Application System/400™

SC41-8084-00

Programming: Performance Tools/400 Guide

Version 2

System and Application Support

Take Note!

Before using this information and the product it supports, be sure to read the general information under "Notices" on page xiii.

First Edition (May 1991)

This edition applies to the licensed programs IBM Performance Tools/400 (Program 5738-PT1) and IBM Operating System/400 (Program 5738-SS1), Version 2 Release 1 Modification 0 and to all subsequent releases and modifications until otherwise indicated in new editions. Make sure you are using the proper edition for the level of the product.

Order publications through your IBM representative or the IBM branch serving your locality. Publications are not stocked at the address given below.

A form for readers' comments is provided at the back of this publication. If the form has been removed, you may address your comments to:

Attn Department 245
IBM Corporation
3605 Highway 52 N
Rochester, MN 55901-7899

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you or restricting your use of it.

Contents

Figures

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates.

Any reference to an IBM licensed program or other IBM product in this publication is not intended to state or imply that only IBM's program or other product may be used.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Commercial Relations, IBM Corporation, Purchase, NY 10577.

The following terms, denoted by an asterisk (*) in this publication, are trademarks of the IBM Corporation in the United States and/or other countries:

Application System/400	AS/400
AT	IBM
OfficeVision	Operating System/400
OS/400	Personal Computer AT
Personal System/2	RPG/400
SAA	Systems Application Architecture
400	

The following terms, denoted by a double asterisk (**) in this publication, are trademarks of other companies as follows:

Intel	Intel Corporation
RM/COBOL	Ryan McFarland Corporation
RM/COBOL-85	Ryan McFarland Corporation
Xerox	Xerox Corporation

This publication could contain technical inaccuracies or typographical errors.

This guide may refer to products that are announced but are not yet available.

Information that has changed since Version 1 Release 3 Modification 0 is indicated by a vertical bar (|) to the left of the change.

This publication 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.

Programming Interface

This guide is intended to provide information about how to use the Performance Tools/400 licensed program to measure and improve the performance of an AS/400 system. It contains complete descriptions of all the *Performance Tools/400 Guide* reports and displays and provides examples of how to use them. Performance Tools/400 contains no programming interfaces for customers.

About This Guide

This guide explains how to use performance tools to collect data about the performance of a system, job, or program. It also explains how to analyze and print the data to help identify and correct any problems.

This guide is intended for application programmers, programmers, and system operators.

What You Should Know

You should be familiar with the information about performance analysis as described in the *Programming: Work Management Guide*, SC41-8078, before using this guide.

In this guide, personal computer means an IBM Personal Computer such as a 5170 Personal Computer AT* or an 8560 Personal System/2*.

The performance estimates presented are approximations which are believed to be sound. The degree of success that you may achieve in the use of IBM equipment and programs is dependent upon a number of factors, many of which are not under IBM's control. Thus IBM neither warrants nor guarantees that you can or will achieve similar results. It is your responsibility to validate the estimates furnished and to determine their relevance to your operation. Any configuration recommended by the capacity planner should be verified with your Market Support personnel since the capacity planner does not consider all attachable devices.

How This Guide Has Changed

The following chapters are new to this edition, or have significant changes:

| • Chapter 3, "Performance Tuning"
| • Chapter 5, "Advisor"
| • Chapter 6, "Displaying Performance Data"
| • Chapter 8, "Printing Performance Reports"
| • Chapter 9, "Performance Graphics"
| • Chapter 13, "A Problem Analysis Case Study"

| Chapter 13 is a performance analysis scenario providing users of the AS/400 system with an initial approach to determining the source of performance problems using available system tools. It can be used as an introduction to performance analysis and augments the other AS/400 performance-related publications. The scenario follows a data processing manager and an IBM systems engineer through the steps they take to isolate and resolve a performance problem.

Changes since the previous edition of the guide are indicated by a vertical line to the left of the change.

Related Printed Information

The following list presents related printed information that may help you as you use this guide.

The manuals below are listed with their full title and base order number. When these manuals are referred to in this guide, a shortened version of the title is used.

- *AS/400 Performance and Capacity Planning Newsletter, GC21-8175*, provides information on how to use the modeling tool (MDLSYS) for an initial proposal and how to create and use measured profiles. The newsletter also provides information for migrating a System/36 or System/38 workload to the AS/400 system.

  Short Title: *Performance and Capacity Planning Newsletter*

- *Licensed Programs and New Release Installation Guide*, SC41-9878, provides the system operator or system administrator with step-by-step procedures for installing the licensed programs from IBM.

  Short Title: *Licensed Programs and New Release Installation Guide*

- *New User's Guide*, SC41-8211, contains information about how to use function keys.

  Short Title: *New User's Guide*

- *Programming: Control Language Reference*, SC41-0030, provides the application programmer or programmer with a description of the AS/400 control language (CL) and its commands.

  Short Title: *CL Reference*

- *Programming: Work Management Guide*, SC41-8078, provides the programmer with information about how to create an initial work management environment and how to change it.

  Short Title: *Work Management Guide*

- *Service: Diagnostic Aids — Volume 1*, LY44-0597, provides the software service representative with an information source about error logs, dumps, and traces.

  Short Title: *Diagnostic Aids — Volume 1*

Chapter 1. Introduction to Performance Tools

The *Performance Tools/400 Guide* explains how to use performance tools to help you accomplish the following:

- Performance measurement
- Capacity planning
- Performance analysis

Before you use performance tools, however, you must establish performance objectives for:

- Throughput and response time for interactive jobs
- Throughput for batch jobs
- Resource utilizations for the system

Realize that by concentrating on one objective, you could adversely affect another. For example, if your users want fast response time, you need to design and operate your system so that your users receive stable response time over a range of system loads. This choice, however, could cause batch jobs to run slower.

Note: Two types of response time are discussed in this guide. **Internal response time** is the AS/400* system host response time. **External response time** is the end-user response time and includes communications time for both locally and remotely attached display stations.

When response time is mentioned, unless stated otherwise, assume that it is **internal response time**.

For additional information, see "Elements of Response Time" on page A-1.

After you set the performance objectives for interactive and batch jobs, install and use performance tools. The primary functions of performance tools, as described in the following sections, are related. Use each function to make the greatest use of your system's performance.

Performance Measurement

When you first start to use performance tools, use the Start Performance Tools (STRPFRT) command to show the IBM* Performance Tools/400 menu. You should collect performance data (using the Start Performance Monitor (STRPFRMON) command), and produce a system report (using the Print System Report (PRTSYSRPT) command), so you have a measure of how your system currently operates. This process is described in "Using Defaults to Print Performance Reports" on page 8-6.

To determine if your performance objectives are met, collect and analyze performance data regularly. In this way, you accumulate a history of system performance. This history is important in managing system performance, especially if your environment is one of growth and change. By monitoring system performance, you may avoid excessive use of your system's resources.

See Chapter 4, "Collecting System Performance Data," for more information.

Capacity Planning

To estimate your system resource utilization as your workload or environment grows, use the model system (MDLSYS) capabilities. MDLSYS is the capacity planning feature for the AS/400 system. By following MDLSYS recommendations, you may be able to maintain satisfactory system performance and system resource utilizations.

Do capacity planning before you make changes, such as adding new applications or altering the system configuration. See Chapter 10, "Capacity Planning and Performance Prediction," for more information.

Performance Analysis

After you review the performance measurements, you might want to see more detailed performance data. Use the Print System Report (PRTSYSRPT) and Print Component Report (PRTCPTRPT) commands to help you decide if further analysis is necessary. Examples of these reports are shown in Chapter 8, "Printing Performance Reports."

To provide more detail, you can also produce reports that use trace data (specify TRACE(*ALL) on the STRPFRMON command). Use the Print Transaction Report (PRTTNSRPT) command to help you do further analysis of performance problems you may be experiencing.

The advisor, described in Chapter 5, "Advisor," analyzes collected performance data and produces conclusions and recommendations for improving system performance. You can have the advisor put the recommendations into effect, or you can use the conclusions and recommendations to help you decide how to adjust system tuning values.

See Chapter 11, "Programmer Performance Utilities," for an explanation and examples of other utilities you can use to analyze the performance of applications on your system. See Chapter 6, "Displaying Performance Data," for an explanation on interactively displaying performance data.

"Summary of Data Collection and Report Commands" on page 4-10 provides a summary of data collection and reporting commands.

Performance Analysis Overview

Performance analysis is a method for investigating, measuring, and correcting deficiencies so that system performance meets the user's expectations. It does not matter much that the *system* is a computer; it could be an automobile or a washing machine. The problem-solving approach is essentially the same:

1. Understand the symptoms of the problem.
2. Use tools to measure and define the problem.
3. Isolate the cause.
4. Correct the problem.
5. Use tools to verify the correction.

Initially, the analyst knows the user is not satisfied with the way the system is working. It is running "too slow," "too noisy," or "too hot," and so on. The analyst, mechanic, or repair person must first understand what the problem really is. The best way to find out is to observe the problem condition firsthand. Can you confirm the user's complaint? If you cannot, get as much information as possible from users who have experienced the problem. Look and listen for the problem descriptions most in common among the users.

The key to success with any performance issue is to have a clear definition of the users' performance criteria. In other words, given the application mix, what do users want from the system in terms of interactive response time, batch throughput, and processing requirements? For example, a system that supports an interactive order entry application may have a response time criterion to ensure that customers do not perceive abnormal delays. Another criterion may require that end-of-day processing be completed by a specific time. Given these requirements, performance objectives can be established around system resource utilization guidelines. With a clear statement of goals and objectives, performance analysis can proceed on a firm basis.

Once the objectives are understood, 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 application mix? 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 an hypothesis that may explain the cause of the problem. The analyst can use certain Operating System/400* (OS/400*) commands and the IBM Performance Tools/400 to measure the system performance. Reviewing the measured data helps to further define the problem and helps to validate or reject the hypothesis. Once the apparent cause or causes have been isolated, a solution can be proposed. Dealing with one solution at a time, programs can be redesigned and tested. 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, one 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 may become a performance bottleneck.

Improvements to system performance, whether to interactive throughput, interactive response time, batch throughput, or some combination, may take many forms, from simply adjusting activity level or pool size to changing the application code itself. Ultimately, however, any improvement can come only through analysis of the critical resources (CPU, main storage, disk, and remote lines) and contention for system and application objects.

Chapter 2. Starting Performance Tools

This chapter explains how to install and set up your performance tools. Information on how to use the Start Performance Tools (STRPFRT) command is also provided.

Installing Performance Tools

To install the performance tools product, 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 a library by this name is on your system, rename it before you install performance tools, using the Rename Object (RNMOBJ) command. This step will ensure the proper operation of the performance tools.

Use the following command to place the performance tools in library QPFR:

```
RSTLICPGM  LICPGM(5738PT1)  DEV(tape-device-name)  OPTION(*BASE)
```

See the *Licensed Programs and New Release Installation Guide* for additional information on installing licensed programs.

Printer File and Output Queues

The performance tools printer files have a default forms size of 8-1/2 x 11 inches, an overflow line number of 60, and a characters-per-inch setting of 10 or 15 (this setting depends on whether the report is 80 or 132 characters wide). If the printer file characteristics you want are different from the supplied printer file characteristics, use the Change Printer File (CHGPRTF) command to alter them. Use of the generic name, QP*, on this command changes all printer files in library QPFR to the new form size.

The default output queue on the performance job description (QPFRJOBD) is QPFROUTQ. Reports, submitted as batch jobs, use this job description as the default. If you want to use a different output queue from the queue established by performance tools, use the Change Job Description (CHGJOBD) command. Specify the output queue you want to use for the OUTQ parameter on the CHGJOBD command.

Start Performance Tools (STRPFRT) Command

Use the STRPFRT command to start performance tools. After you enter the command, the IBM Performance Tools/400 menu appears. From this display, you can either choose one of the menu selections, or enter a command:

```
PERFORM                  IBM Performance Tools/400     System..:  S0000000

Select one of the following:

      1. Select type of status
      2. Collect performance data
      3. Print performance report
      4. Capacity planning/modeling
      5. Programmer performance utilities
      6. Configure and manage tools
      7. Display performance data
      8. System activity
      9. Performance graphics
     10. Advisor

     70. Related commands


Selection or command
===> ____________________________________________________________________

F3=Exit   F4=Prompt   F9=Retrieve   F12=Cancel   F13=User support
F16=System main menu
(C) COPYRIGHT IBM CORP. 1981, 1991.
```

Press F3 (Exit) or F12 (Cancel) to exit the IBM Performance Tools/400 menu.

Enter commands on the command line. Use F4 (Prompt) and F9 (Retrieve) to prompt for or retrieve commands that you enter on the command line.

To review any messages that are returned to you on the message line, position the cursor on the message line and press the Help key for additional detail. Pressing F10 (Display messages in job log) from this detail display allows you to view all of the messages currently in the job log.

Each time you use STRPFRT, the following occurs:

- The library QPFR is added to the library list (between the system and user positions of the library list).
- The IBM Performance Tools/400 menu appears.

When you finish using performance tools, press F3 (Exit). When you do so, the library QPFR is removed from the job's library list.

Once you use the STRPFRT command to start the performance tools, any further attempt to use the command from within the operating environment for performance tools fails. If you try to start the performance tools program when it is already operating from your job, a message appears that indicates that the operating environment for performance tools is already active. Multiple jobs may use performance tools at the same time but only one performance monitor (data collection job) can be active at any given time.

Note: The library QTEMP is used for various performance tool control objects. If you clear this library during use of performance tools, the tools could fail to operate properly.

Displaying the System or Job Status

If you choose option 1 (Select type of status) on the IBM Performance Tools/400 menu, the Select Type of Status display appears:

```
                          Select Type of Status

Select one of the following:

     1. Work with system status
     2. Work with subsystem
     3. Work with current job
     4. Work with submitted job(s)
     5. Work with specified jobs(s)
     6. Work with active job(s)
     7. Work with disk status
```

On the Select Type of Status display, you can use a set of OS/400 commands to provide you with information about the performance of the system or a particular job.

Each option on the Select Type of Status display has a corresponding command associated with it, as shown in the following list. To use a function, such as working with the system status, either enter option 1 on the command line of the Select Type of Status display *or* enter WRKSYSSTS on any command line.

Select Type of Status Option	Corresponding Command
Work with system status	WRKSYSSTS
Work with subsystem	WRKSBS
Work with current job	WRKJOB
Work with submitted job(s)	WRKSBMJOB
Work with specified job(s)	WRKJOB
Work with active job(s)	WRKACTJOB
Work with disk status	WRKDSKSTS

For more information on these commands, see the *CL Reference* manual. Also see Chapter 3, "Performance Tuning," for more information on the OS/400 commands WRKSYSSTS, WRKACTJOB, and WRKDSKSTS.

Chapter 3. Performance Tuning

This chapter describes how to make performance adjustments to your system. There are two approaches you can take when tuning the system:

- The system can make performance adjustments automatically.
- You can make the performance adjustments manually.

"Automatic System Tuning" discusses the approach that most users should take. For an expanded discussion of what occurs during automatic performance adjustment, see "Performance Adjustments" on page 3-27.

The rest of the chapter discusses manual performance adjustments topics, including:

- Performance components
- Performance commands
  - Work with System Status (WRKSYSSTS)
  - Work with Active Jobs (WRKACTJOB)
  - Work with Disk Status (WRKDSKSTS)
- Basic tuning
- Specialized tuning

To make manual performance adjustments and understand the performance values that can be set, see "Performance Components" on page 3-2, "Performance Commands" on page 3-9, and "Basic Tuning" on page 3-14. If you are experienced in performance tuning and have a large environment (more than 100 active jobs), see "Specialized Tuning" on page 3-24.

There are many options for tuning your system. The concepts presented here give you some general guidelines, not all the answers. Each system environment is unique, requiring you to observe performance and make adjustments that are best for your environment.

Automatic System Tuning

The system can set performance values automatically to provide efficient use of system resources. There are two forms of automatic performance adjustments: initial program load (IPL) adjustments and dynamic adjustments. If performance tuning is new to you, you should set up the system to perform dynamic performance adjustments.

Adjusting Performance at Initial Program Load

If you want the system to do initial tuning for you, set system value QPFRADJ to 1. Each time you do an IPL, the system examines the machine configuration information and makes performance adjustments to achieve efficient use of system resources. No further performance adjustments occur until you do an IPL to the system again, select dynamic performance adjustments, or issue CL commands that change the performance values. For a description of what happens during IPL performance adjustments, see "Initial Program Load Performance Adjustments" on page 3-27.

Adjusting Performance Dynamically

If you want the system to dynamically make performance adjustments, set system value QPFRADJ to 2. When the system is started and periodically thereafter, the system examines the machine configuration, the jobs running on the system, storage requirements, and so on, and makes performance adjustments. Performance values settings change periodically to improve resource use on the system. For more details on the dynamic performance adjustments, see "Dynamic Performance Adjustments" on page 3-28.

Performance Components

Achieving efficient system performance requires a proper balance among system resources. Overusing any resource affects performance. You can make the best use of each resource by using proper settings for the work management parameters.

Job States

Jobs running on the system can be in any of the following states:

- Active
- Wait
- Ineligible

An **active job**, exists in main storage and processes work requested by the application. A job in the **wait state** needs a resource that is not available. An **ineligible job** has work to do, but the system is unable to accept more work at that time.

Wait States

When a job is waiting for a resource, it may wait in main storage or it may be removed from main storage until the resource is available. There are two types of waits: short waits and long waits.

A **short wait** is a period of time no longer than two seconds. It occurs in main storage and causes one of the available activity level positions to be unavailable until the requested resource is available or two seconds have passed. Some typical causes of short waits are:

- Sending a WRITE instruction to a display when *NO is specified on the defer write (DFRWRT) parameter
- Sending break messages to work stations
- Specifying *YES on the restore display (RSTDSP) parameter on display files

When using remote lines, avoid short waits because they cause the wait time in main storage to be much longer for a job than if the job was waiting for resources at a local work station.

A **long wait** is a period of time longer than two seconds. When a long wait occurs, the job's position in main storage becomes available for another job. Other examples of long waits are:

- Record lock conflicts
- Distributed data management data requests

A specialized form of long wait, called **key/think wait**, occurs outside the activity level when a job completes a work assignment and returns to request more work. This is a user-specified time period giving the user time to decide what data should be entered and to type this data. When the job receives a new assignment or runs out of time (times out), it attempts to run again. If no activity level space is available, the job becomes ineligible.

The transition from one job state to another is shown in Figure 3-1.

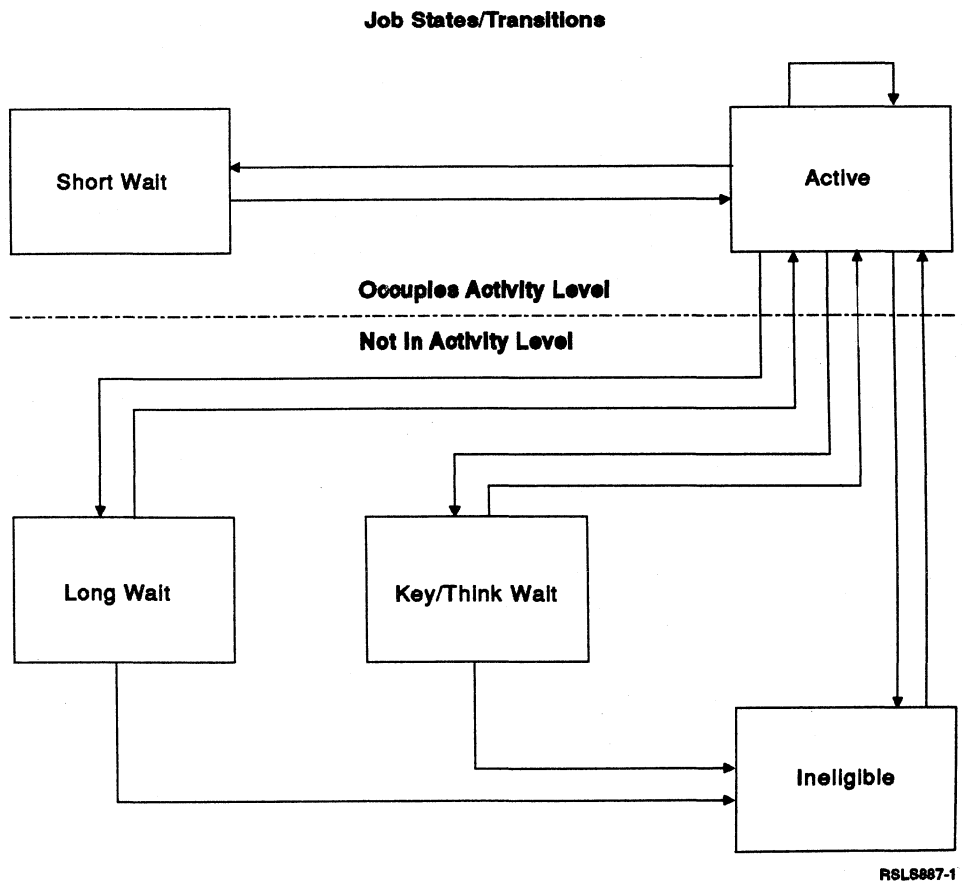

Figure 3-1. Job State Transitions

Activity Levels and Ineligible Queue

For interactive environments, typically there are more jobs running than there is space for them to run. When a job attempts to run, there must be space for the job in main storage. To restrict the number of jobs in main storage at one time, the activity level is specified for each pool in the system. Before a job can become active, an activity level must be available.

If an activity level is available, the job becomes active and begins processing in main storage. If no activity level is available, the job becomes ineligible. When a job becomes ineligible, it is placed in the ineligible queue until an activity level is available.

If the job enters the long wait state by other than a lock conflict, it is placed behind all other jobs of equal priority already on the ineligible queue. This is called a **first-in, first-out priority queue**.

However, if a job becomes ineligible after a long wait caused by a lock conflict, it is placed in front of jobs of equal priority already on the ineligible queue. The most common reasons for this change to normal queue placement are:

- The job entered a long wait as a result of a lock conflict because it was active (referring to objects in main storage) before the conflict occurred. If the wait was short (and many are), you may be able to get the job back into an activity level before all of the objects the job was using are removed from main storage.
- When the job has been granted the lock, it leaves the wait state. If other jobs on the ineligible queue are to use the same object, they must wait until the object is once again available. Therefore, you want jobs holding locks on objects to use them and make them available for other jobs to use. To accomplish this, the job moves ahead of any potential requesters.

By correctly managing the ineligible queue, the system may avoid unnecessary job transitions and disk operations. As a result, throughput and response time are more consistent.

System Objects

Each job running in the system is assigned to a storage pool. When a job is active, it resides in its assigned storage pool. Active jobs refer to many different system objects. When jobs use these objects, they must be in main storage. If they are not in main storage, they must be read into main storage from their locations on disk (auxiliary) storage. Some of the objects used by jobs are:

- Data areas
- File override information
- Device files
- Application codes
- System codes
- Open file information
- Queues
- Logical files
- Subfile work areas
- Physical files
- Program variables

Process Access Groups

A **process access group** (PAG) is a group of job-related objects that can be paged in and out of storage in a single operation when a job (process) enters or leaves a long wait. Although many different object types are found in the PAG, they fall into two main categories: objects that are shared by jobs and objects that are unique to a specific job. When an object is shared, only one copy of the object exists. For example, application code used by 20 jobs resides in main storage in only one place but is used by all the jobs. However, the variables and data in the application do not have the same values for all jobs using the application. These portions of the application and other unique objects are packaged as an object called a process access group (PAG).

When a job waits, the PAG is written to auxiliary storage and the main storage space is available for other jobs. When a job enters an activity level, the PAG is automatically transferred from auxiliary storage to main storage. This activity is similar to swapping. Whenever a job is active, the pages of the PAG that are actually used must be in main storage.

PURGE Parameter

PURGE is a work management tuning parameter. To determine how to get a job's PAG into main storage, the system refers to the value specified on the PURGE parameter in the job class, which is resolved when the job first enters the system. The value for the PURGE parameter is either *YES or *NO.

PURGE (*YES)

The following list describes the characteristics of the PURGE parameter when *YES is specified:

- Adapts to work load and storage size and operates as
  - *YES in limited storage
  - *NO in adequate storage
- Performs better with small main storage

Specifying *YES relieves you of deciding which value to use and is the proper choice for most environments. Figure 3-2 illustrates the manner of transfer when the job class is set to PURGE (*YES).

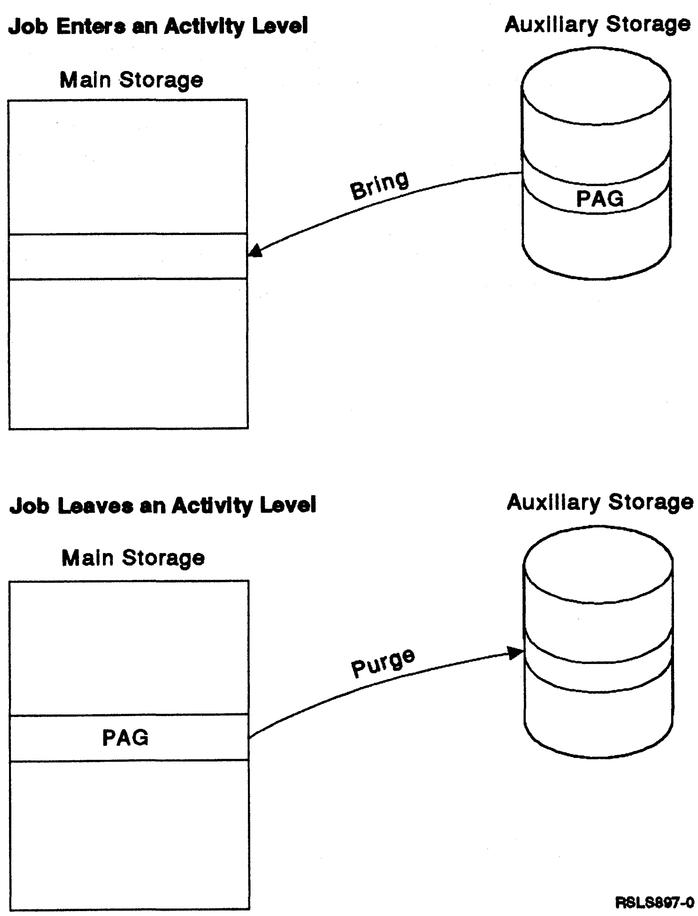

*Figure 3-2. PURGE Parameter when Specifying *YES*

PURGE (*NO)

If your system has enough main storage for the PAGs that are used, the system evaluates the efficiency of automatically reading and writing the PAG. If, when a job enters an activity level, the system determines that the job's PAG is already in main storage, the system does not try to read the entire PAG. Similarly, when the job leaves an activity level, the job's PAG is not automatically written to auxiliary storage. This process is called dynamic PURGE and is specified by PURGE (*NO).

When the job class is set to PURGE (*NO), the system does not read any portion of the PAG until the job requires a part of the PAG that is not in main storage. When this occurs, a small portion of the PAG, starting with the requested data, is read into main storage. When the job leaves an activity level, none of the job's PAG is written to auxiliary storage. The job's PAG remains in main storage until some job currently in an activity level requires more main storage. The system assigns main storage to the job that is currently active. At this point, a small portion of the inactive job's PAG is written to auxiliary storage.

The following list describes the characteristics of the PURGE parameter when *NO is specified:

- May use fewer nondatabase READ operations per transaction
- May cause fewer WRITE operations per transaction
- May transfer fewer pages
- May reduce disk use
- May use less processing unit time per transaction
- Requires more main storage

Figure 3-3 shows the manner of transfer when the job class indicates PURGE (*NO).

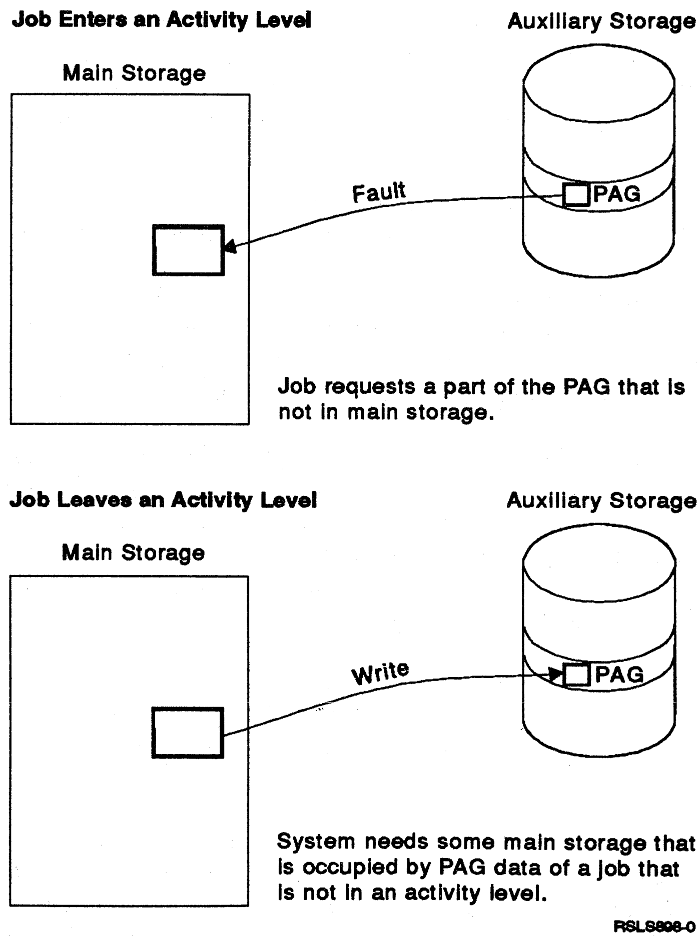

*Figure 3-3. PURGE Parameter when Specifying *NO*

Time Slice

Time slice is another of the work management tuning parameters. The time slice value is specified in the job's class. The value represents the amount of processing unit time a job is allowed to use while processing a transaction. It does not represent the elapsed time of a transaction.

If a job fails to complete a transaction in the specified time slice, one of the following occurs:

- If no jobs of equal or higher priority are on the ineligible queue, the job is given another time slice, remains in main storage, and continues the transaction.
- If jobs of equal or higher priority are on the ineligible queue, the job currently running is removed from main storage and placed on the ineligible queue. A job from the ineligible queue is moved into main storage and processing continues.
- If the job is interactive and a time slice end pool is specified for the job, the job is moved to the time slice end pool.

Long-Running Interactive Transactions

As explained in "Time Slice" on page 3-7, a job is allocated a time slice (an amount of processing time) when it begins to process a transaction. The time slice is used to prevent processing-unit-intensive transactions from using all the resources of the pool in which the transaction is running. When a job fails to complete a transaction within its assigned time slice, one of the following may occur:

- The job moves to another pool.
- The transaction run is temporarily suspended by the system.

If you specify *BASE in the system value QTSEPOOL, the system attempts to reduce the impact of a long-running transaction on other users in the pool. To do this, the system runs the remainder of the transaction in the *BASE pool rather than allowing it to complete in the interactive pool.

At the completion of the long-running transaction, the system returns the job to the interactive pool. This action assumes that less important work is running in *BASE and that the time slice for the interactive jobs is exceeded by only a small percentage of transactions. Usually the transactions that exceed the default time slice value of two seconds for an interactive job are transactions that are characteristic of batch-type activity.

In general, these actions have a positive effect on system performance and *BASE should be specified for this system value. However, the following situations may cause a negative effect on system performance:

- *BASE is very small.

  If the pool is too small, there is not enough storage to contain the work being moved to the pool. When this occurs, the system begins to perform a large number of disk operations. As a result, jobs are unable to perform productive disk requests and performance is poor. If this situation occurs in your environment, add storage to the *BASE pool.

- The activity level in *BASE is not set properly.

  If the activity level is too large, either add storage to the *BASE pool or reduce the activity level. If the activity level is too small, increase the activity level and increase the main storage in *BASE. Jobs running in *BASE should have about 500KB (1KB equals 1024 bytes) per activity level to perform efficiently.

By properly sizing *BASE and choosing an appropriate activity level, moving long-running transactions from the interactive pool to *BASE should provide better system performance. If system performance is not better after several tries, set QTSEPOOL to *NONE and reset the system pool sizes and activity levels to their original values.

Performance Commands

Three system commands are available to help you observe the performance of your system:

- Work with System Status (WRKSYSSTS)
- Work with Disk Status (WRKDSKSTS)
- Work with Active Jobs (WRKACTJOB)

Also, you can use the Performance Tools/400 licensed program to help analyze your performance. 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. This section discusses only the system commands. To observe the system performance, complete the following steps:

1. Use the WRKSYSSTS, WRKDSKSTS, or WRKACTJOB command.
2. Allow the system to collect data for a minimum of 2 to 5 minutes.
3. Press F5 (Refresh) to refresh the display and present the performance data.
4. Tune your system based on the new data.

Then press F10 (Restart) to restart the elapsed time counter.

Working with System Status

The Work with System Status display shows the current status of the system. When tuning the system, make sure that the machine pool is treated separately from the other pools. Type the WRKSYSSTS command on the command line and press the Enter key. The Work with System Status display appears:

```
                        Work with System Status                    SYS01
                                                         02/23/90  09:52:11
% CPU used . . . . . . . :       35.6    Auxiliary storage:
Elapsed time . . . . . . :   00:01:32      System . . . . . . . . :     1803 M
Jobs in system . . . . . :         96      % used . . . . . . . . :   64.3508
% addresses used:                          Total  . . . . . . . . :     1803 M
  Permanent  . . . . . . :      2.805      Current unprotect used :      285 M
  Temporary  . . . . . . :      5.906      Maximum unprotect  . . :      307 M

Type pool size and activity level changes, press Enter.

System    Pool   Reserved   Max    -----DB-----  ---Non-DB---
 Pool   Size (K)  Size (K)  Active  Fault  Pages  Fault  Pages
   1       9639      4517   +++       .0     .0     .3     .4
   2       4000         0     6       .0    5.4    1.8    4.0
   3      35163         0    58       .0     .2    1.1    6.5
   4        350         0     5       .0     .0     .0     .0

                                                                   Bottom
Command
===>
F3=Exit   F4=Prompt          F5=Refresh   F9=Retrieve      F10=Restart
F11=Display transition data  F12=Cancel   F24=More keys
```

Figure 3-4 shows the nondatabase fault rate. Since the machine pool contains objects used system-wide, page faulting in the machine pool affects all jobs on the system. Therefore, it is desirable to maintain a low page fault rate in this pool. The guidelines you should apply are listed in Figure 3-4.

Figure 3-4. Nondatabase Paging Faults

Main Storage Size	Good	Acceptable	Poor
Less than or equal to 12MB	<2	2 – 5	>5
More than 12MB	<1	1 – 3	>3

The only way to affect paging in the machine pool is to adjust the size of the pool (values are represented in the second column of the Work with System Status display).

If you want information about the job state transition data, press F11 to obtain the following display:

```
                         Work with System Status                        SYS01
                                                             02/23/90  09:52:11
% CPU used . . . . . . . :       35.6    Auxiliary storage:
Elapsed time . . . . . . :   00:01:32      System . . . . . . . . :      1803 M
Jobs in system . . . . . :         96      % used . . . . . . . . :   64.3508
% addresses used:                          Total  . . . . . . . . :      1803 M
  Permanent  . . . . . . :      2.805      Current unprotect used :       285 M
  Temporary  . . . . . . :      5.906      Maximum unprotect  . . :       307 M

Type pool size and activity level changes, press Enter.

System     Pool   Reserved   Max    Active->  Wait->  Active->
 Pool    Size (K)  Size (K)  Active    Wait     Inel     Inel
   1        9639      4517    +++       2.6      .0       .0
   2        4000         0      6      77.5      .0       .0
   3       35163         0     58      18.2      .0       .0
   4         350         0      5        .0      .0       .0

                                                                      Bottom
Command
===>
F3=Exit   F4=Prompt      F5=Refresh   F9=Retrieve   F10=Restart
F11=Display pool data    F12=Cancel   F14=Work with Subsystems   F24=More keys
```

As the Work with System Status display shows, all other pools require attention to both database and nondatabase page faulting rates and the job transitions. The guidelines for the faulting rates in these pools are based on the sum of the database and nondatabase faults per second. In addition, you can adjust the processor speed and the PURGE attribute.

Figure 3-5 shows the general guidelines for each pool.

Figure 3-5. Sum of Database and Nondatabase Faulting Levels per Pool

	PURGE(*YES)			PURGE(*NO)		
Model	**Good**	**Acceptable**	**Poor**	**Good**	**Acceptable**	**Poor**
B10 B20 C04 C06 C10 C20 D04 D06 D10	<10	10–15	>15	<15	15–20	>20
B30 B35 B40 B45 C25 D20 D35	<15	15–20	>20	<20	20–25	>25
B50 B60 B70 D25 D45	<20	20–30	>30	<25	25–30	>30
D70 D80	<25	25–35	>35	<30	30–40	>40

In addition, you should observe the total number of faults in all pools. Figure 3-6 shows the general guideline for the total of the database and nondatabase faults per second in all pools.

Figure 3-6. Sum of Database and Nondatabase Paging Faults in All Pools

Model	Good	Acceptable	Poor
B10 B20 C04 C06 C10 C20 D04 D06 D10	<20	20–25	>25
B30 B35 B40 B45 C25 D20 D35	<25	25–35	>35
B50 B60 B70 D25 D45	<30	30–45	>45
D70 D80	<40	40–50	>50

When observing job transitions:

- Wait-to-ineligible transitions need not be 0 all the time. In heavy-use periods, it may be advisable to cause jobs to become ineligible to avoid excessive page fault rates.

  To determine the proper number of wait-to-ineligible transitions, divide the number of wait-to-ineligible transitions by the active-to-wait transitions. Compare your results to the following list.

  Good <.1
  Acceptable .1 – .25
  Poor >.25

  It is usually advisable to complete a transaction during a single time slice. This reduces the number of times the job enters and leaves main storage. Therefore, the active-to-ineligible transitions should be 0. However, long running transactions should not occupy activity levels for the entire transaction.

- You should establish a time slice value that allows 90% of the transactions in your environment to finish in a single time slice. A good starting value is three times the average processing unit per minute.

When using this display, remember that page fault rates are much more important than the job transition values. If you correctly tune the page fault rates, the transition rates usually fall within the guidelines.

You can increase or decrease the pool size or the activity level to get the desired values for pools 2 through 16. The mechanics of these actions are discussed under "Basic Tuning" on page 3-14.

Working with Disk Status

The WRKDSKSTS command shows you your system's disk activity and helps you determine the performance capabilities of your system's disks. The *CL Reference* manual contains a description of the WRKDSKSTS command and formatting information. Type the WRKDSKSTS command on the command line and press the Enter key. The Work with Disk Status display appears:

```
                          Work with Disk Status                    SYS01
                                                          02/23/90  11:04:19
Elapsed time:    00:01:55

             Size    %      I/O   Request   Read  Write    Read  Write    %
Unit  Type    (M)  Used     Rqs  Size (K)    Rqs    Rqs     (K)    (K)  Busy
  1   9332    200  98.9      .4       1.2     .0     .3     1.9    1.0     1
  2   9332    200  60.1      .6       1.0     .4     .2     1.0    1.0     2
  3   9332    200  61.8      .5       1.6     .2     .2     1.4    1.7     2
  4   9332    200  61.5      .6        .9     .3     .3      .9     .8     2
  5   9332    200  61.7      .7       1.9     .3     .4     2.5    1.4     2
  6   9332    200  60.5      .3       1.2     .1     .2     2.0     .7     1
  7   9332    200  63.2      .9       1.2     .3     .5     1.8     .9     3
  8   9332    200  56.3      .8       5.8     .4     .4     3.0    9.1     3
  9   9332    200  56.1      .6       4.2     .2     .3     1.3    6.4     2




                                                                     Bottom
Command
===>
F3=Exit   F5=Refresh   F12=Cancel   F24=More keys
```

Before observing disk status, tune your system according to the paging guidelines described in the topic "Working with System Status" on page 3-9. When viewing the Work with Disk Status display, observe the percent busy data. Each unit (actuator) should be less than 40% busy. If each unit is between 40% and 60% busy, you may experience variable response times. If each unit is more than 60% busy, you do not have enough actuators to provide stable and consistent performance. The **actuator** is the device within an auxiliary storage device that moves the read and write heads. If you have a well-tuned system with actuators that exceed 40% busy, you should increase the number of disk actuators.

It is possible to experience inadequate performance even if only one actuator exceeds the 40% busy guideline. This may be caused by the placement of frequently used data on a single actuator. If this occurs on your system, use the Performance Tools/400 licensed program disk report to determine which data is causing the heavy use. You can save, delete, or restore some objects to improve performance.

An actuator may exceed the 40% guideline for a short period of time. This condition may be caused by a batch job that is accessing data. If the data is not concentrated on a particular actuator, the high level of use should migrate from actuator to actuator while the batch job is running. Also, an actuator in an auxiliary storage pool (ASP) may be used heavily. But typically this is not considered to exceed the guidelines. If you observe this activity, do not change the disk configuration.

Working with Active Jobs

The WRKACTJOB command measures system performance. The following examples show the WRKACTJOB displays.

To view the Work with Active Jobs display, type WRKACTJOB on any command line, and press the Enter key.

```
                              Work with Active Jobs                     SYS01
                                                             02/23/90  11:04:44
CPU %:     .0     Elapsed time:   00:00:00     Active jobs:   67

Type options, press Enter.
  2=Change    3=Hold   4=End    5=Work with   6=Release    7=Display message
  8=Work with spooled files     13=Disconnect...

Opt  Subsystem/Job  User        Type  CPU %  Function        Status
     QBASE          QSYS        SBS      .0                   DEQW
       USER1        USER1       EVK      .0                   ICFW
       USER1        USER1       EVK      .0  *  -PASSTHRU     EVTW
       USER1        USER1       EVK      .0  *  -PASSTHRU     EVTW
       USER1        USER1       EVK      .0  *  -PASSTHRU     EVTW
       USER2        USER2       EVK      .0  *  -PASSTHRU     EVTW
       USER2        USER2       EVK      .0  *  -PASSTHRU     EVTW
       USER3        USER3       EVK      .0  *  -PASSTHRU     EVTW
       DSP020000    PGRENGS     INT      .0  MNU-PROGRAM      DSPW
                                                                        More...
Parameters or command
===>
F3=Exit      F5=Refresh   F10=Restart statistics   F11=Display elapsed data
F12=Cancel   F23=More options   F24=More keys
```

To display information about elapsed data, press F11 to obtain the following display:

```
                              Work with Active Jobs                     SYS01
                                                             02/23/90  11:04:44
CPU %:     .0     Elapsed time:   00:00:00     Active jobs:   67

Type options, press Enter.
  2=Change    3=Hold   4=End     5=Work with    6=Release    7=Display message
  8=Work with spooled files                     13=Disconnect...
                                               --------Elapsed---------
Opt  Subsystem/Job  Type  Pool  Pty      CPU  Int     Rsp  AuxIO  CPU %
     QBASE          SBS     2   0      813.8                  0     .0
       USER1        EVK     2   20       2.1                  0     .0
       USER1        EVK     2   50        .4                  0     .0
       USER1        EVK     2   50        .2                  0     .0
       USER1        EVK     2   50        .2                  0     .0
       USER2        EVK     2   50        .1                  0     .0
       USER3        EVK     2   50        .2                  0     .0
       USER4        EVK     2   50        .2                  0     .0
       DSP020000    INT     3   20      15.7   0      .0      0     .0
                                                                        More...
Parameters or command
===>
F3=Exit      F5=Refresh   F10=Restart statistics   F11=Display status
F12=Cancel   F23=More options   F24=More keys
```

Use both the WRKSYSSTS and the WRKACTJOB commands when attempting to observe your system's performance. With each observation period, you should examine and evaluate the measures of system performance against the goals you have set. Some of the typical measures include:

- Interactive throughput and response time, available from the WRKACTJOB display.
- Batch throughput. Observe the `AuxIO` and `CPU%` values for active batch jobs.
- Spool throughput. Observe the `AuxIO` and `CPU%` values for active writers.

Each time you make tuning adjustments, you should measure and compare all of your key performance measures. Make and evaluate adjustments one at a time.

Basic Tuning

This section describes some of the steps you can take to initially configure the system pool sizes and activity levels to tune your system efficiently. The following topics are discussed:

- Initial machine pool size
- Choosing pool configuration
- Adjustments to pool sizes and activity levels

Initial Machine Pool Size

Maintaining low page fault rates in the machine pool helps the system perform better. When setting the initial machine pool size, use the following:

```
S = M + J + L + F
```

where:

S Initial machine pool size
M Minimum machine pool size (see Figure 3-7)
J Job space (see Figure 3-11 on page 3-18)
L Communications space (see Figure 3-12 on page 3-19)
F Functional space (see Figure 3-13 on page 3-20)

To find the value for the minimum machine pool size, locate your main storage size in Figure 3-7 and use the corresponding value.

Figure 3-7 (Page 1 of 2). Minimum Machine Pool Size

Main Storage Size (MB)	Minimum Machine Pool Size (KB)
4	1175
8	1625
12	2050
16	2400
20	2750
24	3050
28	3350
32	3650
36	3950
40	4250

Figure 3-7 (Page 2 of 2). Minimum Machine Pool Size

Main Storage Size (MB)	Minimum Machine Pool Size (KB)
48	4800
64	6200
72	6900
80	7600
96	9000
112	10300
128	11600
144	12800
160	14000
192	16200
208	17300
224	18400
240	19500
256	20600
272	21700
288	22800
304	23900
320	25000
336	26100
352	27200
384	29400

Figure 3-8 shows the main storage sizes that are available for Models B10 through B70.

Figure 3-8 (Page 1 of 2). Supported Main Storage Sizes for Models B10 through B70

Main Storage Size (MB)	Model								
	B10	B20	B30	B35	B40	B45	B50	B60	B70
4	Y	Y	Y	—	—	—	—	—	—
8	Y	Y	Y	Y	Y	Y	—	—	—
12	Y	Y	Y	Y	Y	Y	—	—	—
16	Y	Y	Y	Y	Y	Y	Y	—	—
20	—	Y	Y	Y	Y	Y	Y	—	—
24	—	Y	Y	Y	Y	Y	Y	—	—
28	—	Y	Y	Y	Y	Y	Y	—	—
32	—	—	—	Y	Y	Y	Y	Y	Y
36	—	—	Y	—	—	—	Y	—	—
40	—	—	—	Y	Y	Y	Y	—	—
48	—	—	—	—	—	—	Y	Y	Y
64	—	—	—	—	—	—	—	Y	Y

Figure 3-8 (Page 2 of 2). Supported Main Storage Sizes for Models B10 through B70									
Main Storage Size (MB)	**Model**								
	B10	**B20**	**B30**	**B35**	**B40**	**B45**	**B50**	**B60**	**B70**
80	—	—	—	—	—	—	—	Y	Y
96	—	—	—	—	—	—	—	Y	Y
128	—	—	—	—	—	—	—	—	Y
160	—	—	—	—	—	—	—	—	Y
192	—	—	—	—	—	—	—	—	Y

Figure 3-9 shows the main storage sizes that are available for Models C04 through C25.

Figure 3-9. Supported Main Storage for Models C04 through C25					
Main Storage Size (MB)	**Model**				
	C04	**C06**	**C10**	**C20**	**C25**
8	Y	Y	Y	Y	Y
12	Y	Y	Y	Y	Y
16	—	Y	Y	Y	Y
20	—	—	Y	Y	Y
24	—	—	—	Y	Y
28	—	—	—	Y	Y
32	—	—	—	Y	Y
36	—	—	—	—	Y
40	—	—	—	—	Y

Figure 3-10 shows the main storage sizes that are available for Models D04 through D80.

Figure 3-10 (Page 1 of 2). Supported Main Storage Sizes for Models D04 through D80											
Main Storage Size (MB)	**Model**										
	D04	**D06**	**D10**	**D20**	**D25**	**D35**	**D45**	**D50**	**D60**	**D70**	**D80**
8	Y	Y	Y	Y	—	Y	—	—	—	—	—
12	Y	Y	Y	Y	—	—	—	—	—	—	—
16	Y	Y	Y	Y	Y	Y	Y	—	—	—	—
20	—	Y	Y	Y	—	—	—	—	—	—	—
24	—	—	Y	Y	Y	Y	Y	—	—	—	—
28	—	—	Y	Y	—	—	—	—	—	—	—
32	—	—	Y	Y	Y	Y	Y	Y	Y	Y	—
36	—	—	—	Y	—	—	—	—	—	—	—
40	—	—	—	Y	Y	Y	Y	—	—	—	—
48	—	—	—	—	Y	Y	Y	Y	Y	Y	—
56	—	—	—	—	Y	Y	Y	—	—	—	—
64	—	—	—	—	Y	—	Y	Y	Y	Y	Y

Figure 3-10 (Page 2 of 2). Supported Main Storage Sizes for Models D04 through D80

Main Storage Size (MB)	Model										
	D04	D06	D10	D20	D25	D35	D45	D50	D60	D70	D80
72	—	—	—	—	—	Y	—	—	—	—	—
80	—	—	—	—	—	—	Y	Y	Y	Y	Y
96	—	—	—	—	—	—	—	Y	Y	Y	Y
112	—	—	—	—	—	—	—	Y	Y	Y	Y
128	—	—	—	—	—	—	—	Y	Y	Y	Y
144	—	—	—	—	—	—	—	—	Y	Y	Y
160	—	—	—	—	—	—	—	—	Y	Y	Y
176	—	—	—	—	—	—	—	—	Y	Y	Y
192	—	—	—	—	—	—	—	—	Y	Y	Y
208	—	—	—	—	—	—	—	—	—	Y	Y
224	—	—	—	—	—	—	—	—	—	Y	Y
240	—	—	—	—	—	—	—	—	—	Y	Y
256	—	—	—	—	—	—	—	—	—	Y	Y
272	—	—	—	—	—	—	—	—	—	—	Y
288	—	—	—	—	—	—	—	—	—	—	Y
304	—	—	—	—	—	—	—	—	—	—	Y
320	—	—	—	—	—	—	—	—	—	—	Y
336	—	—	—	—	—	—	—	—	—	—	Y
352	—	—	—	—	—	—	—	—	—	—	Y
384	—	—	—	—	—	—	—	—	—	—	Y

Job space is set aside in the machine pool for each active job in the system. Use Figure 3-11 on page 3-18 to get a value for this part of the system. However, a more accurate value is derived if you:

- Estimate the maximum number of jobs (including group jobs and PC Support jobs) that are active at the same time.
- Add the value you estimate to the value obtained from Figure 3-7 on page 3-14, because each job requires 1KB in the machine pool.

Figure 3-11. Job Space							
	Models						
Main Storage Size (MB)	**B10 C04 C06**	**C10 D04**	**B20 B30 C20 D06 D10**	**B35 B40 B45 C25 D20 D35**	**B50 D25 D45**	**B60 B70 D50 D60 D70**	**D80**
4	20	—	30	—	—	—	—
8	40	40	60	75	—	—	—
12	50	50	70	80	—	—	—
16	55	60	75	90	100	—	—
20	—	70	80	95	110	—	—
24	—	—	85	100	120	—	—
28	—	—	90	105	130	—	—
32	—	—	95	110	140	160	—
36	—	—	100	115	150	—	—
40	—	—	—	120	160	—	—
48	—	—	—	125	170	240	—
56	—	—	—	130	180	—	—
64	—	—	—	—	185	320	320
72	—	—	—	135	—	—	—
80	—	—	—	—	190	380	400
96	—	—	—	—	—	425	440
112	—	—	—	—	—	445	480
128	—	—	—	—	—	465	520
144	—	—	—	—	—	485	560
160	—	—	—	—	—	500	600
176	—	—	—	—	—	515	640
192	—	—	—	—	—	525	680
208	—	—	—	—	—	550	715
224	—	—	—	—	—	575	750
240	—	—	—	—	—	600	785
256	—	—	—	—	—	625	820
272	—	—	—	—	—	—	855
288	—	—	—	—	—	—	890
304	—	—	—	—	—	—	920
320	—	—	—	—	—	—	950
336	—	—	—	—	—	—	980
352	—	—	—	—	—	—	1010
384	—	—	—	—	—	—	1070

For each line, protocol, line type, and controller, the machine pool must include some additional main storage. Figure 3-12 is given as a simple guide to the additional storage required. These values are based on estimates of the numbers and types of communications lines that may appear for each model and main storage size. However, you can do a better job because you know the communications configuration of the system. To estimate additional main storage:

1. Add 125KB for each line.
2. Add 100KB for each protocol (BSC, SDLC, asynchronous, and so on).
3. Add 25KB for each controller.

Figure 3-12 (Page 1 of 2). Communications Space

	Models					
Main Storage Size (MB)	**B10 C04 C06**	**B20 B30 C10 C20 D04 D06 D10**	**B35 B40 B45 C25 D20 D35**	**B50 D25 D45**	**B60 B70 D50 D60 D70**	**D80**
4	250	250	—	—	—	—
8	350	350	400	—	—	—
12	425	475	525	—	—	—
16	500	550	675	750	—	—
20	—	650	775	825	—	—
24	—	750	850	900	—	—
28	—	900	950	1000	—	—
32	—	1050	1075	1125	1250	—
36	—	1200	1285	1250	—	—
40	—	—	1350	1375	—	—
48	—	—	1450	1500	1750	—
56	—	—	1525	1600	—	—
64	—	—	—	1675	2200	2300
72	—	—	1625	—	—	—
80	—	—	—	1775	2800	3000
96	—	—	—	—	3250	3300
112	—	—	—	—	3425	3600
128	—	—	—	—	3600	3900
144	—	—	—	—	3725	4200
160	—	—	—	—	3850	4500
176	—	—	—	—	3975	4800
192	—	—	—	—	4100	5100
208	—	—	—	—	4200	5400
224	—	—	—	—	4300	5700
240	—	—	—	—	4375	6000
256	—	—	—	—	4450	6300
272	—	—	—	—	—	6550
288	—	—	—	—	—	6800

Figure 3-12 (Page 2 of 2). Communications Space						
	Models					
Main Storage Size (MB)	B10 C04 C06	B20 B30 C10 C20 D04 D06 D10	B35 B40 B45 C25 D20 D35	B50 D25 D45	B60 B70 D50 D60 D70	D80
304	—	—	—	—	—	7050
320	—	—	—	—	—	7300
336	—	—	—	—	—	7500
352	—	—	—	—	—	7700
384	—	—	—	—	—	8100

Finally, certain system functions, when active, require additional space in the machine pool. If you plan to use any of the functions shown in Figure 3-13, you should add the appropriate amount of main storage to the machine pool.

Figure 3-13. Functional Space

Function	Addition to Machine Pool
3270 emulation or remote attachment	50KB
Checksum	5% of main storage size
Save or restore	68KB[1]
Double-byte character set	50KB
X.25	48KB
Token-ring local area network	250KB
OfficeVision/400*	600KB

1 For each save or restore operation occurring at the same time.

Choosing Your Pool Configuration

After you have set the machine pool size, you need to decide what to do with the remaining storage. Your choices include:

- Creating no additional pools
- Creating a separate pool for interactive jobs, for example, the *INTERACT shared pool in the QINTER subsystem
- Creating a separate pool for spooled jobs, for example, the *SPOOL shared pool in the QSPL subsystem

If you choose to create separate pools for QINTER and QSPL, you can add the values from Figure 3-14 and Figure 3-15 to determine the initial values for the size and activity levels for the spooled job. Then use Figure 3-16 on page 3-21 to calculate the *BASE pool sizes and activity levels. Batch jobs (including System/36 environment-started batch jobs) run in the *BASE pool. The remaining storage is the QINTER pool size. To calculate an activity level for the QINTER pool, divide the pool size by the value determined from Figure 3-17 on page 3-22.

Figure 3-14 shows the QSPL pool sizes and activity levels for printers capable of advanced function printing.

Figure 3-14. QSPL Pool Sizes and Activity Levels for Printers Capable of Advanced Function Printing

Number of Writers	Size (KB)	Activity Level
1	500	1
2	1000	2
3	1300	3
4	1500	4
>4	1700	5

Figure 3-15 shows the QSPL pool sizes and activity levels for non-advanced function printers.

Figure 3-15. QSPL Pool Sizes and Activity Levels for Printers Not Capable of Advanced Function Printing

Number of Writers	Size (KB)	Activity Level
1	80	1
2	160	2
3	225	3
4	290	4
>4	350	5

The activity level reflects an estimate of the maximum number of batch jobs (including called programs) that can be active at the same time as the interactive work. Figure 3-16 assumes that interactive work is running at the same time as batch jobs.

*Figure 3-16 (Page 1 of 2). *BASE Pool Sizes and Activity Levels*

Main Storage Size (MB)	Pool Size (KB)/Activity by Model						
	B10 C04 C06	**C10 D04**	**B20 B30 C20 D06 D10**	**B35 B40 B45 C25 D20 D35**	**B50 D25 D45**	**B60 B70 D50 D60 D70**	**D80**
4	500/2	—	500/2	—	—	—	—
8	1000/3	1000/3	1000/3	1000/3	—	—	—
12	1250/3	1250/3	1250/3	1250/3	—	—	—
16	1500/4	1500/4	1500/4	1500/4	1500/4	—	—
20	—	1750/4	1750/4	1750/4	1750/4	—	—
24	—	—	2000/5	2000/5	2000/5	—	—
28	—	—	2350/5	2350/5	2350/5	—	—
32	—	—	2700/5	2700/5	2750/6	2750/6	—
36	—	—	2700/5	3000/5	3000/6	—	—
40	—	—	—	3500/5	3500/6	—	—
48	—	—	—	4000/6	4000/6	4000/6	—
56	—	—	—	4625/6	4625/6	—	—
64	—	—	—	5250/7	5250/7	5250/7	
72	—	—	—	5875/7	—	—	—

Figure 3-16 (Page 2 of 2). *BASE Pool Sizes and Activity Levels							
	Pool Size (KB)/Activity by Model						
Main Storage Size (MB)	**B10 C04 C06**	**C10 D04**	**B20 B30 C20 D06 D10**	**B35 B40 B45 C25 D20 D35**	**B50 D25 D45**	**B60 B70 D50 D60 D70**	**D80**
80	—	—	—	—	6500/7	6500/7	6500/7
96	—	—	—	—	—	8000/7	8000/7
112	—	—	—	—	—	8875/7	8875/7
128	—	—	—	—	—	9750/7	9750/7
144	—	—	—	—	—	10625/7	10625/7
160	—	—	—	—	—	11500/7	11500/7
176	—	—	—	—	—	12750/7	12750/7
192	—	—	—	—	—	14000/7	14000/7
208	—	—	—	—	—	15000/8	15000/8
224	—	—	—	—	—	15900/8	15900/8
240	—	—	—	—	—	16800/8	16800/8
256	—	—	—	—	—	17700/8	17700/8
272	—	—	—	—	—	—	18600/9
288	—	—	—	—	—	—	19500/9
304	—	—	—	—	—	—	20100/9
320	—	—	—	—	—	—	20700/10
336	—	—	—	—	—	—	21300/10
352	—	—	—	—	—	—	21900/10
384	—	—	—	—	—	—	23100/10

Figure 3-17 shows the activity level factor for the QINTER pool:

Figure 3-17. QINTER Activity Level Factor

Main Storage Size (MB)	Activity Level Factor (KB)
≤ 12	450KB
16 — 28	900KB
32 — 48	1600KB
64 — 192	2500KB
208 — 272	3000KB
288 — 384	3500KB

Adjustments to Pool Sizes and Activity Levels

Once you have set initial pool sizes and activity levels, you should begin to observe system performance. When you are observing performance:

- Use both WRKSYSSTS and WRKACTJOB commands.
- Observe at intervals of 5 minutes.
- Observe when the system is handling *normal* workloads (for example, observations during the noon hour may not give meaningful data).

- Apply the page fault rate guidelines shown in Figure 3-5 on page 3-11.

If your observations indicate page fault rates outside the stated guidelines, you should determine the cause of the high page fault rate. Do not make adjustments to tuning if some abnormal condition is causing the high page fault rate. Interactive program compiles, communications error recovery procedures (ERP), open query file (OPNQRYF), application errors, and sign-off activity are examples of irregular activities that may cause a severe hit to performance. Rather than adjusting the system tuning, you should try to remove as many as possible from your operating environment.

If you decide that the high page fault rates in your system occur during normal system operation, adjust the tuning parameters, primarily pool size and activity levels. Tune the machine pool first. You can only reduce machine pool page faulting by increasing the size of the pool. Once you have reached the values shown as *good* in Figure 3-4 on page 3-10, Figure 3-5 on page 3-11, and Figure 3-6 on page 3-11, you can stop.

If your machine pool is experiencing very low page fault rates (<0.4), you should decrease the size of the machine pool. If the page fault rate is this low, you may be affecting work in some other pools.

After tuning the machine pool fault rate, you need to focus on the other pools in the system. Once again, you can adjust size and activity level.

Adjusting Activity Levels

Adjusting the activity level may be the most beneficial way to improve the tuning in pools 2 through 16. The possible problems that can exist in these pools are:

- Low fault rate with an unacceptable value of wait-to-ineligible transitions

  In this situation, the activity level is probably set too low. The paging rate is low because only a few jobs are allowed in the pool at the same time. As a result, the jobs must wait for an activity level before they can run. Increase the activity level by one until the paging rate in the pool and the wait-to-ineligible transitions reach an acceptable level. Be certain to measure and compare each change so you can identify the best activity level.

- A moderate fault rate with an acceptable value of wait-to-ineligible transitions

  This situation indicates that the activity level may be set correctly, but the pool size is too small. To alleviate this condition, configure more storage and, possibly, more activity levels for this pool.

- A high fault rate with either a good or an unacceptable value of wait-to-ineligible transitions

  In this situation, the activity level is probably set too high. In this case, too many jobs are allowed into the pool at the same time. Instead of doing productive work in main storage, the jobs must read information to complete their transactions many times. Demand for space in the pool is so high that jobs are unable to use their pages before another job has overlaid their information. This condition is called **thrashing**. To reduce thrashing, reduce the activity level by one until paging returns to an acceptable rate. Once again, measure and observe performance after each change.

As you adjust the activity level, you may find that you are cycling through the situations described above. Using the data for the WRKSYSSTS and WRKACTJOB commands, select the tuning values that have provided the best performance. If performance is still inadequate, use the Model System (MDLSYS) command of the Performance Tools/400 to evaluate system upgrades to main storage, disk, and the processing unit.

When applying these tuning techniques, change the value of QPFRADJ to 0 and set the subsystem description. Otherwise, the next system IPL changes the settings of the performance values that you have set.

Adjusting Pool Sizes

After your activity level is at its optimum setting, examine the sizes of the pools in your system. If you detect a pool with a low fault rate and a low wait-to-ineligible count, the pool probably has too much storage. Reduce the size of storage-rich pools and add size to the pools with less storage. Reduce size in decrements of 10%. Once again, measure each change. See *Work Management Guide* for information on how to change pool sizes.

Reviewing Performance

Once you have set good tuning values, you should periodically review them to ensure your system continues to do well. 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 may 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.

Specialized Tuning

After you have had some time to observe your system's performance, you may want to consider more specialized tuning for your environment. Some considerations are:

- Using PURGE(*NO) for interactive jobs
- Separating batch work from *BASE
- Using multiple pools for interactive jobs
- Using multiple pools for batch jobs

Specifying PURGE(*NO) for Interactive Jobs

The characteristics of interactive jobs are different when PURGE(*NO) is specified on the job class rather than PURGE(*YES). In environments with sufficient main storage, you may want to specify PURGE(*NO).

Environments that will benefit from running PURGE(*NO) are environments with:

- A lot of main storage
- A lot of transactions

- Transactions that use a lot of processing unit time

You can multiply the number of jobs running in the pool by 250KB to determine the size a pool should be when PURGE(*NO) is specified. If the result of your multiplication is approximately equal to your pool size, set the PURGE attribute to *NO. Use Figure 3-17 on page 3-22 to calculate the initial activity level of the pool.

Separate Batch Work from *BASE

To this point, the assumption has been that batch jobs have shared *BASE storage with system jobs (for example, SCPF, QSYSARB, and subsystem monitors) and system transients (for example, file OPEN and CLOSE). As a result, batch jobs compete for space in the *BASE pool. Batch performance may improve if the jobs are moved to a separate pool. Calculate the size of the pool for each job running at the same time (a suggested starting point is 300KB). Set the activity level to be equal to the number of jobs running at the same time. To make the most efficient use of main storage, you should separate batch jobs only if batch processing is continually active. If batch is not active, the storage allocated to the batch pool is not used. This does not make efficient use of main storage. Figure 3-18 shows the guidelines for the sum of database and nondatabase faults in *BASE after the batch jobs have been removed from *BASE.

*Figure 3-18. Database and Nondatabase Paging Faults in *BASE for All Models with Batch Jobs No Longer in *BASE*

Model Number	Good	Acceptable	Poor
B10, B20, C04, C06, C20, D04, D06, D10	<5	5 – 7	>7
B30, B40, B45, C25, C35, D20, D35	<7	7 – 10	>10
B50, B60, B70, D25, D45, D50, D60	<10	10 – 15	>15
D70, D80	<15	15 – 20	>20

Multiple Pools for Interactive Jobs

There are instances where distinct sets of interactive users should be isolated. Some examples are:

- Critical applications
- Programmers
- Users performing office-type functions exclusively
- Data entry personnel

For programmers and users performing office-type functions exclusively, you are attempting to isolate users who are performing the same functions. Often, the functions performed by these users are different from the functions used by all other users. Also, some of these users may be classified as casual users, and isolating them helps protect their objects. For data entry personnel, you are isolating extremely active users to give the best possible response time for their activity.

If you are considering this approach, you need to determine how many users are to be put in each pool. Then, after setting the necessary routing entries in the subsystem description, you need to calculate a pool size, activity level, and PURGE attribute for each pool. If you are setting a separate pool for casual users:

- Set the PURGE attribute to *YES.
- Calculate the activity level by dividing by four the number of work stations to be routed to the pool.
- Multiply the activity level by the applicable value from Figure 3-17 on page 3-22 to set the pool size.

Each pool should be tuned to the optimal pool size and activity level. Remember, once you have separated work into pools, the pools can no longer share the storage.

Multiple Pools for Batch Jobs

Each batch job may use objects that are different from the objects used by other batch jobs in the same pool. Production batch and program compile do not use the same objects. Long-running jobs may not perform well if sharing a pool with short-running jobs. System/36 environment evokes may disrupt other batch jobs' performance. If you have some of these situations, you may want to set up multiple batch pools. When you separate the pools, set up a pool for each batch job type and use the calculated batch job size for each job in the pool (a suggested starting point is 300KB). Each pool should have an activity level of one and only one active job. Again, if there is no work being performed in these pools, the main storage is not used by the system to support jobs in other pools.

System/36 Environment Tuning

Within the System/36 environment, you should consider several key tuning parameters. In particular, the following may result in performance improvements:

- Set the multiple requester terminal (MRT) delay value to avoid unnecessary start and end activity.
- Minimize the amount of security validation with the MRT security value.
- Use PURGE(*NO) for MRT jobs.
- Limit OCL logging.
- Route evoked jobs to the batch job pools and run them at batch priority.

When setting the MRT delay value, select a value that minimizes the number of times the MRT job starts and ends. The overhead for starting and ending on the AS/400 system is much greater than on the System/36. The MRT delay value specifies the amount of time (in seconds) that the MRT job remains active after the last work station detaches from the job. Delaying the ending of the MRT job may allow another work station to attach to the MRT while it is still active.

The MRT security parameter controls the amount of authority checking that is performed each time a work station attaches to the MRT job. The default value causes authority checking of all objects (files, libraries, and so on) used by the job for each work station that attaches to the MRT. You may set this value so that authority checking is performed only when the MRT is started. Use the default value only in cases where a user has authority to the MRT program but not to all the objects used by the program. These cases are not common.

Since the MRT job has several work stations attached, the PURGE attribute setting may influence the MRT performance. Setting the PURGE attribute to *NO for MRT jobs reduces the number of disk operations performed by the system each time a transaction is performed by any work station that is attached to the MRT.

You should also reduce OCL logging for all System/36 programs. The default value is to log all OCL commands that are run. This creates additional overhead, large job message queues, and large job logs. Set the value to log OCL commands that run by system procedures only, not application procedures. Use application procedure OCL logging only during debugging of the procedure. When an application procedure is running correctly, turn off OCL logging.

Set the priority and pool in which evoked jobs run to reduce interference with the interactive jobs. The nature of the evoked job is often that of a batch job. Therefore, any subsystem supporting evoke operations should have a routing entry that places the evoked jobs in the *BASE pool at a running priority that is the same as batch.

By correctly setting these System/36 running parameters, you may improve performance. Each user's requirements need to be analyzed to find the best setting for each of the parameters described.

Performance Adjustments

The system can adjust performance automatically either at initial program load or dynamically.

Initial Program Load Performance Adjustments

If you would like the system to do initial tuning for you, but you do not want the system to perform dynamic tuning, set system value QPFRADJ to 1. When you IPL, the system examines the machine configuration and the controlling subsystem value. If QPFRADJ is set to 1, the system uses the configuration information to set the initial pool sizes and activity levels. If the controlling subsystem is QBASE or QCTL, the system sets up separate pools for spool and interactive jobs.

The IPL performance adjustments result in changes to the following values:

- Machine pool size (QMCHPOOL system value)
- Base pool activity level (QBASACTLVL system value) if the controlling subsystem is QSYS/QBASE, QSYS/QCTL, QGPL/QBASE, or QGPL/QCTL
- Pool number 2 in subsystem QGPL/QSPL to use shared pool *SPOOL
- Pool size and activity level for shared pool *SPOOL
- Pool number 2 in subsystems QSYS/QBASE and QGPL/QBASE to using shared pool *INTERACT if controlling subsystem is QSYS/QBASE, QSYS/QCTL, QGPL/QBASE, or QGPL/QCTL
- Pool number 2 in subsystems QSYS/QINTER and QGPL/QINTER to using shared pool *INTERACT if controlling subsystem is QSYS/QBASE, QSYS/QCTL, QGPL/QBASE, or QGPL/QCTL
- Pool size and activity level for shared pool *INTERACT

These performance values are set to the values defined in "Basic Tuning" on page 3-14.

If you make any adjustments to the pool size or activity level values, you should consider setting QPFRADJ to 0. Otherwise, the values are reset at the next IPL. If you are satisfied with the QPFRADJ values but change your environment to run night batch jobs, you may wish to leave QPFRADJ set to 1.

Dynamic Performance Adjustments

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 underused storage pools to pools that would benefit from more storage. This tuning also sets activity levels to balance the number of jobs in the pool with the storage allocated for the pool. The tuner uses the guidelines in "Working with System Status" on page 3-9 to adjust the system.

Dynamic tuning changes the following performance values:

- *MACHINE pool size (QMCHPOOL system value)
- *BASE pool activity level (QBASACTLVL system value)
- Pool size and activity level for shared pool *INTERACT
- Pool size and activity level for shared pool *SPOOL
- Pool sizes and activity levels for shared pools *SHRPOOL1-10

If you make any adjustments to the pool size or activity level values, you should consider setting QPFRADJ to 0 or 1. Otherwise, the values are changed by the IPL or periodic performance adjustments.

Minimum Pool Size

When the dynamic tuning support is making changes to storage pool sizes, the dynamic tuning support ensures that the pools are not reduced beyond a minimum size. Figure 3-19 shows the minimum pool sizes for pools that are having some page faults occur.

Figure 3-19. Minimum Pool Size

Share Pool	Minimum Pool Size
*MACHINE	Minimum size determined by vertical licensed internal code (VLIC)
*BASE	Minimum size is value of QBASPOOL system value
*INTERACT	600KB
*SPOOL	80KB
*SHRPOOL 1 – 10	300KB

You can use the Work with System Value (WRKSYSVAL) command to change the minimum size of the *BASE pool. The minimum size used by the dynamic tuning support for pools other than the *BASE pool cannot be changed.

If no page faults occur in a pool, the dynamic tuning support may reduce the size of the *BASE, *INTERACT, *SPOOL, or *SHRPOOL1 through 10 pools to 48KB. This allows other pools to use the storage that is not being used.

Journaling Dynamic Performance Adjustments

This section describes how you can set up a journal to record the changes made to the system by the dynamic tuning support. You can track, or journal, the system so you can see which changes have been made to the system. The changes are recorded in journal QSYS/QPFRADJ. You should know how to perform journal management operations, such as saving a journal receiver, changing journal receivers, and deleting old journal receivers. For more information about journal management, see the *Backup and Recovery Guide*.

To set up the performance adjustment journal, do the following:

1. Create a journal receiver of your choice by using the Create Journal Receiver (CRTJRNRCV) command:

   ```
   CRTJRNRCV  JRNRCV(QSYS/QPFRADJ1)
   ```

   Name the journal receiver QPFRADJ1, or select a similar name you can use to create a naming convention such as QPFRADJ2 or QPFRADJ3, for future journal receivers.

   After you create the first receiver, you can use the CHGJRN JRNRCV(*GEN) command to create additional receivers and attach them to the QSYS/QPFRADJ journal automatically with the correct naming convention.

2. Create the journal QSYS/QPFRADJ by using the Create Journal (CRTJRN) command:

   ```
   CRTJRN  JRN(QSYS/QPFRADJ)  JRNRCV(QSYS/QPFRADJ1)
   ```

   You must use the name QSYS/QPFRADJ, and you must have authority to add objects to QSYS. Specify the name of the journal receiver you created in the previous step and any other options on the command.

Note: If dynamic tuning is active (system value QPFRADJ is 2), there is a slight delay in logging entries to the journal. If you need to start logging immediately, change the system value QPFRADJ to 1, and then change it back to 2. This will stop and restart the dynamic tuning support.

After you create the performance journal, use the following steps to copy the performance changes to a file:

1. Create a file of your choice by using the Create Duplicate Object (CRTDUPOBJ) command:

   ```
   CRTDUPOBJ  OBJ(QAWCTPJE)  FROMLIB(QSYS)  OBJTYPE(*FILE)
     TOLIB(MYLIB)  NEWOBJ(MYFILE)
   ```

   The newly created file is formatted to match the format of the journal entries.

2. Copy the journal entries to the file by using the Display Journal (DSPJRN) command:

   ```
   DSPJRN  JRN(QSYS/QPFRADJ)  ENTTYP(TP)  OUTPUT(*OUTFILE)
     OUTFILE(MYLIB/MYFILE)
   ```

Figure 3-20 lists the various fields (found in file QSYS/QAWCTPJE) in the performance tuning (TP) journal entry:

Figure 3-20. Journal Entry Formats

Field Name	Description	Field Attributes
TPPNAM	Name of shared pool	Character(10)
TPFLG1	Pool changed flag	Character(1)
TPCSIZ	Current pool size	Packed decimal(8,0)
TPCRES	Pool reserve size	Packed decimal(8,0)
TPCACT	Current activity level	Packed decimal(6,0)
TPDFLT	Database page faults per second	Packed decimal(6,2)
TPNFLT	Nondatabase page faults per second	Packed decimal(6,2)
TPWI	Job transitions from wait to ineligible	Packed decimal(6,0)
TPAW	Job transitions from active to wait	Packed decimal(6,0)
TPCJOB	Current number of jobs running in pool	Packed decimal(6,0)
TPAJOB	Average number of jobs running in pool	Packed decimal(6,0)
TPNSIZ	New pool size	Packed decimal(8,0)
TPNACT	New activity level	Packed decimal(6,0)

Chapter 4. Collecting System Performance Data

This chapter describes how to collect data using the Start Performance Monitor (STRPFRMON) command. Other ways of collecting data using performance tools are described in Chapter 11, "Programmer Performance Utilities," and Chapter 7, "System Activity." The figures shown in the sections following "Summary of Data Collection and Report Commands" on page 4-10 show the performance tools data collection commands, and describe when you use each in analyzing the performance of your system.

The STRPFRMON command is important in the overall analysis of your system. Use it to collect data about resources that influence the performance of your system (processing unit, main storage, auxiliary storage, and communications).

The STRPFRMON command is provided with the OS/400 program. The *Work Management Guide* contains additional information on collecting performance data and the database files that result and database file descriptions.

The performance monitor provides for the collection of sample data and trace data.

Sample data is data collected at regular intervals over a specified period of time. You control the time between samples by using the interval parameter on the STRPFRMON command. Sample data is collected and stored in the following database files:

File	Type of Data
QAPMCONF	System configuration data
QAPMSYS	System data
QAPMJOBS	Job data
QAPMDISK	Direct access storage device (DASD) data
QAPMPOOL	Main storage data
QAPMHDLC	High-level data link control (HDLC) data
QAPMASYN	Asynchronous (ASYNC) data
QAPMBSC	Binary synchronous communications (BSC) data
QAPMX25	X.25 data
QAPMECL	Establishment communications link (ECL) or token-ring LAN data
QAPMETH	Ethernet statistics data
QAPMBUS	Bus counter data
QAPMCIOP	Communications controller data
QAPMMIOP	Multifunction controller data
QAPMDIOP	Storage device controller data
QAPMLIOP	Local work station controller (WSC) data
QAPMRESP	Local work station response time data

Use the trace parameter on the STRPFRMON command to control the collection of trace data. **Trace data** is data collected continuously over a period of time you specify. Trace data is stored in the QAPMDMPT database file.

Whenever you use the STRPFRMON command, you collect sample data, but you can also collect trace data. You generally choose to collect trace data to gain additional detailed information about specific jobs and transactions. By collecting trace data, you can often gain insight into other problems involving job contention, program resource use, transaction delays, and so on.

Use the Print System Report (PRTSYSRPT) and Print Component Report (PRTCPTRPT) commands to print the sample data you collect. To review examples of these reports, see "System Report" on page 8-10 and "Component Report" on page 8-28.

Use the Print Transaction Report (PRTTNSRPT), the Print Lock Report (PRTLCKRPT), and the Print Trace Report (PRTTRCRPT) commands to see the data collected through trace. Refer to "Transaction Report" on page 8-40 and to "Print Lock Report (PRTLCKRPT) Command" on page 8-114 to review the information provided from trace data collection.

Some of the commands described in Chapter 11, "Programmer Performance Utilities," make use of trace data collected using the STRPFRMON command. See "Summary of Data Collection and Report Commands" on page 4-10 for more information on the commands that use the trace data.

To free up disk space used to save performance data you no longer need, you can use the Delete Performance Data (DLTPFRDTA) command.

The remaining sections in this chapter describe when and how you collect data using the STRPFRMON command.

When to Collect Performance Data

You should collect performance data regularly to establish a record of system performance. If you know the performance characteristics of your system, you can judge the effect of a change in system workload before you make the change.

Use the STRPFRMON command to collect sample data to establish this performance history. Use the PRTSYSRPT command to produce a report from this data. The report provides an overview of the performance variables that should be tracked in the system's history. To see an example of the System Report, see "System Report" on page 8-10.

Collect sample data regularly, so you can make valid comparisons of information. Collect the data during similar system workloads.

Also collect sample data before and after major changes, such as installing a new application. This data provides valuable information on changes to system resource utilizations as a result of the workload changes. When you review this data, limit your observations to the differences in the active workloads on the system. Look at the utilizations for the processing unit, disk, main storage, and communications. If any of these resources is consistently overcommitted, you should determine the reason. See Figure 10-3 on page 10-6 for more information. The threshold values provided in this table represent overcommitment of the given resource.

To determine the reason for consistent overcommitment of a given resource, collect trace data. The performance monitor trace data provides more detailed information on the jobs, programs, and individual transactions that run on the system. Depending on the resource, you might use some of the additional analysis tools discussed in Chapter 11, "Programmer Performance Utilities."

Use the PRTTNSRPT command, as described in "Transaction Report" on page 8-40, to specify the trace data you want to see. You can determine the elements of transaction response time from these reports. In addition, you can identify the level of resource use for the measured transactions.

The length of time you collect data depends on whether you collect only sample data, or a combination of sample and trace data. The length of time you collect data can generally be longer when you collect only sample data. When you collect sample data, you want it to reflect the changes that occur when the various interactive and batch workloads go from light to moderate to heavy activity, and back to light activity again.

When you collect trace data, you may want to reduce the length of time for the measurement, because the area where the trace data is stored is limited in size, and can contain between 40 000 and 60 000 transactions. The trace collection ends automatically when this area is filled. Thus, the length of time you collect data depends on the activity on your system and the volume of trace records that result. However, because you collect trace data to observe the high utilization of a given resource, you should try to collect data for workload periods that correspond to those in which you have seen excessive utilization occur.

When to End Data Collection

When the performance monitor (STRPFRMON) collection ends, the measurement data can automatically be transferred from the collection area into a database file. If you are collecting only sample data, no additional system overhead is associated with stopping the monitor. You can stop the monitor at any time. If you are collecting trace data, there is additional system overhead and you should *try to schedule the end of the measurement to coincide with a lighter workload period*, or use the Dump Trace (DMPTRC) command to transfer the data at a later time. For example, schedule the performance monitor collection to end at noon if there is a heavy workload measurement during the morning.

Collecting Sample or Trace Data

To collect sample or trace data, follow these steps:

1. Enter the Start Performance Tools (STRPFRT) command on any command line to show the IBM Performance Tools/400 menu.
2. Choose option 2 (Collect performance data) on the IBM Performance Tools/400 menu.

```
PERFORM              IBM Performance Tools/400        System..:

Select one of the following:

       1. Select type of status
       2. Collect performance data
       3. Print performance reports
       4. Capacity planning/modeling
       5. Programmer performance utilities
       6. Configure and manage tools
       7. Display performance data
       8. System activity
       9. Performance graphics
      10. Advisor
```

3. Press the Enter key. The Collect Performance Data display appears.

```
                           Collect Performance Data

The performance monitor status:
     Status . . . . . . . . . : Not running




Select one of the following:

     1. Start collecting data
     2. Stop collecting data
     3. Work with performance collection
```

Note: Only one performance monitor function can be active in the system. The current status of the performance monitor is shown.

4. Choose option 1 (Start collecting data), and press the Enter key. The Start Collecting Data display appears.

```
                           Start Collecting Data

Select one of the following:

     1. Collect data with defaults
     2. Collect data with menus
     3. Collect data with command
```

On this display, there are three ways to start the performance monitor to collect data. Any option you choose results in the collection of performance data using STRPFRMON. Options 1 and 2 are designed for a new user of performance tools (they provide more guidance in starting the performance monitor). These options are discussed in the following sections. If you choose option 3 (Collect data with command), a prompt for the STRPFRMON command appears as though you entered the command and pressed F4 (Prompt).

Using System Defaults to Collect Data

If you choose option 1 (Collect data with defaults) on the Start Collecting Data display, the Collect Data with Defaults display appears.

```
                         Collect Data with Defaults

 Type choices, press Enter.

   Member . . . . . . . .   *GEN_____     Name
     Library  . . . . . .   QPFRDATA__    Name
   Text . . . . . . . . .   ________________________________

   Time duration
     Hours  . . . . . . .   2_______      0-999
     Minutes  . . . . . .   0_______      0-60

 F3=Exit   F12=Cancel
```

1. Type the name of the member and library where you want to store the performance data. *GEN creates a member name based on the date and time. The default for *Library* is QPFRDATA.
2. Type a description for the sample data in the *Text* field.
3. Enter the length of time you want to collect performance data in the *Time duration* field.
4. Press the Enter key, and the data collection process begins. The defaults for the other STRPFRMON command parameters are used.

   See the *CL Reference* manual for more information about the STRPFRMON command parameters.

Using Menus to Collect Data

You can change the defaults for the STRPFRMON command by entering the command directly, or by selecting option 2 (Collect data with menus) on the Start Collecting Data display.

If you choose option 2, the Collect Data with Menus display appears.

```
                         Collect Data with Menus

 Type choices, press Enter to continue.

   Member . . . . . . . .   *GEN_____     Name
     Library  . . . . . .   QPFRDATA__    Name
   Text . . . . . . . . .   ________________________________
```

1. Type the member and library name where you want to store the performance data. The default for *Library* is QPFRDATA.
2. Enter a description of the performance data in the *Text* field, if appropriate.
3. Press the Enter key. The Collect Additional Data display appears.

```
                         Collect Additional Data

 Type options, press Enter.
   1=Collect
  Option    Type of data
    _       Trace Data
    _       Communications Data
```

4. To indicate that you want to collect trace data, type a 1 in the *Option* column.
5. To indicate that you want to collect performance data for communications lines and objects, type a 1 in the *Option* column. Press the Enter key.

   The Set Data Collection Time display appears.

```
                         Set Data Collection Time

 Type choices, press Enter to continue.

   Number of minutes . . . . . . . . . . . . . .  15     5-60 (increments of 5)
     between collections

   End time option . . . . . . . . . . . . . . .  1      1=Elapsed time
                                                         2=Time of day
                                                         3=No maximum
```

6. Type how often data should be collected in the *Number of minutes between collections* field. Fifteen minutes is usually an adequate sample interval for problem analysis. This is the sample rate for sample data collection.
7. Type when you want data collection to stop in the *End time option* field. If you choose option 1 (Elapsed time), go to step 9. If you choose option 2 (Time of day) for this prompt, the Set End Time display appears.

```
                              Set End Time

 Type choices, press Enter to continue.

   Number of days from today . . . . . . . . .   0         0-9

   Time of day . . . . . . . . . . . . . . . .   00:00:00  HH:MM:SS
```

8. Indicate the number of days you want to collect data in the *Number of days from today* field. Also indicate when you want data collection to stop in the *Time of day* field.
9. If you choose option 1 (Elapsed time) for the *End time option* prompt on the Collect Additional Data display, the Set Length of Time to Collect Data display appears.

```
                       Set Length of Time to Collect Data

Type choices, press Enter to continue.

  Length of time to collect data
    Hours . . . . . . .  2       0-999
    Minutes . . . . . .  0       0-60
```

10. Type the length of time you want to collect data in the *Length of time to collect data* fields.

After you press the Enter key, the data collection process begins.

For information on how to print various sample and trace data reports, refer to Chapter 8, "Printing Performance Reports."

Collecting Performance Data Automatically

Note: If you are using automatic data collection for the first time, refer to "Performance Collection Setup" on page 4-10 before you begin.

You can choose to have your system automatically collect performance data on a weekly schedule. Automatic performance collection allows you to select specific days of the week on which automatic data collection is to occur using the OS/400 performance monitor.

On the IBM Performance Tools/400 menu, choose option 2 (Collect performance data) and press the Enter key. The Collect Performance Data display appears.

```
                         Collect Performance Data
                                                        01/26/91  14:16:49
Performance monitor status:
  Status . . . . . . . . . . . . :  Not running






Select one of the following:

     1. Start collecting data
     2. Stop collecting data
     3. Work with performance collection


Selection or command
===> __________________________________________________________________
____________________________________________________________________

F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F12=Cancel
```

Choose option 3 (Work with performance collection) and press the Enter key. The Work with Performance Collection display appears. (You can also use the WRKPFRCOL command to access the Work with Performance Collection display.)

```
                      Work with Performance Collection

 Type options, press Enter.
   1=Add   2=Change   3=Hold   4=Remove   5=Display   6=Release

        Performance
 Opt    Collection      Status     Description
  _     __________
  _     A                RLS
  _     F0900            RLS
  _     S0900            RLS















                                                                      Bottom
 F3=Exit   F5=Refresh   F12=Cancel
```

A list of the existing performance collections is displayed. The *Performance Collection* column specifies the details about when to collect performance data automatically.

To add a new collection, type a 1 (Add) in the *Opt* column and press the Enter key. The Add Performance Collection (ADDPFRCOL) display appears. (You can also use the ADDPFRCOL command to access the Add Performance Collection (ADDPFRCOL) display.)

```
                   Add Performance Collection (ADDPFRCOL)

 Type choices, press Enter.

 Performance Collection . . . . .   __________      Name
 Collection days  . . . . . . . .   _____           *ALL, *MON, *TUE, *WED...
               + for more values    _____
 Collection start time  . . . . .   _______         HHMM
 Text 'description' . . . . . . .   *BLANK

 Library  . . . . . . . . . . . .   QPFRDATA__      Name
 Time interval (in minutes) . . .   15_____         5, 10, 15, 20, 25, 30, 35...
 Stops data collection  . . . . .   *ELAPSED        *ELAPSED, *TIME
 Days from start day  . . . . . .   0______         0-7
 Hour . . . . . . . . . . . . . .   2______         0-168
 Minutes  . . . . . . . . . . . .   0______         0-59
 Data type  . . . . . . . . . . .   *ALL            *ALL, *SYS
 Trace type . . . . . . . . . . .   *NONE           *NONE, *ALL
 Dump the trace . . . . . . . . .   *YES            *YES, *NO
 Job trace interval . . . . . . .   .5__            .5 - 9.9 seconds
                                                                      More...
 F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
 F24=More keys
```

Page down to the next page of the Add Performance Collection (ADDPFRCOL) display.

```
                    Add Performance Collection (ADDPFRCOL)

Type choices, press Enter.

Job types . . . . . . . . . . .   *DFT_          *NONE, *DFT, *ASJ, *BCH...
             + for more values    ____
Local response time:
  Boundary 1 . . . . . . . . . .  *SYS_          Number, *SYS
  Boundary 2 . . . . . . . . . .  _____          Number
  Boundary 3 . . . . . . . . . .  _____          Number
  Boundary 4 . . . . . . . . . .  _____          Number
Job queue . . . . . . . . . . .   QCTL______     Name
  Library . . . . . . . . . . .     *LIBL_______ Name, *LIBL, *CURLIB
Message queue . . . . . . . . .   *NONE_____     Name, *NONE
  Library . . . . . . . . . . .     ___________  Name, *LIBL, *CURLIB




                                                                  Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

These displays are similar to those used with the Start Performance Monitor (STRPFRMON) command. However, on the Add Performance Collection (ADDPFRCOL) display, you can establish a regular schedule for collecting performance data automatically on any of the following days:

***MON** Monday
***TUES** Tuesday
***WED** Wednesday
***THU** Thursday
***FRI** Friday
***SAT** Saturday
***SUN** Sunday
***ALL** All days of the week

You must name your collection and specify the day and time or combination of days and times, to collect performance data. The Performance Monitor needs a minimum of 15 minutes between the stop time for a collection and the start time for any following collection. You can change any of the other parameters.

Option 2 (Change) on the Work with Performance Collection display allows you to change any of the existing collections. You can also use the CHGPFRCOL command to access the Change Performance Collection display. The Change Performance Collection display looks the same as the Add Performance Collection display. You can change any of the fields. Press the Enter key to save your changes.

Option 3 (Hold) on the Work with Performance Collection display allows you to stop a collection from running until it is released with option 6 (Release). These functions allow you to keep your collection specifications, but turn the automatic collections off and on by holding and releasing them.

Option 4 (Remove) on the Work with Performance Collection display removes a collection specification from the system. You cannot use or work with a collection once it has been removed. There is not a confirmation display, so be sure you want to remove the collection before you do so.

Option 5 (Display) on the Work with Performance Collection display shows you a display similar to that of the Add Performance Collection display. Values for the parameters are listed on the display.

Performance Collection Setup

Automatic performance collection requires a batch job (QPFRCOL) that queries the schedule created by the Add Performance Collection function and submits the Start Performance Monitor (STRPFRMON) command at the appropriate times. This job exists as an autostart job entry in the IBM-supplied subsystems QBASE and QCTL. If you are using Version 1 Release 3 or later of one of these controlling subsystems, the batch job is started for you after each IPL. If you are not using the latest release of QBASE or QCTL, you can add an autostart job entry for the performance collection batch job to one of your subsystems. The *Work Management Guide* contains more information about autostart job entries.

When the job QPFRCOL is started and there are no performance collections defined that are in a released status, the job ends. Therefore, anytime you add a collection and there are no existing performance collections in a released status, submit the performance collection job as follows:

```
SBMJOB  JOBD(QGPL/QPFRCOL)  USER(*JOBD)  RQSDTA(*JOBD)  RTGDA(*JOBD)
```

You can type this command on any command line. This batch job runs until all the performance collections are removed or held. You need to resubmit the job if you add performance collections later.

Summary of Data Collection and Report Commands

Figure 4-1 through Figure 4-5 in the following sections present the commands for various levels of data collection. These figures also show the related report commands, show the type of data collected, provide a summary of the information contained in the reports, and describe when you might use these commands.

Refer to the figures indicated for information on the following data collection levels:

- System (Figure 4-1 on page 4-11)
- Job (Figure 4-2 on page 4-13)
- Program (Figure 4-3 on page 4-13)
- File use and structure (Figure 4-4 on page 4-13)
- Application (Figure 4-5 on page 4-14)

If you use the performance tools menus and displays to collect data and produce reports, these figures may help you understand, at a glance, the capabilities of performance tools. If you bypass the menus and displays by entering commands on the available command entry lines, these figures may serve as a reference for the available commands.

System-Level Analysis

System-level data collection and analysis provides you with a comprehensive view of how the system operates. This information ranges from a system operational overview to an analysis of individual transactions. System-level data collection and analysis also provides you with system modeling functions for capacity planning and performance prediction.

Use system-level data to identify what additional collection and analysis should be done.

A summary of system data collection and report commands is shown in Figure 4-1.

Figure 4-1 (Page 1 of 2). System Data Collection and Report Commands

System Data Collection Command	Level of Data	Type of Data	Report Command	Information Shown on the Reports	When to Use the Command
STRPFRMON	Job Disk System	Sample data	ANZPFRDTA	Contention analysis and recommendations	Processing trends System model Workload projection Hardware growth Processing unit Main storage Disk
STRPFRMON	Job Disk System	Sample data	PRTSYSRPT PRTCPTRPT	Workload Utilization Processing unit Disk Main storage Communications Model parameters	Processing trends System model Workload projection Hardware growth Processing unit Main storage Disk
STRPFRMON	System Job Program	Trace data	PRTTNSRPT	Workload Utilization Processing unit Exceptional waits Transaction detail Top ten reports Object contention Batch thread System model parameters Transaction summary and detail	Workload projection Hardware growth Pool configuration Overcommitment Application design File contention Transaction Significance Classification Program use System model Processing trends
STRPFRMON	System Job Program Files Disk	Trace data	PRTTRCRPT	Resources used Exceptions State transitions	Progression of batch jobs traced through time
STRPFRMON	Job Program Files Disk	Sample data	MDLSYS	System performance projections Capacity planning Configuration planning	Before installing When growth is anticipated, either in hardware or workload When a new application is to be installed Performance analysis

Figure 4-1 (Page 2 of 2). System Data Collection and Report Commands

System Data Collection Command	Level of Data	Type of Data	Report Command	Information Shown on the Reports	When to Use the Command
STRPFRMON	Job Program Files Disk	Trace data	PRTLCKRPT	File, record, or object contention by Object name Holding or requesting job Time	To reduce or remove object contention Problem analysis
STRDSKCOL	File Disk Object	Trace data	PRTDSKRPT	I/O requests by Object and Disk unit	When disk has excessively high activity level
STRPFRMON	Job Program Files Disk	Sample data	PRTJOBRPT	Utilization Processing unit Disk Communications Workload	Problem analysis
STRPFRMON	Job Program Files Disk	Sample data	PRTPOLRPT	Utilization Main Storage Workload Subsystem	Problem analysis
STRPFRMON	Job Program Files Disk	Sample data	PRTRSCRPT	Utilization I/O Processing unit Disk	Problem analysis

For more information on the report commands shown in this figure, see Chapter 8, "Printing Performance Reports." The MDLSYS command is described in Chapter 10. The PRTLCKRPT command is described in "Print Lock Report (PRTLCKRPT) Command" on page 8-114. The STRDSKCOL and PRTDSKRPT commands are described in "Analyzing Disk Activity" on page 11-34.

Job Trace Analysis

Job trace analysis enhances the operating system's standard trace job reports and provides a summary of job operation and transaction processing. The primary use for job trace analysis is to determine application flow. You can determine what parts of a job use the most resources, and measure the effect of program changes relative to previous trace data. Do not use job trace analysis to determine accurate job or transaction processing times.

A summary of job trace data collection and report commands is shown in Figure 4-2 on page 4-13.

Figure 4-2. Job Trace Data Collection and Report Commands

System Data Collection Command	Level of Data	Type of Data	Report Command	Information Shown on the Reports	When to Use the Command
STRJOBTRC	Job Program Files	Trace data	PRTJOBTRC ENDJOBTRC	Program name Control flow I/O operations Full/shared opens Exceptions Message handling Disk I/O summary	For program development To identify jobs or programs that perform poorly

For more information about the data collection or report commands, see Chapter 11, "Programmer Performance Utilities."

Program Instruction Analysis

The commands shown in Figure 4-3 are intended primarily for program development use, and would not normally be used in a production environment. The commands provide a detailed analysis of the instruction frequency within an individual program or set of programs.

A summary of program data collection and report commands is shown in Figure 4-3.

Figure 4-3. Program Data Collection and Report Commands

System Data Collection Command	Level of Data	Type of Data	Report Command	Information Shown on the Reports	When to Use the Command
STRSAM STRSAMCOL	Program	Sample data	PRTSAMDTA	Program instruction processing frequency	For program development For code optimization

For more information about the Start Sampled Address Monitor (STRSAM) command, see "Start Sampled Address Monitor (STRSAM) Command" on page 11-11. For more information about the Start Sampled Address Monitor Collection (STRSAMCOL) command, see "Start Sampled Address Monitor Data Collection (STRSAMCOL) Command" on page 11-12. For more information about the Print Sampled Address Monitor Data (PRTSAMDTA) command, see "Print Sampled Address Monitor Data (PRTSAMDTA) Command" on page 11-13.

File Use and Database Structure Analysis

The commands shown in Figure 4-4 provide an overview of the program file use and the database file structure of an application.

Figure 4-4 (Page 1 of 2). File Use and Structure Data Report Commands

System Data Collection Command	Level of Data	Type of Data	Report Command	Information Shown on the Reports	When to Use the Command
	Program File use structure		ANZPGM	Program file	For application use analysis

Figure 4-4 (Page 2 of 2). File Use and Structure Data Report Commands

System Data Collection Command	Level of Data	Type of Data	Report Command	Information Shown on the Reports	When to Use the Command
	Program File use structure		ANZDBF	Physical file structure	For application analysis
	Program File use structure		ANZDBFKEY	Logical file structure	For file analysis

For more information about the Analyze Program (ANZPGM) command, see "Analyze Program (ANZPGM) Command" on page 11-16. For more information about the Analyze Database File (ANZDBF) command, see "Analyze Database File (ANZDBF) Command" on page 11-19. For more information about the Analyze Database File Key (ANZDBFKEY) command, see "Analyze Database File Keys (ANZDBFKEY) Command" on page 11-21.

Job Analysis

Job analysis provides you with a view of the operational environment for all jobs, or a group of jobs, in the system at a given time. Use the information from a specific process analysis to improve the performance of the process. This analysis can help you improve the program environment to reduce the number of the following:

- Open files
- File buffer and work space sizes
- File open placement in a program
- Active programs

A summary of job data collection and report commands is shown in Figure 4-5.

Figure 4-5. Process Data Collection and Report Commands

System Data Collection Command	Level of Data	Type of Data	Report Command	Information Shown on the Reports	When to Use the Command
DSPACCGRP	Job Program Files		DSPACCGRP ANZACCGRP	File use Files used at the same time Open Data Path Buffer size Formats (size and number) I/O counts Duplicates PAG size Active programs	Reduce program size Reduce number of open files Reduce process access group (PAG) I/O Determine group job candidates

For more information about the DSPACCGRP command, see "Display Access Group (DSPACCGRP) Command" on page 11-24.

Chapter 5. Advisor

The advisor provides an easy-to-use way to improve many of the performance characteristics of your system.

The advisor fits into the set of performance tools between automatic system tuning and the more specialized tools (such as the Model System command) and the reports (such as a Print System Report).

Automatic system tuning is a useful method for maintaining the basic conditions for good performance. If it is set to work at each IPL, it resets the basic tuning values to the recommended settings for the system configuration and controlling subsystem. Dynamic automatic system tuning adjusts only pool sizes and activity levels of shared pools based on system activity as measured at user-specified intervals.

The advisor can help you to define specific tuning values and other parts of a processing environment to provide better performance for specific processing conditions on your system.

The advisor analyzes performance data you collect with the performance monitor and it can produce recommendations and conclusions to help improve performance. The advisor might recommend changes to basic system tuning values, and might list conclusions about conditions that could cause performance problems.

You can choose to have the advisor change system tuning values as it recommends, or you can decide to make only the changes you select. You can use the advisor's conclusions to make changes to your system, to guide further performance data collection, or to help you request performance reports containing more information and explanations.

The advisor can help you to improve system performance, but it will not identify or fix all performance problems. The performance information analyzed includes:

- Storage pool sizes and time slice values
- Disk and main storage utilization
- Communications activities
- Input/output processor utilization
- Unusual job activities — exceptions or excessive use of system resources.

The advisor does not:

- Make any recommendations for modifying specific application programs to improve their performance
- Consider the effects on performance of sieze locks, long waits, or trace data

The advisor is a good first tool to use to improve system performance. In many instances, it will be the only tool required to make the improvements you need. This chapter takes you through the process for using the advisor. In general, this process consists of the following steps:

1. Identify when the performance problems occur.
2. Use the performance monitor to collect data during the time periods when each problem seems to occur.

3. Request the advisor to analyze the data.
4. Use the advisor's output to change system tuning values, to guide further data collection, or to request other more detailed performance reports.
5. Observe the effects of any tuning changes, and decide if another cycle through this process is required to further improve performance or to eliminate unwanted side effects.

Notes:

1. The examples in this chapter show how to use the advisor, but they do not contain specific solutions for any performance problems that might exist on a particular AS/400 system.
2. Sometimes an analysis of data collected during normal system operation can help in selecting the advisor recommendations to implement to solve performance problems occuring at other times.

Collecting the Right Performance Data

Before collecting performance data, you should clearly describe the problem to be investigated. From system users' comments or your own experience, you can begin to formulate a description of the problem. The problem description does not need to be overly detailed or technical, just try to simply describe one problem. For example:

- Interactive (or batch) processing seems too slow.
- File updating should go faster.
- At times the entire system seems to be sluggish.

Next, determine when the problem is most likely to occur. Maybe interactive work is slow first thing in the morning. Perhaps batch processing seems slow late in the afternoon.

When you can clearly describe the problem and have determined when it seems to occur, you are ready to collect performance data to be analyzed by the advisor.

If possible, focus on collecting data for one problem at a time. Of course, try to collect the data when the problem seems to occur most often. You can decide later how much of the data you want the advisor to analyze. For more information about when to collect performance data and how much to collect, see the first few pages of Chapter 4, "Collecting System Performance Data."

The performance monitor is used to collect performance data. It can be run using the default data collection values. Information about how to use the performance monitor is in Chapter 4, "Collecting System Performance Data."

Requesting an Analysis

After performance data is collected for the time periods near when a problem seems to occur, you request the advisor to analyze all or part of that data. To start the advisor, you can select option 10 (Advisor) on the IBM Performance Tools/400 menu, or type the Analyze Performance Data (ANZPFRDTA) command on any command line.

Note: To analyze performance data from a library other than QPFRDATA when using the ANZPFRDTA command, type the command and press F4 (Prompt) to change the library name.

```
PERFORM                    IBM Performance Tools/400
                                                            SYSTEM:   RCHAS725
Select one of the following:

     1. Select type of status
     2. Collect performance data
     3. Print performance report
     4. Capacity planning/modeling
     5. Programmer performance utilities
     6. Configure and manage tools
     7. Display performance data
     8. System activity
     9. Performance graphics
    10. Advisor

    70. Related commands


Selection or command
===> 10

F3=Exit   F4=Prompt   F9=Retrieve   F12=Cancel   F13=User support
F16=System main menu
```

The next two steps in requesting a performance data analysis are:

- Select the member containing the performance data to analyze.
- Select the time intervals of data to analyze.

Selecting a Member

When option 10 (Advisor) is selected, or the ANZPFRDTA command is run, the Select Member for Analysis display appears.

```
                          Select Member for Analysis

Library . . . .   QPFRDATA

Type option, press Enter.
  1=Select   5=Display

Option  Member      Text                                          Date      Time
  _     Q900221115                                              11/01/90  22:01:07
  _     Q903611411                                              12/27/90  14:11:25
  _     Q903460708                                              12/12/90  07:08:43
  _     Q903411604                                              12/07/90  16:04:23
  _     Q903411420                                              12/07/90  14:20:22
  _     Q903411411                                              12/07/90  14:11:09







                                                                         BOTTOM
F3=Exit   F12=Cancel   F15=Sort by name   F16=Sort by text
F19=Sort by date/time
(C) COPYRIGHT IBM CORP. 1981, 1991.
```

The most important fields and columns on this display are as follows:

Library The name of the library containing the performance data members listed on this display. To see the contents of another library, type the new library name over the current one, and press the Enter key.

Option 1 Enter a 1 (Select) in this column to select only one member for analysis by the advisor, and press the Enter key.

Option 5 Enter a 5 (Display) to display performance data just as though you had entered the DSPPFRDTA command from the main menu.

Member The name of a member in the named library. Each member contains performance data collected by the performance monitor.

Text Can contain a text description of the member to help identify it.

Date The date when the data in the member was collected.

Time The time when the data in the member was collected.

To request an analysis, select only one member that contains performance data collected during a time when the problem occurred. When you select a member and press the Enter key, a Select Time Intervals to Analyze display appears.

Note: When you return to the Select Member for Analysis display, the 1 typed for the member remains. This is a reminder that you may wish to display this member.

Selecting Time Intervals

```
                        Select Time Intervals to Analyze

Member . . . . . . . :   Q900221115      Library  . . . . . . :   QPFRDATA

Type options, press Enter.
  1=Select

                    Transaction  --CPU Util---  --High---  -Pool Fault-
Opt  Date   Time    Count   Rsp  Tot  Int  Bch  Dsk  Unit  Mch  Usr  ID    Excp
 _   11/01  11:21    595    1.4   98   43   42   22  0019    3   13  04   11599
 _   11/01  11:26    443    1.8   98   25   61   27  0026    2    9  04   12403
 _   11/01  11:31    438    1.0   95   22   62   29  0019    1   11  05   13206
 _   11/01  11:36    496    1.4   89   33   46   21  0012    1   13  04   13511
 _   11/01  11:41    459    1.4   98   22   64   22  0016    2   14  04   11966
 _   11/01  11:46    344    1.2   96   17   68   17  0016    2    9  07   10739
 _   11/01  11:51    394     .8   94   13   73   18  0009    1   14  05    6701
 _   11/01  11:56    402    1.1   99   20   69   15  0026    1   13  04    8219
 _   11/01  12:01    312    1.0   96   18   71   21  0019    2   14  07    9522
 _   11/01  12:06    389    1.0   96   25   63   19  0019    1   13  07   23382
 _   11/01  12:11    281    1.3   93   16   68   23  0019    3   18  07   11487
                                                                          More...
F3=Exit   F5=Refresh   F11=Display histogram   F12=Cancel   F13=Select all
```

The Select Time Intervals to Analyze display lists all the time intervals of performance data collected in the library member selected on the Select Member for Analysis display. To analyze a different member, press F12 (Cancel) to return to the Select Member for Analysis display.

The columns on the Select Time Intervals to Analyze display can help you focus the analysis on time intervals when the suspected performance problem seems to have occurred. If there are no obvious reasons to select only some of the displayed time intervals, you can select them all for analysis by pressing F13 (Select all).

The columns on this display that can help you select particular time intervals to analyze are as follows:

Transaction Count
The total number of transactions during the time interval.

Note: Performance problems often are noticed first when the system is busiest.

Transaction Rsp
The average transaction response time during the time interval.

Note: Long average transaction response times can cause complaints that the system is slow.

CPU Util The percent of the elapsed time interval when the processing unit was busy. For systems with several processing units, this percentage is across all of the processing units. The percent of processing unit utilization during the time interval by various types of work are shown as follows:

Tot	All types of work
Int	Interactive work
Bch	Batch work

Note: These values can help you focus on time intervals when the processing unit utilization is high for one or more types of work. A system might appear to be running slowly when processing unit utilization is high for many minutes.

High Dsk and Unit
The disk arm that was most busy during the time interval is identified in the *Unit* column. The percentage of the time interval the disk was busy is in the *Dsk* column.

Pool Fault
The number (ID) of the user storage pool with the highest page fault rate during the time interval. The numbers of faults for various types of pools are in:

Mch	Machine pool faults
Usr	User pool faults

Note: A page fault occurs when the system must move part of a program, data, or other item into a storage pool. System performance might seem sluggish when the page fault rate is high.

Excp The number of program exceptions during the time interval. These are caused by system-detected conditions (such as address overflows, floating-point errors, and so on).

Note: An analysis of time intervals with peaks in the number of exceptions could be useful in solving performance problems.

When one or more time intervals are selected for analysis, press the Enter key to request the analysis by the advisor.

Using a Histogram

Sometimes a graph of the data for one of the performance values in the data makes it easier to select specific time intervals of data for analysis. To define and display a graph (called a **histogram**), press F11 (Display histogram) on the Select Time Intervals to Analyze display. The display then changes to include the Select Histogram window.

```
                      Select Time Intervals to Analyze
 ............................................
 :              Select Histogram            : ary  . . . . . . :    QPFRDATA
 :                                          :
 :  Type option, press Enter.               :
 :    1=Select                              :
 :                                          :
 :  Opt     View                            :  --High---  -Pool Fault-
 :   _      Transaction count               :  Dsk  Unit  Mch   Usr ID  Excp
 :   _      Transaction response time       :   22  0019    3    13 04  11599
 :   _      Total CPU utilization           :   27  0026    2     9 04  12403
 :   1      Interactive CPU utilization     :   29  0019    1    11 05  13206
 :   _      Batch CPU utilization           :   21  0012    1    13 04  13511
 :   _      High disk utilization           :   22  0016    2    14 04  11966
 :                               More...    :   17  0016    2     9 07  10739
 :  F3=Exit   F12=Cancel                    :   18  0009    1    14 05   6701
 :                                          :   15  0026    1    13 04   8219
 :..........................................:   21  0019    2    14 07   9522
  _   11/01  12:06    389   1.0   96   25   63   19  0019    1    13 07  23382
  _   11/01  12:11    281   1.3   93   16   68   23  0019    3    18 07  11487
                                                                       More...
 F3=Exit   F5=Refresh   F11=Display histogram   F12=Cancel   F13=Select all
```

The *View* column lists the performance values that can be selected to define the Y (vertical) histogram axis. The X (horizontal) histogram axis always shows the time intervals contained in the member.

As an example, to make it easier to see the time intervals where CPU utilization is the highest, you could select one of the CPU utilization views. A sample histogram for Interactive CPU Utilization follows:

```
                      Select Time Intervals from Histogram

 Type a '1' under each interval to select, press Enter.

 Interactive CPU utilization
       48 :
       44 :
       40 :*            *
       36 :*            *
       32 :*  *         *   *
       28 :*  *        **   **   *
       24 :** *     * ***   ***  *
       20 :*****  * * ****  ***  *
       16 :****** ********* ******
       12 :***********************
        8 :***********************
        4 :***********************
        0 :***********************
           11111111111111111111111
           +-------+-------+-------+-------+-------+-------+-------+-------+
         11:21   12:01   12:41
 F3=Exit   F5=Refresh   F11=Display histogram   F12=Cancel   F13=Select all
```

On this example, it is easy to see and select the time intervals of greatest interactive processing unit use. The number 1 is entered to select each time interval to be analyzed. All of the intervals can be quickly selected by pressing F13 (Select all), as shown in the example.

After the Enter key is pressed on the Select Time Intervals to Analyze display or on the Select Time Intervals from Histogram display, the advisor analyzes the performance data for the selected time intervals.

Notes:

1. An analysis of large amounts of performance data can take a long time and could affect system performance for other users.
2. The analysis performed by the advisor includes all of the types of performance data for the selected time intervals, and is not limited to the type of data selected to create the histogram.

Using the Advisor's Results

Depending on the content of the selected performance data, the advisor can produce recommendations, conclusions, and interval conclusions. What these are and how you can use them are explained as you look at the following examples.

When a performance data analysis has completed, the Display Recommendations display shows the results.

```
                        Display Recommendations
                                                       SYSTEM:   RCHAS725
Member . . . . . . . :   Q900221115     Library  . . . . . . :   QPFRDATA
System . . . . . . . :   RCHAS725       Version/Release  . . :   2/ 1.0
Start date . . . . . :   11/01/90       Model  . . . . . . . :   D80
Start time . . . . . :   11:21:13       Serial number  . . . :   10-00725

Type options, press Enter.
  5=Display details

Option    Recommendations and conclusions
  _                     Recommendations
  _       Decrease pool size for listed pools.
  5       Increase pool size for listed pools.
  _       Decrease activity level in listed pools.
  _       Add more main storage.
  _       ASP space capacity exceeded guideline of 80.0%.
  _                      Conclusions
  _       Pool fault rates exceeded guideline.
  _       Pool fault rates below guideline.
  _       SDLC utilizations exceeded 50% guideline.
                                                                    More...
F3=Exit   F6=Print   F9=Tune system   F12=Cancel   F21=Command line
```

Understanding Recommendations

The *Recommendations* section of this display deals with conditions that significantly affect system performance. The recommendations result from comparing the system values and conditions in the analyzed performance data to the basic AS/400 performance guidelines.

The recommendations suggest changes to the basic system tuning values that can improve performance. They also list problems that can be solved by other actions. In this example, the recommendations about changing pool sizes can be carried out by changing system tuning values. But, the recommendation about ASP (auxiliary storage pool) space capacity might require redefining the use of system disk space or adding to system disk capacity. You might need technical assistance to complete this type of recommendation.

To see more details about a recommendation, type 5 in the *Option* column. As an example, the following displays show the details for the example recommendation *Increase pool size for listed pools*.

```
                        Display Detailed Recommendation

Recommendation:

Increase pool size for listed pools.

Detailed recommendation:
PFR2567
Technical description . . . . . . . . :   The following table shows the
pool identifier, the current pool size, and the suggested pool size.

       Pool       From        To       Pool       From        To

          1      10238      12193



Increasing the pool size will reduce the page fault rate which will
                                                                  More...
Press Enter to continue.

F3=Exit   F12=Cancel
```

In this example only pool 1 should be increased in size. The text beginning at the bottom of this display and continuing on the following displays discusses the effects of changing a pool's size.

```
                        Display Detailed Recommendation

Recommendation:

Increase pool size for listed pools.

Detailed recommendation:
improve the response time and throughput of jobs in this pool.

Decreasing the pool size will free storage that may in turn be given to
pools with high fault rates.

Removing a pool will free storage that may in turn be given to pools with
high fault rates.

A pool will be increased by at least ten percent of its current size.
Pools that are decreased will all be decreased by the same percentage,
with ten percent of the current size as the maximum amount of decrease.
For example, if a 1500K pool needs storage, and a 2000K and 1000K pool can
                                                                  More...
Press Enter to continue.

F3=Exit   F12=Cancel
```

Many recommendations include this type of information to help you choose the right changes to make to your system.

Changing System Tuning Values

To see and select the tuning changes related to the recommendations, press F9 (Tune system) on the Display Recommendations screen. A display similar to the following appears:

```
                         Select Tuning Recommendations

                                             Advisor          Current      Data
                        Name/              Recommended        System    Collection
Value To Be Changed     Number                Value           Value        Value
POOLSIZE (K)            *MACHINE              12193            9420        11085
ACTIVITY LEVEL          *BASE                     6               7            6
POOLSIZE (K)            *INTERACT             70755           39683        70755
ACTIVITY LEVEL          *INTERACT                27              21           31
POOLSIZE (K)            *SPOOL                   80              49           80
ACTIVITY LEVEL          *SPOOL                    3               2            3



                                                                          Bottom
Select one of the following:

     1. Tune to advisor's recommendations
     2. Restore system to data collection values
Selection

     _

F3=Exit   F12=Cancel
```

On this Select Tuning Recommendations display you have several choices:

- Select menu option 1 (Tune to advisor's recommendations) to have the advisor make all the changes shown in the *Advisor Recommended Value* column. Usually this is a good choice to make when starting to solve a performance problem.
- Leave the values as they are listed in the *Current System Value* column.
- Select menu option 2 (Restore system to data collection values) to have the advisor set the values as they were when the analyzed performance data was collected (shown in the *Data Collection Value* column).
- Write down the tuning values that fit your needs, and use the appropriate system commands to change the values individually.

Note: The analysis and recommendations are based on the *Data Collection Values*. The *Current System Value* column is there for your reference and in case you want to reset your configuration to what it was at the time of data collection. If the *Advisor Recommended Value* equals the *Data Collection Value*, then the advisor is saying that this is an adequate setting for the workload analyzed. If the *Advisor Recommended Value* does not equal the *Data Collection Value*, then you will see recommendations and conclusions as to what should be changed.

Understanding Conclusions

The *Conclusions* section of the Display Recommendations display lists conditions that could have affected performance when the analyzed data was collected. These conditions can include threshholds reached, save and restore activities, teleprocessing line errors, and so on.

```
                          Display Recommendations
                                                      SYSTEM:   RCHAS725
Member . . . . . . . :   Q900221115     Library . . . . . . :   QPFRDATA
System . . . . . . . :   RCHAS725       Version/Release . . :   2/ 1.0
Start date . . . . . :   11/01/90       Model . . . . . . . :   D80
Start time . . . . . :   11:21:13       Serial number . . . :   10-00725

Type options, press Enter.
  5=Display details

Option      Recommendations and conclusions
  _                          Recommendations
  _         Decrease pool size for listed pools.
  _         Increase pool size for listed pools.
  _         Decrease activity level in listed pools.
  _         Add more main storage.
  _         ASP space capacity exceeded guideline of 80.0%.
  _                            Conclusions
  5         Pool fault rates exceeded guideline.
  _         Pool fault rates below guideline.
  _         SDLC utilizations exceeded 50% guideline.
                                                                More...
F3=Exit   F6=Print   F9=Tune system   F12=Cancel   F21=Command line
```

Some conclusions describe conditions that caused the advisor to make particular recommendations. Other conclusions not related to recommendations can be used as guides for collecting more performance data, or for adjusting the system.

To see more details about a conclusion, type 5 in the *Option* column. The following example is the display showing details for the conclusion *Pool fault rates exceeded guideline* that supports the recommendation to increase the size of pool 1.

```
                          Display Detailed Conclusion

Conclusion:

Pool fault rates exceeded guideline.

Detailed conclusion:
PFR2513
Technical description . . . . . . . . :   The following table shows the
pool identifier, the maximum fault rate over all the intervals, the fault
rate guideline, the number of intervals the guideline was exceeded out of
1 intervals, and the date and time the maximum fault rate occurred. For
pool 2 (*BASE) the guideline is based on the fact that there are no user
jobs running in *BASE.

         ID        Rate       Guide  Intervals      Date       Time

          1         3.6        3.0          3   11/01/91   12:31:04

                                                                More...
Press Enter to continue.

F3=Exit   F12=Cancel
```

In this example, the guideline of 3 faults was exceeded for pool 1 in three of the analyzed time intervals. The maximum fault rate was 3.6.

Understanding Interval Conclusions

The *Interval Conclusions* section of the Display Recommendations display contains the detailed data to support the conclusions for the analyzed time intervals.

```
                         Display Recommendations
                                                         SYSTEM:   RCHAS725
Member . . . . . . . :   Q900221115      Library . . . . . . :   QPFRDATA
System . . . . . . . :   RCHAS725        Version/Release . . :   2/ 1.0
Start date . . . . . :   11/01/90        Model . . . . . . . :   D80
Start time . . . . . :   11:21:13        Serial number . . . :   10-00725

Type options, press Enter.
  5=Display details

Option      Recommendations and conclusions
                           Interval Conclusions
  5         Pool fault rates above guideline.
  _         Total disk I/O was 225.  (179 Reads and 46 Writes)
  _         No performance problems found on listed TRLAN lines.
  _         Total system fault rate above 50.0 guideline.
  _         Thrashing may have occurred.

                                                                   Bottom
F3=Exit   F6=Print   F9=Tune system   F12=Cancel   F21=Command line
```

To see more details about an interval conclusion, type 5 in the *Option* column. The following example is the display showing details for the sample interval conclusion, *Pool fault rates above guideline*, which supports the conclusion that pool fault rates exceeded the guideline.

```
                    Display Detailed Interval Conclusion

Interval conclusion:

Pool fault rates above guideline.

Detailed interval conclusion:
PFR2553
Technical description . . . . . . . . :   The following table shows the
pool identifier, the fault rate, and the time the fault rate exceeded the
guideline.

         Id       Rate       Guide        Date        Time

         1         3.0         3.0    11/01/90    11:21:13
         1         3.1         3.0    11/01/90    12:11:06
         1         3.6         3.0    11/01/90    12:31:04

                                                                   Bottom
Press Enter to continue.

F3=Exit   F12=Cancel
```

In this example we see exactly when, and by how much, the fault rate guideline was exceeded for pool 1 in the analyzed time intervals.

Following is the details display for another type of sample interval conclusion, *Total disk I/O was 225 . . . :*

```
                     Display Detailed Interval Conclusion

 Interval conclusion:

 Total disk I/O was 225.  (179 Reads and 46 Writes)

 Detailed interval conclusion:
 PFR2847
 Cause . . . . . :   This is the sum of all the disk I/O for all the
 selected intervals for all the disk devices.  This does not indicate a
 problem; this data is provided simply for reference.








                                                                  Bottom
 Press Enter to continue.

 F3=Exit   F12=Cancel
```

An interval conclusion like this one provides information but does not support a conclusion or recommendation. It does not report a problem but provides information that can be helpful in understanding how your system is performing.

Tune System by Advisor's Recommendations

After you request a performance analysis, and look over the results, often the next step is to have the advisor tune the system as it recommends. Do this by selecting menu option 1 (Tune to advisor's recommendations) on the Select Tuning Recommendations display.

Next, observe the effects of the changes. Use the performance monitor to collect more performance data during the next time period when you expect the problem to occur. Also, observe the system and watch for the usual symptoms of the problem. Ask users who experienced the problem if they still notice it. Watch for any possible unwanted side effects from the tuning changes. These can occur if the changes are not fully compatible with some of your processing requirements, or if several problems are being worked on.

The first attempt to solve a basic performance problem can be successful. But sometimes the steps described in this chapter must be repeated until the best possible performance is achieved for your system and your processing requirements.

The original problem may continue or new problems may occur. The advisor might have no further recommendations or conclusions that you can use. At this time you could use other performance reports and commands to work on the problem. These are described in Chapter 3, "Performance Tuning."

Sometimes tuning alone will not solve performance problems. To handle the intended work load, a system might need additional main storage, disk storage, or processing speed. The Model System (MDLSYS) command can be used to determine if system processing capacities should be increased. For more information about the MDLSYS command and capacity planning, see Chapter 10, "Capacity Planning and Performance Prediction."

Chapter 6. Displaying Performance Data

This chapter describes how to interactively view the performance data collected by the Start Performance Monitor (STRPFRMON) command. This Display Performance Data function can be used either while the performance monitor is active or after the collection is completed.

Note: The data collection does not need to contain the trace data in order to use this display function. Trace data may be required, however, to further analyze performance problems isolated by this function.

Display Performance Data

To interactively display sample performance data, you can:

- Type the Display Performance Data (DSPPFRDTA) command on any command line using the default value of *SELECT for the member parameter.
- Select option 7 (Display performance data) on the IBM Performance Tools/400 menu.
- Select option 10 (Advisor) on the IBM Performance Tools/400 menu and then select option 5 (Display) for the member parameter or select option 7 and the Select Performance Member display appears.

```
                              Select Performance Member

 Library  . . . .   QPFRDATA

 Type option, press Enter.
   1=Select

 Option  Member      Text                                  Date      Time
   _     Q892621356  Pfr Test Case A 14:00 - 15:00 Purge   09/19/90  13:56:06
   _     Q892611334  Pfr Test Case C 13:30 - 14:30 Purge   09/18/90  13:34:39
   _     Q892581430  Pfr Test Case B 14:30 - 15:30 Purge   09/15/90  14:30:04
   _     Q892571400  Pfr Test Case A 14:00 - 15:00 Purge   09/14/90  14:00:43











                                                                      Bottom
 F3=Exit  F5=Refresh  F12=Cancel  F15=Sort by member  F16=Sort by text
 F19=Sort by date/time
 (C) COPYRIGHT IBM CORP. 1981, 1991.
```

The member name, a text description, and the date and time you collected each set of performance data appear on this display. If you cannot find the data you want to display, use the Roll keys to page through the list of members or use the appropriate function key to sort the sets of performance data. You can sort the data by member name, text description, or by the date and time the member was created. When you find the performance data you want to display, type a 1 in the corresponding *Option* field.

If you are searching for a member located in a library that is different from the one currently listed in the *Library* field at the top of the display, type a new library name in the field and press the Enter key. A list of the performance members available in the library you specified appears. You can then select to display one of them.

To indicate the performance member to display, either make a selection on the Select Performance Member display or enter a valid member name in the MBR parameter of the DSPPFRDTA command. The Display Performance Data function then starts to read the performance database files. All the performance information required by this function is processed now, so there is reasonable response time when moving between displays later.

Note: The initial processing may cause a noticeable delay in presenting the first display.

After all the data is processed, the main display for the Display Performance Data function appears.

```
                              Display Performance Data

Member . . . . . . . .   Q892621356           F4 for list
  Library  . . . . . .     QPFRDATA

Elapsed time . . . . :   00:54:16        Version  . . . . . . :   2
System . . . . . . . :   MACHINE1        Release  . . . . . . :   1.0
Start date . . . . . :   09/19/90        Model  . . . . . . . :   B60
Start time . . . . . :   13:56:48        Serial number  . . . :   10-12883

CPU utilization (interactive)  . . . . . . . . . :         42.01
CPU utilization (other)  . . . . . . . . . . . . :         34.92
Job count  . . . . . . . . . . . . . . . . . . . :           539
Transaction count  . . . . . . . . . . . . . . . :         10836
Transactions per hour  . . . . . . . . . . . . . :         11980
Average response (seconds) . . . . . . . . . . . :          1.74
Disk utilization (percent) . . . . . . . . . . . :          5.24
Disk I/O per second  . . . . . . . . . . . . . . :         138.9


F3=Exit      F4=Prompt   F5=Refresh   F6=Display all jobs   F10=Command entry
F12=Cancel   F24=More keys
```

The *Member, Library*, and *Elapsed time* fields appear on every display where performance data is shown. The *Member* field shows the performance member that contains the data, and the *Library* field shows the library in which the performance database files are located. The *Elapsed time* field displays the length of time that the performance measurement ran or has run (if the Performance Monitor is currently collecting data into this member).

On this display you can change both the *Member* and *Library* fields. If you type a new member name in the *Member* field and press the Enter key, the data in that member appears on the display. If you type a new library name in the *Library* field and press the Enter key, the program tries to locate the member in the specified library. If you press F4 (Prompt) after you enter the library name, the Select Performance Member display uses the specified library to present a list of data collections.

The remaining fields on the Display Performance Data display are as follows:

System The system name of the AS/400 system on which the performance data was collected.

Start date The date on which the performance measurement began.

Start time The time at which the performance measurement began.

Version The version of the OS/400 program that collected the performance data.

Release The release level of the OS/400 program that collected the performance data.

Model The model of the AS/400 system on which the performance data was collected.

Serial number
The serial number of the AS/400 system on which the performance data was collected.

CPU utilization (interactive)
The percentage of the elapsed time during which the processing unit was used by interactive jobs (interactive, PC Support, display station pass-through, distributed data management (DDM), System/36, and multiple requester terminal (MRT)). For multiple-processor systems, this is the average use across all processors.

CPU utilization (other)
The percentage of the elapsed time during which the processing unit was used by all jobs except interactive jobs. For multiple-processor systems, this is the average use for other than interactive.

Job count The number of active jobs that were running in the measurement. This count includes only those jobs that used any CPU resource or performed any I/O activity. This count does not include any licensed internal code tasks.

Transaction count
The number of transactions that occurred during the entire measurement.

Transactions per hour
The average number of transactions that occurred per hour during the measurement.

Average response (seconds)
The average internal response time (in seconds) per transaction.

Disk utilization (percent)
The percentage of the elapsed time during which the disks were utilized. This value represents an average for all disk units.

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.

Disk I/O per second
The average number of disk I/O operations performed each second during the entire measurement.

To help you analyze the performance data, the Display Performance Data function highlights the values on this display which exceed the threshold values listed in Figure 10-3 on page 10-6. Therefore, if the interactive CPU utilization or the disk utilization exceeds the threshold, the field is highlighted on the display.

To access a command line after you start the Display Performance Data function, press F10 (Command entry). This allows you to work from a command entry display without exiting the display function. Once you exit the command entry, you are immediately returned to the Display Performance Data display without having to experience the initial processing delay.

To better understand system performance, you might want to view the data sorted by category. The second set of function keys on this display allows you to group the performance data by subsystem, job type, or interval.

```
                              •
                              •
                              •

F13=Display by subsystem   F14=Display by job type   F15=Display by interval
F24=More keys
```

By categorizing the data, you might be able to isolate a group of jobs that require further analysis. If you do, you can then display the performance statistics for individual jobs.

The next three sections describe the displays that show the performance data separated into the subsystem, job type, and interval categories.

Display Performance Data by Subsystem

If you press F13 on the Display Performance Data display, the Display by Subsystem display appears.

```
                               Display by Subsystem

Member . . . . . . . :   Q892621356     Elapsed time . . . . :    00:54:16
  Library  . . . . . :    QPFRDATA

Type options, press Enter.  Press F6 to display all jobs.
  5=Display jobs

                             CPU         Job          Tns       Average       Disk
Option    Subsystem          Util       Count        Count      Response       I/O
  _       *MACHINE           9.86         130            0           .00     42933
  _       QBASE               .00           1            0           .00        29
  _       QBATCH            17.99          32            0           .00     23337
  _       QCTL               2.35           3          295           .71     10937
  _       QINTER            40.28         355        10221          1.77     79822
  _       QSNADS              .00           0            0           .00         0
  _       QSPL               5.33         145            0           .00     22639

                                                                           Bottom
F3=Exit   F6=Display all jobs   F12=Cancel   F14=Display by job type
F15=Display by interval
```

This display categorizes the performance data according to the subsystem in which the activity occurred. The columns on the Display by Subsystem display are as follows:

Subsystem
The name of a subsystem that existed on the system during the measurement. The subsystem *MACHINE is always present because this represents all entries that do not have a subsystem associated with them, such as licensed internal code tasks.

CPU Util The percentage of the elapsed time during which the processing unit was used by the jobs running in a particular subsystem. For multiple-processor systems, this is the average use across all processors.

Job Count
The number of active jobs running in a particular subsystem during the measurement. This count includes only those jobs that used any CPU resource or performed any I/O activity.

Tns Count
The number of transactions performed by the jobs running in the subsystem.

Average Response
The average internal response time (in seconds) per transaction in the particular subsystem.

Disk I/O The number of physical disk I/O operations performed by the jobs in the subsystem.

From this display you may be able to isolate a single subsystem or group of subsystems that are of particular interest. To view the performance data for the jobs in particular subsystems, type a 5 in the appropriate *Option* fields and press the Enter key. If you do not want to select a particular subsystem, but would rather view the data for all the jobs in the measurement, press F6 (Display all jobs).

Display Performance Data by Job Type

If you press F14 on the Display Performance Data display, the Display by Job Type display appears.

```
                              Display by Job Type

Member . . . . . . . :    Q892621356       Elapsed time . . . . :    00:54:16
  Library  . . . . . :      QPFRDATA

Type options, press Enter.  Press F6 to display all jobs.
  5=Display jobs

                           CPU        Job         Tns     Average      Disk
Option   Job Type          Util      Count       Count   Response       I/O
  _      Autostart          .00          0           0        .00         0
  _      Batch            19.65         69           0        .00     33769
  _      Evoke              .00          0           0        .00         0
  _      Interactive      40.94        354       10502       1.74     81815
  _      LIC               9.66        127           0        .00     42347
  _      Pass-Through       .08          3          14       2.42       195
  _      Sbs Monitor        .97          4           0        .00      2984
  _      System             .20          3           0        .00       586
  _      Writer            4.32        107           0        .00     18001

                                                                      Bottom
F3=Exit   F6=Display all jobs   F12=Cancel   F13=Display by subsystem
F15=Display by interval
```

This display categorizes the performance data according to the job types of the jobs running on your system. The columns on the Display by Job Type display are as follows:

Job Type A type of job present on the system during the measurement.

The values that can appear in this column are as follows:

- Autostart
- Batch
- DDM (distributed data management) Server
- Evoke
- Interactive
- LIC (licensed internal code)
- MRT (multiple requester terminal)
- Pass-through
- PC Support
- Reader
- Sbs (subsystem) Monitor
- System
- System/36
- Writer

CPU Util The percentage of the elapsed time during which the processing unit was used by the jobs of a particular type. For multiple-processor systems, this is the average use across all processors.

Job Count
The number of active jobs of a particular type running during the measurement. This count includes only those jobs that used any CPU resource or performed any I/O activity.

Tns Count
The number of transactions performed by the jobs of a particular type.

Average Response
The average internal response time (in seconds) per transaction for the jobs of a particular type.

Disk I/O The number of physical disk I/O operations performed by the jobs of a particular type.

From this display you may be able to isolate a single job type or group of job types that are of particular interest. To view the performance data for the jobs of particular job types, type a 5 in the appropriate *Option* fields and press the Enter key. If you do not want to select a particular job type, but would rather view the data for all the jobs in the measurement, press F6 (Display all jobs).

Display Performance Data by Interval

If you press F15 on the Display Performance Data display, the Display by Interval display appears.

```
                          Display by Interval

Member . . . . . . . :   Q892621356     Elapsed time . . . . :    00:54:16
  Library  . . . . . :     QPFRDATA

Type options, press Enter.  Press F6 to display all jobs.
  5=Display jobs

                                CPU       Job       Tns    Average      Disk
Option  Date       Time         Util      Count     Count  Response      I/O
  _     09/19/89   14:01:48     71.35      212       957       1.70    12644
  _     09/19/89   14:06:47     82.68      220       982       1.87    12893
  _     09/19/89   14:11:47     87.27      229       955       1.96    14717
  _     09/19/89   14:16:46     63.05      236       921       1.83    11693
  _     09/19/89   14:21:46     55.34      207       823       2.11     9715
  _     09/19/89   14:26:45     54.79      204       885       2.54     9585
  _     09/19/89   14:31:45     49.17      209       806       1.37     9787
  _     09/19/89   14:36:44     58.54      215      1021       1.31    11073
  _     09/19/89   14:41:45     71.72      241      1207       1.56    15103
  _     09/19/89   14:46:44     71.04      222      1089       1.60    15018
                                                                       More...
F3=Exit   F6=Display all jobs    F12=Cancel    F13=Display by subsystem
F14=Display by job type
```

This display categorizes the performance data according to the collection intervals which occurred during the measurement. The columns on the Display by Interval display are as follows:

Date The end date of the data collection interval.

Time The end time of the data collection interval.

CPU Util The percentage of the time in the interval during which the processing unit was used. For multiple-processor systems, this is the average use across all processors.

Job Count
The number of active jobs running during the particular interval. This count includes only those jobs that used any CPU resource or performed any I/O activity. This count does not include any licensed internal code tasks.

Tns Count
The number of transactions performed by all jobs during the interval.

Average Response
The average internal response time (in seconds) per transaction for all jobs during the interval.

Disk I/O The number of physical disk I/O operations performed by all jobs during the interval.

From this display, you may be able to isolate a single interval or group of intervals that are of particular interest. To view the performance data for the jobs in particular intervals, type a 5 in the appropriate *Option* fields and press the Enter key. If you do not want to select a particular interval, but would rather view the data for all the jobs in the measurement, press F6 (Display all jobs).

Display Jobs

If you selected a subsystem on the Display by Subsystem display, selected a job type on the Display by Job Type display, selected an interval on the Display by Interval display, or pressed F6 (Display all jobs) on any of these or the Display Performance Data display, the Display Jobs display appears.

```
                                Display Jobs

Subsystem . . . . :   QINTER              Member  . . . . . :    Q892621356
Elapsed time  . . :   00:54:16              Library . . . . :      QPFRDATA

Type options, press Enter.
  5=Display job detail

                                              Job      CPU       Tns      Avg     Disk
Option  Job          User         Number   Type     Util      Count    Rsp      I/O
  _     DSP01        QSECOFR      022220   INT      1.15        202    1.2      941
  _     DSP100       QPGMR        022213   INT       .70        155    1.4      694
  _     DSP89        QPGMR        022219   INT       .64         75    3.6      674
  _     DSP02        QSYSOPR      022222   INT       .63        127    2.2      956
  _     DSP47        QPGMR        022210   INT       .62        164    1.1      888
  _     DSP09        QSECOFR      022215   INT       .59        129    1.2      661
  _     DSP57        QPGMR        022223   INT       .56        151    2.0      625
  _     DSP38        QPGMR        022217   INT       .56         52    2.7     2177
  _     DSP92        QPGMR        022207   INT       .55        163    2.3      706
  _     DSP03        QSYSOPR      022168   INT       .52         96    1.8     1235
                                                                              More...
F3=Exit          F12=Cancel   F15=Sort by job    F16=Sort by job type
F19=Sort by CPU  F24=More keys
```

This display appears when you request to view the jobs in a particular subsystem. If you request a job type or interval, the *Subsystem* indicator at the top of the display is replaced by a *Job Type* or a *Interval* indicator. Also, if you selected a particular job type, the *Job Type* column does not appear since all the jobs have the same type as indicated by the *Job Type* field at the top of the display. If you request to see all the jobs (by pressing F6 on the Display by Subsystem, the Display by Job Type, or the Display by Interval displays) the appropriate indicator (*Subsystem*, *Job Type*, or *Interval*) appears at the top of the display showing a value of '*ALL' and the *Job Type* column is present. If F6 is pressed from the Display Performance Data display, there is no indicator, such as subsystem, job type, or interval, at the top of the display. Also, in this case, the *Job Type* column would be present.

The columns that can appear on the Display Jobs display are as follows:

Job The job name for the performance data being displayed.

User The user profile associated with the job.

Number
The number assigned to the job.

Job Type
The type of the job. The values that can appear in this column and their meanings are as follows:

ASJ	Automatic start job
BCH	Batch
DDM	DDM server
EVK	Evoke
INT	Interactive

LIC	Licensed internal code task
MRT	Multiple requester terminal
PCS	PC Support
PTH	Pass-through
RDR	Reader
S36	System/36
SBS	Subsystem monitor
SYS	System
WTR	Writer

CPU Util
The percentage of the elapsed time during which the processing unit was used by the job. For multiple-processor systems, this is the average use across all processors.

Tns Count
The number of transactions performed by the job.

Avg Rsp
The average internal response time (in seconds) per transaction for the job.

Disk I/O
The number of physical disk I/O operations performed by the job.

From this display you can select particular jobs for which more detailed information can be displayed.

Display Job Detail

If you type a 5 in the *Option* field next to a job on either the Display Jobs display or the Display Remote Jobs display (see page 6-19), and press the Enter key, the Display Job Detail display appears.

```
                              Display Job Detail

Job . . . . . . . . :   DSP01          Job type . . . . . . :   INT
  User . . . . . . . :    QSECOFR      Subsystem  . . . . . :   QINTER
  Number . . . . . . :    022220       Pool . . . . . . . . :   04
Member . . . . . . . :   Q892621356    Priority . . . . . . :   20
  Library  . . . . . :    QPFRDATA     Elapsed time . . . . :   00:54:16


              CPU        Tns     Average      Disk   Act->  Wait->    Act->
Interval  Seconds      Count   Response       I/O    Wait    Inel     Inel
14:01:48     .000          0         .0         0      .0      .0       .0
14:06:47     .000          0         .0         0      .0      .0       .0
14:11:47     .883          2        3.0       293      .8      .0       .0
14:16:46    4.930         18        2.8       138     3.6      .0       .0
14:21:46   12.935         43        1.7       116     8.6      .0       .0
14:26:45    1.938         25         .4        67     5.0      .0       .0
14:31:45     .693         10         .5        28     2.0      .0       .0
                                                                     More...
Press Enter to continue.

F3=Exit   F11=View 2   F12=Cancel   F15=Sort by interval
F24=More keys
```

The Display Job Detail display provides you with the performance data for a particular job, broken down by collection intervals. This display presents the performance information using three different views, which can be accessed by function keys. F11 shows you the next view in the series.

Besides the common fields *Member, Library*, and *Elapsed time*, which appear on each display, the following fields are shown in the header portion of each view:

Job The job name for the performance data being displayed.

User The user profile associated with the job.

Number The number assigned to the job.

Job Type The type of the job. Refer to the section describing the Display Jobs display for a list of possible values.

Subsystem
The subsystem in which the job ran.

Pool The pool in which the job ran during the first collection interval.

Priority The priority at which the job ran during the first collection interval.

The columns in the first view of the Display Job Detail display are:

Interval The end time for the collection interval.

CPU Seconds
The number of seconds the processing unit was utilized by the job during the collection interval.

Tns Count
The number of transactions performed by this job during the collection interval.

Average Response
The average internal response time (in seconds) per transaction during a collection interval.

Disk I/O The total number of physical I/O operations performed during the collection interval. This value is the sum of the synchronous/asynchronous, database/nondatabase reads and writes shown in View 2.

Act → Wait
The average number of active-to-wait state transitions that occurred each minute during the collection interval.

Wait → Inel
The average number of wait-to-ineligible state transitions that occurred per minute during the collection interval.

Act → Inel
The average number of active-to-ineligible state transitions that occurred per minute during the collection interval.

If you press F11 (View 2) on the first view of the Display Job Detail display, the second view appears.

```
                            Display Job Detail

Job . . . . . . . . :   DSP01          Job type . . . . . . :   INT
  User . . . . . . . :    QSECOFR      Subsystem  . . . . . :   QINTER
  Number . . . . . . :    022220       Pool . . . . . . . . :   04
Member . . . . . . . :  Q892621356     Priority . . . . . . :   20
  Library  . . . . . :    QPFRDATA     Elapsed time . . . . :   00:54:16

          -------Synchronous I/O--------  -------Asynchronous I/O-------
             DB     DB  Non-DB  Non-DB     DB     DB  Non-DB  Non-DB
Interval   Read  Write    Read   Write   Read  Write    Read   Write
14:01:48      0      0       0       0      0      0       0       0
14:06:47      0      0       0       0      0      0       0       0
14:11:47      0      0     293       0      0      0       0       0
14:16:46    103      0      35       0      0      0       0       0
14:21:46    112      0       4       0      0      0       0       0
14:26:45     67      0       0       0      0      0       0       0
14:31:45     28      0       0       0      0      0       0       0
                                                                   More...
Press Enter to continue.

F3=Exit   F11=View 3   F12=Cancel   F15=Sort by interval
F24=More keys
```

The columns in the second view of the Display Job Detail display are as follows:

Interval The end time for the collection interval.

Synchronous I/O - DB Read
The number of synchronous database read operations performed by the job during the collection interval.

Synchronous I/O - DB Write
The number of synchronous database write operations performed by the job during the collection interval.

Synchronous I/O - Non-DB Read
The number of synchronous nondatabase read operations performed by the job during the collection interval.

Synchronous I/O - Non-DB Write
The number of synchronous nondatabase write operations performed by the job during the collection interval.

Asynchronous I/O - DB Read
The number of asynchronous database read operations started by the job during the collection interval.

Asynchronous I/O - DB Write
The number of asynchronous database write operations started by the job during the collection interval.

Asynchronous I/O - Non-DB Read
The number of asynchronous nondatabase read operations started by the job during the collection interval.

Asynchronous I/O - Non-DB Write
The number of asynchronous nondatabase write operations started by the job during the collection interval.

Note: The asynchronous I/O operations are performed by system asynchronous I/O tasks.

If you press F11 (View 3) on the second view of the Display Job Detail display, the third view appears.

```
                              Display Job Detail

 Job . . . . . . . . :   DSP01          Job type . . . . . . :   INT
   User . . . . . . . :     QSECOFR     Subsystem  . . . . . :   QINTER
   Number . . . . . . :     022220      Pool . . . . . . . . :   04
 Member . . . . . . . :   Q892621356    Priority . . . . . . :   20
   Library  . . . . . :     QPFRDATA    Elapsed time . . . . :   00:54:16


           ---Cmn I/O----  -----Logical I/O------  ---Printer----
 Interval    Get     Put    Read   Write   Other   Pages   Lines
 14:01:48      0       0       0       0       0       0       0
 14:06:47      0       0       0       0       0       0       0
 14:11:47      0       0       0       0       0       0       0
 14:16:46      0       0       0    1327       5       0       0
 14:21:46      0       0       0    4078      28       0       0
 14:26:45      0       0       0     383      16       0       0
 14:31:45      0       0       0     133       6       0       0
                                                                    More...
 Press Enter to continue.

 F3=Exit   F11=View 1   F12=Cancel   F15=Sort by interval
 F24=More keys
```

The columns in the third view of the Display Job Detail display are as follows:

Interval The end time for the collection interval.

Cmn I/O—Get
The number of communications read operations performed by the job during the collection interval.

Cmn I/O—Put
The number of communications write operations performed by the job during the collection interval.

Logical I/O—Read
The number of logical database reads performed by the job during the collection interval.

Logical I/O—Write
The number of logical database writes performed by the job during the collection interval.

Logical I/O—Other
The number of other logical database I/O operations (updates, scans, and so on) performed by the job during the collection interval.

Printer Pages
The number of pages printed by the job during the collection interval.

Printer Lines
The number of lines printed by the job during the collection interval.

Display Performance Data for System Resources

When you are on the Display Performance Data display, you may want to view the performance data specifically related to storage pools, disk units, or communications lines, instead of the job-related information previously discussed. The third set of function keys, as shown below, allows you to do this.

```
                        •
                        •
                        •

F19=Display pool detail   F20=Display disk detail
F21=Display communications detail   F24=More keys
```

Display Pool Detail

If you press F19 on the Display Performance Data display, the Display Pool Detail display appears.

```
                           Display Pool Detail

Member . . . . . . . :   Q892621356      Elapsed time . . . . :     00:54:16
  Library  . . . . . :     QPFRDATA

Type options, press Enter.
  5=Display pool intervals

             Size   Act      Tns     Avg     DB        DB   Non-DB  Non-DB
Opt  Pool     (K)   Lvl     Count    Rsp   Faults    Pages  Faults   Pages
 _    01    22528    0          0     .0       .0       .0      .2     2.9
 _    02     4096   30        449    3.2       .6     23.3     3.8    25.3
 _    03     3072   50          0     .0       .3      3.7     4.6    32.6
 _    04    60416   72      10387    1.6      2.4    152.4    14.2    77.1
 _    05     8192   30          0     .0      1.0     54.1     2.1    32.3




                                                                        Bottom
F3=Exit   F11=Display transitions   F12=Cancel   F15=Sort by pool
F24=More keys
```

The Display Pool Detail display presents performance information for each pool in the measurement. Two views are used in order to present all the pool information.

The columns that appear on both views are as follows:

Pool The pool identifier.

Size (K) The size of the storage pool in kilobytes (KB). (KB equals 1024 bytes.) This value is retrieved from the performance data collected in the first interval.

Act Lvl The activity level for the pool. This represents the maximum number of jobs that can be active in the pool at any time.

Tns Count
The number of transactions performed by the jobs running in the pool.

Avg Rsp The average internal response time (in seconds) per transaction in the pool.

The columns that are unique to the first view, as shown above are as follows:

DB Faults The average number of database faults that occurred per second in the pool.

DB Pages The average number of database pages transferred into the pool per second.

Non-DB Faults
The average number of nondatabase faults that occurred per second in the pool.

Non-DB Pages
The average number of nondatabase pages transferred into the pool per second.

If you press F11 (Display transitions), the fault and page information is replaced with state transition data.

```
                              Display Pool Detail

Member . . . . . . . :    Q892621356      Elapsed time . . . . :    00:54:16
  Library  . . . . . :      QPFRDATA

Type options, press Enter.
  5=Display pool intervals

               Size  Act      Tns      Avg  Active->    Wait->     Active->
Opt  Pool       (K)  Lvl    Count      Rsp    Wait    Ineligible  Ineligible
 _    01      22528    0        0       .0    24.4         .0          .0
 _    02       4096   30      449      3.2     7.6         .0          .0
 _    03       3072   50        0       .0   144.5         .0          .0
 _    04      60416   72    10387      1.6   269.0         .0          .0
 _    05       8192   30        0       .0    29.9         .0          .0

                                                                        Bottom
F3=Exit   F11=Display faults and pages   F12=Cancel   F15=Sort by pool
F24=More keys
```

The columns that are unique to the second view are as follows:

Active → Wait
The average number of active-to-wait state transitions that occurred each minute in the pool.

Wait → Ineligible
The average number of wait-to-ineligible state transitions that occurred each minute in the pool.

Active → Ineligible
The average number of active-to-ineligible state transitions that occurred each minute in the pool.

Although the Display Pool Detail display presents the pool information as totals for the entire measurement, you may want to examine the data for a particular pool over time. Using the Display pool intervals option allows you to view the same pool information broken down into the time intervals in which it occurred.

Display Pool Interval

By typing a 5 in the *Option* field next to a pool and pressing the Enter key, the Display Pool Interval display appears with performance information for that pool.

```
                          Display Pool Interval

Pool . . . . . . . . :   04              Member . . . . . . . :   Q892621356
Elapsed time . . . . :   00:54:16          Library  . . . . . :     QPFRDATA

            Size  Act    Tns    Avg    DB       DB   Non-DB  Non-DB
Interval     (K)  Lvl  Count    Rsp  Faults   Pages  Faults   Pages
14:01:48   60416   72    895    1.7     2.7   135.0    12.7    56.8
14:06:47   60416   72    915    1.9     2.2   176.3    13.0    52.7
14:11:47   60416   72    936    1.9     2.4   119.9    14.7    60.7
14:16:46   60416   72    886    1.8     2.1   167.7    14.8    76.6
14:21:46   60416   72    811    2.1     2.6   237.6    13.2    54.5
14:26:45   60416   72    861    1.4     2.0   105.3    13.8    93.6
14:31:45   60416   72    794    1.3     2.1   125.2    12.7    53.6
14:36:44   60416   72    942    1.3     2.1   154.0    13.6    54.0
14:41:45   60416   72   1146    1.5     2.7   159.5    16.4    69.4
14:46:44   60416   72   1077    1.6     2.5   157.5    16.0    64.7
14:51:44   60416   72   1124    1.4     2.8   136.2    15.6   209.2
                                                                     Bottom
Press Enter to continue.

F3=Exit   F11=Display transitions   F12=Cancel   F15=Sort by interval
F24=More keys
```

The Display Pool Interval display presents the same columns of information as the Display Pool Detail display, except that the data is broken down by time intervals. A second view (not shown here) also exists for the Display Pool Interval display, which presents the data for the state transitions.

Display Disk Detail

If you press F20 on the Display Performance Data display, the Display Disk Detail display appears.

```
                              Display Disk Detail

Member . . . . . . . :   Q892621356      Elapsed time . . . . :    00:54:16
  Library  . . . . . :     QPFRDATA

Type options, press Enter.
  5=Display disk intervals

                                                 -----Activity Per Second------
              Size  ASP  CSS     %       %       Read     Read    Write    Write
Opt  Unit      (M)  ID   ID     Used    Busy     Rqs      (K)     Rqs      (K)
 _   0001      427  01   00     60.7     6.5     1.0      2.3      1.8      1.6
 _   0002      427  01   00     60.8     5.5      .8      2.2      1.7      1.1
 _   0003      427  02   00     68.5      .5      .0       .0       .2       .3
 _   0004      427  02   00     68.5      .4      .0       .0       .2       .2
 _   0005      427  01   00     60.8     6.8     1.2      4.6      1.8      1.3
 _   0006      427  01   00     60.8     6.9     1.3      4.4      1.7      1.2
 _   0007      427  01   00     60.8     6.8     1.2      4.0      1.9      1.4
 _   0008      427  02   00     68.5      .5      .0       .0       .2       .3
 _   0009      427  01   00     60.8     7.4     1.2      4.1      2.0      1.5
                                                                         More...
F3=Exit   F12=Cancel   F15=Sort by unit   F22=Sort by % used
F23=Sort by % busy
```

The Display Disk Detail display presents performance information for each disk unit attached to the system on which the data collection was performed. The columns on this display are as follows:

Unit The disk arm identifier (unit number). An 'A' or 'B' following the unit number indicates that the disk unit is mirrored. (For example, 0001A and 0001B are a mirrored pair.)

Size (M) The capacity of the disk in megabytes (MB). (1 MB equals 1048576 bytes.)

ASP ID The auxiliary storage pool to which the disk unit was allocated.

CSS ID The checksum set to which the disk unit was allocated.

% Used The percentage of the disk space used.

% Busy The percentage of the elapsed time the disk arm was busy servicing a disk read or write request.

Read Rqs The number of read requests performed each second by the disk arm.

Read (K) The number of kilobytes (KB) transferred each second by read operations (1 KB equals 1024 bytes)

Write Rqs The number of write requests performed each second by the disk arm.

Write (K) The number of kilobytes transferred each second by write operations.

Although the Display Disk Detail display presents the disk information as totals for the entire measurement, you may want to examine the data for a particular disk unit over time. Using the Display disk intervals option allows you to view the same disk information broken down into the time intervals in which it occurred.

Display Disk Interval

By typing a 5 in the *Option* field next to a disk unit and pressing the Enter key, the Display Disk Interval display appears with performance information for that disk unit.

```
                          Display Disk Interval

Unit . . . . . . . . :   0001            Member . . . . . . . :   Q892621356
Size (M) . . . . . . :   427               Library  . . . . . :     QPFRDATA
Elapsed time . . . . :   00:54:16

                                         -------Activity Per Second-------
           ASP   CSS     %       %        Read     Read     Write     Write
Interval   ID    ID     Used    Busy      Rqs      (K)      Rqs       (K)
14:01:48   01    00     60.6     7.8      1.4      2.6      2.1       1.8
14:06:47   01    00     60.6     5.7       .8      1.4      1.5       1.3
14:11:47   01    00     60.6     4.7       .7      1.8      1.9       1.3
14:16:46   01    00     60.6     4.5       .9      2.7      1.2        .9
14:21:46   01    00     60.6     5.3       .7      1.6      1.4       1.1
14:26:45   01    00     60.7     4.0       .8      1.7       .7        .7
14:31:45   01    00     60.7     5.1       .9      2.0      1.4       1.3
14:36:44   01    00     60.7     5.3       .7      2.0      1.6       1.8
14:41:44   01    00     60.7    12.8      1.5      4.2      3.3       3.1
                                                                        More...
Press Enter to continue.

F3=Exit   F12=Cancel   F15=Sort by interval   F22=Sort by % used
F23=Sort by % busy
```

The Display Disk Interval display presents the same columns of information as the Display Disk Detail display, except that the data is broken down by time intervals.

Note: The *Size (M)* column has been moved to the *Size (M)* field at the top of the display since the size of the disk unit cannot change from one interval to the next.

Display Communications Line Detail

If you press F21 on the Display Performance Data display, the Display Communications Line Detail display appears.

Note: The data in the following figures are examples only and should not be used to tune any system.

```
                    Display Communications Line Detail

Member . . . . . . . :   Q892621356      Elapsed time . . . . :    00:54:16
  Library  . . . . . :     QPFRDATA

Type options, press Enter.
  5=Display remote jobs   7=Display communications interval data

         Line          Line        Line        Tns     Average     Job       %
Option   ID            Type        Speed      Count    Response    Count    Busy
  _      BSC41         BSC           2.4          0        .00        0       .0
  _      BSC20         BSC           2.4          0        .00        0       .0
  _      SDLCLINEA1    SDLC          9.6       1301       1.80       56     52.0
  _      SDLC100       SDLC          9.6       1047       1.22       71     30.6
  _      SDLCB         SDLC          9.6        558       1.16       28     20.3
  _      SDLCLINE02    SDLC          9.6       1168       2.85       45     44.8
  _      SDLCLINE03    SDLC          9.6       1024       1.34       34     30.3
  _      SDLC1500      SDLC          9.6       1556       1.74       45     58.7
  _      SDLC1000      SDLC          9.6       1317       2.35       96     50.4

                                                                        Bottom
F3=Exit   F12=Cancel   F15=Sort by line ID   F20=Sort by transactions
F24=More keys
```

Figure 6-1. Display Communications Line Detail

The Display Communications Line Detail display presents performance information for each communications line attached to the system. The columns on this display are as follows:

Line ID The communications line identifier.

Line Type The type of the communications protocol used on the line. Possible values are as follows:

- **ASYNC** Asynchronous communications
- **BSC** Binary synchronous communication
- **ELAN** Ethernet local area network
- **SDLC** Synchronous data link control
- **TRLAN** Token-Ring local area network
- **X.25** X.25

Line Speed
The speed of the line in kilobits (1000) per second.

Tns Count
The number of transactions that occurred over the line.

Average Response
The average internal response time (in seconds) of jobs attached to this line not including line time.

Job Count
The number of active jobs running on the line. This count includes only those jobs which used any CPU resource or performed any I/O activity.

% Busy The percentage of the elapsed time the line was busy transferring information.

Note: The *% Busy* field is calculated using the line speed listed in the line description. Entering incorrect values in the line description will adversely affect this field.

The Display Communications Line Detail display presents the totals for each line in the measurement. One of the options on this display lets you view performance data for the jobs using a communications line. The other option displays the time interval performance data for a communications line.

Display Remote Jobs

If you type a 5 (Display remote jobs) on the Display Communications Line Detail display in the *Option* field next to a communications line and press the Enter key, the Display Remote Jobs display appears with the performance information for that line listed by job.

```
                              Display Remote Jobs

 Line . . . . . . . . :   SDLCLINEA1      Member . . . . . . . :   Q892621356
 Line type  . . . . . :   SDLC              Library  . . . . . :     QPFRDATA
 Line speed . . . . . :   9.6             Elapsed time . . . . :   00:54:16

 Type options, press Enter.
   5=Display job detail

                                           Job      CPU       Tns     Avg     Disk
 Option  Job         User          Number  Type     Util      Count   Rsp      I/O
   _     DSP92       QPGMR         022207  INT       .55        163   2.3      706
   _     DSP79       QPGMR         022191  INT       .52         71   2.6      734
   _     DSP13       QPGMR         022075  INT       .50        104   2.2      776
   _     DSP41       QSECOFR       022123  INT       .36        108   1.6      704
   _     DSP63       QPGMR         022203  INT       .31         46   2.2      785
   _     DSP85       QSECOFR       022193  INT       .26         58   2.3      527
   _     DSP08       QSYSOPR       022120  INT       .21         53   2.1      397
   _     DSP54       QPGMR         022175  INT       .20         51   1.7      275
   _     DSP50       QPGMR         022172  INT       .17         41   1.9      261
                                                                             More...
 F3=Exit           F12=Cancel    F15=Sort by job    F16=Sort by job type
 F19=Sort by CPU   F24=More keys
```

Figure 6-2. Display Remote Jobs

The line identifier, the type of communications protocol, and the line speed for the selected communications line are shown in the fields *Line, Line type*, and *Line speed* at the top of the display.

The columns for the Display Remote Jobs display are as follows:

Job The remote job name for the performance data being displayed.

User The user profile associated with the job.

Number The number assigned to the job.

Job Type The type of the job. Refer to the section describing the Display Jobs display for a list of possible values.

CPU Util The percentage of the elapsed time during which the processing unit was used by the job. For multiple-processor systems, this is the average use across all processors.

Tns Count
The number of transactions performed by the job.

Avg Rsp The average internal response time (in seconds) per transaction for the job.

Disk I/O The number of physical disk I/O operations performed by the job.

If you type a 5 in the *Option* column, you can display more detailed information for the remote job. This option calls the Display Job Detail display, just as option 5 did from the Display Jobs display. Refer to "Display Job Detail" on page 6-9 for information on the performance data that will be shown.

Display Communications Interval Data

To see a display of performance data for a communications line by time interval, type a 7 (Display communications interval data) in the *Option* field next to the communications line on the Display Communications Line Detail display, and press the Enter key. The resulting Display Communications Interval Data display lists the performance averages and totals for that communications line for the time intervals in the current performance data member.

From the Display Communications Interval Data display you can request data about the jobs using the communications line during any of the listed time intervals. To do this, type a 5 in the *Option* column by the selected time interval.

Each communications protocol has its own type of Display Communications Interval Data display, but all are quite similar. An example and description of this display for each protocol follow.

The display for synchronous data link control (SDLC) follows:

```
                     Display Communications Interval Data

Line ID . . . . . . . . . :   MPLSCHI      Member . . . . . . . . . . : MONDAY
Line type . . . . . . . . :   SDLC           Library . . . . . . . . :   QPFRDATA
Line speed . . . . . . . . :  19.2         Elapsed time . . . . . . . : 00:24:50
Bus number . . . . . . . . :  0
IOP address . . . . . . . :   05

Type options, press Enter,
  5=Display remote jobs

                                           Pct                  Pct
                                         I Frames             Frames
            Itv       Line   I Frames    Trnsmitd   Frames     Recd
Option      End       Util   Trnsmitd    in Error    Recd    in Error
   5     13:08:00      78      1,818          09     1,818         02
         13:23:00      78      1,818          07     1,818         00
         13:38:00      78      1,818          00     1,818         00
         13:53:00      78      1,818          00     1,818         00
         14:08:00      78      1,818          00     1,818         00
                                                                         Bottom
F3=Exit    F11=View 2     F12=Cancel    F15=Sort by itv end
F20=Sort by line util     F24=More keys
```

Figure 6-3. Display Communications Interval Data for SDLC

The following columns are shown when you press F11 (View 2):

```
                          Pct       Congestion
                          Poll     Local   Remote
          Itv      Line   Retry     Not     Not
Option    End      Util   Time     Ready   Ready
  5     13:08:00    78     4.6       09      12
        13:23:00    78     4.6       08      05
        13:38:00    78     4.6       04      00
        13:53:00    78     4.6       00      00
        14:08:00    78     4.6       00      00
                                                                Bottom
F3=Exit   F11=View 1    F12=Cancel   F15=Sort by itv end
F20=Sort by line util   F24=More keys
```

Figure 6-4. Display Communications Interval Data for SDLC - View 2

The line identifier, the type of communications protocol, the line speed, the bus number, and the IOP address for the selected communications line are shown in the fields *Line ID, Line type, Line speed, Bus number*, and *IOP address* at the top of each display.

The data columns for the two SDLC displays are as follows:

Itv End The interval end time (hour and minute).

Line Util The percentage of the elapsed time the line was busy transferring information.

I Frames Trnsmitd
The total number of information frames transmitted.

Percent I Frames Trnsmitd in Error
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).

Frames Recd
The number of frames received, including frames with errors and frames that are not valid.

Percent Frames Recd in Error
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).

Pct Poll Retry Time
The percent of the time interval the line was unavailable while the host system IOP waited for a response from a remote work station controller (or a remote AS/400 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).

Local Not Ready
The percent of all frames received by the host that resulted in receive-not-ready frames transmitted by the host. A large percent often means that the host cannot process data fast enough (congestion).

Remote Not Ready
The percent of all frames transmitted by the host that resulted in receive-not-ready frames transmitted by a remote device. A large percent often means that the remote device cannot process data fast enough (congestion).

The display for X.25 follows:

```
                    Display Communications Interval Data

Line ID . . . . . . . . :    NETWORK      Member . . . . . . . . . . : MONDAY
Line type . . . . . . . :    X.25           Library . . . . . . . . :   QPFRDATA
Line speed . . . . . . . :   19.2         Elapsed time . . . . . . : 00:24:50
Bus number . . . . . . . :   0
IOP address  . . . . . . :   05

Type options, press Enter,
  5=Display remote jobs

                                           Pct                  Pct
                                         I Frames              Frames
           Itv        Line   I Frames    Trnsmitd     Frames     Recd
Option     End        Util   Trnsmitd    in Error      Recd    in Error
   5     13:08:00      78      1,818        08         1,818       08
         13:23:00      68      1,818        06         1,818       00
         13:38:00      58      1,818        03         1,818       00
         13:53:00      48      1,818        00         1,818       00
         14:08:00      38      1,818        00         1,818       00
                                                                          Bottom
F3=Exit    F11=View 2     F12=Cancel    F15=Sort by itv end
F20=Sort by line util     F24=More keys
```

Figure 6-5. Display Communications Interval Data for X.25

The following columns display when you press F11 (View 2):

```
                                Congestion
                              Local    Remote     Reset     Reset
           Itv        Line     Not       Not     Packets   Packets
Option     End        Util    Ready     Ready    Trnsmitd    Recd
   5     13:08:00      78      09         12        12        12
         13:23:00      68      08         05        05        05
         13:38:00      58      04         00        00        00
         13:53:00      48      00         00        00        00
         14:08:00      38      00         00        00        00
                                                                          Bottom
F3=Exit    F11=View 1     F12=Cancel    F15=Sort by itv end
F20=Sort by line util     F24=More keys
```

Figure 6-6. Display Communications Interval Data for X.25—View 2

The line identifier, the type of communications protocol, the line speed, the bus number, and the IOP address for the selected communications line are shown in the fields *Line ID*, *Line type*, *Line speed*, *Bus number*, and *IOP address* at the top of each display.

The following describes the data columns on the two X.25 displays that are different from those described for Figure 6-3 on page 6-20 and Figure 6-4 on page 6-21:

Reset Packets Trnsmitd
The number of reset packets transmitted by the network. **Reset packets** are packets retransmitted because an error occurred.

Reset Packets Recd
The number of reset packets received by the network.

The display for token-ring local area network (TRLAN) follows:

```
                     Display Communications Interval Data

Line ID . . . . . . . . :   LOCALNET1   Member . . . . . . . . . . : MONDAY
Line type . . . . . . . :   TRLAN         Library . . . . . . . :    QPFRDATA
Line speed . . . . . . . :      4000.0  Elapsed time . . . . . . : 00:24:50
Bus number . . . . . . . :  0
IOP address  . . . . . . :  05

Type options, press Enter,
  5=Display remote jobs

                                                  -------- Congestion ---------
                            I Frames   I Frames   Local   Local  Remote  Remote
            Itv      Line   Trnsmitd     Recd      Not     Seq     Not     Seq
Option      End      Util    Per Sec    Per Sec   Ready   Error   Ready   Error
   5     13:08:00     78        842        333      10      08      22      06
         13:23:00     68        742        333      10      08      22      05
         13:38:00     58        642        333      10      08      22      04
         13:53:00     48        542        333      10      08      22      00

                                                                         Bottom
F3=Exit     F11=View 2    F12=Cancel    F15=Sort by itv end
F20=Sort by line util    F24=More keys
```

Figure 6-7. Display Communications Interval Data for TRLAN

The following columns are shown when you press F11 (View 2):

```
                                   Response    ----- Remote LAN -----
            Itv      Line   Frame    Timer     --Frames---    -Frames-     MAC
Option      End      Util   Retry    Ended     Transmitted    Received   Errors
   5     13:08:00     78       4        10              18          17      382
         13:23:00     68       4        09              18          17       20
         13:38:00     58       4        05              18          17      234
         13:53:00     48       4        00              18          17        4

                                                                         Bottom
F3=Exit     F11=View 1    F12=Cancel    F15=Sort by itv end
F20=Sort by line util    F24=More keys
```

Figure 6-8. Display Communications Interval Data for TRLAN - View 2

The line identifier, the type of communications protocol, the line speed, the bus number, and the IOP address for the selected communications line are shown in the fields *Line ID, Line type*, *Line speed*, *Bus number*, and *IOP address* at the top of each display.

The following describes the data columns on the two TRLAN displays that are different from those described for Figure 6-3 on page 6-20 and Figure 6-4 on page 6-21:

I Frames Trnsmitd Per Sec
The number of information frames transmitted per second.

I Frames Recd Per Sec
The number of information frames received per second.

Local Seq Error
The percent of frames received out of order by the system. This can occur when the host system cannot process data fast enough.

Remote Seq Error
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.

Frame Retry
The number of attempts to retransmit a frame to a remote controller.

Response Timer Ended
The number of times the response timer ended waiting for a response from a remote device.

Remote LAN Frames Transmitted
The number of frames transmitted to a LAN connected to the locally attached LAN.

Remote LAN Frames Received
The number of frames received from a LAN connected to the locally attached LAN.

MAC Errors
The number of medium access control (MAC) errors. These can include: internal adapter errors, token errors, a frame discarded, and so on.

The display for Ethernet local area network (ELAN) follows:

```
                    Display Communications Interval Data

Line ID . . . . . . . . :   LOCALNET2   Member . . . . . . . . . . : MONDAY
Line type . . . . . . . :   ELAN          Library . . . . . . . :   QPFRDATA
Line speed . . . . . . . :       4000.0  Elapsed time . . . . . . : 00:24:50
Bus number . . . . . . . :   0
IOP address . . . . . . :   05

Type options, press Enter,
  5=Display remote jobs

                                                  -------- Congestion ---------
                             I Frames   I Frames  --- Local ---   --- Remote --
          Itv       Line  Trnsmitd     Recd      Not     Seq      Not      Seq
Option    End       Util   Per Sec   Per Sec    Ready   Error    Ready    Error
  5    13:08:00      78       842        333      10      08       22       06
       13:23:00      68       742        333      10      08       22       05
       13:38:00      58       642        333      10      08       22       04
       13:53:00      48       542        333      10      08       22       00

                                                                        Bottom
F3=Exit    F11=View 2   F12=Cancel   F15=Sort by itv end
F20=Sort by line util   F24=More keys
```

Figure 6-9. Display Communications Interval Data for ELAN

The following columns are shown when you press F11 (View 2):

```
                                    Response
          Itv         Line   Frame   Timer      MAC
Option    End         Util   Retry   Ended      Errors
  5     13:08:00       78      4        10        128
        13:23:00       68      4        09        337
        13:38:00       58      4        05          5
        13:53:00       48      4        00         32

                                                                        Bottom
F3=Exit    F11=View 1   F12=Cancel   F15=Sort by itv end
F20=Sort by line util   F24=More keys
```

Figure 6-10. Display Communications Interval Data for ELAN—View 2

The line identifier, the type of communications protocol, the line speed, the bus number, and the IOP address for the selected communications line are shown in the fields *Line ID*, *Line type*, *Line speed*, *Bus number*, and *IOP address* at the top of each display.

The columns in this section are the same as described for Figure 6-3 on page 6-20, Figure 6-4 on page 6-21, Figure 6-7 on page 6-23, and Figure 6-8 on page 6-23.

The display for binary synchronous communications (BSC) follows:

```
                    Display Communications Interval Data

Line ID . . . . . . . . :   PNTTOPNT    Member . . . . . . . . . . : MONDAY
Line type . . . . . . . :   BSC           Library . . . . . . . . :   QPFRDATA
Line speed . . . . . . . :    9.6       Elapsed time . . . . . . : 00:24:50
Bus number . . . . . . . :   0
IOP address . . . . . . :   05

Type options, press Enter,
  5=Display remote jobs

                               Total     Pct Data     Total     Pct Data
                               Data       Chars       Data       Chars
            Itv        Line    Chars     Trnsmitd     Chars      Recd       Line
Option      End        Util   Trnsmitd   in Error   Received   in Error    Errors
  5      13:08:00       32    2,341,445        18     34,211         02       383
         13:23:00       32    2,341,445        08     34,211         01       121




                                                                          Bottom
F3=Exit     F11=View 2    F12=Cancel    F15=Sort by itv end
F20=Sort by line util     F24=More keys
```

Figure 6-11. Display Communications Interval Data for BSC

The line identifier, the type of communications protocol, the line speed, the bus number, and the IOP address for the selected communications line are shown in the fields *Line ID, Line type*, *Line speed*, *Bus number*, and *IOP address* at the top of the display.

The following describes the data columns on the BSC display that are different from those described for Figure 6-3 on page 6-20 and Figure 6-4 on page 6-21:

Total Data Chars Trnsmitd
The number of data characters transmitted successfully.

Pct Data Chars Trnsmitd in Error
The percent of data characters transmitted with error.

Total Data Chars Received
The number of data characters received successfully.

Pct Data Chars Recd in Error
The percent of data characters received with error.

Line Errors
The total of all line errors. The errors counted are:

- Negative acknowledgement character received to text sent
- Wrong acknowledgement character to text sent
- Enqueue to text sent
- Invalid format (unrecognized) sent
- Enqueue to acknowledged character sent

Note: Check the line if the number of errors increases greatly over time.

The display for asynchronous data link control (ASYNC) follows:

```
                    Display Communications Interval Data

Line ID . . . . . . . . :   STARTSTOP   Member . . . . . . . . . . : MONDAY
Line type . . . . . . . :   ASYNC         Library . . . . . . . :    QPFRDATA
Line speed . . . . . . . :   0.3         Elapsed time . . . . . . : 00:24:50
Bus number . . . . . . . :   0
IOP address . . . . . . :   05

Type options, press Enter,
  5=Display remote jobs

                                 Bytes        Bytes                   Pct PDUs
           Itv        Line    Transmitted    Received       PDUs      Received
Option     End        Util      Per Sec      Per Sec      Received    in Error
   5     13:08:00      18           211           88           34          08
         13:23:00      18           121           80           34          02




                                                                           Bottom
F3=Exit    F12=Cancel    F15=Sort by itv end    F20=Sort by line util
F24=More keys
```

Figure 6-12. Display Communications Interval Data for ASYNC

The line identifier, the type of communications protocol, the line speed, the bus number, and the IOP address for the selected communications line are shown in the fields *Line ID, Line type*, *Line speed*, *Bus number*, and *IOP address* at the top of the display.

The following describes the data columns on the ASYNC display that are different from those described for Figure 6-3 on page 6-20 and Figure 6-4 on page 6-21:

PDUs Received

The number of protocol data units (PDUs) received during the time interval.

Pct PDUs Received in Error

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.

Display Remote Interval Jobs

This display lists information about the jobs using a communications line during a time interval. To request it, type a 5 (Display remote jobs) in the *Option* column by a time interval on a Display Communications Interval Data display, and press the Enter key. The following display appears:

```
                     Display Interval Remote Jobs

Interval . . . . . . . . :     13:08     Member . . . . . . . . . . : MONDAY
Line ID  . . . . . . . . :   MPLSCHI       Library  . . . . . . . :   QPFRDATA
Line type  . . . . . . . :      SDLC     Elapsed time . . . . . . . : 00:24:50
Line speed . . . . . . . :      19.2
Line utilization . . . . :       78%

                                    Job     CPU     Tns     Avg    Disk
Job          User          Number   Type    Util    Count   Rsp     I/O
DSP15        X07733        030191   DDM     .16       19     .0     230
DSP40        SMITH         030275   INT     .24      240    3.5    1598
DSP43        U5531         030212   DDM     .00        0     .0      76









                                                                   Bottom
F3=Exit    F12=Cancel    F15=Sort by job    F16=Sort by job type
F19=Sort by CPU    F24=More keys
```

The end time for the selected time interval, the line name, line type, line speed, and average use during the time interval are shown in the fields *Interval*, *Line ID*, *Line type*, *Line speed* and *Line utilization* at the top of this display. The column descriptions are the same as for Figure 6-2 on page 6-19.

Chapter 7. System Activity

This chapter describes the functions that allow you to work with performance data for the jobs and licensed internal code (LIC) tasks currently running on the system. These functions provide the ability to interactively view and collect the data in a QAITMON database file using the Work with System Activity (WRKSYSACT) command and to print reports based on the collected data (print activity report). These functions are available as AS/400 commands or through option 8 (System activity) on the IBM Performance Tools/400 menu. If you select option 8, the System Activity menu appears.

```
                              System Activity

 Select one of the following:

      1. Work with system activity
      2. Print activity report
                                      •
                                      •
                                      •
```

Refer to "Work with System Activity" and "Print Activity Report" on page 7-11 for a description of both selections shown on the System Activity menu.

Work with System Activity

The Work with System Activity (WRKSYSACT) command allows you to view and collect performance data in a realtime fashion. This data, which consists of CPU utilizations, synchronous and asynchronous I/O counts, and more, is reported for any job or task that is currently active on the system.

Note: To be considered active, a job or task must use at least one-tenth of 1% (.1%) of the processing unit or perform one I/O operation.

Similar to the performance monitor started by the Start Performance Monitor (STRPFRMON) command, only one call of the Work with System Activity function can be active at one time. If this function is currently active when the WRKSYSACT command is run, you receive a message indicating that the function is already active.

The performance statistics reported by this function represent activity that has occurred during the elapsed time since a previous collection. Notice that this may contrast with other system functions that generally provide cumulative values until specifically reset. In most cases the time interval between data collections ranges from 1 second to several minutes, depending on how often you want to view or collect new data. On systems with very little activity, a subsecond refresh interval may be possible.

Depending on the value specified for the OUTPUT parameter of the Work with System Activity command, the data gathered by this function is processed in one of the following methods:

- Shown on the display station only
- Written to a database file only
- Shown on the display station and written to a database file

Refer to the *CL Reference* manual for further information on this parameter.

When the data is written to a database file only, this function submits a batch job under the name WRKSYSACT When the data is written to a file and shown on the display, the statistics are put in the file each time the display is refreshed. This does not include the data presented on the initial display of the Work with System Activity display. See the file descriptions in Figure 7-1 on page 7-10.

When the data is being directed to the display station (either the first or third method), the first display that appears resembles the following:

```
                         Work with System Activity

 Automatic refresh in seconds  . . . . . . . . . . . . . . . . .      5
 Elapsed time . . . . :   00:00:02       Overall CPU util . . :    74.1

 Type options, press Enter.
   1=Monitor job   5=Work with job
                                                    Total  Total
     Job or                                   CPU   Sync   Async    PAG      EAO
 Opt Task        User        Number   Pty   Util   I/O     I/O    Fault    Excp
  _  DSP25       QPGMR       008124    20   56.1    12      4       0        2
  _  DSP01       QSECOFR     008122    20    6.4     5      2       0        3
  _  DSP07       QSYSOPR     008123    20    4.4     1      0       0        0
  _  DSP04       QSECOFR     008119     1    3.7     0      0       0        0
  _  #A000A                             0    1.0     0      0       0        0
  _  #50009                             0     .6     0      0       0        0
  _  SMAI03                             0     .3     0      0       0        0
  _  ROUT04                             0     .3     0      0       0        0
  _  SMAI01                             0     .3     0      0       0        0
  _  #A0003                             0     .3     0      0       0        0
                                                                          More...
 F3=Exit   F10=Update list    F11=View 2   F12=Cancel
 F19=Automatic refresh        F24=More keys
```

The input-capable field *Automatic refresh in seconds* at the top of the display controls the amount of time between display refreshes when the automatic refresh feature is active. Refer to "Automatic Refresh Mode" on page 7-6 for more information on this field. The second field at the top of the display, *Elapsed time*, reflects the length of time in which the currently shown performance statistics occurred. Described in a different way, this value represents the time between the last display refresh and the next-to-last display refresh.

Note: Because the statistics shown by this function are deltas, Work with System Activity automatically gathers the data twice before displaying the first display. Therefore, the initial *Elapsed time* should be approximately 2 seconds, which means that the statistics shown occurred in the 2 seconds previous to the current display.

Finally, the *Overall CPU util* represents the CPU utilization for the entire system during the elapsed time. This value does not always equal the sum of the individual CPU utilizations shown in the list, since a job or task could use an extremely small amount of processing unit time, thus affecting the overall utilization, but not use enough CPU resource to be included in the list of active jobs. (Refer to the requirements for being considered active at the beginning of this chapter.) The discrepancy in CPU utilizations, however, is small and should have little effect on the usability of this function.

Also, the *Overall CPU util* could exceed 100% on extremely busy systems, because the data collection process does not occur instantaneously. However, you should be aware that overall CPU utilizations slightly over 100% are an acceptable possibility.

For a multiple-processor system, *Overall CPU util* is replaced by a utilization value for each processor in the system. Here is an example of this part of the display for a dyadic system (2 processors in the system):

```
                          Work with System Activity

 Automatic refresh in seconds . . . . . . . . . . . . . . . . .      5
 Elapsed time . . . . . :   00:00:02
 CPU 1 util . . . . . . :     36.9       CPU 2 util  . . . . . :   37.2

 Type options, press Enter.
   1=Monitor job   5=Work with job
```

The sum of the CPU utilization numbers might not equal the sum of the individual utilizations shown in the list, for the same reasons given for the *Overall CPU util* value for a single-processor system.

The two options shown on the Work with System Activity allow you to analyze specific jobs and tasks that appear in the list. Refer to "Monitoring Specific Jobs" on page 7-7 and "Working with Jobs" on page 7-8 for more information on these options.

The jobs and tasks are presented on this display in decreasing order of either CPU utilization or total number of synchronous and asynchronous I/O operations performed. This order is initially controlled by the Sequence (SEQ) parameter on the Work with System Activity command. The default is to sort the jobs and tasks by CPU utilization. Once the function has been started, however, F16 (Sequence by I/O, Sequence by CPU) serves as a switch between the two methods.

The Work with System Activity function uses three different views in order to present all the performance statistics. Pressing F11 shows you the next view in the series and pressing F10 shows you the previous view. The columns in the first view of the Work with System Activity list are:

Job or Task
: The job or task name for the performance statistics being displayed. This column is shown only when INFTYPE(*ALL) is specified. See the note below regarding this column.

User
: The user profile associated with the job. This column is not shown when INFTYPE(*TASK) is specified.

Number
: The number assigned to the job. This column is not shown when INFTYPE(*TASK) is specified.

Pty The priority at which the job or task was running when the performance statistics were collected.

Note: You may notice a low-priority job having its priority temporarily raised to a higher priority. This occurs when the low priority job has a seize/lock on an object to which a higher-priority job needs access. The lower job priority is temporarily increased to release the seize/lock for the other job sooner. The original priority is automatically restored when seize/lock is released.

CPU Util The percentage of the elapsed time during which the processing unit was used by the job or task. For multiple-processor systems, this is the average use across all processors.

Total Sync I/O
The total number of synchronous physical disk I/O operations performed by the job or task during the elapsed time. This value is the sum of the synchronous database/nondatabase reads and writes shown in View 2.

Total Async I/O
The total number of asynchronous physical disk I/O operations started by the job or task during the elapsed time. This value is the sum of the asynchronous database/nondatabase reads and writes shown in View 3.

PAG Fault
The number of process access group (PAG) faults caused by the job or task during the elapsed time.

EAO Excp The number of effective address overflow (EAO) exceptions caused by the job or task during the elapsed time.

Note: As mentioned above, the *Job or Task* column is shown only when INFTYPE(*ALL) is specified. This value for the *Information type* parameter instructs the function to display both jobs and tasks. Specifying INFTYPE(*JOB) causes the *Job or Task* column, to be replaced by the *Job* column because only jobs are to be displayed. Similarly, specifying INFTYPE(*TASK) causes the *Job or Task* column to be replaced by the *Task* column since only tasks are to be displayed. Later sections of this chapter describes how to switch between these information types through the use of function keys.

If you press F11 (View 2) on the first view of the Work with System Activity display, the second view appears.

```
                        Work with System Activity

Automatic refresh in seconds  . . . . . . . . . . . . . . . . .     5
Elapsed time . . . . :    00:00:02       Overall CPU util . . :    74.1

Type options, press Enter.
  1=Monitor job    5=Work with job
                                                     --------Synchronous---------
      Job or                                    CPU     DB      DB   Non-DB  Non-DB
Opt   Task        User        Number   Pty   Util   Read   Write    Read   Write
 _    DSP25       QPGMR       008124    20   56.1      2       5       3       2
 _    DSP01       QSECOFR     008122    20    6.4      3       0       2       0
 _    DSP07       QSYSOPR     008123    20    4.4      0       0       1       0
 _    DSP04       QSECOFR     008119     1    3.7      0       0       0       0
 _    #A000A                             0    1.0      0       0       0       0
 _    #50009                             0     .6      0       0       0       0
 _    SMAI03                             0     .3      0       0       0       0
 _    ROUT04                             0     .3      0       0       0       0
 _    SMAI01                             0     .3      0       0       0       0
 _    #A0003                             0     .3      0       0       0       0
                                                                          More...
F3=Exit   F10=Update list    F11=View 3    F12=Cancel
F19=Automatic refresh        F24=More keys
```

The new columns in the second view of the Work with System Activity display are as follows:

Synchronous DB Read
: The number of synchronous database read operations performed by the job or task during the elapsed time.

Synchronous DB Write
: The number of synchronous database write operations performed by the job or task during the elapsed time.

Synchronous Non-DB Read
: The number of synchronous nondatabase read operations performed by the job or task during the elapsed time.

Synchronous Non-DB Write
: The number of synchronous nondatabase write operations performed by the job or task during the elapsed time.

If you press F11 (View 3) on the second view of the Work with System Activity display, the third view appears.

```
                     Work with System Activity

Automatic refresh in seconds  . . . . . . . . . . . . . . . . .     5
Elapsed time . . . . :   00:00:02       Overall CPU util . . :    74.1

Type options, press Enter.
  1=Monitor job   5=Work with job
                                                  --------Asynchronous--------
     Job or                                 CPU     DB     DB   Non-DB  Non-DB
Opt  Task        User        Number  Pty   Util   Read  Write   Read    Write
 _   DSP25       QPGMR       008124   20   56.1      0      4      0        0
 _   DSP01       QSECOFR     008122   20    6.4      2      0      0        0
 _   DSP07       QSYSOPR     008123   20    4.4      0      0      0        0
 _   DSP04       QSECOFR     008119    1    3.7      0      0      0        0
 _   #A000A                            0    1.0      0      0      0        0
 _   #50009                            0     .6      0      0      0        0
 _   SMAI03                            0     .3      0      0      0        0
 _   ROUT04                            0     .3      0      0      0        0
 _   SMAI01                            0     .3      0      0      0        0
 _   #A0003                            0     .3      0      0      0        0
                                                                       More...
F3=Exit   F10=Update list    F11=View 1   F12=Cancel
F19=Automatic refresh        F24=More keys
```

The new columns in the third view of the Work with System Activity display are as follows:

Asynchronous DB Read
: The number of asynchronous database read operations started by the job or task during the elapsed time.

Asynchronous DB Write
: The number of asynchronous database write operations started by the job or task during the elapsed time.

Asynchronous Non-DB Read
: The number of asynchronous nondatabase read operations started by the job or task during the elapsed time.

Asynchronous Non-DB Write
: The number of asynchronous nondatabase write operations started by the job or task during the elapsed time.

Note: The asynchronous I/O operations are performed by system asynchronous I/O tasks.

Automatic Refresh Mode

Automatic Refresh mode represents an important feature of the Work with System Activity function. Once started, this mode continually updates the display without requiring further user intervention.

To start the Automatic Refresh mode, first enter the desired number of seconds between refreshes in the *Automatic refresh in seconds* field. This value, which has an initial default of 5 seconds, can range from a minimum of 1 second to a maximum of 900 seconds (15 minutes).

Note: Setting the *Automatic refresh seconds* at 5 or greater generally results in the Work with System Activity function using reasonably small amounts of the processing unit, depending on the size of the system being monitored. Setting this value lower than 5 seconds causes this function to use larger amounts of the processing unit, and therefore, is not recommended.

Once you have established the desired refresh interval, pressing F19 (Automatic refresh) starts the automatic refresh mode. Automatic refresh continues to display the same view and type of information that was previously selected. For example, if you had been examining both jobs and tasks using View 1, the display appears as follows.

Note: For a multiple-processor system, the use of each processor in the system is shown instead of the *Overall CPU util*.

```
                         Work with System Activity

Automatic refresh in seconds . . . . . . . . . . . . . . . . :      5
Elapsed time . . . . :    00:00:05        Overall CPU util . . :    19.4


                                                  Total  Total
    Job or                                 CPU    Sync   Async    PAG     EAO
    Task        User        Number  Pty   Util    I/O     I/O    Fault   Excp
    DSP25       QPGMR       008124   20   18.2     16       5       0       4
    DSP04       QSECOFR     008119    1     .9      0       0       0       0
    SMAI03                            0     .2      0       0       0       0






Attn=Terminate automatic refresh
Or press Sys Req and option 2 to exit automatic refresh and WRKSYSACT
```

The Automatic Refresh function attempts to maintain even refresh intervals by compensating for the time required to process, display, and, possibly, write the performance data. Therefore, you may occasionally notice that the elapsed time does not exactly match the value specified for the *Automatic refresh in seconds* field.

To end the Automatic Refresh mode, you must press the Attention (Attn) key to return to the normal operating mode. If you are in Automatic Refresh mode and the Attn key cannot end the function, use the System Request function with option 2 to end the Automatic Refresh and the Work with System Activity functions.

Due to the nature of the Automatic Refresh function, all keys other than the Attn key (and System Request) are disabled while this function is active. Therefore, you must select the desired view, type of information, and sequence before starting the Automatic Refresh function.

Monitoring Specific Jobs

While using the Work with System Activity function, you may wish to view the performance statistics for a set of jobs and tasks on the system. By typing a 1 in the *Opt* column before a list entry, that job or task is selected for monitoring. You may monitor as many as 20 jobs and tasks at a single time. If five jobs and tasks are selected for monitoring, the Work with System Activity display appears as follows:

```
                         Work with System Activity

 Automatic refresh in seconds . . . . . . . . . . . . . . . . .     5
 Elapsed time . . . . :   00:00:02      Overall CPU util . . :    74.1

 Type options, press Enter.
   1=Monitor job   5=Work with job
                                                  Total  Total
      Job or                                CPU   Sync  Async   PAG    EAO
 Opt  Task       User      Number  Pty  Util   I/O    I/O  Fault   Excp
  1   DSP25      QPGMR     008124   20  56.1    12      4      0      2
  1   DSP01      QSECOFR   008122   20   6.4     5      2      0      3
  1   DSP04      QSECOFR   008119    1   3.7     0      0      0      0
  1   SMAI03                         0    .3     0      0      0      0
  1   SMAI01                         0    .3     0      0      0      0







                                                                     Bottom
 F13=Display jobs and tasks   F14=Display jobs only   F15=Display tasks only
 F16=Sequence by I/O          F24=More keys
```

Once you have selected jobs and tasks for monitoring, the Work with System Activity function is placed in a subset mode. While in this mode, you see performance data for only the selected jobs and tasks whenever the display is refreshed. Also in this mode, you can use option 5 (Work with job) on a job and the job remains in the selected group. To remove a single job or task from the selected group (as long as it is not the last or only selected entry), blank out the option field and press the Enter key. This causes a new group to be built from those entries that still have a 1 in the *Opt* field.

To return to normal operating mode, press either F13 (Jobs and tasks), F14 (Jobs only), or F15 (Tasks only). These function keys are the only way to end the monitoring feature without exiting the Work with System Activity function.

Working with Jobs

By typing a 5 in the option field next to a job and pressing the Enter key, the Work with Job (WRKJOB) command is started for that job. Selecting more than one job before pressing the Enter key causes the WRKJOB command to be started multiple times.

Note: Option 5 (Work with job) is valid only with jobs. This function cannot be started for tasks.

Refer to the *CL Reference* manual or the *Work Management Guide* for further information on the Work with Job command.

Displaying Different Information Types

As previously mentioned, you can control the type of information being shown on the display. This control comes through the use of the INFTYPE (Information type) parameter or through the use of F13 (Display jobs and tasks), F14 (Display jobs only), or F15 (Display tasks only).

If you specify INFTYPE(*ALL) on the Work with System Activity command or press F13, statistics for both jobs and tasks are shown. Column headings and function keys similar to the following appear on the Work with System Activity display:

```
                                           Total  Total
    Job or                           CPU   Sync   Async   PAG     EAO
Opt Task       User       Number Pty Util  I/O    I/O     Fault   Excp
 _  DSP25      QPGMR      008124  20 56.1   12      4       0       2
 _  DSP01      QSECOFR    008122  20  6.4    5      2       0       3
 _  SMAI03                         0   .3    0      0       0       0
 _  ROUT04                         0   .3    0      0       0       0
                                   .
                                   .
                                   .
F14=Display jobs only   F15=Display tasks only   F16=Sequence by I/O
F24=More keys
```

If you specify INFTYPE(*JOB) or press F14, statistics for jobs only are shown. Column headings and function keys similar to the following appear on the Work with System Activity display:

```
                                           Total  Total
                                     CPU   Sync   Async   PAG     EAO
Opt Job        User       Number Pty Util  I/O    I/O     Fault   Excp
 _  DSP25      QPGMR      008124  20 56.1   12      4       0       2
 _  DSP01      QSECOFR    008122  20  6.4    5      2       0       3
                                   .
                                   .
                                   .
F13=Display jobs and tasks   F15=Display tasks only   F16=Sequence by I/O
F24=More keys
```

And finally, if you specify INFTYPE(*TASK) or press F15, statistics for tasks only are shown. Column headings and function keys similar to the following appear on the Work with System Activity display:

```
                                           Total  Total
                                     CPU   Sync   Async   PAG     EAO
Opt Task                         Pty Util  I/O    I/O     Fault   Excp
 _  SMAI03                         0   .3    0      0       0       0
 _  ROUT04                         0   .3    0      0       0       0
                                   .
                                   .
                                   .
F13=Display jobs and tasks   F14=Display jobs only   F16=Sequence by I/O
F24=More keys
```

Accessing Work Management Functions

To assist you in analyzing the performance of the system, function keys F20 through F23 have been set up to provide access to several Work Management functions. The third set of function keys appears on the Work with System Activity display as follows:

```
F20=Work with active jobs    F21=Work with system status
F22=Work with subsystems     F23=Work with disk status    F24=More keys
```

F20 starts the Work with Active Jobs (WRKACTJOB) command. F21 starts the Work with System Status (WRKSYSSTS) command. F22 starts the Work with Subsystems (WRKSBS) command, and F23 starts the Work with Disk Status (WRKDSKSTS) command. Refer to the *Work Management Guide* for further information on these commands.

Content of Database File QAITMON

The collected performance data is stored in the file QAITMON located in the library specified by the LIB parameter on the Work with System Activity command. Each performance collection, which is stored in a member determined by the MBR parameter, contains one record for each active job or task in an interval.

Figure 7-1 describes the content of a single record in QAITMON.

Figure 7-1 (Page 1 of 2). File QAITMON

Field Name	Attributes	Description
LVLID	CHAR(7)	The level of the module that collected this data and the level of this file in the form VVRRRFF, where VV = version number, RRR = release number, and FF = file level.
DTETIM	CHAR(13)	The date (CMMDDYY) and time (HHMMSS) that the data was collected.
ITVTIM	PACKED(11,0)	The time between data collections, where one unit equals 4096 microseconds.
CPUTOT	PACKED(11,0)	The total processing unit time used by all tasks and jobs during the interval, where one unit equals 4096 microseconds. For multiple-processor systems, this is the average use by all processors.
NAME	CHAR(10)	The job or task name for this entry.
JOBUSR	CHAR(10)	The user profile associated with a job.
JOBNBR	CHAR(6)	The number assigned to the job.
PTY	CHAR(3)	The priority of the job or task when the data was collected.
CPUDLT	PACKED(11,0)	The processing unit time used by this task or job during the interval, where one unit equals 4096 microseconds. For multiple-processor systems, this is the average use by all processors.
IOTOT	PACKED(11,0)	The total physical I/O operations (synchronous and asynchronous) performed by this job or task.
SDBR	PACKED(11,0)	The number of synchronous database reads.
SNDBR	PACKED(11,0)	The number of synchronous nondatabase reads.
SDBW	PACKED(11,0)	The number of synchronous database writes.
SNDBW	PACKED(11,0)	The number of synchronous nondatabase writes.
ADBR	PACKED(11,0)	The number of asynchronous database reads.
ANDBR	PACKED(11,0)	The number of asynchronous nondatabase reads.

Figure 7-1 (Page 2 of 2). File QAITMON

Field Name	Attributes	Description
ADBW	PACKED(11,0)	The number of asynchronous database writes.
ANDBW	PACKED(11,0)	The number of asynchronous nondatabase writes.
PAGFLT	PACKED(11,0)	The number of process access group faults.
EAOCNT	PACKED(11,0)	The number of effective address overflow exceptions.
JTFLAG	CHAR(1)	A flag indicating whether this record represents a job or task where '00'X = job and '80'X = task.
RSRV1	CHAR(4)	Reserved field.
PERMW	PACKED(11,0)	The number of writes that were for permanent objects.
IOPND	PACKED(11,0)	The number of I/O-pending page faults.
SMSYNC	PACKED(11,0)	The number of waits for asynchronous I/O operations to complete.
OVRTOT	PACKED(11,0)	The total number of binary, decimal, and floating point overflow exceptions.
CPU1	PACKED(11,0)	For multiple-processor systems, the time used in processor one by jobs and tasks during the interval. One unit of time equals 4096 microseconds.
CPU2	PACKED(11,0)	For multiple-processor systems, the time used in processor two by jobs and tasks during the interval. One unit of time equals 4096 microseconds.
CPUCNT	PACKED(3,0)	The number of active processors in the system during data collection.

Print Activity Report

The Print Activity Report (PRTACTRPT) command creates a report using the performance data collected by the Work with System Activity (WRKSYSACT) command. This report is produced in the spooled file QPITACTR.

Depending on the value specified for the Report Type (RPTTYPE) option on the Print Activity Report command, one of two report types, or both, are created. The summary report provides the top 10 listings showing the most CPU-intensive and the most I/O-intensive entries over the entire specified period. The detailed report shows a selected number of entries for each interval in the specified period. These entries are ordered according to a user selected field. Refer to the following sections for more detail on each of these report types. Refer to the *CL Reference* manual for further information on the RPTTYPE parameter.

Summary Activity Report

The Summary Activity Report consists of two sections. The first lists (in decreasing order) the top 10 entries according to CPU utilization during the specified period, and the second lists (also in decreasing order) the top 10 entries according to total I/O activity performed during the specified period. The value used for total I/O is actually the sum of the total synchronous I/O and the total asynchronous I/O. If 10 active jobs or tasks are not present in the specified period, these sections list as many entries as are available.

The following represents a sample Summary Activity Report:

```
                                        System Activity Report                                         04/19/90  15:30:21
                                                                                                                Page    1
                                        Sample Activity Report

Member . . . . :  QAITMON     Report Type . . . . :  SUMMARY     Version . . . . :     2   Started . . . . :  4/19/90  15:28:00
Library . . . . :  QPFRDATA                                       Release . . . . :   1.0   Stopped . . . . :  4/19/90  15:28:57
```

Order by CPU Utilization:

Job or Task	User	Number	Pty	CPU Util	Total Sync I/O	Total Async I/O	PAG Fault	EAO Excp	Synchronous I/O DB Read	Synchronous I/O DB Write	Synchronous I/O Non-DB Read	Synchronous I/O Non-DB Write	Asynchronous I/O DB Read	Asynchronous I/O DB Write	Asynchronous I/O Non-DB Read	Asynchronous I/O Non-DB Write
DSP04	QPGMR	008119	10	58.0	206	378	10	821	116	0	88	2	378	0	0	0
DSP01	QSECOFR	008134	20	12.4	58	3	1	14	0	0	58	0	3	0	0	0
DSP89	QPGMR	008123	20	11.6	65	2	1	2	0	0	65	0	2	0	0	0
DSP09	QSYSOPR	008191	20	5.2	95	7	9	1	4	4	85	2	7	0	0	0
DSP100	QPGMR	008122	20	2.8	7	0	0	0	0	0	7	0	0	0	0	0
DSP06	QSECOFR	008194	20	2.4	7	0	6	0	0	0	5	2	0	0	0	0
QBASE	QSYS	007956	0	1.7	25	0	0	0	0	0	17	8	0	0	0	0
SFTR			0	.7	45	0	0	0	0	0	34	11	0	0	0	0
WRKSYSACT	QSECOFR	008193	1	.5	0	0	0	0	0	0	0	0	0	0	0	0
#50009			0	.1	5	0	0	0	0	0	5	0	0	0	0	0

Order by Total I/O:

Job or Task	User	Number	Pty	CPU Util	Total Sync I/O	Total Async I/O	PAG Fault	EAO Excp	Synchronous I/O DB Read	Synchronous I/O DB Write	Synchronous I/O Non-DB Read	Synchronous I/O Non-DB Write	Asynchronous I/O DB Read	Asynchronous I/O DB Write	Asynchronous I/O Non-DB Read	Asynchronous I/O Non-DB Write
DSP04	QPGMR	008119	10	58.0	206	378	10	821	116	0	88	2	378	0	0	0
DSP09	QSYSOPR	008191	20	5.2	95	7	9	1	4	4	85	2	7	0	0	0
SMP001			0	.1	82	0	0	0	0	7	0	75	0	0	0	0
SMP005			0	.1	82	0	0	0	0	7	0	75	0	0	0	0
SMP003			0	.1	79	0	0	0	0	5	0	74	0	0	0	0
SMP004			0	.1	79	0	0	0	0	7	0	72	0	0	0	0
SMP006			0	.1	79	0	0	0	0	12	0	67	0	0	0	0
SMP002			0	.1	77	0	0	0	0	8	0	69	0	0	0	0
DSP89	QPGMR	008123	20	11.6	65	2	1	2	0	0	65	0	2	0	0	0
DSP01	QSECOFR	008134	20	12.4	58	3	1	14	0	0	58	0	3	0	0	0

Figure 7-2. Sample Summary Activity Report

The header portion of this report contains the following information:

Report title
The title of the report.

Current date and time
The date and time when this report was printed.

Report page number
The page number currently being printed.

User-selected report title
The title specified by the user on the TITLE parameter of the Print Activity Report command.

Member The name of the member in QAITMON that contained the performance data.

Library The library where QAITMON was located.

Report type
The type of report, either summary or detail, being printed.

Version The version of the Performance Tools/400 licensed program that collected the data.

Release The release level of the Performance Tools/400 licensed program that collected the data.

Period start date and time
The start date and time of the period during which the performance statistics being printed were collected.

Period end date and time
The end date and time of the period during which the performance statistics being printed were collected.

The columns in the summary activity report are:

Job or Task
The name of the job or task for which the performance statistics are being printed.

User The user profile associated with the job.

Number The number assigned to the job.

Pty The priority at which the job or task was running when the performance statistics were first collected.

CPU Util The percentage of the specified period during which the processing unit was used by the job or task. For a multiple-processor system, this is the average use across all processors.

Total Sync I/O
The total number of synchronous physical disk I/O operations performed by the job or task during the specified period. This value is the sum of the synchronous database/nondatabase reads and writes.

Total Async I/O
The total number of asynchronous physical disk I/O operations started by the job or task during the specified period. This value is the sum of the asynchronous database/nondatabase reads and writes.

PAG Fault
The number of process access group (PAG) faults caused by the job or task during the specified period.

EAO Excp The number of effective address overflow (EAO) exceptions caused by the job or task during the specified period.

Synchronous I/O DB Read
The number of synchronous database read operations performed by the job or task during the specified period.

Synchronous I/O DB Write
The number of synchronous database write operations performed by the job or task during the specified period.

Synchronous I/O Non-DB Read
The number of synchronous nondatabase read operations performed by the job or task during the specified period.

Synchronous I/O Non-DB Write
The number of synchronous nondatabase write operations performed by the job or task during the specified period.

Asynchronous I/O DB Read
The number of asynchronous database read operations started by the job or task during the specified period.

Asynchronous I/O DB Write
The number of asynchronous database write operations started by the job or task during the specified period.

Asynchronous I/O Non-DB Read
The number of asynchronous nondatabase read operations started by the job or task during the specified period.

Asynchronous I/O Non-DB Write
The number of asynchronous nondatabase write operations started by the job or task during the specified period.

Note: The asynchronous I/O operations are performed by system asynchronous I/O tasks.

Detail Activity Report

For each interval available in the specified period, the Detail Activity Report lists the performance statistics for the number of entries specified by the Number of Jobs (NBRJOBS) parameter. The entries are ordered according to the Sequence (SEQ) parameter.

The following represents a sample Detail Activity Report.

System Activity Report 04/19/90 15:30:21
Page 2

Sample Activity Report

Member . . . . : QAITMON Report Type . . . . : DETAIL Version . . . . : 2 Started . . . . : 4/19/90 15:28:00
Library . . . . : NOVEYDATA Sequence . . . . . : CPU Release . . . . : 1.0 Stopped . . . . : 4/19/90 15:28:57

Time . . . . : 15:28:00 Total CPU Utilization (Percent) . . . . . . : 6.7

					Total	Total			-------Synchronous I/O-----				------Asynchronous I/O------			
Job or Task	User	Number	Pty	CPU Util	Sync I/O	Async I/O	PAG Fault	EAO Excp	DB Read	DB Write	Non-DB Read	Non-DB Write	DB Read	DB Write	Non-DB Read	Non-DB Write
DSP04	QPGMR	008119	10	2.6	19	0	0	0	0	0	19	0	0	0	0	0
SFTR			0	2.4	33	0	0	0	0	0	27	6	0	0	0	0
WRKSYSACT	QSECOFR	008193	1	.8	0	0	0	0	0	0	0	0	0	0	0	0
DSP20	QPGMR	008192	20	.3	3	0	0	0	0	0	3	0	0	0	0	0
SMP001			0	.0	3	0	0	0	0	0	0	3	0	0	0	0
SMP003			0	.0	4	0	0	0	0	0	0	4	0	0	0	0
SMP002			0	.0	4	0	0	0	0	0	0	4	0	0	0	0
SMP004			0	.0	3	0	0	0	0	0	0	3	0	0	0	0
SMP005			0	.0	2	0	0	0	0	0	0	2	0	0	0	0
SMP006			0	.0	3	0	0	0	0	0	0	3	0	0	0	0

Time . . . . : 15:28:05 Total CPU Utilization (Percent) . . . . . . : 34.3

					Total	Total			-------Synchronous I/O-----				------Asynchronous I/O------			
Job or Task	User	Number	Pty	CPU Util	Sync I/O	Async I/O	PAG Fault	EAO Excp	DB Read	DB Write	Non-DB Read	Non-DB Write	DB Read	DB Write	Non-DB Read	Non-DB Write
DSP04	QPGMR	008119	10	30.3	85	71	5	0	55	0	29	1	71	0	0	0
WRKSYSACT	QSECOFR	008193	1	.5	0	0	0	0	0	0	0	0	0	0	0	0
SMP004			0	.4	20	0	0	0	0	1	0	19	0	0	0	0
SMP003			0	.3	23	0	0	0	0	0	0	23	0	0	0	0
SMP006			0	.3	21	0	0	0	0	2	0	19	0	0	0	0
SMP002			0	.3	20	0	0	0	0	1	0	19	0	0	0	0
SMP005			0	.3	21	0	0	0	0	0	0	21	0	0	0	0
SMP001			0	.3	22	0	0	0	0	1	0	21	0	0	0	0
SMAI05			0	.2	0	0	0	0	0	0	0	0	0	0	0	0
SMAI03			0	.2	0	0	0	0	0	0	0	0	0	0	0	0

Figure 7-3. Sample Detail Activity Report

The header portion of this report contains the same information as found on the summary report, except for the *Sequence* field, which defines the order of the entries listed for each interval. The value found in this field corresponds to the value specified for the sequence (SEQ) parameter on the Print Activity Report command.

The body of the Detail Activity Report contains the same columns of information found on the summary report. There are, however, two additional fields associated with the statistics for each interval:

Time The end time for the collection interval.

Total CPU Utilization
The processing unit use for the entire system during the collection interval.

For a multiple-processor system, *Total CPU Utilization* is replaced by a utilization value for each processor in the system.

Chapter 8. Printing Performance Reports

This chapter explains how to print reports from the performance data you collect using the Start Performance Monitor (STRPFRMON) command. Examples of the reports are also shown.

Use the following commands to produce reports:

- Print System Report (PRTSYSRPT)
- Print Component Report (PRTCPTRPT)
- Print Transaction Report (PRTTNSRPT)
- Print Lock Report (PRTLCKRPT)
- Print Job Report (PRTJOBRPT)
- Print Pool Report (PRTPOLRPT)
- Print Resource Report (PRTRSCRPT)
- Print Batch Job Trace Report (PRTTRCRPT)

Note: In this chapter, when response time is mentioned, assume that it is internal response time, unless stated otherwise. Internal response time is the AS/400 host response time.

For information about the elements of response time, see "Elements of Response Time" on page A-1.

The *CL Reference* manual contains additional information about the commands described in this chapter.

Using Menus to Print Performance Reports

To print selected information from the data you collect, choose option 3 (Print performance report) on the IBM Performance Tools/400 menu. The Print Performance Report display appears:

```
                          Print Performance Report

Library . . . . . .   QPFRDATA

Type option, press Enter.
  1=System report   2=Component report   3=Transaction report   4=Lock report
  5=Job report   6=Pool report   7=Resource report   8=Batch job trace report

Option   Member      Text                                    Date      Time
  1      LOCKS                                               12/11/89  12:20:29
  _      DLTTEST18                                           12/11/89  10:10:20
  _      DLTTEST17                                           12/11/89  10:10:02
  _      DLTTEST16                                           12/11/89  10:09:42
  _      DLTTEST15                                           12/11/89  10:09:32
  _      DLTTEST14                                           12/11/89  10:09:22
  _      DLTTEST13                                           12/11/89  10:09:04
  _      DLTTEST11                                           12/11/89  10:08:56
  _      DLTTEST10                                           12/11/89  10:08:49
  _      DLTTEST9                                            12/11/89  10:08:44
  _      DLTTEST8                                            12/11/89  10:08:35
                                                                        More...
F3=Exit   F5=Refresh   F11=Work with your spooled output files   F12=Cancel
F15=Sort by member     F16=Sort by text
```

The member name, a text description, and the date and time you collected each set of performance data appear on this display. If you cannot find the data you want to print a report from, use the appropriate function key to sort the sets of performance data. You can sort them by member name, text description, or by the date and time the member was created. When you find the performance data, indicate the type of report you want by typing the appropriate option.

The System Report supplies you with an overview of how the system is operating. This report provides you with summary information on the workload, resource use, storage pool utilization, disk utilization, and communications. The Component Report supplies you with information about the same components of system performance as a System Report, but at a greater level of detail.

The System Report, Component Report, Job Report, Pool Report, and Resource Report are produced from sample data you collect using the STRPFRMON command. If you collect trace data using the STRPFRMON command, you can produce a Transaction Report, Lock Report, or Batch Job Trace Report from this information.

To select the type of report to print, type one of the following options that corresponds to the type of report:

Option	**Description**
1	System Report
2	Component Report
3	Transaction Report
4	Lock Report
5	Job Report
6	Pool Report
7	Resource Report
8	Batch Job Trace Report

If you choose to print a Transaction Report, press the Enter key and the parameters for the Print Transaction Report (PRTTNSRPT) command appear.

If you choose to print a Lock Report, press the Enter key and the parameters for the Print Lock Report (PRTLCKRPT) command appear.

If you choose to print a Batch Job Trace Report, press the Enter key and the parameters for the Print Batch Job Trace Report (PRTTRCRPT) command appear.

The Select Categories for Report display appears when you select to print one of the following reports:

- System Report
- Component Report
- Job Report
- Pool Report
- Resource Report

```
                         Select Categories for Report

Member . . . . . . . . . . . :   TUEDTA

Type options, press Enter.  Press F6 to print entire report.
  1=Select

Option     Category
  _        Time interval
  _        Job
  _        User ID
  _        Subsystem
  _        Pool
  _        Communications line
  _        Control unit
  _        Functional area




                                                                     Bottom
F3=Exit   F6=Print entire report   F12=Cancel
```

The name of the performance data member you chose on the Print Performance Report display appears at the top of the Select Categories for Report display.

Type a 1 in the *Option* column next to those categories of information for which you want performance data. Press the Enter key.

Note: Normally you include all categories of information in your report. To do this, do not place a 1 in any category. Instead, simply press F6.

By choosing the Time interval option, the Select Time Intervals display appears. This display shows an interactive view of some of the key performance parameters of the data collected.

The member name you typed on the Print Performance Reports display appears in the *Performance data* field. The intervals you defined to collect the performance data appear.

```
                           Select Time Intervals

Library . . . . . :   QPFRDATA      Performance data . . . . . :   TUEDTA

Type options, press Enter.
  1=Select
                                         CPU          High        Pool
                   Transaction     -Utilization-   --Util--   -Fault/Sec-
Opt  Date   Time   Count  Resp     Tot Inter Bch   Dsk Unit   Mch User ID   Excp
 _   12/29  12:39     33   1.5      3     2   0     2 0017     0     1 03     77
 _   12/29  12:44     26    .9      1     1   0     1 0002     0     0 03      7
 _   12/29  12:49     20    .2      1     0   0     1 0009     0     0 03      7








 F3=Exit                      F5=Refresh                  F12=Cancel
 F13=Sort (date/time)         F14=Sort (count)            F24=More keys
```

Use the Select Time Intervals display to choose specific time intervals from the performance data to produce a report. You should select specific time intervals for these reasons:

- To help manage the volume of data associated with the performance measurement. The Select Time Intervals display allows you to interactively select the time intervals of interest. This selection reduces the amount of processing required to produce the requested report, and also reduces the size of the resulting report.
- To create measurement profiles that serve as input to the capacity planning support provided by the Model System (MDLSYS) command. For more information about measurement profiles, see Chapter 10, "Capacity Planning and Performance Prediction."

To select time intervals to print on your report, type a 1 in the *Opt* column next to the appropriate intervals. When you select multiple intervals, they are combined to create a single report, or a single measurement profile.

If it is difficult for you to find the time interval you are interested in printing information for, before making any selections, you can sort the time intervals in a different way. You can choose to sort the time intervals in the following ways:

- Date/time
- Transaction count
- Response time
- Total processing unit time
- Interactive processing unit time
- Batch processing unit time
- Disk utilization
- Machine pool faults
- User pool faults
- Exceptions

If you choose to print the report with only certain categories of information, a display appears for each category. For example, if you choose Pool, the Select or Omit Pools display appears.

Use the Select or Omit Pools display to select pools to include or omit from your report. To use this display, type the number for the pools you want to select or omit. If you do not know the pool numbers to select, press F4 (Prompt) to see a list of pools that were active during the collection of performance data.

```
                              Select or Omit Pools

 Member . . . . . . . . . . . :   TUEDTA

 Type options, press Enter.
   1=Select   2=Omit

 Option      Pool      Text
   _          01       Machine pool
   _          02       Base pool
   _          03
   _          04
                                                                        Bottom








 F12=Cancel
```

Type a 1 in the *Option* column next to the items you want to include in your report. Or type a 2 if you want certain items omitted from your report.

Note: You cannot use both the Select and Omit options at the same time. You must indicate either the items to select or the items to omit.

To include all the items in the report, leave the *Option* column blank.

For each category you choose on the Select Categories for Report display, you must complete one of the following corresponding displays:

- Select Time Intervals
- Select or Omit Pools
- Select or Omit Jobs
- Select or Omit User IDs
- Select or Omit Subsystems
- Select or Omit Communications Lines
- Select or Omit Control Units
- Select or Omit Functional Areas

Once you choose the information you want to appear on your report from the options shown on these displays, the Select Report Options display appears. Following is the version of the display that appears if you did not use the Select Time Intervals display to choose any time intervals:

```
                         Specify Report Options

Type choices, press Enter.

  Report title . . .   New data entry in production______________

  Start:
    Day  . . . . . .   *FIRST__                 *FIRST, MM/DD/YY
    Time . . . . . .   *FIRST__                 *FIRST, HH:MM:SS

  Stop:
    Day  . . . . . .   *LAST___                 *LAST, MM/DD/YY
    Time . . . . . .   *LAST___                 *LAST, HH:MM:SS

  Measured profile:
    Profile  . . . .   *NONE___                 *NONE, name
    Replace  . . . .   N                        Y=Yes, N=No
```

If you so choose, type a report title in the *Report title* field. Specify the start and stop date and time. If you do not specify the start and stop date and time, the report includes data from the first (or only) date that data was collected, to the last (or only) date that data was collected. If you want to keep this data as a capacity planning profile, type the profile name in the *Measured profile* field. Press the Enter key to process and print your report.

Note: The *Measured profile* field appears only if you requested printing of a System Report.

If you made use of the Select Time Intervals display, the following version of the Select Report Options display appears instead:

```
                         Specify Report Options

Type choices, press Enter.

  Report title . . .   ________________________________________

  Measured profile:
    Profile  . . . .   *NONE___                 *NONE, name
    Replace  . . . .   N                        Y=Yes, N=No
```

If you so choose, type a report title in the *Report title* field. If you want to keep this data as a capacity planning profile, type the profile name in the *Measured profile* field. Press the Enter key to process and print your report.

The following section provides an example of how to collect performance data and print a System Report using the default values.

Using Defaults to Print Performance Reports

The example in this section describes how to collect performance data and print a System Report using the default values on the displays shown.

Note: These values are the defaults for the STRPFRMON and PRTSYSRPT commands. For more information about the default values for these commands, see the *CL Reference* manual.

Use the Start Performance Tools (STRPFRT) command to start performance tools. Make the choices shown on the following displays to collect performance data and print a System Report using the system defaults.

1. To start collecting performance data, choose option 2 (Collect performance data) on the IBM Performance Tools /400 menu.

```
PERFORM                    IBM Performance Tools/400
                                                         System:   RCH38366
Select one of the following:

     1. Select type of status
     2. Collect performance data
     3. Print performance report
     4. Capacity planning/modeling
     5. Programmer performance utilities
     6. Configure and manage tools
     7. Display performance data
     8. System activity
     9. Performance graphics
    10. Advisor

    70. Related commands
```

2. Press the Enter key. The Collect Performance Data display appears.

```
                           Collect Performance Data

Select one of the following:

     1. Start collecting data
     2. Stop collecting data
     3. Work with performance collection
```

3. Choose option 1 (Start collecting data).
4. Press the Enter key. The Start Collecting Data display appears.

```
                            Start Collecting Data

Select one of the following:

  1. Collect data with defaults
  2. Collect data with menus
  3. Collect data with command
```

5. Choose option 1 (Collect data with defaults).
6. Press the Enter key. The Collect Data with Defaults display appears.

```
                          Collect Data with Defaults

Type choices, press Enter.

  Member . . . . . . . .   Thursdata_     Name

    Library  . . . . . .     QPFRDATA__   Name

  Text . . . . . . . . .   System performance for Thursday

  Time duration:
    Hours  . . . . . . .   9              0-999
    Minutes  . . . . . .   0              0-60
```

7. Type an appropriate name for the performance data in the *Member* field. In this example, the performance data is called Thursdata.

8. Type an appropriate description for the performance data you want to collect in the *Text* field. In this example, the description is System performance for Thursday.

9. Type how long you want to collect performance data in the *Time duration* fields. In this example, 9 is typed in the *Hours* field so performance data is collected for 9 hours.

10. Press the Enter key to start the data collection process.

11. Press F3 (Exit).

    Wait for the data collection process to end before you complete the next steps. In this example, you would wait for 9 hours before you would continue.

12. Once the data collection process is complete, use the STRPFRT command to start performance tools.

13. Choose option 3 (Print performance report) on the IBM Performance Tools/400 menu. The Print Performance Report display appears.

```
                            Print Performance Report

Library . . . . . .   QPFRDATA

Type option, press Enter.
  1=System report   2=Component report   3=Transaction report   4=Lock report
  5=Job report   6=Pool report   7=Resource report   8=Batch job trace report

Option   Member      Text                             Date       Time
  1      THURSDATA   SYSTEM PERFORMANCE FOR THURSDAY  12/11/90   12:20:29
         DLTTEST18                                    12/11/90   10:10:20
         DLTTEST17                                    12/11/90   10:10:02
```

14. Type a 1 in the *Option* column, next to the performance member called THURSDATA, to indicate that you want to print a System Report from this performance data.

15. Press the Enter key. The Select Categories for Report display appears.

```
                         Select Categories for Report

Member . . . . . . . . . . . :   THURSDATA

Type options, press Enter.  Press F6 to print entire report.
  1=Select

Option     Category
  _        Time interval
  _        Job
  _        User ID
  _        Subsystem
  _        Pool
  _        Communications line
  _        Control unit
  _        Functional area
```

16. Press F6 to indicate that you do not want to restrict the categories for this report.

The Specify Report Options display appears.

```
                          Specify Report Options

 Type choices, press Enter.

   Report title . . .   SYSTEM PERFORMANCE FOR THURSDAY

   Start:
     Day  . . . . . .   *FIRST                     *FIRST, MM/DD/YY
     Time . . . . . .   *FIRST                     *FIRST, HH:MM:SS

   Stop:
     Day  . . . . . .   *LAST                      *LAST, MM/DD/YY
     Time . . . . . .   *LAST                      *LAST, HH:MM:SS

   Measured profile:
     Profile  . . . .   *NONE                      *NONE, name
     Replace  . . . .   N                          Y=Yes, N=No
```

17. Press the Enter key to submit a batch request to print a System Report for the entire data collection period.

18. Press F3 (Exit) to go to the IBM Performance Tools/400 menu.

    The batch request you submit takes a period of time to complete, depending on the amount of data collected. Use the Work with Submitted Job (WRKSBMJOB) command to check the status of the request.

    After the System Report has been produced, you can view it online and direct it to an active writer by following steps 19 through 22.

```
 PERFORM                    IBM Performance Tools/400
                                                          System:   RCH38366
 Select one of the following:

      1. Select type of status
      2. Collect performance data
      3. Print performance report
      4. Capacity planning/modeling
      5. Programmer performance utilities
      6. Configure and manage tools
      7. Display performance data
      8. System activity
      9. Performance graphics
     10. Advisor

     70. Related commands


 Selection or command
 ===> WRKSPLF___________________________________________________________

 F3=Exit   F4=Prompt   F9=Retrieve   F12=Cancel   F13=User support
 F16=System main menu
```

19. Type WRKSPLF on the command entry line, and press the Enter key.

    The Work with All Spooled Files display appears.

```
                         Work with All Spooled Files

Type options, press Enter.
  2=Change   3=Hold   4=Delete   5=Display   6=Release   8=Attributes

                                Device or                    Total    Cur
Opt  File         User          Queue       User Data  Sts   Pages    Page  Copy
 2   QPPTSYSR     USERID        QSYSPRT                RDY       7             1









Parameters for option 2 or command
===> OUTQ(outqname)____________________________________________
F3=Exit   F10=View 3   F11=View 2   F12=Cancel   F24=More keys
```

On this display you could choose option 5, for example, to view the System Report online.

20. For this example, type a 2 under the *Opt* column to change the output queue for the System Report (the QPPTSYSR file). In this example, you might want to move the report to an output queue that has an active writer, so the report prints on the device the writer is associated with.
21. Type the new output queue name. To do this, type OUTQ(outqname) on the command entry line.
22. Press the Enter key. The System Report prints when a device is available.

System Report

To produce the System Report, use the Print System Report (PRTSYSRPT) command, or select option 1 (System report) on the Print Performance Reports display.

The System Report provides an overview of system operation during the data collection period. Produce and save this report periodically so you have a record of the workload and resource utilization (for example, how your system meets the users needs and at what cost). Use the reports to see what processing trends are developing, and to project when you might need to make application, system, or operational changes to accommodate changing workloads.

The System Report consists of the following sections:

- System Report: Workload
- System Report: Resource Utilization
- System Report: Resource Utilization Expansion
- System Report: Storage Pool Utilization
- System Report: Disk Utilization
- System Report: Communications Summary
- Report Selection Criteria
- System Model (MDLSYS) Parameters

Each page of the System Report shows the header information in Figure 8-1 on page 8-11.

```
                                                      System Report                                                11/24/90 10:28:19
                                                        Workload                                                          Page 0001
                                                  Sample System Report
Member  . . . : TEMP          Model/Serial  . . :  B40/10-15005   Main storage . . :    40.0 M   Started  . . . . : 11/02/90 10:01:54
  Library . . : QPFRDATA      System name . . . :  ABCSYSTEM      Version/Release  :    2/ 1.0   Stopped  . . . . : 11/02/90 10:31:51
```

Figure 8-1. System Report: Header Information

This header information consists of the following:

- Report title
- Current date and time
- Report page number
- User-selected report title
- Data member name
- Library name
- Model number
- Serial number
- System name
- Main storage size
- OS/400 version and release level
- Data collection start date and time
- Data collection stop date and time

The remainder of this section shows you the formats of the various reports you can produce, and provides an explanation of the columns in these reports.

Note: For multiple-processor systems, processing unit use for each processor is the average use across all processors.

System Report: Workload

The first part of the Workload section, shown in Figure 8-2, shows the interactive workload of the system.

```
                                                      System Report                                                11/24/90 10:28:19
                                                        Workload                                                          Page 0001
                                                  Sample System Report
Member  . . . : TEMP          Model/Serial  . . :  B40/10-15005   Main storage . . :    40.0 M   Started  . . . . : 11/02/90 10:01:54
  Library . . : QPFRDATA      System name . . . :  ABCSYSTEM      Version/Release  :    2/ 1.0   Stopped  . . . . : 11/02/90 10:31:51
Interactive Workload
     Job               Number            Average          Logical DB        ------ Printer -------        Communications        MRT
     Type           Transactions         Response         I/O Count          Lines         Pages           I/O Count         Max Time
-------------       ------------         --------        -----------       ----------     -------        --------------      --------
PassThru                   3,085              .61             38,200                0           0                     0
MRT                           56              .85              6,795                0           0                     0            56
Total/Average              3,141              .75             44,995                0           0                     0
------------------------------------------------------------------------------------------------------------------------------------
```

Figure 8-2. Workload Section: Interactive Workload

The columns in the Interactive Workload section of this report are as follows:

Job Type Type of interactive job. Possible values are as follows:

- Distributed data management (DDM) server
- Interactive
- Multiple requester terminal (MRT)
- Pass-through
- PC Support server
- System/36

Number Transactions
Total number of interactive transactions processed.

Average Response
Average response time (in seconds) for interactive transactions. The Total/Average interactive response time does not include transactions for DDM server jobs.

Logical DB I/O Count
Number of database file input/output requests sent by programs for each job type.

Note: Logical I/O contrasts with physical 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 result in access to auxiliary storage.

Printer Lines
Number of lines printed.

Printer Pages
Number of pages printed.

Communications I/O Count
Number of communications I/O operations issued.

MRT Max Time
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.

The second part of the Workload section, shown in Figure 8-3, shows the noninteractive workload of the system.

Noninteractive Workload

Job Type	Number of Jobs	Logical DB I/O Count	Printer Lines	Printer Pages	Communications I/O Count	CPU per Logical I/O	Logical I/O /Second
Batch	1	1,227	0	0	0	.01	.6
Spool	1	0	0	0	0	.00	.0
Evoke	25	0	0	0	0	.00	.0
Total/ Average	27	1,227	0	0	0	.01	.6

Total CPU Utilization . . . . . . : 41.9

Figure 8-3. Workload Section: Noninteractive Workload

The columns in the noninteractive workload section of this report are as follows:

Job Type Type of noninteractive job. Possible values are as follows:

- Autostart
- Batch
- Evoke
- Spool

Number of Jobs
Number of jobs.

Logical DB I/O Count
Number of database I/O operations.

Printer Lines
Number of lines printed.

Printer Pages
Number of pages printed.

Communications I/O Count
Number of communications I/O operations.

CPU per Logical I/O
Processing unit time used for each logical disk I/O operation.

Logical I/O /Second
Average number of logical disk I/O operations per second.

Total CPU Utilization
Total processing unit use.

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 a multiple-processor system, *Total CPU Utilization* is replaced by a utilization value for each processor in the system. Here is an example of this part of the display for a dyadic system (2 processors in the system):

```
Total/              27            1,227              0        0               0              .01          .6
 Average CPU utilization . . . . . . . . . . . . . . :   41.9
 CPU 1 utilization . . . . . . . . . . . . . . . . . :   41.7
 CPU 2 utilization . . . . . . . . . . . . . . . . . :   42.2
```

System Report: Resource Utilization

The Resource Utilization section, shown in Figure 8-4, shows the average resource utilization per interactive transaction. Use it to note changes in resource utilization from one measurement period to another and to determine resource utilization trends.

```
                                                     System Report                                         11/24/90 10:28:19
                                                  Resource Utilization                                            Page 0001
                                                  Sample System Report
Member  . . . : TEMP         Model/Serial  . . :  B40/10-15005   Main storage . . :   40.0 M   Started  . . . . : 11/02/90 10:01:54
  Library . . : QPFRDATA     System name . . . :  ABCSYSTEM      Version/Release  :   2/ 1.0   Stopped  . . . . : 11/02/90 10:31:51
                  ------------------- Average Per Transaction -------------------
                                                Sync         Async          Logical
    Job           Response        CPU           Disk          Disk            DB
    Type          Seconds       Seconds         I/O           I/O             I/O
--------------    --------      -------        -----         -----          -------
PassThru               .6           .14          5.9            .0            12.3
Total/Average          .6           .14          5.9            .0            12.3
-----------------------------------------------------------------------------------------------------------------------------------
```

Figure 8-4. Resource Utilization

The columns in the first part of the Resource Utilization section of the System Report are as follows:

Job Type Special interactive job categories. Possible values are as follows:

- Distributed data management (DDM) server
- Interactive
- Multiple requester terminal (MRT)
- Pass-through
- PC Support server
- System/36

Response Seconds
Average response time in seconds per transaction.

CPU Seconds
Average processing unit seconds used per transaction.

Sync Disk I/O
Synchronous disk I/O operations per transaction.

Async Disk I/O
Asynchronous disk I/O operations per transaction.

Logical DB I/O
Average number of logical I/O operations per transaction.

The second part of the Resource Utilization section of the System Report is shown in Figure 8-5.

Job Type	CPU Util	Tns /Hour Rate	Active Jobs per Interval	Disk I/O Per Second								
				Total I/O	Synchronous				Asynchronous			
					DBR	DBW	NDBR	NDBW	DBR	DBW	NDBR	NDBW
PassThru	25.2	6,274	25	10.3	10.0	.0	.2	10.0	.2	10.0	.0	.2
Total/Average	25.2	6,274	51	10.3	10.0	.0	.2	10.0	.2	10.0	.0	.2

Figure 8-5. Resource Utilization

The columns in this part of the report are as follows:

Job Type Special interactive job categories as follows:

- Distributed data management (DDM) server
- Interactive
- Multiple requester terminal (MRT)
- Pass-through
- PC Support server
- System/36

CPU Util Percentage of the system processing unit used by each type of job. For a multiple-processor system this is the total system use divided by the number of processors. Use the total processing unit percentage from this column when referring to the guidelines in Figure 10-3 on page 10-6 for interactive processing unit use. Consistent use at or above the threshold value given in Figure 10-3 on page 10-6 affects the ability of the system to perform any batch, or other low priority work.

Note: You should consider adding the system processing unit percent to the total interactive processing unit percent when applying the above guidelines. In most cases, the majority of the system processing unit percent supports the interactive workload and should be included with the interactive processing unit value. The system processing unit percent appears in Figure 8-6 on page 8-15.

Tns/Hour Rate
Average number of transactions per hour.

Active Jobs Per Interval
Average number of jobs of this type that were active per sample interval.

Disk I/O per Second
Average number of physical disk I/O operations per second.

Total I/O Sum of the read and write operations.

Synchronous DBR
Synchronous database read operations per second.

Synchronous DBW
Synchronous database write operations per second.

Synchronous NDBR
Synchronous nondatabase read operations per second.

Synchronous NDBW
Synchronous nondatabase write operations per second.

Asynchronous DBR
Asynchronous database read operations per second.

Asynchronous DBW
Asynchronous database write operations per second.

Asynchronous NDBR
Asynchronous nondatabase read operations per second.

Asynchronous NDBW
Asynchronous nondatabase write operations per second.

Note: The asynchronous I/O operations are performed by system asynchronous I/O tasks.

System Report: Resource Utilization Expansion

The Resource Utilization Expansion section, shown in Figure 8-6, gives the average resource use per transaction by job type.

```
                                                  System Report                                                    11/24/90 10:28:19
                                          Resource Utilization Expansion                                                   Page 0003
                                              Sample System Report
Member . . . : TEMP         Model/Serial  . . :  B40/10-15005  Main storage . . :    40.0 M  Started  . . . . : 11/02/90 10:01:54
  Library . . : QPFRDATA     System name . . . :  ABCSYSTEM     Version/Release  :    2/ 1.0  Stopped  . . . . : 11/02/90 10:31:51
               ------------------------------------------- Average per Transaction -------------------------------------------
               -------------------- Physical Disk I/O -----------------------    ------- Logical --------    -- Communications --
     Job       ------- Synchronous --------     ------- Asynchronous --------    -----Data Base I/O -----             I/O
     Type       DBR     DBW     NDBR    NDBW     DBR     DBW     NDBR    NDBW    Read     Write    Other      Get         Put
-------------- -----   -----   -----   -----    -----   -----   -----   -----   -----    -----    -----      -----       -----
PassThru         5.7      .0      .1      .0       .0      .0      .0      .0     9.3       .0      3.0        .0          .0
Average          5.7      .0      .1      .0       .0      .0      .0      .0     9.3       .0      3.0        .0          .0
-----------------------------------------------------------------------------------------------------------------------------
```

Figure 8-6. Resource Utilization Expansion

The columns in the first part of the Resource Utilization Expansion section of the System Report are as follows:

Job Type Special interactive job categories. Possible values are as follows:

- Distributed data management (DDM) server
- Interactive
- Multiple requester terminal (MRT)
- Pass-through
- PC Support server
- System/36

Synchronous DBR
Synchronous database read operations per transaction.

Synchronous DBW
Synchronous database write operations per transaction.

Synchronous NDBR
Synchronous nondatabase read operations per transaction.

Synchronous NDBW
Synchronous nondatabase write operations per transaction.

Asynchronous DBR
Asynchronous database read operations per transaction.

Asynchronous DBW
Asynchronous database write operations per transaction.

Asynchronous NDBR
Asynchronous nondatabase read operations per transaction.

Asynchronous NDBW
Asynchronous nondatabase write operations per transaction.

Logical Database I/O Read
Logical database read operations per transaction.

Logical Database I/O Write
Logical database write operations per transaction.

Logical Database I/O Other
Other logical database operations per transaction. This includes operations such as update and delete.

Communications I/O Get
Number of communication get operations per transaction.

Communications I/O Put
Number of communication put operations per transaction.

The second part of the Resource Utilization Expansion section, shown in Figure 8-7, contains CPU and I/O utilization information.

Priority	Job Type	CPU Util	Cum Util	Disk I/O Sync	Disk I/O Async	CPU per I/O Sync	CPU per I/O Async	DIO /Sec Sync	DIO /Sec Async
010	System	.9	.9	202	0	.0800	.0000	.1	.0
020	Interactive	.0	.0	0	0	.0000	.0000	.0	.0
	DDM Server	.0	.0	0	0	.0000	.0000	.0	.0
	PassThru	25.2	26.1	18,247	0	.0244	.0000	10.3	.0
.									
.									
Total/Average				18,449	0			10.4	.0

Figure 8-7. Resource Utilization Expansion (Second Part)

The columns in this part of the report are as follows:

Priority Job priority.

Job Type Special jobs categories. Possible values are as follows:

- Autostart
- Batch
- Distributed data management (DDM) server
- Evoke
- Interactive
- Multiple requester terminal (MRT)
- Pass-through
- PC Support server
- System/36
- Spool
- System

CPU Util Total percentage of processing unit time used by each job type. For a multiple-processor system, this is the total system use divided by the number of processors.

Cum Util Cumulative CPU use (a running total).

Note: This is taken from the individual jobs and may differ slightly from the total processing unit use on the workload page.

Disk I/O Sync
Total number of synchronous disk I/O operations.

Disk I/O Async
Total number of asynchronous disk I/O operations.

CPU per I/O Sync
CPU use per synchronous I/O.

CPU per I/O Async
CPU use per asynchronous I/O.

DIO/Sec Sync
Number of synchronous I/O operations per second.

DIO/Sec Async
Number of asynchronous I/O operations per second.

System Report: Storage Pool Utilization

Use the Storage Pool Utilization section of the System Report, shown in Figure 8-8, along with the guidelines provided in Chapter 5 of the *Work Management Guide*, to help you set the storage pool size and activity level.

```
                                                System Report                                                  11/24/90 10:28:19
                                           Storage Pool Utilization                                                   Page 0004
                                             Sample System Report
Member . . . : TEMP        Model/Serial  . . :  B40/10-15005  Main storage . . :   40.0 M   Started  . . . . : 11/02/90 10:01:54
  Library . . : QPFRDATA   System name . . . :  ABCSYSTEM     Version/Release  :   2/ 1.0   Stopped  . . . . : 11/02/90 10:31:51
                                                                --------- Avg per Second --------    ---- Avg per Minute ----
Pool      Size      Act      CPU        Number        Average    ------ DB -----    ---- Non-DB ---    Act-     Wait-     Act-
 ID        (K)      Lvl      Util        Tns          Response    Fault    Pages     Fault    Pages     Wait     Inel      Inel
----    -------     ---      -----     -------       --------     -----    -----     -----    -----    ------   ------    ------
 01       6,000      00      10.0            0            .0        .0       .0        .0       .1        0        0         0
 02      34,960      25      31.9        3,085            .6       5.2     23.6        .3       .4      104        0         0
Total/   40,960              41.9        3,085            .6       5.2     23.6        .3       .5      104        0         0
Average
```

Figure 8-8. Storage Pool Utilization

The columns in the Storage Pool Utilization section are as follows:

Pool ID Pool identifier.

Size (K) Size of the pool in kilobytes (KB). (KB equals 1024 bytes).

Act Lvl Activity level at the time of the first sample interval.

CPU Util Total processing unit use for jobs that ran in this pool. For a multiple-processor system, this is the average use across all processors.

Note: This is taken from the individual job work control blocks (WCBs) and may differ slightly from the total processing unit use on the workload page.

Number Tns
Total number of transactions processed by jobs in this pool.

Average Response
Average transaction response time in seconds.

DB Fault Average number of database faults per second.

DB Pages Average number of database pages read per second.

Non-DB Fault
Average number of nondatabase faults per second.

Non-DB Pages
Average number of nondatabase pages read per second.

Act-Wait Average number of active-to-wait job state transitions per minute.

Wait-Inel Average number of wait-to-ineligible job state transitions per minute.

Act-Inel Average number of active-to-ineligible job state transitions per minute.

System Report: Disk Utilization

The Disk Utilization section of the System Report, as shown in Figure 8-9, shows the utilization for each disk.

```
                                                    System Report                                       11/24/90 16:01:39
                                                   Disk Utilization                                             Page 0005
                                                 Sample System Report
Member  . . . : TEMP          Model/Serial  . . :  B40/10-15005   Main storage . . :    40.0 M   Started  . . . . : 11/02/90 10:01:54
  Library . . : QPFRDATA      System name . . . :  ABCSYSTEM      Version/Release  :    2/ 1.0   Stopped  . . . . : 11/02/90 10:31:51
                    Size    IOP     IOP     ASP    CSS    --Percent--         Op per    K per     ----- Average Time per I/O -----
Unit      Type      (M)     Util    ID      ID     ID     Full      Util      Second     I/O       Service       Wait      Response
----      ----    -------   ----    ----    ---    ---    ----      ----    ---------  ---------   -------      ------     --------
0001      9335        427   10.0    0-01    01     00     74.0        .1          .06        .8       .015         .005         .020
0002      9335        427   10.0    0-01    01     00     63.7        .1          .02       1.2       .043         .002         .045
0003      9335        427   15.0    0-02    01     00     63.9        .0          .00       1.7       .000         .000         .000
0004      9335        427   15.0    0-02    01     00     64.0        .0          .00       2.1       .000         .000         .000
0005      9335        427   10.0    0-03    01     00     63.9        .0          .01       2.3       .000         .000         .000
0006      9335        427   10.0    0-03    01     00     63.9        .0          .01        .8       .000         .000         .000
0007      9335        427   15.0    0-04    01     00     64.0        .0          .03        .8       .000         .000         .000
0008      9335        427   15.0    0-04    01     00     63.7        .0          .04        .8       .000         .000         .000
Average                                                   65.2        .0          .21       1.0       .000         .000         .000
```

Figure 8-9. Disk Utilization

The columns in the Disk Utilization section are as follows:

Unit Disk arm identifier (unit number). An 'A' or 'B' following the unit number indicates that the disk unit is mirrored. (For example, 0001A and 0001B are a mirrored pair.)

Type Disk type.

Size (M) Disk space capacity in millions of bytes.

IOP Util Percentage of utilization for each input/output processor.

Note: For the multifunction I/O processors, this is utilization due to disk activity only, not communications activity.

IOP ID Input/output processor identification number.

ASP ID Auxiliary storage pool identifier.

CSS ID Checksum set identifier.

Percent Full
Percentage of disk space capacity in use.

Percent Util
Average disk arm utilization (busy). See Figure 10-3 on page 10-6 for a guideline on this utilization. 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.

Op per Second
Average number of disk operations per second.

K per I/O Average number of kilobytes (1024) transferred per disk operation.

Average Service Time
Average disk service time per I/O operation.

Average Wait Time
Average disk wait time per I/O operation.[1]

Average Response Time
Average disk response time per I/O operation.[1]

System Report: Communications Summary

The Communications Summary section of the System Report, as shown in Figure 8-10, shows the use of the communications lines and processors.

```
                                              System Report                                              11/24/90 10:28:19
                                          Communication Summary                                                Page 0006
                                           Sample System Report
Member . . . : TEMP        Model/Serial  . . :  B40/10-15005  Main storage . . :  40.0 M  Started  . . . . : 11/02/90 10:01:54
  Library . . : QPFRDATA   System name . . . :  ABCSYSTEM     Version/Release  :  2/ 1.0  Stopped  . . . . : 11/02/90 10:31:51
                                    Line     Avg  Max    Active          Number                       ------ Bytes per Second -----
BUS/IOP/              Protocol     Speed     Util        Devices      Transactions       Response      Received        Transmitted
  Line                --------    -------    -----       -------      ------------       --------      ---------       -----------
BUS 0  IOP 06 (6110)
  L01ES                  SDLC         9.6     21  65          0                  0             .0          206.8              54.4
  L02ES                  SDLC         9.6     22  64          0                  0             .0          212.0              54.5
  L03ES                  SDLC         9.6     21  51          0                  0             .0          207.9              52.7
BUS 0  IOP 11 (6110)
  RALIN21                SDLC         9.6     44  87          2                  0             .0          109.4             426.7
  RALIN22                SDLC         9.6     39  72          1                  0             .0          104.2             375.4
BUS 2  IOP 02 (6110)
  VMBRIDGE               BSC          9.6      2  44          0                  0             .0           31.1                .2
```

Figure 8-10. Communications Summary

The columns in the Communications Summary section are as follows:

BUS/IOP/Line
The first line lists the BUS ID, IOP ID, and IOP model (in parentheses). Each following indented line lists an attached communciations line ID.

Protocol Line protocol.

- SDLC
- ASYNC
- BSC
- X25
- TRLAN
- ELAN (Ethernet)

[1] These fields are not valid for Version 1, Release 1.0 converted data.

Line Speed
Line speed (1000 bits per second, or bps).

Avg Util The average and maximum percentage of line capacity used during the measured time interval. Use the average percentage of remote line use from this column when referring to the guidelines in Figure 10-3 on page 10-6 for remote line use.

Max Util Consistent use at or above the threshold value given in Figure 10-3 on page 10-6 will affect system performance and cause longer response times or less throughput.

Active Devices
Average number of active devices on the line.

Number Transactions
Number of transactions.

Response Average system response (service) time.

Bytes per Second Received
Average number of bytes received per second.

Bytes per Second Transmitted
Average number of bytes transmitted per second.

Report Selection Criteria

The Report Selection Criteria section, shown in Figure 8-11, gives the selection values you chose to produce the report.

```
                          Report Selection Criteria
Select Parameters

Pools                       - 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16

Jobs                        - 012345/Useridwxyz/Jobname123
                              987654/Useridabcd/Jobname456

User IDs                    - User1       User2       User3       User4      User5      User6
                              User7       User8       User9       User10     User11     User12

Subsystems                  - Subsystem1 Subsystem2 Subsystem3 Subsystem4 Subsystem5 Subsystem6
                              Subsystem7 Subsystem8 Subsystem9 Subsystema Subsystemb Subsystemc

Communications Lines        - Line1       Line2       Line3       Line4      Line5      Line6
                              Line7       Line8       Line9       Line10     Line11     Line12

Control Units               - Ctlr1       Ctlr2       Ctlr3       Ctlr4      Ctlr5      Ctlr6
                              Ctlr7       Ctlr8       Ctlr9       Ctlr10     Ctlr11     Ctlr12

Functional Areas            - Accounting            Payroll               Research
                              Development           ProjectX              MrNolansStaff

                            - No Select parameters were chosen.
```

Figure 8-11. Report Selection Criteria Report: Select Parameters

The SELECT parameters on the Report Selection Criteria section are as follows:

Pools The pools you specify. The values can be from 1 through 16.

Jobs The jobs you specify. The format of the entries is `jobnumber/username/jobname`.

User IDs The user IDs you specify. Each user identification is a 10-character name.

Subsystems
The subsystem names you specify. Each name is a 10-character name.

Communications Lines
The communications line names you specify. Each name is a 10-character name.

Control Units
The controller names you specify. Each name is a 10-character name.

Functional Areas
The functional areas you specify. Each name is a 20-character name. For more information on functional areas, see Chapter 12, "Managing the Performance Tools Configuration."

If you use no SELECT parameters, the message `No SELECT parameters were chosen` appears.

The OMIT parameters on the Report Selection Criteria section of the System Report are shown in Figure 8-12.

```
Omit Parameters

Pools                  - 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16

Jobs                   - 012345/Useridwxyz/Jobname123
                         987654/Useridabcd/Jobname456

User IDs               - User1      User2      User3      User4     User5      User6
                         User7      User8      User9      User10    User11     User12

Subsystems             - Subsystem1 Subsystem2 Subsystem3 Subsystem4 Subsystem5 Subsystem6
                         Subsystem7 Subsystem8 Subsystem9 Subsystema Subsystemb Subsystemc

Communications Lines   - Line1      Line2      Line3      Line4     Line5      Line6
                         Line7      Line8      Line9      Line10    Line11     Line12

Control Units          - Ctlr1      Ctlr2      Ctlr3      Ctlr4     Ctlr5      Ctlr6
                         Ctlr7      Ctlr8      Ctlr9      Ctlr10    Ctlr11     Ctlr12

Functional Areas       - Accounting          Payroll             Research
                         Development         ProjectX            MrNolansStaff

                       - No Omit parameters were chosen.
```

Figure 8-12. Report Selection Criteria: Omit Parameters

The OMIT parameters on the Report Selection Criteria section are as follows:

Pools The pools you specify. The values can be from 1 to 16.

Jobs The jobs you specify. The format of the entries is `jobnumber/username/jobname`

User IDs The user IDs you specify. Each user identification is a 10-character name.

Subsystems
The subsystem names you specify. Each name is a 10-character name.

Communications Lines
The communications line names you specify. Each name is a 10-character name.

Control Units
The controller you specify. Each name is a 10-character name.

Functional Areas
The functional areas you specify. Each name is a 20-character name.

If you did not use OMIT parameters, the message `No OMIT parameters were chosen` appears.

System Model Parameters

If you specify the Measured Profile (MSRPRF) parameter on the Print System Report (PRTSYSRPT) command, or use the *Measured profile* field on the Specify Report Options menu, the System Model Parameters section of the System Report is produced.

The System Model Parameters section lists the contents of the measured profile, which serves as the input to the capacity planning support described in Chapter 10, "Capacity Planning and Performance Prediction." To use the data on the System Model Parameters section as input to the capacity planner (MDLSYS), refer to the profile name shown on the System Model Parameters section. See Chapter 10 for additional information on the use of measured profiles.

The System Model Parameters section of the System Report is separated into the following parts:

- Configuration
- Disk Detail
- Transaction Summary
- Batch/Spool Summary
- Disk Summary
- Display Station Summary
- Interactive Pool Summary

Not all the values shown on the System Model Parameters section are measured. The values fall into these categories:

- Measured values
- Estimated values (based on measured values)
- Modeled values (projections based on measured and estimated values)

Because some values on the System Model Parameters section are not measured, it is important to use your judgment when you select the data for a measured profile. For example, if you select measurement intervals that are not next to each other, or contain diverse types of workloads, the estimated and modeled values contained in the measured profile might not be valid. If you then use this measured profile within the capacity planner, the planning results provided may not be correct.

Figure 8-13 on page 8-23 shows the configuration part of the System Model Parameters section of the System Report.

```
Configuration:
 CPU Model . . . . . . . . . . . . . . . . . . . . . :        B50
 Main storage (MB) . . . . . . . . . . . . . . . . . :         36 MB
 Checksums . . . . . . . . . . . . . . . . . . . . . :          2 1=Active  2=Not active
 Disk mirroring. . . . . . . . . . . . . . . . . . . :          2 1=Active  2=Suspended
 Disk IOPs . . . . . . . . . . . . . . . . . . . . . :          2
 Local workstation IOPs. . . . . . . . . . . . . . . :          1
```

Figure 8-13. System Model Parameters: Configuration

The configuration columns in the System Model Parameters section are as follows:

CPU Model
The processing unit model number.

Main storage (MB)
The total main storage size, as measured in megabytes (1024 squared).

Checksums
Indicates whether the checksums are active.

Disk mirroring
Indicates whether disk mirroring is active.

Disk IOPs The number of disk IOP controllers.

Local work station IOPs
The number of local work station IOPs.

Figure 8-14 shows the disk detail part of the System Model Parameters section of the System Report.

```
Disk Detail:
--------- IOP --------        Number I/Os     ------------- Disk -----------
 Count  Feature   Util        per Second      Feature     Arms    Controllers
  01      2615     .01                 1        9332         2           0
  02      2615     .01                 1        9332         4           0
```

Figure 8-14. System Model Parameters: Disk Detail

The disk detail columns in the System Model Parameters section are as follows:

IOP Count
The relative disk IOP number.

IOP Feature
The disk IOP feature number.

IOP Util The fraction of the time interval the disk IOP was performing I/O operations.

Note: For the multifunction IOPs, this is only the IOP utilization due to disk activity.

Number I/Os per Second
The number of I/Os per second for this particular IOP.

Disk Feature
The type of disk (9332, 9335, and so on).

Disk Arms
The number of disk arms for this IOP.

Disk Controllers
The number of disk storage controllers for this IOP.

Figure 8-15 shows the transaction summary part of the System Model Parameters section of the System Report.

```
Transaction Summary:
 Transactions per hour (local) . . . . . . . . . . :     190 *
 Transactions per hour (remote). . . . . . . . . . :       0 *
 Average response time (seconds) . . . . . . . . . :     2.3
 CPU seconds per transaction . . . . . . . . . . . :   .3023
 Synchronous disk I/O per transaction. . . . . . . :     8.1
 Asynchronous disk I/O per transaction . . . . . . :     1.4
 Permanent writes per transaction. . . . . . . . . :       1
 System disk I/O per transaction . . . . . . . . . :    29.2 *
 System CPU per transaction (seconds). . . . . . . :   .3497 *
 Active display stations (locals). . . . . . . . . :       1 *
 Active display stations (remotes) . . . . . . . . :       0 *
 Exceptional Wait
  Constant (seconds) . . . . . . . . . . . . . . . :     .00 **
  Variable (seconds) . . . . . . . . . . . . . . . :    1.79 **

*  These values are not measured.  They are estimates based on
   the other measured values.

** These values are calculated by the model.
```

Figure 8-15. System Model Parameters: Transaction Summary

The transaction summary columns in the System Model Parameters section are as follows:

Transactions per hour (local)
The interactive transactions per hour attributed to local display stations.

Transactions per hour (remote)
The interactive transactions per hour attributed to remote display stations.

Average response time (seconds)
The average interactive response time.

CPU seconds per transaction
The average processing unit seconds per transaction.

Synchronous disk I/O per transaction
The average number of synchronous physical disk I/O (perations per interactive transaction.

Asynchronous disk I/O per transaction
The average number of asynchronous physical disk I/O operations (er interactive transaction.

Permanent writes per transaction
The average number of permanent write operations per interactive transaction. Permanent writes exclude checksum writes.

System disk I/O per transaction
The total number of physical disk I/O operations attributed (o the system per interactive transaction.

System CPU per transaction (seconds)
The average number of system processing unit seconds per interactive transaction.

Active display stations (locals)
The number of local display stations entering transactions during the measurement period.

Active display stations (remotes)
The number of remote display stations entering transactions during the measurement period.

Exceptional wait
The average exceptional wait time, in seconds, per transaction.

Note: This is a calculated value. If the sum of the constant and variable wait time is greater than one second, you should run STRPFRMON measurements with trace data collection and compare the measured exceptional wait value, which PRTTNSRPT provides, with this calculated value. If the values are significantly different, use the value from PRTTNSRPT, dividing it equally between constant and variable wait time.

Constant The portion of exceptional wait time held constant as throughput increases.

Variable The portion of exceptional wait time that varies as throughput increases.

Figure 8-16 shows the batch/spool summary part of the System Model Parameters section of the System Report.

```
Batch/Spool summary:
 Batch synchronous I/O per second. . . . . . . . :       .4
 Batch asynchronous I/O per second . . . . . . . :       .2
 Batch permanent writes per second . . . . . . . :       .3
 Batch CPU seconds per I/O . . . . . . . . . . . :    .0085
 Number of batch jobs. . . . . . . . . . . . . . :        2
 Batch impact factor . . . . . . . . . . . . . . :   4.9651 **
 Spool I/O per second. . . . . . . . . . . . . . :       .1
 Spool database reads per second . . . . . . . . :      .00
 Spool CPU seconds per I/O . . . . . . . . . . . :    .0069

** These values are calculated by the model.
```

Figure 8-16. System Model Parameters: Batch/Spool Summary

The batch/spool summary columns in the System Model Parameters section are as follows:

Batch synchronous I/O per second
The average number of synchronous physical disk I/O operations per second of batch processing.

Batch asynchronous I/O per second
The average number of asynchronous physical disk I/O operations per second of batch processing.

Batch permanent writes per second
The average number of permanent write operations per second of batch processing. Permanent writes exclude checksum writes.

Batch CPU seconds per I/O
The average number of system processing unit seconds used by all batch jobs for each I/O performed by a batch job.

Number of batch jobs
The average number of active batch jobs. A batch job is considered active if it averages at least one I/O per 5 minutes.

Batch impact factor
Batch workload adjustment for modeling purposes.

Spool I/O per second
The average number of physical disk I/O operations per second of spool processing.

Spool database reads per second
The average number of read operations to database files per second of spool processing.

Spool CPU seconds per I/O
The average number of system processing unit seconds used by all spool jobs for each I/O performed by a spool job.

Figure 8-17 shows the disk summary part of the System Model Parameters section of the System Report.

```
Disk Summary:
 Disk utilization. . . . . . . . . . . . . . . . :      .01 *
 Disk transfer size (KB) . . . . . . . . . . . . :     1.10
 Ratio of write disk I/O to total disk I/O . . . :      .58

* These values are not measured.  They are estimates based on
  other measured values.
```

Figure 8-17. System Model Parameters: Disk Summary

The disk summary columns in the System Model Parameters section are as follows:

Disk utilization
The fraction of the time interval that the disk arms were performing I/O operations.

Disk transfer size (KB)
The average number of kilobytes transferred per disk operation.

Ratio of write disk I/O to total disk I/O
The fraction of the total disk activity that is due to writing data to the disks.

Figure 8-18 shows the display station summary part of the System Model Parameters section of the System Report.

```
Display Station Summary:
 Total fields per transaction. . . . . . . . . . :       75 **
 Total characters per transaction. . . . . . . . :      644 **
 Local workstation IOP utilization . . . . . . . :      .06 *

* These values are not measured.  They are estimates based on
  the other measured values.

** These values are calculated by the model.
```

Figure 8-18. System Model Parameters: Display Station Summary

The display station summary columns in the System Model Parameters section are as follows:

Total fields per transaction
The average number of display station fields either read from or written to per interactive transaction.

Total characters per transaction
The average number of characters either read from or written to display station screens per interactive transaction.

Local work station IOP utilization
The fraction of the time interval the work station I/O processors are busy.

Figure 8-19 shows the interactive pool summary part of the System Model Parameters section of the System Report.

```
Interactive pool summary:
 Pool size (KB). . . . . . . . . . . . . . . . . . . :   24,450 KB
 Percent transactions (purge no) . . . . . . . . . :       16
 Percent transactions (purge yes). . . . . . . . . :       84
 Percent transactions (dynamic no) . . . . . . . . :       84
 Activity level. . . . . . . . . . . . . . . . . . . :       40
 Working set size (KB) . . . . . . . . . . . . . . :      521 **
 Over commitment ratio . . . . . . . . . . . . . . :     .021 **

** These values are calculated by the model.
```

Figure 8-19. System Model Parameters: Interactive Pool Summary

The interactive pool summary columns in the System Model Parameters section are as follows:

Pool size (KB)
The total size in kilobytes of all pools that incurred interactive job activity.

Percent transactions (purge no)
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)
A measure of system main storage utilization. The percent of all interactive transactions that were done with the purge attribute of YES.

Percent transactions (dynamic no)
A measure of system main storage utilization. The percent of all interactive transactions that were done with the purge attribute of dynamic NO.

Activity level
The sum of activity levels for all interactive pools that had interactive job activity running in them.

Working set size (KB)
An estimate of the main storage size, in kilobytes, required by a single interactive job running the transactions described by this report.

Over commitment ratio
The main storage overcommitment ratio (OCR).

Component Report

To produce the Component Report, use the Print Component Report (PRTCPTRPT) command, or select option 2 (Component report) on the Print Performance Reports display.

This series of reports, like the System Report, is also produced from the sample data you collect with the STRPFRMON command. It expands on the detail for each component of system performance shown on the System Report. Data is shown for each sample interval, or in some cases, for each job. Because the report can be lengthy, you may want to use the Select Time Intervals display when requesting this report to select only those measurement intervals of interest to you.

The Component Report consists of the following sections:

- Component Internal Activity
- Job Workload Activity
- Storage Pool Activity
- Disk Activity
- IOP Utilizations
- Local Work Stations — Response Time Buckets
- Exception Occurrence Summary and Interval Counts
- Report Selection Criteria

Each page of the Component Report shows the header information in Figure 8-20.

```
                                                  Component Report                                              04/20/90 10:06:31
                                              Component Report Section                                                 Page 0001
                                              Sample Component Report

Member . . . : GOODSTUF    Model/Serial  . . :  B60/10-12883   Main Storage . . :    96.0 M  Started . . . . :    3/29/90 11:35:54
  Library . . : QPFRDATA   System name . . . :  RCH38366       Version/Release  :    2/ 1.0  Stopped . . . . :    3/29/90 13:35:39
```

Figure 8-20. Component Report: Header Information

This header information consists of:

- Report title
- Current date and time
- Report page number
- User-selected report title
- Data member name
- Library name
- Model number
- Serial number
- System name
- Main storage size
- OS/400 version and release level
- Data collection start date and time
- Data collection stop date and time

The remainder of this section shows you the formats of the various reports you can produce and provides an explanation of the columns in these reports.

Component Interval Activity

The Component Interval Activity section of the Component Report, shown in Figure 8-21, gives the use of the processing unit, disks, and pools at various time intervals.

```
                                   Component Report                                        04/20/90 10:06:31
                               Component Interval Activity                                         Page 0002
                                Sample Component Report

Member . . . : GOODSTUF   Model/Serial  . . :  B60/10-12883  Main Storage . . :   96.0 M  Started . . . . :   3/29/90 11:35:54
  Library . . : QPFRDATA  System name . . . :  RCH38366      Version/Release  :   2/ 1.0  Stopped . . . . :   3/29/90 13:35:39
```

Itv End	Tns /Hour	Rsp /Tns	CPU Utilization Total	CPU Utilization Inter	CPU Utilization Batch	Disk I/O Per Second Sync	Disk I/O Per Second Async	High Utilization Disk	High Utilization Unit	Pool Faults/Sec Mch	Pool Faults/Sec User	Pool Faults/Sec ID	Excp
11:40	1,282	.0	26.0	24.7	1.3	4.9	.0	2	0003	0	4	03	1,401
11:45	1,216	.0	7.7	5.8	1.9	16.7	.0	2	0016	0	2	03	2,598
11:50	903	.1	19.5	18.6	.9	4.8	.0	5	0020	0	10	03	520
11:55	794	.2	2.8	2.5	.3	1.3	.0	2	0023	0	1	03	82

Figure 8-21. Component Interval Activity

The columns in the Component Interval Activity section are as follows:

Itv End Interval end time (hour and minute).

Tns/Hour Number of interactive transactions per hour.

Rsp/Tns Average interactive transaction response time in seconds.

Total CPU Utilization
Percentage of available processing unit time used by interactive jobs, batch jobs, all system jobs, and licensed internal code tasks. For a multiple-processor system this is the average use across all processors.

Inter CPU Utilization
Percentage of available processing unit time used by the following types of jobs:

- Interactive
- Multiple requester terminal (MRT)
- System/36 environment interactive
- Pass-through
- Target distributed data mangement (DDM) servers
- PC Support servers

Note: For a multiple-processor system, this is the average use across all processors.

Batch CPU Utilization
Percentage of available CPU time used by the following types of jobs:

- Batch
- Autostart
- Evoke
- SCPF (Start CPF), spool reader/writer

Note: For a multiple-processor system, this is the average use across all processors.

Sync Disk I/O per Second
Average synchronous disk I/O operations per second.

Async Disk I/O per Second
Average asynchronous disk I/O operations per second.

High Utilization Disk
Percent of utilization of the most utilized disk arm during this interval.

High Utilization Unit
Disk arm that had the most utilization during this interval.

Pool Mch Faults/Sec
Average number of machine pool page faults per second.

Pool User Faults/Sec
Average number of user pool page faults per second, for the user pool with highest fault rate during this interval.

Pool ID Faults
User pool that had the highest page fault rate.

Excp
Total number of program exceptions that occurred (see "Exception Occurrence Summary and Interval Counts" on page 8-37).

Job Workload Activity

The Job Workload Activity section of the Component Report, shown in Figure 8-22, 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 EAO exceptions, the number of PAG faults, the number of arithmetic overflows, and the number of permanent writes for each job.

```
                                              Component Report                                        04/20/90 10:06:31
                                            Job Workload Activity                                           Page 0002
                                           Sample Component Report

Member . . . : GOODSTUF     Model/Serial  . . :  B60/10-12883   Main Storage . . :    96.0 M  Started . . . . :    3/29/90 11:35:54
  Library . . : QPFRDATA    System name . . . :  RCH38366       Version/Release  :    2/ 1.0  Stopped . . . . :    3/29/90 13:35:39
```

Job Name	User Name	Job Number	Typ	Pl	Pty	CPU Util	Tns	Tns /Hour	Rsp	Disk I/O Sync	Disk I/O Async	Disk I/O Logical	Cmn I/O	EAO Excp	PAG Fault	Arith Ovrflw	Perm Write
C12TS1	CLASS23	000492	I	03	20	.4	81	40	.48	0	0	123	0	0	0	0	0
C13T	QSECOFR	000379	I	03	20	.0	0	0	.00	0	0	0	0	0	0	0	0
C13TS1	CLASS33	000493	I	03	20	1.6	185	556	.32	137	0	1293	0	0	0	0	0
C14T	QSECOFR	000382	I	03	20	.0	3	2	.66	399	0	0	0	0	0	0	0
C14TS1	CLASS43	000494	I	03	20	1.1	39	408	.23	141	0	0	0	0	0	0	0
.																	
.																	
.																	
Total/Average						.1	2580	1298	.86	7263	0	51269	0	0	0	0	0

Figure 8-22. Job Workload Activity

The columns in the Job Workload Activity section are as follows:

Job Name
Name of the job.

User Name
Name of the user.

Job Number
The number of the job.

Typ
Type of job.

Possible job type values include the following:

A Autostart
B Batch
C PC Support server

D Distributed data management (DDM) server
E Evoke
I Interactive
L License internal code task
M Subsystem monitor
P Pass-through
R Spool reader
S System
T Multiple requester terminal (MRT)
W Spool writer
X Start system job
3 System/36

Pl Pool that the job ran in.

Pty Priority of the job.

CPU Util Percentage of available processing unit time used by the job. For a multiple-processor system, this is the total use divided by the number of processors.

Note: This processing unit use is calculated by dividing the number of CPU seconds used by the job by the number of seconds the job is actually active. The sum of these utilizations does not necessarily equal the average shown on the Component Interval Activity Report since these utilizations are an average over an entire measurement interval. In addition, the sum of these utilizations does not equal those shown on the System Report since these utilizations are an average over an entire measurement interval.

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).

EAO Excp Number of effective address overflow exceptions.

PAG Fault
Total number of times the program access group (PAG) was referred to, but was not in main storage.

Arith Ovrflw
Number of arithmetic overflow exceptions.

Perm Write
Number of permanent writes.

Storage Pool Activity

The Storage Pool Activity section of the Component Report, shown in Figure 8-23, gives detailed information for each storage pool. This information includes the storage pool activity level, as well as the number of transactions processed in each pool.

```
                                        Component Report                                         04/20/90 10:06:31
                                      Storage Pool Activity                                             Page 0005
                                     Sample Component Report

Member  . . . : GOODSTUF   Model/Serial  . . :  B60/10-12883   Main Storage . . :    96.0 M  Started . . . . :   3/29/90 11:35:54
  Library . . : QPFRDATA   System name . . . :  RCH38366       Version/Release  :    2/ 1.0  Stopped . . . . :   3/29/90 13:35:39

Pool identifier . . . : 01

         Pool                                 Avg                  ---------- Avg per Second -----------   ----- Avg per Minute -----
 Itv     Size      Act        Total          Rsp         CPU      ------ DB ------   ----- Non-DB -----   Act-      Wait-      Act-
 End     (KB)      Level      Tns            Time        Util     Faults    Pages    Faults     Pages     Wait      Inel       Inel
-----   ------     -----     ------          ----        ----     ------    -----    ------     -----     -----     -----      -----
11:40     3000       7           0             .0         2.8        .0        0        .2         1        5         0          0
11:45     3000       7          32            1.4         2.8        .0        0        .0         0        6         0          0
11:50     3000       7           9             .3         4.5        .0        0        .2         1       14         0          0
11:55     3000       7          17             .5         2.1        .0        0        .0         0        7         0          0
```

Figure 8-23. Storage Pool Activity

The POOL IDENTIFIER, shown at the top of the Storage Pool Activity section, specifies the storage pool identifier (the value can be from 1 through 16). A separate Storage Pool Activity section exists for each pool that was in use during the measurement period and was selected on the PRTCPTRPT command.

The columns in the Storage Pool Activity section are as follows:

Itv End Interval end time (hour and minute).

Pool Size (KB)
Initial pool size in kilobytes (1024).

Act Level Initial pool activity level.

Total Tns Number of transactions processed in this pool.

Avg Rsp Time
Average transaction response time.

CPU Util Percentage of available processing unit time used. For a multiple-processor system, this is the total use divided by the number of processors.

DB Faults Pool database faults per second.

DB Pages Pool database pages per second.

Non-DB Faults
Pool nondatabase faults per second.

Non-DB Pages
Pool nondatabase pages per second.

Act-Wait Number of transitions per minute from active state to wait state by processes assigned to this pool.

Wait-Inel Number of transitions per minute from wait state to ineligible state by processes assigned to this pool.

Act-Inel Number of transitions per minute from active state to ineligible state by processes assigned to this pool.

Disk Activity

The Disk Activity section of the Component Report, shown in Figure 8-24, gives the average disk activity per hour, and the disk capacity for each disk.

```
                                                         Component Report                                        04/20/90 10:06:31
                                                           Disk Activity                                                Page 0009
                                                      Sample Component Report

Member . . . : GOODSTUF     Model/Serial  . . :  B60/10-12883   Main Storage . . :     96.0 MB  Started . . . . :   10/31/90 09:56:12
  Library . . : QPFRDATA    System name . . . :  RCH38366       Version/Release  :     2/ 1.0   Stopped . . . . :   10/31/90 11:28:55

                                               --------- Average Disk Activity Per Hour --------
                  Srv      ---- Buffer ---     ------------ Disk Arm Seek Distance -------------      --------- Disk Capacity ---------
Unit     Util     Time     Over       Under          0      1/12       1/6      1/3      2/3     >2/3     Permanent        MB        Percent
----     ----     ----     -----      -----    -------   -------   -------   ------   ------   ------     ---------    -------     -------
0001      2.9     .027       0          0          226       782       174      730      353       24             0         13         6.5
0002      2.7     .023       0          1          286       660       239      312      224       52             0         91        45.5
0003      2.4     .024       0          1          243       772       333       99      335       18             0         90        45.0
0004      1.5     .028       0          0          118       454       237       66      351       25             0         91        45.5
0005      2.5     .026       0          0          210       567       294      150      467       32             0         91        45.5
0006      1.9     .026       1          1          149       462       244      201      335       49             0         91        45.5
                                                                         .
                                                                         .
                                                                         .
                                                                         .
Total/    2.3     .025       1          3        1,235     3,699     1,524    1,562    2,067      202             0        467        38.9
Average
```

Figure 8-24. Disk Activity

The columns in the Disk Activity section are as follows:

Unit Disk arm identifier. An A or B following the unit number indicates that the disk unit is mirrored. (For example, 0001A and 0001B are a mirrored pair.)

Util Drive utilization.

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.

Srv Time Average service time per request in seconds.

Buffer Over
Number of buffer overruns.

Buffer Under
Number of buffer underruns.

Disk Arm Seek Distance
Average seek distance distributions per hour:

0 Number of zero seeks

1/12 Number of seeks between 0 and 1/12 of the disk

1/6 Number of seeks between 1/12 and 1/6 of the disk

1/3 Number of seeks between 1/6 and 1/3 of the disk

2/3 Number of seeks between 1/3 and 2/3 of the disk

>2/3 Number of seeks greater than 2/3 of the disk

Disk Capacity
Average amount of disk space used or available.

Permanent Amount of disk space used for checksum protection.

MB Millions of bytes available on the disk.

Percent Percent of space available on the disk.

Input/Output Processor (IOP) Utilizations

The IOP Utilizations section of the Component Report, shown in Figure 8-25, gives the input/output processor (IOP) utilization for communications, direct access storage devices (DASDs), multifunction (DASD and communication), and local work stations. See Figure 10-3 on page 10-6 for a list of values. Consistent utilization, at or above the threshold value provided in Figure 10-3, of the DASD IOP and multifunction IOP will affect system performance and cause longer response times or less throughput.

Component Report
IOP Utilizations
Sample Component Report

04/20/90 10:06:31
Page 0011

Member . . . : GOODSTUF	Model/Serial . . : B60/10-12883	Main Storage . . :	96.0 M	Started . . . . :	3/29/90 11:35:54
Library . . : QPFRDATA	System name . . . : RCH38366	Version/Release :	2/ 1.0	Stopped . . . . :	3/29/90 13:35:39

Communications IOP's	Utilization	--- OPSTART Msg --- Reverse	Normal	--- Bytes Transmitted --- IOP	System	Restart Queues	BNA Received	Available Storage
BUS 0 IOP 05	2.0	0	573	124,141	575,176	0	0	3,835,536
BUS 1 IOP 02	1.0	0	24	0	4,224	0	0	9,999,999
BUS 2 IOP 02	2.0	0	1,142	771,556	345,574	0	0	9,999,999

DASD IOPs	Utilization	Ops per Sec
BUS 0 IOP 01	1.0	1
BUS 0 IOP 09	2.0	2
BUS 1 IOP 01	1.0	0

Multi-function IOPs	Utilization
BUS 2 IOP 01	2.0

Local Work Station IOPs	Util	Controller Name	-- OPSTART Msg -- Reverse	Normal	Bytes Transmitted IOP	System	Restart Queues	BNA Received	----- Queue Average ----- Wait	Suspend	Active	Twinaxial Util
BUS 0 IOP 03	15.0	CTL01	0	2597	396727	1853723	0	0	7.0	95.0	5.0	7.0
BUS 0 IOP 04	14.0	CTL02	0	2255	790955	2690092	0	0	34.0	64.0	12.0	33.0

Figure 8-25. IOP Utilizations

The columns in the IOP Utilizations section are as follows:

Communications IOPs
ID for communications IOP.

Utilization Percentage of utilization of each communications IOP.

OPSTART Msg Reverse
OPSTART bus unit message received from another bus unit using reverse flow method 2 (always 0).

OPSTART Msg Normal
OPSTART bus unit message received from another bus unit using normal flow.

Bytes Transmitted IOP
Total bytes transmitted from an IOP to the system across the bus.

Bytes Transmitted System
Total bytes transmitted to the IOP from the system across the bus.

Restart Queues
Restart queues sent.

BNA Received
Occurrences of Buffer-not-available message received.

Available Storage
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.

DASD IOPs
ID for each DASD IOP.

Utilization Percentage of utilization of each DASD IOP.

Ops per Sec
Disk operations per second.

Multifunction IOPs
ID for multifunction IOP.

Utilization Percentage of utilization of each multifunction IOP.

Local Work Station IOPs
ID for each local work station IOP.

Util Percentage of utilization of each local work station IOP.

Controller Name
Control unit description name, as specified on the Create Local Work Station Controller command.

OPSTART Msg Reverse
OPSTART bus unit message received from another bus unit using reverse flow method 2.

OPSTART Msg Normal
OPSTART bus unit message received from another bus unit using normal flow.

Bytes Transmitted IOP
Total bytes transmitted from the IOP to the system across the bus.

Bytes Transmitted System
Total bytes transmitted to the IOP from the system across the bus.

Restart Queues
Restart queues bus unit message set to another bus unit.

BNA Received
Occurrences of Buffer-not-available message received.

Queue Average Wait
Wait-for-I/O queue count. The average number of I/O requests on the wait-for-I/O queue at sample time. The number of requests on this queue is checked, or sampled, approximately every 1/2 second. The wait-for-I/O queue holds I/O requests that are being processed or waiting to be processed.

Queue Average Suspend
The average number of elements on the suspend queue at sample time. The number of elements on this queue is checked, or sampled, approximately every 1/2 second.

Queue Average Active
The average number of elements on the active queue at sample time. The number of elements on this queue is checked, or sampled, approximately every 1/2 second. The active queue holds I/O requests

that were sent from the host system and were not yet sent to the wait-for-I/O queue.

Twinaxial Util
Twinaxial utilization. The number of times when the wait-on-I/O queue was sampled and the count was not zero (I/O in progress) divided by the sample count.

Local Work Stations

The Local Work Stations section of the Component Report, shown in Figure 8-26, gives the utilization of each controller, and also shows the range of response times for each device.

```
                                          Component Report                                              04/20/90 10:06:31
                                 Local Work Stations - Response Time Buckets                                   Page 0012
                                          Sample Component Report

Member . . . : GOODSTUF    Model/Serial  . . :  B60/10-12883  Main Storage . . :   96.0 M  Started . . . . :   3/29/90 11:35:54
  Library . . : QPFRDATA    System name . . . :  RCH38366      Version/Release  :   2/ 1.0  Stopped . . . . :   3/29/90 13:35:39

  Ctl/Device       Util     Bus    IOP
---------------    ----     ---    ---
CTL01              15.0      0      3
                                          0-001      001-002    002-004    004-008      >008
                                          -----      -------    -------    -------     -----
      WSA01                                  87           3          3          3          7
      WSA02                                 121           4          0          2          7
      WSA03                                   6           1          0          0          0
      WSA04                                  27           3          1          1          1
                                                                     .
                                                                     .
                                                                     .
                                                                     .
Total Responses                           2,093         117         65         28         42
```

Figure 8-26. Local Work Stations – Response Time Buckets

The columns in the Local Work Stations section are as follows:

Ctl Controller identifier.

Device Device identifier.

Util Controller utilization.

Bus The bus number the controller is attached to.

IOP The IOP number the controller is attached to.

0-001 Number of response times in this range (range presented in seconds).

001-002 Number of response times in this range (range presented in seconds).

002-004 Number of response times in this range (range presented in seconds).

004-008 Number of response times in this range (range presented in seconds).

>008 Number of response times in this range.

Total Responses
The total number of transactions counted by this controller for the report period.

The response time intervals are set up on the system STRPFRMON command. The values for the response times may vary depending on the values you use in the command.

Exception Occurrence Summary and Interval Counts

The Exception Occurrence Summary and Interval Counts section of the Component Report, shown in Figure 8-27, gives the number of exceptions that occurred, and the frequency of these exceptions.

It should be noted that in some cases these exception counts can be high even under normal system operation.

Component Report — 04/20/90 10:06:31
Exception Occurrence Summary and Interval Counts — Page 0013
Sample Component Report

Member . . . : GOODSTUF	Model/Serial . . : B60/10-12883	Main Storage . . :	96.0 M	Started . . . . :	3/29/90 11:35:54
Library . . : QPFRDATA	System name . . . : RCH38366	Version/Release :	2/ 1.0	Stopped . . . . :	3/29/90 13:35:39

Exception Counts

Exception Type	Description	Total
Size	Size	3
Binary Overflow	Binary overflow	0
Decimal Overflow	Decimal overflow	0
Flp Overflow	Floating point overflow	0
Decimal Data	Decimal data	70
Aut Lookup	Authority lookup	7,649
PAG Fault	Process Access Group fault	0
Seize Conflict	Seize conflict	31
Lock Conflict	Lock conflict	50
Verify	Verify	549
EAO Total	Effective Address Overflow total	1,810

Exceptions per Second

Itv End	Size	Binary Overflow	Decimal Overflow	Flp Overflow	Decimal Data	Aut Lookup	PAG Fault	Seize Conflict	Lock Conflict	Verify	EAO Total
11:40	.0	.0	.0	.0	.0	3.1	.0	.8	.1	1.2	3.0
11:45	.0	.0	.0	.0	.0	5.9	.0	.1	.0	.2	.3
11:50	.0	.0	.0	.0	.0	.7	.0	.6	.2	.9	.9
11:55	.0	.0	.0	.0	.0	.1	.0	.0	.1	.1	.4

.
.
.
.

Figure 8-27. Exception Occurrence Summary and Interval Counts

The columns in the Exception Occurrence Summary and Interval Counts section are as follows:

Exception Type

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.

Description

More detailed description of the exception type.

Total Total exception counts for the reporting period.

Itv End Ending time for the sample interval in which STRPFRMON recorded the exception.

Size Decimal data overflow and underflow exceptions per second. An indication of improper field size on numeric calculations.

Binary Overflow
Number of binary overflows per second.

Decimal Overflow
Number of decimal overflows per second.

Flp Overflow
Number of floating point overflows per second.

Decimal Data
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.

Aut Lookup
Number of authority lookups per second.

PAG Fault
Number of faults involving the process access group per second.

Seize Conflict
Number of seize exceptions per second. For more detailed information, you can run the performance monitor with the TRACE(*ALL) option, 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.

Lock Conflict
Number of lock exceptions per second. Database record contention is reflected in this count. For more information, run the performance monitor with the TRACE(*ALL) option and use the PRTTNSRPT and 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.

Verify Number of verify exceptions per second. Verify exceptions occur when a pointer needs to be resolved. When calling by symbolic name, if the name is not resolved, a verify exception results.

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.

EAO Total The count of effective address overflow exceptions per second. Vertical licensed internal code (VLIC) log entries are made if the exception count exceeds the threshold for the system. The threshold is release-dependent. Refer to the *Diagnostic Aids — Volume 1* manual for additional information on the VLIC log. If many VLIC entries are made for this exception, contact your IBM service representative.

Report Selection Criteria

The Report Selection Criteria section of the Component Report, shown in Figure 8-28, gives the selection values you chose to produce the report.

```
                                       Component Report                                             04/20/90 10:06:31
                                   Report Selection Criteria                                                Page 0014
                                    Sample Component Report

Member . . . : GOODSTUF    Model/Serial  . . :  B60/10-12883   Main Storage . . :    96.0 M  Started . . . . :   3/29/90 11:35:54
  Library . . : QPFRDATA    System name . . . :  RCH38366       Version/Release  :    2/ 1.0  Stopped . . . . :   3/29/90 13:35:39

Select Parameters

Pools                      - 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16

Jobs                       - 012345/Useridwxyz/Jobname123
                             987654/Useridabcd/Jobname456

User IDs                   - User1      User2      User3      User4      User5      User6
                             User7      User8      User9      User10     User11     User12

Subsystems                 - Subsystem1 Subsystem2 Subsystem3 Subsystem4 Subsystem5 Subsystem6
                             Subsystem7 Subsystem8 Subsystem9 Subsystema Subsystemb Subsystemc

Communications Lines       - Line1      Line2      Line3      Line4      Line5      Line6
                             Line7      Line8      Line9      Line10     Line11     Line12

Control Units              - Ctlr1      Ctlr2      Ctlr3      Ctlr4      Ctlr5      Ctlr6
                             Ctlr7      Ctlr8      Ctlr9      Ctlr10     Ctlr11     Ctlr12

Functional Areas           - Accounting           Payroll              Research
                             Development          ProjectX             MrNolansStaff

                           - No Select parameters were chosen.

Omit Parameters

Pools                      - 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16

Jobs                       - 012345/Useridwxyz/Jobname123
                             987654/Useridabcd/Jobname456

User IDs                   - User1      User2      User3      User4      User5      User6
                             User7      User8      User9      User10     User11     User12

Subsystems                 - Subsystem1 Subsystem2 Subsystem3 Subsystem4 Subsystem5 Subsystem6
                             Subsystem7 Subsystem8 Subsystem9 Subsystema Subsystemb Subsystemc

Communications Lines       - Line1      Line2      Line3      Line4      Line5      Line6
                             Line7      Line8      Line9      Line10     Line11     Line12

Control Units              - Ctlr1      Ctlr2      Ctlr3      Ctlr4      Ctlr5      Ctlr6
                             Ctlr7      Ctlr8      Ctlr9      Ctlr10     Ctlr11     Ctlr12

Functional Areas           - Accounting           Payroll              Research
                             Development          ProjectX             MrNolansStaff

                           - No Omit parameters were chosen.
```

Figure 8-28. Report Selection Criteria

The SELECT parameters on the Report Selection Criteria Report are as follows:

Pools The pools you specify. Valid values are from 1 through 16.

Jobs The jobs you specify. The format of the entries is:

`jobnumber/username/jobname.`

User IDs The user IDs you specify. Each user ID is a 10-character name.

Subsystems
The subsystem names you specify. Each name is a 10-character name.

Communications Lines
The communications line names you specify. Each name is a 10-character name.

Control Units
The controller names you specify. Each name is a 10-character name.

Functional Areas
The functional areas you specify. Each name is a 20-character name.

If you use no SELECT parameters, the `No Select parameters were chosen` message appears.

The OMIT parameters on the Report Selection Criteria Report are as follows:

Pools The pools you specify. The values can be from 1 through 16.

Jobs The jobs you specify. The format of the entries is: `Jobnumber/username/jobname`.

User IDs The user IDs you specify. Each user identification is a 10-character name.

Subsystems
The subsystem names you specify. Each name is a 10-character name.

Communications Lines
The communication line names you specify. Each name is a 10-character name.

Control Units
The controller names you specify. Each name is a 10-character name.

Functional Areas
The functional areas you specify. Each name is a 20-character name.

If you did not use OMIT parameters, the `No Omit parameters were chosen` message appears.

Transaction Report

To produce the Transaction Report, use the PRTTNSRPT command, or select option 3 (Transaction report) on the Print Performance Reports display. When you use the PRTTNSRPT command, you can choose to print three types of reports using the report type (RPTTYPE) parameter:

- A Job Summary Report (always request this report first)
- A Transaction Report
- A Transition Report

The Transaction and Transition Reports provide detailed information. For this reason, if you choose to print these reports, use the selection values available on the PRTTNSRPT command to select specific jobs, users, or time intervals. In that way you can limit the output to relevant information only.

The PRTTNSRPT command uses the following printer files:

File	Description
QPSPDJS	Job summary report output
QPSPDTS	Transaction report output
QPSPDTD	Transition report output

For information on how to use the PRTTNSRPT command, see the *CL Reference* manual.

The remainder of this section shows you the formats of the various reports you can produce, and provides an explanation of the columns in these reports.

Job Summary Report Option

The Job Summary Report option (RPTTYPE(*SUMMARY)) provides the following sections:

- Job Summary
- System Summary Data
- Distribution of Transactions by CPU/Transaction
- Transaction Significance
- Interactive Transactions by 5 Minute Intervals
- Interactive Throughput by 5 Minute Intervals
- Interactive CPU Utilization by 5 Minute Intervals
- Interactive Response Time by 5 Minute Intervals
- Scatter Diagram of Interactive Transactions by 5-Minute Intervals
- Interactive Program Transaction Statistics
- Summary of Seize/Lock Conflicts by Object
- Report Selection Criteria

The default for the OPTION parameter on the PRTTNSRPT command is *SS. If you leave this default, the following special summary sections print:

- Priority-Jobtype-Pool Statistics
- Job Statistics
- Interactive Program Statistics
- Individual Transaction Statistics
- Longest Seize/Lock Conflicts
- Longest Holders of Seize/Lock Conflicts
- Batch Job Analysis
- Batch Thread Analysis

Transaction Report

The Transaction Report (RPTTYPE(*TNSACT)) provides detailed information about each transaction that occurred in the job:

- Transaction response time
- Name of the program that is active at the time the transaction starts
- Processing unit time use
- Number of I/O requests

To limit the volume of output when you request this report, use the selection parameters on the PRTTNSRPT command to specify specific jobs.

Transition Report

The Transition Report (RPTTYPE(*TRSIT)) provides information similar to that of the Transaction Report, but the data (for example, processing unit time, I/O requests) is shown for each job state transition, rather than just the transitions shown when the job is waiting for work station input. The detail shown in this report helps you to determine the program that ran during a transition, or to determine when an unsatisfied lock request occurred.

Normally, you do not want to print the entire Transition Report because of the volume of output. Instead, you use the PRTTNSRPT command job and time selection parameters to select specific jobs during specific time periods to review.

Job Summary Report

Each page of the Job Summary Report shows the header information in Figure 8-29.

```
                                        Job Summary Report                                               05/17/90  4:06:07
                                          Job Summary                                                          Page 0001

Member  . . . : TJST41      Model/Serial  . . :  B60/10-12883  Main storage . . :   96.0 M  Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2    System name . . . :      RCH38366  Version/Release  :   2/ 1.0  Stopped  . . . . : 05/16/90 15:27:21
```

Figure 8-29. Job Summary Report: Header Information

This header information consists of the following:

- Report title
- Current date and time
- Report page number
- User-selected report title
- Data member name
- Library name
- Model number
- Serial number
- System name
- Main storage size
- OS/400 version and release level
- Data collection start date and time
- Data collection stop date and time

The remainder of this section shows you the formats of the various sections of the Job Summary Report you can produce, and provides an explanation of the columns in these sections.

Job Summary

The Job Summary section of the Job Summary Report shows the following information for each job in the system:

- The name and type of job (for example, interactive, batch)
- The number of transactions in the job
- The average transaction response time
- The average processing unit time per transaction
- The average number of disk I/O requests per transaction
- The number of lock waits
- The number of seize conflicts
- The key+think time per transaction

If the Job Summary section shows jobs that have high response times, high disk I/O activity, high processing unit utilization, or a number of lock requests, use the Transaction Report to investigate further.

An example of the Job Summary section is shown in Figure 8-30.

Job Summary Report 05/17/90 4:06:07
Job Summary Page 0001

Member . . . : TJST41 Model/Serial . . : B60/10-12883 Main storage . . : 96.0 M Started . . . . : 05/16/90 14:48:28
Library . . : TJSWORK2 System name . . . : RCH38366 Version/Release : 2/ 1.0 Stopped . . . . : 05/16/90 15:27:21

Job Name	User Name	*On/Off* Job Number	Pl	Typ	Pty	Prg	Tot Nbr Tns	Response Sec Avg	Response Sec Max	CPU Sec Util	CPU Sec Avg	CPU Sec Max	Average DIO/Transaction Synchronous DBR	NDBR	Wrt	Sum	Max	Async Sum	Async Max	Number Cft Lck	Number Cft Sze	K/T /Tns Sec
SCPF	QSYS	000000	02	X	52							.52					4					
QSYSARB	QSYS	020202	02	S	00							.65					3					
QLUS	QSYS	020203	02	S	20																	
QBASE	QSYS	020205	02	M	00							1.95					42			4		
QSPL	QSYS	020207	02	M	00																	
NET38	QSYS	020363	02	M	00																	
PUTGET	QSECOFR	020365	02	B	15																	
ROUTER	QSECOFR	020366	02	B	15																	
MONITOR	QSECOFR	020367	02	B	15																	
RESTART	QSECOFR	020368	02	B	16																	
SCHWEYER	SCHWEYER	020380	02	I	20																	
INSTALLE	INSTALL	020385	02	I	20																	
INSTALLE	INSTALL	020386	02	I	20																	
DRLIPPS	DRLIPPS	020391	02	I	20	Y	36	.2	3.0	.1	.08	.89		2		2	90					14
DSP01	SFRASER	020398	02	I	20	Y	204	.7	7.0	4.4	.50	6.05		2		2	42			4		11
DLH	DJW	020403	02	I	20	Y	26	.2	.9	.1	.12	.38		2		2	20					9
PRT01	QSPLJOB	020405	03	W	15																	
SFRASER	AMUNDSON	020408	02	I	20	Y	29	.5	2.7	.2	.24	1.03		7		7	103					57
MOREY	MOREY	020410	02	I	20																	
MOREY	MOREY	020412	02	I	20																	
QSNADS	QSYS	020417	02	M	00																	
QZDSTART	QSNADS	020418	02	A	40																	
QROUTER	QSNADS	020419	02	B	40																	
RCHAS374	QSNADS	020420	02	B	40																	
RCHAS374	QGATE	020421	02	B	40																	
RCHAS500	QSNADS	020422	02	B	40																	
RCHVMV	QSNADS	020423	02	B	40																	
RCHVMX2	QSNADS	020424	02	B	40																	
QDIA	QSNADS	020426	02	B	40																	
QDIALOCAL	QSNADS	020427	02	B	19																	
QNFTP	QSNADS	020429	02	B	40																	
QPFRMON	QPGMR	020437	02	B	00					.2		4.78					90					
DSP02	DRLIPPS	*020438*	02	I	20	Y	42	.4	3.2	.2	.14	.55		4	2028	2032	5167	56	304			2
PRTTNSRPT	SFRASER	020439	02	B	50					.6		14.83					311		4			
PRTTNSRPT	SFRASER	020440	02	B	50					.5		12.59					183		304			
QDFTJOBD	SFRASER	020441	02	B	50					29.0		677.63					8		304			
PRTTNSRPT	SFRASER	020442	02	B	50					.5		13.00					195		4			

Figure 8-30. Job Summary Section

The columns in the Job Summary section are as follows:

Job Name
The name of the job that the summary line describes. 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.

User Name
The name of the user who submitted the job.

Job Number
The job number of the job which the summary line describes. An * before the job number indicates the job signed on during the measurement period. An * after the job number indicates the job signed off during the measurement period.

Pl The number of the storage pool the job ran in (2 through 16).

Typ The system job type and subtype. The possible job types are:

- **A** Autostart
- **B** Batch
- **C** Programmable work station application server
- **D** Target distributed data management (DDM) server
- **I** Interactive
- **L** Licensed internal code task
- **M** Subsystem monitor
- **R** Spool reader
- **S** System
- **W** Spool writer
- **X** Start the system

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. Refer to "Changing the Job Type Data" on page 8-85 for information on the CHGJOBTYP command. The possible job subtypes are:

- **T** Multiple-requester terminal (System/36 environment only)
- **E** Evoke (communications batch)
- **P** Print driver job
- **J** Prestart job

Pty The run priority of the job.

Prg The purge attribute of the job.

Tot Nbr Tns
The total number of transactions the PRTTNSRPT program determined from the input data which were accomplished for the job.

Response Sec Avg and Max
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.

CPU Util The percentage of the available processing unit time that was used by the job. For a multiple-processor system, this is the average use across all processors.

CPU Sec Avg and Max
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.

Average DIO/Transaction
The next seven columns provide information about physical disk I/O counts:

Note: 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
The average number of synchronous database read requests per transaction.

Synchronous NDBR
The average number of synchronous nondatabase read requests per transaction.

Synchronous Wrt
The average number of synchronous database and nondatabase write requests per transaction.

Synchronous Sum
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 Max
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.

Async Sum
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).

Async Max
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.

Number Lck Cft
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 PRTLCKRPT command.

Number Sze Cft
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 job that held the object, the name and type of object being held, and what the job was waiting for.

K/T /Tns Sec
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.

In the above descriptions, the statement "If this number is high then..." is used many times. The term "high" is dependent on the application. One example is the number of **lock-waits**. An application that has many users accessing a database at the same time could, under normal conditions, have numerous lock-waits.

You must evaluate each situation individually. If the values are difficult to explain (an application should have very few locks and yet many are reported), then further analysis will be required. The Transaction and Transition Reports can help in this analysis.

System Summary Data

An example of the first set of parts of the System Summary Data section of the Job Summary Report is shown in Figure 8-31.

```
                                               Job Summary Report                                        05/17/90  4:06:07
                                               System Summary Data                                              Page 0002

Member  . . . : TJST41      Model/Serial  . . :  B60/10-12883   Main storage . . :   96.0 M   Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2    System name . . . :      RCH38366   Version/Release  :   2/ 1.0   Stopped  . . . . : 05/16/90 15:27:21

TRACE PERIODS FOR TRACE DATE.
                                       Elapsed
      Started          Stopped         Seconds
      --------         --------       ---------
      14.48.28         15.27.21           2,332

CPU BY PRIORITY FOR ALL JOBS FOR TOTAL TRACE PERIOD.

                          CPU          Cum CPU         CPU
 Pty          CPU         Util           Util           QM
----      ---------      -----        -------       ------
 00          25.353       1.08           1.08         1.01
 15                                      1.08         1.01
 16                                      1.08         1.01
 19                                      1.08         1.01
 20         121.567       5.21           6.29         1.06
 40                                      7.58         1.08
 50         718.033      30.79          38.37         1.62
 52            .519        .02          67.95         3.12
```

Figure 8-31. System Summary Data Section - 1

The Trace Periods for Trace Date part of the System Summary Data section shows the following information:

Started The time of the first record in the trace data, in the form HH.MM.SS (hours, minutes, seconds).

Stopped The time of the last record in the trace data, in the form HH.MM.SS (hours, minutes, seconds).

Elapsed Seconds
The time elapsed (in seconds) for the trace period.

The CPU by Priority for All Jobs for Total Trace Period part on the System Summary Data section shows the following information:

Pty The run priority of the jobs.

CPU The total processing unit seconds used by the jobs with a given priority.

CPU Util The percentage of available processing unit time used by the jobs with a given priority. For a multiple-processor system, this is the average use across all processors.

Cum CPU Util
The cumulative percentage of available processing unit time used by the jobs with a priority higher or equal to the given priority.

CPU QM The simple processing unit queuing multiplier. For an explanation of queuing multiplier, see Appendix B.

An example of the second set of parts of the System Summary Data section of the Job Summary Report is shown in Figure 8-32.

```
                                                   Job Summary Report                                                    05/17/90  4:06:07
                                                  System Summary Data                                                           Page 0003

Member  . . . : TJST41       Model/Serial  . . :  B60/10-12883  Main storage . . :   96.0 M  Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2     System name . . . :      RCH38366  Version/Release  :   2/ 1.0  Stopped  . . . . : 05/16/90 15:27:21

CPU AND DISK I/O PER JOB TYPE FOR ALL JOBS FOR TOTAL TRACE PERIOD.

     Job              Nbr        CPU       CPU    --Disk I/O Requests--    CPU Sec/        Sync I/O
     Type             Jobs     Seconds     Util     Sync        Async      Sync DIO        /Elp Sec         Planning Parameters
---------------       ----     -------     ----   --------    --------     --------        --------         ------------------------
INTERACTIVE             10       121.6      5.2       6156         543        .0000          6855.1         Elapsed Seconds =   2332
BATCH A,B,C,D,X         20       723.3     31.0       2381          25        .0000          7419.5         Tns Selected    =    337
SPOOL WTR/RDR            1          .0       .0          0         999        .0000              .0         TCPU=  .3607   TDIO= 315
SYSTEM JOBS              6         2.6       .1         45          84        .0578              .0         SCPU=  .0610   SDIO=5762
SYSTEM TASKS            84        17.9       .8       1634        1200        .0000          1128.5         BCPU=  .0000   BDIO= 420
---------------       ----     -------     ----   --------    --------     --------        --------         XSUM=          PDIO=   0
   ** TOTALS **        121       865.4     37.1      10216        2851        .0000          6586.7         100.00 Percent Selected

DATA FOR SELECTED TIME INTERVAL (OR TOTAL TRACE PERIOD IF NO TIME SELECTION).

  INTERACTIVE TRANSACTION AVERAGES BY JOB TYPE.

T                                      Avg     CPU/   ----- Sync Disk I/O Rqs/Tns -----  Async   W-I     Excp     Key/   Active  Est
y       Nbr    Nbr     Pct    Tns     Rsp      Tns     DB     DB     NDB    NDB           DIO    Wait    Wait    Think    K/T    Of
p  Prg  Jobs   Tns     Tns   /Hour   (Sec)    (Sec)  Read  Write   Read  Write    Sum   /Tns    /Tns    /Tns     /Tns    /Tns   AWS
-- ---  ----  ------  -----  -----  ------   ------  -----  -----  -----  -----  -----  -----  ------  -------  ------  ------  ---
IM  YES    10    337  100.0    520    .549     .361      0    156      3    156    315    624    .000     .091  13.912  12.044    2

  EXCEPTIONAL WAIT BREAKDOWN BY JOB TYPE.

                         A-I        Short      Short       Seize      Lock       Event      Excs       EM3270     DDM Svr     Other
                         Wait       Wait       WaitX       Wait       Wait       Wait       ACTM       Wait        Wait        Wait
Type         Purge       /Tns       /Tns       /Tns        /Tns       /Tns       /Tns       /Tns       /Tns        /Tns        /Tns
----         -----      ------     ------     ------      ------     ------     ------     ------     ------     --------     ------
 IM           YES         .000       .013       .000        .000       .000       .000       .073       .000        .000        .005
```

Figure 8-32. System Summary Data Section - 2

The CPU and Disk I/O per Job Type for All Jobs for Total Trace Period part on the System Summary Data section shows the following information:

Job Type The type of jobs the summary line is for.

Nbr Jobs The number of jobs.

CPU Seconds
The total available processing unit time used by the jobs during the trace period.

CPU Util The percentage of available processing unit time used by the jobs during the trace period.

Disk I/O Requests
The total number of synchronous and asynchronous disk I/O requests issued by the jobs during the trace period.

CPU Sec /Sync DIO
The average processing unit time in seconds per synchronous disk I/O operation for each type of job.

Sync I/O /Elp Sec
The average number of synchronous disk I/O requests for all jobs, per second of elapsed time used by the jobs.

Planning Parameters
The planning parameters for the jobs selected. The values given for the transactions selected are:

- Total processing unit seconds and disk I/O per transaction (TCPU, TDIO)
- System processing unit seconds and disk I/O per transaction (SCPU, SDIO)
- Batch processing unit seconds and disk I/O per transaction (BCPU, BDIO)
- Checksum I/O and physical disk I/O per transaction (XSUM, PDIO)

The Interactive Transaction Averages by Job Type part of the System Summary Data section shows the following information:

Typ The type and subtype of the jobs

Prg The purge attribute of the jobs.

Nbr Jobs The number of interactive jobs selected for the report.

Nbr Tns The number of transactions.

Note: The values for transaction counts, and other transaction-related information, shown on the reports you produce using the PRTTNSRPT command, may vary from the values shown on the reports you produce using the PRTSYSRPT and PRTCPTRPT commands. These differences are caused because the PRTTNSRPT command uses trace data as input, while the PRTSYSRPT and PRTCPTRPT commands use sample data as input. See Appendix A, "Defining Transaction Boundaries," for additional information.

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.

Pct Tns The percentage of the total transactions within the given trace period with the given purge attribute.

Tns/Hour The hourly rate of transactions based on the length of the run and number of transactions.

Avg Rsp (Sec)
The average transaction response time in seconds.

Cpu/Tns (Sec)
The number of processing unit seconds per transaction.

Sync Disk I/O Rqs/Tns
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).

Async DIO /Tns
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).

W-I Wait/Tns
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.

Excp Wait /Tns
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.

Key/Think /Tns
The average think time and keying time (or the delay time between transaction boundaries), in seconds, for the interactive jobs.

Active K/T /Tns
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.

Est Of AWS
An estimate of the number of active work stations for the trace period. This value is calculated as shown in Figure 8-33.

$$\text{AWS} = \text{TNS/HOUR} \times \frac{(\text{AVGRSP} + \text{ACTIVE KEY/THINK})}{3600}$$

Figure 8-33. Equation for the Estimated Number of Active Work Stations

The Exceptional Wait Breakdown by Job Type part of the System Summary Data section shows the following information:

Type The type and subtype of the jobs.

Purge The purge attribute of the jobs.

A-I Wait /Tns
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.

Short Wait /Tns
The average time, in seconds, of short (active) wait time per transaction. If the value is high, it may be due to the use of DFRWRT(*NO) and RSTDSP(*YES) in the program display files.

Short WaitX /Tns (Short wait extended)
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. The use of DFRWRT(*NO) and/or RSTDSP(*YES) in the display files could be a reason for this value to be high.

Seize Wait /Tns
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.

Lock Wait /Tns
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 PRTLCKRPT command.

Event Wait /Tns
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.

Excs ACTM /Tns
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). In addition, if there are enough activity levels available, and plenty of interactive work of higher priority to do, a job waits more to get its turn at 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 excess activity-level time value is calculated by subtracting from the job's average active time, the sum of the job's active wait time, the product of the job's processing unit time multiplied by the average beginning activity-level value, and the product of the job's number of I/O operations per transaction multiplied by 50 milliseconds. The formula for excessive activity-level time is shown in Figure 8-34.

```
Active Time – [Active Wait + (CPU X Beginning Activity Level) + (Number of I/O X .050)]
```

Figure 8-34. Formula for Excessive Activity-Level Time

EM3270 Wait /Tns
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. Since there are requirements on event-wait processing, not all transition combinations can be detected.

DDM Svr Wait /Tns
The average time, in seconds, that a source distributed data management (DDM) server job has spent waiting for the target system to respond to a request for data per transaction. This value would include line time and time spent by the target system responding to the request for data.

Other Wait /Tns
The average time, in seconds, spent waiting that was not in any of the previous categories per transaction. An example would be time spent waiting during a save/restore operation when the system requested new media (tape or diskette).

An example of the third set of parts of the System Summary Data section of the Job Summary Report is shown in Figure 8-35 on page 8-52.

Job Summary Report 05/17/90 4:06:07
System Summary Data Page 0004

Member . . . : TJST41 Model/Serial . . : B60/10-12883 Main storage . . : 96.0 M Started . . . . : 05/16/90 14:48:28
Library . . : TJSWORK2 System name . . . : RCH38366 Version/Release : 2/ 1.0 Stopped . . . . : 05/16/90 15:27:21

ANALYSIS BY INTERACTIVE TRANSACTION CATEGORIES.

Category	Avg CPU /Tns	CPU Util	Cum CPU Util	----- Sync Disk I/O Rqs/Tns ----- DB Read	DB Write	NDB Read	NDB Write	Sum	Async DIO /Tns	Nbr Tns	Pct Tns	Avg Rsp /Tns	Excp Wait /Tns	Avg K/T /Tns	Est Of AWS
VERY SIMPLE VS	.068	.5								182	54.0	.107	.034	11.344	1
** SIMPLE S	.071	.6	.6			1		1		208	61.7	.124	.037	12.140	1
-Boundary-	.161														
** MEDIUM M	.183	.2	.8			2		2		29	8.6	.302	.067	7.539	
-Boundary-	.219														
** COMPLEX X	1.014	4.3	5.1	1	526	7	526	1060	2104	100	29.7	1.506	.208	13.150	1
VERY COMPLEX VX	5.328	2.1		1		4		5		9	2.7	6.116	.519	5.227	
Total/Avg of **	.361				156	3	156	315	624	337	100.0	.549	.090	12.044	2

ANALYSIS BY INTERACTIVE RESPONSE TIME.

Category	Avg Rsp /Tns	Nbr Tns	Pct Tns	Cum Pct Tns	Avg CPU /Tns	CPU Util	Cum CPU Util	----- Sync Disk I/O Rqs/Tns ----- DB Read	DB Write	NDB Read	NDB Write	Sum	Async DIO /Tns	Excp Wait /Tns	Avg K/T /Tns
Sub-Second	.211	286	84.9	84.9	.132	1.6	1.6			1		1		.039	11.655
1 - 1.999 Sec	1.351	33	9.8	94.7	.852	1.2	2.8	1	1593	4	1593	3191	6372	.270	19.140
2 - 2.999 Sec	2.469	6	1.8	96.5	.710	.2	3.0	9		39		48		.396	5.004
3 - 4.999 Sec	3.501	3	.9	97.4	1.126	.1	3.1			35		35		1.085	5.542
5 - 9.999 Sec	6.116	9	2.7	100.1	5.328	2.1	5.2	1		4		5		.519	5.227
GE 10 Seconds				100.1			5.2								

ANALYSIS BY INTERACTIVE KEY/THINK TIME.

Category	Avg K/T /Tns	Nbr Tns	Pct Tns	Cum Pct Tns	Avg CPU /Tns	CPU Util	Cum CPU Util	----- Sync Disk I/O Rqs/Tns ----- DB Read	DB Write	NDB Read	NDB Write	Sum	Async DIO /Tns	Avg Rsp /Tns	Excp Wait /Tns
LT 2 Seconds	.961	143	42.4	42.4	.264	1.6	1.6		368	2	368	738	1472	.396	.072
2 - 14.999 Sec	5.754	151	44.8	87.2	.469	3.0	4.6	1		4		5		.732	.114
15 - 29.999 Sec	19.772	21	6.2	93.4	.346	.3	4.9			1		1		.470	.069
30 - 59.999 Sec	38.080	7	2.1	95.5	.342	.1	5.0			1		1		.463	.085
60 - 299.999 Sec	122.555	12	3.6	99.1	.239	.1	5.1			1		1		.353	.059
GE 300 Seconds	509.881	3	.9	100.0	.142		5.1			2		2		.246	.044

Figure 8-35. System Summary Data Section - 3

The Analysis by Interactive Transaction Categories part of the System Summary Data section provides a breakdown of the transactions into the categories very simple, simple, medium, and complex categories, relative to their processing unit utilization.

Note: The Total/Avg is only a total or average of the simple, medium, and complex categories. The very simple category is a part of the simple category. The very complex category is a part of the complex category.

These transaction categories depend on the processing unit model. They are introduced here and in some of the following reports as a way to highlight the differences that exist in the work being done on the system.

When you are considering adding new applications, determine the new application's transaction characteristics. For example, determine if a high volume of complex transactions is typical with this new application. By analyzing the transaction characteristics of new applications, you might be able to foresee the need to acquire additional hardware resources for the new application.

If you obtain a new application from a supplier, it is reasonable to ask for information about the application's transaction characteristics.

The Analysis by Interactive Transaction Categories part of the System Summary Data section shows the following information:

Category The category that the transactions fell into based on the processing unit model. The boundary values that were used to separate the transactions are given in the *Avg CPU /Tns* column.

Avg CPU /Tns
The average number of processing unit seconds per transaction that fell in the given category.

CPU Util The percentage of available processing unit time that was used by the transactions in the given category.

Cum CPU Util
The cumulative percentage of available processing unit time used by the transactions.

Sync Disk I/O Rqs/Tns
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).

Async DIO /Tns
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).

Nbr Tns The total number of transactions in a given category.

Pct Tns The percentage of all transactions that fell into a given category.

Avg Rsp /Tns
The average response time per transaction in seconds.

Excp Wait /Tns
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.

Avg K/T /Tns
The average think time and keying time (or the delay time between transaction boundaries), in seconds, for the interactive jobs.

Est Of AWS
An estimate of the number of active work stations for the trace period. 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. This value is calculated as shown in Figure 8-36.

$$AWS = TNS/HOUR \times \frac{(AVGRSP + ACTIVE\ KEY/THINK)}{3600}$$

Figure 8-36. Equation for the Estimated Number of Active Work Stations

The Analysis by Interactive Response Time part of the System Summary Data section, provides transaction information sorted by response time categories and shows the following information:

Category The category that the transactions fell into based on their response time.

Avg Rsp /Tns
The average response per transaction in seconds for the transactions that fell into the given category.

Nbr Tns The total number of transactions in a given category.

Pct Tns The percentage of all transactions that fell into a given category.

Cum Pct Tns
The cumulative percentage of all transactions that have an average response time per transaction equal to or less than the given category.

Avg CPU /Tns
The average number of processing unit seconds per transaction that fell in the given category.

CPU Util The percentage of available processing unit time that was used by the transactions in the given category.

Cum CPU Util
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.

Sync Disk I/O Rqs/Tns
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).

Async DIO /Tns
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).

Excp Wait /Tns
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.

Avg K/T /Tns
The average think time and keying time (or the delay time between transaction boundaries), in seconds, for the interactive jobs.

The Analysis by Interactive Key/Think Time part of the System Summary Data section provides information about the key+think time and shows the following information:

Category The category that the transactions fell into based on their key+think time.

Avg K/T /Tns
The average think time and keying time (or the delay time between transaction boundaries), in seconds, for the interactive jobs.

Nbr Tns The total number of transactions in a given category.

Pct Tns The percentage of all transactions that fell into a given category.

Cum Pct Tns
The cumulative percentage of all transactions that have an average response time per transaction equal to or less than the given category.

Avg CPU /Tns
The average number of processing unit seconds per transaction that fell in the given category.

CPU Util The percentage of available processing unit time that was used by the transactions in the given category.

Cum CPU Util
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.

Sync Disk I/O Rqs/Tns
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).

Async DIO /Tns
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).

Avg Rsp /Tns
The average response per transaction in seconds for the transactions that fell into the given category.

Excp Wait /Tns
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.

Distribution by CPU/Transaction

The Distribution of Transactions by CPU/Transaction section of the Job Summary Report, shown in Figure 8-37 on page 8-57, provides a graphical view of the distribution of simple, medium, and complex transactions. This chart shows the number of transactions versus the processing unit time per transaction in seconds.

```
                                                   Job Summary Report                                        05/17/90  4:06:07
                                       Distribution of Transactions by CPU/Transaction                                Page 0005

Member  . . . : TJST41       Model/Serial  . . :  B60/10-12883  Main storage . . :   96.0 M  Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2     System name . . . :      RCH38366  Version/Release  :   2/ 1.0  Stopped  . . . . : 05/16/90 15:27:21

Number
 Of
 Tns
------
   50+    S
   48+    S
   46+    S
   44+    S
   42+    S
   40+    S
   38+    S
   36+    S
   34+    S S
   32+    S S
   30+    S S
   28+    S S
   26+    S S
   24+    S S
   22+    SSS S
   20+    SSS S
   18+    SSSSS
   16+    SSSSS   S                                                              X
   14+    SSSSS   S                                                              X
   12+   SSSSSS   S                                                              X
   10+   SSSSSSS  Sm                                                             X
    8+   SSSSSSSSSSm  m                                                          X
    6+   SSSSSSSSSSmm mX                                                         X
    4+   SSSSSSSSSSmm mXX   X    X X                        X                    X
    2+   SSSSSSSSSSmm mXX XXX XXXXXXXXX   X  XX             XX X X X             X
    2-   SSSSSSSSSSmmmmXXXXXXXXXXXXXXXXX XX XXXX        X   XX X X X    XX XX    X
        +----+----+----+----+----+----+----+----+----+----+----+----+----+----+ +++
     .000      .150      .300      .450      .600      .750      .900     1.050
                                CPU/Transaction  (Seconds)

Transaction Categories:

     S = Simple Transactions
     m = Medium Transactions
     X = Complex Transactions
```

Figure 8-37. Distribution of Simple, Medium, and Complex Processing Unit Transactions

Transaction Significance

The Transaction Significance section of the Job Summary Report, shown in Figure 8-38 on page 8-58, provides a graphical view of the processing unit use, categorized by simple, medium, and complex transactions. This chart shows the percent of available processing unit time used versus the processing unit time per transaction in seconds.

```
                                                     Job Summary Report                                           05/17/90  4:06:07
                                                   Transaction Significance                                             Page 0006

Member  . . . : TJST41      Model/Serial  . . :  B60/10-12883   Main storage . . :   96.0 M   Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2    System name . . . :      RCH38366   Version/Release  :   2/ 1.0   Stopped  . . . . : 05/16/90 15:27:21

Percent
  CPU
-------
  2.6+                                                                            X
  2.5+                                                                            X
  2.4+                                                                            X
  2.3+                                                                            X
  2.1+                                                                            X
  2.0+                                                                            X
  1.9+                                                                            X
  1.8+                                                                            X
  1.7+                                                                            X
  1.6+                                                                            X
  1.5+                                                                            X
  1.4+                                                                            X
  1.3+                                                                            X
  1.2+                                                                            X
  1.1+                                                                            X
  1.0+                                                                            X
   .9+                                                                            X
   .8+                                                                            X
   .7+                                                                            X
   .6+                                                                            X
   .5+                                                                            X
   .4+                                                                            X
   .3+                                                                            X
   .2+                                                                            X
   .1+            S                                         X  X                  X
   .1-   SSSSSSSSSSmmmmXXXXXXXXXXXXXXXXX XX XXXX       X   XX X X X    XX XX      X
        +----+----+----+----+----+----+----+----+----+----+----+----+----+----+ +++
     .000      .150      .300      .450      .600       .750      .900     1.050
                                CPU/Transaction  (Seconds)

Transaction Categories:

     S = Simple Transactions
     m = Medium Transactions
     X = Complex Transactions
```

Figure 8-38. Percentage of Processing Unit Used by Transaction Categories

Transactions by Intervals

The Interactive Transactions by 5 Minute Intervals section of the Job Summary Report, shown in Figure 8-39 on page 8-59, provides a count of the number of active jobs during a 5-minute interval that performed at least one transaction. It also shows the number of jobs that were signed on and off during the 5-minute intervals. Transaction rates per 5-minute intervals are shown in several different formats.

```
                                          Job Summary Report                                              05/17/90  4:06:07
                                 Interactive Transactions by 5 Minute Intervals                                 Page 0007

Member  . . . : TJST41       Model/Serial  . . :  B60/10-12883   Main storage . . :    96.0 M   Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2     System name . . . :      RCH38366   Version/Release  :    2/ 1.0   Stopped  . . . . : 05/16/90 15:27:21

                             --- Pct Of Tns ---  Pct CPU By  Nbr   Nbr    Sync  Async     Avg     Excp     Pct   Seize  Active  Est
 Itv   Active   Nbr    Tns   --- Categories ---  Categories  Sign  Sign   DIO   DIO       Rsp     Wait     Ex-Wt  Wait   K/T     Of
 End    Jobs    Tns   /Hour  %VS* %S %M %X *%VX  %S %M %X    offs  ons    /Tns  /Tns      /Tns    /Tns     /Rsp   /Tns   /Tns    AWS
-----  ------  -----  -----  ---- --+--+-- ----  ---+--+---  ----  ----  -----  -----  -------  -------  -----  ------  ------  ---
 ***
14.45*      1      9    108   44* 56+00+44  00    00+00+01                  35     28     .550     .124    23            4.333
14.50       4     65    780   52* 66+06+28  05    01+00+08                 999   1999     .664     .110    17           16.569    3
14.55       2     56    672   80* 84+02+14  05    01+00+08                  28    659     .522     .037     7            5.839    1
15.00       2     48    576   52* 54+13+33 *02    01+00+05                 768   7689     .500     .032     6           16.104    2
15.05       3     99   1188   49* 60+11+29 *02    01+01+08      1     1  78999    886     .554     .116    21            3.960    1
15.10       1     15    180   53* 53+07+40 *00    00+00+01                  19     58     .406     .124    31           19.000
15.15       2     17    204   29* 35+06+59 *00    00+00+01                 768   9877     .676     .174    26           32.882    1
15.20       1      8     96   63* 75+12+13 *00    00+00+00                6577    453     .268     .057    21           39.875    1
15.25*      1     20    240   35* 40+20+40 *00    00+00+01                  45     67     .451     .083    18            6.050
     * Denotes Partial Interval Data
```

Figure 8-39. Interactive Transactions by 5 Minute Intervals

The columns in the Interactive Transactions by 5 Minute Intervals section are as follows:

Itv End The interval end time (hour and minute).

Active Jobs
The number of interactive jobs that were active during the interval.

Nbr Tns The number of transactions.

Tns/Hour The hourly rate of transactions based on the length of the run and number of transactions.

Pct Of Tns Categories
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 CPU By Categories
The percentage of available processing unit time used by the the transactions which 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.

Nbr Sign offs
The number of jobs that signed off during the interval.

Nbr Sign ons
The number of jobs that signed on during the interval.

Sync DIO /Tns
The average number of synchronous I/O requests per transaction during the interval.

Async DIO /Tns
The average number of asynchronous I/O requests per transaction during the interval.

Avg Rsp /Tns
The average response time per transaction in seconds.

Excp Wait /Tns
The average exceptional wait time per transaction, in seconds. This value is the sum of those waits listed under the Exceptional Wait

Breakdown by Job Type/Category part of the System Summary Data section.

Pct Ex-Wt /Rsp
The percentage of the response time that is due to exceptional wait.

Seize Wait /Tns
The average seize wait time per transaction in seconds. This is the average amount of time which 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.

Active K/T /Tns
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.

Est Of AWS
An estimate of the number of active work stations for the interval. This value is calculated as shown in Figure 8-40.

$$\text{AWS} = \text{TNS/HOUR} \times \frac{(\text{AVGRSP} + \text{ACTIVE KEY/THINK})}{3600}$$

Figure 8-40. Equation for the Estimated Number of Active Work Stations

Interactive Throughput

The Interactive Throughput by 5 Minute Intervals section of the Job Summary Report, shown in Figure 8-41 on page 8-61, gives simple, medium, and complex transactions relative to the number of transactions according to an interval end time.

```
                                        Job Summary Report                                          05/17/90  4:06:07
                              Interactive Throughput by 5 Minute Intervals                               Page 0008

Member  . . . : TJST41       Model/Serial  . . :  B60/10-12883  Main storage . . :   96.0 M  Started   . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2     System name . . . :      RCH38366  Version/Release  :   2/ 1.0  Stopped   . . . . : 05/16/90 15:27:21

                             Number Of Transactions Per Hour
 Itv
 End
       0      1000      2000      3000      4000      5000      6000      7000
-----  +----+----+----+----+----+----+----+----+----+----+----+----+----+----+
 ***
14.50   XXmSSSSS
14.55   XSSSSSS
15.00   XXmSSS
15.05   XXXmmSSSSSSS
15.10   XS
15.15   XS
15.20   S

Throughput Components:

     S = Simple Transactions
     m = Medium Transactions
     X = Complex Transactions
```

Figure 8-41. Interactive Throughput by 5 Minute Intervals

Interactive CPU Utilization

The Interactive CPU Utilization by 5 Minute Intervals section of the Job Summary Report, shown in Figure 8-42, gives simple, medium, and complex transactions relative to their processing unit utilization.

```
                                        Job Summary Report                                          05/17/90  4:06:07
                              Interactive CPU Utilization by 5 Minute Intervals                          Page 0009

Member  . . . : TJST41       Model/Serial  . . :  B60/10-12883  Main storage . . :   96.0 M  Started   . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2     System name . . . :      RCH38366  Version/Release  :   2/ 1.0  Stopped   . . . . : 05/16/90 15:27:21

                       Percent CPU Utilization
 Itv
 End    0   10   20   30   40   50   60   70   80   90  100
-----   +----+----+----+----+----+----+----+----+----+----+
 ***
14.50   XXXXS
14.55   XXXX
15.00   XXS
15.05   XXXXS
15.10   S
15.15   X
15.20   S

CPU Components:

     S = Simple Transactions
     m = Medium Transactions
     X = Complex Transactions
```

Figure 8-42. Interactive CPU Utilization by 5 Minute Intervals

Interactive Response Time

The Interactive Response Time by 5 Minute Intervals section of the the Job Summary Report, shown in Figure 8-43 on page 8-62, gives the response components relative to the resulting response time.

```
                                              Job Summary Report                                        05/17/90  4:06:07
                                   Interactive Response Time by 5 Minute Intervals                           Page 0010

Member  . . . : TJST41       Model/Serial  . . :  B60/10-12883   Main storage . . :    96.0 M   Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2     System name . . . :       RCH38366   Version/Release  :    2/ 1.0   Stopped  . . . . : 05/16/90 15:27:21
                          Average Response Time (Seconds)
 Itv
 End    0      1.00      2.00      3.00      4.00      5.00      6.00      7.00
-----   +----+----+----+----+----+----+----+----+----+----+----+----+----+----+
 ***
14.50    RRRRRRw
14.55    RRRRR
15.00    RRRRR
15.05    RRRRRw
15.10    RRRw
15.15    RRRRRww
15.20    RRw

Response Components:

     R = CPU + Disk + Wait-to-Ineligible
     w = Exceptional Wait
```

Figure 8-43. Interactive Response Time by 5 Minute Intervals

Scatter Diagram

The Scatter Diagram section of the Job Summary Report, shown in Figure 8-44, gives the average of measured response times for 5-minute intervals compared to transaction rates.

```
                                              Job Summary Report                                        05/17/90  4:06:07
                              Scatter Diagram of Interactive Transactions by 5 Minute Intervals              Page 0011

Member  . . . : TJST41       Model/Serial  . . :  B60/10-12883   Main storage . . :    96.0 M   Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2     System name . . . :       RCH38366   Version/Release  :    2/ 1.0   Stopped  . . . . : 05/16/90 15:27:21

                    Response Time vs Number of Transactions per Hour
                    -----------------------------------------------

     OVFLW
     -----
     7.00+
         -
R        -
E        -
S   6.00+
P        -
O        -
N        -
S   5.00+
E        -
         -
T        -
I   4.00+
M        -
E        -
         -
I   3.00+
N        -
         -
S        -
E   2.00+
C        -
O        -
N        -
D   1.00+
S        - 1     1
         - 1  011    1
         -1
    0.00+----+----+----+----+----+----+----+----+----+----+----+----+----+----+|O
        0        1000      2000      3000      4000      5000      6000      7000|F

                             Number of Transactions per Hour

Legend:      1-9 Indicates the Number of Occurrences.
               *  Indicates more than 9 Occurrences.
               0  Identifies Average of All Occurrences.
```

Figure 8-44. Interactive Transactions by 5 Minute Intervals

Interactive Program Transaction Statistics

The Interactive Program Transaction Statistics section of the Job Summary Report, shown in Figure 8-45, arranges the programs by the number of transactions associated with the programs.

```
                                          Job Summary Report                                    05/17/90  4:06:07
                                     Interactive Program Statistics                                     Page 0017

Member . . . : TJST41      Model/Serial  . . :  B60/10-12883  Main storage . . :   96.0 M  Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2   System name . . . :      RCH38366  Version/Release  :   2/ 1.0  Stopped  . . . . : 05/16/90 15:27:21
```

Rank	Number Tns	Program Name	CPU /Tns	CPU Util	Cum CPU Util	Sync Disk I/O Rqs/Tns DB Read	DB Write	NDB Read	NDB Write	Sum	Async DIO /Tns	Rsp /Tns	Short Wait /Tns	Seize Wait /Tns	Pct Tns	Cum Pct Tns
1	96	QUIINMGR	.300	1.2	1.2	1		5		6		.570	.016		28.5	28.5
2	87	QSUBLDS	.757	2.8	4.1			1		1		.871	.001		25.8	54.3
3	67	QSMCSMSU	.062	.2	4.2			1		1		.163	.010		19.9	74.2
4	35	QMHGSD	.352	.5	4.8			3		3		.683	.025		10.4	84.6
5	22	QSPDSPF	.175	.2	4.9	1		4		5		.321			6.5	91.1
6	17	QPTPRCSS	.226	.2	5.1			3		3		.355	.020		5.0	96.1
7	5	QMHDSMSS	.279	.1	5.2			4		4		.718	.092		1.5	97.6
8	3	QMHDSPJL	.092		5.2							.188			.9	98.5
9	2	QSUEXIT	.066		5.2			4		4		.181			.6	99.1
10	1	*SIGNOFF*	.468		5.2			3		3		.831	.024		.3	99.4
11	1	QMHDSEXT	.352		5.2		76	15	76	167	152	1.052	.111		.3	99.7
12	1	QSMTAPS	.076		5.2			4		4		.309	.089		.3	100.0

Figure 8-45. Interactive Program Transaction Statistics

The columns in the Interactive Program Transaction Statistics section are as follows:

Rank The order of the program according to the number of transactions.

Number Tns
The number of transactions associated with the program.

Program Name
The name of the program.

CPU /Tns The amount of available processing unit time per transaction in seconds.

CPU Util The percentage of available processing unit time used by the program. (This is the primary value to look at on this report.)

Cum CPU Util
The cumulative percentage of available processing unit time used by the listed programs through the current program.

Sync Disk I/O Rqs/Tns
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).

Async DIO /Tns
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).

Rsp /Tns The average response time per transaction in seconds.

Short Wait /Tns
The average time, in seconds, of short (active) wait time per transaction. If the value is high, it may be due to the use of DFRWRT(*NO) and RSTDSP(*YES) in the program display files.

Seize Wait /Tns
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.

Pct Tns The percentage of all transactions that were associated with a program.

Cum Pct Tns
The cumulative percentage of all transactions through the listed program.

Seize/Lock Conflicts by Object

The Summary of Seize/Lock Conflicts by Object section of the Job Summary Report, shown in Figure 8-46, gives information about the locks and seizes associated with objects. The unnamed object, shown as ADDR 00000E00, is the licensed internal code (LIC) database in-use table. It often appears in this report when there are a high number of database file opens and closes.

```
                                          Job Summary Report                                                 05/17/90  4:06:07
                                  Summary of Seize/Lock Conflicts by Object                                          Page 0029

Member  . . . : TJST41        Model/Serial  . . :  B60/10-12883  Main storage . . :   96.0 M  Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2      System name . . . :     RCH38366   Version/Release  :   2/ 1.0  Stopped  . . . . : 05/16/90 15:27:21

                                                    -------- Interactive Waiters ---------    -------- Noninteractive Waiters -------
                                                    ----- Locks -----   ---- Seizes -----     ----- Locks -----   ----- Seizes -----
 Type     Library      File        Member           Number    Avg Sec   Number    Avg Sec     Number    Avg Sec   Number    Avg Sec
------   ----------   ----------   ----------       ------   --------   ------   --------     ------   --------   ------   --------
JOBQ     QGPL         QBATCH                                                                       4       .004
MSGQ     QUSRSYS      SFRASER                            4       .005
----------------------------------------------     ------   --------   ------   --------     ------   --------   ------   --------
* Total Conflicts and Avg Sec/Conflict                   4       .005                              4       .004
                                                    ------              ------
* Total Transactions With Conflicts                      4
                                                    ------              ------
* Averages Per Conflict Transaction                   1.00       .005
```

Figure 8-46. Summary of Seize/Lock Conflicts by Object

The columns in the Summary of Seize/Lock Conflicts by Object section are as follows:

Type The type of seize/lock conflict.

Library The library that contains the object.

File The file that contains the object.

Member The member that was involved in the conflict.

Number Locks
The number of locks attributed to interactive or noninteractive waiters.

Avg Sec Locks
The average length of a lock in seconds attributed to interactive or noninteractive waiters.

Number Seizes
The number of seizes attributed to interactive or noninteractive waiters.

Avg Sec Seizes
The average length of a seize in seconds attributed to interactive or noninteractive waiters.

Special System Information

To include additional system information, specify the *SS value on the OPTION parameter when you use the PRTTNSRPT command.

In general, the information identifies exceptional conditions and events that occur over the measurement period. If you analyze these exceptions, you might find jobs and programs you need to examine. A summary of these sections of the Job Summary Report follows.

Priority-Jobtype-Pool Statistics

The Priority-Jobtype-Pool Statistics section of the Job Summary Report, shown in Figure 8-47, shows the total processing unit seconds and physical I/O requests for each category of priority-jobtype and pool combination recorded during the overall test period. The number of total transactions is shown for job type I only.

```
                                        Job Summary Report                                              05/17/90  4:06:07
                                    Priority-Jobtype-Pool Statistics                                            Page 0012

Member  . . . : TJST41       Model/Serial  . . :  B60/10-12883  Main storage . . :  96.0 M  Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2     System name . . . :      RCH38366  Version/Release  :  2/ 1.0  Stopped  . . . . : 05/16/90 15:27:21

        Job                      CPU        --- Disk I/O Requests ----       Number
Pty     Type      Pool         Seconds         Sync          Async            Tns
---     ----      ----       -----------   -----------   -----------      ----------
00      B         02               4.782            90
00      M         02               1.952            42
00      S         02                .648             3
00      L         01              16.006           612
00      L         02               1.965         31022         61824
15      B         02
15      W         03
16      B         02
19      B         02
20      I         02             121.567          2356           304             336
20      S         02
40      A         02
40      B         02
50      B         02             718.033         22287         41216
52      X         02                .519             4
```

Figure 8-47. Priority-Jobtype-Pool Statistics

The columns in the Priority-Jobtype-Pool Statistics section are as follows:

Pty The run priority of the jobs.

Job Type The job type and subtype.

Pool The pool in which the jobs were running.

CPU Seconds
The total processing unit seconds used by the jobs with a given combination of priority, job type, and pool.

Sync Disk I/O Requests
The total number of synchronous disk I/O requests for the given combination of priority, job type, and pool.

Async Disk I/O Requests
The total number of asynchronous disk I/O requests for the given combination of priority, job type, and pool.

Number Tns
The total number of transactions for the given combination of priority, job type, and pool.

Job Statistics

The Job Statistics section of the Job Summary Report shows the 10 jobs with the:

- Most transactions (shown in Figure 8-48)
- Largest average response time
- Largest average processing unit time per transaction
- Largest synchronous disk I/O per transaction
- Largest asynchronous disk I/O per transaction
- Most seize conflicts
- Most record lock conflicts
- Most active-to-ineligible occurrences
- Most wait-to-ineligible occurrences
- Most event wait occurrences

Job Summary Report 05/17/90 4:06:07
Job Statistics Page 0013

Member . . . : TJST41 Model/Serial . . : B60/10-12883 Main storage . . : 96.0 M Started . . . . : 05/16/90 14:48:28
Library . . : TJSWORK2 System name . . . : RCH38366 Version/Release : 2/ 1.0 Stopped . . . . : 05/16/90 15:27:21

JOBS WITH MOST TRANSACTIONS

Rank	Job Name	User Name	Job Number	Pl	Typ	Pty	Nbr Tns	Rsp /Tns	CPU /Tns	CPU Util	Cum CPU Util	Sync DIO /Tns	Async DIO /Tns	Nbr W-I	Nbr A-I	Nbr Evt	Number Conflict Lck	Number Conflict Sze	Pct Tns	Cum Pct Tns
1	DSP01	SFRASER	020398	02	I	20	204	.699	.504	4.4	4.4	2					4		60.5	60.5
2	DSP02	DRLIPPS	020438	02	I	20	42	.379	.139	.3	4.7	127	245						12.5	73.0
3	DRLIPPS	DRLIPPS	020391	02	I	20	36	.195	.080	.1	4.8	3							10.7	83.7
4	SFRASER	AMUNDSON	020408	02	I	20	29	.461	.239	.3	5.1	8							8.6	92.3
5	DLH	DJW	020403	02	I	20	26	.239	.122	.1	5.2	3							7.7	100.0
6																				
7																				
8																				
9																				
10																				

Figure 8-48. Job Statistics Showing the Most Transactions

The columns in the Job Statistics section are as follows:

Rank The order of the job.

Job Name
The name of the job.

User Name
The name of the user who submitted the job.

Job Number
The job number.

Pl The pool that the job was running in.

Typ The type and subtype of the job.

Pty The priority at which the job was running.

Nbr Tns The number of transactions by the job.

Rsp /Tns The average response time per transaction in seconds for the job.

CPU /Tns The average amount of available processing unit time per transaction in seconds.

CPU Util The percentage of available processing unit time used by the job.

Cum CPU Util
The cumulative percentage of available processing unit time used by all the jobs in the listing through the current job.

Sync DIO /Tns
The average number of synchronous disk I/O operations per transaction.

Async DIO /Tns
The average number of asynchronous disk I/O operations per transaction.

Nbr W-I The number of wait-to-ineligible state transitions by the job. This column shows how many times the job had to wait for an activity-level slot before the system could begin processing the transaction.

Nbr A-I 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 Evt The number of event waits that occurred during the job processing.

Number Lck Conflict
The number of times the job had a lock conflict.

Number Sze Conflict
The number of times the job had a seize conflict.

Pct Tns The percentage of the total transactions that were due to this job.

Cum Pct Tns
The cumulative percentage total transactions that were due to the jobs in the listing through this job.

Interactive Program Statistics

The Interactive Program Statistics section of the Job Summary Report gives additional program information showing the top 10 programs with the largest average:

- Processing unit time per transaction (shown in Figure 8-49)
- Synchronous disk I/O per transaction
- Asynchronous disk I/O per transaction
- Response time per transaction
- Synchronous database reads per transaction
- Synchronous database writes per transaction
- Synchronous nondatabase reads per transaction
- Synchronous nondatabase writes per transaction

```
                                              Job Summary Report                                     05/17/90  4:06:07
                                         Interactive Program Statistics                                      Page 0018

Member . . . : TJST41      Model/Serial  . . :  B60/10-12883  Main storage . . :   96.0 M  Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2   System name . . . :      RCH38366  Version/Release  :   2/ 1.0  Stopped  . . . . : 05/16/90 15:27:21

PROGRAMS WITH HIGHEST CPU/TNS
```

Rank	Number Tns	Program Name	CPU /Tns	CPU Util	Cum CPU Util	----- Sync Disk I/O Rqs/Tns ----- DB Read	DB Write	NDB Read	NDB Write	Sum	Async DIO /Tns	Rsp /Tns	Short Wait /Tns	Seize Wait /Tns	Pct Tns	Cum Pct Tns
1	87	QSUBLDS	.757	2.8	2.8			1		1		.871			25.8	25.8
2	1	*SIGNOFF*	.468		2.8			3		3		.831			.3	26.1
3	35	QMHGSD	.352	.5	3.4			3		3		.683			10.4	36.5
4	1	QMHDSEXT	.352		3.4		76	15	76	167	152	1.052			.3	36.8
5	96	QUIINMGR	.300	1.2	4.6	1		5		6		.570			28.5	65.3
6	5	QMHDSMSS	.279	.1	4.7			4		4		.718			1.5	66.8
7	17	QPTPRCSS	.226	.2	4.8			3		3		.355			5.0	71.8
8	22	QSPDSPF	.175	.2	5.0	1		4		5		.321			6.5	78.3
9	3	QMHDSPJL	.092		5.0							.188			.9	79.2
10	1	QSMTAPS	.076		5.0			4		4		.309			.3	79.5

Figure 8-49. Programs with the Highest Processing Unit Transactions

The columns in the Interactive Program Statistics section are as follows:

Rank The order of the program.

Number Tns
The number of transactions associated with the program.

Program Name
The name of the program.

CPU /Tns The amount of available processing unit time per transaction in seconds.

CPU Util The percentage of available processing unit time used by the program.

Cum CPU Util
The cumulative percentage of available processing unit time used by the listed programs through the current program.

Sync Disk I/O Rqs/Tns
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).

Async DIO /Tns
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).

Rsp /Tns The average response time per transaction in seconds.

Short Wait /Tns
The average time, in seconds, of short (active) wait time per transaction. If the value is high, it may be due to the use of DFRWRT(*NO) and RSTDSP(*YES) in the program display files.

Seize Wait /Tns
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.

Pct Tns The percentage of all transactions that were associated to a program.

Cum Pct Tns
The cumulative percentage of all transactions through the listed program.

Individual Transaction Statistics

The Individual Transaction Statistics section of the Job Summary Report lists the 10 transactions with the least or most:

- Response time (shown in Figure 8-50 on page 8-70)
- Processing unit service time
- Total synchronous disk I/O
- Total asynchronous disk I/O
- Synchronous database reads
- Synchronous database writes
- Synchronous nondatabase reads
- Synchronous nondatabase writes
- Asynchronous database reads
- Asynchronous database writes
- Asynchronous nondatabase reads
- Asynchronous nondatabase writes
- Short-wait-extended time
- Short-wait time
- Lock-wait time
- Excessive activity-level wait time
- Active time
- Effective address overflow exceptions

- Binary overflow exceptions
- Decimal overflow exceptions
- Floating point overflow exceptions
- Process access group fault exceptions
- Checksum I/O operations
- Permanent writes

```
                                    Job Summary Report                                                  05/17/90  4:06:07
                            Individual Transaction Statistics                                                  Page 0021

Member . . . : TJST41     Model/Serial  . . :  B60/10-12883  Main storage . . :   96.0 M  Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2  System name . . . :  RCH38366      Version/Release  :   2 /1.0  Stopped  . . . . : 05/16/90 15:27:21

TRANSACTIONS WITH LONGEST RESPONSE TIMES

Rank      Value        Time        Program       Job Name      User Name       Number      Pool      Type      Priority
----    --------  ------------   ----------   ----------    ----------      ------      ----      ----      --------
  1       6.953   14.51.51.619   QSUBLDS       DSP01         SFRASER         020398       02        I          20
  2       6.859   15.05.57.610   QMHGSD        DSP01         SFRASER         020398       02        I          20
  3       6.684   14.54.45.748   QSUBLDS       DSP01         SFRASER         020398       02        I          20
  4       6.485   14.57.04.456   QSUBLDS       DSP01         SFRASER         020398       02        I          20
  5       6.008   14.58.39.300   QSUBLDS       DSP01         SFRASER         020398       02        I          20
  6       5.691   15.05.08.794   QSUBLDS       DSP01         SFRASER         020398       02        I          20
  7       5.620   15.02.55.162   QSUBLDS       DSP01         SFRASER         020398       02        I          20
  8       5.375   14.53.17.618   QSUBLDS       DSP01         SFRASER         020398       02        I          20
  9       5.367   14.55.59.637   QSUBLDS       DSP01         SFRASER         020398       02        I          20
 10       3.830   14.51.24.024   QSUBLDS       DSP01         SFRASER         020398       02        I          20
```

Figure 8-50. Transactions with the Longest Response Time

The columns in the Individual Transaction Statistics section are as follows:

Rank The order of the transaction according to the data being put in order by importance.

Value The value of the data being compared for the transaction.

Time The time when the transaction completed.

Program The name of the program with which the transaction is associated.

Job Name The name of the job with which the transaction is associated.

User Name The name of the job with which the transaction is associated.

Number The number of the job with which the transaction is associated.

Pool The number of the pool that the transaction was in.

Type The type and subtype of the job.

Priority The priority of the job.

Longest Seize/Lock Conflicts

The Longest Seize/Lock Conflicts section of the Job Summary Report (shown in Figure 8-51 on page 8-71) shows the 20 longest lock or seize conflicts during the trace period.

```
                                         Job Summary Report                                       05/17/90  4:06:07
                                      Longest Seize/Lock Conflicts                                        Page 0027

Member  . . . : TJST41       Model/Serial  . . :  B60/10-12883   Main storage . . :    96.0 M   Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2     System name . . . :      RCH38366   Version/Release  :   2/ 1.0   Stopped  . . . . : 05/16/90 15:27:21

                              Job        User      Job                        Holder- Job Name.. User Name. Number   Pool  Type  Pty
Rank   Value       Time       Name       Name      Number Pl Typ Pty S/L   Object- Type.. Library... File...... Member.... RRN......
---- -------- ------------ ---------- ---------- ------ -- --- --- ---   ---------------------------------------------------------
  1     .006  15.10.13.223 DSP01      SFRASER    020398 02  I   20  L    HOLDER- PRTTNSRPT  SFRASER    020439    02     I      20
                                                                         OBJECT- MSGQ    QUSRSYS    SFRASER
  2     .005  15.25.20.550 DSP01      SFRASER    020398 02  I   20  L    HOLDER- PRTTNSRPT  SFRASER    020442    02     I      20
                                                                         OBJECT- MSGQ    QUSRSYS    SFRASER
  3     .005  15.23.22.018 DSP01      SFRASER    020398 02  I   20  L    HOLDER- QDFTJOBD   SFRASER    020441    02     I      20
                                                                         OBJECT- MSGQ    QUSRSYS    SFRASER
  4     .005  15.10.21.028 DSP01      SFRASER    020398 02  I   20  L    HOLDER- PRTTNSRPT  SFRASER    020440    02     I      20
                                                                         OBJECT- MSGQ    QUSRSYS    SFRASER
  5     .005  15.10.01.483 QBASE      QSYS       020205 02  M   00  L    HOLDER- DSP01      SFRASER    020398    02     I      20
                                                                         OBJECT- JOBQ    QGPL       QBATCH
  6     .005  15.08.06.635 QBASE      QSYS       020205 02  M   00  L    HOLDER- DSP01      SFRASER    020398    02     I      20
                                                                         OBJECT- JOBQ    QGPL       QBATCH
  7     .004  15.25.00.872 QBASE      QSYS       020205 02  M   00  L    HOLDER- DSP01      SFRASER    020398    02     I      20
                                                                         OBJECT- JOBQ    QGPL       QBATCH
  8     .004  15.11.17.183 QBASE      QSYS       020205 02  M   00  L    HOLDER- DSP01      SFRASER    020398    02     I      20
                                                                         OBJECT- JOBQ    QGPL       QBATCH
```

Figure 8-51. Longest Seize/Lock Conflicts

The columns in the Longest Seize/Lock Conflicts section are as follows:

Rank The order of the seize or lock conflict.

Value The number of seconds in which the seize or lock conflict occurred.

Time The time at which the seize or lock conflict occurred.

Job Name
The name of the job that had the conflict.

User Name
The name of the user that had the conflict.

Job Number
The number of the job that had the conflict.

Pl The pool in which the job was running.

Typ The type and subtype of the job.

Pty The priority of the job.

S/L Whether the conflict was a seize (S) or lock (L) conflict.

Holder Job Name
The name of the job that held the object.

Holder User Name
The name of the user that held the object.

Holder Number
The number of the job that held the object.

Holder Pool
The pool that held the job while it was running.

Holder Type
The type and subtype of the holder&csq,s job.

Holder Pty
The priority of the holder's job.

Object Type
The type of seize/lock conflict.

Object Library
The library that contains the object.

Object File
The file that contains the object.

Object Member
The member that was involved in the conflict.

Object RRN
The relative record number of the record involved in the conflict.

Longest Holders of Seize/Lock Conflicts

The Longest Holders of Seize/Lock Conflicts section of the Job Summary Report (see Figure 8-52) shows the holders of the longest lock or seize conflicts during the trace period.

```
                                                  Job Summary Report                                                      05/17/90  4:06:07
                                         Longest Holders of Seize/Lock Conflicts                                                Page 0028

Member  . . . : TJST41       Model/Serial  . . :  B60/10-12883   Main storage . . :   96.0 M   Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2     System name . . . :      RCH38366   Version/Release  :   2/ 1.0   Stopped  . . . . : 05/16/90 15:27:21

                                 Job        User       Job                           ---------------------- Object -----------------------
Rank    Value       Time         Name       Name       Number  Pl  Typ Pty  S/L  Type    Library     File        Member      RRN
---- --------  ------------  ---------- ---------- ------  --  --- ---  ---  ------  ----------  ----------  ----------  ---------
  1      .004  15.11.17.183  QBASE      QSYS       020205  02   M   00   L   JOBQ    QGPL        QBATCH      QNETF
```

Figure 8-52. Longest Holders of Seize/Lock Conflicts

The columns in the Longest Holders of Seize/Lock Conflicts section are as follows:

Rank The order of the seize or lock conflict.

Value The number of seconds in which the seize or lock conflict occurred.

Time The time at that the seize or lock conflict occurred.

Job Name
The name of the job that had the conflict.

User Name
The name of the user that had the conflict.

Job Number
The number of the job that had the conflict.

Pl The pool in which the job was running.

Typ The type and subtype of the job.

Pty The priority of the job.

S/L Whether the conflict was a seize (S) or lock (L) conflict.

Object Type
The type of seize/lock conflict.

Object Library
The library that contains the object.

Object File
The file that contains the object.

Object Member
The member that was involved in the conflict.

Object RRN
The relative record number of the record involved in the conflict.

Batch Job Analysis

The Batch Job Analysis section of the Job Summary Report (see Figure 8-53) shows information on the batch job workload during the trace period.

```
                                          Job Summary Report                                                05/17/90  4:06:07
                                          Batch Job Analysis                                                        Page 0030

Member  . . . : TJST41      Model/Serial  . . :  B60/10-12883  Main storage . . :   96.0 M  Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2    System name . . . :      RCH38366  Version/Release  :   2/ 1.0  Stopped  . . . . : 05/16/90 15:27:21
```

Job Name	User Name	Job Number	Pl	Typ	Pty	Start	Stop	Elapsed Seconds	CPU Seconds	CPU Util	Sync Disk I/O	Async Disk I/O	Synchronous BCPU /DIO	Synchronous DIO/Sec Elp	Synchronous DIO/Sec Act	Synchronous DIO/Sec Ded	Excp Wait Sec
PRTTNSRPT	SFRASER	020439	02	B	50	15.08.06	15.10.14	127.725	14.826	11.6	5311	325		633	313	30	102.318
PRTTNSRPT	SFRASER	020440	02	B	50	15.10.01	15.10.21	19.752	12.586	63.7	183	34		134	376	30	.157
QDFTJOBD	SFRASER	020441	02	B	50	15.11.17	15.23.22	724.887	677.625	93.5	6598	10304		77	90	30	.042
PRTTNSRPT	SFRASER	020442	02	B	50	15.25.01	15.25.21	19.912	12.996	65.3	195	23		710	961	30	.132

Figure 8-53. Batch Job Analysis

The columns in the Batch Job Analysis section are as follows:

Job Name
The name of the job.

User Name
The name of the user who submitted the job.

Job Number
The number of the job.

Pl The pool in which the job ran.

Typ The type and subtype of the job.

Pty The priority of the job.

Start The time the job started.

Stop The time the job ended.

Elapsed Seconds
The number of seconds elapsed from when the job started to when the job ended.

CPU Seconds
The amount of available processor unit time used by the job in seconds.

CPU Util The percentage of the available processor unit time used by the job.

Sync Disk I/O
The number of synchronous disk I/O operations by the job.

Async Disk I/O
The number of asynchronous disk I/O operations by the job.

BCPU / Synchronous DIO
The average number of batch processor unit seconds per synchronous disk I/O operation.

Synchronous DIO / Elp Sec

The number of synchronous disk I/O operations per elapsed second.

Synchronous DIO / Act Sec

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

The 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.

Excp Wait Sec

The total amount of exceptional wait time in seconds for the job. This value is the sum of those waits listed under the Exceptional Wait Breakdown by Job Type /Category part of the System Summary Data section.

Batch Thread Analysis

The Batch Thread Analysis section of the Job Summary Report (see Figure 8-54) shows information on the batch job workload during the trace period according to threads.

```
                                         Job Summary Report                                       05/17/90  4:06:07
                                        Batch Thread Analysis                                            Page 0031

Member  . . . : TJST41      Model/Serial  . . :  B60/10-12883  Main storage . . :   96.0 M   Started  . . . . : 05/16/90 14:48:28
  Library . . : TJSWORK2    System name . . . :      RCH38366  Version/Release  :   2/ 1.0   Stopped  . . . . : 05/16/90 15:27:21
```

Thread	Pty	Number Jobs	Elapsed Seconds	CPU Seconds	Excp Wait	Sync Disk I/O	Async Disk I/O
1	00	1	2332.668	4.782	2320.389	654	435
2	15	1	2332.675		2332.675	654	765
3	15	1	2332.677		2332.677	876	876
4	15	1	2332.684		2332.684	467	654
5	16	1	2332.676		2332.676	988	877
6	19	1	2332.665		2332.665	5463	546
7	40	1	2332.664		2332.664	765	876
8	40	1	2332.665		2332.665	979	765
9	40	1	2332.668		2332.668	243	424
10	40	1	2332.668		2332.668	877	453
11	40	1	2332.669		2332.669	456	776
12	40	1	2332.669		2332.669	788	876
13	40	1	2332.673		2332.673	344	243
14	40	1	2332.675		2332.675	543	5456
15	50	1	19.752	12.586	.157	5183	10304
16	50	3	872.524	705.447	102.492	34104	430912

Figure 8-54. Batch Thread Analysis

The columns in the Batch Job Analysis section are as follows:

Thread The thread of the jobs. Jobs that are in the same thread cannot be running at the same time. A thread is similar to an activity-level.

Pty The priority of the jobs in the thread.

Number Jobs

The number of batch jobs in the thread.

Elapsed Seconds

The number of seconds elapsed.

CPU Seconds

The amount of available processor unit time used by the jobs in the thread in seconds.

Excp Wait The amount of exceptional wait time for the jobs in the thread in seconds.

Sync Disk I/O
The number of synchronous disk I/O operations by the jobs in the thread.

Async Disk I/O
The number of asynchronous disk I/O operations by the jobs in the thread.

Report Selection Criteria

The Report Selection Criteria section of the Job Summary Report, shown in Figure 8-55, gives the selection values you chose to produce the report.

```
                                               Job Summary Report                                               04/20/90 10:06:31
                                             Report Selection Criteria                                                 Page 0014
                                             Sample Transaction Report

Member  . . . : GOODSTUF    Model/Serial  . . :  B60/10-12883  Main Storage . . :   96.0 M  Started . . . . :   3/29/90 11:35:54
  Library . . : QPFRDATA    System name . . . :  RCH38366      Version/Release  :   2/ 1.0  Stopped . . . . :   3/29/90 13:35:39

Select Parameters

Pools                       - 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16

Jobs                        - Jobname1
                              Jobname2
                              Jobnum

User IDs                    - User1      User2      User3      User4     User5      User6
                              User7      User8      User9      User10    User11     User12

Functional Areas            - Accounting          Payroll             Research
                              Development         ProjectX            MrNolansStaff

                            - No Select parameters were chosen.

Omit Parameters

Pools                       - 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16

Jobs                        - 012345/Useridwxyz/Jobname123
                              987654/Useridabcd/Jobname456

User IDs                    - User1      User2      User3      User4     User5      User6
                              User7      User8      User9      User10    User11     User12

Functional Areas            - Accounting          Payroll             Research
                              Development         ProjectX            MrNolansStaff

                            - No Omit parameters were chosen.

Options Selected            - SS Include Special Summary Reports
```

Figure 8-55. Job Summary Report: Report Selection Criteria

The SELECT parameters on the Report Selection Criteria Report are as follows:

Pools The pools you specify. The values can be from 1 through 16.

Jobs The jobs you specify. Either the job number or the jobname is selected.

User IDs The user IDs you specify. Each user identification is a 10-character name.

Functional Areas
The functional areas you specify. Each name is a 20-character name.

If you use no SELECT parameters, the `No Select parameters were chosen` message appears.

The OMIT parameters on the Report Selection Criteria Report are as follows:

Pools The pools you specify. The values can be from 1 through 16.

Jobs The jobs you specify. The format of the entries is: `Jobnumber/username/jobname`.

User IDs The user IDs you specify. Each user identification is a 10-character name.

Functional Areas
The functional areas you specify. Each name is a 20-character name.

If you did not use OMIT parameters, the `No Omit parameters were chosen` message appears.

The options which were selected are also given.

Transaction Report

To produce the Transaction Report, specify RPTTYPE(*TNSACT) on the PRTTNSRPT command. The Transaction Report output has two parts:

- The Transaction Report, which shows data about each transaction in the job
- The report, which shows data about overall job operation

Each page of the Transaction Report contains header information shown in Figure 8-56.

```
                                                  Transaction Report                                          06/19/90  0:46:59
                                                                                                                      Page 0001
Member  . . . : TJST40      Model/Serial  . . :  B60/10-15018  Main storage . . :   96.0 M  Started  . . . . : 06/07/90 08:10:12
  Library . . : QPFRDATA    System name . . . :      RCH38366  Version/Release  :   2/ 1.0  Stopped  . . . . : 06/07/90 08:41:58
Job name  . . : SCPF        User name . . . . :     QSYS       Job number . . . :   000000  TDE/Pl/Pty/Prg . :    004B/02/52/YES
```

Figure 8-56. Transaction Report: Header Information

This header information consists of the following:

- Report title
- Current date and time
- Report page number
- User-selected report title
- Data member name
- Library name
- Model name and serial number
- User-defined system name
- Main storage size in megabytes
- Release number of the data collected
- Data collection start date and time
- Data collection stop date and time
- The name of the job
- The name of the user associated with the job
- The job number
- TDE/Pl/Pty/Prg data, which shows the task dispatching element (TDE) number, the storage pool in which the job ran, the job priority, and the job's initial purge attribute.

An example of the Transaction Report is shown in Figure 8-57 on page 8-77.

```
                                                 Transaction Report                                         06/19/90  0:46:59
                                                                                                                   Page 0004
Member  . . . : TJST40        Model/Serial  . . :  B60/10-15018  Main storage . . :   96.0 M  Started  . . . . : 06/07/90 08:10:12
  Library . . : QPFRDATA      System name . . . :      RCH38366  Version/Release  :   2/ 1.0  Stopped  . . . . : 06/07/90 08:41:58
Job name  . . : QLUS          User name . . . . :     QSYS       Job number . . . :   005203  TDE/Pl/Pty/Prg . :    004F/02/20/NO

         E              CPU   ---- Physical I/O Counts -----  ***** Transaction Response Time (Sec/Tns) ******  -BMPL-
         x              Sec   ----- Synchronous ------ Async   ******  - Activity Level Time -  Inel    Long    C    I Seize
         c   Program    Per    DB   DB  NDB   NDB        Disk    ****                 Short   Seize   Time    Wait    u    n  Hold    Key/
  Time   p    Name      Tns   Read  Wrt Read  Wrt   Sum  I/O      **      Active    Wait     Cft   A-I/W-I Lck/Oth  r    l  Time   Think
-------- - ---------- ------- ---- ---- ---- ---- ---- -----  -------- -------- ------- ------- ------- -------  --   -- ----- ------
08.20.00 Y QWCLMNSR     .121              65    1   66    0     2.060    2.058                                    2               588.8
08.21.40 Y QWCLMNSR     .158              41        41    0   100.320    2.347    .029    .091             97.943  1                  .0
08.21.43   QWCLMNSR     .040               9         9    0      .870     .792    .024                       .054  2                  .0
08.25.02 Y QWCLMNSR     .101              66    1   67    0   200.142    2.067                            198.074  2                  .0
08.25.20 Y QWCLMNSR     .141              34        34    0    19.226    2.578    .013                     16.635  1                  .0
08.25.23   QWCLMNSR     .043              10        10    0     1.095    1.006    .043                       .046  2                  .0

---------------------------------------------
J O B   S U M M A R Y   D A T A   (T O T A L S)
---------------------------------------------

        Average         .101    0    0   38    0   38    0    53.952     .904    .027    .091    .000  62.550               .0  196.2
        Count                                                   6       12       4       1               5                1      3
        Minimum         .040                         9             .870     .055    .013    .091            .046                  .0
        Maximum         .158                        67          200.142    2.185    .043    .091         198.074               588.8
        Total/Job       .604                       227
                                                          0   912.588 Elapsed      .1 Percent CPU Utilization
```

Figure 8-57. Transaction Report

The columns in the Transaction Report are as follows:

Time The time that the transaction ended.

Excp A Y in this column means that the transaction had exceptions. The types of exceptions which are included are process access group exceptions, effective address overflow exceptions, and decimal, binary, and floating point overflow. See the transition report to see which exceptions the transaction had.

Program Name
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.

CPU Sec per Tns
The processing unit time per transaction.

Synchronous Disk I/O Counts
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).

Async DIO /Tns
The sum of the asynchronous DB Read, DB Write, NDB Read, and NDB Write requests (the number of asynchronous I/O requests per transaction).

Transaction Response Time (Sec/Tns)
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).

Activity Level Time
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.

Inel Time A-I/W-I
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).

Long Wait Lck/Oth
The amount of time the job spent waiting for a system resource. An example of a long wait would be a record-lock conflict.

BMPL - Cur and Inl
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).

Note: Multiprogramming level (MPL) is used interchangeably with activity level.

Seize Hold Time
The amount of time that the transaction held up other jobs in the system by a seize or lock on an object.

Key/Think The amount of time spent waiting for the work station user by the program.

Job Summary Data

The Job Summary Data section of the Transaction Report includes averages of the job data. Some of this information is also found in the Job Summary section of the Job Summary Report.

The columns in the report give, for each of the jobs listed, the following information:

Average The average value of the item described in the column for all transactions.

Count The number of occurrences of the item in the column.

Minimum The minimum value of the item that occurred in the column.

Maximum The maximum value of the item that occurred in the column.

Total /Job The total (sum) of the item in the column for the job.

CPU Sec Per Tns
The processing unit time per transaction.

Synchronous Disk I/O Counts
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).

Async DIO /Tns
The sum of the asynchronous DB READ, DB WRITE, NDB READ, and NDB WRITE requests (the number of asynchronous I/O requests per transaction).

Transaction Response Time (Sec/Tns)
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).

Activity Level Time
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.

Inel Time A-I/W-I
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).

Long Wait The amount of time the job spent waiting for a system resource. An example of a long wait would be a record-lock conflict.

BMPL - Cur and Inl
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).

Note: Multiprogramming level (MPL) is used interchangeably with activity level.

Seize Hold Time
The amount of time that the transaction held up other jobs in the system on an object.

Key/Think The amount of time spent waiting for the work station user by the program.

Transition Report

To produce the Transition Report, specify RPTTYPE(*TRSIT) on the PRTTNSRPT command.

The Transition Report is composed of two sections:

- Transition Detail, which shows each state transition made by the job (going from one state to another, such as active-to-ineligible)
- Summary, which shows the same data as the summary output from the Transaction Report

Each page of the Transition Report contains the header information shown in Figure 8-58.

```
                                                  Transition Report                                             06/19/90  0:46:58
                                                                                                                        Page 0001

Member  . . . : TJST40       Model/Serial  . . :  B60/10-15018  Main storage . . :   96.0 M  Started  . . . . : 06/07/90 08:10:12
  Library . . : QPFRDATA     System name . . . :      RCH38366  Version/Release  :   2/ 1.0  Stopped  . . . . : 06/07/90 08:41:58
Job name  . . : SCPF         User name . . . . :    QSYS        Job number . . . :   000000  TDE/Pl/Pty/Prg . :    004B/02/52/NO
```

Figure 8-58. Transition Report: Header Information

This header information consists of the following:

- Report title
- Current date and time
- Report page number
- User-selected report title
- Data member name
- Library name
- Model name and serial number
- User-defined system name
- Main storage size in megabytes
- Release number of the data collected
- Data collection start date and time
- Data collection stop date and time
- The name of the job
- The name of the user associated with the job
- The job number

- The TDE/Pl/Pty/Prg data, which shows the task dispatching element (TDE) number, the storage pool in which the job ran, the job priority, and the job's initial purge attribute.

An example of the Transition Report is shown in Figure 8-59.

```
                                                   Transition Report                                          06/19/90  0:46:58
                                                                                                                     Page 0001

Member  . . . : TJST40       Model/Serial  . . :  B60/10-15018   Main storage . . :   96.0 M   Started  . . . . : 06/07/90 08:10:12
  Library . . : QPFRDATA     System name . . . :     RCH38366    Version/Release  :   2/ 1.0   Stopped  . . . . : 06/07/90 08:41:58
Job name  . . : SCPF         User name . . . . :   QSYS          Job number . . . :   000000   TDE/Pl/Pty/Prg . :    004B/02/52/NO

Job type  . . : X        Elapsed Time -- Seconds           Sync/Async Phy I/O      -MPL-
                         -----------------------       ------------------------- C   I     Last 4 Programs in Invocation Stack
             State Wait   Long   Active   Inel    CPU    DB    DB  NDB   NDB       u   n -------------------------------------------
    Time     W A I Code   Wait    /Rsp*   Wait    Sec    Read  Wrt Read  Wrt  Tot  r   l    Last      Second      Third      Fourth
------------ ----- ---- ------- ------- ------- -------- ---- ---- ---- ---- ---- -- -- ---------- ---------- ---------- ----------
08.10.11.964 *TRACE ON
08.21.49.540 ->A         697.577    .002            .002                                1
08.21.49.974 W<-   EVT              .434            .099    1    2    5    1            1    QWCISCFR   ADR=000000 ADR=000000 ADR=000000
                                                                 9         1                 PAG=    0  XSum=    0 PWrt=   3
 ----------  QWCISCFR               .436*           .101    1    2    5    1    9*           PAG=    0  XSum=    0 PWrt=   3
                                                            0    9    0    1   10

08.25.30.256 ->A         220.282    .001            .001                                1
08.25.33.951 W<-   EVT             3.694            .102    1    2    8    1            1
                                                                10         1                 PAG=    0  XSum=    0 PWrt=   3
                                                                                             EAO=    1  Dec =    0 Bin =    0 Flp =    0
 ----------  QWCISCFR            223.977*           .103    1    2    8    1   12*           PAG=    0  XSum=    0 PWrt=   3
                                                            0   10    0    1   11            EAO=    1  Dec =    0 Bin =    0 Flp =    0

08.26.57.371 ->A          83.420    .002            .002                                1
08.27.43.417  A    TSE            46.047            .083    1              1            1    QMHRTMSS   QMHLOGER   QWCISCFR   ADR=000000
                                                                 9                           PAG=    0  XSum=    0 PWrt=   1
08.28.12.252 W<-   EVT            28.834            .099         4    6    6            1    QWCISCFR   ADR=000000 ADR=000000 ADR=000000
                                                                10         1                 PAG=    0  XSum=    0 PWrt=   7
 ----------  QWCISCFR            158.303*           .184    1    4    6    7   18*           PAG=    0  XSum=    0 PWrt=   8
                                                            0   19    0    1   20

08.41.58.051      /OFF   825.799
08.41.58.051 *TRACE OFF

---------------------------------------------
J O B   S U M M A R Y   D A T A   (T O T A L S)
---------------------------------------------

                         CPU   ---- Physical I/O Counts -----  ***** Transaction Response Time (Sec/Tns) ******  -BMPL-
                         Sec   ----- Synchronous ------ Async   ******  - Activity Level Time -  Inel    Long    C    I Seize
                         Per    DB   DB  NDB   NDB       Disk   ****            Short   Seize   Time     Wait    u    n  Hold   Key/
                         Tns   Read  Wrt Read  Wrt  Sum  I/O     **    Active   Wait     Cft   A-I/W-I Lck/Oth  r    l  Time  Think
                       ------- ---- ---- ---- ---- ---- -----  -------- ------- ------- ------- ------- -------  --  -- ----- ------
Average                  .129    1    3    6    3   13    14   127.572  19.752    .000    .000    .000 151.851           .5     .0
Count                                                               3       4                              2            1      1
Minimum                  .101                        9    10      .436    .434                        83.420           .5     .0
Maximum                  .184                       18    20   223.977  46.047                       220.282           .5     .0
Total/Job                .388                       39    41  1080.288 Elapsed          .0 Percent CPU Utilization
```

Figure 8-59. Transition Report

The columns in the Transition section are as follows:

Time A column heading that shows the time the transaction from one state to another occurred, in the HH.MM.SS.mmm arrangement.

State The three possible states the job can be in are shown in Figure 8-60 on page 8-82.

```
──────►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 type of transition is shown below.  The
        direction and position of the arrow indicates which
        transition was made.  For example:

W A I
─────
W─►         Job went from wait-to-active state (this defines the
            beginning of a transaction in the report).

──►I        Job went from wait-to-ineligible state.

  A◄─       Job went from ineligible-to-active state.

  W         Job went from active-to-wait state but stayed in the
            activity level (for example, a short wait).

  A         Job went from active-to-active state, staying in the
            activity level (for example, a time-slice end occurred
            and no jobs waiting for an activity level).

W◄─         Job went from active-to-wait state (this entry is the end
            of a transaction in the report).
```

Figure 8-60. Possible Job States

Wait Code 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.

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/RESP* 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 --).

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.

SZRL Seize/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 seize conflict to be released.

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.

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).

Elapsed Time—Seconds
Shows the time spent by the job, in the following columns:

Long Wait Elapsed times in the state (such as waiting for the next transaction or lock-wait time).

Active/Rsp
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.

Inel Wait The time the job spent in the ineligible wait state waiting for an activity level.

CPU Sec The processing unit time use by the job in this state.

Physical I/O Counts
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.

Cur MPL The number of jobs holding an activity level in the storage pool.

Cur Inl MPL
The number of jobs waiting for an activity level (ineligible) in the storage pool.

Last 4 Programs in Invocation Stack
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. See Appendix A, "Defining Transaction Boundaries," for more information about the transaction boundary.

Usually, the third or fourth program in the stack is the program shown in the transaction summary PGMNAME data. However, if the *Wait Code* column has a value, the program in the column labeled *Last* 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.

The values ADR=000000 or ADR=UNKWN can also appear as the program name. The ADR=000000 occurs when there was no program active at that level in the job when the trace record was created. ADR=UNKWN indicates that the program did not exist on the system at the time the trace record data was dumped to a database file. This happens if you have deleted (or replaced) the program before the monitor stops. The program names are put into the trace record when the trace data is put into a database file when the monitor ends or when the Dump Trace (DMPTRC) command is used.

---------- (pgmname)
The transaction totals record. For example, ---------- QMHGSD, as shown in Figure 8-59 on page 8-81. This report line occurs each time the job has an active-to-wait transaction. Totals are created for RESP* (response) time, *CPU SECS*, and I/O counts for the transaction.

PAG The number of process access group exceptions.

XSum The number of checksum I/O operations.

PWrt The number of permanent write I/O operations.

EAO The number of effective address overflow exceptions.

Dec The number of decimal overflow exceptions.

Bin The number of binary overflow exceptions.

Flp The number of floating point overflow exceptions.

The summary section of the Transition Report shows the same information as the summary section of the Transaction Report, described in "Job Summary Data" on page 8-78.

Changing the Job Type Data

Use the Change Job Type (CHGJOBTYP) command to change the job type for jobs that appear on the reports you produce using the PRTTNSRPT command.

Note: Using the CHGJOBTYP command will not change any of the information that appears on the System or Component reports.

Use the CHGJOBTYP command for these reasons:

- A job runs as a batch job, but has the characteristics of an interactive job. For example, the batch job runs at the same (or higher) priority as interactive work, and handles requests from interactive users. By changing the job type from batch to interactive, this batch job and the system are more accurately characterized on the Transaction Report.
- A job runs interactively, but has the characteristics of a batch job. In this case, changing the job type from interactive to batch improves the interactive characterization shown on the Transaction Report. For example, if you did not change the job type, the processing unit time per interactive transaction would be distorted.
- To correct a missing job type. If the trace data that the Transaction Report is based on does not contain the job type, performance tools try to correlate the jobs found in the trace data with the job types in the sample data. When this cannot be done, there may be jobs that do not show a job type. Use CHGJOBTYP to assign a job type.

Changing Job Type Values

The possible job type values are:

A Autostart
B Batch
C Programmable work station application server
D Target distributed data management (DDM) server
I Interactive
L Licensed internal code
M Subsystem monitor
R Spool reader
S System
W Spool writer
X Start the system

Operational Considerations

The CHGJOBTYP command reads all records in the trace data file member (in library/QAPMDMPT) to build a subfile that contains the names of all jobs that were traced.

If the trace data has not been sorted into job sequence, the command must sort it before reading the trace records.

Because the command runs interactively, and because it reads (and perhaps sorts) the file, the time the command takes to run can be lengthy. The time could significantly affect other interactive jobs. You would normally run CHGJOBTYP only once per trace file member.

Using the Print Transaction Report Command

The example in this section describes how you can use the PRTTNSRPT command and the Transaction and Transition reports to identify problems and produce detailed information about the problem. Figure 8-61 shows a portion of a Job Summary section obtained by using the RPTTYPE(*SUMMARY) parameter on the PRTTNSRPT command.

Job Summary Report 05/17/90 4:06:07
Job Summary Page 0010

Member . . . : TJST41 Model/Serial . . : B60/10-12883 Main storage . . : 96.0 M Started . . . . : 05/16/90 14:00:28
Library . . : TJSWORK2 System name . . . : RCH38366 Version/Release : 2/ 1.0 Stopped . . . . : 05/16/90 15:27:21

Job Name	User Name	*On/Off* Job Number	Pl	Typ	Pty	Prg	Tot Nbr Tns	Response Sec Avg	Response Sec Max	CPU Sec Util	CPU Sec Avg	CPU Sec Max	Synchronous DBR	Synchronous NDBR	Synchronous Wrt	Synchronous Sum	Synchronous Max	Async Sum	Async Max	Number Cft Lck	Number Cft Sze	K/T /Tns Sec
WS015	USR003	131713	03	1			10	5.3	21.4		.37	.70	8	24	6	38	84				2	58
WS020	QPGMR	131685	05	B								1.49					135			1	8	
WS012	USR005	131714	03	1			8	3.8	8.4		.39	.88	6	17	8	31	57				3	5
WS010	QPGMR	131700	05	B								31.92					1024			2	4	
WS008	USR007	131715	03	1			66	4.9	23.4		.56	1.61	8	26	10	44	168				58	62
WS011	USR001	131718	03	1			70	14.1	216.5		.82	8.12	73	36	33	81	817			6	154	35
WS005	USR009	131719	03	1			10	6.9	15.1		.64	1.77	7	31	5	43	81				12	415
WS019	USR010	131720	03	1			19	7.0	21.4		.32	.95	1	20	4	25	66				2	214
WS017	USR015	131721	03	1			135	3.9	26.7		.26	.93	5	15	5	25	98			1	78	29
WS003	QPGMR	131722	04	1			6	9.5	26.9		.36	.53	1	30	4	35	68			3	3	153
WS001	USR017	131723	03	1			116	3.7	21.3		.19	.79	4	15	6	25	112				37	34

Figure 8-61. Job Summary-Level Report

Consider job WS011 USR001 131718 on the Job Summary-Level Report. The average response time of 14.1 seconds, with 0.82 processing unit seconds per transaction, and maximum values of 216.5 seconds with 8.12 processing unit seconds per transaction for this job, call out the need for further investigation.

Use the RPTTYPE(*TNSACT) parameter on the PRTTNSRPT command to produce transaction information. Specify 131718 on the SLTJOB parameter to produce information for job number 131718 only. Figure 8-62 shows a portion of the output for this request.

Transaction Report 05/18/90 0:46:59
Page 0004

Member . . . : TJST41 Model/Serial . . : B60/10-12883 Main storage . . : 96.0 M Started . . . . : 05/16/90 14:00:28
Library . . : TJSWORK2 System name . . . : RCH38366 Version/Release : 2/ 1.0 Stopped . . . . : 05/16/90 15:27:21
Job name . . : WS011 User name . . . . : USR001 Job number . . . : 131718 TDE/Pl/Pty/Prg . : 04E9/03/20/NO

Time	Excp	Program Name	CPU Sec Per Tns	Synchronous DB Read	Synchronous DB Wrt	Synchronous NDB Read	Synchronous NDB Wrt	Synchronous Sum	Async Disk I/O	****** **** **	Active	Short Wait	Seize Cft	Inel Time A-I/W-I	Long Wait Lck/Oth	BMPL Cur	BMPL Inl	Seize Hold Time	Key/ Think
14.38.25		PGM20A	.242	4		10	9	23		1.108	1.108					11			.000
14.38.28		PGM21A	.472	1		21	7	29		1.975	1.975					13			2.885
14.38.35		PGM22A	.311			7	9	16		1.130	1.130					12			.000
14.38.39		PGM23A	.313	4		17	9	30		1.150	1.150					13			.000
14.38.46		PGM24A	.276			13	7	20		.832	.832					18			1.645
14.38.48		PGM25A	3.477	286		52	24	362		71.105	71.105		.665			18			.000
14.40.01		PGM26A	.446	10		45	9	64		10.124	10.124		5.200			16			.000
14.40.20		PGM27A	.471	5		20	7	32		2.438	2.438					11			.000
14.40.26		PGM28A	.815	3		28	10	41		5.109	5.109					14			2.250
14.40.31		PGM29A	.346	36		25	8	69		3.907	3.907					19			.000

Figure 8-62. Transaction Report

The Transaction Report shows each transaction entered during the measurement period for the job WS011 USR011 131718. Of particular interest on this report is the transaction with very high processing unit seconds per transaction and high I/O counts. A Transition Report shows additional information about these transactions. To produce a Transition Detail Report, use the RPTTYPE(*TRSIT) parameter on the PRTTNSRPT command. Specify 131718 on the SLTJOB parameter. Specify 143825 as the start time and 144301 as the end time on the period parameter. This limits the output to include information only on the 14.38.48 transaction. Figure 8-63 shows a portion of the output from this request.

```
                                         Transition Report                                              05/18/90  4:46:58
                                                                                                                 Page 0011

Member  . . . : TJST41       Model/Serial  . . :  B60/10-12883  Main storage . . :   96.0 M   Started  . . . . : 05/16/90 14:00:28
  Library . . : TJSWORK2     System name . . . :      RCH38366  Version/Release  :   2/ 1.0   Stopped  . . . . : 05/16/90 15:27:21
Job name  . . : WS011        User name . . . . :     USR001     Job number . . . :   131718   TDE/Pl/Pty/Prg . :    04E9/03/20/NO

Job type  . . : I         Elapsed Time -- Seconds          Sync/Async Phy I/O     -MPL-
                          -----------------------        ------------------------ C   I     Last 4 Programs in Invocation Stack
             State Wait   Long   Active   Inel    CPU     DB   DB  NDB   NDB      u   n -------------------------------------------
    Time     W A I Code   Wait    /Rsp*   Wait    Sec    Read  Wrt Read  Wrt  Tot r   l     Last      Second      Third      Fourth
------------ ----- ---- ------- ------- ------- -------- ---- ---- ---- ---- ---- -- -- ---------- ---------- ---------- ----------
14.38.46.958 ->A          6.805                    .074             11            18
14.38.46.789 W<-                   .832            .202              2     7      18
 ----------  PGM20A                .832*           .276    0        13     7   20*

14.38.48.207 ->A           .418                    .060             12            18
14.38.48.905       SZWT      */   1.191/*                             HOLDER-- BY1    MSBY       130822
                                                                       OBJECT- MBR  PMMS    MSIT        MSIT
14.38.59.987       SZRL                                                WAITER- OR27   BYRKS      131285              .011 SECS
                                                                       OBJECT- MBR  PMMS    SPHD        SPHD
14.39.00.321       SZWT      */    .474/*                             HOLDER-- OR27   BYRLS      131285
                                                                       OBJECT- MBR  PMMS    SPHD        SPHD
14.39.08.197   A  TSE             19.989          1.873  169        12            12   QDBGETSQ  PGM20A      PGM40B     PGM1
14.39.59.313 W<-                  51.116          1.544  117        28    24      19   QT3REQIO  QWSGET      PGN20A
 ----------  PGM20A               71.105*         3.477  286        52    24  362*
14.40.01.788 ->A          2.475                    .073             13            16
14.40.03.620       SZWT      */   5.200/*                             HOLDER-- BY1    MSBY       130822
                                                                      OBJECT- MBR   PMMS    MSID       MSID
14.40.11.912 W<-                  10.124           .373   10        32     9      14
 ----------  PGM20A               10.124*          .446   10        45     9   64*
```

Figure 8-63. Transition Report

The Transition Report shows detail about the job transition that occurred during the 14.38.48 transaction. With this detail you can see some of the application programs that were involved in the transaction. For example, PGM20A, PGM40B, and PGM1 are the user application programs used. Even at this level of detail, much processing takes place that is not shown in this report. Only when the job makes a transition from one of the states (active, waiting, or ineligible) is any information captured for this report.

In addition to the state transition information, the Transition Report shows details about the wait time during the transaction. The codes SZWT and SZRL identify seize waits and seize releases. The TSE code identifies time slice ends.

The following information is available on the Transition Report about the 14.38.48 transaction:

- At 14.38.48.207 the transaction became active after a wait of 0.418 seconds (the key+think time).
- The transaction was active for 19.989 seconds, during which there were two seize waits (one for 0.191 and the other for 0.474 seconds).
- The object, and job holding it, are shown for each seize occurrence.
- The job released a seize it had. At 14.39.08.197, following the 19.989 seconds of active time, the transaction exceeded the time slice assigned to the job. Due to the job's priority and the availability of an activity slot, the transaction did not go to the ineligible state. It continued active for another 51.116 seconds. The transaction completed, as indicated by the transition to the wait state.

- The total transaction response time was 71.105 seconds.
- The application program used to label this transaction is PGM20A.
- Other application programs were involved in this transaction, as shown by the programs listed at the time slice end (TSE) transition.

  Additional processing and programs may have been involved in this transaction that are not listed on this report. To learn more about this application, use some of the utilities described in Chapter 11, "Programmer Performance Utilities." For example, use the trace job and sampled address monitor to collect additional data on future occurrences of one of these transactions. (Your ability to collect additional information on transactions assumes that from the information shown on the Transaction Report, you can determine what the end-user transaction is.)

  By analyzing the application and programs involved in the transaction, you can often improve the transaction, and can free resources that improve overall system throughput and response times.

Job Interval Report

To produce the Job Interval Report, use the Print Job Report (PRTJOBRPT) command, or select option 5 (Job report) on the Print Performance Reports display.

This report, like the others described in this chapter, is produced from the sample data collected with the STRPFRMON command. The four major sections of this report 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.

If a value is too large to fit in the allotted space, a 9 is printed in each numeric field in the report.

The heading printed on each page of the report, shown in Figure 8-64, includes information on the performance data set from which the report is printed and information about the machine on which the data was collected.

```
                                             Job Interval Report                                     04/19/90 13:47:50
                                            Interactive Job Summary                                    Page    1
                                           User-Selected Report Title

Member . . . : WASCOMTEST  Model/Serial  . . :  B60/10-12883  Main storage  . . :   96.0 M  Started  . . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG     System name . . . :      TEST#366  Version/Release . :    2/1.0  Stopped  . . . . : 01/18/90 09:34:00
```

Figure 8-64. Job Interval Report: Header Information

The header information consists of the following:

- Report title
- Current date and time
- Report page number
- User-selected report title
- Data member name
- Library name
- Model number
- Serial number
- System name

- Main storage size
- OS/400 version and release level
- Data collection start date and time
- Data collection stop date and time

Interactive Job Summary

The Interactive Job Summary section of the Job Interval Report, shown in Figure 8-65, lists one line for all selected interactive jobs that existed during each selected interval (a total of one line per interval).

```
                                  Job Interval Report                                                  04/19/90 13:47:50
                                Interactive Job Summary                                                      Page     1
                               User-Selected Report Title

Member . . . : WASCOMTEST   Model/Serial  . . :  B60/10-12883  Main storage  . . :   96.0 M   Started  . . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG     System name . . . :      TEST#366  Version/Release . :    2/1.0   Stopped  . . . . : 01/18/90 09:34:00

 Itv    Act       Tns      Rsp/   ------------- Number of I/O --------------    Tns/    CPU    PAG    EAO     XSum    Perm   Arith
 End    Jobs     Count     Tns       Sync      Async     Logical      Cmn       Hour   Util   Fault   Excp    I/O     Write  Ovrflw
-----  -----  ----------  ------  ---------  ---------  ---------  ---------  -------  ----  ------  ------  ------  ------  ------
09:19      2          58    2.43       1075          0       5361          0      239   4.7       0       0       0       0       0
09:34      2          17    1.41        207          0        775          0       68    .5       0       0       0       0       0
```

Figure 8-65. Job Interval Report: Interactive Job Summary Section

The columns in the Interactive Job Summary section of the Job Interval Report are as follows:

Itv End The interval end time (hour and minute).

Act Jobs The number of selected interactive jobs that were active during the interval.

Tns Count
The number of transactions performed by the selected interactive jobs during the interval.

Rsp/Tns The average response time in seconds 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.

Sync The number of synchronous disk I/O operations performed by the selected interactive jobs during the interval.

Async The number of asynchronous disk I/O operations started by the selected interactive jobs during the interval.

Logical The number of logical disk I/O operations performed by the selected interactive jobs during the interval.

Cmn The number of communications I/O operations performed by the selected interactive jobs during the interval.

Tns/Hour The average number of transactions per hour processed by the selected interactive jobs during the interval.

CPU Util The percentage of available processing unit time used by the selected interactive jobs in the interval. For multiple-processor systems, this is the total used divided by the number of processors.

PAG Fault
The number of faults involving the process access group that occurred for the selected interactive jobs during the interval.

EAO Excp The number of effective address overflow exceptions that occurred for the selected interactive jobs during the interval.

XSum I/O The number of checksum I/O operations that occurred for the selected interactive jobs during the interval.

Perm Write
The number of permanent write operations performed for the selected interactive jobs during the interval.

Arith Ovrflw
The number of arithmetic overflow exceptions that occurred for the selected interactive jobs during the interval.

Noninteractive Job Summary

The Noninteractive Job Summary section of the Job Interval Report, shown in Figure 8-66, lists one line for all selected noninteractive jobs that existed during each selected interval (a total of one line per interval).

```
                                          Job Interval Report                                          04/19/90 13:47:50
                                       Noninteractive Job Summary                                           Page    2
                                       User-Selected Report Title

Member  . . . : WASCOMTEST   Model/Serial  . . :  B60/10-12883   Main storage  . . :   96.0 M   Started  . . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG      System name . . . :      TEST#366   Version/Release . :    2/1.0   Stopped  . . . . : 01/18/90 09:34:00

 Itv    Act    CPU    ---- Number of I/O Per Second ----   - CPU/ I/O -   Line    Page    PAG      EAO      XSum     Perm     Arith
 End    Jobs   Util     Sync    Async   Logical     Cmn   Sync   Async   Count   Count   Fault    Excp     I/O      Write    Ovrflw
-----  -----   ----   -------  -------  -------  -------  -----  -----  -------  -----  -------  -------  -------  -------  -------
09:19      2     .3       .2       .0       .0       .0     17      0        0      0        0        0        0        0        0
09:34      2     .0       .0       .0       .0       .0     16      0        0      0        0        0        0        0        0
```

Figure 8-66. Job Interval Report: Noninteractive Job Summary Section

The columns in the Noninteractive Job Summary section of the Job Interval Report are as follows:

Itv End The interval end time (hour and minute).

Act Jobs The number of selected noninteractive jobs that were active during the interval.

CPU Util The percentage of available processing unit time used by the selected noninteractive jobs during the interval. For multiple-processor systems, this is the total used divided by the number of processors.

Sync I/O Per Second
The average number of synchronous disk I/O operations performed per second by the selected noninteractive jobs during the interval.

Async I/O Per Second
The average number of asynchronous disk I/O operations started per second by the selected noninteractive jobs during the interval.

Logical I/O Per Second
The average number of logical disk I/O operations performed per second by the selected noninteractive jobs during the interval.

Cmn I/O Per Second
The average number of communications I/O operations performed per second by the selected noninteractive jobs during the interval.

CPU/Sync I/O
The average number of processing unit milliseconds used for each synchronous disk I/O operation performed by the selected noninteractive jobs during the interval.

CPU/Async I/O
The average number of processing unit milliseconds used for each asynchronous disk I/O operation started by the selected noninteractive jobs during the interval.

Line Count
The number of lines printed by the selected noninteractive jobs during the interval.

Page Count
The number of pages printed by the selected noninteractive jobs during the interval.

PAG Fault
The number of faults involving the process access group that occurred for the selected noninteractive jobs during the interval.

EAO Excp The number of effective address overflow exceptions that occurred for the selected noninteractive jobs during the interval.

XSum I/O The number of checksum I/O operations that occurred for the selected noninteractive jobs during the interval.

Perm Write
The number of permanent write operations performed for the selected noninteractive jobs during the interval.

Arith Ovrflw
The number of arithmetic overflow exceptions that occurred for the selected noninteractive jobs during the interval.

Interactive Job Detail

The Interactive Job Detail section of the Job Interval Report, shown in Figure 8-67, gives detailed information by interval and job. One line is printed for each selected interactive job that existed during each selected interval (generally more than one line per interval).

```
                                              Job Interval Report                                           04/19/90 13:47:50
                                             Interactive Job Detail                                                Page    3
                                           User-Selected Report Title

Member  . . . : WASCOMTEST   Model/Serial  . . :  B60/10-12883  Main storage  . . :    96.0 M  Started  . . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG      System name . . . :      TEST#366  Version/Release . :     2/1.0  Stopped  . . . . : 01/18/90 09:34:00

                                                                 -------- Physical I/O per Transaction  ------             SYNC
 Itv     Job         User      Job                TNS    Rsp     CPU    ---- Synchronous -----  ---- Asynchronous ----  CPU     I/O
 End     Name        Name      Number  PL  Pty   /HR    /Tns    /Tns     DBR   DBW  NDBR  NDBW   DBR   DBW  NDBR  NDBW  Util    /Sec
-----  ----------  ----------  ------  --  ---  -----  ------  ------  ----  ----  ----  ----  ----  ----  ----  ----  ----  -------
09:19  DDWEBER     MEP         006209   3   20    161    2.15    .987    .5    .0  10.6    .0    .0    .0    .0    .0   4.4      1.2
09:19  NOVEY       NOVEY       006232   3   20     78    3.00    .147    .0    .0   1.0    .0    .0    .0    .0    .0    .3       .0
09:34  DDWEBER     MEP         006209   3   20     12     .66    .408    .0    .0   1.6    .0    .0    .0    .0    .0    .1       .0
09:34  NOVEY       NOVEY       006232   3   20     56    1.57    .253   3.6    .0   9.0    .0    .0    .0    .0    .0    .3       .2
```

Figure 8-67. Job Interval Report: Interactive Job Detail Section

The columns in the Interactive Job Detail section of the Job Interval Report are as follows:

Itv End The interval end time (hour and minute).

Job Name
The name of the job.

User Name
The name of the user.

Job Number
The number of the job.

PL The number of the pool in which the job ran.

Pty The priority at which the job ran.

Tns/Hour The average number of transactions per hour for the job during the interval.

Rsp/Tns The average response time (seconds) per transaction.

CPU/Tns 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.

Synchronous DBR
The average number of synchronous database read operations per transaction for the job during the intervals. This is calculated from the synchronous database disk read count divided by the number of transactions processed.

Synchronous DBW
The average number of synchronous database write operations per transaction for the job during the intervals. This is calculated from the synchronous database disk write count divided by the number of transactions processed.

Synchronous NDBR
The average number of synchronous nondatabase read operations per transaction for the job during the intervals. This is calculated from the synchronous nondatabase disk read count divided by the number of transactions processed.

Synchronous NDBW
The average number of synchronous nondatabase write operations per transaction for the job during the intervals. This is calculated from the synchronous nondatabase disk write count divided by the number of transactions processed.

Asynchronous DBR
The average number of asynchronous database read operations per transaction for the job during the intervals. This is calculated from the asynchronous database disk read count divided by the number of transactions processed.

Asynchronous DBW
The average number of asynchronous database write operations per transaction for the job during the intervals. This is calculated from the asynchronous database disk write count divided by the number of transactions processed. This field will not be printed if the job did not process any transaction during the interval.

Asynchronous NDBR
The average number of asynchronous nondatabase read operations per transaction for the job during the intervals. This is calculated

from the asynchronous nondatabase disk read count divided by the number of transactions processed.

Asynchronous NDBW
The average number of asynchronous nondatabase write operations per transaction for the job during the intervals. This is calculated from the asynchronous nondatabase disk write count divided by the number of transactions processed.

CPU Util The percentage of available processing unit time used by the job during the interval. For multiple-processor systems, this is the total used divided by the number of processors.

Sync I/O /Sec
The average number of synchronous disk I/O operations per second performed by the job during the interval.

Noninteractive Job Detail

The Noninteractive Job Detail section of the Job Interval Report, shown in Figure 8-68, gives detailed information by interval and job. One line is printed for each selected noninteractive job that existed during each selected interval (generally more than one line per interval).

```
                                               Job Interval Report                                        04/19/90 13:47:50
                                             Noninteractive Job Detail                                    Page    4
                                            User-Selected Report Title

Member  . . . : WASCOMTEST  Model/Serial  . . :  B60/10-12883  Main storage  . . :   96.0 M  Started  . . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG     System name . . . :      TEST#366  Version/Release . :    2/1.0  Stopped  . . . . : 01/18/90 09:34:00

 Itv      Job        User      Job                            Elapsed   CPU   --- Nbr I/O /Sec ---  - CPU / I/O -  --- Printer ---
 End      Name       Name      Number   Pool   Type   Pty      Time    Util   Sync   Async   Lgl    Sync   Async    Lines   Pages
-----  ----------  ----------  ------   ----   ----   ---    -------   ----   -----  -----   -----  -----  -----   -------  -----
09:19  QNFTP       QSNADS      006181     2      B     40      14:31     .3      0       0      0     22      0        0       0
09:19  QROUTER     QSNADS      006176     2      B     40      14:31     .0      0       0      0      9      0        0       0
09:34  QNFTP       QSNADS      006181     2      B     40      14:58     .0      0       0      0     27      0        0       0
09:34  QROUTER     QSNADS      006176     2      B     40      14:58     .0      0       0      0      9      0        0       0
```

Figure 8-68. Job Interval Report: Noninteractive Job Detail Section

The columns in the Noninteractive Job Detail section of the Job Interval Report are as follows:

Itv End The interval end time (hour and minute).

Job Name
The name of the job.

User Name
The name of the user.

Job Number
The number of the job.

Pool The number of the pool in which the job ran.

Type The type and subtype of the job.

Pty The priority at which the job ran.

Elapsed Time
Elapsed time, which is the amount of time (minutes and seconds) for which the job existed during the interval. This will be the same as the interval length unless the job started or ended during the interval, in which case it will be less.

CPU Util The percentage of available processing unit time used by the job during the interval. For multiple-processor systems, this is the total used divided by the number of processors.

Sync I/O /Sec
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.

Async I/O /Sec
The average number of asynchronous disk I/O operations started per second by the job during the interval. This is calculated from the asynchronous disk I/O count divided by the elapsed time.

Lgl I/O /Sec
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.

CPU/Sync I/O
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.

CPU/Async I/O
The average number of milliseconds of processing unit time taken for each asynchronous disk I/O operation. This is calculated from the milliseconds of the processing unit time used by the job divided by the asynchronous disk I/O count.

Printer Lines
The number of lines printed by the job during the interval.

Printer Pages
The number of pages printed by the job during the interval.

Report Selection Criteria

The Report Selection Criteria section of the Job Interval Report, shown in Figure 8-69 on page 8-95, gives the selection values you chose to produce the report.

```
                                          Job Interval Report                                        04/19/90 13:28:36
                                        Interactive Job Detail                                              Page    3
                                      User-Selected Report Title

Member . . . : WASCOMTEST   Model/Serial  . . :  B60/10-12883   Main storage  . . :    96.0 M   Started  . . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG     System name . . . :      TEST#366   Version/Release . :     2/1.0   Stopped  . . . . : 01/18/90 09:34:00

Select Parameters

Pools                       - 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16

Jobs                        - 012345/Useridwxyz/Jobname123
                              987654/Useridabcd/Jobname456

User IDs                    - User1      User2      User3      User4     User5      User6
                              User7      User8      User9      User10    User11     User12

Subsystems                  - Subsystem1 Subsystem2 Subsystem3 Subsystem4 Subsystem5 Subsystem6
                              Subsystem7 Subsystem8 Subsystem9 Subsystema Subsystemb Subsystemc

Communications Lines        - Line1      Line2      Line3      Line4     Line5      Line6
                              Line7      Line8      Line9      Line10    Line11     Line12

Control Units               - Ctlr1      Ctlr2      Ctlr3      Ctlr4     Ctlr5      Ctlr6
                              Ctlr7      Ctlr8      Ctlr9      Ctlr10    Ctlr11     Ctlr12

Functional Areas            - Accounting           Payroll              Research
                              Development          ProjectX             MrNolansStaff

                            - No Select parameters were chosen.

Omit Parameters

Pools                       - 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16

Jobs                        - 012345/Useridwxyz/Jobname123
                              987654/Useridabcd/Jobname456

User IDs                    - User1      User2      User3      User4     User5      User6
                              nnnnnn     User8      User9      User10    User11     User12

Subsystems                  - Subsystem1 Subsystem2 Subsystem3 Subsystem4 Subsystem5 Subsystem6
                              Subsystem7 Subsystem8 Subsystem9 Subsystema Subsystemb Subsystemc

Communications Lines        - Line1      Line2      Line3      Line4     Line5      Line6
                              Line7      Line8      Line9      Line10    Line11     Line12

Control Units               - Ctlr1      Ctlr2      Ctlr3      Ctlr4     Ctlr5      Ctlr6
                              Ctlr7      Ctlr8      Ctlr9      Ctlr10    Ctlr11     Ctlr12

Functional Areas            - Accounting           Payroll              Research
                              Development          ProjectX             MrNolansStaff

                            - No Omit parameters were chosen.
```

Figure 8-69. Job Interval Report: Report Selection Criteria Section

THE information shown in the Report Selection Criteria section of the Job Interval Report is as follows:

Select Parameters
The criteria used to choose the data records to be included in the report. The criteria are generally specified using a SLTxxx parameter of the PRTJOBRPT command. Only nondefault values (something other than *ALL) are printed. If a parameter is not specified, it does not appear on the report.

Omit Parameters
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 PRTJOBRPT command. Only nondefault values (something other than *NONE) are printed. If a parameter was not specified, it does not appear on the report.

Pools The list of pools selected to be included (SLTPOOLS parameter) or excluded (OMTPOOLS parameter).

Jobs The list of jobs selected to be included (SLTJOB parameter) or excluded (OMTJOB parameter). This does not include jobs selected using the STLFCNARA or OMTFCNARA parameter.

User ID The list of users selected to be included (SLTUSRID parameter) or excluded (OMTUSRID parameter).

Subsystems
The list of subsystems selected to be included (SLTSBS parameter) or excluded (OMTSBS parameter).

Communications Lines
The list of communications lines selected to be included (SLTLINE parameter) or excluded (OMTLINE parameter).

Control Units
The list of control units selected to be included (SLTCTL parameter) or excluded (OMTCTL parameter).

Functional Areas
The list of functional areas selected to be included (SLTFCNARA parameter) or excluded (OMTFCNARA parameter). For more information on functional areas, see Chapter 12, "Managing the Performance Tools Configuration."

Pool Interval Report

To produce the Pool Interval Report, use the Print Pool Report (PRTPOLRPT) command or select option 6 (Pool report) on the Print Performance Reports display.

This report, like the others in this chapter, is produced from the sample data collected by the STRPFRMON command. The Pool Report contains 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.

If a value is too large to fit in the allotted space, a 9 is printed in each numeric field in the report.

The heading printed on each page of the report, shown in Figure 8-70, includes information on the performance data set from which the report is printed and information about the machine on which the data was collected.

```
                                              Pool Interval Report                                    04/19/90 13:28:36
                                               Subsystem Activity                                     Page    1
                                            User-Selected Report Title

Member  . . . : WASCOMTEST  Model/Serial  . . : B60/10-12883   Main storage  . . : 96.0 M     Started . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG     System name . . . :      TEST#366  Version/Release . :  2/1.0     Stopped . . . : 01/18/90 09:34:00
```

Figure 8-70. Pool Interval Report: Header Information

This header information consists of:

- Report title
- Current date and time
- Report page number

- User-selected report title
- Data member name
- Library name
- Model number
- Serial number
- System name
- Main storage size
- OS/400 version and release level
- Data collection start date and time
- Data collection stop date and time

Subsystem Activity

The Subsystem Activity section of the Pool Interval Report, shown in Figure 8-71, gives the performance information on the subsystems during each selected interval. One line is printed for each subsystem and active pool combination that existed during each selected interval.

```
                                                    Pool Interval Report                                    04/19/90 13:28:36
                                                     Subsystem Activity                                            Page    1
                                                 User-Selected Report Title

Member  . . . : WASCOMTEST   Model/Serial  . . : B60/10-12883    Main storage  . . : 96.0 M      Started . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG      System name . . . :       TEST#366  Version/Release . :  2/1.0      Stopped . . . : 01/18/90 09:34:00

                                            --------- Physical I/O per Transaction ---------  ------------- Job Maximums ------------
 Itv   Subsystem        CPU                 ----- Synchronous -----  ---- Asynchronous -----  CPU   Phy
 End   Name         PL  Util      Tns         DBR   DBW  NDBR  NDBW   DBR   DBW  NDBR  NDBW  Util  I/O     Tns    Rsp     A-W  W-I A-I
-----  ----------   --  ----  ----------    ----- ----- ----- -----  ----- ----- ----- -----  ---- ------ ----- ------ ------ --- ---
09:19  QBASE         2    .5          32                  1.5                                  .9     41    30    .16     45   0   0
09:19  QBASE         3  19.3         472       .1         7.4                                 6.1  2,975    82   3.00    113   0   0
09:19  QSNADS        2   1.1           0                                                       .3    423     0    .00    460   0   0

09:34  QBASE         2   3.2           2      6.5       176.5                                79.1    199     1   2.00     37   0   0
09:34  QBASE         3   4.3         310       .2         2.9                                 1.4    378   121   2.25    133   0   0
09:34  QSNADS        2    .2           0                                                       .0    158     0    .00     87   0   0
```

Figure 8-71. Pool Interval Report: Subsystem Activity Section

The columns on the Subsystem Activity section of the Pool Interval Report are as follows:

Itv End The interval end time (hour and minute).

Subsystem Name
The name of the subsystem.

PL The number of the pool in which the subsystem ran.

CPU Util The percentage of available processing unit time used by the selected jobs in the subsystem. For a multiple-processor system, this is the total used divided by the number of processors.

Tns The number of transactions processed for the selected jobs in the subsystem.

Synchronous DBR
The average number of synchronous database read operations per transaction for the jobs in the system during the interval. This is calculated from the synchronous database read count divided by the transactions processed. This field is not printed if the jobs in the system did not process any transactions.

Synchronous DBW
The average number of synchronous database write operations per transaction for the jobs in the system during the interval. This is cal-

culated from the synchronous database write count divided by the transactions processed. This field is not printed if the jobs in the system did not process any transactions.

Synchronous NDBR
The average number of synchronous nondatabase read operations per transaction for the jobs in the system during the interval. 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
The average number of synchronous nondatabase write 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.

Asynchronous DBR
The average number of asynchronous database read operations per transaction for the jobs in the system during the interval. This is calculated from the asynchronous database read count divided by the transactions processed. This field is not printed if the jobs in the system did not process any transactions.

Asynchronous DBW
The average number of asynchronous database write operations per transaction for the jobs in the system during the interval. This is calculated from the asynchronous database write count divided by the transactions processed. This field is not printed if the jobs in the system did not process any transactions.

Asynchronous NDBR
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.

Asynchronous NDBW
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.

Job Maximum CPU Util
The highest percentage of available processing unit time used by a selected job in the subsystem.

Job Maximum Phy I/O
The highest number of physical disk input and output operations by a selected job in the subsystem.

Job Maximum Tns
The highest number of transactions by a selected job in the subsystem.

Job Maximum Rsp
: The highest average response time in seconds per transaction by a selected job in the subsystem. The response time is the amount of time spent waiting for and using the resources divided by the number of transactions processed.

Job Maximum A-W
: The highest number of active-to-wait state transitions by a selected job in the subsystem.

Job Maximum W-I
: The highest number of wait-to-ineligible state transitions by a selected job in the subsystem.

Job Maximum A-I
: The highest number of active-to-ineligible state transitions by v selected job in the subsystem.

Pool Activity

The Pool Activity section of the Pool Interval Report, shown in Figure 8-72, gives the performance information on the storage pools at various time intervals. One line is printed for each active pool that existed during each selected interval.

```
                                                Pool Interval Report                                            04/19/90 13:28:36
                                                  Pool Activity                                                       Page    2
                                              User-Selected Report Title

Member . . . : WASCOMTEST   Model/Serial  . . : B60/10-12883   Main storage  . . : 96.0 M     Started . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG      System name . . . :      TEST#366   Version/Release . :  2/1.0     Stopped . . . : 01/18/90 09:34:00

                                         -------- Physical I/O per Transaction ----------  ------------- Job Maximums ------------
 Itv       Act   Size  CPU               ----- Synchronous -----  ---- Asynchronous -----  CPU   Phy
 End   PL  Lvl   (K)   Util      Tns     DBR   DBW   NDBR  NDBW   DBR   DBW   NDBR  NDBW   Util  I/O    Tns    Rsp      A-W  W-I A-I
 ----- --  ---  -----  ----  ----------  ----- ----- ----- -----  ----- ----- ----- -----  ---- ------ ----- ------ ------ --- ---
 09:19  2    7      7   1.4          32     .1          3.2                                  .9    423    30    .16    460   0   0
 09:19  3   87     77  19.3         472     .1          7.4                                 6.1  2,975    82   3.00    113   0   0

 09:34  2    7      7   3.4           2    8.0        183.5                                79.1    199     1   2.00     87   0   0
 09:34  3   87     77   4.3         310     .2          2.9                                 1.4    378   121   2.25    137   0   0
```

Figure 8-72. Pool Interval Report: Pool Activity Section

The columns on the Pool Activity section of the Pool Interval Report are as follows:

Itv End
: The interval end time (hour and minute).

PL
: The number of the pool.

Act Lvl
: The activity level of the pool during the interval.

Size (K)
: The size of the pool in kilobytes (1024 bytes).

CPU Util
: The percentage of available processing unit time used by the selected jobs in the pool.

Tns
: The total number of transactions processed by the selected jobs in the pool.

Synchronous DBR
: The average number of synchronous database read operations on the disk per transaction for the selected jobs in the pool. It is the total synchronous database reads divided by the total transactions in the pool.

Synchronous DBW
The average number of synchronous database write operations on the disk per transaction for the selected jobs in the pool. It is the total synchronous database writes divided by the total transactions in the pool.

Synchronous NDBR
The average number of synchronous nondatabase read operations on the disk per transaction for the selected jobs in the pool. It is the total synchronous nondatabase reads divided by the total transactions in the pool.

Synchronous NDBW
The average number of synchronous nondatabase write operations on the disk per transaction for the selected jobs in the pool. It is the total synchronous nondatabase writes divided by the total transactions in the pool.

Asynchronous DBR
The average number of asynchronous database read operations on the disk per transaction for the selected jobs in the pool It is the total asynchronous database reads divided by the total transactions in the pool.

Asynchronous DBW
The average number of asynchronous database write operations on the disk per transaction for the selected jobs in the pool. It is the total asynchronous database writes divided by the total transactions in the pool.

Asynchronous NDBR
The average number of asynchronous nondatabase read operations on the disk per transaction for the selected jobs in the pool. It is the total asynchronous nondatabase reads divided by the total transactions in the pool.

Asynchronous NDBW
The average number of asynchronous nondatabase write operations on the disk per transaction for the selected jobs in the pool. It is the total asynchronous nondatabase writes divided by the total transactions in the pool.

Job Maximum CPU Util
The highest percentage of available processing unit time used by a selected job in the pool.

Job Maximum Phy I/O
The highest number of physical disk input and output operations by a selected job in the pool.

Job Maximum Tns
The highest number of transactions by a selected job in the pool.

Job Maximum Rsp
The highest average response time in seconds per transaction by a selected job in the pool. The response time is the amount of time spent waiting for and using the resources divided by the number of transactions.

Job Maximum A-W
The highest number of active-state to wait-state transitions by a selected job in the pool.

Job Maximum W-I
The highest number of wait-state to ineligible-state transitions by a selected job in the pool.

Job Maximum A-I
The highest number of active-state to ineligible-state transitions by a selected job in the pool.

Report Selection Criteria

The Report Selection Criteria section of the Pool Interval Report, shown in Figure 8-73, gives the selection values you chose to produce the report.

```
                                              Pool Interval Report                                      04/19/90 13:28:36
                                            Report Selection Criteria                                         Page    3
                                            User-Selected Report Title

Member  . . . : WASCOMTEST   Model/Serial  . . : B60/10-12883    Main storage  . . : 96.0 M      Started . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG      System name . . . :     TEST#366    Version/Release . :  2/1.0      Stopped . . . : 01/18/90 09:34:00

Select Parameters

Pools                        - 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16

Jobs                         - 012345/Useridwxyz/Jobname123
                               987654/Useridabcd/Jobname456

User IDs                     - User1      User2      User3      User4     User5      User6
                               User7      User8      User9      User10    User11     User12

Subsystems                   - Subsystem1 Subsystem2 Subsystem3 Subsystem4 Subsystem5 Subsystem6
                               Subsystem7 Subsystem8 Subsystem9 Subsystema Subsystemb Subsystemc

Communications Lines         - Line1      Line2      Line3      Line4     Line5      Line6
                               Line7      Line8      Line9      Line10    Line11     Line12

Control Units                - Ctlr1      Ctlr2      Ctlr3      Ctlr4     Ctlr5      Ctlr6
                               Ctlr7      Ctlr8      Ctlr9      Ctlr10    Ctlr11     Ctlr12

Functional Areas             - Accounting          Payroll              Research
                               Development         ProjectX             MrNolansStaff

                             - No Select parameters were chosen.

Omit Parameters

Pools                        - 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16

Jobs                         - 012345/Useridwxyz/Jobname123
                               987654/Useridabcd/Jobname456

User IDs                     - User1      User2      User3      User4     User5      User6
                               nnnnnn     User8      User9      User10    User11     User12

Subsystems                   - Subsystem1 Subsystem2 Subsystem3 Subsystem4 Subsystem5 Subsystem6
                               Subsystem7 Subsystem8 Subsystem9 Subsystema Subsystemb Subsystemc

Communications Lines         - Line1      Line2      Line3      Line4     Line5      Line6
                               Line7      Line8      Line9      Line10    Line11     Line12

Control Units                - Ctlr1      Ctlr2      Ctlr3      Ctlr4     Ctlr5      Ctlr6
                               Ctlr7      Ctlr8      Ctlr9      Ctlr10    Ctlr11     Ctlr12

Functional Areas             - Accounting          Payroll              Research
                               Development         ProjectX             MrNolansStaff

                             - No Omit parameters were chosen.
```

Figure 8-73. Pool Interval Report: Report Selection Criteria Section

The information shown in the Report Selection Criteria section of the Pool Interval Report is as follows:

Select Parameters
The criteria used to select the data records to be included in the report. The criteria are generally specified using a SLTxxx parameter of the PRTPOLRPT command. Only nondefault values (something other than *ALL) are printed. If a parameter was not specified, it does not appear on the report.

Omit Parameters
The criteria used to select the data records to be excluded from the report. The criteria are generally specified using an OMTxxx parameter of the PRTPOLRPT command. Only nondefault values (something other than *NONE) are printed. If a parameter was not specified, it does not appear on the report.

Pools The list of pools selected to be included (SLTPOOLS parameter) or excluded (OMTPOOLS parameter).

Jobs 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.

User IDs The list of users selected to be included (SLTUSRID parameter) or excluded (OMTUSRID parameter).

Subsystems
The list of subsystems selected to be included (SLTSBS parameter) or excluded (OMTSBS parameter).

Communications Lines
The list of communications lines selected to be included (SLTLINE parameter) or excluded (OMTLINE parameter).

Control Units
The list of control units selected to be included (SLTCTL parameter) or excluded (OMTCTL parameter).

Functional Areas
The list of functional areas selected to be included (SLTFCNARA parameter) or excluded (OMTFCNARA parameter). For more information on functional areas, see Chapter 12, "Managing the Performance Tools Configuration."

Resource Interval Report

To produce the Resource Interval Report, use the Print Resource Report (PRTRSCRPT) command, or select option 7 (Resource report) on the Print Performance Reports display.

This report, like others in this chapter, is produced from the sample data collected with the STRPFRMON command. The six major sections of this report 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.

If a value is too large to fit in the allotted space, a 9 is printed in each numeric field in the report.

The heading printed on each page of the report, shown in Figure 8-74, includes information on the performance data set from which the report is printed and information about the machine on which the data was collected.

```
                                          Resource Interval Report                                        04/19/90 13:58:59
                                          Disk Utilization Summary                                        Page    1
                                          User-Selected Report Title

Member  . . . : WASCOMTEST   Model/Serial  . . : B60/10-12883   Main storage  . . : 96.0 M    Started . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG      System name . . . :     TEST#366   Version/Release . :  2/1.0    Stopped . . . : 01/18/90 09:34:00
```

Figure 8-74. Resource Interval Report: Header Information

The header information consists of the following:

- Report title
- Current date and time
- Report page number
- User-selected report title
- Data member name
- Library name
- Model number
- Serial number
- System name
- Main storage size
- OS/400 version and release level
- Data collection start date and time
- Data collection stop date and time

Disk Utilization Summary

The Disk Utilization Summary section of the Resource Interval Report, shown in Figure 8-75, gives detailed disk information by time intervals. Information is shown for all disk arms that are configured on the system. Also, the disk arm with the highest utilization and the disk arm with the highest average seek time for each time interval are shown. Consistent disk arm utilization at or above the threshold value, provided in Figure 10-3 on page 10-6, will affect system performance and cause longer response times and/or less throughput.

```
                                          Resource Interval Report                                        04/19/90 13:58:59
                                          Disk Utilization Summary                                        Page    1
                                          User-Selected Report Title

Member  . . . : WASCOMTEST   Model/Serial  . . : B60/10-12883   Main storage  . . : 96.0 M    Started . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG      System name . . . :     TEST#366   Version/Release . :  2/1.0    Stopped . . . : 01/18/90 09:34:00

             Average       Average      Average      Average                          High      High      High      Disk
  Itv        Physical      Reads        Writes       K per       Avg        High      Util      Srv       Srv       Space
  End        I/O /Sec      /Sec         /Sec         I/O         Util       Util      Unit      Time      Unit      Used
 -----      ---------    ---------    --------     -------      ----       ----      ----      ----      ----     -------
 09:19          10.3          5.0          5.3         1.4        .6        1.5      0009      .039      0033       5,182
 09:34           5.9          3.5          2.3         1.5        .3        1.5      0018      .049      0033       5,186
 -----      ---------    ---------    --------     -------      ----
 Average:        8.0          4.3          3.7         1.5        .5
```

Figure 8-75. Resource Interval Report: Disk Utilization Summary Section

The columns in the Disk Utilization Summary section of the Resource Report are as follows:

Itv End The interval end time (hour and minute).

Average Phys I/O /Sec
The average number of physical disk read and write operations per second made on all disks on the system.

Average Reads/Sec
: The average number of physical disk read operations per second made on all disks on the system.

Average Writes/Sec
: The average number of physical disk write operations per second made on all disks on the system.

Average K per I/O
: The average number of kilobytes (1024) transferred during each disk read or write operation.

Avg Util The average percentage of available time that disks were busy. It is a composite average for all disks on the system.

High Util The percentage of use for the disk arm that has the highest utilization.

High Util Unit
: The disk arm with the highest utilization.

High Srv Time
: 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.

Disk Space Used
: The total disk space used in millions of bytes for the entire system.

Disk Utilization Detail

The Disk Utilization Detail section of the Resource Interval Report, shown in Figure 8-76, gives detailed disk information for the selected time intervals. Information is shown for each disk arm that is configured on the system. Consistent disk arm utilization at or above the threshold value, provided in Figure 10-3 on page 10-6, will affect system performance and cause longer response times and/or less throughput.

```
                                             Resource Interval Report                                          04/19/90 13:58:59
                                             Disk Utilization Detail                                                   Page    2
                                            User-Selected Report Title

Member . . . : WASCOMTEST   Model/Serial  . . : B60/10-12883    Main storage  . . : 96.0 M      Started . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG     System name . . . :     TEST#366    Version/Release . :  2/1.0      Stopped . . . : 01/18/90 09:34:00

        Bus    IOP Id/     ASP   CSS    Itv    ---- Physical I/O per Second ---     K per             Queue    Avg Time per  I/O
Unit    ID     Model       ID    ID     End        Total      Reads      Writes      I/O     Util     Length     Service    Wait
----    ---    ---------   ---   ---   -----   ----------  ---------  ---------   -------   ----    -------    -------    ----
0001      0    01 (6110)   01    00    09:19         .288       .095       .192       4.6    1.0        .24       .032    .003
                                       09:34         .433       .311       .132       3.1    1.1        .01       .053    .006
                                               ----------  ---------  ---------   -------   ----    -------    -------    ----
Unit Average:                                        .366       .204       .162       3.7    1.0        .12       .042    .004

0002      0     01 (6110)  01    00    09:19         .117       .050       .066       6.8     .3        .12       .016    .005
                                       09:34         .094       .060      .034        5.1     .3        .24       .023    .004
                                               ----------  ---------  ---------   -------   ----    -------    -------    ----
Unit Average:                                        .105       .055       .050       6.0     .3        .18       .0019   .004
```

Figure 8-76. Resource Interval Report: Disk Utilization Detail Section

The columns in the Disk Utilization Detail section of the Resource Interval Report are as follows:

Unit The number assigned by the system to identify a specific disk unit or arm.

Bus ID The bus identification number.

IOP ID/Model
The input/output processor identification number and the model number in parentheses.

ASP ID The auxiliary storage pool number.

CSS ID The checksum set number that identifies which checksum set a disk is part of.

Itv End The interval end time (hour and minute).

Total Physical I/O per Second
The average number of physical disk I/O operations performed per second by the disk arm.

Reads per Second
The average number of disk read operations performed per second by the disk arm.

Writes per Second
The average number of disk write operations performed per second by the disk arm.

K per I/O The average number of kilobytes (1024) read or written for each disk I/O operation.

Util The percentage of time the disk arm was being used (busy).

Queue Length
The average number of I/O requests that had to wait in the queue for this unit.

Avg Time per Service
The amount of time a disk arm uses to process a given request.

I/O Wait 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.

Communications Line Detail

A Communications Line Detail section of a Resource Interval Report contains information about the line activity when performance data was collected for the specified member. One detail section is produced for each protocol in use on the lines that data was collected for. Figure 8-77 on page 8-106 through Figure 8-81 on page 8-109 are examples of the detail sections for the communications protocols.

```
                                Resource Interval Report                                          04/19/90 13:58:5
                                Communications Line Detail                                        Page    1
                                User-Selected Report Title

Member . . . : MONDAY      Model/Serial  . . :  B40/XX-XXXXX  Main storage . . :   16.0 M  Started  . . . . : 01/19/90 13:03:19
  Library . . : QPFRDATA    System name . . . :      SYS400    Version/Release  :   1/ 2.0  Stopped  . . . . : 01/19/90 14:57:50

 PROTOCOL = SDLC  (SORT BY INTERVAL)
                                                   Percent                          Percent    Pct     - Congestion --
         Bus/                     Bytes      Total   I Frames   Bytes     Total       Frames     Poll    Local    Remote
 Itv     IOP/      Line     Line Trnsmitd  IFrames  Trnsmitd   Recd      Frames     Received   Retry    Not      Not
 End     Line      Speed    Util Per Sec   Trnsmitd  in Error  per Sec    Recd       in Error   Time    Ready    Ready
 -----  ---------- -------- ---- -------  ---------  --------  -------  ---------   --------   -----   ------   ------
        Bus 0
        IOP 05
         (6110)

 13:08  MPLSCHI        9.6    48     384      2,645        21      384      4,086         08      03       18       00
 13:28  MPLSROCH       9.6    38     384      2,645        18      384      4,086         06      02       06       00
 13:48  SUPPLIER      19.2    22     384      2,645        04      384      4,086         02      18       07       02
```

Figure 8-77. Resource Interval Report: Communications Line Detail - SDLC

This is an example of the report section for communications lines using the synchronous data link control (SDLC) protocol. The data in this example is sorted by the data collection interval end times. The column descriptions follow:

Itv End The time (hour and minute) when the data collection interval ended.

Bus/IOP/Line
The bus ID, IOP ID and feature number, and the line description name.

Line Speed
The line speed in kilobits (1 kilobit = 1000 bits) per second.

Line Util The percent of available line capacity used by transmit and receive operations.

Bytes Trnsmitd per Sec
The average number of bytes transmitted per second.

Total I Frames Trnsmitd
The total number of information frames transmitted.

Percent I Frames Trnsmitd in Error
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).

Bytes Recd per Sec
The average number of bytes received per second.

Total Frames Recd
The number of frames received, including frames with errors and frames that are not valid.

Percent Frames Received in Error
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).

Pct Poll Retry Time
The percent of the time interval the line was unavailable while the IOP waited for a response from a work station controller (or remote AS/400 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).

Local Not Ready
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).

Remote Not Ready
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).

```
                                                 Resource Interval Report                                   04/19/90 13:58:59
                                                 Communication Line Detail                                  Page      1
                                                 User-Selected Report Title

Member . . . : MONDAY      Model/Serial  . . :  B40/XX-XXXXX  Main storage . . :   16.0 M  Started  . . . . : 01/19/90 13:03:19
  Library . . : QPFRDATA   System name . . . :      SYS400    Version/Release  :   1/ 2.0  Stopped  . . . . : 01/19/90 14:57:50

 PROTOCOL = X.25 (SORT BY INTERVAL)

                                                        Percent                           Percent
          Bus/                          Bytes     Total      I/Frames   Bytes     Total      Frames     ---- Reset ----
 Itv      IOP/       Line    Line  Trnsmitd  I/Frames  Trnsmitd    Recd     Frames     Received   --- Packets ---
 End      Line       Speed   Util   per Sec  Trnsmitd  in Error   per Sec    Recd      in Error   Trnsmitd   Recd
 -----  ----------  -------  ----  --------  --------  ---------  -------  ---------  --------   --------  -------
          Bus 0
          IOP 03
           (6150)

13:08   NETWORK       9600    78        384     2,645         12      384      4,086        21         222      222
13:13   NETWORK       9600    78        384     2,645         11      384      4,086        08         222      222
13:18   NETWORK       9600    78        384     2,645         09      384      4,086        00         222      222
```

Figure 8-78. Resource Interval Report: Communications Line Detail - X.25

This is an example of the report section for communications lines using the X.25 protocol. The column descriptions are the same as for Figure 8-77 on page 8-106, except for the following:

Reset Packets Trnsmitd
The number of reset packets transmitted by the network.

Reset Packets Recd
The number of reset packets received by the network. **Reset packets** are packets retransmitted because an error occurred.

```
                                                   Resource Interval Report                                         04/19/90 13:58:59
                                                   Communication Line Detail                                        Page    1
                                                   User-Selected Report Title

Member  . . . : MONDAY       Model/Serial  . . :  B40/XX-XXXXX  Main storage . . :   16.0 M   Started  . . . . : 01/19/90 13:03:19
  Library . . : QPFRDATA     System name . . . :     SYS400     Version/Release  :   1/ 2.0   Stopped  . . . . : 01/19/90 14:57:50

 PROTOCOL = TRLAN (SORT BY INTERVAL)

                                                              ------- Congestion -------
        Bus/                                I/Frames   I/Frames  -- Local ---   -- Remote --             Rsp       Remote LAN
 Itv    IOP/        Line       Line     Trnsmitd     Recd       Not    Seq    Not    Seq    Frame    Timer   --- Frames ---    MAC
 End    Line        Speed      Util     per Sec      per Sec    Ready  Error  Ready  Error  Retry    Ended   Trnsmitd Recd     Errors
 -----  ---------  --------   ----     ---------    ---------  -----  -----  -----  -----  ------   ------  -------- -----    ------
        Bus 1
        IOP 01
         (6240)

13:08   LOCALNET    4000.0     78          382          384     10      8     22      9       4       10        18     17       112
13:23   LOCALNET    4000.0     78          382          384     10      8     22      9       4       10        18     17       110
13:38   LOCALNET    4000.0     78          382          384     10      8     22      9       4       10        18     17        99
```

Figure 8-79. Resource Interval Report: Communications Line Detail - TRLAN

This is an example of the report section for communications lines using the token-ring local area network (TRLAN) protocol. The column descriptions are the same as for Figure 8-77 on page 8-106, except for the following:

I Frames Trnsmitd per Sec
The number of information frames transmitted per second.

I Frames Recd per Sec
The number of information frames received per second.

Local Seq Error
The percent of frames received out of order by the host system. This can occur when the host cannot process data fast enough.

Remote Seq Error
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.

Frame Retry
The number of attempts to retransmit a frame to a remote controller.

Rsp Timer Ended
The number of times the response timer ended waiting for a response from a remote device.

Remote LAN Frames Trnsmitd
The number of frames transmitted to a LAN connected to the locally attached LAN.

Remote LAN Frames Recd
The number of frames received from a LAN connected to the locally attached LAN.

MAC Errors
The number of medium access control (MAC) errors.

```
                                   Resource Interval Report                                    04/19/90 13:58:59
                                   Communication Line Detail                                           Page    1
                                   User-Selected Report Title

Member  . . . : MONDAY      Model/Serial  . . :  B40/XX-XXXXX   Main storage . . :    16.0 M   Started  . . . . : 01/19/90 13:03:19
  Library . . : QPFRDATA    System name . . . :     SYS400      Version/Release  :    1/ 2.0   Stopped  . . . . : 01/19/90 14:57:50

 PROTOCOL = ELAN (SORT BY INTERVAL)

                                                             ------- Congestion -------
          Bus/                           I/Frames    I/Frames  -- Local ---  -- Remote --              Rsp
 Itv      IOP/        Line      Line     Trnsmitd    Recd      Not    Seq    Not    Seq     Frame     Timer
 End      Line        Speed     Util     per Sec     per Sec   Ready  Error  Ready  Error   Retry     Ended
-----  ---------    --------    ----    ---------   ---------  -----  -----  -----  -----   ------    ------
       Bus 1
       IOP 01
        (6240)

13:08  LOCALNET2    4000.0      78          382         384     10       8     22      9        4        10
13:23  LOCALNET2    4000.0      78          382         384     10       8     22      9        4        10
13:38  LOCALNET2    4000.0      78          382         384     10       8     22      9        4        10
```

Figure 8-80. Resource Interval Report: Communications Line Detail - ELAN

This is an example of the report section for communications lines using the Ethernet local area network (ELAN) protocol. The column descriptions are the same as Figure 8-77 on page 8-106 and Figure 8-79 on page 8-108.

```
                                   Resource Interval Report                                    04/19/90 13:58:59
                                   Communication Line Detail                                           Page    1
                                   User-Selected Report Title

Member  . . . : MONDAY      Model/Serial  . . :  B40/XX-XXXXX   Main storage . . :    16.0 M   Started  . . . . : 01/19/90 13:03:19
  Library . . : QPFRDATA    System name . . . :     SYS400      Version/Release  :    1/ 2.0   Stopped  . . . . : 01/19/90 14:57:50

 PROTOCOL = ASYNC (SORT BY INTERVAL)

            Bus/                                   Bytes        Bytes       Total      Pct PDUs
 Itv        IOP/        Line       Line         Transmitted    Received     PDUs       Received
 End        Line        Speed      Util          per Sec       per Sec     Received    in Error
-----    ----------   --------     ----        -----------    --------    --------    --------
       Bus 1
       IOP 01
        (6240)

13:08    STARTSTOP          0.3     18               343          433         343          03
13:23    STARTSTOP          0.3     23               343          433         343          00
```

Figure 8-81. Resource Interval Report: Communications Line Detail - ASYNC

This is an example of the report section for communications lines using the asynchronous (ASYNC) protocol. The column descriptions are the same as for Figure 8-77 on page 8-106, except for the following:

Total PDUs Received
: The number of protocol data units (PDUs) received during the time interval.

Pct PDUs Received in Error
: 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.

Resource Interval Report
Communication Line Detail
User-Selected Report Title

04/19/90 13:58:59
Page 1

Member . . . : MONDAY   Model/Serial . . : B40/XX-XXXXX   Main storage . . : 16.0 M   Started . . . . : 01/19/90 13:03:19
Library . . : QPFRDATA   System name . . . : SYS400   Version/Release : 1/ 2.0   Stopped . . . . : 01/19/90 14:57:50

PROTOCOL = BSC (SORT BY INTERVAL)

Itv End	Bus/ IOP/ Line	Line Speed	Line Util	Bytes Transmitted per Sec	Total Data Characters Transmitted	Pct Data Characters Transmitted in Error	Bytes Received per Sec	Total Data Characters Received	Pct Data Characters Received in Error	Line Errors
	Bus 1 IOP 01 (6240)									
13:08	PNTTOPNT	4.8	30	444	12,345,444	09	104	34,211	01	383
13:23	PNTTOPNT	4.8	24	444	12,345,444	05	104	34,211	00	121

Figure 8-82. Resource Interval Report: Communications Line Detail - BSC

This is an example of the report section for communications lines using the binary synchronous communications (BSC) protocol. The columns descriptions are the same as Figure 8-77 on page 8-106, except for the following:

Total Data Characters Transmitted
The number of data characters transmitted successfully.

Pct Data Characters Transmitted in Error
The percent of data characters transmitted with error.

Total Data Characters Received
The number of data characters received successfully.

Pct Data Characters Received in Error
The percent of data characters received with error.

Line Errors
The total of all detected errors.

Note: Check the condition of the line if this value increases greatly over time.

Communications IOP Utilizations

The Communications IOP Utilizations section of the Resource Interval Report, shown in Figure 8-83, gives communications input/output processor (IOP) utilization.

Resource Interval Report
Communications IOP Utilizations
User-Selected Report Title

04/19/90 13:58:59
Page 7

Member . . . : WASCOMTEST   Model/Serial . . : B60/10-12883   Main storage . . : 96.0 M   Started . . . : 01/18/90 09:04:26
Library . . : QPFRJAG   System name . . . : TEST#366   Version/Release . : 2/1.0   Stopped . . . : 01/18/90 09:34:00

Bus ID	IOP ID/ Model	Itv End	Utilization	-- OPSTART Msg -- Reverse	Normal	--- Bytes Transmitted --- IOP	System	Restart Queues	BNA Received	Avail Local Storage (K)
0	05 (6110)	09:19	1.6	0	95	13,786	79,560	0	0	160
		09:34	2.0	0	201	17,488	247,280	0	0	160
1	02 (6110)	09:19	3.1	0	469	173,922	262,933	0	0	687
		09:34	.7	0	86	1,138	77,551	0	0	687

Figure 8-83. Resource Interval Report: Communications IOP Utilizations Section

The columns in the Communications IOP Utilizations section of the Resource Interval Report are as follows:

Bus ID The bus identification number.

IOP ID/Model
The input/output processor identification number and the model number in parentheses.

Itv End The interval end time (hour and minute).

Utilization The average percentage of available time each IOP was used (busy).

OPSTART Msg Reverse
The number of OPSTART bus unit messages received from another bus unit using the reverse flow method.

OPSTART Msg Normal
The number of OPSTART bus unit messages received from another bus unit using a normal flow method.

IOP Bytes Transmitted
The number of bytes transmitted from the IOP to the system across the bus.

System Bytes Transmitted
The number of bytes transmitted from the system to the IOP across the bus.

Restart Queues
The number of restart queues sent by the IOP.

BNA Received
The number of Buffer-not-available messages received by the IOP.

Avail Local Storage (K)
The number of kilobytes (1024) of free local storage in the IOP.

Disk IOP Utilizations

The Disk IOP Utilizations section of the Resource Interval Report, shown in Figure 8-84, gives input/output processor (IOP) utilization for direct access storage devices (DASDs). Consistent Disk IOP utilization at or above the threshold value, provided in Figure 10-3 on page 10-6, affects system performance and causes longer response times and/or less throughput.

```
                                        Resource Interval Report                                        04/19/90 13:58:59
                                         Disk IOP Utilizations                                                  Page    8
                                       User-Selected Report Title

Member  . . . : WASCOMTEST   Model/Serial  . . : B60/10-12883    Main storage  . . : 96.0 M      Started . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG      System name . . . :     TEST#366    Version/Release . :  2/1.0      Stopped . . . : 01/18/90 09:34:00

Bus      IOP ID/       Nbr       Itv              ------ Disk I/O /Sec ------      --- KB per I/O ----
ID       Model         Arms      End      Util       Reads          Writes         Read        Write
---      ----------    ----     -----     ----    -----------    -----------      -------     -------
  0       01 (6110)       7     09:19       .7          .595           .685          6.9         4.9
                                09:34       .6          .635           .326          4.3         4.4

  0       09 (6110)      16     09:19      2.2         1.878          1.637          9.5         5.8
                                09:34      1.3          .915           .845         10.5         5.7
```

Figure 8-84. Resource Interval Report: Disk IOP Utilizations Section

The columns in the Disk IOP Utilizations section of the Resource Interval Report are as follows:

Bus ID The bus identification number.

IOP ID/Model
The input/output processor identification number and model number in parentheses.

Nbr Arms The number of arms on the disk IOP.

Itv End The interval end time (hour and minute).

Util The percentage of utilization for the disk IOP.

Disk I/O Reads /Sec
The average number of disk read operations per second by the disk IOP.

Disk I/O Writes /Sec
The average number of disk write operations per second by the disk IOP.

KB per I/O Read
The average number of kilobytes (1 KB equals 1024 bytes) transferred per read operation.

KB per I/O Write
The average number of kilobytes transferred per write operation.

Multifunction IOP Utilizations

The Multifunction IOP Utilizations section of the Resource Interval Report, shown in Figure 8-85, gives input/output processor (IOP) utilization for both DASD and communications devices. Consistent utilization at or above the threshold value, shown in Figure 10-3 on page 10-6, affects system performance and causes longer response times and/or less throughput.

```
                                   Resource Interval Report                               04/19/90 13:58:59
                                Multifunction IOP Utilizations                                      Page    9
                                  User-Selected Report Title

Member . . . : WASCOMTEST  Model/Serial  . . : B60/10-12883   Main storage  . . : 96.0 M    Started . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG    System name . . . :    TEST#366    Version/Release . :  2/1.0    Stopped . . . : 01/18/90 09:34:00

Bus     IOP ID/          Itv
ID      Model            End        Utilization
---     -----------      -----      -----------
 0       05 (6110)       09:19           1.6
                         09:34           2.0

 1       02 (6110)       09:19           3.1
                         09:34            .7
```

Figure 8-85. Resource Interval Report: Multifunction IOP Utilizations Section

The columns in the Multifunction IOP Utilizations section of the Resource Interval Report are as follows:

Bus ID The bus identification number.

IOP ID/Model
The IOP identification number and the model number in parentheses.

Itv End The interval end time (hour and minute).

Utilization The percent of utilization for each Multifunction IOP.

Local Work Station IOP Utilizations

The Local Work Station IOP Utilizations section of the Resource Interval Report, shown in Figure 8-86 on page 8-113, gives input output processor (IOP) utilization for local work stations.

```
                                        Resource Interval Report                                    04/19/90 13:58:59
                                     Local Work Station IOP Utilization                                    Page   10
                                        User-Selected Report Title

Member  . . . : WASCOMTEST   Model/Serial  . . : B60/10-12883    Main storage  . . : 96.0 M      Started . . . : 01/18/90 09:04:26
  Library . . : QPFRJAG      System name . . . :     TEST#366    Version/Release . :  2/1.0      Stopped . . . : 01/18/90 09:34:00

Bus      IOP ID/        Work Station     Itv               Active
ID       Model          Controller       End      Util     Wrk Stn     0-001      001-002     002-004     004-008      >008
---      ----------     -----------      -----    ----     -------    -------     -------     -------     -------     -------

 0        03 (6110)     CTL01            09:19    11.0           2         50           3           2                       2
                                         09:34    11.7           3         63           7                                   1

 0        04 (6110)     CTL02            09:19    21.6           7        325          35          13           6            5
                                         09:34    20.3           7        165           9           4           3
                                                                          -------     -------     -------     -------     -------
Total Responses:                                                          603          54          19           9           8
```

Figure 8-86. Resource Interval Report: Local Work Station IOP Utilizations Section

The columns in the Local Work Station IOP Utilizations section of the Resource Interval Report are as follows:

Bus ID The bus identification number.

IOP ID/Model
The input/output processor identification number and the model number in parentheses.

Work Station Controller
The name of the local work station controller.

Itv End The interval end time (hour and minute).

Util The percent of utilization for each local work station IOP.

Active Wrk Stn
The number of work stations with activity.

0-001 The number of times the response time was between zero and 1 second.

001-002 The number of times the response time was between 1 and 2 seconds.

002-004 The number of times the response time was between 2 and 4 seconds.

004-008 The number of times the response time was between 4 and 8 seconds.

>008 The number of times the response time was greater than 8 seconds.

Analyzing Seize/Lock Conflicts

Seizes/locks are system-locking functions that ensure integrity during certain operations. For example, the system uses a seize during logical file maintenance when the underlying physical files are changed.

Conflicts occur when one job has an object lock or seize and another job requests control of the same object. A common example of a lock conflict is when a job reads a record for update and a second job requests a lock for the same record.

If the Print Transaction Report (PRTTNSRPT) job summary output shows a high number for either the number of lock or seize conflicts, look at the Transaction Detail Report and Transition Detail Report to further analyze the situation. You can also use the PRTLCKRPT command to print the Seize/Lock Conflict Report to see what conflicts occurred.

If the PRTTNSRPT command output shows several lock waits, or system throughput is low and the processing unit time and disk use is also low, these conditions could be caused by lock-wait conflicts occurring in jobs due to contention for files, records, or other objects. Analyze the resource management trace data using the PRTLCKRPT command to determine a cause.

You can normally expect to see some conflicts occur for a short period of time on some objects. If you see several lock conflicts occur for nondatabase objects, it may be a normal situation (such as writers and jobs contending for output queues). However, if the locks last a long time (more than 5 to 10 seconds), and they cause objectionable delays to end users, this situation could indicate that you need to make some changes to the operational environment.

If the report shows several database record locks that last for more than 5 to 10 seconds, a program may have read a record for update and continued processing without releasing (writing) the record. This situation is normal in many applications. However, in a heavily loaded system, the job that holds the record lock may reach the end of its time slice while it holds the lock. When this condition occurs, it delays other jobs that need the record.

If the report shows several seizes that last for a period of time (over 1 second), this condition can indicate object contention problems. To ensure the accuracy of the object, the system does not allow access to the object until all the necessary changes are made.

Print Lock Report (PRTLCKRPT) Command

Use the PRTLCKRPT command to get 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 seizes. 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.

To create the trace data for this command, start the performance monitor (use the STRPFRMON command) and specify tracing. After you run the test, stop the monitor. When the monitor job completes, use the PRTLCKRPT command to print the report. Use the same value for both the MBR and LIB parameters.

For information on how to enter the PRTLCKRPT command, see the *CL Reference* manual.

When you use the PRTLCKRPT command, the following file is used as input:

File	Description
QAPMDMPT	Database file that is output from the STRPFRMON command and updated by the PRTTNSRPT command.

Note: Because the PRTLCKRPT command uses trace output from the STRPFRMON command, the STRPFRMON command must be run first.

Following are the output files from the PRTLCKRPT command:

File	Description
QPPTLCK	Printer file
QAPTLCKD	Database file

Note: In the following description, the term *lock* means lock or seize unless otherwise noted.

The PRTLCKRPT command produces a report with several printed copies. An optional detail list of the resource management trace records from QAPMDMPT prints first. This list may be sorted by the times that a lock occurred, the name of the job requesting the lock, the name of the job holding the lock, or the name of the locked object. The list may print four times (once for each of these sequences).

Next, these summaries print detail listings summarized by:

- Requesting job
- Holding job
- Object name

Figure 8-87 shows an example of the detail listing, sorted by time of day (in this case). The report options were selected to include only locks lasting at least two seconds that occurred between 13:33:00 and 13:34:00 (as noted in the footer printed at the bottom of the summary page).

```
11/18/88  13:45:40                               Seize/Lock Wait Statistics by Time of Day                              Page   1
 TOD of   Length                                                                 Object                                   Record
  Wait    of Wait L Requestor's Job Name         Holder's Job Name                Type    Object Name                     Number
--------  ------- - ---------------------------- ---------------------------- ------ ------------------------------- ----------
13.33.15     2033 L SCPF       QSYS       000000 WKS12502   ROBERT     159708 MSGQ   QHST       QSYS
13.33.25     3372   WKS04001   JACKIE     159505 WKS20601   NANCY      159599 MBR    INVMAST    PROD1      INVMAST
13.33.28     2344 L WKS13300   MASTER     159706 SCPF       QSYS       000000 MSGQ   QHST       QSYS
13.33.36    10871   QSYSARB    QSYS       159052 WKS12202   JIMST      159627 CUD    C00B99
13.33.36     2187   WKS03000   TERRI      159701 WKS12003   KAREN      159524 MBR    ORDER01    PROD1      ORDER01
13.33.40     2400   AUT0001    MASTER     159546 WKS01009   WENDY      159387 MBR    CUSTMAST   PROD1      CUSTMAST
13.33.52     3881   WKS04001   JACKIE     159505 WKS12102   JIM        159392 MBR    INVMAST    PROD1      INVMAST
13.33.52     4386   WKS05001   HARRY      159196 WKS12020   JIM        159392 MBR    INVMAST    PROD1      INVMAST
13.33.00     2289   WKS05001   HARRY      159196 WKS04000   NANCY      159527 MBR    INVMAST    PROD1      INVMAST
Member SORTED         Library QPFRDATA       Period from 13.33.00 through 13.34.00       2,000 ms minimum wait
```

Figure 8-87. Example of a Detail Listing

The columns in the detail listing are as follows:

TOD of Wait
The time of day of the start of the conflict.

Length of Wait
The number of milliseconds the requestor waited for the locked object.

L
Whether this is a lock or seize conflict. The column contains an L if lock, blank if seize.

Requestor's Job Name
The name of the job requesting the locked object.

Holder's Job Name
The name of the job holding the lock.

Object Type
The type of the locked object.

Object Name
The name of the locked object.

Record Number
For database file members, the relative record number of the record within the database file member.

Figure 8-88 shows an example of the Requesting Job Summary section of the same report. The other summary sections have a similar format.

```
11/18/87 13:45:40                                    Seize/Lock Wait Statistics Summary                                      Page    5
                                                              Locks                         Seizes
Requestor's Job Name                                    Count        Avg Length        Count        Avg Length
-----------------------------                       ------------- -------------     ------------- -------------
AUTO001     MASTER      159546                                                             1           2,400
WKS13300    MASTER      159706                            1           2,344
QSYSARB     QSYS        159052                                                             1          10,871
WKS03000    TERRI       159701                                                             1           2,187
WKS04001    JACKIE      159505                                                             2           3,627
WKS05001    HARRY       159196                                                             2           3,338
SCPF        QSYS        000000                            1           2,033
Member SORTED           Library QPFRDATA            Period from 13.33.00 through 13.34.00      2,000 ms minimum wait
```

Figure 8-88. Example of Summary by Requesting Job

The columns in the Requesting Job Summary section of the report are as follows:

Requestor's Job Name
The name of the job requesting the locked object (the same as in the detail listing). For the holding job summary report, this column is *Holder's Job Name*. For the object summary report, the two columns are *Object Type* and *Object Name*.

Two columns are repeated, once for locks and once for seizes:

Count The number of locks or seizes that occurred.

Avg Length
The average number of milliseconds a lock or seize was held.

Consider the following points when you use the PRTLCKRPT command:

1. The PRTTNSRPT output may show a high incidence of wait-to-ineligible state transitions in the transaction summary output. If this situation occurs, it could mean that many jobs are waiting for internal system object locks and holding an activity level while waiting. The PRTLCKRPT report may identify these locks.
2. The Detailed Lock Conflicts Report (shown in Figure 8-87 on page 8-115) shows each object lock conflict that meets the specified selection values. Do not assume that each conflict shown for an object lock is associated with a separate request for the object from the program that originally requested it.

   When multiple requests (from multiple jobs) cause contention for an object, the requests are processed in the order received, by job priority. When conflicts occur, multiple lock requests are made by internal programs in behalf of the program that originally made the request, until the lock is granted. These internal requests appear on the summary, resulting in more conflicts than actually occurred from the originating program's viewpoint.

   PRTLCKRPT processing does not analyze the internal lock conflicts and relate them to the original request.

Batch Job Trace Report

To produce the Batch Job Trace Report, use the Print Trace Report (PRTTRCRPT) command. Prior to printing the Batch Job Trace Report, you must use the Start Performance Monitor (STRPFRMON) command with the JOBTRCITV and JOBTYPE options and the Print Transaction Report (PRTTNSRPT) command with the *FILE option. The PRTTNSRPT command creates the QTRJOBT file that the Batch Job Trace Report uses.

The Batch Job Trace Report shows the progression of different job types (for example, batch jobs) traced through time. Resources utilized, exceptions, and state transitions are reported.

Each page of the Batch Job Trace Report shows the header information in Figure 8-89.

```
                                          Batch Job Trace Report                                          06/13/90 14:01:46
                                               Job Summary                                                    Page     1
                                       Sample Batch Job Trace Report

Member . . . :  GOODSTUF  Model/Serial  . . : B60/10-15018  Main storage . . :    96.0 M  Started . . . . :   6/07/90 08:09:22
  Library . . :  QPFRDATA  System name . . . :     RCH38366  Version/Release  :   2/ 1.0  Stopped . . . . :   6/07/90 08:41:55
```

Figure 8-89. Batch Job Trace Report: Header Information

The header information consists of:

- Report title
- Current date and time
- Report page number
- User-selected report title
- Data member name
- Library name
- Model number
- Serial number
- System name
- Main storage size
- OS/400 version and release level
- Data collection start date and time
- Data collection stop date and time

Job Summary

The Job Summary section of the Batch Job Trace Report, shown in Figure 8-90 on page 8-118, gives the number of traces, the number of I/O operations, the number of seize and lock conflicts, the number of EAO exceptions, and the number of state transitions for each batch job.

```
                                      Batch Job Trace Report                                      06/13/90 14:01:46
                                           Job Summary                                             Page    1
                                   Sample Batch Job Trace Report

Member . . . :    GOODSTUF  Model/Serial  . . :  B60/10-15018  Main storage . . :    96.0 M  Started . . . . :    6/07/90 08:09:22
  Library . . :   QPFRDATA  System name . . . :      RCH38366  Version/Release  :    2/ 1.0  Stopped . . . . :    6/07/90 08:41:55

                                                                           --- Physical --     Seize      EAO    --- State ---
   Job          User        Job             -- Job --   Number    CPU     -- I/O Count --    and Lock    Excp   -Transitions-
   Name         Name       Number   Pool    Type  Pty   Traces    Util     Sync     Async    Conflicts   /1000    A-A    A-I
----------   ----------    ------   ----    ----  ---   ------    ----    ------   ------    ---------   -----   -----  -----
C12TS1       CLASS23       005276    02      B     50      395    10.6     1,155      350           95     1.5      27     11
C13T         QSECOFR       005277    02      B     50      428    16.6     1,137      523          123      .5       0      0
C13TS1       CLASS33       005236    02      B     40        1      .0       103       23        1,446      .0     112     89
C14T         QSECOFR       005275    03      B      0       16      .4       518      351           87      .1       5     37
C14TS1       CLASS43       005230    02      B     40        1      .0        49        8          267      .0      33      6
```

Figure 8-90. Job Summary

The columns in the Job Summary section are as follows:

Job Name Name of the job.

User Name Name of the user.

Job Number Number of the job.

Pool Pool in which the job ran.

Job Type Job type and subtype.

Possible job type values include the following:

- **A** Autostart
- **B** Batch
- **C** PC Support server
- **D** Distributed data management (DDM) server
- **M** Subsystem monitor
- **R** Spool reader
- **S** System
- **W** Spool writer
- **X** Start system job

Possible job subtype values include the following:

- **M** Multiple requester terminal (MRT) (System/36 environment only)
- **E** Evoke (communications batch)
- **P** Print driver job
- **J** Prestart job

Job Pty Priority of the job.

Number Traces
Number of traces.

CPU Util Percentage of available processing unit time used by the job. For multiple-processor systems, this is the average use across all processors.

Physical I/O Count
Number of synchronous and asynchronous disk operations (reads and writes).

Seize and Lock Conflicts
Number of seize conflicts and lock waits.

EAO Excp / 1000
Number of effective address overflow exceptions divided by 1000.

State Transitions A-A
Number of active-to-active transitions.

State Transitions A-I
Number of active-to-ineligible transitions.

Chapter 9. Performance Graphics

This chapter describes the functions that allow you to work with performance data in a graphical format. The performance data is collected using the Start Performance Monitor (STRPFRMON) command. The graphs can be displayed interactively, printed, plotted, or saved to a graphics data format (GDF) file for use by other utilities, such as the Business Graphics Utility (BGU).

Note: This chapter does not refer to Capacity Planning graphics. For information about using the capacity planner graphics feature, see Chapter 10, "Capacity Planning and Performance Prediction."

Summary

Two distinct types of graphs can be displayed: performance graphs and historical graphs. Performance graphs use the performance data collected from a single run of the performance monitor. Performance graphs are useful for singling out jobs that are performing poorly or evaluating the activities performed by a user or class of users on the system during a specified period.

Historical graphs use performance data collected from several runs of the performance monitor. Historical data is the summary of the performance data created by the STRPFRMON command. The Create Historical Data (CRTHSTDTA) command is used to summarize the performance data for use by the historical graphs. Historical graphs are used to show how the performance of a system has changed over time.

Use the following steps to display a performance graph:

1. Create a graph format using the Create Graph Format (CRTGPHFMT) command. (Graph formats are reusable.)
2. Collect performance data using the STRPFRMON command.
3. Display the graph using the Display Performance Graph (DSPPFRGPH) command.

Use the following steps to display a historical graph:

1. Create a graph format using the CRTGPHFMT command. (Graph formats are reusable.)
2. Collect performance data using the STRPFRMON command.
3. Create the historical data using the CRTHSTDTA command.
4. Display the graph using the Display Historical Graph (DSPHSTGPH) command.

When you select option 9 (Performance graphics) on the IBM Performance Tools/400 menu, the Performance Tools Graphics menu appears.

```
PERFORMG                     Performance Tools Graphics
                                                            System:    RCH38366
Select one of the following:

     1. Work with graph formats and packages
     2. Work with historical data
     3. Display graphs and packages

    70. Related commands

Selection or command
===> ________________________________________________________________

F3=Exit   F4=Prompt   F9=Retrieve   F12=Cancel   F13=User support
F16=System main menu
(C) COPYRIGHT IBM CORP. 1981, 1990.
```

You can also reach this menu by typing the Start Performance Graphics (STRPFRG) command on the command line of any display. From this menu, you can work with performance data in a graphical format.

Work with Graph Formats and Packages

Graph formats are templates or outlines used by the DSPPFRGPH and the DSPHSTGPH commands to display graphs in a user-defined format. Figure 9-1 shows the 15 predefined graph formats included in QPFRDATA, the IBM-supplied performance library.

Figure 9-1 (Page 1 of 2). QPFRDATA Graph Formats

Graph Format Name	Description
QIBMASYNC	Asynchronous disk I/O per second against time
QIBMCMNIOP	Communications IOP use against time
QIBMCMNLIN	Maximum communciations line use against time
QIBMCPUPTY	Processor unit use of jobs with priorities 0-19, 20-39, 40-59, 60-79, and 80-99 against time
QIBMCPUTYP	Processor unit use of batch, interactive, and system jobs against time
QIBMDSKARM	Disk arm use against time
QIBMDSKIOP	Disk IOP use against time
QIBMLWSIOP	Local work station IOP use against time
QIBMMFCIOP	Multifunction communications IOP use against time
QIBMMFDIOP	Multifunction disk IOP use against time
QIBMPCTDSK	Percentage of disk occupied against time
QIBMRSP	Interactive response time against time
QIBMTOTDSK	Total disk I/O per second against time
QIBMTNS	Transactions per hour against time

Figure 9-1 (Page 2 of 2). QPFRDATA Graph Formats

Graph Format Name	**Description**
QIBMSYNC	Synchronous disk I/O per second against time

Graph packages allow you to group several graph formats into a single entity. This is useful for printing, displaying, or plotting a number of graphs at once. Instead of having to issue several DSPPFRGPH or DSPHSTGPH commands to print several graphs, you can use the package name (one command) to print all of the graphs in the graph package. Also included in QPFRDATA is a predefined graph package, QIBMPKG, which contains the 15 IBM standard graph formats.

If you select option 1 (Work with graph formats and packages) on the Performance Tools Graphics menu, the Work with Graph Formats and Packages display appears.

```
                       Work with Graph Formats and Packages

Library . . . . . . . . . . .    QPFRDATA

Type options, press Enter.
  1=Create graph format    2=Change          3=Copy   4=Delete
  5=Display sample graph   6=Create package  8=Display package contents

        Format/
Option  Package      Type       Text
 _      __________              ________________________________
 _      PACKAGE1     PACKAGE    Graph Package containing format w/ func. areas
 _      QIBMPKG      PACKAGE    IBM Graph Package
 _      FORMAT1      FORMAT     CPU Utilization vs. Time-Functional Areas
 _      NWCTEST      FORMAT
 _      QIBMASYNC    FORMAT     Asynchronous Disk I/O per Second vs. Time
 _      QIBMCMNIOP   FORMAT     Communications IOP Utilization vs. Time
 _      QIBMCPUPTY   FORMAT     CPU Utilization vs. Time (Priority)
 _      QIBMCPUTYP   FORMAT     CPU Utilization vs. Time (Job Type)
 _      QIBMDSKARM   FORMAT     Disk Arm Utilization vs. Time
                                                                  More...

F3=Exit   F5=Refresh   F12=Cancel   F15=Sort by format   F16=Sort by text
```

This display shows you the graph formats and graph packages that exist in the library specified in the *Library* field. The graph format or graph package name, a format or package indicator, and a text description appear on the display. If you cannot find the format or package you want to work with, use the appropriate function key to sort the formats and packages. You can sort them by name, type, or text description. When you find the graph format or package you want to work with, select the function you want to perform by typing the appropriate option in the *Option* field and pressing the Enter key.

If you are searching for a graph format or graph package located in a library that is different from the one currently listed in the *Library* field at the top of the display, type a new library name in the field and press the Enter key. A list of the graph formats and graph packages available in the library you specified appears. You can then select one of them to work with.

Create Graph Format

To create a new graph format, type a 1 (Create graph format), the graph format name, and the description on the first line under the *Option*, *Format/Package*, and *Description* columns, and press the Enter key. The CRTGPHFMT command prompt appears.

Specify how your graphs are displayed by selecting from the following options:

- Titles
- X-axis data
- Y-axis data
- Data type
- Individual line breakdown
- Graph type

Graph Types

The graph types available are:

- Line
- Scatter plot
- Surface
- Floating bar
- Composite bar

Line Graphs

Use line graphs to show change occurring over time. Line graphs can represent increases, decreases, trends, and general fluctuations of quality.

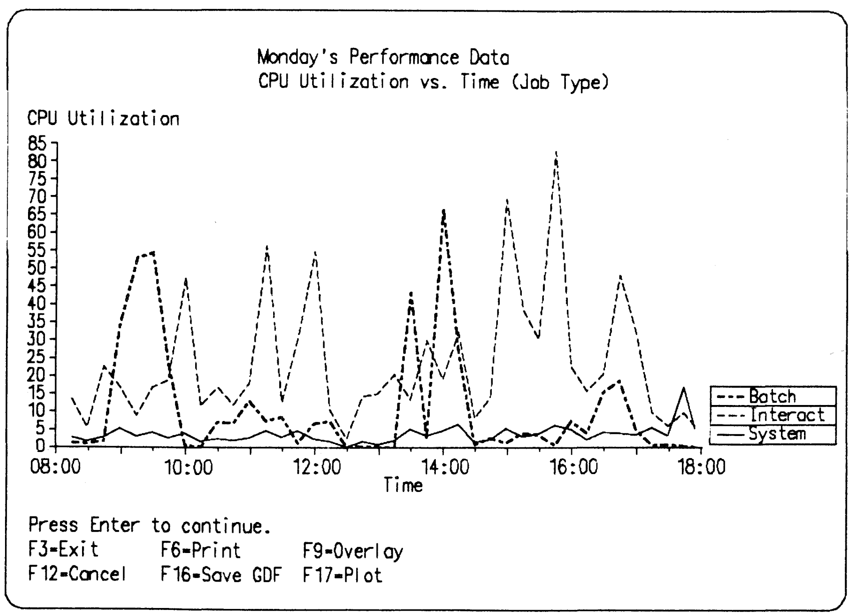

RV2S067-0

Figure 9-2. Line Graph: Data Represented as Lines

Each plotted point is shown by a marker; the plotted points are connected to form a continuous line. Each line is assigned a different color. If lines overlap, the color of the last legend entry at that point is displayed.

Scatter Plots

Scatter plots are similar to line graphs, except that the lines that connect the data points are not drawn.

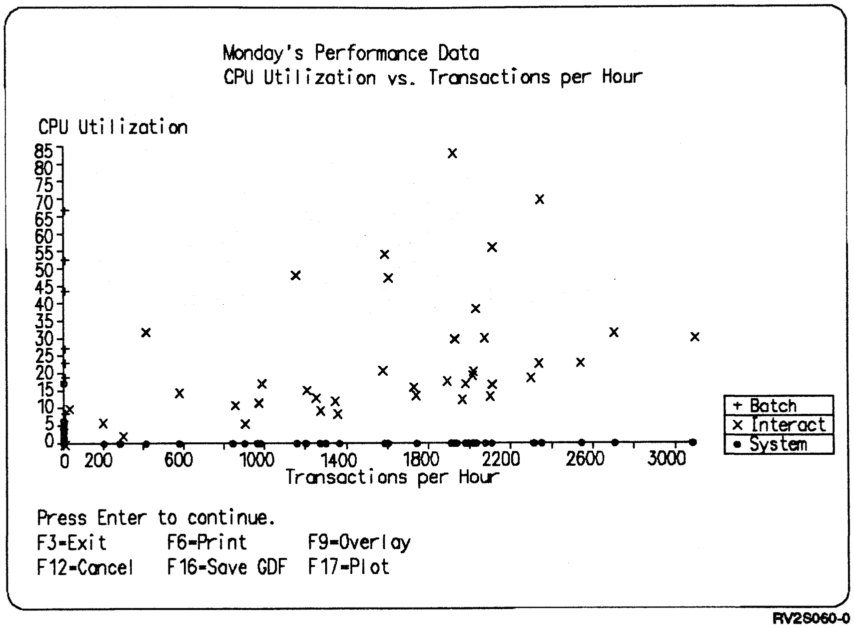

Figure 9-3. Scatter Plot: Data Represented as Markers

Surface Graphs

Use surface graphs like line graphs to show changes occurring over time. Surface graphs emphasize volume by shading the area between the lines and the X-axis if you specify Y (yes) for the area fill option.

Note: If you do not use the area fill option in your surface graph, your graph will be a cumulative line graph. If there is a legend entry with a value of zero to plot, its line covers the line plotted previously because there is no change to plot. Although shading requires more time to display or plot than simply drawing the lines, the area fill option may show more clearly which legend entries represent the different areas, particularly in cases where a line of one color may cover another.

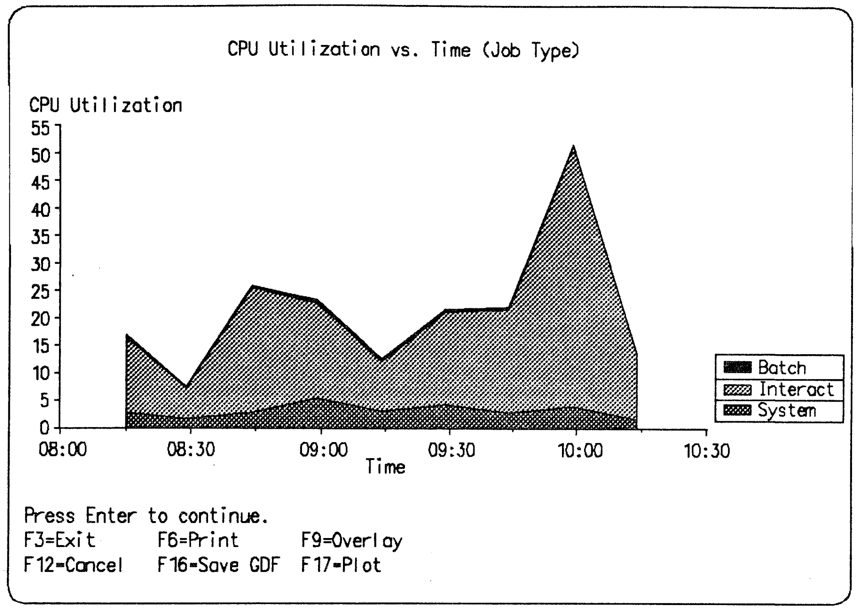

Figure 9-4. Surface Graph: Data Represented as Shaded Regions

Bar Graphs

Use bar graphs to show: changes occurring over time, parts of an entity, relationships between variables, and comparisons.

Use **composite-bar** graphs to show how parts comprise the entity, and how the entity relates to other entities.

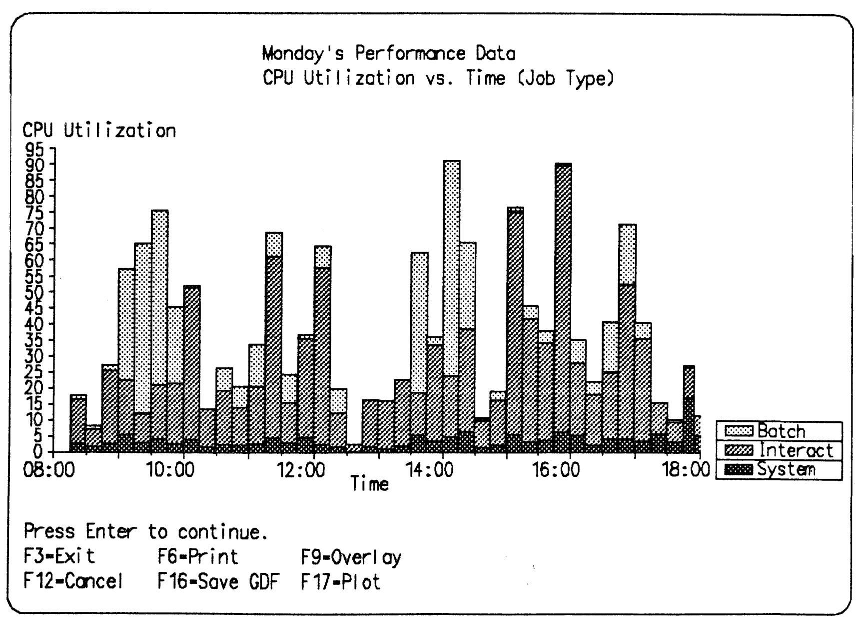

Figure 9-5. Composite-Bar Graph

Floating-bar graphs are similar to composite-bar graphs, except that the first component is not shown. Use floating-bar graphs to show the lower limits of each entity, in addition to the relationship of the elements that comprise the entity.

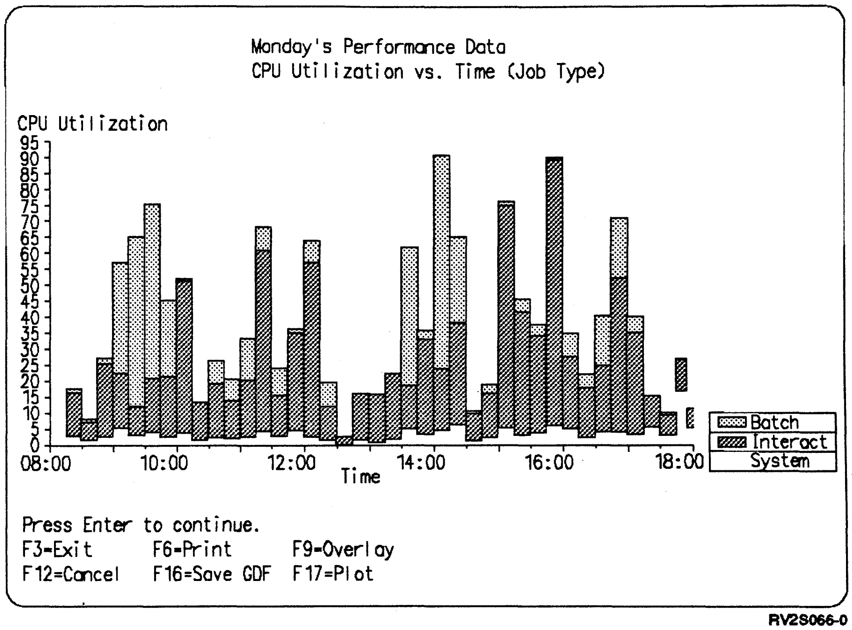

Figure 9-6. Floating-Bar Graph

Data Types

Data types control the number of lines displayed in your graph. They are a means of categorizing the information provided in your graph. For example, if you want the graph CPU Over Time and want a separate line plotted for every priority data type, you would specify *PRIORITY as your data type. You would then be presented with a display that would allow you to enter 1 to 16 priority ranges for plotting in this particular graph. Data types, therefore, control the legend entries in your graph.

Data types available for graphing are:

All Jobs	*ALL (default)
Job Type	*JOBTYPE
Priority	*PRIORITY
Functional Area	*FCNARA
IOP (input/output processor)	*IOP
Disk	*DISK
Communications Lines	*CMNLINE

The data types can be described as follows:

All Jobs (*ALL)
*ALL puts all of the jobs into one group for graphing. *ALL is the default value.

Job Type (*JOBTYPE)
*JOBTYPE includes individual job types, such as interactive, and conglomerate types, such as *ALLINTER (all interactive), *ALLBATCH (all batch), and *ALLSYSTEM (all system). All interactive refers to a job

with a job type of 'I', and includes interactive, PC Support, System/36, MRT, and pass-through jobs.

Priority (*PRIORITY)
*PRIORITY puts jobs into priority ranges. For example, the range 10-20 includes all jobs that have priorities between 10 and 20, inclusive.

Functional Area (*FCNARA)
*FCNARA puts jobs into each of the functional areas that are to be graphed. Functional areas must be unique over the data that is graphed. That is, if a job exists in more than one of the functional areas selected for the graph, an error message is issued indicating that the job exists in more than one functional area. Also, you cannot use functional areas to graph historical data. Refer to Chapter 12, "Managing the Performance Tools Configuration" on page 12-1 for a detailed description of functional areas.

IOP (*IOP)
*IOP allows users to graph maximum and average utilization lines for the particular type of input/output processor.

Disk (*DISK)
*DISK allows users to graph maximum and average utilization lines for the disk arms. It also allows maximum and average lines for the percentage of disk occupied.

Communications Lines (*CMNLINE)
*CMNLINE allows users to graph individual communications line use or the maximum use of all communications lines.

Valid Data Types for Axis Selections

Figure 9-7 shows the possible combinations for X-axis and Y-axis values based on the data type being graphed. For example, if you want to graph Time against Disk IOP Utilization, specify a data type of *IOP.

Figure 9-7 (Page 1 of 2). Valid X-axis and Y-axis Values

	X-Axis										
Y-Axis	**Time**	**CPU Utilization**	**Trans per Hour**	**Total Nbr of Trans**	**Resp Time**	**Sync Disk I/O per Sec**	**Total Sync I/O**	**Async Disk I/O per Sec**	**Total Async I/O**	**Total Disk I/O per Sec**	**Total Disk I/O**
CPU Utilization	X[2]	—	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]
Transactions per Hour	X[2]	X[1]	—	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]
Total Number of Transactions	X[2]	X[1]	X[1]	—	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]
Response Time	X[2]	X[1]	X[1]	X[1]	—	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]
Sync Disk I/O per Second	X[2]	X[1]	X[1]	X[1]	X[1]	—	X[1]	X[1]	X[1]	X[1]	X[1]
Total Sync Disk I/O	X[2]	X[1]	X[1]	X[1]	X[1]	X[1]	—	X[1]	X[1]	X[1]	X[1]

Figure 9-7 (Page 2 of 2). Valid X-axis and Y-axis Values

Y-Axis	X-Axis										
	Time	**CPU Utilization**	**Trans per Hour**	**Total Nbr of Trans**	**Resp Time**	**Sync Disk I/O per Sec**	**Total Sync I/O**	**Async Disk I/O per Sec**	**Total Async I/O**	**Total Disk I/O per Sec**	**Total Disk I/O**
Async Disk I/O per Second	X[2]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	—	X[1]	X[1]	X[1]
Total Async Disk I/O	X[2]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	—	X[1]	X[1]
Total Disk I/O per Second	X[2]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	—	X[1]
Total Disk I/O	X[2]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	X[1]	—
Communications IOP Utilization	X[3]	—	—	—	—	—	—	—	—	—	—
Disk IOP Utilization	X[3]	—	—	—	—	—	—	—	—	—	—
Local Work Station IOP Utilization	X[3]	—	—	—	—	—	—	—	—	—	—
Multifunction Communications IOP Utilization	X[3]	—	—	—	—	—	—	—	—	—	—
Multifunction Disk IOP Utilization	X[3]	—	—	—	—	—	—	—	—	—	—
Disk Arm Utilization	X[4]	—	—	—	—	—	—	—	—	—	—
Disk Percent Occupied	X[4]	—	—	—	—	—	—	—	—	—	—
Communications Line Utilization	X[5]	—	—	—	—	—	—	—	—	—	—

Key:

1. A graph type of *SCATTER and data type of *ALL, *FCNARA, *JOBTYPE, or *PRIORITY are required.
2. A data type of *ALL, *FCNARA, *JOBTYPE, or *PRIORITY is required.
3. A data type of *IOP is required.
4. A data type of *DISK is required.
5. A data type of *CMNLINE is required.

Legends

The legends displayed in the graph are controlled by the data type specified (for example, *JOBTYPE). The maximum number of legend entries you can specify for each data type is as follows:

Data Type	Maximum Legend Entries
All	1
Job Type	16
Priority	16
Functional Area	16
IOP	2
Disk	2
Communications Line	16

Create Graph Package

To create a new graph package, type a 6 (Create graph package), the graph package name, and the text description on the first line under the *Option*, *Format/Package*, and *Text* columns, and press the Enter key. The Create Graph Package display appears.

```
                              Create Graph Package

 Graph package . . :   PACKAGE2
   Library . . . . :     QPFRDATA

 Type options, press Enter.
   1=Select   5=Display sample graph

 Option     Format         Text
   _        FORMAT1        CPU Utilization vs. Time-Functional Areas
   _        FORMAT2        Response Time vs. Time-Functional Areas
   _        QIBMASYNC      Asynchronous Disk I/O per Second vs. Time
   _        QIBMCMNIOP     Communications IOP Utilization vs. Time
   _        QIBMCPUPTY     CPU Utilization vs. Time (Priority)
   _        QIBMCPUTYP     CPU Utilization vs. Time (Job Type)
   _        QIBMDSKARM     Disk Arm Utilization vs. Time
   _        QIBMDSKIOP     Disk IOP Utilization vs. Time
   _        QIBMDSKOCC     Percentage of Disk Occupied vs. Time
   _        QIBMLWSIOP     Local Workstation IOP Utilization vs. Time
   _        QIBMMFCIOP     Multifunction IOP (Comm) Util vs. Time
                                                                  More...

 F3=Exit   F5=Refresh   F12=Cancel   F16=Sort by text
```

On this display, type a 1 (Select) by any graph formats that you want to include in the graph package. If you are unsure about including a graph format in the package, type a 5 (Display sample graph) by the format in question. This displays a sample graph using the format selected. When you have made all of your selections and there are only 1's in the *Option* column, press the Enter key to create the graph package.

Change Graph Formats and Packages

To change an existing graph format or graph package, type a 2 (Change) next to the format or package name on the Work with Graph Formats and Packages display, and press the Enter key. If you are changing a graph format, the Change Graph Format (CHGGPHFMT) command prompt appears. Make your changes and press the Enter key. If you are changing a graph package, the Change Graph Package display appears.

```
                         Change Graph Package

Graph package . . :   PACKAGE1
  Library . . . . :     QPFRDATA
  Text  . . . . . :   Text for package 1

Type options, press Enter.
  1=Select   5=Display sample graph

Option     Format         Text
  1        FORMAT1        CPU Utilization vs. Time-Functional Areas
  1        FORMAT2        Response Time vs. Time-Functional Areas
  1        QIBMASYNC      Asynchronous Disk I/O per Second vs. Time
  _        QIBMCMNIOP     Communications IOP Utilization vs. Time
  _        QIBMCPUPTY     CPU Utilization vs. Time (Priority)
  _        QIBMCPUTYP     CPU Utilization vs. Time (Job Type)
  _        QIBMDSKARM     Disk Arm Utilization vs. Time
  _        QIBMDSKIOP     Disk IOP Utilization vs. Time
  _        QIBMDSKOCC     Percentage of Disk Occupied vs. Time
  _        QIBMLWSIOP     Local Workstation IOP Utilization vs. Time
                                                                More...
F3=Exit   F5=Refresh   F10=Restore list   F12=Cancel
F15=Sort by format     F16=Sort by text
```

On this display, 1's appear next to the graph formats that are already included in the graph package. To eliminate a graph format from the package, replace the 1 with a blank. To add additional graph formats to the package, type a 1 (Select) next to the graph formats you want to include. To display a sample of a graph format, type a 5 (Display sample graph) next to the graph format and press the Enter key. A sample graph using the graph format is displayed.

Note: You cannot change the IBM standard graph formats and graph package (QIBMxxxxxx).

Copy Graph Formats and Packages

To copy a graph format or graph package, type a 3 (Copy) next to the format or package name on the Work with Graph Formats and Packages display and press the Enter key.

```
                  Work with Graph Formats and Packages

Library . . . . . . . . . . .   QPFRDATA

Type options, press Enter.
  1=Create graph format    2=Change          3=Copy   4=Delete
  5=Display sample graph   6=Create package  8=Display package contents
```

Either the Copy Graph Format (CPYGPHFMT) or Copy Graph Package (CPYGPHPKG) command prompt appears. You can copy a graph format or package to another library or into the same library under a different name. A graph format or package that is created in a library cannot have the same name as a graph format or graph package that already exists in the library.

Copying graph formats and packages is useful for changing a base format or package, such as the IBM standard graph formats and package (QIBMxxxxxx).

Note: Option 8 (Display package contents) cannot be specified for graph formats.

Delete Graph Formats and Packages

To delete graph formats and graph packages, type a 4 (Delete) next to the format and package names on the Work with Graph Formats and Packages display, and press the Enter key.

If a graph format you selected to delete is contained in any graph packages, a warning message displays telling you that the format is in a package. If you delete the graph format, the format is also removed from the graph package. If all the graph formats in a graph package are deleted, the package is also deleted.

Note: You cannot delete the IBM standard graph formats and graph package (QIBMxxxxxx).

Display Sample Graph

To display a sample of a graph format, type a 5 (Display sample graph) next to the graph format name on the Work with Graph Formats and Packages display, and press the Enter key. A sample graph using the graph format is displayed.

Note: This option is not valid for graph packages.

Display Package Contents

To display the contents of a graph package, type an 8 (Display package contents) next to the graph package name on the Work with Graph Formats and Packages display, and press the Enter key. The Display Package Contents display appears.

```
                         Display Package Contents

Graph package . . :   PACKAGE1
  Library . . . . :     QPFRDATA

Type options, press Enter.
  5=Display sample graph

Option     Format        Text
  _        FORMAT1       CPU Utilization vs. Time-Functional A: eas
  _        QIBMASYNC     Asynchronous Disk I/O per Second vs. Time
  _        QIBMCMNIOP    Communications IOP Utilization vs. Time




                                                               Bottom

F3=Exit   F5=Refresh   F12=Cancel   F16=Sort by text
```

On this display, type a 5 (Display sample graph) to see a sample graph displayed using the graph format.

Work with Historical Data

The Display Historical Graph (DSPHSTGPH) command uses historical data to show the changes in resource utilization on your system over time. Historical data is a summary of the performance data created by the performance monitor.

Note: Files are created to contain the historical data. For each performance member with historical data, there is a single value for each type of information that can be graphed for each day of the member's performance collection period. Thus, the amount of data is reduced and summarized into the historical files. Once you have historical data for a member, you may choose to delete the performance data (DLTPFRDTA) created through the initial performance data collection to free file storage space.

Since historical graphs can help show trends in your system's performance, it is recommended that you create historical data in a given library for members that are collected at the same time. (For example, you might want to compare data that was all collected on Wednesdays from 8:00 a.m. to 12:00 p.m., whereas you probably would not want a historical graph with one member collected on Wednesday from 8:00 a.m. to 12:00 p.m. and the other on Saturday from 1:00 to 5:00 p.m.)

If you select option 2 (Work with historical data) on the Performance Tools Graphics menu, the Work with Historical Data display appears.

```
                         Work with Historical Data

Library . . . . . .   QPFRDATA

Type options, press Enter.
  1=Create historical data   4=Delete historical data

        Member      Historical
Option  Name          Data     Date      Time
  _     Q900040843    NO       01/04/90  08:43:15
  _     Q900031050    NO       01/03/90  10:51:00
  _     THURSDATA     YES      01/01/90  10:42:48
  _     TESTDATA      YES      01/01/90  10:26:12
  _     JAN011990     NO       01/01/90  09:57:27
  _     Q900010955    NO       01/01/90  09:55:41
  _     MONDAY        YES      12/30/89  11:17:03
  _     Q893632332    YES      12/29/89  23:32:19
  _     Q893621407    YES      12/28/89  14:07:11
  _     Q893621142    NO       12/28/89  11:42:30
  _     Q893611538    NO       12/27/89  15:39:02
                                                                 More...
F3=Exit   F5=Refresh   F11=Display text   F12=Cancel
F15=Sort by member     F16=Sort by text
```

The member name, a historical data indicator, and the date and time you collected each set of performance data appear on this display. To display the member text description, press F11 (Display text). If you cannot find the data you want to work with, use the appropriate function key to sort the sets of performance and historical data. You can sort them by member name, text description, or by the date and time the member was created. When you find the data you want to work with, indicate the function you want to perform by typing the appropriate option.

If you are searching for performance or historical data located in a library that is different from the one currently listed in the *Library* field at the top of the display, type a new library name in the *Library* field and press the Enter key. A list of performance and historical data members available in the library you specified appears. You can then select one of them to work with.

Note: All of the members in the historical data must have unique names. If you create a member that has the same name as a historical data member, you may want to change the name by using the Copy Performance Data (CPYPFRDTA) command to use the new member for historical purposes.

It is best to use the created name option (*GEN) on the STRPFRMON command to make sure that the names of your performance data members are unique.

Create Historical Data

To create historical data for performance members, type a 1 (Create) by the members, and press the Enter key. The Confirm Create of Historical Data display appears.

```
                     Confirm Create of Historical Data

Library . . . . . :   QPFRDATA

Press Enter to confirm your choices for 1=Create.
Press F12=Cancel to return to change your choices.

        Member      Historical
Option  Name           data     Date      Time
  1     Q900040843     NO       01/04/90  08:43:15
  1     Q900031050     NO       01/03/90  10:51:00
  1     MONDAYDATA     YES      01/01/90  10:42:48

                                                                 Bottom
F11=Display text   F12=Cancel
```

On this display, press the Enter key to create historical data for the members. Once historical data has been created for a member, you can delete the original performance data using the Delete Performance Data (DLTPFRDTA) command if the data is not needed for performance analysis, capacity planning, or performance graphing.

Delete Historical Data

To delete the historical data created by the Create Historical Data command, type a 4 (Delete) by members that contain historical data, and press the Enter key. This does not delete the original performance data.

Note: If the performance data for a member no longer exists, you cannot re-create historical data for that member after the historical data has been deleted.

Display Graphs and Packages

You can view or print graphs from your display. You can also store a graph in a GDF file for use by other utilities, such as the BGU. This is done on the Specify Graph Options display.

If you select option 3 (Display Graphs and Packages) on the Performance Tools Graphics menu, the Display Graphs and Packages display appears.

```
                         Display Graphs and Packages

Select one of the following:

     1. Display performance data graphs
     2. Display historical data graphs









Selection or command
===> ____________________________________________________________
_________________________________________________________________
F3=Exit   F4=Prompt   F9=Retrieve   F12=Cancel
```

Two distinct types of graphs can be displayed: performance graphs and historical graphs. Performance graphs use performance data collected from a single run of the performance monitor. Performance graphs are used to single out jobs that are performing poorly or to evaluate which activities were performed by a user or class of users on the system during a specified period.

Historical graphs use performance data collected from several runs of the performance monitor. Historical data is the summary of the performance data created by the STRPFRMON command. The CRTHSTDTA command is used to summarize the performance data for use by the historical graphs. Historical graphs are used to show how the performance of a system has changed over time.

Note: It is best to collect the performance data used for historical graphs over the same period of time. For example, if your normal working day is from 8:00 a.m. to 5:00 p.m., you would not want to create a historical graph to evaluate system performance during working hours using system performance data collected from 5:00 p.m. to 8:00 a.m.

If you want to collect data at predetermined times, use the Work with Performance Collection (WRKPFRCOL) command to schedule the running of the performance monitor. See Chapter 4, "Collecting System Performance Data," for more information on this command.

Display Performance Graphs

If you select option 1 (Display performance data graphs) on the Display Graphs and Packages display, the Select Graph Formats and Packages display appears.

```
                         Select Graph Formats and Packages

Library . . . . . . . . . . .    QPFRDATA

Type options, press Enter.
  1=Select   5=Display sample graph   8=Display package contents

        Format/
Option  Package     Type      Text
  _     NEWPACKAGE  PACKAGE   Graph Package for Job Types
  _     PACKAGE1    PACKAGE   Graph Package containing IOP formats
  _     QIBMPKG     PACKAGE   IBM Graph Package
  _     FORMAT1     FORMAT    CPU Utilization vs. Time-Functional Areas
  _     FORMAT2     FORMAT    Response Time vs. Time-Functional Areas
  _     QIBMASYNC   FORMAT    Asynchronous Disk I/O per Second vs. Time
  _     QIBMCMNIOP  FORMAT    Communications IOP Utilization vs. Time
  _     QIBMCPUPTY  FORMAT    CPU Utilization vs. Time (Priority)
  _     QIBMCPUTYP  FORMAT    CPU Utilization vs. Time (Job Type)
  _     QIBMDSKARM  FORMAT    Disk Arm Utilization vs. Time

                                                                 More...
F3=Exit   F5=Refresh   F12=Cancel   F14=Sort by format   F15=Sort by text
```

This display shows you the graph formats and graph packages that exist in the library you specified. The graph format or graph package name, a format or package indicator, and a text description appear on the display. If you cannot find the format or package you want to use in your performance graph, use the appropriate function key to sort the formats and packages. You can sort them by name, type, or text description. When you find the graph format or package you want to use in your performance graph, type a 1 in the corresponding *Option* field.

If you are searching for a graph format or graph package located in a library that is different from the one currently listed in the *Library* field at the top of the display, type a new library name in the *Library* field, and press the Enter key. A list of graph formats and graph packages available in the library you specified appears. You can then select one of them to use in your performance graph.

Display Sample Graph

To display a sample of a graph format, type a 5 (Display sample graph) next to the graph format, and press the Enter key. A sample graph using the graph format appears.

Note: This option is not valid for graph packages.

Display Graph Package

To display the contents of a graph package, type an 8 (Display package contents) next to the graph package, and press the Enter key. A list of the graph formats contained in the graph package appears.

Note: This option is not valid for graph formats.

Select Performance Data Member

After you select a graph format or graph package to use in your performance graph, the Select Performance Data Member display appears.

```
                          Select Performance Data Member

 Library . . . . . .    QPFRDATA

 Type options, press Enter.
   1=Select

         Member
 Option  Name        Text                                   Date      Time
   _     Q900040843                                         01/04/90  08:43:15
   _     Q900031050                                         01/03/90  10:51:00
   _     MONDATA3    Monday Data-third run                  01/01/90  10:42:48
   _     MONDATA2    Monday Data-second run                 01/01/90  10:26:12
   _     MONDATA1    Monday Data-first run                  01/01/90  09:57:27
   _     Q900010955                                         01/01/90  09:55:41
   _     TESTDATA    Test of System                         12/30/89  11:17:03
   _     Q893632332                                         12/29/89  23:32:19
   _     Q893621407                                         12/28/89  14:07:11
   _     THURSDAY    Thursday's Data                        12/28/89  11:42:30
   _     Q893611538                                         12/27/89  15:39:02
                                                                        More...
 F3=Exit   F5=Refresh      F12=Cancel    F15=Sort by member
 F16=Sort by text
```

The member name, a text description, and the date and time you collected each set of performance data appear on this display. If you cannot find the data you want to display, use the appropriate function keys to sort the sets of performance data. You can sort the data by member name, text description, or by the date and time the member was created. When you find the performance data you want to use in your performance graph, type a 1 in the corresponding *Option* field.

If you are searching for a member located in a library that is different from the one currently listed in the *Library* field at the top of the display, type a new library name in the *Library* field and press the Enter key. A list of the performance members available in the specified library appears. You can then select a member to display.

Select Categories for Performance Graphs

If the graph format or graph package you previously selected does not graph only IOP, disk, or communications line data, the Select Categories for Performance Graphs display appears.

```
                    Select Categories for Performance Graphs

Member . . . . . :   MONDAYDATA
  Library . . . . :    QPFRDATA

Type options, press Enter.  Press F6 to include all data in the graph.
  1=Select

Option     Category
  _        Job
  _        User ID
  _        Subsystem
  _        Pool
  _        Communications line
  _        Control unit
  _        Functional area



                                                                  Bottom
F3=Exit   F6=Include all data in the graph   F12=Cancel
```

Type a 1 in the *Option* column next to the categories of information from which you want performance data. Press the Enter key.

Note: Normally, you include all categories of information in your graph. To do this, do not type a 1 in any category. Instead, simply press F6.

If you choose to display the graph with only certain categories of information, a display appears that allows you to enter selection criteria for each category. This is the same as selecting categories of information to include in performance reports. See Chapter 8, "Printing Performance Reports," for more information on printing performance reports.

Specify Graph Options

When you have chosen the information you want to appear on your performance graph, or if you selected a graph format with IOP, disk, or communications line data type, the Specify Graph Options display appears.

```
                        Specify Graph Options

Type choices, press Enter.

  Graph title . . .    *MBRTEXT

  Graph subtitle . .   CPU Utilization vs. Time

  X-axis range:
    First . . . . .    *AUTO__                 *SAME, *AUTO, Number
    Last . . . . . .   _______                 Number

  Y-axis range:
    First . . . . .    *AUTO__                 *SAME, *AUTO, Number
    Last . . . . . .   _______                 Number

  Area fill . . . .    *NO                     *SAME, *YES, *NO

  Start:
    Day . . . . . .    *FIRST                  *FIRST, MM/DD/YY
    Time . . . . . .   *FIRST                  *FIRST, HH:MM:SS
                                                                More...
F3=Exit   F12=Cancel
```

Figure 9-8. Specify Graph Options

Page down to view the rest of the graph options.

```
                        Specify Graph Options

Type choices, press Enter.

  Stop:
    Day . . . . . .    *LAST                   *LAST, MM/DD/YY
    Time . . . . . .   *LAST                   *LAST, HH:MM:SS

  Output . . . . . .   *______                 *, *PRINT, *PLOT, *OUTFILE








                                                                Bottom
F3=Exit   F12=Cancel
```

On this display you can specify a new title, subtitle, axis ranges, area fill value, start time and date, stop time and date, and output value for your performance graph. If you selected a graph format for your performance graph, the values for the title, subtitle, axis ranges, and area fill defined in the graph format appear. Changing any of the values on the Specify Graph Options display only changes the format for the graph created. The graph format does not change. If you selected a graph package for your performance graph, *SAME appears for the title, subtitle, and axis ranges. *SAME means to leave these values as they are defined in the individual graph formats in the package. If you specify any new values, the new values appear on all of the graphs in the package.

For example, if you type `New Graph Title` for the graph title and the graph package contained three graph formats, the resulting three graphs would have "New Graph Title" as their title.

The area fill option allows you to override the area fill option on the graph format to display a graph more quickly. Filling in (or shading) an area is accomplished by drawing several lines. Densely shaded patterns require more lines. Each line that is drawn takes time. Consequently, the graph displays faster if area fill is not used. If the *area* fill option on the graph format is *YES, then selecting *NO for the *area fill* option causes the area not to be filled.

You may specify the start and stop date and time for the performance data to be shown in the graph. If you do not specify the start and stop date and time, the graph includes data from the first (or only) date that data was collected to the last (or only) date that data was collected.

The output option specifies how the graphs are to be displayed. Valid selections are:

*****
Shows the graph on your display station. Your display station can be either a graphics or nongraphics display station. A graphics display station shows the graph with colors, shading, and so forth. A nongraphics display station shows the graph using characters you choose to represent colors, shading, and so forth. Once your graph is shown, you may define one overlay by pressing F9 (Overlay). An overlay is a graph that is placed on top of the current graph (see "Display Graph Overlay" on page 9-22 for more information).

***PRINT**
Prints graphs on either a graphics or nongraphics printer. Supported graphics printers can be found in the online help information.

Notes:

1. If your system printer is not included in the help information, the graph can still be printed on the device. It will, however, be printed in a nongraphical format.
2. The appearance of graphs printed or displayed by graphical devices can be different from how they appear when printed or displayed by nongraphical devices, especially when *AUTO is specified for the Y (vertical) axis.

***PLOT**
Sends graphs to a plotter (a device for drawing on paper or transparencies). The 6180, 6182, 7371, and 7372 plotters are supported.

***OUTFILE**
Saves graphs in a graphics data format (GDF) file. You can use this file to display the graph on any system supporting the graphical data display manager (GDDM) function or the IBM AS/400 Business Graphics Utility (BGU) licensed program.

Note: Graph packages cannot be sent to a GDF file.

Press the Enter key to display your graph or graphs.

Display Historical Graphs

Historical graphs allow you to graphically see how your system performed during many runs of the performance monitor. This shows you how the performance of your system has changed over time. For example, it can show how processing unit utilization increased or fluctuated.

If you select option 2 (Display historical data graphs) on the Display Graphs and Packages display, the Select Graph Formats and Packages display appears. This is the same display that is shown for displaying performance graphs. (See "Display Performance Graphs" on page 9-16 for more information.) After you select a graph format or graph packge from the Select Graph Formats and Packages display, the Specify Graph Options display appears.

```
                              Specify Graph Options

 Type choices, press Enter.

   Graph title . . .   *BLANK____________________________________

   Graph subtitle . .  *BLANK____________________________________

   X-axis range:
     First . . . . .   *AUTO___                  *SAME, *AUTO, Number
     Last . . . . . .  ________                  Number

   Y-axis range:
     First . . . . .   *AUTO___                  *SAME, *AUTO, Number
     Last . . . . . .  ________                  Number

   Area fill . . . .   *NO                       *SAME, *YES, *NO

   Data library . . .  QPFRDATA__                Name

                                                                More...
 F3=Exit   F12=Cancel
```

Page down to view the rest of the graph options.

```
                              Specify Graph Options

 Type choices, press Enter.

   Start:
     Day . . . . . . . .   *FIRST__           *FIRST, *SELECT, MM/DD/YY

   Stop:
     Day . . . . . . . .   *LAST___           *LAST, MM/DD/YY

   Create historical
   data . . . . . . . . .  *NO_               *YES, *NO

   Output . . . . . . . .  *______            *, *PRINT, *PLOT, *OUTFILE




                                                                Bottom
 F3=Exit   F12=Cancel
```

Most of the values specified on this display are the same as described when specifying graph options for performance graphs (see "Specify Graph Options" on page 9-18 for details). The following values are different:

Data library
The location of the historical data to use in the graph.

Start day
The first day of historical data to include in the graph (type *SELECT here to select from a list of all the members in the specified library).

Stop day
The last day of historical data to include in the graph.

Create historical data
If this is *NO, then only members within the start and stop range that contain historical data are included in the graph. If this is *YES, then all members within the start and stop range are included in the graph, and historical data is created for members that do not have it.

Note: If *YES is typed for *Create historical data*, it might take more time to display the graph.

Display Graph Overlay

Once you have a performance graph or historical graph displayed, you can define one overlay by pressing F9 (Overlay). An overlay is a graph that is placed on top of another graph so that you can see both graphs at the same time. Overlays can help you compare one graph to another as shown below.

You must select a graph format with the same X-axis specified.

If you want to overlay a historical graph, you cannot display a graph format with functional area data type.

Note that when you overlay a graph, there is a maximum of 16 legend entries between the two graphs. Therefore, if you are currently displaying a graph with two legend entries, your overlayed graph may have only a maximum of 14 legend entries (if allowed for the data type in the graph format). See "Legends" on page 9-9 for the maximum number of legend entries for the individual data types. If you are currently displaying a graph with 16 legend entries, you cannot overlay a second graph.

Press F9 (Overlay), and the Select Graph Format display appears. Select the graph format that you want to overlay above the graph that is currently displayed.

```
                          Select Graph Format

Library . . . . . .   QPFRDATA


Type option, press Enter.
  1=Select

Option     Format         Text
  _        QIBMASYNC      Asynchronous Disk I/O per Second vs. Time
  _        QIBMCMNIOP     Communications IOP Utilization vs. Time
  _        QIBMCPUPTY     CPU Utilization vs. Time (Priority)
  _        QIBMCPUTYP     CPU Utilization vs. Time (Job Type)
  _        QIBMDSKARM     Disk Arm Utilization vs. Time
  _        QIBMDSKIOP     Disk IOP Utilization vs. Time
  _        QIBMDSKOCC     Percentage of Disk Occupied vs. Time
  _        QIBMLWSIOP     Local Workstation IOP Utilization vs. Time
  _        QIBMMFCIOP     Multifunction IOP (Comm) Util vs. Time
  _        QIBMMFDIOP     Multifunction IOP (Disk) Util vs. Time
  _        QIBMRSP        Interactive Response Time vs. Time
                                                                 More...

F3=Exit   F5=Refresh      F12=Cancel   F16=Sort by text
```

Select a graph format and press the Enter key, and the Specify Graph Overlay Options display appears.

```
                     Specify Graph Overlay Options

Type choices, press Enter.

  New graph title      *BLANK________________________________________

  New graph subtitle   *BLANK________________________________________

  Y-axis range:
    First . . . . .    *AUTO_____           *SAME, *AUTO, Number
    Last . . . . . .   __________           Number

  Area fill . . . .    *NO_                 *YES, *NO






F3=Exit   F12=Cancel
```

Specify a title and subtitle for your new, overlayed graph on this display. If you do not specify a new title and subtitle, your new graph title and subtitle are left blank.

The Y-axis range value defaults to the value that was specified in the graph format. Here, again, you have the chance to change it. You may choose to have the same range as defined in the graph format (*SAME), you may choose to have it automatically fit the range of values (*AUTO), or you may specify the range yourself by typing in the numbers.

You also select whether to have area fill in the overlayed graph.

After you press the Enter key, your two graphs should be displayed. You can use the function keys on the display to print or plot the overlay or send the overlay format to a GDF file. The following is an example of an overlay graph.

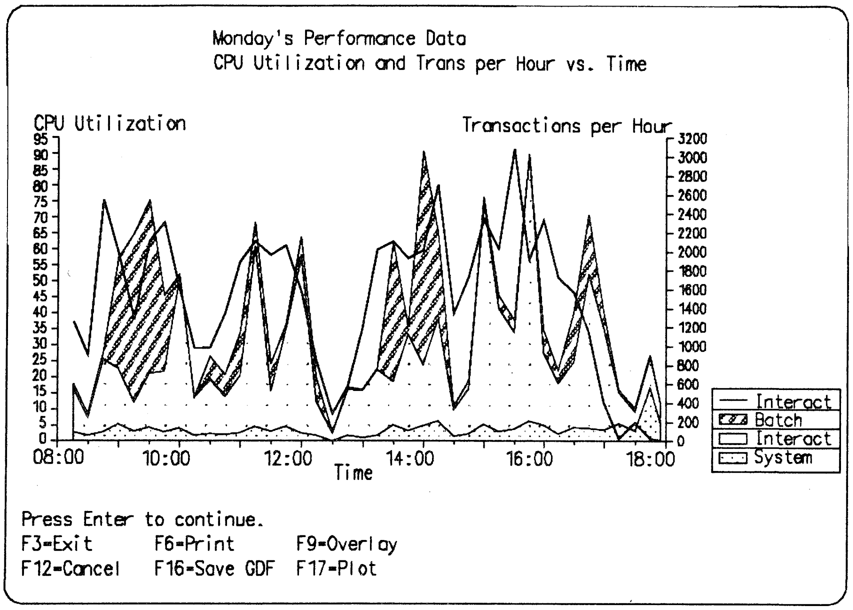

RV2S059-0

Chapter 10. Capacity Planning and Performance Prediction

This chapter describes what capacity planning is, and provides examples to show how you might use this feature.

Note: In this chapter, when response time is mentioned, assume that it is *external response time* unless stated otherwise. External response time is the end-user AS/400 system response time and includes communications time for both locally and remotely attached display stations.

For information about the elements of response time, see "Elements of Response Time" on page A-1.

Using Capacity Planning

Capacity planning is a process you use to predict your hardware configuration for future data processing needs. This process began before you installed your system. Your data processing requirements were estimated using reference data processing accounts, system sizing guides, and recommendations from the Model System (MDLSYS) command, the capacity planner for the AS/400 system.

After your system has been installed, you can continue the capacity planning process using MDLSYS. Perform capacity planning in the following situations:

- The system load grows. When you predict that the system load will increase in the volume of transactions, first run the Start Performance Monitor (STRPFRMON) command to gather performance data showing the current system load. Use this data to make projections as to the capability of the system's present configuration to handle the additional volume.
- You anticipate changes to applications. If there are plans to make functional changes to one or more applications on your system, you need to understand the effects of these changes. To evaluate the effect of the changes, collect and analyze performance data on the current application. Measure the changed application in a prototype environment. The effects of the changes can then be easily determined.
- Current performance does not meet objectives. Conduct performance analysis if the end-user response time is longer, or if the total system throughput is less than the users need. If you conduct capacity planning with current system performance data, you can gain insight into problem areas (processing unit, disk, main storage, and so on).
- An application backlog exists. If you need to install additional applications, conduct performance prediction to estimate their effects on the system. To do this, measure the current applications, and estimate the new applications using IBM-supplied profiles. Doing so provides you with data about the performance effect on the existing applications, as well as information on the performance of the new applications.
- Costs are becoming an issue. Use the output of capacity planning to assist you in estimating what, if any, additional hardware or programming changes should be made to improve performance or throughput. You can also use capacity planning to estimate the effects of changes to hardware as well as to programs. With this data, you can decide how to improve performance while minimizing cost, both in time and money.

- You are installing another AS/400 system. If the new system is to do the same type of work as your current AS/400 system but at a different throughput level, measure your current system and use that data to predict what system configuration will meet the new throughput level.

  If the new system will do completely different work, use IBM-supplied profiles to predict the system configuration.

Model System (MDLSYS) Command

Use MDLSYS, the capacity planner for the AS/400 system to identify the processing limits of your system. You can also use it to identify which system upgrades are needed to accommodate additional work. This upgrading might include any of the following parts of the system:

- Processing unit (CPU)
- Disk arms
- Main storage
- Disk input/output processors (IOPs)
- Disk storage controllers
- Display stations
- Local work station IOPs
- Remote work station controllers
- Remote lines
- Communications IOPs

When a transaction enters the system, it can be serviced by any or all of these parts. If one is busy, the transaction must wait to be serviced. A wait at any point could cause throughput and response time degradation. MDLSYS identifies where contention occurs and predicts what would happen to the workload if you upgraded these parts. Even after upgrading, the capacity planning process continues.

MDLSYS uses these sources of input:

- Workload description
- Performance objectives
- System configuration

The MDLSYS outputs include:

- Hardware upgrade recommendations (if appropriate)
- Performance predictions
- System configuration
- Printed reports
- Graphs

The system configuration is listed as both input and output for MDLSYS because you can describe your configuration as input, or you can ask MDLSYS to find a configuration to fit your projected workload.

Model System (MDLSYS) Input

The following sections describe the sources of input for MDLSYS.

Workload Description

The workload description you supply as input can be one or more capacity planning measured profiles, predefined (IBM-supplied) profiles, or combinations of these.

Note: MDLSYS supports only 5250-type devices that are attached either locally through twinaxial cable or remotely through SDLC communications lines. Review Appendix A, "Defining Transaction Boundaries," for more information on tools limitations and operational considerations.

A **capacity planning measured profile** is a collection of measured data from the Print System Report (PRTSYSRPT) command that describes how your system performs. This data consists of internal response time, throughput, processing unit time per transaction, number of disk I/Os per transaction, and so on. See "Selecting Data for Capacity Planning Measured Profiles" on page 10-11 for more information. See Appendix D, "Capacity Planner System/36 Migration Utility" for information about using System/36 measured data as input to MDLSYS.

Note: The capacity planner supports the AS/400 and System/38 environments for measured profiles used as input to MDLSYS. The System/36 environment can only be used as input to MDLSYS when there is no MRT activity present in the measured profile.

A **predefined** (IBM-supplied) **profile** contains the same type of data as the measured profile, but describes application programs that are representative of the common applications developed by users. With these predefined profiles, you can determine the performance effect of a new application on your system without having to write the new application.

One of the predefined profiles is OfficeVision/400. This workload represents the IBM OfficeVision/400 licensed program. For OfficeVision/400, MDLSYS only models the workload using OfficeVision/400 on 5250-type terminals. It does not model the workload using OfficeVision/400 on a personal computer (sometimes referred to as a programmable work station). An OfficeVision/400 profile consists of 11 predefined functions such as working with notes, mail, calendars, directories, and word processing. You can mix these functions in any combination to approximate the office workload on your system. You can select one of several workload types that provide default mixes of these functions. In addition, three user types provide default key+think times. Refer to Appendix F, "Capacity Planner Office Workload Description," for more information.

Another predefined profile is Repeatable Approach to Measuring Performance-COBOL (RAMP-C). IBM uses this set of programs to determine service levels for comparing system performance. RAMP-C is composed of four interactive transactions that range from simple to complex, depending on the amount of processing unit, disk, and main storage the transaction uses when it runs. You can mix these transactions in any combination to match what you believe to be the complexity of your interactive application. Refer to Appendix C, "Capacity Planner RAMP-C Workload Description," for more information.

Another predefined profile is a batch workload. Use it to predict the effect of batch processing on the interactive performance. These batch programs range from processor-intensive to disk-intensive. Processor-intensive batch programs do not significantly affect interactive performance (assuming batch has a lower priority than interactive and it runs in its own storage pool). A disk-intensive batch program may increase the interactive response time if the disk utilization is high. Another one of the batch program types is the data/text merge option of the OfficeVision/400 licensed program.

Note: Use the batch workloads to predict the effect of batch processing on the interactive performance. Do not use the batch workloads to predict how much batch throughput your system can do. You can see the relative changes in the batch throughput as your system workload or configuration changes.

A fourth predefined workload is spooled printing. Use this predefined workload to predict the effect of spool processing on the interactive performance. Printers can be attached to a local or remote work station controller. If the total number of lines printed per minute is large, the controller and remote line utilization can become large, causing the interactive response time to increase.

Interactive Performance Objectives

Performance objectives include the following:

- Average external response time (measured in seconds)
- Average interactive throughput (measured in transactions per hour)

External response time is the amount of time it takes the system to process one transaction. A transaction is defined as the work done by the system when you press the Enter key or a function key. The end of the transaction occurs when a new display appears, and the keyboard is unlocked, waiting for more input.

Interactive throughput is the number of transactions the system can process per unit of time.

You may specify response time or throughput objectives, or both. If you do so, the evaluator considers this information when recommending changes. These user-defined objectives may result in MDLSYS selecting a system configuration that is well below the guideline values and possibly over-designed. That is why user-defined response time and throughput objectives are optional, and it is *not* recommended that they be entered for the initial system analysis.

The factors that affect the performance objectives are as follows:

- Key+think time
- Number of active display stations

Key+think time is the user time between transactions (typing, thinking, idle time, and so on). **Number of active display stations** is the number of display stations logged onto the system and doing work (creating transactions).

The need to specify performance objectives depends on the type of work you do on your system. To make predictions, however, MDLSYS needs the key+think time and number of display stations values. If you do not know these values, you can calculate them on the Composite Workload Objectives display.

MDLSYS uses the equation shown in Figure 10-1.

$$\text{Hourly throughput} = \frac{3600 \times (\text{Number of display stations})}{\text{External Response time} + (\text{Key} + \text{think time in seconds})}$$

Figure 10-1. The MDLSYS Equation

You do not need to supply a response time objective, because MDLSYS can estimate a conservative response time. Therefore, if you know two of these three values, MDLSYS can calculate the third for you.

To estimate the key+think time for your typical work station operator (casual, interrupted, steady), estimate the average number of characters that are keyed per second or per hour. After you determine this value, use Figure 10-2 to estimate the key time, and add it to your estimate for the think time.

Figure 10-2. Key Rates and Key Times

User Type	Keys/Second	Keys/Hour	Key Time for 10 Char.	Key Time for 20 Char.	Key Time for 30 Char.	Key Time for 40 Char.	Key Time for 50 Char.
Casual	2/second	7200/hour	5.0	10.0	15.0	20.0	25.0
Interrupted	4/second	14,400/hour	2.5	5.0	7.5	10.0	12.5
Steady	5/second	18,000/hour	2.0	4.0	6.0	8.0	10.0

The system configuration is based on these objectives. The number of active display stations and key+think time both directly affect the throughput, which in turn directly affects all the system resource utilizations. In addition, the number of active display stations and working set size (main storage required per user) directly determine how much main storage you need.

System Configuration

If you are using measured profiles, the noncommunication system configuration information is captured from the measured system. Enter the measured data exactly as it is reported. Do not change the system configuration or number of display stations in the measured profile input section. Only do this on the Work with Planning Results output displays or the Composite Workload Objectives display. Since MDLSYS does not capture the measured communication configuration, it makes an estimate. You might have to change this to match your existing communications configuration.

If you do not use measured profiles, MDLSYS creates an initial, minimum system configuration estimate.

MDLSYS allows you to change the following devices and system values for an existing measured profile on the composite workload objectives or on the output displays. You can also create a measured profile by specifying the appropriate devices and system values.

- Processing unit (CPU) model number
- Number of disk arms
- Disk feature number

- Main storage size
- Number of disk IOPs
- Disk IOP feature
- Number of disk storage controllers
- Number of work stations
- Number of local work station IOPs
- Main storage interactive pool size
- Purge option for interactive jobs
- Checksums protection
- Disk mirroring option

Model System (MDLSYS) Output

The following sections describe the output for MDLSYS.

Hardware Upgrade Recommendations

MDLSYS, the capacity planner, has an evaluator as part of its analysis. MDLSYS assumes you have already done program design analysis and now want to investigate the system resources. The evaluator makes recommendations on hardware changes needed to improve and maintain performance.

To make these recommendations, the evaluator uses a threshold value, a guideline value, and a minimum value, which are unique for each resource (see Figure 10-3). These values are compared against the modeled utilizations of the key system resources.

Figure 10-3. MDLSYS Resource Evaluator Values

Resource Description	Minimum	Guideline	Threshold
Interactive CPU Utilization[1]	0.25	0.60	0.70
Disk Arm Utilization	0.20	0.35	0.45
Disk IOP Utilization	0.10	0.50	0.65
Disk Storage Controller Utilization	0.05	0.20	0.25
Local WS IOP Utilization[2]	0.10	0.45	0.60
Multiple Function IOP Utilization	0.10	0.50	0.65
Communications IOP Utilization[3]	0.10	0.45	0.60
Remote WS Controller Utilization	0.10	0.40	0.50
Remote Line Utilization	0.10	0.40	0.50

Notes:

1 This is the CPU Utilization of all activity with priority equal to or greater than the interactive jobs.

2 The modeled Local WS IOP Utilization is the processor utilization that affects response time. It is not the same as the measured Local WS IOP Utilization.

3 The modeled Communications IOP Utilization does not include the utilization due to low priority polling since this does not affect response time. Therefore, it might not match the measured utilization.

If a system resource utilization is above the threshold value, your response times could be high and unstable. Therefore, this resource must be upgraded or additional resources must be added. If the evaluator notes a threshold condition, MDLSYS automatically upgrades that resource when you press the Enter key.

Note: The evaluator messages appear on the top half of the first MDLSYS Work with Planning Results output display.

The guideline value is a warning condition. It tells you that you are approaching the problem area or threshold value. You should consider investigating possible solutions to an approaching problem. If you select the evaluator's recommendation of the hardware upgrade as the proper solution to the problem, MDLSYS does the upgrade if you press the Enter key. This is an option, not a requirement. You determine if the hardware upgrade is appropriate for your situation. If your resource utilizations are all below the guideline values, the system is considered balanced and will have both stable response time and some room for growth.

If a resource utilization is below the minimum value, the resource is underutilized. If your throughput and response time objectives are met, you may consider a hardware removal. By pressing the Enter key, MDLSYS does the recommended removal for you.

You may also specify response time or throughput objectives, or both. If you do so, the evaluator takes this into consideration when recommending changes. These user-defined objectives may result in MDLSYS selecting a system configuration that is well below the guideline values and could be over-designed. That is why user-defined response time and throughput objectives are optional, and it is *not* recommended that they be entered for the initial system analysis.

Performance Prediction

MDLSYS produces the following performance predictions:

- Average hardware utilizations (processing unit, disk, controllers, and so on)
- Average internal and external interactive response times
- Average interactive, batch, and spooled print throughputs
- Average disk I/Os per interactive and batch transactions

Compare your performance objectives with the following performance predictions:

- Response time per transaction to the processor power available
- Throughput to the response time expectations of the users
- Increased interactive throughput to a corresponding decrease in batch throughput

Use the performance predictions to decide if hardware or programming changes are needed. Generally, if your programs are written efficiently and response time is slow, you need to change or add hardware or change system parameters such as the pool configuration, pool size, or activity level. If, however, you are not sure whether your programs are written efficiently, use MDLSYS to determine if it is more feasible to change or add hardware, or to investigate the effect of program redesign by changing your workload characteristics.

System Configuration

MDLSYS displays all the system configuration information as part of the output. On the output displays, you may also change the following system configuration information. If you make a change to any of these values, your change will take the place of any evaluator recommendations.

- Processing unit (CPU) model number
- Main storage size
- Number of disk arms
- Disk feature number
- Disk IOP feature
- Number of disk IOPs
- Number of disk storage controllers
- Number of local work station IOPs
- Number of communications SDLC lines
- Number of multiple function IOPs
- Communications line speed
- Number of communications IOPs
- Communications IOP feature
- Number of remote work station controllers
- Main storage interactive pool size
- Purge option for interactive jobs
- Checksums protection
- Disk mirroring protection

MDLSYS Printed Reports

You can create printed output reports from the MDLSYS output displays by pressing the Print function key. You may include any or all of the following items in the printed report:

- Measured profiles in the current response file
- Predefined profiles in the current response file
- Objectives (including system functions)
- Planning results
- Growth

In addition, on the Work with Measured Workload Profiles display, you may select the print option to print any of the measured profiles listed (with the exception of System/36 migration profiles).

MDLSYS Graphs

It is often easier to spot a trend or compare rates of change by inspecting a graph rather than a series of numbers. MDLSYS provides a means to graphically represent the information provided.

You may save graph file data from the Work with Planning Results output displays. You may work with your graph files by pressing the Work with Graph Files function key from the Work with Planning Results output displays. You may work with previously saved graph files by pressing the Work with Graph Files function key from the Work with Capacity Planning Response Files display.

Using MDLSYS graphs, you may choose from bar graphs, line graphs, or surface graphs showing a number of the most common comparisons. The possible X-axis and Y-axis parameters are:

- X-axis:
  - Interactive throughput in transactions per hour
  - Active display stations
  - Average external response time in seconds
- Y-axis:
  - Average response time broken down by external and internal
  - CPU utilization due to interactive, system and spool, and batch activity
  - Disk arm utilization due to interactive, system and spool, and batch activity
  - Disk I/O per second due to interactive, system and spool, and batch activity
  - Disk I/O per second due to batch processing
  - Overcommitment ratio

You may print or plot a line graph or surface graph. If you are working with a bar graph, you may print it, but you cannot plot it on a plotter.

What to Expect from MDLSYS

MDLSYS is a combination of an analytic model and an evaluator. As such, the results depend entirely on the accuracy of the input, the analytic model, and the rule base for the evaluator.

The MDLSYS analytic model is based on the principles of queuing theory. Because of this, certain assumptions or conditions must exist for the analytic model to be accurate. Some primary assumptions are:

- The workload is reasonably homogeneous. That is, one particular program is not causing a resource limitation.
- The workload is steady-state. This means that the number of interactive users and batch tasks is relatively constant.
- When Purge-Y or N is indicated, it is assumed that this is a system-wide parameter rather than one assigned (specified) by task.
- When checksums or disk mirroring protection is indicated, it is assumed that this is a system-wide parameter.
- All printing uses the print spool functions of the system.
- There are no severe resource limitations (bottlenecks) in the system. That is, none of the primary devices (processing unit, disk, and so on) are excessively utilized.
- All display stations specified as active are signed on and work is being performed on them.
- The average service time of each disk request is the same. All disk requests are spread evenly across all disk drives.
- MDLSYS supports only 5250-type devices that are attached either locally through twinaxial cable or remotely through SDLC communications lines.

- One storage pool is defined for each of the three classes of jobs (interactive, batch, and spool). Jobs of different classes cannot share a storage pool, nor can a class have more than one storage pool assigned to it. *BASE is assumed to be the batch pool. Its size, the sizes of the spool and machine pools, and the initial pool activity levels are set using guidelines described in "Basic Tuning" on page 3-14.
- The activity to the local work station IOPs is spread evenly across all IOPs, and the utilization of each IOP is the same. The service time and response time depend on the application.
- The activity to the remote communications lines is spread evenly across all lines and all controllers; the utilization of each line and controller is the same. The service time and response time is dependent on the application. All lines have the same line speed.
- The job types are in the following priority order:
  1. Spool
  2. Interactive
  3. Batch

As the system approaches saturation (high utilization) of one or more resources, the accuracy of the analytic model diminishes.

An analytic model is best used to predict when a resource will become saturated. Knowing when a resource will become saturated allows you to plan your workloads and hardware upgrades.

Note: MDLSYS reports the average response time. The actual response times vary, based on the applications you run.

Figure 10-4 depicts the range of results you can expect from MDLSYS. At the higher utilizations, the range of the results is greater. Because MDLSYS is an analytic model, the results are not precise. For example, calculated average response times of 2 or 3 seconds are considered valid, when in actuality users could be experiencing an average response time of 2.4 seconds.

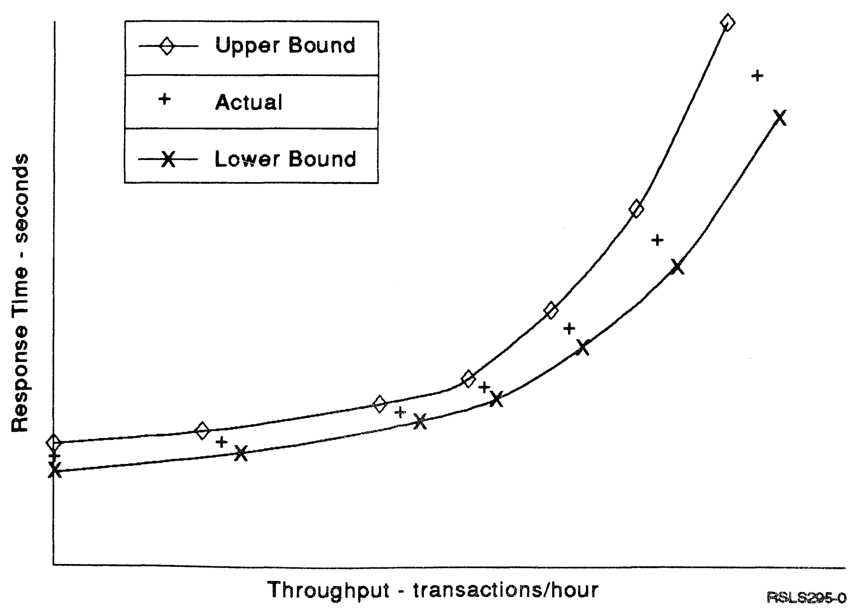

Figure 10-4. Range of Results Using MDLSYS

MDLSYS does not calculate the disk storage requirements for programs, files, and libraries. MDLSYS calculates the number of disk arms required to achieve acceptable performance based solely on the number of disk I/Os per second.

Selecting Data for Capacity Planning Measured Profiles

First use the STRPFRMON command to collect input data for capacity planning. Then use the Print System Report (PRTSYSRPT) command to produce the measured profile, using the measure profile (MSRPRF) parameter. The example "Adding Throughput to Current Workload" on page 10-13 outlines these steps. Consider the following information to help ensure valid results.

Collect data for a time period representative of a period of average load, as well as a period of heavy load. Once you collect the data, review it to ensure that it accurately represents the system, and that it is useful for capacity planning.

Always review the System Report's program summary sections to ensure that the mix of work represents the mix that you consider most important.

When you are satisfied that the data is representative, consistent, and useful for future capacity planning sessions, you can use the PRTSYSRPT command to specify that a profile be created.

Create measured profiles for a continuous period of time. Include all users in a single measured profile, or divide them into multiple profiles. Multiple profiles may be important if the workloads and performance characteristics of the users differ significantly. Multiple profiles are also important if you plan to use the growth function of the capacity planner, and the users have different growth rates.

Whenever you use the measurement interface to MDLSYS, either by itself or combined with other applications, the following assumptions and limitations are present:

- The data collection is done on a well-tuned system. This means that all pool sizes and activity levels are optimally set, secondary paging (thrashing) does not occur, and excess storage is not present. See Chapter 3, "Performance Tuning," for help with this system tuning.
- When data collection is done, all interactive jobs have the same priority, and the time slices are large enough to minimize active-to-ineligible transitions (shown as A-I in the Interactive Jobs Summary).
- If the system console is signed on but not doing productive work, you might consider omitting this specific user when creating a measured profile.
- If data collection is done when there is significant batch or spool activity, MDLSYS makes certain assumptions about this activity. One of these assumptions is that all batch jobs run at a lower priority than all interactive jobs and all spool jobs run at a higher priority than all interactive jobs. If this assumption is wrong, or if you feel that your results are inaccurate with batch and spool activity included, try creating a profile with this batch and spool activity omitted and use that profile. On most trace listings, however, there is usually a small amount of batch or spool activity, even when no batch runs are present, because the performance monitor gathers the low-level data.

- MDLSYS supports only 5250-type devices attached either locally through twinaxial cable or remotely through SDLC communications lines. If there are any other types of interactive devices that are active when the measurement is taken, their activity is not included in the total interactive activity that is captured and used in the measured profile. Examples of these other types of devices are 3270 and ASCII devices or personal computers attached through a local area network.
- Default display characteristics are the same as RAMP-C. You can, however, change these characteristics.
- As input, MDLSYS only supports the AS/400 and System/38 environments for measured profiles. The System/36 environment can only be used as input to MDLSYS when there is no MRT activity in the measured profile.
- MDLSYS does not bring in the remote line, remote work station controller, or communications IOP characteristics with the measured profile. You must change the MDLSYS estimates for these values on the Work with Planning Results display to match your environment.
- MDLSYS does not include local work station input/output processors (IOPs) with minimal utilizations.
- The capacity planner assumes no tape or diskette is attached to the disk IOP. For measured profiles added from the Work with Measured Profiles display, any disk IOPs with no disk arms are removed. Any disk arms not allocated to the system auxiliary storage pool (ASP) are removed.
- If data collection is done when applications that do field-by-field processing such as RM/COBOL-85** are running, omit these applications as MDLSYS may give inaccurate results. See "Operational Considerations" on page A-3 for more information on this limitation.

When you collect performance data to use for capacity planning, consider what abnormal situations might occur that would skew the data. Abnormal situations might include programming development activity (such as compiles), system operator activity (such as library save/restore during peak interactive times), higher-than-normal transaction rates, periods of low activity in which response time is significantly better than normal, and batch jobs running at the same priority as interactive jobs.

When using a measured profile as a model, the average response time can be affected by many values (disk seek times, service times, and so forth). MDLSYS assumes values which may vary from the actual values. This may have an impact on predicting the response.

Capacity Planning Examples

The capacity planning examples shown in the following sections demonstrate some of the ways you can use the capacity planning feature:

- Adding throughput to the current workload
- Projecting future system needs as a result of overall system growth and individual workload growth
- Adding new applications, including office, to a known workload
- Adding batch and spool workloads
- Graphing results
- Saving response files and graph files

Note: The MDLSYS results you see might be different from the results in these examples because the models are continuously updated.

MDLSYS works with four types of files. The differences among these files are explained below.

A **response file** contains all the information used by MDLSYS. It is a collection of the responses you entered while working with MDLSYS and includes system configuration data, workload descriptions, and any measured data that you are using. A list of existing response files appears on the Work with Capacity Planning Response File display, the second display you see when entering MDLSYS.

A **measured profile** results from a measurement you made on the system. See the "System Model Parameters" on page 8-22 and "Selecting Data for Capacity Planning Measured Profiles" on page 10-11 for information on how to create these profiles. The existing measured profiles are listed on the Work with Measured Workload Profiles display. You must create a response file before you can include a measured profile in the MDLSYS analysis.

A **configuration file** contains a description of a specific system configuration. This file can be used to include a specific configuration for analysis without manually changing multiple system configuration values or overlaying the current workload characteristics. You can do this on the Work with Planning Results display of MDLSYS.

A **graph file** contains planning results information based on the response file in use at the time the graph file was saved. It also contains some formatting information for defining the graph's characteristics. You may change the formatting information, but not the planning results data.

Note: The first time response files and configuration files from a previous release are used, they are automatically converted by MDLSYS to the record layout that corresponds to the current release. Once this has been done, the converted response files and configuration files cannot be moved back to the previous release and used.

Any existing measured profiles are converted to the new release while the Convert Performance Data (CVTPFRDTA) command is running.

If you are working with converted data, it is recommended that you convert the base performance data and then recreate the associated measured profiles.

Because graph file formats have not changed in Version 2, graph file conversion is not a concern.

Adding Throughput to Current Workload

If you plan to increase the number of display stations, but plan on running the same workload per user, there are several steps you need to take to plan for the increased throughput.

First, use the STRPFRMON command to measure your current system performance. Data is collected about the internal response time, throughput, resource utilizations, and system configuration parameters. For more information on the STRPFRMON command, refer to Chapter 4, "Collecting System Performance

Data." See the example following "Using Defaults to Print Performance Reports" on page 8-6 for more information on the menus involved.

Next, use the PRTSYSRPT command to build a measurement profile based on the collected data. The measurement profile is used in predicting system performance for projected needs. For more information on the PRTSYSRPT command, refer to "System Report" on page 8-10. See the example following "Using Defaults to Print Performance Reports" on page 8-6 for more information on the menus involved.

Then follow these steps to add throughput to the current workload:

1. Enter the MDLSYS command to begin working with the capacity planner. The display that shows the capacity planner disclaimer appears.

2. Press the Enter key. The Work with Capacity Planning Response File display appears.

```
                         Work with Capacity Planning Response Files

 Library . . . .   QPFRDATA__   Name

 Type options, press Enter.
 1=Create          2=Change                             3=Save
 4=Delete          5=Work with planning results         6=Grow results

 Opt   File        Description                              Date       Time
  1    EXAMPLE1    Add work stations to a measured workload









 F3=Exit                       F4=Work with graph files          F10=Sort by name
 F12=Cancel                    F13=Sort by description           F14=Sort by date
```

3. On this display, select option 1 (Create) to create a response file.

4. Type the file name in the *File* field and a description of the file in the *Description* field. In this example, EXAMPLE1 is entered as the file name, and Add work stations to a measured workload is entered as the description.

5. Press the Enter key. A response file member named EXAMPLE1 is created, and the Work with Response File Components display appears.

```
                         Work with Response File Components

 Response file . . . . . :   EXAMPLE1

 Type options, press Enter.
   2=Change        4=Delete

  Opt   Component        Type       Description
   _    Functions         4         System Functions
   2    Objectives        5         Workload Objectives









 Type:     1=Measured  2=Predefined  3=User  4=System functions
           5=Objectives    6=S/36 migration profile

 F3=Exit              F6=Measured profiles         F9=Add predefined profile
 F11=Planning results  F12=Cancel                  F14=Growth
```

6. To include measured profiles produced by the PRTSYSRPT command, press F6 (Measured profiles). The Work with Measured Workload Profiles display appears.

```
                         Work with Measured Workload Profiles

 Library . . . .    QPFRDATA___    Name

 Type options, press Enter.
        1=Add        4=Delete              6=Print

   Opt Profile    Type Description                          Date     Time
    1  COMPLEXINQ 1    COMPLEX INQUIRY WORKLOAD             10/02/90 24:22:21
    1  SIMPLEINQ  1    SIMPLE INQUIRY WORKLOAD              10/02/90 24:22:21








 Type:        1=Measured    6=S/36 migration profile

 F3=Exit           F6=Define profile information          F12=Cancel
 F16=Get S/36 migration profile                            F24=More keys
```

Profiles created using the PRTSYSRPT command are shown on this display. In this example, the current workload consists of both a simple and a complex inquiry. The PRTSYSRPT command was run twice, once to produce the COMPLEXINQ measured profile and once to produce SIMPLEINQ.

7. To select the two profiles, type a 1 (Add) in the *Opt* column next to each of the measured profiles, and press the Enter key. The measured profiles are now included in the response file. The Work with Response File Components display appears.

```
                    Work with Response File Components

 Response file . . . . . :   EXAMPLE1

 Type options, press Enter.
   2=Change        4=Delete

  Opt    Component        Type        Description
   _     Functions         4          System Functions
   2     Objectives        5          Workload Objectives

   _     COMPLEXINQ        1          COMPLEX INQUIRY WORKLOAD
   _     SIMPLEINQ         1          SIMPLE INQUIRY WORKLOAD




 Type:     1=Measured   2=Predefined   3=User   4=System functions
           5=Objectives    6=S/36 migration profile

 F3=Exit              F6=Measured profiles            F9=Add predefined profile
 F11=Planning results F12=Cancel                      F14=Growth
```

8. To enter the workload objectives for COMPLEXINQ and SIMPLEINQ, select option 2 (Change) for the Workload Objectives, and press the Enter key. The Composite Workload Objectives display appears.

```
                        Composite Workload Objectives

 Type information, press Enter.

                                                   ---------Objectives-------
              Dsp      Active     Key+Think        Response       Throughput
 Workload     Type      Dsps      Time (Seconds)   Time (Seconds) Trans/Hour
 COMPLEXINQ    1        __10       __15.0           ____.0         _______0
 COMPLEXINQ    2        ___0       _____0           ____.0         _______0
 SIMPLEINQ     1        __10       ___8.0           ____.0         _______0
 SIMPLEINQ     2        ___0       _____0           ____.0         _______0







 Dsp(Display station) type:  1=Local 2=Remote

 F3=Exit        F6=Calculate         F12=Cancel
```

This display shows that the COMPLEXINQ and SIMPLEINQ workloads were both measured at 10 local display stations. The average key+think time and throughput (measured in transactions per hour) are also shown. (You do not enter these values in this example.)

9. To enter this input as the workload objectives for the capacity planner, press the Enter key. The Work with Response File Components display appears.

```
                    Work with Response File Components

Response file . . . . . :   EXAMPLE1

Type options, press Enter.
  2=Change        4=Delete

 Opt  Component       Type       Description
  _   Functions        4         System Functions
  _   Objectives       5         Workload Objectives

  _   COMPLEXINQ       1         COMPLEX INQUIRY WORKLOAD
  _   SIMPLEINQ        1         SIMPLE INQUIRY WORKLOAD




Type:     1=Measured  2=Predefined  3=User  4=System functions
          5=Objectives    6=S/36 migration profile

F3=Exit              F6=Measured profiles          F9=Add predefined profile
F11=Planning results  F12=Cancel                   F14=Growth
```

10. To ensure that the modeled results are similar to the measured results, calibrate the model. To do this, press F11 (Planning results). The Work with Planning Results display appears.

Note: Throughout this capacity planning example, it is assumed that you press the Enter key to accept all upgrades recommended by the evaluator messages (required and optional).

```
                         Work with Planning Results                   V2R1M0

The current configuration meets your performance objectives. If you are
satisfied with the current configuration, you may press F14 (Growth) or F4
(Work with Graph Files) to run or start that function.






                               ----Objectives----  ---------Predicted----------
             Dsp    Active   Response   Thruput  --Response Time---   Thruput
Workload    Type     Dsp       Time    Trans/Hr  Internal  External  Trans/Hr
Total/Avg             20                            1.4       1.6      5,154
COMPLEXINQ    1       10        .0          0       2.8       2.9      1,326
SIMPLEINQ     1       10        .0          0       1.0       1.1      3,828


Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit                F9=Next configuration       F11=Next view
F12=Cancel                                         F24=More keys
```

This display shows you the modeled results using the measured workload objectives as a basis.

The predicted results should be similar to your measured results. If they are not, review the assumptions listed in "What to Expect from MDLSYS" on page 10-9.

11. Press F9 (Next configuration) on the Work with Planning Results display to go to the next display of system configuration data.

Note: When you use F9, you only page forward on the upper half of this display.

```
                    Work with Planning Results                       V2R1M0

 Type changed values for the base machine configuration, press Enter.
                                 Util                                    Util
 CPU model . . . . . .   D35     .48    Communications IOPs  . .   0     .00
 Main storage . . . . .    8 MB          Comm IOP feature . . . .   ___
 Disk IOPs  . . . . . .   0      .00    Remote lines . . . . . .    0    .00
 Disk arms  . . . . . :   4      .34    Remote WS controller . .    0    .00
 Local WS IOPs  . . . .   2      .09    Remote line speed  . . .     0
 Multifunction IOPs . :   1      .35
 Display station multiplier:    Base . . . .  .50  Increment . . . .   .25

                              ----Objectives----  ---------Predicted----------
           Dsp    Active   Response   Thruput   --Response Time---   Thruput
 Workload  Type    Dsp       Time    Trans/Hr   Internal  External  Trans/Hr
 Total/Avg           20                            1.4       1.6       5,154
 COMPLEXINQ  1       10        .0        0         2.8       2.9       1,326
 SIMPLEINQ   1       10        .0        0         1.0       1.1       3,828


 Note: Performance estimates - Press the Help key to see disclaimer.
 F3=Exit                F9=Next configuration        F11=Next view
 F12=Cancel                                          F24=More keys
```

12. Once the model is calibrated so that the modeled results reflect the measured results, press F6 (Change response file) to change the workload objectives to add display stations. The Work with Response File Components display appears.

```
                    Work with Response File Components

 Response file . . . . . :   EXAMPLE1

 Type options, press Enter.
   2=Change        4=Delete

  Opt   Component        Type       Description
   _    Functions          4        System Functions
   2    Objectives         5        Workload Objectives

   _    COMPLEXINQ         1        COMPLEX INQUIRY WORKLOAD
   _    SIMPLEINQ          1        SIMPLE INQUIRY WORKLOAD




 Type:     1=Measured  2=Predefined  3=User  4=System functions
           5=Objectives    6=S/36 migration profile

 F3=Exit                F6=Measured profiles         F9=Add predefined profile
 F11=Planning results   F12=Cancel                   F14=Growth
```

13. Use option 2 (Change) to change the workload objectives. Press the Enter key. The Composite Workload Objectives display appears.

```
                         Composite Workload Objectives

 Type information, press Enter.

                                                      ---------Objectives-------
                Dsp      Active        Key+Think      Response        Throughput
  Workload      Type     Dsps        Time (Seconds)   Time (Seconds)  Trans/Hour
 COMPLEXINQ     1        __10          15.0             __.0          ________0
 COMPLEXINQ     2        ___0          ___0             __.0          ________0
 SIMPLEINQ      1        __15          _8.0             __.0          ____5,267
 SIMPLEINQ      2        ___0          ___0             __.0          ________0







 Dsp(Display station) type:  1=Local 2=Remote

 F3=Exit       F6=Calculate        F12=Cancel

  Calculated throughput.
```

14. To determine the effect of adding five SIMPLEINQ display stations, enter 15 display stations as shown. If the five additional display station operators enter information at the same rate as the 10 original display station operators, the key+think time for SIMPLEINQ remains the same. The throughput, however, changes.

15. To calculate this value, position the cursor in the SIMPLEINQ *Throughput* field and press F6 (Calculate value).

    Note: Use F6 to calculate the key+think time, active display stations, or throughput. You cannot use F6 to calculate the value for the *Response Time Objective* field. If you do not enter a value in the *Response Time Objective* field, MDLSYS makes a conservative estimate of this and uses it in the calculation.

16. To enter this input as the new workload objectives for the capacity planner, press the Enter key. The Work with Response File Components display appears.

```
                    Work with Response File Components

Response file . . . . . :   EXAMPLE1

Type options, press Enter.
  2=Change       4=Delete

 Opt   Component        Type       Description
  _    Functions          4        System Functions
  _    Objectives         5        Workload Objectives

  _    COMPLEXINQ         1        COMPLEX INQUIRY WORKLOAD
  _    SIMPLEINQ          1        SIMPLE INQUIRY WORKLOAD




Type:     1=Measured  2=Predefined  3=User  4=System functions
          5=Objectives    6=S/36 migration profile

F3=Exit              F6=Measured profiles          F9=Add predefined profile
F11=Planning results  F12=Cancel                   F14=Growth
```

17. Press F11 (Planning results) to model the new workload. The Work with Planning Results display appears.

```
                         Work with Planning Results                 V2R1M0

The disk use equals or exceeds its recommended threshold. Press Enter
to add more disk. If you accept the evaluator recommendation, other resources
may be upgraded automatically to accommodate the recommendation.







                              ----Objectives----  ---------Predicted---------
            Dsp    Active    Response   Thruput   --Response Time---  Thruput
Workload    Type    Dsp        Time    Trans/Hr   Internal  External  Trans/Hr
Total/Avg            25                             3.7       3.9      3,179
COMPLEXINQ   1       10         .0          0       3.7       3.8      1,286
SIMPLEINQ    1       15         .0      5,267       3.8       3.9      1,893


Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit                 F9=Next configuration       F11=Next view
F12=Cancel                                          F24=More keys
```

The results of the modeled system appear. This time, however, the proposed workload rather than the original measured workload is modeled.

The configured hardware does not give the required performance. The evaluator message shows that more disk arms are required to achieve a balanced system.

18. To accept this change, press the Enter key. Disk arms are automatically added to the existing configuration, and the analysis is run again.

19. When the analysis is complete, the Work with Planning Results display appears with an evaluator message that shows the results of the new anal-

ysis. Continue this iterative process until the hardware configuration achieves your workload objectives. Once it does, press F9 (Next configuration) to page through the system hardware configuration data on the Work with Planning Results display.

```
                       Work with Planning Results                    V2R1M0

Type changed values for the base machine configuration, press Enter.
                              Util                                      Util
CPU model  . . . . . .   D35   .48    Communications IOPs  . .   0       .00
Main storage . . . . .     8 MB       Comm IOP feature . . . .   ___
Disk IOPs  . . . . . .   1     .19    Remote lines . . . . . .     0     .00
Disk arms  . . . . . :   6     .19    Remote WS controller . .     0     .00
Local WS IOPs  . . . .   2     .17    Remote line speed  . . .      0
Multifunction IOPs . :   1     .29
Display station multiplier:   Base  . . . .   .50  Increment . . . .   .25

                           ----Objectives----  ---------Predicted----------
           Dsp    Active   Response   Thruput  --Response Time---   Thruput
Workload   Type    Dsp       Time    Trans/Hr  Internal  External  Trans/Hr
Total/Avg           25                            2.0       2.2     3,377
COMPLEXINQ  1       10         .0          0      2.0       2.2     1,365
SIMPLEINQ   1       15         .0      5,267      2.1       2.2     2,012


Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit                F9=Next configuration        F11=Next view
F12=Cancel                                          F24=More keys
```

By adding two disk arms, the workload objectives in this example are met.

20. Press F11 (Next view) to page through the modeled system performance characteristics at various throughput levels for the same hardware configuration. The *Display station multiplier* field shows the throughput levels. The *Base* field shows the starting point (25 x 0.50 = 13). The *Increment* field shows the increment display station value (25 x 0.25 = 6).

Note: When you use F11, you only page forward on the bottom half of the Work with Planning Results display.

```
                       Work with Planning Results                    V2R1M0

Type changed values for the base machine configuration, press Enter.
                              Util                                      Util
CPU model  . . . . . .   D35   .48    Communications IOPs  . .   0       .00
Main storage . . . . .     8 MB       Comm IOP feature . . . .   ___
Disk IOPs  . . . . . .   1     .19    Remote lines . . . . . .     0     .00
Disk arms  . . . . . :   6     .19    Remote WS controller . .     0     .00
Local WS IOPs  . . . .   2     .17    Remote line speed  . . .      0
Multifunction IOPs . :   1     .29
Display station multiplier:   Base  . . . .   .50  Increment . . . .   .25

Throughput   --Avg Rsp Time--  --Active Dsp--  Key+   --Utilizations--
Trans/Hour   Local     Remote  Local   Remote  Think  CPU  Disk   Comm   OCR
    2,288     1.1         .0     13        0   19.4  .24   .09    .00   .33
    3,308     1.3         .0     19        0   19.4  .36   .15    .00   .48
    4,296     1.6         .0     25        0   19.4  .48   .21    .00   .64
    5,233     2.0         .0     31        0   19.4  .58   .28    .00   .79
    6,202     2.7         .0     38        0   19.4  .70   .38    .00   .97
    6,696     4.3         .0     44        0   19.4  .77   .54    .00  1.12
Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit                F9=Next configuration        F11=Next view
F12=Cancel                                          F24=More keys
```

The first set of values that appear on the bottom half of the display are the response times at various throughputs, the active number of display stations, average key+think time, and system resource utilizations. The main storage over commitment ratio (OCR) is a measure of the system main storage utilization.

Note: The values displayed at the various throughput levels are those that are saved when you create graph files (see page 10-13) and can be displayed using the graph options (see page 10-8).

21. If you press F11 (Next view), the next Work with Planning Results display appears with another view of system resources and characteristics at the various throughput levels.

```
                    Work with Planning Results                    V2R1M0

Type changed values for the base machine configuration, press Enter.
                               Util                                    Util
CPU model  . . . . . .  D35    .48    Communications IOPs  . .   0      .00
Main storage . . . . .    8 MB        Comm IOP feature . . . .
Disk IOPs  . . . . . .   1     .19    Remote lines . . . . . .     0    .00
Disk arms  . . . . . :   6     .19    Remote WS controller . .     0    .00
Local WS IOPs  . . . .   2     .17    Remote line speed  . . .      0
Multifunction IOPs . :   1     .29
Display station multiplier:   Base  . . . .   .50  Increment . . . .   .25

Throughput  --Interactive--  -System/Spool--  -----Batch----  -----Total----
Trans/Hour  CPU  Disk   I/O  CPU  Disk   I/O  CPU  Disk  I/O  CPU  Disk  I/O
     2,288  .23   .06    12  .02   .03     6  .00   .00    0  .24   .09   18
     3,308  .33   .10    19  .03   .05    10  .00   .00    0  .36   .15   29
     4,296  .44   .14    27  .04   .07    14  .00   .00    0  .48   .21   41
     5,233  .54   .19    35  .05   .09    17  .00   .00    0  .58   .28   52
     6,202  .65   .26    45  .06   .12    21  .00   .00    0  .70   .38   66
     6,696  .71   .36    53  .07   .18    26  .00   .00    0  .77   .54   80
Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit                 F9=Next configuration        F11=Next view
F12=Cancel                                           F24=More keys
```

22. If you press F11 (Next view) again, the next Work with Planning Results display appears.

```
                    Work with Planning Results                    V2R1M0

Type changed values for the base machine configuration, press Enter.
                               Util                                    Util
CPU model  . . . . . .  D35    .48    Communications IOPs  . .   0      .00
Main storage . . . . .    8 MB        Comm IOP feature . . . .
Disk IOPs  . . . . . .   1     .19    Remote lines . . . . . .     0    .00
Disk arms  . . . . . :   6     .19    Remote WS controller . .     0    .00
Local WS IOPs  . . . .   2     .17    Remote line speed  . . .      0
Multifunction IOPs . :   1     .29
Display station multiplier:   Base  . . . .   .50  Increment . . . .   .25

Throughput  Utilization    Interactive    --System--   Active          Response
Trans/Hour  CPU    Disk    CPU     I/O    CPU   I/O   Displays   OCR     Time
     2,288  .24     .09    .23      12    .02     6         13   .33      1.1
     3,308  .36     .15    .33      19    .03    10         19   .48      1.3
     4,296  .48     .21    .44      27    .04    14         25   .64      1.6
     5,233  .58     .28    .54      35    .05    17         31   .79      2.0
     6,202  .70     .38    .65      45    .06    21         38   .97      2.7
     6,696  .77     .54    .71      53    .07    26         44  1.12      4.3

Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit                 F9=Next configuration        F11=Next view
F12=Cancel                                           F24=More keys
```

23. If you press F11 again, the next Work with Planning Results display appears.

```
                    Work with Planning Results                         V2R1M0

Type changed values for the base machine configuration, press Enter.
                             Util                                      Util
CPU model  . . . . . .  D35   .48    Communications IOPs  . .   0      .00
Main storage . . . . .    8 MB       Comm IOP feature . . . .
Disk IOPs  . . . . . .   1    .19    Remote lines . . . . . .     0    .00
Disk arms  . . . . . :   6    .19    Remote WS controller . .     0    .00
Local WS IOPs  . . . .   2    .17    Remote line speed  . . .      0
Multifunction IOPs . :   1    .29
Display station multiplier:   Base  . . . .   .50  Increment . . . .   .25

Throughput   --Avg Rsp Time--   ---Spool Throughput---   Batch Throughput
Trans/Hour   Local     Remote   Local LPM   Remote LPM      IOs/Second
     2,288     1.1         .0           0            0               0
     3,308     1.3         .0           0            0               0
     4,296     1.6         .0           0            0               0
     5,233     2.0         .0           0            0               0
     6,202     2.7         .0           0            0               0
     6,696     4.3         .0           0            0               0
Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit                 F9=Next configuration         F11=Next view
F12=Cancel                                            F24=More keys
```

Projecting Future System Needs as a Result of Overall System Growth

Follow these steps to plan your future system needs:

1. To determine what system upgrades are required to sustain a projected growth rate, press F24 (More keys) on the Work with Planning Results display, and then press F14 (Growth). The Define Growth Rate display appears.

```
                            Define Growth Rate

Response file . . . . . :   EXAMPLE1

Type your yearly interactive growth rate.  Optionally, type growth periods in
  months. Each period is the number of months from now. To grow by workload
  press F14.

  Growth rate (%) . . . . . . . .   10   1-100

  Period 1 (months) . . . . . . .   12   0-99
  Period 2 (months) . . . . . . .   36   0-99
  Period 3 (months) . . . . . . .   60   0-99





F3=Exit                                   F12=Cancel
F13=Work with response file               F14=Grow by workload
```

2. Enter an estimate for your yearly interactive processing growth rate. Base this estimate on the following factors:

   - Gross sales
   - Acquisitions of, or mergers with, other companies

- Addition of new processes/functions
- Number of employees (especially those who use data processing equipment)

3. You may enter three growth periods (in months) to be calculated. If you do not enter growth periods, MDLSYS estimates the next three time periods when a hardware upgrade is required. In this example, the growth rate is 10%, and the growth periods are 12, 36, and 60 months.

   Note: If you are interested in analyzing the effect of growth for a particular application, press F14. An example of this analysis is shown in "Projecting Future System Needs as a Result of Individual Workload Growth" on page 10-34.

4. Press the Enter key, and the Display Growth Analysis display appears.

```
                        Display Growth Analysis                    V2R1M0

Response file . . . . . :   EXAMPLE1      Growth rate . . . . . . :   010 %

                  -- Present--  -- 12 Months--  -- 36 Months--  -- 60 Months--
                         Util            Util            Util            Util
CPU model          D35    .48      D35    .53      D45    .47      D45    .51
Main storage         8               8              16              16
Disk IOPs            1    .19        1    .23        1    .27        1    .35
Disk controllers     0    .00        0    .00        0    .00        0    .00
Disk arms            6    .19        6    .25        6    .27        6    .31
Local WS IOPs        2    .19        2    .21        2    .23        2    .26
Display stations    25              27              33              40
Multifunction IOP    1    .29        1    .48        1    .29        1    .29

Inter trans/hour      4,296          4,784           5,081           5,538
Inter response time     1.6            1.8             1.9             2.1
Batch I/Os per sec        0              0               0               0
Spooled print LPM         0              0               0               0

Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit        F9=Additional hardware       F12=Cancel
F13=Work with response file
Resource changes for each growth period are highlighted.
```

5. Press F9 (Additional hardware) for the next Display Growth Analysis display.

```
                        Display Growth Analysis                      V2R1M0

Response file . . . . . :   EXAMPLE1      Growth rate . . . . . . :   010 %

                   -- Present--  -- 12 Months--  -- 36 Months--  -- 60 Months--
                          Util            Util            Util            Util
CPU model          D35     .48      D35    .53      D45    .47      D45    .51
Comm IOPs            0     .00        0    .00        0    .00        0    .00
Remote lines         0     .00        0    .00        0    .00        0    .00
Line speed           0                0               0               0
Remote WS ctl        0     .00        0    .00        0    .00        0    .00

Inter trans/hour      4,296          4,784           5,081           5,538
Inter response time     1.6            1.8             1.9             2.1
Batch I/Os per sec        0              0               0               0
Spooled print LPM         0              0               0               0


Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit        F9=Additional hardware       F12=Cancel
F13=Work with response file
Resource changes for each growth period are highlighted.
```

The Display Growth Analysis displays show the hardware required to sustain a 10% annual growth rate for the number of months (12, 36, and 60) specified. Changes from one growth period to the next are highlighted on the display.

In this example, the capacity planner calculated that the Model D35 must be upgraded to a Model D45 with 16MB at 36 months to sustain the 10 percent growth rate. No changes are required at 12 months or 60 months except for the additional display stations.

Adding New Applications to a Known Workload

Assume that a new interactive application is being designed, and you want to determine the effect it will have on system performance. You can define this workload by using predefined profiles. In this example, use RAMP-C, an IBM-supplied profile. RAMP-C is defined as a combination of transaction types with distinct workload characteristics. You can vary the mixture of these transaction types to characterize different workloads.

1. From the Display Growth Analysis display you were on in the previous example, press F12 (Cancel) twice, and F6 (Change response file) once to go to the Work with Response File Components display.

```
                    Work with Response File Components

Response file . . . . . :   EXAMPLE1

Type options, press Enter.
  2=Change        4=Delete

 Opt    Component        Type       Description
  _    Functions          4        System Functions
  _    Objectives         5        Workload Objectives

  _    COMPLEXINQ         1        COMPLEX INQUIRY WORKLOAD
  _    SIMPLEINQ          1        SIMPLE INQUIRY WORKLOAD




Type:     1=Measured  2=Predefined  3=User  4=System functions
          5=Objectives    6=S/36 migration profile

F3=Exit              F6=Measured profiles        F9=Add predefined profile
F11=Planning results  F12=Cancel                 F14=Growth
```

2. Press F9 (Add predefined profile). The Select Predefined Profiles to Add display appears.

```
                      Select Predefined Profiles to Add

 Type options, press Enter.
   1=Select

      Predefined
 Opt  Profile         Description
 _    OFFICE1         IBM Supplied (SAA OfficeVision/400)
 1    RAMP-C          IBM Supplied (RAMP-C)
 _    BATCH           IBM Supplied (BATCH)
 _    SPOOL           IBM Supplied (SPOOL)









 F3=Exit        F12=Cancel
```

3. Type a 1 (Select) in the *Opt* column for RAMP-C and press the Enter key. The Specify Number of Predefined Profiles to Add display appears.

```
                  Specify Number of Predefined Profiles to Add

 Type choices, press Enter.

   Number of predefined profiles for:
     RAMP-C       . . . . . . . . . . . . .   1    0-5













 F3=Exit     F12=Cancel
```

4. In this example, assume that one RAMP-C-defined profile is sufficient to characterize the new application. Type a 1 and press the Enter key. The Define Ramp-C Mix display appears.

```
                          Define Ramp-C Mix

Type information, press Enter.

Predefined profile. . . . . . . . . . . . .   DEPTADD____

                                          Throughput          Key+Think
Transaction Type                            Ratio           Time (Seconds)
Inquiry                                      35                 _7.5
Simple updates                               25                 10.0
Multiple entry updates                       30                 18.0
Complex processing                           10                 25.0







F3=Exit       F12=Cancel
```

To characterize the new application, adjust the throughput ratio and the key+think time for the individual RAMP-C transaction types. See Appendix C, "Capacity Planner RAMP-C Workload Description," for more information. For simplicity, assume that the default values are sufficient.

Note: The default values are typical of an average data processing application.

Throughput ratio is the relative use of this transaction type within the application. If a program consists of one transaction type, that transaction type has a relative throughput of one. If a program consists of more than one transaction type, the relative use of each transaction type is defined as the relative number of times the transaction type is run during an interval of time. Thus, for a program with four transaction types, the first of which is run twice as often as the other three, the relative throughput may be expressed equivalently as 2,1,1,1 or 40,20,20,20.

5. Type the profile name in the *Predefined profile* field, and press the Enter key. The Work with Response File Components display appears.

```
                    Work with Response File Components

 Response file . . . . . :   EXAMPLE1

 Type options, press Enter.
   2=Change        4=Delete

  Opt    Component       Type       Description
   _     Functions        4         System Functions
   2     Objectives       5         Workload Objectives

   _     COMPLEXINQ       1         COMPLEX INQUIRY WORKLOAD
   _     SIMPLEINQ        1         SIMPLE INQUIRY WORKLOAD
   _     DEPTADD          2         IBM Supplied (RAMP-C)



 Type:     1=Measured  2=Predefined  3=User  4=System functions
           5=Objectives    6=S/36 migration profile

 F3=Exit              F6=Measured profiles          F9=Add predefined profile
 F11=Planning results F12=Cancel                    F14=Growth
```

6. To enter the workload objectives for DEPTADD, type a 2 in the *Opt* column for Objectives, and press the Enter key. The Composite Workload Objectives display appears.

```
                      Composite Workload Objectives

 Type information, press Enter.

                                                   ---------Objectives-------
              Dsp      Active      Key+Think       Response        Throughput
 Workload     Type     Dsps        Time (Seconds)  Time (Seconds)  Trans/Hour
 COMPLEXINQ    1       __10         _24.2            __.0          _______0
 COMPLEXINQ    2       ___0         _24.2            __.0          _______0
 SIMPLEINQ     1       __15         _17.1            __.0          ___2,259
 SIMPLEINQ     2       ___0         _17.0            __.0          _______0
 DEPTADD       1       ___8         _13.0            __.0          ___1,724
 DEPTADD       2       ___0         _13.0            __.0          _______0




 Dsp(Display station) type:  1=Local 2=Remote

 F3=Exit        F6=Calculate         F12=Cancel

  The throughput was calculated.
```

7. To continue capacity planning for the new application, type an 8 in the *Active Dsps* column for DEPTADD. Typing an 8 in this column assumes that eight display station operators will use the new application.

8. Now place the cursor in the *Throughput* field and press F6 (Calculate) to calculate the throughput objective.

   Note: If you know the throughput objective for this new configuration, you could enter that value in the *Throughput* field. To calculate the associated key+think time, put the cursor in the *Key+Think Time* field and press F6.

9. Press the Enter key. The Work with Response File Components display appears.

```
                     Work with Response File Components

Response file . . . . . :   EXAMPLE1

Type options, press Enter.
  2=Change        4=Delete

 Opt    Component       Type         Description
  _    Functions         4          System Functions
  _    Objectives        5          Workload Objectives

  _    COMPLEXINQ        1          COMPLEX INQUIRY WORKLOAD
  _    SIMPLEINQ         1          SIMPLE INQUIRY WORKLOAD
  _    DEPTADD           2          IBM Supplied (RAMP-C)




Type:     1=Measured  2=Predefined  3=User  4=System functions
          5=Objectives    6=S/36 migration profile

F3=Exit               F6=Measured profiles          F9=Add predefined profile
F11=Planning results  F12=Cancel                    F14=Growth
```

10. To model the effect of the new application, press F11 (Planning results). The Work with Planning Results display shows the modeled results.

With the addition of the new application, the evaluator message states that the current system configuration does not give the required performance. The capacity planner determines the additional system resources that are required to meet your performance objectives. The message communicates these recommendations at the top of the display. If you accept the proposed changes by pressing the Enter key, the capacity planner models the new configuration. Use this iterative process until the configuration achieves your performance objectives.

Note: All the displays are not shown for this iterative process in this example. Only the final display is shown here.

```
                          Work with Planning Results                  V2R1M0
The current configuration has some resource uses that exceed guidelines
but are less than thresholds; also, the interactive throughput is
higher than your objective. You may return to the Work with Response File
Components display, choose option 2 (Change Objectives), and use the
predicted response times to recalculate the objectives, the number of display
stations, or the key+think times for each application. This recommendation is
optional.  If you are satisfied with the current configuration, you may press
F14 (Growth) or F4 (Work with graph files) to start that function.


                          ----Objectives----  ---------Predicted----------
            Dsp   Active  Response   Thruput  --Response Time---   Thruput
Workload   Type    Dsp      Time    Trans/Hr  Internal  External  Trans/Hr
Total/Avg            33                          1.5       1.7       6,260
COMPLEXINQ   1       10        .0          0     2.5       2.7       1,337
SIMPLEINQ    1       15        .0      2,259     1.5       1.7       2,875
DEPTADD      1        8        .0      1,724      .9       1.1       2,048


Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit                  F9=Next configuration       F11=Next view
F12=Cancel                                           F24=More keys
```

11. Once you complete this process, press F9 (Next configuration) to page through the system configuration data.

```
                         Work with Planning Results                    V2R1M0

 Type changed values for the base machine configuration, press Enter.
                                 Util                                     Util
 CPU model  . . . . . .   D35_   .59    Communications IOPs  . .  _0      .00
 Main storage . . . . .   _16 MB        Comm IOP feature . . . .  ___
 Disk IOPs  . . . . . .   _1     .27    Remote lines . . . . . .  ___0    .00
 Disk arms  . . . . . :    6     .28    Remote WS controller . .  ___0    .00
 Local WS IOPs  . . . .   _2     .27    Remote line speed  . . .  ____0
 Multifunction IOPs . :    1     .31
 Display station multiplier:    Base  . . . .  _.50  Increment . . . .  _.25

                              ----Objectives----  ---------Predicted----------
              Dsp    Active   Response   Thruput  --Response Time---   Thruput
 Workload     Type    Dsp       Time     Trans/Hr Internal  External  Trans/Hr
 Total/Avg              33                           1.5       1.7     6,260
 COMPLEXINQ    1        10        .0          0      2.5       2.7     1,337
 SIMPLEINQ     1        15        .0      2,259      1.5       1.7     2,875
 DEPTADD       1         8        .0      1,724       .9       1.1     2,048

 Note: Performance estimates - Press the Help key to see disclaimer.
 F3=Exit               F9=Next configuration       F11=Next view
 F12=Cancel                                        F24=More keys
```

To satisfy your performance objectives with the addition of the new workload, you would have to upgrade your 8MB Model D35 to a 16MB Model D35.

If this upgrade is not acceptable, you may want to review the design of the new application. Once you characterize the new design, perform the entire process again.

Adding Batch and Spool Workloads

The process for adding batch and spool workloads is similar to adding a new application. For this reason, not all of the displays for this example are shown unless they are unique to the addition of batch and spool.

1. From the Work with Planning Results display, press F6 (Change response file) to change the existing response file. The Work with Response File Components display appears.

2. Press F9 (Add predefined profile) to add the batch and spool predefined profiles. The Select Predefined Profiles to Add display appears.

```
                     Select Predefined Profiles to Add
Type options, press Enter.
  1=Select

     Predefined
Opt  Profile          Description
_    OFFICE1          IBM Supplied (SAA OfficeVision/400)
_    RAMP-C           IBM Supplied (RAMP-C)
1    BATCH            IBM Supplied (BATCH)
1    SPOOL            IBM Supplied (SPOOL)









F3=Exit        F12=Cancel
```

3. Type a 1 in the *Opt* column for BATCH and SPOOL.

4. Press the Enter key. The Define Spool Print Jobs display appears.

```
                          Define Spool Print Jobs

Type information, press Enter.

  Printer             Number of        Total Lines Printed     Average Line
  Type                Printers            Per Minute              Length
 Local                ____1             _____500                 80
 Remote               ____0             _______0                 80










F3=Exit        F12=Cancel
```

In this example, assume one local printer is being added that you expect to drive at 500 lines per minute.

Note: Type the actual lines-per-minute value you expect, not the rate specified for the printer.

5. Type the data as shown on the previous display, and press the Enter key. Because you are also adding the batch predefined profile, the Define Batch Jobs display appears.

```
                         Define Batch Jobs

Type choices, press Enter.

  Concurrent batch jobs . . . . . . . . . .   0    0-9

  -OR-

  Specify type:
    Sequential (consecutive) . . . . . . . .  1    0-9
    Indexed sequential . . . . . . . . . . .  0    0-9
    Random by relative record
      number (direct)  . . . . . . . . . . .  0    0-9
    Indexed random   . . . . . . . . . . . .  0    0-9
    OfficeVision/400 data/text merge . . . .  0    0-9






F3=Exit        F12=Cancel
```

On this display you enter the number of each type of batch job. To view a description of each type, press the Help key.

6. In this example, add one sequential batch job, and press the Enter key.

The Work with Response File Components display appears.

```
                  Work with Response File Components

Response file . . . . . :  EXAMPLE1

Type options, press Enter.
  2=Change       4=Delete

 Opt    Component       Type       Description
  _    Functions         4        System Functions
  _    Objectives        5        Workload Objectives

  _    COMPLEXINQ        1        COMPLEX INQUIRY WORKLOAD
  _    SIMPLEINQ         1        SIMPLE INQUIRY WORKLOAD
  _    DEPTADD           2        IBM Supplied (RAMP-C)
  _    SPOOL             2        IBM Supplied (SPOOL)
  _    BATCH             2        IBM Supplied (BATCH)


Type:    1=Measured  2=Predefined  3=User  4=System functions
         5=Objectives    6=S/36 migration profile

F3=Exit               F6=Measured profiles        F9=Add predefined profile
F11=Planning results  F12=Cancel                  F14=Growth
```

This display shows that the SPOOL and BATCH profiles have been added to the response file.

7. To model the effect of adding SPOOL and BATCH, press F11 (Planning results). The Work with Planning Results display appears.

```
                     Work with Planning Results                      V2R1M0

The local workstation IOP use equals or exceeds its recommended
guideline.  Press Enter to add a local workstation IOP.
If you accept the evaluator recommendation, other resources may be upgraded
automatically to accommodate the recommendation.  This recommendation is
optional.  If you are satisfied with the current configuration, you may press
F14 (Growth) or F4 (Work with Graph files) to start that function.




                             ----Objectives----  ---------Predicted----------
            Dsp    Active   Response   Thruput   --Response Time---   Thruput
Workload    Type    Dsp       Time     Trans/Hr  Internal  External  Trans/Hr
Total/Avg             33                             1.5       1.8      6,237
COMPLEXINQ   1        10         .0          0       2.6       2.8      1,332
SIMPLEINQ    1        15         .0      2,259       1.5       1.8      2,864
DEPTADD      1         8         .0      1,724        .9       1.1      2,041

Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit                 F9=Next configuration        F11=Next view
F12=Cancel                                           F24=More keys
```

The evaluator message states that the local work station IOP exceeds the recommended guidelines. The utilization is below the threshold value, so the addition of another local work station IOP is optional.

8. To accept the recommendations of the evaluator message, press the Enter key.

9. When you are satisfied with the configuration, press F9 (Next configuration) to view the system configuration data sections of the Work with Planning Results display.

```
                     Work with Planning Results                      V2R1M0

Type changed values for the base machine configuration, press Enter.
                            Util                                       Util
CPU model . . . . . .   D35    .86   Communications IOPs . .   0       .00
Main storage . . . . .   16  MB      Comm IOP feature . . . .  ___
Disk IOPs  . . . . . .   1     .20   Remote lines . . . . . .    0     .00
Disk arms  . . . . . :   6     .33   Remote WS controller . .    0     .00
Local WS IOPs  . . . .   3     .27   Remote line speed  . . .     0
Multifunction IOPs . :   1     .33
Display station multiplier:   Base . . . .   .50  Increment . . . .   .25

                             ----Objectives----  ---------Predicted----------
            Dsp    Active   Response   Thruput   --Response Time---   Thruput
Workload    Type    Dsp       Time     Trans/Hr  Internal  External  Trans/Hr
Total/Avg             33                             1.5       1.8      6,237
COMPLEXINQ   1        10         .0          0       2.6       2.8      1,332
SIMPLEINQ    1        15         .0      2,259       1.5       1.8      2,864
DEPTADD      1         8         .0      1,724        .9       1.1      2,041

Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit                 F9=Next configuration        F11=Next view
F12=Cancel                                           F24=More keys
```

By accepting all the recommendations of the evaluator messages, the addition of BATCH and SPOOL adds an additional local work station IOP.

Projecting Future System Needs as a Result of Individual Workload Growth

This example starts where the previous example ended.

1. From the Work with Planning Results display, press F24 (More keys). Now press F14 (Growth). The Define Growth Rate display appears.

```
                              Define Growth Rate

Response file . . . . . :   EXAMPLE1

Type your yearly interactive growth rate.  Optionally, type growth periods in
  months. Each period is the number of months from now. To grow by workload
  press F14.

  Growth rate (%) . . . . . . . .  _10   1-100

  Period 1 (months) . . . . . . .  12    0-99
  Period 2 (months) . . . . . . .  36    0-99
  Period 3 (months) . . . . . . .  60    0-99







F3=Exit                                    F12=Cancel
F13=Work with response file                F14=Grow by workload
```

2. Because you want to predict future needs based on the growth of an individual interactive workload, rather than the uniform growth of the entire system, press F14 (Growth by workload). The Define Growth Rate by Workload display appears.

```
                         Define Growth Rate By Workload

Response file . . . . . :   EXAMPLE1

Type your yearly interactive growth rate by workload and the growth
  periods in months.  Each period is the number of months from now.

  Period 1 (months)  . . . . . . .  12     0-99
  Period 2 (months)  . . . . . . .  36     0-99
  Period 3 (months)  . . . . . . .  60     0-99

                                  Growth Rate
  Workload                           0-100%
  COMPLEXINQ                          _10
  SIMPLEINQ                           _10
  DEPTADD                             __0





F3=Exit        F12=Cancel        F13=Work with response file
```

3. Enter the data as shown on the previous display. This data assumes that only the complex and simple inquiry workloads will increase by 10% annually, and that you want to calculate growth at 12, 36, and 60 months. Press the Enter key. The Display Growth Analysis display appears.

```
                          Display Growth Analysis                    V2R1M0

Response file . . . . . :   EXAMPLE1       Growth by workload

                   -- Present --  -- 12 Months--  -- 36 Months--  -- 60 Months--
                             Util            Util            Util            Util
CPU model            D35      .86    D35      .88    D45      .76    D45      .81
Main storage          16               8              16              16
Disk IOPs              1      .30      1      .32      1      .41      1      .42
Disk controllers       0      .00      0      .00      0      .00      0      .00
Disk arms              6      .33      6      .36      6      .40      6      .40
Local WS IOPs          3      .27      3      .29      3      .33      3      .42
Display stations      33              36              44              53
Multifunction IOP      1      .33      1      .35      1      .32      1      .39

Inter trans/hour        4296            4784            5081            5538
Inter response time      1.6             1.8             1.9             2.1
Batch I/Os per sec         0               0               0               0
Spooled print LPM          0               0               0               0

Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit        F9=Additional hardware       F12=Cancel
F13=Work with response file
Resource changes for each growth period are highlighted.
```

4. Press F9 (Additional hardware) for the next Display Growth Analysis display.

```
                          Display Growth Analysis                    V2R1M0

Response file . . . . . :   EXAMPLE1       Growth by workload

                   -- Present--   -- 12 Months--  -- 36 Months--  -- 60 Months--
                             Util            Util            Util            Util
CPU model            D35      .86    D35      .88    D45      .71    D45      .81
Comm IOPs              0      .00      0      .00      0      .00      0      .00
Remote lines           0      .00      0      .00      0      .00      0      .00
Line speed             0               0               0               0
Remote WS ctl          0      .00      0      .00      0      .00      0      .00

Inter trans/hour        4296            4784            5081            5538
Inter response time      1.6             1.8             1.9             2.1
Batch I/Os per sec         0               0               0               0
Spooled print LPM          0               0               0               0



Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit        F9=Additional hardware       F12=Cancel
F13=Work with response file
Resource changes for each growth period are highlighted.
```

No upgrades are required at 12 months or 60 months. Capacity planning shows that you should upgrade from an 16MB Model D35 to a 16MB Model D45 at 36 months to sustain individual growth rates of 10% for both the complex and simple inquiry workloads. Also shown on the bottom of this display are batch throughput (in I/Os per second) and spool throughput (in lines per minute).

Adding Office to a Known Workload

If you plan to add display stations, which will run IBM OfficeVision/400 to your known workload, you must follow several steps to accommodate the new workload. You must first measure the current system performance, build a measured profile, and calibrate the model to ensure the modeled results are similar to the measured results. For an example, see "Adding Throughput to Current Workload" on page 10-13.

When the model is calibrated to your known workload, do the following steps to add the Office workload. This example begins where "Adding Throughput to Current Workload" on page 10-13 ends.

1. From the Work with Planning Results display, press F6 (Change response file) to change the response file. The Work with Response File Components display appears.

```
                         Work with Response File Components

 Response file . . . . . :  EXAMPLE2

 Type options, press Enter.
   2=Change       4=Delete

  Opt   Component      Type      Description
   _    Functions       4        System Functions
   _    Objectives      5        Workload Objectives

   _    COMPLEXINQ      1        COMPLEX INQUIRY WORKLOAD
   _    SIMPLEINQ       1        SIMPLE INQUIRY WORKLOAD




 Type:     1=Measured  2=Predefined  3=User  4=System functions
           5=Objectives    6=S/36 migration profile

 F3=Exit               F6=Measured profiles          F9=Add predefined profile
 F11=Planning results  F12=Cancel                    F14=Growth
```

2. Press F9 (Add predefined profiles). The Select Predefined Profiles to Add display appears.

```
                         Select Predefined Profiles to Add
Type options, press Enter.
  1=Select

     Predefined
Opt  Profile          Description
 1   OFFICE1          IBM Supplied (SAA OfficeVision/400)
 _   RAMP-C           IBM Supplied (RAMP-C)
 _   BATCH            IBM Supplied (BATCH)
 _   SPOOL            IBM Supplied (SPOOL)








F3=Exit        F12=Cancel
```

3. Type a 1 in the *Opt* column for OFFICE1 and press the Enter key. The Specify Number of Predefined Profiles to Add display appears.

```
                  Specify Number of Predefined Profiles to Add

Type choices, press Enter.

  Number of predefined profiles for:
    OFFICE1      . . . . . . . . . . . . .  1            0-5










F3=Exit     F12=Cancel
```

4. In this example, assume that one office profile is sufficient to characterize the new workload. Type a 1 and press the Enter key. The Define Office Mix display appears.

```
                              Define Office Mix

  Enter a new profile name. Change other fields if necessary.


  Predefined profile name . . . . . . ADDOFFICE


  Workload Type . . . . . . . . 1     0=User defined
                                      1=IBM Office Benchmark V2
                                      2=Secretarial
                                      3=Managerial
                                      4=Professional
                                      5=Correspondence Center
                                      (Note: Data/text merge - select
                                      BATCH predefined profile)

  User Type:  . . . . . . . . . 2     0=User defined
                                      1=Casual
                                      2=Interrupted
                                      3=Steady

  F3=Exit     F12=Cancel
```

Type the profile name in the *Predefined profile name* field. Specify the workload type and user type that most closely represents the workload to be added. In this case, workload type 1, IBM Office Benchmark, with a user type of 2 (Interrupted) is selected. See Appendix F, "Capacity Planner Office Workload Description," for more information about the office workloads.

5. Press the Enter key. The Define Office Mix Details display appears.

```
                             Define Office Mix Details

 Profile. . . . . . . . . . . . . :    ADDOFFICE
 Workload type. . . . . . . . . . :    IBM Office Benchmark V2

 Change fields if necessary, press Enter.

                                          --Thruput Ratio-- --Key+Think Time--
 Function                              Func         Trans  User            Sec
 Update One Calendar. . . . . . . . .    4            28     2            83.5
 Update Group Calendar. . . . . . . .    0             0     2            27.2
 View One Calendar. . . . . . . . . .    4            12     2           207.9
 Process Heavy Mail . . . . . . . . .    5           160     2            17.8
 Process Light Mail . . . . . . . . .    5            95     2            30.4
 Create Small Document. . . . . . . .    2            30     2            34.6
 Revise Small Document. . . . . . . .    2            30     2            39.3
 Revise Large Document. . . . . . . .    0             0     2            11.7
 View Directory Entry . . . . . . . .    6            36     2            99.2
 End/Begin Office . . . . . . . . . .    1             2     2           308.4  +

 User(User Type):  0=User defined   1=Casual   2=Interrupted   3=Steady

 F3=Exit       F12=Cancel
```

You may adjust the function throughput ratio and user type or the key+think time for the individual office functions.

Notes:

a. The default values are based on the IBM Office Benchmark Version 2 workload type that was selected on the previous display. If you change any of the defaulted values, the workload type becomes user-defined.

b. From this display, page down to see the last office function.

6. Press the Enter key. The Define Office Mix Details display appears with an informational message at the bottom of the display. The message gives the average number of transactions, per function, for this function mix. This value is necessary to translate all MDLSYS transaction throughput information into office function throughput information.

```
                              Define Office Mix Details

Profile. . . . . . . . . . . . . :    ADDOFFICE
Workload type. . . . . . . . . . :    IBM Office Benchmark V2

Change fields if necessary, press Enter.

                                      --Thruput Ratio--  --Key+Think Time--
Function                             Func       Trans  User          Sec
Update One Calendar. . . . . . . . .   4          28     2           83.5
Update Group Calendar. . . . . . . .   0           0     2           27.2
View One Calendar. . . . . . . . . .   4          12     2          207.9
Process Heavy Mail . . . . . . . . .   5         160     2           17.8
Process Light Mail . . . . . . . . .   5          95     2           30.4
Create Small Document. . . . . . . .   2          30     2           34.6
Revise Small Document. . . . . . . .   2          30     2           39.3
Revise Large Document. . . . . . . .   0           0     2           11.7
View Directory Entry . . . . . . . .   6          36     2           99.2
End/Begin Office . . . . . . . . . .   1           2     2          308.4 +

User(User Type):  0=User defined   1=Casual   2=Interrupted   3=Steady

F3=Exit        F12=Cancel
 The average number of transactions per function is  13.2.
```

7. Press the Enter key. The Work with Response File Components display appears.

```
                       Work with Response File Components

Response file . . . . . :   EXAMPLE2

Type options, press Enter.
  2=Change        4=Delete

 Opt    Component        Type        Description
  _    Functions          4         System Functions
  _    Objectives         5         Workload Objectives

  _    COMPLEXINQ         1         COMPLEX INQUIRY WORKLOAD
  _    SIMPLEINQ          1         SIMPLE INQUIRY WORKLOAD
  _    ADDOFFICE          2         IBM Supplied (OfficeVision/400)



Type:     1=Measured  2=Predefined  3=User  4=System functions
          5=Objectives    6=S/36 migration profile

F3=Exit              F6=Measured profiles           F9=Add predefined profile
F11=Planning results  F12=Cancel                    F14=Growth
```

8. To enter the workload objectives for the office profile, type a 2 in the *Opt* column for objectives, and press the Enter key. The Composite Workload Objectives display appears.

```
                         Composite Workload Objectives

 Type information, press Enter.

                                                        ---------Objectives-------
               Dsp      Active       Key+Think         Response         Throughput
  Workload     Type      Dsps       Time (Seconds)    Time (Seconds)    Trans/Hour
 COMPLEXINQ     1       __10          15.0              ___.0           ______0
 COMPLEXINQ     2       ___0          __.0              ___.0           ______0
 SIMPLEINQ      1       __15          _8.0              ___.0           __5,267
 SIMPLEINQ      2       ___0          __.0              ___.0           ______0
 ADDOFFICE      1       ___5          46.6              ___.0           ______0
 ADDOFFICE      2       ___0          46.6              ___.0_          ______0





Dsp(Display station) type:  1=Local 2=Remote

 F3=Exit       F6=Calculate              F12=Cancel
```

9. Type a 5 in the *Active Dsps* column for ADDOFFICE. This indicates that five display station operators will use this Office application mix.

10. Press the Enter key. The Work with Response File Components display appears.

```
                        Work with Response File Components

 Response file . . . . . :  EXAMPLE2

 Type options, press Enter.
   2=Change        4=Delete

  Opt     Component       Type      Description
   _     Functions         4       System Functions
   _     Objectives        5       Workload Objectives

   _     COMPLEXINQ        1       COMPLEX INQUIRY WORKLOAD
   _     SIMPLEINQ         1       SIMPLE INQUIRY WORKLOAD
   _     ADDOFFICE         2       IBM Supplied (OfficeVision/400)



 Type:     1=Measured  2=Predefined  3=User  4=System functions
           5=Objectives    6=S/36 migration profile

 F3=Exit               F6=Measured profiles          F9=Add predefined profile
 F11=Planning results  F12=Cancel                    F14=Growth
```

11. Press F11 (Planning results) to model the effect of the new office workload. The Work with Planning Results display shows the modeled results.

    With the addition of the office workload, the capacity planner predicts the system performance and evaluates any additional system resources that may be required to meet your performance objectives. The evaluator messages are shown at the top of the display. Press the Enter key to accept the proposed changes, and the capacity planner will model the new configuration. Repeat this process until the configuration achieves your performance objectives. You may change the configuration by pressing F9 to page through the system configuration data.

Note: All the displays are not shown for the process in this example. Only the final display is shown here.

```
                         Work with Planning Results                      V2R1M0

The current configuration meets your performance objectives. If you are
satisfied with the current configuration, you may press F14 (Growth) or F4
(Work with graph files) to run or start that function.







                                 ----Objectives----  ---------Predicted----------
             Dsp    Active     Response    Thruput   --Response Time---    Thruput
Workload     Type    Dsp         Time     Trans/Hr   Internal  External   Trans/Hr
Total/Avg              30                               2.5       2.7       3,681
COMPLEXINQ    1        10           .0          0       2.5       2.7       1,339
SIMPLEINQ     1        15           .0          0       2.6       2.8       1,973
ADDOFFICE     1         5           .0          0       2.0       2.2         369

Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit                 F9=Next configuration        F11=Next view
F12=Cancel                                           F24=More keys
```

Graphing Results

After you have a set of planning results, you may follow the steps below to see a graphic representation of those results.

Note: The following example starts where the previous example ended.

1. From the Work with Planning Results display, press F4 (Work with graph files) to work with graph files. The Work with Capacity Planning Graph Files display appears.

```
                    Work with Capacity Planning Graph Files

Library . . . .   QPFRDATA__    Name

Type options, press Enter.
  1=Create graph           3=Save graph file               4=Delete
  5=Work with graph        6=Combine graphs

Opt   File        Description                              Date      Time
 1    GRAPH1____  From EXAMPLE2, add office______________










F3=Exit                   F5=Refresh                       F10=Sort by name
F12=Cancel                F13=Sort by description          F14=Sort by date
```

2. Select option 1 (Create Graph) on this display to create a graph file from current planning results data.

3. Type the file name in the *File* field and a description of the file in the *Description* field. In this example, GRAPH1 is entered as the file name, and From EXAMPLE2, add office is entered as the description.

4. Press the Enter key. A graph file member named GRAPH1 is created, and the Work with Graph Page 1 of 2 display appears.

```
                               Work with Graph                   Page 1 of 2

Graph file . . . . . . . . .:  GRAPH1

Type choices, press Enter.

 Graph type          2  1=Bar graph
                        2=Line graph
                        3=Surface graph
 X axis              1  1=Interactive throughput
                        2=Active display stations
                        3=External response time
                        4=All
 Y axis              1  1=Response time
                        2=CPU utilization
                        3=Disk arm utilization
                        4=Disk I/O per second
                        5=Batch disk I/O per second
                        6=Over commitment ratio
                        7=All


F3=Exit                    F5=Save graph file                F12=Cancel
```

At this point, you have the option of working with bar graphs, line graphs, or surface graphs. Line and surface graphs require a display station with graphics capability for display.

5. Type a 2 in the *Graph type* field, a 1 in the *X axis* field, and a 1 in the *Y axis* field.

   If you type 4 for the *X axis* field or 7 for the *Y axis* field, all of the valid combinations for that axis are created. This combination of graphs may be displayed, printed, or plotted as a set.

6. Press the Enter key. The Work with Graph Page 2 of 2 display appears.

```
                              Work with Graph                   Page 2 of 2

Graph file . . . . . . . . .:   GRAPH1

To change graph definition, type information, press Enter.
Use F14 to display, F17 to plot, F18 to print selected graphs.

 Graph title. . . . . . . . . .   Work stations added____________________
 X axis range:
    Begin value . . . . . . . .   ____0  Number
    Increment . . . . . . . . .   __2.0  Number
 Y axis range:
    Begin value . . . . . . . .   ____0  Number
    Increment . . . . . . . . .   __0.4  Number






F3=Exit            F5=Save graph file              F6=Calculate
F12=Cancel         F24=More keys
```

7. Type a title in the *Graph title* field. In this example, `Work stations added` is entered as a title. If you change the X-axis range or Y-axis default range values, press the Enter key to verify the new values. For this example, leave the capacity planner default calculated values for the X-axis range and Y-axis range as they are.

8. Press F14 (Display). A display of the graph appears.

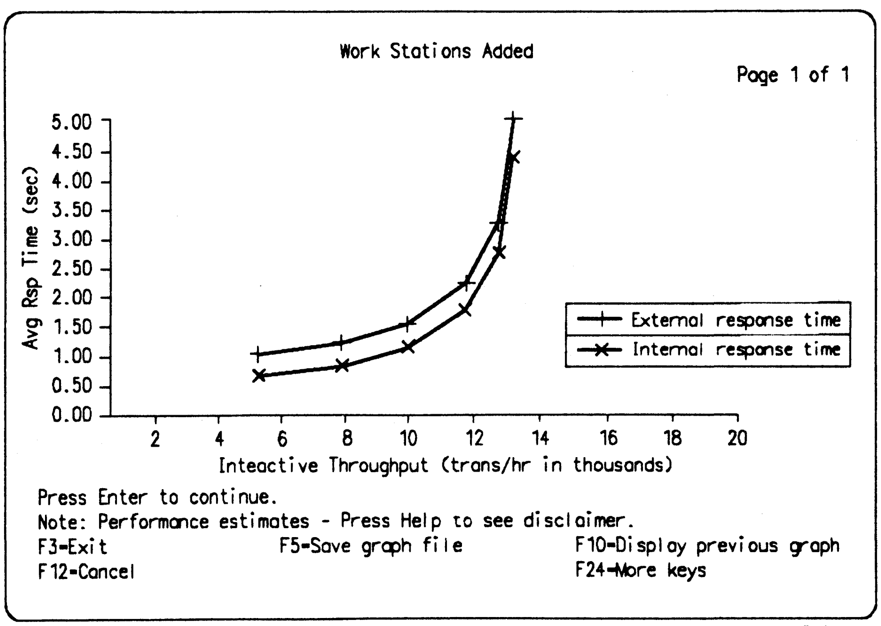

RV2S064-1

The display that you see has a title, scale factors, a legend, the graph itself, and system information. You may print or plot this graph at this point as well.

9. Press the Enter key to return to the Work with Graph Page 2 of 2 display.

```
                              Work with Graph                      Page 2 of 2

Graph file . . . . . . . . .:  GRAPH1

To change graph definition, type information, press Enter.
Use F14 to display, F17 to plot, F18 to print selected graphs.

 Graph title. . . . . . . . . .  Work stations added_____________________
 X axis range:
    Begin value . . . . . . . .  ____0  Number
    Increment . . . . . . . . .  __2.0  Number
 Y axis range:
    Begin value . . . . . . . .  ____0  Number
    Increment . . . . . . . . .  __0.4  Number







F3=Exit             F5=Save graph file                F6=Calculate
F12=Cancel          F24=More keys
```

10. Press F5 (Save graph file). The Save Current Graph File display appears.

```
                         Save Current Graph File

Library . . . . .    QPFRDATA__   Name

Type information, press Enter.

  Graph file . . . . . . .   Graph1____     Name
  Description  . . . . . .   From EXAMPLE2, add office____________
  Replace  . . . . . . . .   N              Y=Yes, N=No











F12=Cancel
```

The graph file is discarded unless you save it before leaving capacity planning graphics.

11. Press the Enter key to save your graph file. The Work with Graph Page 2 of 2 display appears again.

```
                            Work with Graph                     Page 2 of 2

Graph file . . . . . . . . . :  GRAPH1

To change graph definition, type information, press Enter.
Use F14 to display, F17 to plot, F18 to print selected graphs.

 Graph title. . . . . . . . . .  Work stations added_____________________
 X axis range:
    Begin value . . . . . . . .  ____0  Number
    Increment . . . . . . . . .  __2.0  Number
 Y axis range:
    Begin value . . . . . . . .  ____0  Number
    Increment . . . . . . . . .  __0.4  Number






F3=Exit             F5=Save graph file                F6=Calculate
F12=Cancel          F24=More keys
Graph file GRAPH1 saved in library QPFRDATA.
```

12. Press F12 (Cancel) twice to return to the Work with Capacity Planning Graph files display.

13. Press F12 (Cancel) once more to leave capacity planning graphics.

Leaving the Capacity Planner

Once you have completed the capacity planning analysis, press F3 (Exit). If you are not within capacity planning graphics, the following display appears.

```
                          Exit Capacity Planner

Select one of the following:

1. Save response file and exit
2. Exit without saving
3. Resume session











Selection
   _

F12=Cancel
```

To save the response file you have created, type a 1 (Save response file and exit) and press the Enter key. To exit without saving the response file, type a 2

(Exit without saving) and press the Enter key. To return to the display where you pressed F3 (Exit), type a 3 (Resume session) and press the Enter key.

If you elect to save the response file, the Save Current Response File display appears.

```
                              Save Current Response File

  Library . . . .     QPFRDATA__    Name

 Type information, press Enter.

   Response file . . . . .   EXAMPLE2__      Name
   Description . . . . . .   Add office to measured workload______
   Replace . . . . . . . .   N               Y=Yes     N=No









  F12=Cancel
```

Make any changes you want and press the Enter key.

Note: To save a response file during the capacity planning analysis, press F17 (Save response file) on certain displays, such as the Work with Planning Results display.

If you are working on a graph file within capacity planning graphics and press F3 (Exit), the following display appears.

```
                              Exit Capacity Planner

  Select one of the following:

  1. Save response file and exit
  2. Save graph file and exit
  3. Save response file and graph file and exit
  4. Exit without saving
  5. Resume session










  Selection
     _

  F12=Cancel
```

To save the capacity planning response file only, type a 1 (Save response file and exit) and press the Enter key. To save the capacity planning graph file only, type a 2 (Save graph file and exit) and press the Enter key. To save both the response file and the graph file, type a 3 (Save response file and graph file and exit) and press the Enter key. To exit without saving either file, type 4 (Exit without saving) and press the Enter key. To return to the display where you pressed F3 (Exit), type a 5 (Resume session) and press the Enter key.

If you type 1 to save the response file, the Save Current Response File display appears. If you type 2 to save the current graph file, the Save Current Graph File display appears. If you type 3 to save both files, the Save Current Response File display appears, followed by the Save Current Graph File display.

Chapter 11. Programmer Performance Utilities

This chapter describes the commands that you access with option 5 (Programmer performance utilities) on the IBM Performance Tools/400 menu. When you choose option 5, the Programmer Performance Utilities display appears.

```
                    Programmer Performance Utilities
Select one of the following:

     1. Work with job traces
     2. Work with program run statistics
     3. Select file and access group utilities
     4. Analyze disk activity
```

The utilities shown on the Programmer Performance Utilities display provide you with support for the detailed performance analysis of applications when you are working to understand or improve the performance of those applications.

See "Summary of Data Collection and Report Commands" on page 4-10 for an overview of the commands you use with performance tools, their data collection requirements, and their intended uses.

Job Traces

If you select option 1 (Work with job traces) on the Programmer Performance Utilities display, the Work with Job Traces display appears.

```
                         Work with Job Traces

 Select one of the following:

     1. Start job trace
     2. Stop job trace
     3. Print job trace reports
```

On this display you can choose to start or stop a job trace. After you collect the trace data, you can print job trace reports that show information about input/output (I/O) operations, file use, transaction timing, job flow, and so on.

The options in the Job Trace display and the corresponding commands are as follows:

Job Traces	**Corresponding Command**
Start Job Trace	STRJOBTRC
Stop Job Trace	ENDJOBTRC
Print Job Trace Reports	PRTJOBTRC

For more information on job traces, see "Analyzing Job Flow and Transaction Performance" on page 11-4.

Program Run Statistics

If you choose option 2 (Work with program run statistics) on the Programmer Performance Utilities display, the Work with Program Run Statistics display appears.

```
                    Work with Program Run Statistics

Select one of the following:

     1. Define environment
     2. Start data collection
     3. Stop data collection
     4. Remove environment
     5. Print run statistics report
```

On this display you can define a data collection environment, collect program operating statistics, remove the data collection environment, and create a report based on the data collected. The report shows the relative time spent in each high-level language instruction in a program.

The options in the Work with Program Run Statistics display and the corresponding commands are as follows:

Program Run Statistics	**Corresponding Command**
Define Environment	STRSAM
Start Data Collection	STRSAMCOL
Stop Data Collection	ENDSAMCOL
Remove Environment	ENDSAM
Print Run Statistics Report	PRTSAMDTA

For more information on program run statistics, see "Analyzing Program Instruction Run Time" on page 11-10.

File and Process Access Group (PAG) Utilities

If you choose option 3 (Select file and access group utilities) on the Programmer Performance Utilities menu, the Select File and Access Group Utilities display appears.

```
                 Select File and Access Group Utilities

Select one of the following:

     1. Analyze program/file use
     2. Analyze physical/logical file relationships
     3. Analyze file key structure
     4. Collect/display access group data
     5. Analyze access group data
```

On this display you can choose to create reports that show the program-to-file use, the physical-to-logical file relationships, the file key structure, or access group data. You can also use this display to determine if the application programs use shared display and database files, if the files are ordered by their frequency of use, if the files remain open but have no activity, or if programs free their static storage or keep it active.

Notes:

1. Before you use option 3, be sure that the processing for option 2 has completed. The output from option 2 is used as input for this function.
2. Option 5 is dependent on data collected by option 4. So you must take option 4 first.

The options in the Select File and Access Group Utilities display and the corresponding commands are as follows:

File and Access Group Utilities	**Corresponding Command**
Analyze Program/File Use	ANZPGM
Analyze Physical/Logical File Relationships	ANZDBF
Analyze File Key Structure	ANZDBFKEY
Collect/Display Access Group Data	DSPACCGRP
Analyze Access Group Data	ANZACCGRP

For more information on file and PAG utilities, see "Analyzing the Relationship of Programs and Database Files" on page 11-16 and "Analyzing Process Information" on page 11-24.

Analyze Disk Activity

If you choose option 4 (Analyze disk activity) on the Programmer Performance Utilities menu, the Analyze Disk Activity display appears.

```
                          Analyze Disk Activity

 Select one of the following:

      1. Start disk data collection
      2. Stop disk data collection
      3. Print disk activity report
```

On this display, you can start or stop storage management trace data collection. After you collect the data, you can print a report that shows information about I/O requests by object and by disk unit.

The options in the Disk Activity Analysis display and the corresponding commands are as follows:

Disk Activity Analysis	**Corresponding Command**
Start Data Collection	STRDSKCOL
Stop Data Collection	ENDDSKCOL
Print Disk Activity Report	PRTDSKRPT

For more information on disk activity analysis, see "Analyzing Disk Activity" on page 11-34.

Analyzing Job Flow and Transaction Performance

Use the job trace commands to collect trace information about a job. You can do this while the job runs in the normal production environment, or you can set up a special test for a job or program and trace how it runs. Once you collect the trace information, print the reports (there are two summary reports and one detail report). The summary reports allow you to determine the overall performance of the job without analyzing the detail report. Use the summary reports to guide you through the detail report.

Do not produce the detailed job analysis until you define which program or job you want to analyze.

Job trace analysis enhances the operating system's standard trace job reports and provides a summary of job operation and transaction processing. The primary use for job trace analysis is to determine how a job processes. You can determine what parts of a job use the most resources, and measure the effect of program changes relative to previous trace data. Do not use job trace analysis to determine accurate job or transaction processing times.

Start Job Trace (STRJOBTRC) Command

Use the STRJOBTRC command to start the job trace function. The End Job Trace (ENDJOBTRC) and Print Job Trace (PRTJOBTRC) commands provide summary and detail reports of the job trace data.

For information on how to enter the STRJOBTRC command, see the *CL Reference* manual.

Consider the following points when you use STRJOBTRC:

- The job trace function usually changes the paging characteristics of a job. Therefore, the trace reports may not show representative times for program operation.
- To cancel the job trace without saving any of the collected data, use the TRCJOB SET(*CNL) command.
- The job trace function issues a Start Service Job (STRSRVJOB) command if a job other than the current job is specified on the STRJOBTRC command.

End Job Trace (ENDJOBTRC) Command

Use the ENDJOBTRC command to do the following:

- Stop the job trace and direct the trace data into a user-defined database file member.
- You may start the PRTJOBTRC command to print the reports that analyze the trace data. These analysis reports provide an estimate of the response and processing times. They also show the number of database reads, non-database reads, and write I/O operations.

For information on how to enter the ENDJOBTRC command, see the *CL Reference* manual.

The database file QAPTTRCJ is created as output when you use the ENDJOBTRC command.

The printer file created by this command is the same as that created by the PRTJOBTRC command, as described in "Print Job Trace (PRTJOBTRC) Command."

Print Job Trace (PRTJOBTRC) Command

Use the PRTJOBTRC command to print a report of all, or a selected part, of the job trace data. The job trace data that prints comes from the database file member that was created when you ran the ENDJOBTRC command.

For information on how to enter the PRTJOBTRC command, see the *CL Reference* manual.

The following printer files are the output when you use the PTRJOBTRC command:

File	Description
QPPTTRC1	First part of the summary report (Trace Analysis Summary)
QPPTTRC2	Second part of the summary report (Trace Analysis I/O Summary)
QPPTTRCD	Trace record detail report (Trace Job Information)

Both the Trace Analysis Summary Report and the Trace Analysis I/O Summary Report show the job trace data detail by transaction. On these reports, two lines for each transaction show all the trace records for that transaction. A transaction boundary is determined by consecutive trace records with these characteristics:

- The first trace record indicates a call to the program specified by the end of transaction (ENDTNS) parameter.
- The second trace record indicates a return to the program specified by the start of transaction (STRTNS) parameter.

The default ENDTNS and STRTNS parameters cause the trace records to be shown by work station transactions on these reports. A transaction begins when a user presses the Enter key, or otherwise responds to a program prompt, and ends when the program next requests input from the work station. You can change these parameters in order to summarize other types of transactions, such as record processing (useful when tracing a batch job), or communications I/O.

The summary reports show you the number and types of I/O operations that occurred for each transaction, the number of full and shared file opens and closes, the number of subfile operations, and the number of messages that occurred in the transaction. Messages may be the result of normal operation or they may be due to program actions that you can avoid (full open/close, duplicate keys in a file, or incorrect subfile processing).

The summary reports also contain a reference to the detail report. Every detail record has a sequence number in it. The summaries show the starting and ending detail report sequence numbers for each transaction summarized. The detail report program can be limited to a range of sequence numbers. This feature allows you to run the summaries, then print only the detail you are interested in.

The collection of trace data takes a certain amount of processing time, the amount of which can vary depending on such factors as system load and model. This overhead time is included in the trace data on which the PRTJOBTRC command reports. The command attempts to subtract the overhead time from the reported figures, leaving only the time used for program processing. Due to the variability of the overhead time, this adjustment may not be accurate. This adjustment is an estimate only. Therefore, do not use reported processing times as an absolute measure of the response time of a program or set of programs.

Figure 11-1 shows an example of the Trace Analysis Summary Report.

```
  Title                                          TRACE ANALYSIS SUMMARY                                12/01/87
   FILE-QAJOBTRC   LIBRARY-QPFRDATA   MBR-QAJOBTRC                               JOB- BYSINN    .VLLXR239  .003368
                                            P H Y S I C A L   I / O
              SECONDS    CPU SECONDS    DB READS  NON-DB RDS        WRITES  WAITS  SEQUENCE
WAIT-ACT       4.852         .009                                                        16
ACTIVE         1.425         .788                       33              5      1        108
WAIT-ACT       4.093         .017                        3                              112
ACTIVE          .247         .110                        7              5      1        119
WAIT-ACT       3.736         .009                                                       123
ACTIVE          .658         .572                        8              5      1        180
WAIT-ACT       1.793         .005                                                       184
ACTIVE          .512         .193                       19              3      2        206
WAIT-ACT       4.195         .009                                                       210
AVERAGE         .711         .426                       18              5      1          4
TOTAL          2.842        1.703                       70             18      5
```

Figure 11-1. Trace Analysis Summary Report

The header of the Trace Analysis Summary Report shows the following values:

Title The title specified on the command.

FILE The name of the database file containing the trace data.

LIBRARY The library the database file is in.

MBR The database file member containing the trace data.

JOB The name of the job that was traced.

The columns in the detailed section of the Trace Analysis Summary Report are as follows:

ACTIVE or WAIT-ACT
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.

SECONDS
The approximate time the job was waiting or active.

CPU SECONDS
The approximate processing unit time used for the transaction. If the value is zero (or blank), you may have chosen the wrong value for the model parameter.

DB READS
The number of physical database reads that occurred.

NON-DB RDS
The number of physical nondatabase reads that occurred.

WRITES The number of physical writes that occurred.

WAITS The number of waits that occurred.

SEQUENCE
The job trace sequence number in the detail report that this summary line refers to.

AVERAGE and TOTAL
Averages and totals for the fields described above. 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.

Figure 11-2 shows an example of the Trace Analysis I/O Summary Report.

```
   Title                                           TRACE ANALYSIS I/O SUMMARY                       12/01/87
    FILE-QAJOBTRC   LIBRARY-QPFRDATA   MBR-QAJOBTRC                             JOB- BYSINN    .VLLXR239  .003368
                               P R O G R A M     *******  PROGRAM DATA BASE  I/O *******    FULL     SHARE    SUBFILE
             SECONDS  SEQNCE NAME      CALL INIT GETDR GETSQ GETKY  GETM   PUT  PUTM   UDR OPN CLS OPN CLS READS WRITES  MSGS
WAIT-ACT       4.852      16
ACTIVE         1.425     108 QPTPAGD0      1                                                         1               11
WAIT-ACT       4.093     112
ACTIVE          .247     119
WAIT-ACT       3.736     123
ACTIVE          .658     180                                                                                         11
WAIT-ACT       1.793     184
ACTIVE          .512     206                                                                     1
WAIT-ACT       4.195     210
AVERAGE         .711       4                                                                                          6
TOTAL          2.842                       1                                                     1   1               22
```

Figure 11-2. Trace Analysis I/O Summary Report

The columns in the Trace Analysis I/O Summary Report are as follows:

Title The title specified on the command.

FILE The name of the database file containing the trace data.

LIBRARY The library the database file is in.

MBR The database file member containing the trace data.

JOB The name of the job that was traced.

WAIT-ACT
The time that the job was inactive, probably due to typing or think time by the user.

ACTIVE The time the job was processing.

SECONDS
The time the job was waiting or active.

SEQNCE The job trace sequence number in the detail computer printout that this summary line refers to.

PROGRAM NAME
The name of the last program called that was not in the library QSYS before the end of a transaction.

PROGRAM CALL
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 INIT
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 DATABASE I/O
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

The values for OPENS and CLOSES in the programs are as follows:

FULL OPN
The number of full opens, for all types of files.

FULL CLS
The number of full closes, for all types of files.

SHARE OPN
The number of shared opens for all types of files.

SHARE CLS
The number of shared closes for all types of files.

The valid values for *Subfile I/O* are as follows:

SUBFILE READS
The number of subfile reads.

SUBFILE WRITES
The number of subfile writes.

MSGS The number of messages sent to the job during each transaction.

The Trace Job Information Report, shown in Figure 11-3, has essentially the same format as the system-supplied trace job output. The *Diagnostic Aids — Volume 1* manual contains additional information on trace jobs.

```
Sample Job Trace Report                               TRACE JOB INFORMATION                                   PAGE     1
FILE-QAJOBTRC    LIBRARY-QPFRDATA   MBR-QAJOBTRC         JOB- BYSINN     .VLLXR239  .003368
TIME             SEQNBR  FUNCTION  PROGRAM     LIBRARY     ENTRY EXIT  INV        CPU       DB        NON-DB    WRITTEN     WAITS
15 04 26 225
       .         000001  RETURN    QPTTRCJ1    QPFR        0077  00CF  03        .012                   1
15 04 26 262
       .         000002  CALL      QCLRTNE     QSYS        0001  002D  04                               1
15 04 26 296
       .         000003  XCTL      QCLCLNUP    QSYS        0001  0048  04        .012
15 04 26 307
       .         000004  RETURN    QPTTRCJ1    QPFR        00D0  00D0  03        .008
15 04 26 316
       .         000005  RETURN    QCMD        QSYS        016C  0153  02        .012
15 04 26 330
       .         000006  CALL      QMHRCMSS    QSYS        0001  037E  03        .012                   1
15 04 26 363
       .         000007  CALL      QMHGSD      QSYS        0001  00F5  04        .012
15 04 26 372
       .         000008  CALL      QMHRTMSS    QSYS        0001  0136  05        .008
15 04 26 383
       .         000009  RETURN    QMHGSD      QSYS        00F6  0397  04        .016
15 04 26 397
       .         000010  CALL      QWSPUT      QSYS        0001  08A6  05        .028
15 04 26 429
       .         000011  XCTL      QWSGET      QSYS        0001  027E  05        .012
15 04 26 440
       .         000012  CALL      QT3REQIO    QSYS        0001  0055  06        .061                   5          3          1
15 04 26.445     000013   T3-ENTRY
15 04 26.447     000014   T3REQIO-REQIO
15 04 31.285     000015   T3DEQ-DEQ
```

Figure 11-3. Trace Job Information Report

The columns in the Trace Job Information Report are as follows:

TIME The time of day for the trace entry. The time is sequentially given in hours, minutes, seconds, and fractions of a second.

SEQNBR The number of the trace entry.

FUNCTION This causes the trace entry to be recorded. The possible trace entries are as follows:

Trace Entry	Description
Call	Call external.
Data	A data trace.
Event	Event handler.
EXTXHINV	External exception handler.
EXTXHRET	Call termination because of a return from an exception.
INTXHINV	Internal exception handler.
INTXHRET	Return from an exception.
INVEXIT	Call because of a call exit routine.
ITERM	Intervening call termination.
ITRMXRSG	Call termination because of a resignaling exception.
PTRMTPP	Process termination.
PTRMUNHX	Termination because of an unhandled exception.
Return	Return external.

RSMTRC	Trace resumed.
SSPTRC	Trace suspended.
XCTL	Transfer control.

PROGRAM The name of the program for the entry.

LIBRARY The library name that contains the program associated with the trace entry.

ENTRY The instruction in the program where the program was given control. This is true when a program is nonobservant and observant.

EXIT The instruction number in the program where the program gave up control.

INV The call level of the program.

CPU 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.

DB The number of physical database reads that occurred for the entry.

NON-DB The number of physical nondatabase reads that occurred for the entry.

WRITTEN The number of physical writes that occurred for the entry.

WAITS The number of waits that occurred for the entry.

The read and write counts do not include any asynchronous I/O operations. The counts indicate the number of I/O requests (either single or multiple page) sent to the device, and describe the request queuing at the device.

Analyzing Program Instruction Run Time

Use the Sampled Address Monitor to determine the relative amount of processing time spent in different parts of a program or set of programs. This tool does not provide data about total processing time used by a program, but it can identify high-leverage areas of a program to improve.

For example, suppose you are interested in program SPTEST. You want to know where time is being spent in the program. SAM allows you to define a window (using the Start SAM (STRSAM) command) over SPTEST, which is subdivided into smaller units called **panes**. You also specify the rate at which sampling is to be done (using the Start SAM Collection (STRSAMCOL) command).

At each sampling, SAM determines in which pane of the defined window for SPTEST that operation is occurring, if any. If the program is running an instruction in a pane, the counter corresponding to that pane is increased by 1. When the sampled address monitor is stopped (End SAM Collection (ENDSAMCOL) command), the counters for the panes in the window are put into a database file where a data reduction program (Print SAM Data (PRTSAMDTA) command) can access them.

Because instruction operation is sampled only once every millisecond at the smallest sampling interval, a program or function within the program must be run many times to get meaningful output.

While sampling is active, system performance is degraded exponentially relative to the inverse of the sampling interval. This degradation is approximately 40% for the shortest interval of 1 millisecond. It decreases as the sampling interval increases: 10% for a 5-millisecond interval, 5% for 10 milliseconds, and so on.

Note: When SAM is active, it continues to operate until you turn it off (ENDSAMCOL and ENDSAM). The function continues to operate even across an initial program load (IPL) unless it is turned off. Given the degradation that occurs when SAM is active, it is *very important to end your use of SAM by using an ENDSAMCOL and ENDSAM*.

If the system (or the subsystem that SAM is running in) stops abnormally, and ENDSAMCOL and ENDSAM have not completed, the SAM might still be in use when the system starts again. You should enter ENDSAMCOL and ENDSAM, when the system is available, to ensure that SAM is not active.

You can perform detailed program statement analysis when you have defined which programs you want to analyze. Use the SAM report to guide you to the program instructions that use the most processing unit time. Once you determine which program instructions these are, see if changes can be made to improve the program's performance. Perhaps you can redefine the data types, change the algorithms, or change the array indexing in the program.

A typical program analysis session might run like this:

1. STRSAM PGM(lib/pgmname).
2. STRSAMCOL (1-millisecond sample intervals).
3. Run whatever it is to be measured. (You may have to do several repeated runs of a given code path to get enough sample points to reflect meaningful output.)
4. ENDSAMCOL MBR(dbmbr) TITLE('SAM test data').
5. ENDSAM. (This can also be done after the PRTSAMDTA command.)
6. PRTSAMDTA MBR(dbmbr).

Start Sampled Address Monitor (STRSAM) Command

Use the STRSAM command to prepare the environment (programs and windows) to start the sampling of instruction addresses. Once you define the environment, it remains defined until you issue the ENDSAM command, even across initial program loads.

No detectable change to the performance of the system occurs when the SAM environment is defined, unless you issue the STRSAMCOL command and do not issue the ENDSAMCOL command. The STRSAM command must be preceded by an ENDSAM command if you have defined but not deleted any windows (using ENDSAM). If you recompile the program, you will have to issue the ENDSAMCOL command and then issue the STRSAMCOL command again.

For information on how to enter the STRSAM command, see the *CL Reference* manual.

Start Sampled Address Monitor Data Collection (STRSAMCOL) Command

Use the STRSAMCOL command to start the data collection process for the group of windows you last defined with the STRSAM command. Instruction sample statistics accumulate for each defined window until you issue the STRSAMCOL command with the CLEAR (*YES) option (the default) or until you issue the ENDSAMCOL command, which temporarily suspends, but does not necessarily stop accumulation.

If the sampling interval is set to 1 millisecond, a hardware interrupt occurs every millisecond to determine if the machine is processing an instruction within one of the defined panes. If it is, the occurrence is recorded in an internal SAM data collection area, and normal instruction processing continues.

Note: Sample data is collected for all active jobs in the system, not only the job that the SAM commands were issued in. In other words, the counters are incremented even if another job runs a program you selected for sampling.

For information on how to enter the STRSAMCOL command, see the *CL Reference* manual.

End Sampled Address Monitor Data Collection (ENDSAMCOL) Command

Use the ENDSAMCOL command to stop the sample data collection and store the accumulated internal data for processing later by the PRTSAMDTA command or report creation. This command does not change the windows you defined by the previous STRSAM command, nor does it clear the collected data for that group. To keep the data and use it cumulatively, issue the STRSAMCOL command with the CLEAR(*NO) option.

Note: To use the ENDSAMCOL command, you must have issued the STRSAM and STRSAMCOL commands (and you must not have issued an ENDSAM command).

For information on how to enter the ENDSAMCOL command, see the *CL Reference* manual.

When you use the ENDSAMCOL command, the following are created as output:

File	Description
QAPTSAMH	Database file (input to the PRTSAMDTA command)
QAPTSAMV	Database file (input to the PRTSAMDTA command)

End Sampled Address Monitor (ENDSAM) Command

Use the ENDSAM command to delete the sampled address monitor data collection environment, which includes any existing internal data, and the current window definitions from the previous STRSAM command.

The ENDSAM command is not valid if you enter it while in the data collection mode (you issued the STRSAMCOL command, but not the ENDSAMCOL command).

For information on how to enter the ENDSAM command, see the *CL Reference* manual.

Print Sampled Address Monitor Data (PRTSAMDTA) Command

Use the PRTSAMDTA command to format and print the sample data saved by the ENDSAMCOL command.

For information on how to enter the PRTSAMDTA command, see the *CL Reference* manual.

The following input files are used by the PRTSAMDTA command:

File	Description
QAPTSAMH	Database file that is the output from the ENDSAMCOL command
QAPTSAMV	Database file that is the output from the ENDSAMCOL command

When you use the PRTSAMDTA command, the print file QPPTSAM is created as output.

Figure 11-4 shows an example of the SAM report. Blank lines have been removed to conserve space.

```
V2R1M0  91/05/24                                       SAM Example                                         05/17/88 Page    1
Summary of Run Information
Run start:  5/17/88 14:49:04     Stop: 14:49:44     CPU:        30,305 milliseconds
Inter-sample time:       1 milliseconds     Sum of counts:        19,675     Idle counts:          4,516
Number of windows:  1     Maximum count:          3,086
Serial number: 10-01013       Model: B60

V2R1M0  91/05/24                                       SAM Example                                         05/17/88 Page    2
Summary of Window Information
No   Panes   PSize      No Outside         Maximum            Sum  Address       Window ID
--  ------  ------  -------------  -------------  -------------  ------------  --------------------------------
 1     337       2            471          3,086         14,688  00AB49000458  QPTSAMTC.QPFRTEMP      WINDOW  01

V2R1M0  91/05/24                                       SAM Example                                         05/17/88 Page    3
Window:  1      Start: 00AB49000458     ID: QPTSAMTC.QPFRTEMP      WINDOW  01
Program: QPTSAMTC        Library: QPFRTEMP        Compile:  5/09/88 11:20:53
HLL Range    HLL %     Count          Scale Factor: 40
ILOOP         0.22        32 *
JLOOP         4.56       670 *****************
550          62.27      9146 ********************************************************************************
560          24.09      3538 ********************************************************************************
570           6.28       922 ************************
580           2.59       380 **********
```

Figure 11-4. Sampled Address Monitor Report

The first line of each page of the report includes:

- Information that identifies the current system version
- Data set title specified on the ENDSAMCOL command
- Date of the report
- Page number

The Summary of Run Information section of this report shows summary information for the SAM session, or run. Following are the columns in this section.

Run start Date and time the session began. This value corresponds to the time the STRSAMCOL command ran.

Stop Time the session ended. This value corresponds to the time the ENDSAMCOL command ran.

CPU Number of milliseconds of all processing time used during the session. This value includes all jobs running on the system, not just your own.

Inter-sample time
Number of milliseconds between samples.

Sum of counts
Total number of instruction samples that occurred within the defined windows during the session.

Idle counts
Number of instruction samples when the processor was idle (not running any job on the system).

Number of windows
Number of windows set up by the STRSAM command.

Maximum count
Largest number of instruction samples that occurred in one pane.

Serial number
Processor serial number.

Model Processor model code.

The Summary of Window Information section of this report summarizes information for each defined window.

No Window number. If the program requires several windows, adjust the percentages for each window to show their actual value for the program.

Panes Number of panes in this window.

PSize Size, in bytes, of the panes in this window.

No Outside
Number of instruction samples that occurred outside of this window.

Maximum Maximum number of instruction samples that occurred in any of the panes in this window.

Sum Sum of all the instruction samples that occurred in all of the panes for this window.

Address Window virtual address.

Window ID
Sampled program name and window number.

The remaining pages of the report summarize the sample counts by instruction. The sample report was created using the PRTSAMDTA command's RPTTYPE(*HLL) parameter, so samples are grouped by high-level-language instruction, or program statement. When using RPTTYPE(*HLL), the program must have been compiled without removing observability. (To see greater levels of detail, specify *MI or *FULL instead of *HLL. Additional columns print for these levels, and are briefly explained here, but are seldom useful except to systems programmers.)

HLL Range
The high-level-language statement number that the sample data is for. Statements with no samples are not listed.

Note: Always consult a current printout of the program to determine the specific high-level-language statement that the SAM output refers to. Use the DETAIL(*SERVICE) option of the Display Object Description (DSPOBJD) command to find out if a program object, the compiler list you use, and the current source member match at the date and time levels.

When an instruction sample is taken, the system uses the current value of the machine's model-unique licensed internal code instruction address register (IAR) to determine which window and pane to increase. For most instructions, but not all (such as branch, move long, and compare long), the IAR points to the instruction following the one currently running. Therefore, depending on the pane size you use and the current instruction, it is possible that the counts can differ by one instruction or one pane. The PRTSAMDTA command does not take this variation into account.

As a result, if the high-level-language instruction is one that translates to a single machine instruction, the results shown with the PRTSAMDTA command may be misleading. However, instructions like these usually do not cause performance problems and the variation is tolerable.

HLL % Amount of processing time used by the HLL statement number compared to the other HLL statements in the window.

In this example, statement 550 was the active statement for 62.27% of the time the sampling range was in window number 1.

Note: Because the lower-level instructions may overlap a pane if the pane size is greater than 2, the high-level-language percentage may not be exact. However, it is close enough to be useful in that you can determine within one or two high-level-language statements which ones are being used the most.

For a program that spans several windows, the percentages shown are relative only to the window in which they are reported. To determine the percentage of use for one statement compared to all windows over the program, determine the total counts for all windows and divide this total into the count for the specific instruction.

MI N The machine instruction (MI) number. It appears only if the report type is *MI or *FULL.

MI % The MI percentage of operation in the window. It appears only if the report type is *M1 or *FULL.

Offset The instruction offset in the window. It appears only if the report type is *FULL.

% The instruction percentage of operation in the window. It appears only if the report type is *FULL.

Count The number of samples taken while running the instruction. This value is followed by a chart that can be used to visually determine the time spent processing the instruction compared to the time spent processing the other instructions.

Analyzing the Relationship of Programs and Database Files

Use the Analyze Program (ANZPGM) command and the Analyze Database File (ANZDBF) command to print an overview of the programs and files used in an application. The commands provide reports showing program-to-file use and physical and logical file relationships in the libraries.

Use the Analyze Database File Keys (ANZDBFKEY) command to print an overview of the key structure of logical files in an application.

These commands provide you with a file and program use overview, and key definition detail. It may be that your files or programs have changed since they were first written and the file use has changed. For example, there may now be more logical files over your physical files than the application currently needs. This situation can cause performance degradation, especially if many key field changes or record adds occur. Remove any unneeded logical views.

Although you may use these commands infrequently, it is recommended that you use them periodically to get a good understanding of the program-to-file relationships, and of the logical file structure used in the applications.

Analyze Program (ANZPGM) Command

Use the ANZPGM command to produce reports showing program-to-file and file-to-program relationships.

For information on how to enter the ANZPGM command, see the *CL Reference* manual.

When you use the ANZPGM command, the following print files are created as output:

File	**Description**
QPPTANZP	The program-to-file relationship report (Program-to-File Cross-Reference)
QPPTANZP	The file-to-program relationship report (File-to-Program Cross-Reference)

Figure 11-5 shows an example of the Program-to-File Cross-Reference Report.

```
12/01/87 13:37:09                  Program to File Cross-Reference                                          Page     1
                                                                                                            File Usage
                                                                                                 Record     1=In 2=Out
Library    Program    Program Text Description                         Object      Library     Format     4=Upd 8=?
---------- ---------- ------------------------------------------------ ----------- ----------- ---------- ----------
QPFR       OLDPTCHGJR                                                  QAPMDMPT    *LIBL                       1
                                                                       QDPTJTYP    QPFR        SFL             3
                                                                       QDPTJTYP    QPFR        SFLCTL          3
                                                                       QDPTJTYP    QPFR        QDPTF1          3
                                                                       QDPTJTYP    QPFR        QDPTF2          3
                                                                       QDPTJTYP    QPFR        QDPTF3          3
                                                                       QDPTJTYP    QPFR        HELP1           3
                                                                       QTRIDX      QPFRDATA    IDXREC          6
                                                                       QJTYP1      QPFRTEMP    IDXREC          1
                                                                       QJTYP2      QPFRTEMP    IDXREC          6
           OLDPTTNSRP                                                  *FILE                                   8
                                                                       QAPMJOBS                                8
                                                                       QSYSPRT                                 8
                                                                       QTRIDX                                  8
                                                                       QTRINTD                                 8
                                                                       QTRINTU                                 8
                                                                       QTRJOBI                                 8
                                                                       QTRJOBO                                 8
                                                                       QAPMDMPT    &LIB                        8
                                                                       QAPMJOBS    &LIB                        8
                                                                       QTRIDX      &LIB                        8
                                                                       QTRJSUM     &LIB                        8
                                                                       QTRTSUM     &LIB                        8
                                                                       QDDSSRC     *LIBL                       1
                                                                       QPTMPLST    *LIBL
                                                                       QPTTRIDX    *LIBL
                                                                       QDDSSRC     QPFR                        1
                                                                       &TRCJOBS    QTEMP                       8
```

Figure 11-5. ANZPGM Program-to-File Cross-Reference Report

The ANZPGM Program-to-File Cross-Reference Report shows the following columns:

Library and Program
The name of the program that uses the file shown.

Program Text Description
The program's text description, if it was provided at program creation.

Object and Library
The name of the object that the program refers to, and the name of the library the object is in.

Record Format
The name of the formats in the file used by the program in the file being referred to.

File Usage
The manner in which the file is used by the program (1—input, 2—output, 4—update, 8—unknown, or any of the OR'd combinations of these, such as 3—input-output, 6—output-update). See the Display Database Relations (DSPDBR) command and the Display Program References (DSPPGMREF) command in the *CL Reference* manual for additional information about other values you can use in this field.

Figure 11-6 shows an example of the ANZPGM File-to-Program Cross-Reference Report.

12/01/87 13:37:15 File to Program Cross Reference Page 1

Library	Object	Record Format	Library	Program	Program Text Description	File Usage 1=In 2=Out 4=Upd 8=?
			QPFR	QCPBB3		
				QCPBB4		
				QCPCTL		
				QCPDRIVE		
				QCPEVAL		
				QCPGSP		
				QCPINIT		
				QCPMDL		
				QCPMGR		
				QCPSETUP		
				QCPUSER		
				QCPXTRCT		
				QMNADDTO		
				QMNGOMNU		
				QMNMAIN		
				QPTBATCH		
				QPTCPTRP		
				QPTCPTSL		
				QPTCPTWK		
				QPTLCKQ		
				QPTPGMX2		
				QPTSAME		
				QPTSAMEC		
				QPTSAMS		
				QPTSAMSC		
				QPTSLECT		
				QPTSYSRP		
				QPTSYSSL		
				QPTSYSWK		
				QPTTRCJ1	STRJOBTRC CPP	2
				QPTTRIDX		
				QPTTST1		
			QPFRTEMP	QPTSAMX	Sampled Address Monitor test program	
				QPTSYSRP		
	*FILE		QPFR	OLDPTTNSRP		8
	*FILE			QPTTNSRB		8
	*FILE			QPTTNSRP		8
	*FILE		QPFRTEMP			8
	*NONE		QPFR	QPTTRCJ0	ENDJOBTRC CPP	
	*NONE			QPTTRCJ1	STRJOBTRC CPP	
	QAJOBTRC			QPTTRCRP	ENDJOBTRC CPP	8
	QAPMDMPT			QPTCHGJT		8

Figure 11-6. ANZPGM File-to-Program Cross-Reference Report

The ANZPGM File-to-Program Cross-Reference Report shows the following columns:

Library and Object
The name and library the file is in.

Record Format
The names of the record formats in the file.

Library and Program
The names and library of the programs that use the file.

Program Text Description
The program text description.

File Usage
The manner in which the file is used (1—input, 2—output, 4—update, 8—unknown, and OR'd combinations of these values).

Analyze Database File (ANZDBF) Command

Use the ANZDBF command to print reports detailing physical and logical file relationships.

For information on how to enter the ANZDBF command, see the *CL Reference* manual.

When you use the ANZDBF command, the following files are created as output:

File	**Description**
QPPTANZD	The print file that has the physical-to-logical database file relationships report (Database Relation Cross-Reference).
QPPTANZD	The print file that has the logical-to-physical database file relationships report (Logical File Listing).
QAPTAZDR	The database file that serves as input to the ANZDBFKEY command.

Figure 11-7 shows an example of the ANZDBF Database Relation Cross-Reference Report.

```
12/01/87 14:29:31      Database Relation Cross Reference               Page     1
Type                                                          Depncy
P=Phy                              Depnd  Dependent   Dependent   Type
L=Lgl  File        Library     Count  File        Library     D/A
-----  ----------  ----------  -----  ----------  ----------  ------
  L    QAOFCALL    QOFCFLS
       QAOFCL
       QAOFDL
       QAOFDS1L
       QAOFDS2L
       QAOFDS3L
       QAOFDS4L
       QAOFDS5L
       QAOFDS6L
       QAOFENRL
       QAOFGMTL
       QAOFMAIL
       QAOFOAIL
       QAOFOAUL
       QAOFOGGL
       QAOFOGL
       QAOFOTML
       QAOFOWNL
  P    QAOFASGP
       QAOFBSTP
       QAOFCP                      2  QAOFCALL    QOFCFLS        D
                                   2  QAOFCL      QOFCFLS        D
       QAOFDDP                     1  QAOFOWNL    QOFCFLS        D
       QAOFDHP
       QAOFDOCP
       QAOFDP                      7  QAOFDL      QOFCFLS        D
                                   7  QAOFDS1L    QOFCFLS        D
                                   7  QAOFDS2L    QOFCFLS        D
```

Figure 11-7. ANZDBF Database Relation Cross-Reference Report

The ANZDBF Database Relation Cross-Reference Report has the following columns:

Type The file type (P-Physical, L-Logical).

File The name of the file.

Library The library containing the file.

Depnd Count
The number of logical files dependent on this file.

Dependent File
The names of each dependent logical file.

Dependent Library
The library the dependent logical files are in.

Depncy Type D/A
D—Data share dependency. A—Access share dependency.

The entries in the *Type*, *File*, and *Library* columns are left blank if they are the same as the previous line.

Figure 11-8 shows an example of the ANZDBF Logical File Report.

```
12/01/87 14:29:34                 Logical File Listing                         Page    1
                              Depncy                              Type
Dependent     Dependent       Type                                P=Phy
File          Library         D/A      File         Library       L=Lgl
----------    ----------      ------   ----------   ----------    -----
QAOFCALL      QOFCFLS           D      QAOFCP                       P
QAOFDL        QOFCFLS           D      QAOFDP
QAOFENRL      QOFCFLS           D      QAOFENRP                     P
QAOFGMTL      QOFCFLS           D      QAOFGP                       P
QAOFOAUL      QOFCFLS           D      QAOFOAUP                     P
QAOFOGGL      QOFCFLS           D      QAOFOGP                      P
QAOFMAIL      QOFCFLS           D      QAOFRMLP                     P
            7 records processed
```

Figure 11-8. ANZDBF Logical File Report

The ANZDBF Logical File Report shows the following:

Dependent File
The names of each dependent logical file.

Dependent Library
The library the dependent logical files are in.

Depncy Type D/A
D—Data share dependency. A—Access share dependency.

File The name of the physical file.

Library The library containing the physical file.

Type The physical file type.

Analyze Database File Keys (ANZDBFKEY) Command

Use the ANZDBFKEY command to print a report showing the key structure of logical files.

For information on how to enter the ANZDBFKEY command, see the *CL Reference* manual.

When you use the ANZDBFKEY command, the following input file is used:

File	Description
QAPTAZDR	Database file that is the output from the ANZDBF command.

Note: Because the ANZDBFKEY command uses the output from the ANZDBF command as its input, be sure the ANZDBF command is finished before you use the ANZDBFKEY command. The ANZDBFKEY command tests the existence of the ANZDBF output file and, if the file does not exist, the program ends.

When you use the ANZDBFKEY command, the following files are created as output:

File	Description
QPPTANZK	Print file for the access path and record selection report (Key Fields and Select/Omit Listing).
QPPTANKM	Print file for the logical file key report (Analysis of Keys for Database Files).

The information provided in these reports may suggest ways of combining logical files, for physical files with a number of logical files over them. This process of combining reduces the total number of logical files the system must maintain.

For example, consider an application that uses these two logical views of the same physical file:

- Logical file FILEA with key FIELD1
- Logical file FILEB with keys FIELD1 and FIELD2

In this case, it is likely that you could delete FILEA and use FILEB instead.

Reducing the number of logical views an application uses can help the performance of the application and of the system.

Figure 11-9 gives an example of the ANZDBFKEY Key Fields and Select/Omit Listing.

This report lists the access path and selection (logical files only) values based on the output produced by the Display File Description (DSPFD) command with a single line for each key field or selection rule.

```
12/01/87 14:35:02      Key Fields and Select/Omit Listing             Page    1
    File       Library         Order  Path Type Unique Maintenance
PHY QAOFCP     QOFCFLS         FIFO   KEYED       N    *IMMED
        Based on          Format      Key Field  Seq    Sign  Zone  Alt
                                       NAME
                                       JDATE            SIGN
                                       STIME            SIGN
                                       SEQ              SIGN
                                       EXT              SIGN
                                       GMTGNO           SIGN
    File       Library         Order  Path Type Unique Maintenance
LGL QAOFCALL   QOFCFLS         FIFO   KEYED       N    *IMMED
        Based on          Format      Key Field  Seq    Sign  Zone  Alt
    QAOFCP     QOFCFLS    CALRC1       MJDATE           SIGN
                                       MTIME            SIGN
                                       NAME
    **Record Selection**  Format      Field      S/O Comp Values
                          CALRC1      MJDATE      S   GT   +0
                                      MTIME       A   GT   +0
                                      EXT         A   LE   +2
                                                  O   AL
```

Figure 11-9. ANZDBFKEY Key Fields and Select/Omit Listing

In the ANZDBFKEY Key Fields and Select/Omit Listing Report, the first output line shows the following:

File The file name and, to the left of the name, the file type—physical (PHY) or logical (LGL).

Library The name of the library in which the file is contained.

Order Ascending or descending sequence for the keys (LIFO, FIFO).

Path Type The type of access path (ARRIVAL, KEYED, or SHARED).

Unique Whether unique keys are used (Y or N).

Maintenance
*IMMED, *RBLD, or *DLY.

The second output line shows the following:

Based On The physical file name.

Format The format name in the logical file.

Key Field The name of the key field (can be one or more lines).

Seq The key sequence (blank is ascending, DES is descending).

Sign The key sign (blank, SIGN, or ABSV).

Zone The zone/digit specified (blank, ZONE, or DIGIT).

Alt The alternative collating sequence (YES or blank).

If record selection is used, the third output line shows the following:

Format The logical file format name.

Field The select/omit field name.

S/O Whether to select (S) or omit (O).

Comp The compare relation EQ, GT, LT, AL (all),....

Values The values to compare against.

Printer File QPPTANKM lists the file names, and for logical files, the key fields for each format in descending order from major key to minor key.

You can use this list to find ways to combine logical files, when physical files have many logical files over them. By combining files, you can reduce the number of logical views an application requires, and the total number of logical files the system must maintain. Having fewer files to maintain can improve the performance of the application and of the system.

Figure 11-10 shows an example of the ANZDBFKEY Analysis of Keys for Database Files Report.

```
12/01/87 14:35:02                        Analysis of Keys for Database Files                                   Page 1
   Physical File QAOFCP      Library QOFCFLS
                        Logical  Maint                                                                        No.
File        Library     Format        ******* Key Fields Major to Minor *******                               Keys S/O
QAOFCP      QOFCFLS                 I NAME      JDATE      STIME      SEQ        EXT        GMTGNO             6
QAOFCALL    QOFCFLS     CALRC1      I MJDATE    MTIME      NAME                                                3  YES
QAOFCL      QOFCFLS     MTGREC      I GMTGNO    NAME       JDATE      STIME      SEQ        EXT                6
```

Figure 11-10. ANZDBFKEY Analysis of Keys for Database Files Report

The columns in the ANZDBFKEY Analysis of Keys for Database Files Report are as follows:

Physical File
The name of the physical file.

Library The physical file library.

File The logical files over the physical file.

Library The library the file is in.

Logical Format
The logical file format name.

Maint Maintenance. Specify I (immediate), R (rebuild), or D (delay).

Key Fields Major to Minor
Up to seven key fields.

No. Keys The number of key fields in the file.

S/O Whether select/omit is specified for key. YES indicates it is specified.

Analyzing Process Information

Use the Display and Analyze Access Group (DSPACCGRP and ANZACCGRP) commands to examine and analyze the process access group (PAG).

Use DSPACCGRP to report on PAG data for selected jobs. Use ANZACCGRP to further analyze the DSPACCGRP output.

Process access group analysis provides you with a view of the operational environment for all jobs, or a group of jobs, in the system at a given time. Use the information from process analysis to tune your system. When you tune your system, you improve the program environment, causing a reduction in the number of the following:

- Open files
- File buffer and work space sizes
- File open placement in a program
- Active programs

Display Access Group (DSPACCGRP) Command

Use the DSPACCGRP command to see the following:

- Size of a job's PAG
- Open files
- I/O count for all files
- Program data storage use
- Active program

For information on how to enter the DSPACCGRP command, see the *CL Reference* manual.

When you use the DSPACCGRP command, the following output files are created:

File	Description
QPPTPAGD	Print file
QAPTPAGD	Database file (input to the ANZACCGRP command)

The following DSPACCGRP display appears only when you specify OUTPUT(*).

```
MM/DD/YY                  Process Access Group Display                DSPACCGRP
HH:MM:SS                 'Title specified on the command'
   Job: job-name         User:  user id     Number: nnnnnn     Type: I
    PAG Size: nnnnn          Number of files: nn       I/O count:          nn

FILE         LIBRARY     MEMBER        SIZE    I/O COUNT
filename(s)  libname     member       value     value


PROGRAM NAME           STORAGE SIZE          STATUS
  pgmname(s)             value               FREE/ACTIVE






F3=Exit       F5=Refresh        PRINT=QPPTPAGD
```

When the PAG information is being displayed, the function key options to DSPACCGRP are as follows:

1. Press F3 (Exit) or the Enter key to end the program.
2. Press F5 (Refresh) to display the PAG again to show any changes that occurred since the current display was shown. If the job cannot be displayed because its structure is changing, or the job ended, a message appears and the command ends.
3. Press the Print key to print the displayed job's PAG data. The format is the same as the format used when you select OUTPUT(*PRINT).

The fields on the Process Access Group display are as follows:

Title The value you specified on the command to identify the output.

Job, User, and Number
The name for the job you display, your user ID, and the job number is shown.

Type The type of job (for example, I=interactive, B=batch, and so on).

PAG Size The total size of the job's PAG that is in use. This value includes the size needed for the files in the job, and the storage size used for variables (constants, arrays, parameter lists).

Number of files
A count of the number of files open in the job, such as database files, spool files, and display files.

I/O count A count of the number of I/O operations for all files currently open in the job.

For each open file, the following information is shown:

File The name of the file.

Library The library the file is in.

Member The member name for database files (includes spool files).

Size The size of the open data path for the file. This value varies depending on the number of record formats in the file and buffer size.

I/O count A count of all I/O operations issued to the file since the last full open.

The following information is shown for each program that has been called. It is possible that a program is no longer active in a job. To determine whether it is inactive, use option 11 on the DSPJOB command. Once a program is no longer active, it is removed from the DSPJOB command's output list.

Program Name
The name of the program.

Storage Size
The amount of storage used by the program for variables and arrays.

Status The ACTIVE or FREE status of the storage for the program.

Figure 11-11 shows an example of the DSPACCGR Process Access Group Information Report. The fields on this report are the same as those explained for the display.

```
12/01/87              Process Access Group Information               Page   1
15:01:59
   Job: BYSINN          User: VLLXR239      Number: 003368      Type: I
    PAG Size:   164,352
FILE        LIBRARY     MEMBER        SIZE    I/O COUNT
----------  ----------  ----------  --------  ---------
QDGENDSP    QSYS                      9216        68
QRZLHEL3    QSYS        QRZLHEL3      3072         1
QDPTPAGD    QPFR                      4096         3
          Number of files:  3     Total I/O count:          72

PROGRAM NAME         STORAGE SIZE         STATUS
------------         ------------         ------
  QCADRV                     80           ACTIVE
  QSCTJOB                    32           ACTIVE
  QSCTEVTH                   80           FREE
  QSPHNCMD                  112           ACTIVE
  QSCPTREC                   32           ACTIVE
  QSPSBMJB                   96           ACTIVE
  QWTCCJOB                   16           ACTIVE
  QSPSBMEX                   32           ACTIVE
```

Figure 11-11. Process Access Group Information

When you use the DSPACCGRP command, be aware of the following:

1. When a program returns, free its storage automatically (RPG LR frees the storage; any CL program return frees the program's storage) or use another program to free the storage explicitly by using an RPG FREE instruction, a CL Reclaim Resources instruction, or a COBOL CANCEL instruction. (This approach assumes that the application design does not use frequent repeated calls to the program.)

   Leave programs active only if you call them frequently (at least once every third transaction). This approach saves space, and reduces the number of PAG pages that are read into main storage but not referred to. (This approach assumes that the application design does not use frequent repeated calls to the program.)

2. Free space in the PAG is reused by any program if it is not followed by any active entries. If the free space has ACTIVE space after it, it is reused by a call to the same program. If there is ACTIVE space following FREE space, you have large holes of unused FREE space. These holes are undesirable because the wasted space in the PAG causes additional processing time during PAG I/O. You can remove these holes by calling those programs that remain active before calling other programs.

3. To reduce these holes further, include subroutines in the main program instead of calling an external program. This method might mean duplicating code in several programs, but it can save space and significantly reduce the time to activate the routine.

4. If the job you display is changing (such as calling and returning from programs or opening and closing files), DSPACCGRP may try again several times to get an accurate view of the PAG. DSPACCGRP tries five times before it stops, and displays a message telling you to try the command again.

5. The command shows up to 999 file names and then stops. It tells you that more files exist, but could not be displayed.

6. The I/O count (shown for each file) might give you some unexpected results. For example, if you use subfiles on a display, several PUTs of records to the subfile are included in the I/O count. This is in addition to the I/O operation against the subfile control record that causes the subfile to display.

7. PAG SIZE is the total in-use size of the PAG. Because there can be a number of inactive programs (programs that have returned), the total allocated size of the PAG can be larger than the size currently in use. For example, a job that was previously using interactive debug, but is now doing some other function, could have an allocated size much larger than the in-use size.

8. To reduce the amount of main storage needed for the job PAG, group files that are most frequently used. To do this, order the specification (or OPENs, if explicit open is used) of files based on their I/O count.

   DSPACCGRP shows the files in the order the system opens them (the RPG implicit OPEN opens the files in reverse order from their order in the program). It does not matter whether the most frequently used files are specified or explicitly opened first or last, as long as they are grouped together by frequency of use.

Analyze Process Access Group (ANZACCGRP) Command

Use the ANZACCGRP command to determine whether:

- Application programs use shared display and database files
- Files are ordered by their frequency of use
- Files are remaining open but have no activity
- Programs are freeing their static storage or keeping it active

All of these conditions can affect system performance, especially if the jobs have very large PAGs.

Consider the size of the PAG when you decide whether to use the PURGE(*NO) job attribute when running an application. An application that requires fast response time and has frequent and steady operator interaction would be a candidate for running in a PURGE(*NO) environment.

The availability of main storage to dedicate to this application is the critical factor in the decision to use PURGE(*NO). In general, the larger the PAG, the larger the amount of main storage that would need to be dedicated.

By using the DSPACCGRP command job selection, the ANZACCGRP command allows you to analyze the PAG of a single job, a set of jobs, or all jobs in the system. It is most useful in examining a large number of PAGs collected using the Job(*All), Job(*Int), or Job(generic name) options of the DSPACCGRP command.

For information on how to enter the ANZACCGRP command, see the *CL Reference* manual.

The ANZACCGRP command uses the QAPTPAGD database file (the output from the DSPACCGRP command) as input.

Note: Because the ANZACCGRP command uses the output from the DSPACCGRP command, be sure to run the DSPACCGRP command first with parameter OUTPUT(*FILE) or OUTPUT(*BOTH).

File	Description
QSYSPRT	Print file

The ANZACCGRP command produces a four-part report:

- Environment Summary
- Job Summary
- File Summary
- Program Summary

Figure 11-12 shows an example of the ANZACCGRP Environment Summary. This section shows information for four categories of job types.

```
                                       ANZACCGRP SUMMARY REPORT
                                            Sample Report
DATA FROM: QPFRDATA  /QAPTPAGD      MBR: QAPAGDTA
 DATE : 12/14/87   TIME : 09:58:57
 E N V I R O N M E N T    S U M M A R Y
----------------------------------------
```

JOB TYPE	NUMBER OF JOBS	AVERAGE # FILES/JOB	AVG # DUP FILES/JOB	AVG # DSP FILES/JOB	AVERAGE IO COUNT/JOB	AVERAGE PAG SIZE
INTERACTIVE	4	2		2	44	207872
BATCH/AUTOSTART	4				188	87808
READER/WRITER	1	3			35	126464
OTHER	6	5			75	135253
TOTAL/AVERAGE	15	3		2	94	141380

Figure 11-12. Environment Summary

The following columns appear on the Environment Summary:

JOB TYPE
The type of jobs the summary line is for (Interactive, Batch/Autostart, Reader/Writer, or Other).

NUMBER OF JOBS
The number of jobs of the corresponding type.

AVERAGE # FILES/JOB
The average number of open files.

AVG # DUP FILES/JOB
The average number of files open more than once for each job. If this number is high, it indicates that files are being opened with SHARE(*NO).

AVG # DSP FILES/JOB
The average number of open display files for each interactive job. If this number is high, it indicates that display files are defined SHARE(*NO) or that several display files are used by an application. If there are not a large number of display file formats that take a lot of PAG space, consider using SHARE(*YES) and combining display files.

AVERAGE I/O COUNT/JOB
The average number of file operations for all open files, not opened since the job started.

AVERAGE PAG SIZE
The average PAG size for each job type.

Figure 11-13 shows an example of the ANZACCGRP Job Summary. This section shows information for all selected jobs.

ANZACCGRP SUMMARY REPORT
Sample Report

DATA FROM: QPFRDATA /QAPTPAGD MBR: QAPAGDTA
DATE : 12/14/87 TIME : 09:58:57
J O B S U M M A R Y

JOB NAME	USER NAME	JOB NUMBER	TP	PAG SIZE	TOTAL FILES USED	DUP. FILE USED	DSP. FILE USED	TOTAL I/O COUNT	SIZE OF FILES	SIZE OF DUP FILES	SIZE OF DSP FILES	ACT PGM	MULT PGM CALL	TOTAL PSSA SIZE	DUP PSSA SIZE
SCPF	QSYS	000000	X	123392											
QSPL	QSYS	000958	M	78336								1		4672	
QSYSARB	QSYS	000952	S	284160	34			453	95744						
QLUS	QSYS	000953	S	82944											
QBASE	QSYS	000956	M	164352								2		4864	
PRINT46D	QSPLJOB	000959	W	126464	3		1	35	10240		4608	3		112	
WS03	CLERK1	000962	I	334848	3	2	3	65	59904	39936	59904	2		272	
WS02	TJONES	000963	I	184832	3		2	102	23552		20992	1		80	
WS05	CLERK2	000965	I	201216	1		1	9	16384		16384	2		272	
NET38	QSYS	000966	M	78336								1		4672	
PUTGET	QSECOFR	000968	B	99840	2		1	752	17408		3584	1		80	
ROUTER	QSECOFR	000969	B	78336								1		80	
MONITOR	QSECOFR	000970	B	94720								1		80	
RESTART	QSECOFR	000971	B	78336								1		80	
WS04	JSMITH	000976	I	110592	1		1	2	19968		19968				

Figure 11-13. ANZACCGRP Job Summary

The Job Summary contains the following columns:

JOB NAME/USER NAME/JOB NUMBER
The name of the job, name of the user, and job numbers.

TP The type of job. The various types of codes and descriptions follow:

Job Type Code	Job Type Description
A	Autostart
B	Batch
I	Interactive
M	Other
R	Reader
S	Other
W	Writer
X	Other

PAG SIZE The gross size of the PAG for the job.

TOTAL FILES USED
The number of all open files for the job (including display and spool files).

DUP. FILE USED
The number of files with more than one ODP in the PAG. If this number is not zero, it indicates that files are being opened SHARE(*NO). If possible, the files should be combined and created with SHARE(*YES).

DSP. FILE USED
The number of display files open for the job. If this number is high, it indicates that the display files are being opened SHARE(*NO) or that the application has multiple display files. If possible, display files should be combined and created with SHARE(*YES).

TOTAL I/O COUNT
: The count of file operations (including OPEN and CLOSE) for all open files in the job. This, however, is not the total I/O count for the job.

SIZE OF FILES
: The amount of space in the PAG that is occupied by ODPs for all open files.

SIZE OF DUP FILES
: The amount of space in the PAG that is occupied by ODPs for files opened with SHARE(*NO). The first ODP for each duplicate file is not included in this total. This value represents the amount of space you will save if you use SHARE(*YES).

SIZE OF DSP FILES
: The amount of space in the PAG that is occupied by ODPs for display files.

ACT PGM The number of programs in the program stack that are active.

MULT PGM CALL
: The number of active programs that appear more than once in the program stack.

TOTAL PSSA SIZE
: The amount of space in the PAG that is occupied by variables defined by the active programs in the program stack.

DUP PSSA SIZE
: The amount of space in the PAG that is occupied by variables defined by the active programs that appear more than once in the program stack. The first program static storage area (PSSA) entry for each duplicate program is not included in this total. The value represents the amount of space that could be saved if applications could be designed to avoid calling the program multiple times.

Figure 11-14 shows an example of the ANZACCGRP File Summary. This section shows information for all open files.

```
                         ANZACCGRP SUMMARY REPORT
                              Sample Report
DATA FROM: QPFRDATA  /QAPTPAGD     MBR: QAPAGDTA
 DATE : 12/14/87    TIME : 09:58:57
 F I L E    S U M M A R Y
-------------------------

                                      FILE NUMBER  NUMBER OF     TOTAL     AVERAGE
 FILE NAME   LIBRARY      MEMBER      TYPE OF JOBS TIMES OPEN  I/O COUNT  ODP SIZE
----------- ---------- ----------     ---- ------- ---------- ----------- --------
 QRZPHELM    QSYS       QRZPHELM       DB       1         1                  2560
 QRZPHVPD    QSYS       QRZPHVPD       DB       1         1              8   2560
 QRZPHDEL    QSYS       QRZPHDEL       DB       1         1                  2560
 QRZPHLOC    QSYS       QRZPHLOC       DB       1         1                  2560
 QRZPPTF     QSYS       QRZPPTF        DB       1         1                  2560
 QRZPLCG     QSYS       QRZPLCG        DB       1         1                  2560
 QRZPTCBL    QSYS       QRZPTCBL       DB       1         1                  2560
 QRZPTOEM    QSYS       QRZPTOEM       DB       1         1                  2560
 QRZPTSPC    QSYS       QRZPTSPC       DB       1         1                  2560
 QRZLHEL1    QSYS       QRZLHEL1       DB       1         1            135   3072
 QRZLHEL2    QSYS       QRZLHEL2       DB       1         1                  3072
 QRZLHEL3    QSYS       QRZLHEL3       DB       2         2            246   3072
 QRZLHEL4    QSYS       QRZLHEL4       DB       1         1              8   3072
 QRZLHEL5    QSYS       QRZLHEL5       DB       1         1                  3072
 QRZLHEL7    QSYS       QRZLHEL7       DB       1         1                  3072
 QRZLHVP1    QSYS       QRZLHVP1       DB       1         1             59   2560
 QRZLHVP2    QSYS       QRZLHVP2       DB       1         1                  3072
 QRZLLOC1    QSYS       QRZLLOC1       DB       1         1                  2560
 QRZLTEL1    QSYS       QRZLTEL1       DB       1         1                  3072
 QRZLTEL2    QSYS       QRZLTEL2       DB       1         1                  3072
 QRZLHEDR    QSYS       QRZLHEDR       DB       1         1                  3072
 QRZLDSUA    QSYS       QRZLDSUA       DB       1         1                  3072
 QRZLDSU1    QSYS       QRZLDSU1       DB       1         1                  3072
 QRZLPATH    QSYS       QRZLPATH       DB       1         1                  2560
 QRZLPTH1    QSYS       QRZLPTH1       DB       1         1                  3072
 QRZLLCG1    QSYS       QRZLLCG1       DB       1         1                  3072
 QRZLLCG2    QSYS       QRZLLCG2       DB       1         1                  2560
 QADBXREF    QSYS       QADBXREF       DB       1         1                  2560
 QADBFDEP    QSYS       QADBFDEP       DB       1         1                  2560
 QADBLDEP    QSYS       QADBLDEP       DB       1         1                  2560
 QADBLDNC    QSYS       QADBLDNC       DB       1         1                  3072
 QADBXDIC    QSYS       QADBXDIC       DB       1         1                  2560
 QADBXFIL    QSYS       QADBXFIL       DB       1         1                  3072
 QADBXOWN    QSYS       QADBXOWN       DB       1         1                  3072
 QPSPLPRT    QSYS                      DS       1         1             12   4608
 Q04079N001  QSPL       Q979578132     DB       1         1             20   2560
 QDSPUIM     QSYS                      DS       2         4*            67  19968
 QDGENDSP    QSYS                      DS       2         2             46  16384
 QDDSPEXT    QSYS                      DS       1         1              8   4608
 QAPTPAGD    QPFRDATA   QAPAGDTA       DB       1         1             57   2560
 NETFILE     QNET                      DS       1         1            745   3584
 NET38SPOOL  QNETSPOOL  VN6A633564     DB       1         1              7  13824
```

Figure 11-14. ANZACCGRP File Summary

The ANZACCGRP File Summary contains the following columns:

FILE NAME/LIBRARY/MEMBER
Name of the open file ODP.

FILE TYPE
One of the following:

DB Database or spooled file
DS Display or ICF file
SP Spooled file being created

NUMBER OF JOBS
The total number of jobs using the file.

NUMBER OF TIMES OPEN
The total number of times the file has been opened. If this number is larger than the number of jobs (indicated by an *), the file has been created with SHARE(*NO).

TOTAL I/O COUNT
The total number of file operations (including OPEN and CLOSE) that have been performed by all jobs for currently open files. If this number is very low, the applications should be changed to open the file only when needed.

AVERAGE ODP SIZE
The average amount of space in the PAG occupied by each open occurrence of the file. In general, this number represents the ODP size for the file.

Figure 11-15 shows an example of the ANZACCGRP Program Summary. This section shows information for all active use programs. Only programs with a status of active are included. A status of free indicates that the program has returned (RPG with LR on, COBOL with a STOP RUN).

```
              ANZACCGRP SUMMARY REPORT
                   Sample Report
DATA FROM: QPFRDATA  /QAPTPAGD     MBR: QAPAGDTA
 DATE : 12/14/87   TIME : 09:58:57
 P R O G R A M    S U M M A R Y
------------------------------
                 NUMBER      NUMBER OF      AVERAGE
 PROGRAM NAME    OF JOBS    INVOCATIONS    PSSA SIZE
 ------------    -------    -----------    ---------
  QWTMEJIN          3            3            4672
  QMHUNMSG          3            3             192
  QCADRV            8            8              80
  QSPWTRM1          1            1              16
  QSPPRTWT          1            1              16
```

Figure 11-15. ANZACCGRP Program Summary

The ANZACCGRP Program Summary contains the following columns:

PROGRAM NAME
The name of the program.

NUMBER OF JOBS
The number of jobs using the program.

NUMBER OF INVOCATIONS

The total number of times the program appears. If this value is larger than the number of jobs, it indicates that the program has been called more than once by a job.

AVERAGE PSSA SIZE

The average size in the PAG of the variables defined in the program. In general, this number represents the PSSA size for the program.

Analyzing Disk Activity

When a disk unit is busy more than half the time, disk I/O time is lengthened and can adversely affect overall system performance. If one or two disk units are consistently more than 50% busy and other disk units are less busy, and if the high activity is due primarily to database file I/O, it may be possible to balance the use and improve performance by distributing the heavily used files among disk units.

Use the disk activity analysis commands to collect and print information about active objects and disk units, in order to identify those that are candidates for adjustment. This should be done if an imbalance, as described above, is evident in the Disk Utilization section of the report produced by the PRTSYSRPT command or the output of the WRKDSKSTS command.

Disk activity data collection uses one of the licensed internal code traces (the main storage management trace) controlled by the Trace Internal (TRCINT) command. The consequences of this are:

- Data collection will not start if any vertical licensed internal code trace is currently running (such as that started by the STRPFRMON command).
- Due to the volume of trace data recorded, limit collection time to no more than 5 minutes (1 minute should be adequate).
- The overhead involved while tracing is active can cause significant performance degradation.
- The initial phase of ending collection causes high machine pool paging and possible throughput or response time degradation for all jobs in the system.

To collect and analyze disk activity data, perform these steps:

1. Run the Start Disk Data Collection (STRDSKCOL) command to start main storage management tracing.
2. Wait a few minutes (1 to 5 minutes) while the data is collected. You should not run infrequently performed functions that influence disk activity, such as program development or batch jobs, during this time.
3. Run the End Disk Data Collection (ENDDSKCOL) command to end tracing and prepare the collected data for analysis.
4. Run the Print Disk Report (PRTDSKRPT) command to produce a report from the collected data. You may combine this step with the previous step, since ENDDSKCOL can produce the report directly.
5. Analyze the report for problems. Following are some common problems and their corresponding appropriate corrective actions.

If one file is particularly active and only a few records are in use, the application should be investigated. Consider using a data area or data queue in place of the records in the database file, or increase or remove the FRCRATIO on the file using the Change Physical File (CHGPF) command.

If one file is very active and many records are in use, distributing the file more evenly over several disk units (as described below) may help performance.

If multiple files on a particular disk unit consistently have high activity, moving one or more of the files (as described below) to another unit may reduce the activity on the busy unit.

To move a physical file from one disk unit to another, or to distribute it among multiple disk units, follow these steps:

1. Rename the file using the Rename Object (RNMOBJ) command.
2. Create a new version of the file, using the Create Physical File (CRTPF) command, with the parameters described below.
3. Copy the data from the original file to the new file, using the Copy File (CPYF) command.
4. Delete and recreate any logical files that are over the physical file, using the Delete File (DLTF) and Create Logical File (CRTLF) commands.
5. Delete the original (renamed) file using the DLTF command.

   When disk space is limited, or if a current backup of the file exists, the following steps can be used to perform the same function:

   a. Save the file to tape or diskette using the Save Object (SAVOBJ) command (if not already done).
   b. Delete the physical file and any logical files.
   c. Restore the file using the Restore Object (RSTOBJ) command.
   d. Recreate any logical files.

The placement of a database file can be controlled through the use of CRTPF command's SIZE, ALLOCATE, and UNIT parameters. The effects of these parameters are explained below.

To place a file on a particular disk unit, specify the parameter UNIT(nn). The system will allocate the file to the specified unit, regardless of the SIZE and ALLOCATE parameters, if the unit has adequate space available. This applies to the initial allocation and to the extents. After the CRTPF command completes, check the messages to see if the specific UNIT allocation was successful.

To distribute a file over several disk units, specify the parameters SIZE, ALLOCATE(*YES), and UNIT(*ANY). The SIZE value has three parts: the initial number of records, the number of records in any additional extents, and the number of extents. Specify an initial number of records and the number of records in an extent so that the number of records multiplied by one more than the record length is less than 32KB (32 768). If this figure is not less than 32KB, the file allocation will be to the unit with the largest amount of free space.

Use caution when following this plan to distribute data among several units. Files that have a lot of records inserted into them on a regular basis are poor candidates for a small initial allocation and small extent size. This is because of the additional overhead necessary to get an additional extent every time a few more records are inserted.

Note: When changing the allocation of a file, you may adversely affect the time for a save/restore of the object, especially if you are reallocating the file to a smaller extent size.

Start Disk Data Collection (STRDSKCOL) Command

Use the STRDSKCOL command to start collecting main storage management trace data for later analysis of disk activity. The command fails if a mutually exclusive trace is already running (such as from a prior uninterrupted use of the STRDSKCOL or STRPFRMON command). If the command is successful, a message is displayed showing the date and time that tracing was started.

The main storage management trace quickly accumulates large amounts of data, and also significantly affects the performance of the system. Because of this, collection time should be kept to a minimum.

For information on how to enter the STRDSKCOL command, see the *CL Reference* manual.

End Disk Data Collection (ENDDSKCOL) Command

Use the ENDDSKCOL command, no more than 1 to 5 minutes after using the STRDSKCOL command, to end collection of main storage management trace data. If the command is successful, a message is displayed showing the date and time, tracing is stopped, and the accumulated data is prepared for later analysis using the PRTDSKRPT command.

If the command fails (such as when a library that is not valid is specified for storage of the data), or if the parameter DSPDSKDTA (*NO) is used, data collection is stopped but the data isn't saved. In this case, you may run the ENDDSKCOL command again with corrected parameters (before running any other type of trace); a warning message (tracing not active) will be displayed, but the previously collected data will be saved.

For information on how to enter the ENDDSKCOL command, see the *CL Reference* manual.

The ENDDSKCOL command generally takes several minutes to run, and is not recommended for interactive processing (don't use JOBD (*NONE)).

The ENDDSKCOL command writes the collected data to the following file:

File	Description
QAPTDSKD	Database file (input to the PRTDSKRPT command)

Print Disk Activity Report (PRTDSKRPT) Command

Use the PRTDSKRPT command to print a two-part report from the storage management trace data collected by the STRDSKCOL and ENDDSKCOL commands. The first part of the report, Object Activity, shows I/O requests for selected disk units and objects. The second part of the report, Disk Unit Activity, shows I/O requests for all disk units.

Objects and disk units can be selectively included in the report. It is recommended that you use the PRTSYSRPT or WRKDSKSTS results to select which to include, in order to reduce the volume of the report.

For information on how to enter the PRTDSKRPT command, see the *CL Reference* manual.

The PRTDSKRPT command uses the following files:

File	Description
QAPTDSKD	Database file (output from the ENDDSKCOL command)
QPPTDSK	Printer file (report created)

Figure 11-16 shows an example Disk Activity Report. Blank lines have been removed to conserve space.

```
Run:  3/01/88 10:56:39             ----------Title----------        Page    1

Data File:    QPFRDATA/QAPTDSKD QADSKDTA ----------------text-----------------
Units:        *ALL or list of up to 10 unit numbers
Objects:      *ALL or list of up to 10 lib/obj mbr (multiple lines possible)
Object Class: *DB, *USER, or *ALL

Object Activity Detail  -or-  Object Activity Summary
                                                             Record   Object
Unit    Rqs    Length  Object     Library    Member          Number    Type
----    ----   ------- ---------- ---------- ----------     -------   -------
  7      10       16   LWY_SOURCE LWY        VARLEN_MT                 FILE
  7      17      227   Q04079N003 QSPL       Q771766258                FILE

  8       1        1   QAOSDP08   QUSRSYS    QAOSDP08                  FILE
  8       6        9   LWY_SOURCE LWY        VARLEN_1                  FILE

  9       1        1   QAOSDP04   QUSRSYS    QAOSDP04                  FILE
  9       4        4   LWY_SOURCE LWY        VARLEN_1                  FILE

 10      22       28   LWY_SOURCE LWY        VARLEN_1                  FILE
 10      21      182   Q04079N003 QSPL       Q771766258                FILE

 11       4        4   QAOFODP    QUSRSYS    QAOFODP                   FILE
 11       6        6   QHST88061A QSYS       QHST88061A                FILE

 12     129      243   QAOFCP     QUSRSYS    QAOFCP                    FILE
 12       6       15   QAOFENRP   QUSRSYS    QAOFENRP                  FILE

 13      15      173   LWY_SOURCE LWY        VARLEN_1                  FILE
 13       1        2   NET38PUNCH QNETSPOOL  MEMBER2206                FILE
```

Figure 11-16 (Part 1 of 2). Disk Activity Report

```
14      9        27  QAOSDP01   QUSRSYS    QAOSDP01            FILE
14      3         4  QAOFCP     QUSRSYS    QAOFCP              FILE
14      5        46  LWY_SOURCE LWY        VARLEN_1            FILE

17     10        70  LWY_SOURCE LWY        VARLEN_1            FILE

18      8        39  LWY_SOURCE LWY        VARLEN_1            FILE
18      4        29  LWY_SOURCE LWY        VARLEN_MT           FILE
18      5         6  NET38PUNCH QNETSPOOL  MEMBER2206          FILE
18     20        31  QHST88061A QSYS       QHST88061A          FILE

19     12        21  QAOFCP     QUSRSYS    QAOFCP              FILE
19      8        48  LWY_SOURCE LWY        VARLEN_1            FILE

20      6         6  QAOSDP01   QUSRSYS    QAOSDP01            FILE
20     13        64  LWY_SOURCE LWY        VARLEN_1            FILE
20      1         1  NET38PUNCH QNETSPOOL  MEMBER2206          FILE
20      3         3  QHST88061A QSYS       QHST88061A          FILE

21      9        12  QAOSDP04   QUSRSYS    QAOSDP04            FILE
21      7        47  LWY_SOURCE LWY        VARLEN_1            FILE
21      5        10  LWY_SOURCE LWY        VARLEN_MT           FILE

22     28        28  Q04079N003 QSPL       Q771766258          FILE
22      1         1  QAOSDP01   QUSRSYS    QAOSDP01            FILE
22      1         1  QTRPLS     QPATOOL3   JUN03               FILE
22      4         3  QTRPLS     QPATOOL3   AUG13               FILE
22      4         4  LWY_SOURCE LWY        VARLEN_MT           FILE
22     28       510  Q04079N003 QSPL       Q771766258          FILE

Run:  3/01/88 10:56:39          ----------Title----------        Page    2

Data File:    QPFRDATA/QAPTDSKD QADSKDTA ----------------text-----------------

                          Disk Unit Activity

       Number of     Percent of    Number of    Percent of
Unit  I/O Requests  I/O Requests     Pages        Pages
----  ------------  ------------   ---------    ----------
  1          14         0.6              0         0.0
  7         158         6.8          2,677         3.4
  8         131         5.6          2,228         2.8
  9         157         6.8          3,242         4.1
 10         143         6.2          1,698         2.1
 11         120         5.2          1,921         2.4
 12         322        14.0          2,585         3.2
 13         120         5.2          5,662         7.2
 14         181         7.8         39,926        50.8
 17          84         3.6            843         1.0
 18         151         6.5          1,104         1.4
 19         146         6.3          2,549         3.2
 20         131         5.6          1,068         1.3
 21         196         8.5          2,195         2.7
 22         246        10.6         10,833        13.7

          2,300                     78,531
```

Figure 11-16 (Part 2 of 2). Disk Activity Report

The heading printed on all pages of the Disk Activity Report includes:

Run The date and time when the storage management trace data collection was ended (through the ENDDSKCOL command).

Title The title specified on the TITLE parameter.

Data File The file where the data is stored: library, file, member, and text. It is selected by the MBR and LIB parameters.

The heading printed on the first part of the report, Object Activity, also shows the values used to select items for inclusion in the section.

Units *ALL or a list of the disk units selected.

Objects *ALL or a list of the objects (library, object, member) selected. This list may extend over several lines.

Object Class
Class of objects selected, including the following:

***DB** Database physical files

***USER** Some common OS/400 objects as well as database physical files

***ALL** All objects

The first part of the report, Object Activity, shows I/O requests sorted by disk unit and object. For a summary report (Object Activity Summary), all activity for the same object on a disk unit is combined into a single line. For a detail report (Object Activity Detail), the I/O requests are listed separately.

Unit The disk arm identifier (unit number).

Rqs The total number of I/O requests for the object on this unit. For a summary report, this is the number of unique requests that was issued for records from the database file member described in the remainder of the line.

Length The maximum possible number of 512-byte pages for the I/O request. This length may be significantly different from the actual amount of data transferred into or out of main storage. This is especially true for system objects such as Process Access Groups (PAGs), the Work Control Block Table (WCBT), and the QSYSOPR, QCONSOLE, and QHST message queues. This value may represent the maximum length request or the allocated (not the in-use) size of the object.

Object The name of the object.

Library For user objects and some system objects, the library where the object is located. For some system objects this field is undefined and should be ignored.

Member The database file member name.

Record Number
For a summary report, this field is blank. Otherwise, for database file members, it is the relative record number of the last logical record in the data block represented by the trace record.

Object Type OS/400 object type. See the *CL Reference* manual for a list of most of the possible object types.

The second part of the report, Activity by Disk Unit, summarizes activity by disk unit. The report includes one line for each disk unit that showed activity during the trace period.

Unit
The disk arm identifier (unit number).

Number of I/O Requests
The total number of I/O requests to the unit during the trace period. The number of requests will be larger than the unit's individual sums in the first part of the report unless all objects are included. The last entry in this column is a total of the preceding entries.

Percent of I/O Requests
The unit's I/O requests as a percentage of all I/O requests during the trace period.

Number of Pages
The maximum possible number of 512-byte pages that could have been transferred to and from the unit according to the selection values specified. This length may be significantly different from the actual amount of data transferred into or out of main storage, especially for PAGs, the WCBT, and some message queues such as QHST, QCONSOLE, and QSYSOPR. The last entry in this column is a total of the preceding entries.

Percent of Pages
The maximum possible number of pages that could have been transferred to/from this unit as a percentage of all pages that could have been transferred to and from all units during the trace period.

Chapter 12. Managing the Performance Tools Configuration

If you choose option 6 (Configure and manage tools) on the IBM Performance Tools/400 menu, the Configure and Manage Tools display appears.

```
                        Configure and Manage Tools

 Select one of the following:

      1. Work with functional areas
      2. Delete performance data
      3. Copy performance data
      4. Convert performance data

Selection or command
===> ________________________________________________________________

F3=Exit   F4=Prompt   F9=Retrieve   F12=Cancel
```

From this display you can manage or change the objects used in the performance tools.

Work with Functional Areas

Functional areas provide a way to define and save selection values that you use on the System and Component Reports. For example, you might save a set of jobs or users as a functional area. Then, each time you use the Print System Report (PRTSYSRPT) and Print Component Report (PRTCPTRPT) commands, you specify the name of the functional area to use. Functional areas also work with the Print Job Report (PRTJOBRPT), Print Pool Report (PRTPOLRPT), and Display Performance Graph (DSPPFRGPH) commands. Specify these names on the Select Functional Areas (SLTFCNARA) and the Omit Functional Areas (OMTFCNARA) parameters.

If you choose option 1 (Work with functional areas) on the Configure and Manage Tools display, the Work with Functional Areas display appears.

```
                         Work with Functional Areas

 Library . . . . . : QPFRDATA

 Type option, press Enter.
   1=Create       2=Change       3=Copy       4=Delete
 Option   Functional Area          Text
   _      _________________        ________________________
   _      PAYROLL                  PAYROLL DEPARTMENT
   _      SHIPPING                 SHIPPING DEPARTMENT









 F3=Exit     F12=Cancel
```

This display shows the functional areas that exist in the library you specified. To create a new functional area, type the name and description on the first line under the *Functional Area* and *Text* columns, and press the Enter key. To select an existing functional area, type a 2 (Change), 3 (Copy), or a 4 (Delete) in the *Option* column next to the functional area of your choice.

If you choose to create or change a functional area, the Work with Functional Area display appears.

```
                         Work with Functional Area

 Functional Area . . . . . :    PAYROLL

 Type new/changed values, press Enter.
   4=Delete

 Option   Job Name   User ID          Option   Job Name   User ID
   _      W505       SUE                _      ________   ________
   _      W500       TOM                _      ________   ________
   _      W502                          _      ________   ________
   _                 JOHN               _      ________   ________
   _      WS*_____   ________           _      ________   ________
   _      ________   ________           _      ________   ________
   _      ________   ________           _      ________   ________
   _      ________   ________           _      ________   ________
   _      ________   ________           _      ________   ________




 F3=Exit     F12=Cancel
```

On this display you specify the job name and the user ID (or both) you want to include in the functional area. If you choose, you can specify only a job name, only a user ID, a generic job name (of the form yyy*), or a generic user name.

Thus, WS* would include all jobs that have a job name starting with WS and any user ID name. A blank is taken as any value.

If you want to delete a job name or a user ID, type a 4 (Delete) next to the appropriate name or ID and press the Enter key. To add new job names or user IDs, type them at the end of the list, and press the Enter key.

If you were to produce a System Report using the PRTSYSRPT command with the SLTFCNARA (PAYROLL) keyword, the report would show data pertaining to the jobs and users shown on the Work with Functional Area display.

Delete Performance Data

Use option 2 (Delete performance data) on the Configure and Manage Tools display to delete performance data that you no longer need on your system. When you choose option 2, the Delete Performance Data display appears.

```
                          Delete Performance Data

Library . . . . . .   QPFRDATA__

Type option, press Enter.
  4=Delete
Option   Member       Text                                   Date       Time
  _      XYZ                                                 12/15/87   14:05:55
  _      PERFTESTC4   2 hours w/ 5 minute intervals          12/15/87   08:05:48
  _      PERFTESTC3   Duration of 2 hours                    12/14/87   09:21:44
  _      PERFTESTC2                                          12/11/87   14:42:46

                                                                         Bottom
F3=Exit   F5=Refresh   F12=Cancel   F15=Sort by member   F16=Sort by text
```

The members that appear on this display are those used on the Start Performance Monitor (STRPFRMON) command for the keyword MBR when data was collected. To delete a member from this list, type a 4 (Delete) next to the appropriate member and press the Enter key. The member you delete is deleted from the following data collection files:

QAPGSUMD	QAPMECL	QAPMSBSD	QTRTSUM
QAPMASYN	QAPMETH	QAPMSYS	
QAPMBSC	QAPMHDLC	QAPMTSK	
QAPMBUS	QAPMIOBS	QAPMX25	
QAPMCIOP	QAPMJOBS	QAPTLCKD	
QAPMCONF	QAPMLIOP	QTRIDX	
QAPMDIOP	QAPMMIOP	QTRJOBT	
QAPMDISK	QAPMPOOL	QTRJSUM	
QAPMDMPT	QAPMRESP	QTRSLWT	

Copy Performance Data

Use option 3 (Copy performance data) on the Configure and Manage tools display to make copies of performance data members. When you choose option 3, the Select Performance Member display appears.

```
                          Select Performance Member

 Library . . . . . .   QPFRDATA__

 Type option, press Enter.
   1=Select

 Option   Member       Text                               Date       Time
   1      MONDAY       Performance Data for Monday        10/25/90   16:25:39
   1      TUESDAY      Performance Data for Tuesday       10/26/90   13:55:08
   1      WEDNESDAY    Performance Data for Wednesday     10/27/90   13:50:15
   _      THURSDAY     Performance Data for Thursday      10/28/90   12:00:34
   _      FRIDAY       Performance Data for Friday        10/29/90   10:05:46
   _      TESTRUN      Test run of system                 10/19/90   20:31:42
   _      Q892910958                                      10/18/90   09:58:45
   _      Q892902009                                      10/17/90   20:09:23




 F3=Exit    F12=Cancel    F15=Sort by member    F16=Sort by text
 F19=Sort by date/time
                                     (C) COPYRIGHT IBM CORP. 1981, 1990.
```

The members that appear on this display are those used on the Start Performance Monitor (STRPFRMON) command for the keyword MBR when data was collected. To copy a member or members from the list, type a 1 (Select) next to the appropriate member(s) and press the Enter key. The Copy Performance Data display appears.

```
                          Copy Performance Data Member

 Type choices, press Enter.

 --------Copy From--------    --------Copy To--------
 Member         Library       Member         Library
 MONDAY         QPFRDATA      MONDAY         NEWLIB
 TUESDAY        QPFRDATA      TUESDAY        NEWLIB
 WEDNESDAY      QPFRDATA      WEDNESDAY      NEWLIB
                                                                Bottom




 F3=Exit    F12=Cancel
```

This display shows you the members you selected to copy and where they are to be copied to. For each member listed, type the name of the new member and the library that contains it in the *Copy To* entries of the screen, and then press the Enter key. When the copy completes, you have exact copies of the old performance members in the new performance members for the following files:

QAPMASYN	QAPMDMPT	QAPMMIOP	QAPMX25
QAPMBSC	QAPMECL	QAPMPOOL	
QAPMBUS	QAPMETH	QAPMRESP	
QAPMCIOP	QAPMHDLC	QAPMSBSD	
QAPMDIOP	QAPMIOBS	QAPMSYS	
QAPMDISK	QAPMLIOP	QAPMTSK	

Convert Performance Data (CVTPFRDTA) Command

Use option 4 (Convert performance data) on the Configure and Manage Tools display to convert performance data collected prior to Version 2 Release 1. The data is converted to the file formats needed by Version 2 Release 1 Performance Tools/400.

```
                    Configure and Manage Tools

 Select one of the following:

      1. Work with functional areas
      2. Delete performance data
      3. Copy performance data
      4. Convert performance data










Selection or command
===> 4

F3=Exit   F4=Prompt   F9=Retrieve   F12=Cancel
```

When you select option 4, the Convert Performance Data (CVTPFRDTA) display appears. You can also use the CVTPFRDTA command to select the CVTPFRDTA display.

```
                    Convert Performance Data (CVTPFRDTA)

 Type choices, press Enter.

 From library . . . . . . . . . .   __________    Name
 To library . . . . . . . . . . .   __________    Name
 Job Description  . . . . . . . .   *USRPRF___    Name, *USRPRF, *NONE
   Library  . . . . . . . . . . .   __________    Name, *LIBL, *CURLIB









                                                                  Bottom
 F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
 F24=More keys
```

The Convert Performance Data (CVTPFRDTA) command converts down level performance data to the formats needed to be processed by the current release of the performance measurement/analysis tools. First, the release level on which the data was collected is determined. Then, all members of all files that need conversion are converted to the appropriate format.

The following files must be present for the conversion to take place:

QAPMCIOP	QAPMDISK	QAPMPOOL
QAPMCONF	QAPMJOBS	QAPMRESP
QAPMDIOP	QAPMLIOP	QAPMSYS

The following files are copied, or converted if necessary, if they are present:

QACPCNFG	QAPMBSC	QAPMHDLC
QACPPROF	QAPMBUS	QAPMMIOP
QACPRESP	QAPMDMPT	QAPMX25
QAITMON	QAPMECL	QAPTAPGP
QAPMASYN	QAPMETH	

The conversion can be done in the library in which the current data resides, or in a different library. If the conversion is done in the same library, the current data is replaced by the new data. If the conversion is done in a different library, the new data exists in the new library while the current data continues to exist in the current library.

Note: If a different library is specified for the new data, those files in the current library that do not need conversion are copied to the new library.

To convert performance data collected prior to Release 3.0, complete the following items on the display.

From library
Specifies the library that contains the data being converted.

To library
Specifies the library that contains the converted data.

Job Description
Specifies the job description used to submit the file-conversion job for batch processing.

The possible job description values are:

***USRPRF**
The job description defined for the submitting job's user profile.

job-description-name
Specify the name of the job description to be used.

***NONE**
A batch job is not submitted. Processing continues interactively while the user waits.

The possible library values are:

***LIBL**
The library list is used to locate the job description.

***CURLIB**
The current library for the job is used to locate the job description. If no current library entry exists in the library list, QGPL is used.

library-name
The library where the job description is located.

Note: The user's work station is not available for other use during this time, which can be significant for long jobs.

Chapter 13. A Problem Analysis Case Study

This chapter provides users of the AS/400 system with an initial approach to determining the source of performance problems using available system tools (both operating system functions and additional licensed programs).

This example provides an introduction to performance analysis and shows you some performance analysis techniques. The scenario describes how a user who is experienced in performance analysis assists a company in resolving a performance concern. You can see how the expert identifies the real problem, isolates the cause, and provides the recommended solution.

You should read "Performance Analysis Overview" on page 1-2 before reading this case study.

Note: Although the technique used represents just one of many different approaches to performance analysis, and the problem described is a small subset of the real-life possibilities, the example is designed to provide initial guidance in developing an overall strategy for performance problem analysis. The names of people and the events described in this chapter are fictitious, and any likeness to actual people is purely coincidental. Because customer applications and requirements vary, IBM makes no representation or warranty that the methodology described herein will solve or eliminate unique customer performance problems.

Introduction to Performance Analysis

Performance problem analysis is a methodology for investigating, measuring, and correcting deficiencies so that system performance meets the user's expectations. It does not matter much that the "system" is a computer; it could be an automobile or a washing machine. The problem-solving approach is essentially the same:

1. Understand the symptoms of the problem.
2. Use tools to measure and define the problem.
3. Isolate the cause.
4. Correct the problem.
5. Use tools to verify that the problem is corrected.

Initially, the analyst knows the user is not satisfied with the way the system is working. For example, it may be running too slow, too noisy, too hot, and so on. The analyst, mechanic, or repair person must first understand what the problem really is. The best way to find out is to observe the problem condition personally. Can the analyst confirm the user's complaint? If the analyst cannot, he should get as much information as possible from those users who have experienced the problem. What are the most common problem descriptions?

The key to success with any performance issue is to have a clear definition of the users' performance criteria. In other words, given the application mix, what do the users want from the system in terms of interactive response time, batch throughput, and processing requirements? For example, a system that supports an interactive order-entry application may have a response time criterion to ensure that customers do not perceive abnormal delays. Another criterion may require that end-of-day processing be completed by a specific time. Given these requirements, the analyst can establish performance objectives around system resource utilization guidelines. With a clear statement of goals and objectives, performance analysis can proceed on a firm basis.

When the objectives are understood, it is important to assess whether the hardware configuration is adequate to support the workload. Is there enough processing unit capacity? Is main storage sufficient for the application mix? It the analyst answers these questions first, perhaps through capacity planning modeling techniques, needless effort can be avoided 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 certain OS/400 commands and Performance Tools/400 to measure the system performance. The analyst should review the measured data to further define the problem and to validate or reject the hypothesis. When the apparent cause or causes have been isolated, a solution can be proposed. The analyst should deal with one solution at a time. Then changes can be made and tested. 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, one must recognize the interrelationship between the critical system resources and attempt to balance these resources, namely processing unit, disk, main storage, and, for communications, remote lines. Each of these resources may become a performance problem.

Improvements to system performance, whether to interactive throughput, interactive response time, batch throughput, or some combination, may take many forms from simply adjusting activity level or pool size to changing the application code itself. Ultimately, however, any improvement will come only through analysis of the critical resources (processing unit, main storage, disk, and remote lines) and contention for system and application objects.

The Case Study

This scenario starts with a company called Armstrong Sporting Goods, Inc. (a fictitious company). As a distributor of sporting equipment throughout the southeastern United States, Armstrong selected the AS/400 system as a means for automating much of their order-entry, accounting, inventory, and shipping operations. High quality customer service is critical to the continued success of this company.

The Players

Sue Miller is the new data processing (DP) manager for Armstrong and is the person who provided the IBM support team with most of the information regarding the perceived performance problem. Having been the DP manager for just a short time, Sue is anxious to establish her credibility with the company by quickly addressing an end-user concern over system performance.

Bob Williams is the assigned IBM systems engineer. He has been asked by Sue to assist the DP staff in resolving the performance issue. In this scenario, he is the expert.

As you read through the rest of this story, you can look over Bob's shoulder as he observes the symptoms of Armstrong's performance problem and then proceeds to isolate the cause. Along the way, you will be introduced to additional people who prove to be instrumental in getting the problem resolved.

The Configuration

Here is the configuration of the system in this story:

- One AS/400 Model D45, 16 megabytes main storage
- Two communications lines
  - A 2400 baud electronic customer support switched line
  - A spare line that is currently unused
- One 9347 Tape Unit
- One 6262 Printer
- Twenty-four 3197 Display Stations
- Two 4224 Printers
- Four 9332 Model 600 Disk Units
- OS/400 Version 2 Release 1 licensed program

With this information as the base for our example, let's begin!

The Problem

It was 9:00 a.m. Monday morning, and Sue Miller had just finished introducing herself as the new data processing manager to Bob Williams, a systems engineer from IBM. The two of them were in her office to review the systems management procedures currently in place at Armstrong. During the discussion, Sue mentioned that no formal attempt was being made to monitor the system's performance on a regular basis. Other activities, such as system backup and change management, had already been addressed by Armstrong, but Sue felt that she needed to have a better understanding of how well their current system was handling the daily demands of the company. This interest was actually prompted by concerns brought up at a recent meeting with the department managers. Some of the end users had complained that the system was running too slow and at times appeared to "go to sleep."

Bob was happy to hear that Sue wanted to start developing a performance management strategy for Armstrong. He remarked that he had worked with several companies in the past who unfortunately waited until a serious situation occurred before starting to make an effort to better understand their system requirements. Without historical information to compare past performance, the problem analysis became much more difficult.

Bob and Sue then continued talking the rest of the morning about other topics of interest involving the data processing department. At the close of their discussion, Bob suggested another meeting to further investigate the source of the performance concerns. In the meantime, Sue was asked to do the following:

1. Read Chapter 3, "Performance Tuning" to better understand performance guidelines and basic tuning techniques. This would help prepare Sue for the upcoming analysis activities that Bob would assist her with.
2. Use the error reporting functions, such as the Print Error Log (PRTERRLOG) command, to see if the system is experiencing hardware problems. Although this should be a part of normal systems management, all perfor-

mance analysis activities should first ensure that the system is running error free.

3. Install the Performance Tools/400 license program that Armstrong had purchased two weeks earlier. This licensed program would assist them in their investigation of the problem.
4. Survey the end users to find out who was experiencing unsatisfactory system performance and of what type (that is, interactive response time, batch throughput, and so forth). The performance objectives for those end users should then be determined and put in writing.

With that, Bob left with the agreement that they would get together on Friday morning.

Review

Sue is new to Armstrong and is not familiar with the system's performance history. She must quickly learn the objectives of the end users. How important is it that certain display station operators receive subsecond response time, and is it realistic given the requirements of the application? Are there any critical batch jobs that must be finished by the end of the day? These kinds of questions need to be answered for Sue to determine if a problem exists.

As Bob mentioned, even though the AS/400 system provides software tools to monitor performance, both in the operating system and in optional licensed programs, many companies do not track their system's resource usage. Using the AS/400 system's ability to continuously collect performance data, a business can review workload trends on a periodic basis.

At this point, Sue suspects a problem exists because of informal remarks by some of the end users. She has no solid evidence describing the problem and, therefore, cannot give Bob any concrete information to work with. We are not even sure if there is a performance problem. This is usually where many analysis experiences begin. Bob decided that before proceeding with the analysis, the customer should first review basic tuning guidelines, make sure an important software product he relies on (the Performance Tools/400) is ready to use, check to see if the system is running without hardware problems, and gather more information from the end users.

Checking the System's Performance

On Friday morning, Bob returned to Armstrong to begin working with Sue on analyzing the system's performance. Two days earlier, Sue met with all of the department managers where she asked them to survey their staffs on how well they thought the system was performing. The results were to be returned to her by the following Monday. Also, a review of the system error reports did not indicate that the system was having hardware problems.

Bob felt that the first step in analyzing system performance was to review data from the system interactively using the control language (CL) commands:

- Work with System Status (WRKSYSSTS)
- Work with Active Jobs (WRKACTJOB)
- Work with Disk Status (WRKDSKSTS)

Using these commands, he could quickly see if the system was able to handle the requests for processing unit, disk, and main storage adequately at that instant in time. He cautioned that because the results changed dynamically with the workload, he could not determine for sure that the system had all the capacity it needed all the time. Also, it was important that the time selected to run these commands did not include work not normally running (for example, excessive sign-ons and sign-offs). Sue assured Bob that now would be a good time to look at the system.

The following illustrations show the results of the commands and how Bob interpreted them. See Chapter 3, "Performance Tuning" for many of the guidelines Bob followed. (All output was generated using the OS/400 Version 2 Release 1 licensed program.)

First Bob issued the Work with System Status (WRKSYSSTS) command.

```
                         Work with System Status                    SYS400
                                                        01/12/91  09:31:43
% CPU used . . . . . . . :          55.7    Auxiliary storage:
Elapsed time . . . . . . :      00:09:31      System . . . . . . . . :     2400 M
Jobs in system . . . . . :           102      % used . . . . . . . . :   57.5494
% addresses used:                             Total  . . . . . . . . :     2400 M
  Permanent  . . . . . . :         2.483      Current unprotect used :      326 M
  Temporary  . . . . . . :          .026      Maximum unprotect  . . :      328 M

Type pool size and activity level changes, press Enter.

System    Pool   Reserved    Max   -----DB-----  ---Non-DB---
 Pool   Size (K)  Size (K)  Active  Fault  Pages  Fault  Pages
   1      3710      1914     +++       .0     .0     .2    1.6
   2      1500         0       4       .9    2.9    1.8    4.6
   3       290         0       4       .0    1.3     .0     .4
   4     10884         0      12      3.2   27.2    3.9   24.1

                                                                  Bottom
Command
===>
F3=Exit   F4=Prompt            F5=Refresh    F9=Retrieve    F10=Restart
F11=Display transition data    F12=Cancel    F24=More keys
```

- The overall processing unit use was 55.7 percent and did not reflect an excessively busy system.
- The elapsed time for measurement was greater than 5 minutes but less than 15 minutes—a good choice when looking for valid data that is not skewed by short surges of activity or long periods that tend to average out problems.
- The number of jobs in the system at first appeared high to Sue, but Bob explained that this number reflected all the jobs the system was keeping track of, even if they had finished but still had output yet to print (for example, job logs).
- The fault rate of the machine storage pool (always system pool 1) did not exceed 1 fault per second, indicating that pool 1 was large enough. (More information on performance tuning is available in the *Work Management Guide*.)
- The fault rate of the rest of the machine storage pools (system pools 2 through 4) was not too heavy (database + nondatabase < 10) and the total faults of all the pools was less than 15. In general, main storage did not appear to be overcommitted.

Bob pressed F11 to select the second view.

```
                         Work with System Status                     SYS400
                                                           01/12/91  14:07:43
% CPU used . . . . . . . :        55.7   Auxiliary storage:
Elapsed time . . . . . . :    00:09:31     System . . . . . . . . :    2400 M
Jobs in system . . . . . :         102     % used . . . . . . . . :   57.5494
% addresses used:                          Total  . . . . . . . . :    2400 M
  Permanent  . . . . . . :       2.483     Current unprotect used :     326 M
  Temporary  . . . . . . :        .026     Maximum unprotect  . . :     328 M

Type pool size and activity level changes, press Enter.

System    Pool    Reserved    Max   Active->  Wait->   Active->
 Pool   Size (K)  Size (K)  Active    Wait     Inel      Inel
   1       3710      1914     +++       .0       .0        .0
   2       1500         0       4      2.5       .0        .0
   3        290         0       4      8.3       .0        .0
   4      10884         0      12      2.1      2.0        .0

                                                                    Bottom
Command
===>
F3=Exit   F4=Prompt           F5=Refresh   F9=Retrieve   F10=Restart
F11=Display transition data   F12=Cancel   F24=More keys
```

- The ratio of Wait->Ineligible to Active->Wait for the interactive pool (system pool 4) was approximately 2 percent and confirmed to Bob that the activity level was set properly. He added that many customers set their activity levels so that the Wait->Ineligible is always zero. The level could be too high, causing major problems during exceptionally busy periods.
- Active->Ineligible for system pool 4 was zero. Usually, any value greater than zero in the interactive storage pool is a good indication that jobs are exceeding their time-slice values and may be candidates for submission to batch for processing.

  Bob then issued the Work with Active Jobs (WRKACTJOB) command.

```
                            Work with Active Jobs                    SYS400
                                                           01/19/91  09:47:30
CPU %:    57.2     Elapsed time:    00:08:46     Active jobs:    35

Type options, press Enter.
  2=Change   3=Hold   4=End   5=Work with   6=Release   7=Display message
  8=Work with spooled files   13=Disconnect ...

Opt  Subsystem/Job  User        Type  CPU %  Function        Status
 __  QBATCH         QSYS        SBS      .0                   DEQW
 __    ARPOST       ACT07       BCH    24.4  PGM-AR320        RUN
 __  QCMN           QSYS        SBS      .0                   DEQW
 __  QCTL           QSYS        SBS      .0                   DEQW
 __    DSP01        QSYSOPR     INT      .0  MNU-MAIN         DSPW
 __  QINTER         QSYS        SBS     1.0                   DEQW
 __    DSP02        SHIP01      INT     1.3  MNU-SM001        DSPW
 __    DSP09        SERV02      INT      .0  CMD-WRKSPLF      DSPW
 __    DSP10        SERV03      INT      .7  PGM=CS110        RUN
                                                                     More...
Parameters or command
===>
F3=Exit      F5=Refresh   F10=Restart statistics    F11=Display elapsed data
F12=Cancel   F23=More options   F24=More keys
```

- The active job count was 35. When divided into the number of jobs in the system (102/35=3.0), the result showed that Armstrong was doing a good job of cleaning up the job logs and keeping the number of jobs the system tracked to a minimum.

- No interactive jobs were using an excessive amount of processing unit use (more than 2 percent).

  Next, Bob issued the Work with Disk Status (WRKDSKSTS) command.

```
                          Work with Disk Status                      SYS400
                                                            01/19/91  10:03:59
Elapsed time:   00:09:11

             Size    %     I/O   Request   Read  Write   Read  Write    %
Unit  Type    (M)   Used   Rqs   Size (K)   Rqs   Rqs     (K)   (K)   Busy
  1   9332    300   67.6   3.3      1.7     2.7    .5     1.6   2.1    11
  2   9332    300   56.1    .6       .9      .4    .2     1.0    .7     2
  3   9332    300   55.9   1.1      1.0      .9    .1     1.0   1.5     3
  4   9332    300   56.0   2.4      2.9     1.9    .5     1.3   8.2     9
  5   9332    300   56.1   1.2      1.5     1.1    .1     1.6    .5     4
  6   9332    300   56.1   1.5      4.1     1.3    .1     4.5   1.0     5
  7   9332    300   56.0    .2      3.0      .2    .0     3.0    .0     1
  8   9332    300   56.2   3.7       .5     3.5    .1      .5   1.0    13

                                                                     Bottom
Command
===> ________________________________________________________________________
F3=Exit   F5=Refresh   F12=Cancel   F24=More keys
```

- Except for disk unit 1 (load source unit), all the other units had approximately the same amount of space used, indicating an evenly distributed system, and none of those units were over 75 percent full.
- No one drive was busy more than 13 percent of the time, and they were well under the threshold of 40 percent.

The net result of Bob's initial observation of the system showed that the system was responding well to the workload at that moment in time. Sue again commented that this period of the day was a good representation of Armstrong's normal demands on the system.

Bob felt he had a good idea of what the overall system was doing, but he planned to later validate his findings by using the advisor, option 10 on the Performance Tools menu. The **advisor** is a tool that can be run over data gathered by the performance monitor to provide conclusions and recommendations about system performance. In the meantime, with the feedback from the different departments still pending, Bob suggested using another means of gathering performance data from the system. This could be done by starting the performance monitor using the Work with Performance Collection (WRKPFRCOL) command. With this command, information similar to that provided by the earlier commands and additional detailed data on job processing could be collected over a number of days (that is, with multiple data collections) without operator intervention. Another way to start the performance monitor would be to use the Start Performance Monitor (STRPFRMON) command, but this command does not allow for automatic data collection, which is what Bob and Sue were after. The collected data could then be reviewed through commands provided by Performance Tools/400.

To begin the process of collecting performance data with the WRKPFRCOL command, Bob explained that first they would have to decide when to run the performance monitor and for how long. Without knowing exactly when the problem was occurring, Bob suggested they collect data for the entire first shift with 15-minute collection intervals. Using this size interval time would help to identify peak workloads that deserved further attention. Any longer intervals might mask a problem. Because Bob and Sue were still not sure of the type of problem that existed, the data collection would not include trace information. **Tracing** is a process by which information about each job state transition is recorded in a special table. The data may later be dumped to a database file, which can then be processed by Performance Tools/400. Tracing can result in a large amount of collected data, which could affect system performance when dumped. Normally, a more selective collection process can be used during problem isolation to limit the amount of data.

Sue entered the WRKPFRCOL command and specified the values as shown in the following display:

```
                    Work with Performance Collection

Type options, press Enter.
  1=Add   2=Change   3=Hold   4=Remove   5=Display   6=Release


        Performance
Opt     Collection      Status     Description
 1      PERFPROB__










                                                            Bottom
F3=Exit   F5=Refresh   F12=Cancel
```

The following display was shown next. Sue typed the values that are highlighted.

```
                    Add Performance Collection (ADDPFRCOL)

Type choices, press Enter.
Performance Collection . . . . . > PERFPROB      Name
Collection days  . . . . . . . .   *MON          *ALL, *MON, *TUE, *WED...
              + for more values    *TUE
Collection start time  . . . . .   0800          HHMM
Text 'description' . . . . . . .   All day collection
Library  . . . . . . . . . . . .   QPFRDATA      Name
Text 'description' . . . . . . .   Unknown Performance Problem
Time interval (in minutes) . . .   15            5, 10, 15, 20, 25, 30, 35...
Stops data collection  . . . . .   *ELAPSED      *ELAPSED, *TIME
Days from current day  . . . . .   0             0-7
Hour . . . . . . . . . . . . . .   10            0-168
Minutes  . . . . . . . . . . . .   0             0-59
Data type  . . . . . . . . . . .   *ALL          *ALL, *SYS
Trace type . . . . . . . . . . .   *NONE         *NONE, *ALL
Dump the trace . . . . . . . . .   *YES          *YES, *NO
Job trace interval . . . . . . .   .5            .5 - 9.9 seconds
Job types  . . . . . . . . . . .   *DFT          *NONE, *DFT, *ASJ, *BCH...
              + for more values
                                                                    More...
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Sue added a collection for Monday and Tuesday from 8:00 a.m until 6:00 p.m. A collection period of ten hours would include data on sign-offs and batch jobs at the end of the day. Bob commented to Sue that because this was the first time that they were adding a performance collection, they would have to submit the performance collection job after using the ADDPFRCOL command by typing the following command:

```
SBMJOB  JOBD(QGPL/QPFRCOL)  RQSDTA(*JOBD)
```

This command ensures that the performance monitor is started at the right times.

Bob left and planned to return at the end of the week to review the output with her. By that time, Sue should have collected data over two days and had time to go over the results of her end-user survey.

Review

Bob lacked information on who was experiencing the performance problem, so he decided to take some preliminary steps in understanding how well the system was responding to the daily workload. He did this by using the standard system commands, which dynamically show usage of main storage, processing unit, and disk. The important point that Bob wanted Sue to understand was that these commands only displayed this information for a very specific point in time and could not be used to represent the system's performance under all the different workloads it had to handle. This was a quick means of looking for obvious resource problems.

Bob referred to Chapter 3, "Performance Tuning" on page 3-1 and could not see anything that would indicate that the system was running slow. The next step was to capture data over a longer period of time using the Work with Performance Collection (WRKPRFCOL) command. Maybe the problem was occurring at a specific time of the day.

In most everyday situations, performance data could be collected over large periods of time to get a good idea of system activity and trends. Sampling intervals of longer duration (20 to 30 minutes) are fine for normal system tracking, but

Bob and Sue are investigating a possible problem. Shorter intervals (10 to 15 minutes) would help to highlight a problem.

Still critical to Bob's investigation was the result of Sue's survey. They still did not know what kind of a problem they were facing. It is important to thoroughly define the problem.

Reviewing the End-User Survey Results

On Tuesday, Sue received the final survey results from all the end users. Following is a copy of the survey form Sue distributed.

Armstrong Computer End-User Survey

On a scale of 1 to 5, please rate how well the computer system meets your needs in the following categories:

1 = Excellent, 2 = Satisfactory, 3 = Average,
4 = Needs some improvement, 5 = Needs much improvement

1. Availability of the computer ___
2. Interactive response time ___
3. Timeliness of printed output requests ___
4. Timeliness of batch run requests ___

For those items answered with a 4 or 5, please indicate any concerns you might have.

____
____
____

(We will follow up this survey with personal interviews for those who would like to help the data processing department improve its services to all the end users.)

Thirty-seven surveys were returned. Sue decided to concentrate on only the returned forms that indicated a 4 of 5 in any of the categories. She noticed that only two of the surveys had reflected a dissatisfaction with the system, and both were from the order-entry department. Also, the only category with negative responses was number two, *Interactive response time*. One of the two negative surveys included the following comment: Ever since the new procedure, which allowed customers to call in their orders, was put into use, the system seemed to take a very long time before the entry display appeared.

Sue met briefly with the order-entry department to discuss their survey responses and to better understand their performance requirements. During the meeting, Sue learned that the department's daily workload included both batch and interactive processing. Their batch jobs ran mostly in the evenings unattended and were not presenting a problem. The interactive jobs, however, were experiencing much longer response times than the department's objective of 2 seconds. Sue reviewed some basic application requirements, such as the average number of database read operations per transaction, and could not readily determine the source of the problem. Sue then decided that it would be better to review her findings with Bob on Friday.

Analyzing System Performance

On Friday morning, Bob arrived to analyze the collected performance data. First, Sue updated him on the results of the survey. Bob was very interested in the concerns of the order-entry department and commented that they would investigate the order-entry application. First, he would like to analyze the system performance once more using the advisor. Below is the sequence of events and displays that Bob used to perform the system analysis using the advisor.

Bob started at the Performance Tools/400 menu and selected option 10. He then selected the appropriate library and member and pressed the Enter key, which took him to the Select Time Intervals to Analyze display.

```
                     Select Time Intervals to Analyze

Member . . . . . . . .    PERFPROB      Library  . . . . . . :   QPFRDATA

Type options, press Enter
  1=Select

                     ---Tns----  --CPU Util---  High Util  --Pool Fault/Sec--
Opt  Date   Time   Cnt    Rsp   Tot  Int  Bch  Dsk  Unit  Mch  Usr  ID  Excp
 _   01/15  08:15  309    .84   16   10    3    1  0001    0    0  03  1920
 _   01/15  08:29  266    .46    6    3    1    1  0001    0    0  03  1015
 _   01/15  08:44  635    .87   24   15    5    1  0001    0    0  03  1174
 _   01/15  08:59  494    .92   53   30   15    1  0001    0    0  03  1229
 _   01/15  09:14  318    .70   62   32   20    1  0001    0    0  03  1103
 _   01/15  09:29  526    .89   71   40   25    1  0001    0    0  03  1573
 _   01/15  09:44  574    .73   43   20   15    1  0001    0    0  03  1668
 _   01/15  09:59  399    .94   48   20   19    1  0001    0    0  03  1350
 _   01/15  10:14  243   4.45   11    5    2    1  0001    0    0  03  1920
 _   01/15  10:29  246   1.49   24   15    3    1  0001    0    0  03  1834

                                                                      Bottom
F3=Exit   F5=Refresh   F11=Display histogram   F12 Cancel   F13=Select all
```

Bob suggested to Sue that they analyze all the intervals at this stage to get an idea of overall system performance. Sue agreed and Bob pressed F13 (Select all) and pressed the Enter key.

```
                        Display Recommendations

 Member . . . . . . . :   PERFPROB       Library  . . . . . . :  QPFRDATA
 System . . . . . . . :   SYS400         Version/release  . . :  2/1.0
 Start date . . . . . :   07/01/91       Model  . . . . . . . :  D45
 Start time . . . . . :   08:00:01       Serial number  . . . :  XX-XXXX
 Type options, press Enter.
   5=Display details
 Option     Recommendations and conclusions
                         Recommendations
   _        Examine error logs for indications of problems.
                           Conclusions
   _        Pool 3 fault rate is well below guidelines of 25.0
   _        Pool 4 fault rate is well below guidelines of 25.0
   _        Pool 2 W->I transition zero. Fault rate within guidelines.
   _        No performance problems were detected in system data file.
   _        No performance problems found on SDLC line MCLINE
   _        No performance problems found with DIOP(s)
   _        No performance problems found with CIOP(s)
                      Interval Conclusions
   _        ASP 1 arm % busy ranged from 21.9% on arm 0008 to 10.2% on arm 0004.
   _        Total system I/O during all selected intervals was 436203 .
                                                                       More
 F3=Exit   F6=Print   F9=Tune system   F12=Cancel   F21=Command line
```

The Display Recommendations display showed Bob and Sue that the system was performing within the guidelines and that no system-related problems or errors where affecting the performance of the system. After having completed the system analysis using the advisor, which confirmed his analysis earlier in the week, he mentioned to Sue that another way to quickly analyze system data and view trends was to use performance graphics.

Following is the sequence of events that Bob specified to produce the graphs. Bob went to the Performance Tools/400 menu and selected Option 9 (Performance graphics). Then the following display appeared:

```
 PERFORMG                Performance Tools/400 Graphics
                                                          System:    SYS400

 Select one of the following:

      1. Work with graph formats and packages
      2. Work with historical data
      3. Display graphs and packages

     70. Related commands








 Selection or command
  ===> 3_____________________________________________________________
 F3=Exit   F4=Prompt   F9=Retrieve   F12=Cancel   F13=User support
 F16=System main menu
```

Bob explained to Sue that Performance Tools contains numerous preformatted graphs for customers to user. Option 1 allows the user to work with the graph formats and packages, and option 2 allows the user to create historical data from data collected over different monitor runs (for example, once a week for a month). Historical data summarizes performance members so you can display each member as a point on the historical graph. Then a user can view system performance trends in a graphical format. Because Armstrong had previously not been collecting performance data, Sue agreed to set up a collection schedule for once a week to establish some historical data. Bob suggested they use the IBM-supplied graph formats to show performance graphs (rather than historical graphs), so they selected option 1 (Display performance data graphs).

```
                         Display Graphs and Packages

Select one of the following:

     1. Display performance data graphs
     2. Display historical data graphs











Selection or command
 ===> 1
F3=Exit   F4=Prompt   F9=Retrieve   F12=Cancel
```

Bob pointed out that the QIBMxxx formats are supplied by IBM. He commented that a good graph to begin with is the processing unit use versus time (by job type), so they selected the QIBMCPUTYP member and pressed the Enter key.

```
                         Select Graph Formats and Packages

Library  . . . . . . . . . . .    QPFRDATA

Type options, press Enter.
  1=Select   5=Display sample graph   8=Display package contents

        Format/
Option  Package     Type      Description
  _   QIBMPKG      PACKAGE   IBM GRAPH PACKAGE
  _   QIBMASYNC    FORMAT    Asynchronous Disk I/O per Second vs. Time
  _   QIBMCMNIOP   FORMAT    Communications IOP Utilization vs. Time
  _   QIBMCPUPTY   FORMAT    CPU Utilization vs. Time (Priority)
  1   QIBMCPUTYP   FORMAT    CPU Utilization vs. Time (Job Type)
  _   QIBMDSKARM   FORMAT    Disk Arm Utilization vs. Time
  _   QIBMDSKIOP   FORMAT    Disk IOP Utilization vs. Time
  _   QIBMDSKOCC   FORMAT    Percentage of Disk Occupied vs. Time
  _   QIBMLWSIOP   FORMAT    Local Workstation IOP Utilization vs. Time
  _   QIBMMFCIOP   FORMAT    Multifunction IOP (Comm) Util vs. Time
  _   QIBMMFDIOP   FORMAT    Multifunction IOP (Disk) Util vs. Time
                                                                     More
F3=Exit   F10=Restore list   F12=Cancel   F14=Sort by type   F15=Sort by name
F16=Sort by Description
```

On the following display Bob selected the performance data member to be graphed.

```
                    Select Performance Data Member

Library . . . . . .   QPFRDATA

Type options, press Enter.
  1=Select

        Member
Option  Name        Description                          Date      Time
  1     PERFPROB                                          07/01/91  14:33:24












                                                                  Bottom
F3=Exit   F5=Refresh      F12=Cancel   F15=Sort by member
F16=Sort by Description   F19=Sort by date/time
```

On the following display Bob pressed F6 (Include all data) and proceeded to the next display containing the graph.

```
               Select Categories for Performance Graphs

Member  . . . . . :   PERFPROB
  Library . . . . :     QPFRDATA

Type options, press Enter.  Press F6 to include all data in the graph.
  1=Select

Option      Category
  _         Job
  _         User ID
  _         Subsystem
  _         Pool
  _         Communications line
  _         Control unit
  _         Functional area





                                                                  Bottom
F3=Exit   F6=Include all data   F12=Cancel
```

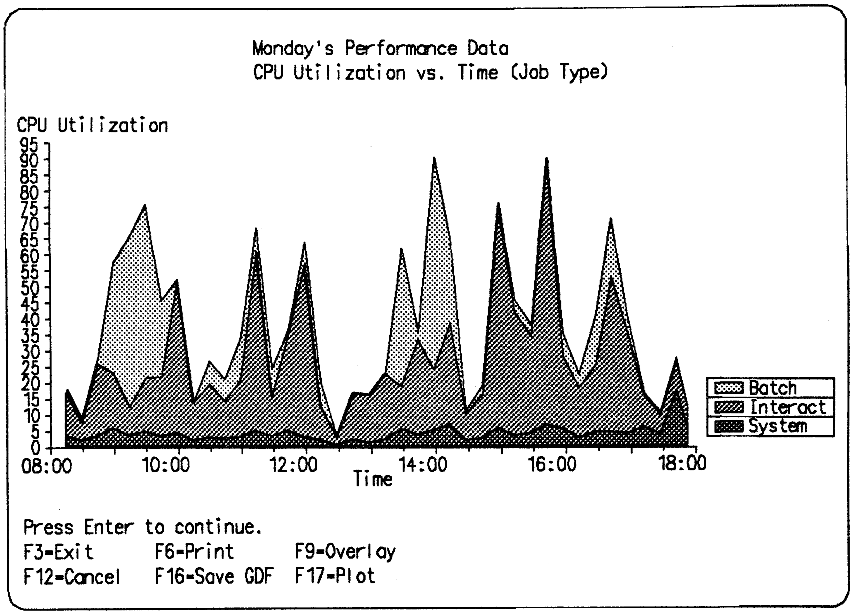

RV2S057-0

Here Bob commented to Sue that he would only show the first two and one-half hours of the collected data to give her a quick idea of what it would look like. He did this by changing the start and stop parameters to produce the following displays. The following are the graphs that Bob and Sue elected to look at (that is, they followed the same previous steps to use the formats QIBMRSP, QIBMDSKARM, and QIBMDSKIOP).

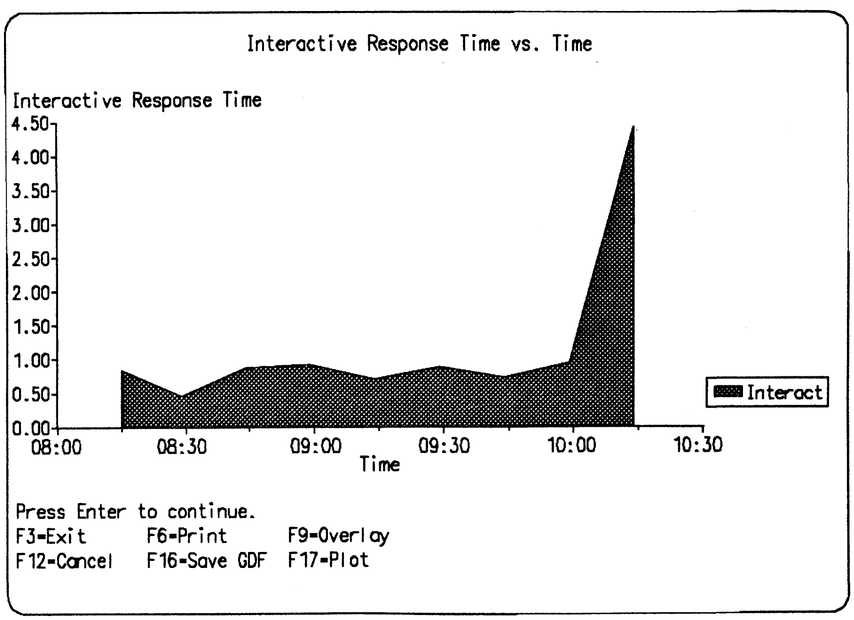

RV2S058-0

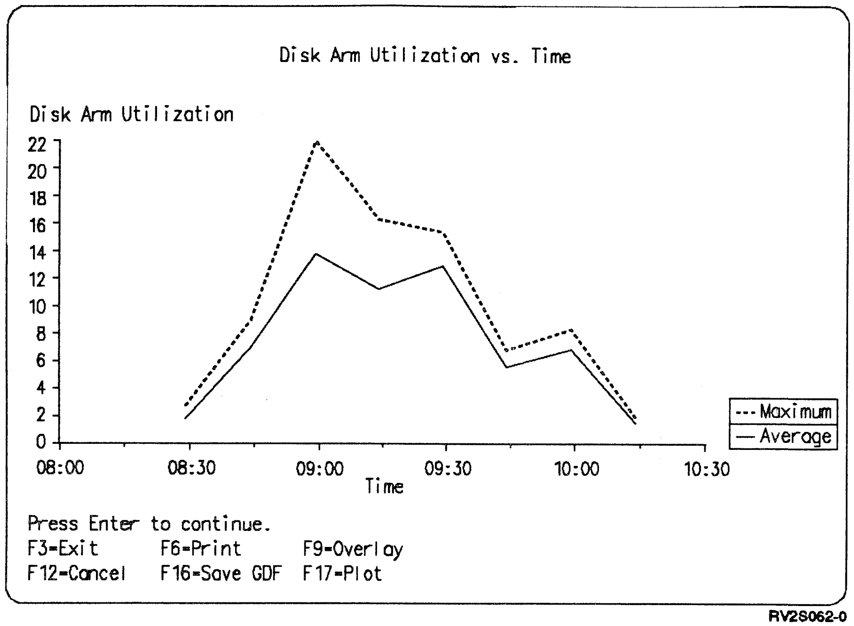

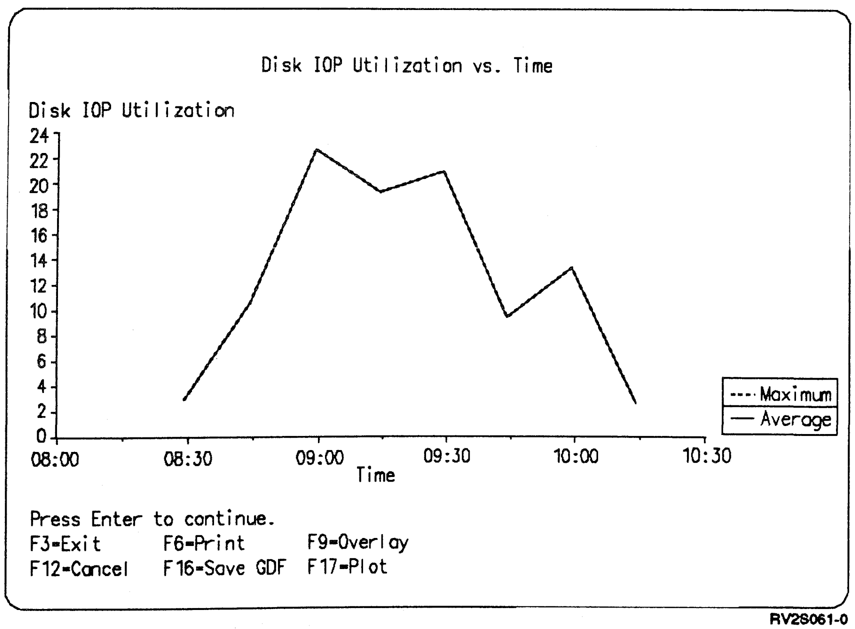

Bob then explained the graphs that they had produced and commented to Sue that the processing unit use, disk arm use, and disk IOP use showed no resource problems, and that the graphs were a quick way to pick up those types of problems without having to analyze the reports. The interactive response time graph, however, did show an abnormality just after 10:00 a.m., which should be investigated further. While the graphs gave a clear overview of how the system was performing, in Armstrong's situation, more detailed analysis was required of the gathered data.

Another way to review the collected data was to use the Display Performance Data (DSPPFRDTA) command. They could quickly see a summary of all the data interactively and isolate data of interest, which they could explore further. Following is the sequence of steps Bob used to perform further analysis.

```
                         Display Performance Data

Member . . . . . . . .   PERFPROB            F4 for list
  Library  . . . . . .    QPFRDATA

Elapsed time . . . . :   09:53:52       Version  . . . . . . :   1
System . . . . . . . :   SYS400         Release  . . . . . . :   3.0
Start date . . . . . :   01/15/91       Model  . . . . . . . :   B40
Start time . . . . . :   08:00:01       Serial number  . . . :   XX-XXXXX

CPU utilization (interactive)  . . . . . . . . . :          23.55
CPU utilization (other)  . . . . . . . . . . . . :          14.78
Job count  . . . . . . . . . . . . . . . . . . . :            172
Transaction count  . . . . . . . . . . . . . . . :          15098
Transactions per hour  . . . . . . . . . . . . . :           1525
Average response (seconds) . . . . . . . . . . . :           1.45
Disk utilization (percent) . . . . . . . . . . . :           7.29
Disk I/O per second  . . . . . . . . . . . . . . :           20.9



F3=Exit      F4=Prompt   F5=Refresh   F6=Display all jobs   F10=Command entry
F12=Cancel   F24=More keys
```

1. As in the earlier Work with Active Job (WRKACTJOB) and Work with Disk Status (WRKDSKSTS) commands, Bob found that the overall use of the processing unit and disk was not exceptionally high. Also, Sue agreed with Bob that the average response time of 1.45 seconds was acceptable. (This value reflects the internal response time of the system and does not include line transmission time, which is usually not a big difference for local work stations.)

   Bob then pressed F15 (Display by interval), which is available after pressing F24 (More keys).

```
                           Display by Interval

Member . . . . . . . :   PERFPROB        Elapsed time . . . . :    09:53:52
  Library  . . . . . :    QPFRDATA

Type options, press Enter.  Press F6 to display all jobs.
  5=Display jobs

                              CPU       Job       Tns   Average     Disk
Option  Date      Time        Util     Count     Count  Response     I/O
 _      01/15/91  08:15:00   16.11       19       309       .84      486
 _      01/15/91  08:29:59    6.97       16       226       .46     2897
 _      01/15/91  08:44:57   24.97       25       635       .87    11705
 _      01/15/91  08:59:56   53.18       28       494       .92    16719
 _      01/15/91  09:14:54   62.45       24       318       .70    17373
 _      01/15/91  09:29:53   71.60       31       526       .89    20635
 _      01/15/91  09:44:51   43.06       29       574       .73     9642
 _      01/15/91  09:59:49   48.08       19       399       .94     9409
 5      01/15/91  10:14:47   11.97       15       243      4.45     3076
 _      01/15/91  10:29:45   24.45       23       246      1.49    12556
                                                                        Bottom
F3=Exit   F6=Display all jobs   F12=Cancel   F13=Display by subsystem
F14=Display by job type
```

The Display by Interval display showed that the system was performing well for most of the users. Bob quickly rolled through all the displays, searching for the intervals where the average response times seemed noticeably higher than the average from the previous Display Performance Data display. Bob explained that if an intermittent response time problem existed, the shorter sampling interval should help to highlight it. This logic is not foolproof, he added, because high transaction counts could still reduce the average response time and mask a problem.

2. Bob found several intervals where the average response exceeded the 1.45 second average significantly. He reviewed the data to see who was having the worst response times by:

   a. Selecting option 5 (Display jobs) on the Display by Interval display
   b. Pressing F24 (More keys)
   c. Pressing F21 (Sort by response)

```
                              Display Jobs

Interval . . . . :    10:14:47          Member  . . . . . :    PERFPROB
Elapsed time  . . :   09:53:52            Library . . . . :      QPFRDATA

Type options, press Enter.
  5=Display job detail

                                               Job      CPU      Tns     Avg     Disk
Option  Job        User         Number   Type     Util     Count    Rsp      I/O
  5     DSP18      ORDENTRY01   014273   INT       .55        17   15.6      169
  5     DSP19      ORDENTRY02   014274   INT      1.55        21   13.4      252
  _     DSP38      CREDIT03     014343   INT       .71         6    3.0      389
  _     DSP14      RECV01       014337   INT       .04         1    2.0       54
  _     DSP34      SALES02      014339   INT       .32        11    1.4      243
  _     DSP41      CREDIT01     014285   INT      1.93        24    1.3      493
  _     DSP11      SHIPPING01   014289   INT       .34         8    1.3      251
  _     DSP01      QSYSOPR      014276   INT      2.10        51     .8      832
  _     DSP22      SALES01      014322   INT       .55        28     .7      311
  _     DSP40      ACTRCV01     014329   INT      2.32        62     .3      216
                                                                          Bottom
F3=Exit           F12=Cancel   F15=Sort by job    F16=Sort by job type
F19=Sort by CPU   F24=More keys
```

DSP18 and DSP19 had very high average responses, but the total number of disk I/Os for each of these jobs did not appear to be high. Sue confirmed that these were the order-entry users that had been complaining.

3. Bob entered option 5 (Display job detail) on the Display Jobs display for both these jobs to further investigate them.

```
                           Display Job Detail

Job  . . . . . . . . :   DSP18          Job type . . . . . . :    INT
  User . . . . . . . :     ORDENTRY01   Subsystem  . . . . . :    QINTER
  Number . . . . . . :     014273       Pool . . . . . . . . :    04
Member . . . . . . . :   PERFPROB       Priority . . . . . . :    20
  Library  . . . . . :     QPFRDATA     Elapsed time . . . . :    09:53:52


             CPU        Tns    Average     Disk    Act->   Wait->    Act->
Interval  Seconds     Count   Response      I/O     Wait    Inel      Inel
10:14:47    2.070        17      15.6       169      .9       .0        .0






                                                                     Bottom
Press Enter to continue.

F3=Exit    F11=View 2    F12=Cancel    F15=Sort by interval    F24=More keys
```

The Display Job Detail display allowed Bob to review the job's resource requirements in greater detail. There are actually three views that make up the total detail picture.

As Bob scanned the interval data for DSP18, he remarked that the I/O counts per transaction did not justify the high response time average. Also, the Wait->Inel and the Act->Inel were both zero, indicating that the job was obtaining and holding an activity level when needed.

DSP19 showed the same situation.

Bob proceeded to review all the detail information for those two jobs. The following conclusions were drawn:

- Both jobs were experiencing extremely wide variations in average response times.
- These variations were occurring between 9:30 a.m. and 4:00 p.m. on both of the days that data was collected.
- Resource utilizations (processing unit, disk, and main storage) were not excessively high at those times.

Bob mentioned to Sue that these two jobs were definite candidates for further investigation. The sample data, however, would not give them the detail to determine the actual cause of the erratic response times. They would need to capture another type of data using the Start Performance Monitor command. Trace data would give them greater detail on individual transactions, such as the program that was most likely running. First, though, they would meet with the order-entry personnel to get more information on how they use the system and what types of problems they were experiencing.

Sue made arrangements for all of them to meet.

Review

Bob reviewed the system performance by using the advisor to confirm his earlier conclusion about system performance. Bob and Sue then used the graphics to quickly see if there were intervals that showed particularly bad response times and high use of system resources (higher than the guidelines). This helped Bob and Sue focus on certain time intervals rather than the whole time period.

Bob decided that another way to quickly review the daily workload and its effect on system resources was to use the Display Performance Data (DSPPFRDTA) command. Rather than scan a printed report, he could interactively scan the collected sample data and isolate individual jobs that might be experiencing poor performance.

Normally, sample data could be collected with longer intervals (20 to 30 minutes) over longer periods (possibly all day) and be used to track the system's performance trends. This would enable a company to better manage its system's resources and perhaps prevent major performance surprises.

Though Sue informed Bob of the particular jobs to investigate, Bob decided to initially use the Display Performance Data command to review the overall system statistics. He then proceeded to focus on individual jobs. Bob could have just as easily selected only the order-entry jobs to look at. With no previous data to look at, Bob wanted to view all the jobs to get a feel for Armstrong's system usage.

Even with this type of data, more detail on what a job is doing must be gathered using the trace paramenter of the Start Performance Monitor command if the cause of the problem is to be isolated. Tracing, however, can generate a great deal of data and could affect the end users when dumping the trace tables. Tracing should only be used when in problem analysis mode and for shorter periods of time than when collecting just sample data. Bob wants to talk to the end users to help him understand the problem and hopefully trace the system at the most opportune time.

Understanding the Symptoms of the Problem

The two order-entry clerks met with Bob and Sue. Bob questioned them on the types of operations they performed, the problems with the system, and the types of additions or changes made that may be related to their problems. The following items were the results of their discussion:

1. Karen and Tim work in the order-entry department as clerks, processing the orders that are mailed in daily. They have their own private work station area, which restricts them from viewing each other's activities. Early in the morning, the orders are broken into two stacks. Each of the clerks takes a stack and types it into its own separate transaction file on the system. At the end of the day, a batch program posts both transaction files to the main order file. Basically, the orders are not officially in the system until the day after they are typed.

2. Armstrong gives its largest customers the ability to call in urgent orders. Normally, only Karen has the authority to take the call and run a menu option that allows her to enter the order directly into the order file without being first held in a transaction file. This type of action usually occurs about twelve to twenty times a day and requires very short interactive response times because the customer is on the telephone as the order is entered. In the past, Karen has had no problem with completing an entire telephone order in under 40 seconds.

3. Recently, Armstrong had changed its policy, allowing all of its customers the ability to call in urgent orders and inquire about order status. This has caused the number of telephone orders to increase to a point where, now, both Karen and Tim are authorized to take telephone orders and enter them directly into the order file. Each of them currently averages 40 calls a day. It seems to them that the same menu option that took less than a second to bring up the display can now take 30 to 40 seconds. This caused serious problems with customers waiting on the telephone.

Bob suggested a plan to help find the cause of the intermittent response time problem. Because transferring collected trace data to a database file might affect all of the users on the system, his plan involved controlling both the amount of time that the monitor ran and when the trace data would be dumped to a file.

Sue would run the Start Performance Monitor (STRPFRMON) command with tracing active for one hour, and she would have the option to dump the trace table later. At the end of each run, she would call Karen or Tim and ask if the problem occurred. If it had, Sue would give Bob a call, and, at the end of the day, explicitly dump the trace table with the Dump Trace (DMPTRC) command. If the problem had not occurred, Sue would restart the monitor for another hour (the trace table is overlaid with the new data). The problem happened often enough, so it should only take a few attempts to capture the necessary data.

They all agreed that this would be the best approach to resolve the problem without affecting the rest of the users. They would start the procedure that afternoon. Bob made arrangements to be back on Monday morning.

The following shows how Sue entered the Start Performance Monitor (STRPFRMON) command that afternoon:

```
                   Start Performance Monitor (STRPFRMON)

Type choices, press Enter.

Member . . . . . . . . . . . . .  oeproblem     Name, *GEN
Library  . . . . . . . . . . . .  QPFRDATA      Name
Text 'description' . . . . . . .  Order Entry Problem - Trace On

Time interval (in minutes) . . .  5             5, 10, 15, 20, 25, 30, 35...
Stops data collection  . . . . .  *ELAPSED      *ELAPSED, *TIME, *NOMAX
Days from current day  . . . . .  0             0-9
Hour . . . . . . . . . . . . . .  1             0-999
Minutes  . . . . . . . . . . . .  0             0-99
Data type  . . . . . . . . . . .  *ALL          *ALL, *SYS
Trace type . . . . . . . . . . .  *ALL          *NONE, *ALL
Dump the trace . . . . . . . . .  *NO           *YES, *NO
Job trace interval . . . . . . .  .5            .5 - 9.9 seconds
Job types  . . . . . . . . . . .  *DFT          *NONE, *DFT, *ASJ, *BCH...
               + for more values

                                                                  More...
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display        F24=More keys
```

Notice that *Trace type* was changed to *ALL and *Dump the Trace* was changed to *NO.

During the first monitor run that afternoon, the problem did not occur. After the second monitor run ended, Sue received word that Tim experienced two major response time problems when trying to enter a telephone order. No more monitor runs would be needed. At the end of the day (after most of the users had signed off), Sue issued the Dump Trace (DMPTRC) command to prepare the necessary data for Bob. The command looked like this:

```
                              Dump Trace (DMPTRC)

 Type choices, press Enter.

 Member . . . . . . . . . . . . .   oeproblem     Name
 Library  . . . . . . . . . . . .   QPFRDATA      Name
 Job queue  . . . . . . . . . . .   QCTL          Name, *NONE
   Library  . . . . . . . . . . .     QSYS        Name, *LIBL, *CURLIB
 Text 'description' . . . . . . .   *BLANK

                                                                     Bottom
 F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
 F24=More keys
```

With that accomplished, Sue was ready for Bob's visit on Monday.

Analyzing the Data

Bob returned that Monday. After Sue related the activities of Friday afternoon, the next step was to begin analyzing the data. Because the Display Performance Data command could only show sample data, Bob chose to print a job summary report using the Print Transaction Report (PRTTNSRPT) command. To reduce the amount of printed output to be analyzed, the report was limited to only the order-entry jobs.

```
                     Print Transaction Report (PRTTNSRPT)

 Type choices, press Enter.

 Member . . . . . . . . . . . . . > oeproblem     NAME
 Report title . . . . . . . . . . > Order Entry Problem - Trace On
 Report type  . . . . . . . . . .   *SUMMARY      *SUMMARY, *TNSACT, *TRSIT...
               + for more values
 Time period for report:

   Starting time  . . . . . . . .   *FIRST        TIME, *FIRST

   Ending time  . . . . . . . . .   *LAST         TIME, *LAST


                            Additional Parameters

 Library  . . . . . . . . . . . .   QPFRDATA      NAME
 Report option  . . . . . . . . .   *SS           *SS, *SI, *OZ, *EV, *HV, ' '
                                    *EV
                                                                     More...
 F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
 F13=How to use this display        F24=More keys
```

```
                      Print Transaction Report (PRTTNSRPT)

Type choices, press Enter.

Select jobs . . . . . . . . . .   *ALL          Character value, *ALL
              + for more values
Omit jobs . . . . . . . . . . . > *NONE         Character value, *NONE
              + for more values
Select users . . . . . . . . . .  ordentry*     Name, generic*, *ALL
              + for more values
Omit users . . . . . . . . . . .  *NONE         Name, generic*, *NONE
              + for more values
Select pools . . . . . . . . . .  *ALL          1-16, *ALL
              + for more values
Omit pools . . . . . . . . . . .  *NONE         1-16, *NONE
              + for more values
Select functional areas . . . .   *ALL
              + for more values
Omit functional areas . . . . .   *NONE

                                                                   Bottom
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display        F24=More keys
```

The following pages show selected sections from the Print Transaction Report output with areas that Bob highlighted and discussed with Sue.

Job Summary

The job summary for the order-entry department showed the following report and results:

```
                                   Job Summary Report                                       07/08/91 10:33:42
                                      Job Summary                                                  Page 0001
                               Order-Entry Problem - Trace On

Member . . . : OEPROBLEM   Model/Serial  . . :  D45/XX-XXXXX  Main storage . . :  16.0 M  Started  . . . . : 07/05/91 14:03:19
  Library . . : QPFRDATA   System name . . . :      SYS400    Version/Release  :  2/ 1.0  Stopped  . . . . : 07/05/91 14:57:50

                    *On/Off*    T  P  P  Tot  Response Sec        CPU Sec          ---- Average DIO/Transaction -----  Number  K/T
   Job       User      Job       y  t  r  Nbr  -------------  ------------------  ------ Synchronous ----- --Async--    Cft    /Tns
   Name      Name     Number  Pl p  y  g  Tns    Avg    Max   Util   Avg    Max   DBR  NDBR  Wrt  Sum  Max  Sum  Max  Lck Sze  Sec
---------- ---------- ------  -- -- -- - ----  ------ ------  ---- ------ ------  ---- ---- ---- ---- ---- ---- ----  ------ -----
DSP18      ORDENTRY01 031288  02 I  20 N   26     3.2   38.2    .4    .24    .42     4    8    3   15   51    1   13    2   2    85
DSP19      ORDENTRY02 031289  02 I  20 N   31      .3    2.0    .5    .26    .87     3    9    2   14   29                         74
```

The following information was extracted from this report:

- The average response time for ORDENTRY01 was 3.2 seconds and at least one of its transactions lasted 38.2 seconds. ORDENTRY01 is Tim's user profile. What was happening to that job?
- The average processing unit time per transaction was .24 seconds with at least one transaction using .42 processor seconds. These times could not be the reason for the poor performance.
- ORDENTRY01 showed an average of 15 disk I/O operations per transaction and a worst case of 51 disk I/Os per transaction. Using .05 seconds as an average disk I/O service time, these numbers do not justify the exceptionally long response times.

Sue asked Bob about the two in the *Lock Conflict* column. Bob remarked that this value indicated the number of times that ORDENTRY01 needed to wait for an object being held by another job.

This first page of the report indicated to Bob that Tim was definitely experiencing poor response times, especially since his workload was similar to Karen's. Bob needed further information on what components of the response time were causing the problem.

System Summary Data

Bob scanned further down the report to look at interactive transaction averages and exceptional-wait breakdown by job type.

DATA FOR SELECTED TIME INTERVAL (OR TOTAL TRACE PERIOD IF NO TIME SELECTION).

INTERACTIVE TRANSACTION AVERAGES BY JOB TYPE.

Typ	Prg	Nbr Jobs	Nbr Tns	Pct Tns	Tns /Hour	Avg Rsp (Sec)	CPU/ Tns (Sec)	Sync Disk I/O Rqs/Tns DB Read	DB Write	NDB Read	NDB Write	Sum	Async DIO /Tns	W-I Wait /Tns	Excp Wait /Tns	Key/ Think /Tns	Active K/T /Tns	Est Of AWS
I	NO	2	57	100.0	62	1.613	.253	3	2	8	1	14	0	.000	1.314	79.092	55.254	

EXCEPTIONAL WAIT BREAKDOWN BY JOB TYPE.

Type	Purge	A-I Wait /Tns	Short Wait /Tns	Short WaitX /Tns	Seize Wait /Tns	Lock Wait /Tns	Event Wait /Tns	Excs ACTM /Tns	EM3270 Wait /Tns	DDM Svr Wait /Tns	Other Wait /Tns
I	NO	.000	.033	.000	.003	1.277	.000	.001	.000	.000	.

The following performance information was extracted from this report:

- The average response time for 57 transactions was 1.613 seconds. This does not appear to be too high.
- However, the amount of processing unit time and disk time per transaction do not justify the 1.613 response time.
- Of the 1.613 seconds, 1.314 is spent in what is known as exceptional wait. The Excp Wait/Tns time is that portion of response time that cannot be attributed to processing unit or disk usage and is caused by contention for internal system resources (for example, waiting for a message queue). Normally, this value should be less than 10 percent of the total average response time.
- Almost all of the exceptional wait time is being spent in the Lock Wait category. (Remember Sue's question?)

Bob saw further data supporting the existence of a problem. He explained to Sue that these high numbers still reflected averages.

Analysis by Interactive Response Time

The next section Bob looked at in the report (Analysis by Interactive Response Time) would help define the makeup of the transactions.

ANALYSIS BY INTERACTIVE RESPONSE TIME.

Category	Avg Rsp /Tns	Nbr Tns	Pct Tns	Cum Pct Tns	Avg CPU /Tns	CPU Util	Cum CPU Util	Sync Disk I/O Rqs/Tns DB Read	DB Write	NDB Read	NDB Write	Sum	Async DIO /Tns	Excp Wait /Tns	Avg K/T /Tns
Sub-Second	.332	51	89.5	89.5	.229	.2	.2	1	1	2		4		.037	51.979
1 - 1.999 Sec	1.512	3	5.3	94.8	.498		.2	7		12	9	28	1		10.028
2 - 2.999 Sec	2.274	1	1.8	96.6	.419		.2	2	25	1	23	51	13		1.185
3 - 4.999 Sec				96.6			.2								
5 - 9.999 Sec				96.6			.2								
GE 10 Seconds	36.664	2	3.5	100.1	.091		.2			2		2		36.497	233

The following information regarding response time was extracted from this trace report:

- Of the 57 transactions measured, only two were greater than 10 seconds, and together they averaged 36.664 seconds.
- Almost all of that time (36.497 seconds) was spent as exceptional wait time. Remember that exceptional wait time is nonproductive time. What were those two transactions doing?

This section of the report allows Bob to help evaluate performance versus objectives. Bob sees that both jobs are actually getting excellent service most of the time. Two transactions, however, seem to be the source of the high averages.

Individual Transaction Statistics

Bob needed to find out more about those two transactions, so he scanned further down the report to the Individual Transaction Statistics section.

```
                                         Job Summary Report                                            07/08/91 18:33:42
                                   Individual Transaction Statistics                                          Page 0021
                                     Order-Entry Problem - Trace On

Member . . . : OEPROBLEM    Model/Serial  . . :  D45/XX-XXXXX   Main storage . . :   16.0 M   Started  . . . . : 07/05/91 14:03:19
  Library . . : QPFRDATA    System name . . . :     SYS400      Version/Release  :   2/ 1.0   Stopped  . . . . : 07/05/91 14:57:50

TRANSACTIONS WITH LONGEST RESPONSE TIMES

Rank       Value        Time        Program       Job Name       User Name       Number       Pool      Type      Priority
----     --------   ------------   ----------    ----------     ----------       ------       ----      ----      ---------
  1       38.157    14.23.27.921   ORD110        DSP18          ORDENTRY01       031288        02        I            20
  2       35.171    14.32.08.618   ORD110        DSP18          ORDENTRY01       031288        02        I            20
  3        2.274    14.36.11.625   QUIINMGR      DSP18          ORDENTRY01       031288        02        I            20
  4        1.951    14.41.22.705   QUIINMGR      DSP19          ORDENTRY02       031289        02        I            20
  5        1.543    14.05.56.163   QUIINMGR      DSP18          ORDENTRY01       031288        02        I            20
  6        1.041    14.05.47.886   QUIINMGR      DSP18          ORDENTRY01       031288        02        I            20
  7         .777    14.35.55.734   QUIINMGR      DSP18          ORDENTRY01       031288        02        I            20
  8         .567    14.33.08.820   QUIINMGR      DSP19          ORDENTRY02       031289        02        I            20
  9         .562    14.35.40.131   QUIINMGR      DSP18          ORDENTRY01       031288        02        I            20
 10         .491    14.29.15.071   QUIINMGR      DSP19          ORDENTRY02       031289        02        I            20
```

This section lists the individual transactions of various statistics (longest response time, processing unit, service time, and so on).

- ORDENTRY01 had two very long response times during the collection period, one at 14:23:27 (38.157 seconds) and the other at 14:32:08 (35.171 seconds).
- Bob noticed that, at both of these times, the program involved in the transaction was ORD110.

Transactions with Longest Lock Wait Time

Bob then looked at the Transactions with the longest lock wait time.

```
TRANSACTIONS WITH LONGEST LOCK WAIT TIME

Rank       Value        Time        Program       Job Name       User Name       Number       Pool      Type      Priority
----     --------   ------------   ----------    ----------     ----------       ------       ----      ----      ---------
  1       37.822    14.23.27.921   ORD110        DSP18          ORDENTRY01       031288        02        I            20
  2       34.977    14.32.08.618   ORD110        DSP18          ORDENTRY01       031288        02        I            20
  3
  4
  5
  6
  7
  8
  9
 10
```

At the same time as those long transactions, ORDENTRY01 experienced extremely long lock waits. In fact, almost the entire time spent in the transactions was spent waiting on locks. Again, program ORD110 was involved.

Bob and Sue now had an idea of what was causing the problem. But what kind of lock was it and why couldn't ORDENTRY01 get that lock? More questions needed answering.

Longest Seize/Lock Conflicts

Bob's next step was to go to the Longest Seize/Lock Conflicts section of the Job Summary Report.

```
                                            Job Summary Report                                                07/08/91 10:33:42
                                        Longest Seize/Lock Conflicts                                                  Page 0026
                                       Order-Entry Problem - Trace On

Member . . . : OEPROBLEM   Model/Serial  . . :  D45/XX-XXXXX   Main storage . . :   16.0 M   Started  . . . . : 07/05/91 14:03:19
  Library . . : QPFRDATA   System name . . . :      SYS400     Version/Release  :   2/ 1.0   Stopped  . . . . : 07/05/91 14:57:50

                                Job        User      Job                       Holder- Job Name.. User Name. Number   Pool  Type  Pty
Rank   Value       Time         Name       Name      Number Pl Typ Pty S/L  Object- Type.. Library... File...... Member.... RRN......
---- -------- ------------ ---------- ---------- ------ -- --- --- ---  ------------------------------------------------------------
  1   37.819 14.23.28.135 DSP18      ORDENTRY01 031288 02  I   20  L   HOLDER- DSP19      ORDENTRY02 031289   02    I     2
                                                                       OBJECT- DS      OELIB      ORDCTL     ORDCTL
  2   34.974 14.32.08.691 DSP18      ORDENTRY01 031288 02  I   20  L   HOLDER- DSP19      ORDENTRY02 031289   02    I     20
                                                                       OBJECT- DS      OELIB      ORDCTL     ORDCTL
  3     .090 14.32.43.670 DSP18      ORDENTRY01 031288 02  I   20  S   HOLDER- DSP19      ORDENTRY02 031289   02    I     20
                                                                       OBJECT- DS      OELIB      ORDCTL     ORDCTL
  4     .089 14.24.05.959 DSP18      ORDENTRY01 031288 02  I   20  S   HOLDER- DSP19      ORDENTRY02 031289   02    I     20
                                                                       OBJECT- DS      OELIB      ORDCTL     ORDCTL
```

This section lists the longest seize/lock conflicts in descending order with the time it happened, the requesting job, the holding job, and the held object.

- The two transactions with the long response times for ORDENTRY01 are listed here as the two longest instances of a lock conflict. The times coincide with those earlier in the report.
- The holding job (preventing ORDENTRY01 from obtaining the necessary lock) in both instances was ORDENTRY02 (Karen's interactive job).
- The lock request is for a file called ORDCTL in library OELIB.

Bob narrowed the problem to a conflict between the two jobs ORDENTRY01 and ORDENTRY02. However, Bob wanted to get a little more information on the transactions that both ORDENTRY01 and ORDENTRY02 were running during the lock conditions. Further detail on the transactions in question could be explored by running another Print Transaction Report, this time asking for transition detail information. This report normally produces a great deal of output. The report could be efficiently reviewed by selecting only the jobs and times involved with the problem.

Bob entered the Print Transaction Report (PRTTNSRPT) command to get the following display:

```
                     Print Transaction Report (PRTTNSRPT)

Type choices, press Enter.

Member . . . . . . . . . . . . . > OEPROBLEM     Name
Report title . . . . . . . . . . > 'Order Entry Problem - Transitional Report'
Report type  . . . . . . . . . .   *TRSIT        *SUMMARY, *TNSACT, *TRSIT...
               + for more values
Time period for report:

  Starting time  . . . . . . . .   142000        TIME, *FIRST

  Ending time  . . . . . . . . .   143500        TIME, *LAST


                          Additional Parameters

Library  . . . . . . . . . . . .   QPFRDATA      Name
Report option  . . . . . . . . .   *SS           *SS, *SI, *OZ, *EV, *HV, ' '
               + for more values
                                                                     More...
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display        F24=More keys
```

```
                     Print Transaction Report (PRTTNSRPT)

Type choices, press Enter.

Select jobs  . . . . . . . . . .   *ALL          Character value, *ALL
               + for more values
Omit jobs  . . . . . . . . . . . > *NONE         Character value, *NONE
               + for more values
Select users . . . . . . . . . .   ordentry*     Name, generic*, *ALL
               + for more values
Omit users . . . . . . . . . . .   *NONE         Name, generic* *NONE
               + for more values
Select pools . . . . . . . . . .   *ALL          1-16, *ALL
               + for more values
Omit pools . . . . . . . . . . .   *NONE         1-16, *NONE
               + for more values
Select functional areas  . . . .   *ALL
               + for more values
Omit functional areas  . . . . .   *NONE

                                                                      Bottom
F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display        F24=More keys
```

Notice that this time the output had been reduced to only showing information about the ORDENTRY01 and ORDENTRY02 jobs between 14:20:00 and 14:35:00.

Transition Report for ORDENTRY01

The following shows sections of the Transition Report resulting from the PRTTNSRPT command just issued.

```
                                                  Transition Report                                          07/08/91 11:41:45
                                           Order-Entry Problem - Trace On                                           Page 0001

Member  . . . : OEPROBLEM    Model/Serial  . . :  D45/XX-XXXXX  Main storage . . :   16.0 M   Started  . . . . : 07/05/91 14:20:00
  Library . . : QPFRDATA     System name . . . :     SYS400     Version/Release  :   2/ 1.0   Stopped  . . . . : 07/05/91 14:35:00
Job name  . . : DSP18        User name . . . . :    ORDENTRY01  Job number . . . :   031288   TDE/Pl/Pty/Prg . :    01B4 02 20

Job type  . . : I        Elapsed Time -- Seconds          Sync/Async Phy I/O     -MPL-
                         -----------------------       ------------------------ C  I     Last 4 Programs in Invocation Stack
              State Wait  Long   Active   Inel    CPU    DB    DB  NDB   NDB      u  n -------------------------------------------
    Time      W A I Code  Wait    /Rsp*   Wait    Sec   Read  Wrt Read  Wrt  Tot r  l    Last       Second       Third      Fourth
------------ ----- ---- ------- ------- ------- -------- ---- ---- ---- ---- ---- -- -- ---------- ---------- --------------------

*****   NO  Is The New Purge Attribute.      -

14.23.27.921 ->A                                                                   1
14.23.28.135       LKWT     */  37.819/*                      HOLDER-- DSP19       ORDENTRY02 031289
                                                               OBJECT- DS      OELIB       ORDCTL       ORDCTL
14.23.28.136 W<-   LKW             .215         .078                               1   QDBGETSQ   ORD110      QUIMNDRV   QUICMENU
14.24.05.957 ->A          37.822                                                   2
14.24.05.959       SZWT     */    .089/*                      HOLDER-- DSP19       ORDENTRY02 031289
                                                               OBJECT- DS      OELIB       ORDCTL       ORDCTL
14.24.06.077 W<-                   .120         .014                               2   QT3REQIO   QWSGET      ORD110     QUMNDRV
 ----------  ORD110              38.157*         .092    0    0    4    0   4*


<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<   Later in the report...   >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>


14.32.08.618 ->A         167.297                                                   1
14.32.08.691       LKWT     */  34.974/*                      HOLDER-- DSP19       ORDENTRY02 031289
                                                               OBJECT- DS      OELIB       ORDCTL       ORDCTL
14.32.08.692 W<-   LKW             .075         .075                               1   QDBGETSQ   ORD110      QUIMNDRV   QUICMENU
14.32.43.669 ->A          34.977                                                   2
14.32.43.670       SZWT     */    .090/*                      HOLDER-- DSP19       ORDENTRY02 031289
                                                               OBJECT- DS      OELIB       ORDCTL       ORDCTL
14.32.43.788 W<-                   .119         .014                               2   QT3REQIO   QWSGET      ORD110     QUIMNDRV
 ----------  ORD110              35.171*         .089    0    0    0    0   0*
14.33.24.806 ->A          41.018   .001         .001                               1
                                                              2
14.33.24.918   W                   .112         .019                               1   QT3REQIO   QWSCLOSE    QDMCLOSE   ORD110
                                                                                       PAG=    0  XSum=    0  PWrt=    2
14.33.24.949   A                   .030         .001                               1
14.33.25.012   W                   .064         .057                               1   QT3REQIO   QWSPUT      QUIMNDRV   QUICMENU
14.33.25.072   A                   .059         .001                               1
14.33.25.087 W<-                   .015         .014                               1   QT3REQIO   QWSGET      QUIMNDRV   QUICMENU
 ----------  QUIMNDRV             .281*         .093    0    2    0    0   2*      PAG=    0  XSum=    0  PWrt=    2
```

First, Bob scanned the Transition Report for ORDENTRY01 and noticed:

- At 14.23.28.135, ORDENTRY01 went into a lock wait for 37.819 seconds because of a request for file ORDCTL, over which job ORDENTRY02 had a lock.
- ORD110 appeared to be the order-entry program asking for the file. Programs starting with the letter Q (for example, QDBGETSQ) are normally IBM-supplied system service routines.
- The same condition appeared to be happening at 14.32.08.691.

Transition Report for ORDENTRY02

Next, Bob scanned the Transition Report for ORDENTRY02 and noted that at the times surrounding the lock waits for ORDENTRY01 (14.23.17.455 to 14.24.05.954 and 14.31.48.059 to 14.32.43.665), ORDENTRY02 was also running ORD110.

```
                                        Transition Report                                      07/08/91 11:31:45
                                   Order-Entry Problem - Trace On                                      Page 0003

Member  . . . : OEPROBLEM    Model/Serial  . . :  D45/XX-XXXXX  Main storage . . :   16.0 M   Started  . . . . : 07/05/91 14:20:00
  Library . . : QPFRDATA     System name . . . :     SYS400     Version/Release  :   2/ 1.0   Stopped  . . . . : 07/05/91 14:35:00
Job name  . . : DSP19        User name . . . . :    ORDENTRY02  Job number . . . :   031289   TDE/Pl/Pty/Prg . :    01C3 02 20

Job type  . . : I       Elapsed Time -- Seconds            Sync/Async Phy I/O      -MPL-
                        -----------------------         ------------------------ C  I     Last 4 Programs in Invocation Stack
            State Wait   Long   Active   Inel    CPU    DB    DB  NDB   NDB      u  n ------------------------------------------
    Time    W A I Code   Wait   /Rsp*    Wait    Sec   Read  Wrt Read  Wrt  Tot r  l   Last      Second       Third      Fourth
------------ ----- ---- ------- ------- ------- -------- ---- ---- ---- ---- ---- -- -- ---------- ---------- --------------------

*****    NO  Is The New Purge Attribute.       -

14.22.44.088   A                  .027           .001                             1
14.22.44.145   W                  .057           .055                             1    QT3REQIO    QWSPUT     QUIMNDRV   QUICMENU
14.22.44.222   A                  .077                                            1
14.22.44.237 W<-                  .016           .014                             1    QT3REQIO    QWSGET     QUIMNDRV   QUICMENU
 ----------  QUIMNDRV             .350*          .088    0    2    0    1    3*       PAG=    0  XSum=     0 PWrt=   3
14.23.17.455 ->A        33.217    .001           .001                             1
14.23.17.538 W<-                  .083           .082                             1    QT3REQIO    QWSGET     ORD110     QUIMNDRV
 ----------  ORD110               .084*          .083    0    0    0    0    0*
14.24.05.885 ->A        48.347                                                    1
14.24.05.954       LKRL                                         WAITER- DSP18      ORDENTRY01 031288      37.819 SECS
                                                                OBJECT- DS     OELIB      ORDCTL     ORDCTL
14.24.06.049       SZRL                                         WAITER- DSP18      ORDENTRY01 031288        .089 SECS
                                                                OBJECT- DS     OELIB      ORDCTL     ORDCTL
14.24.06.061   W                  .176           .018                             2    QT3REQIO    QWSCLOSE   QDMCLOSE   ORD110
                                                                                       PAG=    0  XSum=     0 PWrt=   3
14.24.06.094   A                  .033           .001                             1
14.24.06.151   W                  .057           .055                             1    QT3REQIO    QWSPUT     QUIMNDRV   QUICMENU
14.24.06.228   A                  .077                                            1
14.24.06.243 W<-                  .016           .015                             1    QT3REQIO    QWSGET     QUIMNDRV   QUICMENU
 ----------  QUIMNDRV             .359*          .089    0    2    0    1    3*       PAG=    0  XSum=     0 PWrt=   3

<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<  Later in the report...   >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>

 ----------  QUIMNDRV             .219*          .135    0    0    0    0    0*
14.31.48.059 ->A       148.910                                                    1
14.31.48.138 W<-                  .079           .077                             1    QT3REQIO    QWSGET     ORD110     QUIMNDRV
 ----------  ORD110               .079*          .077    0    0    0    0    0*
14.32.43.598 ->A        55.460    .001           .001                             1
14.32.43.665       LKRL                                         WAITER- DSP18      ORDENTRY01 031288      34.974 SECS
                                                                OBJECT- DS     OELIB      ORDCTL     ORDCTL
14.32.43.760       SZRL                                         WAITER- DSP18      ORDENTRY01 031288        .090 SECS
                                                                OBJECT- DS     OELIB      ORDCTL     ORDCTL
                                                              2         1
14.32.43.773   W                  .175           .018                             2    QT3REQIO    QWSCLOSE   QDMCLOSE   ORD110
                                                                                       PAG=    0  XSum=     0 PWrt=   3
14.32.43.806   A                  .034           .001                             1
14.32.43.865   W                  .058           .057                             1    QT3REQIO    QWSPUT     QUIMNDRV   QUICMENU
14.32.43.942   A                  .077                                            1
14.32.43.957 W<-                  .015           .014                             1    QT3REQIO    QWSGET     QUIMNDRV   QUICMENU
 ----------  QUIMNDRV             .360*          .091    0    2    0    1    3*       PAG=    0  XSum=     0 PWrt=   3
```

Bob felt that exploring program ORD110 might help them understand why the locks were occurring. Sue took Bob to the data processing department to talk with Armstrong's lead programmer.

Review

Bob and Sue together determined that they would first select only the order-entry jobs when producing the transaction summary reports. They could do this only because they had a good idea of the jobs in question. Under different circumstances, using all of the jobs as input to the report may be necessary. Limit the number of transactions to analyze whenever possible.

Through the different sections of the report, Bob was able to isolate not only the job in trouble, but also the individual transactions, times, and programs involved in the problem.

It is important to note that Bob did not stop at finding the job having a problem. It is much more critical that the cause of the problem be found. ORDENTRY01 is the job preventing ORDENTRY02 from obtaining service. By looking at job transition information and matching times, the suspected program (ORD110) was identified.

With this information, Bob and Sue could now approach the application developers for a solution.

Finding the Cause and Correcting the Problem

Mike Brown was the senior applications programmer for Armstrong. Although not directly responsible for the order-entry application code, he could get the necessary documentation on ORD110 for Bob so that the record lock problem could be analyzed.

1. Mike and Bob reviewed the program's flowchart and source code and were able to determine that ORD110 is an RPG/400* program that opens four files:

CUSMSTL	Customer master file, input only
ORD110D	Work station display file, input and output
ORDFILL	Open orders file, output only (add)
ORDCTL	Order control file, update

2. When *Enter a new order* is selected, ORD110 gets the single control record from ORDCTL, which contains the next order number. Every order must have a unique order number.
3. The order-entry clerk responds to a prompt from display file ORD110D, asking the clerk for the customer number. This customer number is then used by the program to get customer information from the CUSMSTL logical file, which in turn is presented to the clerk.
4. The clerk enters the necessary order data. When finished, the data is added to ORDFILL as a new order.
5. Finally, the order number field of the control record is incremented by one and written back to ORDCTL. This allows the next order entered to have the next higher order number.

To Bob and Mike, the record lock conflict for ORDCTL was very obvious. With only one clerk using ORD110, the lock on the control record for update did not present any problem. Armstrong's original policy of having Karen as the only authorized user of ORD110 ensured that only one clerk would use ORD110. The other orders received through the mail would not be assigned an order number until the night time batch job.

With the change in policy allowing multiple clerks to access ORD110, two clerks could now attempt to enter an order at the same time. Only one clerk, though, could have the ORD110D display available to them because they would first need an exclusive lock on the control record. This record would be locked for the entire order process. The requesting job's display would be inhibited while the holding job completed its order. Because the process only lasted about 30 seconds, the control record was released before another requesting job timed out (the default wait time on a record lock is 60 seconds). Had the time-out occurred, a function check would have alerted the data processing department to a lock problem much sooner.

Mike quickly created a coding correction for ORD110 such that the reading, incrementing, and updating of the control record would be done at the end of the order process. This would allow the records to be locked and released in an instant and allow other jobs to do the same. Later on, a more efficient technique, such as using a data area to store the control information, could be further explored.

Bob suggested to Sue that she run the performance monitor again for the next day to measure the results of the change. The tracing option would be off and the monitor should run for the whole day. The order-entry department was to notify her if the response time situation occurred again. At a later date, Bob would return to work with Sue on developing some system monitoring practices that Armstrong should use with Performance Tools/400.

Final Review

The case study you just read is an example of one person's approach to solving a typical application performance problem. The methodology was based on several logical steps:

1. Understand the symptoms of the problem

   Initially, Bob was made aware of a problem with very little information to help him to solve it. His first actions involved using commands to determine how well the system was reacting to the overall workload. By isolating the users having problems and talking to them, he was able to identify their objectives and substantiate the existence of a problem. The information he collected through the interview with the order-entry personnel was critical in effectively analyzing the situation.

2. Use tools to measure and define the problem

   Performance Tools/400 proved instrumental in determining not only what jobs were part of the problem, but also what programs were involved and at what times. Problems like poor response time have a definite cause and, in most cases, the available system tools can help capture and report the vital information. Selecting specific times and jobs enabled Bob to reduce the amount of data that had to be analyzed.

3. Isolate the cause and correct the problem

   Bob and Mike carefully analyzed the problem and examined the application and database design to develop a solution. They also ensured that the solution did not produce negative effects for other jobs or cause incorrect data in the business operations.

4. Use tools to verify the problem is corrected

   As mentioned earlier, Sue was to run the Start Performance Monitor (STRPFRMON) command the next day to measure the results of the change. If new problems appeared, the above steps would be repeated until the solution became acceptable.

Armstrong's story is an example of a single, isolated problem. In some cases, a system may have many different problems occurring at the same time. Prioritize the problems to select which items to investigate first. When those problems are resolved, go after the next in line until the situation no longer justifies the time and effort.

Another situation may be that a big problem is the result of an accumulation of many little design flaws. Some poor programming techniques may not affect one user much, but if multiplied by many jobs running at the same time, the result can be dramatic.

Finally, the fact may be that the resources are seriously overcommitted and that it is time for a model upgrade or another disk controller. Use the capacity planning option of Performance Tools/400 to help you determine the additional resources needed to meet the performance objectives.

Learn the proper usage of the tools available to you, and start to put into place a strategy that will help you get the most out of your AS/400 system.

Appendix A. Defining Transaction Boundaries

Performance tools reports show different values for transaction service time and resource use, depending on what command you use to analyze the performance data. These values vary because of differences in the data collected by the commands, and can supply different values for the transaction boundary start and end times. Be careful when you analyze and compare data collected for the same run using different tools.

Elements of Response Time

The elements of end-user (external) response time to interactive transactions are composed of communications time (input and output) and host (internal) response time, as shown in Figure A-1. For locally attached display stations, communications time includes the local Work Station Input/Output Processor (IOP) time. For remotely attached display stations, communications time includes communications line time, communications IOP time, and Remote Work Station Controller time as appropriate.

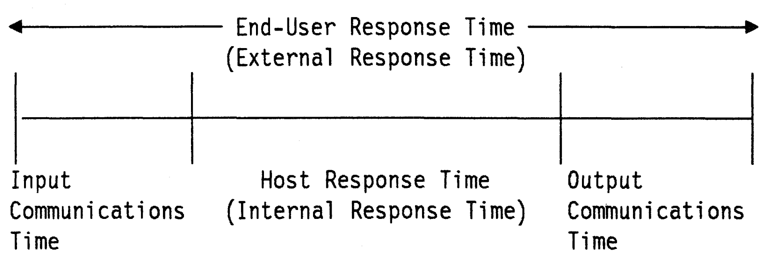

Figure A-1. Elements of Interactive Response Time

The AS/400 system interactive response time values reported by the Work with Active Job (WRKACTJOB), Print System Report (PRTSYSRPT), Print Component Report (PRTCPTRPT), and Print Transaction Report (PRTTNSRPT) commands refer only to the host (internal) response time. (An exception to this is the Local Work Station Report, shown in Figure 8-26 on page 8-36. This report does factor in local work station IOP time.)

For locally attached displays, the communications time is usually less than 1 second. For remote displays, the communications time may be longer. To approximate the actual time, use the line speed and number of characters sent and received, assuming that the line is not heavily loaded. If the line is heavily loaded, the external response time increases due to the queueing time. Review the line utilization and data transmission values on the System Report, shown in "System Report: Communications Summary" on page 8-19 to determine line component to approximate line time.

You can also use the Model System (MDLSYS) capacity planning feature to estimate external response times at local and remote display stations, because MDLSYS projects both internal and external response times. MDLSYS supports only 5250-type devices that are attached either locally through twinaxial cable or remotely through SDLC communications lines.

The host response time can, however, be shown in more detail, as in Figure A-2.

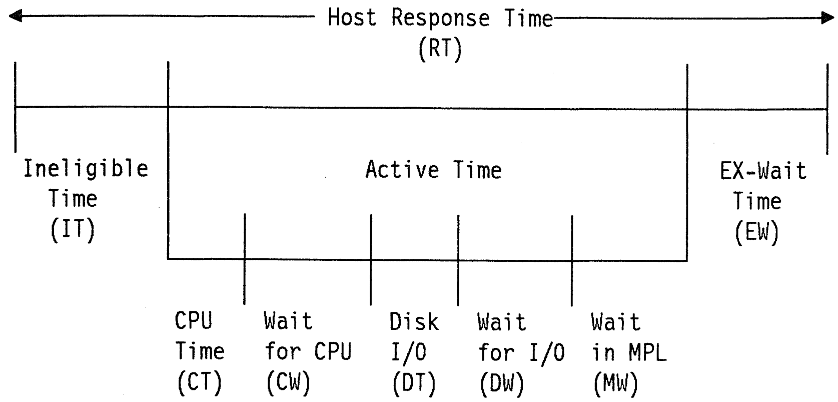

Figure A-2. Elements of Host Response Time

Note: Multi-Programming Level (MPL) is a term used interchangeably with Activity Level.

The average ineligible time, processing unit time, wait in MPL time, and exceptional wait time per transaction are available directly from the output of the PRTTNSRPT command.

Differences in the Transaction Response Reports

Figure A-3 on page A-3 compares the ways that the Print Job Trace (PRTJOBTRC) command, the PRTTNSRPT command, the PRTSYSRPT command, and the WRKACTJOB command determine transaction boundaries.

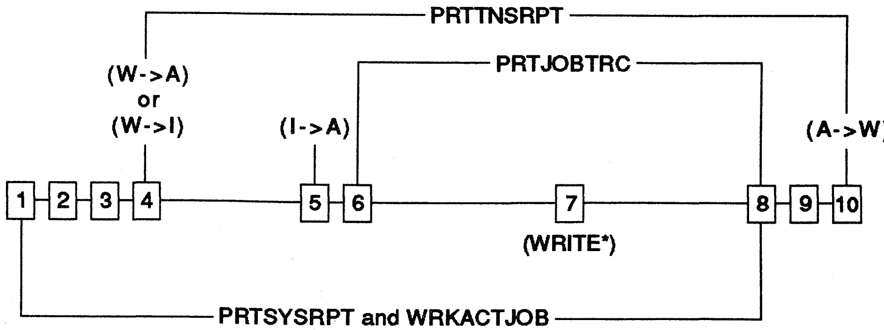

Work Station I/O Manager
(1) External I/O request received (PRTSYSRPT start)
(2) Licensed internal code processing complete
(3) Job put into activity level or ineligible state
(4) Trace record generated (PRTTNSRPT start)

OS/400 System Application
(5) Ineligible time complete (I-A)
(6) Return to QWSGET (Start of transaction on job trace)
(7) Write to Work Station

Work Station I/O Manager
(8) Call QT3REQIO (End of transaction on job trace, Transaction response times, PRTSYSRPT transaction end)
(9) Job goes to IOM to wait on I/O (PRTSYSRPT transaction end)
(10) A-W trace recorded (PRTTNSRPT transaction end)

RV2S065-0

Figure A-3. Comparison of Transaction Boundary Definitions

PRTSYSRPT and WRKACTJOB define a transaction from the time it is processed by the licensed internal code I/O manager (licensed internal code IOM) until the system work station I/O program QT3REQIO is called to wait for input.

PRTTNSRPT defines a transaction from the time trace records are produced at the beginning when the job state changes from wait-to-active or wait-to-ineligible (the start) until the job goes to a long wait (active-to-wait).

These commands include the time the job spent in the ineligible state waiting for an activity level as part of the transaction response time.

PRTJOBTRC defines a transaction from the time the job becomes eligible (for example, it is granted an activity level) within the system work station input program (QWSGET), until the system work station I/O program QT3REQIO is called to wait for input.

Note: This command does not include the time spent in the ineligible state waiting for an activity level in the transaction boundary definition.

Operational Considerations

Limitations exist in the system's ability to detect certain types of transactions.

When you review performance reports, be aware of when your system workload consists of any of the following types of work:

- Programmable work station servers
- Distributed data management (DDM) servers
- 3270 emulation jobs
- Finance terminals

- Pass-through jobs

Transaction-type data (such as the data collected for throughput and response time) is unavailable sometimes, and in some instances (such as for finance types of work), cannot be associated with the individual jobs or terminals that originated the transactions.

When you find that differences exist between the sample data reports (PRTSYSRPT or PRTCPTRPT) and the trace data report (PRTTNSRPT), it is often due to the presence of one or more of these types of work. Use the Select/omit option on the reporting commands to remove these types of jobs so the information shown on the reports is more representative of your environment.

You may find that the performance tools transaction information is inaccurate for applications such as RM/COBOL-85 for the AS/400 licensed program that do field-by-field processing. (Field-by-field processing implies that for every field in which data is entered, there is processing by the CPU as the field is exited.) The tools report each field processed as a transaction. Since these 'field' transactions may not do much processing other than return to the screen to enable the next field to be entered, the transaction information is skewed. When all of the fields on the screen have been entered, what would be viewed as a normal transaction occurs, that is, all of the information is processed.

If the transaction information is skewed due to field by field processing, it cannot be used as input to MDLSYS. MDLSYS uses the transaction information to establish its base information. It then uses the base information to predict the AS/400 modes, response time, transactions, and utilizations for a given number of work stations. If the transaction information is skewed, MDLSYS may give incorrect results. Another piece of information that can be incorrect is the working set size that is calculated when a profile is generated for MDLSYS by the system report. The working set size is calculated using the amount of CPU used as a basis for the amount of storage required. Since a small amount of CPU may have been used to go to the next field, the working set size may be grossly understated.

Appendix B. Model System (MDLSYS) Queuing Equations

When you use MDLSYS capacity planning, you should have a basic knowledge of the queuing equations used in the program. You should also understand the relationship between system resource utilization and system performance. This section presents a basic explanation of the queuing equation concepts used in capacity planning.

The MDLSYS queuing equations combine the input data, which shows the uses of the system by interactive and batch processing, into a set of estimates to show the contention for the system resources of processing unit (CPU), main storage, and disk.

When multiple jobs simultaneously need a resource, contention occurs. The result is that a job takes longer to complete its task than it would in a dedicated environment where no other jobs are contending for the resource. In many instances, where resource contention occurs, each requester has the same chance as any other to use the resource. This occurs, for example, in the system when several interactive jobs need to use the processing unit, disk, or main storage, and these jobs run at the same priority.

When contention occurs, the jobs that cannot get the resource must wait for a period of time until it becomes available. This resource wait time increases the total time required to use the resource (for example, from the time it is requested until the time it is released). The amount of time a job must wait for the resource (the **queuing time**) depends on the utilization of the resource.

Assume that in a noncontention (or dedicated) environment, an interactive transaction normally uses a half-second of processing unit time before it completes. If the processing unit is 50% utilized (on the average) by a number of jobs, all with the same priority, you can use queuing theory to calculate that the job requires 1 second to complete (the total time is the sum of the processing unit service time and the processing unit queuing time). The job still uses a half-second of processing unit time, but because of the queuing effect, it takes twice as long to use that half-second. This multiplying effect, called the queuing multiplier (QM), can be estimated by the equation shown in Figure B-1.

$$QM = \frac{1}{(1 - u)}$$

where u is the utilization of the resource
(50% utilization is a u of 0.5)

Figure B-1. Estimating the Queuing Multiplier

From this equation you can see that the lower the server utilization, the lower the value of QM.

Other values for QM for a given *u* are shown in Figure B-2.

u	QM
.1	1.1
.2	1.25
.3	1.43
.4	1.66
.5	2
.6	2.5
.7	3.33
.8	5
.9	10

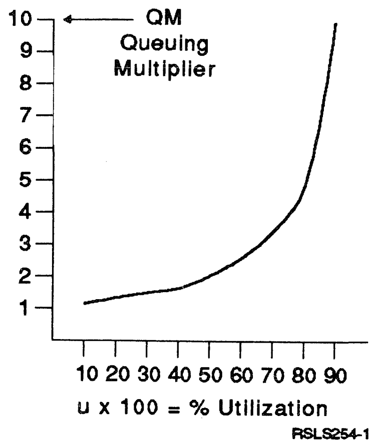

Figure B-2. Queuing Multiplier for Given Utilization of the Resource

The values shown for the QM indicate that the higher the utilization of a resource, the larger the QM value. Consequently, it takes longer for each requester to obtain service from the resource. In any system, the higher the utilization of a resource, the longer each request waits to use that resource.

When total time for a single-server, single-queue service (such as a disk drive) is calculated, the QM equation is used. When multiple-server, single-queue service time (such as the activity level slots within a storage pool) or single-server, multiple-queue service time (such as the AS/400 processing unit) are calculated, the equations for the QM are different. However, the principle and the effect are the same.

Note: The value for the QM is independent of the speed of operation of the server (the resource); the QM value is the same for a given utilization of a server.

The preceding example forms the basis for the calculations performed by capacity planning. Using system measurements as input, capacity planning output closely parallels system performance data. In some cases the output may not match, but it is close enough to be useful.

The primary function of MDLSYS is to apply queuing equations to the input data to determine the contention for system resources. The output shows the total system utilization of these resources, including an estimate of batch utilization of resources (using the assumption that batch runs at a lower priority than interactive).

MDLSYS does not consider separately the application design, file and record contention, transaction types, or variations in interactive load over a period of time. Instead, the input data of transaction processing unit time and disk I/O provide the effect of the design, the system load, and the system configuration. For this reason, any system input to capacity planning should be from a system that is properly tuned and running under expected conditions of loading.

Response Time Variation

The transaction response times on a system vary to some degree around the average transaction response time. At high utilizations, interactive response times vary significantly around the average value, and in some cases may fluctuate widely.

You need to consider the effect of the response time variations from the average. If you are achieving your objectives for interactive response time, you need to understand the effect of server utilization and queuing on your objectives.

If you install additional interactive workload on a system that is currently providing satisfactory response time, the additional workload could cause the total interactive processing unit utilization to increase. Depending on the amount of additional load, both the average response time for current applications and the amount of variation from the average may increase.

Model System (MDLSYS) Implementation

MDLSYS tries to relate throughput to resource utilization. As throughput increases, resource utilization increases. As both increase, resource contention, requester wait times, and queuing times increase. As more work stations and applications are added, the system gets busier and utilization increases.

The following equations are intended to give a high-level view of the modeling that takes place. These are not all-inclusive of the process MDLSYS uses.

A major function of MDLSYS is to calculate resource wait time.

The general equation for internal response time in MDLSYS is given in Figure B-3.

```
R = Rc + Rd + Rm + Rx    where   R  = average internal response time
                                 Rc = processing unit time component
                                      (service time plus wait time)
                                 Rd = disk component
                                      (service time plus wait time)
                                 Rm = main storage component
                                      (activity level plus wait time)
                                 Rx = exceptional wait component
                                      (excessive wait time)
```

Figure B-3. Response Time in MDLSYS

MDLSYS converts the input data regarding resource utilization and service time into wait times for each component. In general, a response component is calculated using one of the equations shown in Figure B-4.

```
Rn = Service time + wait time
              or
Rn = Service time × queuing multiplier
```

Figure B-4. Response Component

The concept of the queuing multiplier (QM) was discussed earlier. The formula that demonstrates the value is shown in Figure B-5.

$$\frac{1}{(1 - u)}$$

where u is the server utilization divided by 100 (for example, 50% = 0.5)

Figure B-5. Value of the Queuing Multiplier

Processing Unit Queuing

The formula that calculates processing unit queuing time for interactive transactions in MDLSYS is shown in Figure B-6.

Queuing for Interactive Transactions

$$\frac{1}{1 - U((n - 1)/n)}$$

where n = activity level or active display stations, whichever is smaller
U = CPU utilization for system, spool, and interactive tasks

Figure B-6. Processing Unit (CPU)

Disk Queuing

The formula that calculates disk queuing time in MDLSYS is shown in Figure B-7.

$$\frac{1}{1 - U((n - 1)/n)}$$

where n = activity level or active display stations, whichever is smaller
U = drive utilization

Figure B-7. Disk Queuing

The formula for the QM tends to lessen the effect of high arrival rates at a disk, thereby giving a more accurate value for disk queuing than the simpler 1 / (1-U) equation.

Main Storage Utilization

The formulas that calculate activity level queuing time in MDLSYS are shown in Figure B-8 and Figure B-9.

$$U = \frac{(Rc + Rd + Ex) \times A}{n}$$

where
Rc = processing unit response
Rd = disk response time
Ex = exceptional wait time
A = arrival rate (transactions/seconds)
n = number of activity levels
U = activity level utilization

Figure B-8. Activity Levels in MDLSYS

The equation that determines the QM for an activity level is shown in Figure B-9.

$$\frac{1}{1 - \left(\frac{U(n-1)}{n}\right)^{AL}}$$

where
U = (see previous calculation)
n = number of active display stations
AL = activity level

Figure B-9. Queuing Multiplier for an Activity Level

Appendix C. Capacity Planner RAMP-C Workload Description

RAMP-C is a generic interactive commercial application that is divided into transaction types called classes 1, 2, 3, and 4. Each transaction type represents a different complexity of workload. Figure C-1 on page C-2 lists the characteristics for each transaction type.

You can mix these transaction types in any combination to resemble an application that you would like to analyze. Each mixture of transaction types becomes a different RAMP-C predefined profile in MDLSYS. For example, if your application is very simple, the Inquiry transaction type might be sufficient to describe that application. This would become one of your RAMP-C predefined profiles. If another application has simple menus, in addition to other displays, with increased complexity or file activity, use the Inquiry or Simple Update transaction types with the complex ones, such as Multiple Entry Update, to create another predefined profile.

IBM designed the RAMP-C generic application to measure and compare system performance capabilities. The transaction type mixture that MDLSYS initially defaults to is the most common mixture used in marketing performance literature. This default mix of the transaction types is considered representative of a typical data processing application. To have the RAMP-C predefined profile resemble your commercial application, change the default mixture accordingly.

A sample end user equivalent of each of the four transaction types follows:

Transaction Type	**End User Equivalent**
Inquiry (Class 1)	Looking up a part price using the part number as a key.
Simple Updates (Class 2)	Entering an accounts receivable or entering an order consisting of one line item.
Multiple Entry Updates (Class 3)	Entering an order consisting of six line items.
Complex Processing (Class 4)	Entering an order consisting of 10 line items or updating 10 items in an inventory file.

Figure C-1 shows the RAMP-C characteristics, according to class.

Figure C-1. RAMP-C Characteristics (by Class)

RAMP-C Characteristics	**Inquiry (Class 1)**	**Simple Updates (Class 2)**	**Multiple Entry Updates (Class 3)**	**Complex Processing (Class 4)**
Database Files (per Program)	2	3	3	5
Read by Key	1	3	11	25
Read Sequential	--	2	5	5
Update	--	1	6	10
Add (to Single Logging File)	1	1	1	1
Display Files	1	1	1	1
Fields				
Input	1	5	--	--
Constant	36	36	12	12
Output	23	18	8	8
Input/Output	--	5	30	54
Characters				
Input	8	40	0	0
Constant	425	439	154	154
Output	120	80	21	21
Input/Output	0	40	230	414
COBOL Statements Run (Average per Transaction)	99	131	345	625
Program Size (COBOL Statements)	403	462	560	652

Appendix D. Capacity Planner System/36 Migration Utility

When migrating an application from System/36 to the AS/400 System/36 environment, the performance depends on the System/36 application structure. The Capacity Planner System/36 Migration Utility uses the System/36 measured performance data, and through modeling, translates the data into AS/400 System/36 environment performance data for the AS/400 capacity planner (MDLSYS). MDLSYS is then run to determine an AS/400 configuration that meets the anticipated performance needs.

The Capacity Planner System/36 Migration Utility translates System/36 measured performance data into AS/400 performance data for MDLSYS based on the following assumptions:

- The Capacity Planner Migration Utility should only be used to measure and model those applications that will be migrated to the AS/400 System/36 environment.
- There is no Multiple Requester Terminal (MRT) program bottleneck on the System/36. If there is a bottleneck, the performance tool's results may be incorrect because of incorrect workload overheads. The incorrect results are usually optimistic, which means the predicted external response times are low.
- The batch activity on the System/36 is at a minimum.
- The local and remote throughput is proportional to the number of active local and remote work station operations measured on the System/36 using the System Measurement Facility (SMF).
- The number of active local and remote display stations are proportional to the number of local and remote work station operations.
- The assumptions in Chapter 10, "Capacity Planning and Performance Prediction" still apply. The information in "What to Expect from MDLSYS" on page 10-9 and in "Selecting Data for Capacity Planning Measured Profiles" on page 10-11 is of particular importance.

The ability of the Capacity Planner System/36 Migration Utility to translate System/36 measured performance data into AS/400 System/36 environment performance data for MDLSYS is limited by the following:

- The System/36 migration package is used for migration to the System/36 environment on the AS/400 system.
- MDLSYS is not a disk storage capacity sizer. It does not predict the disk storage capacity needed, but it predicts the number of disk arms needed for performance.
- MDLSYS does not model the wait time for the MRT programs. Too many display stations attached to a MRT program on the System/36 environment may result in excessive wait times, which are not shown in the model results.

Migration Utility Instructions for the System/36 MIGUTL Procedure

Use the Migration Utility to help select a comparable AS/400 configuration when you are migrating from the System/36 to the AS/400 System/36 environment.

Before you use the Migration Utility on the System/36, be sure program temporary fix (PTF) level DK3700 is installed. This PTF allows additional, necessary counters to be captured by SMF. To check whether the PTF level is correct, type `DUMP PTF,CRT` on the System/36 and check the DK level in the top right portion of the display.

The MIGUTL procedure performs the following functions:

- Runs SMF to collect system measurement data.

  Note: SMF runs for 2 hours and requires approximately 2000 blocks of disk space for the files.

- Runs the SMFDATA procedure.
- Saves the file on diskettes with a VOLID of MIGUTL. The save checks for the correct VOLID of MIGUTL and issues an error if the correct diskette is not in slot 1.
- Deletes the data collection files created by SMF.

If the measurement is run with remote communications such as DDM, ICF, and so on, the modeled response time per transaction for your interactive display stations may be high because system communications activity is included with the interactive workload. However, the total system workload will be characterized.

Run the Migration Utility during a peak workload period for the workload to be migrated. To use the Migration Utility, follow these instructions:

1. Format three diskettes (either 8 inch or 5-1/4 inch) with a VOLID of MIGUTL by entering the following command:

   ```
   INIT MIGUTL,,FORMAT2,S1,S3
   ```

   Note: The diskettes may be removed while MIGUTL is running but they must be replaced when SMF stops running (2 hours).

2. Enter the following command at any workstation:

   ```
   // EVOKE MIGUTL
   ```

   This makes the job a batch job.

   Note: Should any stops occur while running SMF, select option 2 so the other job steps continue. Stops may indicate that the SMF data file is full or the SMF task is not active.

3. Have the diskettes taken to the AS/400 system for capacity planning with MDLSYS.

When the measurement is running, there should be a significant amount of interactive work present. For example, do not take the measurement over the lunch hour when few users are active. Otherwise, the measurement data will not be representative of the workload that will be migrated to the AS/400 system.

Example of Migrating System/36 Measured Data

This example shows you how to use the Capacity Planner System/36 Migration Utility to translate System/36 measured data, through modeling, into AS/400 System/36 environment performance data for MDLSYS.

Note: This example was created early in the development of the capacity planner. The numbers shown in this example are for demonstration purposes only and do not represent any real results.

1. Follow the instructions in "Migration Utility Instructions for the System/36 MIGUTL Procedure" on page D-2.
2. Put the first diskette, with the VOLID of MIGUTL, into the AS/400 diskette drive. Verify the name of the diskette unit on the system.
3. Enter the MDLSYS command to begin working with the capacity planner. The display stating the capacity planner disclaimer appears.
4. Press the Enter key. The Work with Capacity Planning Response File display appears.

```
                         Work With Capacity Planning Response Files

 Library . . . .   QPFRDATA     Name

 Type options, press Enter.
 1=Create          2=Change                           3=Save
 4=Delete          5=Work with planning results       6=Grow results

 Opt   File        Description                             Date      Time
  1    EXAMPLE36   Add S/36 migration profile to MDLSYS














 F3=Exit                     F4=Work with graph files          F10=Sort by name
 F12=Cancel                  F13=Sort by description           F14=Sort by date
```

5. On this display, select option 1 (Create) to create a response file.
6. Type a file name in the *File name* field, and a description of the file in the *Description* field. In this example, EXAMPLE36 is entered as the file name, and Add S/36 Migration profile to MDLSYS is entered as the description.

7. Press the Enter key. A response file member named EXAMPLE36 is created, and the Work with Response File Components display appears.

```
                    Work with Response File Components

 Response file . . . . . :   EXAMPLE36

 Type options, press Enter.
   2=Change        4=Delete

  Opt    Component        Type       Description
   _     Functions         4         System Functions
   _     Objectives        5         Workload Objectives










 Type:     1=Measured  2=Predefined  3=User  4=System functions
           5=Objectives    6=S/36 migration profile

 F3=Exit              F6=Measured profiles          F9=Add predefined profile
 F11=Planning results  F12=Cancel                   F14=Growth
```

8. To include the System/36 migration profile, press F6. The Work with Measured Workload Profile display appears.

```
                    Work with Measured Workload Profiles

 Library . . . .   QPFRDATA___   Name

  Type options, press Enter.
        1=Add         4=Delete              6=Print

   Opt Profile    Type Description                          Date      Time
    _ COMPLEXINQ 1   COMPLEX INQUIRY WORKLOAD            10/02/90 14:22:21
    _ SIMPLEINQ  1   SIMPLE INQUIRY WORKLOAD             10/02/90 14:22:21









  Type:         1=Measured    6=S/36 migration profile

  F3=Exit           F6=Define profile information          F12=Cancel
  F16=Get S/36 migration profile                           F24=More keys
```

Profiles created using the PRTSYSRPT command are shown on this display. If System/36 migration profiles had been previously created, these would also appear.

9. To add your new System/36 migration profile, press F16. The Get System/36 Migration Profile display appears.

```
                    Get System/36 Migration Profile

Type information, press Enter.

  Diskette device . . . . . . .  DKT01       Name

  Migration profile:
    Profile . . . . . . . . . .  MIGRATE36   Name
    Description . . . . . . . .  System/36 to System/36 Environment Example
    Library . . . . . . . . . .  QPFRDATA    Name
    Replace . . . . . . . . . .  N           Y=Yes, N=No











F3=Exit         F12=Cancel
```

10. Type the name of the diskette unit, a profile name, and descriptive text. In this example, `Migrate36` is the profile name, and `System/36 to System/36 Environment Example` is the description. Verify that the diskette is in the diskette drive. Press the Enter key. The Work with Measured Workload Profile display appears again.

```
                    Work with Measured Workload Profiles

Library . . . .   QPFRDATA     Name

 Type options, press Enter.
        1=Add         4=Delete               6=Print

  Opt Profile    Type Description                          Date      Time
   _ COMPLEXINQ 1    COMPLEX INQUIRY WORKLOAD             10/02/90 14:22:21
   _ SIMPLEINQ  1    SIMPLE INQUIRY WORKLOAD              10/02/90 14:22:21
   1 MIGRATE36  6    System/36-System/36 Environment Example 10/16/90 18:32:34








 Type:          1=Measured     6=S/36 migration profile

 F3=Exit            F6=Define profile information           F12=Cancel
 F16=Get S/36 migration profile                             F24=More keys
```

11. To select the System/36 migration profile, type a 1 (Add) in the *Opt* column next to the MIGRATE36 profile, and press the Enter key. The migration profile is now included in the response file. The Work with Response File Components display appears again.

```
                        Work with Response File Components

 Response file . . . . . :  EXAMPLE36

 Type options, press Enter.
   2=Change        4=Delete

  Opt    Component        Type      Description
  _      Functions         4        System Functions
  2      Objectives        5        Workload Objectives

  _      MIGRATE36         6        System/36-System/36 Environment Example





 Type:    1=Measured  2=Predefined  3=User  4=System functions
          5=Objectives    6=S/36 migration profile

 F3=Exit              F6=Measured profiles          F9=Add predefined profile
 F11=Planning results F12=Cancel                    F14=Growth
```

12. To enter the workload objectives for MIGRATE36, select option 2 (Change component) for the Workload Objectives and press the Enter key. The Composite Workload Objectives display appears.

```
                        Composite Workload Objectives

 Type information, press Enter.

                                                   ---------Objectives-------
             Dsp     Active      Key+Think         Response        Throughput
 Workload    Type     Dsp      Time (Seconds)    Time (Seconds)   Trans/Hour
 MIGRATE36   1       ____31        _14.8             ___.0           ___0
 MIGRATE36   2       ____ 0        _14.8             ___.0           ___0








 Dsp(Display station) type:  1=Local 2=Remote

 F3=Exit       F6=Calculate              F12=Cancel
```

This display shows the active display stations from the measured System/36 data for the MIGRATE36 profile. The calculated average key+think time is also shown.

13. To enter this input as the workload objectives for the capacity planner, press the Enter key.

    The Work with Response File Components display appears again.

```
                    Work with Response File Components

Response file . . . . . :   EXAMPLE36

Type options, press Enter.
  2=Change       4=Delete

 Opt    Component        Type       Description
  _    Functions          4        System Functions
  _    Objectives         5        Workload Objectives

  _    MIGRATE36          6        S/36 to 36 Environment Example





Type:     1=Measured  2=Predefined  3=User  4=System functions
          5=Objectives    6=S/36 migration profile

F3=Exit               F6=Measured profiles          F9=Add predefined profile
F11=Planning results  F12=Cancel                    F14=Growth
```

14. To perform capacity planning for the migration from the System/36 to the AS/400 System/36 environment, press F11. The Work with Planning Results display appears.

    Note: Throughout this capacity planning example, it is assumed that you press the Enter key to accept all upgrades recommended by the evaluator messages (required and optional).

```
                              Work with Planning Results

The current configuration meets your performance objectives. If you are
satisfied with the current configuration, you may press F14=Growth or F4=Work
with Graph Files to run or start that function.






                             ----Objectives----  ---------Predicted----------
            Dsp    Active    Response   Thruput  --Response Time---   Thruput
Workload    Type    Dsp        Time    Trans/Hr  Internal  External  Trans/Hr
Total/Avg            31                               1.1       1.2     6,987
MIGRATE36    1       31          .0          0        1.1       1.2     6,987



Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit                 F9=Next configuration        F11=Next view
F12=Cancel                                           F24=More keys
```

This display shows you the modeled results. The current configuration is a balanced system.

15. Press F9 (Next configuration) to page forward to the next display of system configuration data.

    Note: When you use F9, you only page forward on the upper half of this display.

```
                    Work with Planning Results

Type changed values for the base machine configuration, press Enter.
                              Util                                      Util
CPU model . . . . . .    D06   .57    Communications IOPs . .   0       .00
Main storage . . . . .    12 MB       Comm IOP feature . . . .
Disk IOPs . . . . . .    0     .00    Remote lines . . . . . .    0     .00
Disk arms . . . . . :    4     .29    Remote WS controller . .    0     .00
Local WS IOPs . . . .    2     .12    Remote line speed  . . .     0
Multifunction IOPs . :   1     .44
Display station multiplier:   Base . . . .    .50  Increment . . . .    .25

                           ----Objectives----  ---------Predicted----------
           Dsp    Active   Response   Thruput  --Response Time---   Thruput
Workload   Type    Dsp       Time    Trans/Hr  Internal  External  Trans/Hr
Total/Avg            31                           1.1       1.2      6,987
MIGRATE36    1       31         .0          0     1.1       1.2      6,987



Note: Performance estimates - Press the Help key to see disclaimer.
F3=Exit                 F9=Next configuration        F11=Next view
F12=Cancel                                           F24=More keys
```

The base machine configuration appears.

Through modeling, this is the anticipated performance for migrating from System/36 to the above Model D06 configuration of the AS/400 system.

Appendix E. Correlation of the System/36 and AS/400 System Performance Parameters

The differences that exist between the AS/400 system and System/36, both hardware and software, make it difficult to map performance parameters from one system to the other. Figure E-1 shows the correlation for some of the key System/36 performance parameters. This information is intended for those new AS/400 system users, who have a background in System/36 performance analysis.

Figure E-1 (Page 1 of 2). Correlation of the System/36 and AS/400 System Performance Parameters

System/36 Parameter	AS/400 Parameter
Main storage processor Control storage processor	These parameters correlate to the processor utilization reported by performance tools. See "System Report: Resource Utilization" on page 8-13 to see the format of this report. This is a portion of the System Report that is produced when you use the Print System Report (PRTSYSRPT) command.
Disk (1-4)	The System/36 disk information correlates to the disk utilization. See "System Report: Disk Utilization" on page 8-18 to see the format of this report. This is a portion of the System Report that is produced when you use the PRTSYSRPT command. There is also a Disk Activity Report that is a portion of the Component Report. You produce the Component Report when you use the Print Component Report (PRTCPTRPT) command.
Communications lines (1-8)	See "System Report: Communications Summary" on page 8-19 to see the format of this report. This is a portion of the System Report that is produced when you use the PRTSYSRPT command. More detailed information is shown in the Communications Report, which is a portion of the Component Report. You produce the Component Report when you use the PRTCPTRPT command. Details about communications lines are in a Resource Interval Report: Communications Line Summary report produced by the Print Resource Report (PRTRSCRPT) command (see "Resource Interval Report" on page 8-102 for more information). Communications lines information can be displayed by using the Display Performance Data (DSPPFRDTA) command.
Translated transfer calls/loads	Nothing on the AS/400 system correlates to these counters.

Figure E-1 (Page 2 of 2). Correlation of the System/36 and AS/400 System Performance Parameters

System/36 Parameter	AS/400 Parameter
User Area Disk Activity (UADA)	The closest correlation is shown in "System Report: Storage Pool Utilization" on page 8-17. This report is a portion of the System Report that is produced when you use the PRTSYSRPT command. The faults per second statistics represent the measure of main storage utilization. More detailed information can be found in "Storage Pool Activity" on page 8-32. This report is a portion of the Component Report that is produced when you use the PRTCPTRPT command.
Number of disk record waits	On the AS/400 system, the corresponding measure is given by seize/lock statistics. To obtain seize/lock statistics, you must collect trace data when running the performance monitor. To do this, specify TRACE(*ALL) on the Start Performance Monitor (STRPFRMON) command. This information is shown as a portion of the special summary reports that are part of the Transaction Report. To produce it, leave the OPTION(*SS) value on the PRTTNSRPT command. A more detailed analysis can also be made by using the Print Lock Report (PRTLCKRPT). See Chapter 11, "Programmer Performance Utilities" for information on this command.
Work station controller utilization	The local work station controller is included in the IOP Utilizations Report. This report is a portion of the Component Report that is produced when you use the PRTCPTRPT command. The IOP Utilizations Report includes the controller time in the response times that it shows.
File placement	File placement on the AS/400 system has little correlation to that of System/36. The AS/400 operating system, using algorithms to internally partition files, automatically spreads individual files across all available disk drives to balance disk utilization. This removes the need to do any file placement analysis.
Disk seek operations greater than one-third of the disk	Because of the file placement actions taken by the AS/400 operating system, there is essentially no user control that can be applied to this measure (other than adding auxiliary disk space). The statistic can be found as part of the Component Report.

Appendix F. Capacity Planner Office Workload Description

The office workload is characterized by several functions. A range of functions is performed by an average professional office software user and is not intended to be all-inclusive. The workload is based on the IBM Systems Application Architecture* (SAA*) OfficeVision/400 licensed program. MDLSYS only models the office workload of 5250-type display stations attached to the AS/400 system. It does not characterize the programmable-work-station-type environments.

Note: The characteristics of the office functions in Version 1 Release 3 and Version 2 Release 1 of MDLSYS changed to conform with a new version of the IBM Office Benchmark. Users of previous releases of the MDLSYS will see several differences in the office functions. (See "Capacity Planner Office Workload Descriptions Prior to Version 1 Release 3" on page F-6 for descriptions of the previous versions of these functions.)

- The names of most of the functions have changed.
- The work done by the functions has changed.
- A new function (Signoff/Signon System) has been added. This increases the number of office functions from 10 to 11.

Because of these changes, you could see a performance improvement in your office workload from previous releases of MDLSYS. You should consider changing your office workload to adjust for these new office characteristics.

Each function consists of one or more related steps, and begins and ends with the initial office display, menu, or starting position. The functions include the following:

- **Update One Calendar**. Similar to Update Calendar (Single) in the office functions prior to Version 1 Release 3.
  - View the user's calendar.
  - Add and delete a meeting.
- **Update Group Calendar**. Similar to Update Calendar (Group) in the office functions prior to Version 1 Release 3.
  - View the group's calendars.
  - Add and delete a meeting.
- **View One Calendar**. Similar to View Calendar (Single) in the office functions prior to Version 1 Release 3.
  - View the user's calendar.
- **Process Heavy Mail**. Similar to Work with Mail in the office functions prior to Version 1 Release 3.
  - Create a note and mail it to other users.
  - View note.
  - Delete note.
  - View the list of mail items received.
  - View a note, forward it to other users, and delete it.
  - Print document.

- **Process Light Mail**. Similar to Create and Send Note in the office functions prior to Version 1 Release 3.
  - Create a note and mail it to other users.
  - View note.
  - Delete note.
  - View the list of mail items received.
  - View note.
  - Delete note.
- **Create Small Document**. Similar to Word Processing Create in the office functions prior to Version 1 Release 3.
  - Get prestored document format.
  - Type in document text.
  - Save, print, and delete the document.
- **Revise Small Document**. Similar to Word Processing Simple Revise in the office functions prior to Version 1 Release 3.
  - Get an existing document.
  - Move a paragraph once.
  - Delete and insert a paragraph.
  - End the edit session.
  - Save the document.
  - Mail the changed document to other users.
- **Revise Large Document**. Similar to Word Processing Complex Revise in the office functions prior to Version 1 Release 3.
  - Get an existing document.
  - Go to a particular page.
  - Find and replace a word.
  - Find a phrase and copy a sentence.
  - Delete a sentence.
  - Go to another page.
  - Delete, get, and move different paragraphs.
  - Check spelling on a particular page.
  - Paginate, save, and delete a document.
- **View Directory Entry**. Similar to View Directory in the office functions prior to Version 1 Release 3.
  - Search the user directory for a specific user name and view the person's telephone number.
- **End/Begin Office**. Similar to End Begin Office in the office functions prior to Version 1 Release 3.
  - Leave the office main menu.
  - Wait.
  - Begin office (go back to the office main menu).
- **Signoff/Signon System**. Not in the office functions prior to Version 1 Release 3.
  - Signoff the system.
  - Wait.
  - Signon the system.

Each function may take several transactions (Enter keys or command keys) to be processed. The ratio of transactions per function is different for the various functions. See Figure F-1 for a representation of this relationship.

Figure F-1. Transactions per Office Function (Based on the IBM Office Benchmark Version 2 Office Functions)

Office Function	**Transactions per Function**
Update One Calendar	7
Update Group Calendar	17
View One Calendar	3
Process Heavy Mail	32
Process Light Mail	19
Create Small Document	15
Revise Small Document	15
Revise Large Document	56
View Directory Entry	6
End/Begin Office	2
Signoff/Signon System	3

An office profile is defined by assigning a function ratio and a key+think time to each of the functions. MDLSYS converts the function ratio to the appropriate transaction ratios. MDLSYS also displays the calculated average number of transactions per selected function mix. You can use this value to calibrate the transaction throughput information that MDLSYS calculates.

Each of the five workload types provides a set of default function ratios for the various functions. These defaults are listed in the workload type tables that follow. These default workload types span the range of typical office users, from word processing persons to managers. You can select another workload type, user defined, to define your own function ratios for the various tasks. Use this option if none of the default workload types represent your usage.

Each of the three user types (steady, interrupted, and casual) provides a default key+think time for the functions. A steady user type is word processing personnel typing at a rapid pace with few breaks. An interrupted user type performs one-half the processing of the steady user type, and a casual user type performs one-third the processing of a steady user type. You can select another user type, such as user defined, to define your own key+think times.

Figure F-2 shows office function mixes for the IBM Office Benchmark version 2 workload type.

Figure F-2. IBM Office Benchmark Workload Type (Version 2)

Office Function	**Function Ratio**	**Transaction Ratio**
Update One Calendar	4	28
View One Calendar	4	12
Process Heavy Mail	5	160
Process Light Mail	5	95
Create Small Document	2	30
Revise Small Document	2	30
View Directory Entry	6	36
End/Begin Office	1	2
Signoff/Signon System	1	3
Totals	30	396

Figure F-3 shows office function mixes for the secretarial workload type.

Figure F-3. Secretarial Workload Type (Based on the IBM Office Benchmark Version 2 Office Functions)

Office Function	**Function Ratio**	**Transaction Ratio**
Update Group Calendar	1	17
Process Heavy Mail	1	32
Revise Large Document	1	56
View Directory Entry	1	6
Totals	4	111

Figure F-4 shows office function mixes for the IBM managerial workload type.

Figure F-4. Managerial Workload Type (Based on the IBM Office Benchmark Version 2 Office Functions)

Office Function	**Function Ratio**	**Transaction Ratio**
View One Calendar	2	6
Process Heavy Mail	1	32
Process Light Mail	1	19
View Directory Entry	2	12
Totals	6	69

Figure F-5 shows office function mixes for the professional workload type.

Figure F-5. Professional Workload Type (Based on the IBM Office Benchmark Version 2 Office Functions)

Office Function	**Function Ratio**	**Transaction Ratio**
Update One Calendar	1	7
View One Calendar	1	3
Process Heavy Mail	1	32
Create Small Document	1	15
Revise Small Document	1	15
View Directory Entry	1	6
Totals	6	78

Figure F-6 shows office function mixes for the correspondence center workload type.

Figure F-6. Correspondence Center Workload Type (Based on the IBM Office Benchmark Version 2 Office Functions)

Office Function	**Function Ratio**	**Transaction Ratio**
Create Small Document	5	75
Revise Small Document	3	45
Revise Large Document	2	112
Totals	10	232

Data/Text Merge Workload Type

Data/text merge creates multiple letters by merging a shell document (a outline format of a letter into which variable information can be added) with a fill-in document (the data or information to be merged). The data/text merge job is the continuous merging and spooling of the shell document with the fill-in document. Data/text merge can be done interactively or through batch. Only the batch environment is provided by MDLSYS.

To add data/text merge to your workload, select BATCH on the Select Predefined Profiles to Add display. The Define Batch Jobs display appears. Type the number of office data/text merge jobs and press the Enter key.

Capacity Planner Office Workload Descriptions Prior to Version 1 Release 3

This section lists the office function definitions of the IBM Office Benchmark for all releases prior to Version 1 Release 3. These are provided mainly for the user who has run previous versions of MDLSYS to model office applications, and wants to compare definitions of any release prior to Version 1 Release 3 to the new definitions.

- Update Calendar (Single)
  - View the user's calendar.
  - Add and delete a meeting.
- Update Calendar (Group)
  - View a group's calendars.
  - Add and delete a meeting.
- View Calendar (Single)
  - View the user's calendar.
- Work with Mail
  - View the list of mail items received.
  - View a note and send a reply to it.
  - View a document and print it.
  - View a note and forward it to other users.
- Create and Send Note
  - Create a note and mail it to other users.
- Word Processing Create
  - Create a letter.
  - Type the inside address and block copy.
  - Type several pages of the letter.
  - Create an envelope for the letter.
  - Page through the letter.
  - Save, print, and delete the document.
- Word Processing Simple Revise
  - Get an existing document.
  - Insert and delete a paragraph.
  - Move a paragraph several times.
  - End the edit session.
  - Put the document into the user's document folder.
  - Mail the revised document to other users, and delete it.
- Word Processing Complex Revise
  - Get a document.
  - Go to a particular page.
  - Find and replace a word several times.
  - Find a phrase and copy a sentence.
  - Delete a sentence.
  - Move and delete paragraphs.
  - Go to another page.
  - Delete several sentences.
  - Delete, get, and move different paragraphs.
  - Check spelling on a particular page.
  - Paginate, save, and delete the document.
- View Directory
  - Search the user directory for a specific user name and view the person's telephone number.
- End/Begin
  - End and begin the office software program or environment.

Figure F-7 shows the ratio of transactions (Enter keys or command keys) per function.

Figure F-7. Transactions per Office Function

Office Function	Transactions per Function
Update Calendar (Single)	7
Update Calendar (Group)	17
View Calendar (Single)	3
Work with Mail	20
Create and Send Note	5
Word Processing Create	41
Word Processing Simple Revise	26
Word Processing Complex Revise	56
View Directory	6
End/Begin	2

Figure F-8 shows office function mixes for the IBM Office Benchmark workload type.

Figure F-8. IBM Office Benchmark Workload Type

Office Function	Function Ratio	Transaction Ratio
Update Calendar (Single)	2	14
View Calendar (Single)	2	6
Work with Mail	4	80
Create and Send Note	6	30
Word Processing Simple Revise	2	52
View Directory	2	12
End/Begin	1	2
Totals	19	196

Figure F-9 shows office function mixes for the secretarial workload type.

Figure F-9. Secretarial Workload Type

Office Function	Function Ratio	Transaction Ratio
Update Calendar (Group)	1	17
Work with Mail	1	20
Word Processing Complex Revise	1	56
View Directory	1	6
Totals	4	99

Figure F-10 shows office function mixes for the managerial workload type.

Figure F-10. Managerial Workload Type

Office Function	Function Ratio	Transaction Ratio
View Calendar (Single)	2	6
Work with Mail	1	20
Create and Send Note	1	5
View Directory	2	12
Totals	6	43

Figure F-11 shows office function mixes for the professional workload type.

Figure F-11. Professional Workload Type

Office Function	Function Ratio	Transaction Ratio
Update Calendar (Single)	1	7
View Calendar (Single)	1	3
Work with Mail	1	20
Word Processing Create	1	41
Word Processing Simple Revise	1	26
View Directory	1	6
Totals	6	103

Figure F-12 shows office function mixes for the correspondence center workload type.

Figure F-12. Correspondence Center Workload Type

Office Function	Function Ratio	Transaction Ratio
Word Processing Create	5	205
Word Processing Simple Revise	3	78
Word Processing Complex Revise	2	112
Totals	10	395

Glossary

This glossary includes terms and definitions from:

- The *American National Dictionary for Information Processing Systems*, copyright 1982 by the Computer and Business Equipment Manufacturers Association (CBEMA). Copies may be purchased from the American National Standards Institute, 1430 Broadway, New York, New York 10018. Definitions are identified by the symbol (A) after the definition.
- The *Information Technology Vocabulary*, developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Committee (ISO/IEC JTC1/SC1). Definitions of published segments of the vocabularies are identified by the symbol (I) after the definition; definitions from draft international standards, draft proposals, and working papers in development by the ISO/IEC JTC1/SC1 vocabulary subcommittee are identified by the symbol (T) after the definition, indicating final agreement has not yet been reached among participating members.

active group job. A group job that was not suspended by the Transfer to Group Job (TFRGRPJOB) command.

active record. An active subfile record or any record format that is currently shown on a display. Contrast with *inactive record*.

activity level. A characteristic of a subsystem that specifies the maximum number of jobs that can compete at the same time for the processing unit.

advisor. In Performance Tools/400, a tool used to analyze data collected by the Start Performance Monitor (STRPFRMON) command. The advisor analyzes a performance data collection and produces a list of recommendations and conclusions to improve system performance.

American National Standard Code for Information Interchange (ASCII). The code developed by the American National Standards Institute for information exchange among data processing systems, data communications systems, and associated equipment. The ASCII character set consists of 7-bit control characters and symbolic characters, plus one parity bit.

application. (1) A particular business task, such as inventory control or accounts receivable. (2) In CSP/AE, the collection of CSP/AE objects that together can be run on an AS/400 system. An application consists of a program object, up to five map group objects (depending on how many different devices are supported), and any number of table objects. (3) In SQL, a program or set of programs that perform a task; for example, a payroll application.

area fill. The shading used inside a bar on a graph and below the lines of a surface graph.

AS/400 Business Graphics Utility Version 2 (BGU). See *IBM AS/400 Business Graphics Utility Version 2 (BGU)*.

AS/400 Office. See *IBM SAA OfficeVision/400 Version 2*.

AS/400 PC Support. See *IBM PC Support/400 Version 2*.

ASCII. See *American National Standard Code for Information Interchange (ASCII)*.

ASP. See *auxiliary storage pool (ASP)*.

asynchronous communications. A method of communications supported by the operating system that allows an exchange of data with a remote device, using either a start-stop line or an X.25 line. Asynchronous communications includes the file transfer support and the interactive terminal facility support.

asynchronous disk I/O. In Performance Tools/400, a disk access operation that is not expected to complete before program operation can continue. Contrast with *synchronous disk I/O*.

asynchronous I/O. A series of input/output operations that are being done separately from the job that requested them.

attribute. (1) A characteristic or trait of one or more items. (2) In AS/400 Business Graphics Utility, the characteristics that make up the chart format. (3) In an SQL database design, a characteristic of an entity; for example, the telephone number of an employee is one of that employee's attributes.

authority lookups. In Performance Tools/400, the process whereby the licensed internal code determines whether a particular user ID is authorized to access a specific object.

autostart job. A job doing repetitive work or one-time initialization work that is associated with a particular subsystem. The autostart jobs associated with a subsystem are automatically started each time the subsystem is started.

auxiliary storage. All addressable disk storage other than main storage. Contrast with *main storage*.

auxiliary storage pool (ASP). A group of units defined from the disk units that make up auxiliary storage.

ASPs provide a means of isolating certain objects on specific disk units to prevent the loss of data due to disk media failures on other disk units. See also *system ASP* and *user ASP*.

axis. In AS/400 Business Graphics Utility and GDDM, one of the intersecting horizontal or vertical scales where data values are plotted on a chart.

bar graph. In Performance Tools/400, a graph consisting of several bars of equal width. The value of the dependent variable is indicated by the height of each bar.

base pool. A storage area that contains all unassigned main storage on the system and whose minimum size is specified in the system value QBASPOOL. The system-recognized identifier is *BASE.

batch. Pertaining to a group of jobs to be run on a computer sequentially with the same program with little or no operator action. Contrast with *interactive*.

batch device. Any device that can read serial input or write serial output, or both, but cannot be used to communicate interactively with the system. Examples of batch devices are printers, magnetic tape units, or diskette units.

batch job. A predefined group of processing actions submitted to the system to be performed with little or no interaction between the user and the system. Contrast with *interactive job*.

batch subsystem. A part of main storage where batch jobs are processed.

BGU. See *IBM AS/400 Business Graphics Utility Version 2 (BGU)*.

binary synchronous communications (BSC). A data communications line protocol that uses a standard set of transmission control characters and control character sequences to send binary-coded data over a communications line. See also *synchronous data link control (SDLC)*.

bps. Bits per second.

BSC. See *binary synchronous communications (BSC)*.

buffer. (1) A routine or an area of storage that corrects for the different speeds of data flow or timings of events, when transferring data from one device to another. (2) A portion of storage used to hold input or output data temporarily.

bus. One or more conductors used for transmitting signals or power.

Business Graphics Utility (BGU). See *IBM AS/400 Business Graphics Utility Version 2 (BGU)*.

character field. An area that is reserved for information that can contain any of the characters in the character set. Contrast with *numeric field*.

chart. In AS/400 Business Graphics Utility, displayed, printed, or plotted output that compares one or more sets of variable data in chart form. The types of charts are bar, line, pie, surface, histogram, Venn diagram, and text.

checksum protection. (1) A function that protects data stored in the system auxiliary storage pool from being lost because of the failure of a single disk. When checksum protection is in effect and a disk failure occurs, the system automatically reconstructs the data when the system program is loaded after the device is repaired. (2) In TCP/IP, the sum of a group of data associated with the group and used for error checking purposes.

checksum set. Units of auxiliary storage defined in groups to provide a way for the system to recover data if a disk failure occurs when checksum protection is in effect.

CL. See *control language (CL)*.

command. (1) A statement used to request a function of the system. A command consists of the command name abbreviation, which identifies the requested function, and its parameters. (2) In SDLC, a frame transmitted by a primary station. Asynchronous balanced mode stations send both commands and responses. Contrast with *response*. (3) In SNA, any field set in the transmission header (TH), request header (RH), or request unit (RU) that states an action or that starts a protocol.

communications configuration. The physical placement of communications controllers, the attachment of communications lines, and so forth; and the configuration descriptions that describe the physical configuration to the system and describe how the configuration will be used by the system. See also *line configuration*, *controller configuration*, and *device configuration*.

communications controller. The I/O processor card in the card enclosure.

communications line. The physical link (such as a wire or a telephone circuit) that connects one or more work stations to a communications controller, or connects one controller to another.

composite bar graph. In Performance Tools/400, a bar graph in which multiple vertical axis values for the same horizontal axis value are stacked one on top of another. See also *floating bar graph*.

computer graphics. The use of a computer to produce images of relationships, such as charts, and two- or three-dimensional drawings by means of dots, lines, curves, and so forth.

configuration. The physical and logical arrangement of devices and programs that make up a data processing system. See also *communications configuration*, *line configuration*, *controller configuration*, and *device configuration*.

configure. (1) To describe the interconnected arrangement of the devices, programs, communications, and optional features installed on a system. (2) To describe setting up auxiliary storage pools and checksum protection.

control language (CL). The set of all commands with which a user requests system functions.

controller. A device that coordinates and controls the operation of one or more input/output devices (such as work stations) and synchronizes the operation of such devices with the operation of the system as a whole.

controller configuration. The process of creating configuration descriptions for the local (device configuration) and remote (communications configuration) controllers that make up a data processing system. See also *line configuration* and *device configuration*.

data area. A system object used to communicate data, such as CL variable values between the programs within a job and between jobs. The system-recognized identifier for the data area is *DTAARA.

database. All the data files stored in the system.

database file. One of several types of the system object type *FILE kept in the system that contains descriptions of how input data is to be presented to a program from internal storage and how output data is to be presented to internal storage from a program. See also *physical file* and *logical file*.

DDM. See *distributed data management (DDM)*.

DDM server. A job on the host system that processes DDM requests from other systems by accessing the database on the host system on behalf of other jobs on other systems.

default. (1) A value that is automatically supplied or assumed by the system or program when no value is specified by the user. (2) In DDS, the value specified by the user with the DFT or DFTVAL keyword in DDS. (3) In SQL, a predetermined value, attribute, or option that is supplied by the system when no value is specified by the user. For example, the default of a column is blanks if the data type is character, or zeros if the data type is numeric.

device configuration. The physical placement of display stations, printers, and so forth; and the configuration descriptions that describe the physical configuration to the system and describe how the configuration will be used by the system. See also *line configuration* and *controller configuration*.

device file. One of several types of the system object type *FILE. A device file contains a description of how data is to be presented to a program from a device or how data is to be presented to the device from the program. Devices can be display stations, printers, diskette units, tape units, or remote systems.

disk. A direct-access storage medium with magnetically recorded data.

disk drive. A device for moving and controlling the disk.

diskette. A thin, removable magnetic disk in a protective jacket.

display file. (1) A device file to support a display station. (2) In BASIC, any file that has the keyword DISPLAY specified in the OPEN statement for the file.

display station. A device that includes a keyboard from which an operator can send information to the system and a display screen on which an operator can see the information sent to or the information received from the system.

display station pass-through. A communications function that allows a user to sign on to one system (either an AS/400 system, System/38, or System/36) from another system (either an AS/400 system, System/38, or System/36) and use that system's programs and data. Sometimes called pass-through.

distributed data management (DDM). A function of the operating system that allows an application program or user on one system to use database files stored on remote system. The systems must be connected by a communications network, and the remote systems myst also be using DDM.

dynamic. Pertaining to events occurring at run time, or during processing.

EAO exception. See *effective address overflow (EAO) exception*.

effective address overflow (EAO) exception. In Performance Tools/400, a condition in which the licensed internal code must make address adjustments not made above the machine interface.

Ethernet. A type of local area network that is supported by the Operating System/400 licensed program. OS/400 Ethernet provides support for the Digital Equipment Corporation, Intel Corporation, and Xerox standard (Ethernet Version 2) and the IEEE 802.3 standard. These local area networks use

Carrier Sense Multiple Access with Collision Detection (CSMA/CD) as the media access method.

exceptional wait. In Performance Tools/400, 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.

FIFO. See *first-in first-out (FIFO)*.

finance device. A device, such as the 4700 Finance Communications System devices and the 3694 Document Processor, that performs functions specifically related to the finance industry. The 3180, 3270, and 5250 work stations are not finance devices.

first-in first-out (FIFO). A queuing technique in which the next request to be processed from a queue is the request of highest priority that has been on the queue for the longest time.

floating bar graph. In Performance Tools/400, a graph that shows bars detached from either line. See also *composite bar graph*.

format. (1) A defined arrangement of such things as characters, fields, and lines, usually used for displays, printouts, files, or documents. (2) The arrangement or layout of fields in a record. (3) The arrangement or layout of data on a storage medium, such as disk, tape, or diskette. (4) To set the block size for the 9332 Disk Unit, either automatically by the system or specifically by the user. (5) In BASIC, a representation of the correct form of a command or statement. (6) To arrange information on a page, in a file, or on a display screen. (7) In PC Support/400, to prepare a diskette so that it can be used by a personal computer.

GDDM. See *graphical data display manager (GDDM)*.

GDF file. See *graphics data format (GDF) file*.

graph. (1) See *chart*. (2) In Performance Tools/400, the displayed, printed, or plotted output that represents the horizontal and vertical axis variables specified by the user for a collection of performance data.

graphical data display manager (GDDM). A function of the operating system that processes both text and graphics for output on a display, printer, or plotter.

graphics. (1) Pictures and illustrations. (2) Pertaining to charts, tables, and their creation. See also *computer graphics*.

graphics data format (GDF) file. A picture definition in a coded order format used internally by GDDM and, optionally, providing the user with a lower-level programming interface than the GDDM application programming interface.

hardware. Physical equipment, rather than programs, procedures, rules, and associated information.

HDLC. See *high-level data link control (HDLC)*.

high-level data link control (HDLC). A form of communications line control that uses a specified series of bits rather than control characters to control data transmission over a communications line.

I/O. See *input/output*.

IBM AS/400 Business Graphics Utility Version 2 (BGU). The IBM licensed program that can be used to design, plot, display, and print business charts.

IBM office benchmark. In Performance Tools/400, an IBM-specified mixture of office tasks (for example, viewing a calendar, working with mail, working with documents) that is used to compare performance measurements on different systems. The mixture of these tasks remains constant from system to system (such as from a System/38 to an AS/400 system), but the way in which the tasks are carried out changes from system to system.

IBM Operating System/400 Version 2 (OS/400). Pertaining to the IBM licensed program that can be used as the operating system for the AS/400 system.

IBM PC Support/400 Version 2. The IBM licensed program that provides system functions to an attached personal computer.

IBM Performance Tools/400 Version 2. The IBM licensed program that allows a user to collect performance data, display performance data, print performance reports, and perform capacity planning.

IBM SAA OfficeVision/400 Version 2. The IBM licensed program that allows users to prepare, send, and receive mail; schedule items on calendars; maintain directories of names and addresses; file and retrieve docuemnts; and create and maintain distribution lists. SAA OfficeVision/400 also provides word processing functions and the capability to work on behalf of other users.

inactive record. An inactive subfile record or any record format that is not currently shown on a display. Contrast with *active record*.

independent work station. See *programmable work station (PWS)*.

initial program load (IPL). The process that loads the system programs from the system auxiliary storage, checks the system hardware, and prepares the system for user operations.

input/output. Data provided to the computer or data resulting from computer processing.

input/output processor (IOP). A functional unit or the part of an I/O controller that processes programmed instructions and controls one or more input/output devices or adapters.

interactive. Pertaining to the dialog-like exchange of information between people and a computer. Contrast with *batch*.

interactive job. A job started for a person who signs on to a work station. Contrast with *batch job*.

IOP. See *input/output processor (IOP)*.

IPL. See *initial program load (IPL)*.

job. (1) A unit of work to be done by a computer. (2) In the SAA OfficeVision/400 calendar function, an item that schedules a control language (CL) command to run at any date and time.

job log. A record of requests submitted to the system by a job, the messages related to the requests, and the actions performed by the system on the job. The job log is maintained by the system program.

join logical file. A logical file that combines (in one record format) fields from two or more physical files. See also *logical file*.

LAN. See *local area network (LAN)*.

legend. In AS/400 Business Graphics Utility and GDDM, an explanatory list of the symbols, lines, and shaded areas on a chart.

library. (1) A system object that serves as a directory to other objects. A library groups related objects, and allows the user to find objects by name. The system-recognized identifier for the object type is *LIB. (2) The set of publications for a system.

licensed internal code. The layered architecture below the machine interface (MI) and above the machine, consisting of the model-independent and the model-unique licensed internal code. The licensed internal code carries out many functions, including, but not limited to storage management, pointers and addressing, program management functions, exception and event management, data functions, I/O managers, security functions, and proprietary system design information.

licensed program. A separately orderable program, supplied by IBM, that performs functions related to processing user data. Examples of licensed programs are PC Support/400, SAA COBOL/400, AS/400 Application Development Tools, SAA OfficeVision/400, and so on.

line configuration. The process of creating configuration descriptions for the lines that make up a data processing system. See also *controller configuration* and *device configuration*.

line graph. In Performance Tools/400, a graph in which plotted points (each optionally represented by a marker) are joined by straight or curved lines. See also *scatter plot*.

local area network (LAN). The physical connection that allows the transfer of information among devices located on the same premises.

local work station. A work station that is connected directly to the system without a need for data transmission functions.

lock. The process by which integrity of data is ensured by preventing more than one user from accessing or changing the same data or object at the same time.

lock state. A condition defined for an object that determines how it is locked, how it is used (read or write), and whether the object can be shared (used by more than one job).

logical file. A description of how data is to be presented to or received from a program. This type of database file contains no data, but it defines record formats for one or more physical files. See also *join logical file* and *database file*. Contrast with *physical file*.

logical I/O. In Performance Tools/400, the operation used to pass a buffer of data from data management to the data management code of an application program.

main storage. All addressable storage where programs are run. Contrast with *auxiliary storage*.

main storage pool. A division of main storage, which allows the user to reserve main storage for processing a job or group of jobs, or to use the pools defined by the system. Contrast with *auxiliary storage pool*.

MB. See *megabyte*.

megabyte. A unit of measure for storage capacity. For main storage, 1 megabyte = 1 048 576 bytes (1024 x 1024); for auxiliary storage (disk, diskette, and tape), 1 megabyte = 1 000 000 bytes (1000 x 1000).

member. Different sets of data, each with the same format, within one database file.

migration. The process of moving data and source from one computer system to another without converting the data.

mirrored pair. Two units that contain the same data and are referred to by the system as one entity (one unit).

numeric field. An area that is reserved for a particular unit of information and that can contain only the digits 0 through 9. Contrast with *character field*.

ODP. See *open data path (ODP)*.

Office. See *IBM SAA OfficeVision/400 Version 2*.

OfficeVision/400. See *IBM SAA OfficeVision/400 Version 2*.

open. The function that connects an object of type *FILE to a program for processing.

open data path (ODP). A control block created when a file is opened. An ODP contains information about the merged file attributes and information returned by input or output operations. The ODP only exists while the file is open.

operating system. A collection of system programs that control the overall operation of a computer system.

Operating System/400. See *IBM Operating System/400 Version 2 (OS/400)*.

OS/400. See *IBM Operating System/400 Version 2 (OS/400)*.

overlay. In Performance Tools/400, a graph that is placed on top of another graph so that a user can view both graphs at the same time.

packed field. A field that contains data in the packed decimal format.

PAG. See *process access group (PAG)*.

page fault. An exception that occurs when a program refers to data or programs that are marked as not in main storage.

PASA. See *program automatic storage area (PASA)*.

pass-through. See *display station pass-through*.

PC Support. See *IBM PC Support/400 Version 2*.

performance monitor. A function of the operating system that observes system and device activity, and records these observations in a database file.

physical disk I/O. In Performance Tools/400, a disk operation for reading or writing data.

physical file. A description of how data is to be presented to or received from a program and how data is actually stored in the database. A physical file contains one record format and one or more members. See also *database file*. Contrast with *logical file*.

plot. In GDDM, to represent graphically on paper using a plotting device.

plotter. In AS/400 Business Graphics Utility, a device for drawing a chart on paper or transparencies.

pool. A division of main or auxiliary storage. See also *base pool* and *storage pool*.

pool database faults. In Performance Tools/400, the total number of interruptions to jobs that were required to transfer data into the pool to permit the program to process the database data.

pool database pages. In Performance Tools/400, the total number of pages of database data transferred from auxiliary storage to the pool to permit the program to run.

pool nondatabase faults. In Performance Tools/400, the total number of interruptions to jobs (not necessarily assigned to this pool) that were required to transfer data into the pool to permit the machine interface instruction to access the nondatabase data.

pool nondatabase pages. In Performance Tools/400, the total number of pages of nondatabase data transferred from auxiliary storage to the pool to permit the program to run.

previous release. The last required release of the system (such as Release 1.0) prior to the current release (such as Release 2.0), including any modification levels (such as Release 1.0 Modification Level 1 or Modification Level 2) that were not required.

process access group (PAG). A group of job-related objects that may be paged in and out of storage in a single operation when a job (process) enters or leaves a long wait.

program automatic storage area (PASA). A system object that contains call level information for each program on the program stack. The PASA can also contain space for program variables, which is allocated when the program object is called.

program static storage area (PSSA). A system object that contains static variable data for programs on the program stack. The PSSA contains space for program variables that is activated when the program object is activated. The PSSA is contained in the process access group (PAG).

program temporary fix (PTF). A temporary solution to, or bypass of, a defect in a current release of a licensed program.

programmable work station (PWS). A work station that has some degree of processing capability and allows the user to change its functions.

PSSA. See *program static storage area (PSSA)*.

PTF. See *program temporary fix*.

purge. In Performance Tools/400, a job attribute that specifies whether a job is to be marked eligible to be moved out of main storage to auxiliary storage when entering a long wait or leaving the activity level.

queue. A list of messages, jobs, files, or requests waiting to be read, processed, printed, or distributed in a predetermined order.

read operation. An input operation that obtains data from a file or device and passes it to a program.

record. A group of related data, words, or fields treated as a unit, such as one name, address, and telephone number.

relative record number. A number that specifies the relationship between the location of a record and the beginning of a database file, member, or subfile. For example, the first record in a database file, member, or subfile has a relative record number of 1.

remote. Pertaining to a device, system, or file that is connected to another device, system, or file through a communications line.

RM/COBOL-85 for the AS/400 Version 2. The IBM licensed program that supports the American National Standard COBOL X3.23-1985 Standard and the American National Standard COBOL X3.23-1974 Standard, and includes unique extensions that aid in writing interactive work station applications. This licensed program is the AS/400-specific version of the RM/COBOL-85 implementation, a high-level language available on a number of computing systems.

scatter plot. In Performance Tools/400, a variety of line graph in which only the marked points, and not their joining lines, are drawn. See also *line graph*.

SDLC. See *synchronous data link control (SDLC)*.

server. (1) A computer that shares its resources with other computers in the network. (2) In a local area network, a data station that provides services to other data stations.

shared file. A file whose open data path can be shared between two or more programs processing in the same job. See *open data path (ODP)*.

SNA. See *Systems Network Architecture (SNA)*.

spool. The system function of putting files or jobs into disk storage for later processing or printing.

spooled file. A file that holds output data waiting to be processed, such as information waiting to be printed. Also known as *spooled output file*.

spooled output file. See *spooled file*.

spooling writer. The general name to refer to the function of the diskette writer and printer writer.

storage pool. A logical division of storage reserved for processing a job or group of jobs.

subsystem. An operating environment, defined by a subsystem description, where the system coordinates processing and resources.

surface graph. In Performance Tools/400, a graph similar to a line graph, except that no markers appear, and the areas between successive lines are shaded.

synchronous data link control (SDLC). (1) A form of communications line control that uses commands to control the transfer of data over a communications line. (2) A communications discipline conforming to subsets of the Advanced Data Communication Control Procedures (ADCCP) of the American National Standards Institute (ANSI) and High-Level Data Link Control (HDLC) of the International Standards Organization (ISO), for transferring synchronous, code-transparent, serial-by-bit information over a communications line. Transmission exchanges may be duplex or half-duplex over switched or nonswitched lines. The configuration of the connection may be point-to-point, multipoint, or loop. Compare with *binary synchronous communications (BSC)*.

synchronous disk I/O. In Performance Tools/400, a disk access operation that must complete before program operation can continue. Contrast with *asynchronous disk I/O*.

system ASP. The auxiliary storage pool where system programs and data reside. It is the storage pool used if a storage pool is not defined by the user. See also *auxiliary storage pool* and *user ASP*.

system configuration list. A list of devices that are provided with the system.

Systems Network Architecture (SNA). In IBM networks, the description of the layered logical structure, formats, protocols, and operational sequences that are used for transmitting information units through networks, as well as controlling the configuration and operation of networks.

system object. A machine object classification. Any of the machine objects shipped with the system or any of the operating system objects created by the system.

throughput. (1) The measure of the amount of work performed by a computer over a period of time, for example, number of jobs per day. (2) In data communications, the total traffic between stations over a period of time.

time slice. The amount of processor time (specified in milliseconds) allowed for a job before other waiting jobs of equal priority are allowed to process data.

token-ring network. A local area network that sends data in one direction throughout a specified number of locations by using the symbol of authority for control of the transmission line, called a token, to allow any sending station in the network (ring) to send data when the token arrives at that location.

TRLAN. Abbreviation in the commands, parameters, and options for IBM Token-Ring Network. See also *token-ring network*.

twinaxial cable. A cable made of two twisted wires inside a shield that is used on the 5250 family devices.

| **user ASP**. One or more auxiliary storage pools used to isolate some system objects from the other system objects stored in the system ASP. See also *auxiliary storage pool (ASP)* and *system ASP*.

user ID. See *user identification (user ID)*.

user identification (user ID). (1) The name used to associate the user profile with a user when a user signs on the system. (2) The first part of a two-part network name used in the system distribution directory and in the office applications to uniquely identify a user. The network name is usually the same as the user profile name, but does not need to be.

window. A part of the display screen with visible boundaries in which information is displayed.

work station. A device used to transmit information to or receive information from a computer; for example, a display station or printer.

work station controller (WSC). An I/O controller card in the card enclosure that provides the direct connection of local work stations to the system.

| **writing**. The action of making a recording of data on an external storage device or other data medium.

WSC. See *work station controller (WSC)*.

| **X.25**. A CCITT Recommendation that defines the physical level (physical layer), link level (data link layer), and packet level (network layer) of the OSI reference model. An X.25 network is an interface between data terminal equipment (DTE) and data circuit-terminating equipment (DCE) operating in the packet mode, and connected to public data networks by dedicated circuits. X.25 networks use the connection-mode network service.

3270 display emulation. The function of the operating system 3270 device emulation support that converts 3270 data streams intended for a 3278 display station into data streams that can be recognized by a display station attached to the AS/400 system.

5250 display station. Any display station, attached by twinaxial cable, that uses 5250 data streams.

Index

D

F

G

K

L

T

U

Readers' Comments

Application System/400™
Programming:
Performance Tools/400 Guide
Version 2

Publication No. SC41-8084-00

Use this form to tell us what you think about this manual. If you have found errors in it, or if you want to express your opinion about it (such as organization, subject matter, appearance) or make suggestions for improvement, this is the form to use.

To request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should talk to your IBM representative or to your IBM authorized remarketer. This form is provided for comments about the information in this manual and the way it is presented.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Be sure to print your name and address below if you would like a reply.

Name

Address

Company or Organization

Phone No.

Readers' Comments
SC41-8084-00

Cut or Along

Fold and Tape

Please do not staple

Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

ATTN DEPT 245
IBM CORPORATION
3605 HWY 52 N
ROCHESTER MN 55901-7899

Fold and Tape

Please do not staple

Fold and Tape

SC41-8084-00

Cut or Along

Readers' Comments

Application System/400™
Programming:
Performance Tools/400 Guide
Version 2

Publication No. SC41-8084-00

Use this form to tell us what you think about this manual. If you have found errors in it, or if you want to express your opinion about it (such as organization, subject matter, appearance) or make suggestions for improvement, this is the form to use.

To request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should talk to your IBM representative or to your IBM authorized remarketer. This form is provided for comments about the information in this manual and the way it is presented.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Be sure to print your name and address below if you would like a reply.

Name	Address
Company or Organization	
Phone No.	

Readers' Comments
SC41-8084-00

IBM®

Cut
Alo

Fold and Tape | **Please do not staple** | Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

ATTN DEPT 245
IBM CORPORATION
3605 HWY 52 N
ROCHESTER MN 55901-7899

Fold and Tape | **Please do not staple** | Fold and Tape

SC41-8084-00

Cut
Alo

IBM®

Program Number: 5738-PT1

Printed in Denmark by N. Olaf Møller

SC41-8084-00