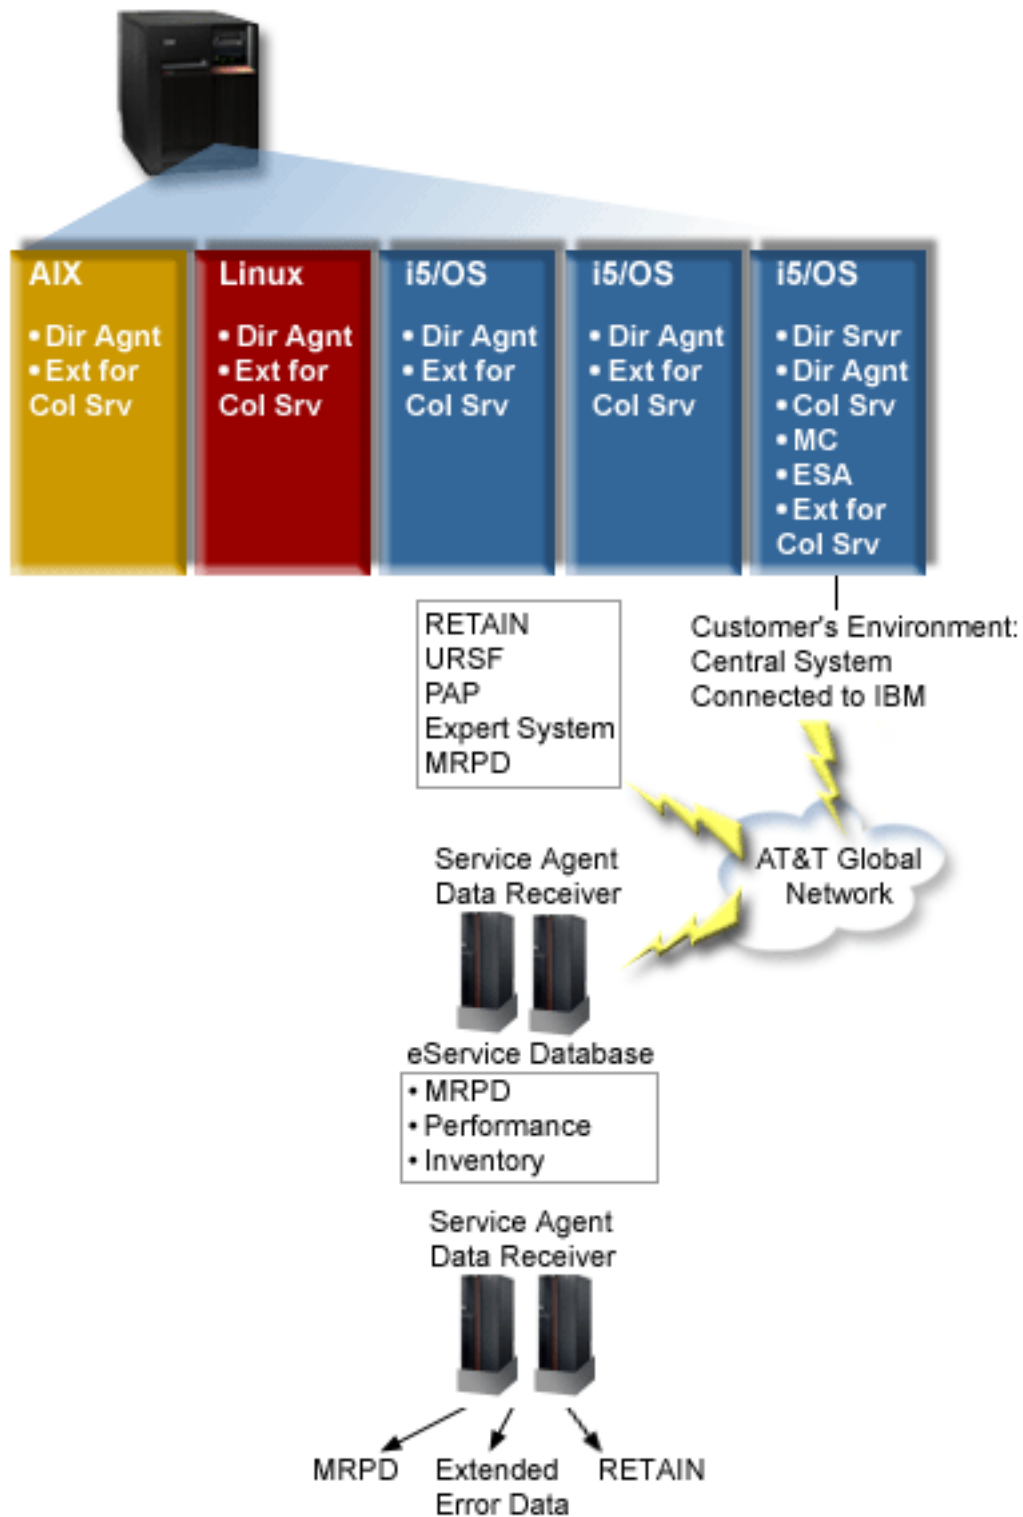
To address these problems, Collection Services, with IBM Director Multiplatform, can now retrieve data about CPU usage and number of processors available from your server partitions regardless of the operating system running on them. PM Agent summarizes the data before it gets shipped to IBM. Providing a cross-partition view of CPU utilization helps you and IBM do a much better job of managing your system resources. This includes balancing workload across the current set of processors as well as being able to plan for the purchase of more or faster processors when necessary.

The graphic below illustrates how the collection of CPU utilization data across logical partitions works. The central system has the IBM Director Server installed on an i5/OS partition that is running Collection Services with the \*LPAR category selected. Each of the other partitions must have the IBM Director Agent installed and configured so that IBM Director Server can collect performance data from them. Each partition must also have the Director Multiplatform extension for Collection Services installed.

IBM Director Server retrieves the CPU utilization data for each partition, including itself, at regular intervals and stores that data in the Collection Services \*MGTCOL object. The data is then processed and written to the QAPMLPAR database file. Finally, PM Agent collects and aggregates the data and prepares to transmit it to IBM.

Although this graphic shows Management Central and IBM Electronic Service Agent™ (ESA) set up to transmit data on the same partition as the IBM Director Server and Collection Services, the transmission mechanism to IBM could actually be running on a completely different system and still be set up to gather the cross-partition data from PM Agent and send it to IBM, as usual.





Key
Dir Srvr = IBM Director Server
Dir Agnt = IBM Director Agent
Col Srv = Collection Services
MC = Management Central
ESA = IBM Electronic Service Agent
Ext for Col Srv = Director Multiplatform extension for Collection Services
RETAIN = Remote technical assistance information network
URSF = Universal remote support facility
MRPD = Machine Reported Product Data

## Set up data collection across logical partitions

The following procedure provides you with an overview of the steps you must complete to collect performance data across logical partitions:

1. Ensure your IP network is properly configured for all logical partitions on the same physical system.
2. Ensure you are running a supported operating system on each logical partition for which you want to collect performance data:
  - i5/OS, Version 5 Release 4 or later
  - AIX 5L™, version 5.3
  - Red Hat Enterprise Linux AS, version 3.0, for IBM PowerPC®
  - SUSE LINUX Enterprise Server 8 for IBM pSeries® and IBM iSeries
  - SUSE LINUX Enterprise Server 9 for IBM pSeries and IBM iSeries
3. Ensure that you have applied the following Collection Services fixes to the logical partition that acts as your management server:
  - SI12971
  - SI13838 (superseded by SI16328)
  - SI15131 (superseded by SI16499)
  - SI16328 (Linux support)
  - SI16499 (AIX support)

For the latest information about Collection Services cross-partition support for Linux operating systems, see the informational APAR II13986.

Go to Fix Central for the latest PTF fixes.

4. Use the Virtualization Engine™ to install IBM Director Server on the i5/OS partition that you want to act as the management server. Consider the management server the central control point that communicates to managed systems, devices, and Collection Services. When the Virtualization Engine installation wizard is complete, IBM Director Server and IBM Director Agent are installed on the i5/OS partition that you want to act as the management server.
5. Install IBM Director Agent on the logical partitions that you want to be managed by IBM Director Server. These logical partitions must be on the same physical system as the logical partition where IBM Director Server is installed.
6. Install IBM Director Console on the system that you want to function as your Director Multiplatform management console.
7. Complete the required configuration steps:
  - a. Authorize users for i5/OS on the management partition.
  - b. Start Director Multiplatform on each logical partition.
  - c. Start IBM Director console on your management console.

- d. In IBM Director Console, add each logical partition on which you want to monitor performance by right-clicking in the Group Contents pane and selecting **New → IBM Director Systems**.
  - e. After you have added each partition, request access to manage the logical partition. In the Group Contents pane, right-click the partition and select **Request Access**.
8. If you have IBM Director version 5.20 or later installed on the partition, skip this step. On the i5/OS management partition, install the Director Multiplatform extension for Collection Services by copying the necessary files for Collection Services from the Collection Services directory to the appropriate Director Multiplatform directory. The Collection Service files are ColSrvLparDataExt.TWGExt, ColSrvLparDataSubagt.TWGSubagent, and ColSrvDir.jar. Copy the Collection Services files using the following commands:
- ```

CPY
OBJ('/qibm/proddata/os400/collectionservices/lib/ColSrvLparDataExt.TWGExt')
  TODIR('/qibm/userdata/director/classes/extensions')

CPY
OBJ('/qibm/proddata/os400/collectionservices/lib/ColSrvLparDataSubagt.TWGSubagent')
  TODIR('/qibm/userdata/director/classes/extensions')

CPY
OBJ('/qibm/proddata/os400/collectionservices/lib/ColSrvDir.jar')
  TODIR('/qibm/userdata/director/classes')

```
9. If you have IBM Director version 5.20 or later installed on the partition, skip this step. Distribute the Collection Services files from the management partition to the i5/OS partitions from which you plan to collect performance data. You can do this by File Transfer Protocol (FTP) with the binary option, or by mapping a drive and copying the files to the file system, or by any other distribution mechanism that you might have in place. You can access the files on the i5/OS management partition in the directory /qibm/proddata/os400/collectionservices/lib.
- a. Distribute ColSrvLparDataExt.TWGExt to the Director Multiplatform extensions directory /qibm/userdata/director/classes/extensions, on the i5/OS partition that you want to manage.
  - b. Distribute ColSrvLparDataSubagt.TWGSubagent to the Director Multiplatform extensions directory /qibm/userdata/director/classes/extensions, on the i5/OS partition that you want to manage.
  - c. Distribute ColSrvDir.jar to Director Multiplatform classes directory /qibm/userdata/director/classes, on the i5/OS partition that you want to manage.
10. If you have IBM Director version 5.20 or later installed on the partition, skip this step. On each Linux partition, install the Director Multiplatform extension for Collection Services by installing the Collection Services RPM file ColSrvDirExt.rpm.
- a. Distribute the Collection Services RPM file from the management partition to the Linux partitions from which you plan to collect performance data. You can do this by File Transfer Protocol (FTP) with the binary option, or by mapping a drive and copying the files to the file system, or by any other distribution mechanism that you might have in place. You can use Qshell to access the RPM file in the i5/OS management partition directory /qibm/proddata/os400/collectionservices/lib/ColSrvDirExt.rpm.
  - b. On each Linux partition, run the following command from the directory where the RPM file exists:
 

```
rpm -Uhv --force ColSrvDirExt.rpm
```
11. If you have IBM Director version 5.20 or later installed on the partition, skip this step. On each AIX partition, install the Director Multiplatform extension for Collection Services by installing the Collection Services package aix-ColSrvDirExt.bff.
- a. Distribute the Collection Services package file from the management server to the AIX partitions from which you plan to collect performance data. You can do this by File Transfer Protocol (FTP) with the binary option, or by mapping a drive and copying the files to the file system, or by any other distribution mechanism that you might have in place. You can use Qshell to access the package file in the i5/OS management partition directory /qibm/proddata/os400/collectionservices/lib/aix-ColSrvDirExt.bff.

- b. On each AIX partition, run the following command from the directory where the BFF file exists:  

```
installp -Fac -d aix-ColSrvDirExt.bff ColSrvDirExt
```
12. In IBM Director Console, update the collection inventory on each partition by right-clicking the partition and selecting **Perform Inventory Collection**.
13. Activate PM Agent.
14. Use the following Start Performance Collection (STRPFCOL) command to start Collection Services:  

```
STRPFCOL CYCCOL(*YES)
```

You can use the following tools to track performance data across partitions:

- Collection Services
- PM for System i5
- IBM Systems Workload Estimator

#### **Related concepts**

Configure TCP/IP

“IBM Performance Management for System i5” on page 88

The IBM Performance Management for System i5 (PM for System i5) offering automates the collection, archival, and analysis of system performance data and returns reports to help you manage system resources and capacity.

Management Central

#### **Related tasks**

“Activating PM Agent” on page 89

PM Agent is a part of the operating system and you must activate it to use its collecting capabilities.

Partitioning the server

“Sending PM Agent data with Service Agent over Extreme Support (Universal Connection)” on page 90

PM Agent uses Collection Services to gather the nonproprietary performance and capacity data from your server. After you have collected this data, you can use Electronic Service Agent over Extreme Support to send the data to IBM.

#### **Related information**

Performance data files: QAPMLPAR



Install IBM Director Server



Installing IBM Director Agent

Installing IBM Director Console



Authorizing users for i5/OS



Starting IBM Director console



Start IBM Director



Virtualization Engine



Domino 6 for iSeries Best Practices Guide (Workload Estimator)

## **Finding wait statistics for a job, task, or thread**

During the running of a job, task, or thread, conditions arise that cause that process to wait (for example, while the system resolves a lock or hold on a required object).

Collection Services can collect data on the cause and duration of the time a process spends waiting. This data is reported in the Collection Services database files QAPMJOBWT and QAPMJOBWTD.

**Note:** To query the QAPMJOBWTD file, the CCSID of your job must be set to the CCSID of the primary language installed on the system (not to 65535 binary data).

| Another tool that shows job wait statistics is IBM i5/OS Job Watcher. IBM i5/OS Job Watcher returns  
| real-time information about a selected set of jobs, threads, and Licensed Internal Code (LIC) program  
| tasks. At specified time intervals, IBM i5/OS Job Watcher samples anywhere from one thread per job to  
| all threads per job. IBM i5/OS Job Watcher gathers a variety of performance data, including detailed wait  
| statistics for jobs, tasks, and threads.

| There are 32 wait buckets which accumulate wait state data. These static wait buckets, used by both  
| Collection Services and IBM i5/OS Job Watcher, provide a stable view of the wait state data. In Collection  
| Services, data from these buckets is reported in files QAPMJOBWT and QAPMJOBWTG. In Job Watcher,  
| data from these buckets is reported in QAPYJWTD and QAPYJWSTS.

#### **Related concepts**

“IBM i5/OS Job Watcher” on page 56

IBM i5/OS Job Watcher provides for the collection of job data for any or all jobs, threads, and tasks on the system. It provides call stacks, SQL statements, objects being waited on, Java JVM statistics, wait statistics and more which are used to diagnose job related performance problems.

#### **Related information**

Performance data files: QAPMJOBWT

Performance data files: QAPMJOBWTD

Performance data files: QAPMJOBWTG

A jobs life

Work management job attributes

## **Understanding disk consumption by Collection Services**

The amount of disk resource Collection Services consumed varies greatly depending on the settings that you use.

For illustration purposes, assume that Collection Services is used daily and cycles at midnight, causing each \*MGTCOL object to contain one day’s worth of data collection. Next, establish a base size for one day’s worth of data collection by using the default properties for Collection Services. A standard plus protocol profile with an interval value of 15 minutes can collect 500 MB of data in the \*MGTCOL object. The size actually collected per day using the default properties can vary greatly depending on system size and usage. The 500 MB example might represent a higher-end system that is heavily used.

| Interval rate | Intervals per collection | Multiplier | Size in MB |
|---------------|--------------------------|------------|------------|
| 15 minutes    | 96                       | 1          | 500        |

The size of one day’s worth of data is directly proportional to the number of intervals collected per collection period. For example, changing the interval rate from 15 minutes to 5 minutes increases the number of intervals by a factor of 3 and increases the size by the same factor.

| Interval rate | Intervals per collection | Multiplier | Size in MB |
|---------------|--------------------------|------------|------------|
| 15 minutes    | 96                       | 1          | 500        |
| 5 minutes     | 288                      | 3          | 1500       |

To continue this example, the following table shows the size of one \*MGTCOL object produced each day by Collection Services at each interval rate, using the default standard plus protocol profile.

| Interval rate | Intervals per collection | Multiplier | Size in MB |
|---------------|--------------------------|------------|------------|
| 15 minutes    | 96                       | 1          | 500        |
| 5 minutes     | 288                      | 3          | 1500       |
| 1 minutes     | 1440                     | 15         | 7500       |
| 30 seconds    | 2880                     | 30         | 15000      |
| 15 Seconds    | 5760                     | 60         | 30000      |

The size of the \*MGTCOL object, in this example, can vary from 500 MB to 30 GB depending on the rate of collection. You can predict a specific system's disk consumption for one day's collection interval through actual observation of the size of the \*MGTCOL objects created, using the default collection interval of 15 minutes and the standard plus protocol profile as the base and then using the multiplier from the above table to determine the disk consumption at other collection intervals. For example, if observation of the \*MGTCOL object size reveals that the size of the object for a day's collection is 50 MB for 15-minute intervals, then you could expect Collection Services to produce \*MGTCOL objects with a size of 3 GB when collecting data at 15-second intervals.

**Note:** Use caution when considering a collection interval as frequent as 15 seconds. Frequent collection intervals can adversely impact system performance.

## Retention period

The retention period also plays a significant role in the amount of disk resource that Collection Services consumes. The default retention period is one day. However, practically speaking, given the default values, the \*MGTCOL object is deleted on the third day of collection past the day on which it was created. Thus, on the third day of collection there is two days' worth of previously collected data plus the current day's data on the system. Using the table above, this translates into having between 1 GB and 1.5 GB of disk consumption at 15-minute intervals, and 60 to 90 GB of disk consumption at 15-second intervals on the system during the third day and beyond.

The formula to calculate disk consumption based on the retention period value is:

$(\text{Retention period in days} + 2.5) * \text{Size of one day's collection} =$   
Total Disk Consumption

**Note:** 2.5 corresponds to two days of previous collection data, and an average of the current day (2 days + 1/2 day).

Using the above tables and formula, a retention period of 2 weeks gives you a disk consumption of 8.25 GB at 15-minute intervals and 495 GB at 15-second intervals for the example system.

It is important to understand the disk consumption by Collection Services to know the acceptable collection interval and retention period for a given system. Knowing this can ensure that disk consumption will not cause system problems. Remember to consider that a system monitor or a job monitor can override a category's collection interval to graph data for a monitor. A system administrator must ensure that monitors do not inadvertently collect data at intervals that cause excess data consumption.

## Collecting ARM performance data

You can use Collection Services to collect Application Response Measurement (ARM) performance data.

The ARM APIs collect the performance data for ARM transactions. (The ARM APIs are a set of APIs developed by the Open Group to allow applications to report the progress of application transactions.) These transactions are reported in the QAPMARMTRT and QAPMUSRTNS database files.



Enterprise Workload Manager (EWLM) is a robust performance management tool that allows you to view the performance of work that runs in your business environment. Furthermore, EWLM allows you to assign performance goals to specific work. This allows you to monitor application-level transactions separate from operating system processes or monitor the work that a partition processes as a whole entity. EWLM encourages middleware and third-party software vendors to instrument their applications with the Open Group Application Response Measurement 4.0 (ARM) APIs. EWLM uses the ARM APIs to collect detailed performance data from ARM applications.

To learn more about ARM APIs, visit The Open Group Web site at [www.theopengroup.org/arm](http://www.theopengroup.org/arm).

## IBM i5/OS Job Watcher

IBM i5/OS Job Watcher provides for the collection of job data for any or all jobs, threads, and tasks on the system. It provides call stacks, SQL statements, objects being waited on, Java JVM statistics, wait statistics and more which are used to diagnose job related performance problems.

IBM i5/OS Job Watcher is similar in sampling function to the CL commands WRKACTJOB and WRKSYSACT in that each refresh computes delta information for the ending interval. The data collected from the jobs, threads, or tasks being watched is done so in a non-intrusive manner.

IBM i5/OS Job Watcher can be configured and managed through the IBM Systems Director Navigator for i5/OS Performance interface or CL commands.

### Related tasks

“Managing IBM i5/OS Job Watcher” on page 72

Manage IBM i5/OS Job Watcher using IBM Systems Director Navigator for i5/OS.

### Related reference

Add Job Watcher Definition (ADDJWDFN)

See the Add Job Watcher Definition (ADDJWDFN) command for information about specifying the performance data that is to be collected during a Job Watcher collection.

End Job Watcher (ENDJW)

See the End Job Watcher (ENDJW) command for information about ending a Job Watcher collection.

Remove Job Watcher Definition (RMVJWDFN)

See the Remove Job Watcher Definition (RMVJWDFN) command for information about removing a Job Watcher definition from the system.

Start Job Watcher (STRJW)

See the Start Job Watcher (STRJW) command for information about starting a Job Watcher collection.

## IBM i5/OS Disk Watcher

IBM i5/OS Disk Watcher provides for the collection of disk performance data to diagnose disk related performance problems.

IBM i5/OS Disk Watcher allows you to obtain data concerning I/O operations to disk units, along with frequently needed run-time data to determine which objects, files, processes, threads, and tasks are being accessed. This tool surfaces data beyond what is provided by such tools as WRKDSKSTS, WRKSYSSTS, and WRKSYSACT. Disk Watcher provides a mechanism to use short and longer duration traces to collect disk I/O data along with the associated task and object name.

Some potential uses of this tool are:

- Evaluating the performance of I/O operations on multi-path disk units
- Evaluating the performance of I/O queuing
- Determining how performance may be improved by re-spreading data across units
- Determining the optimal placement of devices, IOAs, or buses



| IBM i5/OS Disk Watcher can be configured and managed through the IBM Systems Director Navigator  
| for i5/OS Performance interface or CL commands.

| **Related tasks**

| “Managing IBM i5/OS Disk Watcher” on page 71  
| Manage IBM i5/OS Disk Watcher using IBM Systems Director Navigator for i5/OS.

| **Related reference**

| Add Disk Watcher Definition (ADDDWDFN)  
| See the Add Disk Watcher Definition (ADDDWDFN) command for information about adding a Disk  
| Watcher definition from the system.  
| End Disk Watcher (ENDDW)  
| See the End Disk Watcher (ENDDW) command for information about ending a Disk Watcher  
| collection.  
| Remove Disk Watcher Definition (RMVDWDFN)  
| See the Remove Disk Watcher Definition (RMVDWDFN) command for information about removing a  
| Disk Watcher definition from the system.  
| Start Disk Watcher (STRDW)  
| See the Start Disk Watcher (STRDW) command for information about starting a Disk Watcher  
| collection.

| **Related information**

| Disk Watcher data files

## Performance explorer

Performance explorer collects more detailed information about a specific application, program or system resource, and provides detailed insight into a specific performance problem. This includes the capability both to perform several types and levels of traces and to run detailed reports.

Performance explorer is a data collection tool that helps the user identify the causes of performance problems that cannot be identified by collecting data using Collection Services or by doing general trend analysis. Two reasons to use performance explorer include:

- Isolating performance problems to the system resource, application, program, procedure, or method that is causing the problem
- Analyzing the performance of applications

The AS/400® Performance Explorer Tips and Techniques book provides additional examples of the performance explorer functions and examples of the enhanced performance explorer trace support.

Performance explorer is a tool that helps find the causes of performance problems that cannot be identified by using tools that do general performance monitoring. As your computer environment grows both in size and in complexity, it is reasonable for your performance analysis to gain in complexity as well. The performance explorer addresses this growth in complexity by gathering data on complex performance problems.

**Note:** Performance explorer is the tool you need to use after you have tried the other tools. It gathers specific forms of data that can more easily isolate the factors involved in a performance problem; however, when you collect this data, you can significantly affect the performance of your system.

This tool is designed for application developers who are interested in understanding or improving the performance of their programs. It is also useful for users knowledgeable in performance management to help identify and isolate complex performance problems.

### Related concepts

“User-defined transactions” on page 40

Collection Services and performance explorer collect performance data that you define in your applications.

### Related information



AS/400 Performance Explorer Tips and Techniques book



Performance Tools PDF

## Performance explorer concepts

Performance explorer works by collecting detailed information about a specified system process or resource. This topic explains how performance explorer works, and how best to use it.

Performance explorer has advantages for people who need detailed performance analysis. Using performance explorer you can:

- Determine what is causing a performance problem on the system down to the level of user, job, file, object, thread, task, program, module, procedure, statement, or instruction address.
- Collect performance information on user-developed and system software.
- Do a detailed analysis on one job without affecting the performance of other operations on the system.
- Analyze data on a system other than the one on which it was collected. For example, if you collect data on a managed system in your network, you can send it to the central site system for analysis.

Like Collection Services, performance explorer collects data for later analysis. However, they collect very different types of data. Collection Services collects a broad range of system data at regularly scheduled intervals, with minimal system resource consumption. In contrast, performance explorer starts a session that collects trace-level data. This trace generates a large amount of detailed information about the resources consumed by an application, job, or thread. Specifically, you can use Performance Explorer to answer specific questions about areas like system-generated disk I/O, procedure calls, Java method calls, page faults, and other trace events. It is the ability to collect very specific and very detailed information that makes the performance explorer effective in helping isolate performance problems. For example, Collection Services can tell you that disk storage space is rapidly being consumed. You can use performance explorer to identify what programs and objects are consuming too much disk space, and why.


When performance explorer is running, it creates only the files that are needed for the collection.

**Note:** You can collect performance explorer data and Collections Services data at the same time.

## How performance explorer works

The following figure should help you become familiar with the normal path through the performance explorer. For details on each of these steps, see *Configure performance explorer*. The figure shows a basic work cycle that consists of the following steps:

1. Define a performance explorer data collection. You can also add a filter to limit the amount of data collected by specifying a compare value for specific events.
2. Start the performance explorer to collect the data based on your definition.
3. Run your program, command, or workload.
4. End the collection, which saves the collected data to a set of database files.
5. Create and print reports from the database files.



To learn more about performance explorer, refer to any of the following performance explorer topics.

#### **Related concepts**

“Collection Services” on page 24

Collection Services provides for the collection of system management data. It is the primary collector of system data.

#### **Related tasks**

“Configuring performance explorer” on page 64

To collect detailed trace information, you need to tailor performance explorer to work optimally with the application process from which the trace is being taken.

### **Performance explorer definitions:**

The parameters and conditions that determine what data performance explorer collects and how it collects it are configured and stored using performance explorer definitions. This topic explains how to use these definitions and provides a sample illustrating a simple definition.

To collect performance explorer data, you need to tell performance explorer what data to gather. You do this by using the Add Performance Explorer Definition (ADDPEXDFN) command to create a performance explorer definition. After the definition is completed and saved, you are ready to continue to the next task in the cycle of work.

Before creating a new definition, consider what kinds of information you want and the amount of detail you need. The performance explorer provides the following types of data collection:

#### **Statistics type definitions**

Identifies applications and IBM programs or modules that consume excessive CPU use or that perform a high number of disk I/O operations. Typically, you use the statistical type to identify programs that should be investigated further as potential performance bottlenecks.

- Good for first order analysis of i5/OS programs, procedures, and MI complex instructions.
  - Gives number of invocations
  - Gives both inline and cumulative CPU usage in microseconds
  - Gives both inline and cumulative number of synchronous and asynchronous I/O
  - Gives number of calls made
- Works well for short or long runs

- Size of the collected data is fairly small and constant for all runs
- Run time collection overhead of ILE procedures may be a problem due to the frequency of calls. Although run time is degraded, the collected statistics are still accurate because Performance Explorer removes most of the collection overhead from the data.
- Uses combined or separated data areas. The MRGJOB parameter on the ADDPEXDFN command specifies whether all program statistics are accumulated in one data area, or kept separate (for example, one data area for each job).

The statistics can be structured in either a hierarchical or flattened manner.

- A hierarchical structure organizes the statistics into a call tree form in which each node in the tree represents a program procedure run by the job or task.
- A flattened structure organizes the statistics into a simple list of programs or procedures, each with its own set of statistics.

Here is an example of a performance explorer statistics definition called MYSTATS that will show CPU and disk resource usage on a per program or procedure level.

```
ADDPEXDFN DFN(MYSTATS) /* The name of the definition. */
TYPE(*STATS) /* The type of definition */
JOB(*ALL) /*All Jobs */
TASKS(*ALL) /*All tasks */
DTAORG(*FLAT) /* Do not keep track of who calls who */
```

## Profile type definitions

Identifies high-level language (HLL) programs, modules, procedures, and statements that consume excessive CPU utilization based on source program statement numbers.

- Program profile (specify TYPE(\*PROFILE) and PRFTYPE(\*PGM) on the ADDPEXDFN command)
  - Gives detailed breakdown of where you are spending time within a set of programs within a specific job.
  - Can summarize the data by program, module, procedure, statement, or instruction.
  - Size of collection is fairly small and constant regardless of length of run.
  - Limit of 16 MI programs means that you should use this as a second order analysis tool.
  - Can vary overhead by changing sample interval. An interval of 2 milliseconds seems a good first choice for benchmarks.
  - No restrictions on pane size due to the number of programs specified or the size of the programs specified.

Here is an example of a performance explorer program profile definition called PGMPROF that will show usage for a particular procedure.

```
ADDPEXDFN DFN(PGMPROF) /* The name of the definition. */
TYPE(*PROFILE) /* The type of definition */
JOB(*ALL) /*All Jobs */
PGM((MYLIB/MYPGM MYMODULE MYPROCEDURE)) /* The name of the program to monitor. */
INTERVAL(1) /* 1-millisecond samples will be taken. */
```

- Job profile (specify the following on the ADDPEXDFN command: TYPE(\*PROFILE) and PRFTYPE(\*JOB))
  - Gives detailed breakdown of where you are spending time in the set of jobs or tasks of the collection.
  - Size of collection is relatively small but not constant. The size increases as the length of the run increases.
  - Can profile all jobs and tasks on the system or can narrow the scope of data collected to just a few jobs or tasks of interest.
  - Can vary overhead by changing sample interval. An interval of 2 milliseconds seems a good first choice for benchmarks.

Here is an example of a performance explorer job profile definition called ALLJOBPROF that will show usage for all your jobs.

```
ADDPDXDFN DFN(ALLJOBPROF) /* The name of the definition. */
TYPE(*PROFILE) /* The type of definition */
PRFTYPE(*JOB) /* A job profile type will be monitored. */
JOB(*ALL) /*All Jobs */
TASKS(*ALL) /*All tasks */
INTERVAL(1) /* 1-millisecond samples will be taken. */
```

## Trace definitions

Gathers a historical trace of performance activity generated by one or more jobs on the system. The trace type gathers specific information about when and in what order events occurred. The trace type collects detailed reference information about programs, Licensed Internal Code (LIC) tasks, i5/OS job, and object reference information.

- Some common trace events are:
  - Program and procedure calls and returns
  - Storage, for example, allocate and deallocate.
  - Disk I/O, for example, read operations and write operations.
  - Java method, for example, entry and exit.
  - Java, for example, object create and garbage collection.
  - Journal, for example, start commit and end commit.
  - Synchronization, for example, mutex lock and unlock or semaphore waits.
  - Communications, for example, TCP, IP, or UDP.
- Longer runs collect more data.

Here is an example of a performance explorer trace definition called DISKTRACE that will show usage for all disk events.

```
ADDPDXDFN DFN(DISKTRACE) /* The name of the definition. */
TYPE(*TRACE) /* The type of definition */
JOB(*ALL) /*All Jobs */
TASKS(*ALL) /*All tasks */
TRCTYPE(*SLTEVT) /* Only selected individual events and machine instructions
are included in the trace definition */
SLTEVT(*YES) /* *SLTEVT allows you to specify individual machine instructions
and events to be specified in addition to the categories of events
available with the TRCTYPE parameter. */
DSKEVT((*ALL)) /* All disk events are to be traced. */
```

Here is an example of a performance explorer trace definition called HEAPEVENTS.

```
ADDPDXDFN DFN(HEAPEVENTS) /* The name of the definition. */
TYPE(*TRACE) /* The type of definition */
JOB(*ALL) /*All Jobs */
TASK(*ALL) /*All tasks */
MAXSTG (100000) /*Maximum storage. Set to 100000 because the default of
10000 KB is often too small for the large number of heap events that can be
generated when tracing all jobs and all tasks.*/
TRCTYPE(*HEAP) /* Selects all heap events from the STGEVT
(storage events) parameter. */
```

### Related concepts

“Performance explorer reports” on page 63

After you have collected performance data with a performance explorer session, you can view it by running the included reports or by querying the database files directly.

### Related tasks

“Configuring performance explorer” on page 64

To collect detailed trace information, you need to tailor performance explorer to work optimally with the application process from which the trace is being taken.

#### Related reference

Add Performance Explorer Definition (ADDPEXDFN) command

#### Performance explorer database files:

The data that performance explorer collects is stored in performance explorer database files.

The following table shows the performance explorer (PEX) data files collected by the system when using data collection commands. Type the Display File Field Description (DSPFFD) command as follows to view the contents for a single file:

DSPFFD FILE(xxxxxxxxxx)

where xxxxxxxxx is the name of the file that you want to display.

| Type of information contained in file                              | File name  |
|--------------------------------------------------------------------|------------|
| Trace Resources Affinity                                           | QAYPEAFN   |
| Auxiliary storage management event data                            | QAYPEASM   |
| Auxiliary storage pool (ASP) information data                      | QAYPEASPI  |
| Base event data                                                    | QAYPEBASE  |
| Basic configuration information                                    | QAYPECFG   |
| Communications event data                                          | QAYPECMN   |
| Disk event data                                                    | QAYPEDASD  |
| Disk server event data                                             | QAYPEDSRV  |
| Event type and subtype mapping                                     | QAYPEEVENT |
| File Serving event data                                            | QAYPEFILSV |
| Configured filter information                                      | QAYPEFTRI  |
| Performance measurement counter (PMC) selection                    | QAYPEFQCFG |
| Heap event data                                                    | QAYPEHEAP  |
| Hardware monitor data                                              | QAYPEHMON  |
| Hardware monitor total data                                        | QAYPEHTOT  |
| Performance explorer Java event data                               | QAYPEJVA   |
| Performance explorer Java class information data                   | QAYPEJVC   |
| Performance explorer Java method information data                  | QAYPEJVM   |
| Performance explorer Java name information data                    | QAYPEJVNI  |
| Licensed Internal Code (LIC) bracketing data                       | QAYPELBRKT |
| Machine interface (MI) complex instructions collected on           | QAYPELCPLX |
| Jobs collected on                                                  | QAYPELJOB  |
| Licensed Internal Code (LIC) modules to collect data on            | QAYPELLIC  |
| Metrics to collect data on                                         | QAYPELMET  |
| Machine interface (MI) program, module, or procedures collected on | QAYPELMI   |
| Task names to collect data on                                      | QAYPELNAMT |
| Task number to collect data on                                     | QAYPELNUMT |

| Type of information contained in file                        | File name  |
|--------------------------------------------------------------|------------|
| Configured tasks                                             | QAYPELTASK |
| Machine interface (MI) program bracketing data               | QAYPEMBRKT |
| Machine interface (MI) complex instructions mapping          | QAYPEMICPX |
| Addresses of machine interface (MI) pointer                  | QAYPEMIPTR |
| Machine interface (MI) user event data                       | QAYPEMIUSR |
| Portable Application Solutions Environment (PASE) event data | QAYPEPASE  |
| Page fault event data                                        | QAYPEPGFLT |
| Program profile data                                         | QAYPEPPANE |
| Licensed Internal Code (LIC) address resolution mapping      | QAYPEPROCI |
| Resource management process event data                       | QAYPERMPM  |
| Resource management seize lock event data                    | QAYPERMSL  |
| Reference information                                        | QAYPEREF   |
| Miscellaneous resolution data                                | QAYPERINF  |
| Database level indicator                                     | QAYPERLS   |
| General information                                          | QAYPERUNI  |
| Segment address range (SAR) data                             | QAYPESAR   |
| Segment address resolution mapping                           | QAYPESEGI  |
| Basic statistics data                                        | QAYPESTATS |
| Synchronization event data                                   | QAYPESYNC  |
| Process and task resolution mapping                          | QAYPETASKI |
| Trace job equivalent event data                              | QAYPETBRKT |
| Common trace data for all events                             | QAYPETIDX  |
| Trace index data (by time and task)                          | QAYPETIDXL |
| Trace index data (by time)                                   | QAYPETID2L |
| Task switch event data                                       | QAYPETSKSW |
| User-defined bracketing hook data                            | QAYPEUSRDF |

## Migration of performance explorer database files

The performance explorer (PEX) database files change from release to release, as new events and new data are added to the files. When you migrate to a new release of i5/OS, if the system finds incompatible PEX database files, it moves these files to the QPEXD $\nu$ rmxx library, where  $\nu$ rm=version. The system displays a status message that indicates that the files are being moved. After the files are moved, the system displays a completion message that indicates whether the move succeeded or failed. If the move fails, the system displays the Incompatible repository message.

### Related concepts

“Performance explorer reports”

After you have collected performance data with a performance explorer session, you can view it by running the included reports or by querying the database files directly.

### Performance explorer reports:

After you have collected performance data with a performance explorer session, you can view it by running the included reports or by querying the database files directly.



Performance explorer gathers detailed information about a program or job's behavior and performance and stores this information in performance explorer database files. You can query these files with SQL, or by running one of several reports. You can generate four different reports with performance explorer: Statistics, Profile, Trace, and Base reports. See Performance explorer definitions for information on why you would use a particular definition to generate one of these reports. Each report is discussed in detail in the Performance Tools.

You can create and print performance explorer reports by using the Print Performance Explorer Report (PRTPEXRPT) command. Use the OUTFILE parameter when you want to customize your Trace Report. The following commands are examples for printing reports for each type of performance explorer data:

- Print a \*STATS report sorting by the CPU time used  
PRTPEXRPT MBR(MYSTATS) LIB(MYLIB) TYPE(\*STATS) STATSOPT(\*CPU)
- Print a profile report summarized by procedure  
PRTPEXRPT MBR(MYPROFILE) LIB(MYLIB) TYPE(\*PROFILE)  
PROFILEOPT(\*SAMPLECOUNT \*PROCEDURE)
- Print a trace sorted by task ID  
PRTPEXRPT MBR(MYTRACE) LIB(MYLIB) TYPE(\*TRACE) TRACEOPT(\*TASK)

Performance explorer stores its collected data in the QAVPETRCI file, which is located in the QPFR library. Type the following command to view the contents for a single record:

```
DSPFFD FILE(QPFR/QAVPETRCI)
```

#### **Related concepts**

“Performance explorer definitions” on page 59

The parameters and conditions that determine what data performance explorer collects and how it collects it are configured and stored using performance explorer definitions. This topic explains how to use these definitions and provides a sample illustrating a simple definition.

#### **Related reference**

“Performance explorer database files” on page 62

The data that performance explorer collects is stored in performance explorer database files.

Print Performance Explorer Report (PRTPEXRPT) command

#### **Related information**



Performance Tools PDF

## **Configuring performance explorer**

To collect detailed trace information, you need to tailor performance explorer to work optimally with the application process from which the trace is being taken.

To configure performance explorer, follow these steps:

1. Create a session definition that informs the system which performance data you want to collect. On the Add Performance Explorer Definition (ADDPEXDFN) display, specify the collection type and a name for the definition. This definition is stored as a database member by that name in the QAPEXDFN file in library QUSRSYS. The name that you specify is used on the Start Performance Explorer (STRPEX) command.
2. (Optional) Add a filter (Add PEX Filter (ADDPEXFTR) command). A performance explorer filter identifies the performance data that is to be collected during a performance explorer session, and is meant to limit the amount of data collected by specifying a compare value for specific events.
3. Start collecting data (Start Performance Explorer (STRPEX) command). A job may be in more than one performance explorer collection if the \*PMCO event is not being collected. If the \*PMCO event is being collected, then a job can be in more than one collection only if all the collections have the same interval specification (ADDPEXDFN INTERVAL() parameter). You can specify a definition and optional filter on the STRPEX command.

4. Run your command, program, or workload for data that you want to analyze.
5. Stop collecting the data and save it to database files for analysis. Use the End Performance Explorer (ENDPEX) command to stop the collection.
6. Analyze the performance data. The Print Performance Explorer Report (PRTPEXRPT) command provides unique reports for each type of data (statistical, profile, trace profile, or trace).  
The following are other options for analysis:
  - Write your own queries for the set of database files.
  - Use iDoctor for System i5 - PEX Analyzer. Use iDoctor for System i5 is a set of software performance analysis tools and associated services that extend your ability to evaluate the health of your system by gathering detailed information and providing automated, graphical analysis of this data.
  - Use the IBM Performance Trace Data Visualizer for iSeries (PTDV). PTDV is a Java application that you can use for performance analysis of applications. PTDV works with the Performance Explorer to allow you to view program flows and get details such as CPU time, wall clock time, number of cycles, and number of instructions, summarized by trace, job, thread, and procedures.
7. To end the performance explorer session, use the End Performance Explorer (ENDPEX) command.

All of the performance explorer commands can be accessed with one of the following methods:

- The command interface. Type the commands from the command line. All the commands are part of the i5/OS operating system.
- The Performance Tools menu options.

#### **Related concepts**

“Performance explorer concepts” on page 58

Performance explorer works by collecting detailed information about a specified system process or resource. This topic explains how performance explorer works, and how best to use it.

“Performance explorer definitions” on page 59

The parameters and conditions that determine what data performance explorer collects and how it collects it are configured and stored using performance explorer definitions. This topic explains how to use these definitions and provides a sample illustrating a simple definition.

#### **Related reference**

Add PEX filter (ADDPEXFTR) command

Start Performance Explorer (STRPEX) command

Print Performance Explorer Report (PRTPEXRPT) command

### **Ending performance explorer:**

To end the performance explorer session, use the End Performance Explorer (ENDPEX) command.

The End Performance Explorer (ENDPEX) command performs the following actions on the collected data:

- Places the collected data in files QAYPExxx in the library that you specify. Use OPTION(\*END) and DTAOPT(\*LIB) to do this. The database member name for all the QAYPExxx files uses the session name as the default unless you specify a name for the DTAMBR parameter. You can specify RPLDTA(\*NO) to not overwrite existing data with the new data or RPLDTA(\*YES) to overwrite the existing data with the new data. Unless you are a very sophisticated user, use RPLDTA(\*NO).
- Places the collected data into a single IBM-defined file. Use OPTION(\*END) and DTAOPT(\*MGTCOL) to do this. Typically, you would use \*MGTCOL only under the direction of an IBM service representative. Specifying the \*MGTCOL value on the DTAOPT parameter saves the collection information into a management collection object. The management collection object option should be used only if the data is going to be shipped to IBM. The performance tools can analyze only the database files.

- Discards the collected data. Use `OPTION(*END)` if you want to save the data or `DTAOPT(*DLT)` to discard any collected data. You do this when you determine the collected data cannot be used. For example, one of the suspected jobs did not start as expected. If you choose the `*DLT` option, the collected performance data for the session is never saved.
- Suspends the collection session but does not end it. Use `OPTION(*SUSPEND)` to do this. You can later start the data collection again by issuing the `STRPEX` command with `OPTION(*RESUME)` for the specific session ID.

**Note:** If you forget the active collection session name, use the `ENDPEX SSNID(*SELECT)` command.

## Performance Trace Data Visualizer

Performance Trace Data Visualizer for iSeries is a tool for processing, analyzing, and viewing Performance Explorer collection data residing in PEX database files.

For more information, go to the Performance Trace Data Visualizer Web site.

## | IBM Systems Director Navigator for i5/OS Performance interface

| The IBM Systems Director Navigator for i5/OS Performance interface allows you to display and manage performance data by bringing together a variety of performance information into one place.

| The Performance interface is part of the IBM Systems Director Navigator for i5/OS. The Performance interface allows you to display summarized data in many different types of reports and graphs. This can provide you with a quick picture of what is happening on your system. The interface allows you to start with a summarized view and then drill down into the data for more details.

### | Related information

| IBM Systems Director Navigator for i5/OS

## | Managing collections

| View performance data and manage your performance collections using IBM Systems Director Navigator for i5/OS.

### | Viewing a collection:

| To view a collection, follow these steps.

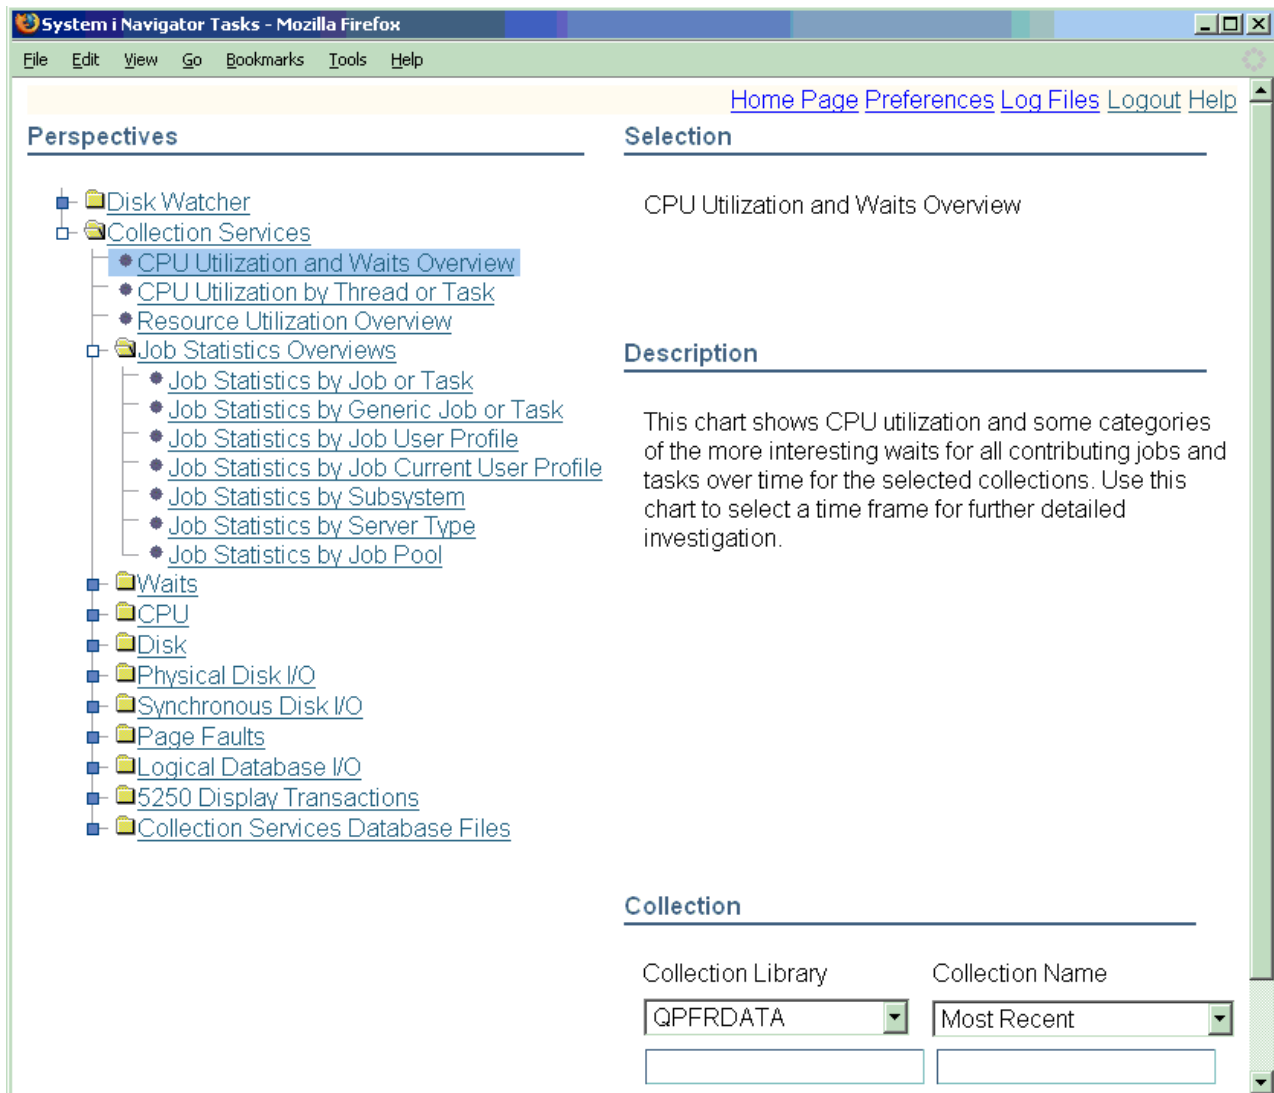
1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
2. Select **Investigate Data**.
3. Select the collection library.
4. Select the collection name.
5. Expand the content package by clicking the solid box.
6. Keep expanding the nodes in the tree until you navigate to the perspective that you would like to use.
7. Select the perspective.
8. Click **OK**.

### | *Investigate Data:*

| Investigate Data allows you to view data for each of the collectors (Collection Services, IBM i5/OS Job Watcher, and IBM i5/OS Disk Watcher).

| Each collector has an associated content package that defines the perspectives and views of the data. A content package is similar to a spreadsheet application's template file. The collections provide the data to be rendered.

The following shows an example of the perspectives list page. Perspectives are shown in the navigation tree.



The following shows an example of a view with a chart. You can quickly navigate to a previous view or to the Perspectives list by clicking on links in the navigation trail. For example, clicking on **Perspective List** will return you to the list of perspectives navigation tree.

You can take several actions on the chart using the toolbar buttons. The toolbar buttons are described as follows.



This toolbar button allows you to select one or more items on the chart. Some views allow actions to be performed on selected items.



This toolbar button allows you to pan the chart.



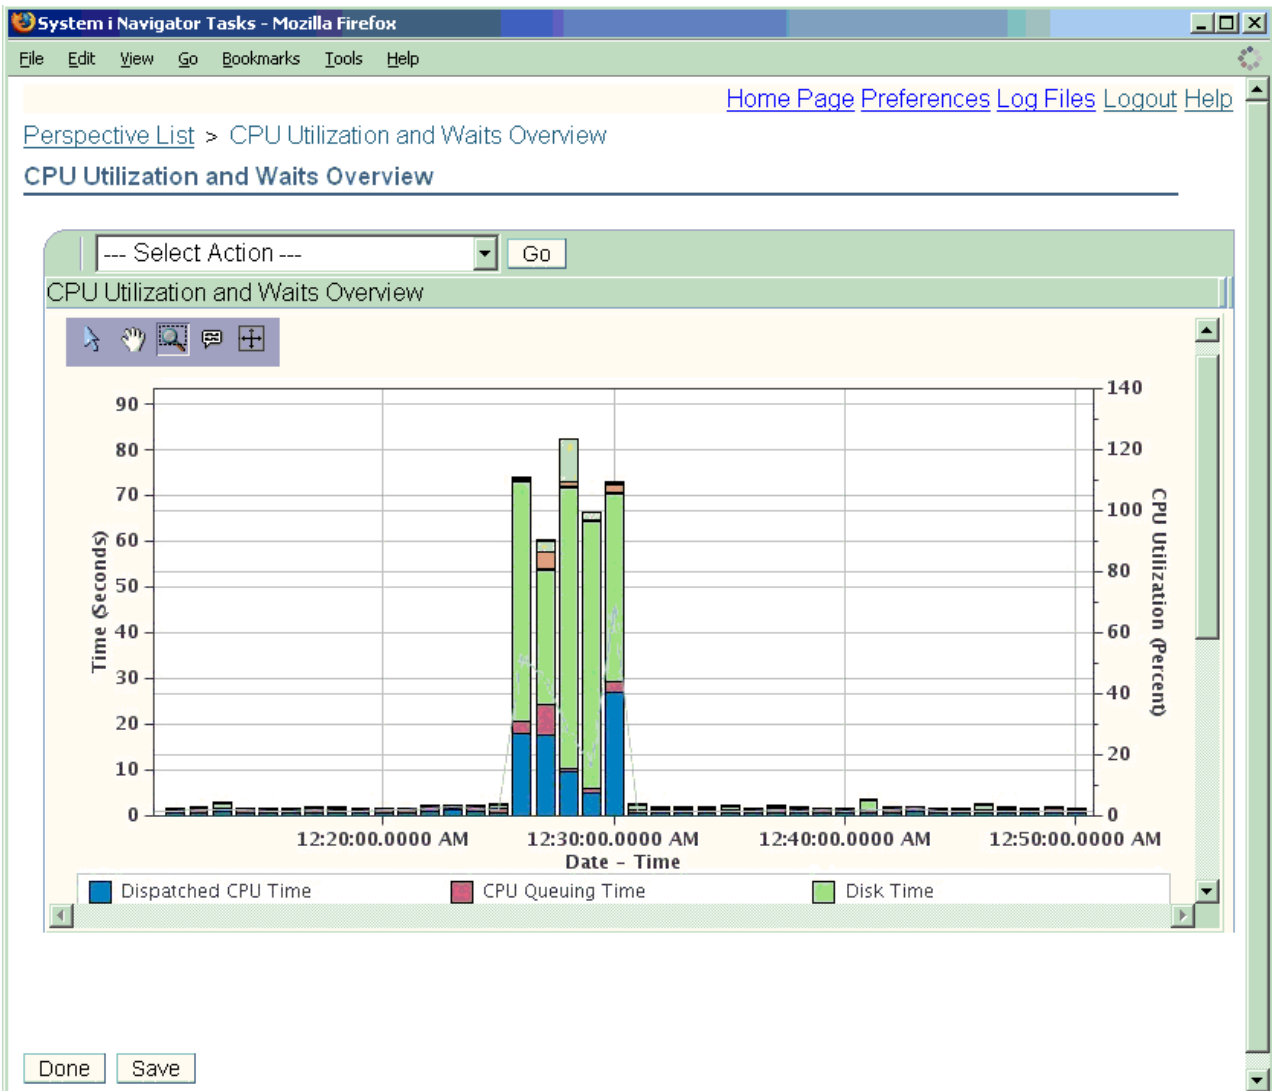
This toolbar button allows you to define an area of the chart and zoom into that area.



This toolbar button allows information to be displayed as you move the cursor over data on the chart.



This toolbar button allows you to zoom out to fit the screen.



You may need to install some or all of the following, depending upon the level of function that you need.

- 5761-SS1 Option 46 - Viewer function (required).
- 5761-PT1 Option 1 - IBM i5/OS Disk Watcher function.
- 5761-PT1 Option 3 - IBM i5/OS Job Watcher function.

### Copying a collection:

To copy a collection, follow these steps.

1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
2. Click **Collections**.
3. Select the collection that you want to copy.

- | 4. From the **Select Actions** menu, select **Copy**.

#### | **Deleting a collection:**

- | To delete a collection, follow these steps.
- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Collections**.
- | 3. Select the collection that you want to delete.
- | 4. From the **Select Actions** menu, select **Delete**.

#### | **Saving a collection:**

- | To save a collection, follow these steps.
- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Collections**.
- | 3. Select the collection that you want to save.
- | 4. From the **Select Actions** menu, select **Save**.

#### | **Restoring a collection:**

- | To restore a collection, follow these steps.
- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collections**.
- | 4. Click **Restore Collection**.

#### | **Converting a collection:**

- | To convert a collection that was collected in the previous release, follow these steps.
- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Collections**.
- | 3. Select the collection that you want to convert.
- | 4. From the **Select Actions** menu, select **Convert**.

#### | **Viewing collection properties:**

- | To create database files, follow these steps.
- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Collections**.
- | 3. Select the collection that you want to view the properties for.
- | 4. From the **Select Actions** menu, select **Properties**.

### | **Managing Collection Services**

- | Manage Collection Services using IBM Systems Director Navigator for i5/OS.

#### | **Starting Collection Services:**

- | Start Collection Services by doing the following.
- | To start Collection Services, follow these steps:
- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.

- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collectors** .
- | 4. Expand **Collection Services** .
- | 5. Click **Start Collection Services** .

#### | **Stopping Collection Services:**

| Stop Collection Services by doing the following.

| To stop Collection Services, follow these steps:

- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collectors** .
- | 4. Expand **Collection Services** .
- | 5. Click **Stop Collection Services** .

#### | **Configuring Collection Services:**

| Configure Collection Services by doing the following.

| To configure Collection Services, follow these steps:

- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collectors**.
- | 4. Expand **Collection Services**.
- | 5. Click **Configure Collection Services**. Specifying that summary data should be generated when the collection is cycled facilitates quicker processing of the performance database data by analysis tools.

#### | **Cycling Collection Services:**

| Cycle Collection Services by doing the following.

| To cycle Collection Services, follow these steps:

- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collectors** .
- | 4. Expand **Collection Services** .
- | 5. Click **Cycle Collection Services** .

#### | **Creating database files:**

| To create database files, follow these steps.

- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Collection Manager**.
- | 3. Select a Collection Services \*MGTCOL object based collection that you want to create performance data for.
- | 4. From the **Select Actions** menu, select **Create Database Files**.

#### | **Related tasks**

| “Creating database files from Collection Services data” on page 26

| Use this information to manually or automatically create database files from Collection Services data.



## | **Displaying Collection Services status:**

| Display Collection Services status by doing the following.

| To display Collection Services status, follow these steps:

- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collectors** .
- | 4. Expand **Collection Services** .
- | 5. Click **Collection Services Status** .

## | **Managing IBM i5/OS Disk Watcher**

| Manage IBM i5/OS Disk Watcher using IBM Systems Director Navigator for i5/OS.

### | **Related concepts**

| “IBM i5/OS Disk Watcher” on page 56

| IBM i5/OS Disk Watcher provides for the collection of disk performance data to diagnose disk related performance problems.

## | **Starting Disk Watcher:**

| Start Disk Watcher by doing the following.

| To start Disk Watcher, follow these steps:

- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collectors** .
- | 4. Expand **Disk Watcher** .
- | 5. Click **Start Disk Watcher** .

## | **Stopping Disk Watcher:**

| Stop Disk Watcher by doing the following.

| To stop Disk Watcher, follow these steps:

- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collectors** .
- | 4. Expand **Disk Watcher** .
- | 5. Click **Stop Disk Watcher** .

## | **Adding a Disk Watcher definition:**

| Add a Disk Watcher definition by doing the following.

| To add a Disk Watcher definition, follow these steps:

- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collectors** .
- | 4. Expand **Disk Watcher** .
- | 5. Click **Add Disk Watcher Definition** .



### | **Deleting a Disk Watcher definition:**

| Delete a Disk Watcher definition by doing the following.

| To delete a Disk Watcher definition, follow these steps:

- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collectors** .
- | 4. Expand **Disk Watcher** .
- | 5. Click **Disk Watcher Definitions** .
- | 6. Select the Disk Watcher definition that you want to delete.
- | 7. From the **Select Actions** menu, select **Delete**.

### | **Displaying the properties of a Disk Watcher definition:**

| Display the properties of a Disk Watcher definition by doing the following.

| To display the properties of a Disk Watcher definition, follow these steps:

- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collectors** .
- | 4. Expand **Disk Watcher** .
- | 5. Click **Disk Watcher Definitions** .
- | 6. Select the Disk Watcher definition that you want to display the properties for.
- | 7. From the **Select Actions** menu, select **Properties**.

### | **Managing IBM i5/OS Job Watcher**

| Manage IBM i5/OS Job Watcher using IBM Systems Director Navigator for i5/OS.

#### | **Related concepts**

- | “IBM i5/OS Job Watcher” on page 56
- | IBM i5/OS Job Watcher provides for the collection of job data for any or all jobs, threads, and tasks on the system. It provides call stacks, SQL statements, objects being waited on, Java JVM statistics, wait statistics and more which are used to diagnose job related performance problems.

### | **Starting Job Watcher:**

| Start Job Watcher by doing the following.

| To start Job Watcher, follow these steps:

- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collectors** .
- | 4. Expand **Job Watcher** .
- | 5. Click **Start Job Watcher** .

### | **Stopping Job Watcher:**

| Stop Job Watcher by doing the following.

| To stop Job Watcher, follow these steps:

- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.

- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collectors** .
- | 4. Expand **Job Watcher** .
- | 5. Click **Stop Job Watcher** .

#### | **Adding a Job Watcher definition:**

| Add a Job Watcher definition by doing the following.

| To add a Job Watcher definition, follow these steps:

- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collectors** .
- | 4. Expand **Job Watcher** .
- | 5. Click **Add Job Watcher Definition** .

#### | **Deleting a Job Watcher definition:**

| Delete a Job Watcher definition by doing the following.

| To delete a Job Watcher definition, follow these steps:

- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collectors** .
- | 4. Expand **Job Watcher** .
- | 5. Click **Job Watcher Definitions** .
- | 6. Select the Job Watcher definition that you want to delete.
- | 7. From the **Select Actions** menu, select **Delete**.

#### | **Displaying the properties of a Job Watcher definition:**

| Display the properties of a Job Watcher definition by doing the following.

| To display the properties of a Job Watcher definition, follow these steps:

- | 1. Select **Performance** from your IBM Systems Director Navigator for i5/OS window.
- | 2. Click **Show All Performance Tasks**.
- | 3. Expand **Collectors** .
- | 4. Expand **Job Watcher** .
- | 5. Click **Job Watcher Definitions** .
- | 6. Select the Job Watcher definition that you want to display the properties for.
- | 7. From the **Select Actions** menu, select **Properties**.

## **System i 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 business-to-business (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 System i 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 System i Navigator, or on your PC. In fact, you could even turn your PC off. System i 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 System i Navigator for Wireless.

System i 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.

##### **Related concepts**

“Collection Services” on page 24

Collection Services provides for the collection of system management data. It is the primary collector of system data.

##### **Related reference**

“Network performance” on page 19

The network design, hardware resources, and traffic pressure often have a significant effect on the performance of e-business applications. You can use this topic for information on how to optimize network performance, and tune server communication resources.

#### **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. Find out 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. The settings in the monitor properties override 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 starting 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.

### Related concepts

"Collection Services" on page 24

Collection Services provides for the collection of system management data. It is the primary collector of system data.

"System i Navigator graph history" on page 86

Graph history provides a graphical display of performance data collected by Collection Services over a specified period of time.

## Configuring a monitor

You can configure a monitor in System i Navigator. Use this information to learn how to set up a monitor and how to configure it to take the best advantage of the available options.

System monitors are highly interactive tools that you can use to gather and display real-time performance data from your endpoint systems. Creating a new monitor is a quick and easy process that begins at the New Monitor window:

1. In System i Navigator, expand Management Central, select **Monitors**, right-click **System**, and then select **New Monitor**.
2. 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.
3. 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.
4. 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.
5. 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.
6. 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.
7. Select your systems and groups. Use the New Monitor-Systems and Groups page to select the endpoint systems or system groups where you want to start a monitor.

After you have created your monitor, right-click the monitor name and select **Start** to run the monitor and begin working with monitor graphs.

## Monitor 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 configure 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.

Table 1.

| Metric groups                                        | Metric description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CPU Utilization                                      | <p>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"><li>• CPU Utilization (Average)</li><li>• CPU Utilization (Interactive Jobs)</li><li>• CPU Utilization (Interactive Feature)</li><li>• CPU Utilization (Database Capability)</li><li>• CPU Utilization (Secondary Workloads)</li><li>• CPU Utilization Basic (Average)</li></ul> <p>To learn more about these metrics and how to use them, see the online help available on the <b>General</b> tab of the New Monitor window or the Monitor Properties window in System i Navigator.</p> |
| Interactive Response Time (Average and Maximum)      | The response time that interactive jobs experience on your system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Transaction Rate (Average)                           | The number of transactions per second completed by all jobs on your system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Transaction Rate (Interactive)                       | <p>The number of transactions per second completed on your system by the following types of jobs:</p> <ul style="list-style-type: none"><li>• Interactive</li><li>• Multiple requester terminal (MRT)</li><li>• System/36™ environment interactive</li><li>• Pass-through</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

Table 1. (continued)

| Metric groups                                         | Metric description                                                                                               |
|-------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| 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. After you become familiar with the Management Central metrics, which metrics you select 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.

### System and job monitor interaction with Collection Services

Collection Services is both a valuable tool for performance analysis as a stand-alone application and as a utility used by other applications for the collection of performance data.

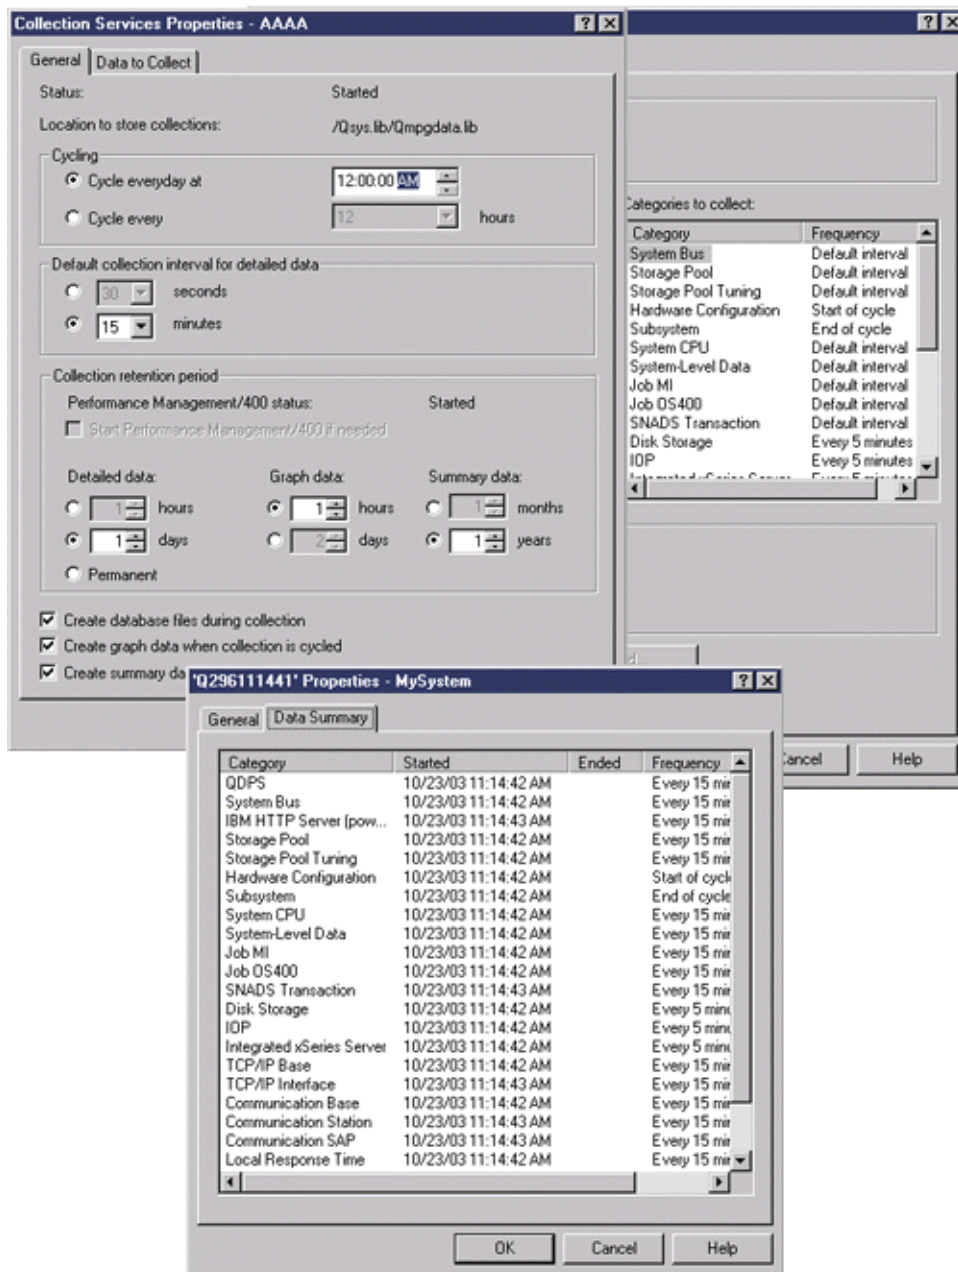
- | Sometimes, performance analysis causes confusion when you are trying to determine which application is
- | responsible for activity you may see on your system. One easy rule to remember for this issue is that
- | even if it looks like those other applications are busy, there is one and only one Collection Services data
- | collection occurring on the system at any given time.

The following scenarios explain the different combinations between system monitors and job monitors and Collection Services and what Collection Services displays.

### Collection Services is collecting data using the default values

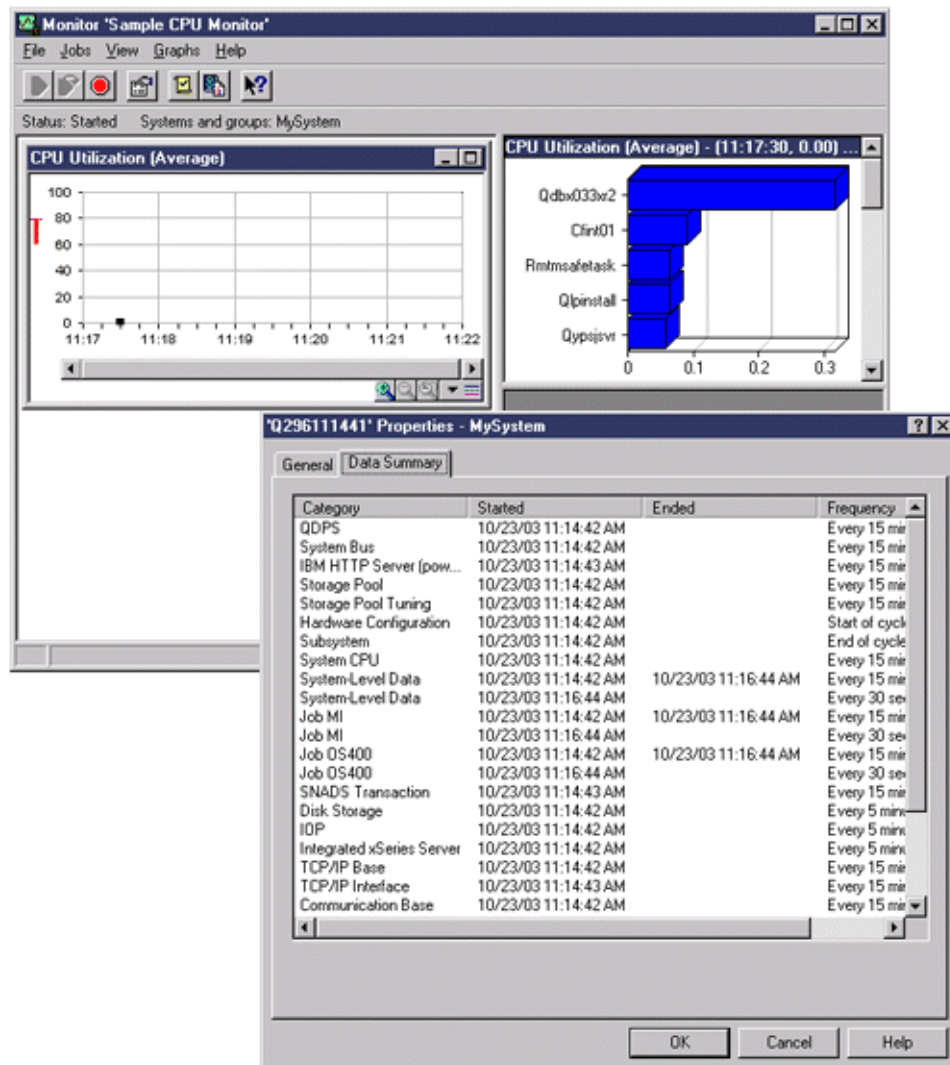
In this scenario, there are no system monitors or job monitors active on the system. When viewing the Collection Services properties page and the \*MGTCOL object properties view, you see something similar to the following.





## Both Collection Services and a system monitor are started

This scenario shows that Collection Services had already started at some point, and later someone started a system monitor to collect CPU Utilization (Average) metric data at 30-second intervals. Notice in the \*MGTCOL object properties view that the collection interval for System Level Data, Job MI Data, and Job OS Data categories changed from 15 minutes to 30 seconds. This demonstrates that the same \*MGTCOL object is being used, and only those categories necessary to calculate information for a given metric were changed to collect at the new interval.



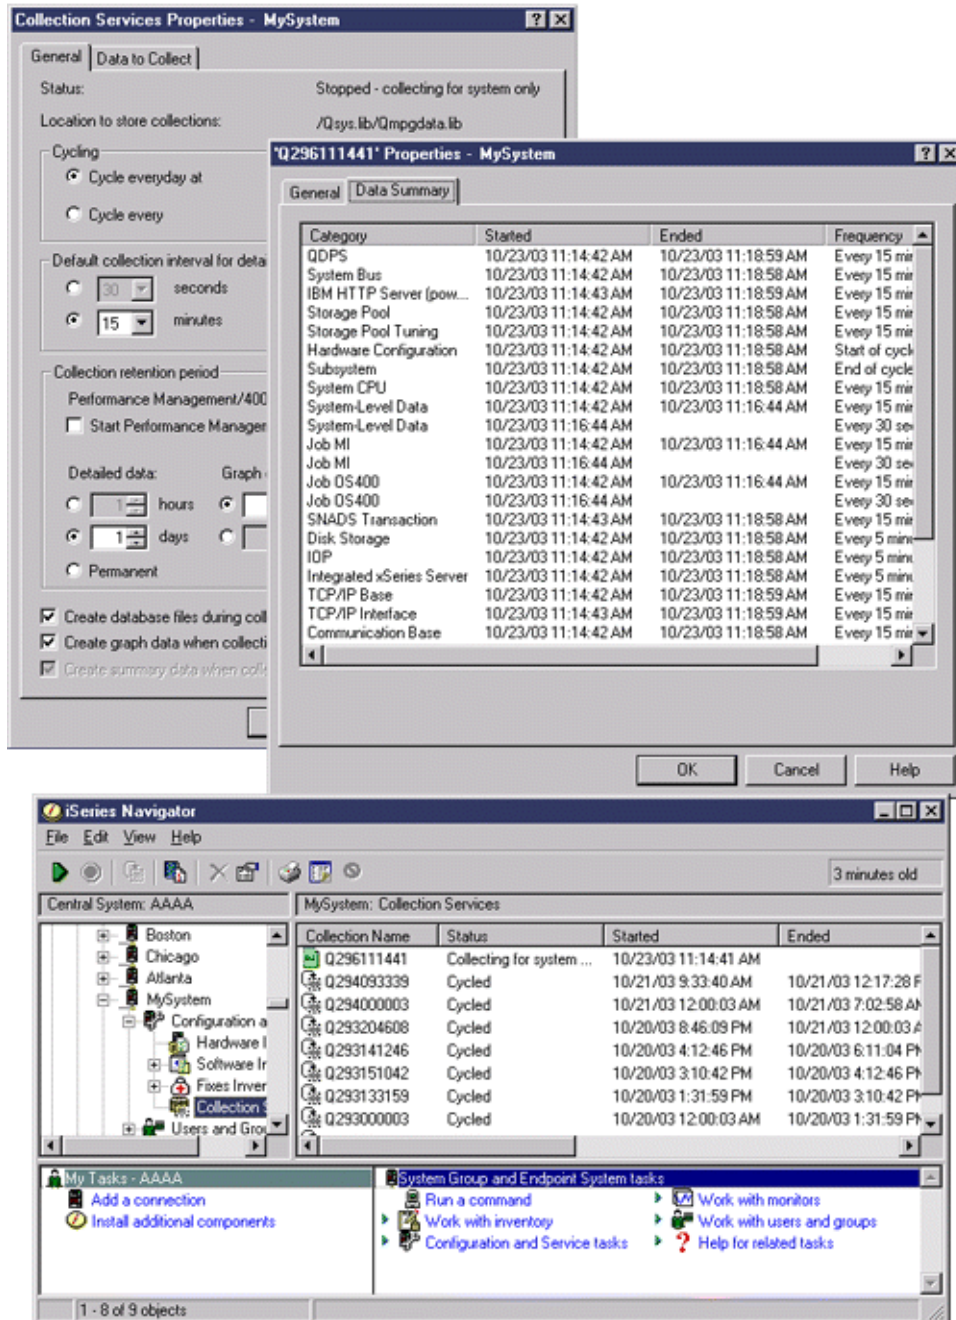
## Collection Services stopped and system monitor remains started

In this scenario, Collection Services was stopped and the system monitor remains started and continues to collect data necessary to calculate the graph metrics.

Observe the following:

- The Collection Services properties page shows a status of **System collection stopped. Collecting for system only**.
- The \*MGTCOL object properties page shows that data collection has ended for all categories except for those necessary to calculate the graph metric data.
- The Collection Services list view shows the \*MGTCOL object with a status of **Collecting for system....** This might be confusing; therefore, to get the status of Collection Services, look at the Collection Services Properties page.





## Scenarios: System i Navigator monitors

Use this information to see 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 System i Navigator provide a powerful set of tools for researching and managing system performance. For an overview of the types of monitors provided by System i Navigator, see [System i 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 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 System i 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:  
`SNDRMSG 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 system, 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 System i 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 system, 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 193.

1. In System i 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 193.

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

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

DCL &MSGTXT *CHAR 132

RCVMMSG 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.

### Related concepts



Work with notification

See the Work with notification topic for information about how to use the notification function of Advanced Job Scheduler.

### Scenario: Message monitor:

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

### Situation

Your company has several systems, 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 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 System i 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 systems 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.

## System i Navigator for Wireless

System i Navigator for Wireless allows you to monitor performance data over a wireless connection, using a personal digital assistant (PDA), Internet-ready telephone, or traditional Web browser. The System i Navigator for Wireless uses the performance data collected by Collection Services.

With your wireless device, you can:

- Run commands across multiple systems
- Start and view system, job, and message monitors
- Work with jobs and messages from the monitors (hold, release, end, reply, get details)
- Manage Integrated xSeries Server

For an overview of how System i Navigator for Wireless can help you get started with remote monitoring, see the topic [System i Navigator for Wireless](#).

For complete and up-to-date information about remote monitoring, see the [System i Navigator for Wireless](#) home page.

#### **Related information**

[System i Navigator for Wireless](#) topic

## **System i Navigator graph history**

Graph history provides a graphical display of performance data collected by Collection Services over a specified period of time.

Graph history provides a graphical view of performance data collected over days, weeks, months, or years with Collection Services. You do not need to have a system monitor running to view performance data. As long as you use Collection Services to collect data, you can view the Graph History window.

**Note:** For more information about monitoring system performance, see the [Track performance data](#) topic.

#### **Related concepts**

“Collection Services” on page 24

Collection Services provides for the collection of system management data. It is the primary collector of system data.

“Tracking performance” on page 6

Tracking your system performance over time allows you to plan for your system’s growth and ensures that you have data to help isolate and identify the cause of performance problems. Learn which applications to use and how to routinely collect performance data.

### **| 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 System i Navigator to activate Performance Management Agent (PM Agent) over multiple systems. When you activate PM Agent, 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 Agent enabled, the graph data field supports 1 to 7 days. With PM Agent 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 Agent, you can specify one to seven days. If you do start PM Agent, you can specify 1 to 30 days. The default is one hour.

- **Graph history 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 Agent 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.
- **Related concepts**  
“Collection Services” on page 24  
Collection Services provides for the collection of system management data. It is the primary collector of system data.
- **Related tasks**  
“Activating PM Agent” on page 89  
PM Agent is a part of the operating system and you must activate it to use its collecting capabilities.

## Viewing graph history

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

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

1. Follow the System i 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 System i 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.

**Related concepts**



“Collection Services” on page 24

Collection Services provides for the collection of system management data. It is the primary collector of system data.

## | **IBM Performance Management for System i5**

| The IBM Performance Management for System i5 (PM for System i5) offering automates the collection, archival, and analysis of system performance data and returns reports to help you manage system resources and capacity.

| The PM for System i5 offering includes the Performance Management Agent (PM Agent). The PM Agent (formerly known as PM/400®) is a function of the operating system that provides automated collection of nonproprietary Collection Services data, reduces the data, and sends the data to IBM. All collection sites are network secure, and the time of the transfer is completely under your control. When you send your data to IBM, you eliminate the need to store all the trending data yourself. IBM stores the data for you and provides you with a series of reports and graphs that show your server’s growth and performance. You can access your reports electronically using a traditional browser.

| This offering, when used with the IBM Systems Workload Estimator, allows you to better understand how your business trends relate to the timing of required hardware upgrades, such as central processing unit (CPU) or disk. The IBM Systems Workload Estimator can size a systems consolidation or evaluate upgrading a system with logical partitions, by having PM Agent send the data for multiple systems or partitions to the IBM Systems Workload Estimator.


| PM Agent uses less than 1 percent of your CPU. It uses approximately 58 MB of disk space, which depends on your hardware model and the size of your collection intervals.

### | **Related concepts**

| “Collection Services” on page 24

| Collection Services provides for the collection of system management data. It is the primary collector of system data.

### | **Related information**

|  PM for System i Web site

| See the PM for System i Web site for more information about PM for System i.

## | **PM Agent concepts**

| Learn about the functions and benefits PM Agent can provide and about important implementation considerations.

| PM Agent uses Collection Services to gather the nonproprietary performance and capacity data from your server and then sends the data to IBM. This information can include CPU utilization and disk capacity, response time, throughput, application and user usage. When you send your data to IBM, you eliminate the need to store all the trending data yourself. IBM stores the data for you and provides you with a series of reports and graphs that show your server’s growth and performance. You can access your reports electronically using a traditional browser.

| The most important requirement for establishing an accurate trend of the system utilization, workload, and performance measurements is consistency. Ideally, performance data should be collected 24 hours per day. Because of the relationship between PM Agent and Collection Services, you need to be aware of the implications that can occur when you are using PM Agent.

| Here are some guidelines to help you define your collections when you are using PM Agent:

- | • **Collect data continuously with Collection Services.**

PM Agent satisfies this requirement by collecting data 24 hours a day with Collection Services. PM Agent collects performance data at 15-minute intervals. PM Agent uses the 15-minute interval default, but does not change what the interval is set to. A 15-minute interval is the recommended interval.

- **Select the Standard plus protocol profile.**

Standard plus protocol is the default value for the collection profile. The collection profile indicates what data is collected. The data categories in the standard plus protocol profile correspond to the \*ALL value for the DATA parameter on the Start Performance Monitor (STRPFRMON) command. The collection does not cycle (unless required to do so for other reasons). This action is done to gather enough information for PM Agent reports.

- **Avoid making interim changes to collection parameters when PM Agent is active.**

For example, when you activate PM Agent, the **Create database files during collection** field is checked as the default value. If you change this, PM Agent changes it back to the default value on the hour. The change takes effect immediately. The collection does not cycle (unless required to do so for other reasons).

**Related reference**

“Collection Services collection profiles” on page 28

Descriptions of the Collection Services collection profiles. The collection profiles define what is collected.

## **Configuring PM Agent**

To start using PM Agent, you need to activate it, set up a transmission method, and customize the data collection and storage.

PM Agent automates the collection of performance data through Collection Services. You can specify which library to put the data in as long as the library resides on the base auxiliary storage pool (ASP). The library should not be moved to an independent auxiliary storage pool because an independent auxiliary storage pool can be varied off, which stops the PM Agent collection process. PM Agent creates the library during activation if the library does not already exist.

To begin using PM Agent, you need to perform the following tasks:

### **Activating PM Agent:**

PM Agent is a part of the operating system and you must activate it to use its collecting capabilities.

You must start PM Agent to take advantage of its data collecting capabilities. You can start PM Agent by using any one of the following methods:

#### **Use System i Navigator**

Use System i Navigator to activate PM Agent over multiple systems. When you activate PM Agent, you can view data that was collected days ago, weeks ago, or months ago. You go beyond the real-time monitor capabilities. Without PM Agent enabled, you can view up to the last 7 days of data. With PM Agent enabled, you choose the length of time to retain the data.

To start PM Agent from System i Navigator, do the following:

1. In System i Navigator, expand the system where you want to start PM Agent.
2. Expand **Configuration and Service**.
3. Right-click **Collection Services**.
4. Select **PM Agent**.
5. Select **Start**.
6. Select the systems on which you want to start PM Agent.
7. Click **OK**.

## | Reply to message CPAB02A in the QSYSOPR message queue

| When the QSYSWRK subsystem starts, this message asks if you want to activate PM Agent.

- | 1. From the character-based interface, reply with a G to the message in QSYSOPR "Do you want to activate PM Agent? (I G C)." QSYSOPR message queue receives the message that PM Agent is activated.
- | 2. Update your contact information. Issue the **GO PM400** command and specify option 1.

## | Issue the Configure PM Agent (CFGPMAGT) command

| From the character-based interface, you can issue the Configure PM Agent (CFGPMAGT) command.

| You can proceed to the next step in the setup process, which is to determine which transmission method to use to send data to IBM.

### | **Related concepts**

| "System i Navigator graph history" on page 86

| Graph history provides a graphical display of performance data collected by Collection Services over a specified period of time.

| "Determining which PM Agent transmission method to use"

| Determine how you want to send data. You can either gather your data with the Management Central inventory function and send the data with Electronic Service Agent (Extreme Support) or you can have PM Agent gather the data and send it over the SNA protocol.

### | **Related tasks**

| "Deactivating PM Agent" on page 101

| Learn how you can stop PM Agent.

## | Determining which PM Agent transmission method to use:

| Determine how you want to send data. You can either gather your data with the Management Central inventory function and send the data with Electronic Service Agent (Extreme Support) or you can have PM Agent gather the data and send it over the SNA protocol.

| The PM Agent transmission process takes advantage of the network configuration that you perform with Management Central to set up a central system and endpoint systems. However, you can use the character-based interface to configure PM Agent. Choose which transmission method you want to use:

- | • Send data with the Electronic Service Agent over Extreme Support
  - | If you choose this transmission method, you need to configure PM Agent to have data gathered by the Management Central inventory function.
  - | IBM recommends you use ESA over Extreme Support.
- | • Send data with SNA protocol
  - | If you choose this transmission method, you need to configure PM Agent by using the character-based interface. PM Agent gathers your data and transmits it by using SNA.

| Once you have implemented which transmission method you want to use, you are ready to do the other tasks to manage PM Agent.

### | **Related concepts**

| Management Central

### | **Related reference**

| "Managing PM Agent" on page 101

| Now that you have set up your network, you can perform a variety of tasks with PM Agent.

| *Sending PM Agent data with Service Agent over Extreme Support (Universal Connection):*

| PM Agent uses Collection Services to gather the nonproprietary performance and capacity data from your server. After you have collected this data, you can use Electronic Service Agent over Extreme Support to send the data to IBM.

| Here are the steps to follow to configure for PM Agent:

| 1. Activate PM Agent.

| You must start PM Agent to take advantage of its data collecting capabilities.

| 2. Set up your Management Central network.

| Define which server is your central system and which servers are your endpoint systems. You can use this network hierarchy to send your data from your endpoint systems to a central location before sending the data to IBM.

| 3. Connect to IBM to transmit your data with the Universal Connection.

| This is the connection that Management Central will use to transmit the PM Agent data to IBM.

| 4. Gather PM Agent performance data.

| Use the Management Central inventory function to gather your data.

| 5. Send your data to IBM.

| Use the Electronic Service Agent, which is available under Extreme Support in the Management Central hierarchy, to send your data to IBM. The Electronic Service Agent uses the Universal Connection.

| You can also send data with the SNA protocol.

| Once you have configured PM Agent, you are ready to do the other tasks to manage PM Agent.

| **Related concepts**

| Universal Connection

| **Related tasks**

| Electronic Service Agent over Extreme Support

| "Activating PM Agent" on page 89

| PM Agent is a part of the operating system and you must activate it to use its collecting capabilities.

| Set up your Management Central network

| "Gathering PM Agent performance data"

| You can use Management Central to gather your PM Agent performance data.

| "Sending data with SNA protocol" on page 92

| If you choose not to take advantage of sending data with the Electronic Service Agent over Extreme Support, you can still use the character-based interface to transmit data.

| **Related reference**

| "Managing PM Agent" on page 101

| Now that you have set up your network, you can perform a variety of tasks with PM Agent.

| *Gathering PM Agent performance data:*

| You can use Management Central to gather your PM Agent performance data.

| Perform the following tasks:

| 1. Activate PM Agent.

| 2. Configure the Universal Connection.

| 3. Set up your Management Central network.

| 4. Verify that the Electronic Service Agent is installed on your system or accessible from your system.

| To gather PM Agent performance data on an endpoint system or system group, follow these steps:

1. In System i Navigator, expand **Management Central**.
2. Expand **Endpoint Systems** or **System Groups**.
3. Right-click an endpoint system or a system group, and select **Inventory**.
4. Select **Collect**.
5. Select one or more inventories to collect. In this case, you would select **PM Agent performance data**.
6. If you want an action to run on the central system when collection completes, select the action from the list.
7. Click **OK** to start collecting the data immediately or click **Schedule** to specify when to collect the data.

Once you have configured your servers, you are ready to do the other tasks to manage PM Agent.

#### **Related tasks**

“Activating PM Agent” on page 89

PM Agent is a part of the operating system and you must activate it to use its collecting capabilities.

Configured the Universal Connection

Set up your Management Central network

Verified that the Electronic Service Agent is installed on your system or accessible from your system

#### **Related reference**

“Managing PM Agent” on page 101

Now that you have set up your network, you can perform a variety of tasks with PM Agent.

*Sending data with SNA protocol:*

If you choose not to take advantage of sending data with the Electronic Service Agent over Extreme Support, you can still use the character-based interface to transmit data.

PM Agent asks you a series of questions about the configuration and use of your servers. The Configure PM Agent display asks you questions about how you want your servers to send and receive PM Agent performance data. The first part of the process involves setting up your network. The second part asks you how you want to transmit your data. When you use the character-based interface, you can use a direct dial line to transmit data.

Follow these steps to send data with SNA:

1. Activate PM Agent
 

You must start PM Agent to take advantage of its data collecting capabilities.
2. Select which network configuration you want to use.
 

Determine which network configuration you will use to transmit data. Choose how you connect to IBM by using a direct dial line, an existing Internet Service Provider (ISP), or a virtual private network (VPN). If you want to use ISP or VPN, you must configure a Universal Connection.

If you decide to use the direct dial line to report data to IBM, you have several choices as to how you configure your network. Select which configuration is appropriate for your network, and perform the steps outlined for that particular configuration from the Configure PM Agent display:

  - As a single server that sends its data directly to IBM.
  - As a host server, which means that you want your server to receive performance data from other servers (remote servers) and then forward the data to IBM. The host server cannot be at a release level that is earlier than other servers. In other words, the host server must be at the same release level or later than other servers.
  - As a remote server, which means that you can send performance data to a host server. You identify on the Configure PM Agent display that you need a remote server, and then use option 5 (Work with remote iSeries systems) from the PM Agent menu to define your remote servers.
3. Work with remote servers.

If you choose to set up your network for a host server, you need to identify those servers that will send their data to your host server. You can ignore this step if you are using a single server or a remote server.

#### 4. Customize PM Agent.

After you have configured your network, you need to establish the global parameters for the operation of the PM Agent software. You need to define the PM Agent data telephone number if you would like to connect to IBM with a direct dial line.

Once you have configured your servers, you are ready to do the other tasks for managing PM Agent.

#### **Related concepts**

Virtual private network (VPN)

#### **Related tasks**

“Sending PM Agent data with Service Agent over Extreme Support (Universal Connection)” on page 90

PM Agent uses Collection Services to gather the nonproprietary performance and capacity data from your server. After you have collected this data, you can use Electronic Service Agent over Extreme Support to send the data to IBM.

“Activating PM Agent” on page 89

PM Agent is a part of the operating system and you must activate it to use its collecting capabilities.

“Setting a direct dial line for PM Agent” on page 100

For most locations, PM Agent tries to select the correct data telephone number for your location.

Configure a Universal Connection

“Configuring PM Agent network for a single server”

A single server sends its data directly to IBM.

“Configuring PM Agent network for a host server” on page 94

A host server receives performance data from other servers and then forwards the data to IBM .

“Configuring PM Agent network for a remote server” on page 95

A remote server sends its performance data to a host server.

“Working with remote servers” on page 95

In some sites, a host server in a network sends the required performance data to IBM for processing.

“Customizing PM Agent” on page 98

Now that you have set up your network, you may need to customize PM Agent to fit your needs.

#### **Related reference**

“Managing PM Agent” on page 101

Now that you have set up your network, you can perform a variety of tasks with PM Agent.

*Configuring PM Agent network for a single server:*

A single server sends its data directly to IBM.

Here are the steps that you need to follow to configure PM Agent for a single server only if PM Agent gathers data and transmits data over SNA. From the Configure PM Agent (CFGPMAGT) display on your server:

1. Type **CFGPMAGT** from the command line.
2. Specify **2 = Send data with SNA protocol** for the **Select connectivity option for sending performance data to IBM** field.
3. Specify **0 = No** for the **Receive performance data** field.
4. If you specify **2 = Send data with SNA protocol** for the **Select connectivity option for sending performance data to IBM**, you see additional information that indicates whether the appropriate communications objects exist. If the objects do not exist, PM Agent creates the communications objects for you for transmission. Respond appropriately to the additional displays.



5. Type your company's contact information on the Work with Contact Information display making sure you fill out all the mandatory fields, otherwise the configuration will be incomplete.

If you decide that the single server setup is not what you want, you can choose another SNA configuration option.

Once you have configured your servers, you are ready to do the other tasks to manage PM Agent.

**Related tasks**

"Sending data with SNA protocol" on page 92

If you choose not to take advantage of sending data with the Electronic Service Agent over Extreme Support, you can still use the character-based interface to transmit data.

**Related reference**

"Managing PM Agent" on page 101

Now that you have set up your network, you can perform a variety of tasks with PM Agent.

*Configuring PM Agent network for a host server:*

A host server receives performance data from other servers and then forwards the data to IBM .

Here are the steps that you need to follow to configure PM Agent for a host server only if PM Agent gathers data and transmits data over SNA:

1. From the Configure PM Agent display on your host server

- Type **CFGPMAGT** from the command line.
- Specify **2 = Send data with SNA protocol** for the **Select connectivity option for sending performance data to IBM** field.
- Specify **1 = Yes** for the **Receive performance data** field.

2. From the Work with Remote System i systems display on your host server

- Press F6 (Create) to identify which servers will send their data to your host server.
- Complete the fields and press Enter.

**Note:** The following situation occurs only if PM Agent gathers data and transmits data over SNA. If you have a network of systems, it is recommended that you use the Universal Connection and Management Central in System i Navigator to gather and transmit the data for those systems.

PM Agent automatically schedules the transmission of data from the primary server to IBM the day after data is received from a remote server. If the automatic scheduling does not fit your work management scheme, you can manually schedule the transmission of the data from the primary server.

Here is a tip that you should keep in mind when scheduling the transmission of your data. Throughout the week, evenly schedule the transmission of data to the primary server. This action minimizes the performance impact on the primary server. For example, in a network of twelve servers, you might have three groups of four systems. You can schedule each group to send their data on Monday, Wednesday, and Friday. This evenly distributes the amount of data that is sent to the primary server.

If you decide that the host server setup is not what you want, you can choose another SNA configuration option.

Once you have configured your servers, you are ready to do the other tasks to manage PM Agent.

**Related tasks**

"Sending data with SNA protocol" on page 92

If you choose not to take advantage of sending data with the Electronic Service Agent over Extreme Support, you can still use the character-based interface to transmit data.

**Related reference**

- | "Managing PM Agent" on page 101
- | Now that you have set up your network, you can perform a variety of tasks with PM Agent.
- | *Configuring PM Agent network for a remote server:*
- | A remote server sends its performance data to a host server.
- | Here are the steps that you need to follow to configure PM Agent for a remote server only if PM Agent gathers data and transmits data over SNA. From the Configure PM Agent display (CFGPMAGT) on your remote server, do these steps:
- | 1. Type **CFGPMAGT** from the command line.
- | 2. Specify **2 = Do not send performance data** for the **Select connectivity option for sending performance data to IBM** field.
- | 3. Specify **0 = No** for the **Receive performance data** field.
- | **Note:** If you have a network of systems, it is recommended that you use the inventory function of System i Navigator to gather your data and then transmit the data for those systems over the Universal Connection.
- | If you decide that the remote server setup is not what you want, you can choose another SNA configuration option.
- | Once you have configured your servers, you are ready to do the other tasks to manage PM Agent.
- | **Related tasks**
- | "Sending data with SNA protocol" on page 92
- | If you choose not to take advantage of sending data with the Electronic Service Agent over Extreme Support, you can still use the character-based interface to transmit data.
- | **Related reference**
- | "Managing PM Agent" on page 101
- | Now that you have set up your network, you can perform a variety of tasks with PM Agent.
- | *Working with remote servers:*
- | In some sites, a host server in a network sends the required performance data to IBM for processing.
- | When you use a host server network, you have the other servers in the network send their performance data to this host server for transmission to IBM. To set up your network to use a host server, you must identify the other remote servers and set the schedule for their data transmission. The Work with Remote System i systems display enables you to define these other servers.
- | **Notes:**
- | 1. You do not have to use this display if you are setting up your network as a remote server or as a single server. You perform this task only if PM Agent gathers data and transmits data over SNA.
- | 2. If you have a network of systems, it is recommended that you use the inventory function of System i Navigator to gather your data and then transmit the data for those systems over the Universal Connection.
- | Follow these steps to define your remote servers:
- | 1. Type **GO PM400** from the command line.
- | 2. Type a 5 (Work with remote System i systems) from the PM Agent Menu and press Enter. You do not see a remote server displayed initially. You must create a new remote location.
- | 3. Create a new remote location by pressing F6 (Create).



4. Record the values for the following information. Use the Display Network Attributes (DSPNETA) command to display these values from the remote system.
  - Local network ID
  - Default local location
 The Work with Remote System i systems display shows a list of remote servers. This list includes the status for the servers (active or inactive) and the descriptions for each server.
5. Create or change the description for a remote site server by using the PM Agent Remote Site Maintenance display or the Change Remote Site System i display. The remote location name must be unique between remote servers.

PM Agent automatically schedules the transmission of data from the primary server to IBM the day after data is received from a remote server. If the automatic scheduling does not fit your work management scheme, you can manually schedule the transmission of the data from the primary server. To manually schedule the transmission of data, see the PM Agent scheduler.

The PM Agent software assumes that you defined the Advanced Peer-to-Peer Networking® (APPN) link between the server that receives data (the host server) and the server that sends data (the remote server). If your system has the system value QCRTAUT (Create default public authority) set to \*EXCLUDE or \*USE, you should see Create a device description for a remote server for information on how to define your controller descriptions. If your network does not meet these assumptions, see Non-APPN network considerations for information about creating device pairs to support the connection to each remote server.

Once you have defined your remote servers, you are ready to customize PM Agent to use a specific line connection.

- Related tasks**
- “Scheduling jobs with PM Agent” on page 102  
Learn how to schedule jobs with PM Agent.
  - “Creating a device description for PM Agent” on page 97  
You can create a device description for PM Agent.
  - “Working with remote servers in an APPC network”  
The primary server receives PM Agent data from other servers and then sends the data to IBM. The remote server sends PM Agent data to the primary server.
  - “Customizing PM Agent” on page 98  
Now that you have set up your network, you may need to customize PM Agent to fit your needs.

*Working with remote servers in an APPC network:*

The primary server receives PM Agent data from other servers and then sends the data to IBM. The remote server sends PM Agent data to the primary server.

The following information assumes that the controllers that are referred to have previously been defined.

You need to create device pairs to support the connection to each remote server only if PM Agent gathers data and transmits data over SNA.

1. Use the Create Device Description (APPC) (CRTDEVAPPC) command. On the remote server, type CRTDEVAPPC. Press F4 to prompt for the parameters, and define the values with the following information:

|                               |                                                                                         |
|-------------------------------|-----------------------------------------------------------------------------------------|
| <i>Table 2. Remote system</i> |                                                                                         |
| DEVD(Q1PLOC)                  | Specifies the name of the device description.                                           |
| RMTLOCNAME(Q1PLOC)            | Specifies the name of the remote location.                                              |
| ONLINE(*YES)                  | Specifies whether this device is varied online when the system is started or restarted. |

*Table 2. Remote system (continued)*

|                      |                                                                                                                                         |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| LCLLOCNAME(Q1PRMxxx) | Specifies the local location name. Q1PRMxxx matches the RMTLOCNAME of the primary server, where xxx is unique for each remote location. |
| CTL(yyyyyyy)         | Specifies the name of the attached controller, where yyyyyyy is a controller that attaches to the primary server.                       |
| MODE(Q1PMOD)         | Specifies the mode name.                                                                                                                |
| APPN(*NO)            | Specifies if the device is APPN-capable.                                                                                                |

- Specify the following information on the primary server. At the command line, type CRTDEVAPPC. Press F4 to prompt for the parameters, and define the values with the following information:

*Table 3. Primary server*

|                      |                                                                                                                                                                        |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DEV(D(Q1PRMxxx)      | Specifies the name of the device description. The name that is used here matches the device description name for the remote system.                                    |
| RMTLOCNAME(Q1PRMxxx) | Specifies the name of the remote location. The name that is used here matches the LCLLOCNAME value of the remote server, where xxx is unique for each remote location. |
| ONLINE(*YES)         | Specifies whether this device is varied online when the system is started or restarted.                                                                                |
| LCLLOCNAME(Q1PLOC)   | Specifies the local location name. This value matches the RMTLOCNAME of the remote server.                                                                             |
| CTL(aaaaaa)          | Specifies the name of the attached controller, where aaaaaa is a controller that attaches to the remote server.                                                        |
| MODE(Q1PMOD)         | Specifies the mode name.                                                                                                                                               |
| APPN(*NO)            | Specifies if device is APPN-capable.                                                                                                                                   |

- Vary on the devices (Vary Configuration (VRYCFG) command) after you define the APPC devices. On the remote server, type VRYCFG. Press F4 to prompt for the parameters.

*Table 4. Vary on remote system*

|                |                                             |
|----------------|---------------------------------------------|
| CFGOBJ(Q1PLOC) | Specifies the configuration object.         |
| CFGTYPE(*DEV)  | Specifies the type of configuration object. |
| STATUS(*ON)    | Specifies the status                        |

- Type option 5 on the PM Agent Menu to add Q1PRMxxx as a remote server. See Work with remote servers for instructions on how to add a remote server.

Now that you have finished configuring PM Agent, see Manage PM Agent for other tasks that you can perform with PM Agent.

#### **Related tasks**

“Working with remote servers” on page 95

In some sites, a host server in a network sends the required performance data to IBM for processing.

#### **Related reference**

“Managing PM Agent” on page 101

Now that you have set up your network, you can perform a variety of tasks with PM Agent.

#### *Creating a device description for PM Agent:*

You can create a device description for PM Agent.

The following steps are necessary on each remote server that has the Create default public authority (QCRTAUT) system value set to \*EXCLUDE or \*USE. If QUSER does not have \*CHANGE authority to device description Q1PLOC, remote transmissions will fail. These steps ensure that the device will not be created or deleted automatically.

**Note:** This task is necessary only if PM Agent gathers data and transmits data over SNA.

- | If you allow the device to be created automatically, the device description is created with PUBLIC \*EXCLUDE or \*USE authority, depending on the value set for QCRTAUT. Whether a device can be created or deleted automatically is controlled by the controller.
- | For systems that are not configured to use APPN, see Work with remote servers in a non-APPN environment for information on how to create the device description.
- | The following information assumes that the controller that will be used to communicate with the host server was defined previously on the remote server.

| On the *remote server*, re-create device description Q1PLOC:

```
| VRYCFG  CFGOBJ(Q1PLOC)
|         CFGTYPE(*DEV)
|         STATUS(*OFF)
| DLTDEVD  DEVD(Q1PLOC)
|
| CRTDEVAPPC DEVD(Q1PLOC)
|           RMTLOCNAME(Q1PLOC)
|           ONLINE(*NO)
|           LCLLOCNAME(name of remote system)
|           RMTNETID(remote netid of primary (or central) system)
|           CTL(name of controller that the device will be attached to)
|           AUT(*EXCLUDE)
|
| CRTOBJAUT OBJ(Q1PLOC)
|           OBJTYPE(*DEVD)
|           USER(QUSER)
|           AUT(*CHANGE)
|
| VRYCFG  CFGOBJ(Q1PLOC)
|         CFGTYPE(*DEV)
|         STATUS(*ON)
```

#### | **Related tasks**

- | “Working with remote servers in an APPC network” on page 96
- | The primary server receives PM Agent data from other servers and then sends the data to IBM. The remote server sends PM Agent data to the primary server.

#### | **Related reference**

- | Create Controller Description (APPC) (CRTCTLAPPC) command
- | Change Controller Description (APPC) (CHGCTLAPPC) command
- | Display Controller Description (DSPCTLD) command
- | “Managing PM Agent” on page 101
- | Now that you have set up your network, you can perform a variety of tasks with PM Agent.

#### | **Customizing PM Agent:**

- | Now that you have set up your network, you may need to customize PM Agent to fit your needs.
- | The Work with PM Agent Customization display provides you with the ability to:

#### | **Establish global parameters for the operation of PM Agent software**

- | The global parameters allow you to customize the following items. See the online help for a description of these fields:

- | • Priority limits
- | • Trend and shift schedules

#### | **Define your PM Agent data telephone number**

| Outside the United States and Canada, you must provide PM Agent with the telephone number of the IBM location that will receive your data. For most locations, PM Agent tries to select the correct telephone data number for your location when you initiate the configure PM Agent process.

#### | **Vary a line off and on with PM Agent**

| The PM Agent Line Control display allows PM Agent to vary the line off, transmit the PM Agent data, and then put the line back in the connect pending state.

| To customize the global parameters, do the following steps:

- | 1. Type **GO PM400** from the command line.
- | 2. Type a 3 from the PM Agent menu to display the Work with PM Agent Customization display and press Enter.

| If you are using Collection Services to gather your PM Agent data, you should take into account some data collection considerations for PM Agent.

| See manage PM Agent for other tasks that you can perform with PM Agent.

#### | **Related reference**

| "Managing PM Agent" on page 101

| Now that you have set up your network, you can perform a variety of tasks with PM Agent.

| *Verifying the PM Agent telephone number:*

| If your server is using a direct dial connection to IBM , you must verify that the PM Agent telephone number is correct. The telephone number must also contain the correct prefixes for your line.

| **Note:** This is for SNA transmissions only.

| To check the format of the telephone number of the electronic customer support line, do the following steps:

- | 1. Type  
| DSPDTAARA DTAARA(QUSRSYS/QESTELE)  
  
| and press Enter.
- | 2. Determine the connection number prefix found in offset 0. For example, if offset 0 is 'T9:1800xxxxxx' the prefix is T9:.
- | 3. Type  
| DSPDTAARA DTAARA(QUSRSYS/Q1PGTELE)  
  
| and press Enter.
- | 4. Offset 0 (zero) is the dialing string that will be used. (The other numbers will not be used.)
- | 5. If you use your ECS line to order PTFs, you can compare the format in offset 0 (zero) to the format used for the ECS line, CALL QESPHONE, make a note of the string being used, and compare it to the value found in step 2.  
| The telephone numbers will be different but the prefix should be the same (that is, SST9:1800..., SST:1800...and so forth).

| If you need to change your telephone number, use the Change Data Area (CHGDTAARA) command:

| Type **CHGDTAARA**, where DTAARA is Q1PGTELE, LIB is QUSRSYS, the substring starting position is \*ALL, and the New value is 'SST:18005475497'

| **Note:** The new value should be your dialing prefix, followed by 18005475497 for U.S.A and Canada.

| Now that you have completed your PM Agent configuration, see manage PM Agent for the tasks that you can perform.

| **Related reference**

| "Managing PM Agent" on page 101

| Now that you have set up your network, you can perform a variety of tasks with PM Agent.

| *Setting a direct dial line for PM Agent:*

| For most locations, PM Agent tries to select the correct data telephone number for your location.

| You should always verify that the PM Agent data telephone number is correct. If you do not have information that contains the PM Agent data telephone number and the PM Agent support number, contact your local IBM support personnel. They can provide you with the proper telephone numbers.

| **Note:** This telephone number is not required if you are transmitting data through the Universal Connection. You need this number only if you are using the direct dial line.

| To define the PM Agent data telephone number or to change the number, do the following steps:

- | 1. Type **GO PM400** from the command line.
- | 2. Type a 3 from the PM Agent Menu to display the Work with PM Agent Customization display and press Enter.
- | 3. On this display, scroll forward until you see the section of the display that shows you the telephone number fields.
- | 4. Type the correct dialing sequence in the **IBM PM Agent phone number** field. For many IBM modems, you are required to use the colon (:) character for dial tone.

| **Related tasks**

| "Verifying the PM Agent telephone number" on page 99

| If your server is using a direct dial connection to IBM , you must verify that the PM Agent telephone number is correct. The telephone number must also contain the correct prefixes for your line.

| *Varying a line off and on with PM Agent:*

| Sometimes the line that PM Agent uses is in the connect pending state. This state does not allow PM Agent to access the line to transmit data.

| The PM Agent Line Control display allows PM Agent to vary off the line, transmit the data, and then put the line back in the connect pending state. When you use this display, you can change the PM Agent transmission task (Q1PCM1) to check for line status and vary off the appropriate line. Once the transmission is complete, the same line is placed in a connect pending state.

| **Note:** This task is necessary only if PM Agent gathers data and transmits data over SNA.

| To vary off and on a line, do the following steps:

- | 1. Start the PM Agent line monitoring function by typing **CFGPMLIN** from the command line. You should see the Configure PM Agent line display.
- | 2. Read the warning that is shown on the first display and press Enter.
- | 3. Define the line, controller, and device combinations that PM Agent needs to vary off.
- | 4. Use the prompt **Do you want PM Agent automatic line control active?** as a master control switch for the function. If you specify **YES**, the PM Agent function is active. If you specify **NO**, the function is disabled.

- | If you specify **NO**, you do not need to define the line control list again when you specify **YES**. You can vary off and on a line by specifying the line only. You can vary off and on a line, controller, and device by specifying all three descriptions.
- | 5. Verify the line, controller, and device that you defined. Press Enter to see a summary of your choices.
- | 6. Press Enter to confirm your choices or press F12 to return to the previous display to change your entries.

| You can also set up PM Agent line control by using the Configure PM Agent (CFGPMAGT) command.

| **Related reference**

| "Managing PM Agent"

| Now that you have set up your network, you can perform a variety of tasks with PM Agent.

| **Managing PM Agent**

| Now that you have set up your network, you can perform a variety of tasks with PM Agent.

| After you have set up your network to use PM Agent, you can perform the following tasks:

| **Related reference**

| End PM Agent (Q1PENDPM) API

| **Deactivating PM Agent:**

| Learn how you can stop PM Agent.

| To stop PM Agent from running, you can use either of the following methods:

| **With System i Navigator**

| Perform the following steps:

- | 1. In System i Navigator, expand the system where PM Agent is running.
- | 2. Expand **Configuration and Service**.
- | 3. Right-click **Collection Services**.
- | 4. Select **PM Agent**.
- | 5. Select **Stop**.
- | 6. Select the systems on which you want to stop PM Agent.
- | 7. Click **OK**.

| **With an API**

| Use the End PM Agent (Q1PENDPM) API to deactivate PM Agent.

| **Related tasks**

| "Activating PM Agent" on page 89

| PM Agent is a part of the operating system and you must activate it to use its collecting capabilities.

| **Changing PM Agent contact information:**

| Learn how to change your contact information from the original settings.

| During the configuration of the PM Agent software, you identified the contact person and provided mailing information for your organization. If at a later time, you need to update the information, use the Work with Contact Information option to change that information. To change the contact information, do the following steps:

- | 1. Type **GO PM400** from the command line.

2. Type a 1 from the PM Agent Menu and press Enter. The Work with Contact Information display appears.
3. Change the contact information, as appropriate, and press Enter.

### Scheduling jobs with PM Agent:

Learn how to schedule jobs with PM Agent.

Integral to the PM Agent software is a scheduler that automatically starts the jobs that are necessary to support the PM Agent performance data collection and analysis.

Part of the PM Agent software activation process includes starting a job that is called Q1PSCH. This job, in turn, starts other jobs as shown in the following table:

To access the PM Agent scheduled jobs, do the following:

1. Type **GO PM400** from the command line.
2. Type a 2 from the PM Agent Menu and press Enter. The Work with Automatically Scheduled Jobs display appears.
3. You can change the status for each job from active to inactive. Type a 2 (Change) next to the job that you want to change and press Enter. You are shown the Change Automatically Scheduled Jobs display.

The following table shows you a list of the possible PM Agent jobs.

| PM Agent scheduled jobs |               |                                                                                                                                                                                                   |
|-------------------------|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Job                     | Schedule      | Function                                                                                                                                                                                          |
| Q1PTEST                 | At activation | Verifies that PM Agent is activated and then goes to inactive status.                                                                                                                             |
| Q1PCM1                  | Weekly        | Transmits the reduced performance data to IBM. This job is active only if you are using a direct dial line.                                                                                       |
| Q1PCM2                  | Daily         | Varies communications offline.                                                                                                                                                                    |
| Q1PPMSUB                | Hourly        | Ensures that Collection Services is collecting data.                                                                                                                                              |
| Q1PDR                   | Daily         | Performs data reduction and purges performance data.                                                                                                                                              |
| Q1PPG                   | Monthly       | Purges reduced performance data.                                                                                                                                                                  |
| Q1PCM3                  | As needed     | Varies communications offline after direct dial transmission fails to vary lines off.                                                                                                             |
| Q1PCM4                  | As needed     | Accesses the PM Agent data from remote servers. This job is started only if you have added remote systems by using option 5 from the PM Agent Menu.                                               |
| Q1PPMCHK                | Every 4 hours | Verifies that data collection is active.                                                                                                                                                          |
| Q1PMONTH                | Monthly       | Allows for monthly transmission if you require an additional transmission during the month. The default value is set to inactive. This job is available only if you are using a direct dial line. |



## | **Omitting items from IBM Performance Management for System i5 analysis:**

| Learn how to omit jobs, users, and communications lines when performing an analysis with IBM Performance Management for System i5 (PM for System i5).

| The PM for System i5 software application summary includes an analysis of items for batch jobs, users, and communication lines. However, some jobs, users, or communication lines are not appropriate for such an analysis. For example, you may want to exclude jobs with longer than normal run times, such as autostart jobs, in the run-time category.

| You can omit groups of batch jobs and users from the analysis by using the generic omit function. For example, to omit all jobs starting with MYAPP specify: MYAPP\*

| To work with omissions, do the following steps:

- | 1. Type **GO PM400** from the command line.
- | 2. Type a 4 from the PM Agent Menu and press Enter. The Work with Omissions display appears.
- | 3. Type the appropriate option number depending on which item you want to omit.
  - | • Type a 1 to work with jobs.
  - | • Type a 2 to work with users.
  - | • Type a 3 to work with communications lines.
- | 4. Type a 1 in the appropriate field to omit either a user or a job from a particular category. In the case of a communications line, type the name of the line and then type a 1 in the appropriate field.

## | **Stopping PM Agent momentarily:**

| Learn how you can stop PM Agent momentarily.

| If you need to stop PM Agent from verifying that Collection Services is collecting data, you can use the scheduler job to change the date to a future date for the Q1PPMSUB job.

- | 1. Type **GO PM400** from the command line.
- | 2. Type a 2 (Work with automatically scheduled jobs).
- | 3. Type a 2 (Change) next to the Q1PPMSUB job.
- | 4. Change the date or time to a future date and time.
- | 5. Press Enter. This change will momentarily stop PM Agent from verifying that Collection Services is collecting data. You must end what is currently being collected.

| **Note:** PM Agent will not start, cycle, or change Collection Services until the date and time to which you set the Q1PPMSUB job has been reached.

### | **Related tasks**


| “Scheduling jobs with PM Agent” on page 102  
| Learn how to schedule jobs with PM Agent.

## | **Viewing PM Agent status:**

| Learn how to use System i Navigator or the PM Agent menu to display PM Agent status.

| You can use either System i Navigator or the PM Agent Menu on your system to display the status of PM Agent. Use the PM Agent Status dialog to view the overall status of PM Agent on one or more systems or groups. For example, you are shown details as to whether PM Agent is active. Use the PM Agent Menu to view the Collection Services status, PM Agent scheduler status, the performance data release, the last transmission attempt, performance data members, and the performance data size.

| To view the overall status for PM Agent from System i Navigator, do the following steps:

- | 1. In System i Navigator, expand an endpoint system or a system group.
  - | 2. Expand **Configuration and Service**.
  - | 3. Right-click **Collection Services**.
  - | 4. Select **Performance Agent**.
  - | 5. Select **Status**.
- | To view the detailed status for PM Agent from the PM Agent menu, do the following steps:
- | 1. Type **GO PM400** from the command line.
  - | 2. Type a 6 from the command line and press Enter. See the online help for descriptions of each field.
- | **Viewing IBM Performance Management for System i reports:**
- | See examples of the IBM Performance Management for System i reports and explanations of how to interpret those reports.
- | The output of the IBM Performance Management for System i offering is a set of management reports and graphs. The purpose of the reports and graphs is to give management a clear understanding of the current performance of their servers and an accurate growth trend. To view the reports and learn about some of their benefits and uses, visit the IBM Performance Management for System i Web site.
- | **Related concepts**
- | “System i Navigator graph history” on page 86
- | Graph history provides a graphical display of performance data collected by Collection Services over a specified period of time.
- | **Related information**
- |  PM for System i5 Web site

## IBM Systems Workload Estimator

The IBM Systems Workload Estimator is a Web-based sizing tool for System i, System p, and System x. You can use this tool to size a new system, to size an upgrade to an existing system, or to size a consolidation of several systems.

The Workload Estimator allows measurement input to best reflect your current workload; one method is by using data from Performance Management for System i. The Workload Estimator also provides a variety of built-in workloads to reflect your emerging application requirements. Virtualization can be used to yield a more robust solution. The Workload Estimator will provide current and growth recommendations for processor, memory, and disk that satisfy the overall client performance requirements.

### Related information



IBM Systems Workload Estimator

See the IBM Systems Workload Estimator Web site to run the online version of the Workload Estimator.

## Performance Tools

- | The Performance Tools licensed program includes many supplemental features that supplement or extend the capabilities of the basic performance tools that are available in the operating system.

### Related concepts

- | “Collection Services” on page 24
- | Collection Services provides for the collection of system management data. It is the primary collector of system data.

“Manager and Agent feature comparison” on page 106

You can use the Manager and Agent features to efficiently divide required functions of Performance Tools over a distributed environment. This topic contains a description of these two features, the functions they each contain, and information about how to use them most effectively.

#### **Related information**



Performance Tools PDF

CL commands for performance

### **Performance Tools concepts**

Describes a variety of tools to help you collect and analyze performance information. Find detailed information about exactly which tools perform which functions and how they work.

#### **Related concepts**

“Collection Services” on page 24

Collection Services provides for the collection of system management data. It is the primary collector of system data.

### **Functions provided in Performance Tools:**

Performance Tools includes a variety of applications for collecting, analyzing, and reporting performance data. Knowing which functions are available, and which are best suited for a given task can be complex. This topic describes the functions included in this licensed program.

Performance Tools includes reports, interactive commands, and other functions. For example, Performance Tools includes the following:

| <b>Tool</b>              | <b>Description</b>                                                                                                                                                                                                                                                                                                          |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Display Performance Data | The Display Performance Data graphical user interface allows you to view performance data, summarize the data into reports, display graphs to show trends, and analyze the details of your system performance all from within System i Navigator.                                                                           |
| Reports                  | The reports organize Collection Services performance data and trace data in a logical and useful format. The reports are discussed in detail in the Performance Tools book.                                                                                                                                                 |
| Graphics function        | The Performance Tools graphics function allows you to work with performance data in a graphical format. You can display the graphs interactively, or you can print, plot, or save the data to a graphics data format (GDF) file for use by other utilities. This tool is discussed in detail in the Performance Tools book. |
| IBM i5/OS Job Watcher    | The Job Watcher function of the IBM Systems Director Navigator for i5/OS Performance interface is included in Performance Tools.                                                                                                                                                                                            |
| IBM i5/OS Disk Watcher   | The Disk Watcher function of the IBM Systems Director Navigator for i5/OS Performance interface is included in Performance Tools.                                                                                                                                                                                           |

#### **Related concepts**

“IBM Systems Director Navigator for i5/OS Performance interface” on page 66

The IBM Systems Director Navigator for i5/OS Performance interface allows you to display and manage performance data by bringing together a variety of performance information into one place.

#### **Related tasks**

“Performance Tools plug-in” on page 106

You can view system resource utilization data in System i Navigator. You can view the data, graph it, and summarize it into reports. Find information about how to access this function here.

#### **Related reference**

Work with System Activity (WRKSYSACT) command

#### **Related information**

Performance Tools reports



## **Manager and Agent feature comparison:**

You can use the Manager and Agent features to efficiently divide required functions of Performance Tools over a distributed environment. This topic contains a description of these two features, the functions they each contain, and information about how to use them most effectively.

Performance Tools is available with two separately installable features. This topic explains the differences between the two features to help you decide which feature is more appropriate for your applications.

### **Manager feature**

The Performance Tools Manager feature is a full-function package, intended to be used on the central site system in a distributed environment or on a single system. If you require analysis of trace data, viewing data graphically, viewing system activity in real time, or managing and tracking system growth, the Manager feature of the Performance Tools licensed program is more useful. The Manager feature also contains the IBM i5/OS Disk Watcher function of the of the IBM Systems Director Navigator Performance interface.

### **Agent feature**

The Performance Tools Agent feature, with a subset of the Manager function, is a lower-priced package with the more basic functions. In a distributed environment, the Agent feature works well for managed systems in the network because the data can be sent to the Manager if detailed analysis is required. It is also an effective tool for sites that need a reasonable level of self-sufficiency but have no expert skills available.

The Agent feature of Performance Tools provides functions to simplify the collection, management, online display, data reduction, and analysis of performance data. The Performance explorer reporting function and its associated commands are included in the base option in the Performance Tools for i5/OS licensed program and, therefore, are available with either the Manager feature or the Agent feature. The major Performance Tools functions not contained in the Agent feature are performance and trace reports, performance utilities (job traces and the select file utilities), system activity monitoring, and performance graphics.

### **Related concepts**

“IBM Systems Director Navigator for i5/OS Performance interface” on page 66

The IBM Systems Director Navigator for i5/OS Performance interface allows you to display and manage performance data by bringing together a variety of performance information into one place.

## **Performance Tools plug-in:**

You can view system resource utilization data in System i Navigator. You can view the data, graph it, and summarize it into reports. Find information about how to access this function here.

Performance Tools can display performance data from the Display Performance Data graphical user interface (GUI), which is a plug-in for System i Navigator. From the GUI, you can view performance data, summarize the data into reports, display graphs to show trends, and analyze the details of your system performance.

### **Metrics**

System i Navigator displays performance metrics over a selected interval of time. The performance metrics you can view in the Graphs pane of the Display Performance Data GUI include:

- Transaction Count
- Transaction Response Time
- Total CPU Utilization

- | • Interactive CPU Utilization
- | • Batch CPU Utilization
- | • Interactive Feature Utilization
- | • High Disk Utilization
- | • Machine Pool Page Faults/Second
- | • User Pool Page Faults/Second
- | • Exceptions

| The Details pane allows you to view detailed performance data for the selected time interval in a variety of ways. To analyze your system performance, you can view job data, subsystem data, pool data, or disk unit data.

## | **Reports**

| In addition to viewing graphs and detail data, you can also print reports from the Display Performance Data GUI. Performance reports allow you to research areas of the system that are causing performance problems. You can run different reports to see where your system resources are being used. Printing reports in Performance Tools is available only when option 1 (Manager feature) of Performance Tools for i5/OS (5761-PT1) is installed on the central system.

| The reports you can print from the Display Performance Data GUI include:

- | • System
- | • Component
- | • Job
- | • Pool
- | • Resource

## | **Accessing through System i Navigator**

| The Display Performance Data GUI is a plug-in for System i Navigator. If you have already installed the plug-in, it can be accessed from System i Navigator by following these steps:

- | 1. In System i Navigator, expand **My Connections** (or your active environment).
- | 2. Expand the server that contains the performance data you want to view.
- | 3. Expand **Configuration and Service**.
- | 4. Right-click **Collection Services**, select **Performance Tools**, and then select **Performance Data**.
- | 5. Select the performance data file that you want to display.
- | 6. Click **Display**.

| For more information on how to use the Display Performance Data GUI in System i Navigator, see the System i Navigator online help.

### | **Related concepts**

| “Manager and Agent feature comparison” on page 106

| You can use the Manager and Agent features to efficiently divide required functions of Performance Tools over a distributed environment. This topic contains a description of these two features, the functions they each contain, and information about how to use them most effectively.

### | **Related tasks**

| “Installing the Performance Tools plug-in to System i Navigator” on page 111

| You can install the Performance Tools plug-in to System i Navigator to view system resource utilization data.

## | **Reporting CPU utilization:**

Find out how the total CPU that is consumed across virtual processors is reported.

Prior to V5R3, processor utilization was calculated as a percentage of the available CPU time. Collection Services reported, in the performance database files, the time used on each processor along with elapsed interval time. Users of this data, such as the Performance Tools reports and displays, needed to add up the time used on each processor to get the total system CPU that was consumed. The available CPU time was calculated as the number of processors in the partition multiplied by the duration of the data collection interval. Finally, the CPU time was divided by the calculated available time to get the utilization percentages.

The problem with the previous methodology is that all users of the data assumed whole virtual processors and depended on no changes to the configured capacities. Logical partitions with partial processor capacities and the capability to do dynamic configuration no longer worked with this methodology. Temporary solutions for minimizing the impacts of these problems included scaling the utilization of the system processors to what would be reported for a whole number of processors and cycling Collection Services when the configuration changed. Because the individual job CPU time was not scaled, the additional time was accounted for by reporting it as being consumed by HVLPTASK. The HVLPTASK task did not actually use CPU, but CPU time was shown to be consumed by HVLPTASK for accounting purposes. The CPU time charged to HVLPTASK scaled the amount of work that was done by real jobs, which resulted in the system CPU percent utilization going from 0 to 100 in direct proportion to the amount of customer work that was performed.

In V5R3, Collection Services reports the total CPU that is consumed and the total CPU that is available to the partition within the interval. The concept of HVLPTASK and CPU scaling to whole virtual processors in shared processor environments does not exist. Collection Services no longer cycles the collection when the configuration changes.

Collection Services now reports the total processor time that is consumed by the partition along with the amount of processor time that was available to be consumed within the partition, regardless of the number of virtual processors that are configured, the partition units that are configured, or how they changed during the interval. To calculate utilization, users of this data divide the reported CPU consumed by the available capacity. This method of calculating CPU utilization eliminates the increasingly error-prone task of computing available CPU time. CPU utilization that is calculated with these new metrics is accurate regardless of how many processing units (whole or fractional) exist, when the processing units changed, or how often the units changed.

Several reasons account for this change in calculating CPU utilization. One reason is that with scaling utilization for jobs or groups of jobs appeared to be much smaller than would be anticipated. This concept is demonstrated in the example that follows. Another reason is that a configuration change could make CPU reporting not valid. Traditionally, the number of CPUs was based on the value that was configured at the beginning of a collection and an IPL was needed to change it. When dynamic configuration was introduced, Collection Services cycled the collection to handle the configuration changes, which assumed that changes would be infrequent. However, the more frequent the change, the more cycling occurs. If the changes are too frequent, collecting performance data is not possible. Lastly, even if the proper configuration data were reported and used for every interval, you would not know what happened between the time the interval started and until it completed. Utilization would still be calculated incorrectly in any interval where there was one or more configuration changes.

### Example

Partition A has a capacity of 0.3 processor units and is defined to use one virtual processor. The collection interval time is 300 seconds. The system is using 45 seconds of CPU (15 seconds by interactive jobs and 30 seconds by batch jobs). In this example, the available CPU time is 90 seconds (.3 of 300 seconds). The total CPU utilization is 50%.



| Prior to V5R3, when the numbers were scaled, system CPU usage is reported as 150 seconds. 150 seconds  
 | divided by 300 seconds of interval time results in 50% utilization. The interactive utilization is 15 seconds  
 | divided by 300 seconds, which is 5%. The batch utilization is 30 seconds divided by 300 seconds, which is  
 | 10%. The HVLPTASK is getting charged with 35% utilization (150 seconds minus 45 seconds), or 105  
 | seconds divided by 300 seconds. These percentages give us a total of 50%.

| Beginning in V5R3, the 45 seconds of usage is no longer scaled but is reported as is. The calculated CPU  
 | time that is derived from the reported consumed CPU time divided by the reported available capacity is  
 | 50% (45 seconds divided by 90 seconds). The interactive utilization percentage is 17% (15 seconds divided  
 | by 90 seconds). The batch utilization percentage is 33% (30 seconds divided by 90 seconds).

| Release                | Total CPU | Interactive | Batch | HVLPTASK |
|------------------------|-----------|-------------|-------|----------|
| OS/400 V5R2 or earlier | 50%       | 5%          | 10%   | 35%      |
| OS/400® V5R3 or later  | 50%       | 17%         | 33%   | N/A      |

## | **Considerations**

| In V5R3 and later, the Convert Performance Data (CVTPFRDTA) command performs normally. However,  
 | the data in the converted files is changed to be consistent with the unscaled system CPU data  
 | (QAPMSYSCPU database file). The results should be the same as if the data were collected on a V5R3 or  
 | later system, but the data is different than the values that existed in the files in a prior release.

| The existing and unchanged tools that calculate CPU utilization do not show the correct results for  
 | shared processor partitions or partitions that have had configuration changes during data collection. This  
 | includes those tools that use the performance database as well as those that use the QPMLPFRD API.

| You can copy a V5R3 or later management collection object (\*MGTCOL) to a prior release and generate  
 | the database files. However, you should be aware of the following:

- | • The reported CPU data remains unscaled (shared processor environments). This means that the total  
 | system CPU that is reported by the tools using virtual processors (including Performance Tools) is not  
 | correct.
- | • A management collection object (\*MGTCOL) that spans configuration changes will result in an  
 | inaccurate calculation of the percentage of CPU during those intervals after the change occurred.

## | **Reporting configured capacity:**

| Find out where the information for configured capacity is recorded.

| The partition capacity values are determined initially when the partition is started and depends on the  
 | capacity resources available at the time. These initial values can be altered through configuration changes  
 | while the partition is active.

| Logical partitions (LPARs) enable some partitions to exceed their configured capacity in certain situations.  
 | During these times, the processor utilization metrics of these partitions can be greater than 100% of the  
 | configured capacity.

| The usage and capacity information is recorded in the QAPMSYSTEM database file. The virtual processor  
 | information is recorded in the QAPMSYSCPU database file. The following values summarize this  
 | information:

## | **Virtual processors**

| The number of processors that are assigned to a logical partition that is sharing processor  
 | capacity of the shared processor pool. This value determines the number of concurrent processors



that could be active in the logical partition. This value is included in the QAPMSYSCPU performance database file in the field (or column) named SCTACT.

**Shared processor pool capacity available**

Total processor capacity in the shared processor pool available for use by shared processor logical partitions. This value is included in the QAPMSYSTEM performance database file in a column named SYSPLA. If partitions configured as uncapped compete for available shared pool capacity in excess of the guaranteed amount, the distribution of processor capacity is determined by the uncapped weight assigned to the logical partition.

**Shared processor capacity used**

Total amount of shared processor capacity used by all active shared processor logical partitions. Total amount of CPU used within the shared pool by all partitions that share the pool. This value is included in the QAPMSYSTEM performance database file in a column named SYSPLU.

**Partition guaranteed capacity**

Processor capacity configured to a shared processor logical partition from the shared processor pool. This value is included in the QAPMSYSTEM performance database file in a column named SYSCTA. The 5250 OLTP capacity configured is recorded in column named SYIFTA.

**Partition processor utilization**

Total CPU time used by a logical partition. In a shared processor logical partition with uncapped capacity, this value may exceed the guaranteed capacity if there is unused capacity in the shared processor pool. This value is included in the QAPMSYSTEM performance database file in a column named SYSPTU. The 5250 OLTP capacity used is recorded in column named SYIFUS. The maximum processor capacity in a partition is determined by the number of virtual processors configured.

**Partition available capacity**

The amount of processor capacity that could have been used by the logical partition. This value is included in the QAPMSYSTEM performance database file in a column named SYSUTA. This is the processor capacity utilized (SYSPTU) plus the unused capacity in the shared processor pool (SYSPLA), subject to the following limits:

- The minimum is the configured (guaranteed) capacity.
- The maximum is the capacity based on the number of virtual processors assigned to the partition and pool.

**Related information**

Performance data files: QAPMSYSTEM

Performance data files: QAPMSYSCPU

**5250 online transaction processing (OLTP):**

This topic describes 5250 online transaction processing and what jobs or threads are associated with this work.

*Online transaction processing (OLTP)* refers to a type of interactive application in which requests submitted by users are processed as soon as they are received. The following are examples of OLTP processing:

- The system interactions through a 5250 session, a pass-through job, or a Telnet job.
- A workstation-based request from a Domino mail or calendar, or a browser-based application.

System i Access jobs use both interactive and batch, depending on the function. Before V5R3, these jobs were included in the CA4 category and listed as interactive. The distributed data management (DDM) server jobs were also listed as interactive.

After V5R3, the Performance Tools licensed program is updated to better distribute the workloads, depending on which processor capacity feature that the CPU cycles were charged against. The interactive CPU reporting refers to those jobs whose CPU is allocated to the 5250 OLTP processor capacity. The

| System i Access jobs are listed in the appropriate sections of the Performance Tools reports. In addition, the DDM jobs moved from the Interactive workload section of the reports to the Non-interactive workload section.

## | **Installing and configuring Performance Tools**

| See this topic for installation and setup instructions.

| To install Performance Tools, you need a user profile with save system (\*SAVSYS) authority. You can use the system operator profile to obtain this authority.

| Performance Tools must run in a library named QPFR. If you have a library by this name on your system, use the Rename Object (RNMOBJ) command to rename it before you install Performance Tools. This step will ensure the proper operation of Performance Tools.

| Use the following command to place the Performance Tools in library QPFR:

| RSTLICPGM LICPGM(xxxxPT1) DEV(NAME) OPTION(\*BASE)

| You must then perform one of the following:

| • If you have purchased the Manager feature, use the following command:

| RSTLICPGM LICPGM(xxxxPT1) DEV(*tape-device-name*) OPTION(1)

| • If you have purchased the Agent feature, use the following command:

| RSTLICPGM LICPGM(xxxxPT1) DEV(NAME) OPTION(2)

| • In addition to installing either the Manager or the Agent feature, if you have purchased IBM i5/OS Job Watcher, use the following command:

| RSTLICPGM LICPGM(xxxxPT1) DEV(*tape-device-name*) OPTION(3)

| If you have several CD-ROMs to install, the following situation may occur. After installing the first one, you may receive a message saying that the licensed program is restored but no language objects were restored. If this occurs, insert the next CD-ROM and enter the following:

| RSTLICPGM LICPGM(xxxxPT1) DEV(NAME) RSTOBJ(\*LNG) OPTION(\*BASE)

| Another method for installing the Performance Tools program is to type GO LICPGM and use the menu options.

| Performance Tools is a processor-based program. The usage type is concurrent, and the program is installed with a usage limit \*NOMAX.

| This program is discussed in detail in the Performance Tools book.

### | **Related information**

|  Performance Tools PDF

### | **Installing the Performance Tools plug-in to System i Navigator:**

| You can install the Performance Tools plug-in to System i Navigator to view system resource utilization data.

| Before you install the Performance Tools plug-in, you should first ensure that the Performance Tools (5761-PT1) licensed program is installed. To install the Performance Tools plug-in, do the following:

- | 1. In System i Navigator, right-click **My Connections** and select **Install Options** → **Install Plug-ins**.
- | 2. On the Install Plug-ins panel, select the system that you want to install the plug-in from (must be a system that has the 5761-PT1 product installed) and click **OK**.
- | 3. Enter your i5/OS user profile name and password on the Password panel and click **OK**. (The prompt may ask for the Windows® password, but it needs to be the i5/OS user profile password.)

**Note:** Some Windows operating systems may require the Windows and i5/OSuser profile passwords to match.

4. A scan for plug-ins on the selected system will occur next. When the Plug-in Selection panel appears, check the Performance Tools checkbox and click **Next**.
5. The first time the System i Navigator is used after installing a plug-in, the System i Navigator scan panel will appear. Click **Scan Now**. If you do not click the Scan Now button, the plug-in that was just installed will be disabled and will not show up in System i Navigator.

## Performance Tools reports

Performance Tools reports provide information on data that has been collected over time. Use these reports to get additional information about the performance and use of system resources.

The Performance Tools reports provide an easy way for you to look at your collected data and isolate performance problems. After you have collected performance data over time, you can print the reports to see how and where system resources are being used. The reports can direct you to specific application programs, users, or inefficient workloads that are causing slower overall response times.

Collection Services provides data for most of the Performance Tools reports with the exception of the Transaction, Lock, and Trace reports. You must use the Start Performance Trace (STRPFTRC) and End Performance Trace (ENDPFTRC) commands to collect the trace information for those three reports.

### Overview of Performance Tools reports:

The following list describes each report and gives a brief overview as to why you would use a particular report.

*Table 5. Overview of Performance Tools reports*

| Report           | Description                                                                                                                                                                                                                                                                                                | What is shown                                                                                                                                   | How you use the information                         |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------|
| System Report    | Uses Collection Services data to provide an overview of how the system is operating. The report contains summary information on the workload, resource use, storage pool utilization, disk utilization, and communications. Run and print this report often to give you a general idea of your system use. | System workload. The report includes the database capabilities data.                                                                            | Workload projection                                 |
| Component Report | Uses Collection Services data to provide information about the same components of system performance as a System Report, but at a greater level of detail. This report helps you find which jobs are consuming high amounts of system resources, such as CPU, disk, and so on.                             | Resource use, communications, system and user jobs. The report includes the database capabilities data and the Interactive Feature utilization. | Hardware growth and configuration processing trends |

Table 5. Overview of Performance Tools reports (continued)

| Report                 | Description                                                                                                                                                                                                                                                                                                                                                                                                      | What is shown                                                                                                 | How you use the information                                                                   |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| Transaction Reports    | Uses trace data to provide detailed information about the transactions that occurred during the performance data collection.                                                                                                                                                                                                                                                                                     | Workload and utilization of CPU, disk, main storage, transaction workload, object contention                  | Workload projection, pool configuration, application design, file contention, and program use |
| Lock Report            | Uses trace data to provide information about lock and seize conflicts during system operation. With this information you can determine if jobs are being delayed during processing because of unsatisfied lock requests or internal machine seize conflicts. These conditions are also called waits. If they are occurring, you can determine which objects the jobs are waiting for and the length of the wait. | File, record, or object contention by time; the holding job or object name; the requesting job or object name | Problem analysis. Reduction or elimination of object contention.                              |
| Batch Job Trace Report | Uses trace data to show the progression of different job types (for example, batch jobs) traced through time. Resources utilized, exceptions, and state transitions are reported.                                                                                                                                                                                                                                | Job class time-slice end and trace data                                                                       | Problem analysis and batch job progress                                                       |
| Job Interval Report    | Uses Collection Services data to show information on all or selected intervals and jobs, including detail and summary information for interactive jobs and for noninteractive jobs. Because the report can be long, you may want to limit the output by selecting the intervals and jobs you want to include.                                                                                                    | Jobs by interval                                                                                              | Job data                                                                                      |
| Pool Interval Report   | Uses Collection Services data to provide a section on subsystem activity and a section on pool activity. Data is shown for each sample interval. Because the report can be long, you may want to limit the output by selecting the intervals and jobs you want to include.                                                                                                                                       | Pools by interval                                                                                             | Pool data                                                                                     |

Table 5. Overview of Performance Tools reports (continued)

| Report                   | Description                                                                                                                                                                                                  | What is shown         | How you use the information |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|-----------------------------|
| Resource Interval Report | Uses Collection Services data to provide resource information on all or selected intervals. Because the report can be long, you may want to limit the output by selecting the intervals you want to include. | Resources by interval | System resource use         |

Performance explorer and Collection Services are separate collecting agents. Each one produces its own set of database files that contain grouped sets of collected data. You can run both collections at the same time.

#### Sample System Report - Workload:

The Workload section of the system report displays the interactive and non-interactive workload of the system.

The first part of the Workload section of the System Report displays the Interactive Workload of the system. The second part of the Workload section displays the Non-Interactive Workload of the system.

|                                          |                     |                      |                      |                      |                          |                          |                    |
|------------------------------------------|---------------------|----------------------|----------------------|----------------------|--------------------------|--------------------------|--------------------|
|                                          |                     |                      |                      | System Report        |                          | 6/26/04 16:06            |                    |
|                                          |                     |                      |                      | Workload             |                          | Page 0                   |                    |
| Member . . . .                           | PNT6PERF            | Model/Serial . .     | 825/10-5M0FM         | Main storage . .     | 8192.0 MB                | Started . . . .          | 04/07/04 19:11     |
| Library . . .                            | CARR098R01          | System name . .      | CARREGT              | Version/Release .    | 5/ 4.0                   | Stopped . . . .          | 04/07/04 20:15     |
| Partition ID .                           | 000                 | Feature Code . .     | 7415-2472-7415       | Int Threshold . .    | 100.00 %                 |                          |                    |
| Virtual Processors:                      | 32                  | Processor Units :    | 32.0                 |                      |                          |                          |                    |
| QPFRA DJ . . .                           | 0                   | QDYNPTYSCD . . .     | 1                    | QDYNPTYADJ . . . : 1 |                          |                          |                    |
| Interactive Workload                     |                     |                      |                      |                      |                          |                          |                    |
| Job Type                                 | Number Transactions | Average Response     | Logical DB I/O Count | ----- Printer Lines  | ----- Pages              | Communications I/O Count | MRT Max ime        |
| Interactive                              | 3,242               | .65                  | 16,734               | 12,910               | 339                      | 0                        | 0                  |
| DDM Server                               | 0                   | .00                  | 864,667              | 443                  | 23                       | 1,596,096                | 0                  |
| PassThru                                 | 6,645               | .68                  | 343,262              | 1,119,009            | 27,769                   | 240                      | 0                  |
| Total                                    | 9,887               |                      | 1,224,663            | 1,132,362            | 28,131                   | 1,596,336                |                    |
| Average                                  |                     | .67                  |                      |                      |                          |                          |                    |
| -----                                    |                     |                      |                      |                      |                          |                          |                    |
| Non-Interactive Workload                 |                     |                      |                      |                      |                          |                          |                    |
| Job Type                                 | Number Of Jobs      | Logical DB I/O Count | ----- Printer Lines  | ----- Pages          | Communications I/O Count | CPU Per Logical I/O      | Logical I/O/Second |
| Batch                                    | 18,151              | 1,030,253,068        | 18,656,603           | 544,032              | 1,531,738                | .0001                    | 95,526.4           |
| Spool                                    | 70                  | 1,066                | 14,933               | 369                  | 0                        | .0285                    | .0                 |
| AutoStart                                | 56                  | 426,047              | 1,692,060            | 41,502               | 178,288                  | .0008                    | 39.5               |
| COLLECTION                               | 1                   | 2,910                | 0                    | 0                    | 0                        | .0171                    | .2                 |
| SQL                                      | 192                 | 3,252,232            | 3,519                | 88                   | 0                        | .0003                    | 301.5              |
| MGMTCENTRAL                              | 2                   | 12,229               | 0                    | 0                    | 0                        | .0046                    | 1.1                |
| Total                                    | 18,903              | 1,033,969,357        | 20,367,115           | 585,991              | 1,713,007                |                          |                    |
| Average                                  |                     |                      |                      |                      |                          | .0003                    | 95,871.0           |
| Average CPU Utilization . . . . . : 61.0 |                     |                      |                      |                      |                          |                          |                    |
| CPU 1 Utilization . . . . . : 55.4       |                     |                      |                      |                      |                          |                          |                    |
| CPU 2 Utilization . . . . . : 57.9       |                     |                      |                      |                      |                          |                          |                    |
| CPU 3 Utilization . . . . . : 61.5       |                     |                      |                      |                      |                          |                          |                    |
| CPU 4 Utilization . . . . . : 62.2       |                     |                      |                      |                      |                          |                          |                    |
| CPU 5 Utilization . . . . . : 62.0       |                     |                      |                      |                      |                          |                          |                    |
| CPU 6 Utilization . . . . . : 60.1       |                     |                      |                      |                      |                          |                          |                    |
| CPU 7 Utilization . . . . . : 61.7       |                     |                      |                      |                      |                          |                          |                    |
| CPU 8 Utilization . . . . . : 63.1       |                     |                      |                      |                      |                          |                          |                    |
| CPU 9 Utilization . . . . . : 55.4       |                     |                      |                      |                      |                          |                          |                    |
| CPU 10 Utilization. . . . . : 56.0       |                     |                      |                      |                      |                          |                          |                    |
| CPU 11 Utilization. . . . . : 59.9       |                     |                      |                      |                      |                          |                          |                    |
| CPU 12 Utilization. . . . . : 60.6       |                     |                      |                      |                      |                          |                          |                    |
| CPU 13 Utilization. . . . . : 60.9       |                     |                      |                      |                      |                          |                          |                    |
| CPU 14 Utilization. . . . . : 62.5       |                     |                      |                      |                      |                          |                          |                    |
| CPU 15 Utilization. . . . . : 63.7       |                     |                      |                      |                      |                          |                          |                    |
| CPU 16 Utilization. . . . . : 64.1       |                     |                      |                      |                      |                          |                          |                    |

```

CPU 17 Utilization. . . . . : 54.7
CPU 18 Utilization. . . . . : 57.3
CPU 19 Utilization. . . . . : 59.8
CPU 20 Utilization. . . . . : 60.6
CPU 21 Utilization. . . . . : 61.6
CPU 22 Utilization. . . . . : 62.9
CPU 23 Utilization. . . . . : 63.9
CPU 24 Utilization. . . . . : 64.7
CPU 25 Utilization. . . . . : 57.0
CPU 26 Utilization. . . . . : 55.2
CPU 27 Utilization. . . . . : 66.2
CPU 28 Utilization. . . . . : 61.1
CPU 29 Utilization. . . . . : 62.4
CPU 30 Utilization. . . . . : 63.2
CPU 31 Utilization. . . . . : 66.2
CPU 32 Utilization. . . . . : 66.4

```

```

Total CPU Utilization (Interactive Feature) . . : .0
Total CPU Utilization (Database Capability) . . : 51.6

```

### Sample Component Report - Job Workload Activity:

The Job Workload Activity section of the Component Report gives the total number of transactions, the transactions per hour, the average response time, the number of disk operations, the number of communications operations, the number of PAG faults, the number of arithmetic overflows, and the number of permanent writes for each job.

The values that display in the report header reflect the configuration metrics obtained from the QAPMCONF file when the collection started. These values might change for each interval within a collection period due to dynamic changes in logical partition configuration.

| Component Report                                               |                  |                                                                                            |        |          |             |                   |     |   |      |       |               |                   |                  | 10/02/03 17:12:15 |              |                 |               |
|----------------------------------------------------------------|------------------|--------------------------------------------------------------------------------------------|--------|----------|-------------|-------------------|-----|---|------|-------|---------------|-------------------|------------------|-------------------|--------------|-----------------|---------------|
| Job Workload Activity                                          |                  |                                                                                            |        |          |             |                   |     |   |      |       |               |                   |                  | Page 6            |              |                 |               |
| Perf data from 14:00 to 16:00 at 1 min                         |                  |                                                                                            |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| Member . . . : Q275140000 Model/Serial . . : 890/10-3907F      |                  |                                                                                            |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| Main storage . . : 56.4 GB Started . . . . : 10/02/03 14:00:00 |                  |                                                                                            |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| Library . . : PTLIBV5R3 System name . . : ABSYSTEM             |                  |                                                                                            |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| Version/Release : 5/ 3.0 Stopped . . . . : 10/02/03 16:00:00   |                  |                                                                                            |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| Partition ID : 003 Feature Code . : 7427-2498-7427             |                  |                                                                                            |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| Int Threshold . . : .00 %                                      |                  |                                                                                            |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| Virtual Processors: 4 Processor Units : 4.0                    |                  |                                                                                            |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| Job Name                                                       | User Name/Thread | Job Number                                                                                 | T<br>p | P<br>l y | CPU<br>Util | DB<br>Cpb<br>Util | Tns |   |      |       | -----<br>Sync | Disk I/O<br>Async | -----<br>Logical | Cmn<br>I/O        | PAG<br>Fault | Arith<br>Ovrflw | Perm<br>Write |
| ADMIN                                                          | QTMHHTTP         | 955725                                                                                     | B      | 02 25    | .02         | .0                | 0   | 0 | .000 | 14771 | 615           | 0                 | 0                | 0                 | 0            | 0               | 2787          |
| ADMIN                                                          | QTMHHTTP         | 955727                                                                                     | B      | 02 25    | .00         | .0                | 0   | 0 | .000 | 24    | 0             | 0                 | 0                | 0                 | 0            | 0               | 2             |
| ADMIN                                                          | QTMHHTTP         | 955728                                                                                     | B      | 02 25    | .00         | .0                | 0   | 0 | .000 | 0     | 0             | 165               | 0                | 0                 | 0            | 0               | 0             |
| ADMIN                                                          | QTMHHTTP         | 956347                                                                                     | B      | 02 25    | .14         | .0                | 0   | 0 | .000 | 959   | 343           | 1349              | 0                | 0                 | 0            | 0               | 736           |
| AMQALMPX                                                       | QMCM             | 955751                                                                                     | B      | 02 35    | .00         | .0                | 0   | 0 | .000 | 0     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| AMQPCSEA                                                       | QMCM             | 955757                                                                                     | B      | 02 35    | .00         | .0                | 0   | 0 | .000 | 0     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| AMQRMPPA                                                       | QMCM             | 955773                                                                                     | B      | 02 35    | .01         | .0                | 0   | 0 | .000 | 14    | 0             | 2                 | 0                | 0                 | 0            | 0               | 0             |
| AMQRRMFA                                                       | QMCM             | 955752                                                                                     | B      | 02 35    | .00         | .0                | 0   | 0 | .000 | 1     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| AMQZDMAA                                                       | QMCM             | 955753                                                                                     | B      | 02 35    | .00         | .0                | 0   | 0 | .000 | 0     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| AMQZLAA0                                                       | QMCM             | 955755                                                                                     | B      | 02 20    | .02         | .0                | 0   | 0 | .000 | 7     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| AMQZLAA0                                                       | QMCM             | 955774                                                                                     | B      | 02 20    | .00         | .0                | 0   | 0 | .000 | 2     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| AMQZXMA0                                                       | QMCM             | 955749                                                                                     | B      | 02 20    | .00         | .0                | 0   | 0 | .000 | 1     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| CFINT01                                                        |                  | L 01 00                                                                                    |        |          | .26         | .0                | 0   | 0 | .000 | 0     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| CFINT02                                                        |                  | L 01 00                                                                                    |        |          | .06         | .0                | 0   | 0 | .000 | 0     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| CFINT03                                                        |                  | L 01 00                                                                                    |        |          | .08         | .0                | 0   | 0 | .000 | 0     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| CFINT04                                                        |                  | L 01 00                                                                                    |        |          | .08         | .0                | 0   | 0 | .000 | 0     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| CFINT05                                                        |                  | L 01 00                                                                                    |        |          | .00         | .0                | 0   | 0 | .000 | 0     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| CFINT06                                                        |                  | L 01 00                                                                                    |        |          | .00         | .0                | 0   | 0 | .000 | 0     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| COLDQT                                                         |                  | L 01 82                                                                                    |        |          | .00         | .0                | 0   | 0 | .000 | 0     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| CPUTEST                                                        | WLCPU            | 953645                                                                                     | B      | 02 51    | .00         | .0                | 0   | 0 | .000 | 0     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| CPUTEST                                                        | WLCPU            | 953647                                                                                     | B      | 02 51    | .00         | .0                | 0   | 0 | .000 | 0     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| CPUTEST                                                        | WLCPU            | 953648                                                                                     | B      | 02 51    | .00         | .0                | 0   | 0 | .000 | 0     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| CPUTEST                                                        | WLCPU            | 953649                                                                                     | B      | 02 51    | .00         | .0                | 0   | 0 | .000 | 0     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| CPUTEST                                                        | WLCPU            | 953650                                                                                     | B      | 02 51    | .00         | .0                | 0   | 0 | .000 | 0     | 0             | 0                 | 0                | 0                 | 0            | 0               | 0             |
| Job Name                                                       |                  | -- Job name                                                                                |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| User Name/Thread                                               |                  | -- User name or secondary thread identifier                                                |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| Job Number                                                     |                  | -- Job number                                                                              |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| Typ                                                            |                  | -- Job type                                                                                |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| Pl                                                             |                  | -- Pool that the job ran in                                                                |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| Pty                                                            |                  | -- Priority of the job                                                                     |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| CPU Util                                                       |                  | -- Percentage of available CPU time used by the job. This is the average of all processors |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| DB Cpb Util                                                    |                  | -- Percentage of database capability used by the job to perform database processing        |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| Tns                                                            |                  | -- Total number of transactions for the job                                                |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |
| Tns /Hour                                                      |                  | -- Transactions per hour                                                                   |        |          |             |                   |     |   |      |       |               |                   |                  |                   |              |                 |               |

```

Rsp                -- Average interactive transaction response time in seconds
Sync Disk I/O      -- Number of synchronous disk operations (reads and writes)
Async Disk I/O     -- Number of asynchronous disk operations (reads and writes)
Logical Disk I/O   -- Number of logical disk operations (Get, Put, Upd, Other)
Cmn I/O            -- Number of communications operations (Get, Put)
PAG Fault          -- Number of faults involving the Process Access Group
Arith Ovrflw       -- Number of arithmetic overflow exceptions
Perm Write         -- Number of permanent writes
.
.
.
Column            Total            Average
-----
CPU Util          98.740 *
DB Cpb Util       82.3
Tns               2,099
Tns /Hour         1,043
Rsp              1.610
Sync Disk I/O     304,001
Async Disk I/O    1,906,898
Logical Disk I/O  6,257,174
Cmn I/O           0
PAG Fault         0
Arith Ovrflw      3
Perm Write        1,980,564
* ---- Average based on the total elapsed time for the selected intervals

```

## Printing the performance reports:

You can print reports using the performance data that you collected.

**Note:** If your trace data and sample data are both in the current library, you can use F20 to toggle between the two Print Performance Report displays.

After you have collected your data, you must create a set of performance data files from the performance information stored in a management collection (\*MGTCOL) object. Use the Create Performance Data (CRTPFRTDA) command. After you have created the data files, you can request to print your reports.

Use the following commands to print reports for sample data that you collected with Collection Services:

- Print System Report (PRTSYSRPT)
- Print Component Report (PRTCPTRPT)
- Print Job Interval Report (PRTJOBRPT)
- Print Pool Report (PRTPOLRPT)
- Print Resource Report (PRTRSCRPT)

Use the following commands to print reports for trace data that you collected with the Start Performance Trace (STRPFRTRC) and Trace Internal (TRCINT) commands:

- Print Transaction Report (PRTTNSRPT)
- Print Lock Report (PRTLCKRPT)
- Print Job Trace Report (PRTTTCRPT)

**Note:** You must use the End Performance Trace (ENDPFRTRC) command to stop the collection of performance trace data and then optionally write performance trace data to a database file before you can print the Transaction reports.

## Related information

CL commands for performance

## Performance Report columns:

Each report includes columns of information. Look here for descriptions of that information.

>8.0 (Component) The number of times the response time was greater than 8 seconds.



| **%Write Cache Overruns**  
 | (Component) Percent of Write Cache Overruns during the collection interval.

| **----- (pgmname)**  
 | (Transaction) The transaction totals record. For example, ----- QUYLIST,. This report line  
 | occurs each time the job has an active-to-wait transaction. Totals are created for Rsp\* (response  
 | time), *CPU Secs*, and I/O counts for the transaction.

| **A-I Wait /Tns**  
 | (Transaction) The average time, in seconds, of active-to-ineligible wait time per transaction. If this  
 | value is high, it may be because the time-slice value is set too low for many of the interactive  
 | jobs. Consider increasing the time slice-value.

| **Aborts Recd**  
 | (Resource Interval) The number of frames received that contained HDLC abort indicators. This  
 | indicates that the remote equipment ended frames before they were complete.

| **Act Jobs**  
 | (Job Interval) The number of selected jobs (interactive or noninteractive, depending on the report  
 | section) that were active during the interval.

| **Act Level**  
 | (Component) Initial pool activity level.

| **Act Lvl**  
 | (System, Pool Interval) Activity level. For the Pool Activity section of the Pool Interval Report, the  
 | activity level of the pool during the interval. For the Storage Pool Utilization section of the  
 | System Report, the activity level at the time of the first sample interval.

| **Act-Inel**  
 | (System, Component) Average number of active-to-ineligible job state transitions per minute.

| **Act-Wait**  
 | (System, Component) Number of transitions per minute from active state to wait state by  
 | processes assigned to this pool.

| **ACTIVE**  
 | (Job Trace) The time the job was processing.

| **Active Devices**  
 | (System) Average number of active devices on the line.

| **Active display stations (local or remote)**  
 | (System) The number of local or remote display stations entering transactions during the  
 | measurement period.

| **Active Jobs**  
 | (Transaction) The number of interactive jobs that were active during the interval.

| **Active Jobs Per Interval**  
 | (System) Average number of jobs of this type that were active per sample interval.

| **Active K/T /Tns**  
 | (Transaction) An average think time and keying time (or the delay time between the end of one  
 | transaction and the start of the next transaction), in seconds, for the active work stations  
 | (described under Est of AWS). Active K/T /TNS delay time differs from Key/Think /TNS delay  
 | time in that any delay time greater than 600 seconds has been rounded to 600 seconds. This  
 | technique is used to reduce the effect of very casual users (those who may do intermittent work  
 | or leave their work stations for long periods of time) on the estimate of active work stations.

| **Active Wrk Stn**  
 | (Resource Interval) The number of work stations with activity.

- | **Active/Rsp**  
 | (Transaction) The time the job spends (either waiting or active) during transaction processing,  
 | while it holds an activity level.
- | **Activity level**  
 | (System) The sum of activity levels for all interactive pools that had interactive job activity  
 | running in them.
- | **Activity Level Time**  
 | (Transaction) A breakdown of the transaction time spent *ACTIVE*, waiting on a *SHORT WAIT*,  
 | and waiting on a *SEIZE/CFT* (seize conflict). The *SHORT WAIT* and *SEIZE CFT* time are included  
 | under *ACTIVITY LEVEL TIME*, because the activity-level slot is not given up during these times.  
 | Note that the seize conflict time is included in the active time, not added to it to get  
 | transaction/response time, as is the case for waiting time.
- | **Arith Ovrflw**  
 | (Component, Job Interval) The number of arithmetic overflow exceptions that occurred for the  
 | selected interactive jobs during the interval.
- | **ASP ID**  
 | (System, Resource Interval) Auxiliary storage pool identifier.
- | **ASP Rsc Name**  
 | (System, Resource) Identifies the ASP resource name to which the disk unit was allocated at  
 | collection time.
- | **Async** (System, Component, Transaction, Job Interval) The number of asynchronous disk I/O operations  
 | started by the selected interactive jobs during the interval. The job that starts the I/O operation  
 | may continue processing without having to wait for the I/O operation to complete. The I/O  
 | operation is completed by a background system test.
- | **Async DIO /Tns**  
 | (Transaction) The sum of the averages of the asynchronous DB READ, DB WRITE, NDB READ,  
 | and NDB WRITE requests (the average number of asynchronous I/O requests per transaction for  
 | the job).
- | **Async Disk I/O**  
 | (System, Component, Transaction) Number of asynchronous disk input/output operations per  
 | transaction.
- | **Async Disk I/O per Second**  
 | (Component) Average asynchronous disk I/O operations per second.
- | **Async Disk I/O Requests**  
 | (Transaction) The total number of asynchronous disk I/O requests for the given combination of  
 | priority, job type, and pool.
- | **Async I/O /Sec**  
 | (Job Interval) The average number of asynchronous disk I/O operations started per second by the  
 | job during the interval. This is calculated by dividing the asynchronous disk I/O count by the  
 | elapsed time.
- | **Async I/O Per Second**  
 | (Job Interval) The average number of asynchronous disk I/O operations started per second by the  
 | selected noninteractive jobs during the interval.
- | **Async Max**  
 | (Transaction) Listed under Average DIO/Transaction, the maximum number of asynchronous  
 | DBR, NDBR, and WRT I/O requests encountered for any single transaction by that job. If the job  
 | is not an interactive or autostart job type, the total disk I/O for the job is listed here.

| **Async Sum**

| (Transaction) Listed under Average DIO/Transaction, the sum of the averages of the  
| asynchronous DBR, NDBR, and WRT requests (the average number of asynchronous I/O requests  
| per transaction for the job).

| **Asynchronous DBR**

| (System, Job Interval, Pool Interval) The average number of asynchronous database read  
| operations on the disk per transaction for the job during the intervals. This is calculated by  
| dividing the asynchronous database read count by the transactions processed. This field is not  
| printed if the jobs in the system did not process any transactions. For the Resource Utilization  
| section of the System Report, it is the number of asynchronous database read operations per  
| second.

| **Note:** The asynchronous I/O operations are performed by system asynchronous I/O tasks.

| **Asynchronous DBW**

| (System, Job Interval) The average number of asynchronous database write operations on the disk  
| per transaction for the selected jobs during the interval. This is calculated by dividing the  
| asynchronous database write count by the transactions processed. This field is not printed if the  
| jobs in the system did not process any transactions. For the Resource Utilization section of the  
| System Report, it is the number of asynchronous database read operations per second.

| **Note:** The asynchronous I/O operations are performed by system asynchronous I/O tasks.

| **Asynchronous disk I/O per transaction**

| (System) The average number of asynchronous physical disk I/O operations per interactive  
| transaction.

| **Asynchronous NDBR**

| (System, Job Interval, Pool Interval) The average number of asynchronous nondatabase read  
| operations per transaction for the jobs in the system during the interval. This is calculated from  
| the asynchronous nondatabase read count divided by the transactions processed. This field is not  
| printed if the jobs in the system did not process any transactions. For the Resource Utilization  
| section of the System Report, it is the asynchronous nondatabase read operations per second.

| **Note:** The asynchronous I/O operations are performed by system asynchronous I/O tasks.

| **Asynchronous NDBW**

| (System, Job Interval, Pool Interval) The average number of asynchronous nondatabase write  
| operations per transaction for the jobs in the system during the interval. This is calculated from  
| the asynchronous nondatabase write count divided by the transactions processed. This field is not  
| printed if the jobs in the system did not process any transactions. For the Resource Utilization  
| section of the System Report, it is the number of asynchronous nondatabase write operations per  
| second.

| **Note:** The asynchronous I/O operations are performed by system asynchronous I/O tasks.

| **Avail Local Storage (K)**

| (Resource Interval) The number of kilobytes of free local storage in the IOP.

| **Available Storage**

| (Component) Available local storage (in bytes). The average number of bytes of available main  
| storage in the IOP. The free local storage is probably not joined because it has broken into small  
| pieces.

| **Average**

| (Transaction) The average value of the item described in the column for all transactions.

| **AVERAGE**

| (Job Trace) Averages for the fields. The entry on the AVERAGE line in the SEQUENCE column

| shows the number of STRTNS and ENDTNS pairs encountered. For an interactive job, this is the  
| number of transactions entered while the trace was on if the default STRTNS and ENDTNS  
| values were used.

| **Average Disk Activity Per Hour**  
| (Component) See Disk Arm Seek Distance

| **Average DIO/Transaction**  
| (Transaction) Seven columns of information about physical disk I/O counts. Physical I/O  
| contrasts with logical I/O shown elsewhere in these reports. A logical I/O is a request sent at the  
| program level that might result in an access to auxiliary storage (DASD). A physical I/O refers to  
| those requests that actually result in access to auxiliary storage.

- | • Synchronous DBR
- | • Synchronous NDBR
- | • Synchronous Wrt
- | • Synchronous Sum
- | • Synchronous Max
- | • Async Sum
- | • Async Max

| **Average K per I/O**  
| (Resource Interval) The average number of kilobytes transferred during each disk read or write  
| operation.

| **Average Phys I/O /Sec**  
| (Resource Interval) The average number of physical disk read and write operations per second  
| made on all disks on the system.

| **Average Reads/Sec**  
| (Resource Interval) The average number of physical disk read operations per second made on all  
| disks on the system.

| **Average Response**  
| (System) Average response time (in seconds) for interactive transactions. The Total/Average  
| interactive response time does not include transactions for DDM server jobs.

| **Average Response Time**  
| (System) Average disk response time per I/O operation.

| **Average Response Time (seconds)**  
| (System) The average interactive response time.

| **Average Service Time**  
| (System) Average disk service time per I/O operation. This is the amount of time a request  
| would take if there were no contention.

| **Average Wait Time**  
| (System) Average disk wait time per I/O operation. Normally due to contention.

| **Average Writes/Sec**  
| (Resource Interval) The average number of physical disk write operations per second made on all  
| disks on the system.

| **Avg CPU /Tns**  
| (Transaction) The average number of processing unit seconds per transaction that fell in the given  
| category.

| **Avg K/T /Tns**  
| (Transaction) The average think time and keying time (or the delay time between transaction  
| boundaries), in seconds, for the interactive jobs.

|  |                                          |                                                                                                                                                                                                                                                                                                                   |
|--|------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <b>Avg Length</b>                        |                                                                                                                                                                                                                                                                                                                   |
|  | (Lock)                                   | The average number of milliseconds a lock or seize was held.                                                                                                                                                                                                                                                      |
|  | <b>Avg Rsp (Sec)</b>                     |                                                                                                                                                                                                                                                                                                                   |
|  | (Transaction)                            | The average transaction response time in seconds.                                                                                                                                                                                                                                                                 |
|  | <b>Avg Rsp /Tns</b>                      |                                                                                                                                                                                                                                                                                                                   |
|  | (Transaction)                            | The average response per transaction (in seconds) for the transactions that fell into the given category.                                                                                                                                                                                                         |
|  | <b>Avg Rsp Time</b>                      |                                                                                                                                                                                                                                                                                                                   |
|  | (Component)                              | Average transaction response time.                                                                                                                                                                                                                                                                                |
|  | <b>Avg Sec Locks</b>                     |                                                                                                                                                                                                                                                                                                                   |
|  | (Transaction)                            | The average length of a lock in seconds attributed to interactive or noninteractive waiters.                                                                                                                                                                                                                      |
|  | <b>Avg Sec Seizes</b>                    |                                                                                                                                                                                                                                                                                                                   |
|  | (Transaction)                            | The average length of a seize in seconds attributed to interactive or noninteractive waiters.                                                                                                                                                                                                                     |
|  | <b>Avg Time per Service</b>              |                                                                                                                                                                                                                                                                                                                   |
|  | (Resource Interval)                      | The amount of time a disk arm uses to process a given request.                                                                                                                                                                                                                                                    |
|  | <b>Avg Util</b>                          |                                                                                                                                                                                                                                                                                                                   |
|  | (System, Resource Interval)              | On the Disk Utilization Summary of the Resource Report, the average percentage of available time that disks were busy. It is a composite average for all disks on the system. On the Communications Summary of the System Report, the average percentage of line capacity used during the measured time interval. |
|  | <b>Batch asynchronous I/O per second</b> |                                                                                                                                                                                                                                                                                                                   |
|  | (System)                                 | The average number of asynchronous physical disk I/O operations per second of batch processing.                                                                                                                                                                                                                   |
|  | <b>Batch CPU seconds per I/O</b>         |                                                                                                                                                                                                                                                                                                                   |
|  | (System)                                 | The average number of system processing unit seconds used by all batch jobs for each I/O performed by a batch job.                                                                                                                                                                                                |
|  | <b>Batch CPU Utilization</b>             |                                                                                                                                                                                                                                                                                                                   |
|  | (Component)                              | Percentage of available processing unit time used by the jobs that the system considers to be batch.                                                                                                                                                                                                              |
|  | <b>Note:</b>                             | For a multiple-processor system, this is the average use across all processors.                                                                                                                                                                                                                                   |
|  | <b>Batch impact factor</b>               |                                                                                                                                                                                                                                                                                                                   |
|  | (System)                                 | Batch workload adjustment for modeling purposes.                                                                                                                                                                                                                                                                  |
|  | <b>Batch permanent writes per second</b> |                                                                                                                                                                                                                                                                                                                   |
|  | (System)                                 | The average number of permanent write operations per second of batch processing.                                                                                                                                                                                                                                  |
|  | <b>Batch synchronous I/O per second</b>  |                                                                                                                                                                                                                                                                                                                   |
|  | (System)                                 | The average number of synchronous physical disk I/O operations per second of batch processing.                                                                                                                                                                                                                    |
|  | <b>BCPU / Synchronous DIO</b>            |                                                                                                                                                                                                                                                                                                                   |
|  | (Transaction)                            | The average number of batch processor unit seconds per synchronous disk I/O operation.                                                                                                                                                                                                                            |
|  | <b>Bin</b>                               | (Transaction) The number of binary overflow exceptions.                                                                                                                                                                                                                                                           |
|  | <b>Binary Overflow</b>                   |                                                                                                                                                                                                                                                                                                                   |
|  | (Component)                              | Number of binary overflows per second.                                                                                                                                                                                                                                                                            |

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <b>BMPL - Cur and Inl</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Transaction) The number of jobs currently in the activity level (beginning current multiprogramming level), and the number of jobs on the ineligible queue (beginning ineligible multiprogramming level) for the storage pool that the job ran in when the job left the wait state (the beginning of the transaction).                                                                                                                                                               |
|  | <b>Note:</b> Multiprogramming level (MPL) is used interchangeably with activity level.                                                                                                                                                                                                                                                                                                                                                                                                |
|  | <b>Bundle Wait Count</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|  | (Component) Total number of times the tasks and jobs waited for journal bundles to be written to disk.                                                                                                                                                                                                                                                                                                                                                                                |
|  | <b>Bundle Wait Pct</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|  | (Component) Percentage of time (relative to the interval elapsed time) spent waiting for journal bundles to be written to disk.                                                                                                                                                                                                                                                                                                                                                       |
|  | <b>Bundle Writes System</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|  | (Component) Number of bundle writes to internal system journals. A bundle write is a group of journal entries which are deposited together by the system.                                                                                                                                                                                                                                                                                                                             |
|  | <b>Bundle Writes User</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Component) Number of bundle writes to user-created journals. A bundle write is a group of journal entries which are deposited together by the system.                                                                                                                                                                                                                                                                                                                                |
|  | <b>Bytes per Second Received</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|  | (System) Average number of bytes received per second.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|  | <b>Bytes per Second Transmitted</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|  | (System) Average number of bytes transmitted per second.                                                                                                                                                                                                                                                                                                                                                                                                                              |
|  | <b>Bytes Recd per Sec</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Resource Interval) The average number of bytes received per second.                                                                                                                                                                                                                                                                                                                                                                                                                  |
|  | <b>Bytes Trnsmitd per Sec</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|  | (Resource Interval) The average number of bytes transmitted per second.                                                                                                                                                                                                                                                                                                                                                                                                               |
|  | <b>Category</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|  | (Transaction) A group of transactions categorized together. In the Analysis by Interactive Transaction Category, the transactions are categorized by the processing unit model. The boundary values that are used to separate the transactions are given in the <i>Avg CPU /Tns</i> column. For the Analysis by Interactive Response Time, they are categorized by their response time. For the Analysis by Interactive Key/Think Time, they are categorized by their key/think time. |
|  | <b>Cache Hit Statistics</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|  | (Component) Statistics data about use of cache including:                                                                                                                                                                                                                                                                                                                                                                                                                             |
|  | <ul style="list-style-type: none"> <li>• The percent of Device Cache Read Hit for each arm.</li> <li>• The percent of Controller Cache Read Hit for each arm.</li> <li>• The percent of efficiency of write cache</li> </ul>                                                                                                                                                                                                                                                          |
|  | <b>Device read</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|  | Device Read is the number of Device Cache Read Hits (DSDCRH) divided by number of Device Read Operations (DSDROP), expressed as a percent                                                                                                                                                                                                                                                                                                                                             |
|  | <b>Controller read</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|  | Controller Read is the number Controller Cache Read Hits (DSCCRH) divided by number of Read Commands (DSRDS), expressed as a percent.                                                                                                                                                                                                                                                                                                                                                 |
|  | <b>Write efficiency</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|  | Write efficiency is the difference between Write Commands (DSWRTS) and Device Write Operations (DSDWOP) divided by Write Commands (DSWRTS), expressed as a percent.                                                                                                                                                                                                                                                                                                                   |

|  |                                 |                                                                                                                                                                                                                                                        |
|--|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <b>EACS Read</b>                |                                                                                                                                                                                                                                                        |
|  |                                 | The percent of read hits by the Extended Adaptive Cache Simulator.                                                                                                                                                                                     |
|  | <b>EACS Resp</b>                |                                                                                                                                                                                                                                                        |
|  |                                 | The percent of response time improvement by the Extended Adaptive Cache Simulator.                                                                                                                                                                     |
|  | <b>Capped</b>                   |                                                                                                                                                                                                                                                        |
|  |                                 | (System) Indicates whether the partition was capped or uncapped at the end of each interval.                                                                                                                                                           |
|  |                                 | This column is only printed for the i5/OS partition collecting performance data.                                                                                                                                                                       |
|  | <b>Channel</b>                  |                                                                                                                                                                                                                                                        |
|  |                                 | (Resource Interval) The B-channel used by the IDLC line. (special condition)                                                                                                                                                                           |
|  | <b>Cmn</b>                      | (Job Interval) The number of communications I/O operations performed by the selected interactive jobs during the interval.                                                                                                                             |
|  | <b>Cmn I/O</b>                  |                                                                                                                                                                                                                                                        |
|  |                                 | (Component) Number of communications operations (Get, Put).                                                                                                                                                                                            |
|  | <b>Cmn I/O Per Second</b>       |                                                                                                                                                                                                                                                        |
|  |                                 | (Job Interval) The average number of communications I/O operations performed per second by the selected noninteractive jobs during the interval.                                                                                                       |
|  | <b>Collision Detect</b>         |                                                                                                                                                                                                                                                        |
|  |                                 | (Resource Interval) The number of times that the terminal equipment (TE) detected that its transmitted frame had been corrupted by another TE attempting to use the same bus.                                                                          |
|  | <b>Commit Ops</b>               |                                                                                                                                                                                                                                                        |
|  |                                 | (Component) Commit operations performed. Includes application and system-provided referential integrity commits.                                                                                                                                       |
|  | <b>Communications I/O Count</b> |                                                                                                                                                                                                                                                        |
|  |                                 | (System) Number of communications I/O operations.                                                                                                                                                                                                      |
|  | <b>Communications I/O Get</b>   |                                                                                                                                                                                                                                                        |
|  |                                 | (System) Number of communication get operations per transaction.                                                                                                                                                                                       |
|  | <b>Communications I/O Put</b>   |                                                                                                                                                                                                                                                        |
|  |                                 | (System) Number of communication put operations per transaction.                                                                                                                                                                                       |
|  | <b>Communications Lines</b>     |                                                                                                                                                                                                                                                        |
|  |                                 | (System, Component, Job Interval, Pool Interval) For the Report Selection Criteria, the list of communications lines selected to be included (SLTLINE parameter) or excluded (OMTLINE parameter). These are the communications line names you specify. |
|  | <b>Control Units</b>            |                                                                                                                                                                                                                                                        |
|  |                                 | (System, Component, Job Interval, Pool Interval) The list of control units selected to be included (SLTCTL parameter) or excluded (OMTCTL parameter). These are the controller names you specify.                                                      |
|  | <b>Count</b>                    | (Transaction, Lock) The number of occurrences of the item in the column. For example, in a lock report, it is the number of locks or seizes that occurred.                                                                                             |
|  | <b>CPU</b>                      | (Transaction) The total processing unit seconds used by the jobs with a given priority.                                                                                                                                                                |
|  | <b>CPU</b>                      | (Job Trace) The approximation of the CPU used on this trace entry. This is a calculated value based on the time used and the CPU model being run.                                                                                                      |
|  | <b>CPU /Tns</b>                 |                                                                                                                                                                                                                                                        |
|  |                                 | (Transaction, Job Interval) The amount of available processing unit time per transaction in seconds.                                                                                                                                                   |
|  | <b>CPU Model</b>                |                                                                                                                                                                                                                                                        |
|  |                                 | (System) The processing unit model number.                                                                                                                                                                                                             |



- | **CPU per I/O Async**  
| (System) CPU use per asynchronous I/O.
- | **CPU per I/O Sync**  
| (System) CPU use per synchronous I/O.
- | **CPU per Logical I/O**  
| (System) Processing unit time used for each logical disk I/O operation.
- | **CPU QM**  
| (Transaction) The simple processing unit queuing multiplier.
- | **CPU Sec**  
| (Transaction) The processing unit time used by the job in this state.
- | **CPU Sec /Sync DIO**  
| (Transaction) The ratio of CPU seconds divided by synchronous disk I/O requests for each type of job.
- | **CPU Sec Avg and Max**  
| (Transaction) The average processing unit time per transaction for the job and the largest processing unit time used for a transaction in the job. If the job is not an interactive or autostart job type, then only the total processing unit time for the job is listed under the MAX column heading.
- | **CPU Sec per Tns**  
| (Transaction) The processing unit time per transaction.
- | **CPU Seconds**  
| (System, Transaction, Component) Average processing unit seconds used per transaction. For System Summary Data, it is the total available processing unit time used by the jobs during the trace period. For Priority-Jobtype-Pool Statistics, it is the total processing unit seconds used by the jobs with a given combination of priority, job type, and pool. For Batch Job Analysis, it is the amount of available processor unit time used by the job in seconds. For Concurrent Batch Job Statistics, it is the amount of available processor unit time used by the jobs in the job set in seconds.
- | **CPU SECONDS**  
| (Job Trace) The approximate processing unit time used for the transaction.
- | **CPU seconds per transaction**  
| (System) The average processing unit seconds per transaction.
- | **CPU Util**  
| (System, Component, Transaction, Job Interval, Pool Interval, Batch Job Trace) Percentage of available processing unit time used. For multiple-processor systems, this is the total utilization divided by the number of processors.
- | **CPU Util per Transaction**  
| (Component) The result of the CPU Utilization divided by the total number of transactions for the job.
- | **CPU Utilization (Batch)**  
| The percentage of available CPU time that is used by batch jobs. This is the average of all processors.
- | **CPU Utilization (Interactive)**  
| The percentage of available CPU time that is used by interactive jobs. This is the average of all processors.
- | **CPU Utilization (Total)**  
| The percentage of available CPU time that is used by interactive and batch jobs. This is the average of all processors.

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <b>Note:</b> For uncapped partitions, the Total CPU utilization might exceed 100 percent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|  | <b>CPU/Async I/O</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|  | (Job Interval) The average number of milliseconds of processing unit time taken for each asynchronous disk I/O operation. This is calculated by dividing the milliseconds of the processing unit time the job used by the asynchronous disk I/O count.                                                                                                                                                                                                                                                                                                                                                                 |
|  | <b>CPU/Sync I/O</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|  | (Job Interval) The average number of milliseconds of processing unit time taken for each synchronous disk I/O operation. This is calculated from the milliseconds of the processing unit time used by the job divided by the synchronous disk I/O count.                                                                                                                                                                                                                                                                                                                                                               |
|  | <b>CPU/Tns</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|  | (Transaction) The average number of processing seconds per transaction for the job during the interval. This is calculated from the amount of processing unit time used divided by the number of transactions processed.                                                                                                                                                                                                                                                                                                                                                                                               |
|  | <b>Cpu/Tns (Sec)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|  | (Transaction) The number of processing unit seconds per transaction.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|  | <b>Ctl</b> (Component) Controller identifier.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|  | <b>Cum CPU Util</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|  | (Transaction) The cumulative percentage of available processing unit time used by the transactions that have an average response time per transaction equal to or less than the given category. For example, in CPU by Priority for All Jobs for Total Trace Period (System Summary Data), it is the unit time used by the jobs with a priority higher or equal to the given priority.                                                                                                                                                                                                                                 |
|  | <b>Cum Pct Tns</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|  | (Transaction) Cumulative CPU percent per transaction. For system summary data, it is the cumulative CPU percentage of all transactions that have an average response time per transaction equal to or less than the given category. For Interactive Program Transactions Statistics, it is the cumulative CPU percentage of all transactions through the listed program. For Job Statistics section, it is the cumulative CPU percentage of total transactions through the listed job. For Interactive Program Statistics section, it is the cumulative CPU percentage of all transactions through the listed program. |
|  | <b>Cum Util</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|  | (System) Cumulative CPU use (a running total).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|  | <b>Note:</b> This is taken from the individual jobs and may differ slightly from the total processing unit use on the workload page.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|  | <b>Cur Inl MPL</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|  | (Transaction) The number of jobs waiting for an activity level (ineligible) in the storage pool.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|  | <b>Cur MPL</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|  | (Transaction) The number of jobs holding an activity level in the storage pool.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|  | <b>Current User</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|  | (Job) The user under which the job was running at the end of each interval.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|  | <b>DASD Ops/Sec</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|  | (Component) Disk operations per second.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|  | <b>DASD Ops Per Sec Reads</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|  | (Resource) Number of reads per second                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|  | <b>DASD Ops Per Sec Writes</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|  | (Resource) Number of writes per second                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

|  |                                                     |                                                                                                                                                                                                                                                                                                           |
|--|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <b>Datagrams Received</b>                           |                                                                                                                                                                                                                                                                                                           |
|  | (Component)                                         | The total number of input datagrams received from interfaces. This number includes those that were received in error.                                                                                                                                                                                     |
|  | <b>DB</b>                                           | (Job Trace) The number of physical database reads that occurred for the entry.                                                                                                                                                                                                                            |
|  | <b>DB Cpb Util</b>                                  |                                                                                                                                                                                                                                                                                                           |
|  | (Component)                                         | The percentage of database capability that is used to perform database processing.                                                                                                                                                                                                                        |
|  | <b>DB Fault</b>                                     |                                                                                                                                                                                                                                                                                                           |
|  | (System, Component)                                 | Average number of database faults per second.                                                                                                                                                                                                                                                             |
|  | <b>DB Pages</b>                                     |                                                                                                                                                                                                                                                                                                           |
|  | (System, Component)                                 | Average number of database pages read per second.                                                                                                                                                                                                                                                         |
|  | <b>DB Read</b>                                      |                                                                                                                                                                                                                                                                                                           |
|  | (Transaction)                                       | When listed in Physical I/O Counts column, it is the number of database read requests while the job was in that state. When listed in the Sync Disk I/O Rqs/Tns column, it is the average number of synchronous database read requests per transaction.                                                   |
|  | <b>DB READS</b>                                     |                                                                                                                                                                                                                                                                                                           |
|  | (Job Trace)                                         | The number of physical database reads that occurred.                                                                                                                                                                                                                                                      |
|  | <b>DB Write</b>                                     |                                                                                                                                                                                                                                                                                                           |
|  | (Transaction)                                       | When listed in the Sync Disk I/O Rqs/Tns column, it is the average number of synchronous database write requests per transaction.                                                                                                                                                                         |
|  | <b>DB Wrt</b>                                       |                                                                                                                                                                                                                                                                                                           |
|  | (Transaction)                                       | When listed in the Physical I/O Counts column, it is the number of database write requests while the job was in that state. When listed in the Synchronous Disk I/O Counts column, it is the number of synchronous database write requests per transaction.                                               |
|  | <b>DDM I/O</b>                                      |                                                                                                                                                                                                                                                                                                           |
|  | (Component, Job Interval)                           | The number of logical database I/O operations for a distributed data management (DDM) server job.                                                                                                                                                                                                         |
|  | <b>DDM Svr Wait /Tns</b>                            |                                                                                                                                                                                                                                                                                                           |
|  | (Transaction)                                       | The average time, in seconds, that a source distributed data management (DDM) server job spent waiting for the target system to respond to a request for data per transaction. This value includes line time and time spent by the target system responding to the request for data.                      |
|  | <b>Dec</b>                                          | (Transaction) The number of decimal overflow exceptions.                                                                                                                                                                                                                                                  |
|  | <b>Decimal Data</b>                                 |                                                                                                                                                                                                                                                                                                           |
|  | (Component)                                         | Data exception count per second. A data exception occurs when data that is not valid is detected by arithmetic instructions. Examples are signs or digit codes that are not valid in decimal instructions, or an insufficient number of farthest left zeros in multiply instructions.                     |
|  | <b>Decommit Ops</b>                                 |                                                                                                                                                                                                                                                                                                           |
|  | (Component)                                         | Decommit operations performed. Includes application and system-provided referential integrity decommits.                                                                                                                                                                                                  |
|  | <b>Decimal Overflow</b>                             |                                                                                                                                                                                                                                                                                                           |
|  | (Component)                                         | Number of decimal overflows per second.                                                                                                                                                                                                                                                                   |
|  | <b>Description</b>                                  |                                                                                                                                                                                                                                                                                                           |
|  | (Component)                                         | More detailed description of the exception type.                                                                                                                                                                                                                                                          |
|  | <b>Detected Access Transmission Error (DTSE) In</b> |                                                                                                                                                                                                                                                                                                           |
|  | (Resource Interval)                                 | The number of times the network termination 1 (NT1) end point notified the terminal equipment (TE) of an error in data crossing the ISDN U interface from the line transmission termination (LT) to the NT1 end point. The NT1 end point reports the errors to the TE through the maintenance channel S1. |

| **Detected Access Transmission Error (DTSE) Out**  
 | (Resource Interval) The number of times the network termination 1 (NT1) end point notified the  
 | terminal equipment (TE) of an error in data crossing the ISDN U interface from the NT1 end  
 | point to the LT. The NT1 end point reports the errors to the TE through the maintenance channel  
 | S1.

| **Device**  
 | (Component) Device identifier.

| **DIO/Sec Async**  
 | (System) Number of asynchronous I/O operations per second.

| **DIO/Sec Sync**  
 | (System) Number of synchronous I/O operations per second.

| **Disk Arm Seek Distance**  
 | (Component) Average seek distance distributions per hour:

|  |                |                                                  |
|--|----------------|--------------------------------------------------|
|  | <b>0</b>       | Number of zero seeks                             |
|  | <b>1/12</b>    | Number of seeks between 0 and 1/12 of the disk   |
|  | <b>1/6</b>     | Number of seeks between 1/12 and 1/6 of the disk |
|  | <b>1/3</b>     | Number of seeks between 1/6 and 1/3 of the disk  |
|  | <b>2/3</b>     | Number of seeks between 1/3 and 2/3 of the disk  |
|  | <b>&gt;2/3</b> | Number of seeks greater than 2/3 of the disk     |

| **Disk Arms**  
 | (System) The number of disk arms for this IOP.

| **Disk Capacity**  
 | (Component) Average amount of disk space used or available.

|  |                |                                          |
|--|----------------|------------------------------------------|
|  | <b>MB</b>      | Millions of bytes available on the disk. |
|  | <b>Percent</b> | Percent of space available on the disk.  |

| **Disk Controllers**  
 | (System) The number of disk storage controllers for this IOP.

| **Disk Feature**  
 | (System) The type of disk (9332, 9335, and so on).

| **Disk I/O Async**  
 | (System, Component) Total number of asynchronous disk I/O operations.

| **Disk I/O Logical**  
 | (Component) The number of logical disk operations, such as gets and puts.

| **Disk I/O per Second**  
 | (System) Average number of physical disk I/O operations per second.

| **Disk I/O Reads /Sec**  
 | (Resource Interval) The average number of disk read operations per second by the disk IOP.

| **Disk I/O Requests**  
 | (Transaction) The total number of synchronous and asynchronous disk I/O requests issued by the  
 | jobs during the trace period.

| **Disk I/O Sync**  
 | (System, Component) Total number of synchronous disk I/O operations.

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <b>Disk I/O Writes /Sec</b>                                                                                                                                                                                                                                                                                                                                                                                                                            |
|  | (Resource Interval) The average number of disk write operations per second by the disk IOP.                                                                                                                                                                                                                                                                                                                                                            |
|  | <b>Disk IOPs</b>                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|  | (System) The number of disk IOP controllers.                                                                                                                                                                                                                                                                                                                                                                                                           |
|  | <b>Disk mirroring</b>                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|  | (System) Indicates whether disk mirroring is active.                                                                                                                                                                                                                                                                                                                                                                                                   |
|  | <b>Disk Space Used</b>                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|  | (Resource Interval) The total disk space used in gigabytes for the entire system.                                                                                                                                                                                                                                                                                                                                                                      |
|  | <b>Disk transfer size (KB)</b>                                                                                                                                                                                                                                                                                                                                                                                                                         |
|  | (System) The average number of kilobytes transferred per disk operation.                                                                                                                                                                                                                                                                                                                                                                               |
|  | <b>Disk utilization</b>                                                                                                                                                                                                                                                                                                                                                                                                                                |
|  | (System) The fraction of the time interval that the disk arms were performing I/O operations.                                                                                                                                                                                                                                                                                                                                                          |
|  | <b>Dsk CPU Util</b>                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|  | (System, Resource Interval) The percentage of CPU used by the disk unit.                                                                                                                                                                                                                                                                                                                                                                               |
|  | <b>Dtgm Req Transm Dscrd</b>                                                                                                                                                                                                                                                                                                                                                                                                                           |
|  | (Component) The percentage of IP datagrams that are discarded because of the following reasons:                                                                                                                                                                                                                                                                                                                                                        |
|  | • No route was found to transmit the datagrams to their destination.                                                                                                                                                                                                                                                                                                                                                                                   |
|  | • Lack of buffer space.                                                                                                                                                                                                                                                                                                                                                                                                                                |
|  | <b>Dtgm Req for Transm Tot</b>                                                                                                                                                                                                                                                                                                                                                                                                                         |
|  | (Component) The total number of IP datagrams that local IP user-protocols supplied to IP in requests for transmission.                                                                                                                                                                                                                                                                                                                                 |
|  | <b>Elapsed Seconds</b>                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|  | (Transaction, Component) The elapsed time in seconds. For the Batch Job Analysis section of the Transaction Report, it is the number of seconds elapsed from when the job started to when the job ended. For the Concurrent Batch Job Statistics section of the Transaction Report, it is the total elapsed time of all jobs in that job set.                                                                                                          |
|  | <b>Elapsed Time</b>                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|  | (Job Interval) The amount of time (minutes and seconds) for which the job existed during the interval. This is the same as the interval length unless the job started or ended during the interval, in which case it is less.                                                                                                                                                                                                                          |
|  | <b>Elapsed Time--Seconds</b>                                                                                                                                                                                                                                                                                                                                                                                                                           |
|  | (Transaction) Shows the time spent by the job, in the following columns:                                                                                                                                                                                                                                                                                                                                                                               |
|  | <b>Long Wait</b>                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|  | Elapsed times in the state (such as waiting for the next transaction or lock-wait time).                                                                                                                                                                                                                                                                                                                                                               |
|  | <b>Active/Rsp</b>                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|  | During transaction processing, the time the job spends (either waiting or active) while it holds an activity level. At the end of a transaction (on the transaction totals line), this is the time the job spent processing the transaction in an activity level, for long waits caused by locks, and in the ineligible state.                                                                                                                         |
|  | <b>Inel Wait</b>                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|  | The time the job spent in the ineligible wait state waiting for an activity level.                                                                                                                                                                                                                                                                                                                                                                     |
|  | <b>EM3270 Wait /Tns</b>                                                                                                                                                                                                                                                                                                                                                                                                                                |
|  | (Transaction) The average, in seconds, of the time spent waiting on the host system communications for Systems Network Architecture (SNA) and binary synchronous communications (BSC) 3270DE per transaction. Program logic is required to determine if the emulation program is communicating with the display or the host processing unit. Because there are requirements on event-wait processing, not all transition combinations can be detected. |

|  |                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <b>ENTRY</b>                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|  |                                      | (Job Trace) The instruction in the program where the program was given control.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|  | <b>EORn</b>                          | (Transaction) Listed in the Wait Code column, End of response time for transaction n. These codes are in the wait code column, but they are not wait codes. They indicate transaction boundary trace records.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|  | <b>EOTn</b>                          | (Transaction) Listed in the Wait Code column, End of transaction for transaction for type n. These codes are in the wait code column, but they are not wait codes. They indicate transaction boundary trace records.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|  | <b>Estimated Exposr AP Not Jrnld</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|  |                                      | (Component) System-estimated access path recovery time exposure in minutes if no access paths were being journaled by the system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|  | <b>Estimated Exposr Curr System</b>  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|  |                                      | (Component) System-estimated access path recovery time exposure in minutes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|  | <b>Est Of AWS</b>                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|  |                                      | (Transaction) An estimate of the number of active work stations for the trace period or interval. Any delay time greater than 600 seconds has been rounded to 600 seconds. This technique is used to reduce the effect of very casual users (those who may do intermittent work or leave their work stations for long periods of time) on the estimate of active work stations.                                                                                                                                                                                                                                                                                                                                               |
|  | <b>Event Wait /Tns</b>               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|  |                                      | (Transaction) The average time, in seconds, of the event-wait time per transaction. Often requests made by a job that runs on the system are made to asynchronous jobs. These asynchronous jobs use an event to signal completion of the request back to the requester. The event-wait time is the time the requesting job waits for such a signal.                                                                                                                                                                                                                                                                                                                                                                           |
|  | <b>EVT</b>                           | (Transaction) Listed in the Wait Code column, Event Wait. This is a long wait that occurs when waiting on a message queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|  | <b>Exception Type</b>                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|  |                                      | (Component) Type of program exception that results from the internal microprogram instructions being run in internal microprogram instructions procedure. Because these exceptions are monitored at a low level within the system, it is difficult to associate these exceptions with specific end-user operations. The counts are meaningful when the processing unit time required to process them affects system performance. A variation in the counts may indicate a system change that could affect performance. For example, a large variation in seize or lock counts may indicate a job scheduling problem or indicate that contention exists between an old application and a new one that uses the same resources. |
|  | <b>Note:</b>                         | To see the seize and lock counts, you should collect the trace data by using the Start Performance Trace (STRPFRTRC) command. Run the Print Transaction Report (PRTTNSRPT) to list the objects and jobs that are holding the locks.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|  | <b>Exceptional wait</b>              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|  |                                      | (System) The average exceptional wait time, in seconds, per transaction. An <i>exceptional wait</i> is that portion of internal response time that cannot be attributed to the use of the processor and disk. An exceptional wait is caused by contention for internal resources of the system, for example, waiting for a lock on a database record.                                                                                                                                                                                                                                                                                                                                                                         |
|  | <b>Constant</b>                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|  |                                      | The portion of exceptional wait time held constant as throughput increases.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|  | <b>Variable</b>                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|  |                                      | The portion of exceptional wait time that varies as throughput increases.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|  | <b>Excp</b>                          | (Component, Transaction) For the Component Report, it is the total number of program exceptions that occurred per second. For the Transaction Report, a Y in this column means that                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

| the transaction had exceptions. The types of exceptions that are included are process access group  
| exceptions, and decimal, binary, and floating point overflow. See the Transition Report to see  
| which exceptions the transaction had.

| **Excp Wait**

| (Transaction) The amount of exceptional wait time for the jobs in the job set in seconds.

| **Excp Wait /Tns**

| (Transaction) The average exceptional wait time, in seconds, per transaction. This value is the  
| sum of those waits listed under the Exceptional Wait Breakdown by Job Type part.

| **Excp Wait Sec**

| (Transaction) The total amount of exceptional wait time in seconds for the job.

| **Excs ACTM /Tns**

| (Transaction) The average time, in seconds, of the excess activity level time per transaction (for  
| example, time spent in the active state but not using the processing unit). If enough activity levels  
| are available and there is plenty of interactive work of higher priority to do, a job waits longer  
| for processing unit cycles. If the value is greater than .3, look at jobs that correspond to particular  
| applications for more information. By looking at these jobs, you might be able to determine  
| which application's jobs are contributing most to this value. Use the Transaction and Transition  
| Reports for these jobs for additional information. The formula for excessive activity-level time is  
| shown below:

| Active Time - [  
| (multiplier X CPU X Beginning Activity Level) +  
| (Number of synchronous disk I/O operations X .010)]

| **Note:** If the beginning activity level is greater than 1, the multiplier equals 0.5. If the beginning  
| activity level is any other value, the multiplier equals 1.

| **EXIT** (Job Trace) The instruction number in the program where the program gave up control.

| **Expert Cache**

| (System, Component) Directs the system to determine which objects or portions of objects should  
| remain in a shared main storage pool based on the reference patterns of data within the object.  
| Expert cache uses a storage management tuner, which runs independently of the system dynamic  
| tuner, to examine overall paging characteristics and history of the pool. Some values that you  
| might see in this column are associated with the Work with Shared Pools (WRKSHRPOOL)  
| command:

- | • 0=\*FIXED, which indicates the system does not dynamically adjust the paging characteristics of  
| the storage pool. The system uses default values.
- | • 3=\*CALC, which indicates the system dynamically adjusts the paging characteristics of the  
| storage pool for optimum performance.

| **Exposed AP System Journalled**

| (Component) The number of exposed access paths currently being journaled by the system.

| **Exposed AP System Not Journalled**

| (Component) The number of exposed access paths currently not being journaled by the system.

| **/F** (System, Resource Interval) The line speed of the protocol reported as full duplex. This indicator  
| applies to the line speeds for an Ethernet (ELAN) token-ring (TRLAN) line, or an asynchronous  
| transfer mode line.

| **Far End Code Violation**

| (Resource Interval) The number of unintended code violations detected by the network  
| termination 1 (NT1) end point for frames transmitted to the NT1 end point on the interface for  
| the T reference point. The NT1 end point reports a violation to the termination equipment (TE)  
| through the maintenance channel S1.



- | **Faults** (System) A value that represents the total page faults that occurred for each job type or job priority during the collection. This is the same value as shown in the JBTFLLT field of the QAPMJOB or QAPMJOB file.
- | **File** (Transaction) The file that contains the object.
- | **Flp** (Transaction) The number of floating point overflow exceptions.
- | **Flp Overflow**  
(Component) Number of floating point overflows per second.
- | **Frame Retry**  
(Resource Interval) The number of attempts to retransmit a frame to a remote controller.
- | **Frames Received Pct Err**  
(Resource Interval) The percentage of frames received in error. Errors can occur when the host system has an error or cannot process received data fast enough.
- | **Frames Received Total**  
(Resource Interval) The total number of frames received including frames with errors and frames that are not valid.
- | **Frames Transmitted Pct Err**  
(Resource Interval) The percentage of frames retransmitted due to error.
- | **Frames Transmitted Total**  
(Resource Interval) The total number of frames transmitted.
- | **FULL CLS**  
(Job Trace) The number of full closes for all types of files.
- | **FULL OPN**  
(Job Trace) The number of full opens for all types of files.
- | **FUNCTION**  
(Job Trace) This causes the trace entry to be recorded. The possible trace entries are as follows:

Table 6.

| Function ID | Description                                     |
|-------------|-------------------------------------------------|
| DATA        | Data trace record                               |
| CALL        | Call external                                   |
| XCTL        | Transfer control                                |
| EVENT       | Event handler invocation                        |
| EXTXHINV    | External exception handler invocation           |
| INTXHINV    | Internal exception handler invocation           |
| INTXHRET    | Return from internal exception handler          |
| INVEXIT     | Invocation exit                                 |
| RETURN      | Return external                                 |
| ITRMXRSG    | Invocation ended due to resignaling exception   |
| EXTXHRET    | Return external or from a procedure instruction |
| PTRMTTP     | Termination phase end                           |
| PTRMUNX     | End process due to an unhandled exception       |
| NOTUSED     | This type is a non-valid trace type             |
| ITERM       | Invocation ended                                |
| CANCLINV    | Cancel invocation instruction                   |

| **Functional Areas**  
 | (System, Component, Transaction, Job Interval, Pool Interval) For Report Selection Criteria, the  
 | list of functional areas selected to be included (SLTFCNARA parameter) or excluded  
 | (OMTFCNARA parameter).

| **/H** (System, Resource Interval) The line speed of the protocol reported as half duplex. This indicator  
 | applies to the line speeds for an Ethernet (ELAN) token-ring (TRLAN) line, or an asynchronous  
 | transfer mode line.

| **HDW** (Transaction) Listed in the Wait Code column, Hold Wait (job suspended or system request). The  
 | job released a lock it had on the object named on the next detail line of the report (OBJECT --).  
 | The job that was waiting for the object is named on this line (WAITER --) along with the amount  
 | of time the job spent waiting for the lock to be released.

| **High Srv Time**  
 | (Resource Interval) The highest average service time in seconds for a disk arm in the system.

| **High Srv Unit**  
 | The disk arm with the highest service time.

| **High Util**  
 | (Resource Interval) The percentage of use for the disk arm that has the highest utilization.

| **High Util Unit**  
 | (Component, Resource Interval) The disk arm with the highest utilization.

| **High Utilization Disk**  
 | (Component) Percent of utilization of the most utilized disk arm during this interval.

| **High Utilization Unit**  
 | (Component) Disk arm that had the most utilization during this interval.

| **Holder Job Name**  
 | (Transaction) The name of the job that held the object.

| **Holder Number**  
 | (Transaction) The number of the job that held the object.

| **Holder Pool**  
 | (Transaction) The pool that held the job while it was running.

| **Holder Pty**  
 | (Transaction) The priority of the holder's job.

| **Holder Type**  
 | (Transaction) The type and subtype of the holder's job.

| **Holder User Name**  
 | (Transaction) The name of the user that held the object.

| **Holder's Job Name**  
 | (Lock) The name of the job holding the lock.

| **I Frames Recd per Sec**  
 | (Resource Interval) The number of information frames received per second.

| **I Frames Trnsmitd per Sec**  
 | (Resource Interval) The number of information frames transmitted per second.

| **I/O Wait**  
 | (Resource Interval) The amount of time in which a given I/O request is ready to be processed,  
 | but the disk arm is not yet available to perform the request.

|  |                                   |                                                                                                                                                                                                                                                                                                                                                                             |
|--|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <b>ICMP Messages Error</b>        |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Component)                       | This is the number of Internet Control Message Protocol (ICMP) messages that the entity received but determined that the messages had errors or are messages that the entity did not send due to problems.                                                                                                                                                                  |
|  | <b>ICMP Messages Received</b>     |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Component)                       | This is the total number of Internet Control Message Protocol (ICMP) messages that the entity received.                                                                                                                                                                                                                                                                     |
|  | <b>ICMP Messages Sent</b>         |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Component)                       | This is the total number of Internet Control Message Protocol (ICMP) messages that the entity attempted to send.                                                                                                                                                                                                                                                            |
|  | <b>Incoming Calls Pct Retry</b>   |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Resource Interval)               | The percentage of incoming calls that were rejected by the network.                                                                                                                                                                                                                                                                                                         |
|  | <b>Incoming Calls Total</b>       |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Resource Interval)               | The total number of incoming call attempts.                                                                                                                                                                                                                                                                                                                                 |
|  | <b>Inel Time A-I/W-I</b>          |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Transaction)                     | The amount of time the job spent in the ineligible state, either coming from time slice end (active-to-ineligible) or from the wait state (wait-to-ineligible).                                                                                                                                                                                                             |
|  | <b>Inel Wait</b>                  |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Transaction)                     | Listed in the Elapsed Time--Seconds column, the time the job spent in the ineligible wait state waiting for an activity level.                                                                                                                                                                                                                                              |
|  | <b>Int Feat Util</b>              |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Component)                       | The percentage of Interactive Feature that is used by all jobs.                                                                                                                                                                                                                                                                                                             |
|  | <b>Inter CPU Utilization</b>      |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Component)                       | Percentage of available processing unit time used by the jobs that the system considers to be interactive.                                                                                                                                                                                                                                                                  |
|  | <b>Note:</b>                      | For a multiple-processor system, this is the average use across all processors.                                                                                                                                                                                                                                                                                             |
|  | <b>INV</b>                        | (Job Trace) The call level of the program.                                                                                                                                                                                                                                                                                                                                  |
|  | <b>IOP</b>                        | (Component) Input/output processor (IOP) Resource name and model number for each communications IOP, DASD IOP, local workstation IOP, and multifunction IOP. Communications IOP is the percent of CPU used in the IOP. The percent does not necessarily mean that the IOP is doing any data transfers. Some of the percent can be attributed to overhead of an active line. |
|  | <b>IOP Name/Line</b>              |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (System, Resource Interval)       | Input/output (IOP) processor resource name and model number line.                                                                                                                                                                                                                                                                                                           |
|  | <b>IOP Name(Model)</b>            |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Resource Interval)               | The input/output processor (IOP) identification and the model number in parentheses.                                                                                                                                                                                                                                                                                        |
|  | <b>IOP Name</b>                   |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (System, Component)               | Input/Output processor (IOP) resource name.                                                                                                                                                                                                                                                                                                                                 |
|  | <b>IOP Name Network Interface</b> |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Resource Interval)               | The IOP name of the network interface.                                                                                                                                                                                                                                                                                                                                      |
|  | <b>IOP Processor Util Comm</b>    |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Component, Resource)             | Utilization of IOP due to communications activity.                                                                                                                                                                                                                                                                                                                          |
|  | <b>IOP Processor Util LWSC</b>    |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Component, Resource)             | Utilization of IOP due to local workstation activity.                                                                                                                                                                                                                                                                                                                       |
|  | <b>IOP Processor Util DASD</b>    |                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Component, Resource)             | Utilization of IOP due to DASD activity.                                                                                                                                                                                                                                                                                                                                    |

- | **IOP Processor Util Total**  
 | (Component, Resource Interval) The total percent of utilization for each local workstation, disk,  
 | and communications IOP.
- | **IOP Util**  
 | (System) For the Disk Utilization section of the System Report, it is the percentage of utilization  
 | for each input/output processor (IOP).
- | **Note:** For the multifunction I/O processors, this is utilization due to disk activity only, not  
 | communications activity. For the System Model Parameter section it is the fraction of the  
 | time interval the disk IOP was performing I/O operations.
- | **Itv End**  
 | (Component, Transaction, Job Interval, Pool Interval, Resource Interval) The time (hour and  
 | minute) when the data was collected. For the Exception Occurrence Summary and Interval  
 | Counts of the Component Report, it is the ending time for the sample interval in which  
 | Collection Services recorded the exception.
- | **Job Maximum A-I**  
 | (Pool Interval) The highest number of active-state to ineligible-state transitions by a selected job  
 | in the pool or subsystem.
- | **Job Maximum A-W**  
 | (Pool) The highest number of active-to-wait state transitions by a selected job in the pool or  
 | subsystem.
- | **Job Maximum CPU Util**  
 | (Pool Interval) The highest percentage of available processing unit time used by a selected job in  
 | the pool or subsystem.
- | **Job Maximum Phy I/O**  
 | (Pool Interval) The highest number of physical disk input and output operations by a selected job  
 | in the pool or subsystem.
- | **Job Maximum Rsp**  
 | (Pool Interval) The highest response time in seconds per transaction by a selected job in the pool  
 | or subsystem. The response time is the amount of time spent waiting for and using the resources  
 | divided by the number of transactions.
- | **Job Maximum Tns**  
 | (Pool Interval) The highest number of transactions by a selected job in the pool or subsystem.
- | **Job Maximum W-I**  
 | (Pool Interval) The highest number of wait-state to ineligible-state transitions by a selected job in  
 | the pool or subsystem.
- | **Job Name**  
 | (Component, Transaction, Job Interval, Batch Job Trace) Name of the job. In the Job Summary  
 | Report of the Transaction Report, a job (identical job name, user name, and job number) appears  
 | multiple times in this list if the job uses the system Reroute Job (RRTJOB) command.
- | **Job Number**  
 | (Component, Transaction, Job Interval, Batch Job Trace) The number of the job which the  
 | summary line describes. In the Transaction Report, an asterisk (\*) before the job number indicates  
 | the job signed on during the measurement period. An asterisk (\*) after the job number indicates  
 | the job signed off during the measurement period.
- | **Job Pty**  
 | (Batch Job Trace) Priority of the job.
- | **Job Set**  
 | (Transaction) The number of job sets is the number of batch jobs that could be active at any time

during the trace period. If two jobs run sequentially, they show up as two jobs in the same job set. If two jobs run concurrently, they show up in two different job sets.

## Job Type

(All Reports except where noted for the Transaction Report) Job type and subtype. Possible job type values include the following:

**A** Autostart

**B** Batch

**BD** Batch immediate (Transaction only)

**Note:** The batch immediate values are shown as BCI on the Work with Active Job display and as BATCHI on the Work with Subsystem Job display.

**BE** Batch evoke (Transaction only)

**BJ** Batch pre-start job (Transaction only)

**C** Programmable workstation application server, which includes 5250 emulation over APPC and System i Access host servers running either APPC or TCP/IP. A job is reported as a System i Access server if any of the following items are true:

- Incoming APPC evoke requests one of the server program names. This also applies to the pre-started jobs for the QSERVER, QCMN, and QSYSWRK subsystems that are already waiting for the named program.
- Incoming IP port number corresponds to one of the service name-description-port-numbers. This also applies to the pre-started jobs for the QSERVER, QCMN, and QSYSWRK subsystems that are already waiting for the assigned IP port number.
- Incoming IPX socket number corresponds to one of the service name-description-port-numbers. This also applies to the pre-started jobs for the QSERVER, QCMN, and QSYSWRK subsystems that are already waiting for the assigned IPX port number.
- Incoming 5250 display emulation jobs that come from APPC data streams sent by 5250 emulation under OS/2<sup>®</sup> Communications Manager or WARP equivalent.

**D** Target distributed data management (DDM) server

**I** Interactive. Interactive includes twinaxial data link control (TDLC), 5250 remote workstation, and 3270 remote workstation. For the Transaction Report, this includes twinaxial data link control (TDLC), 5250 remote workstation, 3270 remote workstation, SNA pass-through, and 5250 Telnet.

**L** Licensed Internal Code task

**M** Subsystem monitor

**P** SNA pass-through and 5250 Telnet pass-through. On the Transaction Report, these jobs appear as I (interactive).

**R** Spool reader

**S** System

**W** Spool writer, which includes the spool write job, and if Advanced Function Printing<sup>™</sup> (AFP<sup>™</sup>) is specified, the print driver job.

**WP** Spool print driver (Transaction only)

**X** Start system job

Possible job subtype values include the following:

**D** Batch immediate job

**E** Evoke (communications batch)

|  |          |                                                                |
|--|----------|----------------------------------------------------------------|
|  | <b>J</b> | Pre-start job                                                  |
|  | <b>P</b> | Print driver job                                               |
|  | <b>T</b> | Multiple requester terminal (MRT) (System/36 environment only) |
|  | <b>3</b> | System/36                                                      |

Noninteractive job types include:

- | • Autostart
- | • Batch
- | • Evoke
- | • System i Access-Bch
- | • Server
- | • Spool
- | • Distributed data management (DDM) server

Special interactive job categories include:

- | • Interactive
- | • Multiple requester terminal (MRT)
- | • Pass-through
- | • System/36

| **Jobs** (System, Component, Transaction, Pool Interval, Job Interval) The jobs you specify. The format of the entries is jobnumber/username/jobname. For the Report Selection Criteria report, it is the list of jobs selected to be included (SLTJOB parameter) or excluded (OMTJOB parameter). This does not include jobs selected by using the STLFCNARA or OMTFCNARA parameter.

| **K per I/O** (System, Resource Interval) The average number of kilobytes (1024 bytes) read or written for each disk I/O operation.

| **K/T /Tns Sec** (Transaction) The average delay time, or time spent keying and thinking between transactions for the job, in seconds. The value represents the time interval between active-to-wait and wait-to-active or wait-to-ineligible job state transitions.

| **KB per I/O Read** (Resource Interval) The average number of kilobytes (1 KB equals 1024 bytes) transferred per read operation.

| **KB per I/O Write** (Resource Interval) The average number of kilobytes (1024 bytes) transferred per write operation.

| **KB Received/Second** (System, Component) The total number of kilobytes (1024) received per second on the specified interface when it was active on the selected intervals, which includes framing characters.

| **KB Transmitted/Second** (System, Component) The total number of kilobytes (1024) transmitted per second from the specified interface when it was active on the selected intervals, which includes framing characters.

| **KBytes Transmitted IOP** (Component, Resource Interval) Total kilobytes transmitted from an IOP to the system across the bus.

| **KBytes Transmitted System** (Component, Resource Interval) Total kilobytes transmitted to the IOP from the system across the bus.

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <b>Key/Think</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|  | (Transaction) The amount of time spent waiting for the work station user by the program.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|  | <b>Key/Think /Tns</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|  | (Transaction) The average think time and keying time (or the delay time between transaction boundaries), in seconds, for the interactive jobs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|  | <b>L</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|  | (Lock) Whether this is a lock or seize conflict. The column contains an L if lock, blank if seize.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|  | <b>LAPD Pct Frames Recd in Error</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|  | (Resource Interval) The percentage of frames received in error (applies to D-channel only). Errors can occur when the host system has an error or cannot process received data fast enough.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|  | <b>LAPD Pct Frames Trnsmitd Again</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|  | (Resource Interval) The percentage of frames retransmitted due to error (applies to D-channel only).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|  | <b>LAPD Total Frames Recd</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|  | (Resource Interval) The total number of frames received including frames with errors and frames that are not valid (applies to D-channel only).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|  | <b>LAPD Total Frames Trnsmitd</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|  | (Resource Interval) The total number of frames transmitted (applies to D-channel only).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|  | <b>Last 4 Programs in Invocation Stack</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|  | (Transaction) The last four programs in the program stack. For example, at the start of a transaction (such as when the work station operator presses the Enter key), you see the program names QT3REQIO, QWSGET, and the program that issued a read operation. At the end of the transaction (such as when the program writes to the display), you see QT3REQIO, QWSPUT, and the program that wrote the display. Usually, the third or fourth program in the stack is the program shown in the transaction summary PGMNAME data. However, if the <i>Wait Code</i> column has a value, the program in the column labeled <i>Last</i> is the one that caused the trace record. If there is no program name in a column, the program name was the same as the previous one in the column, and the name is omitted. |
|  | <b>Length of Wait</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|  | (Lock) The number of milliseconds the requester waited for the locked object.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|  | <b>Lgl I/O /Sec</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|  | (Job Interval) The average number of logical disk I/O operations performed per second by the job during the interval. This is calculated from the logical disk I/O count divided by the elapsed time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|  | <b>Library</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|  | (System, Transaction) The library that contains the object.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|  | <b>LIBRARY</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|  | (Job Trace) The library name that contains the program associated with the trace entry.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|  | <b>Line Count</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|  | (Job Interval) The number of lines printed by the selected noninteractive jobs during the interval.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|  | <b>Line Descriptn</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|  | (Resource Interval) Line description name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|  | <b>Line Errors</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|  | (Resource Interval) The total of all detected errors. Check the condition of the line if this value increases greatly over time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|  | <b>Line Speed</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|  | (System, Resource Interval) The line speed in kilobits (1 kilobit = 1000 bits) per second.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|  | <b>Line Type/Line Name</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|  | (Component, System) The type and name of the line description that is used by the interface. For                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



| interfaces that do not use a line descriptions, the Line Name field will be shown as \*LOOPBACK,  
| \*OPC, or \*VIRTUALIP with no Line Type specified.

| **Line Util**

| (Resource Interval) The percent of available line capacity used by transmit and receive operations.

| **Line Util Trans/Recd**

| (Resource Interval) The percent used of the data transmission capacity of the communications  
| line. The number of bits transmitted and the number of bits received, during the interval, divided  
| by the line speed.

| **LKRL** (Transaction) Lock Released. The job released a lock it had on the object named on the next detail  
| line of the report (OBJECT --). The job that was waiting for the object is named on this line  
| (WAITER --) along with the amount of time the job spent waiting for the lock to be released.

| **LKW** (Transaction) Listed in the Wait Code column, Lock Wait. If there are a number of these, or you  
| see entries with a significant length of time in the ACTIVE/RSP\* column, additional analysis is  
| necessary. The LKWT report lines that precede this LKW report line show you what object is  
| being waited on, and who has the object.

| **LKWT**

| (Transaction) Listed in the Wait Code column, Lock Conflict Wait. The job is waiting on a lock  
| conflict. The time (\* / time /\*) is the duration of the lock conflict and, though not equal to the  
| LKW time, should be very close to it. The holder of the lock is named at the right of the report  
| line (HOLDER --). The object being locked is named on the next report line (OBJECT --).

| **Local End Code Violation**

| (Resource Interval) The number of times an unintended code violation was detected by the  
| terminal equipment (TE) for frames received at the interface for the ISDN S/T reference point.

| **Local Not Ready**

| (Resource Interval) The percent of all receive-not-ready frames that were transmitted by the host  
| system. A large percentage often means the host cannot process data fast enough (congestion).

| **Local work station IOP utilization**

| The fraction of the time interval the work station I/O processors are busy.

| **Local work station IOPs**

| (System) The resource name and model number for each local workstation IOP.

| **Lock Conflict**

| (Component) Number of lock exceptions per second. Database record contention is reflected in  
| this count. For more information, issue the Start Performance Trace (STRPFRTTC) command and  
| use the Print Transaction Report (PRTTNSRPT) and Print Lock Report (PRTLCKRPT) commands.  
| This count could be very high, even under normal system operation. Use the count as a monitor.  
| If there are large variations or changes, explore these variations in more detail.

| **Lock Wait /Tns**

| (Transaction) The average time, in seconds, of the lock-wait time per transaction. If the value is  
| high, investigate with the transaction detail calculation and the Print Lock Report (PRTLCKRPT)  
| command.

| **Logical**

| (Job Interval) The number of logical disk I/O operations performed by the selected interactive  
| jobs during the interval.

| **Logical Database I/O Other**

| (System) Other logical database operations per transaction. This includes operations such as  
| update and delete.

| **Logical Database I/O Read**

| (System) Logical database read operations per transaction.

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <b>Logical Database I/O Write</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|  | (System) Logical database write operations per transaction.                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|  | <b>Logical DB I/O</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|  | (System) Average number of logical I/O operations per transaction.                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|  | <b>Logical DB I/O Count</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|  | (System) Number of times an internal database I/O read, write, or miscellaneous function was called. This does not include I/O operations to readers, writers, or I/O operations caused by the Copy Spooled File (CPYSPLF) command or the Display Spooled File (DSPSPLF) command. If you specify SEQONLY(*YES), you see numbers that show each block of records read or written, not the number of individual records read or written. Miscellaneous functions include the following: updates, deletes, force-end-of-data, and releases. |
|  | <b>Logical Disk I/O</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|  | (Component) Number of logical disk operations (Get, Put, Update, Other).                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|  | <b>Logical I/O /Second</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|  | (System) Average number of logical disk I/O operations per second.                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|  | <b>Logical I/O Per Second</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|  | (Job Interval) The average number of logical disk I/O operations performed per second by the selected noninteractive jobs during the interval.                                                                                                                                                                                                                                                                                                                                                                                           |
|  | <b>Long Wait</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|  | (Transaction) The time the job spent waiting for a system resource. An example of a long wait would be a record-lock conflict. Also listed in the Elapsed Time--Seconds column, it is the elapsed time in the state (such as waiting for the next transaction or lock-wait time).                                                                                                                                                                                                                                                        |
|  | <b>Long Wait Lck/Oth</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|  | (Transaction) The amount of time the job spent waiting for a system resource. An example of a long wait would be a record-lock conflict.                                                                                                                                                                                                                                                                                                                                                                                                 |
|  | <b>Loss of Frame Alignment</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|  | (Resource Interval) The number of times a time period equivalent to two 48-bit frames elapsed without detecting valid pairs of line code violations.                                                                                                                                                                                                                                                                                                                                                                                     |
|  | <b>MAC Errors</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|  | (Resource Interval) The number of medium access control (MAC) errors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|  | <b>Main storage (MB)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|  | (System) The total main storage size, as measured in megabytes. These codes are in the wait code column, but they are not wait codes. They indicate transaction boundary trace records.                                                                                                                                                                                                                                                                                                                                                  |
|  | <b>Max Util</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|  | (System) Consistent use at or above the threshold value given will affect system performance and cause longer response times or less throughput.                                                                                                                                                                                                                                                                                                                                                                                         |
|  | <b>Maximum</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|  | (Transaction) The maximum value of the item that occurred in the column.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|  | <b>Member</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|  | (System, Transaction) For the System Report, this is the name of the performance data member that was specified on the TOMBR parameter of the Create Performance Data (CRTPFRTDA) command. For the Transaction Report, the member that was involved in the conflict.                                                                                                                                                                                                                                                                     |
|  | <b>Minimum</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|  | (Transaction) The minimum value of the item that occurred in the column.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|  | <b>MRT Max Time</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|  | (System) The time spent waiting, after MRTMAX is reached, by jobs routed to a multiple requester terminal.                                                                                                                                                                                                                                                                                                                                                                                                                               |

|       **Note:** No value appears in this column if job type is not MRT.

| **MSGs**  
|       (Job Trace) The number of messages sent to the job during each transaction.

| **MTU size (bytes)**  
|       (System) The size of the largest datagram that can be sent or received on the interface. The size is  
|       specified in octets (bytes). For interfaces that are used for transmitting network datagrams, this is  
|       the size of the largest network datagram that can be sent on the interface.

| **Nbr A-I**  
|       (Transaction) The number of active-to-ineligible state transitions by the job. This column shows  
|       the number of times that the job exceeded the time-slice value assigned to the job, and had to  
|       wait for an activity-level slot before the system could begin processing the transaction. If a value  
|       appears in this column, check the work that the job was doing, and determine if changes to the  
|       time-slice value are necessary.

| **Nbr Disk Units**  
|       (System) The number of disk units assigned to the reported partition.

| **Nbr Evt**  
|       (Transaction) The number of event waits that occurred during the job processing.

| **Nbr Jobs**  
|       (Transaction) The number of jobs.

| **Nbr Sign offs**  
|       (Transaction) The number of jobs that signed off during the interval.

| **Nbr Sign ons**  
|       (Transaction) The number of jobs that signed on during the interval.

| **Nbr Tns**  
|       (Transaction) The number of transactions in a given category.

|       **Note:** The values for transaction counts and other transaction-related information shown on the  
|       reports you produce using the Print Transaction Report (PRTTNSRPT) command may vary  
|       from the values shown on the reports you produce using the Print System Report  
|       (PRTSYSRPT) and Print Component Report (PRTCPTTRPT) commands. These differences  
|       are caused because the PRTTNSRPT command uses trace data as input, while the  
|       PRTSYSRPT and PRTCPTTRPT commands use sample data as input.

|       If there are significant differences in the values for transaction-related information shown on these  
|       reports, do not use the data until you investigate why these differences exist.

| **Nbr W-I**  
|       (Transaction) The number of wait-to-ineligible state transitions by the job. This column shows  
|       how many times the job had to wait for a transaction.

| **NDB Read**  
|       (Transaction) Listed in Physical I/O Counts column, it is the number of nondatabase read  
|       requests while the job was in that state. Listed in the Sync Disk I/O Rqs/Tns column, it is the  
|       average number of synchronous nondatabase read requests per transaction.

| **NDB Write**  
|       (Transaction) Listed in the Sync Disk I/O Rqs/Tns column, it is the average number of  
|       synchronous nondatabase write requests per transaction.

| **NDB Wrt**  
|       (Transaction) Listed in Physical I/O Counts column, the number of nondatabase write requests  
|       while the job was in that state. Listed under Synchronous Disk I/O Counts column, it is the  
|       number of synchronous nondatabase write requests per transaction.

- | **NON-DB**
- | (Job Trace) The number of physical nondatabase reads that occurred for the entry.
- | **Non-DB Fault**
- | (System, Component) Average number of nondatabase faults per second.
- | **Non-DB Pages**
- | (System, Component) Average number of nondatabase pages read per second.
- | **NON-DB RDS**
- | (Job Trace) The number of physical nondatabase reads that occurred.
- | **Non SMAPP**
- | (Component) Journal deposits not directly related to SMAPP (System Managed Access Path Protection).
- | **Non-SSL Inbound Connect**
- | (System) The number of non-SSL inbound connections accepted by the server.
- | **Non-Unicast Packets Received**
- | (System) The total number of non-unicast packets delivered to a higher-layer protocol for packets received on the specified interface.
- | **Non-Unicast Packets Sent**
- | (System) The total number of packets that higher-level protocols requested to be transmitted to a non-unicast address; therefore, this number includes those packets that were discarded or were not sent as well as those packets that were sent.
- | **Number**
- | (Transaction) The number of the job with which the transaction is associated.
- | **Number I/Os per Second**
- | (System) The number of I/Os per second for this particular IOP.
- | **Number Jobs**
- | (Transaction) The number of batch jobs in the job set.
- | **Number Lck Cft**
- | (Transaction) The number of lock-wait (including database record lock) state conflicts that occurred during the job processing. If this number is high, look at the Transaction and Transition Reports for the job to see how long the lock-wait state conflicts were lasting. In addition, you can do further investigation using the reports produced when you use the Print Lock Report (PRTLCKRPT) command.
- | **Number Lck Conflict**
- | (Transaction) The number of times the job had a lock conflict.
- | **Number Locks**
- | (Transaction) The number of locks attributed to interactive or noninteractive waiters.
- | **Number of batch jobs**
- | (System) The average number of active batch jobs. A batch job is considered active if it averages at least one I/O per 5 minutes.
- | **Number of Jobs**
- | (System) Number of jobs.
- | **Number of Packets Received with Errors**
- | (System) The total number of packets that were received with errors or discarded for other reasons. For example, a packet could be discarded to free up buffer space.
- | **Number Seizes**
- | (Transaction) The number of seizures attributed to interactive or noninteractive waiters.

| **Number Size Cft**  
 | (Transaction) The number of seize/lock conflicts that occurred during the job processing. If this  
 | number is high, look at the Transaction and Transition Reports for the job to see how long the  
 | conflicts lasted, the qualified name of the job that held the object, the name and type of object  
 | being held, and what the job was waiting for.

| **Number Size Conflict**  
 | (Transaction) The number of times the job had a seize conflict.

| **Number Tns**  
 | (System, Transaction) Total number of transactions processed. For example, in the System Report  
 | it is the total number of transactions processed by jobs in this pool. In the Transaction Report it is  
 | the number of transactions associated with the program.

| **Number Traces**  
 | (Batch Job Trace) Number of traces.

| **Number Transactions**  
 | (System) Total number of transactions processed.

| **Object File**  
 | (Transaction) The file that contains the object.

| **Object Library**  
 | (Transaction) The library that contains the object.

| **Object Member**  
 | (Transaction) The member that was involved in the conflict.

| **Object Name**  
 | (Lock) The name of the locked object.

| **Object RRN**  
 | (Transaction) The relative record number of the record involved in the conflict.

| **Object Type**  
 | (Transaction, Lock) The type of the locked object. The following are possible object types:

|     **AG**     Access group

|     **CB**     Commit block

|     **CBLK**   Commit block

|     **CD**     Controller description

|     **CLS**    Class

|     **CMD**    Command

|     **CTLD**   Controller description

|     **CTX**    Context

|     **CUD**    Control unit description

|     **CUR**    Cursor

|     **DEVD**  
 |         Device description

|     **DS**     Data space

|     **DSI**    Data space index

|     **DTAARA**  
 |         Data area

|  |               |                                    |
|--|---------------|------------------------------------|
|  | <b>EDTD</b>   | Edit description                   |
|  | <b>FILE</b>   | File                               |
|  | <b>JOB</b>    | Job description                    |
|  | <b>JOBQ</b>   | Job queue                          |
|  | <b>JP</b>     | Journal port                       |
|  | <b>JRN</b>    | Journal                            |
|  | <b>JRNRCV</b> |                                    |
|  |               | Journal receiver                   |
|  | <b>JS</b>     | Journal space                      |
|  | <b>LIB</b>    | Library                            |
|  | <b>LIND</b>   | Line description                   |
|  | <b>LUD</b>    | Logical unit description           |
|  | <b>MBR</b>    | Member                             |
|  | <b>MEM</b>    | Database file member               |
|  | <b>MSGF</b>   | Message file                       |
|  | <b>MSGQ</b>   |                                    |
|  |               | Message queue                      |
|  | <b>ND</b>     | Network description                |
|  | <b>OCUR</b>   |                                    |
|  |               | Database operational cursor        |
|  | <b>OUTQ</b>   |                                    |
|  |               | Output queue                       |
|  | <b>PGM</b>    | Program                            |
|  | <b>PROG</b>   | Program                            |
|  | <b>PRTIMG</b> |                                    |
|  |               | Print image                        |
|  | <b>QDAG</b>   |                                    |
|  |               | Composite piece - access group     |
|  | <b>QDDS</b>   |                                    |
|  |               | Composite piece - data space       |
|  | <b>QDDSI</b>  |                                    |
|  |               | Composite piece - data space index |
|  | <b>QTAG</b>   | Temporary - access group           |
|  | <b>QTDS</b>   | Temporary - data space             |
|  | <b>QTDSI</b>  |                                    |
|  |               | Temporary - data space index       |
|  | <b>SBSD</b>   | Subsystem description              |
|  | <b>TBL</b>    | Table                              |

#### Omit Parameters

(System, Component, Transaction, Job Interval, Pool Interval) The criteria used to choose the data records to be excluded from the report. The criteria are generally specified using an OMTxxx

| parameter of the command. Only nondefault values (something other than \*NONE) are printed.  
 | If a parameter was not specified, it does not appear on the report.

| **Op per Second**  
 | (System) Average number of disk operations per second.

| **Other Wait /Tns**  
 | (Transaction) The average time, in seconds, spent waiting that was not in any of the previous  
 | categories per transaction. For example, the time spent waiting during a save/restore operation  
 | when the system requested new media (tape or diskette).

| **Outgoing Calls Pct Retry**  
 | (Resource Interval) The percentage of outgoing calls that were rejected by the network.

| **Outgoing Calls Total**  
 | (Resource Interval) The total number of outgoing call attempts.

| **Over commitment ratio**  
 | (System) The main storage over commitment ratio (OCR).

| **PAG** (Transaction) The number of process access group faults.

| **PAG Fault**  
 | (Component, Job Interval) In the Exception Occurrence Summary of the Component Report, it is  
 | the total number of times the program access group (PAG) was referred to, but was not in main  
 | storage. The Licensed Internal Code no longer uses process access groups for caching data.  
 | Because of this implementation, the value will always be 0 for more current releases. In the  
 | Exception Occurrence Summary of the Component Report, it is the number of faults involving  
 | the process access group per second.

| **Page Count**  
 | (Job Interval) The number of pages printed by the selected noninteractive jobs during the interval.

| **Pct CPU By Categories**  
 | (Transaction) The percentage of available processing unit time used by the transactions that fell  
 | into the various categories. See the ANALYSIS by Interactive Transaction Categories part of the  
 | System Summary Data Section for an explanation of the categories.

| **Pct Data Characters Received in Error**  
 | (Resource Interval) The percent of data characters received with error.

| **Pct Data Characters Transmitted in Error**  
 | (Resource Interval) The percent of data characters transmitted with error.

| **Pct Datagrams Error**  
 | (Component) The percentage of datagrams that were discarded due to these errors:  
 | • The IP address in the destination field of the IP header was not a valid address to be received  
 | at this entity.  
 | • The protocol was unknown or unsupported.  
 | • Not enough buffer space.

| **Pct Error Responses**  
 | (Component) Percentage of responses in error.

| **Pct Ex-Wt /Rsp**  
 | (Transaction) The percentage of the response time that is due to exceptional wait.

| **Pct ICMP Messages Error**  
 | (Component) This is the number of Internet Control Message Protocol (ICMP) messages that the  
 | entity received but determined that the messages had errors or are messages that the entity did  
 | not send due to problems.



### **Pct Of Tns Categories**

(Transaction) The percentage of all transactions that fell into the various categories. See the Analysis by Interactive Transaction Categories part of the System Summary Data Section for an explanation of the categories.

### **Pct Packets Received Error**

(System) The percentage of packets that were received with errors or discarded for other reasons. For example, a packet could be discarded to free up buffer space.

### **Pct Packets Sent Error**

(System) The percentage of packets that were not sent because of errors or discarded for other reasons. For example, a packet could be discarded to free up buffer space.

### **Pct PDUs Received in Error**

(Resource Interval) The percent of protocol data units (PDUs) received in error during the time interval. These errors can occur if the host system has errors or cannot receive data fast enough (congestion).

**Note:** A protocol data unit (PDU) for asynchronous communications is a variable-length unit of data that is ended by a protocol control character or by the size of the buffer.

### **Pct Poll Retry Time**

(Resource Interval) The percent of the time interval the line was unavailable while the IOP waited for a response from a work station controller (or remote system) that was in disconnect mode.

**Note:** To minimize this lost time:

- Vary on only the controllers that are turned on.
- Turn on all controllers.
- Use the Change Line Description (SDLC) (CHGLINSDLC) command to set the connect poll timer to a small value (reduces wait time).
- Use the Change Controller Description (CHGCTLxxxx) command (where xxxx is APPC, FNC, RWS, or RTL, as appropriate) to set the NDMPOLLTMR value to a large value (increases time between polls).

### **Pct Tns**

(Transaction) The percentage of the total transactions. For the System Summary section of the Job Summary Report, the transactions are within the given trace period with the given purge attribute. For the Interactive Program Transaction Statistics section of the Job Summary Report, the percentage of transactions that were associated with a program. For the Job Statistics section, it is the percentage of total transactions that were due to this job. For the Interactive Program Statistics section, it is all transactions that were associated to a program.

### **Pct UDP Datagrams Error**

(Component) The percentage of User Datagram Protocol (UDP) datagrams for which there was no application at the destination port or that could not be delivered for other reasons.

### **Percent Errored Seconds**

(Resource Interval) The percentage of seconds in which at least one Detected Access Transmission (DTSE) in or out error occurred.

### **Percent Frames Received in Error**

(Resource Interval) The percent of all received frames that were received in error. Errors can occur when the host system has an error or cannot process received data fast enough (congestion).

### **Percent Full**

(System) Percentage of disk space capacity in use.

- | **Percent I Frames Trnsmitd in Error**
- | (Resource Interval) The percent of transmitted information frames that required retransmission. Retransmissions can occur when a remote device has an error or cannot process received data fast enough (congestion).
- | **Percent Severely Errored Seconds**
- | (Resource Interval) The percent of seconds in which at least three Detected Access Transmission (DTSE) in or out errors occurred.
- | **Percent transactions (dynamic no)**
- | (System) A measure of system main storage utilization. The percent of all interactive transactions that were done with the purge attribute of dynamic NO.
- | **Percent transactions (purge no)**
- | (System) A measure of system main storage utilization. The percent of all interactive transactions that were done with the purge attribute of NO.
- | **Percent transactions (purge yes)**
- | (System) A measure of system main storage utilization. The percent of all interactive transactions that were done with the purge attribute of YES.
- | **Percent Util**
- | (System) Average disk arm utilization (busy). Consistent use at or above the threshold value provided for disk arm utilization affects system performance, which causes longer response times or less throughput.
- | **Note:** The percent busy value is calculated from data measured in the I/O processor. When comparing this value with percent busy reported by the Work with Disk Status (WRKDSKSTS) command, some differences may exist. The WRKDSKSTS command estimates percent busy based on the number of I/O requests, amount of data transferred, and type of disk unit.
- | The system-wide average utilization does not include data for mirrored arms in measurement intervals for which such intervals are either in resuming or suspended status.
- | **Perm Size**
- | (Component) Kilobytes placed within the permanent area; these are traditional journal entries which can be retrieved and displayed.
- | **Perm Write**
- | (Component, Job Interval) The number of permanent write operations performed for the selected jobs during the interval.
- | **Permanent writes per transaction**
- | (System) The average number of permanent write operations per interactive transaction.
- | **Physical I/O Count**
- | (Transaction, Batch Job Trace) For the Job Summary section of the Batch Job Trace Report, the number of synchronous and asynchronous disk operations (reads and writes). For the Transition Report, the next five columns provide information about the number of synchronous and asynchronous disk I/O requests while the job was in the given state. The first line is the synchronous disk I/O requests, and the second line is the asynchronous disk I/O requests.
- | **DB Read**
- | The number of database read requests while the job was in that state.
- | **DB Wrt**
- | The number of database write requests while the job was in that state.
- | **NDB Read**
- | The number of nondatabase read requests while the job was in that state.

|       **NDB Wrt**  
|       The number of nondatabase write requests while the job was in that state.

|       **Tot**     The total number of DB Read, DB Wrt, NDB Read, and NDB Wrt requests.

| **Physical Writes**  
|       (Component) Physical journal write operations to disk.

| **PI**     (Component, Transaction, Job Interval, Pool Interval) The number of the pool in which the  
|       subsystem or job ran.

| **Pool**    (Transaction, Job Interval, Batch Job Trace) The number of the pool containing the transaction (for  
|       example, in which the job ran.)

| **Pool ID**  
|       (System) Pool identifier.

| **Pool ID Faults**  
|       (Component) User pool that had the highest page fault rate.

| **Pool Mch Faults/Sec**  
|       (Component) Average number of machine pool page faults per second.

| **Pool size (MB)**  
|       (Component) For the Storage Pool Activity section of the Component Report it is the initial pool  
|       size in megabytes.

| **Pool User Faults/Sec**  
|       (Component) Average number of user pool page faults per second, for the user pool with highest  
|       fault rate during this interval.

| **Pools**   (System, Component, Transaction, Job Interval, Pool Interval) In the Report-Selection Criteria  
|       section, the list of pools selected to be included (SLTPOOLS parameter) or excluded (OMTPOOLS  
|       parameter). Otherwise, the pools you specify. The values can be from 1 through 64.

| **Prg**     (Transaction) The purge attribute of the jobs.

| **Printer Lines**  
|       (System, Job Interval) The number of lines printed by the job during the interval.

| **Printer Pages**  
|       (System, Job Interval) The number of pages printed by the job during the interval.

| **Priority**  
|       (System, Transaction) The priority of the job.

| **Program**  
|       (Transaction) The name of the program with which the transaction is associated.

| **PROGRAM**  
|       (Job Trace) The name of the program for the entry.

| **PROGRAM CALL**  
|       (Job Trace) The number of non-QSYS library programs called during the step. This is not the  
|       number of times that the program named in the PROGRAM NAME field was called.

| **PROGRAM DATABASE I/O**  
|       (Job Trace) The number of times the IBM-supplied database modules were used during the  
|       transaction. The database module names have had the QDB prefix removed (PUT instead of  
|       QDBPUT). The type of logical I/O operation performed by each is as follows:

|       **GETDR**  
|       Get direct

|       **GETSQ**  
|       Get sequential

|       **GETKY**  
|           Get by key

|       **GETM**  
|           Get multiple

|       **PUT, PUTM**  
|           Add a record

|       **UDR**   Update, delete, or release a record

| **PROGRAM INIT**  
|       (Job Trace) The number of times that the IBM-supplied initialization program was called during the transaction. For RPG programs this is QRGXINIT, for COBOL it is QCRMAIN. Each time the user program ends with LR (RPG) or END (COBOL), the IBM-supplied program is also called. This is not the number of times the program named in the PROGRAM NAME field was initialized. QCRMAIN is used for functions other than program initialization (for example, blocked record I/O, some data conversions).

| **Program Name**  
|       (Transaction) For the Job Summary section of the Transaction Report, the name of the program in control at the start of the transaction. Other programs may be used during the transaction. For the Transaction Report section, the name of the program active at the start of the transaction. If ADR=UNKNWN (address unknown) is shown under the column, the program was deleted before the trace data was dumped to the database file. If ADR=000000 is shown under the column, there was not enough trace data to determine the program name, or there was no program active at that level in the job when the trace record was created.

| **PROGRAM NAME**  
|       (Job Trace) The name of the last program called that was not in the library QSYS before the end of a transaction.

| **Protocol**  
|       (System) Line protocol.

- |       • SDLC
- |       • ASYNC
- |       • BSC
- |       • X25
- |       • TRLAN
- |       • ELAN (Ethernet)
- |       • IDLC
- |       • DDI
- |       • FRLY
- |       • PPP

| **Pty**       (Component, Transaction, Job Interval) Priority of the job. For the Concurrent Batch Job Statistics section of the Transaction Report, it is the priority of the jobs in the job set.

| **PU**        (System) Partition capacity. The number of processor units assigned to the reported partition.

| **Purge**     (Transaction) The purge attribute of the jobs.

| **PWrt**     (Transaction) The number of permanent write I/O operations.

| **Queue Length**  
|       (Resource Interval) The average number of I/O requests that had to wait in the queue for this unit.

| **Rank**     (Transaction) The order. For the Job Summary section, it is the order of the program according to the number of transactions. For the Job Statistics section, it is the order of the job. For the

| Interactive Program Statistics section, it is the order of the program. For the Individual Transaction Statistics section, it is the order of the transaction according to the data being put in order by importance. For the Largest Seize/Lock Conflicts section, it is the order of the seize or lock conflict.

| **Ratio of write disk I/O to total disk I/O**  
 | (System) The fraction of the total disk activity that is due to writing data to the disks.

| **Reads per Second**  
 | (Resource Interval) The average number of disk read operations performed per second by the disk arm.

| **Receive CRC Errors**  
 | (Resource Interval) The number of received frames that contained a cycle redundancy check (CRC) error. This indicates that the data was not received error free.

| **Record Number**  
 | (Lock) For database file members, the relative record number of the record within the database file member.

| **Remote LAN Pct Frames Recd**  
 | (Resource Interval) The number of frames received from a local area network (LAN) connected to the locally attached LAN.

| **Remote LAN Pct Frames Trnsmitd**  
 | (Resource Interval) The number of frames transmitted to a local area network (LAN) connected to the locally attached LAN.

| **Remote Not Ready**  
 | (Resource Interval) The percentage of all receive-not-ready frames that were received by the host system. A large percentage often means the remote device cannot process data fast enough (congestion).

| **Remote Seq Error**  
 | (Resource Interval) The percent of frames received out of order by a remote device or system. This can occur when the remote device or system cannot process data fast enough.

| **Req type**  
 | (Component) The type of request being reported.

| **Requests received**  
 | (System, Component) The number of requests of all types received by the server.

| **Requestor's Job Name**  
 | (Lock) The name of the job requesting the locked object (the same as in the detail listing).

| **Reset Packets Recd**  
 | (Resource Interval) The number of reset packets received by the network. **Reset packets** are packets retransmitted because an error occurred.

| **Reset Packets Trnsmitd**  
 | (Resource Interval) The number of reset packets transmitted by the network.

| **Response**  
 | (System) Average system response (service) time.

| **Response Sec Avg and Max**  
 | (Transaction) The average (AVG) and maximum (MAX) transaction response time, in seconds, for the job. The average response time is calculated as the sum of the time between each pair of wait-to-active and active-to-wait transitions divided by the number of pairs that were encountered for the job. The MAX response time is the largest response time in the job.

| **Response Seconds**  
 | (System) Average response time in seconds per transaction.

- | **Responses sent**
- | (System, Component) The number of responses of all types sent by the server.
- | **Rsp** (Component) Average interactive transaction response time in seconds.
- | **Rsp Time**
- | (Component, Resource Interval) The average external response time (in seconds). For the Local
- | Work Station IOP Utilizations section of the Resource Interval Report, it is the response time for
- | work stations on this controller. For the Remote Work Stations section of the Component Report,
- | it is the response time for this work station.
- | **Rsp Timer Ended**
- | (Resource Interval) The number of times the response timer ended waiting for a response from a
- | remote device.
- | **Rsp/Tns**
- | (Component, Transaction, Job Interval) The average response time (seconds) per transaction. For
- | the Job Summary section of the Job Interval Report, it is the response time per transaction for the
- | selected interactive jobs during the interval (the amount of time spent waiting for or using the
- | system resources divided by the number of transactions processed). This number will not be
- | accurate unless at least several seconds were spent processing transactions.
- | **S/L** (Transaction) Whether the conflict was a seize (S) or lock (L) conflict.
- | **SECONDS**
- | (Job Trace) The approximate time the job was waiting or active.
- | **Segments Pct Rtrns**
- | (Component) The percentage of segments retransmitted. This number is the TCP segments that
- | were transmitted and that contain one or more previously transmitted octets (bytes).
- | **Segments Rcvd per Second**
- | (Component) The number of segments received per second. This number includes those received
- | in error and those received on currently established connections.
- | **Segments Sent per Second**
- | (Component) The number of segments sent per second. This number includes those sent on
- | currently established connections and excludes those that contain only retransmitted octets
- | (bytes).
- | **Seize and Lock Conflicts**
- | (Batch Job Trace) Number of seize conflicts and lock waits.
- | **Seize Conflict**
- | (Component) Number of seize exceptions per second. For more detailed information, issue the
- | Start Performance Trace (STRPFRTTC) command, and use the PRTTNSRPT or PRTLCKRPT
- | commands. This count could be very high, even under normal system operation. Use the count as
- | a monitor. If there are large variations or changes, explore these variations in more detail.
- | **Seize Hold Time**
- | (Transaction) The amount of time that the transaction held up other jobs in the system by a seize
- | or lock on an object.
- | **Seize Wait /Tns**
- | (Transaction) The average time, in seconds, for all seize-lock conflicts that occur during an
- | average transaction. More than one seize-lock conflict can occur during a single transaction for
- | the same job. If this number is high, investigate those jobs with seize conflicts. The Transaction
- | Report lists each conflict that occurs, the name of the holder, and the name of the object held. For
- | the Transaction by 5-Minute Intervals section of the Job Summary Report, it is the average seize
- | wait time per transaction in seconds. This is the average amount of time that the transactions
- | spent in a seize/lock conflict. If this number is high, look at the Transaction and Transition
- | Reports for the jobs that are causing the excessive wait time.

| **Select Parameters**  
 | (System, Component, Transaction, Job Interval, Pool Interval) The criteria used to choose the data records to be included in the report. The criteria are generally specified using an SLTxxx parameter of the command. Only nondefault values (something other than \*ALL) are printed. If a parameter is not specified, it does not appear on the report.

| **SEQNBR**  
 | (Job Trace) The number of the trace entry.

| **SEQNCE or SEQUENCE**  
 | (Job Trace) The job trace sequence number in the detail report that this summary line refers to.

| **Sequence Error**  
 | (Resource Interval) The number of frames received that contained sequence numbers indicating that frames were lost.

| **Server job name**  
 | (System) The server job number. Identifies the child job for the server.

| **Server job user**  
 | (System) The server job user. Identifies the child job for the server.

| **Server name**  
 | (System) The server job name. Identifies the child job for the server.

| **Server start date/time**  
 | (System) The most recent start or restart time in format mm/dd/yy hh:mm:ss

| **Short Frame Errors**  
 | (Resource Interval) The number of short frames received. A short frame is a frame that has fewer octets between its start flag and end flag than are permitted.

| **Short Wait /Tns**  
 | (Transaction) The average time, in seconds, of short (active) wait time per transaction. For the Interactive Program Statistics section, if the value is high, it may be due to the use of data queues or to the use of DFRWRT(\*NO) or RSTDSP(\*YES) in the program display files.

| **Short WaitX /Tns (Short wait extended)**  
 | (Transaction) The average time, in seconds, of wait time per transaction that resulted due to a short (active) wait that exceeded 2 seconds, and caused a long wait transition to occur. The activity level has been released but this time is still counted against your total response time. Waits on data queues or the use of DFRWRT(\*NO) and/or RSTDSP(\*YES) in the display files could be reasons for this value to be high.

| **Size** (Component) Decimal data overflow and underflow exceptions per second. An indication of improper field size on numeric calculations.

| **Size (MB)**  
 | (System) The size of the pool in megabytes.

| **Size (GB)**  
 | (Pool Interval) The size of the pool in gigabytes.

| **Size (M)**  
 | (System) Disk space capacity in millions of bytes.

| **SHARE CLS**  
 | (Job Trace) The number of shared closes for all types of files.

| **SHARE OPN**  
 | (Job Trace) The number of shared opens for all types of files.

| **SMAPP ReTune**  
 | (Component) System-managed access path protection tuning adjustments.



- | **SMAPP System**  
| (Component) SMAPP-induced journal entries deposited in system-provided (default) journals.
- | **SMAPP User**  
| (Component) SMAPP-induced journal entries deposited in user-provided journals.
- | **SOTn** (Transaction) Listed in the Wait Code column, Start of transaction n. These codes are in the wait  
| code column, but they are not wait codes. They indicate transaction boundary trace records.
- | **Spool CPU seconds per I/O**  
| (System) The average number of system processing unit seconds used by all spool jobs for each  
| I/O performed by a spool job.
- | **Spool database reads per second**  
| (System) The average number of read operations to database files per second of spool processing.
- | **Spool I/O per second**  
| (System) The average number of physical disk I/O operations per second of spool processing.
- | **Srv Time**  
| (Component) Average disk service time per request in seconds not including the disk wait time.
- | **SSL Inbound Connections**  
| (System) AThe number of SSL inbound connections accepted by the server.
- | **Start** (Transaction) The time the job started.
- | **Started**  
| (Transaction) The time of the first record in the trace data, in the form HH.MM.SS (hours,  
| minutes, seconds).
- | **State** (Transaction) The three possible job states are:  
| • **W**--(Wait state) not holding an activity level.  
| • **A**--(Active or wait state) holding an activity level.  
| • **I**--(Ineligible state) waiting for an activity level.  
|  
| The table below shows the possible job state transitions. For example, from **W** to **A** is **yes**, which  
| means it is possible for a job to change from the *wait* state to the *active* state.

| Table 7.

|            |   | To state |     |     |
|------------|---|----------|-----|-----|
|            |   | A        | W   | I   |
| From state | A | yes      | yes | yes |
|            | W | yes      |     | yes |
|            | I | yes      |     |     |

- | **State Transitions A-A**  
| (Batch Job Trace) Number of active-to-active transitions.
- | **State Transitions A-I**  
| (Batch Job Trace) Number of active-to-ineligible transitions.
- | **Stop** (Transaction) The time the job ended.
- | **Stopped**  
| (Transaction) The time of the last record in the trace data, in the form HH.MM.SS (hours,  
| minutes, seconds).
- | **SUBFILE READS**  
| (Job Trace) The number of subfile reads.

| **SUBFILE WRITES**  
 | (Job Trace) The number of subfile writes.

| **Subsystem Name**  
 | (Pool Interval) The name of the subsystem.

| **Subsystems**  
 | (System, Component, Pool Interval) For the System Report, the subsystem names you specify.  
 | Each name is a 10-character name. For the Component Report, the list of subsystems selected to  
 | be included (SLTSBS parameter) or excluded (OMTSBS parameter).

| **Sum** (Transaction) Listed in the Sync Disk I/O Rqs/Tns column, the sum of the averages of the  
 | synchronous DB READ, DB WRITE, NDB READ, and NDB WRITE requests (the average number  
 | of synchronous I/O requests per transaction for the job).

| **SWX** (Transaction) Listed in the Wait Code column, Short Wait Extended. The short wait has exceeded  
 | a 2-second limit and the system has put the transaction into a long wait. This long wait must be  
 | charged to the transaction response time. In most cases, this active-to-wait transaction does not  
 | reflect a transaction boundary.

| **Sync** (Job Interval) The number of synchronous disk I/O operations performed by the selected  
 | interactive jobs during the interval.

| **Sync DIO /Tns**  
 | (Transaction) The average number of synchronous I/O requests per transaction during the  
 | interval.

| **Sync Disk I/O**  
 | (System, Component, Transaction) Synchronous disk I/O operations.

| **Sync Disk I/O per Second**  
 | (Component) Average synchronous disk I/O operations per second.

| **Sync Disk I/O Requests**  
 | (Transaction) The total number of synchronous disk I/O requests for the given combination of  
 | priority, job type, and pool.

| **Sync Disk I/O Rqs/Tns**  
 | (Transaction) The next five columns provide information about the number of synchronous disk  
 | I/O requests per transaction:

| **DB Read**  
 | The average number of synchronous database read requests per transaction.

| **DB Write**  
 | The average number of synchronous database write requests per transaction.

| **NDB Read**  
 | The average number of synchronous nondatabase read requests per transaction.

| **NDB Write**  
 | The average number of synchronous nondatabase write requests per transaction.

| **Sum** The sum of the averages of the synchronous DB READ, DB WRITE, NDB READ, and  
 | NDB WRITE requests (the average number of synchronous I/O requests per transaction  
 | for the job).

| **Sync I/O /Elp Sec**  
 | (Transaction) The average number of synchronous disk I/O requests for all jobs, per second of  
 | elapsed time used by the jobs.

| **Sync I/O /Sec**  
 | (Job Interval) The average number of synchronous disk I/O operations performed per second by  
 | the job during the interval. This is calculated from the synchronous disk I/O count divided by  
 | the elapsed time.

| **Sync I/O Per Second**  
 | (Job Interval) The average number of synchronous disk I/O operations performed per second by  
 | the selected noninteractive jobs during the interval.

| **Synchronous DBR**  
 | (System, Transaction, Job Interval, Pool Interval) The average number of synchronous database  
 | read operations. It is the total synchronous database reads divided by the total transactions. For  
 | the Pool Interval and Job Interval Reports, it is calculated per transaction for the job during the  
 | intervals. For the System Report, it is calculated per second. For the Transaction (Job Summary) it  
 | is calculated per transaction. Listed under Average DIO/Transaction, the average number of  
 | synchronous database read requests per transaction. This field is not printed if the jobs in the  
 | system did not process any transactions.

| **Synchronous DBW**  
 | (System, Transaction, Job Interval, Pool Interval) The average number of synchronous database  
 | write operations. It is the total synchronous database writes divided by the total transactions. For  
 | the Pool Interval and Job Interval Reports, it is calculated per transaction for the job during the  
 | intervals. For the System Report, it is calculated per second. For the Transaction (Job Summary) it  
 | is calculated per transaction. Listed under Average DIO/Transaction, the average number of  
 | synchronous database read requests per transaction. This field is not printed if the jobs in the  
 | system did not process any transactions.

| **Synchronous DIO / Act Sec**  
 | (System, Transaction) The number of synchronous disk I/O operations per active second. The  
 | active time is the elapsed time minus the wait times.

| **Synchronous DIO / Ded Sec**  
 | (Transaction) The estimated number of synchronous disk I/O operations per second as if the job  
 | were running in dedicated mode. Dedicated mode means that no other job would be active or in  
 | contention for resources in the system.

| **Synchronous DIO / Elp Sec**  
 | (Transaction) The number of synchronous disk I/O operations per elapsed second.

| **Synchronous Disk I/O Counts**  
 | (Transaction) The next five columns provide information about the number of synchronous disk  
 | I/O requests per transaction:

| **DB Read**  
 | The number of synchronous database read requests per transaction.

| **DB Wrt**  
 | The number of synchronous database write requests per transaction.

| **NDB Read**  
 | The number of synchronous nondatabase read requests per transaction.

| **NDB Wrt**  
 | The number of synchronous nondatabase write requests per transaction.

| **Sum** The sum of the synchronous DB Read, DB Wrt, NDB Read, and NDB Wrt requests (the  
 | number of synchronous I/O requests per transaction).

| **Synchronous disk I/O per transaction**  
 | (System, Transaction) The average number of synchronous physical disk I/O operations per  
 | interactive transaction.

| **Synchronous Max**  
 | (Transaction) The maximum number of synchronous DBR, NDBR, and WRT I/O requests  
 | encountered for any single transaction by that job. If the job is not an interactive or autostart job  
 | type, the total disk I/O for the job is listed here.

| **Synchronous NDBR**  
 | (System, Transaction, Job Interval, Pool Interval) The average number of synchronous  
 | nondatabase read operations per transaction for the jobs in the system during the interval. For the  
 | Transaction Report, the operations on the disk per transaction for the selected jobs in the pool.  
 | This is calculated from the synchronous nondatabase read count divided by the transactions  
 | processed. This field is not printed if the jobs in the system did not process any transactions.

| **Synchronous NDBW**  
 | (System, Job Interval, Pool Interval) The average number of synchronous nondatabase write  
 | operations on the disk per transaction for the selected jobs in the pool. For the System Report, it  
 | is the operations per transaction for the jobs in the system during the interval. This is calculated  
 | from the synchronous nondatabase write count divided by the transactions processed. This field  
 | is not printed if the jobs in the system did not process any transactions.

| **Synchronous Sum**  
 | (Transaction) The sum of the averages of the synchronous DBR, NDBR, and WRT requests (the  
 | average number of synchronous I/O requests per transaction for the job).

| **Synchronous wrt**  
 | (Transaction) The average number of synchronous database and nondatabase write requests per  
 | transaction.

| **System CPU per transaction (seconds)**  
 | (System) The average number of system processing unit seconds per interactive transaction.

| **System disk I/O per transaction**  
 | (System) The total number of physical disk I/O operations attributed to the system per  
 | interactive transaction.

| **System Starts**  
 | (Component) The number of start journal operations initiated by the system.

| **System Stops**  
 | (Component) The number of stop journal operations initiated by the system.

| **System Total**  
 | (Component) The total number of journal deposits resulting from system-journaled objects. These  
 | are the deposits performed by system-managed access path protection (SMAPP).

| **System ToUser**  
 | (Component) The number of journal deposits resulting from system-journaled objects to  
 | user-created journals.

| **SZWG**  
 | (Transaction) Listed in the Wait Code column, Seize Wait Granted. The job was waiting on a seize  
 | conflict. The original holder released the lock that it had on the object, and the lock was then  
 | granted to the waiting job. The job that was waiting for the object is named on this line (WAITER  
 | --) along with the amount of time the job spent waiting for the seize conflict to be released. The  
 | object that is held is named on the next line of the report (OBJECT --).

| **SZWT** (Transaction) Listed in the Wait Code column, Seize/Lock Conflict Wait. The job is waiting on a  
 | seize/lock conflict. The time (\* / time / \*) is the duration of the seize/lock conflict, and is included  
 | in the active time that follows it on the report. The holder of the lock is named at the right of the  
 | report line (HOLDER --). The object being held is named on the next report line (OBJECT --).

| **Teraspace EAO**  
 | (Component) Listed in the Exception Occurrence summary and Interval Counts. A teraspace

effective address overflow (EAO) occurs when computing a teraspace address that crosses a 16-boundary. A quick estimate indicates that a 1% performance degradation would occur if there were 2,300 EAOs per second.

### Thread

(Job Summary, Transaction, Transition) A thread is a unique flow of control within a process. Every job has an initial thread associated with it. Each job can start one or more secondary threads. The system assigns the thread number to a job as follows:

- The system assigns thread IDs sequentially. When a job is started that uses a job structure that was previously active, the thread ID that is assigned to the initial thread is the next number in the sequence.
- The first thread of a job is assigned a number.
- Any additional threads from the same job are assigned a number that is incremented by 1. For example:

| Job Name  | User Name/<br>Thread | Job Number |
|-----------|----------------------|------------|
| QJVACMSRV | SMITH                | 023416     |
| QJVACMSRV | 00000006             | 023416     |
| QJVACMSRV | 00000007             | 023416     |
| QJVACMSRV | 00000008             | 023416     |

A thread value greater than 1 does not necessarily mean the job has had that many threads active at the same time. To determine how many threads are currently active for the same job, use the WRKACTJOB, WRKSBSJOB, or WRKUSRJOB commands to find the multiple three-part identifiers with the same job name.

### Threads active

(System) The number of threads doing work when the data was sampled.

### Threads idle

(System) The number of idle threads when the data was sampled.

### Time

(Transaction) The time when the transaction completed, or when a seize or lock conflict occurred. Also, a column heading that shows the time the transition from one state to another occurred, in the HH.MM.SS.mmm arrangement.

### TIME

(Job Trace) The time of day for the trace entry. The time is sequentially given in hours, minutes, seconds, and microseconds.

### Tns

(Component, Pool Interval) The total number of transactions processed by the selected jobs in the pool or subsystem.

### Tns Count

(Component, Job Interval) The number of transactions performed by the selected interactive jobs during the interval.

### Tns/Hour

(Component, Transaction, Job Interval) The average number of transactions per hour processed by the selected interactive jobs during the interval.

### Tns/Hour Rate

(System) Average number of transactions per hour.

### TOD of Wait

(Lock) The time of day of the start of the conflict.

### Tot

(Transaction) Listed in Physical I/O Counts column, the total number of DB Read, DB Wrt, NDB Read, and NDB Wrt requests.

### Tot Nbr Tns

(Transaction) The total number of transactions the PRTTNSRPT program determined from the input data that were accomplished for the job.

| **Total** (Component) Total exception counts for the reporting period.

| **TOTAL**

| (Job Trace) Totals for the fields.

| **Total /Job**

| (Transaction) The total (sum) of the items in the column for the job.

| **Total characters per transaction**

| (System) The average number of characters either read from or written to display station screens per interactive transaction.

| **Total CPU Sec /Sync DIO**

| (Transaction) The ratio of total CPU seconds divided by the total synchronous disk I/O requests.

| **Total CPU Utilization**

| (System, Component) Percentage of available processing unit time used by the partition. For a multiple-processor system, this is the average use across all processors. For dedicated partitions, *Total CPU Utilization* is replaced by a utilization value for each processor in the partition. Here is an example of this part of the display for a dedicated partition with two processors:

|                                     |      |
|-------------------------------------|------|
| Average CPU utilization . . . . . : | 41.9 |
| CPU 1 utilization . . . . . :       | 41.7 |
| CPU 2 utilization . . . . . :       | 42.2 |

| In shared processor partitions, individual CPU utilization rows are not printed.

| **Note:** This value is taken from a system counter. Other processing unit uses are taken from the individual job work control blocks (WCBs). These totals may differ slightly. For uncapped partitions, *Total CPU utilization* might exceed 100 percent.

| **Total CPU Utilization (Database Capability)**

| (System) Shows you the DB2 Universal Database™ activity on your systems. This field applies to all systems running V4R5 or later and includes all database activity, including all SQL and data I/O operations.

| **Total CPU Utilization (Interactive Feature)**

| (System) The CPU Utilization (Interactive Feature) shows 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.

| **Total Data Characters Received**

| (Resource Interval) The number of data characters received successfully.

| **Total Data Characters Transmitted**

| (Resource Interval) The number of data characters transmitted successfully.

| **Total Datagrams Requested for Transmission**

| (Component) The percentage of IP datagrams that are discarded because of the following reasons:

- | • No route was found to transmit the datagrams to their destination.
- | • Lack of buffer space.

| **Total fields per transaction**

| (System) The average number of display station fields either read from or written to per interactive transaction.

| **Total Frames Recd**

| (Resource Interval) The number of frames received, including frames with errors and frames that are not valid.

| **Total I Frames Trnsmitd**

| (Resource Interval) The total number of information frames transmitted.

|  |                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <b>Total I/O</b>                           | (System) Sum of the read and write operations.                                                                                                                                                                                                                                                                                                                                                                                                          |
|  | <b>Total PDUs Received</b>                 | (Resource Interval) The number of protocol data units (PDUs) received during the time interval.                                                                                                                                                                                                                                                                                                                                                         |
|  | <b>Note:</b>                               | A protocol data unit (PDU) for asynchronous communications is a variable-length unit of data that is ended by a protocol control character or by the size of the buffer.                                                                                                                                                                                                                                                                                |
|  | <b>Total Physical I/O per Second</b>       | (Resource Interval) The average number of physical disk I/O operations performed per second by the disk arm.                                                                                                                                                                                                                                                                                                                                            |
|  | <b>Total Responses</b>                     | (Component, Resource Interval) The total number of transactions counted along with the average response time for all active work stations or devices on this controller for the report period.                                                                                                                                                                                                                                                          |
|  | <b>Total Seize/Wait Time</b>               | (Component) The response time in milliseconds for each job.                                                                                                                                                                                                                                                                                                                                                                                             |
|  | <b>Total Tns</b>                           | (Component) Number of transactions processed in this pool.                                                                                                                                                                                                                                                                                                                                                                                              |
|  | <b>Transaction Response Time (Sec/Tns)</b> | (Transaction) The response time in seconds for each transaction. This value includes no communications line time. Response times measured at the work station exceed this time by the data transmission time (the time required to transmit data from the work station to the processing unit and to transmit the response data back to the work station from the processing unit).                                                                     |
|  | <b>Transactions per hour (local)</b>       | (System) The interactive transactions per hour attributed to local display stations.                                                                                                                                                                                                                                                                                                                                                                    |
|  | <b>Transactions per hour (remote)</b>      | (System) The interactive transactions per hour attributed to remote display stations.                                                                                                                                                                                                                                                                                                                                                                   |
|  | <b>Transient Size</b>                      | (Component) Kilobytes placed within the journal transient area; these are hidden journal entries produced by the system.                                                                                                                                                                                                                                                                                                                                |
|  | <b>Transmit/Receive/Average Line Util</b>  | (Resource Interval) In duplex mode, the percentage of transmit line capacity used, the percentage of receive line capacity used, and the average of the transmit and receive capacities.                                                                                                                                                                                                                                                                |
|  | <b>TSE</b>                                 | (Transaction) Listed in the Wait Code column, Time Slice End. The program shown in the stack entry labeled LAST is the program that went to time slice end.                                                                                                                                                                                                                                                                                             |
|  | <b>Type</b>                                | (Component, Transaction) The system job type and subtype. The Component Report allows only one character in this column. The Transaction Report allows two characters. The Transaction Report reports the job type and job subtype directly from the QAPMJOB fields. The Component Report takes the job type and job subtype values and converts it to a character that may or may not be the value from the QAPMJOB field. The possible job types are: |
|  | <b>A</b>                                   | Autostart                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|  | <b>B</b>                                   | Batch                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|  | <b>BD</b>                                  | Batch immediate (Transaction only)                                                                                                                                                                                                                                                                                                                                                                                                                      |
|  | <b>Note:</b>                               | The batch immediate values are shown as BCI on the Work with Active Job display and as BATCHI on the Work with Subsystem Job display.                                                                                                                                                                                                                                                                                                                   |
|  | <b>BE</b>                                  | Batch evoke (Transaction only)                                                                                                                                                                                                                                                                                                                                                                                                                          |



|           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>BJ</b> | Batch pre-start job (Transaction only)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>C</b>  | <p>Programmable work station application server, which includes 5250 emulation over APPC and System i Access host servers running either APPC or TCP/IP. See the Host server administration topic for more information. A job is reported as a System i Access server if any of the following items are true:</p> <ul style="list-style-type: none"> <li>• Incoming APPC evoke requests one of the server program names. This also applies to the pre-started jobs for the QSERVER, QCMN, and QSYSWRK subsystems that are already waiting for the named program.</li> <li>• Incoming IP port number corresponds to one of the service name-description-port-numbers. This also applies to the pre-started jobs for the QSERVER, QCMN, and QSYSWRK subsystems that are already waiting for the assigned IP port number.</li> <li>• Incoming IPX socket number corresponds to one of the service name-description-port-numbers. This also applies to the pre-started jobs for the QSERVER, QCMN, and QSYSWRK subsystems that are already waiting for the assigned IPX port number.</li> <li>• Incoming 5250 display emulation jobs that come from APPC data streams sent by 5250 emulation under OS/2 Communications Manager or WARP equivalent.</li> </ul> |
| <b>D</b>  | Target distributed data management (DDM) server                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>I</b>  | Interactive. For the Component Report, this includes twinaxial data link control (TDLC), 5250 remote workstation, and 3270 remote workstation. For the Transaction Report, this includes twinaxial data link control (TDLC), 5250 remote workstation, 3270 remote workstation, SNA pass-through, and 5250 Telnet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>L</b>  | Licensed Internal Code Task                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>M</b>  | Subsystem monitor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>P</b>  | SNA pass-through and 5250 Telnet pass-through. On the Transaction Report, these jobs appear as I (interactive).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>R</b>  | Spool reader                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>S</b>  | System                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>W</b>  | Spool writer, which includes the spool write job, and if Advanced Function Printing (AFP) is specified, the print driver job.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>WP</b> | Spool print driver (Transaction only)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>X</b>  | Start the system                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

The possible job subtypes are:

|          |                                                                |
|----------|----------------------------------------------------------------|
| <b>D</b> | Batch immediate job                                            |
| <b>E</b> | Evoke (communications batch)                                   |
| <b>J</b> | Pre-start job                                                  |
| <b>P</b> | Print driver job                                               |
| <b>T</b> | Multiple requester terminal (MRT) (System/36 environment only) |
| <b>3</b> | System/36                                                      |

**Notes:**

1. Job subtypes do not appear on the Component Report.
2. If the job type is blank or you want to reassign it, use the Change Job Type (CHGJOBTYP) command to assign an appropriate job type.

**Type** (System, Transaction, Job Interval) One of the transaction types listed in the description of the DTNTY field.

| **(System)**  
| The disk type.

| **(Transaction)**  
| The type and subtype of the job.

| **(Transaction)**  
| For the Seize/Lock Conflicts by Object section, the type of seize/lock conflict.

| **UDP Datagrams Received**  
| (Component) The total number of User Datagram Protocol (UDP) datagrams delivered to UDP  
| users.

| **UDP Datagrams Sent**  
| (Component) The total number of User Datagram Protocol (UDP) datagrams sent from this entity.

| **Uncap CPU Avail**  
| (Component) Percentage of CPU time available to a partition in the shared processors pool during  
| the interval in addition to its configured CPU. This value is relative to the configured CPU  
| available for the particular partition.

| **Unicast Packets Received**  
| (System) The total number of subnetwork-unicast packets delivered to a higher-layer protocol.  
| The number includes only packets received on the specified interface.

| **Unicast Packets Sent**  
| (System) The total number of packets that higher-level protocols requested to be transmitted to a  
| subnetwork-unicast address. This number includes those packets that were discarded or were not  
| sent.

| **Unit** (System, Component, Resource Interval) The number assigned by the system to identify a specific  
| disk unit or arm. An 'A' or 'B' following the unit number indicates that the disk unit is mirrored.  
| (For example, 0001A and 0001B are a mirrored pair.)

| **Unit Name**  
| The resource name of the disk arm.

| **User ID**  
| (System, Component, Transaction, Job Interval, Pool) The list of users selected to be included  
| (SLTUSRID parameter) or excluded (OMTUSRID parameter).

| **User Name**  
| (Component, Transaction, Job Interval, Batch Job Trace) Name of the user involved (submitted the  
| job, had a conflict, and so on.)

| **User Name/Thread**  
| (Component, Transaction) If the job information contains a secondary thread, then this column  
| shows the thread identifier. If the job information does not contain a secondary thread, then the  
| column shows the user name. The system assigns the thread number to a job as follows:

- | • The system assigns thread IDs sequentially. When a job is started that uses a job structure that  
| was previously active, the thread ID that is assigned to the initial thread is the next number in  
| the sequence.
- | • The first thread of a job is assigned a number.
- | • Any additional threads from the same job are assigned a number that is incremented by 1. For  
| example:

| Job Name  | User Name/<br>Thread | Job Number |
|-----------|----------------------|------------|
| QJVACMSRV | SMITH                | 023416     |
| QJVACMSRV | 00000006             | 023416     |
| QJVACMSRV | 00000007             | 023416     |
| QJVACMSRV | 00000008             | 023416     |

A thread value greater than 1 does not necessarily mean the job has had that many threads active at the same time. To determine how many threads are currently active for the same job, use the WRKACTJOB, WRKSBSJOB, or WRKUSRJOB commands to find the multiple three-part identifiers with the same job name.

**User Starts**  
(Component) The number of start journal operations initiated by the user.

**User Stops**  
(Component) The number of stop journal operations initiated by the user.

**User Total**  
(Component) The total number of journal deposits resulting from system-journaled objects.

**Util** (Component, Resource Interval) The percent of utilization for each local work station, disk, or communications IOP, controller, or drive.

**Note:** The system-wide average utilization does not include data for mirrored arms in measurement intervals for which such intervals are either in resuming or suspended status.

**Util 2** (Component, Resource) Utilization of co-processor.

**Value** (Transaction) For the Individual Transaction Statistics section of the Job Summary report, it is the value of the data being compared for the transaction. For the Longest Seize/Lock Conflicts section, it is the number of seconds in which the seize or lock conflict occurred.

**Verify** (Component) Number of verify exceptions per second. Verify exceptions occur when a pointer needs to be resolved, when blocked MI instructions are used at security levels 10, 20, or 30, and when an unresolved symbolic name is called. This count could be very high, even under normal system operation. Use the count as a monitor. If there are large variations or changes, explore these variations in more detail.

**VP** (System) The number of virtual processors active in the reported partition.

**Vrt Shr Proc Pool ID**  
(System) Virtual Shared Processor Pool ID. This column is only printed for the i5/OS partition.

**W-I Wait/Tns**  
(Transaction) The average time, in seconds, of wait-to-ineligible time per transaction. This value is an indication of what effect the activity level has on response time. If this value is low, the number of wait-to-ineligible transitions probably has little effect on response time. If the value is high, adding additional interactive pool storage and increasing the interactive pool activity level should improve response time. If you are unable to increase the interactive pool storage (due to limited available storage), increasing the activity level may also improve response time. However, increasing the activity level might result in excessive faulting within the storage pool.

**Wait Code**  
(Transaction) The job state transition that causes the trace record to be produced. The values can be as follows:

**EVT** Event Wait. A long wait that occurs when waiting on a message queue.

**EOTn** End of transaction for transaction for type n. These codes are in the wait code column, but they are not wait codes. They indicate transaction boundary trace records.

**EORn** End of response time for transaction n. These codes are in the wait code column, but they are not wait codes. They indicate transaction boundary trace records.

**Error Responses**  
(Component) The number of responses in error.

**HDW** Hold Wait (job suspended or system request).

**LKRL** Lock Released. The job released a lock it had on the object named on the next detail line

of the report (OBJECT --). The job that was waiting for the object is named on this line (WAITER --) along with the amount of time the job spent waiting for the lock to be released.

**LKW** Lock Wait. If there are a number of these, or you see entries with a significant length of time in the ACTIVE/RSP\* column, additional analysis is necessary. The LKWT report lines that precede this LKW report line show you what object is being waited on, and who has the object.

**LKWT**

Lock Conflict Wait. The job is waiting on a lock conflict. The time (\* / time \*) is the duration of the lock conflict and, though not equal to the LKW time, should be very close to it. The holder of the lock is named at the right of the report line (HOLDER --). The object being locked is named on the next report line (OBJECT --).

**SOTn** Start of transaction n. These codes are in the wait code column, but they are not wait codes. They indicate transaction boundary trace records.

**SWX** Short Wait Extended. The short wait has exceeded a 2-second limit and the system has put the transaction into a long wait. This long wait must be charged to the transaction response time. In other words, this active-to-wait transaction does not reflect a transaction boundary.

**SZWG**

(Transaction) Listed in the Wait Code column, Seize Wait Granted. The job was waiting on a seize conflict. The original holder released the lock that it had on the object, and the lock was then granted to the waiting job. The job that was waiting for the object is named on this line (WAITER --) along with the amount of time the job spent waiting for the seize conflict to be released. The object that is held is named on the next line of the report (OBJECT --).

**SZWT** Seize/Lock Conflict Wait. The job is waiting on a seize/lock conflict. The time (\* / time \*) is the duration of the seize/lock conflict, and is included in the active time that follows it on the report. The holder of the lock is named at the right of the report line (HOLDER --). The object being held is named on the next report line (OBJECT --).

**TSE** Time Slice End. The program shown in the stack entry labeled LAST is the program that went to time slice end. Every time a job uses 0.5 seconds of CPU time (0.2 seconds on the faster processors) between long waits, the system checks if there are jobs of equal priority on the CPU queue. If there are, then the next job with equal priority is granted the CPU and the interrupted job is moved to the queue as the last of equal priority. The job, however, retains its activity level. This is an internal time slice end. When a job reaches the external time slice value, there can be a job state transition from active to ineligible if another job is waiting for an activity level. When a job is forced out of its activity level, its pages are liable to be stolen by other jobs, and cause additional I/O when the job regains an activity level. The IBM-supplied default values of 2 seconds for interactive jobs and 5 seconds for batch jobs may often be too high, especially for the high-end processors. As an initial value, set the time slice at 3 times the average CPU seconds per transaction.

**WTO** Wait Timed Out. The job has exceeded the wait time-out limit defined for a wait (such as a wait on a lock, a message queue, or a record).

**WAITS**

(Job Trace) The number of waits that occurred.

**WAIT-ACT**

(Job Trace) In the Job Trace Analysis Summary, this is the time between the ENDTNS and STRTNS programs is labeled WAIT-ACT. If you were tracing an interactive job and used the default STRTNS and ENDTNS parameters, this value is the time taken to process the transaction.

| In the Job Trace Analysis I/O Summary, this is the time that the job was inactive, probably due to typing or think time by the user.

| **Wait-Inel**  
| (System, Component) Average number of wait-to-ineligible job state transitions per minute.

| **Work Station Controller**  
| (Resource Interval) The name of the remote work station controller.

| **WRITES**  
| (Job Trace) The number of physical writes that occurred.

| **Writes per Second**  
| (Resource Interval) The average number of disk write operations performed per second by the disk arm.

| **WRITTEN**  
| (Job Trace) The number of physical writes that occurred for the entry.

| **WTO** (Transaction) Listed in the Wait Code column, Wait Timed Out. The job has exceeded the wait time-out limit defined for a wait (such as a wait on a lock, a message queue, or a record).

| **0.0-1.0** (Component, Resource Interval) The number of times the response time was between 0 and 1 second.

| **1.0-2.0** (Component, Resource Interval) The number of times the response time was between 1 and 2 seconds.

| **2.0-4.0** (Component, Resource Interval) The number of times the response time was between 2 and 4 seconds.

| **4.0-8.0** (Component, Resource Interval) The number of times the response time was between 4 and 8 seconds.

| **Performance Report header:**

| Each report, regardless of the type or section, contains information in the header of the report that identifies characteristics of the data. Look here for descriptions of the header information.

| **Report title**  
| Identifies the type of performance report on the first line. The second line identifies the section of the report.

| **Current date and time**  
| Indicates the date and time the report was printed.

| **Report page number**  
| Identifies the page of the report.

| **Perf data from *time to time at interval***  
| Indicates the time period over which the data was collected and at what interval.

| **User-selected report title**  
| Indicates the name assigned to the report by a user.

| **Member**  
| Indicates the performance data member used in the report. This name corresponds to the name used on the MBR parameter of the Create Performance data (CRTPFRTDA) command.

| **Library**  
| Identifies the library where the performance data used for a particular report is located.

| **Model/Serial**  
| Indicates the model and serial number of the server on which the performance data for the report was collected. The serial number can be 10 characters.

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <b>Main storage size</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|  | Indicates the size of the main storage on the server on which the performance data was collected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|  | <b>Started</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|  | Indicates the date and time Collection Services started collecting performance data for the report.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|  | Depending on whether or not you select specific intervals or a specific starting time, you could see the following:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|  | <ul style="list-style-type: none"> <li>• If you specify no intervals at which to run the report, the start date and time is the date and time at which the data was collected.</li> <li>• If you specify specific intervals at which to run the report, the start date and time is the date and time at which the data was collected.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                          |
|  | <b>Note:</b> For the System Report only, you should consult the Report Selection Criteria section to find out which intervals were selected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|  | <b>Stopped</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|  | The date and time Collection Services stopped collecting performance data for this report.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|  | Depending on whether or not you select specific intervals or a specific ending time, you could see the following:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|  | <ul style="list-style-type: none"> <li>• If you specify no intervals at which to run the report, the stop date and time is the date and time at which the data was collected.</li> <li>• If you specify specific intervals at which to run the report, the stop date and time is the date and time at which the data was collected.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                            |
|  | <b>Note:</b> For the System Report only, you should consult the Report Selection Criteria section to find out which intervals were selected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|  | <b>System name</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|  | Indicates the name of the server on which the performance data was collected for the report.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|  | <b>Version/Release level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|  | x/ x.0 indicates which version and release level of the operating system the server was running at the time the performance data was collected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|  | <b>Partition ID</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|  | Identifies the ID of the partition on which the collection was run. This change accommodates the logical partition implementation. Here are some of the values that you might see:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|  | <ul style="list-style-type: none"> <li>• If your system is not partitioned (which is the default) or you used Collection Services to collect and print the performance data for the primary partition of a logical partition system, this value is 00.</li> <li>• If you collected data with the Start Performance Monitor (STRPFRMON) command in a previous release, the value for the partition ID is 00.</li> <li>• If you used Collection Services to collect and print the performance data in any secondary partition of a logical partition system, this value is the same as the partition ID that is shown on the Work with System Partitions display under the Start Service Tools (STRSST) command.</li> </ul> |
|  | <b>Feature Code</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|  | Identifies the Interactive feature code value for the server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|  | <b>Int Threshold</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|  | Indicates the percent of the total system CPU for interactive work that was used during the collection period. The value is obtained from the QAPMCONF file (GKEY IT) and reflects the configuration metric obtained when the collection started. You should be aware that this value may change for each interval within a collection period due to dynamic changes in logical partition configuration.                                                                                                                                                                                                                                                                                                                  |
|  | <b>Virtual Processors</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|  | The number of virtual processors configured for the partition. The value is obtained from the                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |



QAPMCONF file (GKEY 13) and reflects the configuration metric obtained when the collection started. You should be aware that this value may change for each interval within a collection period due to dynamic changes in logical partition configuration.

#### **Processor Units**

The number of processor units allocated to the partition. The value is obtained from the QAPMCONF file (GKEY PU) and reflects the configuration metric obtained when the collection started. You should be aware that this value may change for each interval within a collection period due to dynamic changes in logical partition configuration.

Processing units are a unit of measure for shared processing power across one or more virtual processors. One shared processing unit on one virtual processor accomplishes approximately the same work as one dedicated processor. One shared processing unit on two virtual processors accomplishes approximately half the work of two dedicated processors.

#### **Column headings**

Each report also has several columns that make up the information of the report. Some are specific to a particular report and others are consistent between reports. For short descriptions of these columns, see the Performance Report columns page.

## **Extended Adaptive Cache**

You can use Extended Adaptive Cache to improve system performance by collecting disk usage data, and then using those statistics to create a cache, effectively reducing the physical I/O requests for the disk.

Improve your system performance with Extended Adaptive Cache. Extended Adaptive Cache improves both the I/O subsystem and system response times by reducing the number of physical I/O requests that are read from disk. Extended Adaptive Cache generates statistical information for the data and then uses a mix of management strategies to determine which data to cache.

**Note:** In V5R4, Extended Adaptive Cache is supported on older systems. Extended Adaptive Cache uses the feature code, #4331 1.6 GB Read Cache.

### **Extended Adaptive Cache Concepts**

Explore Extended Adaptive Cache. Find information about planning, restrictions, and important considerations before you begin to use this tool.

Improve system performance with Extended Adaptive Cache, an advanced read cache technology that improves both the I/O subsystem and system response times by reducing the number of physical I/O requests that are read from disk. Extended Adaptive Cache not only improves the performance of database-read actions, but of all read actions. This includes read actions that are generated by other system components such as the Integrated xSeries Server. It also works effectively in storage subsystems that have device parity protection or mirrored protection. Extended Adaptive Cache has proven to be highly effective on many types of workloads.

### **How the Extended Adaptive Cache works**

Extended Adaptive Cache is integrated into the I/O subsystem. It operates at the disk subsystem controller level and does not affect the system processor. The storage I/O adapter manages the Extended Adaptive Cache by using a Read Cache Device (such as a solid state disk) to provide the cache memory.

Extended Adaptive Cache generates statistical information for the data, and then uses a mix of management strategies to determine which data to cache. The management of the cache is performed automatically within the I/O adapter and is designed to cache data by using a predictive algorithm. The algorithm considers how recently and how frequently the host has accessed a predetermined range of data.



The design of Extended Adaptive Cache was based on specific data management strategies of the system. Whether the disks are device parity protected, mirrored, or unprotected, the data stored on the disks has a tendency to occur in bands. This means that there are physically contiguous areas of disk storage where data is actively read, physically contiguous areas that are frequently written to, physically contiguous areas that are both actively read and written to, or physically contiguous areas of storage that are not frequently accessed.

This "banding" of data is accounted for in the Extended Adaptive Cache design. The goal is to cache bands characterized as read/write and read-only. A band that is characterized as write-only, while cached in the storage subsystem write cache, remains largely unaffected by Extended Adaptive Cache. Extended Adaptive Cache is also designed to not harm the performance of large blocks of data that are either sequentially written or sequentially read. In this instance, the pre-fetch capability of the disks, as well as other caches in the system, ensures a quick response time.

#### **Related concepts**

Device parity protection

Mirrored protection

#### **Restrictions and considerations for Extended Adaptive Cache:**

See what components Extended Adaptive Cache requires and learn more about what to expect.

Before you begin using Extended Adaptive Cache, you should do some initial planning to take into account any restrictions or considerations that may pertain to your computing environment.

#### **Restrictions**

To use Extended Adaptive Cache, your system must have the following:

- One or more storage I/O adapters that support Extended Adaptive Cache (CCIN 2780 for systems running V5R2 or later).
- Performance Tools for i5/OS licensed program for viewing the reported information.

Extended Adaptive Cache is automatically enabled on supported I/O adapters. There is no controlled on or off switch. Once the I/O adapter has been inserted into the subsystem, Extended Adaptive Cache is activated. It takes approximately an hour for Extended Adaptive Cache to monitor the data flow and populate the read cache memory. After an hour of running Extended Adaptive Cache, your system should show improved performance (depending on your current workload) and increased I/O throughput.

There are no restrictions for using Extended Adaptive Cache with regard to device parity protection and mirrored protection for other disks under the I/O adapter. Finally, Extended Adaptive Cache is designed specifically to complement Expert Cache, and may be used with or without it.

#### **Considerations**

Using the Extended Adaptive Cache allows you to attain a significant decrease in I/O response time and increase in system I/O throughput in most environments. As is the general case with caches, the system configuration and workload influence the effectiveness of Extended Adaptive Cache. Extended Adaptive Cache performs at the storage subsystem level. It caches data for the set of disks that are within that specific subsystem. Therefore, it is logical to add Extended Adaptive Cache to the most active and performance-critical storage subsystems within the system. Extended Adaptive Cache is not considered a pre-fetch type cache and therefore will not interfere with the read-ahead capabilities in the disk.

The larger the area of disk storage that is actively receiving I/O requests, the more selective Extended Adaptive Cache is about deciding when to bring new data into cache. This adaptive ability allows Extended Adaptive Cache to be effective on many workload types and sizes.

### **Related concepts**

Device parity protection

Mirrored protection

“Starting Extended Adaptive Cache”

To start Extended Adaptive Cache and increase your system’s performance, purchase the Read Cache Device.

### **Starting Extended Adaptive Cache:**

To start Extended Adaptive Cache and increase your system’s performance, purchase the Read Cache Device.

Once the Read Cache Device has been inserted into a disk slot on the subsystem, Extended Adaptive Cache will be activated. There is no user-controlled on or off switch. It takes approximately an hour for Extended Adaptive Cache to monitor the data flow and populate the Read Cache Device. After an hour of running Extended Adaptive Cache, your system should show improved performance (depending on your current workload) and increased I/O throughput.

### **Related concepts**

“Restrictions and considerations for Extended Adaptive Cache” on page 166

See what components Extended Adaptive Cache requires and learn more about what to expect.

## **Read Cache Device**

After obtaining the performance data from Extended Adaptive Cache Simulator and deciding that you want Extended Adaptive Cache to improve your system’s performance, you must purchase a Read Cache Device (RCD). Extended Adaptive Cache is automatically enabled through the RCD.

To begin using Extended Adaptive Cache, you must have:

- One or more storage I/O adapters that support Extended Adaptive Cache (CCIN 2748 for systems running V4R4 or later, or CCIN 2778 for systems running V5R1 or later, or CCIN 2757 for systems running the latest release of V5R2).
- A Read Cache Device for each storage I/O adapter that Extended Adaptive Cache is to be activated on (CCIN 6731 for systems running V4R4 or later).

Because Extended Adaptive Cache is automatically enabled through the RCD, there is no controlled on or off switch. The RCD may be added without system interruption through concurrent maintenance. The RCD resides in an internal disk slot and works with all other disk types and capacities. Be aware that all data in the Extended Adaptive Cache is also guaranteed to be on the disks. In the unlikely event of an RCD failure, there will be no data loss.

### **Related concepts**

Concurrent maintenance

## **Intelligent Agents**

The Intelligent Agents console for System i Navigator provides system administrators with an easy way to manage one or more Agent Building and Learning Environment (ABLE) agents running on a single system or across different systems.

Intelligent agents are Java-based software components that are capable of learning certain behaviors over time through complex autonomic algorithms. Intelligent agents can have many different capabilities, from simply monitoring for certain events to more complex actions like analyzing network problems, preventing unplanned system restarts, or managing storage. Although the goal of using agents is to simplify the system administrators tasks through autonomic computing, system administrators still need a way of starting, stopping, responding to, and monitoring the actions of their agents.

The Intelligent Agents console for System i Navigator provides system administrators with an easy way to manage one or more ABLE agents running on a single system or across different systems. After the agent console connects to the agent services that exist across your domain, you can monitor and work with any number of preconfigured agents on any of the systems in your domain.

## **Intelligent Agent concepts**

The Intelligent Agents console uses ABLE agents running on or across a distributed agent platform. Find out more about ABLE agents, and the agent services that make up the distributed platform.

### **ABLE agents:**

The Intelligent Agents console for System i Navigator works with Agent Building and Learning Environment (ABLE) agents.

ABLE agents are Java objects capable of automating tasks through the use of rule-based reasoning and learning certain behaviors over time by using data mining algorithms contained in the ABLE component library. ABLE is a Java framework and toolkit used for building multiagent intelligent autonomic systems, and provides specific support for developing agents that work with the System i Navigator Intelligent Agent platform and console. Intelligent agents developed using ABLE can have the following capabilities:

- Learn from experience and predict future states
- Analyze metric data using classification and clustering algorithms to detect complex states and diagnose problems
- Interface with other autonomic components via web services
- Reason using domain-specific Java application objects
- Use powerful machine reasoning, including: Boolean forward and backward chaining, predicate logic (Prolog), Rete'-based pattern match, and fuzzy systems
- Have autonomous (proactive) behavior and goals
- Correlate events into situations, reason, and take actions

The ABLE toolkit contains several examples of how to design your own agent, and a template agent is included that you can use as a model when developing your own agent. To create an agent that can be fully managed from the console, the agent should extend the `AbleEServerDefaultAgent` example.

### **Related concepts**

"Developing agents" on page 170

Create and customize your own agent to perform the tasks that you want. The Agent Building and Learning Environment (ABLE) toolkit and its associated documentation provide a working development environment and a template agent that can be used as a guide for developing your own agents.

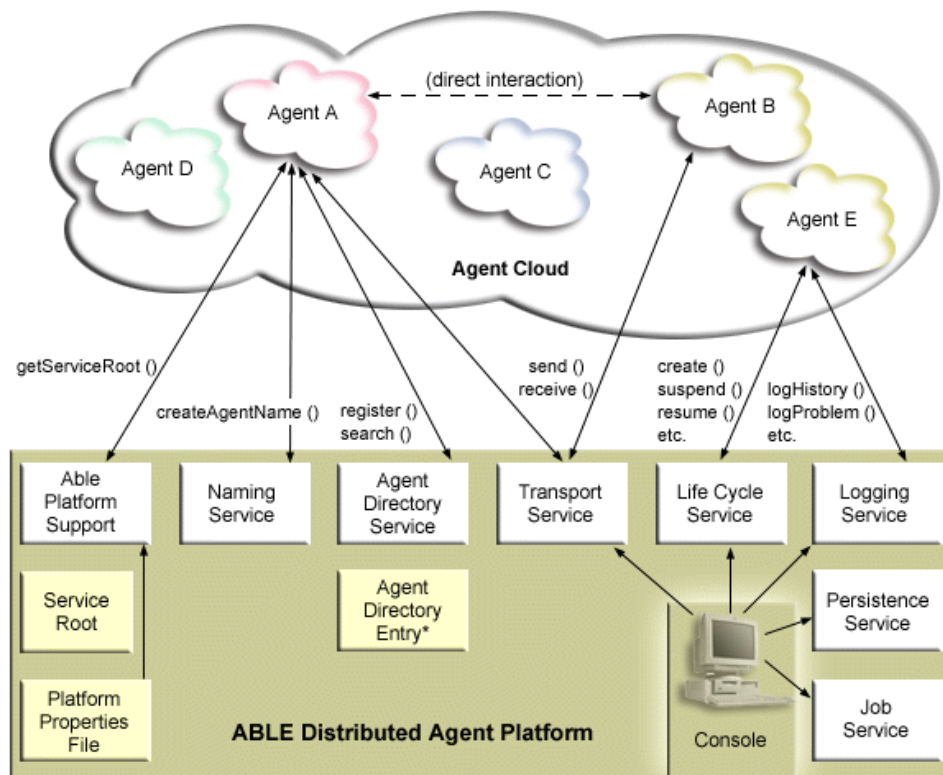
### **Agent platform:**

Agent Services live on your system or across your distributed platform, and are responsible for the life cycle, security, and behavior of your agent.

The Intelligent Agents console in System i Navigator requires that an agent platform be configured on your system, or across a distributed network. An agent platform is nothing more than a set of Java Virtual Machines, or agent pools, that run the services and agents of the platform. The platform is defined by a preferences file called `ableplatform.preferences`. This file lists the location (system and port) of each agent pool (JVM), the services that will run on or across the platform, and the agents that are allowed to run in the platform. If security is configured, the preferences file also lists the Kerberos user and service principals used to authenticate each service, agent, and user that is part of the platform.

Agent services, which can exist on any of the systems across your distributed platform, are responsible for the life cycle, security, and behavior of your agent. Agents running on the same system or distributed

agents running across different systems use the defined set of platform services for different tasks such as getting a unique name, looking up other agents in a directory, logging, and passing messages to another agent.



The following services are made available to the agents running on or across a platform and to the users connected to the platform:

- **Naming service**

This service provides the creation of a globally unique name among all other pieces in the distributed platform. The naming service also provides security for the platform when security is turned on. Kerberos is used when starting the platform to authenticate all services, pools, and users. Throughout the life of the platform, this service also acts as the trusted third party to secure all interactions between the platform's agents, services, and users.

- **Directory service**

When an agent wants to make itself known to other services and agents across the platform, it creates an agent description and registers this description to the directory service. After the agent is registered, descriptions can be modified and removed.

- **Lifecycle service**

This service is used to manage agents. Agents can be created, started, suspended, resumed, and destroyed through this service.

- **Transport service**

This service provides locators for parts of the platform. Interagent communication is also made available by this service.

- **Logging service**

A running agent may encounter a problem that requires outside intervention. The logging service creates and logs requests, and handles the corresponding answers that are sent back to it from the request. The progress of an agent can also be logged to this service for others to view.

- **Job service**

The different services and jobs of the platform register their job entry to this service. This service provides critical information about the platform when the platform is running in the i5/OS operating system.

- **Persistence service**

Services and agents may use this service to save valuable information. The naming, directory, lifecycle, logging and job services can be backed up and stored in a database when the persistence service is configured.

## **Developing agents**

Create and customize your own agent to perform the tasks that you want. The Agent Building and Learning Environment (ABLE) toolkit and its associated documentation provide a working development environment and a template agent that can be used as a guide for developing your own agents.

ABLE is a Java framework, component library, and productivity tool kit for building intelligent agents using machine learning and reasoning.

You can use the ABLE toolkit to develop your own hybrid intelligent agents. This Java framework has its own rule language called ABLE rule language (ARL) and its own GUI-based interactive development environment, the ABLE Agent Editor; both are provided to assist in the construction of ABLE agents.

### **ABLE 2.0**

Both the ABLE toolkit and complete ABLE documentation are available to download in .zip packages.

The System i Navigator Intelligent Agents console is included with a template agent that you can use as a guideline for developing agents to work with the console. The source code for AbleEserverTemplateAgent is stored in ableplatform.jar, located in QIBM/ProdData/OS400/Able.

AbleEserverTemplateAgent makes use of many of the features available when developing agents using the ABLE framework. It demonstrates how an agent would create a set of capabilities that could be managed through the console. It includes a Customize panel that can be used to alter agent settings and an About panel that is used to display information about the agent. It also shows how an agent uses the logging service to log requests and history entries that can be displayed and responded to through the console.

## **Agent capabilities**

The EServerTemplateAgent agent has the following capabilities:

- **Time monitor**

The agent watches for minute and hour changing events and take action. There are four different situations that the agent follow depending on what the capability is set to, or how the user responds to a request if one is logged:

1. Log the change without telling the time.
2. Log the change including the time as a long.
3. Log the change including the time in MM/DD/YY format
4. Do nothing

- **Duplicate request**

The agent watches for multiple hour and minute change requests. There are two different situations that the agent follows with this capability if a duplicate is found.

1. Create a duplicate request
2. Do not create a duplicate request

## Customization panel

The agent supplies a customization panel that allows you to adjust the interval at which the agent checks if the minute or hour has changed.




Figure 1. An example use of the Customization panel

## About panel

The agent supplies an about panel that allows you to provide detailed information about the agent.

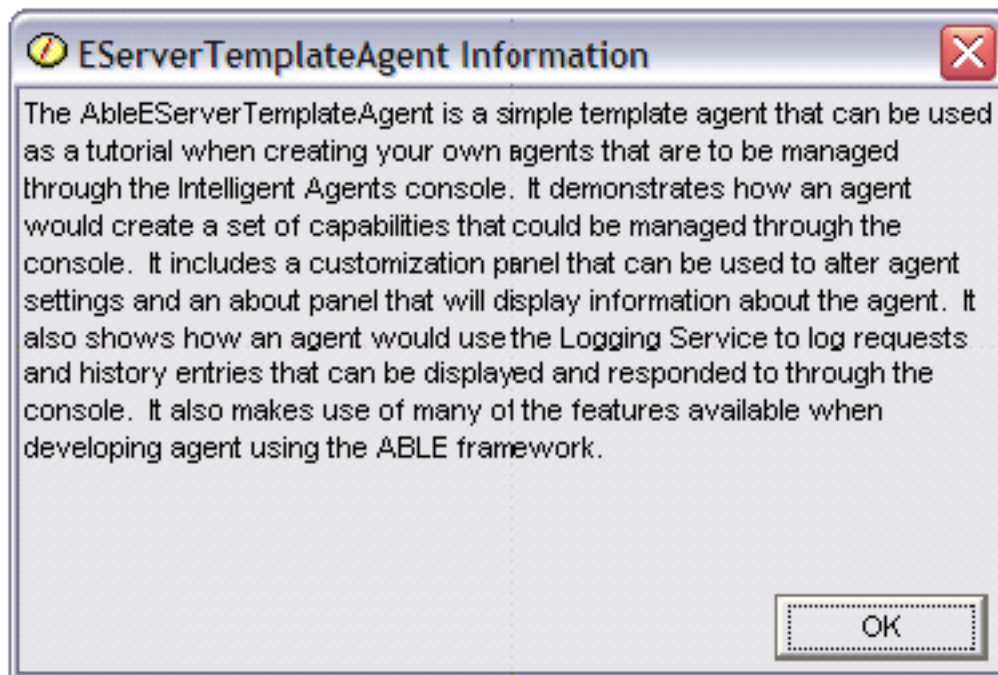


Figure 2. Viewing the template agent's about panel

## Agent Learning and Building Environment 2.0:

Agent Learning and Building Environment (ABLE) is a Java framework, component library, and productivity tool kit for building intelligent agents using machine learning and reasoning.

Both the ABLE 2.0 Toolkit and the ABLE Documentation bundle are available to download as .zip packages:

- ABLE 2.0 Toolkit: AbleAll\_2.0.0.zip

This 6 MB zipped package contains the ABLE Java framework, component library, and tool kit.

- ABLE Documentation: doc.zip



This 12 MB zipped package contains complete ABLE documentation, including an FAQ, the README, license agreement, JavaDoc, and more. Also included in doc.zip is a second zipped package (Able-Class.zip) that contains several exercises and presentations designed to help you develop ABLE agents.

## Set up your agent environment

Before you can begin managing your agents with the Intelligent Agents console, you will need to configure your agents and agent services (the agent platform) to run on or across the systems in your environment. A secure environment requires Kerberos and additional platform configuration.

The System i Navigator Intelligent Agents console functions by connecting to an agent platform running on your system, or across a distributed network. The agent platform defines the agent pools (JVMs) that your agent services and agents will run in. Before you begin setting up your agent platform, you need to determine your security preferences. A secure platform requires that you configure Kerberos.

### Related concepts

“Agent platform” on page 168

Agent Services live on your system or across your distributed platform, and are responsible for the life cycle, security, and behavior of your agent.

## Configuring your agent platform:

This topic provides a brief overview of the agent platform and also provides detailed configuration steps for modifying the platform preferences file. Before you begin using the Intelligent Agents console in System i Navigator, you first need to configure the agent platform.

## Agent platform overview

To manage agents using the intelligent agents console, you must first define, secure, and start an agent platform that the console will connect to. An agent platform is nothing more than a set of Java Virtual Machines, or agent pools, that run the services and agents of the platform. The `ableplatform.preferences` and `able.preferences` files are used to define a platform.

In its simplest form, with security turned off, `ableplatform.preferences` defines:

- The location (system and port) of each Pool.
- The services that will run in the platform.
- The agents that are allowed to run in the platform.

Once the agent platform is set up, the services that run on or across the platform allow an agent to receive a unique name, look up other agents in a directory, log history or requests, pass messages to one another, or control the state of an agent.

## Defining the agent platform

To begin configuring your platform, you must define your agent pools, agent services, permitted agents, and add Kerberos security principals by modifying the `ableplatform.preferences` file.

The default location of `ableplatform.preferences` is `QIBM/ProdData/OS400/Able`.

### Notes:

1. Multiple platforms can be configured, and you need to ensure that your platform does not reside at the same location as an existing platform using the same port. See the Start the agent platform topic for more details.
2. When you open the file and begin making changes to the content, understand that small errors and misspellings will cause the agent platform to fail, and there is currently no easy



way to debug your mistakes. Avoid commenting out properties that are unused, commenting out an unused property can cause the platform to fail. For example, if you choose to run the platform with security turned off, do not comment out the principal properties through the file.

The following code samples taken from `ableplatform.preferences` provide examples of how to modify the platform preferences. To configure your platform, follow these steps:

## 1. Define agent pools.

A platform is nothing more than a set of distributed Java Virtual Machines. Each JVM is called an agent pool, and each JVM or pool can host multiple services and agents (an agent pool does not have to host services, it could be used to run just agents). You must define the location of each of your Java Virtual Machines (agent pools) in the preferences file by specifying the IP address (fully qualified system name) and port. Also, specify an alias (any unique name) for each agent pool. When security is turned on, you must associate a service principal with each agent pool. For more information about using Kerberos service principals, see “Securing your agent environment” on page 175. The following is an example of how a set of agent pools could be defined:

```
#-----
# Java Virtual Machines
#-----
AgentPool.1.Alias      = Pool1
AgentPool.1.Address    = systemname.ibm.com
AgentPool.1.Port       = 55551
AgentPool.1.Principal  = servicePrincipal1

AgentPool.2.Alias      = Pool2
AgentPool.2.Address    = systemname.ibm.com
AgentPool.2.Port       = 55552
AgentPool.2.Principal  = servicePrincipal1

AgentPool.3.Alias      = Pool3
AgentPool.3.Address    = systemname.ibm.com
AgentPool.3.Port       = 55553
AgentPool.3.Principal  = servicePrincipal2
#-----
```

## 2. Define agent services.

Define the agent services that you want to run on the platform, and specify the alias of the agent pool you want them to run in. Each agent service must point to a factory. The *factory* is a Java Class that creates the agent service. The Persistence service is used to restart a platform to its previous state. Specify to turn persistence on or off. If you turn persistence on, you must specify a Database, Table, and Schema so that persistence has a location to store backed up data on. You can also specify a value for the PersistenceRetry property. If the persistence service fails and you specified a value of 5000 for the PersistenceRetry property, it will attempt to retry every 5000 milliseconds. The following code example shows how three different services, Directory, Logging, and Persistence, could be defined:

```
Services=Agent-Directory-Service,Agent-Logging-Service,
Persistence-Service

Agent-Directory-Service.AgentPool      = Pool1
Agent-Directory-Service.Factory =
com.ibm.able.platform.RMIVerifiableDirectoryServiceFactory
Agent-Directory-Service.Persistence    = off
Agent-Directory-Service.PersistenceDatabase = *LOCAL
Agent-Directory-Service.PersistenceTable  = qahadir
Agent-Directory-Service.PersistenceSchema = QUSRSYS
Agent-Directory-Service.PersistenceRetry  = 5000

Agent-Logging-Service.AgentPool      = Pool1
Agent-Logging-Service.Factory =
com.ibm.able.platform.RmiAgentLoggingServiceFactory
Agent-Logging-Service.Persistence    = off
Agent-Logging-Service.PersistenceDatabase = *LOCAL
```

```

Agent-Logging-Service.PersistenceTable      = qahalog
Agent-Logging-Service.PersistenceSchema    = QUSRSYS
Agent-Logging-Service.PersistenceRetry     = 5000
Agent-Logging-Service.Properties           = history-log-max : 100

```

**Note:** You can specify to control performance by adding a history-log-max property to the Logging service. If you specify history-log-max=100, each agent keeps only its 100 most current history logs.

```

Persistence-Service.AgentPool      = Pool1
Persistence-Service.Factory =
com.ibm.able.platform.RmiPlatformPersistenceServiceFactory
Persistence-Service.Properties =
persistence-driver : com.ibm.db2.jdbc.app.DB2Driver,
persistence-protocol : jdbc,
persistence-subProtocol : db2,
blob-type : BLOB,
persistence-dbFlushTime : 1000,
persistence-dbResetAll : off

```

The Persistence service provides backup and recovery for your agent platform. To use persistence with agent services running on or across your platform, you need to define several Persistence-Service.Properties:

- **persistence-driver**  
Defines the JDBC driver that the persistence service will use. By default the persistence-driver is set to use the integrated DB2® driver.
- **persistence-protocol and subProtocol**  
Defines the database protocol that the persistence service will use. By default the protocol is set to jdbc and the subProtocol is set to db2.
- **blob-type**  
Defines the BLOB type associated with the JDBC driver you are using. The default for DB2 is set to BLOB, but if you choose to use a different database, like CloudScape for example, you would define BLOB type as blob-type : LONG VARBINARY.
- **persistence-dbFlushTime**  
Specifies the rate at which you want the persistence service to flush data to the database in milliseconds.
- **persistence-dbResetAll**  
Specifies if all previously persisted data will be cleared from the database when you restart the platform.

### 3. Defining permitted agents.

You must define all of the agents that you want to allow access to the platform and the agent services running on or across the platform. The following is an example of how an agent could be defined. More details about each agent property are listed after the following example:

```

Agent.1.Alias=Agent1
Agent.1.AutonomyLevel=Medium
Agent.1.ClassName=
com.ibm.able.platform.examples.EServerTemplateAgent
Agent.1.ConstructorArgs=String:agentName
Agent.1.EligiblePrincipals=principalAlias1, principalAlias2
Agent.1.EligibleAgentPools=pool1, pool2, pool3
Agent.1.InitArgs=
Agent.1.LastChangedDate=January 11, 2003 11:11am
Agent.1.Type=Tester1
Agent.1.Vendor=IBM1
Agent.1.Version=1.1

```

- **Alias**  
Provide a unique name for your agent. This name will be used by the agent console.
- **AutonomyLevel**

Specify the agents initial autonomy level. A user can change this setting from the console. Determine the level of independence you want to associate with your agent, and set the automation level accordingly. The higher you set the automation level, the less your agent will request permission to take an action. If you set an agent to High automation, it will perform most actions without requesting a response first. If you are concerned about an agent's behavior, you may want to lower the automation level (increasing the frequency by which the agent requests permission to take action), by changing the setting to Medium automation.

- **ClassName**

Specifies the the actual agent Java Class.

- **ConstructorArgs**

Allows you to provide arguments in the properties file that you want to pass to your agent.

- **EligiblePrincipals**

When security is turned on, you must define who has authority to start an instance of your agent by associating one or more user principal aliases with each agent. For more information about using Kerberos service principals, see "Securing your agent environment."

- **EligibleAgentPools**

Specify the alias of one or more agent pools that you want to use to run your agents on the platform.

- **InitArgs**

Allows you to pass in any Init arguments to your agent from the preferences file.

#### 4. Secure your agent platform.

After you have defined your agent pools, agent services, and permitted agents, you may want to configure security on the platform. For more information on Kerberos principals, trust levels, and how they are used and defined to secure the agent platform, see "Securing your agent environment."

After you have defined your agent pools, agent services, and permitted agents, and optionally set up security, you need to start the agent platform.

##### **Related concepts**

"Agent platform" on page 168

Agent Services live on your system or across your distributed platform, and are responsible for the life cycle, security, and behavior of your agent.

##### **Related tasks**

"Starting the agent platform" on page 180

After you define the agent platform and optionally secure your platform, you need to start all the Java Virtual Machines associated with your agent services using i5/OS CL commands.

#### **Securing your agent environment:**

It is strongly recommended that you use Kerberos user and service principals to authenticate users, agent pools, and agent services to one another on or across a secure platform or distributed platform.

Platform security can be turned on or off. If you choose to run on or across a platform that has security turned off, anyone can deregister or modify another person's agent descriptions. Anyone can change the capabilities or state of any agent. Anyone can remove or answer any requests, even if they are not their own. Agents can potentially take destructive actions when being used incorrectly or by the wrong user. To ensure that agents are used the way they were intended, security features have been added to the infrastructure of the platform.

When security is turned on, agents and services can authenticate and authorize every action that is taken on or across the platform. An agent can only deregister or alter its own agent description. An agent must authorize all answered requests and capability changes. A certain authority level is required to alter the

state of an agent. The use of an agent can be limited to certain users and locations. When security is turned on, every action that occurs can be traced back to a known user so that platform authentication and authorization can occur.

If you choose to secure your agent platform, you can turn security on by changing the Security property to Security=on in the able.preferences file that defines your platform.

#### *Configuring your platform to use Kerberos:*

The intelligent agent platform uses Kerberos principals to authenticate users and services throughout the agent platform. Kerberos protocol, developed by Massachusetts Institute of Technology, allows a principal (a user or service) to prove its identity to another service within an insecure network.

Authentication of principals is completed through a centralized server called a key distribution center (KDC). The KDC authenticates a user with a Kerberos ticket. These tickets prove the principal's identity to other services in a network. After a principal is authenticated by these tickets, they can exchange encrypted data with a target service.

The platform uses Kerberos to authenticate user sign on and initial platform startup. To use Kerberos to secure your platform, you must either find an existing KDC, or create a working KDC that all parts of the platform use. Every system running a piece of the platform and every PC running a console that connects to this platform must be configured to use this KDC. You need to list all Kerberos principals in the ableplatform.preferences file that are used by the platform to authenticate users and services. Each platform Java Virtual Machine (agent pool) has a service principal associated with it, and each user logging onto the platform from a console needs a user principal. All of these principals need to be added to the KDC.

#### 1. Find or create a usable Kerberos key distribution center (KDC).

The agent platform does not require a KDC on i5/OS. A KDC running on any platform will work. If you cannot find an existing KDC to use, you can create your own. In V5R3 or later, i5/OS supports a Kerberos server in i5/OS PASE. You can configure and manage a Kerberos server from your system. To configure a Kerberos server in i5/OS PASE, complete the following steps:

- a. In a character-based interface, type `call QP2TERM`. This command opens an interactive shell environment that allows you to work with i5/OS PASE applications.
- b. At the command line, enter `export PATH=$PATH:/usr/krb5/sbin`. This command points to the Kerberos scripts that are necessary to run the executable files.
- c. At the command line, enter `config.krb5 -S -d iseriesa.myco.com -r MYCO.COM`. This command updates the `krb5.config` file with the domain name and realm for the Kerberos server, creates the Kerberos database within the integrated file system, and configures the Kerberos server in i5/OS PASE. You are prompted to add a database master password and a password for the `admin/admin` principal, which is used to administer the Kerberos server.
- d. At the command line, enter `/usr/krb5/sbin/start.krb5` to start the servers.

#### 2. Configure systems in your agent environment to use Kerberos.

After you create a Kerberos server (KDC), you need to individually configure each client PC that will attempt to connect to the secure platform, and each system in your agent platform to point to your Kerberos server (KDC).

##### • **Configure your client PC**

To configure a client PC, you need to create a text file called `krb5.conf` in the security folder of the JVM that runs your System i Navigator intelligent agents console located here (where C: is the drive where your System i Access driver is installed):

C:\Program Files\IBM\Client Access\JRE\Lib\Security

The krb5.conf file tells all JVMs started from this JRE which KDC to use when dealing with Kerberos. The following is an example of what a generic krb5.conf file might look like if the KDC realm is KDC\_REALM.PASE.COM and is found on system1.ibm.com:

```
[libdefaults]
default_realm          = KDC_REALM.PASE.COM
default_tkt_enctypes   = des-cbc-crc
default_tgs_enctypes   = des-cbc-crc

[realms]
KDC_REALM.PASE.COM = {
    kdc = system1.rchland.ibm.com:88
}

[domain_realm]
.rchland.ibm.com = KDC_REALM.PASE.COM
```

- **Configure your system**

To point your system to your KDC, you need to modify the following file:

/QIBM/userdata/OS400/networkauthentication/ krb5.conf

The krb5.conf file tells all JVMs started from this JRE which KDC to use when dealing with Kerberos. The following is an example of what a generic krb5.conf file might look like on the server if the KDC realm is KDC\_REALM.PASE.COM and is found on system1.ibm.com:

```
??(libdefaults??)
    default_realm = KDC_REALM.PASE.COM
??(appdefaults??)
??(realms??)
    KDC_REALM.PASE.COM = {
        kdc = system1.rchland.ibm.com:88
    }
??(domain_realm??)
    system1.rchland.ibm.com = KDC_REALM.PASE.COM
```

### 3. Acquire Kerberos user and service principals.

After you configure a KDC, you need to create the user and service principals you plan to use to secure the platform, and register these principals to the KDC:

#### Service principals:

Each agent pool (JVM) defined in the ableplatform.preferences file must have a service principal associated with it. Service principals are specific to the system that they run on, so they must include that system name and be in the following format: ServicePrincipalName/systemName@KDCRealm. Each of the agent pools on the platform can use the same service principal, or you can specify that each pool use its own service principal. If each of your agent pools has different authority levels, then different principals should be used for each different authority level.

#### User principals:

Each user that you want to allow to connect to the secure platform through the console needs a user principal. User principals can be associated with each agent definition listed in the ableplatform.preferences file. A user principal can connect to a platform from the console, regardless of the system the console is running on. Because of this, a user principal only needs to include the principal name and the KDC realm that the principal belongs to: UserPrincipalName@KDCRealm.

You need to add a principal to the KDC for each service and user principal that your platform uses. To add your principals to your KDC if you are using the integrated KDC on the server, follow these steps:

- In a character-based interface, type `call QP2TERM`.
- At the command line, enter `export PATH=$PATH:/usr/krb5/sbin`. This command points to the Kerberos scripts that are necessary to run the executable files.
- At the command line, type `kadmin -p admin/admin`, and press **Enter**.

- d. Sign in with administrator's password.
- e. Enter the following at a command line:
  - To add service principals for pools running on a system:  
`addprinc -pw secret servicePrincipalName/fully qualified host name@REALM`
  - To add user principals:  
`addprinc -pw secret jonesm`. This creates a principal for a user to log in from a console.
  - To add service principals for pools running on a PC:  
`addprinc -requires_preauth -e des-cbc-crc:normal -pw host/pc1.myco.com`.

If you are using the integrated KDC, see the following topics for more information on how to add principals to your KDC:

- If you are adding service principals for pools that run on a system, see [Add i5/OS principals to the Kerberos server](#).
  - If you are adding user principals or service principals for pools that run on a PC, see [Create Host principals for Windows 2000 workstations and users](#)
4. Add service principals to each keytab file.

When you start up a secure platform, each agent pool uses the principal that it was defined to start with, and uses it to authenticate itself. This requires each pool JVM to have access to valid Kerberos credentials for the principal it is using. The i5/OS Start Agent Services (STRAGTSRV) command handles this, as long as there is an entry in the keytab file for the principal that is being used.

To add an entry to the keytab file for each service principal when you are running the integrated KDC on a system, do the following:

- a. In a character-based interface, type STRQSH. This command starts the qsh shell interpreter.
- b. Enter the following command (where *ServicePrincipal* is the name of the service principal you want to add, *system@KDCRealm* is the fully qualified system name and Kerberos realm, and *thePassword* is the password associated with your service principal): `keytab add ServicePrincipal/system@KDCRealm -p thePassword`

After you set up your KDC and create your user and service principals, you need to configure security in your `ableplatform.preferences` file.

#### Related tasks

[Configure a Kerberos server in i5/OS PASE](#)

[Configure network authentication](#)

["Configuring platform security"](#)

Before you begin, ensure that you have configured your Kerberos key distribution center (KDC).

#### Configuring platform security:

Before you begin, ensure that you have configured your Kerberos key distribution center (KDC).

When security is turned on, `ableplatform.preferences` acts as a policy file for the security of the platform it defines. The following steps provide examples for how principals, trust levels, and permissions could be configured:

1. Define user and service principals.  
 After you acquire user and service principals, and register them with your KDC, you need to add these principals to the `ableplatform.preferences` file. When security is turned on, a user must be defined with a valid Kerberos user principal to gain access to the platform, and all agent services and agent pools must have a valid Kerberos service principal assigned to them. Add the user or service principals you have registered with your KDC, and specify an alias for each principal (the alias can be any unique name you want to use).



```
#-----
# Principals
#-----
Principal.1.Alias      = servicePrincipal1
Principal.1.Principal  = name1/systemName@REALM

Principal.2.Alias      = servicePrincipal2
Principal.2.Principal  = name2/systemName@REALM

Principal.3.Alias      = userPrincipal1
Principal.3.Principal  = name1@REALM

Principal.4.Alias      = userPrincipal2
Principal.4.Principal  = name2@REALM
```

## 2. Define trust levels.

After you add user and service principals, you need to define the trust level associated with each principal. A trust level is associated with a principal to help define the capabilities of a user or service on the platform. Associating a trust level with a principal is also a way to group principals. The same trust level can be associated with multiple user and service principals. Add the principal alias you assigned to your service and user principals in step 1 (comma delineated), to the trust level you want to associate it with, and provide a unique name for trust level alias.

```
#-----
# Trust Levels
#-----
TrustLevel.1.Alias      = HighlyTrusted
TrustLevel.1.Principals = servicePrincipal1,userPrincipal1

TrustLevel.2.Alias      = SomewhatTrusted
TrustLevel.2.Principals = servicePrincipal2,userPrincipal2
```

## 3. Associate service principals with agent pools.

A distributed platform can span multiple ports on multiple systems. Each agent pool defines where one part (JVM) or the platform will run. Each agent pool entry contains an alias, an IP address, a port, and a service principal alias. The principal alias specifies what service principal this pool is associated with. Add the service principal alias that you defined above to associate it with your agent pool.

```
#-----
# Agent Pools (Java Virtual Machines)
#-----
AgentPool.1.Alias      = Pool1
AgentPool.1.IPAddress  = systemname.ibm.com
AgentPool.1.Port       = 55551
AgentPool.1.Principal  = servicePrincipal1

AgentPool.2.Alias      = Pool2
AgentPool.2.IPAddress  = systemname.ibm.com
AgentPool.2.Port       = 55552
AgentPool.2.Principal  = servicePrincipal1

AgentPool.3.Alias      = Pool3
AgentPool.3.IPAddress  = systemname.ibm.com
AgentPool.3.Port       = 55553
AgentPool.3.Principal  = servicePrincipal2
```

## 4. Define agent startup authority.

Define which users have the capability to start each of the agents defined on your secure platform. Add one or more user principal aliases to the EligiblePrincipal parameter.

```
#-----
# Permitted Agents
#-----
Agent.1.Alias=Agent1
Agent.1.AutonomyLevel=Medium
Agent.1.ClassName=com.ibm.able.platform.examples.EServerTemplateAgent
Agent.1.ConstructorArgs=String:AgentName1
```



```

Agent.1.EligiblePrincipals=userPrincipal1,userPrincipal2
Agent.1.EligibleAgentPools=Pool2,Pool3
Agent.1.InitArgs=
Agent.1.LastChangedDate=January 11, 2003 11:11am
Agent.1.Type=Tester1
Agent.1.Vendor=IBM1
Agent.1.Version=1.1

```

## 5. Define the algorithm and provider.

You need to define the algorithm and provider of the KeyPairs the platform will use. By default, the preferences file will contain the following setting:

```

#-----
# Cryptography parameters
#-----
CryptographyAlgorithm = DSA
CryptographyProvider  = IBMJCE

```

After you add the necessary security data to the ableplatform.preferences file, save your changes. Turning on security for the platform once it is correctly configured is as simple as opening the able.preferences file that defines your platform, and changing the Security property to Security=on. If you are running an unsecured platform, you need to end and restart the agent platform for security changes to take effect.

### Related tasks

“Configuring your platform to use Kerberos” on page 176

The intelligent agent platform uses Kerberos principals to authenticate users and services throughout the agent platform. Kerberos protocol, developed by Massachusetts Institute of Technology, allows a principal (a user or service) to prove its identity to another service within an insecure network.

“Starting the agent platform”

After you define the agent platform and optionally secure your platform, you need to start all the Java Virtual Machines associated with your agent services using i5/OS CL commands.

## Starting the agent platform:

After you define the agent platform and optionally secure your platform, you need to start all the Java Virtual Machines associated with your agent services using i5/OS CL commands.

Because the platform is made up of one or more Java Virtual Machines, to start the platform you need to start all of the JVMs that make up the platform.

1. Use the Start Agent Services (STRAGTSRV) command to start the agent platform.
2. Use the End Agent Services (ENDAGTSRV) command to stop the agent platform.

**Note:** If you have trouble starting or ending the agent platform, you can turn on tracing for the startup programs by adding or setting the QAHA\_TRACE system environment variable to '1'. This will create log files in QUSRSYS/QAAHALOG. Files named QSBR<job number>, QSBE<job number>, and QEND<job number> will be created for each QAHASBMTER, QAHASBMTEE, and QAHAPLTEND job that has run.

### Related tasks

“Configuring platform security” on page 178

Before you begin, ensure that you have configured your Kerberos key distribution center (KDC).

“Configuring your agent platform” on page 172

This topic provides a brief overview of the agent platform and also provides detailed configuration steps for modifying the platform preferences file. Before you begin using the Intelligent Agents console in System i Navigator, you first need to configure the agent platform.

“Securing your agent environment” on page 175

It is strongly recommended that you use Kerberos user and service principals to authenticate users, agent pools, and agent services to one another on or across a secure platform or distributed platform.

### Related reference

Start Agent Services (STRAGTSRV) command

See the Start Agent Services (STRAGTSRV) command for more information about starting the agent services.

End Agent Services (ENDAGTSRV) command

See the End Agent Services (ENDAGTSRV) command for more information about ending the agent services.

## Managing agents

Use the agent console to connect to your domain and begin managing your agents. Find out how to control the level of automation associated with your agents, and how to easily respond to requests and track agent history.

The Intelligent Agents console is a powerful management tool that allows you to work with your agents, and ensure that they are behaving in a manner that meets your expectations. To display the Intelligent Agents node in System i Navigator, select **View** → **Intelligent Agents** from the main menu.

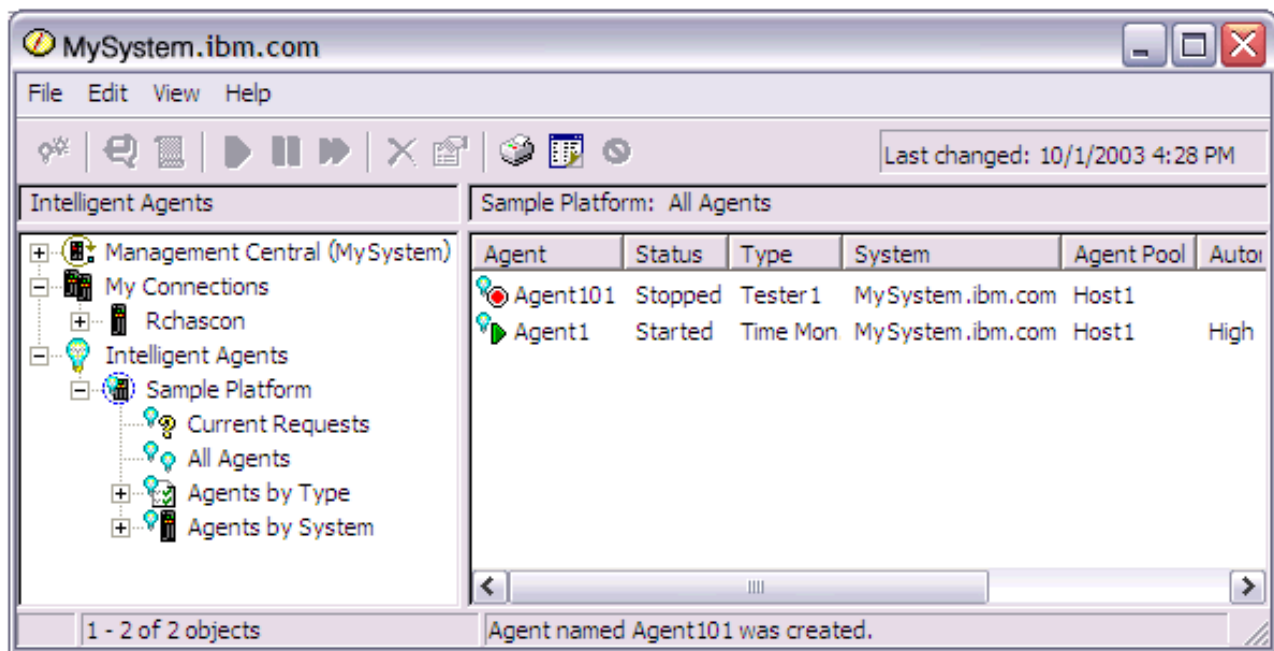


Figure 3. Working with agents in System i Navigator

After you set up your agent environment, you can begin working with the agent console by connecting to your host system (or systems) and creating an instance of an agent to run on that system. Use the console to start, stop, suspend, delete, respond to, and view the history of the agents running on your system or systems. You can also use the console to set up limitations on what actions an agent can perform automatically and what actions require permission.

### Automating agents:

The agent console gives you the capability to control and customize an agent's behavior by associating a level of automation with that agent.

The Intelligent Agents console provides a way for you to control the automated actions an agent can take.

To view an agent's capabilities, and change an agent's automation setting in System i Navigator, follow these steps:

1. Expand **Intelligent Agents**.

2. Expand your intelligent agent's platform.
3. Select **All Agents**.
4. Right-click the agent you want to work with and select **Properties**.
5. Select the **Automation** tab to display the agent's currently configured automation level.
6. Click **Capabilities** to display a list of the actions this agent can take, and the automation level associated with these capabilities.

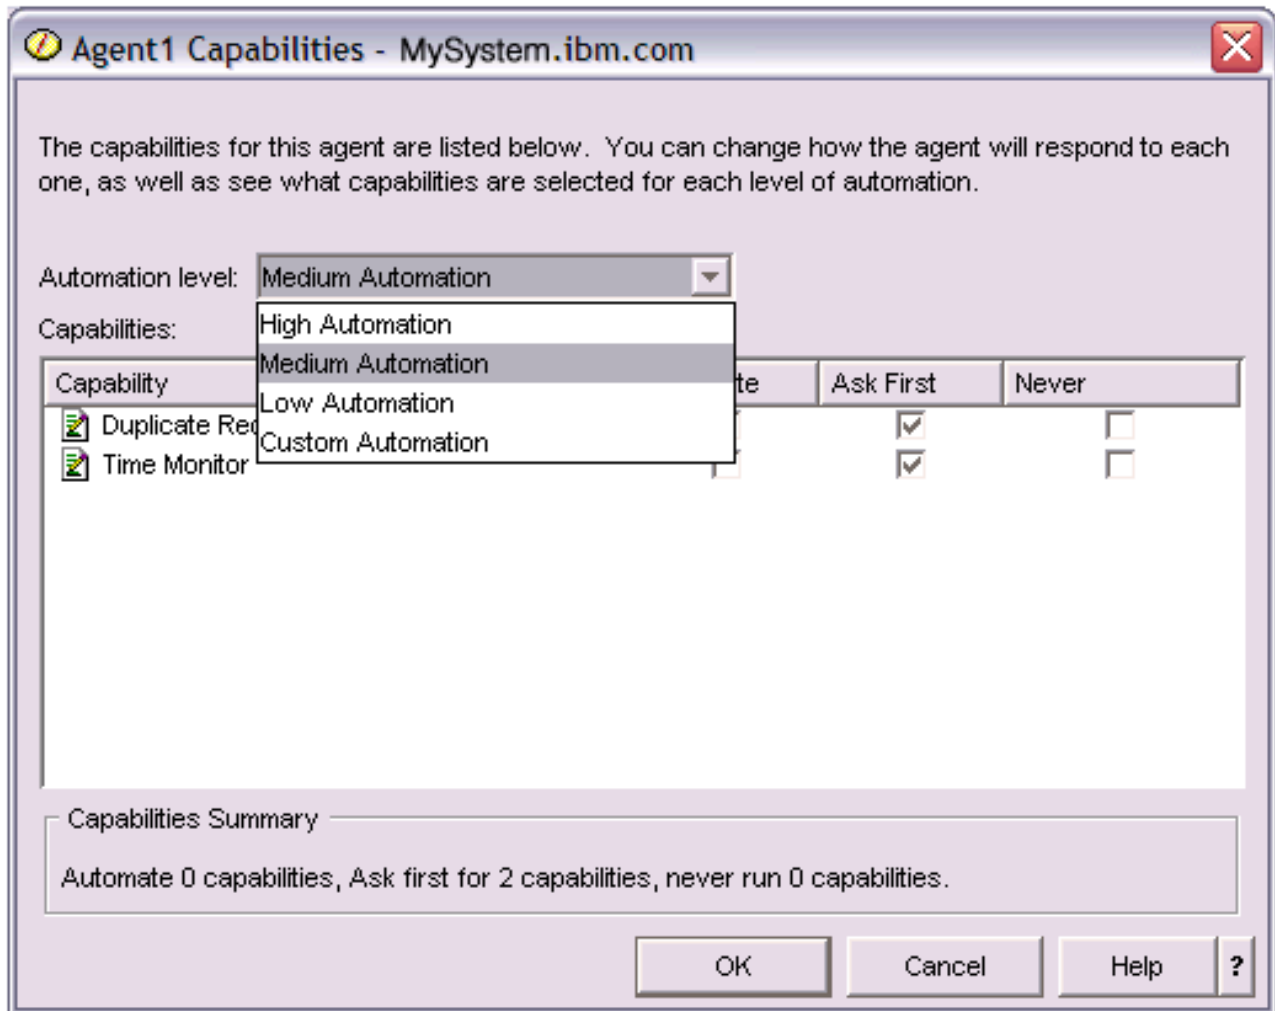


Figure 4. Viewing the automation level associated with the capabilities of a TimeMonitor agent

Every agent has a set of capabilities that define what kinds of actions that agent can perform. The agent console displays an agent's available capabilities associated with the agent's corresponding automation level. Each automation level setting (High automation, Medium automation, Low automation, and Custom automation) will change the states (Automate, Ask first, Never ask) of the available capabilities for the agent.

For example, if an agent has the capability to clear log files when full, when you change the level of automation from **High automation** to **Medium automation**, the agent's capability changes from a state of **Automate** to a state of **Ask first**. The agent now requests permission before it deletes a log file.

Specifying an agent's automation level will determine if an agent automatically performs an action, asks before performing an action, or never performs an action. The possible automation values are:

- **High automation**

The agent will perform most actions automatically, but will ask before performing certain destructive actions. Depending on the agent, certain actions may require that the agent always request outside intervention before it performs that action, even when set to **High automation**.

- **Medium automation**

The agent will perform some actions automatically, and will ask before performing some actions. Depending on the agent, certain actions may require that the agent always request outside intervention before it performs that action, even when set to **Medium automation**.

- **Low automation**

The agent will rarely perform any actions automatically. The agent will almost always request outside intervention before it performs any action.

- **Custom automation**

The agent will automate, ask first, or never perform actions according to how the capabilities are manually configured.

### **Communicating with an agent:**

Easily track and respond to agents that are requesting confirmation or permission to take action.

If the automation setting associated with an agent's capability is set to **Ask first**, before an agent performs an action, the agent will request a response from a user. Some agents will always request a response, regardless of their current automation setting. When an agent requests a response or is waiting to perform an action, the agent's Status field displays **Needs response**.

To respond to an agent in System i Navigator:

1. Expand **Intelligent Agents**.
2. Expand your intelligent agents platform.
3. Select **All Agents**.
4. Right-click the agent and select **Respond**.
5. Select the response you want to work with and click the **Respond** button. The agent will display the problem it is currently seeking a response for.
6. Select a response from the list of possible responses in the **Response** field, and click **OK**.




Figure 5. Responding to your agent's request

You can also view a list of all current requests by selecting **Current Requests** under the main **Intelligent Agents** menu.

#### Viewing agent history:

The agent console logs a history of all your agent's actions.

The agent console allows you to view the history of an agent's requests and actions. The history does not display current requests, only requests and actions that have been responded to. The history log is limited to 1000 entries, and will clear the oldest entry for each new entry that exceeds 1000.

To view an agent's history in System i Navigator, follow these steps:

1. Expand **Intelligent Agents**.
2. Expand your intelligent agents platform.
3. Select **All Agents**.
4. Right-click the agent that you want to view the history of, and select **History**.

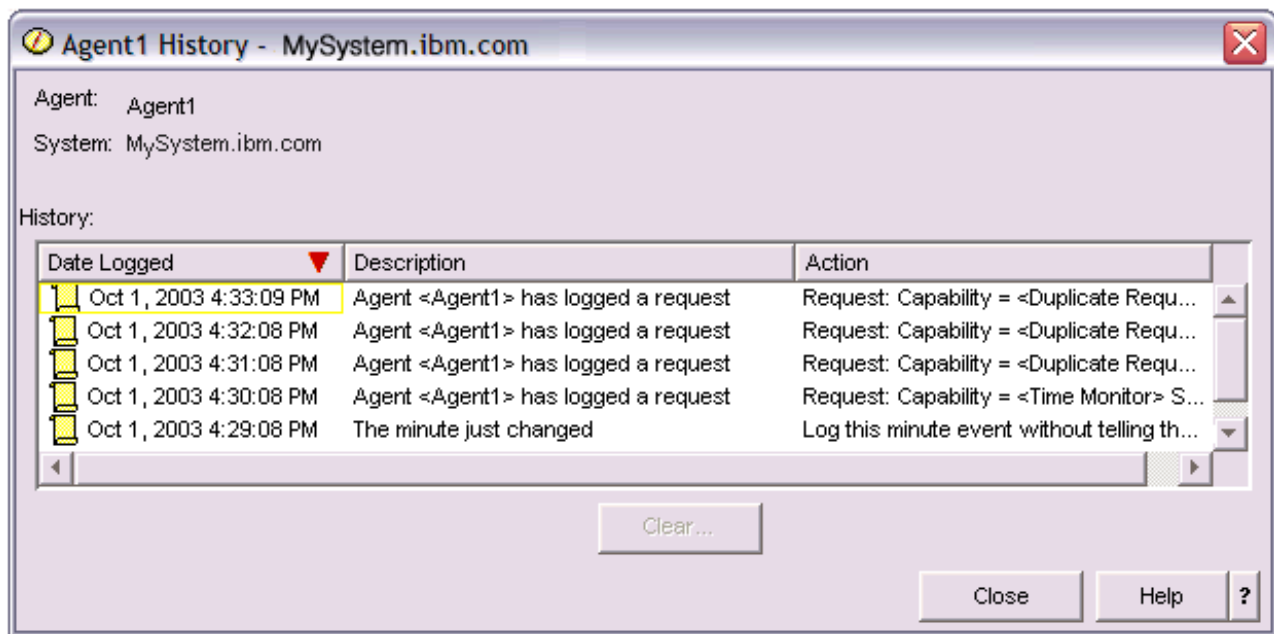


Figure 6. Viewing the history of an agent's requests and actions

## Scenarios: Performance

One of the best ways to learn about performance management is to see examples that illustrate how you can use these applications or tools in your business environment.

### Scenario: Improving system performance after an upgrade or migration

In this scenario, you have just upgraded or migrated your system and it now appears to be running slower than before. This scenario will help you identify and fix your performance problem.

#### Situation

You recently upgraded your system to the newest release. After completing the upgrade and resuming normal operations, your system performance has decreased significantly. You would like to identify the cause of the performance problem and restore your system to normal performance levels.

#### Details

Several problems may result in decreased performance after upgrading the operating system. You can use the performance management tools included in i5/OS and Performance Tools licensed program (5761-PT1) to get more information about the performance problem and to narrow down suspected problems to a likely cause.

1. Check CPU utilization. Occasionally, a job will be unable to access some of its required resources after an upgrade. This may result in a single job consuming an unacceptable amount of the CPU resources.
  - Use WRKSYSACT, WRKSYSSTS, WRKACTJOB, or System i Navigator system monitors to find the total CPU utilization.

- If CPU utilization is high, for example, greater than 90%, check the amount of CPU utilized by active jobs. If a single job is consuming more than 30% of the CPU resources, it may be missing file calls or objects. You can then refer to the vendor, for vendor-supplied programs, or the job's owner or programmer for additional support.
2. Start a performance trace with the STRPFRTTC command, and then use the system and component reports to identify and correct the following possible problems:
    - If the page fault rate for the machine pool is higher than 10 faults/second, give the machine pool more memory until the fault rate falls below this level.
    - If the disk utilization is greater than 40%, look at the waiting and service time. If these values are acceptable, you may need to reduce workload to manage priorities.
    - If the IOP utilization is greater than 60%, add an additional IOP and assign some disk resource to it.
    - If the page faults in the user pool are unacceptably high, you might want to automatically tune performance.
  3. Run the job summary report and refer to the Seize lock conflict report. If the number of seize or lock conflicts is high, ensure that the access path size is set to 1TB. If the seize or lock conflicts are on a user profile, and if the referenced user profile owns many objects, reduce the number of objects owned by that profile.

## 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 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 System i 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:  
`SNDDMSG 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: Message monitor

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

## Situation

Your company has several systems, 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 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 System i 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 systems 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.

## 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 system, 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 System i 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 system, 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 193.

1. In System i 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 193.

```
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.

---





## Related information for Performance

Listed here are the product manuals and IBM Redbooks (in PDF format), Web sites, and information center topics that relate to the Performance topic. You can view or print any of the PDFs.








### Manuals

- Performance Tools for iSeries   
This book provides the programmer with the information needed to collect data about the system, job, or program performance. It also includes tips for printing and analyzing performance data to identify and correct inefficiencies that might exist as well as information about the Manager and Agent features.
- System i Performance Capabilities Reference   
This reference provides highly technical information about server performance useful for performance benchmarking, capacity planning, and planning for server performance.



### IBM Redbooks

- Performance Management for IBM eServer™ iSeries and pSeries: A Systems Management Guide   
The topics in this IBM Redbooks publication support the growing need and recommendation to treat IBM Performance Management as a fundamental yet critical component of a systems management strategy. It explains how you can make the Performance Management offering a part of your systems management strategy, how you use the interactive offerings, and how you navigate in the components of the offering.
- IBM eServer iSeries Performance Management Tools   
Learn about the complete array of IBM iSeries performance management tools! This IBM Redpaper is designed to help you understand the different performance management tools at the IBM i5/OS V5R3M0 level, that are available to you and when to use them.
- AS/400 HTTP Server Performance and Capacity Planning   
The Internet and Web browser-based applications have had a profound effect on how organizations distribute information, perform business processes, service customers, and reach new markets. This book is intended for System i programmers, network and system management professionals, and other information technologists who are responsible for designing, developing, and deploying Web-based applications and information systems.
- AS/400 Performance Explorer Tips and Techniques 

This document provides descriptions and detailed examples of the performance explorer capabilities that were available for V3R6. Specific application examples and reports are provided.

- AS/400 Performance Management   
This document describes a methodology for performance management. It includes setting up performance objectives, collecting and reviewing performance data, tuning of resources, and capacity planning. Performance guidelines and application design tips are also provided.
- DB2 UDB/WebSphere Performance Tuning Guide   
This document provides an overview of WebSphere Application Server architecture and its main components and introduces some of its key application tuning parameters and system tuning parameters.
- IBM eserver iSeries Universal Connection for Electronic Support and Services   
This document introduces Universal Connection. It also explains how to use the variety of support tools that report inventories of software and hardware on your machine to IBM so you can get personalized electronic support, based on your system data.
- Java and WebSphere Performance on IBM eserver iSeries Servers   
This document provides tips, techniques, and methodologies for working with Java and WebSphere Application Server performance-related issues.
- Lotus® Domino for AS/400: Performance, Tuning, and Capacity Planning   
This document describes a methodology for performance management. It includes setting up performance objectives, collecting and reviewing performance data, tuning of resources, and capacity planning. Performance guidelines and application design tips are also provided.
- Managing OS/400 with Operations Navigator V5R1, Volume 1: Overview and More   
This volume presents an overview of Operations Navigator V5R1. It covers such things as managing jobs, subsystems, job queues, and memory pools; monitoring system performance metrics; jobs and messages; and Collection Services.
- Managing OS/400 with Operations Navigator V5R1, Volume 5: Performance Management   
This volume builds on the monitor, graph history, and Collection Services capabilities described in Volume 1. This book shows how to use these functions in an application environment.

## Web sites

- Performance Management for IBM System i Resource Library  ([www.ibm.com/servers/eserver/iseriess/perfmgmt/resource.html](http://www.ibm.com/servers/eserver/iseriess/perfmgmt/resource.html))  
This library holds a collection of performance reference materials, white papers, benchmark reports, and trade press articles written by the System i performance experts.
- Performance Management for IBM System i  ([www.ibm.com/servers/eserver/iseriess/perfmgmt/](http://www.ibm.com/servers/eserver/iseriess/perfmgmt/))  
Performance Management provides the capabilities for customers to understand and manage the performance of their computing environments. Read about the latest Performance Management functions and tools on this Web site.


## Saving PDF files

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

1. Right-click the PDF in your browser (right-click the link above).
2. Click the option that saves the PDF locally.

3. Navigate to the directory in which you would like to save the PDF.
4. Click **Save**.

## Downloading Adobe Acrobat Reader

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

---

## 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

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.

---

## Programming Interface Information

This Performance publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM i5/OS.

---

## Trademarks

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

Advanced Function Printing  
Advanced Peer-to-Peer Networking  
AFP  
AIX  
AIX 5L  
AS/400  
Balance  
DB2  
DB2 Universal Database  
Domino  
Electronic Service Agent  
eServer  
i5/OS  
IBM  
iSeries  
Lotus  
NetServer  
OS/2  
OS/400  
PowerPC  
pSeries  
Redbooks  
System i  
System i5  
System p  
System x  
System/36  
Virtualization Engine  
WebSphere  
xSeries

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

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.

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

---

## Terms and conditions

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

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

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

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

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

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

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





Printed in USA