CA Performance Management for OpenVMS

Performance Agent Administrator Guide

r3.1

Contents

Chapter 10: Parameter Editor Commands 155

Chapter 11: Use the DECwindows Motif Interface 191

Chapter 1: Overview

This chapter introduces CA Performance Management for OpenVMS products and describes how the Performance Agent supports them.

This section contains the following topics:

CA Performance Management for OpenVMS Products

CA Performance Management for OpenVMS products are OpenVMS layered products designed to reduce the time and effort required to manage and monitor OpenVMS system performance and to plan for future resource requirements. These products can be used with standalone systems and with OpenVMS clusters.

CA Performance Management for OpenVMS includes Performance Manager and Performance Agent, which share a common database and basic set of utilities. Any one component may provide these utilities on behalf of the other components of the same version.

Performance Manager

The Performance Manager makes recommendations for improving system performance. It does this by analyzing system data through the application of expert system technology, identifying specific conditions causing performance degradation, and presenting detailed evidence to support its conclusions. It also provides real-time displays of performance data using either DECwindows Motif or character cell displays. The user can interactively view and investigate system performance problems and resource usage. The following functions are included in the Performance Manager:

- Performance knowledge base and rules compiler
- Performance analysis and reports
- Real-time displays of performance data
- Graphing
- Data Exporting facility
- System resource limitations identified when they exist for the workload

- System parameter settings identified that may be adding to system overhead or degrading system performance
- Trends evaluated in system performance
- Effects of changes in workload and configuration characteristics evaluated

To deal with system performance effectively you must understand the workload and the capabilities and limitations of system resources. Generally, any attempt to improve system performance requires specific performance goals stated in measurable terms.

Performance Agent

The Performance Agent supports the Performance Manager by gathering, managing, classifying, and archiving OpenVMS system data, and providing the following functions:

- Data collection and storage
- Data archiving
- Dump reports
- Disk analysis
- PC sampling
- Real-time file activity display

The Performance Manager analyzes the collected data to produce a report that helps the system manager monitor system activities and make performance evaluations by identifying performance problems quickly. This report also identifies potential bottlenecks, including the specific device on which the bottleneck occurs.

The Performance Manager provides a facility to convert the data collected in the Performance Manager database into graphs. These pictures of system performance metrics are a robust source of information that enables a system manager to gain a better understanding of resource utilization and overall workload demands. By defining families of workloads with the Performance Agent's parameter editor, the graphs can represent site-specific user groups.

Performance Agent supports the Performance Manager with the following features:

- Provides support for stand-alone OpenVMS systems as well as clusters
- Allows automatic collection of system performance data easily defined by the user

- Archives data for presentation of any length of time in both report and graph format
- Groups users into workloads and families of workloads to allow for the reporting and modeling of system resources in terms of departments, projects, or any site-specific need
- Collects and reports system program counter data to analyze CPU activity
- Collects and reports disk space use
- Displays file activity on a disk

Intended Audience

The Performance Agent software and this guide are designed specifically for the following users of OpenVMS systems:

- System managers
- System programmers
- Computer operators
- Software Services personnel
- Field Service engineers involved with system management

This guide assumes that the user has knowledge of the following products:

- OpenVMS operating system
- OpenVMS system management

Conventions

The following conventions are used in this guide:

Convention	Meaning
UPPERCASE	Words in uppercase indicate a command, the name of a file, or an abbreviation for a system privilege.
lowercase	In format descriptions, words in lowercase indicate parameters or arguments you must supply.
user input system prompts system displays	In interactive examples, this typeface indicates input entered by the user, a system prompt, or displayed system text.

Convention	Meaning
. . .	Vertical ellipsis points indicate the omission of information from an example or command format. The information is omitted because it is not important to the topic being discussed.
[ ]	In format descriptions, brackets indicate optional elements. You can choose none, one, or all of the options. (Brackets are not optional, however, in the syntax of a directory name in a OpenVMS file specification.)
, . . .	In format descriptions, horizontal ellipsis points indicate additional parameters, values, or other information that you can enter separated by commas.
{ }	In format descriptions, braces indicate required elements. You must choose one of the elements.
|	In format descriptions, a vertical bar indicates required elements. You must choose one of the elements.
( )	In format descriptions, parentheses delimit the parameter or argument list.
Ctrl+X	In procedures, a sequence such as Ctrl+X indicates that you must press the key labeled Ctrl while you press the key or a pointing device button.
MB1, MB2, MB3	MB1 indicates the left mouse button. MB2 indicates the middle mouse button. MB3 indicates the right mouse button.(You can redefine the mouse buttons.)

Chapter 2: Control Data Collection

This chapter describes the following procedures:

- How the Performance Agent software operates
- How to tailor data collection to your own needs
- How to provide for Performance Real-time reporting

This chapter includes the following topics:

- Explain Data Collection
- Control Data Collection
- Explain the Primary Performance Agent
- Explain Alternate Performance Agents
- Support Real-time Performance Agents
- Customize Collection Definitions
- Start the Data Collection Processes
- Use the Database Directory

This section contains the following topics:

Data Collection

Running as a detached process, a Performance Agent collects a variety of system and process data. The software uses both sampling and event-driven techniques in collecting data.

More than one Performance Agent can be running and collecting data on a system. The primary Performance Agent must be running before you start any alternate Performance Agents. Sampling the data produced by the primary Performance Agent, alternate Performance Agents act as clients.

When the Performance Agent software is installed a *schedule file* is created which contains parameters governing the primary Performance Agent. These parameters comprise the *collection definition* named CPD. Use the ADVISE COLLECT SHOW ALL command to display the schedule file. For an example display, see Examining the Schedule File in this chapter.

The use of alternate Performance Agents enables you to select the *classes* of data you want to collect and the interval at which you want to collect the data. You customize data collection by creating or modifying a collection definition in the schedule file. A collection definition, its associated parameters, and data are identified by the name you give to the definition.

You can create alternate collectors using the ADVISE COLLECT ADD command or the DECwindows interface. The collection definitions you create will then appear in the schedule file. For examples of a schedule file, see Examining the Schedule File in this chapter.

A *Real-time* Performance Agent is a special Performance Agent for dynamic viewing. There are two kinds of Real-time Performance Agents. One type of Real-time Performance Agent is started on a remote or local node using DECnet. The other type accesses the data currently being recorded to a disk file either by the primary Performance Agent or by a specified alternate Performance Agent. If the data is currently being recorded, you can view this data in real time, but you must have the SYSLCK privilege. See the chapters covering Real-time displays in the *Performance Manager Administrator Guide* and the *SDK Guide*.

Required Privileges

You need the following privileges to run a data collection process:

- IMPERSONATE
- SYSLOCK
- CMKRNL
- SYSPRV
- SYSNAM
- WORLD
- EXQUOTA
- GRPNAM
- PSWAPM
- ALTPRI

Non-paged Pool Requirements

The primary data collection process allocates non-paged pool resources depending on the number of processes that the system supports and the number of disk drives. For details on setting the non-paged pool requirements, see the *Installation Guide*.

The primary data collection process allocates resources in the following amounts as shown in the following table:

Bytes	Allocations
38400	(1)
Multiply MAXPROCESSCNT by 480	(1)
392	As needed on CPD process image queue overflow (2)
Multiply MAXPROCESSCNT by 4	(1) per active disk
Multiply maximum configured disks by 184	(1)
7680	As needed on VMS I/O event buffer overflow (1)

(1) I/O buffer overruns may occur in heavily I/O-intensive environments. See the appendix "Error Messages and Recovery Procedures" for descriptions of the logicals PSDC$IO_FLUSH_FREQUENCY and PSDC$IO_PACKET_MAX. These logicals control the frequency of overruns and specify the maximum number of non-paged pool allocations due to overruns, respectively.

(2) Process image queue overruns increase as the system image activation rate increases.

Control Data Collection

You control data collection through a *collection definition*. A collection definition is a set of parameters or attributes describing the following data:

- Which data to record
- When to collect the data
- Where to store the data
- How long to keep the data

Options for Collecting Data

You have many options for setting the data collection parameters in a collection definition. Some of these options are different for the CPD (primary) collection definition than for the alternate collection definitions. The following table describes the parameters and their default values:

Parameter	Default Value
Schedule Specifies the weekly schedule for data collection for all nodes in a cluster.	7 days per week, 24 hours
Beginning Specifies the time when the data collection process is to start collecting data. The primary Performance Agent will not serve data for use by alternate or Real-time Performance Agents until the BEGINNING time has past.	For the CPD, the default value is the date on which you installed the software. For alternate collection definitions, the default is today's date.
Ending Specifies the date and time on which the data collection process is to stop.	January 1, 2999
Nodes Specifies the list of the nodes on which data is to be collected	For the CPD definition, the default list includes all the nodes. For alternates, you must specify nodes.
Hot files queue length Determines whether hot file data is collected for the CPD collection definition. The value is a minimum threshold for the queue length of a disk that triggers the recording of hot file data.	For the CPD, 0.33. Alternates do not collect this data.
Interval Specifies how often the data collection data is summarized into a set of data records for reporting.	For the CPD, fixed at 120 seconds. You can specify this for alternate collection definitions.
Delete after Specifies the retention period for how long files are retained in the database before being automatically deleted.	7 days for the CPD 99,999 days for alternates

Parameter	Default Value
Classes The Classes parameter specifies which daily data is recorded. You can choose from process, I/O, and system metrics. See Selecting Data by Classes in this chapter.	All classes
Include or Exclude Lists Specify inclusion and exclusion lists to include or exclude data on specific processes, PIDs, devices, or volumes.	Created by user
Data base directory path Specifies the path to the location for data files for an individual node or entire collection definition.	PSDC$DATABASE Can be specified by individual node
Working set Specifies the working set quota (WSQUOTA)for the process that executes the Performance Agent. The number must be greater than 1024.	2048 Can be specified for a collection definition or by individual node
Free space Specifies a threshold value that is the minimum number of free blocks needed on the database disk for recording data. When the number of free blocks drops below this minimum, a Performance Agent will hibernate (stop recording data). Once the free space becomes available, data recording will resume.	2500 Can be specified by individual node

Select Data by Classes

The class parameter enables you to establish the classes of data that are to be collected. By default all classes (ALL_METRICS) of data are collected. The choices include the following classes:

- SYSTEM_METRICS
- IO_METRICS
- PROCESS_METRICS
- ALL_METRICS

The following table shows the classes of data you can select and lists the types of data collected for each class:

Option Choice	Type of Data Recorded
SYSTEM_METRICS	Parameter records (PAR_REC)
	Metrics records (MET_REC)
	Communication records (COM_REC)
	CPU records (CPU_REC
IO_METRICS	Disk records (DSK_REC)
	Tape records (MAG_REC)
	Communication records (COM_REC)
	Configuration records (CFG_REC)
	Hot File records (FIL_REC)
PROCESS_METRICS	Process records (PRO_REC)
	Metrics records (MET_REC)
	CPU records (CPU_REC)
ALL_METRICS	Process records (PRO_REC)
	Disk records (DSK_REC)
	Tape records (MAG_REC)
	Parameter records (PAR_REC)
	Metrics records (MET_REC)
	Communication records (COM_REC)
	Configuration records (CFG_REC)
	CPU records (CPU_REC)
	Hot File records (FIL_REC)

Explain the Primary Performance Agent

The primary data collection process (which is controlled by the default collection definition, CPD) runs as a detached process. This process is created during system startup if you include the following command in either of the following files, depending on which version of OpenVMS that you are running:

```
SYS$MANAGER:SYSTARTUP_V5.COM
SYS$MANAGER:SYSTARTUP_VMS.COM
```

The Performance Agent startup file is created during installation and is placed in the SYS$STARTUP directory.

You cannot modify the name, interval, and data classes of the CPD collection definition. The primary Performance Agent is the only Performance Agent that collects and records hot file data. To view the CPD collection definition, enter the following command:

```
$ ADVISE COLLECT SHOW SCHEDULE CPD
```

If you are collecting data for a cluster environment, you must run the primary data collection process on all nodes in the cluster to be able to use the full analysis and modeling capabilities of the CA software. This is the only way to have a complete representation of the cluster performance characteristics.

The primary data collection process collects the following performance data for a node at 2-minute intervals and for durations specified by the CPD collection definition within the schedule file:

- A subset of current system SYSGEN parameter settings
- System-wide usage statistics
- Per-process workload statistics
- Hot file statistics
- Utilization statistics by device, which include:
  - Bus
  - Controller
  - Disk
  - Tape
  - Terminal

The Schedule File and Collection Definitions

During installation, the Performance Agent software creates a schedule file containing the default (primary) collection definition named CPD. To modify the CPD collection definition, use the appropriate ADVISE COLLECT commands.

You can enter commands to the Performance Agent software using the ADVISE COLLECT command or by invoking the DECwindows interface with the ADVISE/INTERFACE commands:

```
$ ADVISE/INTERFACE=MOTIF
OR
$ ADVISE/INTERFACE=DECWINDOWS
```

For more information about the DECwindows interface, see the chapter Use the DECwindows Motif Interface (see page 191).

The Performance Agent software collects data according to collection definitions in the PSDC$DATABASE:PSDC$SCHEDULE.DAT file. While the software provides a standard collection definition, the CPD, you can create your own collection definitions.

When Changes Take Affect

The Performance Agent software reads the schedule file once an hour. If you want to inform the Performance Agent immediately of your changes to the collection parameters or to the weekly schedule, you must stop and restart the Performance Agent.

Examine the Schedule File

Before you modify a collection definition, you might want to examine it. To display all information in the schedule file, enter the following command:

```
$ ADVISE COLLECT SHOW ALL
```

CPD, the Default Collection Definition

When you enter the previous command, the system displays information similar to the following example:

```
                = Performance Schedule File =

================================================================
Collection Name //  Attributes  // Weekly Schedule
================================================================
CPD                                     | Monday  0-24
        Delete After   10 Days   | Tuesday  0-24
        HotFile Que.   0.33      | Wednesday 0-24
        Interval/sec  120.0      | Thursday 0-24
        Start Date  24-AUG-1992          | Friday  0-24
        End Date    1-JAN-2999   | Saturday 0-24
        Classes: All                     | Sunday  0-24
        Def Path: PSDC$DATABASE:
        Def Minimum Free Space: 10000
        Def Working Set Size:  4000
----------------------------------------------------------------
-Node---DC Status-------DskFreeSpace--WSquota---Path----------
----------------------------------------------------------------
GALLO  Running      10000    4800   PSDC$DATABASE:
YQUEM  Running      10000    4800   PSDC$DATABASE:
LATOUR Running      10000    4800   PSDC$DATABASE:
```

User-Defined Collection Definitions

If you created any collection definitions, the system will also display them. The following example shows a user-defined collection definition named DB_DISKS:

```
                = Performance Schedule File =

================================================================
=Collection Name //   Attributes    //  Weekly Schedule
================================================================
=DB_DISKS                          | Monday    10-12
        Delete After    7 Days             | Tuesday   10-12
        HotFile Que.    0.33               | Wednesday 10-12
        Interval/sec   30.0                | Thursday  10-12
        Start Date   28-JAN-1993           | Friday    10-12
        End Date      1-JAN-2999           | Saturday
        Classes: I/O                       | Sunday
        Def Path: PSDC$DB_DATABASE:
        Def Minimum Free Space: 2500
        Def Working Set Size:   2048
Volumes to include: ACCTDSK, DOCDSK, ADMINDSK
```

To view an individual collection definition, enter a command similar to the following command:

```
$ ADVISE COLLECT SHOW SCHEDULE DB_DISKS
```

For an explanation of the status of collection definitions, see Check the Data Collection Status (see page 33) in this chapter.

How Primary Data Collection Is Synchronized in a Cluster

The primary data collection process, controlled by the CPD collection definition, gathers information from the system's internal data structures. Some data collection is event driven, some is sampled periodically, and some is sampled at the end of each interval. At the end of each interval, all the information is time-stamped and written to the daily data file with a file type of CPD. All intervals have a fixed duration of two minutes. The time-stamp identifies each interval's data records.

Because each node in a cluster has its own clock and is set independently, there tends to be some variance, great or small, between the clocks. To make an analysis of the data meaningful and effective, the intervals recorded on each node must represent the identical period of real time.

To accomplish this synchronization, the first primary data collection process to start in a cluster initializes a logical clock with the node's system time. Using the Lock Manager, that clock is distributed to all other data collection processes as they start in a cluster. This synchronization strategy guarantees that data sampled independently on separate nodes in a cluster correlate accurately.

Alternate Performance Agents

You can create alternate Performance Agents to collect specific sets of data for special purposes. For example, you can specify an interval of 10 to 15 seconds with only system metrics data to be collected. You could run this Performance Agent for short periods of time during unusually heavy system demands to analyze these demands.

You can also create a Performance Agent to record data based on a longer interval such as 1800 to 3600 seconds. This Performance Agent would produce smaller data files that would act as an alternative to archiving large data files. See the note following this section.

The primary Performance Agent now acts as a data server to all other data collection processes. You can define additional collection processes to collect subsets of data from classes at different intervals and record this data for performance analysis or for dynamic viewing:

- For data files—When you create an alternate collection definition, you may optionally then start a Performance Agent that runs as a detached process using this collection definition.
- For dynamic viewing—You activate Real-time Performance Agents with the Real-time display commands, currently available with the Performance Agent software. For more information on Real-time data displays, see the *Performance Manager Administrator Guide*.

Note: The primary data collection process must be running for any of the other Performance Agents to have access to current data. The primary Performance Agent doesn't have to be actually collecting data to act as a data server (the data collection process may be in a waiting state due to a scheduled wait).

If you start an alternate Performance Agent or Real-time Performance Agent while the primary Performance Agent is not running, each Performance Agent writes a message in the PSDC$DC.LOG file and exits. This will also happen if the primary Performance Agent is stopped intentionally while other Performance Agents are running.

Support Real-time Performance Agents

The Performance Agent software enables you to view dynamic displays of data through the Real-time Monitor. The Performance Agent software can provide data for these displays through DECnet or through TCP/IP. This section describes how you can configure the network to provide access to the data. You can provide access for users through either TCP/IP or DECnet. The following sections provide examples of the commands you can use to provide access.

For information on the Real-time Monitor, see the *Performance Manager Administrator Guide*.

For more information about networking and system security, see HP's OpenVMS documentation.

Set Up the TCP/IP Access and User Account

To provide access to Real-time reporting through a privileged account, complete the following procedures:

- Verify the RSH Service configuration on the monitored node.
- Set the user account up with the necessary privileges on the node to be monitored.
- Define proxy access from the monitor node to the node to be monitored.

The user account on the monitored node must have the necessary privileges and quotas in the user authorization file (UAF) to enable the data collector process to run successfully. These privileges and quotas are the same as those defined for the primary Performance Agent. Preferably, the user account name on the monitored node should be the same as the account name on the monitoring node.

Verify the configuration of RSH

Remote access to Real Time reporting relies upon the TCP/IP RSH service enabled on the monitored node. To verify this prerequisite, for example using HP TCP/IP Services, run the following command and check the output for the service state "enabled":

```
$ TCPIP SHOW SERVICE RSH

Service       Port  Proto    Process          Address      State

RSH            514  TCP      TCPIP$RSH        0.0.0.0       Enabled
```

Use the utility SYS$MANAGER:TCPIP$CONFIG.COM to enable the service as required.

Create or modify the privileged user account

This account must have the necessary privileges and quotas in the user authorization file (UAF) to enable the data collector process to be run successfully. These privileges and quotas are the same as those defined for the primary Performance Agent. The account must have the same name as the account where the Real-Time Monitor will run on the monitor node.

To create a privileged account

1. Set your default directory to SYS$SYSTEM and run the Authorize Utility:

   ```
   $ SET DEFAULT SYS$SYSTEM
   $ RUN AUTHORIZE
   ```

2. If the system UAF has no record for the requested account, add the record and corresponding identifiers to the rights database:

   ```
   UAF> ADD account_name /DIR=[account_dir] /DEV=device_name -
   _UAF> /UIC=[xx,yy] /FLAGS=RESTRICTED /LGICMD=_NL: -
   _UAF> /NOPWDEXPIRED /PASS= password -
   _UAF> /NETWORK/NOBATCH/NOINTERACTIVE
   ```

3. Set necessary quotas and privileges for the requested account:

   ```
   UAF> MODIFY account_name /DEFPRIV=(CMKRNL, SYSNAM, GRPNAM,-
   _UAF> DETACH, PSWAPM, ALTPRI, TMPMBX, WORLD, EXQUOTA, NETMBX,-
   _UAF> SYSLCK)/PRIV=(CMKRNL, SYSNAM, GRPNAM, DETACH, PSWAPM, -
   _UAF> ALTPRI, TMPMBX, WORLD, EXQUOTA, NETMBX, SYSLCK) -
   _UAF> /WSQUO=2048 /WSEXTENT=16000 /PGFLQUO=50000
   UAF> EXIT
   $ CREATE/DIR device_name:[account_dir] /OWNER_UIC=[xx,yy]
   ```

   In the preceding example, you must substitute real values in place of account_name, account_dir, device_name and xx,yy, the UIC.

Set Up TCP/IP Proxy Access

With TCP/IP proxy access you allow only specified remote users to start a Real-time agent on the local node where the Agent is installed.

To create a proxy account in the service (TCPIP Services, MultiNet and TCPware)

1. For example, execute the following command for TCPIP Services:

   ```
   $ TCPIP ADD PROXY account_name /HOST=remote_node -
   /REMOTE_USER=account_name
   ```

   Substitute the real account name and remote name in place of account_name and remote_name where the Real-Time Monitor will run to monitor the local node.

2. To check that the proxy account works, execute the following command from the account account_name of the remote_node.

   ```
   $ RSHELL monitored_node DIR
   ```

3. For MultiNet, be sure all login command procedures and the system login command procedure contain these lines at the top of the file:

   ```
   $ verify = 'f$verify(0)'
   $ if f$mode() .eqs. "OTHER" then exit
   ```

Set Up the DECnet Phase IV Object and a User Account

To provide access to Real-time reporting through a privileged account:

- Create an account with the necessary privileges on the node to be monitored.
- Define the DECnet object.

This account must have the necessary privileges and quotas in the user authorization file (UAF) to enable the Performance Agent process to be run successfully. These are the same privileges and quotas as defined for the main Performance Agent.

The following commands are an example of how you can create a privileged account, perform the following procedures:

1. Set your default directory to SYS$SYSTEM and run the Authorize Utility (AUTHORIZE).

   ```
   $ SET DEF SYS$SYSTEM
   $ RUN AUTHORIZE
   ```

2. Add a user record to the system UAF and corresponding identifiers to the rights database with the following command:

```
UAF> ADD PSDC$SERVER /DIR=[PSDC$SERVER] /DEV=device_name -
_UAF> /DEFPRIV=(CMKRNL, SYSNAM, GRPNAM, IMPERSONATE, PSWAPM, -
_UAF> ALTPRI, TMPMBX, WORLD, EXQUOTA, NETMBX, SYSLCK) -
_UAF> /PRIV=(CMKRNL, SYSNAM, GRPNAM, IMPERSONATE, PSWAPM, -
_UAF> ALTPRI, TMPMBX, WORLD, EXQUOTA, NETMBX, SYSLCK) -
_UAF> /UIC=[xx,yy] /FLAGS=(RESTRICTED,NODISUSER)/LGICMD=_NL: -
_UAF> /NOPWDEXPIRED /PASS= password -
_UAF> /WSQUO=2048 /WSEXTENT=16000 /PGFLQUO=50000 -
_UAF> /NETWORK/NOBATCH/NOINTERACTIVE
UAF> EXIT
$ CREATE/DIR device_name:[PSDC$SERVER] /OWNER_UIC=[xx,yy]
```

   In the preceding example, you must substitute real values in place of *device name* and *xx,yy*, the UIC.

3. You must add a DECnet task, PSDC$SERVER, associated with object 0. Enter the command as follows:

```
$ RUN SYS$SYSTEM:NCP
NCP> DEFINE OBJECT PSDC$SERVER NUMBER 0 FILE SYS$SYSTEM:-
PSDC$SERVER.COM USER PSDC$SERVER PASSWORD <password>
```

   The password must match the password in the UAF entry for PSDC$SERVER.

```
NCP> SET OBJECT PSDC$SERVER ALL
```

   To display information about the DECNET object, enter the following command:

```
NCP> SHOW OBJECT PSDC$SERVER CHARACTERISTICS
NCP> EXIT
```

Create a DECnet Proxy Account

Proxy access is an alternative to the unlimited access provided by the DECnet object with a privileged account. You can use the proxy account to provide limited access. You can allow only specified remote users to start a Real-time Performance Agent on the local node where the Performance Agent is installed.

To create a proxy account

1. Set your default directory to SYS$SYSTEM and run the Authorize Utility (AUTHORIZE).

```
$ SET DEF SYS$SYSTEM
$ RUN AUTHORIZE
```

2. Add a user record to the system user authorization file (UAF) and corresponding identifiers to the rights database with the following command:

```
UAF> ADD PSDC$SERVER /DIR=[PSDC$SERVER] /DEV=device_name-
_UAF> /DEFPRIV=(CMKRNL, SYSNAM, GRPNAM, IMPERSONATE, PSWAPM,-
_UAF> ALTPRI, TMPMBX, WORLD, EXQUOTA, NETMBX,-
_UAF> SYSLCK)/PRIVILEGE=(CMKRNL, SYSNAM, GRPNAM,IMPERSONATE,-
_UAF> PSWAPM, ALTPRI, TMPMBX, WORLD, EXQUOTA,-
_UAF> NETMBX, SYSLCK)/UIC=[xx,yy]-
_UAF> /FLAGS=(RESTRICTED,NODISUSER)/LGICMD=NL:-
_UAF> /WSQUO=2048 /WSEXTENt=20000 /PGFLQUO=50000-
_UAF> /NETWORK/NOBATCH/NOINTERACTIVE
_UAF> EXIT
$ CREATE/DIR device_name:[PSDC$SERVER] /OWNER_UIC=[xx,yy]
```

   In the preceding example, you must substitute real entries in place of *device name* and *xx,yy*, the UIC.

3. Create the proxy database if one does not already exist. Enter the following command to create and initialize the network proxy authorization file, NETPROXY.DAT:

```
UAF> CREATE/PROXY
```

4. Add the users to the network proxy authorization file:

```
UAF> ADD/PROXY MYNODE::JONES PSDC$SERVER /DEFAULT
UAF> EXIT
```

   The preceding command provides JONES, the user on the remote node MYNODE, access through the proxy account PSDC$SERVER.

5. You must define a DECnet task, PSDC$SERVER, associated with object 0. Enter the command as follows:

```
$ RUN SYS$SYSTEM:NCP
NCP> DEFINE OBJECT PSDC$SERVER NUMBER 0 FILE SYS$SYSTEM: -
PSDC$SERVER.COM
```

   Note: Do not enter any USER or PASSWORD for this object.

```
NCP> SET OBJECT PSDC$SERVER ALL
```

   For more information on the Authorize Utility, see HP's OpenVMS documentation.

Set Up the DECnet Phase V Session Control

You need to create a privileged account and define the session control application in order to provide access to real-time.

To provide access to Real-time using DECnet Phase V session control

1. Set your default directory to SYS$SYSTEM and run the Authorize Utility (AUTHORIZE).

   ```
   $ SET DEF SYS$SYSTEM
   $ RUN AUTHORIZE
   ```

2. Add a user record to the system UAF and corresponding identifiers to the rights database with the following command:

   ```
   UAF> ADD PSDC$SERVER /DIR=[PSDC$SERVER] /DEV=device_name-
   _UAF> /DEFPRIV=(CMKRNL, SYSNAM, GRPNAM, IMPERSONATE, PSWAPM,-
   _UAF> ALTPRI, TMPMBX, WORLD, EXQUOTA, NETMBX, SYSLCK)-
   _UAF> /PRIV=(CMKRNL, SYSNAM, GRPNAM, IMPERSONATE, PSWAPM,-
   _UAF> ALTPRI, TMPMBX, WORLD, EXQUOTA, NETMBX, SYSLCK)-
   _UAF> /UIC=[xx,yy] /FLAGS=(RESTRICTED,NODISUSER)/LGICMD=_NL:-
   _UAF> /NOPWDEXPIRED /PASS= password-
   _UAF> /WSQUO=2048 /WSEXTENT=16000 /PGFLQUO=50000-
   _UAF> /NETWORK/NOBATCH/NOINTERACTIVE
   _UAF> EXIT
   $ CREATE/DIR device_name:[PSDC$SERVER] /OWNER_UIC=[xx,yy]
   ```

   In the preceding example, you must substitute real values in place of *device name* and *xx,yy*, the UIC.

3. Login to the SYSTEM account and run the network configuration procedure.

   ```
   $ @SYS$MANAGER:NET$CONFIGURE application_delete "PSDC$SERVER" 1
   $ @SYS$MANAGER:NET$CONFIGURE APPLICATION_ADD "PSDC$SERVER"
   "PSDC$SERVER|[xx,yy]" "{NAME=PSDC$SERVER,NUMBER=0}
   |SYS$SYSTEM:PSDC$RT_V5.EXE |FALSE|FALSE|FALSE|FALSE|TRUE||" 1
   ```

   Note: In the code above, the second and third lines need to be entered as one continuous line. They appear as two lines here due to the page margin constraints.

Customize Collection Definitions

Changing or creating collection definitions allows you to tailor data collection to suit your needs. This section describes the DCL commands you use. If you have the DECwindows Motif interface, you can also create or modify collection definitions. For information, see the chapter Use the DECwindows Motif Interface (see page 191).

Change Weekly Collection Schedule

With the ADVISE COLLECT ADD or MODIFY command, use the /SCHEDULE qualifier to specify by the time of day when the data collection process is to collect data. This schedule applies to all nodes in a collection definition. At the end of the week, the schedule repeats.

You can change the schedule using the ADVISE COLLECT MODIFY command and /SCHEDULE qualifier with the new day and range of hours. See the /SCHEDULE qualifier of the ADVISE COLLECT MODIFY command in the chapter "Performance Agent Commands."

For example, you might enter the following command:

```
$ ADVISE COLLECT MODIFY DB_DISKS -
_$ /SCHEDULE=(WEEKDAYS=(10-11,14-15),NOMONDAY,NOFRIDAY,NOWEEKENDS)
```

For the collection definition specified, this command would have the following effect:

- No data collection on Monday.
- Midweek data collection occurs from 10:00 to 11:00 a.m. and from 2:00 to 3:00 p.m.
- No data collection on Friday.
- The data collection process does not collect data on Saturday or Sunday.

Keep in mind that this schedule will be used by all the nodes in this collection definition.

Change the Database Directory

Use the /PATH qualifier to specify a new database directory for a given node in a collection definition. This affects not only the recording of data but also the archiving of data.

To change the database perform the following procedures

1. Stop the data collection process associated with the collection definition. If you are changing the CPD collection definition, also stop the archiving process (if executing).
2. Copy or move (with the RENAME command) the daily data and history files, if applicable, to the new directory to keep all the files together.
3. After changing the database directory, start the data collection process.

Add, Change, and Reorder Node Names in a Collection Definition

Performance Agent software writes the current cluster node names into the CPD collection definition in the schedule file. If a node is not a member of the cluster during installation, it is not included in this definition. If nodes are added to the cluster, you must add them to the collection definition. The Performance Agent uses the local node name, or list of node names in the case of a cluster, to schedule data collection. To add or remove a node name or to reorder the existing list, use the ADVISE COLLECT MODIFY command.

For example, the following command adds the node NEWSYS to the CPD collection definition:

```
$ ADVISE COLLECT MODIFY CPD NEWSYS
```

The following command adds all current nodes to the CPD collection definition:

```
$ ADVISE COLLECT MODIFY CPD *
```

The following command removes node SYSNOT from the CPD collection definition:

```
$ ADVISE COLLECT MODIFY/REMOVE_NODES=SYSNOT CPD
```

The following command reorders the nodes in all the collection definitions:

```
$ ADVISE COLLECT MODIFY/REORDER_NODES=(CNODE,BNODE,ANODE)
```

By default, the Performance Agent prepares reports according to the order of node names specified in the schedule file. You can reorder the list of names to cause a specific reporting order in your reports.

Change the Beginning and Ending Dates for Data Collection

The beginning and ending dates stored in the collection definition define a window of time during which the data collection process collects data according to the weekly schedule. When you installed the Performance Agent software, it specified beginning and ending dates for the CPD collection definition. The beginning date is the date and time you installed the Performance Agent software; the ending date is set to January 1, 2999. To modify this default collection window, use the ADVISE COLLECT MODIFY command. See the /BEGINNING and /END qualifiers of the ADVISE COLLECT MODIFY in the chapter "Performance Agent Commands."

For example, if you did not want to collect data under the collection definition DB_DISKS over the year-end holidays, assuming the current date is December 18, 2008, you would enter the following command:

```
$ ADVISE COLLECT MODIFY DB_DISKS/BEGINNING=04-JAN-2008:00:00
```

This command defines a window from December 18, 2008 to January 4, 2008 (assuming the ending time had not been changed) and stops the data collection process for DB_DISKS from collecting data until January 4, 2008. As long as the ending date is in the future, the data collection process starts collecting data on January 4, 2008. If the ending date has passed, the data collection process terminates.

Create Your Own Collection Definitions

To collect a subset of data tailored to your own requirements, create a collection definition with the ADVISE COLLECT ADD command. The following command creates a collection definition called SYS_SUMMARY:

```
$ ADVISE COLLECT ADD SYS_SUMMARY-
_$ /INTERVAL = 3600-
_$ /CLASS = SYSTEM_METRICS-
_$ /PATH = PSDC$SUMMARY_DB
```

The preceding command defines a collection definition for recording system metrics data only on an hourly basis. The data files will be placed in the directory specified by PSDC$SUMMARY_DB. These data files may be an alternative to archiving data if a subset of system metric data is sufficient.

Notice that in the preceding example command, nodes are not specified. To collect data you must specify nodes when creating a collection definition.

```
= Performance Schedule File
Collection Name //   Attributes    //  Weekly Schedule
SYS_SUMMARY                           | Monday    0-24
            Delete After      7 Days  | Tuesday   0-24
            HotFile Que.     0.33     | Wednesday 0-24
            Interval/sec  3600.0      | Thursday  0-24
            Start Date    11-APR-1997 | Friday    0-24
            End Date       1-JAN-2999 | Saturday  0-24
            Classes: Metrs            | Sunday    0-24
            Def Path: PSDC$SUMMARY_DB:
            Def Minimum Free Space: 2500
            Def Working Set Size:   2048
CA Performance Agent Module version is: Vx.x
```

Create a New Schedule File

If for any reason you need to create a new schedule file, follow these steps:

To create a new schedule file

1. Delete the old schedule file, if it exists.
2. Enter the command ADVISE COLLECT MODIFY CPD. This command will create and display a new schedule file. This file will contain the CPD collection definition only.

Start the Data Collection Processes

To start a data collection process, enter the following command:

```
$ ADVISE COLLECT START collection-definition-name
```

If you do not specify a collection definition, the CPD collection definition is started.

To assure proper recognition of the system configuration by the Performance Agent software, you must manually stop the data collection processes prior to binding or dissolving volume, stripe, and shadow sets, or reconfiguring hardware. For more information, see Stopping the Data Collection Processes in this chapter. Start the data collection process afterwards.

Only one primary, or CPD, data collection process per node can be active at a given time. The process runs under the user identification code (UIC) of [1,4], has the process name of PSDC$DC_SERVER, and a process base priority of 15.

Monitor Data Collection

You can monitor data collection by displaying Performance Agent status and error log information.

Check the Data Collection Status

To view the status of the data collection process on the nodes in the cluster, enter the following command:

```
$ ADVISE COLLECT SHOW STATUS
```

When you use this command, the system displays information similar to the following example:

```
Collection Name CPD
-------------------------------------------------------------
-Node---DC Status-------DskFreeSpace--WSquota---Path----------
-------------------------------------------------------------
GALLO   Running            10000       2048     PSDC$DATABASE:
YQUEM   Running            10000       2048     PSDC$DATABASE:
LATOUR  Running             9999       2048     PSDC$DATABASE:
XUKV12  Down                2500       2048     PSDC$DATABASE:
```

If you do not have the SYSLCK privileges, the software will not display the status but displays NoPrivs instead.

The following table lists the Performance Agent status messages and their meanings:

DC Status	Meaning
Running	Recording data
SchedWait	Waiting due to schedule
FreeSpaceWait	Waiting for disk space
Stopped	Status of node unknown
NoPrivs	No privilege to view status

Check the Data Collection Error Log

If an error occurs during the data collection process, the Performance Agent records it in a file called PSDC$DATABASE:PSDC$DC.LOG. This log file is common to all data collection processes in the cluster. You can print or display this ASCII file from a terminal. Each entry in the file contains the following data:

- Node name
- Date

- Time stamp
- Error message

The Performance Agent creates the PSDC$DC.LOG file the first time it encounters an error. All subsequent errors will be appended to it, unless you delete the file. Deleting the PSDC$DC.LOG file does not interfere with the normal execution of data collection.

Stop the Data Collection Processes

At times, you will need to stop and restart the data collection processes. The data collection processes must be shut down prior to a scheduled system shutdown using the PSDC$SHUTDWN.COM file. The PSDC$SHUTDWN.COM file is created by the installation of the Performance Agent and placed in the SYS$MANAGER directory. Be sure that this file is invoked from the system shutdown command file.

To stop a specific data collection process, you must first be logged on to the node on which that process is running or use the SYSMAN Utility. To stop the primary CPD data collection process, which has the process name PSDC$DC_SERVER, enter the following command:

```
$ ADVISE COLLECT STOP/WAIT CPD
```

Stopping the primary data collection process causes all the alternate data collection processes to stop.

To stop any alternate data collection process, use a command similar to the following command:

```
$ ADVISE COLLECT STOP DB_DISKS
```

The CPD Performance Agent must execute a termination handler procedure to return some of its non-paged pool. To terminate, it must perform the following actions:

- Dissolve all pending links with other processes running on the system.
- Relinquish several vectors that it uses to measure performance.

There is a normal delay for the termination handler to complete. The result is a delay between the time you issued the command to stop the CPD Performance Agent and the time its process disappears. You can synchronize the command completion with the actual termination by using the /WAIT qualifier.

Use the Database Directory

A Performance Agent stores a node's collected data in a central database directory specified by the collection definition. Created during installation, the default database directory is assigned the logical name PSDC$DATABASE. You can specify an alternative location for a node's data file by modifying the collection definition.

If you install the Performance Agent software in a cluster environment, the PSDC$DATABASE directory should be on a permanently mounted disk, accessible to every node in the cluster because the data collection schedule file, which is common to all nodes, is in this directory. The PSDC$DATABASE directory contains the following files:

- Schedule file
- Parameters file
- Holiday file (optional)
- History files (optional)
- Performance Agent error log file (created when an error occurs)

The PSDC$DATABASE directory also contains the daily data files unless you specify an alternate directory in the collection definition.

Note: If the path is changed, the data files and history files for that node must be moved from the original location to this new location.

If a database disk has fewer than the required number of free blocks, any data collection processes stop recording data in the database. The primary Performance Agent continues to collect data but stops recording the data until there are once again enough free blocks.

If you enter the ADVISE COLLECT SHOW ALL command, the required space is displayed under "DskFreeSpace" and the Performance Agent status shows as "FreeSpaceWait." The Performance Agent then hibernates until sufficient free space becomes available.

To adjust the free-space threshold, use the following command format:

```
ADVISE COLLECT MODIFY collection-name * /FREE_SPACE = number
```

If this threshold is already low, it is better that you free up additional space on the disk.

Normally you never want to run out of database disk space because this will cause gaps in your data and limit the use of this data in reports, graphs, or modeling. Use the collection parameters and collection schedule to control the amount of data written to the database directory. Archiving also contributes to the use of disk space; see Chapter 4 for more information.

Optionally, you can set up a secondary database on another device and modify some of the nodes' database paths to point to these directories.

Use the Daily Data Files (CPD)

The daily data files from the primary Performance Agent are used by other modules of the CA product set. The Performance Agent analyzes the data and creates report and graphs.

The data collection processes create a data file each day for each node in the cluster. The file name contains the origin and date of the data. It has the following format:

```
path:PSDC$nodename_yyyymmmdd.c-d-n
```

path

This is the directory specified in the collection definition.

nodename

This is the node name.

yyyymmmdd

This is the date in 4-digit year 3-letter month 2-digit day format.

c-d-n

This is the name of the collection definition.

Example file name:

```
PSDC$DATABASE:PSDC$GALLO_20080CT26.CPD;1
```

The primary (CPD) data file contains one data record for each interval. The interval is two minutes. Each interval record contains subrecords for each major area of performance data. The following table shows the subrecords in the file:

Subrecord	Function	Record Count
Time Stamp	Contains a time stamp for the interval.	1

Subrecord	Function	Record Count
Metrics	Contains system-wide performance metrics.	1
Parameters	Contains values of some SYSGEN parameters.	1
Process	Contains data for processes and the images that were activated by the processes. Information pertaining to working set sizes, fault rates, UAF parameters, and so forth are contained in this subrecord.	1 record per image executed and 1 record per process terminated
Disk	Contains information about disk activities.	1 per disk device
Tape	Contains information about tape activities.	1 per tape device
Configuration	Contains the node names that are currently part of the cluster membership and/or local node and the SCS statistics for each member.	1 per SCS virtual circuit
CPU	Contains information about the percent of time spent in processor modes.	1 per CPU
Communication	Contains the operation counts for all terminal controllers.	1 per terminal controller
Hot files	Contains hot file activity.	0 to 40 per disk

For more information on Interval Sub record contents, see the chapter Dump Reports (see page 53).

Store and Delete Daily Data Files

Each night at midnight (00:00:00), each data collection process closes the current file, deletes any expired files, and creates the next day's file. Use the ADVISE COLLECT SHOW SCHEDULE command to display the number of days occurring between file creation and deletion. The default expiration time is seven days. Use the ADVISE COLLECT MODIFY/DELETE_AFTER command to modify this time. For more information about the ADVISE COLLECT command, see the chapter "Performance Agent Commands."

Chapter 3: Characterize Workloads

This chapter describes the following attributes:

- Defining Workloads
- Creating User-Group Workload Definitions
- Defining User-Group Families
- Creating Transaction Workload Definitions
- Defining Transaction Families
- Using Workload Families with CA Data

The Performance Agent allows you to group system resource use into manageable units called workloads. Workloads can represent the resource use of users who are part of an organization, such as accounting or finance. Categorizing your system activity in this manner creates workloads focused on the *user* aspect of system resource use. Alternatively, workloads can be groups of similar images that represent system resources used by particular applications, such as order entry or inventory control. Categorizing your system activity in this manner creates workloads focused on the *transaction* consuming system resources.

Workloads can thus be used to view your computer resource use from two different perspectives, by users or by tasks being performed. This double mapping of computer resources allows great flexibility in reporting, graphing, and archiving.

This section contains the following topics:

Define Workloads

The following rules and conditions apply to all workloads:

- Workload definitions are created, modified, and deleted using the Parameter Editor or DECwindows interface. For more information on defining and editing workloads, see the chapters "Parameter Editor Commands" and "Using the DECwindows Motif Interface."
- Workload definitions are stored in the parameters file:

  ```
  PSDC$DATABASE:PSDC$PARAMS.DAT
  ```

- Workload definitions can contain an unlimited number of image names or combinations of image names with a list of user criteria.
- An unlimited number of workload definitions can be defined in the parameters file.
- Wildcards can be used when specifying image names, user names, process names, account names, and UICs.
- Character limits are as follows: process names (15 characters), user names (12 characters), image names (39 characters), account names (8 characters.)
- The device, directory, and file type cannot be included in an image name.

Create User-Group Workload Definitions

User-group workload definitions contain groups of users who have something in common. For example, users might belong to the same department, be on the same project, or have similar jobs. User-group workload definitions facilitate the reporting and graphing of computer resources used by the user group.

The following are examples of Parameter Editor commands to define new user-group workload definitions:

1. To start the Parameter Editor, enter:

   ```
   $ ADVISE EDIT
   ```

2. To add the SECRETARIES workload definition with users Johnson, Whelms, and Hears, enter:

   ```
   PSDC-EDIT> ADD/WORKLOAD SECRETARIES-
   _PSDC-EDIT> /USERNAMES=(JOHNSON,WHELMS,HEARS)
   ```

3. To add the WRITERS workload definition with all users in UIC group 100, enter:

   ```
   PSDC-EDIT> ADD/WORKLOAD WRITERS /UIC=[100,*]
   ```

Define User-Group Families

After you created the user-group workload definitions in the parameters file, you can collect them together into a user-group family. Grouping the workload definitions into a family is essential because it is a family name that you specify when the workload definitions are used in reporting, graphing, and archiving. The family name is a convenient and concise way of specifying a collection of user-group workload definitions.

At installation, Performance Agent provides a user-group family, MODEL_USERGROUPS, which cannot be deleted from the parameter file. The workload definition names of this family are listed in the following bullets:

- OPERATOR-Users in the OPERATOR workload
- SYSTEM_USER-Users in the SYSTEM workload
- DECNET-Users in the DECNET workload
- Other-User names that are not defined in one of the other three workloads of this family

The following Parameter Editor examples define new or existing user-group families:

1. To define the MY_USERS family, which includes the WRITERS and SECRETARIES user-group workload definitions and the preceding list of default Performance Agent user-group workload definitions, enter:

```
PSDC-EDIT> ADD/FAMILY MY_USERS/WORKLOAD=(WRITERS,-
_PSDC-EDIT> SECRETARIES,SYSTEM_USER,OPERATOR,DECNET)
```

2. To add new workload definitions to the user-group family, MODEL_USERGROUPS, enter:

```
PSDC-EDIT> MODIFY/FAMILY MODEL_USERGROUPS/ADDWORKLOAD=-
_PSDC-EDIT> (WRITERS,SECRETARIES)
```

The order in which you specify the workload definitions is important when defining a user-group family. Performance Agent allocates process resource use for the user groups specified on a first-match basis. For example, if you specify user TONES in both user groups WRITERS and SECRETARIES (based on example 2), all resource use by TONES falls into the WRITERS user-group workload definition because it is listed first. The resource use is not split between the two workloads. If you create workload definitions that overlap, specify the workload definition with the more specific matching criteria before the workload definition with the more general matching criteria in the user-group family.

Besides the workload definitions specified, all families are treated as if they contain a workload definition of Other. All resource use not matching specified workload definitions is placed in the Other workload.

Create Transaction Workload Definitions

Transaction workload definitions contain groups of images with similar resource use. When creating a transaction workload definition, verify that the images included in the workload definition have similar service demands. Specifying images with similar service demands ensures that the reported system-wide image response times for transaction workload definitions are meaningful. For example, you might group all images that perform editing into a transaction workload definition. Similarly, all images that perform inquiries or database updates might be grouped into their respective transaction workload definition.

At installation, Performance Agent provides the transaction workload definitions SYSMAN, COMPILES, UTILITIES, EDITORS, and NETWORK. These workload definitions are used as the default if you do not define and use your own. Performance Agent requires that every transaction workload definition have at least one image name.

The following Parameter Editor examples create new transaction workload definitions:

1. To add the database update workload definition named DB_UPDATES with images CST_UPD, INV_UPD, and ORD_UPD, enter:

   ```
   PSDC-EDIT> ADD/WORKLOAD DB_UPDATES/IMAGENAMES=-
   _PSDC-EDIT> (CST_UPD,INV_UPD,ORD_UPD)
   ```

   To narrow the focus of a workload definition to specific images used by certain users, specify combinations of user criteria and image names. Specify that both must match before the resource use is allocated for the workload.

2. To add the writers editing workload definition (assuming all writers are in UIC group 200), enter

   ```
   PSDC-EDIT> ADD/WORKLOAD WRITER_EDITS/UIC=[200,*]-
   _PSDC-EDIT> /IMAGENAMES=(TPU,EDT,LSEDIT) /MATCH=BOTH
   ```

When you define a transaction workload, you have a choice of two types of transactions. You can specify image terminations or terminal responses. This choice determines how the software counts the number of transactions for a particular workload.

These options are as follows:

- Image terminations—The software counts the number of image terminations to measure the transactions.
- Terminal responses—The software counts the number of terminal QIO read completions for terminal devices to measure the transactions.

When you choose the terminal response option, the Performance Manager component provides information on terminal response time. This CA component breaks down workloads into smaller workloads based on differing levels of resource demands. This feature is called *clustering*.

To select transactions based on terminal response, use the /TRANSACTION_UNITS qualifier in the Parameter Editor. The format for this is

```
/TRANSACTION_UNITS=TERMINAL_RESPONSES
```

To select transactions based on image terminations, use the /TRANSACTION_UNITS qualifier in the Parameter Editor. The format for this is

```
/TRANSACTION_UNITS=IMAGE_TERMINATIONS
```

Define Transaction Families

After you define the transaction workload definitions in the parameters file, you can group transaction workload definitions together into a transaction family. The family name is a convenient and concise way of specifying a collection of transaction workload definitions. Grouping the workload definitions into a family is essential because it is a family name that you specify when the workload definitions are used in reporting, graphing, and archiving.

During installation, Performance Agent provides a default transaction family, MODEL_TRANSACTIONS, which you cannot delete from the parameter file. Following is a list of the workload definitions of this family:

- SYSMAN
- COMPILES
- UTILITIES
- EDITORS
- NETWORK
- OTHER (generated during use of a family)

The appendix The Default Workload Families (see page 235) contains the lists of the images included within each of these workload definitions.

The following Parameter Editor examples define new transaction families:

1. To define a new transaction family, MY_TRANSACTIONS, which includes the WRITER_EDITS and DB_UPDATES transaction workload definitions, enter:

   ```
   PSDC-EDIT> ADD/FAMILY MY_TRANSACTIONS/WORKLOAD=-
   _PSDC-EDIT> (WRITER_EDITS,DB_UPDATES)
   ```

2. To add your new workload definitions to the default transaction family, MODEL_TRANSACTIONS, enter:

```
PSDC-EDIT> MODIFY/FAMILY MODEL_TRANSACTIONS/ADDWORKLOAD=-
_PSDC-EDIT> (WRITER_EDITS,DB_UPDATES)
```

The order in which you specify the workload definitions is important when defining a transaction family. Performance Agent allocates process resource use to the various transaction workloads on a first-match basis. For example, if you specify image EDT in both workload definitions UTILITIES and WRITER_EDITS, all resource use of the image EDT falls into the WRITER_EDITS transaction workload definition because it is listed first in the MY_TRANSACTIONS transaction family.

However, if the MODEL_TRANSACTIONS transaction family is used, all resource use of the image EDT falls into the UTILITIES transaction workload definition because workload definitions added to families are placed at the end of the current list of workload definitions in the family. The resource use is not split between the two workloads. If you create workload definitions that overlap, specify the workload definition with the more specific matching criteria before the workload definition with the more general matching criteria.

Besides the workload definitions specified, all transaction families contain another workload definition: OTHER. This definition is generated during the processing of any family. Any transaction activity not matching specified workload definitions is placed in OTHER.

Use Workload Families with CA Data

Generally, when workload families are relevant to a CA command, the qualifier used is /CLASSIFY_BY. Because there are two major types of workload families, the /CLASSIFY_BY qualifier is followed by a keyword of USERGROUP, TRANSACTION, or both and then any workload family name. For example, /CLASSIFY_BY=USERGROUP=MY_USERS.

Archiving

In describing how daily data is to be summarized into archived files, a list of workload families can be supplied in the definition of a history file descriptor. The archiving process summarizes process data into each of the workload families specified by a history file descriptor. The specified families are the only ones that can be used when the archived data is being used as the source for reports and graphs. For more information about archiving, see the chapter History Files and Archiving (see page 47).

Graphing

The Performance Manager component includes a graphing module for displaying data. Predefined sets of graphs are available in categories of Top Users, Top Images, and so forth. One category is Top Workloads. To activate the set of Top Workload graphs, a workload family must be specified with the graph command. This also provides access to the custom workload graphs. For more information about graphing, see the *Performance Manager Administrator Guide*.

Performance Evaluation - Process Statistics

The Performance Manager component includes a report of all process activity contained in the daily collection files. By default, it shows all image activity by processing mode - interactive, batch, or network. To reformat this report in terms of your workload user groups, a workload family name must be provided in the report request.

For more information about analysis reporting, see the *Performance Manager Administrator Guide*.

Dump Archived Data

If your archived data contains families, you may select one of those families with the /CLASSIFY qualifier when dumping archived data. For more information about dump reporting of archived data, see the chapters Dump Reports (see page 53) and Performance Agent Commands (see page 119).

DECwindows Interface

Both the Performance Agent and Performance Manager components include a DECwindows interface. When selecting the desired analysis period and processing options, a menu of the currently available workload families is provided. Selecting an entry in the menu affects the presentation of the data for any graphs, dumps, or performance evaluations. For more information about DECwindows, see the chapter Use the DECwindows Motif Interface (see page 191), or the *Performance Manager Administrator Guide*.

Chapter 4: History Files and Archiving

This chapter describes the following procedures:

- Setting Up a History File Reduction Scheme
- Archiving CA Data
- Using History Files

The Performance Agent software provides data archiving options to archive and reduce the performance data that the primary Performance Agent collects daily, for use in long-term performance studies. When Performance Agent archives the daily data files, the resultant files are called *history files*. You can use history files as the data source for generating performance evaluation reports, dump reports, models, auxiliary analysis, and graphs. Hot file data is not archived.

Performance Agent creates the history files according to a user-specified archiving scheme, as defined in the parameters file. Multiple history files and data reduction schemes are allowed, and each is identified by a unique name. Performance Agent provides two default history file archiving schemes: MONTHLY_MODEL and MONTHLY_USER. You can use, modify, or delete these, or create your own to suit your needs.

To view the elements of the history file definitions, issue the SHOW/HISTORY * command from the Parameter Editor. A file similar to the following example appears:

```
History Descriptor         History Descriptor Parameters

 -------------------  --------------------------------------

MONTHLY_USER        Granularity = Monthly;          (No Periodicity)
                    Interval    = 60 minutes;       Delete_after = 180 days;
                     No modelling data is saved.
                    Usergroup classification families:
                    EACH_USER
                    Normal Schedule
                              Monday              0-24
                              Tuesday             0-24
                              Wednesday           0-24
                              Thursday            0-24
                              Friday              0-24
                              Saturday            0-24
                              Sunday              0-24
                    Holiday Schedule
                              Monday              0-24
                              Tuesday             0-24
                              Wednesday           0-24
                              Thursday            0-24
                              Friday              0-24
                              Saturday            0-24
                              Sunday              0-24
```

```
MONTHLY_MODEL       Granularity = Monthly;          (No Periodicity)
                    Interval    = 60 minutes;       Delete_after = 180 days;
                    Modelling data is saved.
                    Normal Schedule
                              Monday                10-12, 14-16
                              Tuesday               10-12, 14-16
                              Wednesday             10-12, 14-16
                              Thursday              10-12, 14-16
                              Friday                10-12, 14-16
                              Saturday
                              Sunday
                    Holiday Schedule
                              Monday
                              Tuesday
                              Wednesday
                              Thursday
                              Friday
                              Saturday
                              Sunday
```

This section contains the following topics:

Set Up a History File Reduction Scheme

You create a scheme for reducing data in the parameters file. Add or modify history file reduction schemes with the Parameter Editor as described in the chapter "Parameter Editor Commands."

The reduction scheme must specify the following information:

- History description name—A history description 1 to 20 character name.
- Schedule—Specifies the time periods of the daily data for which data is archived, thereby allowing a subset of each day's data to be archived.
- Holiday schedule—Specifies the time periods for which data is archived on holidays. Performance Agent uses this schedule when the archived data falls on a day specified in PSDC$DATABASE:PSDC$HOLIDAYS.TXT. You can edit this file using any text editor.
- Granularity—The time span an individual history file represents. Granularity must be daily, weekly, monthly, quarterly, or yearly. Smaller portions of the historical data can be deleted, moved, or manipulated with a smaller granularity.

- Periodicity—Specifies the calendar period (typical day, week, month, or quarter) into which the daily data for a longer span of time is summarized. For instance, archived data with monthly granularity and weekly periodicity produces typical week data for the month; a daily periodicity would produce a typical day for the month. The value of periodicity must be less than the value of granularity. Periodicity must be daily, weekly, monthly, or quarterly.
- No Periodicity—Causes each day to be maintained individually. In effect, the value of periodicity is equal to the value of granularity.
- Interval of data resolution—The size, in minutes, of each interval record. Any even division of two minutes into a day is valid.
- Retention—The number of days Performance Agent keeps the history file beyond the time span it represents. The archive command first deletes history files that reach retention threshold and then archives the data. A history file representing a period of time before the retention period is subsequently deleted by the next archiving run. Therefore, to archive old daily files, extend the retention period to encompass their time frames.
- Model data switch—When enabled, causes Performance Agent to archive additional data, permitting subsequent modeling from the history file, and reclassification of the workload. All unique PIDs, user names, image names, account names, process names, UICs, and any combination of these names are kept separate.
- List of families—The name of the user group families by which process and image records are summarized if the model data switch is not enabled.

The daily Performance Agent data files and the archived data files reside in the database disk and directory as specified for the node in the CPD collection definition. A unique file exists for each node and for each unit of time defined by the granularity.

The size of a history file depends on the data reduction scheme. The most important element is the size interval you define. For example, if the interval is 60 minutes, the number of data records in the file is significantly smaller than if the interval were four minutes.

Also, if you reduce your history files according to a previously defined workload, the attributes of the workload definition affect the size of the resultant history file. For instance, if the workload definition was created with /UNIQUE_BY=image or /UNIQUE_BY=user qualifiers, your history files are larger because there is more data in a workload consisting of unique images or users.

Archive Performance Agent Data

When you are satisfied with the history file descriptor definitions you set up, you can archive data in two of the following ways:

- Manually begin the archiving process using the ADVISE ARCHIVE command. For a complete description of the command and its qualifiers, see the chapter "Performance Agent Commands."
- Submit a command procedure, such as the PSDC$ARCHIVE.COM, midnight to add the previous day's data to the history files. During product installation, PSDC$ARCHIVE.COM is placed in the PSDC$EXAMPLES area. Tailor this template command procedure and include it in your system's startup command procedure to assure automatic data archiving.

Use History Files

You can use history files as the data source for generating performance evaluation reports, analysis reports without factory rules, dump reports, models, and graphs. To specify a particular history file, use the /HISTORY_DESCRIPTOR qualifier when generating reports or models.

History files cannot be used to generate an analysis report from the Factory rules. Other rules, normally evaluated at analysis report time, are evaluated at archive time. This means you can generate an analysis report from history data using auxiliary rules.

The performance evaluation report output is different in that it includes a new section on the archived rule occurrences. This section includes the following data for each rule that had an occurrence:

- The rule identifier
- The number of rule occurrences
- The history interval time of the first and last rule occurrence
- The archiving process, for each history file descriptor found in the parameters file, reads the daily data file and create, update, or supersede the history file for each node. For a given history descriptor, the Periodicity attribute determines how the archiving process deals with daily data.

If no value is specified for Periodicity, each set of daily data is processed and added to the history file. Created to cover a specific time span, the history file is then updated with each succeeding set of daily data. In this case, updating the history file does not change the version number of the file.

If you specify a value for Periodicity in the history descriptor, daily data is used in calculating new values to supersede the data stored in the history file. For each set of daily data processed, archiving creates a new version of the history file for that time span. In this case, use the Parameter File Editor to limit the number of file versions you want to retain.

If you archive using the two default history file descriptors, MONTHLY_MODEL and MONTHLY_USER, then the archiving process summarizes data into history files according to each history definition. So it summarizes data into intervals representing one hour, 24-hours per day, and create one history file per calendar month for each history definition.

The file name is in the following format:

```
path:PSDC$nodename_dd-mmm-yyyy_dd-mmm-yyyy.name
```

path

Represents the directory specified in the collection definition.

nodename

Is the name of the node from which the daily data was archived.

dd-mmm-yyyy

Is the date representing the start date of the historical data file.

dd-mmm-yyyy

Is the date representing the end date of the historical data file.

name

Is the unique history file descriptor name identifying the database specification in the parameters file.

If the archiving run is terminated prematurely, a subsequent archiving or reporting command employs the journal file. The most recent version of the history file is used to produce a valid history file to update or to create the requested report.

While updating the history files, the archiving process creates a journal file for each history file being updated. This process names each journal file to correspond with the history file using a file type of *.name_JOU*. If the history file is invalid, the archiving process uses the journal file and the invalid file to recover.

No journal files are created for history files that have a history file descriptor periodicity attribute. In this case, the original history file is retained and no new version is created. The journal files are automatically deleted when no longer needed.

Chapter 5: Dump Reports

This chapter contains sample Performance Agent dump reports and information to help you interpret them.

This chapter describes the following topics:

- Reports for Selected Data or Time Periods
- Interpreting the Dump Reports

The Performance Agent dump reports allow you to view data extracted from database files for a specified time period. You can specify a collection definition or accept the default CPD collection definition. You can view items in a Performance Agent data record by using the ADVISE COLLECT REPORT DUMP_ALL command. The DUMP_ALL report provides the following subrecords:

- Metrics subrecords
- Parameters subrecords
- Rule subrecords
- Process subrecords
- Disk subrecords
- Tape subrecords
- Communications subrecords
- Configuration and SCS subrecords
- CPU subrecords
- Hot file subrecords (only in daily data)

The Performance Agent repeats this information for each record and for each node if more than one node is selected for reporting.

This section contains the following topics:

Reports for Selected Data or Time Periods (see page 54)
Interpret the Dump Reports (see page 55)

Reports for Selected Data or Time Periods

A single dump record represents a substantial amount of data because one record of the database file equals all of the individual data items for one interval during the scheduled collection time. A sizeable report with many printed pages of output can be generated from a one-record DUMP_ALL report that covers an interval of two minutes.

Subsets of Data

To view a subset of one record, you can select one of the following report types:

```
DUMP_COMMUNICATIONS
DUMP_CONFIGURATION
DUMP_CPUS
DUMP_DATACELLS
DUMP_DISKS
DUMP_HOTFILES
DUMP_METRICS
DUMP_PARAMETERS
DUMP_PROCESSES
DUMP_RULES
DUMP_TAPES
```

Time Periods

Although you usually request Analysis and Performance Evaluation Reports for time intervals of many hours or days, dump reports should cover a much shorter time span.

You can specify the time period to be covered in the command or you can specify a *dates file*. A dates file enables you to specify a series of days and times for reporting. A dates file is simply an ASCII text file you create that has a filename with a .txt extension and follows a format similar to the following example:

```
16-APR-2008 10:30,16-APR-2008 10:34
17-APR-2008 10:30,17-APR-2008 10:34
18-APR-2008 10:30,18-APR-2008 10:34
```

Note that you use one line for each day in ascending order. Any truncated time specification has the value of 0. The following command example shows the /DATES qualifier:

```
$ ADVISE COLLECT REPORT DUMP_ALL/DATES=dates.txt
```

Interpret the Dump Reports

Dump reports are 132 columns wide, but the samples in this chapter are scaled down to fit on a smaller page. The following sample report was generated by issuing the command:

```
$ ADVISE COLLECT REPORT DUMP_ALL/BEG=16-OCT-2008:12:00-
_$ /END=16-OCT-2008:12:02/OUT=DUMP_RPT/NODE=LATOUR
```

The DUMP_ALL report is the accumulation of all dump report types of recorded data. Data cells and rules are, technically, not recorded data. The rules record is a comma-separated list of rules that fired unless the /HISTORY qualifier is included in the command. The descriptions of all data items follow the sample reports.

Dump Report: Heading

The following example shows a sample heading on a dump report. The table after the example describes the information in the dump report heading.

```
Dump Report                        VM6 (HP rx2600  (1.40GHz/1.5M)                        Page   1
                                                                                    PSDC V3.1-0805

                                  Tuesday 18-JUL-2008 09:00 to 09:02



**  Metrics Record     Start Time Stamp:18-JUL-2008 09:00:00 Data Version: V2.2-55     VMS Version: V8.2-1

                                 End Time Stamp:  18-JUL-2008 09:02:00   Uptime in Seconds: 120
```

Item	Description
Dump Report	Title of the Performance Agent report
VM6 (HP rx260o (1.40GHx/1.5M)	Node (type of system)
Page 1	Page number
PSDC Vn.n-yym	Version of the Performance Agent
Metrics Record	Type of record
Time Stamp	Date and time data was recorded
Data Version	Version of the Performance Agent that recorded the information
VMS Version	Version of VMS running during the time period

Dump Report: Metrics Record

The following example shows a sample metrics record dump report. The table after the example describes the headings in the metrics record section of the dump report:

```
Dump Report                                I64BOX (HP rx2600  (1.40GHz/1.5M)                                      Page   1
                                                                                                               PSDC V3.0-0610
                                            Tuesday 27-JUN-2006 09:00 to 09:10

**      Metrics Record             Start Time Stamp:  27-JUN-2006 09:00:00   Data Version: V2.2-55          VMS Version: V8.2-1
                                     End Time Stamp:  27-JUN-2006 09:02:00   Uptime in Seconds: 120

--------------------------------------------------------  Process Averages  ------------------------------------------------------
Total Processes            36.09  Page Fault Wait              0.00  Free Page Wait               0.00  Collided Page Wait          0.00
Miscellaneous Wait          0.00  Current Executing            1.00  Common Event Flag            0.00
Local Event Flag Wait       7.03  Lcl Evt Flg Outswap Wait     0.00  Hibernate Wait              28.07  Hibernate Outswap Wait      0.00
Suspend Wait                0.00  Suspend Outswapped Wait      0.00  Computable                   0.00  Computable Outswapped       0.00
Interactive Processes       2.00  Batch processes              0.00  Network Processes            1.00
------------------------------------------------------  CPU Time Percentages  --------------------------------------------------
Kernel Mode                 0.25  Executive Mode               0.09  Supervisor Mode              0.00  User Mode                   0.05
Compatability Mode          0.00  Idle Time                   99.56  Interrupt Stack              0.06  MP Sychronization           0.00
---------------------------------------------------------  Paging and I/O  ----------------------------------------------------
Total PgFlts/Sec            4.33  Free List PgFlts/Sec         0.35  Modified List PgFlts/Sec     0.02  Demand Zero PgFlts/Sec      2.16
System PgFlts/Sec           0.00  Global PgFlts/Sec            1.38  Write In Prog PgFlts/Sec     0.00  Image Activations/Sec       0.03
Pages Read/Sec              0.64  Read Operations/Sec          0.43  Pages Written/Sec            0.00  Write Operations/Sec        0.00
Direct I/O /Sec             1.37  Buffered I/O /Sec           26.70  Free List Page Cnt         196150  Modified List Page Cnt      1138
Mailbox Reads/Sec           0.13  Mailbox Writes/Sec           0.13  Logical Name Transl/Sec    124.92  Inswaps/Sec                 0.00
Split Transfers/Sec         0.00  Trnsfr W/O Wind Turn/Sec     0.84  File Window Misses/Sec       0.05  Image Terminations/Sec      0.03
Transition PgFlts/Sec       0.00  Bad Page PgFlts/Sec          0.00  Outswaps/Sec                 0.00  HeadInSwaps/sec             0.00
HeadOutSwaps/sec            0.00
------------------------------------------------------  Cache Hit/Miss Rates  --------------------------------------------------
Directory Hits             23.02  Directory Misses             0.00  Quota Hits                   0.00  Quota Misses                0.00
File ID Hits                0.10  File ID Misses               0.00  Extent Hits                  0.21  Extent Misses               0.00
File Header Hits            1.03  File Header Misses           0.03  Direct. Datablock Hits      23.04  Direct. Datablock Misses    0.00
Storage Bit Map Hits        0.00  Storage Bit Map Misses       0.00
---------------------------------------------------------  XQP and Other  -----------------------------------------------------
Open File Rate/Sec          0.21  Erase QIOs/Sec               0.10  Vol Synch Lck/Sec            0.19  Vol Synch Lck Wts/Sec       0.00
Vol&Dir Synch Lck/Sec      24.27  Vol&Dir Synch Lck Wts/Sec    0.01  Access Lck/Sec              23.29  Cch Free Space Wts/Sec      0.03
File Sys CPU Time %         0.09  Open Files                 398.00
-------------------------------------------------------------  DECnet  --------------------------------------------------------
Arrive Local Packet/Sec     0.00  Depart Local Packet/Sec      0.00  Transit Packet/Sec           0.00  Transit Congest Loss/Sec    0.00
Receive Buff Fail/Sec       0.00
-------------------------------------------------------  Lock Manager Rates  ---------------------------------------------------
New Local Req/Sec          48.24  New Incoming Req/Sec         0.00  New Outgoing Req/Sec         0.00
Local Conver Req/Sec        1.61  Incoming Conver Req/Sec      0.00  Outgoing Conver Req/Sec      0.00
Local Dequeue Req/Sec      48.26  Incoming Dequeue Req/Sec     0.00  Outgoing Dequeue Req/Sec     0.00
Enqueue Req Waiting/Sec     0.02  Enqueue Req Not Queued/Sec   0.00
---------------------------------------------------  Asynchronous System Traps  ------------------------------------------------
Local Blocking ASTs Queued  0.01  Incoming Blocking ASTs       0.00  Outgoing Blocking ASTs       0.00
------------------------------------------------------  Directory Operations  --------------------------------------------------
Incoming Directory Oper.    0.00  Outgoing Directory Oper.     0.00
-----------------------------------------------------------  Deadlocks  -------------------------------------------------------
Incoming Deadlock Messages  0.00  Outgoing Deadlock Messages   0.00  Deadlock Searches            0.00  Deadlock Finds              0.00
---------------------------------------------------------  Pool & Memory  -----------------------------------------------------
NP Pool Maximum         18374656  NP Pool Blocks                  1  NP Free Blocks =< 32 byte       0  NP Free Bytes           10881536
Largest NP Block        10881536  Smallest NP Block        10881536
PG Pool Maximum         10354688  PG Pool Blocks                 48  PG Free Blocks =< 32 byte      31  PG Free Bytes            6551824
Largest PG Block         6549568  Smallest PG Block              16  User Memory Pages          229441  Page Size                   8192
--------------------------------------------------------  Resource Packets  ---------------------------------------------------
SRP Maximum                    0  SRPs in Use                     0  IRP Maximum                     0  IRPs in Use                    0
LRP Maximum                    0  LRPs in Use                     0  RDT Maximum                  8299  Average RDTs in Use         0.00
```

```
Dump Report                                    I64BOX (HP rx2600  (1.40GHz/1.5M)                                          Page   2
                                                                                                                      PSDC V3.0-0610
                                                Tuesday 27-JUN-2006 09:00 to 09:10

**      Metrics Record            Start Time Stamp:  27-JUN-2006 09:02:00   Data Version: V2.2-55         VMS Version: V8.2-1
                                    End Time Stamp:  27-JUN-2006 09:04:00   Uptime in Seconds: 120

Lock ID Table Maximum       4095  Lock IDs in Use             3097  Resource Table Maximum      2048  Resources in Use           3081
-------------------------------------------------------- MSCP Statistics ----------------------------------------------------------
Original Buffers            0.00  Free Buffers             1024.00  Smallest Buffer Allowed   128.00  Free Pool Bytes         1024.00
Buffer Wait Queue           0.00  Wait Queue Highwater Mark   0.00  Split Transfers/Sec         0.00  I/O Fragments/Sec          0.00
Operation Count/Sec         0.00  Read Count/Sec              0.00  Write Count/Sec             0.00
-------------------------------------------------------- Paging Statistics  --------------------------------------------------------
Total Paging Pages      262600.0  Available Paging Pages  262600.0
-------------------------------------------------------- CPU and I/O ---------------------------------------------------------------
Pagewait                    0.00  Swapwait                    0.00  PSWPwait                    0.00  CPUIOidle                  0.00
CPUonly                     0.00  IOonly                      0.00  MultiIO                     0.00  CPUIO                      0.00
VBSSCPU                     0.00
-------------------------------------------------------- Virtual I/O Cache ---------------------------------------------------------
Current Size               20400  Maximum Size               20400  I/O Reads/Sec               0.73  Cache Rd Hits/Sec          0.69
Free Pages                   864  In Use Pages               19536  I/O Writes/Sec              0.08  Files Retained              460
I/O WrtBypass/Sec           0.00  I/O RdBypass/Sec            0.00  Misses-Blksz LT 33             5  Misses-Blksz 33-64            0
Misses-Blksz 65-127            0  Misses-Blksz 128-255           0  Misses-Blksz GT 255            0
-------------------------------------------------------- Process Averages  ---------------------------------------------------------
Total Processes            36.00  Page Fault Wait             0.00  Free Page Wait              0.00  Collided Page Wait         0.00
Miscellaneous Wait          0.00  Current Executing           1.00  Common Event Flag           0.00
Local Event Flag Wait       7.01  Lcl Evt Flg Outswap Wait    0.00  Hibernate Wait             27.99  Hibernate Outswap Wait     0.00
Suspend Wait                0.00  Suspend Outswapped Wait     0.00  Computable                  0.00  Computable Outswapped      0.00
Interactive Processes       2.00  Batch processes             0.00  Network Processes           1.00
------------------------------------------------------ CPU Time Percentages  ------------------------------------------------------
Kernel Mode                 0.23  Executive Mode              0.08  Supervisor Mode             0.00  User Mode                  0.04
Compatability Mode          0.00  Idle Time                  99.60  Interrupt Stack             0.05  MP Sychronization          0.00
--------------------------------------------------------- Paging and I/O  ---------------------------------------------------------
Total PgFlts/Sec            0.00  Free List PgFlts/Sec        0.00  Modified List PgFlts/Sec    0.00  Demand Zero PgFlts/Sec     0.00
System PgFlts/Sec           0.00  Global PgFlts/Sec           0.00  Write In Prog PgFlts/Sec    0.00  Image Activations/Sec      0.00
Pages Read/Sec              0.00  Read Operations/Sec         0.00  Pages Written/Sec           0.00  Write Operations/Sec       0.00
Direct I/O /Sec             0.12  Buffered I/O /Sec          24.68  Free List Page Cnt        196173  Modified List Page Cnt     1138
Mailbox Reads/Sec           0.00  Mailbox Writes/Sec          0.00  Logical Name Transl/Sec   118.19  Inswaps/Sec                0.00
Split Transfers/Sec         0.00  Trnsfr W/O Wind Turn/Sec    0.32  File Window Misses/Sec      0.01  Image Terminations/Sec     0.00
Transition PgFlts/Sec       0.00  Bad Page PgFlts/Sec         0.00  Outswaps/Sec                0.00  HeadInSwaps/Sec            0.00
HeadOutSwaps/Sec            0.00
------------------------------------------------------ Cache Hit/Miss Rates  ------------------------------------------------------
Directory Hits             22.53  Directory Misses            0.00  Quota Hits                  0.00  Quota Misses               0.00
File ID Hits                0.00  File ID Misses              0.00  Extent Hits                 0.00  Extent Misses              0.00
File Header Hits            0.34  File Header Misses          0.00  Direct. Datablock Hits     22.53  Direct. Datablock Misses   0.00
Storage Bit Map Hits        0.00  Storage Bit Map Misses      0.00
--------------------------------------------------------- XQP and Other  ----------------------------------------------------------
Open File Rate/Sec          0.03  Erase QIOs/Sec              0.00  Vol Synch Lck/Sec           0.00  Vol Synch Lck Wts/Sec      0.00
Vol&Dir Synch Lck/Sec      22.70  Vol&Dir Synch Lck Wts/Sec   0.00  Access Lck/Sec             22.58  Cch Free Space Wts/Sec     0.00
File Sys CPU Time %         0.08  Open Files                398.00
------------------------------------------------------------- DECnet  -------------------------------------------------------------
Arrive Local Packet/Sec     0.00  Depart Local Packet/Sec     0.00  Transit Packet/Sec          0.00  Transit Congest Loss/Sec   0.00
Receive Buff Fail/Sec       0.00
```

```
Dump Report                            I64BOX (HP rx2600  (1.40GHz/1.5M)                                Page    3
                                                                                                    PSDC V3.0-0610
                                        Tuesday 27-JUN-2006 09:00 to 09:10

**      Metrics Record               Start Time Stamp:  27-JUN-2006 09:04:00   Data Version: V2.2-55          VMS Version: V8.2-1
                                       End Time Stamp:  27-JUN-2006 09:06:00   Uptime in Seconds: 120

-------------------------------------------------------  Lock Manager Rates  ------------------------------------------------------
New Local Req/Sec          45.29  New Incoming Req/Sec          0.00  New Outgoing Req/Sec          0.00
Local Conver Req/Sec        0.22  Incoming Conver Req/Sec       0.00  Outgoing Conver Req/Sec       0.00
Local Dequeue Req/Sec      45.29  Incoming Dequeue Req/Sec      0.00  Outgoing Dequeue Req/Sec      0.00
Enqueue Req Waiting/Sec     0.01  Enqueue Req Not Queued/Sec    0.00
---------------------------------------------------  Asynchronous System Traps  -------------------------------------------------
Local Blocking ASTs Queued 0.01  Incoming Blocking ASTs         0.00  Outgoing Blocking ASTs        0.00
------------------------------------------------------  Directory Operations  ----------------------------------------------------
Incoming Directory Oper.    0.00  Outgoing Directory Oper.      0.00
-----------------------------------------------------------  Deadlocks  -----------------------------------------------------------
Incoming Deadlock Messages 0.00  Outgoing Deadlock Messages    0.00  Deadlock Searches             0.00  Deadlock Finds              0.00
---------------------------------------------------------  Pool & Memory  ---------------------------------------------------------
NP Pool Maximum         18374656  NP Pool Blocks                   1  NP Free Blocks =< 32 byte      0  NP Free Bytes           10881536
Largest NP Block        10881536  Smallest NP Block         10881536
PG Pool Maximum         10354688  PG Pool Blocks                  48  PG Free Blocks =< 32 byte     31  PG Free Bytes            6551824
Largest PG Block         6549568  Smallest PG Block               16  User Memory Pages         229441  Page Size                   8192
-------------------------------------------------------  Resource Packets  --------------------------------------------------------
SRP Maximum                    0  SRPs in Use                      0  IRP Maximum                    0  IRPs in Use                    0
LRP Maximum                    0  LRPs in Use                      0  RDT Maximum                 8299  Average RDTs in Use         0.00
Lock ID Table Maximum       4095  Lock IDs in Use               3098  Resource Table Maximum      2048  Resources in Use            3082
-------------------------------------------------------  MSCP Statistics  ---------------------------------------------------------
Original Buffers            0.00  Free Buffers               1024.00  Smallest Buffer Allowed   128.00  Free Pool Bytes          1024.00
Buffer Wait Queue           0.00  Wait Queue Highwater Mark     0.00  Split Transfers/Sec         0.00  I/O Fragments/Sec           0.00
Operation Count/Sec         0.00  Read Count/Sec                0.00  Write Count/Sec             0.00
-------------------------------------------------------  Paging Statistics  -------------------------------------------------------
Total Paging Pages      262600.0  Available Paging Pages    262600.0
-------------------------------------------------------  CPU and I/O ---------------------------------------------------------------
Pagewait                    0.00  Swapwait                      0.00  PSWPwait                    0.00  CPUIOidle                   0.00
CPUonly                     0.00  IOonly                        0.00  MultiIO                     0.00  CPUIO                       0.00
VBSSCPU                     0.00
-------------------------------------------------------  Virtual I/O Cache --------------------------------------------------------
Current Size               20400  Maximum Size                 20400  I/O Reads/Sec               0.06  Cache Rd Hits/Sec           0.04
Free Pages                   864  In Use Pages                 19536  I/O Writes/Sec              0.03  Files Retained               461
I/O WrtBypass/Sec           0.00  I/O RdBypass/Sec              0.00  Misses-Blksz LT 33             2  Misses-Blksz 33-64             0
Misses-Blksz 65-127            0  Misses-Blksz 128-255             0  Misses-Blksz GT 255            0
-------------------------------------------------------  Process Averages  --------------------------------------------------------
Total Processes            36.00  Page Fault Wait               0.00  Free Page Wait              0.00  Collided Page Wait          0.00
Miscellaneous Wait          0.00  Current Executing             1.01  Common Event Flag           0.00
Local Event Flag Wait       7.00  Lcl Evt Flg Outswap Wait      0.00  Hibernate Wait             27.99  Hibernate Outswap Wait      0.00
Suspend Wait                0.00  Suspend Outswapped Wait       0.00  Computable                  0.00  Computable Outswapped       0.00
Interactive Processes       2.00  Batch processes               0.00  Network Processes           1.00
----------------------------------------------------  CPU Time Percentages  ------------------------------------------------------
Kernel Mode                 0.22  Executive Mode                0.09  Supervisor Mode             0.00  User Mode                   0.04
Compatability Mode          0.00  Idle Time                    99.60  Interrupt Stack             0.05  MP Sychronization           0.00
--------------------------------------------------------  Paging and I/O  --------------------------------------------------------
Total PgFlts/Sec            0.00  Free List PgFlts/Sec          0.00  Modified List PgFlts/Sec    0.00  Demand Zero PgFlts/Sec      0.00
System PgFlts/Sec           0.00  Global PgFlts/Sec             0.00  Write In Prog PgFlts/Sec    0.00  Image Activations/Sec       0.00
Pages Read/Sec              0.00  Read Operations/Sec           0.00  Pages Written/Sec           0.00  Write Operations/Sec        0.00
```

Item	Description
Total Processes	The average number of processes for the given interval
Page Fault Wait	The average number of processes in the scheduler state PFW, (Page Fault Wait)
Free Page Wait	The average number of processes in the scheduler state FPW, (Free Page Wait)
Collided Page Wait	The average number of processes in the scheduler state COLPG, (Collided Page Wait)
Miscellaneous Wait	The average number of processes in the scheduler state MWAIT (Miscellaneous Wait)
Current Executing	The average number of processes in the scheduler state CUR
Common Event Flag	The average number of processes in the scheduler state CEF, (Common Event Flag wait)

Item	Description
Local Event Flag Wait	The average number of processes in the scheduler state LEF, (Local Event Flag Wait)
Lcl Evt Flg Outswap Wait	The average number of processes in the scheduler state LEFO, (Local Event Flag Wait Outswapped)
Hibernate Wait	The average number of processes in the scheduler state HIB, (Hibernate)
Hibernate Outswap Wait	The average number of processes in the scheduler state HIBO, (Hibernate Outswapped)
Suspend Wait	The average number of processes in the scheduler state SUSP, (Suspended)
Suspend Outswapped Wait	The average number of processes in the scheduler state SUSPO, (Suspended Outswapped)
Computable	The average number of processes in the scheduler state COM, (Computable)
Computable Outswapped	The average number of processes in the scheduler state COMO, (Computable Outswapped)
Interactive Processes	The average number of interactive processes.
Batch	The average number of batch processes.
Network	The average number of network processes.
Kernel Mode	Percentage of the CPU time used in Kernel Mode
Executive Mode	Percentage of the CPU time used in EXEC Mode
Supervisor Mode	Percentage of the CPU time used in Supervisor Mode
User Mode	Percentage of the CPU time used in USER Mode
Compatibility Mode	Percentage of the CPU time used in COMPAT Mode
Idle Time	Percentage of the CPU time spent idle
Interrupt Stack	Percentage of the CPU time spent processing interrupts
MP Synchronization	Percentage of the CPU time used in MP_SYNCH Mode

Item	Description
Total PgFlts/Sec	Average number of page faults per second
Free List PgFlts/Sec	Average number of page faults from the free list per second
Modified List PgFlts/Sec	Average number of page faults from the modified list per second
Demand Zero PgFlts/Sec	Average number of demand zero page faults per second
System PgFlts/Sec	Average number of system page faults per second
Global PgFlts/Sec	Average number of global page faults per second
Write In Prog PgFlts/Sec	Average number of write-in-progress page faults per second
Image Activations/Sec	Average number of image activations per second
Pages Read/Sec	Average number of pages read in from disk due to page faults per second
Read Operations/Sec	Average number of disk read I/Os due to page faults per second
Pages Written/Sec	Average number of pages written to page files, swap files, and mapped image sections per second
Write Operations/Sec	Average number of write I/O operations to page files, swap files, and mapped image sections per second
Direct I/O /Sec	Average number of direct I/O operations per second
Buffered I/O /Sec	Average number of buffered I/O operations per second
Free List Page Cnt	Average number of pages on the free list
Modified List Page Cnt	Average number of pages on the modified list
Mailbox Reads/Sec	Average number of mailbox read operations per second
Mailbox Writes/Sec	Average number of mailbox write operations per second
Logical Name Transl/Sec	Average number of logical name translations per second

Item	Description
Inswaps/Sec	Average number of in-swap operations per second
Split Transfers/Sec	Average number of split I/O operations per second
Trnsfr W/O Wind Turn/Sec	The number of times the executive I/O subsystem successfully maps a virtual to logical segment without needing to invoke XQP services per second
File Window Misses/Sec	The number of times the XQP updates the Window Control Block per second
Image Terminations/Sec	Average number of image terminations per second
Transition PgFlts/Sec	The number of transition state faults per second
Bad Page PgFlts/Sec	The number of page faults resolved from the bad page list
Outswaps/Sec	The number of process outswaps
HeadInSwaps/Sec	The number of process headers swapped in
HeadOutSwaps/Sec	The number of process headers swapped out
Directory Hits	The average number of times per second that directory file control blocks were found in the directory cache
Directory Misses	The average number of times per second that directory file control blocks were not found in the directory cache
Quota Hits	The average number of times per second that entries were found in the quota cache
Quota Misses	The average number of times per second that entries were not found in the quota cache
File ID Hits	The average number of times per second that file identifiers were found in the FID cache
File ID Misses	The average number of times per second that file identifiers were not found in the FID cache
Extent Hits	The average number of times per second that file extents were found in the extent cache
Extent Misses	The average number of times per second that file extents were not found in the extent cache

Item	Description
File Header Hits	The average number of times per second that file headers were found in the file header cache
File Header Misses	The average number of times per second that file headers were not found in the file header cache
Direct. Datablock Hits	The average number of times per second that directory data blocks were found in the directory cache
Direct. Datablock Misses	The average number of times per second that directory data blocks were not found in the directory cache
Storage Bit Map Hits	The average number of times per second that entries were found in the bitmap cache
Storage Bit Map Misses	The average number of times per second that entries were not found in the bitmap cache
Open File Rate/Sec	Average number of file opens per second
Erase QIOs/Sec	Average number of erase I/O operations per second
Vol Synch Lck/Sec	XQP volume synch lock operations per second
Vol Synch Lck Wts/Sec	The number of times the XQP had to wait for vol synch lock per second
Vol&Dir Synch Lck/Sec	XQP Directory and volume synch lock operations per second
Vol&Dir Synch Lck Wts/Sec	The number of times the XQP had to wait for a directory and vol synch lock, per second
Access Lck/Sec	The average XQP access lock rate per second
Cch Free Space Wts/Sec	The average number of times the XQP had to wait for cache free space per second
File Sys CPU Time %	Percentage of the CPU time used while executing in the file system
Open Files	The average number of open files.
Arrive Local Packet/Sec	DECNET arriving local packets per second
Depart Local Packet/Sec	DECNET departing local packets per second
Transit Packet/Sec	DECNET transit packets handled per second
Transit Congest Loss/Sec	DECNET transit congestion losses per second

Item	Description
Receive Buff Fail/Sec	DECNET receiver buffer failure rate
New Local Req/Sec	Average number of local lock ENQ requests per second managed by the local node
New Incoming Req/Sec	Average number of remote lock ENQ requests per second managed by the local node
New Outgoing Req/Sec	Average number of local lock ENQ requests per second managed by a remote node
Local Conver Req/Sec	Average number of local lock conversion requests per second managed by the local node
Incoming Conver Req/Sec	Average number of remote lock conversion requests per second managed by the local node
Outgoing Conver Req/Sec	Average number of local lock conversion requests per second managed by a remote node
Local Dequeue Req/Sec	Average number of local lock DEQ requests per second managed by the local node
Incoming Dequeue Req/Sec	Average number of local lock DEQ requests per second managed by a remote node
Outgoing Dequeue Req/Sec	Average number of local lock DEQ requests per second managed by a remote node
Enqueue Req Waiting/Sec	Average number of local lock ENQs or conversions per second that are placed in the wait queue
Enqueue Req Not Queued/Sec	Average number of local lock ENQs or conversions per second that are not granted and also not placed in the wait queue
Local Blocking ASTs Queued	Average number of local lock requests per second that trigger a blocking AST on the local node
Incoming Blocking ASTs	Average number of remote lock requests per second that trigger a blocking AST on the local node
Outgoing Blocking ASTs	Average number of local lock requests per second that trigger a blocking AST on a remote node
Incoming Directory Oper.	Average number of incoming lock directory functions per second

Item	Description
Outgoing Directory Oper.	Average number of outgoing lock directory functions per second
Incoming Deadlock Messages	Average number of incoming deadlock messages per second
Outgoing Deadlock Messages	Average number of outgoing deadlock messages per second
Deadlock Searches	Average number of deadlock searches per second
Deadlock Finds	Average number of deadlock finds per second
NP Pool Maximum	Average number of non-paged pool bytes, used and free
NP Pool Blocks	Average number of non-paged pool contiguous portions
NP Free Blocks =< 32 byte	Average number of free non-paged pool contiguous portions whose size is less than or equal to 32 bytes
NP Free Bytes	Average number of free non-paged pool bytes
Largest NP Block	The size of the largest non-paged pool contiguous portion
Smallest NP Block	The size of the smallest non-paged pool contiguous portion
PG Pool Maximum	Average number of paged pool bytes, used and free
PG Pool Blocks	Average number of paged pool contiguous portions
PG Free Blocks =< 32 byte	Average number of free paged pool contiguous portions whose size is less than or equal to 32 bytes
PG Free Bytes	Average number of free paged pool bytes
Largest PG Block	The size of the largest paged pool contiguous portion
Smallest PG Block	The size of the smallest paged pool contiguous portion
User Memory Pages	The number of physical pages less the system working set, non-paged pool, and lookaside lists
Page Size	Size of memory pages

Item	Description
SRP Maximum	Average number of small request packets, used and unused
SRPs in Use	Average number of small request packets in use
IRP Maximum	Average number of I/O request packets, used and unused
IRPs in Use	Average number of I/O request packets in use
LRP Maximum	Average number of large request packets, used and unused
LRPs in Use	Average number of large request packets in use
RDT Maximum	Average number of request descriptor table entries, used and unused
Average RDTs in Use	Average number of request descriptor table entries in use
Lock ID Table Maximum	Average number of lock ID table entries, used and free
Lock IDs in Use	Average number of lock ID table entries in use
Resource Table Maximum	Number of entries in the resource hash table
Resources in Use	Number of known resources
Original Buffers	Number of MSCP buffers
Free Buffers	Number of free MSCP buffers
Smallest Buffer Allowed	Minimum size of an MSCP buffer
Free Pool Bytes	Total number of free MSCP buffer bytes
Buffer Wait Queue	Average number of requests waiting for MSCP buffers
Wait Queue Highwater Mark	Maximum number of requests waiting for MSCP buffers
Split Transfers/Sec	Average number of requests per second that had to be split
I/O Fragments/Sec	Average number of contiguous portions of MSCP buffer space
Operation Count/Sec	Average number of MSCP I/O operations per second

Item	Description
Read Count/Sec	Average number of MSCP read I/O operations per second
Write Count/Sec	Average number of MSCP write I/O operations per second
Total Paging Pages	Total number of pages in all paging files
Available Paging Pages	Average number of free pages in all paging files
Pagewait	The percentage of time that the CPU was idle and at least one disk device had paging I/O in progress
Swapwait	The percentage of time that the CPU was idle and at least one disk device had swapping I/O in progress
PSWPwait	The percentage of time that the CPU was idle and at least one disk had either page I/O or swap I/O in progress
CPUIOidle	The percentage of time that the CPU and all disk devices were idle
CPUonly	The percentage of time that a CPU was busy and no disk device was busy
IOonly	The percentage of time that the CPU or all CPUs in a multiprocessor system were idle and at least one disk device was busy
MultioIO	The percentage of time that two or more of the disk devices were busy
CPUIO	The percentage of time that both the CPU and at least one disk device were busy
VBSSCPU	Average percent of CPU time in the Interrupt Stack spent on behalf of Virtual Balance Set Slot transitions only.
Current Size	Virtual I/O Cache free pages
Maximum Size	Virtual I/O Cache maximum size in pages
I/O Reads/Sec	Virtual I/O Cache read I/O rate
Cache Rd Hits/Sec	Virtual I/O Cache read hit rate
Free Pages	Number of Virtual I/O free pages
In Use Pages	Number of Virtual I/O pages in use

Item	Description
I/O Writes/Sec	Virtual I/O Cache write I/O rate
Files Retained	Number of Virtual I/O Cache files retained
I/O WrtBypass/Sec	Virtual I/O Cache writes bypassing rate - writes bypassing the cache/second
I/O RdBypass/Sec	Virtual I/O Cache reads bypassing rate - reads bypassing the cache/second
Misses-Blksz LT 33	The number of read operations with a block size less than 33 that bypassed the eXtended File Cache for the current interval for the local node.
Misses-Blksz 33-64	The number of read operations with a block size from 33 to 64 that bypassed the eXtended File Cache for the current interval for the local node.
Misses-Blksz 65-127	The number of read operations with a block size from 65 to 127 that bypassed the eXtended File Cache for the current interval for the local node.
Misses-Blksz 128-255	The number of read operations with a block size from 128 to 255 that bypassed the eXtended File Cache for the current interval for the local node.
Misses-Blksz GT 255	The number of read operations with a block size greater than 255 that bypassed the eXtended File Cache for the current interval for the local node.

Dump Report: Parameter Record

The following example shows a sample parameters record dump report:

```
Dump Report                              I64BOX (HP rx2600  (1.40GHz/1.5M)                              Page   1
                                                                                                      PSDC V3.0-0610
                                          Thursday 29-JUN-2006 09:00 to 09:02

**   Parameters Record            Start Time Stamp:  29-JUN-2006 09:00:00   Data Version: V2.2-55          VMS Version: V8.2-1
                                    End Time Stamp:  29-JUN-2006 09:02:00   Uptime in Seconds: 120

NPAGEDYN              18325504  NPAGEVIR             100032512  PAGEDYN               10354688  GBLPAGES              3562717
GBLSECTIONS              1620  SPTREQ                    0  IRPCOUNT                     0  IRPCOUNTV                   0
SRPCOUNT                    0  SRPCOUNTV                 0  SRPSIZE                      0  FREELIM                   726
LRPCOUNT                    0  LRPCOUNTV                 0  LRPSIZE                      0  FREEGOAL                 2824
GROWLIM                   726  BORROWLIM               726  SYSMWCNT                 12143  MPW_HILIMIT             65524
MPW_LOLIMIT              2118  MPW_THRESH              4236  MPW_WAITLIMIT            65780  MPW_WRTCLUSTER             64
DORMANTWAIT                 2  DEFPRI                    4  MAXPROCESSCNT              708  BALSETCNT                 706
QUANTUM                     5  PFRATL                    0  PFRATH                       8  MMG_CTLFLAGS                3
PFCDEFAULT                 64  MINWSCNT                 20  WSINC                     2400  WSDEC                    4000
WSMAX                  784384  AWSMIN                  512  AWSTIME                      5  LONGWAIT                   30
SWPRATE                   500  SWPOUTPGCNT             512  SWPALLOCINC                64  ACP_MAPCACHE              353
ACP_HDRCACHE             1412  ACP_DIRCACHE           1412  ACP_WORKSET                  0  ACP_DINDXCACHE            353
ACP_FIDCACHE               64  ACP_EXTCACHE             64  ACP_QUOCACHE               708  ACP_EXTLIMIT              100
LOCKDIRWT                   0  LOCKIDTBL              1792  RESHASHTBL                2048  DEADLOCK_WAIT              10
IOTA                        1  PHYSICALPAGES        262096  PIXSCAN                     40  MSCP_CREDITS               32
MPW_LOWAITLIMIT         65268  MPW_IOLIMIT               4  SCSBUFFCNT                  50  SCSCONNCNT                  6
SCSRESPCNT                300  SCSMAXDG                576  SCSMAXMSG                  216  POOLCHECK                   0
MULTIPROCESSING             3  VBSSENA                   0  CACHE_STATE                 16  GBLPAGFIL              197852
MULTITHREAD                 2  NPAG_INTERVAL            30  CLUSTER_CREDITS             32  NPAG_AGGRESSIVE            100
NPAG_GENTLE               100  LCKMGR_MODE               0  SMP_CPUS                    -1  LOAD_SYS_IMAGES              7
PQL_DWSDEFAULT           5936  PQL_MWSDEFAULT        16384  PQL_DWSQUOTA             11872  PQL_MWSQUOTA             32768
PQL_DWSEXTENT          784384  PQL_MWSEXTENT         784384

User Command: ADVISE COLLECT REPORT DUMP_PARAMETER/BEGIN=29-JUN-2006 09:00:00.00/END=29-JUN-2006 09:02:00.00/OUT=DUMP_PARAMETER.OUT
```

The items identified in the parameters record are SYSGEN parameters. For more information about these parameters, see the OpenVMS documentation on system generation.

Dump Report: Rules Record

The rules record section of the dump report shows the rule ID and the number of occurrences of a particular rule in a time segment. Rules records are stored in an archive file and are available in a dump when the /HISTORY qualifier is specified. When /HISTORY is not specified, in place of a full record you see a comma-separated list of the rule IDs that fired. For more information about rules see the *Performance Manager Administrator* Guide.

The following example shows a sample rules record dump report:

```
**Rules Record                  Start Time Stamp:  02-NOV-1992 01:00:00   Data Version: V1.1-36          VMS Version: V5.4-3
                                  End Time Stamp:  02-NOV-1992 02:00:00   Uptime in Seconds: 3600

Rule ID   Occurrences
-------   -----------
I0160               1
R0096               7
C0010               1

**Rules Record                  Start Time Stamp:  02-NOV-1992 02:00:00   Data Version: V1.1-36          VMS Version: V5.4-3
                                  End Time Stamp:  02-NOV-1992 03:00:00   Uptime in Seconds: 3600

Rule ID   Occurrences
-------   -----------
R0096               3
R0095               4


User Command: ADVISE COLL REPO DUMP_RULES/BEG=2-NOV-1992 01:00:00.00/END=2-NOV-1992 03:00:00.00/HIST=MONTHLY_USER/OUT=DUMP_RULES2-
              /NODE=YQUEM
```

Dump Report: Process Record

The following example shows a sample process record dump report:

```
**Process Record                Start Time Stamp:  07-JAN-1993 12:00:00   Data Version: V1.1-38          VMS Version: V5.4-3
                                  End Time Stamp:  07-JAN-1993 12:02:00   Uptime in Seconds: 120

User: DECNET_MAIL   Image: DSA111:[SYS1.SYSCOMMON.][SYSEXE] NETSERVER
     UIC: [000376,000374]  Pid: 28002F2B  Mpid: 28002F2B    Account: CNB         Process name: SERVER_00B1
     Cpu tics:      13     Uptime:      41.5200 secs.                    VA space:   5168           RWM:   00000000
     CurPri:  5     Gblpgs:       76  WSdef:     818  Bufio/sec:       0   SofPgFlt:    262 /CpuSec   State:       CUR     BPri:  4
     Ac:  1 Tm:  0  Privpgs:     288  WSquo:    1636  Dirio/sec:       0   HrdPgFlt:     15 /CpuSec   ImgAct: 0.024/sec      IBN:  2
     Li:  0 Lo:  0  WSlist:      818  WSext:    9216  Dskio/sec:    0.26   Dsk thpt:   0.77 Pgs/sec   Compct:        0     AWSA:  0
     Top Dev: DSA16                      Top IO/Sec:    0.19              2nd Dev: DSA111                        2nd IO/Sec:    0.07
     Tapio/sec:     0.0   Tap thpt:    0.0 Pgs/sec   RespTim:      0.0    RespDur:      0.0    RspWait:    200.0
     Trmio/sec:    0.00   Trm thpt:   0.00 Chr/sec   ThnkTim:      0.0    CmdWait:      0.0    ProcType 00000004

User: SYSTEM        Image:   Null
     UIC: [000000,000000]  Pid: 28000100  Mpid: 28000100    Account:             Process name: NULL
     Cpu tics:       0     Uptime:     120.0000 secs.                    VA space:      1           RWM:   00000000
     CurPri:  0     Gblpgs:       0  WSdef:       1  Bufio/sec:       0   SofPgFlt:      0 /CpuSec   State:       CUR     BPri:  0
     Ac:  0 Tm:  0  Privpgs:       0  WSquo:       1  Dirio/sec:       0   HrdPgFlt:      0 /CpuSec   ImgAct: 0.000/sec      IBN:  0
     Li:  0 Lo:  0  WSlist:        0  WSext:       1  Dskio/sec:    0.00   Dsk thpt:   0.00 Pgs/sec   Compct:        0     AWSA:  1
     Top Dev:                            Top IO/Sec:    0.00              2nd Dev:                               2nd IO/Sec:    0.00
     Tapio/sec:     0.0   Tap thpt:    0.0 Pgs/sec   RespTim:      0.0    RespDur:      0.0    RspWait:      0.0
     Trmio/sec:    0.00   Trm thpt:   0.00 Chr/sec   ThnkTim:      0.0    CmdWait:      0.0    ProcType 00000000

User: SYSTEM        Image:   Swapper
     UIC: [000001,000004]  Pid: 28000101  Mpid: 28000101    Account:             Process name: SWAPPER
     Cpu tics:       0     Uptime:     120.0000 secs.                    VA space:      1           RWM:   00000000
     CurPri: 16     Gblpgs:       0  WSdef:       1  Bufio/sec:       0   SofPgFlt:      0 /CpuSec   State:       HIB     BPri: 16
     Ac:  0 Tm:  0  Privpgs:       0  WSquo:       1  Dirio/sec:       0   HrdPgFlt:      0 /CpuSec   ImgAct: 0.000/sec      IBN:  0
     Li:  0 Lo:  0  WSlist:        0  WSext:       1  Dskio/sec:    0.00   Dsk thpt:   0.00 Pgs/sec   Compct:        0     AWSA:  1
     Top Dev:                            Top IO/Sec:    0.00              2nd Dev:                               2nd IO/Sec:    0.00
     Tapio/sec:     0.0   Tap thpt:    0.0 Pgs/sec   RespTim:      0.0    RespDur:      0.0    RspWait:      0.0
     Trmio/sec:    0.00   Trm thpt:   0.00 Chr/sec   ThnkTim:      0.0    CmdWait:      0.0    ProcType 00000000

User: DECNET        Image: DSA111:[SYS1.SYSCOMMON.][SYSEXE] NETACP
     UIC: [000001,000004]  Pid: 28000115  Mpid: 28000115    Account: <start>     Process name: NETACP
     Cpu tics:      54     Uptime:     120.0000 secs.                    VA space:  11611           RWM:   00000000
     CurPri: 10     Gblpgs:       0  WSdef:     818  Bufio/sec:       0   SofPgFlt:      0 /CpuSec   State:       HIB     BPri:  8
     Ac:  0 Tm:  0  Privpgs:    8678  WSquo:    9000  Dirio/sec:       0   HrdPgFlt:      0 /CpuSec   ImgAct: 0.000/sec      IBN:  0
     Li:  0 Lo:  0  WSlist:    11918  WSext:   15000  Dskio/sec:    0.00   Dsk thpt:   0.00 Pgs/sec   Compct:        0     AWSA:  0
     Top Dev:                            Top IO/Sec:    0.00              2nd Dev:                               2nd IO/Sec:    0.00
     Tapio/sec:     0.0   Tap thpt:    0.0 Pgs/sec   RespTim:      0.0    RespDur:      0.0    RspWait:      0.0
     Trmio/sec:    0.00   Trm thpt:   0.00 Chr/sec   ThnkTim:      0.0    CmdWait:      0.0    ProcType 00000008

User Command: ADVISE COLL REPO DUMP_PROC/BEG=7-JAN-1993 12:00:00.00/END=7-JAN-1993 12:02:00.00/OUT=DUMP_PROC/NODE=LATOUR
```

The following table describes the headings in the process record section of the dump report:

Item	Description
User	User name associated with the PID

Item	Description
Image directory	Location from which image came
Image	Name of the image the user is executing
UIC	The UIC of the process
Pid	Process ID number
Mpid	The master process ID number or owner process ID number
Account	The account name
Process name	The process name
Cpu tics	CPU time charged to the image in 10-millisecond units. This metric equals the same number with its decimal point moved two places to the right. For example, 100 tics equals 1.00 second.
Uptime	Up time per image in seconds.
VA space	Peak virtual address space used
RWM	Resource wait mask
CurPri	Current priority
Gblpgs	Average count of global pages
WSdef	Default working set size
Bufio/sec	Buffered I/O rate
SofPgFlt	Rate of total page faults per CPU second
State	State of process
BPri	Base priority
Ac	Image activation flag, with 0 = no activation and 1 = activation
Tm	Image termination flag, with 0 = no termination and 1 = termination
Privpgs	Average count of private pages
WSquo	Working set quota
Dirio/sec	Direct I/O rate
HrdPgFlt	Rate of hard page faults per CPU second
ImgActs	Rate of image activations per second
IBN	Process type, with 0 = Interactive, 1 = Batch, and 2 = Network

Item	Description
Li	Login flag, with 0 = no login and 1 = login
Lo	Logout flag, with 0 = no logout and 1 = logout
WSlist	Average working set size
WSext	Working set extent
DskIO/sec	Disk I/O per second
Dsk thpt	I/O throughput in bytes per second
Compct	Percent of uptime in COM(O) state, sampled every 5 seconds
AWSA	Automatic working set adjustment flag, with 0 = AWSA enabled and 1 = AWSA disabled
Top Dev	Disk or tape drive with the highest operation rate from this process
Top IO/Sec	Operation rate per second, to the disk or tape drive with the highest operation rate, from this process
2nd Dev	Disk or tape drive with the second highest operation rate from this process
2nd IO/Sec	Operation rate per second, to the disk or tape drive with the second highest operation rate, from this process
Tapio/sec	Number of I/O operations per second that the process issued to tape devices
Tap thpt	Number of bytes per second transferred to and from tape devices for the process
RespTim	Response time. The time from the completion of a terminal input to the next terminal input or output or to the end of the interval or image termination. The unit of time is milliseconds.
RespDur	Response time in milliseconds. The time from the completion of a terminal input to the next terminal input or to the end of the interval or image termination. This response time plus think time equals 100% of the interval.
RespWait	Either the duration of response time that continues into the next interval (if PRO_M_RT_ENDS is set) or the duration of the most recently completed response time within the interval (if PRO_M_RT_ENDS is clear.) The unit of time is milliseconds.
Trmio/sec	Number of read I/O operations per second issued to the process's terminal.

Item	Description
Trm thpt	Number of characters per second read from the process's terminal (typing rate).
Thnktim	Think time. The time from the start of a terminal input to the completion of that terminal input or to the end of the interval or image termination. The unit of time is milliseconds.
CmdWait	Either the duration of think time that continues into the next interval (if PRO_M_TT_END is set) or the duration of the most recently completed think time within the interval (if PRO_M_RT_ENDS is clear). The unit of time is milliseconds.
ProcType	The process type indicates interactive, batch, network, detached, and/or subprocess. This hexadecimal number is derived from the value of the bits set for the process. Bit 0 interactive bit 3 detached bit 1 batch bit 4 subprocess bit 2 network For example, a value of 1 indicates an interactive process; a value of 12 (hex) indicates a process that is both a batch process and a subprocess.

Dump Report: Disk Records

The following example shows a sample disk records dump report.

```
**Disk Records                  Start Time Stamp:  07-JAN-1993 12:00:00   Data Version: V1.1-38          VMS Version: V5.4-3
                                  End Time Stamp:  07-JAN-1993 12:02:00   Uptime in Seconds: 120

            ---Serv Node---  Allo -Controller-- --Adapter---  PDT ----I/O Oper/SEC----  --MSCP--  Avg  Thruput     Serv Sft Free
Volume Name  Name  Num Type  Clss Type Unit UTy Type Num TR   Flg Totl Pag Swp Rd  Spl Totl P/S  QLen Byt/Sec Bus Time Err Blocks
-----------  ----- --- ----  ---- ---- ---- --- ---- --- --   --- ---- --- --- --- --- ---- --- ---- ------- --- ---- --- -------
BOONE              1  VAX    0  DSA    0  30    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0   51618
WINO               1  VAX    0  DSA    1  30    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00      17   0   45   0  233352
PAGE3              1  VAX    0  DSA    2  21    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0   40800
MOET               1  VAX    0  DSA    5  30    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0   37812
REDWINE            1  VAX    0  DSA    6  30    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0  816810
PAGE2              1  VAX    0  DSA    7  21    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0   22916
MADDOG             1  VAX    0  DSA    8  30    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0  185799
MERLOT             1  VAX    0  DSA    9  30    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0   39990
PORT               1  VAX    0  DSA   10  21    0   0  4    1  0.5 0.0 0.0 0.4 0.0  0.0 0.0 0.01     341   1   25   0  367821
SHERRY             1  VAX    0  DSA   12  21    0   0  4    1  0.1 0.0 0.0 0.1 0.0  0.0 0.0 0.00     107   0   37   0  691809
MIDNIGHT_ENG       1  VAX    0  DSA   13  30    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0  230703
PAGE1A             1  VAX    0  DSA   14  21    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0   40868
BULLSBLOOD         1  VAX    0  DSA   15  30    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0   89529
CHABLIS2           1  VAX    0  DSA   16  30    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00      13   0   33   0  244968
RIPPLE             1  VAX    0  DSA   17  21    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0   97152
RHONE              1  VAX    0  DSA   18  30    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0  255315
ROSE               1  VAX    0  DSA   23  39    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00      21   0   30   0  196650
RUM                1  VAX    0  DSA   24  39    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00      30   0   20   0  784371
THUNDERBIRD        1  VAX    0  DSA   25  39    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0   78741
TOKAY              1  VAX    0  DSA   26  39    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0  482370
MODEXCH            1  VAX    0  DSA   27  39    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0  279810
MOSEL              1  VAX    0  DSA   28  39    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0  413589
BORDEAUX           1  VAX    0  DSA   29  39    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0  416133
ZINFANDEL          1  VAX    0  DSA   30  39    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00     315   0   30   0  278130
BURGUNDY           1  VAX    0  DSA   31  39    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0  160659
COOLER             1  VAX    0  DSA   32  39    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0  107655
CHAMPAGNE          1  VAX    0  DSA   33  64    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0  632064
RIESLING           1  VAX    0  DSA   37  64    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0  367275
SKUNKWORKS         1  VAX    0  DSA   38  64    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0  574482
KORBEL             1  VAX    0  DSA   39  64    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0      57
BELLET             1  VAX    0  DSA   40  64    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0  603039
CHAMBORD           1  VAX    0  DSA   41  30    0   0  4    1  0.0 0.0 0.0 0.0 0.0  0.0 0.0 0.00       0   0    0   0  196896

User Command: ADVISE COLL REPO DUMP_DISK/BEG=7-JAN-1993 12:00:00.00/END=7-JAN-1993 12:02:00.00/OUT=DUMP_DISK/NODE=LATOUR
```

For reports covering specific disks, see the description of the /FILTER qualifier for the ADVISE COLLECT REPORT DUMP_keyword command in the chapter Use the DECwindows Motif Interface (see page 191).

The following table describes the headings in the disk records of the dump report.

Item	Description
Volume Name	Volume name of the disk
Serv Node Name	Name of the node on which the volume is served
Serv Node Num	Node number of serving node
Serv Node Type	Hardware type of the serving node
Allo Clss	Allocation class
Controller Type	Controller type
Controller Unit	Unit number of the controller
Controller UTy	Type of unit
Adapter Type	Adapter type
Adapter Num	Adapter number
Adapter TR	Adapter TR slot number
PDT Flg	A byte field indicating CI, UDA, NI, BVP, DSSI, and SWIFT ports
I/O Oper/SEC Totl	Total I/O operations per second
I/O Oper/SEC Pag	Rate of paging operations per second
I/O Oper/SEC Swp	Rate of swap operations per second
I/O Oper/SEC Rd	Rate of read operations per second
I/O Oper/SEC Spl	Rate of split operations per second
MSCP Totl	Total MSCP operations per second
MSCP P/S	MSCP served paging or swapping operations per second
Avg QLen	Average queue length
Thruput Byt/sec	Number of bytes per second transferred from the disk
Bus	Percentage of time that I/O operations are outstanding
Serv Time	Service time: The average number of milliseconds between the start of an I/O operation to the disk and its completion
Sft Err	Count of soft errors
Free Blocks	Count of free blocks

The following example shows a sample disk records part 2 dump report:

```
**Disk Records Part 2              Start Time Stamp:  07-JAN-1993 12:00:00   Data Version: V1.1-38         VMS Version: V5.4-3
                                     End Time Stamp:  07-JAN-1993 12:02:00   Uptime in Seconds: 120
                                                                                                                  CMMHHVSDMCC MS
              ---------------------------Cache Sizes-------------                                  Hsc Disk Config  Shadow   PAESBOTUNLD2SR
Volume Name  DirIdx DirData  Map  Quota FileId FileExt FileHeader        Cache Name          HscNod Port Req    Root    YSMCSLRATUPPCV
------------ ------ ------- ----- ----- ------ ------- ---------- ----------------------- ------ ---- --- -------- --------------

BOONE          4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 8166FD10 01001000110010
WINO           4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 8167C880 01001000110010
PAGE3          4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 81603350 01001000110010
MOET           4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 81689990 01001000110010
REDWINE        4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 816962E0 01001000110010
PAGE2          4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 815F6AB0 01001000110010
MADDOG         4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 81697030 01001000110010
MERLOT         4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 816AFBD0 01001000110010
PORT           4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 816BC7D0 01001000110010
SHERRY         4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 81689800 01001000110010
MIDNIGHT_ENG   4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 816DB580 01001000110010
PAGE1A         4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 815E9040 01001000110010
BULLSBLOOD     4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 816F2960 01001000110010
CHABLIS2       4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 815DC4F0 01001000110010
RIPPLE         4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 816FF4B0 01001000110010
RHONE          4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 81710CE0 01001000110010
ROSE           4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 81648290 01001000110010
RUM            4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 816F1200 01001000110010
THUNDERBIRD    4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 8172A090 01001000110010
TOKAY          4096    4096  4096   220   4096    4096       4096 _DSA111:XQPCACHE                 0    0 8172A220 01001000110010
MODEXCH        4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 81748EE0 01001000110010
MOSEL          4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 81654C00 01001000110010
BORDEAUX       4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 81748CB0 01001000110010
ZINFANDEL      4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 8162EEC0 01001000110010
BURGUNDY       4096    4096  4096   220   4096    4096       4096 _DSA111:XQPCACHE                 0    0 8176ED90 01001000110010
COOLER         4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 8163B920 01001000110010
CHAMPAGNE      4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 81748B00 01001000110010
RIESLING       4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 8177D760 01001000110010
SKUNKWORKS     4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 81661690 01001000110010
KORBEL         4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 81797710 01001000110010
BELLET         4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 8176F3F0 01001000110010
CHAMBORD       4096    4096  4096     0   4096    4096       4096 _DSA111:XQPCACHE                 0    0 817A47D0 01001000110010


User Command: ADVISE COLL REPO DUMP_DISK/BEG=7-JAN-1993 12:00:00.00/END=7-JAN-1993 12:02:00.00/OUT=DUMP_DISK/NODE=LATOUR
```

The following table shows the heading for the disk records, part 2, in the dump report:

Item	**Description**
Volume Name	Volume name of the disk
Cache Sizes Dir Idx	Actual size of directory index cache
Cache Sizes Dir Data	Actual size of the directory data cache
Cache Sizes Map	Actual size of the map cache

Item	Description
Cache Sizes Quota	Actual size of the quota cache
Cache Sizes File Id	Actual size of the file ID cache
Cache Sizes File Extent	Actual size of the file extent cache
Cache Sizes File Header	Actual size of the file header cache
Cache Name	Name of the extended QIO processor cache
HSC Disk Config HscNod	Name of the HSC node
HSC Disk Config Port	Port number. Port values are not collected for shadow set members. By assigning the logical name PSDC$nodename_unit# to an equivalence string of REQ#_PORT#, the Performance Agent will use these for any shadow set members, and collect them.
HSC Disk Config Req	Requestor number. Requestor values are not collected for shadow set members. By assigning the logical name PSDC$nodename_unit# to an equivalence string of REQ#_PORT#, the Performance Agent will use these for any shadow set members, and collect them.
Shadow Root	The UCB address of the master
Device Status Bits	
CPY	Shadow set copy in progress.
MAS	Device is master or virtual unit of a shadow set, stripe set, or volume set.
MEM	Device is member physical unit of a shadow set, stripe set, or volume set.
HSC	Device is an HSC shadow set.
HBS	Device is a host based shadow set.
VOL	Device is a volume set.
STR	Device is a stripe set.
DUA	Device is dual ported.
MNT	Device is mounted.
CLU	Device is clusterwide.
CDP	Device has two UCBs.
2P	Device has alternate path.

Item	Description
MSC	Device is accessed via MSCP.
SRV	Device is MSCP served.

Dump Report: Tape Record

The following example shows a sample tape records dump report. The following table describes the headings in the tape record section of the dump report.

```
**Tape Records                  Start Time Stamp:  07-JAN-1993 12:00:00   Data Version: V1.1-38          VMS Version: V5.4-3
                                  End Time Stamp:  07-JAN-1993 12:02:00   Uptime in Seconds: 120

             -------Servicing Node-------  Alloc  -----Controller------  --------Adapter--------  PDT   I/O   Soft
Volume Name    Name      Number  HW Type   Class   Type   Unit  UType   Type  Number  TR Slot  Flag  Sec   Errors
-----------  --------    ------  -------   -----  ------  ----  -----   ----  ------  -------  ----  ----  ------
Not Mounted  LATOUR           1      VAX       0    RTFA     0    255      0       0        0     0  0.00       0
Not Mounted  LATOUR           1      VAX       0    RTFB     0    255      0       0        0     0  0.00       0
Not Mounted  LATOUR           1      VAX       0    RTFC     0    255      0       0        0     0  0.00       0
Not Mounted  LATOUR           1      VAX       0    RTFD     0    255      0       0        0     0  0.00       0
Not Mounted  LATOUR           1      VAX       0    RTFE     0    255      0       0        0     0  0.00       0
Not Mounted  LATOUR           1      VAX       0    RTFF     0    255      0       0        0     0  0.00       0
Not Mounted  LATOUR           1      VAX       0    RTFG     0    255      0       0        0     0  0.00       0
Not Mounted  LATOUR           1      VAX       0    RTFH     0    255      0       0        0     0  0.00       0
Not Mounted  BLUE             12     HS70      2     MUA     6      6      0       0        4     1  0.00       0
Not Mounted  JULIO            15     HS70      2     MUA    21     36      0       0        4     1  0.00       0
Not Mounted  JULIO            15     HS70      2     MUA    23     36      0       0        4     1  0.00       0
Not Mounted  JULIO            15     HS70      2     MUA    25     36      0       0        4     1  0.00       0
Not Mounted  JULIO            15     HS70      2     MUA    27     36      0       0        4     1  0.00       0

User Command: ADVISE COLL REPO DUMP_TAPE/BEG=7-JAN-1993 12:00:00.00/END=7-JAN-1993 12:02:00.00/OUT=DUMP_TAPE/NODE=LATOUR
```

Item	Description
Volume Name	Volume name of the tape
Servicing Node Name	Name of the node on which the volume is served
Servicing Node Number	Node number
Servicing Node HW Type	Hardware type of the node
Alloc Class	Allocation class
Controller Type	Controller type
Controller Unit	Unit number of the controller
Controller Utype	Type of unit

Item	Description
Adapter Type	Adapter type
Adapter Number	Adapter number
Adapter TR Slot	Adapter TR slot number
PDT Flag	A byte field indicating CI, UDA, NI, BVP, DSSI, and SWIFT ports.
I/O Sec	Operations count in I/Os per second
Soft Errors	Count of soft errors

Dump Report: Communications Record

The following example shows a sample communication records dump report. The following table describes the headings in the communications record section of the dump report.

```
**Communication Records          Start Time Stamp:  07-JAN-1993 12:00:00   Data Version: V1.1-38          VMS Version: V5.4-3
                                   End Time Stamp:  07-JAN-1993 12:02:00   Uptime in Seconds: 120

        Controller      Op Cnt
        ----------      ------
          TXA             0.0
          VTA             0.0
          LTA             0.0
          RTA             0.6
          RTB             0.0
          RTFA            0.0
          RTFB            0.0
          RTFC            0.0
          RTFD            0.0
          RTFE            0.0
          RTFF            0.0
          RTFG            0.0
          RTFH            0.0

User Command: ADVISE COLL REPO DUMP_COMM/BEG=7-JAN-1993 12:00:00.00/END=7-JAN-1993 12:02:00.00/OUT=DUMP_COMM/NODE=LAIOUR
```

Item	Description
Controller	Three-letter code describing the device type for each controller in the system
Op Cnt	Total operations rate for the associated controller

Dump Report: Configuration and SCS Record

The following example shows a sample configuration and SCS records dump report. The following table describes the headings in the report.

```
**Configuration & SCS Records      Start Time Stamp:  07-JAN-1993 12:00:00   Data Version: V1.1-38          VMS Version: V5.4-3
                                     End Time Stamp:  07-JAN-1993 12:02:00   Uptime in Seconds: 120

Nodes in                                  Adapter Port  Datagram Rates  Msg Rates  Blck Rates  ---Kilobyte Rates---  ---Wait Rates--  Flags
Cluster       Hardware Name               Name    Path  Send Recv Disc  Send Recv  Send Reqst   Send   Recv   Map    Snd Crd Buf Dsc  MSRVNC
--------  ------------------------  -------  ----  ---- ---- ----  ---- ----  ---- -----  -------  -----  -------  ------- -------  ------
LATOUR    VAX 8700                  CIBCI    PAA0     0    0    0     0    0     0     0     0.00   0.00        0     0.00    0.00  100101
BLUE      ES70                              PAA0     0    0    0     0    0     0     0     0.00   0.00        0     0.00    0.00  100000
GALLO     VAX 8700                  CIBCA-B  PAA0     0    0    0     6    5     0     0     0.00   0.01        0     0.01    0.00  100101
YQUEM     VAX 6000-440              CIBCA-B  PAA0     0    0    0     6    9     0     0     0.00   0.00        0     0.00    0.00  100101
ERNEST    ES70                              PAA0     0    0    0     0    1     0     0     0.00   0.00        0     0.00    0.00  100000
NUN       ES70                              PAA0     0    0    0     0    0     0     0     0.00   0.00        0     0.00    0.00  100000
JULIO     ES70                              PAA0     0    0    0     1    1     0     0     0.00   0.00        0     0.00    0.00  100000

User Command: ADVISE COLL REPO DUMP_CONF/BEG=7-JAN-1993 12:00:00.00/END=7-JAN-1993 12:02:00.00/OUT=DUMP_CONF/NODE=LATOUR
```

Item	Description
Nodes in cluster	Node names for active cluster
Kilobyte Rates Send	Kilobyte send rate
Hardware Name	Hardware description of node
Kilobyte Rates Recv	Kilobyte receive rate
Adapter Name	Hardware description of adapter for VAX ports
Kilobyte Rates Map	Kilobyte map rate
Port Path	Device name of the port
Wait Rates Snd Crd	Send credit waits per rate
Datagram Rates Send	Datagram send rate
Wait Rates Buf Dsc	Buffered descriptor waits per second
Datagram Rates Recv	Datagram receive rate

Item	Description
Flags MHRVNC	Flags, where 1=yes, 0=no M = Cluster member H = HSC R = Port is on DSSI V = VAX node N = Port is on NI C = Port is on CI
Datagram Rates Disc	Datagram discard rate
Msg Rates Send	Message send rate
Msg Rates Recv	Message receive rate
Blck Rates Send	Block send rate
Blck Rates Reqst	Block request rate

Dump Report: CPU Record

The following example shows a sample CPU records dump report. The following table describes the headings in the CPU record section of the dump report.

```
**CPU Records                    Start Time Stamp:  07-JAN-1993 12:00:00   Data Version: V1.1-38          VMS Version: V5.4-3
                                   End Time Stamp:  07-JAN-1993 12:02:00   Uptime in Seconds: 120

 Physical          Pri Vec  ------------------------Percentage of Time-------------------------
    ID       Run   Cpu Prc   Kernel   Exec   Super   User   Intrpt   Compat   Idle   Synch  Vector
----------   ---   --- ---   ------   ----   -----   ----   ------   ------   ----   -----  ------
         1     1     1   0      5.5    0.3     0.0    3.1      5.2      0.0   85.9     0.0     0.0

User Command: ADVISE COLL REPO DUMP_CPUS/BEG=7-JAN-1993 12:00:00.00/END=7-JAN-1993 12:02:00.00/OUT=DUMP_CPUS/NODE=LAIOUR
```

Item	Description
Physical ID	Physical CPU ID Number for Multiprocessor Systems
Run	Equals 1 if active, 0 if inactive
Pri Cpu	Equals 1 if primary CPU
Vec Prc	Equals 1 if vector processor

Item	Description
Percentage of Time Kernel	Percentage of time spent in kernel mode
Percentage of Time Exec	Percentage of time spent in executive mode
Percentage of Time Super	Percentage of time spent in supervisor mode
Percentage of Time User	Percentage of time spent in user mode
Percentage of Time Intrpt	Percentage of time spent on the interrupt stack
Percentage of Time Compat	Percentage of time spent in compatibility mode
Percentage of Time Idle	Percentage of time idle
Percentage of Time Synch	Percentage of time spent waiting for a resource protected by a spin lock
Percentage of Time Vector	Percentage of time in vector array processing

Dump Report: Hot Files Record

The following example shows a sample hot files record dump report. Hot file records are not available in history data. The following table describes the headings in the hot files record section of the dump report.

```
**Hot Files Record              Start Time Stamp:  25-AUG-2006 09:00:00   Data Version: V2.2-55          VMS Version: V8.2
                                  End Time Stamp:  25-AUG-2006 09:02:00   Uptime in Seconds: 120

                                                      User's  Srvc  Bytes I/Os Reads Split Pagop Swpop MscpOp Directory specification
   Device                File Name                      PID    /Opr  /Sec  Sec  /Sec  /Sec  /Sec  /Sec   /sec
-------------- ------------------------------------ -------- ----- ----- ----- ----- ----- ----- ----- ------ ------------------------
$3$DKA0        LIC98.DAT;3                          21404B81  0.00  0.00 0.00  0.00  0.00  0.00  0.00   0.00 [SYS0.CA_LIC]
$3$DKA0        ACCOUNTNG.DAT;1                      21401013  4.00  4.27 0.01  0.00  0.00  0.00  0.00   0.00 [SYS0.SYSMGR]
$1$DKC402      SCHEDULER$SHELL.COM;40               2140502C  0.00  0.00 0.00  0.00  0.00  0.00  0.00   0.00 [NSCHED.COM]
$1$DKC402      VERMONT_CREAMERY.LOG;1               21403039  1.90 13619 13.3 13.29  0.01  0.00  0.00   0.00 [NSCHED.DATA]
$1$DKC402      VSS.DAT;10                           21403039  4.10  1899 0.75  0.70  0.00  0.00  0.00   0.00 [NSCHED.DATA]

User Command: ADVISE COLLECT REPORT DUMP_HOTFILES/BEGIN=25-AUG-2006 09:00:00.00/END=25-AUG-2006 09:02:00.00/OUT=DUMP_HOTFILES.LOG
```

Item	Description
Device	Name of the disk device
Reads/Sec	Read operations per second
File Name	File specification of hot file, which may be the FID if OpenVMS could not translate the FID
Split/Sec	Split operations per second
User's PID	File user's process ID
Turns/Sec	Window turns per second
Pagop	Sec operations per second

Item	Description
Srvc/Opr	Service time per operation in milliseconds
Swpop/Sec	Swapping operations per second
Bytes/Sec	Bytes per second
MscpOp/Sec	MSCP operations per second
I/Os Sec	I/O operations per second
Directory Specification	Directory specification for file

Dump Report: Data Cells Record

Each item is a numeric data cell in domain LOCAL. This dump report is not included when DUMP_ALL is specified because it is primarily useful when debugging user-written rules. For the definitions of data cells, see the *Performance Manager Administrator Guide*.

The following example shows a sample data cells record dump report.:

```
Dump Report                          YQUEM (HP rx4640  (1.50GHz/6.0M)                                   Page   1
                                                                                                 PSDC V3.1-080
                                      Thursday 24-JAN-2008 10:00 to 10:02

NUMBER_OF_PROCESSES                     57.00
NUMBER_OF_USERS                          6.00
INTERACTIVE_COUNT                       10.00
BATCH_COUNT                              0.00
NETWORK_COUNT                            2.00
NUMBER_OF_INSWAPPED_PROCESSES           57.00
NUMBER_OF_OUTSWAPPED_PROCESSES           0.00
FREE_BALANCE_SET_SLOTS                2941.00
PROCESSES_IN_COLPG                       0.00
PROCESSES_IN_MWAIT                       0.24
PROCESSES_IN_CEF                         0.00
PROCESSES_IN_PFW                         0.00
PROCESSES_IN_LEF                        24.77
PROCESSES_IN_LEFO                        0.00
PROCESSES_IN_HIB                        29.99
PROCESSES_IN_HIBO                        0.00
PROCESSES_IN_SUSP                        0.00
PROCESSES_IN_SUSPO                       0.00
PROCESSES_IN_FPG                         0.00
PROCESSES_IN_COM                         0.00
PROCESSES_IN_COMO                        0.00
PROCESSES_IN_CUR                         2.00
VBS_INTSTK                               0.00
INTERRUPT_STACK                          0.59
ACTIVE_PROCESSORS                        4.00
PRIMARY_INTERRUPT_STACK                  0.04
PRIMARY_IDLE                            81.74
NONPRIMARY_IDLE                         71.83
KERNEL                                   0.17
EXEC                                     0.02
SUPER                                    0.00
USER                                    24.91
COMPAT                                   0.00
IDLE                                    74.31
```

Sampling Techniques of Dump Report Items

Data collection is timer-driven as well as event-driven. Most data items are collected exclusively at the end of each major sampling interval. However, a Performance Agent collects a subset of items for the metrics record and process record at the end of each minor sampling interval.

A Performance Agent collects the items from the process record at the end of each major interval (as specified by the interval parameter in the collection definition) and whenever any image terminates. Several data items in the process records, collected during minor interval sampling (every five seconds), are moving averages. These averages represent the data for the duration of the major interval or interval portion, during which a specific image was in effect. The Performance Agent resets all moving averages at each image termination.

Chapter 6: Analyze the System Program Counter

This chapter provides instructions for performing the following tasks using CA Performance Agent System Program Counter (PC) Analysis:

- Collecting system-wide PC data using the ADVISE COLLECT SYSTEM_PC command
- Reporting on the collected data using the ADVISE COLLECT REPORT SYSTEM_PC command
- Using System PC data to investigate how system CPU time is spent
- Using System PC data to isolate the cause of High Interrupt Stack activity

System-wide PC (program counter) statistics can be used as a detailed CPU measurement tool. Whenever a performance investigation indicates that CPU usage is high or that the CPU may be a bottleneck, system-wide PC statistics can pinpoint exactly where (PC) and how [IPL (interrupt priority level) or access mode] the CPU is being used.

System-wide PC statistics can be used to answer the following questions:

- Why is kernel mode time so high?
  - Is the time attributable to any particular process?
  - If so, is the time being spent calling system services or is a lot of code executed in kernel mode?
  - What amount of kernel time is accounted for in driver code or FDT routines? (FDT routines perform the device-dependent processing of an I/O request.)
  - What amount of CPU time is spent in each driver and in RMS routines?
- A real-time system is unable to meet its processing requirements but CPU utilization is low. The following questions may come to mind:
  - Are real-time interrupts being blocked by code executing at elevated IPL?
  - Are time-critical parts of user-mode code being stretched by elevated IPL code executing in response to device interrupts, fork processing, I/O post-processing, or AST delivery?

If you encounter any of the following symptoms, you may want to use the PC Analysis facility to further investigate CPU usage:

- Lost CPU time indicated by a value greater than 10% for the Pg+Swp Wait field
- High Interrupt Stack activity indicated by a value of the Inter Stack field in the tabular report greater than 20%
- Excessive CPU usage indicated by a value of less than 15% Total Idle Time on the tabular report

The following sections tell you how to collect and report system-wide PC data, and provide examples of evaluating System PC reports to investigate the symptoms. For detailed descriptions of the PC Analysis commands and reports, see Description of the Program Counter Analysis Report (see page 93) in this chapter.

This section contains the following topics:

Use the System PC Analysis Facility

An overview of the system-wide PC data collection and reporting procedure is shown in the following figure. The first step is to collect System PC data which is written to a log file. The next step is to run a report of System PC data.

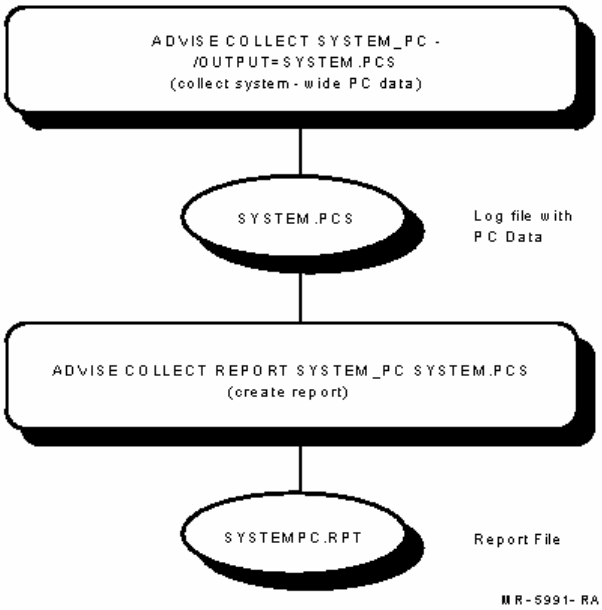

To begin collecting system-wide PC data, you must give a command of the following type:

```
$ ADVISE COLLECT SYSTEM_PC/qualifiers
```

To use this command, you need either the SETPRV account privilege or all the following privileges:

- ALTPRI
- CMKRNL
- PSWAPM
- WORLD

Run as a Separate Process

A PC data collection image runs as part of your process and proceeds to collect PC data as specified by the command qualifiers. Since data is collected by your own process, you are not free to give other commands at your terminal while PC data is being collected. For this reason, it may be preferable to use a batch command file or to spawn another process for this command.

Limit the Size of Data Files

The Performance Agent System PC Analysis facility collects a significant amount of data. Therefore, log files can get rather large. It is recommended that collections be limited to under 15 minutes for most investigations.

As another example, the following command starts collecting system-wide PC statistics for all active processes at the default collection interval of 1 tick (10 milliseconds) and writes data to the PC log file SYSTEM.PCS in directory [ACCOUNT] on disk device DISK$. The collection terminates after the default period of 15 minutes.

```
$ ADVISE COLLECT SYSTEM_PC/OUTPUT=DISK$:[ACCOUNT]SYSTEM.PCS
```

There can be only one process on your system at any given time collecting system-wide PC data. If such a process already exists when you give the above command, no new PC data collection image is started, and you receive the message:

```
%PSDC-F-PCSAMPACT, PC sampling is currently active
```

You can stop your system-wide PC collection process at any time by entering CTRL+Y, or with the DCL STOP command issued from another process.

```
$ ADVISE COLLECT SYSTEM_PC/OUTPUT=PC.DAT Interrupt;
 CTRL+Y was typed here
$
```

When you type a CTRL+Y, the PC log file is closed. To resume the collection process, type the DCL CONTINUE command before entering a command that runs a different image. If the CONTINUE command is used, no new version of the log file is created.

You can inquire on the system-wide PC data collection process by typing CTRL+T at the terminal from which it is running, or by using the DCL command SHOW PROCESS from another terminal.

Collect PC Statistics for a Single Process

To collect PC statistics for a single process, use the /IDENTIFICATION qualifier with the ADVISE COLLECT SYSTEM_PC command and specify the PID (Process ID) for the selected process. You can determine the PID by using the DCL commands SHOW SYSTEM or SHOW USERS. This results in a much smaller PC log file and somewhat different reports.

In the following command PC statistics are collected for the process whose ID is 00000054:

```
$ ADVISE COLLECT SYSTEM_PC/IDENTIFICATION=54/OUTPUT=PC54LOG.PCS
```

Generate Reports from a System-Wide PC Log File

To generate reports from a system-wide PC log file, you must give a command of the following type:

```
$ ADVISE COLLECT REPORT SYSTEM_PC/qualifiers pclogfile
```

This command accepts one parameter, the name of the PC log file, and certain qualifiers as explained below. No special privileges are needed to use this command. The /OUTPUT qualifier can be used to give the file specification for the report; if this is omitted, the report is generated as SYSTEMPC.RPT in the user default directory.

The command below generates a report, SYSTEMPC.RPT, from the log file SYSTEM.PCS. The report contains statistics for a PC location, processor access mode, IPL level, interrupt stack activity, and processes active during the collection period.

```
$ ADVISE COLLECT REPORT SYSTEM_PC SYSTEM.PCS
```

The following command causes the same report to be generated as PC.RPT:

```
$ ADVISE COLLECT REPORT SYSTEM_PC/OUTPUT=PC.RPT SYSTEM.PCS
```

Investigate System CPU Time

If there is excessive system CPU utilization on your system, use the Real-time displays to monitor system performance. For more information, see the *Performance Manager Administrator Guide*.

When a period of excessive system time occurs, run System PC collections.

Generate and examine the report to determine who is using the CPU, at what access modes, and for what reason. The following figure is an example of a System PC Analysis report segment. The following report shows processor usage by process, the number of PC samples, and the total CPU and system time for each process:

```
System PC Sampling
                  Page   8

Processor Utilization
Performance Agent Vx.x

Tuesday 09-APR-2008 09:56:26 to 10:00:12

Processor Utilization by Process

Process Name     Samples     %     Total Time   System Time     System    Processor Utilization by Access Mode
                                  (Seconds)     (Seconds)         Time %   KRNL      EXEC    SUPR    USER    CMPT    IPL>0    IPL>2

--------------------    -------    -------   -----------    ------------           ------     -----        ----       ----       ---
-        ----        -----     -----

Interrupt Stack            0365    11.37%   103.65            103.65               100%   100%       0%       0%       0%       0%
100%   100%

NULL CPU TIME      9397   43.22%     393.97           393.97               100%   100%       0%       0%       0%       0%
100%   100%

SWAPPER                127        0.14%         1.27           0.91                72%   100%       0%       0%       0%       0%
100%    98%

BIEGEL                       1    0.00%         0.01           0.01                100%   100%       0%       0%       0%
0%      100%     0%

DECW$SESSION           3    0.00%         0.03           0.03                100%   100%       0%       0%       0%       0%
100%   100%
```

```
CONFIGURE                42    0.05%         0.42           0.24                57%   100%       0%       0%       0%       0%
100%    26%

DECW$MWM                  1    0.00%         0.01           0.01               100%   100%       0%       0%       0%       0%
100%   100%

ERRFMT                   22    0.02%         0.22           0.21                95%   100%       0%       0%       0%       0%
100%   100%

CACHE_SERVER              4    0.00%         0.04           0.04               100%   100%       0%       0%       0%       0%
100%   100%

CLUSTER_SERVER           28    0.03%         0.28           0.28               100%   100%       0%       0%       0%       0%        64%
57%

OPCOM                    14    0.02%         0.14           0.13                93%   100%       0%       0%       0%       0%
71%     50%

AUDIT_SERVER             60    0.07%         0.60           0.59                98%   100%       0%       0%       0%       0%
93%     92%

JOB_CONTROL             174    0.19%         1.74           1.53                88%    72%      18%       0%      10%       0%
61%     53%

SHADOW_SERVER            2    0.00%         0.02           0.02               100%   100%       0%       0%       0%       0%      100%
100%

VUE$GUEST_2              0    0.00%         0.00           0.00                 0%       0%       0%       0%       0%       0%
0%       0%

DECW$TE_D812             4    0.00%         0.04           0.04               100%   100%       0%       0%       0%       0%
75%     75%

HPSQ23                   4    0.00%         0.04           0.04               100%   100%       0%       0%       0%
0%        75%    75%

Linda_3                  1    0.00%         0.01           0.01               100%   100%       0%       0%       0%
0%      100%   100%

NETACP                 447    0.49%         4.47           3.47                78%   100%       0%       0%       0%
0%        78%      73%

Torrey <1>               1    0.00%         0.01           0.01               100%   100%       0%       0%       0%
0%      100%       0%

GUEST                    1    0.00%         0.01           0.01               100%   100%       0%       0%       0%
0%      100%   100%

MAIL_33853               1    0.00%         0.01           0.01               100%   100%       0%       0%       0%
0%      100%   100%

SCAGT MOM                1    0.00%         0.01           0.01               100%   100%       0%       0%       0%       0%
100%       0%

BOUSQUET_1               0    0.00%         0.00           0.00                 0%       0%       0%       0%       0%
0%          0%       0%

PID=2060AAD6           754    0.83%         7.54           5.85                78%    73%       9%         1%     17%       0%
67%     56%
```

Description of the Program Counter Analysis Report

This section describes the parts of the report.

General Information

General Information is always part of any report generated with ADVISE COLLECT REPORT SYSTEM_PC. The following example shows an example of this output, which is the first page of the system-wide PC data report:

```
System PC Sampling
                                                                Page    1

Experiment Information
                                    Performance Agent VX.X

Tuesday 09-APR-2008 09:56:26 to 10:00:12

Experiment name:

Report date:                          Tuesday 09-APR-2008 10:12:11

Data collection start time:   09-APR-2008 09:56:26.66

Data collection stop time:    09-APR-2008 10:00:12.45

Total elapsed time:              0 00:03:45.79

Sampling interval:               200 ms
```

Header Information

Data Collection Start Time

The time when the PC sampling process began.

Data Collection Stop Time

The time when the PC sampling process ended.

Total Elapsed Time

The clock time for which the PC sampling process was active.

Sample Interval

The interval in milliseconds at which the PC and PSL were sampled.

Processes at Start Time

Process Name

The name of the process.

An entry of the form "PID = nnnnnnnn" where nnnnnnnn is the process identification in hexadecimal may also appear here in cases where a process is created while data sampling is in progress. The collection image writes process PIDs and their associated names to the log file before starting data collection. Therefore, a process created while data collection is in progress does not have an associated name in the log file, so the PID is displayed instead.

PID

The process identification in hexadecimal.

Drivers at Start Time

Driver Name

The name of each device driver or other code loaded into the OpenVMS executive on the system.

Start

The starting system virtual address (in hexadecimal)for the driver code.

End

The ending system virtual address (in hexadecimal)for the driver code.

Executive Layout

This section lists the start and end addresses of OpenVMS images, as well as other loadable code, such as CLUSTELOA, SYSLOA, SCSLOA, and DCL. The following is an example of the Executive Layout section:

VMS Executive Layout

Image	Start	End
BASE_IMAGE	80002000	80004000
SYSLDR_DYN	802EB200	802F0200
DCLP$RMS_EXTENSION	802F0800	802F1A00
RECOVERY_UNIT_SERVICES	802F1C00	802F2400
RMS	80182C00	801AAB00
SYS$NETWORK_SERVICES	8020F000	8020F200
SYS$TRANSACTION_SERVICES	8020F800	8022E800
CPULOA	8022F800	80232400
LMF$GROUP_TABLE	80233200	80234600
SYSLICENSE	80234A00	80236200
.	.	.
.	.	.
.	.	.

Drivers, CRBs, Loadable Code

Image	Start	End
PSDCTIMER	8049CA10	8049DDA0
DFSCDRIVER	80476810	8047EDC0
DFSRDRIVER	8046F5F0	80476810
RTTDRIVER	8046D840	8046E3E0
.	.	.
.	.	.
.	.	.
DMB CRB	80881300	80881378
CPA CRB	80846EC0	80846F70
MPA CRB	802EB884	802EB90C
NLA CRB	802EB884	802EB90C
.	.	.
.	.	.
.	.	.
SYSLOA	804FD650	80500E10
PEEMUL	804F6DD0	804F93CE
VAXEMUL	804F93D0	804FCA4F
Pool	80420600	8049C2000
P1 Vectors	7FFEDE00	7FFEFE00

Drivers, CRBs, Loadable Code

Image	Start	End
CLI (DCL)	7FF75C00	7FF8D9EF
F11BXQP	7FF9B200	7FFD5600

Processor Usage by Process

Unless CLASS=(NODEF,PROCESS) is specified with the ADVISE COLLECT REPORT SYSTEM_PC command, this statistics page is always present. Use of the /IDENTIFICATION qualifier does not affect the content of this part of the report.

The example in this section shows an example of this output which is the second section of the system-wide PC data report. Usage information for each active process is given as follows:

Process Name

The name of the process. An entry of the form "PID = nnnnnnnn" where nnnnnnnn is the Process Identification Number may also appear here in cases where a process is created while data sampling is in progress. This is because the collection image writes process PIDs and their associated names to the log file before starting data collection; a process created while data collection is in progress does not have an associated name in the log file, so the PID is displayed instead.

Samples

The number of PC samples accrued to the process.

%-The percentage of all PC samples accrued to the process.

Total Time (seconds)

The total CPU time (in seconds) accrued to the process, which is equal to (Number of Samples x Sample Interval).

System Time (seconds)

The CPU time in seconds accrued to the process while the PC was in system virtual address space.

System Time %

The percentage of process CPU time spent in system virtual address space.

Process by Access Mode

The percentage of process CPU time spent in each access mode (kernel, executive, supervisor, user, compatibility mode) and the percentage of process CPU time spent at IPL > 0 and IPL > 2.

```
System PC Sampling
                                        Page   8

Processor Utilization
               Performance Agent VX.X

Tuesday 09-APR-2008 09:56:26 to 10:00:12

Processor Utilization by Process
```

```
Process Name     Samples     %     Total Time System Time     System    Processor Utilization by Access Mode
                                                (Seconds)  (Seconds)       Time %  KRNL  EXEC  SUPR  USER  CMPT  IPL>0  IPL>2

---------------          -------     -------     ---------      ---------           ------     ----      ----      ----      ----
----     -----    -----

Interrupt Stack       10365   11.37%     103.65     103.65            100%   100%    0%     0%      0%     0%   100%    100%

NULL CPU TIME      39397   43.22%     393.97     393.97            100%   100%    0%     0%      0%     0%   100%    100%

SWAPPER                 127     0.14%         1.27         0.91             72%   100%    0%     0%      0%     0%   100%     98%

BIEGEL                        1     0.00%         0.01         0.01           100%   100%    0%     0%      0%     0%   100%
0%

DECW$SESSION            3     0.00%         0.03         0.03            100%   100%    0%     0%      0%     0%   100%    100%

CONFIGURE               42      0.05%         0.42         0.24             57%   100%    0%     0%      0%     0%   100%      26%

DECW$MWM                1      0.00%         0.01         0.01           100%   100%    0%     0%      0%     0%   100%    100%

ERRFMT                      22       0.02%        0.22         0.21             95%   100%    0%     0%      0%     0%   100%
100%

CACHE_SERVER           4       0.00%        0.04         0.04           100%   100%    0%     0%      0%     0%   100%    100%

CLUSTER_SERVER      28       0.03%       0.28         0.28           100%   100%     0%     0%      0%     0%    64%       57%

OPCOM                       14       0.02%       0.14         0.13             93%   100%    0%      0%      0%     0%    71%
50%

AUDIT_SERVER            60       0.07%       0.60         0.59             98%   100%    0%     0%      0%     0%    93%      92%

JOB_CONTROL           174       0.19%       1.74         1.53             88%     72%  18%     0%     10%    0%    61%      53%

SHADOW_SERVER         2       0.00%       0.02         0.02           100%   100%    0%     0%      0%     0%   100%   100%

VUE$GUEST_2               0       0.00%       0.00         0.00              0%       0%    0%     0%      0%     0%       0%
0%

DECW$TE_D812            4       0.00%       0.04         0.04           100%   100%    0%     0%      0%     0%    75%      75%

HPSQ23                      4       0.00%       0.04         0.04           100%   100%    0%     0%      0%     0%    75%
75%

Linda_3                       1       0.00%       0.01         0.01           100%   100%    0%     0%      0%     0%   100%
100%

NETACP                    447       0.49%       4.47         3.47             78%   100%    0%     0%      0%    0%       78%
73%

Torrey <1>                    1       0.00%       0.01         0.01           100%   100%    0%     0%      0%    0%    100%
0%

GUEST                         1       0.00%       0.01         0.01           100%   100%    0%    0%      0%    0%     100%
100%

MAIL_33853                  1       0.00%       0.01         0.01           100%   100%    0%    0%       0%    0%     100%   100%

SCAGT MOM                 1       0.00%       0.01         0.01           100%   100%    0%    0%       0%    0%      100%     0%
```

BOUSQUET_1	0	0.00%	0.00	0.00	0%	0%	0%	0%	0%	0%	0%	0%
PID=2060AAD6	754	0.83%	7.54	5.85	78%	73%	9%	1%	17%	0%	67%	56%

Processor Utilization by IPL

Unless /CLASS=(NODEF,PROCESS) is specified with the ADVISE COLLECT REPORT SYSTEM_PC command, this graph is always present. Use of the /IDENTIFICATION qualifier does not affect the content of this part of the report.

The example labeled *CPU Utilization by IPL* Report *Example* shows an example of this output which is the third section of the system-wide PC data report. This section of the report presents a breakdown of CPU time spent at each Interrupt Priority Level (IPL) while the PC was in system virtual address space. Percentages are shown in both numerical and graphical format.

Processor Usage by IPL for Interrupt Stack

Unless /CLASS(NODEF,PROCESS) and/or /IDENTIFICATION is specified with the ADVISE COLLECT REPORT SYSTEM_PC command, this statistics page is always present. Use of the /IDENTIFICATION qualifier produces a similar report page that pertains to a specific process rather than the Interrupt Stack (see the next section).

The example labeled "CPU Utilization by IPL for Interrupt Stack" shows an example of this output, which is the fourth section of the system-wide PC data report. This section of the report presents a breakdown of percentage CPU time spent at each Interrupt Priority Level (IPL) while on the Interrupt Stack. There is no Process Context while on the Interrupt Stack. Percentages are shown in both numerical and graphical format.

Processor Usage by IPL for a Specific Process

This section of the report is obtained by using the /IDENTIFICATION qualifier to indicate a specific process as the filter process. This produces a report page similar to that for Interrupt Stack (see the previous section), but for a process (Process Context).

The example labeled “CPU Utilization by IPL for a Process Report Example” shows an example of this output, which is the fourth section of the system-wide PC report. This section of the report presents a breakdown of percentage CPU time spent at each Interrupt Priority Level (IPL) while the PC was in system virtual address space for a specific process. Percentages are shown in both numerical and graphical format.

CPU Utilization by IPL Report

```
                                    CPU utilization by IPL (Percent)
                                       (All active CPUs combined)
                                          While PC in S0 Space

IPL       %     0         10        20        30        40        50        60        70        80        90       100
---      ---    +---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+
 0        2%    |**                                                                                                   |
 1        0%    |                                                                                                     |
 2        0%    |                                                                                                     |
 3       90%    |******************************************************************************************           |
 4        0%    |                                                                                                     |
 5        0%    |                                                                                                     |
 6        0%    |                                                                                                     |
 7        0%    |                                                                                                     |
 8        2%    |**                                                                                                   |
 9        0%    |                                                                                                     |
10        0%    |                                                                                                     |
11        0%    |                                                                                                     |
12        0%    |                                                                                                     |
13        0%    |                                                                                                     |
14        0%    |                                                                                                     |
15        0%    |                                                                                                     |
16        0%    |                                                                                                     |
17        0%    |                                                                                                     |
18        0%    |                                                                                                     |
19        0%    |                                                                                                     |
20        4%    |****                                                                                                 |
21        0%    |                                                                                                     |
22        0%    |                                                                                                     |
23        0%    |                                                                                                     |
24        0%    |                                                                                                     |
25        0%    |                                                                                                     |
26        0%    |                                                                                                     |
27        0%    |                                                                                                     |
28        0%    |                                                                                                     |
29        0%    |                                                                                                     |
30        0%    |                                                                                                     |
31        0%    |                                                                                                     |
                +---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+
                0         10        20        30        40        50        60        70        80        90       100
```

CPU Utilization by IPL for Interrupt Stack

```
System PC Sampling                          Hanley 200 intervals                    Page   8
Interrupt Stack IPL Percentages                                                     RSDC   V2.1-48
                                  Thursday 11-JUN-1992 10:04:15 to 10:04:15

                                CPU utilization by IPL (Percent)
                                (All active CPUs combined)
                                Pbc Interrupt Stack

IPL     %     0         10        20        30        40        50        60        70        80        90        100
---     ---   +---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+
  0     0%    |                                                                                                   |
  1     0%    |                                                                                                   |
  2     0%    |                                                                                                   |
  3     0%    |                                                                                                   |
  4     0%    |                                                                                                   |
  5     0%    |                                                                                                   |
  6     0%    |                                                                                                   |
  7     0%    |                                                                                                   |
  8     25%   |************************                                                                           |
  9     0%    |                                                                                                   |
 10     0%    |                                                                                                   |
 11     0%    |                                                                                                   |
 12     0%    |                                                                                                   |
 13     0%    |                                                                                                   |
 14     0%    |                                                                                                   |
 15     0%    |                                                                                                   |
 16     0%    |                                                                                                   |
 17     0%    |                                                                                                   |
 18     0%    |                                                                                                   |
 19     0%    |                                                                                                   |
 20     75%   |***************************************************************************                        |
 21     0%    |                                                                                                   |
 22     0%    |                                                                                                   |
 23     0%    |                                                                                                   |
 24     0%    |                                                                                                   |
 25     0%    |                                                                                                   |
 26     0%    |                                                                                                   |
 27     0%    |                                                                                                   |
 28     0%    |                                                                                                   |
 29     0%    |                                                                                                   |
 30     0%    |                                                                                                   |
 31     0%    |                                                                                                   |
             +---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+
             0         10        20        30        40        50        60        70        80        90        100
```

CPU Utilization by IPL for a Process Report

```
System PC Sampling                                                                                   Page  16
Interrupt Stack IPL Percentages                                                                 PSDC    V2.1-48
                                        Tuesday 09-MAR-1993 09:56:26 to 10:00:12


                                         CPU utilization by IPL (Percent)
                                           (All active CPUs combined)
                                                Proc LUNDGREN

    IPL       %     0        10        20        30        40        50        60        70        80        90       100
    ---      ---    +---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+
      0      44%    |********************************************                                                        |
      1       0%    |                                                                                                    |
      2       5%    |*****                                                                                               |
      3       0%    |                                                                                                    |
      4       0%    |                                                                                                    |
      5       0%    |                                                                                                    |
      6       0%    |                                                                                                    |
      7       0%    |                                                                                                    |
      8      45%    |*********************************************                                                       |
      9       0%    |                                                                                                    |
     10       0%    |                                                                                                    |
     11       3%    |***                                                                                                 |
     12       0%    |                                                                                                    |
     13       0%    |                                                                                                    |
     14       0%    |                                                                                                    |
     15       1%    |*                                                                                                   |
     16       0%    |                                                                                                    |
     17       0%    |                                                                                                    |
     18       0%    |                                                                                                    |
     19       0%    |                                                                                                    |
     20       0%    |                                                                                                    |
     21       0%    |                                                                                                    |
     22       0%    |                                                                                                    |
     23       0%    |                                                                                                    |
     24       0%    |                                                                                                    |
     25       0%    |                                                                                                    |
     26       0%    |                                                                                                    |
     27       0%    |                                                                                                    |
     28       0%    |                                                                                                    |
     29       0%    |                                                                                                    |
     30       0%    |                                                                                                    |
     31       0%    |                                                                                                    |
                    +---------+---------+---------+---------+---------+---------+---------+---------+---------+---------+
                    0        10        20        30        40        50        60        70        80        90       100
```

PC Samples by System Image

Use of the /IDENTIFICATION qualifier produces a report page in which statistics under the heading “Filter” pertain to a specific process, rather than the Interrupt Stack. If /IDENTIFICATION is omitted, Filter statistics are for the interrupt stack (no process context).

The example labeled “PC Samples by System Image Report Example” is an example of a PC Sample by System Image report with filter statistics for the interrupt stack.

System Module

This is the name of the system module. The entry "Unknown System Address" can also appear here, meaning that one or more PC samples were found in system virtual address space that did not fall into the range of known address ranges. In addition, the report image reads the system virtual address ranges for the RMS code and for all device drivers and adapter handlers loaded into non-paged pool at the time the collection process was initiated.

Module Samples

The number of PC samples collected in the system virtual address range of the specified system module.

% Total Module Samples

The percentage of PC samples collected in the system virtual address range of the specified system module.

Filter: Process Name

If the /IDENTIFICATION qualifier is used when generating the report, Process Name is that of a specified process. Otherwise, INTERRUPT STACK appears in this heading.

Filter Samples

The number of PC samples collected in the system virtual address range of the specified system module while the process "process name" was the currently executing process. If /IDENTIFICATION was not used when generating the report, the pseudo-process name "INTERRUPT STACK" is listed, and the statistics indicate the number of samples collected while in the system module on the interrupt stack.

% of Module Samples

The percentage of PC samples collected in the system virtual address range of the specified system module while the process "process name" was current. This is calculated as follows:

```
(Number of filter samples for module)
------------------------------------ x 100
(Number of samples for module)
```

If /IDENTIFICATION was not used when generating the report, the pseudo-process name "INTERRUPT STACK" is listed, and the statistics indicate the percentage of samples collected while in the system module on the interrupt stack.

% Total Filter Samples

The percentage of PC samples collected in the system virtual address range of the specified system module while the process "process name" was current. This is calculated as follows:

```
(Number of filter samples for module)
------------------------------------ x 100
(Total number of filter samples)
```

If /IDENTIFICATION was not used when generating the report, the pseudo-process name "INTERRUPT STACK" is listed, and the statistics indicate the percentage of samples collected while in the system module on the interrupt stack.

PC Samples by System Image Report

PC Samples by System Image

			'--- Filter: Interrupt Stack ---'		
System Image	Module Samples	% Total Module Samples	Filter Samples	% Of Module Samples	%Total Filter Samples
MESSAGE_ROUTINES	2	1.43%	0	0.00%	0.00%
EXCEPTION	1	0.71%	0	0.00%	0.00%
PAGE_MANAGEMENT	1	0.71%	0	0.00%	0.00%
EVENT_FLAGS_AND_ASTS	1	0.71%	0	0.00%	0.00%
PROCESS_MANAGEMENT	126	90.00%	0	0.00%	0.00%
SYSTEM_PRIMITIVES	1	0.71%	1	100.00%	12.50%
GABDRIVER	1	0.71%	1	100.00%	12.50%
NETDRIVER	1	0.71%	0	0.00%	0.00%
[illegible]DRIVER	1	0.71%	1	100.00%	12.50%
ESDRIVER	2	1.43%	2	100.00%	25.00%
PRIDRIVER	3	2.14%	3	100.00%	37.50%

PC Samples by System Module

Use of the /IDENTIFICATION qualifier produces a report page in which statistics under the heading "Filter" pertain to a specific process, rather than the Interrupt Stack. If /IDENTIFICATION is omitted, Filter Statistics are for the Interrupt Stack (no Process Context).

The example labeled "PC Samples by System Module Report Example" shows an example of this output. Note that some system modules are omitted, so the modules present do not add up to the totals shown. This section of the report presents a breakdown of CPU time spent in each system module. If the /IDENTIFICATION qualifier was used when generating the report, the Filter statistics are for a specific process rather than INTERRUPT STACK.

System Module

This is the name of the system module. The entry "Unknown System Address" can also appear here, meaning that one or more PC samples were found in system virtual address space that did not fall into the range of known address ranges. Universal entry points into the OpenVMS EXECUTIVE are collected from loadable system images. In addition, the report image reads the system virtual address ranges for the RMS code, MSCP server code, and for all device drivers and adapter handlers loaded into non-paged pool at the time the collection process was initiated.

Module Samples

The number of PC samples collected in the system virtual address range of the specified system module.

% Total Module Samples

The percentage of PC samples collected in the system virtual address range of the specified system module.

Filter: Process Name

If the /IDENTIFICATION qualifier is used when generating the report, Process Name is that of a specified process. Otherwise, INTERRUPT STACK appears in this heading.

Filter Samples

The number of PC samples collected in the system virtual address range of the specified system module while the process "process name" was the currently executing process. If /IDENTIFICATION was not used when generating the report, the pseudo-process name "INTERRUPT STACK" is listed, and the statistics indicate the number of samples collected while in the system module on the interrupt stack.

% of Module Samples

The percentage of PC samples collected in the system virtual address range of the specified system module while the process "process name" was current. This is calculated as follows:

```
(Number of filter samples for module)
 ------------------------------------ x 100
 (Number of samples for module)
```

If /IDENTIFICATION was not used when generating the report, the pseudo-process name "INTERRUPT STACK" is listed, and the statistics indicate the percentage of samples collected while in the system module on the interrupt stack.

% Total Filter Samples

The percentage of PC samples collected in the system virtual address range of the specified system module while the process "process name" was current. This is calculated as follows:

```
(Number of filter samples for module)
 ------------------------------------ x 100
 (Total number of filter samples)
```

If /IDENTIFICATION was not used when generating the report, the pseudo-process name "INTERRUPT STACK" is listed, and the statistics indicate the percentage of samples collected while in the system module on the interrupt stack.

PC Samples by System Module Report Example

PC Samples by System Module

EXEC Image	System Module	Module Samples	% Total Module Samples	!--- Filter: Interrupt Stack ---! Filter Samples	% Of Module Samples	%Total Filter Samples
MESSAGE_ROUTINES						
	EXE$EPACL	2	1.43%	0	0.00%	0.00%
EXCEPTION						
	EXE$EXCEPTION	1	0.71%	0	0.00%	0.00%
PAGE_MANAGEMENT						
	MMG$LOCK_SYSTEM_PAG	1	0.71%	0	0.00%	0.00%
EVENT_FLAGS_AND_ASTS						
	SCH$IEWLVL	1	0.71%	0	0.00%	0.00%
PROCESS_MANAGEMENT						
	SCH$PRESCHED	126	90.00%	0	0.00%	0.00%
SYSTEM_PRIMITIVES						
	EXE$ICRRDSPTH	1	0.71%	1	100.00%	12.50%

Run Statistics:

Total Samples: 179
Total Overflows: 0
Total Processes: 25
Input Source: PC Sampler Logfile V3.2

Command Line:

```
AIVIS COLLECT REPORT SYSTEM_PC HAWLEY20.PCS
```

Image Performance Statistics:

```
ELAPSED:    0 00:00:08.10  CPU: 0:00:05.92  BUFIO: 14  DIRIO: 7  FAULTS: 1768
```

Chapter 7: Display File Activity

This chapter provides the following information about the CA Performance Agent File Activity display:

- Introduction to the File Activity display
- Instructions for invoking the File Activity display
- Description of File Activity display interactive commands

When an I/O bottleneck is due to high activity on a disk, the information from the File Activity display can help you to improve system performance. The File Activity display identifies the files on a disk that are most frequently accessed. In addition, it displays the percentage of read operations, I/O rate, and usage for these files. Depending on the workload, these files may be moved to less busy disks to reduce the I/O on the excessively active disk. Active files may also be marked for allocation of RMS GLOBAL buffers.

This section contains the following topics:

Invoke the File Activity Display (see page 107)
File Activity Interactive Commands (see page 109)

Invoke the File Activity Display

The privileges required to use the ADVISE COLLECT DISPLAY FILES command are CMKRNL, TMPMBX, and SYSPRV.

The following command invokes the File Activity display:

```
$ ADVISE COLLECT DISPLAY FILES/qualifiers[,...]
```

There are two qualifiers used with the command as shown in the following table:

Command Qualifier	Defaults
/DISK=([disk,...]	/DISK=SYS$SYSDEVICE
/INTERVAL=[seconds]	/INTERVAL=10

Use the /DISK qualifier to specify the names of the disks whose activity you wish to monitor. Select disks that show chronically high I/O rates or response times.

Use the /INTERVAL qualifier to specify the interval at which file activity information is displayed. If no interval is specified, an interval of 10 seconds is assumed.

File activity data is collected at a rate computed as the display rate divided by ten. It is always greater than or equal to one second, and less than or equal to five seconds. Due to the small collection interval, collecting data for a number of disks over a long period may affect overall system performance.

The following command displays active files for disks DUA0 and SYS$DISK: at the default rate of ten seconds.

```
$ ADVISE COLLECT DISPLAY FILES/DISK=(DUA0,SYS$DISK)
```

The following example is an example of the File Activity display where files are listed in descending order according to their read/write operation rate:

```
                PSDC File Activity Display on Node MASHUA
Rate  Read%  Usage  Filename    10-MAR-1993 13:41:02    Active Files:    4
-------------------------------------------------------------------------
      11.76   100      1  SYS$SYSDEVICE:[SYSCOMMON.SYSLIB]CMSSHR.EXE;1
       4.90   100      1  MH$:[SYSEXE]PAGEFILE_MASHUA.SYS;1
       2.94   100      2  SYS$SYSDEVICE:[SYSCOMMON.SYSEXE]CMS.EXE;1
       1.96   100      1  SYS$SYSDEVICE:[SYSCOMMON.SYSLIB]SMGSHR.EXE;1
```

The following table describes fields in the File Activity display:

Field	Description
Rate	The number of read and write operations per second for the file
Read %	The percentage of Read I/Os attributed to file reads
Usage	The number of concurrent references to the file
Name	The full name of the active file, including its device (which may be a user-defined logical)
Active Files	The total number of active files on the specified disk(s) with a rate greater than 0

If a disk dismounts while the File Activity Display is running, a message displays and collection discontinues for that disk. If the disk is remounted, it must be respecified before its performance statistics are collected and displayed.

File Activity Interactive Commands

While displaying file activity, you can enter commands to control the display and its characteristics. The following table describes the File Activity display interactive subcommands:

Subcommand	Meaning
EXIT or CTRL+Z	Exits from the display
HELP	Displays HELP information for field descriptions and subcommands
SET DISK	Defines the disks for which statistics are displayed
SET INTERVAL	Defines the rate at which disk statistics are displayed
SHOW DISK	Shows the names of disks for which statistics are displayed
SHOW INTERVAL	Shows the value of the current display interval
SHOW ALL	Shows the value of the current display interval and the names of disks for which statistics are displayed
SPAWN	Spawns a subprocess
WRITE [file-spec]	Copies the screen contents to a file. If no file is specified, the default of PSDC$DISPLAY_FILES.RPT is used

Chapter 8: Disk Space Analysis

This chapter provides general descriptions and examples. For a description of the command syntax, see the chapter Performance Agent Commands (see page 119).

When an I/O bottleneck is due to disk fragmentation, the information from the Disk Space Analysis utility can help you to improve system performance. Use the Disk Space report to determine if there is disk fragmentation, the disks on which it occurs, and the files on the disk that are most affected. After you determine the disks that are fragmented, you can compress them using the OpenVMS BACKUP utility.

The Disk Space report can also help you determine if you are getting maximum utilization from your disks by showing any inconsistencies between how a volume is initialized and mounted and how it is used. For example, you can see if the file extension size is too small or too large in relation to the size of most files on the disk. If the extension size is inappropriate, adjusting it may optimize disk utilization.

This section contains the following topics:

Disk Space Analysis Utility

Use the ADVISE COLLECT REPORT DISK_SPACE command to obtain information about disk volumes. The following are prerequisites when using this utility:

- The volume must be in FILES-11 ODS-2 or ODS-5 format.
- The volume must be mounted without the /FOREIGN qualifier.
- You must have READ access to the files [0,0]BITMAP.SYS and [0,0]INDEXF.SYS on the volume to be analyzed. Generally, this is satisfied by the SYSPRV privilege.

The report command is of the following type:

```
$ ADVISE COLLECT REPORT DISK_SPACE/qualifier device-name,...
```

This command accepts one or more parameters and a qualifier (explained below). The parameters are the names of disk devices on which to report. Physical or logical device names may be specified.

The Performance Agent software contains a knowledge base including rules for evaluating disks. These rules may recommend generating this report for specific devices. You may also want to create this report for disks that are more than 70% full. Before giving the ADVISE COLLECT REPORT command, you can display a list of disk devices on your system using the DCL command SHOW DEVICES.

During the analysis, the volume is locked to prevent storage allocation. The files BITMAP.SYS and INDEXF.SYS are read and the volume is unlocked. The command does not write to the volume being analyzed unless the report file you specify resides on that volume.

By default, the disk information is written to the device defined by the logical SYS$OUTPUT. Using the /OUTPUT qualifier, you can specify the output report name.

The following command writes a disk volume report for disks COOLER and BURGUNDY to the file DISKSPACE.RPT in the current directory:

```
$ ADVISE COLLECT REPORT DISK_SPACE/OUTPUT=DISKSPACE.RPT-
_$ COOLER,BURGUNDY
```

The following command writes a disk volume information report for device HSC000$DUA0 to the file DUA0.RPT:

```
$ ADVISE COLLECT REPORT DISK_SPACE/OUTPUT=DUA0.RPT HSC000$DUA0
```

The Disk Space report contains four sections: Detailed Volume Analysis, Summary of Free Storage, Summary of Allocated Space, and Files with Extension Headers.

Disk Fragmentation Investigation

There are three ways to read the Disk Space report to evaluate disk fragmentation. One way is to compare the size of most of the current files with the size of most available space to determine if fragmentation occurs as new files are created on the disk. Another way is to look at the number of files on a disk and the size and number of extents per file to determine if files on the disk are fragmented.

When a file being written exceeds the space allocated to it, an additional group of contiguous blocks is allocated at another location on the disk. This is called an extent. The following figure shows a file with extents. The third way is to look at the number of files with extension headers to determine if the disk is fragmented.

File with Two Extents

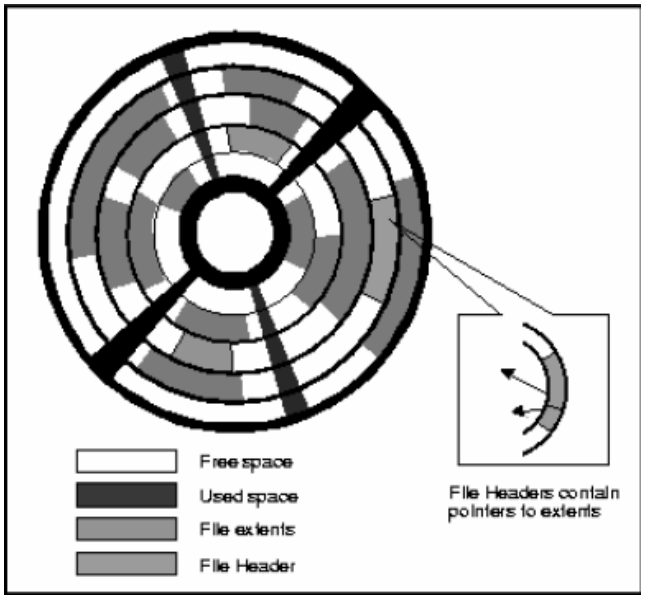

To determine if fragmentation will occur as new files are created, examine the Free Storage Utilization section of the Disk Space report. Look at the free storage extent size closest to 80%. Now examine the Allocated Space Utilization section of the Disk Space report. Look at the allocated extent size closest to 80%. If the size of most space allocated (files created) is smaller than most free storage size (space available), little or no fragmentation is likely to occur on that disk. If the size of most files created is larger than the free extent size, fragmentation may occur as new files are created.

In the following example, in the 82.5 percentile, free storage extent sizes range from 900 to 1500 blocks. In the example labeled "Allocated Space Utilization - Report Example", in the 84.6 percentile, allocated extent sizes range from 60 to 90 blocks. Therefore, the size of free storage is more than enough to accommodate the new files as they are created, and disk fragmentation is not likely to occur if future disk usage follows current usage.

```
Disk Analysis              _NUN$DUA32:  (COOLER)                  Page   1
Summary of Free Storage                          Performance Agent    VX.X
Monday 11-MAY-2008 14:50

Free Storage  Extent Sizes     No. Extents     Cum % Space
--------------------------     -----------     -----------

>=         3,  <         6             201           1.3
>=         6,  <         9              88           2.4
>=         9,  <        15              91           4.5
>=        15,  <        30             110           9.2
>=        30,  <        60              80          16.0
>=        60,  <        90              42          22.5
>=        90,  <       150              22          27.8
>=       150,  <       300              36          43.6
>=       300,  <       600              16          57.8
>=       600,  <       900              10          74.4
>=       900,  <      1500               3          82.5
>=      1500,  <      3000               2          91.0
>=      3000,  <      6000               1         100.0
>=      6000,  <      9000               0         100.0
>=      9000,  <     15000               0         100.0
>=     15000,  <     30000               0         100.0
>=     30000,  <     60000               0         100.0
>=     60000,  <     90000               0         100.0
>=     90000,  <    150000               0         100.0
>=    150000                             0         100.0

Total free blocks  =      46284.
No. of extents     =        702.
Mean blocks/extent =         66.
Smallest extent    =          3.
Largest extent     =       4158.
```

To evaluate file fragmentation, examine the Allocated Space section of the Disk Space report. Look at the number of extents allocated, the total number of files, and the mean number of extents per file. A value of 1 for mean number of extents per file indicates that most files have only one extent; therefore, they are not fragmented. A disk with a small number of files and a large number of extents per file indicates that the files are fragmented. In the following example, the mean number of extents per file is 2 and there are a large number (10052) of small files, indicating that these files may be fragmented.

In some cases, it may not be necessary to use the BACKUP utility on a disk. If a heavily used file has more than one extension header and appears to be the source of the device's bottleneck, copying the file to a contiguous area may be sufficient to improve performance.

```
Disk Analysis              _NUN$DUA32:  (COOLER)                  Page   3
Summary of Allocated Space                        Performance Agent    VX.X
Monday 11-MAY-2008 14:50

Space Allocated per Header     No. Headers     Cum % Headers
--------------------------     -----------     -------------
>=         3,  <         6           2979           29.7
>=         6,  <         9           1162           41.2
>=         9,  <        15           1116           52.3
>=        15,  <        30           1423           66.5
>=        30,  <        60           1296           79.4
>=        60,  <        90            522           84.6
>=        90,  <       150            742           92.0
>=       150,  <       300            448           96.4
>=       300,  <       600            202           98.5
>=       600,  <       900             34           98.8
>=       900,  <      1500             37           99.2
>=      1500,  <      3000             44           99.6
>=      3000,  <      6000             18           99.8
>=      6000,  <      9000             12           99.9
>=      9000,  <     15000              7          100.0
>=     15000,  <     30000              3          100.0
>=     30000,  <     60000              0          100.0
>=     60000,  <     90000              0          100.0
>=     90000,  <    150000              0          100.0
>=    150000                            0          100.0

Minimum allocated extent =          3.
Maximum allocated extent =      19191.
Total allocated blocks   =     844731 ( 94.8% of volume).
Total used      blocks   =     777058 ( 92.0% of allocated).
No. extents allocated    =      18333.
Mean alloc blocks/extent =         46.
Total no. of files       =      10052.
Mean alloc blocks/file   =         84.
Mean no. extents/file    =          2.
No. extension headers    =         26.
No. multi-volume files   =          0.
No. directories          =        435.
```

To further determine if a disk is fragmented, look at the number of extension headers in the Allocated Space section of the Disk Space report. When the number of extension headers is nonzero, there are one or more seriously fragmented files. A single header holds about 80 extent pointers; therefore, a file that needs an extension header is severely fragmented. In the preceding example, there are 26 extension headers, which indicates this disk is heavily fragmented.

Fragmented File with an Extension Header

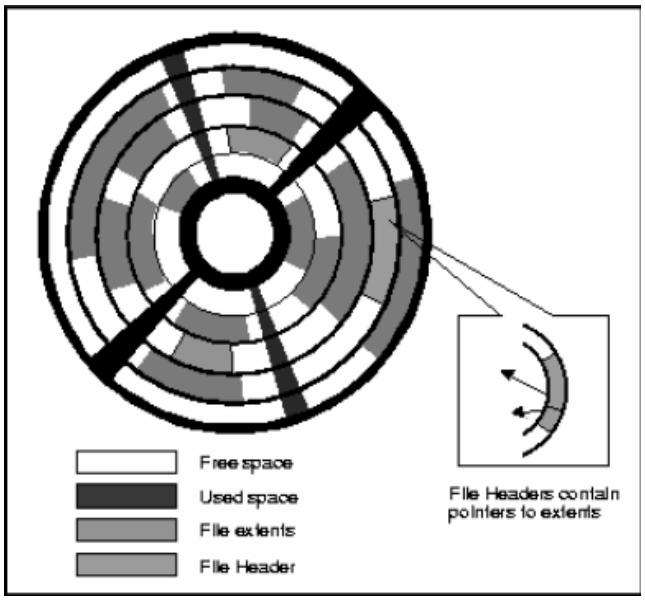

The Disk Space report lists any files with extension headers, as shown in the following example:

```
Disk Analysis                  _NUN$DUA32: (COOLER)                    Page   5
Files with Extension Headers                          Performance Agent    VX.X
Monday 11-MAY-2008 14:50

File name                                                          Ext. headers
-----------------------------------------------------------------  ------------

[JOHNSON.CA_QRB]REP010_RAW_OLD.TXT;1                                          2
[SMITH.V11]TEST_RT.DAT;1                                                      5
[JONES.WORK]MTF_MAIN.UIL;30                                                   1
[LINCOLN.UGUIDE]UGUIDE.SAV;2                                                  9
[SMITH.V11]PSPABLDKB.LIS;1                                                    1

[DESMOND.BOOKS]AA_WHOLE.PS;3                                                  4
[LINCOLN.UGUIDE]PSCP_MAST.PS;1                                                2
[WALTERS.V11-38]PSDC$OLB.OLB;3                                                2
```

Determine if Disk Use Corresponds to Initialization

To determine if disk usage corresponds to the way it is initialized and mounted, examine the Detailed Volume Analysis section. Note the actual values of items affected by initializing and mounting, and compare them to the values they were assigned when the disk was initialized and mounted.

```
Disk Analysis              _NUN$DUA32:  (COOLER)                  Page   1
Detailed Volume Analysis                          Performance Agent    VX.X
Monday 11-MAY-2008 14:50
Items preceded by 'I','M' or 'S' are controlled by Initialize,
Mount or Sysgen.

(I  ) Volume name  is 'COOLER      '.
(I  ) Serial number is          0.
(I  ) Creation date was  8-JAN-1998 14:54:05.19.
(I  ) Volume owner is 'SYSTEM      '.
(IM ) Owner uic is [00001,000004].
(I  ) Format type  is 'DECFILE11B  '.
(IM ) Volume protection is [RWED,RWED,RWED,RWED].
(IMS) Default data checking is NOREAD-CHECK, NOWRITE-CHECK.
(I  ) Structure level is   2, version   1.
(I  ) Allocation cluster size is     3 blocks.
(I  ) Index file bitmap is located at LBN     445614.
(IM ) Default file extension is     5 blocks.
(IM ) Default window size is   7 retrieval pointers.
(I  ) Maximum number of files allowed is     111384.
(IMS) Default number of cached directories is    3.

Volume size is     891072 blocks with
51 blocks/track,
14 tracks/cylinder,
1248 cylinders/volume.
```

In the Detailed Volume Analysis section, items preceded by an *I* are controlled by disk initialization. Items preceded by an *M* are controlled by disk mounting, and items preceded by an *S* are controlled by system parameters. In the preceding example, the item Default file extension is preceded by both *I* and *M*. Therefore, disk initialization and mounting control this item.

Compare the value of the default file extension to the value of mean allocated blocks per extent in the Allocated Space section. If the value of the default file extension is significantly larger than the mean allocated blocks, then files are allocated more blocks than they actually need. Reinitializing the disk and reducing the default file extension size may improve disk utilization.

Chapter 9: Performance Agent Commands

This chapter describes the ADVISE COLLECT and ADVISE ARCHIVE commands, their qualifiers, keywords, and options. For a complete description of the ADVISE EDIT commands, see the chapter Parameter Editor Commands (see page 155).

At installation, the CA Performance Agent software adds these commands to the DCL command table. The ADVISE EDIT command invokes the Parameter Editor for working with archiving, workload, and rules definitions. To initiate a Performance Agent function, issue the ADVISE COLLECT command and the appropriate option to perform the desired task, for example, ADVISE COLLECT START.

This section contains the following topics:

ADVISE ARCHIVE

This command summarizes the daily data files into history data files.

Format

```
ADVISE ARCHIVE
```

Description

The ADVISE ARCHIVE command causes the Performance Agent software to read and summarize CA daily data (from the CPD collection definition), history files defined in the parameters file.

Performance Agent starts the update with the date specified with the /BEGINNING qualifier or the earliest date for which a history file needs updating and for which there is daily data. Performance Agent stops the update on the date indicated with the /ENDING qualifier or up to the last full calendar day.

History files for all defined history descriptors are updated, unless a subset is specified with the history descriptor qualifier. Also, by default, all nodes in the cluster are updated (as defined in the schedule file) unless a subset is specified with the /NODE_NAMES qualifier.

The history file name is as follows:

`path:PSDC$`*`nodename_dd-mmm-yyyy_dd-mmm-yyyy.name`*

path

Represents the directory specified in the collection definition for the given node.

nodename

This is the name of the node from which the data is archived.

dd-mmm-yyyy

This is the second occurrence represents the end date of the historical data file.

name

This is the unique history file descriptor name identifies the database specification in the parameters file.

Performance Agent creates a journal file for each open history file. Performance Agent gives the journal file a name that corresponds to the history file and a file extension *.name_JOU*. If the history file is incomplete, Performance Agent uses the journal file and the incomplete file to recover. If the history file descriptor has a specified period attribute and an archiving failure, the original version of the history file is retained and no new version is created.

Use the sample command procedure PSDC$EXAMPLES:PSDC$ARCHIVE.COM to do daily Performance Agent archiving when submitted to batch.

Qualifiers

/BEGINNING=date

Specifies the date on which archiving is to begin. Date represents both the date and time in standard DCL format.

Used exclusively for the first invocation of the ADVISE ARCHIVE command, with respect to unique history file descriptors. This qualifier marks the start of historical logging of CA daily data. The data defaults are as follows:

- Date of data last archived.
- The previous day, in the absence of previously archived data.

/ENDING=date

Marks the end of historical logging of CA Performance Management daily data. The default is the date of the last full calendar day.

Date represents both the date and time in standard DCL format.

/HISTORY_DESCRIPTOR=history-descriptor-name[,. . . .]

Specifies the history file descriptor and associated history files to be updated. By default, all history files are updated if you do not specify the /HISTORY_DESCRIPTOR qualifier.

/NODE_NAMES=nodename[,. . . .]

You can specify a list of node names for which each node's daily data is archived. The default is to archive data from all nodes in the CA schedule file defined under the CPD collection definition. This qualifier is not typically needed.

/RULES=file-spec/NORULES

Specifies an auxiliary rules file to incorporate rule occurrences into the archived data. The /RULES qualifier overrides automatic augmenting to the CA knowledge base, if established.

ADVISE COLLECT

Provides the mechanism for communicating with the Performance Agent, modifying the schedule file, and viewing collected data.

Format

```
ADVISE COLLECT option
```

The ADVISE COLLECT command options are described individually in this chapter. If you do not specify an option the command defaults to ADVISE COLLECT SHOW ALL, which displays Performance Agent status and the contents of the schedule file. The following table lists all the ADVISE COLLECT command options:

Option	Function
ADD	Adds a collection definition
DELETE	Deletes a collection definition
DISPLAY	Displays file activity
MODIFY	Changes parameters for data collection
REPORT	Generates dump reports
SHOW	Displays information about the Performance Agent, the schedule file, and the version of the Performance Agent component being used
START	Starts data collection on this node
STOP	Stops data collection
SYSTEM_PC	Invokes the system Program Counter sampling utility

ADVISE COLLECT ADD

This command adds a collection definition.

Format

```
ADVISE COLLECT ADD collection-definition-name nodename[,. . . ]
```

Description

The ADVISE COLLECT ADD command creates a new collection definition in the schedule file, establishing the scheduling and collection parameters.

Parameters

collection-definition-name

Specifies the name to be used for the collection definition. The maximum number of characters is twenty.

nodename

Specifies the node or list of nodes on which you want to collect data. You can enter the wild card (*) to specify all current cluster members.

You can use the positional /PATH, /FREE_SPACE, and /WORKING_SET qualifiers after each node to set parameters individually.

Command Qualifiers

/BEGINNING=date

Specifies the date on which the daily data collection is to begin. Date represents both the date and time in standard DCL format.

If the date is in the future, any node running the data collection process begins collecting data when the beginning date is reached. If the beginning date is in the past, any node running the data collection process collects data provided that the ending date has not been reached and that the weekly data collection schedule indicates that data should be collected. The default value for the /BEGINNING qualifier is the date that the Performance Agent software was installed or the date on which the schedule file was created.

/ENDING=date

Specifies the date on which the daily data collection is to end. The ending date represents both the date and time in standard DCL format. By default, /ENDING is set to 1-JAN-2010. When the date specified with the /ENDING qualifier is reached, the data collection process is suspended.

If the ending date is in the future, all nodes listed in the schedule file collect data according to the data collection schedule and the /BEGINNING qualifier. If the ending date is in the past, the data collection is suspended and is not resumed until you change the ending date.

/CLASSES= list of classes

Specifies the classes of daily data that are kept. The default is ALL.

- ALL_METRICS
- PROCESS_METRICS

- IO_METRICS
- SYSTEM_METRICS

Option Choice	Type of Data Recorded
SYSTEM_METRICS	Parameter Records (PAR_REC)
	Metrics Records (MET_REC)
	Communication Records (COM_REC)
	CPU Records (CPU_REC)
IO_METRICS	Disk Records (DSK_REC)
	Tape Records (MAG_REC)
	Communication Records (COM_REC)
	Configuration Records (CFG_REC)
	Hot File Records (FIL_REC)
PROCESS_METRICS	Process Records (PRO_REC)
	Metrics Records (MET_REC)
	CPU Records (CPU_REC)
ALL_METRICS	Process Records (PRO_REC)
	Disk Records (DSK_REC)
	Tape Records (MAG_REC)
	Parameter Records (PAR_REC)
	Metrics Records (MET_REC)
	Communication Records (COM_REC)
	Configuration Records (CFG_REC)
	CPU Records (CPU_REC)
	Hot File Records (FIL_REC)

```
/EXCLUDE= ([NO]PROCESSES=(process_name,. . . ),
[NO]PIDS=(pid,. . . ),
[NO]DEVICES=(device_name,. . . ),
[NO]VOLUMES=(volume_name,. . . ))
```

This qualifier allows you to change or add items to the exclude list associated with the collection definition. The main "CPD" collection definition does not allow exclusion items to be specified, however any alternate collection definition does.

If you choose to exclude specific processes for data collection, those processes are omitted from the collection. All other processes have data collected. Processes may be identified by either the PID, or the Process Name, but not both. Processes may not be excluded if they are being included, so the /EXCLUDE and /INCLUDE qualifier are mutually exclusive for process data.

If you choose to exclude specific devices or volumes for data collection, those devices or volumes are omitted from the collection. All other devices or volumes have data collected. Disk devices may be identified by either the device name, or the volume name, but not both. Devices may not be 'excluded' if they are being included, so the /EXCLUDE and /INCLUDE qualifier are mutually exclusive for device data.

Negating the exclude keywords PROCESSES, PIDS, DEVICES, or VOLUMES allows you to remove all or any entries from the exclude list.

/DELETE_AFTER=number_of_days

Specifies the number of days that daily data is kept before the data collection process automatically deletes it. Number_of_days must be an integer from 1 to 99999. The default value is 7 for the CPD collection definition. The default value is 99,999 for alternate collection definitions.

```
/INCLUDE= ([NO]PROCESSES=(process_name,. . . ),
[NO]PIDS=(pid,. . . ),
[NO]DEVICES=(device_name,. . . ),
[NO]VOLUMES=(volume_name,. . .))
```

This qualifier allows you to change or add items to the list associated with the collection definition. You can specify items for inclusion for all but the CPD collection definition.

If you choose to include specific processes for data collection, only those processes are targeted for collection. All other processes do not have any data collected. Processes may be identified by either the PID, or the Process Name, but not both. Processes may not be 'excluded' if they are being included, so the /EXCLUDE and /INCLUDE qualifier are mutually exclusive for process data.

If you choose to include specific devices or volumes for data collection, only those devices or volumes are targeted for collection. All other devices or volumes so not have any data collected. Disk devices may be identified by either the device name, or the volume name, but not both. Devices may not be 'excluded' if they are being included, so the /EXCLUDE and /INCLUDE qualifier are mutually exclusive for device data.

Negating the include keywords PROCESSES, PIDS, DEVICES, or VOLUMES allows you to remove all or any entries from the include list.

/INTERVAL= seconds

Specifies the interval at which daily data is collected. The default value is 120 seconds. This applies only to alternate collection definitions. The interval is fixed for the primary Performance Agent.

/SCHEDULE=({day=(hour-range)[,. . .]|NO day}[,. . .])

Revises the weekly schedule for the specified collection definition. You can use this qualifier to suspend data collection to control the size of the daily data files.

day

Set as SUNDAY, MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, EVERYDAY, WEEKDAYS, or WEEKENDS.

hour-range

Is specified as m-n, where m and n are numbers from 0 to 24, and m is less than n. More than one hour range can be specified for a given day. Hour-range is mutually exclusive with the NO option.

If a day keyword is omitted, the schedule for that day is not changed. Data collection can be suspended for individual days of the week by negating the keyword (for example, NOSUNDAY) or for all of the days of the week by specifying the NOEVERYDAY keyword. The values [NO]WEEKDAYS and [NO]WEEKENDS enable or disable all data collection on weekdays and weekends.

You must specify an hour range for any day keyword. Do not include an hour range if you are specifying a negated day keyword, such as NOMONDAY.

Less inclusive keyword values override more inclusive values. For example, MONDAY=10-12 overrides EVERYDAY=8-17 for Monday however, the Performance Agent collects data from 8:00 a.m. to 5:00 p.m. for all of the other days of the week.

Positional Qualifiers

/WORKING_SET

Specifies the working set quota for the Performance Agent process. The minimum value is 1024. The default working set quota is 2048, and working set extent is 20K. You can override the defaults by specifying different values for quota.

Increase the working set quota when the Performance Agent page faults heavily.

Following the node name parameter in the command, it specifies an integer value for that node only.

Following the collection definition name, it specifies the default value for the entire collection definition.

/FREE_SPACE=n

Specifies the minimum amount of disk space required to keep the Performance Agent running. If the number of blocks falls below this value the Performance Agent notifies you in the Performance Agent log file (PSDC$DATABASE:PSDC$LOG) and pauses until sufficient disk space is released. The default value is 2500 and the range is 0-999,999.

Following the node name parameter in the command, it specifies an integer value for that node only.

Following the collection definition name, it specifies the default value for the entire collection definition.

/PATH=device:[directory]

Following the node name parameter in the command, it specifies a string value for the location for the data files for that node only.

Following the collection definition name, it specifies the default value for the location for the data files for the entire collection definition.

Specify a device and directory or a logical name that indicates a disk directory, which contains the daily collection data files.

Example

```
$ ADVISE COLLECT ADD PRIMETIME/PATH = $DISK1:[PSDC_DATA]-
_$ /SCHEDULE=(WEEKDAYS = 8-17,NOWEEKENDS)-
_$ /INTERVAL = 300/FREE_SPACE = 50000-
_$ /CLASSES = (SYSTEM_METRICS,IO_METRICS)-
_$ BERT/PATH=BERT$DISK:[BERTS_PSDC_DATA]-
_$ /WORKING_SET=1536/FREE_SPACE=10000,-
_$ ERNIE/PATH=ERNIE$DISK:[ERNIE_PSDC_DATA]-
_$ /FREE_SPACE = 25000,BIGBRD
```

This command creates a collection definition with entries for the following three nodes:

Node	Path	Free space	Working set
BERT	BERT$DISK:	10000	2048
ERNIE	ERNIE$DISK:	25000	1536
BIGBRD	PSDC$DATABASE:	50000	2048

- Weekday data collection occurs between 8:00 a.m. and 5:00 p.m.
- Data is not collected on Saturday and Sunday.
- Data is recorded every 300 seconds (5 minutes).
- Only data included in the System metrics and I/O metrics classes are collected.

ADVISE COLLECT DELETE

This command deletes a collection definition.

Format

```
ADVISE COLLECT DELETE collection-definition-name
```

Description

The ADVISE COLLECT DELETE command removes a collection definition from the schedule file.

Parameters

collection-definition-name

Specifies the name of the collection definition to be deleted. The following command displays the collection definitions in the schedule file:

```
$ ADVISE COLLECT SHOW SCHEDULE PRIMETIME
= Performance Agent Schedule File =
=====================================================================
Collection Name //   Attributes    //  Weekly Schedule
=====================================================================
PRIMETIME                                | Monday    8-17
        Delete After   99999 Days                | Tuesday   8-17
        HotFile Que.   0.33                      | Wednesday 8-17
        Interval/sec   300.0                     | Thursday  8-17
        Start Date     2-MAR-1993                | Friday    8-17
        End Date       1-JAN-2010                | Saturday
        Classes: I/O   Metrs                     | Sunday
        Def Path: $DISK1:[PSDC_DATA]
        Def Minimum Free Space: 50000
        Def Working Set Size:   2048
---------------------------------------------------------------------
-Node---DC Status-------DskFreeSpace--WSquota---Path---------------
---------------------------------------------------------------------
BERT    Down               10000       1536     BERT$DISK:[BERTS_PSDC_DATA]
ERNIE   Down               25000       2048     ERNIE$DISK:[ERNIE_PSDC_DATA]
BIGBRD  Down               50000       2048     $DISK1:[PSDC_DATA]
CA Performance Agent Module version is: VX.X
```

Example

```
$ ADVISE COLLECT DELETE PRIMETIME
```

ADVISE COLLECT DISPLAY FILES

This command displays the most frequently accessed files on a disk.

Format

```
ADVISE COLLECT DISPLAY FILES
```

Description

Use the ADVISE COLLECT DISPLAY FILES command to display hot file activity on your screen. While displaying file activity, you can enter commands to control the display and its characteristics. The following table describes the File Activity display interactive subcommands.

Subcommand	Meaning
EXIT or CTRL+Z	Exits from the display
HELP	Displays HELP information for field descriptions and subcommands
SET DISK	Defines the disks for which statistics are displayed
SET INTERVAL	Defines the rate at which disk statistics are displayed
SHOW DISK	Shows the names of disks for which statistics are displayed
SHOW INTERVAL	Shows the value of the current display interval
SHOW ALL	Shows the value of the current display interval and the names of disks for which statistics are displayed
SPAWN	Spawns a subprocess
WRITE [file-spec]	Copies the screen contents to a file. If no file is specified, the default of PSDC$DISPLAY_FILES.RPT is used

For more information, see the chapter Display File Activity (see page 107).

Qualifiers

/DISK = (disk[,. . .])

Specifies the disks on which you wish to monitor file activity with either the physical device name or equivalent logical name. If you do not specify the /DISK qualifier, Performance Agent displays file activity on SYS$SYSDEVICE.

/INTERVAL= seconds

Specifies the interval at which file activity is displayed. The /INTERVAL value must be in the range of 1 to 3600 seconds. If you do not specify an interval, the default value is 10 seconds.

Example

```
$ ADVISE COLLECT DISPLAY FILES/DISK = $1DUA7/INTERVAL = 120
```

ADVISE COLLECT MODIFY

This command controls data collection.

Format

```
ADVISE COLLECT MODIFY collection-definition-name nodename[,. . .]
```

Description

The ADVISE COLLECT MODIFY command enables you to change parameters of a collection definition in the schedule file. Any active Performance Agents in a cluster reads the schedule file and institute changes at that time.

Parameters

collection-definition-name

Specifies the name of the collection definition.

nodename

Specifies the node or list of nodes on which you want to collect data. You can enter the wild card (*) to specify all current cluster members.

You can use the positional /PATH, /FREE_SPACE, and /WORKING_SET qualifiers after each node to set parameters individually.

Command Qualifiers

/BEGINNING=date

Specifies the date on which the daily data collection is to begin. Date represents both the date and time in standard DCL format.

If the date is in the future, any node running the data collection process begins collecting data when the beginning date is reached. If the beginning date is in the past, any node running the data collection process collects data provided that the ending date has not been reached and that the weekly data collection schedule indicates that data should be collected. The default value for the /BEGINNING qualifier is the date that the Performance Agent software was installed or the date on which the schedule file was created.

/ENDING=date

Specifies the date on which the daily data collection is to end. The ending date represents both the date and time in standard DCL format. By default, /ENDING is set to 1-JAN-2010. When the date specified with the /ENDING qualifier is reached, the data collection process is suspended.

If the ending date is in the future, all nodes listed in the schedule file collect data according to the data collection schedule and the /BEGINNING qualifier. If the ending date is in the past, the data collection is suspended and is not resumed until you change the ending date.

/CLASSES= list of classes

Specifies the classes of daily data that are kept. The default is ALL.

- ALL_METRICS
- PROCESS_METRICS
- IO_METRICS
- SYSTEM_METRICS

Option Choice	Type of Data Recorded
SYSTEM_METRICS	Parameter Records (PAR_REC) Metrics Records (MET_REC) Communication Records (COM_REC) CPU Records (CPU_REC)
IO_METRICS	Disk Records (DSK_REC) Tape Records (MAG_REC) Communication Records (COM_REC) Configuration Records (CFG_REC) Hot File Records
PROCESS_METRICS	Process Records (PRO_REC) Metrics Records (MET_REC) CPU Records (CPU_REC)
ALL_METRICS	Process Records (PRO_REC) Disk Records (DSK_REC) Tape Records (MAG_REC) Parameter Records (PAR_REC) Metrics Records (MET_REC) Communication Records (COM_REC) Configuration Records (CFG_REC) CPU Records (CPU_REC) Hot File Records (FIL_REC)

```
/EXCLUDE= ([NO]PROCESSES=(process_name,. . .),
[NO]PIDS=(pid,. . . ),
[NO]DEVICES=(device_name,. . .),
[NO]VOLUMES=(volume_name,. . .))
```

This qualifier allows you to change or add items to the exclude list associated with the collection definition. The main CPD collection definition does not allow exclusion items to be specified, however any alternate collection definition does.

If you choose to exclude specific processes for data collection, those processes are omitted from the collection. All other processes have data collected. Processes may be identified by either the PID, or the Process Name, but not both. Processes may not be 'excluded' if they are being included, so the /EXCLUDE and /INCLUDE qualifier are mutually exclusive for process data.

If you choose to exclude specific devices or volumes for data collection, those devices or volumes are omitted from the collection. All other devices or volumes have data collected. Disk devices may be identified by either the device name, or the volume name, but not both. Devices may not be 'excluded' if they are being included, so the /EXCLUDE and /INCLUDE qualifier are mutually exclusive for device data.

Negating the exclude keywords PROCESSES, PIDS, DEVICES, or VOLUMES allows you to remove all or any entries from the exclude list.

/DELETE_AFTER=number_of_days

Specifies the number of days that daily data is kept before the data collection process automatically deletes it. Number_of_days must be an integer from 1 to 99999. The default value is 7 for the CPD collection definition. The default value is 99,999 for alternate collection definitions.

/HOTFILE_QUEUE=value

OR/NOHOTFILE_QUEUE

Specifies the minimum average number of jobs in a queue for any disk before the hottest files for that disk are collected. This applies to the CPD collection definition only as alternate Performance Agents do not collect hot file data.

The queue size is averaged over a 2-minute time sample for each disk. If at the end of the 2 minutes, the number of jobs in the queue is greater than or equal to the set value, then the primary Performance Agent records the top five hottest files in the following six categories:

- Operations count
- Throughput

- Paging
- Swapping operations
- Split I/O rate
- Service time
- Read operations

The default is 0.33. The valid values are any positive floating point decimal numbers less than 100.00 and greater than or equal to 0.00.

The /NOHOTFILE_QUEUE qualifier disables the collection of hot files and as a side effect, reduces the Performance Agent overhead and disk space used.

```
/INCLUDE= ([NO]PROCESSES=(process_name,. . . ),
[NO]PIDS=(pid,. . .),
[NO]DEVICES=(device_name,. . .),
[NO]VOLUMES=(volume_name,. . .))
```

This qualifier allows you to change or add items to the list associated with the collection definition. You can specify items for inclusion for all but the CPD collection definition.

If you choose to include specific processes for data collection, only those processes are targeted for collection. All other processes do not have any data collected. Processes may be identified by either the PID, or the Process Name, but not both. Processes may not be 'excluded' if they are being included, so the /EXCLUDE and /INCLUDE qualifier are mutually exclusive for process data.

If you choose to include specific devices or volumes for data collection, only those devices or volumes are targeted for collection. All other devices or volumes do not have any data collected. Disk devices may be identified by either the device name, or the volume name, but not both. Devices may not be 'excluded' if they are being included, so the /EXCLUDE and /INCLUDE qualifier are mutually exclusive for device data.

Negating the include keywords PROCESSES, PIDS, DEVICES, or VOLUMES allows you to remove all or any entries from the include list.

/INTERVAL=seconds

Specifies the interval at which daily data is collected. The default value is 120 seconds. This applies only to alternate collection definitions. The interval is fixed for the primary Performance Agent.

/REMOVE_NODES=(nodename,[,nodename . . .])

This qualifier allows you to remove nodes associated with a particular Collection Definition. Wildcards may be used to remove all nodes or a single node for all Collection Definitions.

/REORDER_NODES=(nodename[,. . .])

Use this qualifier to change the order of the nodes listed in a particular collection definition or all collection definitions by using the wildcard (*) for the collection definition in the command.

/SCHEDULE=({day=(hour-range)[,. . .]|NOday}[,. . .])

Revises the weekly schedule for the specified collection definition. You can use this qualifier to suspend data collection to control the size of the daily data files.

day

Set as SUNDAY, MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, EVERYDAY, WEEKDAYS, or WEEKENDS.

hour-range

Is specified as m-n, where m and n are numbers from 0 to 24, and m is less than n. More than one hour range can be specified for a given day. Hour-range is mutually exclusive with the NO option.

If a day keyword is omitted, the schedule for that day is not changed. Data collection can be suspended for individual days of the week by negating the keyword (for example, NOSUNDAY) or for all of the days of the week by specifying the NOEVERYDAY keyword. The values [NO]WEEKDAYS and [NO]WEEKENDS enable or disable all data collection on weekdays and weekends.

You must specify an hour range for any day keyword. Do not include an hour range if you are specifying a negated day keyword, such as NOMONDAY.

Less inclusive keyword values override more inclusive values. For example, MONDAY=10-12 overrides EVERYDAY=8-17 for Monday however, the Performance Agent collects data from 8:00 a.m. to 5:00 p.m. for all of the other days of the week.

Positional Qualifiers

/WORKING_SET

Specifies the working set quota for a Performance Agent process. The minimum value is 1024. The default working set quota is 2048, and working set extent is 20K. You can override the defaults by specifying different values for quota.

Increase the working set quota when the Performance Agent page faults heavily.

Following the nodename parameter in the command, it specifies an integer value for that node only.

Following the collection definition name, it specifies the default value for the entire collection definition.

/FREE_SPACE

Specifies the minimum amount of disk space required to keep the Performance Agent running. If the number of blocks falls below this value the Performance Agent notifies you in the Performance Agent log file (PSDC$DATABASE:PSDC$LOG)and pauses until sufficient disk space is released. The default value is 2500 and the range is 0 -- 999,999.

Following a nodename parameter in the command, it specifies an integer value for that node only.

Following the collection definition name, it specifies the default value for the entire collection definition.

/PATH=device:[directory]

Following a nodename parameter in the command, it specifies a string value for the location for the data files for that node only.

Following the collection definition name, it specifies the default value for the location for the data files for the entire collection definition.

Specify a device and directory, or a logical name that indicates a disk directory, which contains the daily collection data files.

If you modify the path, you must move or rename all the existing data and history files to maintain access to all files.

Example

```
$ ADVISE COLLECT MODIFY PRIMETIME/SCHEDULE=(WEEKDAYS = 9-18)-
_$ /REORDER_NODES=(BERT,BIGBRD,ERNIE)-
_$ BERT/WORKING_SET = 1800, BIGBRD/FREE_SPACE = 1000
```

The previous command has the following scheduling effects:

- Weekday data collection occurs between 9:00 a.m. and 6:00 p.m.
- The order of listing of the nodes is changed.

```
                  = Performance Agent Schedule File =
================================================================================
Collection Name //   Attributes          //  Weekly Schedule
================================================================================
PRIMETIME                                 | Monday    9-18
         Delete After  99999 Days         | Tuesday   9-18
         HotFile Que.    0.33             | Wednesday 9-18
         Interval/sec   300.0             | Thursday  9-18
         Start Date     2-MAR-2008        | Friday    9-18
         End Date       1-JAN-2010        | Saturday
         Classes: I/O   Metrs             | Sunday
         Def Path: $DISK1:[PSDC_DATA]
         Def Minimum Free Space: 50000
         Def Working Set Size:   2048
--------------------------------------------------------------------------------
-Node---DC Status-------DskFreeSpace--WSquota---Path---------------------
--------------------------------------------------------------------------------
BERT    Down               10000       1800   BERT$DISK:[BERTS_PSDC_DATA]
BIGBRD  Down                1000       2048     $DISK1:[PSDC_DATA]
ERNIE   Down               25000       2048  ERNIE$DISK:[ERNIE_PSDC_DATA]
```

ADVISE COLLECT REPORT DISK_SPACE

This command generates reports for analysis of disk utilization.

Format

```
ADVISE COLLECT REPORT DISK_SPACE device_name[,. . .])
```

Description

Use the ADVISE COLLECT REPORT command to produce Performance Agent Disk Space Analysis reports.

Parameters

device_name

Specifies the name of the disk devices to be covered by the report. Multiple disks may be specified. Either physical device names or logical names may be specified.

Qualifiers

/OUTPUT=file-spec

Writes requested reports to the specified file name. If you do not specify the /OUTPUT qualifier, Performance Agent writes the report to the device defined by the logical SYS$OUTPUT.

ADVISE COLLECT REPORT DUMP_keyword

This command generates displays of raw data using daily and historical data.

Format

```
ADVISE COLLECT REPORT DUMP_report_keyword[,. . .]
```

Description

Use the ADVISE COLLECT REPORT command to produce Performance Agent dump reports. The Performance Agent can generate dump reports using either daily or historical data. Dump reports consist of data extracted from the CA database files for a specified time period. You can specify more than one dump report within the same command line.

For a sample of each dump report, see the chapter "Dump Reports". The following is a keyword list for the types of dump reports available:

```
DUMP_ALL
DUMP_COMMUNICATIONS
DUMP_CONFIGURATION
DUMP_CPUS
DUMP_DATACELLS
DUMP_DISKS
DUMP_HOTFILES
DUMP_METRICS
DUMP_PARAMETERS
DUMP_PROCESSES
DUMP_RULES
DUMP_TAPES
```

You can request different types of dump reports in one command.

Qualifiers

/BEGINNING=date

Specifies the beginning date and time of a reporting period. The date and time format is the standard DCL format, either absolute or relative. If you do not specify the /BEGINNING qualifier, Performance Agent uses 00:00:00 on the same day for which the ending date and time is specified. If you do not specify an /ENDING qualifier, Performance Agent uses 00:00:00 of the current day as the default beginning time.

/COLLECTION_DEFINITION=name

Specifies the name of the collection definition, and therefore, the set of data to be used for the report. If you omit this, the report uses the CPD collection definition.

This qualifier is exclusive with the /HISTORY_DESCRIPTOR qualifier.

To view the Collection Definitions available, use the DCL command ADVISE COLLECT SHOW ALL.

/CLASSIFY_BY=USERGROUP=family_name

Used with the /HISTORY_DESCRIPTOR qualifier to specify which user group family to include in the dump report. The specified family must be saved in the history data. For more information, see the Parameter Editor ADD/HISTORY command.

/DATES=dates-file-spec

Specifies the dates and times to be covered by the report from the file you have created. When you create this file, enter dates in the format using one line per day:

DD-MMM-YYYY HH:MM:SS, DD-MMM-YYYY HH:MM:SS

Use the .txt extension when you name the file.

/ENDING=date

Specifies the ending date and time of a reporting period. The date and time format is the standard DCL format, either absolute or relative. If you do not specify an /ENDING qualifier, Performance Agent defaults to 23:59:59 on the same day for which you specified the beginning date and time or to the current time if the reporting period is for the current day.

/FILTER = filter_keyword

Enables you to select processes or disks for dump reports. The strings specified with this qualifier are case sensitive.

Keyword	Meaning
USERNAME=(string[,. . .])	With the DUMP_PROCESSES report, generates a report on all processes matching the string.
IMAGENAME=(string[,. . .])	With the DUMP_PROCESSES report, generates a report on all images matching the string. Do not specify the device, directory, file extension (exe), or version number.
PROCESSNAME=(string[,. . .])	With the DUMP_PROCESSES report, generates a report on all process records matching the string.
UIC=(uic[,. . .])	With the DUMP_PROCESSES report, generates a report on all process records with the UIC matching the string. You can use an asterisk (*) in either the group or user field of the specified UIC to include all match variations.
PID=(pid[,. . .])	With the DUMP_PROCESSES report, generates a report on all process records matching the specified PID.
VOLUMENAME=(string[,. . .])	With the DUMP_DISK report, generates a report on all entries for all disk records with the volume name matching the string. Do not specify any trailing colons.
DEVICENAME=(string[,. . .])	With the DUMP_DISK report, generates a report on all entries for all disk records with the device name matching the string. Do not specify any trailing colons.

/HISTORY_DESCRIPTOR=history-descriptor-name

Allows you to select history data from the CA database. By default, Performance Agent uses daily data files to supply data for reporting. However, by specifying the name of a history file descriptor, you can select historical data instead.

The history file descriptor must be defined in the parameters file.

/NODE_NAMES=(nodename[,. . .])

Generates a dump report for a specific node or set of nodes. If you do not specify a value for the /NODE_NAMES qualifier, Performance Agent uses the list of nodes specified in the collection definition.

/OUTPUT=file-spec

Writes requested reports to the specified file name. If you do not specify the /OUTPUT qualifier, Performance Agent writes the report to the device defined by the logical SYS$OUTPUT. The default file type is .LIS.

/RULES[=file-spec]

/NORULES

Only available when the Performance Manager is also installed, the /RULES qualifier loads information from the rules file to establish user-defined hardware scaling factors. The file-spec must point to an auxiliary knowledge base which has previously been compiled with the ADVISE PERFORMANCE COMPILE command. The default file type is .KB. If the NORULES qualifier is specified no augmentation of the factory rules occur.

/SCHEDULE=({day=(hour-range)[,. . .]|NOday}[,. . .])

Specifies that a subset of CA data is to be used (or not used) to generate dump reports. By default, Performance Agent selects all data between the /BEGINNING time and the /ENDING time.

day

Set as SUNDAY, MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, EVERYDAY, WEEKDAYS, or WEEKENDS.

hour-range

Is specified as m-n, where m and n are numbers from 0 to 24, and m is less than n. More than one hour range can be specified for a given day. Hour-range is mutually exclusive with the NO option.

If you omit a day keyword, the data for that day is selected. You can prevent data selection for individual days of the week by negating the keyword (for example, NOSUNDAY) or for all of the days of the week by specifying the NOEVERYDAY keyword. Similarly, you can use the values [NO]WEEKDAYS and [NO]WEEKENDS to enable or disable data selection for weekdays and weekends.

You must specify an hour range for any day keyword. Do not include an hour range if you are specifying a negated day keyword, such as NOMONDAY.

Less inclusive keyword values override more inclusive values. For example, MONDAY=10-12 overrides EVERYDAY=8-17 for Monday; however, the Performance Agent collects data from 8:00 a.m. to 5:00 p.m. for all of the other days of the week.

ADVISE COLLECT REPORT SYSTEM_PC

This command generates reports for evaluating system CPU use.

Format

```
ADVISE COLLECT REPORT SYSTEM_PC
```

Description

Use the ADVISE COLLECT REPORT SYSTEM_PC command to summarize the data in the system PC file and produce a report of system PC and PSL information by process.

Parameters

pc-file-spec

This required parameter specifies the file containing the collected PC samples.

Qualifiers

CLASS =(item[,. . .])

Specifies which classes of statistics are included in the report. You can select any of the following class items:

Keyword	Meaning
ALL	All statistics (same as DEFAULT_STATISTICS)
DEFAULT_STATISTICS	Report PROCESS, MODULE, IMAGE, IPL, DRIVER, EXEC
DRIVER	Report Device Drivers and System Code
EXEC	Report OpenVMS Executive Layout
PROCESS[:(CPU_UTILIZATION, ALL)]	Report PC samples by PROCESS
MODULE	Report by MODULE name within EXEC
IMAGE	Report by IMAGE name within EXEC
IPL	Generate IPL usage graphs

All of the class item keywords are also available in the negated form to indicate that reporting of a particular class of data is not desired. In addition, support is provided for allowing:

```
/CLASS=(ALL[,negated-keyword])
```

This allows an "all but these" capability.

```
If the /CLASS keyword is omitted from the command line, then
/CLASS=DEFAULT_STATISTICS is assumed.
```

/IDENTIFICATION=pid qualifier

Specifies the filter process identification (PID) for report generation purposes.

If the /IDENTIFICATION qualifier is present in the command line, then a process identification value must be specified.

If the /IDENTIFICATION qualifier is omitted from the command line, then /NOIDENTIFICATION is assumed and the interrupt stack is used as the filter for the report.

/OUTPUT=file-spec

Writes requested reports to the specified file name. If you do not specify the /OUTPUT qualifier, Performance Agent writes the report to the device defined by the logical SYS$OUTPUT.

ADVISE COLLECT SHOW

This command displays information about the data collection process.

Format

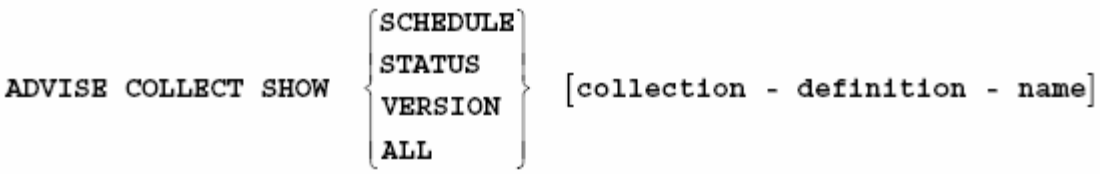

Description

The ADVISE COLLECT SHOW command displays either the schedule file settings, the status of the data collection process and nodes or both, or the version of the Performance Agent component being used.

Parameters

Schedule

The following command displays the contents of the schedule file:

```
$ ADVISE COLLECT SHOW SCHEDULE
```

When you specify SCHEDULE, the following items are displayed:

- Beginning date
- Ending date
- Hot file queue length
- Number of days after which daily data files are deleted
- List of nodes
- Weekly data collection schedule

- Performance Agent working set size
- Performance Agent free space needed
- Path

STATUS

The following command displays the current status of the nodes and the Performance Agents on those nodes:

```
$ ADVISE COLLECT SHOW STATUS PRIMETIME
```

Each node in the schedule file is listed along with the node's status. The status of the data collection process is one of the following:

- Running—The data collection process is running and collecting data.
- Stopped—The Performance Agent is not running on a node in the collection definition. The status of that specific node is unknown.
- SchedWait—The data collection process is running, but data is not being collected due to the schedule specified in the schedule file.
- FreeSpaceWait—The Performance Agent suspends all collection if there are fewer blocks than specified for free space needed on the disk to which the Performance Agent writes data. The DCL command ADVISE COLLECT MODIFY nodename/FREE_SPACE_NEEDED establishes this value.
- NoPrivs—The user issuing the ADVISE COLLECT SHOW command does not have the SYSLCK privilege and is unable to learn the status.

VERSION

The following command displays the version of the Performance Agent component being used.

```
$ ADVISE COLLECT SHOW VERSION
```

ALL

The following command displays all the information listed by the SCHEDULE, STATUS, and VERSION keywords.

```
$ ADVISE COLLECT SHOW ALL
```

The ALL keyword is the default; therefore, specifying ADVISE COLLECT is equivalent to ADVISE COLLECT SHOW ALL.

ADVISE COLLECT START

This command starts the Performance Agent.

Format

```
ADVISE COLLECT START [collection-definition-name]
```

Description

The ADVISE COLLECT START command starts a data collection process using the collection definition specified on the node from which you issue the command. If the collection definition is omitted, the primary Performance Agent is started.

If you are starting a Performance Agent based on an alternate collection definition, the Performance Agent based on the CPD collection definition must already be running.

At installation, a command procedure called UPM$AGENT$STARTUP.COM is created and placed in SYS$STARTUP. This command file contains the ADVISE COLLECT START command, which when invoked from SYS$MANAGER:SYS$STARTUP_V5.COM automatically starts the data collection process.

Parameters

collection-definition-name

Specifies the name of a collection definition on which the Performance Agent process is to be started. If you do not specify a name, the CPD collection definition is started by default.

ADVISE COLLECT STOP

This command stops the Performance Agent.

Format

```
ADVISE COLLECT STOP [collection-definition-name]
```

Description

The ADVISE COLLECT STOP command terminates the data collection process on the node from which you issue the command. Normally, the process does not terminate immediately.

At installation, a command procedure called UPM$AGENT$SHUTDOWN.COM is created and placed in SYS$MANAGER. This command file contains the ADVISE COLLECT STOP command, which when invoked from SYS$MANAGER:SYSHUTDWN.COM terminates the data collection process at normal system shutdown.

Parameters

collection-definition-name

Specifies the name of the collection definition to be stopped. The following command displays the collection definitions in the schedule file:

```
$ ADVISE COLLECT SHOW SCHEDULE
```

Qualifier

/WAIT

/NOWAIT

The /WAIT qualifier causes the DCL command to wait until the data collection process is fully terminated. The default /NOWAIT causes the DCL command to return without waiting.

ADVISE COLLECT SYSTEM_PC

This command invokes the system Program Counter sampling utility.

Format

```
ADVISE COLLECT SYSTEM_PC
```

Description

Use the ADVISE COLLECT SYSTEM_PC command to record system-wide Program Counter (PC) and Processor Status Longword (PSL) usage metrics.

Qualifiers

/BEGINNING=time

Specifies the time when PC data collection starts. Specify an absolute time, a delta time, or combination. If the /BEGINNING qualifier is omitted or the time is omitted, data collection begins immediately.

/COMMENTS="string"

Specifies an optional comment string 1 to 80 bytes long written to the PC log file and printed on the first page of a report. If you enter the /COMMENT qualifier, you must specify a string. If you do not enter /COMMENT, an 80 byte blank comment string is recorded.

/ENDING=time

Specifies the time when PC data collection stops. Specify an absolute time, a delta time, or combination. If the /ENDING qualifier is omitted or the time is omitted, data collection terminates 15 minutes after starting.

/INTERVAL=ticks

Specifies the amount of elapsed time between PC data collection events. The /INTERVAL value must be in the range from 1 to 100 ten millisecond ticks. If the /INTERVAL qualifier is omitted or the value is omitted, the default value of one ten millisecond tick is used.

/IDENTIFICATION=pid

PC data is collected only for the process ID specified. If you enter the /IDENTIFICATION qualifier, you must specify a process identification value. If the /IDENTIFICATION qualifier is omitted, all PC data is collected regardless of the process that generated it.

/OUTPUT=file-spec

Specifies the name of the log file for the PC data. If the /OUTPUT is omitted or no file is specified, SYSTEM.PCS is used.

ADVISE EDIT

This command invokes the CA Parameter Editor.

Format

```
ADVISE EDIT
```

Description

The ADVISE EDIT command invokes the CA Parameter Editor Utility. After you have invoked the utility, you see the PSDC-EDIT>prompt. At this prompt you can issue utility commands. For the commands that can be used within the Parameter Editor, see chapter "Parameter Editor Commands."

ADVISE/INTERFACE

This command invokes the DECwindows interface.

Format

```
ADVISE/INTERFACE[={DECWINDOWS|MOTIF}]
```

Use the ADVISE/INTERFACE command to access the DECwindows interface. The command can determine which window interfaces and which components of the CA Performance Management software are available on your system. The software displays the DECwindows interface based upon the following order of precedence:

1. MOTIF-PSPA
2. MOTIF-PSDC

The ADVISE/INTERFACE=MOTIF command initiates the Performance Manager (PSPA) Motif interface if present or, if not, the Performance Agent (PSDC) Motif interface.

Chapter 10: Parameter Editor Commands

The Parameter Editor is an editor for the CA Performance Management parameters file. The DCL command ADVISE EDIT invokes the utility and allows you to specify the following items:

- Workload definitions and families
- Data archiving schemes
- Automatic augmentation of the Performance Manager knowledge base (Requires the Performance Manager component.)
- File version limits

This section contains the following topics:

ADVISE EDIT

This command invokes the Performance Management Parameter Edit Utility.

Format

```
ADVISE EDIT
```

Description

The ADVISE EDIT command invokes the Parameter Edit Utility. After you invoke the editor, you see the PSDC-EDIT> prompt. At this prompt, you can enter the commands listed in the following table. Updates to the parameters file are not written to disk until you successfully end the editing session with the EXIT command.

Command	Function
ADD/FAMILY	Adds workload family of workload definitions
ADD/HISTORY	Adds a history file descriptor defining data reduction scheme for historical data
ADD/WORKLOAD	Adds a workload definition
DELETE/FAMILY	Deletes a workload family definition
DELETE/HISTORY	Deletes a history file descriptor
DELETE/WORKLOAD	Deletes a workload definition
EXIT	Saves edits and returns user to DCL command level
MODIFY/FAMILY	Modifies an existing workload family
MODIFY/HISTORY	Modifies an existing history file descriptor
MODIFY/WORKLOAD	Modifies an existing workload definition
QUIT	Returns user to DCL command level, without saving edits
SET [NO]AUTO_AUGMENT	Specifies a user knowledge base to augment the Performance Manager knowledge base
SET VERSION_LIMIT	Sets the number of file versions for parameters file and history files
SHOW/ALL	Displays all information in the parameters file
SHOW/AUTO_AUGMENT	Displays current setting for knowledge base augmentation
SHOW/FAMILY	Displays workload family definitions

Command	Function
SHOW/HISTORY	Displays attributes of history file descriptors
SHOW/VERSION_LIMIT	Displays current setting of version limit for parameters file and history files
SHOW/WORKLOAD	Displays workload definitions

ADD/FAMILY

This command defines a new workload family in the parameters file. You can specify the workload family when generating Performance Evaluation Reports, graphs.

Format

```
ADD/FAMILY name
```

Parameter

name

Name is assigned to the workload family. The name can be from 1 to 20 alphanumeric characters in length.

Qualifier

/WORKLOAD=(workload-definition[,. . .])

Specifies the names of the workload definitions in the workload family you are adding. Any workload definitions you specify must already exist in the parameters file.

workload-definitions

Is a 1 to 20 character alphanumeric string that represents a workload.

The order of the workload definition influences how the Performance Agent classifies process data. During the workload classification, Performance Agent compares the process data with each workload definition by order of the workload family. If workload definitions overlap, the Agent uses the first workload definition in the classification family that matches. Therefore, place more explicit workload definitions before more general workload definitions. If the Agent process data does not match workload definitions, the Agent reports the process data in the workload definition OTHER.

Examples

```
>PSDC-EDIT> ADD/FAMILY
DEPARTMENTS/WORKLOAD=(DEPT_1,DEPT_2,DEPT_3,SYSTEMUSER,OPERATOR)
```

This command creates the workload family named DEPARTMENTS. After defined, this workload classification scheme can be used to generate models, graphs, and Performance Evaluation reports.

To classify the process data by departments, specify the /CLASSIFY_BY=USERGROUP=DEPARTMENTS qualifier.

```
PSDC-EDIT> ADD/FAMILY KEY_APPLICATIONS/WORKLOAD=(APP_X,APP_Y,COMPILES,BACKUP,
PAYROLL,MAIL)
```

This command creates the workload family named KEY_APPLICATIONS. After defined, this workload classification scheme may be used to generate models and Performance Evaluation reports.

To classify the process data by key applications, specify the /CLASSIFY_BY=USERGROUP=KEY_APPLICATIONS qualifier.

ADD/HISTORY

This command adds a history file descriptor to the Performance Management parameters file.

Format

```
ADD/HISTORY name
```

Parameter

name

The history file descriptor name is a 1 to 20-character alphanumeric string. Use this name with the /HISTORY qualifier to generate reports, graphs, and models from historical data.

Description

The ADD/HISTORY command adds a history file descriptor to the Performance Management parameters file. A history file descriptor defines how the Performance Agent software summarizes daily data. You can have multiple history file descriptors. History file descriptors defined in the parameters file control the archival process for daily Performance data. For a discussion on how to archive data, see ADVISE ARCHIVE.

Performance Agent summarizes process data in one of two ways: unclassified, as defined with /MODEL_DATA, or by the workload families specified with the /CLASSIFY_BY qualifier. When you specify /MODEL_DATA, the archived data retains enough process-level detail to defer the classification until the archived data is used, rather than classifying it when it is archived.

If you omit both /MODEL_DATA and /CLASSIFY_BY history file descriptor elements, Performance Agent archives and reports on data based on interactive, batch, and network workloads.

Qualifiers

/CLASSIFY_BY=USERGROUP=(workload-family[,. . .])

Specifies that Performance Agent store process data in the history file, summarized by the workload families indicated. The /CLASSIFY_BY qualifier cannot be used with the /MODEL_DATA qualifier.

After Performance Agent archives data with a given classification scheme, all subsequent reporting must use one of the classification workload families specified.

/DELETE_AFTER=number_of_days

Specifies the retention period for a history file, in days. The default is 180 days.

/GRANULARITY={DAILY | WEEKLY | MONTHLY | QUARTERLY | YEARLY}

Specifies the time span represented in an individual history file. The default value is monthly.

Granularity	Definition
DAILY	Hours 00:00 to 23:59
WEEKLY	Monday through Sunday
MONTHLY	The first day through the last day of the month
QUARTERLY	January through March, April through June, July through September, October through December
YEARLY	January 1 through December 31

/HOLIDAY_SCHEDULE=(day=(hour-range)[,. . .]|

NOday[,. . .])

Selects a part of or all of the daily data for holidays to be stored in the history file.

The template holiday file PSDC$EXAMPLES:PSDC$HOLIDAYS.TXT enables your holiday schedule. You can edit and move this file to the PSDC$DATABASE area. When using the /HOLIDAY_SCHEDULE qualifier, use the following parameters:

day

Set as SUNDAY, MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, EVERYDAY, WEEKDAYS, or WEEKENDS.

hour-range

Is specified as m-n, where m and n are numbers from 0 to 24, and m is less than n. More than one hour range can be specified for a given day. Hour-range cannot be used with the NOday option.

/INTERVAL=minutes

Specifies the time period over which Performance Agent averages daily data records into a single history data record in the history file. The default is 60 minutes. The interval must be 2, 4, 6, 8, 10, 12, 16, 18, 20, 24, 30, 32, 36, 40, 48, 60, 72, 80, 90, 96, 120, 144, 160, 180, 240, 288, 360, 480, 720, or 1440 for Performance Management data.

When larger interval values are specified, less disk space is used by the history file. However, this also causes performance data to be summarized over longer periods of time, reducing peaks and valleys of resource use in follow-on processing.

/MODEL_DATA

Indicates that modeling data is to be saved in the history file. The /MODEL_DATA qualifier allows you to generate a baseline model from this history data. This qualifier cannot be used with the /CLASSIFY_BY qualifier. By default, the /MODEL_DATA qualifier is disabled.

History files generated when the /MODEL_DATA qualifier is enabled provide you total and unrestricted ability to generate subsequent reports or graphs with any existing or future classification scheme defined in the parameters file.

/PERIODICITY ={DAILY | WEEKLY | MONTHLY | QUARTERLY}

/NOPERIODICITY

Specifies the calendar period into which the daily files are summarized. You can summarize time spans of up to a year into an average typical day, week, month, or quarter.

You can summarize a week into a typical day; a month into a typical day or week; a quarter into a typical day, week, or month; and a year into a typical day, week, month or quarter.

Consider using graph averaging in place of periodicity. Graph averaging requires more CPU time while periodicity requires more disk space. For more information on graph averaging, see the Performance Manager Administrator Guide.

Valid periodicity values depend upon the selected granularity. The value of the periodicity must be less than the value of granularity. The default is NOPERIODICITY.

Periodicity	Definition
DAILY	Average day
WEEKLY	7 average days
MONTHLY	31 average days
QUARTERLY	93 average days

The following table shows the effect of periodicity on a file defined with yearly granularity:

Periodicity	Definition
DAILY	Averages each daily record into the same typical day. Performance Management archived data is summarized into one typical day record.
WEEKLY	Averages daily data into a typical week for a year. Performance Management archived data represents an average week in the year.
MONTHLY	Averages daily data into a typical month for a year. Performance Management archived data represents an average month in the year.
QUARTERLY	Averages daily data into a typical quarter for a year. Performance Management archived data represents an average quarter in the year.

/SCHEDULE=({day=(hour-range)[,. . .]|Noday}[,. . .])

Selects the daily data to be stored in the history file. You can select all or just some portion of the daily data. When using the /SCHEDULE qualifier:

day

Set as SUNDAY, MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, EVERYDAY, WEEKDAYS, or WEEKENDS.

hour-range

Is specified as m-n, where m and n are numbers from 0 to 24, and m is less than n. More than one hour range can be specified for a given day. Hour-range cannot be used with the NOday option.

Example

```
PSDC-EDIT> ADD/HISTORY/SCHEDULE=(WEEKDAYS=13-15,NOWEEKENDS)/MODEL_DATA PEAKLOAD
```

This example shows you how to create a new history file descriptor named PEAKLOAD.

ADD/WORKLOAD

This command creates a workload definition in the parameters file.

Format

```
ADD/WORKLOAD name
```

Parameter

name

The workload definition is a 1 to 20-character alphanumeric string which represents a workload.

Description

The ADD/WORKLOAD command creates a workload definition in the parameters file. A workload definition is a list of images or users or both that governs how Performance Agent groups process data. Users can be specified in terms of user names, UIC codes, PIDs, account names, or process names. Mixed combinations of user specifications are not allowed.

You must specify either a list of users or images or both, or use the /UNIQUE qualifier to specify that each user or each image creates a workload definition.

Qualifiers

/ACCOUNTNAMES=(accountname[,. . .])

Defines a list of users by their account field from their UAF (User Authorization File) record that Performance Agent checks when evaluating whether the records of a given process belong to a workload. Wildcards are allowed; account names are limited to eight characters to ensure a match because the UAF account name field is eight characters.

/IMAGENAMES=(imagename[,. . .])

Defines a list of image names that Performance Agent checks when evaluating whether the records of a given process belong to a workload. The image name must be 39 or fewer characters and must not include the device directory or file type. Wildcard characters are allowed.

/INCLUDE=(process_type=[NO] INTERACTIVE, [NO] BATCH, [NO] NETWORK [NO] DETACHED), BASE_PRIORITY=([NO] MINIMUM=M, [NO] MAXIMUM=N)

Allows you to select process activity based on the process status (batch, interactive, network, or detached) or on the process Base Priority, or both.

/MATCH={EITHER | BOTH}

Specifies the criteria a process record must meet before Performance Agent can include the process data in the workload. If you specify EITHER, Performance Agent includes process data in the workload if either the user criteria or image name for the process data is defined in the user criteria or image list. If you specify BOTH, Performance Agent includes process data in the workload if both the user list and image name for the process data match the workload definition.

Performance Agent uses the /MATCH qualifier in conjunction with the /USERNAMES, /UICS, /PIDs, /ACCOUNTNAMES, /PROCESSNAMES, and /IMAGENAMES qualifiers. By default, the value of /MATCH is EITHER is the default.

/PIDs=(nnnnnnnn[,. . .])

Defines a list of PIDs that Performance Agent checks when evaluating whether the records of a given process belong to a workload. If the process record's PID matches one of the PIDs in the list, the process is included in the workload if you specify /MATCH=EITHER. Otherwise, the imagename must also match one of the imagenames for the process to be considered in the workload.

/PROCESSNAMES=(processname[,. . .])

Defines a list of process names that Performance Agent checks when evaluating whether the records of a given process belong to a workload. Wildcards are allowed; process names are limited to 15 characters.

/TRANSACTION_UNITS={IMAGE_TERMINATIONS |TERMINAL_RESPONSES}

Specifies how to count the number of transactions for the workload definition.

If you specify TERMINAL_RESPONSES, the number of transactions is determined by the number of terminal QIO read completions for terminal devices.

/UNIQUE_BY={ACCOUNTNAME | USERNAME | IMAGENAME | PID | PROCESSNAME | UIC_GROUP}

Generates a new workload for the occurrence of each unique image, unique user name, unique process name, unique UIC group, or unique account name. This qualifier cannot be used with the /MATCH, /IMAGENAMES, /UICS, /USERNAMES, /ACCOUNTNAMES, or /PROCESSNAMES qualifiers.

/UICS=(uic[,. . .])

Defines a list of UIC specifications that Performance Agent checks when evaluating whether the records of a given process belongs to a workload. You can use the wildcard in the UIC specification in place of the user number or group number or both.

/USERNAMES=(username[,. . .])

Defines a list of user names that Performance Agent checks when evaluating whether the records of a given process belong to a workload. Wildcards are allowed; usernames are 1 to 12 characters.

Example

```
PSDC-EDIT> ADD/WORKLOAD ENGINEERS /USERNAMES=(SMITH,JONES,BROWN)
```

This command creates and adds a workload definition named ENGINEERS.

DELETE/FAMILY

This command deletes a workload family definition from the parameters file.

Format

```
DELETE/FAMILY name
```

Parameter

name

A workload family name is a 1 to 20-character alphanumeric string that represents a set of workloads.

Description

The DELETE/FAMILY command deletes a workload family definition from the parameters file. You can delete a workload family definition if it is not referenced by history file descriptors.

You cannot delete the workload families MODEL_USERGROUPS and MODEL_TRANSACTIONS because they are needed as default families when you develop a model of the system workload from Performance Management data. You may, however, modify these families, see MODIFY/FAMILY.

Example

```
>PSDC-EDIT> DELETE/FAMILY DEPARTMENTS
```

This command deletes the workload family named DEPARTMENTS.

DELETE/HISTORY

This command deletes a history file descriptor from the parameters file.

Format

```
DELETE/HISTORY name
```

Parameter

name

The history file descriptor name is a 1 to 20-character alphanumeric string that represents an archiving scheme for history files.

Description

The DELETE/HISTORY command deletes a history file descriptor from the parameters file. Performance Agent requires you to delete any history files generated by that definition before you delete the history file descriptor.

1. Use the ADVISE COLLECT SHOW STATUS CPD command to determine the disk and directory location for data and history files.
2. Identify the data and history files associated with the history file descriptor.
3. Delete the files using the DELETE/CONFIRM command.
4. Delete the history file descriptor.

Example

```
$ ADVISE COLLECT SHOW STATUS CPD
$ DIRECTORY $DUA1:PSDC$*_%%-%%%-%%%%_%%-%%%-%%%%.PEAKLOAD;*
$ DELETE/CONFIRM $DUA1:PSDC$*_%%-%%%-%%%%_%%-%%%-%%%%.PEAKLOAD;*
$ ADVISE/EDIT
PSDC-EDIT> DELETE/HISTORY PEAKLOAD
```

This command deletes the history files and the history file descriptor named PEAKLOAD.

DELETE/WORKLOAD

This command deletes a workload definition from the parameters file.

Format

```
DELETE/WORKLOAD name
```

Parameter

name

A workload name is a 1 to 20-character alphanumeric string that represents a workload or grouping of your system's activity.

Description

The DELETE/WORKLOAD command deletes a workload definition from the parameters file. You can delete a workload definition if it is not referenced by a workload family. If a workload family references a workload definition, you must modify the family by removing the reference, and then delete the workload.

Example

```
PSDC_EDIT> DELETE/WORKLOAD ENGINEERS
```

This command deletes the workload definition named ENGINEERS.

EXIT

This command saves changes made during the current editing session, if any, and returns you to the DCL command level. You can also save edits and exit by pressing Ctrl+Z.

Format

```
EXIT
```

MODIFY/FAMILY

This command changes workload names associated with a family in the parameters file.

Format

```
MODIFY/FAMILY name
```

Parameter

name

A workload family definition is a 1 to 20-character alphanumeric string that represents a set of workloads.

Description

The MODIFY/FAMILY command changes workload definitions associated with a family in the parameters file.

The order in which workload definitions are listed in the workload family influences how Performance Agent classifies process data. During workload classification, Performance Agent applies process data to the first workload definition matched. If the workload definitions overlap their selection criteria, then Performance Agent uses the first workload that matches the process data. Therefore, place more explicit workload definitions before more general workload definitions.

Qualifiers

/ADDWORKLOAD=(workload-definition[,. . .])

Augments the list of workload definitions previously included in the family's definition. Any workload definition included in the list must already be defined in the parameters file. The workload definitions are added at the end of the list of workload definitions.

/WORKLOAD=(workload-definition[,. . .])

/NOWORKLOAD=(workload-definition[,. . .])

Replaces the names of workload definition that make up the workload family. If you specify NOWORKLOAD, the specified workload definitions are deleted from the current list of workload definitions. To reorder the workload definitions within a family, you must specify all the workload definitions in the order desired.

Example

```
PSDC-EDIT> MODIFY/FAMILY
MODEL_USERGROUPS/ADDWORKLOAD=(ENGINEERS,MANAGERS,SECRETARIES)
```

This command adds workload definitions to the MODEL_USERGROUPS workload family.

MODIFY/HISTORY

This command changes either the schedules or retention period of a history file.

Format

```
MODIFY/HISTORY name
```

Parameter

name

The history file descriptor name is a 1 to 20-character alphanumeric string that represents an archiving scheme for history files. The MODIFY/HISTORY command changes either the schedules or retention period of a history file. Other history file attributes cannot be modified because the consistency of the history file data would be compromised.

Qualifiers

/DELETE_AFTER=number_of_days

Specifies the retention period for a history file, in days. The default is 180 days.

/HOLIDAY_SCHEDULE=(day=(hour-range)[,. . .]|

NOday[,. . .])

Selects a part of or all of the daily data for holidays to be stored in the history file. When using the /HOLIDAY_SCHEDULE qualifier, use the following parameters:

day

Set as SUNDAY, MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, EVERYDAY, WEEKDAYS, or WEEKENDS.

hour-range

Is specified as m-n, where m and n are numbers from 0 to 24, and m is less than n. More than one hour range can be specified for a given day. Hour-range cannot be used with the NOday option.

/SCHEDULE=(day=(hour-range)[,. . .]|NOday[,. . .])

Selects a part of or all of the daily data you store in the history file. When using the /SCHEDULE qualifier, use the following parameters:

day

Set as SUNDAY, MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, EVERYDAY, WEEKDAYS, or WEEKENDS.

hour-range

Is specified as m-n, where m and n are numbers from 0 to 24, and m is less than n. More than one hour range can be specified for a given day. Hour-range cannot be used with the NOday option.

Example

```
PSDC-EDIT> MODIFY/HISTORY FULL_USER/DELETE_AFTER=99999
```

This command modifies the history descriptor FULL_USER to have a new retention period of 99999 days (infinite retention).

MODIFY/WORKLOAD

This command changes a workload definition in the parameters file.

Format

```
MODIFY/WORKLOAD name
```

Parameter

Name

A workload name is a 1 to 20-character alphanumeric string that represents a workload, or grouping of your system's activity.

Description

The MODIFY/WORKLOAD command changes a workload definition in the parameters file. A workload definition is a list of image names or users that govern how Performance Agent groups process data into workloads.

Some qualifiers are mutually exclusive. For example, you cannot use /USERNAMES with /ADDUSERNAMES or /IMAGENAMES with /ADDIMAGENAMES.

Qualifiers

/ACCOUNTNAMES=(accountname[,. . .])

Replaces the existing list of account names for this workload definition.

/NOACCOUNTNAMES=(accountname[,. . .])

Removes the specified account names for this workload definition.

/ADDACCOUNTNAMES=(accountname[,. . .])

Adds account names to the existing workload definition.

/ADDIMAGENAMES=(imagename[,. . .])

Adds image names to the existing workload definition.

/ADDPIDS=(nnnnnnnn[,. . .])

Adds the list of PIDs to the existing list of PIDs in the workload definition. The /PIDS and /ADDPIDS qualifiers are mutually exclusive.

/ADDPROCESSNAMES=(processname[,. . .])

Adds process names to the existing workload definition.

/ADDUSERNAMES=(username[,. . .])

Adds user names to the existing workload definition.

/ADD_UICS=(uic[,. . .])

Adds UICs to the UIC list for the workload definition.

/IMAGENAMES=(imagename[,. . .])

Replaces the existing list of image names for this workload definition.

/NOIMAGENAMES[=(imagename[,. . .])]

Removes the specified image names or all if none are specified.

/INCLUDE=(process_type=[NO] INTERACTIVE, [NO] BATCH, [NO] NETWORK [NO] DETACHED), BASE_PRIORITY=([NO] MINIMUM=M, [NO] MAXIMUM=N)

Allows you to select process activity based on the process status (batch, interactive, network, or detached) or on the process Base Priority, or both.

/MATCH={EITHER | BOTH}

Indicates that a process record must match the user name or the image name or both to be considered in the workload.

/PIDS=(nnnnnnnn[,. . .])

Replaces the existing list of PIDs in the workload definition. The /PIDS and /ADDPIDS qualifiers are mutually exclusive.

/NOPIDS=(nnnnnnnn[,. . .])

Removes the specified PIDs from the existing list of PIDs in the workload definition. If specified without any PIDs, all PIDs are removed from the list.

/PROCESSNAMES=(processname[,. . .])

Replaces the existing list of process names for this workload definition.

/NOPROCESSNAMES=(processname[,. . .])

Removes the specified process names for this workload definitions.

/TRANSACTION_UNITS={IMAGE_TERMINATIONS |TERMINAL_RESPONSES}

Replaces the way you count the number of transactions for the workload definition.

/UNIQUE_BY={USERNAME |IMAGENAME |PROCESSNAME |UIC_GROUP|ACCOUNTNAME}

Generates a new workload for the occurrence of each unique image, unique user, or unique image/user combination. This qualifier cannot be used with the /MATCH, /IMAGENAMES, and /USERNAMES qualifiers.

/USERNAMES=(username[,. . .])

Replaces the existing list of user names for this workload definition.

/NOUSERNAMES[=(username[,. . .])]

Removes the specified users, or all if none are specified.

Examples

```
PSDC-EDIT> MODIFY/WORKLOAD DEPT_1/USERNAMES=(LEWIS,PAGE,FRANKS,HODGES)
```

The previous command modifies the workload definition DEPT_1 to have a new list of user names, superseding the old list. This workload definition can be used in a workload family to describe a portion of the system's workload, namely the activity in Department 1, including the activity of all processes with the usernames matching any of those listed: LEWIS, PAGE, FRANKS, HODGES.

```
PSDC-EDIT> MODIFY/WORKLOAD APP_X/NOIMAGENAMES=(X$PROG2A)
```

The previous command modifies the workload definition APP_X, removing the image name X$PROG2A from the list of image names. This workload definition can be used in a workload family to describe a portion of the system's workload, namely the activity of application X, which includes the activity of all processes with the image names matching any of those identified in the workload definition.

QUIT

This command exits the Parameter Editor, discarding changes.

Format

```
QUIT
```

Description

Exits the Parameter Editor and does not save any changes made during the current editing session.

SET AUTO_AUGMENT

This command establishes automatic augmentation of an auxiliary knowledge base for performance analysis and archiving.

Format

```
SET AUTO_AUGMENT[=file-spec]
SET NOAUTO_AUGMENT
```

Description

The SET AUTO_AUGMENT command establishes automatic augmentation of an auxiliary knowledge base. The initial setting is no automatic augmentation of the Performance Manager knowledge base.

file-spec is the name of the compiled auxiliary rules file.

When you generate an analysis report or archive the daily data, Performance Manager uses the auxiliary rules in addition to the factory rules if automatic augmentation is enabled.

The file-spec is "remembered" and need not be supplied subsequently when setting auto augmentation on.

Example

```
PSDC-EDIT> SET AUTO_AUGMENT=PSDC$DATABASE:SITE_RULES.KB
```

This command establishes PSDC$DATABASE:SITE_RULES.KB to augment the factory rules during subsequent analysis reporting.

SET VERSION_LIMIT

This command establishes the setting of the file version limit on the Performance Management parameters file and history files.

Format

```
SET VERSION_LIMIT=number
```

Description

The SET VERSION_LIMIT command establishes the setting of the file version limit on the Performance Management parameters file and history files. The version limit for these files defines the number of file generations that reside in the PSDC$DATABASE area. The initial setting is zero, indicating no version limits.

Note: Setting a version limit is particularly useful in reducing the disk space requirements for history files that have the PERIODICITY attribute. Each daily data file archived into a history file results in a new version of the history file, which supersedes the former. The version limit restricts the number of simultaneous history files of this type and reduces disk space requirements.

Example

```
PSDC-EDIT> SET VERSION_LIMIT=5
```

This command establishes a limit of five file versions of the Performance Management parameters file and history files.

SHOW/ALL

This command displays all parameters in the parameters file.

Format

```
SHOW/ALL
```

Description

The SHOW/ALL command displays all parameters in the parameters file relating to workloads, workload families, history file descriptors, auxiliary knowledge base (if any), and version limit.

Qualifiers

/OUTPUT=filename

Specifies that the desired information is to be written to an output file. If you do not enter the qualifier, or if you enter /OUTPUT without a file specification, the output is sent to the current process default output stream or device, identified by the logical name SYS$OUTPUT.

Example

```
PSDC-EDIT> SHOW/ALL
     Workload                Selection Criteria
-------------------  ------------------------------------
SYSTEM_USER
                             Include Process Type:Interactive, Batch, Network, Detached
                   Base Priority Minimum: 0,   Maximum: 31
                   Transaction Unit based on: Image Termination
          Usernames: SYSTEM
OPERATOR
                             Include Process Type:Interactive, Batch, Network, Detached
                   Base Priority Minimum: 0,   Maximum: 31
                   Transaction Unit based on: Image Termination
          Usernames: OPERATOR
DECNET
                             Include Process Type:Interactive, Batch, Network, Detached
                   Base Priority Minimum: 0,   Maximum: 31
                   Transaction Unit based on: Image Termination
          Usernames: DECNET
SYSMAN
                             Include Process Type:Interactive, Batch, Network, Detached
                   Base Priority Minimum: 0,   Maximum: 31
                   Transaction Unit based on: Image Termination
          Imagenames:APLIC, ARRAY, BACKUP, BUTTON, CALC$MAIN, CALNOTICE, CDU,
                    CLEAR, CLR, CMS, CONFIGURE, CSP, DBMMON, DIRFMT, DQS$SMB,
                    DTM$FILTER, DVI2LN3, EPC$REGIS, ERRFMT, EVL, HISTORY,
                    HOSTCHECK, INSTALL, JOBCTL, JBC$JOB_CONTROL, MONITOR,
                    NOTICE, OPCOM, PAVN, PLOT, PROCNAM, PROTS, PSAC$MAIN,
                    PSCP$MAIN, PSDC$COLLECT_SYSTEMPC, PSDC$DCOMM, PSDC$DC_V5,
                    PSDC$DECWINDOWS, PSDC$DSKANL, PSDC$EDIT, PSDC$FILES,
                    PSDC$FUNCTIONS, PSDC$MOTIF, PSDC$REPORT_SYSTEMPC,
                    PSDC$RT_V4, PSDC$WINDOW, PSPA$ADVISOR, PSPA$BLDKB,
                    PSPA$DECWINDOWS, PSPA$GRAPH, PSPA$MOTIF,
                    PSPA$RT_CHARCELL, PSPA$RT_MOTIF, QUEMAN, SCHED,
                    SETRIGHTS, STARTUP, SYSGEN, VAXSIM, WHAT, WHYBOOT
COMPILES
                             Include Process Type:Interactive, Batch, Network, Detached
                   Base Priority Minimum: 0,   Maximum: 31
                   Transaction Unit based on: Image Termination
          Imagenames:BASIC, BLISS32, FORTRAN, LINK, MACRO32, PASCAL, VAXC
UTILITIES
                             Include Process Type:Interactive, Batch, Network, Detached
                   Base Priority Minimum: 0,   Maximum: 31
                   Transaction Unit based on: Image Termination
          Imagenames:COPY, CREATE, CREATEFDL, DELETE, DIFF, DIRECTORY, DTM,
                    ENOTES, LNGSPLCOR, LOGINOUT, LPS$SMB, MAIL, NOTES$MAIN,
                    NOTES$SERVER, PHONE, QUOTE_V0, RECOLOR, RENAME, REPLY,
                    RUNOFF, SEARCH, SET, SETP0, SHOW, SHWCLSTR, SORTMERGE,
                    SSU, SUBMIT, TYPE, VMSHELP, VTXPAD
EDITORS
                             Include Process Type:Interactive, Batch, Network, Detached
                   Base Priority Minimum: 0,   Maximum: 31
                   Transaction Unit based on: Image Termination
          Imagenames:EDT, EMACS, EMACSSHR, LSEDIT, SED, TECO32, TEX, TPU
NETWORK
                             Include Process Type:Interactive, Batch, Network, Detached
                   Base Priority Minimum: 0,   Maximum: 31
                   Transaction Unit based on: Image Termination
          Imagenames:ELF, FAL, FILESERV, LATCP, LATSYM, NCP, NETACP,
                    NETSERVER, NM$DAEMON, NM$QUEMAN, NML, REMACP, RTPAD
EACH_USER    Unique Workload for each Username
                    Include Process Type:Interactive, Batch, Network, Detached
                    Base Priority Minimum: 0,   Maximum: 31
EACH_UIC_WL  Unique Workload for each UIC Group
                    Include Process Type:Interactive, Batch, Network, Detached
                    Base Priority Minimum: 0,   Maximum: 31
```

```
Workload  Family              Workload Member(s
-------------------  ------------------------------------
MODEL_USERGROUPS     SYSTEM_USER, OPERATOR, DECNET
MODEL_TRANSACTIONS   SYSMAN, UTILITIES, EDITORS, COMPILES, NETWORK
EACH_UIC_GROUP       EACH_UIC_WL
EACH_USER            EACH_USER

History Descriptor         History Descriptor Parameters
-------------------           ------------------------------------
MONTHLY_USER       Granularity = Monthly;          (No Periodicity)
                   Interval    = 60 minutes;       Delete_after = 180 days;
                   No modelling data is saved.
                   Usergroup classification families:
                   EACH_USER
                   Normal Schedule
                        Monday       0-24
                        Tuesday      0-24
                        Wednesday 0-24
                        Thursday     0-24
                        Friday          0-24
                        Saturday     0-24
                        Sunday       0-24
                   Holiday Schedule
                        Monday       0-24
                        Tuesday      0-24
                        Wednesday 0-24
                        Thursday     0-24
                        Friday          0-24
                        Saturday     0-24
                        Sunday       0-24
MONTHLY_MODEL      Granularity = Monthly;          (No Periodicity)
Interval    = 60 minutes;       Delete_after = 180 days;
Modelling data is saved.
                   Normal Schedule
                        Monday       10-12, 14-16
                        Tuesday      10-12, 14-16
                        Wednesday 10-12, 14-16
                        Thursday     10-12, 14-16
                        Friday          10-12, 14-16
                        Saturday
                        Sunday
                   Holiday Schedule
                        Monday
                        Tuesday
                        Wednesday
                        Thursday
                        Friday
                        Saturday
                        Sunday
                   Data files accumulated (in CPD paths as *.MONTHLY_MODEL)P
                        SDC$DATABASE:PSDC$YQUEM_01-FEB-2008_28-FEB-2008
                        PSDC$DATABASE:PSDC$YQUEM_01-MAR-2008_31-MAR-2008
                        (C)PSDC$DATABASE:PSDC$GALLO_01-FEB-2008_28-FEB-2008
Performance Agent and Performance Advisor knowledge base auto augmenting is DISABLED.
Auto_augment file spec:
PSDC Parameter File version_limit is 32767
```

SHOW/AUTO_AUGMENT

This command displays the setting for automatic augmenting to the Performance Manager knowledge base.

Format

```
SHOW/AUTO_AUGMENT
```

Description

The SHOW/AUTO_AUGMENT command displays the setting for automatic augmenting to the Performance Manager knowledge base. If you enable AUTO_AUGMENT, the Performance Agent displays the file specification of the auxiliary knowledge base.

Qualifiers

/OUTPUT=filename

Specifies that the desired information is to be written to an output file. If you do not enter the qualifier, or if you enter /OUTPUT without a file specification, the output is sent to the current process default output stream or device, identified by the logical name SYS$OUTPUT.

Example

```
PSDC-EDIT> SHOW/AUTO_AUGMENT
PSDC and PSPA knowledge base auto augmenting is DISABLED.
Auto_augment file spec:
```

The display produced by SHOW/AUTO_AUGMENT command in this example reveals that automatic augmenting to the Performance Manager knowledge base is not enabled.

SHOW/FAMILY

This command displays all information pertaining to a workload family.

Format

```
SHOW/FAMILY
```

Parameter

name

A workload family name is a 1 to 20-character alphanumeric string that contains a set of workload definitions. The wildcard character may be used in place of the name to display all families.

Description

The SHOW/FAMILY command displays the family name and the workload definitions that are members of the specified family or for all families.

Qualifiers

/OUTPUT=filename

Specifies that the desired information is to be written to an output file. If you do not enter the qualifier, or if you enter /OUTPUT without a file specification, the output is sent to the current process default output stream or device, identified by the logical name SYS$OUTPUT.

Example

```
PSDC-EDIT> SHOW/FAMILY *
Workload  Family               Workload Members
 -------------------  ------------------------------------
MODEL_USERGROUPS      SYSTEM_USER, OPERATOR, DECNET
MODEL_TRANSACTIONS    SYSMAN, UTILITIES, EDITORS, COMPILES, NETWORK, OLTP
EACH_USER             EACH_USER
```

SHOW/HISTORY

This command displays all information pertaining to the history file descriptor.

Format

```
SHOW/HISTORY name
```

Parameter

name

The history file descriptor name is a 1 to 20-character alphanumeric string that represents a reduction scheme for history files. You can use the wildcard character in place of the name to display all history file descriptors.

Description

The SHOW/HISTORY command displays information for one or all the history descriptors. The information includes the descriptor name, the archival scheme, the history file retention period, whether modeling data was saved, the history schedules, and any archived files.

Qualifiers

/OUTPUT=filename

Specifies that the desired information is to be written to an output file. If you do not enter the qualifier, or if you enter /OUTPUT without a file specification, the output is sent to the current process default output stream or device, identified by the logical name SYS$OUTPUT.

Example

```
PSDC-EDIT> SHOW/HISTORY *

History Descriptor         History Descriptor Parameters
-------------------  -------------------------------------
MY_HIST         Granular ity = Monthly;         NoPeriodicity
                       Interval    = 60 minutes;       Delete_after = 180 days;
                       Modeling data is saved.
                       Normal Schedule
                              Monday    0-24
                              Tuesday   0-24
                              Wednesday 0-24
                              Thursday  0-24
                              Friday    0-24
                              Saturday  0-24
                              Sunday    0-24
                       Holiday Schedule
                              Monday    0-24
                              Tuesday   0-24
                              Wednesday 0-24
                              Thursday  0-24
                              Friday    0-24
                              Saturday  0-24
                              Sunday    0-24
```

This command displays all history descriptor parameters.

SHOW/VERSION_LIMIT

This command displays the version limit for the Performance Management parameters file and history files.

Format

```
SHOW/VERSION_LIMIT
```

Description

The SHOW/VERSION_LIMIT command displays the number of parameters and history file generations that can be kept as a backup for the most recent file generations.

Qualifiers

/OUTPUT=filename

Specifies that the desired information is to be written to an output file. If you do not enter the qualifier, or if you enter /OUTPUT without a file specification, the output is sent to the current process default output stream or device, identified by the logical name SYS$OUTPUT.

Example

```
PSDC-EDIT> SHOW/VERSION_LIMIT
PSDC parameter file version_limit is 5
PSDC-EDIT>
```

The previous command displays a limit of five versions for the Performance Management parameters file and history files.

SHOW/WORKLOAD

This command displays all information pertaining to the named workload definition.

Format

```
SHOW/WORKLOAD name
```

Parameter

name

A workload definition is a 1 to 20-character alphanumeric string that represents a workload or grouping of your system's activity. The wildcard character can be used in place of the name to display all workload definitions.

Description

The SHOW/WORKLOAD command displays the workload definition, the members (users or image names or both), and the selection criteria for the indicated workload definition.

Qualifiers

/OUTPUT=filename

Specifies that the desired information is to be written to an output file. If you do not enter the qualifier, or if you enter /OUTPUT without a file specification, the output is sent to the current process default output stream or device, identified by the logical name SYS$OUTPUT.

Example

```
PSDC-EDIT> SHOW/WORKLOAD *

Workload                      Selection Criteria
--------------------       ---------------------------------------
SYSTEM_USER
Include Process Type:Interactive, Batch, Network, Detached
Base Priority Minimum: 0,   Maximum: 31
Transaction Unit based on: Image Termination
         Usernames: SYSTEM
OPERATOR
Include Process Type:Interactive, Batch, Network, Detached
Base Priority Minimum: 0,   Maximum: 31
Transaction Unit based on: Image Termination
         Usernames: OPERATOR
DECNET
                              Include Process Type:Interactive, Batch, Network, Detached
Base Priority Minimum: 0,   Maximum: 31
SYSMAN
Include Process Type:Interactive, Batch, Network, Detached
Base Priority Minimum: 0,   Maximum: 31


Transaction Unit based on: Image Termination
Imagenames:         APLIC, ARRAY, BACKUP, BUTTON, CALC$MAIN, CALNOTICE, CDU,
CLEAR, CLR, CMS, CONFIGURE, CSP, DBMMON, DIRFMT, DQS$SMB,
DTM$FILTER, DVI2LN3, EPC$REGIS, ERRFMT, EVL, HISTORY,
HOSTCHECK, INSTALL, JOBCTL, MONITOR, NOTICE, OPCOM, PAVN,
PLOT, PROCNAM, PROTS, PSAC$MAIN, PSCP$MAIN, PSDC$DCOMM,
PSDC$DC_V5, PSDC$DECWINDOWS, PSDC$EDIT, PSDC$FUNCTIONS,
PSDC$WINDOW, PSPA$ADVISOR, PSPA$BLDKB, PSPA$DECWINDOWS,
PSPA$GRAPH, QUEMAN, SCHED, SETRIGHTS, STARTUP, SYSGEN,
VAXSIM, WHAT, WHYBOOT
COMPILES


Include Process Type:Interactive, Batch, Network, Detached
Base Priority Minimum: 0,   Maximum: 31
Transaction Unit based on: Image Termination
Imagenames:         BASIC, BLISS32, FORTRAN, LINK, MACRO32, PASCAL, VAXC
UTILITIES
Include Process Type:Interactive, Batch, Network, Detached
Base Priority Minimum: 0,   Maximum: 31
Transaction Unit based on: Image Termination
Imagenames:         COPY, CREATE, CREATEFDL, DELETE, DIFF, DIRECTORY, DTM,
ENOTES, LNGSPLCOR, LOGINOUT, LPS$SMB, MAIL, NOTES$MAIN,
NOTES$SERVER, PHONE, QUOTE_V0, RECOLOR, RENAME, REPLY,
RUNOFF, SEARCH, SET, SETP0, SHOW, SHWCLSTR, SORTMERGE,
SSU, SUBMIT, TYPE, VMSHELP, VTXPAD
EDITORS


Include Process Type:Interactive, Batch, Network, Detached
Base Priority Minimum: 0,   Maximum: 31
Transaction Unit based on: Image Termination
Imagenames:         EDT, EMACS, EMACSSHR, LSEDIT, SED, TECO32, TEX, TPU
NETWORK
Include Process Type:Interactive, Batch, Network, Detached
Base Priority Minimum: 0,   Maximum: 31
Transaction Unit based on: Image Termination


Imagenames:         ELF, FAL, FILESERV, LATCP, LATSYM, NCP, NETACP,
NETSERVER, NM$DAEMON, NM$QUEMAN, NML, REMACP, RTPAD
EACH_USER                           Unique Workload for each Username
Include Process Type:Interactive, Batch, Network, Detached
Base Priority Minimum: 0,   Maximum: 31
EACH_UIC_WL                         Unique Workload for each UIC Group
Include Process Type:Interactive, Batch, Network, Detached
Base Priority Minimum: 0,   Maximum: 31
```

Chapter 11: Use the DECwindows Motif Interface

This chapter describes how to use the DECwindows Motif interface features available with the Performance Agent software. For information on performance analysis features, see the *Performance Manager Administrator Guide*.

This chapter describes how to do the following procedures:

- Start the DECwindows Motif interface
- Customize Collection Definitions
- Produce dumps of collected data
- Edit parameters
- Exit from the DECwindows Motif interface

This chapter also includes the following topics:

- Starting the DECwindows Motif Interface
- Controlling the Interface
- Selecting Data for Viewing Dump Reports
- Displaying a Dump Report
- Customizing
- Viewing the Main Window

This section contains the following topics:

Start the DECwindows Motif Interface

You can use the windowing interface even if the Performance Agent is not installed or running on your workstation. The windowing interface can be directed to your workstation by setting host from a workstation to a node on which the Performance Agent is installed and by issuing this command:

```
$ SET DISPLAY/CREATE/NODE=MYNODE
```

To start the CA DECwindows Motif interface, enter the following command:

```
$ ADVISE/INTERFACE=MOTIF
```

The window interface displays its main window, from which you select the activity you want to perform, as shown in the following main window example:

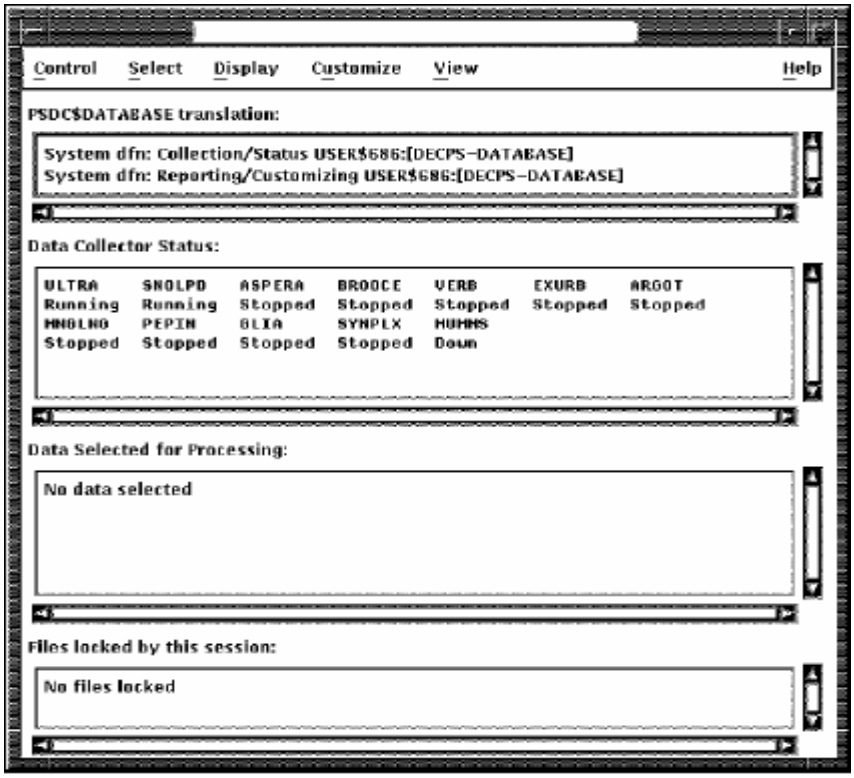

This Main Window for the Performance Agent software allows you to do the following activities:

- Control the CA interface
- Display dumps of collected data

- Customize data collection, PSDC$DATABASE definition, and the parameters file
- View or remove specified main window sections

The Main Window displays the following information:

PSDC$DATABASE translation

Displays the following directory information:

- Collection/Status—Displays the system-wide definition of the CA database. This directory contains the data files produced by the local data collection process and the schedule file for controlling collection.
- Reporting/Customizing—Displays the process definition of the CA database, if it exists. If not, the job definition, group definition, or system definition is displayed.

If you want to reference CA data stored in a directory other than the system directory, you can specify an alternate directory. This alternate directory might contain CA data files from another cluster, or archived data, or both.

Performance Agent Status

The status of the data collection process can be one of the following:

- Running
- Stopped
- Down
- Waiting due to schedule
- Waiting for disk space
- Database path not available

For more information on Performance Agent status, see the chapter Performance Agent Commands (see page 119).

Data Selected for Processing

Displays the start and end time of the collection period, the processing options chosen, and the nodes selected for analysis. Until you select data, the message "No data selected" is displayed.

Files locked by this session

Displays one of the following messages:

- Schedule file
- Parameters file
- No files locked

The Schedule file is locked when you customize the Performance Agent and is unlocked when you complete your changes.

The Parameters file is locked when you customize parameters. When you complete your changes a message box appears asking you if you want to release the file. For more information, see Customizing Parameters in this chapter .

Control the Motif Interface

Pull down the Control menu and release on the menu item you want. The Control menu allows you to do the following actions:

- Read the Parameters file
- Write the Parameters file
- Load binary graph data
- Save binary graph data
- Quit the DECwindows Motif session

Read the Parameter File

When you release on the Read Parameter File . . . menu item, you are prompted with the following message:

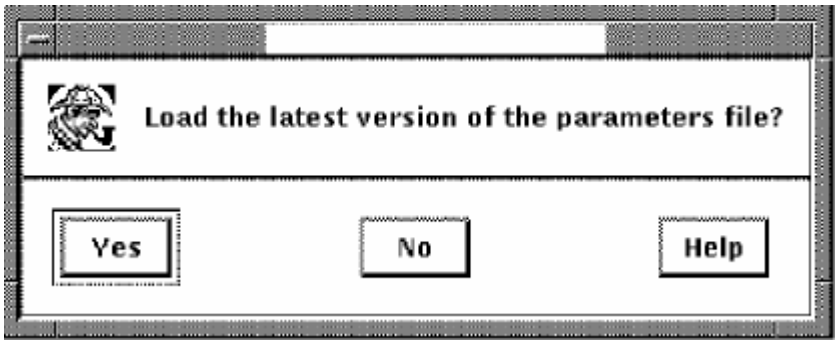

If you respond by clicking Yes, any changes to the parameter file that you have not saved are lost.

Write the Parameter File

Release on the Write Parameter File . . . menu item if you want to save changes you made to the file. If you have not made changes, this option is desensitized.

Load binary graph data

Release on the Load Binary Graph Data menu item to load a selection of graph data that has been previously saved to a specified file. If you have already analyzed a period of time during your current windowing session, this selection replaces those display options with the graph options of the binary data.

Save binary graph data

Release on the Save Binary Graph Data menu item to save the currently selected graph data to a specified file in a binary format. All graph data points are saved, and they can be reloaded using the load option. An informational dialog box confirms the save operation.

Quit the Session

To end a DECwindows Motif session, click the Quit menu item. If you have any outstanding changes to the Parameters file, you are asked if you want to save them. Clicking Yes saves your changes. Clicking No does not. The interface proceeds to end of job.

Select Data for Viewing Dump Reports

Choose the Specific Data . . . menu item from the Select menu to specify data to be included in a dump report. The DECwindows Motif interface displays the Data Selection dialog box.

The following example shows the Data Selection window:

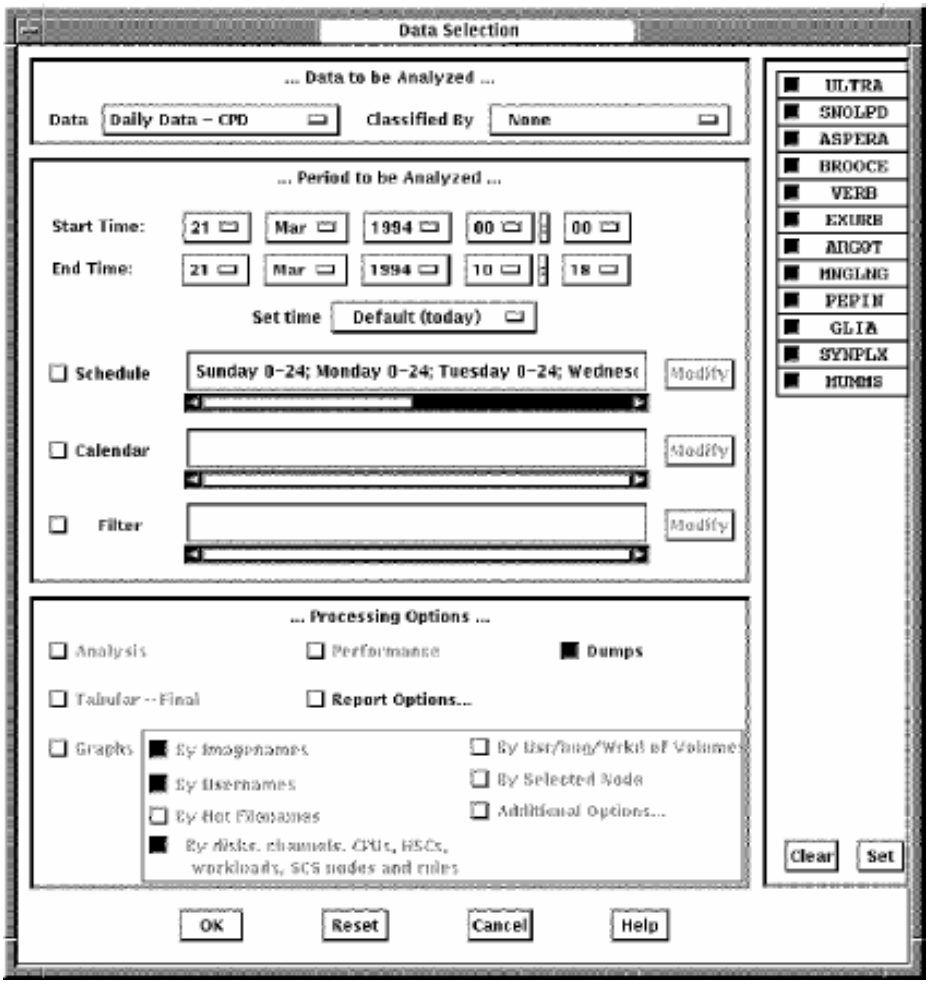

The Data Selection dialog box allows you to do the following actions:

- Choose type and classification of data
- Set the beginning and ending date and time of the reporting period
- Choose processing and report options
- Choose nodes for which data is to be reported

You can enter any of the following parameters in this dialog box:

Data to be Analyzed, Data—Specifies the source of data to be analyzed. Press MB1 on the Data selection box and an option menu appears. The menu lists history file descriptors and daily data. Release on the item you want. The option menu disappears. The menu item you chose is now the current data source. The default data source is daily data.

Note: Changing the source of data can change how it is classified. If you choose a history file descriptor, the Classify By options are changed to reflect those specified by the descriptor's definition.

If you choose daily data, Classify By options is ignored for viewing dump reports. For this reason the source of the data should always be chosen ahead of the classification of the data.

Classify By

This option is relevant to displaying history file data that is based on specific workload families. Specifies how the Performance Agent is to classify data in the dump report. Press MB1 on the Classify By selection box and an option menu appears. The menu lists workload families. Release on the family name you want. The default option is None.

Period to be Analyzed

Hold MB1 down on any of the date and time fields to see all available choices. Release MB1 on the desired value. Press and hold MB2 to advance through the possible values. Press and hold MB3 to move back through the values.

Set time

Press MB1 on the Set time option button and an option menu appears.

You can enter any of the following parameters:

- Default (today)-Specifies midnight to now.
- Yesterday-Specifies yesterday, from midnight to midnight.
- Most recent hour-Specifies the last 60 minutes.
- First hour of today-Specifies the time period of 00:00 to 01:00.
- Advance by a day-Increment the beginning and end dates by one.

- Backup by a day-Decrement beginning and end dates by one.
- Specify text . . .-Release on the Specify text . . . menu item to specify the beginning and end day and time from the keyboard. Performance Agent displays the Time Selection Box, as shown in the following example:

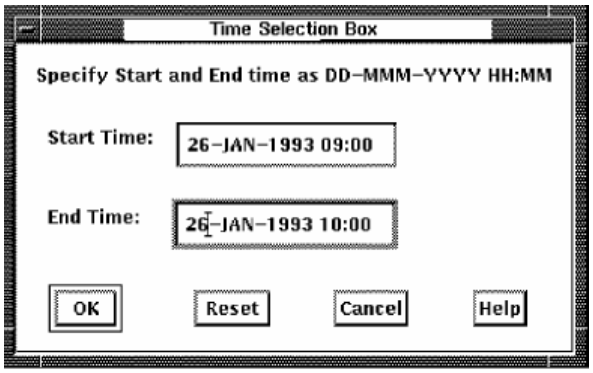

The day and time can be typed in accordance to the format shown in the window.

Clicking the OK button applies the start and end times and removes the dialog box from the screen. The Reset button restores the start and end times to those displayed in the Data Selection box.

Clicking the Cancel button removes the dialog box without changing the time fields currently displayed in the Data Selection box.

Schedule—Use the Schedule Selection dialog box to modify the current selection schedule. This schedule is applied at data selection time to select specific hours of the week for analysis. Click the Schedule button and then click Modify to access the Schedule selection box, as shown in the following example:

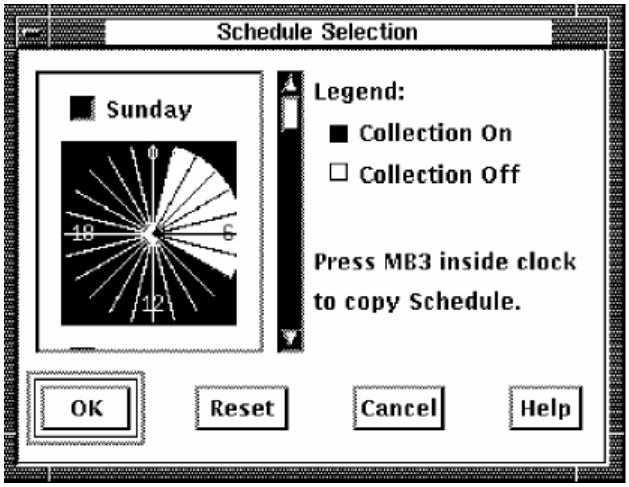

- Use the scroll bar at the right of the clock to move to each day of the week. The 24 hour clock for each day is displayed beneath the day.

- To turn data collection on or off for the entire day, click the toggle button above the clock.
- To select data collection for specified hours during a day, click MB1 on each hour section. As indicated by the Legend, light and dark clock sections indicate when data is collected.

  You can also select time periods by pressing MB1 and dragging the pointer through the hours on

To copy a schedule from one day to another

1. Press and hold MB3 inside the clock face of the schedule you wish to copy. A pop-up menu is displayed. Release on Cut.
2. Scroll to the day you wish to change and press MB3 inside the clock face. Release on the appropriate Paste menu item. The clock displays the new schedule.
3. Click OK after you complete your selections.

   The Schedule Selection box is removed and the new schedule is displayed.
4. Click Reset to erase your changes and revert to your previous selections.
5. Click Cancel to remove the Schedule Selection box.

Calendar

You can specify a monthly schedule for selecting data for analysis. To access the Dates File Selection box, click the Calendar button and then the Modify button. To enter the file of date ranges, click the Load Date ranges box and enter the file, as shown in the following example:

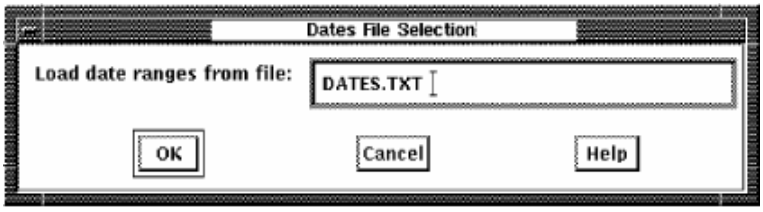

For example, if you wanted to view dump reports for all the Mondays in January but wanted to substitute a Tuesday in place of a Monday that was a holiday, you could do this using a dates file with the following format:

```
04-JAN-2008 10:00, 04-JAN-2008 10:08
11-JAN-2008 10:00, 11-JAN-2008 10:08
19-JAN-2008 10:00, 19-JAN-2008 10:08
25-JAN-2008 10:00, 25-JAN-2008 10:08
```

Using the preceding list, dump reports are created for the times specified for the four days listed, three Mondays and Tuesday the 19th.

After entering the file, click OK. The box is removed and the Start Time and End Time displays are updated and cannot be changed.

When you no longer wish to use the dates file, click the Calendar button. The Start Time and End Time buttons are reactivated.

Filter

This option is available with the Performance Manager. See the Performance Manager Administrator Guide for information about using a filter to select a subset of data for reporting and graphing.

Processing Options

Although many options are available (Analysis, Performance, and Graphs), with Performance Agent only the Dumps option is allowed. Dump specifies that the unanalyzed data are made available for your inspection.

For more information about other types of reports, see the CA Performance Management for OpenVMS Performance Manager Administrator Guide.

Node Control and Toggle Buttons

By default, reporting for all nodes in a collection definition is enabled. Click a node's toggle button to include or exclude the node from the Performance Agent processing. Clear and Set buttons are available for adjusting all toggles.

Control Buttons

The OK button applies your selection, removes the data selection box and activates a Work in Progress dialog box, which includes a real time display of the progress of the analysis procedure. In the case of a Dumps-only selection, no preprocessing of the data is done. The Work in Process immediately shows the procedure as 100% complete.

The Cancel button closes the Data Selection dialog box and resets all widgets.

The Reset button sets all widgets back to their settings of either the startup defaults, or the last OK'd selection.

Display a Dump Report

To initiate a display of the dump reports through the DECwindows Motif interface

- Pull down the Display menu and select a node name that appears in the list provided from the Dump [-> submenu.

  Releasing on a node name opens the Dump window. By default, the configuration record for the first interval record is displayed. From this point the available options are the following:

  – Selecting a different record type for display

  – Examining the same record type for the next interval record

  – Closing the dump window

The example shows a dump report:

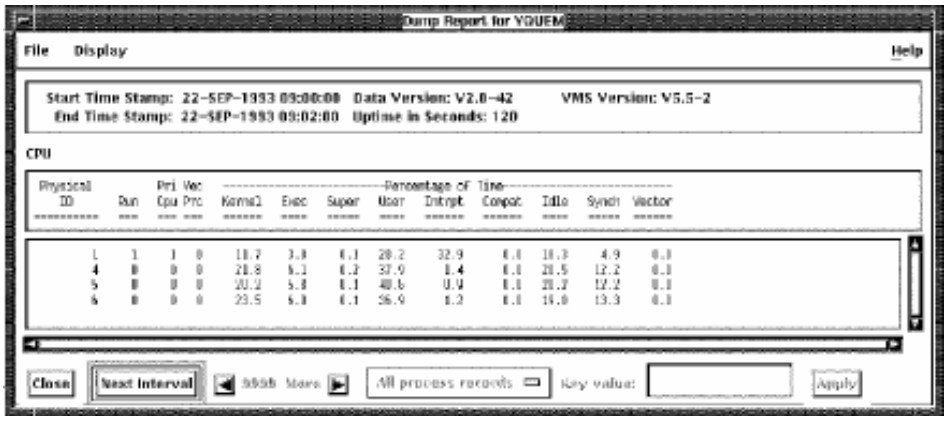

To select a different record type for display

- Press MB1 on the Display pull-down menu in the Dump Report window.

  A list of the possible record types appears in the menu. A set toggle button appears to the left of any currently viewed records. Releasing on a set toggle button removes that record type from the Dump window, while releasing on an unselected record type causes that record to be added to the Dump window. The space required to view the different record types varies considerably. It is left to the user to configure the window as best fits on the screen.

To view the data for the next interval for the selected record types

- Click the Next Interval button.

  The ability to view the previous interval is not provided; only forward motion through the file is available. This is because the Performance Agent files are accessed sequentially only. You can restart to the beginning of the interval by closing the Dump window and reselecting the same node.

To close the Dump window

- Click the Close button at the bottom of the Dump window.
  OR
  Release on the Close item within the File pull down menu.

Filter Process Data

When process records are being viewed, additional filtering options are available through the widgets at the bottom of the window. Because there may be hundreds of process records for a single interval record, a mechanism is provided for looking at subsets of the data. By default, four process subrecords are displayed. The *xxx* More label indicates how many more process records exist for this interval and the arrow buttons on either side allow you to advance or back up through them.

To select a subset of process records, a key type must be selected, then a key value, and then the Apply button used. To select a key type, press MB1 down on the option menu, which is set to "All process records" and release on the desired key type. Then type in the key value. When you click the Apply button, the window is updated to display only those process records that match the key settings. The *xxx* More label are changed to properly reflect the number of process records that match the selected key.

If process records are not currently viewed in the Dump window, the widgets associated with selecting subsets are disabled.

Customize

Pull down the Customize menu and release on the item you want.

The Customize menu allows you to do the following actions:

- View and change Performance Agent settings
- Specify or redefine the PSDC$DATABASE definition
- View and change CA parameters

Customize Data Collection

To view Performance Agent settings, pull down the Customize menu and release on the Data Collection... menu item. CA DECwindows Motif displays the Performance Agent Collection Definition dialog box. This is shown in the following figure.

The primary data collection process collects performance data according to parameters in the CPD collection definition within the schedule file. If no other collection definitions exist, the CPD parameters are displayed in the Collection Definitions box, as shown in the following example:

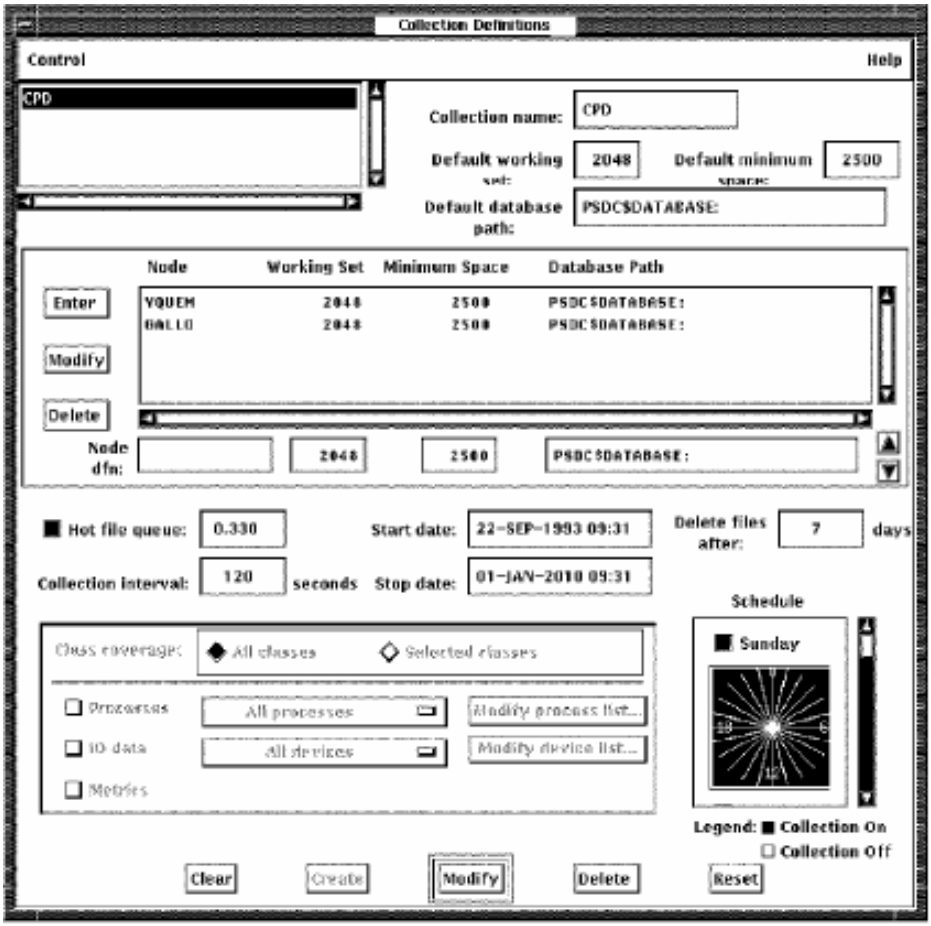

The Collection Definition dialog box allows you to do the following actions:

- Modify, create, or remove collection definitions
- Change collection parameters
- Change collection schedules
- Add or remove nodes

Your changes to the Collection Definition file are applied when you pull down the Control menu and release on Save schedule file or Exit. If you change settings and then decide not to use the changes, pull down the Control menu and release on Load schedule file to start over with the latest version.

When you are done with this window, pull down the Control menu and release on Exit.

Before writing out a new schedule file or releasing the lock, a message box appears with one of the following messages:

- (If modifications were made:) Do you want to write out your changes to the schedule file and release the lock?
- (If no modifications were made:) You have locked access to the schedule File, do you want to release it to other users?
  - Press *YES* to release the lock. Your modifications, if any, are written out.
  - Press *NO* to keep the lock. Your modifications, if any, are not written out.

Modify or Create a Collection Definition

The names of the collection definitions are listed in the box at the upper left-hand corner. To select a collection definition, click the collection definition. The definition and current parameters appear. Use the Tab key to move between parameter entries or point to the entry and click MB1.

Click the buttons at the bottom of the window to do the following actions:

- Clear—Removes all the entries from display and restores all default values
- Create—Adds the collection definition you created to the list

  The software displays an error message if it finds an error in your entries
- Modify—Updates the collection definition with your changes
- Delete—Removes the collection definition from the list and restores all default values
- Reset—Returns all the parameters on display to their original values

Change Collection Definition Parameters

You can enter any of the following parameters:

Collection Definition List

The box in the upper left-hand corner contains the names of the current collection definitions. When you click an entry, the parameters for that collection definition appear.

Collection name

The Collection name is a text field that shows the name of the definition currently displayed. When creating a new definition, use this field to enter the name. Names can be up to 20 characters in length.

Default working set

Enter the working set quota. This value is a decimal number greater than 1024 that sets both the working set quota and the working set extent if the value is higher than the default values. By default, the working set quota is 2048 and the working set extent is 20K. You can override this default by specifying values for individual nodes.

Default minimum space

Enter the minimum number of blocks of free disk space needed on each database disk for each node in the definition. Data collection suspends recording if there are fewer blocks of free disk space available. You can override this default by specifying values for individual nodes.

Default database path

The default database path for all nodes in the collection definition. You can override this default by specifying values for individual nodes, as shown in the following example:

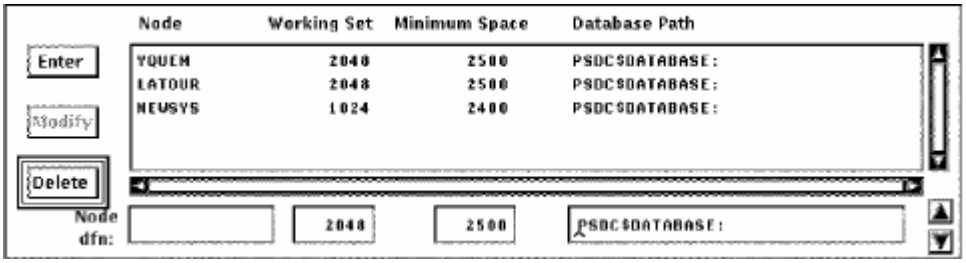

Node Definition Section

This displays the nodes in a collection definition. Default values are displayed. Changes that you make here take effect when you click Create or Modify buttons at the bottom of the window.

To add a node, enter the name and make any modifications you want to the default values shown to the right. Then press Return or click Enter.

To modify a node entry, double click it. The definition appears under the list. Make the changes you wish and then press Return or click Modify to the left of the box.

To delete one or more nodes, click the entries in the list, and then click Delete to the left of the box.

To change the order of the nodes in the list, click the node to be moved, and then the arrows in the right-hand corner. Only one node can be moved at a time.

The changes you make in the node definition section do not modify the collection definition until you click Create or Modify to update the entire collection definition.

Enter—Enter allows a new node to be added to the node list. When you type in a new node name, you can click Enter or press Return to add it to the list.

Modify—Modify takes changes you make to a node definition and puts the changes into the list. Modify is disabled until you select a node from the list by double clicking the node. Changes to the collection definition take effect when you click Modify at the bottom of the window.

Delete-Delete removes any selected, or highlighted, nodes from the list. The nodes are removed from the collection definition when you click Modify at the bottom of the window, as shown in the following example:

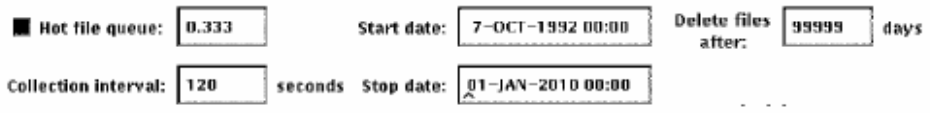

Hot file queue—To enable hot file collection in the CPD collection definition, the toggle button to the left must be on. This allows you to specify the queue length in the text entry box. This pertains only to the CPD collection definition because only the primary Performance Agent collects hot file data.

Enter the minimum average queue size that a disk must have to start collecting the hottest files for that disk. This value is a decimal number less than 100.00 and greater than or equal to 0.00. The default is 0.33.

If the toggle button is off, no hot file data is collected.

If the toggle button is on, you must specify a queue length or accept the default value.

Collection interval—The Collection interval specifies the number in seconds in a sampling interval. At the end of this time period, data is recorded into an interval record.

The CPD collection definition has an interval of 120 seconds which cannot be modified.

You can specify an interval from 1 to 3600 seconds for other definitions.

Delete files after—Enter the number of days that data is to be retained. The default is seven days for the primary Performance Agent. Data files are automatically deleted from the database when they are older than the specified number of days.

For alternate Performance Agents, the default is 99,999 days which enables you to control how many data files you retain.

This value should be coordinated with any data archiving and the amount of free space required for the database area. See Chapter 4 for information on archiving.

Start Date—The date and time on which the collection definition goes into effect. A data collection process hibernates until this date and time.

The format is DD-MMM-YYYY HH:MM

End Date—The date and time on which a data collection process Is to terminate.

The format is DD-MMM-YYYY HH:MM

Changing the Collection Schedule—A 24-hour clock is displayed. A bar on the right side of the clock allows you to scroll to each day of the week. By default, data collection is set ON for each hour of the day, every day of the week.

To set the clock

1. Set or reset the square toggle button above the clock to turn data collection on or off for an entire day.

   To turn off data collection for a specific hour, point to the hour on the clock and click MB1. Holding MB1 down and dragging the pointer around the clock sets data collection to the value of the initial hour setting for a series of hours.

2. Drag the slider on the scroll bar to display the collection schedule for each day of the week or click the up or down stepping arrows.

To duplicate a day's schedule

1. Press and hold MB3 inside the clock. A pop-up menu is displayed. Release on the Cut menu item.
2. Scroll to another day and press MB3 inside the clock. Release on the Paste menu item. The clock displays the copied schedule.

Class Coverage Section

The following figure shows the class selection fields:

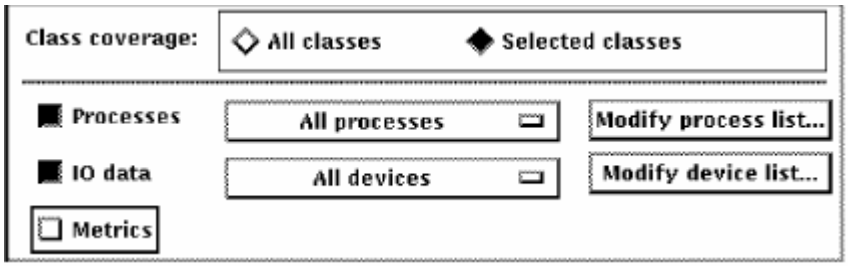

The Class Coverage section allows you to specify which major areas of performance data should be collected. While the CPD cannot be limited, other collection definitions can be limited to only those classes of data needed for special-purpose analysis.

The toggle buttons enable you to choose all of the following or selected classes:

All Classes

Turn on this toggle button to collect all classes of data.

Selected Classes

Turn on this toggle button to choose which classes of data you wish to collect: processes; IO data; or metrics.

Processes

Turn on this toggle button to collect Process data or to collect process data for specified processes.

IO data

Turn on this toggle button to collect IO data or to collect IO data on specified devices.

Metrics

The Metrics field enables you to choose whether or not to collect a summary set of system metrics including such data items as memory utilization, CPU, and I/O parameters.

All processes

When you select the collection of process data, the All Processes option is collected by default. If you modified the process list, the menu is revised to reflect the type of coverage list (include or exclude) and the count of specified processes. Clicking this menu displays the correct list. This menu provides a convenient way to view the coverage list.

Modify process list...

Click this button to specify a list of processes in the Collection Coverage List box. You specify a list of processes to be included in collection or excluded from collection.

All devices

When you select the collection of IO data, the All Devices option is collected by default. If you modified the device list, the menu is revised to reflect the type of coverage list (include or exclude) and the count of specified devices. Clicking this menu displays the correct list. This menu provides a convenient way to view the coverage list.

Modify device list...

Click this button to specify a list of devices in the Collection Coverage List box. You specify a list of devices to be included in collection or excluded from collection.

Coverage lists

For Process and Disk classes, you can collect data for specific lists of items or exclude lists of items, as shown in the following example:

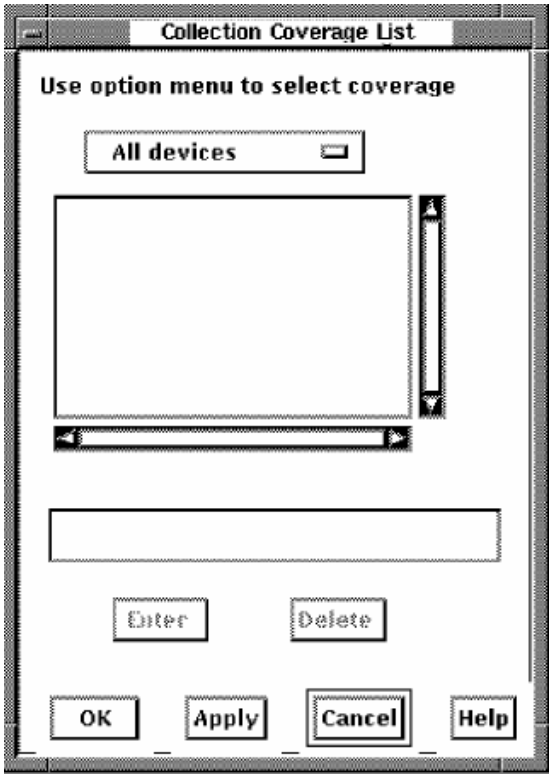

To create a list of processes

1. Push down and hold on the uppermost box in this window.

   The displays All processes or one of the other options.

2. When the menu appears, pull down to the menu item you want and release.

3. Click the text entry field to activate the cursor.

To add a process name

- Type a name and then press Return or click the Enter button.

  The name appears in the list box and be cleared from the text entry field.

To delete a name

1. Click the name in the name list box

2. Click the Delete button.

   The name is removed from the list box. To remove multiple names, click each name and then click the Delete button.

To create a list of devices

1. Push down and hold on the uppermost box in this window.

   This which displays All devices or one of the other options.

2. When the menu appears, pull down to the menu item you want and release.

3. Click the text entry field to activate the cursor.

To add a device name

- Type a name and press Return or click the Enter button.

  The name appears in the list box and be cleared from the text entry field.

To delete a name

- Click the name in the name list box and then click the Delete button.

  The name is removed from the list box. To remove multiple names, click the names and then click the Delete button.

The following figure shows a list of processes for which no data is to be collected:

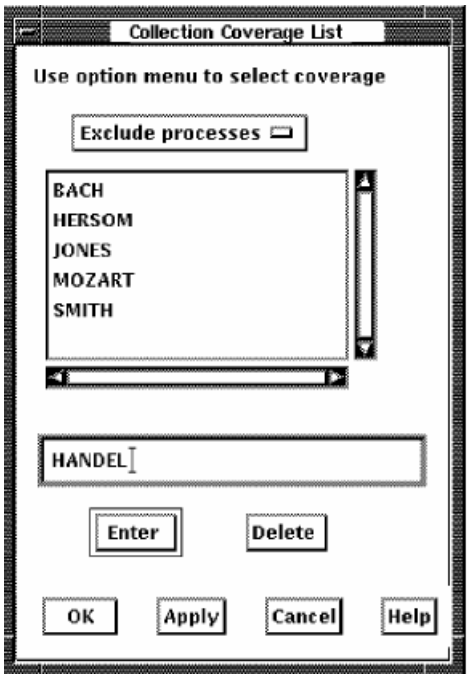

The type of list being managed cannot be converted. For example, an include processes list cannot become an exclude processes list. All list entries must be deleted before the coverage list type can be changed.

Customize PSDC$DATABASE Definition

To specify or redefine a PSDC$DATABASE definition, pull down the Customize menu and release on the PSDC$DATABASE definition . . . menu item. DECwindows Motif displays the Set Database dialog box. This box allows you to redirect editing and review of dump reports to an alternate database area. You can also redirect performance analysis to an alternate area if you have the Performance Manager software, as shown in the following example:

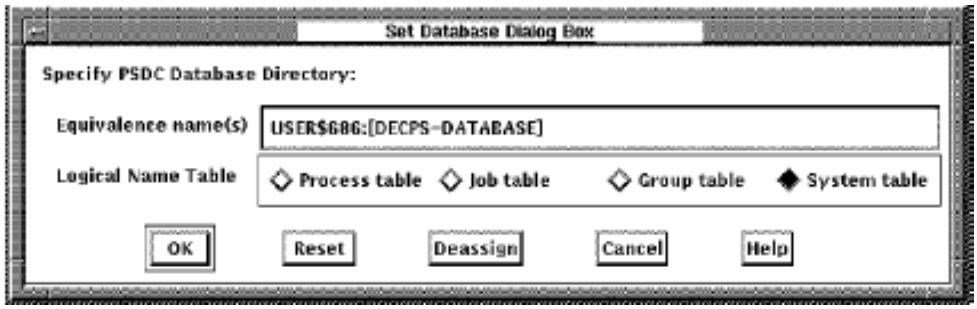

The translation of the lowest level definition is displayed, along with a toggle setting indicating the logical name table in which it is defined.

To remove a definition, click deassign. The dialog box is removed, and the definition at the next highest level goes into effect, possibly reloading a schedule and parameter file.

To create a new definition, click a toggle to specify which logical name table the definition should be placed in and type in the new definition. Click OK when done. The dialog box is removed and any new schedule and parameter files are loaded. There must be at least one definition of PSDC$DATABASE defined for the interface to run.

Customize Parameters

The DECwindows Motif Interface allows you to view and change the CA parameters file. The Parameters file contains workload definitions, family definitions, history file descriptors, and auxiliary knowledge base information.

For more discussions of these definitions, see the chapter Parameter Editor Commands (see page 155).

Start Parameter Editing

To edit parameters

1. Pull down the Customize menu in the Main Window.
2. Click the pointer on the Parameters menu item.
3. Choose a menu item from that submenu.

   Only one user at a time is allowed to edit data in the parameter file PSDC$PARAMS.DAT. The file is locked by anyone using any editor. If the file is locked by another user when you initiate any of the parameter menu's submenus, a message box appears explaining that read-only access to the file is allowed. You are able to view the contents of the file, but any requests to change data are denied. Should the file become unlocked during the course of your DECwindows Motif session, a second message box appears asking you if you would now like update capabilities.
4. Indicate your choice by clicking either Yes or No.

After you gain updated access to the file, it is unavailable to other users. Whenever you close one of the parameter editor's dialog boxes, you are asked if you would now like to release the file, as shown in the following example:

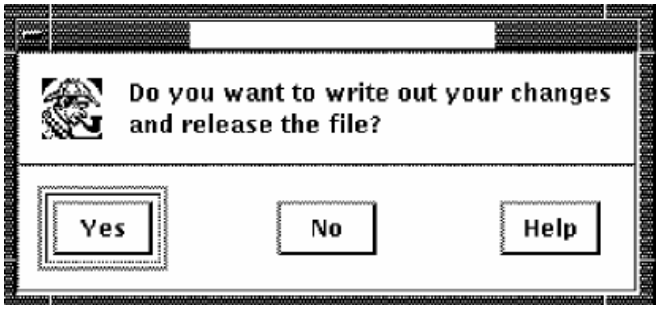

5. When you complete your set of changes to the file, click Yes in response to this request.

   Otherwise, as long as your DECwindows Motif session remains active, even though you may be doing other tasks, such as graphing, you still have the file locked.

   A reminder of the status of the parameter file appears in the main window.

Workload Definitions

To create or modify workload definitions in the parameters file, pull down the parameter submenu and choose the Workload Definitions menu item. CA DECwindows Motif displays the Workload Definitions dialog box shown in the following figure.

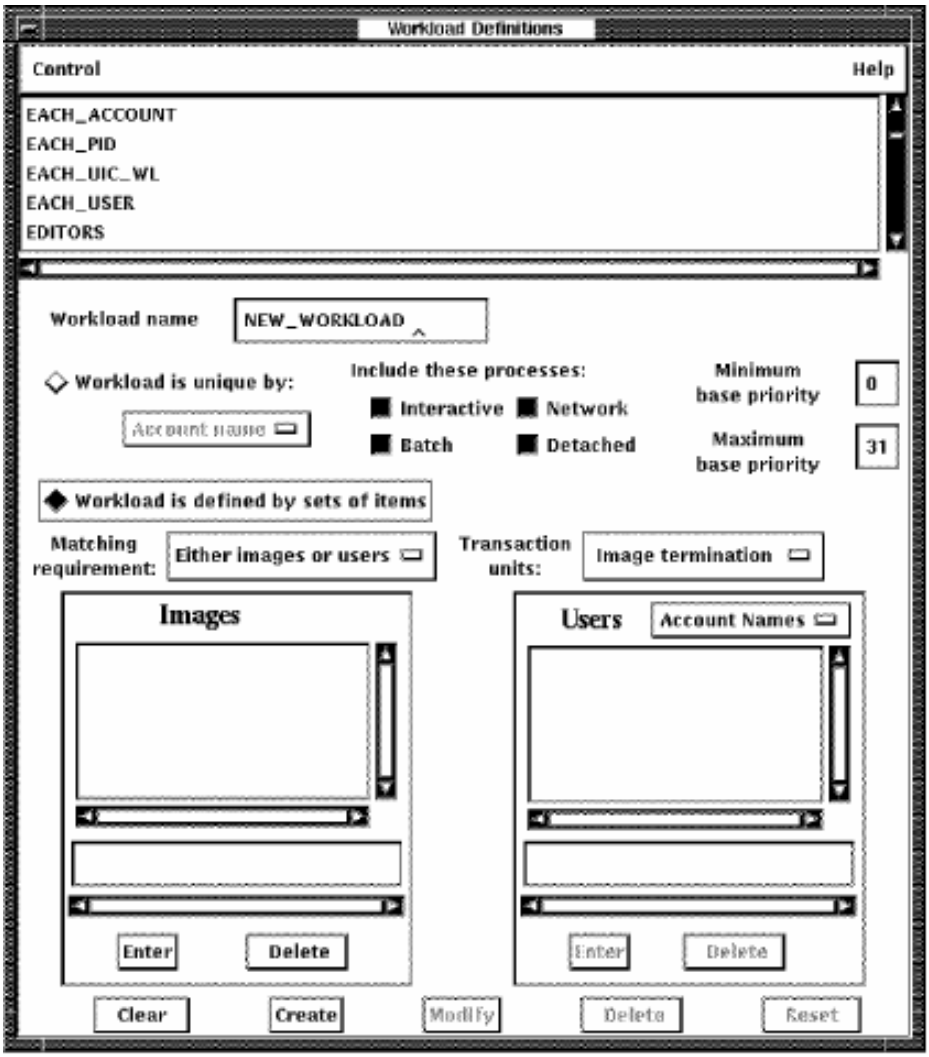

The menu bar contains Control and Help menus. A list of defined workload definitions appears at the top of the dialog box. To close the dialog box, pull down the Control menu and click the Exit menu item.

The Workload Definitions dialog box allows you to do the following actions:

- Create a workload definition
- Delete a workload definition
- Modify a workload definition

Create a Workload Definition

You use the following options to create workload definitions:

Workload name

Enter the name of the workload definition. The workload definition's name is limited to 20 characters.

Workload is unique by

Enables you to specify a category for workload summarization. A workload are defined for each unique element of the category you choose.

Click the Workload is unique by toggle button and then click MB1 on the box beneath this which displays Account name when first accessed. The option menu appears displaying the following items:

- Account name
- Process name
- Image name
- UIC group
- User name
- PID

Click MB1 on the category you want.

Include these processes

Set the toggle buttons to specify the processing modes to be included by the workload definition.

- Interactive
- Batch
- Network
- Detached

Process base priority

Enter the minimum and maximum values for the process Base Priority to be included in the definition. Values can range from 0 to 31 and the minimum value must be less than or equal to the maximum.

Workload is defined by sets of items

Selects the alternative to "unique" by criteria. Allows you to enter lists of user criteria, or images, or both for defining a workload.

Matching requirement

This is an option menu with two entries. Press MB1 on the current setting to view the choices.

Either images or users

Indicates that the Performance Agent matches either the image names or the user criteria of a process record to include the process data in the workload.

Both images and users

Indicates that the Performance Agent must match both the image names and the user criteria of a process record to include the process data in the workload.

Transaction units

This is an option menu with two entries. Press MB1 on the current setting to view the choices. Click either image termination or terminal responses to indicate how response time should be evaluated. This affects the workload frequency when building a model. See the ADD/WORKLOAD (see page 163) command for information about transaction units.

Images

Click the text entry field to activate the text insertion cursor.

To add an Image name

1. Type a name and then press Return or click the Enter button.

   The name appears in the list box and be cleared from the text entry field.

2. If you wish to preserve lowercase characters, enclose image names in double quotes when you enter them.

To delete an Image name

1. Click the name in the image name list box and then click the Delete button.

   The name is removed from the list box.

2. To remove multiple names, click all their names and then click the Delete button.

A list of image names can be provided through a file. Use the at sign (@) as the first character to indicate that the text is to be interpreted as a file name. The default directory is assumed if not supplied, as is a file type of .DAT. The format of the file must be a series of image names separated by white space or commas. Supply only the filename field; do not include the file type.

Image names can contain wildcard characters. Image names can contain up to 39 characters. If you wish to preserve lowercase characters, enclose image names in double quotes when you enter them.

Users

Press MB1 on the Users option menu to view the categories available. Release MB1 on the entry indicating the type of user you want to create. You cannot create a list until you make this selection.

To add a user field

- Type the appropriate string and then press the Return key or click the Enter button.

  The field appears in the list box and is cleared from the text entry field.

To delete a user field

1. Click the entry in the user list box and then click the Delete button.

   The entry is removed from the list box.

2. To remove multiple entries, click all entries to be deleted the fields and then click the Delete button.

A list of user entries can be provided through a file. Use the at sign (@) as the first character to indicate that the text is to be interpreted as a file name. The default directory is assumed if not supplied, as is a file type of .DAT. The format of the file must be a series of user fields separated by white space or commas.

User names can contain wildcard characters. User names longer than 12 characters are truncated to 12 characters to ensure a match because the Performance Agent collects up to 12 user name characters only. Account names can be up to 8 characters in length, and process names up to 15. If you wish to preserve lowercase characters, enclose your entries in double quotes.

User criteria can be specified in terms of UICs, account names, process names, or user names.

A UIC group can be indicated by using an asterisk for the user number, ([200,*]).

Control Buttons

The Clear button removes all dialog box entries.

The Create button adds the workload name to the list and clears the entries.

Delete a Workload Definition

To delete a workload definition

1. Click a workload definition name.

   The dialog boxes are updated to show the current definition field settings.

2. Click the Delete control button to remove the workload definition.

   Performance Agent removes the workload definition name from the workload list and clears the definition fields.

If the definition is not deleted, a message box displays explaining why the request cannot be executed. A failure can occur when a workload family has been defined in terms of this workload. A list of those families displays. Modify the workload family to remove the reference to this workload definition.

Modify Workload Definitions

To modify a workload definition, click a workload definition name. The dialog box displays the current definition values. Modify settings as you wish and click Modify at the bottom of the box.

Workload Family Definitions

To define or modify workload families, pull down the Parameters submenu and choose the Workload Families . . . menu item. Performance Agent displays the Workload Family Definitions dialog box shown in the following figure.

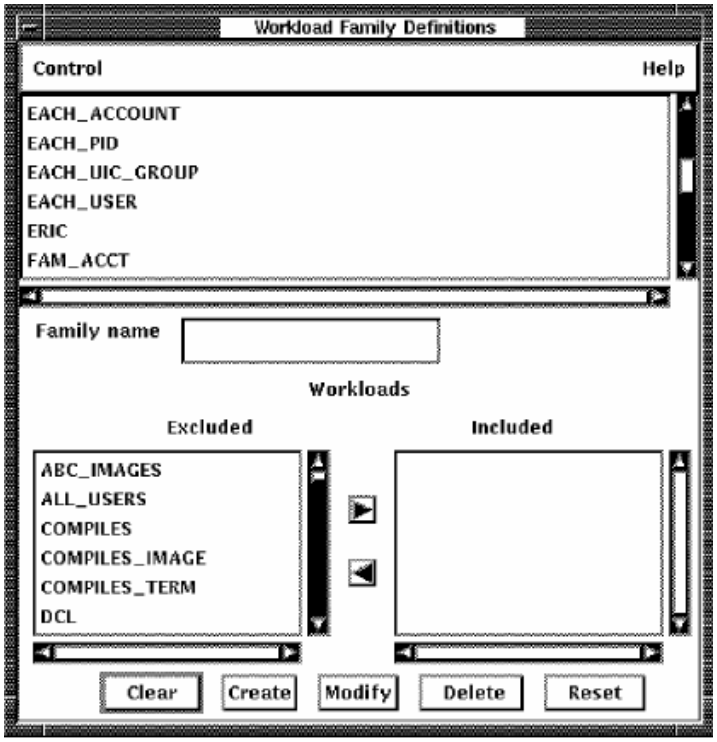

The menu bar contains Control and Help menus. A list of workload family names appears at the top of the dialog box. To close the dialog box, pull down the Control menu, drag the pointer to the Exit menu item and release MB1.

The Workload Family Definitions dialog box allows you to do the following actions:

- Create a workload family
- Delete a workload family
- Modify a workload family

Create a Workload Family

You can enter any of the following:

Family Name

Enter the name of the workload family. The family name is limited to 20 characters.

Workload Specification

The Workloads Excluded list box contains a list of workload definitions. To add a workload definition to a workload family, click the workload name and click the Include transfer button (right arrow).

To indicate a position within this included workload list, click an existing entry in the Workloads Included list box. All new workload definitions are placed ahead of this entry. To deselect a position entry, click it again. All new entries are placed at the end.

To add multiple workload definitions

1. Click all their names contained in the Excluded list box.
2. Click the Include transfer button.

The Workloads Included list box contains a list of workload definitions in the new workload family. To remove an entry from the INCLUDED list box, click the name and click the Exclude transfer button (left arrow). To delete multiple workload definitions, click all their names and then click the Exclude transfer button.

Control Buttons

To remove all entries and cancel the definition, click the Clear button.

To add a family definition to the defined list and clear the entries

- Click the Create button.

Delete a Workload Family

To delete a workload family

1. Click a defined family name.

   The family name and the Workload Included list box display the definition of the selected family.

2. Click the Delete button to remove the family definition

Performance Agent removes the family name from the family list and clears the definition fields. If the family name is not deleted, a message box is displayed to explain why the request could not be executed. A failure occurs if a history file descriptor references this family. The message box specifies the name of any such history descriptors. Because they are the default families for building models, you should not delete the families, MODEL_USERGROUPS, or MODEL_TRANSACTIONS.

Modify a Workload Family

To modify a workload family, click a family name.

The family name and the Workload Included list box display the definition of the selected family.

You can enter any of the following:

Family Name

When you modify a family name, the CA software assumes that you want to create a workload family based on the displayed definitions.

Workload Specifications

The Workloads Excluded list box contains a list of workload definitions that are not part of this family. To add a workload, click the workload definition name and click the Include transfer button (right arrow). To add multiple workload definitions, click all their names contained in the Excluded list box. Then click the Include transfer button.

To indicate a position within this included workload list, click an existing entry in the Included list box. All new workload definitions are placed ahead of this entry. To deselect a position entry, click it again. All new entries are placed at the end.

The Workloads Included list box contains the names of the workload definitions in the family. To remove an entry from the Included list box, click the name and click the Exclude transfer button.

Control Buttons

The Clear button removes all entries.

The Modify button applies the changes and clears the entries.

The Delete button removes the selected family name and its definition.

The Reset button redisplays all entries for the current family definition.

History File Descriptors

To create, modify, or delete history file descriptors in the parameters file, pull down the Parameter submenu and choose the History File Descriptors menu item. Performance Agent displays the History File Descriptors dialog box shown in the following figure:

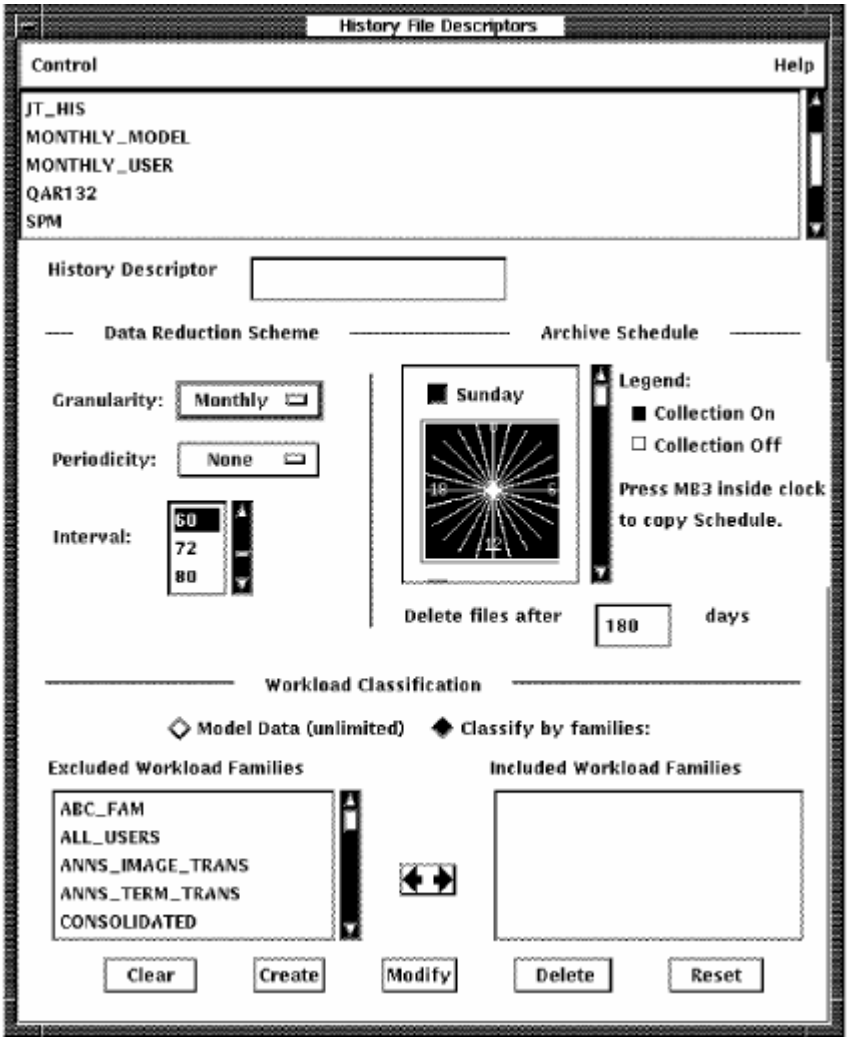

The menu bar contains Control and Help menus. A list of history file descriptors appears at the top of the dialog box. To close the dialog box, pull down the Control menu and click the Exit menu item.

The History File descriptor dialog box allows you to do the following:

- Create a history file descriptor
- Delete a history file descriptor
- Modify a history file descriptor

For a description of history file descriptors, see the chapter Parameter Editor Commands (see page 235).

Create a History File Descriptor

From this screen, you can enter any of the following parameters:

History Descriptor

Enter the name of the history file descriptor. The descriptor name is limited to 20 characters.

Data Reduction Scheme

By default, the granularity value is monthly. Click and hold the pointer on monthly and the granularity pop-up menu displays a list of value options. Release the mouse button when the cursor is on your choice. The pop-up menu disappears and your selection is displayed.

By default, the periodicity value is None. Click and hold the pointer on the current value and the periodicity pop-up menu displays a list of value options. Release the mouse button when the cursor is on your choice. The pop-up menu disappears and your selection is displayed.

By default, the interval value, the time period over which Performance Agent averages daily data records into a single history data record, is 60 minutes. To display valid entries, click the up or down arrow. When your selection is displayed, click the value. Your choice is highlighted.

Archive Schedule

A 24-hour clock is displayed. A bar on the right side of the clock allows you to scroll to each day of the week. By default, archiving is set on for 24 hours a day, seven days a week, including holidays.

To set the clock

1. Set or reset the square toggle button above the clock to turn archiving on or off for an entire day.

   To turn archiving off for a specific hour, point to the hour on the clock and click MB1. Holding down MB1 and dragging the pointer around the clock turns off archiving for a series of contiguous hours.

2. Drag the slider on the scroll bar to display the archiving schedule for each day of the week or click the up or down stepping arrows.

To duplicate a day's schedule

1. Press and hold MB3 inside the clock. A pop-up menu is displayed. Release on the Cut menu item.
2. Scroll to another day and press MB3 inside the clock. Release on the Paste menu item. The clock displays the copied schedule.

Workload Classification

By default, Performance Agent stores process data in the history file summarized by workload families. To save modeling data in the history file, click the model data (unlimited) button. When model data is enabled, no workload families can be selected.

If you choose Classify by Families without specifying the workload families, process data are summarized into four records representing interactive, batch, network, and detached processing. All other process data are lost such as process data based on image name, account name, etc.

Workload Families Excluded

The Workload Families Excluded list box contains a list of workload families. To add a workload family to a history file definition, click a workload family name and click the transfer button.

To add multiple workload families, click all their names contained in the Excluded list box. Then click the transfer button.

Workload Families Included

The Workload Families Included list box contains a list of workload family names. These names specify the workload families that define the new history descriptor. To remove an entry from the Included list box, click the name and click the transfer button.

Control Buttons

The Clear button removes all entries and cancels the definitions.

The Create button adds the family definitions to the defined list and clears the entries.

Delete a History File Descriptor

To delete a history file descriptor

1. Click a defined history file descriptor name.

   The dialog box displays the current definition values.

2. Click the Delete button to remove the history file descriptor.

   Performance Agent removes the descriptor name from the history file descriptor list and clears the definition fields. If you do not delete the history file descriptor, a message box displays explaining why your request was not executed. A failure occurs when history files are created from this definition. A list of these history file names displays.

3. Delete the files and then remove the definition.

Modify a History File Descriptor

To modify a history file descriptor, click a defined history file descriptor name. The dialog box displays the current definition values. You can modify only the archive schedule.

Control Buttons

The Clear button removes all entries and cancels the modifications.

The Modify button applies the changes and clears the entries.

The Delete button removes a selected family definition.

The Reset button redisplays the current family definition.

Parameter Settings

To view parameter settings, pull down the Parameters submenu and choose the Parameter Settings menu item. The software displays the Parameter Settings dialog box, as shown in the following example:

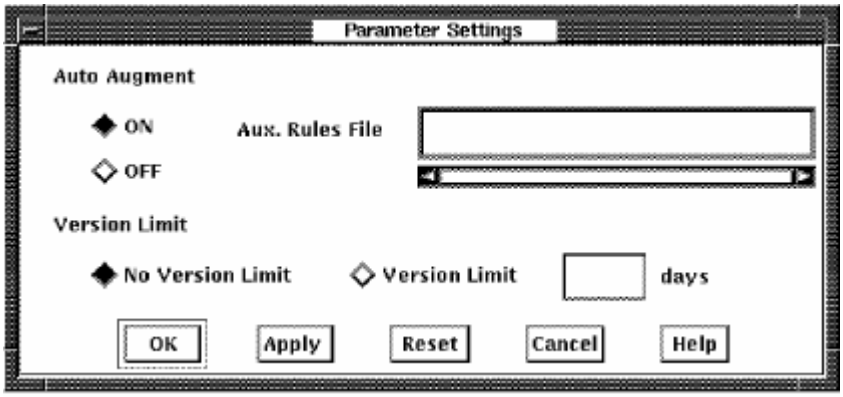

Highlighted buttons indicate the current settings. Your changes in the Parameter Settings box are applied when you click OK or Apply. Click the Reset button to restore the last settings that were applied. Click the Cancel button to close the dialog box without applying any changes.

From the Parameter Settings dialog box, you can set any of the following:

Auto Augment

Click the ON button to establish automatic augmentation of an auxiliary knowledge base for analysis. The initial setting is OFF.

When auto augment is set on, the dialog box allows you to enter the file specification of the compiled auxiliary rules file.

For more information, see SET AUTO AUGMENT in the chapter Parameter Editor Commands (see page 155).

Version Limit

Enter a decimal number to modify the file version limit on the parameters file and history files. The initial setting is 180.

View the Main Window

Pull down the View menu and choose the Main Window sections you want to display.

Depending on your processing mode, not all the sections in the Main window may be relevant or of interest. To allow for smaller windows and the elimination of distracting sections, the View menu contains a series of toggle buttons that can be set or reset to add or remove sections of the Main window. Resetting a toggle button causes a section to disappear; setting it restores the section. Do not remove the File Locks section because you can affect file access by other users.

Appendix A: Error Messages and Recovery Procedures

When the software sends error messages, you can view a list of messages or display more information and advice on how to correct the problem. This appendix describes message format and explains how to view the list of messages and get help with errors.

The following example illustrates the parts of a sample Performance Agent message:

This section contains the following topics:

Parts of a Performance Agent Message

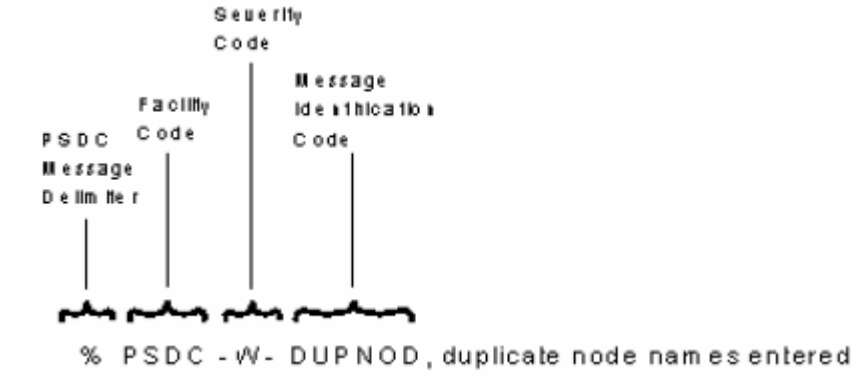

The following table defines the severity codes that are assigned to messages:

Severity Code	Explanation
I	INFORMATIONAL - The Performance Agent software sometimes provides additional information about an action.
W	WARNING - The command may have performed some, but not all, of a requested action; verify the command or output.

Severity Code	Explanation
E	ERROR - The output or program result is incorrect, but the PSDC software will attempt to continue the execution.
S	SEVERE - A fatal error occurred; the PSDC software terminates execution of the request.
F	FATAL - A fatal error occurred; the PSDC software terminates execution of the request.

Get More Information

To display messages from the DCL level (command line), type the following command:

```
$ HELP ADVISE COLLECT ERRORS identification-code
```

Each description includes an explanation and recovery procedure.

To view the list of all messages, type the following command:

```
$ HELP ADVISE COLLECT ERRORS
```

Messages are listed alphabetically by the identification code that precedes each message. The following example shows the explanation provided by the application:

```
$ HELP ADVISE COLLECT ERRORS DUPNOD
ADVISE
COLLECT
Errors
DUPNOD
Duplicate node names entered
Explanation: The same node name was specified more than once with the /NODE_NAMES
qualifier.
User Action: With a /NODE_NAMES qualifier, specify a node name only once.
```

Appendix B: Performance Agent Logical Names

Performance Agent logical names begin with the prefix PSDC$. This appendix lists those names and describes how they are used to control various aspects of the Performance Agent module.

PSDC$DATABASE

A system logical name defined by PSDC$STARTUP.COM if not previously defined. The database area may contain:

Item	Logical Name
Data collection error log	PSDC$DC.LOG
Holidays file	PSDC$HOLIDAYS.TXT
History files	PSDC$nodename_dd-mmm-yyyy_dd-mmm-yyyy.name
History journal files	PSDC$nodename_dd-mmm-yyyy_dd-mmm-yyyy.name_JOU
Daily data files	PSDC$nodename_yyymmmdd.coll-def-name
Parameters file	PSDC$PARAMS.DAT
Schedule file	PSDC$SCHEDULE.DAT

PSDC$EXAMPLES

A system logical name defined by PSDC$STARTUP.COM if not previously defined. This area may contain:

Item	Logical Name
Template command procedure to generate daily reports	PSDC$DAILY.COM
Template command procedure to archive data daily	PSDC$ARCHIVE.COM)
DECwindows setup file	PSDC$RESOURCES

PSDC$hscname_hscunitnumber

The argument is two numbers separated by an underscore (for example, 7_3), which assigns the requestor and port numbers to a shadow set member. This logical provides channel assignments for shadow set members (disks) that are not captured by default. Non-shadow set disks have their channel assignments collected correctly without the need for a logical name; to set up logical names see the procedure following these bullets.

PSDC$NODE

A system logical name of up to six characters like a DECnet node name. It is defined by PSDC$STARTUP.COM if the system is not a member of a cluster system and if it has not previously been defined.

PSDC$IO_PACKET_MAX

This logical name can be used when it is a good idea to preordain the use of non-paged pool; this logical name will bound the amount available to PSDC for I/O and event buffer flushing.

The argument is an integer that specifies the maximum number of non-paged pool allocations, where each allocation is 2008 bytes, for buffering OpenVMS I/O event records. Normally a non-paged pool allocation occurs whenever the system wide device I/O rate causes a deficiency in the number of these buffers available.

PSDC$SPMTIMER

When this logical name is defined and set to 0, extended CPU statistics are not collected. When set to 1, these statistics are collected. If this logical is not defined, all statistics are collected by default.

PSDC$PAGE_FILE_QUOTA

Increase the amount of Page File Quota available to the Performance Agent image.

To set up the logical names

1. Enter the SHOW DEVICE command for a shadow set disk:

```
Device              Device          Error   Volume         Free  Trans Mnt
Name                Status          Count    Label        Blocks Count Cnt
$2$DUA0:   (ERNEST) ShadowSetMember     0 (member of $2$DUS111:)
$2$DUA16:  (ERNEST) ShadowSetMember     0 (member of $2$DUS111:)
$2$DUS111: (ERNEST) Mounted             0 COGNAC          131362   408   3
```

2. On the HSC serving the shadow set (ERNEST), enter the SHOW DISKS command:

```
Unit    Req   Port   Type   State                  Version
0     6      2    RA82   Online, HostAccess      MC - 33, HV - 2
Member of Shadow Unit 111
16    6      3    RA82   Online, HostAccess      MC - 33, HV - 2
Member of Shadow Unit 111
111 (virtual)     RA82   Online, HostAccess      Member Count: 2
```

3. Form the logical name from the HSC nodename and the unit number and the equivalence string from the "req" and "port" columns:

```
$ DEFINE/SYSTEM/EXECUTIVE_MODE PSDC$ERNEST_0 "6_2"
$ DEFINE/SYSTEM/EXECUTIVE_MODE PSDC$ERNEST_16 "6_3"
```

These might be added to PSDC$STARTUP.COM; however, a reinstallation of Performance Management might purge this file.

4. Restart the Performance Agent to load the mapping information from the new logical names.

5. Optionally verify the correct collection with the command ADVISE COLLECT REPORT=DUMP_DISKS.

If the channel assignments are not captured by default and you do not use this scheme to assign HSC channel information to shadow set disks, PSDC will not detect HSC channel bottlenecks for shadow sets, nor will graphs of channel data be available for shadow set I/Os, nor will a model configuration of the disk channels be correct.

Appendix C: The Default Workload Families

The CA software provides two default workload families defined in the PSDC$PARAMS.DAT file. They are MODEL_USERGROUPS and MODEL_TRANSACTIONS. This appendix lists the contents of these workload families.

This section contains the following topics:

Users in the MODEL_USERGROUPS Family

Workload Definition Name	User Names
SYSTEM_USER	SYSTEM
OPERATOR	OPERATOR
DECNET	DECNET

Images in the MODEL_TRANSACTIONS Family

Workload Definition Name	Image Names
SYSMAN	APLIC, ARRAY, BACKUP, BUTTON, CALC$MAIN, CALNOTICE, CDU, CLEAR, CLR, CMS, CONFIGURE, CSP, DBMMON, DIRFMT, DQSSMB, DTMFILTER, DVI2LN3, EPC$REGIS, ERRFMT, EVL, HISTORY, HOSTCHECK, INSTALL, JOBCTL, JBC$JOB_CONTROL, MONITOR, NOTICE, OPCOM, PAVN, PLOT, PROCNAM, PROTS, PSAC$MAIN, PSCP$MAIN, PSDC$COLLECT_SYSTEMPC, PSDC$DCOMM, PSDC$DC_V5, PSDC$DECWINDOWS, PSDC$DSKANL, PSDC$EDIT, PSDC$FILES, PSDC$FUNCTIONS, PSDC$MOTIF, PSDC$REPORT_SYSTEMPC, PSDC$RT_V4, PSDC$WINDOW, PSPA$ADVISOR, PSPA$BLDKB, PSPA$DECWINDOWS, PSPA$GRAPH, PSPA$MOTIF, PSPA$RT_CHARCELL, PSPA$RT_MOTIF, QUEMAN, SCHED, SETRIGHTS, STARTUP, SYSGEN, VAXSIM, WHAT, WHYBOOT
SYSMAN	APLIC, ARRAY, BACKUP, BUTTON, CALC$MAIN, CALNOTICE, CDU, CLEAR, CLR, CMS, CONFIGURE, CSP, DBMMON, DIRFMT, DQSSMB, DTMFILTER, DVI2LN3, EPC$REGIS, ERRFMT, EVL, HISTORY, HOSTCHECK, INSTALL, JOBCTL, JBC$JOB_CONTROL, MONITOR, NOTICE, OPCOM, PAVN, PLOT, PROCNAM, PROTS, PSAC$MAIN, PSCP$MAIN, PSDC$COLLECT_SYSTEMPC, PSDC$DCOMM, PSDC$DC_V5, PSDC$DECWINDOWS, PSDC$DSKANL, PSDC$EDIT, PSDC$FILES, PSDC$FUNCTIONS, PSDC$MOTIF, PSDC$REPORT_SYSTEMPC, PSDC$RT_V4, PSDC$WINDOW, PSPA$ADVISOR, PSPA$BLDKB, PSPA$DECWINDOWS, PSPA$GRAPH, PSPA$MOTIF, PSPA$RT_CHARCELL, PSPA$RT_MOTIF, QUEMAN, SCHED, SETRIGHTS, STARTUP, SYSGEN, VAXSIM, WHAT, WHYBOOT
COMPILES	BASIC, BLISS32, FORTRAN, LINK, MACRO32, PASCAL, VAXC
UTILITIES	COPY, CREATE, CREATEFDL, DELETE, DIFF, DIRECTORY, DTM, ENOTES, LNGSPLCOR, LOGINOUT, LPS$SMB, MAIL, NOTES$MAIN, NOTES$SERVER, PHONE, QUOTE_V0, RECOLOR, RENAME, REPLY, RUNOFF, SEARCH, SET, SETP0, SHOW, SHWCLSTR, SORTMERGE, SSU, SUBMIT, TYPE, VMSHELP, VTXPAD

Workload Definition Name	Image Names
SYSMAN	APLIC, ARRAY, BACKUP, BUTTON, CALC$MAIN, CALNOTICE, CDU, CLEAR, CLR, CMS, CONFIGURE, CSP, DBMMON, DIRFMT, DQSSMB, DTMFILTER, DVI2LN3, EPC$REGIS, ERRFMT, EVL, HISTORY, HOSTCHECK, INSTALL, JOBCTL, JBC$JOB_CONTROL, MONITOR, NOTICE, OPCOM, PAVN, PLOT, PROCNAM, PROTS, PSAC$MAIN, PSCP$MAIN, PSDC$COLLECT_SYSTEMPC, PSDC$DCOMM, PSDC$DC_V5, PSDC$DECWINDOWS, PSDC$DSKANL, PSDC$EDIT, PSDC$FILES, PSDC$FUNCTIONS, PSDC$MOTIF, PSDC$REPORT_SYSTEMPC, PSDC$RT_V4, PSDC$WINDOW, PSPA$ADVISOR, PSPA$BLDKB, PSPA$DECWINDOWS, PSPA$GRAPH, PSPA$MOTIF, PSPA$RT_CHARCELL, PSPA$RT_MOTIF, QUEMAN, SCHED, SETRIGHTS, STARTUP, SYSGEN, VAXSIM, WHAT, WHYBOOT
COMPILES	BASIC, BLISS32, FORTRAN, LINK, MACRO32, PASCAL, VAXC
UTILITIES	COPY, CREATE, CREATEFDL, DELETE, DIFF, DIRECTORY, DTM, ENOTES, LNGSPLCOR, LOGINOUT, LPS$SMB, MAIL, NOTES$MAIN, NOTES$SERVER, PHONE, QUOTE_V0, RECOLOR, RENAME, REPLY, RUNOFF, SEARCH, SET, SETP0, SHOW, SHWCLSTR, SORTMERGE, SSU, SUBMIT, TYPE, VMSHELP, VTXPAD
EDITORS	EDT, EMACS, EMACSSHR, LSEDIT, SED, TECO32, TEX, TPU
NETWORK	ELF, FAL, FILESERV, LATCP, LATSYM, NCP, NETACP, NETSERVER, NM$DAEMON, NM$QUEMAN, NML, REMACP, RTPAD
Other	User names or image names that are not defined in one of the other workloads

Glossary

Actual or real workload

The actual or real workload is the workload the computer system actually performs. Ideally, the actual workload is identical to the business workload.

Alternate Performance Agent

The alternate performance agent is a data collection process based upon user-created specifications (a collection definition) that records performance data. An alternate Performance Agent records user-selected data from the data collected by the primary Performance Agent.

Analysis report

The analysis report is a report that identifies the effects of system parameter settings, hardware configurations, workload mixes, and applications on the performance of any cluster node or the entire cluster. The Analysis report provides conclusions with supporting evidence, and makes recommendations.

Analysis summary

The analysis summary is a short summary appearing in the analysis report after each node analysis. The analysis summary contains the following information:

- Number of Performance Manager records analyzed for the specific reporting period
- Number of Performance Manager records that satisfied any rules
- Number of Performance Manager records that did not satisfy any rules
- Number of conclusions generated for the node being analyzed

Archive

Archive is the act of compressing Performance Management daily data records into history file that can be used in place of daily data to create reports, models, and graphs.

Auxiliary knowledge base

The auxiliary knowledge base is a collection of user-defined site-specific rules and thresholds compiled with the rules compiler. The Auxiliary knowledge base rules are used to augment the factory-supplied rules.

Auxiliary rules

Auxiliary rules are the source file containing rules and thresholds that collectively make up the auxiliary knowledge base.

Baseline load

The baseline load is a Performance Management measurements of your existing system or cluster. The baseline load data is stored in a model input file (.MDL).

Baseline model

A baseline model a model generated from historic or daily Performance Management data. The baseline model output provides a workload characterization report. An unmodified. MDL file, the baseline load, represents Performance Management measurements of your existing system. The model generated from this file is known as the baseline model.

Buffered I/O operation

A buffered I/O operation occurs each time an intermediate system buffer is used in place of the process context buffer.

Business

The business is all of the work that is performed on the hardware you are evaluating.

Business workload

Business workload is the work that the business expects its computer to perform is the business workload.

Collection definition

Collection definition is a group of specifications within the schedule file that control a data collection process. These specifications determine which data to collect, when to collect the data, and where to store the data for an individual Performance Agent.

Combined CPU Usage Chart

The Combined CPU Usage chart is a histogram providing information regarding the use of cluster CPU resources.

Computer system

The computer system is all of the hardware on which business work is performed.

Conclusions

Conclusions are text displayed in an analysis report. See Rule conclusion (see page 246).

Context block

The context block is an argument block to the Callable Interface For Data Extraction. It contains control information, status information, and addresses of the daily data subrecords. The subrecord fields are also known as data cells.

CPD

The CPD collection definition is provided by default in the schedule file. This set of parameters controls the primary data collector process.

CPU branch explicit

In a modeling context the CPU branch explicit is the probable distribution of load across CPUs for a workload or transaction class. Probabilities must sum to 1. This data is included in the model input file.

Custom graph

A custom graph is a graph type that allows you to specify which Performance Manager data items to graph.

Daily data files

Daily data files are files created by the Performance Agent, one for each node in the cluster, every day. The file name has the following format:

```
pathname$DATABASE:PSDC$nodename_yyyymmmdd.coll-def-name
```

Data cell

A data cell is the basic unit of data used to create analysis reports. This data is either retrieved directly from a field in a subrecord of a daily data record or derived from it. The data cell is typically used as a variable in a rule expression.

Data collection error log

The data collection error log is an ASCII file common to the entire cluster called PSDC$DATABASE:PSDC$DC.LOG. Errors that occur during data collection are recorded in the data collection error log file.

Data collection schedule

The data collection schedule is a user-defined schedule by which the Performance Agent records data for nodes in the cluster.

Data collection synchronization

Data collection synchronization is a method used to correlate intervals measured on different nodes within a cluster; those intervals must represent the same real time to make analysis effective.

Data files

Data files are the files from which reports and graphs are constructed. See Daily data files (see page 241) and History files (see page 47).

Data record

The data record contains OpenVMS performance data written by the Performance Agent. A Performance Agent writes one data record for each interval.

Database directory

The database directory is a directory located on a permanently mounted disk, accessible to every node in the cluster. The data collection process may write the daily data files to this area. Performance Management references this area via the system-wide logical name PSDC$DATABASE.

Device

In the context of the Performance Management, a device is a tape or disk drive. It performs measurable units of work.

Direct I/O

The number of direct I/O operations performed per second. This figure is tallied at the $QIO application interface layer.

Disk branch by source

In a modeling context, the probable distribution of load across disks by originating CPU for a workload or a transaction class is Disk branch by source. Probabilities must sum to 1. This data is included in the model input file.

Disk I/O

Disk I/O is the number of I/O operations per second for the device. This figure is tallied at the physical device driver layer.

Dump report

The dump report contains formatted output of data fields for each record of a Performance Management daily data file or history file.

Evidence

Evidence is text displayed in an analysis report. See Rule evidence (see page 246).

Factory rules

Factory rules are the performance rules supplied with the Performance Manager in the PSPA$EXAMPLES area in the file named PSPA$KB.VPR.

Family name

Family name is the identifier for a group of workloads, also known as a workload family.

Family type

There are two family types:

- Usergroup families contain workload definitions based on user criteria
- Transaction families contain workload definitions based on image and process data

Specify either family with the /CLASSIFY_BY qualifier to control use of the family for the reporting facility.

Granularity

Granularity is a Performance Management parameter file element that specifies for each history file descriptor how often a history file is created.

Hard page fault

Hard page fault is an event that occurs each time a process references a virtual page that is not in its working set and requires a read operation from disk.

Histogram

Histogram is an (ASCII) graphic chronological chart showing resource use.

History database

History database contains reduced data from the daily data files, which resides in the history files.

History file

A History file contains data archived from the daily data files. The number of history files created and maintained depends upon the number of history file descriptors and the associated granularity. The file name has the following format:

```
pathname:PSDC$nodename_dd-mmm-yyyy_dd-mmm-yyyy.name.
```

History file descriptor

History file descriptor contains the description that the Performance Agent uses to determine how to archive data to the file.

Holidays file

The holidays file contains a list of holiday dates. The Performance Agent uses the holiday schedule on these dates. The file name is PSDC$HOLIDAYS.TXT.

Hot files

Hot files are the most frequently accessed files on each disk. The primary (main) Performance Agent collects hot file data when the length of a disk queue exceeds the HOTFILE_QUEUE setting.

Interaction

An interaction with a device is a two-step process. First, a unit of work to be done (job, user, process, and so forth) enters the queue of the device. Then, the unit of work is serviced by the device (in a manner dependent on the queuing discipline) and departs the device. The concept may be generalized to the system as a whole. The set of device interactions required to process each unit of work is called a transaction.

Interval

In the context of a daily data file, the interval for writing records is the elapsed time between recording.

Journal file

Performance Agent creates one history journal file for each history file. The Journal files are used by the update process in conjunction with the daily data files to recreate corrupt or deleted history files. The file name has the format: pathname:PSDC$nodename_dd-mmm-yyyy_dd-mmm-yyyy.name_JOU
Do not delete these files. If they exist, they are needed.

Knowledge base

Knowledge base is a file consisting of Performance Manager performance rules used to analyze daily or historic Performance Manager data. It is created with the ADVISE PERFORMANCE COMPILE command from a USER_RULES.VPR file. This data is known as the auxiliary knowledge base.

Measured workload

The measured workload is the workload that can be observed on the system. Ideally, the measured workload is identical to the actual workload, but specified in different terms.
The measured workload is specified by the resource demands that it places on the actual system. These demands, or loadings, are given as the service times of the users at each device along with the number of expected transactions at each device.

MSCP

Mass Storage Control Protocol is a software protocol used to communicate between a host and an intelligent storage controller, or between hosts. In the latter case, one or more hosts serves mass storage devices to other hosts.

OpenVMS cluster

An OpenVMS Cluster is a highly integrated organization of AlphaServer and HP Integrity server system - or VAX and HP AlphaServer system-applications, operating systems, and storage devices.

OTHER

In a modeling context, OTHER is a workload name. All workload families have the catch all workload OTHER to absorb process data that does not match the selection criteria of any defined workload.

Parameters file

The parameters file, PSDC$PARAMS.DAT resides in the PSDC$DATABASE area and serves as a repository for workload characterizations, history file descriptors, and other Performance Management parameters.

Performance Agent

Performance Agent is a detached process that collects and records performance data for each node in the cluster according to a weekly schedule.

Performance evaluation report

The performance evaluation report is a statistical report that helps you determine whether changes that you implemented (based on recommendations in the Analysis Report) improved or degraded system performance.

Periodicity

Periodicity is a parameter file element for a history file descriptor that specifies how often the averaging cycle is restarted.

Predefined graph

A Predefined graph is a graph in that the Performance Manager defines the metrics plotted.

Queue length

Queue length is the average number of outstanding requests, either waiting for or receiving service.

Queuing network model

A queuing network model is a mathematical abstraction of a system in which the computer system is represented as a network of queues. Each queue in the network is evaluated analytically.

Raw data file

See Daily data file (see page 241).

Recommendation

A recommendation is text presented in an analysis report that offers system tuning advice based on rules firing is recommendation. See Rule conclusion (see page 246).

Residence time

Residence time is the time, in seconds, between image activation and image termination is residence time. Average time that a request spends wShile waiting for and receiving service.

Response time

Response time is the elapsed time between the arrival of a request and the moment of completion. In the context of modeling, the time between the moment of request arrival at a device and request completion at a device is response time.

Rule

A rule is one or more rule conditions that are evaluated when Performance Manager Analysis reports are generated. Rules are applied to daily or historic data. If all conditions for a rule are true then there is a rule occurrence. Rules are defined to expose areas of potential system problems. See Rule firing (see page 246).

Rule conclusion

A rule conclusion is a rule element. The rule conclusions are Performance Manager recommendations based on the conclusion text element of a rules file rule construct.
See Rule element (see page 246).

Rule condition

A rule condition is a rule element. A rule condition is made up of one or more rule expressions. Describes the circumstances which must be true to cause a rule occurrence.

Rule elements

The seven rule elements can exist as the following rule constructs:

- Brief conclusion element
- Conclusion element
- Domain element
- Evidence element
- Occurrence element
- Rule condition element
- Rule ID element

Rule evidence

A rule evidence is a rule element. Data satisfying a rule occurrence in a Performance Manager Analysis report. The rule evidence consists of data cell names and values. Typically these data cells are some of those contained in the rule expression.

Rule firing

Rule firing occurs after all the data has been processed, when creating a report, the Performance Manager examines the number of rule occurrences for each rule. If a rule occurrence threshold is met for a particular rule, the rule is said to fire. For each rule that fires an entry is made in the report. The entry may include evidence and conclusions.

Rule identifier

A rule identifier is a rule element. Rule identifier is a five-character alphanumeric code enclosed in braces, one alphabetic character followed by four numeric characters, for example, {M0010} which uniquely identifies the rule.

Rule occurrence

Each time all the rule conditions for a given rule are true, there is one rule occurrence. See Rule firing (see page 246).

Rules compiler

The rules compiler generates a knowledge base file, (name.KB), from a rules source file (name.VPR)

Rules file

Rules file is a compiled knowledge base file (name.KB). A performance rules source file (name.VPR).

Rules file constructs

The following rules file constructs can exist in a rules file:

- Comment construct
- Disable construct
- Literal construct
- Rule construct
- Threshold construct.

Saturation

Saturation is the arbitrary point at which response time at a device becomes substantially higher than the service time.

Schedule file

A schedule file is a file (PSDC$SCHEDULE.DAT) that resides in the PSDC$DATABASE area and controls the recording of Performance Management daily data is Schedule file. A set of parameters (a collection definition) specifies which data to collect, when to collect the data, where to store the data, and how long to keep the data files. The software provides a default set of parameters (the CPD collection definition) and users can create additional collection definitions.

Sensitivity analysis

A sensitivity analysis is a type of "what if" analysis. The act of making small changes in a system model to disclose the areas that are sensitive to these changes.

Service time

Service time is the time required to service a request at a device excluding queuing delay is service time.

Shadow set

A shadow set is one or more compatible physical disk volumes connected together for volume shadowing and represented by a virtual unit. Thus, the term shadow set refers to both the physical unit and the virtual unit.

Soft page fault

Soft page fault occurs each time a process references a virtual page that is not in its working set, but is in memory.

Split I/O

Split I/O is the number of additional physical disk I/O operations required to complete a single user's I/O request, which could not be satisfied in a single I/O to a device.

Summarized

Summarized refers to archived daily data. See History file (see page 243).

Throughput

Throughput is the number of work units processed per unit of time. The Modeling summary report gives system-wide throughput in terms of transactions per second. The modeling detail report gives throughput in terms of I/Os per second for each component.

TPS

Transactions Per Second is the rate at which transactions arrive to a system.

Transaction

A transaction is a quantifiable unit of work. A transaction typically delineates a single processing step in computer systems.

Transaction class

Transaction class is a group of related transactions. They may be related by the function they perform, by the users who initiate them, or by other quantities you define. Transactions may also be determined by the system resource demands. Performance Management generally refers to a transaction class as a workload.

Transaction class workload

Transaction class workload is a workload that contains at least one image name is transaction class workload.

Transaction workload family

Transaction workload family is set of transaction-based workload definitions.

Usergroup workload family

Usergroup workload family is a set of user-based workload definitions.

User defined graph

See Custom graph (see page 241).

Utilization

Utilization is the percentage of a resource's capacity being used.

VUP

VUP is a CPU VAX unit of processing rate. The CPU VUP rate measures the power of a system compared to a VAX 11/780. A VAX 11/780 has a VUP rate of 1.0.

Wait Time

Wait time, or queuing time, is the average time each request spends waiting in a queue for service. During this time, the request accomplishes no useful work. Wait time can be derived specifically for each device in the system or for the system as a whole.

Workload

A workload can be one or more users and/or image names. Workload definitions are identified and stored in the Performance Management parameters file. The goal is to express the system's total workload in manageable and meaningful units that Performance Management can report against.
A *Transaction* workload definition contains at least one image name. Typically, this workload definition contains image names with similar resource demands.
A *Usergroup* workload contains a user specifier (account name, user name, process name or UIC code). Typically this workload contains a group of users who belong to the same business unit such as a department.
Transaction and *Usergroup* are terms applied to a workload by the /CLASSIFY_BY qualifier.

Workload family name

Workload family name is a name identifying a group of workloads.

Workload name

Workload name is a name identifying a workload.

Workload selection criteria

The workload selection criteria is the criteria by which process data is assigned to a workload. A workload is selected when process data matches on either or both of a user specifier or image name.

XFC

The Extended File Cache (XFC) is a clusterwide virtual block file system data cache data cache provided with OpenVMS versions 7.3 and up as a replacement for the Virtual I/O cache.

Index

D

E

F

G

H

I

J

L

M

N

O

P

Q

R

S

T

U

W