
Availability Manager User's Guide

Order Number: AA-RNSJA-TE

April 2001

This guide explains how to use Availability Manager software to detect and correct system availability problems.

Revision/Update Information:	This is a new manual.
Operating System:	Data Analyzer: Windows NT 4.0, SP 3 or higher; Windows 2000; OpenVMS Version 7.1 or later Data Collector: OpenVMS Alpha and VAX Version 6.2 or later
Software Version:	Availability Manager Version 1.4

**Compaq Computer Corporation
Houston, Texas**

© 2001 Compaq Computer Corporation

Compaq, VAX, VMS, and the Compaq logo Registered in U.S. Patent and Trademark Office.

OpenVMS is a trademark of Compaq Information Technologies Group, L.P. in the United States and other countries.

Microsoft, Windows, Windows NT, and Windows 95 are trademarks of Microsoft Corporation in the United States and other countries.

Motif, OSF/1, and UNIX are trademarks of The Open Group in the United States and other countries.

All other product names mentioned herein may be the trademarks of their respective companies.

Confidential computer software. Valid license from Compaq required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Compaq products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

ZK6552

The Compaq *OpenVMS* documentation set is available on CD-ROM.

This document was prepared using DECdocument, Version 3.3-1b.

Contents

Preface	ix
1 Overview	
1.1 What Is the Availability Manager?	1-1
1.2 How Does the Availability Manager Work?	1-2
1.3 How Does the Availability Manager Identify Performance Problems?	1-4
1.3.1 Collecting and Analyzing Data	1-5
1.3.1.1 Types of Data Collection	1-5
1.3.1.2 Events and Data Collection	1-7
1.3.1.3 Data Collection Intervals	1-7
1.3.2 Posting Events	1-8
1.3.2.1 Thresholds and Occurrences	1-8
1.4 How Does the Availability Manager Maintain Security?	1-10
1.4.1 Data Analyzer Password Security	1-10
1.4.2 OpenVMS Data Collector Security Features	1-11
2 Getting Started	
2.1 Starting the Data Collector	2-1
2.2 Starting the Data Analyzer	2-2
2.2.1 How to Start the Data Analyzer on an OpenVMS Alpha Node	2-2
2.2.2 How to Start the Data Analyzer on a Windows NT Node	2-2
2.3 Using the Application Window	2-3
2.3.1 Other Window Components	2-4
2.3.2 Displaying More Information	2-4
2.4 Selecting a Group of Nodes	2-5
2.4.1 Group Names That Are Displayed by Default	2-5
2.4.2 How to Display Other Nodes	2-6
2.5 Displaying Information about Nodes	2-6
2.5.1 Choosing a Node	2-6
2.5.2 Specifying Data to Collect on OpenVMS Nodes	2-7
2.5.3 Sorting Data	2-8
2.6 Getting Help	2-9
3 Getting Information about Nodes	
3.1 Node Panes	3-1
3.1.1 OpenVMS Node Pane	3-1
3.1.2 Windows NT Node Pane	3-2
3.2 Node Data Pages	3-3
3.2.1 Node Summary	3-4

3.2.2	CPU Modes and Process Summaries	3-4
3.2.2.1	Windows NT CPU Modes	3-5
3.2.2.2	OpenVMS CPU Modes Summary	3-6
3.2.2.3	OpenVMS CPU Modes Detail	3-7
3.2.2.4	OpenVMS CPU Process Summary	3-8
3.2.3	Memory Summaries	3-8
3.2.3.1	Windows NT Memory Summary	3-9
3.2.3.2	OpenVMS Memory Summary	3-10
3.2.4	OpenVMS I/O Summaries and Page/Swap Files	3-11
3.2.4.1	OpenVMS I/O Summaries	3-11
3.2.4.2	OpenVMS I/O Page/Swap Files	3-14
3.2.5	Disk Summaries	3-15
3.2.5.1	OpenVMS Disk Status Summary	3-15
3.2.5.2	OpenVMS Single Disk Summary	3-16
3.2.5.3	OpenVMS Disk Volume Summary	3-17
3.2.5.4	Windows NT Logical Disk Summary	3-19
3.2.5.5	Windows NT Physical Disk Summary	3-20
3.2.6	OpenVMS Lock Contention	3-20
3.2.7	OpenVMS Cluster Summary	3-22
3.2.7.1	Summary Pane	3-24
3.2.7.2	Cluster Members Pane	3-24
3.2.7.3	System Communications Architecture (SCA) Virtual Circuit (VC) Data	3-25
3.2.7.4	SCS Connections Data	3-27
3.2.7.5	LAN Virtual Circuit Details Data	3-28
3.2.8	OpenVMS Network Interconnect for the System Communication Architecture (NISCA) Pages	3-28
3.2.8.1	Channel Selection Data	3-29
3.2.8.2	Virtual Circuit (VC) Closures Data	3-30
3.2.8.3	Packets Discarded Data	3-31
3.2.8.4	Transmit Data	3-32
3.2.8.5	Receive Data	3-33
3.2.8.6	Congestion Control Data	3-34
3.2.9	OpenVMS Single Process	3-36
3.2.9.1	Process Information	3-36
3.2.9.2	Working Set	3-37
3.2.9.3	Execution Rates	3-38
3.2.9.4	Process I/O Quotas	3-39
3.2.9.5	Wait States	3-40
3.2.9.6	Job Quotas	3-42

4 Getting Information about Events

4.1	Displaying Event Information	4-2
4.1.1	Data in the Events Pane	4-2
4.1.2	Events Pane Menu Options	4-2
4.2	Criteria Signaling and Displaying an Event	4-2
4.3	Displaying Additional Event Information	4-5

5 Performing Fixes on OpenVMS Nodes

5.1	Understanding Fixes	5-1
5.2	Performing Fixes	5-3
5.2.1	Node Fixes	5-4
5.2.1.1	Crash Node	5-4
5.2.1.2	Adjust Quorum	5-5
5.2.2	Performing Process Fixes	5-5
5.2.3	General Process Fixes	5-6
5.2.3.1	Delete Process	5-6
5.2.3.2	Exit Image	5-7
5.2.3.3	Suspend Process	5-8
5.2.3.4	Resume Process	5-9
5.2.3.5	Change Process Priority	5-10
5.2.4	Process Memory Fixes	5-11
5.2.4.1	Purge Working Set	5-11
5.2.4.2	Adjust Working Set	5-12
5.2.5	Process Limits Fixes	5-13
5.2.5.1	Direct I/O Count Limit	5-14
5.2.5.2	Buffered I/O Count Limit	5-15
5.2.5.3	AST Queue Limit	5-16
5.2.5.4	Open File Limit	5-17
5.2.5.5	Lock Queue Limit	5-18
5.2.5.6	Time Queue Entry Limit	5-19
5.2.5.7	Subprocess Creation Limit	5-19
5.2.5.8	I/O Byte	5-21

6 Customizing the Availability Manager

6.1	Specifying Groups or Individual Nodes to Monitor	6-2
6.2	Changing the Group Membership of a Node	6-4
6.2.1	Changing the Group of an OpenVMS Node	6-4
6.2.2	Changing the Group of a Windows NT Node	6-4
6.3	Customizing OpenVMS Data Collection	6-5
6.4	Customizing OpenVMS Data Filters	6-7
6.4.1	OpenVMS CPU Filters	6-7
6.4.2	OpenVMS Disk Status Filters	6-8
6.4.3	OpenVMS Disk Volume Filters	6-10
6.4.4	OpenVMS I/O Filters	6-11
6.4.5	OpenVMS Lock Contention Filters	6-12
6.4.6	OpenVMS Memory Filters	6-13
6.4.7	OpenVMS Page/Swap File Filters	6-14
6.5	Customizing Events	6-15
6.6	Customizing Security Features	6-16
6.6.1	Changing Data Analyzer Passwords	6-16
6.6.1.1	Changing a Data Analyzer Password for an OpenVMS Data Collector Node	6-17
6.6.1.2	Changing a Data Analyzer Password for a Windows NT Data Collector Node	6-17
6.6.2	Changing Security Triplets on OpenVMS Data Collector Nodes	6-18
6.6.2.1	Understanding OpenVMS Security Triplets	6-18
6.6.2.2	How to Change a Security Triplet	6-19
6.6.2.3	How the Availability Manager Ensures Security	6-21
6.6.3	Changing a Password on a Windows NT Data Collector	6-21

A CPU Process States

B Tables of Events

C OpenVMS Events by Type of Data Collection

Index

Figures

1-1	Availability Manager Node Configuration	1-3
1-2	Requesting and Receiving Information	1-4
1-3	Application Window	1-5
1-4	Data Collection Page	1-6
1-5	Sample Node Summary Page	1-6
1-6	Sample Event Customization Page	1-8
1-7	Flow Chart of Event Testing	1-9
1-8	Availability Manager Password Matching	1-10
2-1	Application Window	2-3
2-2	Group Pane	2-5
2-3	Node Pane	2-6
2-4	OpenVMS Node Summary Page	2-7
2-5	Data Collection Page	2-8
3-1	OpenVMS Node Pane	3-1
3-2	Windows NT Node Pane	3-2
3-3	Node Summary Page	3-4
3-4	Windows NT CPU Modes Page	3-5
3-5	OpenVMS CPU Modes Summary Pane	3-6
3-6	OpenVMS CPU Modes Detail Pane	3-7
3-7	OpenVMS CPU Process Summary Pane	3-8
3-8	Windows NT Memory Page	3-9
3-9	OpenVMS Memory Page	3-10
3-10	OpenVMS I/O Summaries Page	3-12
3-11	OpenVMS I/O Page/Swap Files Page	3-14
3-12	OpenVMS Disk Status Summary	3-15
3-13	OpenVMS Single Disk Summary Page	3-17
3-14	OpenVMS Disk Volume Summary Page	3-18
3-15	Windows NT Logical Disk Summary Page	3-19
3-16	Windows NT Physical Disk Summary Page	3-20
3-17	OpenVMS Lock Contention Page	3-21
3-18	OpenVMS Cluster Summary Page	3-23
3-19	Virtual Circuit Data Page	3-25
3-20	LAN Virtual Circuit Details	3-29
3-21	Channel Selection Data Page	3-30
3-22	Virtual Connect (VC) Closures Data Page	3-31
3-23	Packets Discarded Data Page	3-32

3-24	Transmit Data Page	3-33
3-25	Receive Data Page	3-34
3-26	Congestion Control Data Page	3-35
3-27	Process Information Page	3-36
3-28	Working Set Page	3-37
3-29	Execution Rates Page	3-39
3-30	Process Quotas Page	3-40
3-31	Wait States Page	3-41
3-32	Job Quotas Page	3-42
4-1	Application Window	4-1
4-2	Sample Event Customization Page	4-3
4-3	OpenVMS Data Collection Page	4-4
4-4	OpenVMS Node Pane	4-4
5-1	Crash Node Page	5-4
5-2	Adjust Quorum Page	5-5
5-3	Change Process Priority Page	5-6
5-4	Delete Process Page	5-7
5-5	Exit Image Page	5-8
5-6	Suspend Process Page	5-9
5-7	Resume Process Page	5-10
5-8	Change Process Priority Page	5-11
5-9	Purge Working Set Page	5-12
5-10	Adjust Working Set Page	5-13
5-11	Direct I/O Count Limit Page	5-14
5-12	Buffered I/O Count Limit Page	5-15
5-13	AST Queue Limit Page	5-16
5-14	Open File Limit Page	5-17
5-15	Lock Queue Limit Page	5-18
5-16	Time Queue Entry Limit Page	5-19
5-17	Subprocess Creation Limit Page	5-20
5-18	I/O Byte Page	5-21
6-1	Customize Menu in Application Window	6-1
6-2	Customize Menu in Node Pane	6-2
6-3	Group/Node Lists Customization Page	6-3
6-4	OpenVMS Data Collection Page	6-5
6-5	OpenVMS CPU Filters Page	6-8
6-6	OpenVMS Disk Status Filters Page	6-9
6-7	OpenVMS Disk Volume Filters Page	6-10
6-8	OpenVMS I/O Filters Page	6-11
6-9	OpenVMS Lock Contention Filters Page	6-12
6-10	OpenVMS Memory Filters Page	6-13
6-11	OpenVMS Page/Swap File Filters Page	6-14
6-12	Event Customization Page	6-15
6-13	OpenVMS Security Customization Page	6-17
6-14	Windows NT Security Customization Page	6-18

Tables

1-1	Data Collection Intervals	1-7
3-1	OpenVMS Node Data	3-2
3-2	Windows NT Node Data	3-2
3-3	Node Data Page Menu Bar	3-3
3-4	I/O Data Displayed	3-12
3-5	Summary Pane Data	3-24
3-6	Cluster Member Data	3-24
3-7	Virtual Circuit Data	3-26
3-8	SCS Connections Data	3-27
3-9	LAN Virtual Circuit Details Data	3-28
4-1	Events Pane Data	4-2
5-1	Accessing Availability Manager Fixes	5-1
5-2	Summary of Problems and Matching Fixes	5-2
6-1	Data Collection Choices	6-6
6-2	Data Collection Intervals	6-6
6-3	Security Triplet Verification	6-21
A-1	CPU Process States	A-1
B-1	OpenVMS Events	B-1
B-2	Windows NT Events	B-12
C-1	OpenVMS Threshold Events	C-1
C-2	OpenVMS Nonthreshold Events	C-4

Preface

Intended Audience

This guide is intended for system managers who install and use Compaq Availability Manager software. It is assumed that the system managers who use this product are familiar with Windows terms and functions.

Document Structure

This guide contains the following chapters and appendixes:

- Chapter 1 provides an overview of Availability Manager software, including security features.
- Chapter 2 describes how to start the Availability Manager, use the main Application window, select a group of nodes and individual nodes, and use online help.
- Chapter 3 describes how to select nodes and display node data.
- Chapter 4 describes how to display and interpret events.
- Chapter 5 describes how to take a variety of corrective actions, called **fixes**, to improve system availability.
- Chapter 6 describes the tasks you can perform to filter, select, and customize the display of data
- Appendix A contains a table of CPU process states, which are referred to in Section 3.2.2.4 and in Section 3.2.9.1.
- Appendix B contains a table of OpenVMS and Windows NT events that can be displayed in the Events pane discussed in Chapter 4.
- Appendix C describes the events that can be signaled for each type of OpenVMS data that is collected.

Related Documents

The following manuals provide additional information:

- *OpenVMS System Manager's Manual* describes tasks for managing an OpenVMS system. It also describes installing a product with the POLYCENTER Software Installation utility.
- *OpenVMS System Management Utilities Reference Manual* describes utilities you can use to manage an OpenVMS system.
- *OpenVMS Programming Concepts Manual* explains OpenVMS lock management concepts.

For additional information about Compaq *OpenVMS* products and services, access the Compaq website at the following location:

<http://www.openvms.compaq.com/>

Reader's Comments

Compaq welcomes your comments on this manual. Please send comments to either of the following addresses:

Internet	openvmsdoc@compaq.com
Mail	Compaq Computer Corporation OSSG Documentation Group, ZKO3-4/U08 110 Spit Brook Rd. Nashua, NH 03062-2698

How to Order Additional Documentation

Use the following World Wide Web address to order additional documentation:

<http://www.openvms.compaq.com/>

If you need help deciding which documentation best meets your needs, call 800-282-6672.

Conventions

The following conventions are used in this guide:

Ctrl/ <i>x</i>	A sequence such as Ctrl/ <i>x</i> indicates that you must hold down the key labeled Ctrl while you press another key or a pointing device button.
PF1 <i>x</i>	A sequence such as PF1 <i>x</i> indicates that you must first press and release the key labeled PF1 and then press and release another key or a pointing device button.
Return	In examples, a key name enclosed in a box indicates that you press a key on the keyboard. In text, a key name is not enclosed in a box. In the HTML version of this document, this convention appears as brackets, rather than a box.
...	Horizontal ellipsis points in examples indicate one of the following possibilities: <ul style="list-style-type: none">• Additional optional arguments in a statement have been omitted.• The preceding item or items can be repeated one or more times.• Additional parameters, values, or other information can be entered.
.	Vertical ellipsis points indicate the omission of items from a code example or command format; the items are omitted because they are not important to the topic being discussed.
()	In command format descriptions, parentheses indicate that you must enclose the options in parentheses if you choose more than one.

[]	In command format descriptions, brackets indicate optional elements. You can choose one, none, or all of the options. (Brackets are not optional, however, in the syntax of a directory name in an OpenVMS file specification or in the syntax of a substring specification in an assignment statement.)
{ }	In command format descriptions, braces indicate required elements; you must choose one of the options listed.
bold text	This typeface represents the introduction of a new term. It also represents the name of an argument, an attribute, or a reason. In the HTML version of this Conventions table, this convention appears as <i>italic text</i> .
<i>italic text</i>	Italic text indicates important information, complete titles of manuals, or variables. Variables include information that varies in system output (Internal error <i>number</i>), in command lines (<i>/PRÓDUCER=name</i>), and in command parameters in text (where <i>dd</i> represents the predefined code for the device type).
UPPERCASE TEXT	Uppercase text indicates a command, the name of a routine, the name of a file, or the abbreviation for a system privilege.
Monospace type	Monospace type indicates code examples and interactive screen displays. In the C programming language, monospace type identifies the following elements: keywords, the names of independently compiled external functions and files, syntax summaries, and references to variables or identifiers introduced in an example.
numbers	All numbers in text are assumed to be decimal unless otherwise noted. Nondecimal radices—binary, octal, or hexadecimal—are explicitly indicated.

This chapter provides the following information:

- What the Availability Manager is
- How the Availability Manager works
- How the Availability Manager identifies possible performance problems
- How the Availability Manager maintains security

1.1 What Is the Availability Manager?

The Availability Manager is a system management tool that allows you to monitor, from an OpenVMS node or a Windows NT node, one or more OpenVMS nodes on an extended local area network (LAN).

The Availability Manager helps system managers and analysts target a specific node or process for detailed analysis. This tool collects system and process data from multiple OpenVMS nodes simultaneously; it analyzes the data and uses a graphical user interface (GUI) to display the output.

An older version of the tool, DECamds, uses a Motif GUI to display information about OpenVMS nodes. The newer version, called the Availability Manager, uses a Java GUI to display information about OpenVMS nodes on an OpenVMS or a Windows NT node.

The main Application window of the Availability Manager is divided into three sections that display different types of information about the nodes you are monitoring. Based on its analysis of the data, the Availability Manager notifies you immediately if any node you are monitoring is experiencing a performance problem, especially one that affects the node's accessibility to users. At a glance, you can see whether a problem is a persistent one that warrants further investigation and correction. The Availability Manager also maintains an event log file, where it logs every event displayed in the main Application window. (See Section 1.3 for details.)

An important advantage of the Availability Manager is that it uses its own network protocol; unlike most performance monitors, it does not rely on TCP/IP or any other standard protocol. Therefore, even if a standard protocol is unavailable, the Availability Manager can continue to operate.

You can customize the Availability Manager to meet the requirements of your particular site. For example, you can change the severity levels of the events that are displayed and escalate their importance.

The Availability Manager helps improve OpenVMS system and OpenVMS Cluster availability by providing the following features:

Overview

1.1 What Is the Availability Manager?

Availability	Alerts users to resource availability problems; provides capabilities to improve availability.
Centralized management	Provides centralized management of remote nodes within an extended local area network (LAN).
Intuitive interface	Provides an easy-to-learn and easy-to-use graphical user interface (GUI).
Correction capability	Allows real-time intervention, including adjustment of node and process parameters, even when remote nodes are hung.
Customization	Adjusts to site-specific requirements through a wide range of customization options.
Scalability	Makes it easier to monitor multiple OpenVMS nodes.

1.2 How Does the Availability Manager Work?

The Availability Manager utilizes two types of nodes for monitoring OpenVMS systems:

- One or more OpenVMS Data Collector nodes, which run the software that collects the data on the OpenVMS nodes being monitored.
- An OpenVMS or a Windows NT Data Analyzer node, which contains the software that analyzes the data collected from the monitored OpenVMS nodes.


The Data Analyzer and Data Collector nodes communicate over an extended LAN using an IEEE 802.3 Extended Packet format protocol. Once a secure connection is established, the Data Analyzer instructs the Data Collector to gather specific system and process data.

Although you can run the Data Analyzer as a member of a monitored cluster, it is typically run on a system that is not a member of the cluster being monitored. You can have more than one Data Analyzer application executing in a LAN, but only one Data Analyzer at a time should be running on each system.

Figure 1–1 shows a possible configuration of Data Analyzer and Data Collector nodes.

1.2 How Does the Availability Manager Work?

Figure 1–1 Availability Manager Node Configuration



In Figure 1–1, the Data Analyzer can monitor nodes A, B, and C across the network. The password on node D does not match the password of the Data Analyzer; therefore, the Data Analyzer cannot monitor node D.

For information about password security, see Section 1.4.

Requesting and Receiving Information

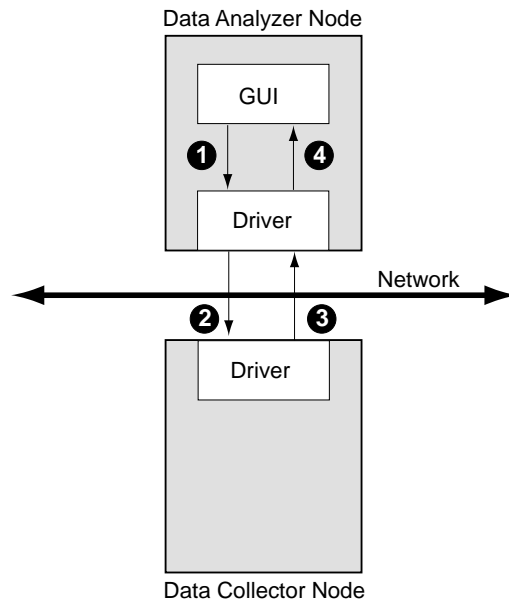
After installing the Availability Manager software, you can begin to request information from one or more Data Collector nodes.

Requesting and receiving information requires the Availability Manager to perform a number of steps, which are shown in Figure 1–2 and explained after the figure.

Overview

1.2 How Does the Availability Manager Work?

Figure 1–2 Requesting and Receiving Information



VM-0358A-AI

The following steps correspond to the numbers in Figure 1–2.

- ❶ The GUI communicates users' requests for data to the driver on the Data Analyzer node.
- ❷ The Data Analyzer driver sends users' requests across the network to a driver on a Data Collector node.
- ❸ The Data Collector driver transmits the requested information over the network to the driver on the Data Analyzer node.
- ❹ The Data Analyzer driver passes the requested information to the GUI, which displays the data.

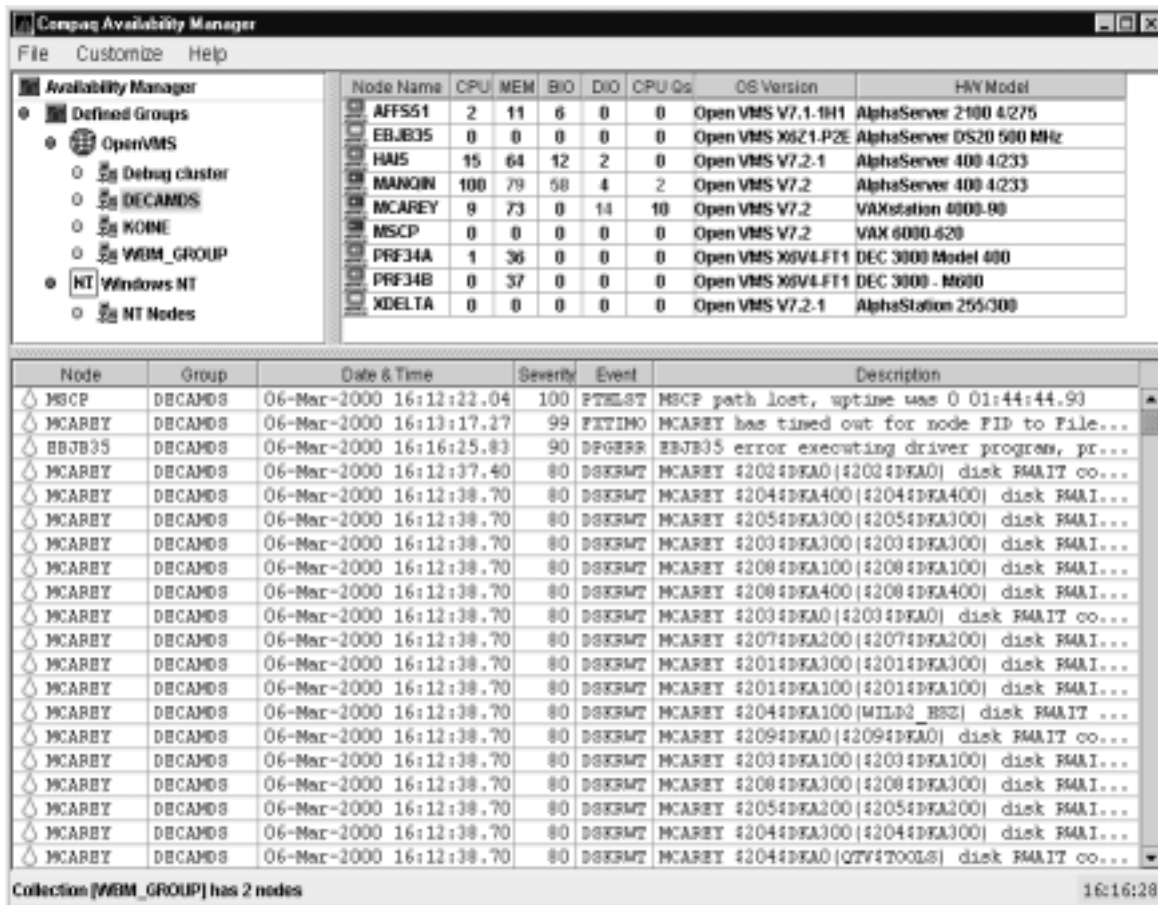
In step 4, the Availability Manager also checks the data for any events that should be signaled. The following section explains in more detail how data analysis and event detection work.

1.3 How Does the Availability Manager Identify Performance Problems?

When the Availability Manager detects problems on your system, it uses a combination of methods to bring these problems to the attention of the system manager. If no data display is open for a particular node, the Availability Manager reduces the data collection interval so that data can be analyzed more closely. Performance events are also signaled in the Events pane in the lower portion of the Application window (Figure 1–3).

1.3 How Does the Availability Manager Identify Performance Problems?

Figure 1–3 Application Window



The following topics are related to detecting and signaling problems:

- Collecting and analyzing data
- Signaling events

1.3.1 Collecting and Analyzing Data

This section explains how the Availability Manager collects and analyzes data. It also defines terms related to data collection and analysis.

1.3.1.1 Types of Data Collection

You can use the Availability Manager to collect data either as a background activity or as a foreground activity.

- **Background** data collection

When you enable background collection of a specific type of data on a specific node, the Availability Manager collects that data whether or not any windows are currently displaying data for that node.

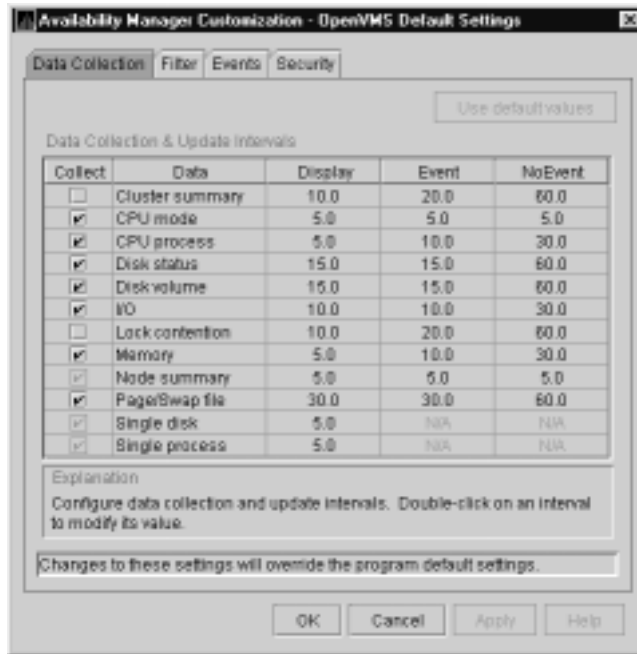
To enable background data collection, select the check box for a specific type of data on a Data Collection page (Figure 1–4). (If the Customize window is for all OpenVMS nodes, you set defaults for all nodes. If the window is for one node, you set collection properties for a single node.)

Overview

1.3 How Does the Availability Manager Identify Performance Problems?

Chapter 6 contains instructions for customizing data collection properties.

Figure 1–4 Data Collection Page

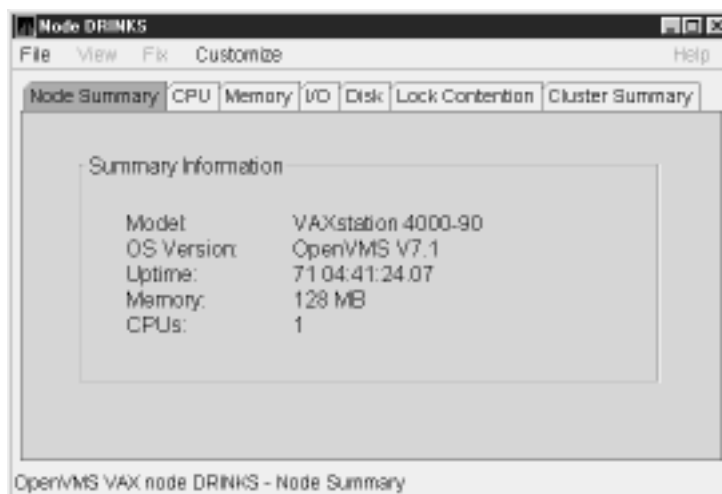


By default, node summary data is always collected.

- **Foreground** data collection

Foreground data collection occurs automatically when you open any data page for a specific node. To open a node data page, double-click a node name in the Node pane of the Application window (Figure 1–3). The Node Summary page is displayed by default, and you can select tabs to display other data pages for that node. Figure 1–5 is an example of a Node Summary page.

Figure 1–5 Sample Node Summary Page



1.3 How Does the Availability Manager Identify Performance Problems?

Foreground data collection for all data types begins automatically when any node data page is displayed. Data collection ends when all node data pages have been closed.

Chapter 3 contains instructions for selecting nodes and displaying node data.

1.3.1.2 Events and Data Collection

An **event** is a problem or potential problem associated with resource availability. Users can customize criteria for events. Events are associated with types of data collected. For example, collection of CPU data is associated with the PRCCUR, PRCMWT, and PRCPWT events. (Appendix B describes events, and Appendix C describes the events that each type of data collection can signal.)

As data is collected, the Availability Manager evaluates it and signals an event whenever the data meets the user-specified criteria. These criteria are called thresholds and occurrences and are explained in Section 1.3.1.3.

Customization features are explained in detail in Chapter 6.

1.3.1.3 Data Collection Intervals

Data collection intervals, which are displayed on the Data Collection page (Figure 1–4), specify the frequency of data collection.

Table 1–1 describes each interval.

Table 1–1 Data Collection Intervals

Interval (in seconds)	Description
Display	How often data should be collected as a foreground activity.
Event	How often data should be collected as a background activity if any events have been posted for that type of data.
NoEvent	How often data should be collected as a background activity if no events have been posted for that type of data.

The following list indicates how the Availability Manager determines which collection interval to use for a particular type of data:

- The Availability Manager starts background data collection at the NoEvent interval (for example, every 75 seconds). If no events have been posted for that type of data, the Availability Manager starts a new collection cycle every 75 seconds.
- If events have been posted for that type of data, the Availability Manager starts a new collection cycle at the Event interval. This rate of data collection continues until all events for that type of data have been removed.
- If there is a page open for a specific node, the Availability Manager starts a new collection cycle at the Display interval. This rate of data collection is used until the display is closed.

Overview

1.3 How Does the Availability Manager Identify Performance Problems?

1.3.2 Posting Events

The Availability Manager posts events when data values exceed user-defined thresholds and occurrences. Threshold and occurrence values are displayed on event customization pages similar to the one shown in Figure 1–6.

Figure 1–6 Sample Event Customization Page



1.3.2.1 Thresholds and Occurrences

The Availability Manager uses the **threshold** value as a criterion for posting an event. In many cases, if a condition exceeds that value, the Availability Manager displays a message in the Events pane of the Application window (see Figure 1–3). Some thresholds are used in more complex tests.

An **occurrence** (or trigger) for a specific event is the number of consecutive data collections that must exceed the event threshold before the Availability Manager signals the event in the Events pane of the Application window and logs it in the Event Log file.

For example, the disk status data that the Availability Manager collects includes the error count on a disk. If you select the Disk Status check box on the Data Collection page (Figure 1–4) and the error count exceeds the threshold value of 15 on the Event Customization page (Figure 1–6) for more than one data collection, an event is posted.

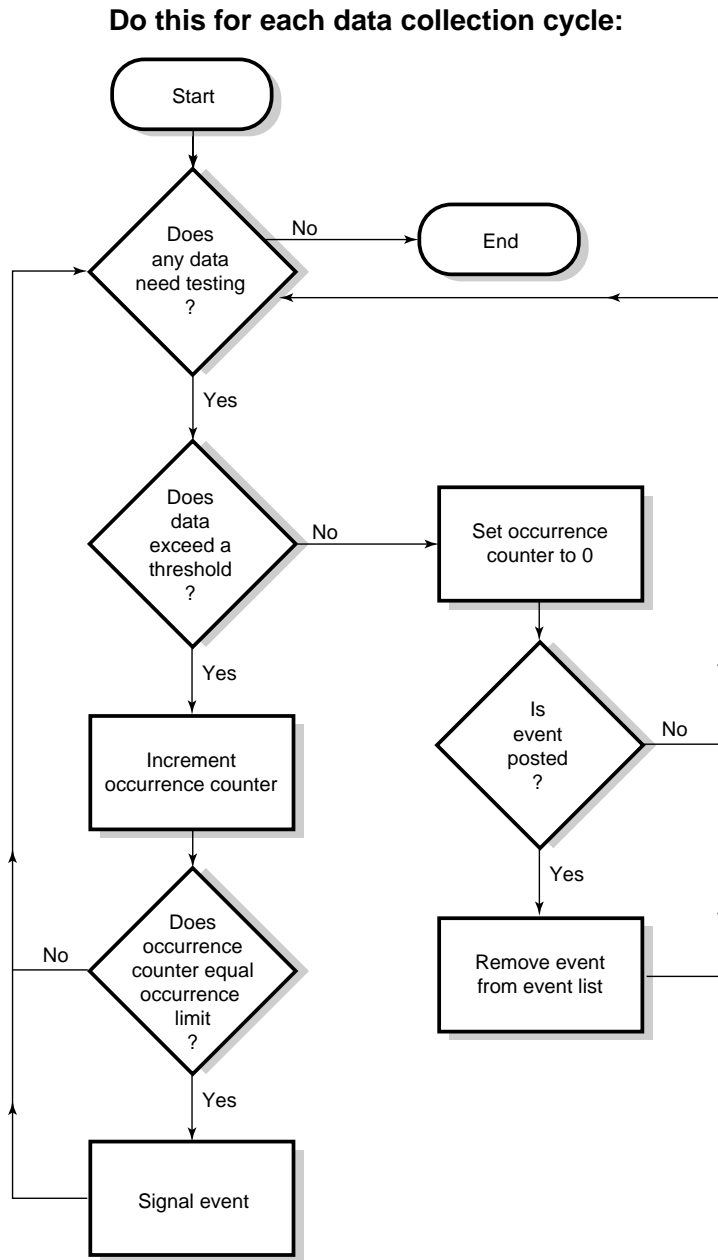
Chapter 6 explains how users can change default values for event thresholds and occurrences.

The Availability Manager evaluates every data collection for events. Any time a data value in a data collection exceeds a threshold, an occurrence counter is incremented. Whenever the occurrence count matches the Occurrence value on the Event Customization page (Figure 1–6), the event is signaled.

1.3 How Does the Availability Manager Identify Performance Problems?

If, at any time during data collection, the data does *not* exceed the threshold, the occurrence counter is set to 0, and the event is removed from the Events pane. Figure 1-7 depicts this sequence.

Figure 1-7 Flow Chart of Event Testing



VM-0480A-AI

Overview

1.4 How Does the Availability Manager Maintain Security?

1.4 How Does the Availability Manager Maintain Security?

The Availability Manager uses passwords to maintain security. These passwords have somewhat different appearances on Windows NT Data Analyzer nodes and on OpenVMS Data Analyzer and Data Collector nodes. On Windows NT Data Analyzer nodes, passwords are up to 8 characters long. On OpenVMS Data Analyzer and Data Collector nodes, passwords are part of a three-part security code called a **security triplet**.

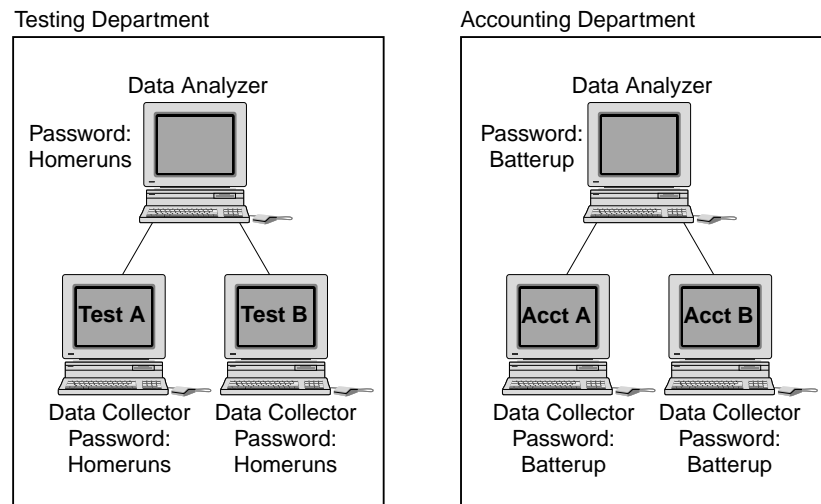
The following sections explain these security methods further.

1.4.1 Data Analyzer Password Security

For monitoring to take place, the password on a Data Analyzer node must match the password section of the security triplet on each OpenVMS Data Collector node. (A Windows NT Data Analyzer checks only the password part of each OpenVMS Data Collector security triplet. OpenVMS Data Collectors impose other security measures, which are explained in Section 1.4.2.)

Figure 1–8 illustrates how you can use passwords to limit access to node information. The Testing Department's Data Analyzer, whose password is HOMERUNS, can access only OpenVMS Data Collector nodes with the HOMERUNS password as part of their security triplets. The same is true of the Accounting Department Data Analyzer, whose password is BATTERUP; it can access only OpenVMS Data Collector nodes with the BATTERUP password as part of their security triplets.

Figure 1–8 Availability Manager Password Matching



VM-0100A-AI

The Availability Manager sets a default password when you install the Data Analyzer. To change that password, you must use a customization option that is explained in Chapter 6.

1.4 How Does the Availability Manager Maintain Security?

1.4.2 OpenVMS Data Collector Security Features

OpenVMS Data Collector nodes have the following security features:

- **Private LAN transport**

The Availability Manager protocol is based on the 802.3 Extended Packet Format (also known as SNAP). The IEEE Availability Manager protocol values are as follows:

```
Protocol ID:      08-00-2B-80-48
Multicast Address: 09-00-2B-02-01-09
```

If you filter protocols for bridges or routers in your network, you need to add these values to your network protocols.

- **Availability Manager data-transfer security**

Each OpenVMS node running as a Data Collector has a file containing a list of security triplets. For Data Analyzer and Data Collector nodes to exchange data, the passwords on the nodes must match.

In addition, note the following:

- Data Collector nodes that have read access allow Windows NT Data Analyzer nodes to view system data.
- Data Collector nodes that have write access also allow a matching Windows NT Data Analyzer node to perform fixes or modifications to the running system.

Chapter 6 explains security triplets and how to edit them.

- **OpenVMS file protection and process privileges**

When the Availability Manager is installed, it sets directory and file protections on its system-level directories so that only the SYSTEM account can read the files in these directories. For additional security on these system-level directories and files, you can create access control lists (ACLs) to restrict and set alarms on write access to the security files. For more information about creating ACLs, see the *OpenVMS Guide to System Security*.

Getting Started

Note

Before you start this chapter, be sure to read the explanation of data collection, events, thresholds, and occurrences in Chapter 1.

This chapter provides the following information:

- How to start the Availability Manager Data Collector
- How to start the Availability Manager Data Analyzer
- How to use the main Application window
- How to select a group of nodes (other than the default)
- How to display basic node data

For information about installing the Availability Manager on OpenVMS or Windows NT systems, refer to the appropriate installation instructions. You can access these instructions from the Availability Manager web pages at the following URL:

<http://www.openvms.compaq.com/openvms/products/availman/>

The installation instructions for OpenVMS systems also include an explanation of how to install and use both DECamds and the Availability Manager on the same system.

2.1 Starting the Data Collector

Beginning with OpenVMS Version 7.2, the files needed to run the Data Collector on OpenVMS nodes are shipped with the OpenVMS operating system. However, if you want the very latest Data Collector software, you need to install it from the Availability Manager Version 1.4 kit. (Prior versions of OpenVMS require that you install the software from either the DECamds Version 7.2-1B-1 or the Availability Manager Version 1.4 kit.)

Once the Data Collector is running on a node, you can monitor that node using DECamds or the Availability Manager.

To use the Data Collector on a particular node, do either of the following:

- Run the following command procedure:

```
$ @SYS$STARTUP:AMDS$STARTUP START
```

- Add the following command to the SYSTARTUP_VMS.COM command file in the SYSSMANAGER directory:

```
@SYS$STARTUP:AMDS$STARTUP START
```

Getting Started

2.2 Starting the Data Analyzer

2.2 Starting the Data Analyzer

This section describes what you need to do after the Availability Manager Data Analyzer is installed. Starting the Data Analyzer is somewhat different on OpenVMS and on Windows NT systems. However, on both systems, starting the Data Analyzer automatically starts the Java graphical user interface (GUI) that allows you to view information that is collected from the Data Collector nodes.

The following sections contain the sequence of steps required to start the Data Analyzer on an OpenVMS node and on a Windows NT node.

2.2.1 How to Start the Data Analyzer on an OpenVMS Alpha Node

Make sure the Data Analyzer has been installed on the OpenVMS Alpha node from which you want to monitor other nodes. To starting the Data Analyzer, perform the following steps:

1. Run the Java setup program:

```
$ @SYS$MANAGER:JAVA$SETUP
```

You do not need to run this program again until you log out of your process. The system displays an error message if you need to run the program.

2. Enter the following command:

```
$ avail/avail
```

The Availability Manager then displays the main Application window, which is shown in Figure 2-1.

Refer to the *Installation Instructions* for qualifiers you can use with the AVAIL/AVAIL command.

2.2.2 How to Start the Data Analyzer on a Windows NT Node

Make sure the Data Analyzer has been installed on the Windows NT node from which you want to monitor other nodes. The steps for starting the Data Analyzer are somewhat different, depending on whether you have Compaq OpenVMS Management Tools for Windows NT (OMT) installed or not. Follow the series of steps that apply to you.

Starting the Data Analyzer Without OMT Installed

To start the Data Analyzer, follow these steps:

1. Choose the following options beginning with the Windows NT Start menu:
Start —> **Programs**
2. Choose **Availability Manager**.
3. Choose **Data Analyzer Startup**.

The Availability Manager then displays the main Application window, which is shown in Figure 2-1.

Starting the Data Analyzer with OMT Installed

To start the Data Analyzer, follow these steps:

1. Choose the following options beginning with the Windows NT Start menu:
Start —> **Programs**
2. Choose **OpenVMS Management Tools**.
3. Choose **Availability Manager**.

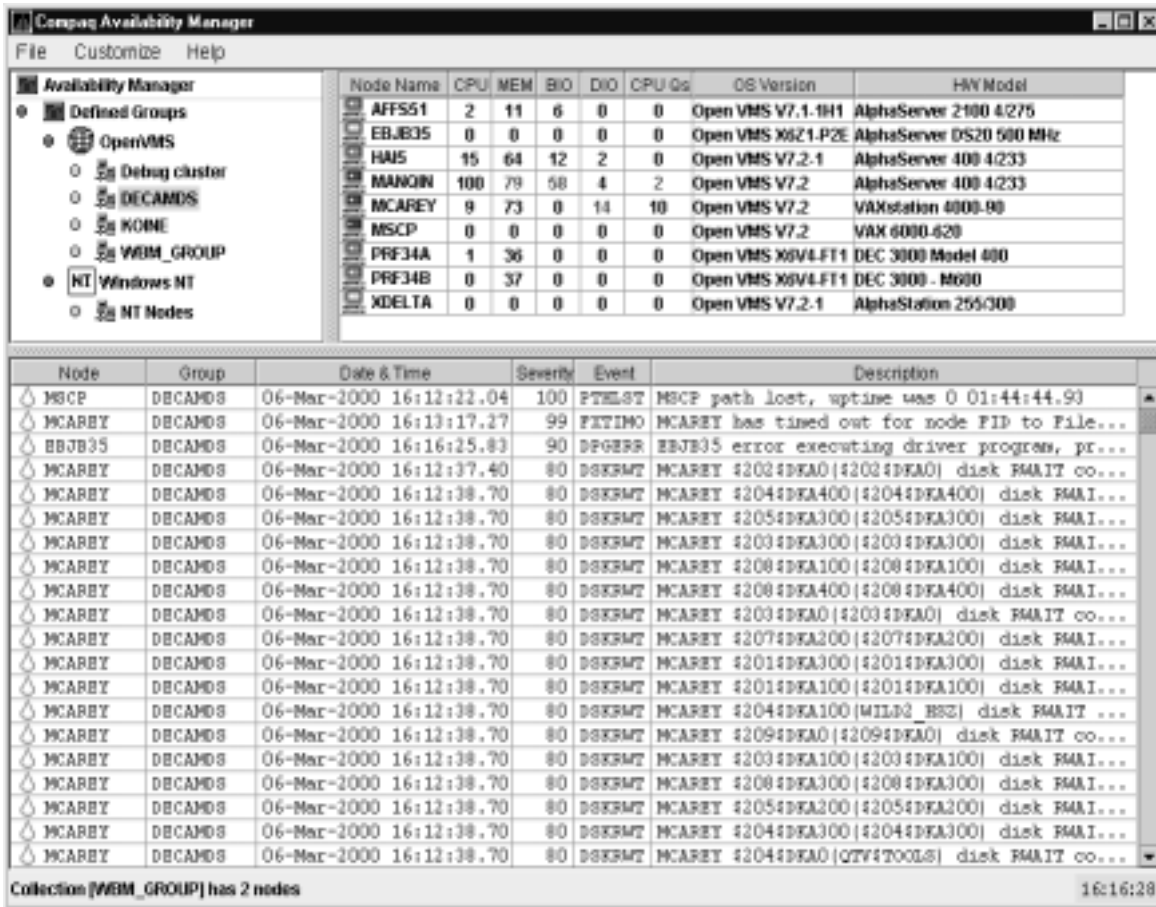
4. Choose **Data Analyzer Startup**.

The Availability Manager then displays the main Application window, which is shown in Figure 2–1.

2.3 Using the Application Window

Figure 2–1 shows the Availability Manager Application window.

Figure 2–1 Application Window



The Application window is divided into three segments, or **panes**:

- **Group pane**

The Group pane is at the top left of the window. In this pane, you select the group of nodes that you want to monitor.

- **Node pane**

The Node pane occupies most of the top right of the window. In this pane, the Availability Manager displays a list of the nodes in the group you have selected, along with summary data for each node. Chapter 3 contains more information about using the Node pane.

Figure 2–1 shows the node information that is displayed when you select a group of OpenVMS nodes. Somewhat different information is displayed for a group of Windows NT nodes. For more information about this, see Chapter 3.

Getting Started

2.3 Using the Application Window

- **Event pane**

The Event pane occupies the entire bottom of the window. In this pane, the Availability Manager displays events that occur on all of the nodes being monitored on your system. An event must reach a certain level of severity to be displayed. You can customize the severity levels at which events are displayed (see Chapter 6). For more information about displaying events, see Chapter 4.

You can change the size of the panes as well as the width of specific fields in the Application window and also the borders between the fields by clicking on a border and dragging it. Scroll bars indicate whether you are displaying all or part of a screen. For example, clicking a right arrow on a scroll bar allows you to view the rightmost portion of a screen.

2.3.1 Other Window Components

In addition to panes, the Application window also includes the following components (see Figure 2-1):

Title bar

The title bar runs across the top of the window and contains the heading **Compaq Availability Manager**.

Menu bar

The menu bar, immediately below the title bar, contains the following menu options:

- **File**

The File menu contains the Exit option, which allows you to stop the Data Analyzer and close the window.

- **Control**

The Control menu contains options to control the collection and display of data. These options are explained in Chapter 3.

- **Customize**

The Customize menu contains options that allow you to customize various aspects of the Availability Manager. These options are explained in Chapter 6.

- **Help**

The Help menu offers different types of online help for the Availability Manager. These options are explained in Section 2.6.

Status bar

The status bar runs across the bottom of the window. It displays the name of the selected group and the number of nodes in that group.

2.3.2 Displaying More Information

In the Application window, you can do the following at any time:

- Click a field to select it.
- Double-click most fields to display a page containing information specific to that field.
- Right-click a field to display a popup menu.

2.4 Selecting a Group of Nodes

To monitor nodes in the Application window, you must select the group that contains those nodes. You select groups in the Group pane (Figure 2–2).

Figure 2–2 Group Pane



Groups are set up during installation and are user definable. You might define groups by function, type of hardware, or geographical location.

For example, if you were to set up groups of nodes by geographical location, you might assign nodes A and B to a group called Dallas and nodes C, D, and E to a group called Denver. When you select a group, the Availability Manager displays only the nodes in that group, as shown in the following table:

Group Selected	Nodes Displayed
Dallas	Node A Node B
Denver	Node C Node D Node E

Compaq recommends that you define a cluster as its own group.

2.4.1 Group Names That Are Displayed by Default

When you start the Availability Manager, the names of groups of nodes that have the Availability Manager Data Collector installed are displayed. By default, the group name of the first node detected is highlighted. (A list of the nodes in that group is displayed in the Node pane.)

Under the Availability Manager heading in the Group pane is a list of one group—or possibly the following two:

- OpenVMS
These are groups in an OpenVMS system group that you can monitor.
- Windows NT
These are groups in a Windows NT system group that you can monitor.

Getting Started

2.4 Selecting a Group of Nodes

2.4.2 How to Display Other Nodes

To display the names of nodes other than the nodes in the default group, click the name of that group in the Group pane (see Figure 2–2). In the Node pane of the Application window (see Figure 2–1), the Availability Manager displays the nodes in the group you have selected.

If you want to display the names of more groups in the Group pane, however, you need to use a customization option. See Section 6.1 for instructions.

2.5 Displaying Information about Nodes

The Node pane of the Application window allows you to focus on resource usage activity at a high level and to display more specific data whenever you want. This section explains the basic use of the Node pane. For more detailed information, see Chapter 3.

Within the group of nodes you select, the Availability Manager displays all the nodes with which that group can communicate. Figure 2–3 shows a list of OpenVMS nodes.

Figure 2–3 Node Pane

Node Name	CPU	MEM	BIO	DIO	CPU Qs	OS Version	HW Model
CMOVEQ	0	36	0	0	0	Open VMS V7.2-1	AlphaStation 200 4/166
DBGAVC	3	19	23	76	0	Open VMS V7.1-1H2	AlphaServer 2100A 5/300
DIRNKS	0	44	1	0	0	Open VMS V7.1	VAXstation 4000-90
LOWFAT	0	38	0	0	0	Open VMS V6.2-1H3	AlphaServer 1000 4/200
MAWK	14	81	2	166	0	Open VMS V7.1	VAXstation 4000-60
PEROIT	1	31	1	0	0	Open VMS V7.1	VAXstation 4000-60
REDSOL	0	33	0	0	0	Open VMS V7.1-1H2	Digital Personal WorkStation
SELAVY	1	23	0	0	0	Open VMS V6.2	VAXstation 4000-60
SSPADE	6	51	3	0	0	Open VMS V7.1	VAXstation 4000-90

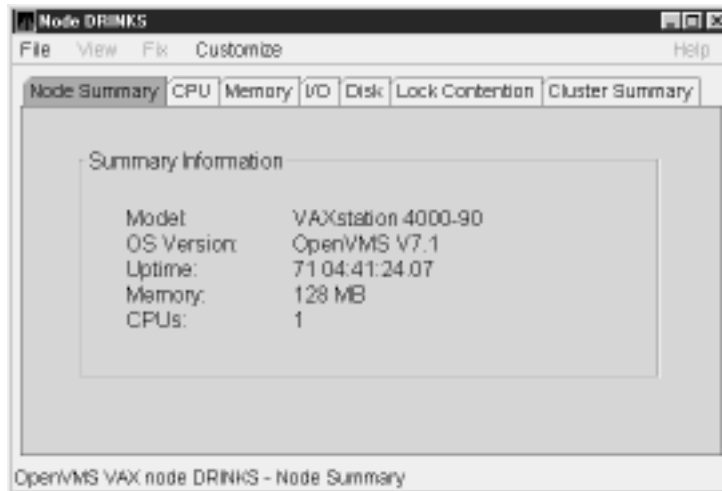
Each node name has an icon next to it. The icon colors represent the following:

- Red Security check was successful. However, a threshold has been exceeded (noted in red), and an event has been posted.
- Yellow Node security check is in progress or has failed.
- Green Security check was successful; data is being collected.
- Black Path to node has been lost.

2.5.1 Choosing a Node

To choose a node on a Windows NT or an OpenVMS system, double-click a node name in the Node pane to highlight the name of the node and display the Node Summary page (Figure 2–4). Alternatively, you can right-click a node name to display a popup menu; to display the Node Summary page, select the **Display...** option.

Figure 2-4 OpenVMS Node Summary Page



The data displayed on this page is explained in detail in Chapter 3.

At the top of the Node Summary page are tabs that correspond to most of the types of node data displayed in the Node pane. When you click a tab on the Node Summary page, the Availability Manager displays most of the same pages that are displayed when you double-click a field in the Node pane (see Figure 2-3).

2.5.2 Specifying Data to Collect on OpenVMS Nodes

For OpenVMS nodes, if you want background data collection (and the associated event detection), you must **turn on** data collection for each type of data you want to collect. On Windows NT nodes, background data collection is always enabled and cannot be turned off.

To turn on various types of data collection, follow these steps:

1. In the Application window, click the **Customize** menu.
2. Click **Customize OpenVMS...**
3. Click the **Data Collection** tab.

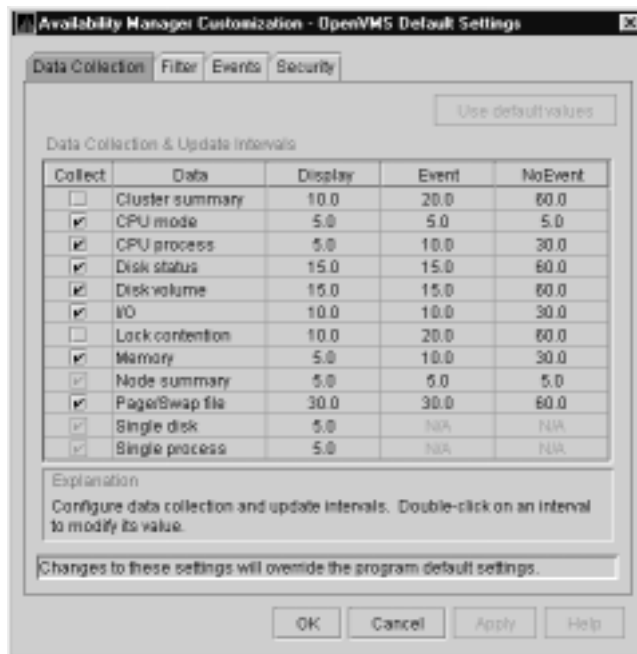
The Availability Manager displays the Data Collection page (Figure 2-5).

Background and foreground collection are explained in detail in Chapter 1.

Getting Started

2.5 Displaying Information about Nodes

Figure 2–5 Data Collection Page



The following types of data are collected by default:

- Node summary
- Single disk
- Single process

For example, to stop collecting node summary data, clear the check box for “Node summary” in the “Collect” column. You cannot, however, turn off the collection of single disk and single process data. These types of data are collected by default when you open a Single Disk Summary page or a Process Information page, respectively.

To turn on a type of data collection, select the check box for that type in the “Collect” column. Table 6–1 identifies the page where each type of data that is collected appears.

On the Data Collection page, you can also change the intervals at which data is collected. Collection intervals explained in Chapter 6.

2.5.3 Sorting Data

You can sort data in the following OpenVMS displays:

- Events pane of the Application window (Figure 1–3)
- CPU Process Summary pane (Figure 3–7)
- Memory page (Figure 3–9)
- Bottom pane of I/O Summary page (Figure 3–10)
- Disk Status Summary page (Figure 3–12)
- Disk Volume Summary page (Figure 3–14)

Getting Started

2.5 Displaying Information about Nodes

Depending on the field, you can sort data alphabetically or numerically. An alphabetical sort is performed using ASCII character values; for example, dollar signs (\$) precede letters in the sort order.

To sort the values in a field, double-click the appropriate column heading. To reverse the sort order, double-click the column heading again.

2.6 Getting Help

To obtain online help, click the **Help** menu on the Application window menu bar. Then choose one of the following options:

Menu Option	Description
Availability Manager Help	Information about using the Availability Manager.
Getting Started	A special online version of help for getting started using this tool.
Availability Manager Release Notes	Last-minute information about the software and how it works.
About Availability Manager...	Information about this Availability Manager release (such as the version number).

Getting Information about Nodes

Note

Before you start this chapter, be sure to read the explanation of data collection, events, thresholds, and occurrences in Chapter 1.

This chapter describes the node data that the Availability Manager displays by default and more detailed data that you can choose to display. Node summary data is the only data displayed by default. (Differences are noted whenever information displayed for OpenVMS nodes differs from that displayed for Windows NT nodes.)

On OpenVMS systems, you can turn on and off the collection of individual types of data, and you can specify several interval times for the collection and display of events. These are discussed in Chapter 6.

3.1 Node Panes

After you select a group of nodes in the Group pane, the Availability Manager automatically displays data for each node within that group on the Node pane of the Application window (Figure 1–3). The following sections describe the data displayed for OpenVMS and Windows NT Node panes.

3.1.1 OpenVMS Node Pane

Figure 3–1 shows an example of an OpenVMS Node pane.

Figure 3–1 OpenVMS Node Pane

Node Name	CPU	MEM	BIO	DIO	CPU Qs	OS Version	HW Model
CMOVEQ	0	36	0	0	0	Open VMS V7.2-1	AlphaStation 200 4/166
DBGAVC	3	19	23	76	0	Open VMS V7.1-1H2	AlphaServer 2100A 5/300
DRINKS	0	44	1	0	0	Open VMS V7.1	VAXstation 4000-90
LOWFAT	0	38	0	0	0	Open VMS V6.2-1H3	AlphaServer 1000 4/200
MAWK	14	81	2	166	0	Open VMS V7.1	VAXstation 4000-60
PEROIT	1	31	1	0	0	Open VMS V7.1	VAXstation 4000-60
REDSQL	0	33	0	0	0	Open VMS V7.1-1H2	Digital Personal WorkStation
SELAVY	1	23	0	0	0	Open VMS V6.2	VAXstation 4000-60
SSPADE	6	51	3	0	0	Open VMS V7.1	VAXstation 4000-90

Data displayed in red on your screen indicates that the amount is above the threshold set for that field.

For each OpenVMS node and group it recognizes, the Availability Manager displays the data shown in Table 3–1. The table lists the event related to each type of data; events that exceed either the default thresholds or the defaults that you set are displayed in red on the screen.

Getting Information about Nodes

3.1 Node Panes

Refer to Section 6.5 for information about setting event thresholds. Appendix B describes OpenVMS and Windows NT events.

Table 3–1 OpenVMS Node Data

Data	Description of Data	Related Event
Node Name	Name of the node being monitored.	n/a
CPU	Percentage of CPU usage of all processes on the node.	n/a
MEM	Percentage of space in memory that all processes on the node use.	LOMEMORY
BIO	Buffered I/O rate of processes on the node.	HIBIOR
DIO	Direct I/O usage of processes on the node.	HIDIOR
CPU Qs	Number of processes in one of the following states: MWAIT, COLPG, PFW, FPG.	n/a
OS Version	Version of the operating system on the node.	n/a
Hardware Model	Hardware model of the node.	n/a

3.1.2 Windows NT Node Pane

Figure 3–2 shows a sample Windows NT Node pane. From the group selected, the Availability Manager displays all the nodes with which it can communicate.

Figure 3–2 Windows NT Node Pane

Node Name	CPU	MEM	DIO	Processes	Threads	Events	Semaphores	Mutexes	Sections	OS Version	HW Model
PYROMAN	1	61	1	13	125	273	97	9	129	Windows NT 4.0	DEC-321064
STELLA	1	50	0	20	168	354	96	19	213	Windows NT 4.0	DEC-321064
UG1996	1	80	0	97	152	464	68	19	203	Windows NT 4.0	DEC-321064

For each Windows NT node in the group you select, the Availability Manager displays the data shown in Table 3–2.

Table 3–2 Windows NT Node Data

Data	Description
Node Name	Name of the node being monitored.
CPU	Percentage of CPU usage of all the processes on the node.
MEM	Percentage of memory that is in use.
DIO	Direct I/O usage of processes on the node.
Processes	Number of processes on the node.
Threads	Number of threads on the node. A thread is a basic executable entity that can execute instructions in a processor.
Events	The number of events on the node. An event is used when two or more threads want to synchronize execution.

(continued on next page)

Table 3–2 (Cont.) Windows NT Node Data

Data	Description
Semaphores	The number of semaphores on the node. Threads use semaphores to control access to data structures that they share with other threads.
Mutexes	The number of mutexes on the node. Threads use mutexes to ensure that only one thread executes a section of code at a time.
Sections	The number of sections on the node. A section is a portion of virtual memory created by a process for storing data. A process can share sections with other processes.
OS Version	Version of the operating system on the node.
HW Model	Hardware model of the node.

3.2 Node Data Pages

The following sections describe node data pages, which you can display in either of two ways:

- Double-click a data item in the Node pane to display an associated page.
- Double-click a node name on the Node pane to display the Node Summary page (Figure 3–3). You can then click other tabs on the Node Summary page to display the same detailed data that you display by double-clicking a data item in the Node pane.

The menu bar on each node data page contains the options described in Table 3–3.

Table 3–3 Node Data Page Menu Bar

Menu Option	Description	For More Information
File	Contains the Close option, which you can choose to exit from the pages.	n/a
View	Contains options that allow you to view data from another perspective.	See specific pages.
Fix	Contains options that allow you to resolve various resource availability problems and improve system performance.	Chapter 5
Customize	Contains options that allow you to organize data collection and analysis and to display data by filtering and customizing Availability Manager data.	Chapter 6

The following sections describe these node data pages:

Node Data Page (or Pages)	Reference
Node Summary	Section 3.2.1
CPU Modes and Process Summaries	Section 3.2.2
Memory Summaries	Section 3.2.3

Getting Information about Nodes

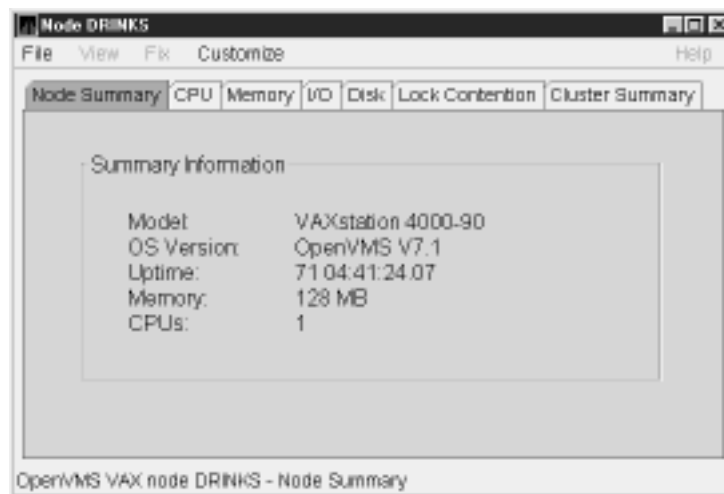
3.2 Node Data Pages

Node Data Page (or Pages)	Reference
OpenVMS I/O Summary and Page/Swap Files	Section 3.2.4
Disk Summaries	Section 3.2.5
OpenVMS Lock Contention	Section 3.2.6
OpenVMS Cluster Summary	Section 3.2.7
OpenVMS Network Interconnect for the System Communications Architecture (NISCA)	Section 3.2.8
OpenVMS Single Process	Section 3.2.9

3.2.1 Node Summary

When you double-click a node name, operating system (OS) version, or hardware model in an OpenVMS or Windows NT Node pane, the Availability Manager displays the Node Summary page (Figure 3–3).

Figure 3–3 Node Summary Page



On this page, the following information is displayed for the node selected:

Data	Description
Model	System hardware model name.
OS Version	Name and version of the operating system.
Uptime	Time (in days, hours, minutes, and seconds) since the last reboot.
Memory	Total amount of physical memory (in megabytes) found on the system.
CPUs	Number of active CPUs on the node.

3.2.2 CPU Modes and Process Summaries

By clicking the **CPU** tab, you can display CPU panes that contain more detailed statistics about CPU mode usage and process summaries than the Node Summary does. You can use the CPU panes to diagnose issues that CPU-intensive users or CPU bottlenecks might cause. For OpenVMS nodes, you can also display information about specific CPU processes.

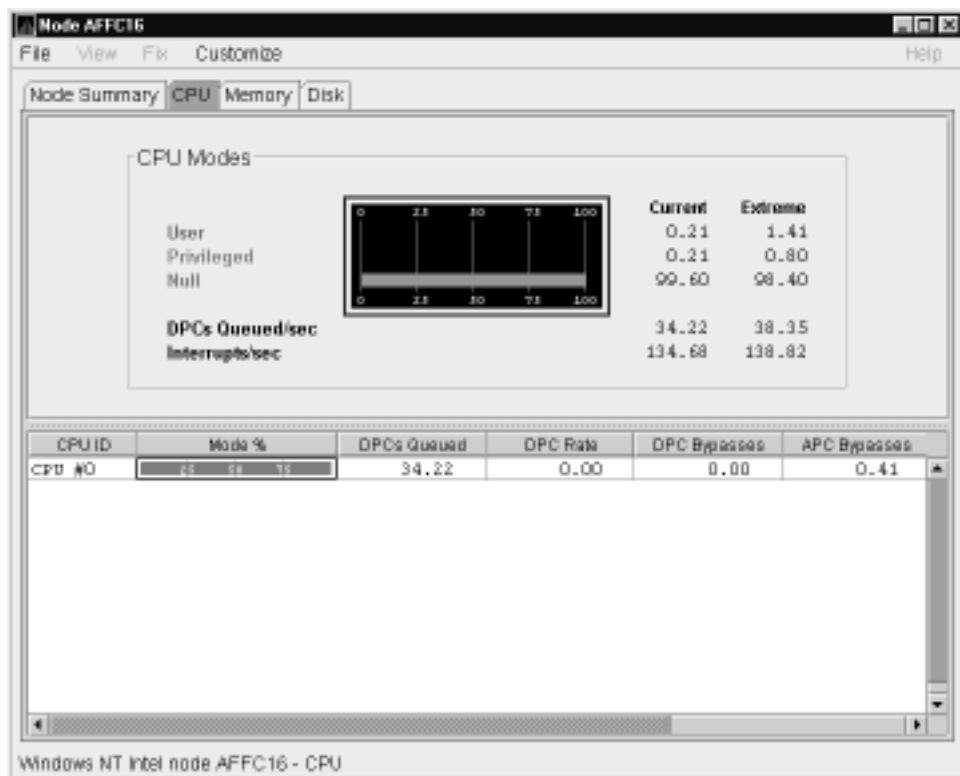
When you double-click a value under the CPU or CPU Qs heading on either an OpenVMS or a Windows NT Node pane, or when you click the **CPU** tab, the Availability Manager displays the CPU Modes Summary in the top pane and, by default in the bottom pane, CPU Modes Detail. You can use the **View** menu to select the CPU Process Summary in the bottom pane (see Section 3.2.2.4).

CPU modes summaries and process summary panes are described in the following sections. Note that there are differences between the pages displayed for OpenVMS and Windows NT nodes.

3.2.2.1 Windows NT CPU Modes

Figure 3–4 shows a sample Windows NT CPU Modes page.

Figure 3–4 Windows NT CPU Modes Page



The top pane of the Windows NT CPU Modes page is a summary of Windows NT CPU usage, listed by type of mode.

On the left, the following CPU modes are listed:

- User
- Privileged
- Null

On the graph, values that exceed thresholds are displayed in red. To the right of the graph are current and extreme amounts for each mode.

Current and extreme amounts are also displayed for the following values:

- Deferred procedure calls (DPCs) queued per second

Getting Information about Nodes

3.2 Node Data Pages

- Interrupts occurring per second

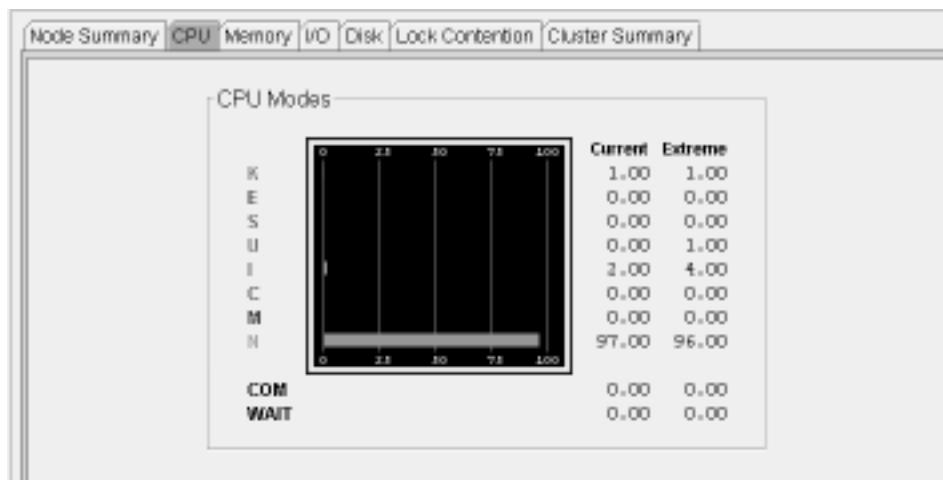
The bottom pane of the Windows NT CPU Modes contains modes details. The following data is displayed:

Data	Description
CPU ID	Decimal value representing the identity of a processor in a multiprocessing system. On a uniprocessor, this value is always CPU #00.
Mode %	Graphical representation of the percentage of active modes on that CPU. The color displayed matches the mode color on the graph on the top pane.
DPCs Queued	Rate that deferred procedure call (DPC) objects are queued to this processor's DPC queue.
DPC Rate	Average rate that DPC objects are queued to this processor's DPC queue per clock tick.
DPC Bypasses	Rate that dispatch interrupts were short-circuited.
APC Bypasses	Rate that kernel asynchronous procedure call (APC) interrupts were short-circuited.

3.2.2.2 OpenVMS CPU Modes Summary

Figure 3-5 shows a sample OpenVMS CPU Modes summary, which is the top pane of the CPU Modes page.

Figure 3-5 OpenVMS CPU Modes Summary Pane



The CPU Modes summary pane shows OpenVMS CPU usage by type of mode. On symmetric multiprocessing (SMP) nodes, percentages are averaged across all the CPUs and are displayed as a single value.

The following table shows the alphabetic symbols that represent CPU modes:

Symbol	Mode
K	Kernel
E	Executive

Symbol	Mode
S	Supervisor
U	User
I	Interrupt
C	Compatibility
M	Multiprocessor synchronization
N	Null

The bars in the graph represent the percentage of CPU cycles used. Values that are lower than the thresholds are displayed in green; values that exceed thresholds are displayed in red. To the right of the graph are current and extreme amounts for each mode.

Below the graph, the Availability Manager displays the state of certain CPU process queues. The amounts displayed for COM are the sums of the queue lengths of processes in the COM and COMO states. The amounts displayed for WAIT are the sums of the queue lengths of processes in the miscellaneous WAIT, COLPG, CEF, PFW, and FPG states. (See Appendix A for explanations of CPU process states.)

3.2.2.3 OpenVMS CPU Modes Detail

The bottom pane of the CPU Modes page contains CPU modes details, as shown in Figure 3–6.

Figure 3–6 OpenVMS CPU Modes Detail Pane

CPU ID	State	Mode %	PID	Process Name	Capabilities
CPU #000	Run		2020016F	CTMI_00010001	PRIMARY RUN QUORUM
CPU #002	Run		20200174	CTMI_00040001	RUN QUORUM

In the OpenVMS CPU Modes Detail pane, the following data is displayed:

Data	Description
CPU ID	Decimal value representing the identity of a processor in a multiprocessing system. On a uniprocessor, this value is always CPU #00.
State	One of the following CPU states: Boot, Booted, Init, Rejected, Reserved, Run, Stopped, Stopping, or Timeout.
Mode %	Graphical representation of the percentage of active modes on that CPU. The color displayed coincides with the mode color on the graph on the top pane.
PID	Process identifier (PID) value of the process that is using the CPU. If the PID is unknown to the console application, the internal PID (IPID) is listed.
Process Name	Name of the process active on the CPU. If no active process is found on the CPU, the name is listed as *** None ***.
Capabilities	One or more of the following CPU capabilities: Primary, Quorum, Run, or Vector.

Getting Information about Nodes

3.2 Node Data Pages

3.2.2.4 OpenVMS CPU Process Summary

To display the OpenVMS CPU Process Summary pane at the bottom of the CPU page, select CPU Process Summary from the **View** menu. Figure 3–7 shows a sample OpenVMS CPU Process Summary pane.

Figure 3–7 OpenVMS CPU Process Summary Pane

PID	Process Name	Priority	State	Rate	Wait	Time
000002B2	TCP/IP\$RFX_B6440	6/ 4	LEP	0.00	0.00	0 00:00:10.85
00000065	DECW\$TR_0065	4/ 4	LEP	0.00	0.00	0 00:10:03.47
000001A7	FTA18:	5/ 4	LEP	0.00	0.00	0 01:42:37.31
00000068	VUB\$SYSTEM_11	6/ 4	LEP	0.00	0.00	0 00:33:22.86
00000069	VUB\$SYSTEM_12	4/ 4	LEP	0.00	0.00	0 00:11:57.39
0000006A	FTA13:	7/ 4	HIB	0.00	0.00	0 00:00:03.44
0000006B	FTA14:	7/ 4	HIB	0.00	0.00	0 00:00:01.84
0000006E	FTA16:	8/ 4	HIB	0.00	0.00	0 00:00:12.48
00000331	VUB\$SYSTEM_10	6/ 4	LEP	1.25	0.00	0 02:37:07.49
000001F4	FTA19:	4/ 4	HIB	0.00	0.00	0 00:00:01.80
000002B5	VUB\$SYSTEM_6	4/ 4	LEP	0.32	0.00	4 04:20:40.62
000000F6	TCP/IP\$PORTM	10/ 8	LEP	0.00	0.00	0 00:00:00.64
00000077	FTA17:	0/ 0	CUR	82.28	0.00	21 07:19:07.30
000003B8	DEA400CACP	10/ 8	HIB	0.00	0.00	0 00:00:00.32

The OpenVMS CPU Process Summary pane displays the following data:

Data	Description
PID	Process identifier, a 32-bit value that uniquely identifies a process.
Process Name	Name of the process active on the CPU.
Priority	Computable (<i>xx</i>) and base (<i>yy</i>) process priority in the format <i>xx/yy</i> .
State	One of the process states listed in Appendix A.
Rate	Percentage of CPU time used by this process. This is the ratio of CPU time to elapsed time. The CPU rate is also displayed in the bar graph.
Wait	Percentage of time the process is in the COM or COMO state.
Time	Amount of actual CPU time charged to the process.

When you double-click a PID on the lower part of an OpenVMS CPU Process (Figure 3–7), Memory (Figure 3–9), or I/O (Figure 3–10) page, the Availability Manager displays an OpenVMS Single Process, where you can click tabs to display specific data about one process.

This data includes a combination of data elements from the CPU Process, Memory, and I/O pages, as well as data for specific quota utilization, current image, and queue wait time. These pages are described in Section 3.2.9.

3.2.3 Memory Summaries

The Memory summary pages contain statistics about memory usage on the node you select. The Memory summary pages displayed for OpenVMS and Windows NT nodes are somewhat different, as described in the following sections.

Getting Information about Nodes

3.2 Node Data Pages

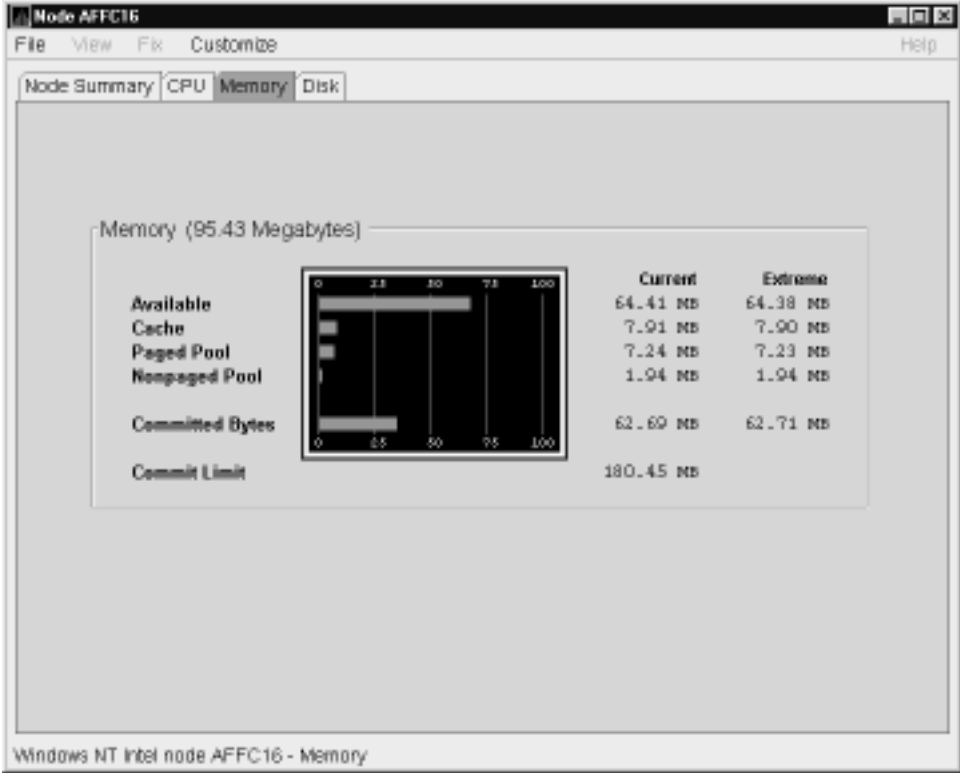
3.2.3.1 Windows NT Memory Summary

You can display the Windows NT Memory in either of the following ways:

- Double-click a node, and then click the **Memory** tab.
- Double-click a value under the MEM heading.

The Availability Manager then displays the Windows NT Memory page (Figure 3–8).

Figure 3–8 Windows NT Memory Page



The Current and Extreme amounts on the page display the data shown in the following table. The table also indicates what the graph amounts represent.

Data	Description
Available	Size (in bytes) of the virtual memory currently on the zeroed, free, and standby lists. Zeroed and free memory are ready for use, with zeroed memory cleared to zeros. Standby memory is removed from a process's working set but is still available. The graph shows the percentage of physical memory that is available for use.
Cache	Number of bytes currently in use by the system cache. The system cache is used to buffer data retrieved from disk or LAN. The system cache uses memory not in use by active processes on the computer. The graph shows the percentage of physical memory devoted to the cache.

Getting Information about Nodes

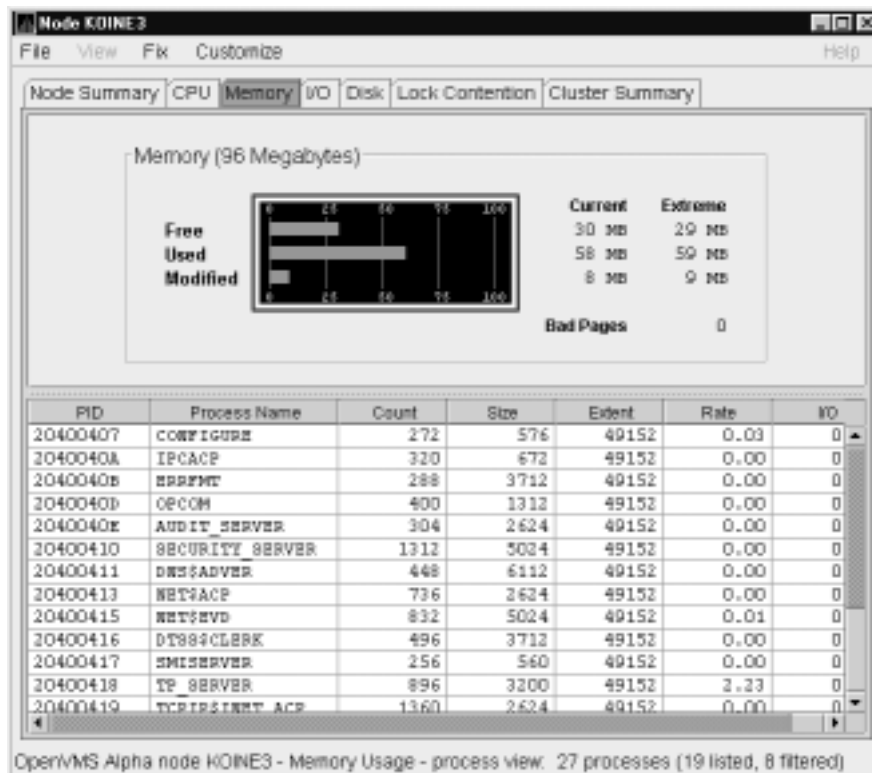
3.2 Node Data Pages

Data	Description
Paged Pool	Number of bytes in paged pool, a system memory area where operating system components acquire space as they complete their tasks. Paged pool pages can be paged out to the paging file when the system does not access them for long periods of time. The graph shows the percentage of physical memory devoted to paged pool.
Nonpaged Pool	Number of bytes in nonpaged pool, a system memory area where operating system components acquire space as they complete their tasks. Nonpaged pool pages cannot be paged out to the paging file; instead, they remain in memory as long as they are allocated. The graph shows the percentage of physical memory devoted to nonpaged pool.
Committed Bytes	Amount of available virtual memory (the Commit Limit) that is in use. Note that the commit limit can change if the paging file is extended. The graph shows the percentage of the Commit Limit used by the Committed Bytes.
Commit Limit	Size (in bytes) of virtual memory that can be committed without having to extend the paging files. If the paging files can be extended, this limit can be raised.

3.2.3.2 OpenVMS Memory Summary

When you double-click a value under the MEM heading in an OpenVMS Node pane, or if you click the **Memory** tab, the Availability Manager displays the OpenVMS Memory page (Figure 3–9).

Figure 3–9 OpenVMS Memory Page



Getting Information about Nodes

3.2 Node Data Pages

The graph in the top pane shows memory distribution (Free, Used, Modified, Bad) as absolute values, in megabytes of memory. Current and extreme values are also listed for each type of memory distribution. (Free memory uses the lowest seen value as its extreme.)

The bottom pane displays the following data:

Data	Description
PID	Process identifier. A 32-bit value that uniquely identifies a process.
Process Name	Name of the process.
Count	Number of physical pages or pagelets of memory that the process is using for the working set count.
Size	Number of pages or pagelets of memory the process is allowed to use for the working set size. The operating system periodically adjusts this value based on an analysis of page faults relative to CPU time used.
Extent	Number of pages or pagelets of memory in the process's working set extent (WSEXTENT) quota as defined in the user authorization file (UAF). Number of pages or pagelets cannot exceed the value of the system parameter WSMAX.
Rate	Number of page faults per second for the process.
I/O	Rate of I/O read attempts necessary to satisfy page faults (also known as page read I/O or the hard fault rate).

When you double-click a PID on the lower part of an OpenVMS CPU Process (Figure 3-7), Memory (Figure 3-9), or I/O (Figure 3-10) page, the Availability Manager displays an OpenVMS Single Process, where you can click tabs to display specific data about one process. This data includes a combination of data from the CPU Process, Memory, and I/O pages, as well as data for specific quota utilization, current image, and queue wait time. These pages are described in Section 3.2.9.

3.2.4 OpenVMS I/O Summaries and Page/Swap Files

By clicking the **I/O** tab on any OpenVMS node data page, you can display a page that contains summaries of accumulated I/O rates. In the top pane, this summary covers all CPUs; in the bottom pane, this summary is per process.

From the **View** menu, you can also choose to display (in the bottom pane) a list of page and swap files.

3.2.4.1 OpenVMS I/O Summaries

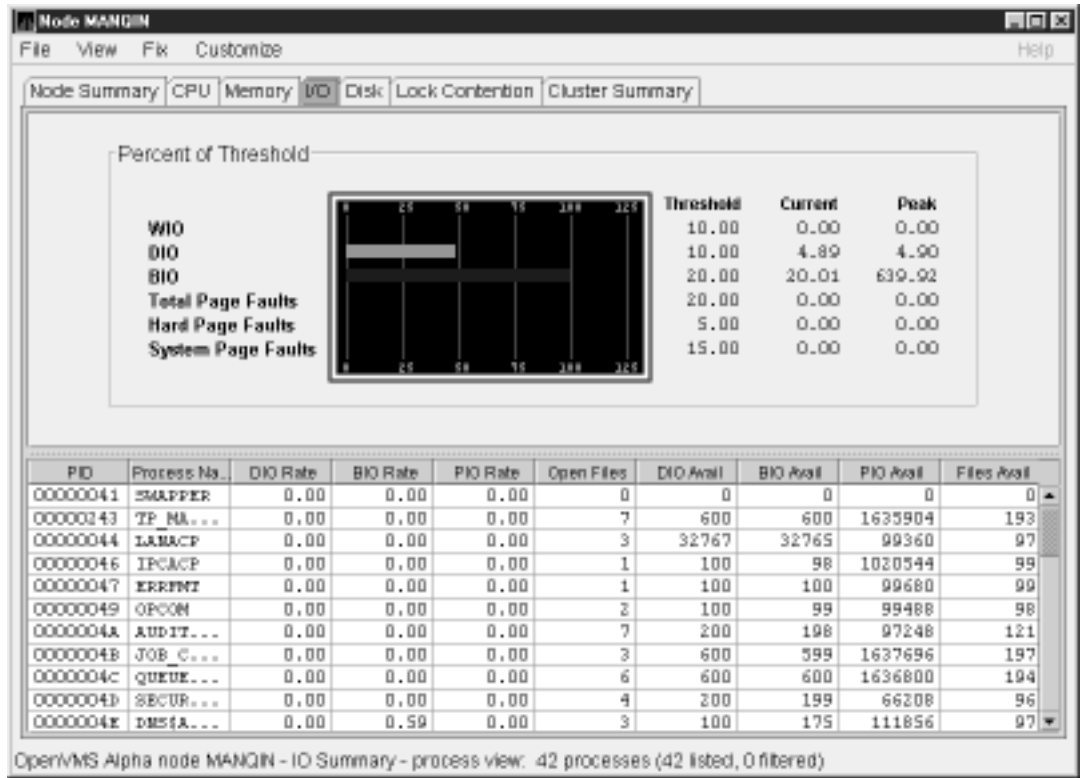
The OpenVMS I/O Summaries display the rate (per second) at which I/O transfers take place, including paging write I/O (WIO), direct I/O (DIO), and buffered I/O (BIO). In the top pane, the summary is for all CPUs; in the bottom pane, the summary is for one process.

When you double-click a data item under the DIO or BIO heading on the Node pane, or if you click the **I/O** tab, the Availability Manager displays, by default, the OpenVMS I/O Summaries (Figure 3-10).

Getting Information about Nodes

3.2 Node Data Pages

Figure 3–10 OpenVMS I/O Summaries Page



The graph in the top pane represents the percentage of thresholds for the types of I/O shown in Table 3–4. The table also shows the event that is related to each data item. Refer to Section 6.5 for information about setting event thresholds.

Table 3–4 I/O Data Displayed

Type of I/O	I/O Description	Related Event
Paging Write I/O Rate (WIO)	Rate of write I/Os to one or more paging files.	HIPWIO
Direct I/O Rate (DIO)	Transfers are from the pages or pagelets containing the process buffer that the system locks in physical memory to the system devices.	HIDIOR
Buffered I/O Rate (BIO)	Transfers are for the process buffer from an intermediate buffer from the system buffer pool.	HIBIOR
Total Page Faults	Total of hard and soft page faults on the system, as well as peak values seen during an Availability Manager session.	HITLTP
Hard Page Faults	Total of hard page faults on the system.	HIHRDP
System Page Faults	Page faults that are taken from kernel processes.	HISYSP

In addition, current and peak values are listed for each type of I/O. Values that exceed thresholds set by the events indicated in the table are displayed in red on the screen. Appendix B describes OpenVMS and Windows NT events.

Getting Information about Nodes

3.2 Node Data Pages

To the right of the graph, the following values are listed:

Value	Description
Threshold	Defined in Event Properties.
Current	Current value or rate.
Peak	Highest value of rate seen since start of data collection.

The bottom pane displays summary accumulated I/O rates on a per-process basis. The following data is displayed:

Data	Description
PID	Process identifier. A 32-bit value that uniquely identifies a process.
Process Name	Name of the current process.
DIO Rate	Direct I/O rate. The rate at which I/O transfers occur between the system devices and the pages or pagelets that contain the process buffer that the system locks in physical memory.
BIO Rate	Buffered I/O rate. The rate at which I/O transfers occur between the process buffer and an intermediate buffer from the system buffer pool.
PIO Rate	Paging I/O rate. The rate of read attempts necessary to satisfy page faults (also known as page read I/O or the hard fault rate).
Open Files	Number of open files.
DIO Avail	Direct I/O limit remaining. The number of remaining direct I/O limit operations available before the process reaches its quota. DIOLM quota is the maximum number of direct I/O operations a process can have outstanding at one time.
BIO Avail	Buffered I/O limit remaining. The number of remaining buffered I/O operations available before the process reaches its quota. BIOLM quota is the maximum number of buffered I/O operations a process can have outstanding at one time.
BYTLM	The number of buffered I/O bytes available before the process reaches its quota. BYTLM is the maximum number of bytes of nonpaged system dynamic memory that a process can claim at one time.
Files	Open file limit remaining. The number of additional files the process can open before reaching its quota. The FILLM quota is the maximum number of files that can be opened simultaneously by the process, including active network logical links.

When you double-click a PID on the lower part of an OpenVMS CPU Process (Figure 3-7), Memory (Figure 3-9), or I/O (Figure 3-10) page, the Availability Manager displays an OpenVMS Single Process, where you can click tabs to display specific data about one process. This data includes a combination of data from the CPU Process, Memory, and I/O pages, as well as data for specific quota utilization, current image, and queue wait time. These pages are described in Section 3.2.9.

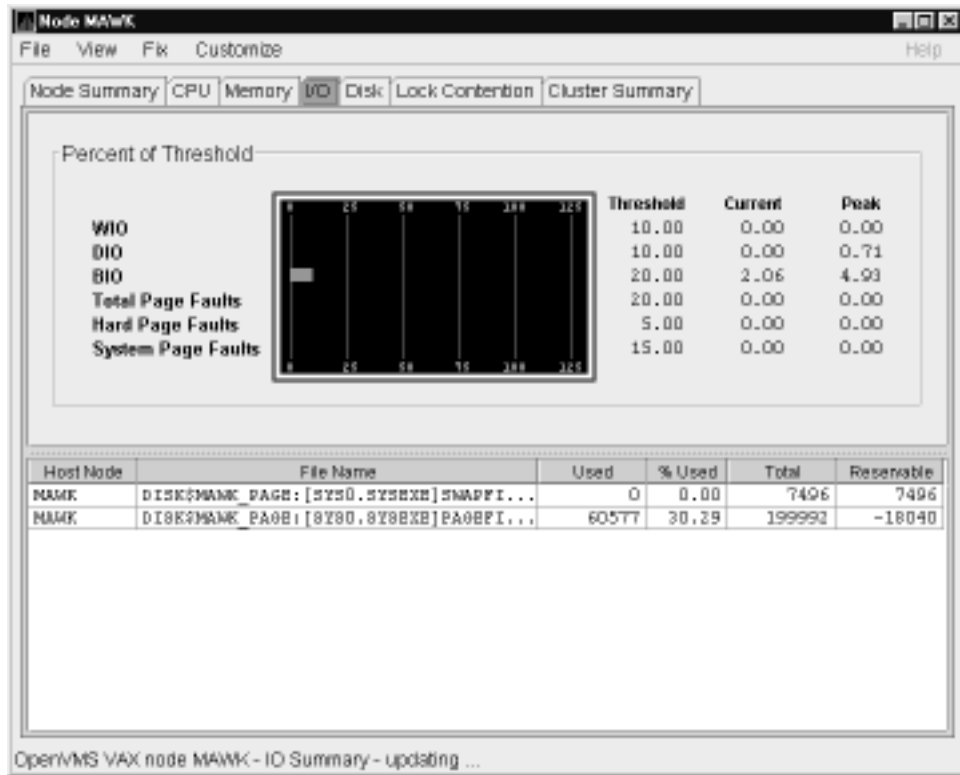
Getting Information about Nodes

3.2 Node Data Pages

3.2.4.2 OpenVMS I/O Page/Swap Files

Click I/O Page/Swap Files on the I/O page **View** menu to select this option. The Availability Manager displays an OpenVMS I/O Page/Swap Files page (Figure 3–11).

Figure 3–11 OpenVMS I/O Page/Swap Files Page



The top pane displays the same information displayed in the OpenVMS I/O Summary (Figure 3–10). The bottom pane displays the following data:

Data	Description
Host Name	Name of the node on which the page or swap file resides.
File Name	Name of the page or swap file. For secondary page or swap files, the file name is obtained by a special AST to the job controller on the remote node. The Availability Manager makes one attempt to retrieve the file name.
Used	Number of used blocks in the file.
% Used	Of the available blocks in each file, the percentage that has been used.
Total	Total number of blocks in the file.
Reservable	The number of reservable blocks in each page or swap file currently installed. Reservable blocks are blocks that might be logically claimed by a process for future physical allocation. A negative value indicates that the file might be overcommitted. Although a negative value is not an immediate concern, it indicates that the file might become overcommitted if physical memory becomes scarce.

3.2.5 Disk Summaries

The **Disk** tab allows you to display disk pages that contain data about availability, count, and errors of disk devices on the system. OpenVMS disk data displays differ from those for Windows NT nodes, as described in the following sections.

On OpenVMS pages, the **View** menu lets you choose the following disk summaries:

- Disk Status Summary
- Disk Volume Summary

Also, on the Disk Status Summary, you can double-click a device name to display a Single Disk Summary page.

3.2.5.1 OpenVMS Disk Status Summary

When you click the **Disk** tab on the OpenVMS Node Summary page, the Availability Manager displays the default disk page, the OpenVMS Disk Status Summary (Figure 3–12). This page displays disk device data, including path, volume name, status, and mount, transaction, error, and resource wait counts.

Figure 3–12 OpenVMS Disk Status Summary

Device Name	Host Path	Volume Name	Status	Error	Trans	Mount	RWait
CMOVEQ4DKA0	CMOVEQ	CMOVEQ_SYS	Mounted	0	372	1	0
CMOVEQ4DKA100	CMOVEQ	REDSQL_63978	Mounted	0	1	1	0
DBSAVC4DKB0	DBSAVC	AXPYMSQ2OPT3	Mounted	0	1	1	0
DBGAVC4DKB100	DBGAVC	ALPHATOOLS1	Mounted	0	1	1	0
DBSAVC4DKB200	DBSAVC	ALPHADEB090	Mounted	0	1	1	0
DBGAVC4DKB300	DBGAVC	YAXDEB062	Mounted	0	1	1	0
DBSAVC4DKB500	DBSAVC	ALPHATOOLS	Mounted	0	1	1	0
DBGAVC4DKC100	DBGAVC	ALPHADEB062	Mounted	0	1	1	0
DBSAVC4DKC200	DBSAVC	ALPHADEB061	Mounted	0	1	1	0
DBGAVC4DKC500	DBGAVC	DCPI1	Mounted	0	1	1	0
SELAVY4DKA100	SELAVY	YDEBUS	Mounted	0	1	1	0
SELAVY4DKA700	SELAVY	SELAVY_2698	Mounted	0	1	1	0
PEROIT4DKA0	PEROIT	YAXDEB064	Mounted	0	1	1	0
PEROIT4DKA300	PEROIT	YAXDEB065	Mounted	0	1	1	0
PEROIT4DKA400	PEROIT	YAXDEB06	Mounted	0	1	1	0
PEROIT4DKA700	PEROIT	PEROIT_64467	Mounted	0	1	1	0
NAMK4DKA300	NAMK	NAMK_PASE	Mounted	0	1	1	0
NAMK4DKA400	NAMK	FREEMARY20	Mounted setLick	0	1	1	0
DRINKS4DKA0	DRINKS	YAXDEB067	Mounted	0	1	1	0
DRINKS4DKA100	DRINKS	DCPS	Mounted	0	1	1	0
DRINKS4DKA500	DRINKS	L0MPAT_USER	Mounted	0	1	1	0
DRINKS4DKA700	DRINKS	DRINKS_2977	Mounted	0	1	1	0
REDSQL4DKA0	REDSQL	REDSQL_USR	Mounted	0	1	1	0
SSPADE4DKA0	SSPADE	YAXYMSY060	Mounted	0	1	1	0
SSPADE4DKA200	SSPADE	XDEBUS	Mounted	0	1	1	0
SSPADE4DKA300	SSPADE	YAXDEB063	Mounted	0	1	1	0

OpenVMS Alpha node CMOVEQ - Disk Survey - Status: 45 disks (27 listed, 18 filtered)

Getting Information about Nodes

3.2 Node Data Pages

This summary displays the following data:

Heading	Description
Device Name	Standard OpenVMS device name that indicates where the device is located, as well as a controller or unit designation.
Host Path	Primary path (node) from which the device receives commands.
Volume Name	Name of the mounted media.
Status	One or more of the following disk status values: <ul style="list-style-type: none"> Alloc Disk is allocated to a specific user. CluTran Disk status is uncertain because of a cluster state transition in progress. Dismount Disk in process of dismounting; may be waiting for a file to close. Foreign Disk is mounted with the /FOREIGN qualifier. Invalid Disk is in an invalid state (most likely Mount Verify Timeout). MntVerify Disk is waiting for a mount verification. Mounted Disk is logically mounted by a MOUNT command. Offline Disk is no longer physically mounted in device drive. Online Disk is physically mounted in device drive. Shadow Set Member Disk is a member of a shadow set. Unavailable Disk is set to unavailable. Wrong Volume Disk was mounted with the wrong volume name. Wrtlck Disk is mounted and write locked.
Error	Number of errors generated by the disk (a quick indicator of device problems).
Trans	Number of in-progress file system operations for the disk.
Mount	Number of nodes that have the specified disk mounted.
Rwait	Indicator that a system I/O operation is stalled, usually during normal recovery from a connection failure or during volume processing of host-based shadowing.

3.2.5.2 OpenVMS Single Disk Summary

To collect single disk data and display the data on the Single Disk Summary, double-click a device name on the Disk Status Summary. Figure 3-13 is an example of a Single Disk Summary page. The display interval of the data collected is 5 seconds.

Figure 3–13 OpenVMS Single Disk Summary Page

Node	Status	Errors	Trans	RWait	Free	QLen	OpRate
CMOVEQ	Mounted	0	1	0	41430	0.00	0.00
DB0AVC	Mounted	0	1	0	40833	0.00	0.00
DRINKS	Mounted	0	1	0	40833	0.00	0.00
LOWEAT	Mounted	0	1	0	41571	0.00	0.00
MAMC	Mounted	0	1	0	40881	0.00	0.00
PEROIT	Mounted	0	1	0	40833	0.00	0.00
REDSQG	Mounted	0	1	0	41289	0.00	0.00
SAL8BC	Mounted	0	1	0	41289	0.00	0.00
SELAVY	Mounted	0	1	0	41571	0.00	0.00
88PADB	Mounted	0	1	0	40833	0.00	0.00

This summary displays the following data:

Data	Description
Node	Name of the node.
Status	Status of the disk: mounted, online, offline, and so on.
Errors	Number of errors on the disk.
Trans	Number of in-progress file system operations on the disk (number of open files on the volume).
Rwait	Indication of an I/O stalled on the disk.
Free	Number of free disk blocks on the volume. An (M) after the free block count indicates this node holds the lock on the volume that DECamsd uses to obtain the true free block count on the volume. Other nodes might not have accessed the disk, so their free block count might not be up to date.
QLen	Average number of operations in the I/O queue for the volume.
OpRate	Each node's contribution to total OpRate for the disk.

3.2.5.3 OpenVMS Disk Volume Summary

When you select the Status Summary option from the **View** menu on the OpenVMS Node Summary page, the Availability Manager displays the OpenVMS Disk Status Summary (Figure 3–14). This page displays disk volume data, including path, volume name, disk block utilization, queue length, and operation rate.

Getting Information about Nodes

3.2 Node Data Pages

Figure 3–14 OpenVMS Disk Volume Summary Page

This summary displays the following data:

Data	Description
Device Name	Standard OpenVMS device name that indicates where the device is located, as well as a controller or unit designation.
Host Path	Primary path (node) from which the device receives commands.
Volume Name	Name of the mounted media.
% Used	Percentage of the number of volume blocks in use in relation to the total volume blocks available.
Free	Number of blocks of volume space available for new data.
Queue	Average number of I/O operations pending for the volume (an indicator of performance; less than 1.00 is optimal).
OpRate	Operation rate for the most recent sampling interval. The rate measures the amount of activity on a volume. The optimal load is device specific.

Windows NT Logical and Physical Disk Summaries

On Windows NT nodes, the **View** menu lets you choose the following summaries:

- Logical Disk Summary
- Physical Disk Summary

3.2.5.4 Windows NT Logical Disk Summary


A **logical disk** is the user-definable set of partitions under a drive letter. The Windows NT Logical Disk Summary displays logical disk device data, including path, label, percentage used, free space, and queue statistics.

To display the Logical Disk Summary page, follow these steps:

1. Double-click a node name in the Node pane to display the Windows NT Node Summary.
2. Click the **Disk** tab on the Windows NT Node Summary.

The Availability Manager displays the Windows NT Logical Disk Summary page (Figure 3–15).

Figure 3–15 Windows NT Logical Disk Summary Page



This summary displays the following data:

Data	Description
Disk	Drive letter, for example, <i>c:</i> ; or <i>Total</i> , which is the summation of statistics for all the disks.
Path	Primary path (node) from which the device receives commands.
Label	Identifying label of a volume.
Type	File system type; for example, FAT or NTFS.
% Used	Percentage of disk space used.
Free	Amount of free space available on the logical disk unit.
Current Queue	Number of requests outstanding on the disk at the time the performance data is collected. It includes requests in progress at the time of data collection.
Average Queue	Average number of both read and write requests that were queued for the selected disk during the sample interval.
Transfers/Sec	Rate of read and write operations on the disk.
KBytes/Sec	Rate data is transferred to or from the disk during write or read operations. The rate is displayed in kilobytes per second.
% Busy	Percentage of elapsed time that the selected disk drive is busy servicing read and write requests.

Getting Information about Nodes

3.2 Node Data Pages

3.2.5.5 Windows NT Physical Disk Summary

A **physical disk** is hardware used on your computer system. The Windows NT Physical Disk Summary displays disk volume data, including path, label, queue statistics, transfers, and bytes per second.

To display the Windows NT Physical Disk Summary, follow these steps:

1. Click the **View** menu on the Windows NT Logical Disk Summary.
2. Click the Physical Disk Summary menu option.

The Availability Manager displays the Windows NT Physical Disk Summary page (Figure 3–16).

Figure 3–16 Windows NT Physical Disk Summary Page

Disk	Path	Current Queue	Average Queue	Transfers/Sec	KBytes/Sec	% Busy	% Read Busy	% Write Busy
0	AFFC53	0	0.00	0.20	0.70	0.03	0.00	0.03
Total	AFFC53	0	0.00	0.20	0.70	0.03	0.00	0.03

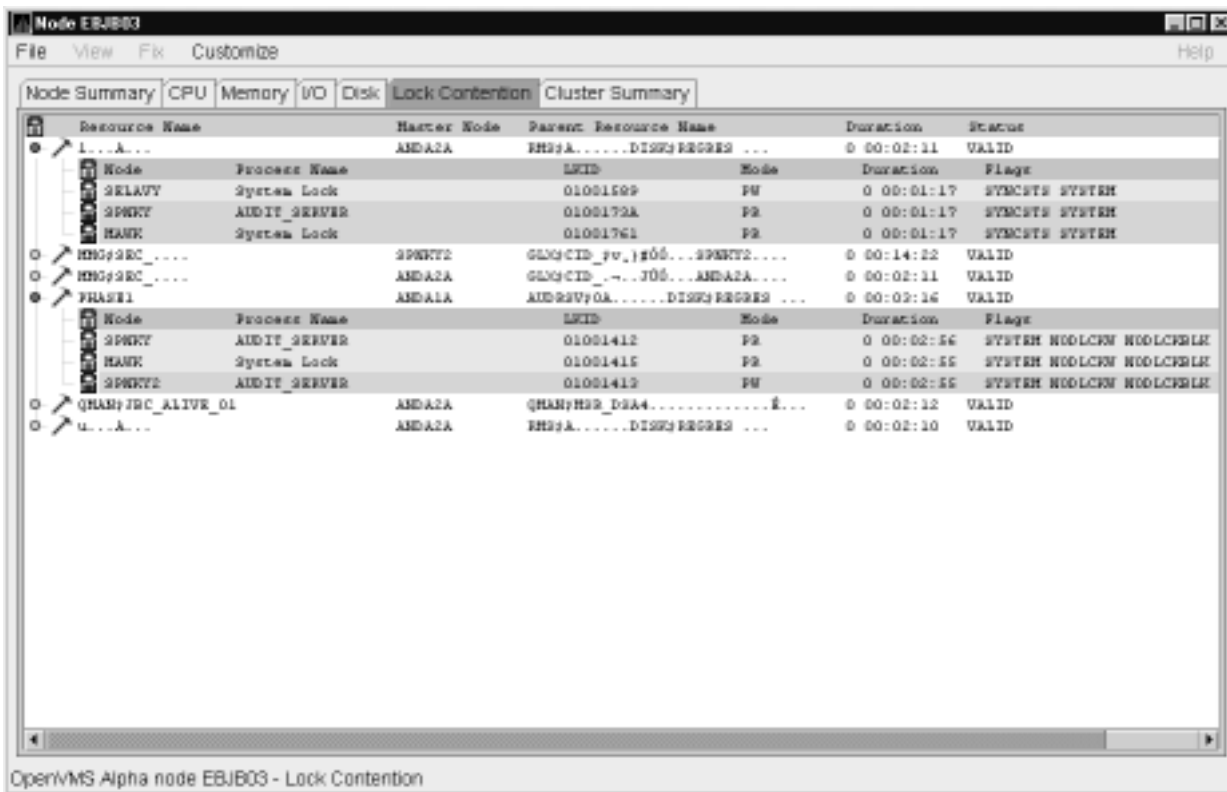
This page displays the following data:

Data	Description
Disk	Drive number, for example, 0, 1, 2 or <i>Total</i> , which is the summation of statistics for all the disks.
Path	Primary path (node) from which the device receives commands.
Current Queue	Number of requests outstanding on the disk at the time the performance data is collected; it includes requests in service at the time of data collection.
Average Queue	Average number of read and write requests that were queued for the selected disk during the sample interval.
Transfers/Sec	Rate of read and write operations on the disk. The rate is displayed in kilobytes per second.
KBytes/Sec	Rate bytes are transferred to or from the disk during read or write operations. The rate is displayed in kilobytes per second.
% Busy	Percentage of elapsed time the selected disk drive is busy servicing read and write requests.
% Read Busy	Percentage of elapsed time the selected disk drive is busy servicing read requests.
% Write Busy	Percentage of elapsed time the selected disk drive is busy servicing write requests.

3.2.6 OpenVMS Lock Contention

When you select the **Lock Contention** tab on the OpenVMS Node Summary page, the Availability Manager displays the OpenVMS Lock Contention. This page, shown in Figure 3–17, displays each resource in the group you have selected for which a potential lock contention problem exists.

Figure 3–17 OpenVMS Lock Contention Page



This page displays the following data:

Data	Description
Resource Name	Resource name associated with the SENQ system service call. An ellipsis (...) indicates an unprintable character.
Master Node	Node on which the resource is mastered.
Parent Resource	Name of the parent resource. If no name is displayed, the resource listed is the parent resource.
Duration	Amount of time elapsed since the Availability Manager first detected the contention situation.
Status	Status of the lock. See the \$ENQ(W) description in the <i>OpenVMS System Services Reference Manual</i> .

When you double-click a data item under these headings, the Availability Manager displays the following data:

Data	Description
Node	Node name on which the lock is granted.
Process Name	Name of the process owning the blocking lock.
LKID	Lock ID value (which is useful with SDA).
Mode	One of the following modes at which the lock is granted or requested: EX, CW, CR, PW, PR, NL.

Getting Information about Nodes

3.2 Node Data Pages

Data	Description
Duration	Length of time the lock has been in the current queue (since the console application found the lock).
Flags	Flags specified with the <code>SENQ(W)</code> request.

Data is displayed in one of three colors:

Color	Meaning
Green	Granted
Yellow	Converting
Pink	Waiting

To interpret the information displayed on the OpenVMS Lock Contention Summary, you should understand OpenVMS lock management services. For more information, see the *OpenVMS System Services Reference Manual*.

Note

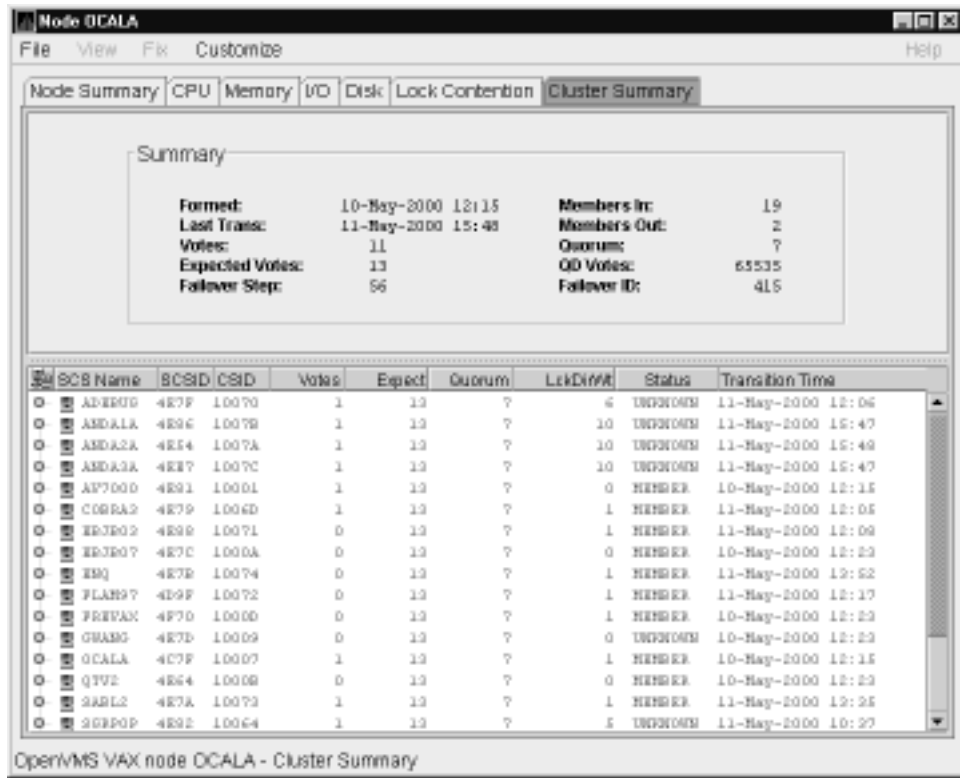
Lock contention data is accurate only if every node in an OpenVMS Cluster environment is in the same group. Multiple clusters can share a group, but clusters cannot be divided into different groups without losing accuracy.

3.2.7 OpenVMS Cluster Summary

When you click the **Cluster Summary** tab on an OpenVMS Node Summary page, the Availability Manager displays the OpenVMS Cluster Summary page (Figure 3-18).

This page contains cluster interconnect information for an entire cluster as well as detailed information about each node in the cluster, including System Communications Services (SCS) connections and LAN virtual circuits for individual nodes. The data items shown on the page correspond to data that the Show Cluster utility (`SHOW CLUSTER`) displays for the `SYSTEMS` and `MEMBERS` classes.

Figure 3–18 OpenVMS Cluster Summary Page



The two panes in the Cluster Summary page display the following information:

- The Summary pane (top) displays summary information about the entire cluster.
- The Cluster Members pane (bottom) displays detailed information about each node in the cluster, including its System Communication Architecture (SCA) connections with other nodes.

The following sections discuss these two panes.

Notes about Pane Displays

Following are notes about the display of data in the panes:

- The Availability Manager signals the LOVOTE event when the difference between the cluster's quorum and votes *is less than* the threshold for the event:

```
LOVOTE, 'node' VOTES count is close to or below QUORUM
```

The default threshold for LOVOTE is 1.

- You can change collection intervals by using the **Customize** menu on the main Application window. Choose the **Customize OpenVMS...** option, and then click the **Data Collection** tab. On the page displayed, you can select data that you want to collect and change the default collection intervals.

Getting Information about Nodes

3.2 Node Data Pages

3.2.7.1 Summary Pane

Table 3–5 describes the Summary pane data (see Figure 3–18).

Table 3–5 Summary Pane Data

Data	Description
Formed	Date and time the cluster was formed.
Last Trans	Date and time of the most recent cluster state transition.
Votes	Total number of quorum votes being contributed by all cluster members and by the quorum disk.
Expected Votes	Number of votes expected to be contributed by all members of the cluster as determined by the connection manager. This value is based on the maximum EXPECTED_VOTES system parameter and the maximized value of the VOTES system parameter.
Failover Step	Current failover step index. Shows which step in the sequence of failover steps the failover is currently executing.
Members In	Number of cluster members to which the Availability Manager has a connection.
Members Out	Number of cluster members to which the Availability Manager either has no connection or has lost connection.
Quorum	Number of votes required to keep the cluster above quorum.
QD Votes	Number of votes given to the quorum disk. A value of 65535 means there is no quorum disk.
Failover ID	Failover instance identification. Unique ID of a failover sequence that indicates to system managers whether a failover has occurred since the last time they checked.

3.2.7.2 Cluster Members Pane

The Cluster Members pane lists all the nodes in the cluster and provides detailed information about each one. The first level of information in the Cluster Members pane contains cluster member data (see Figure 3–18), which is described in Table 3–6.

Table 3–6 Cluster Member Data

Data	Description
SCS Name	System Communications Services (SCS) name for the node (system parameter SCSNODE).
SCSID	SCS identification for the node (system parameter SCSYSTEMID).
CSID	Cluster system identification.
Votes	Number of votes the member contributes.
Expect	Member's expected votes as set by the EXPECTED_VOTES system parameter.
Quorum	Number of votes that must be present for the cluster to function and permit user activity.
LckDirWt	Lock manager distributed directory weight as determined by the LCKDIRWT system parameter.

(continued on next page)

Table 3–6 (Cont.) Cluster Member Data

Data	Description
Status	Current cluster member status: MEMBER, UNKNOWN, or BRK_NON (break nonmember).
Transition Time	Time cluster member had last transition.

3.2.7.3 System Communications Architecture (SCA) Virtual Circuit (VC) Data

SCA VC data includes information about members' virtual circuits to other members of the cluster. More than one virtual circuit indicates more than one path to the other member.

The data displayed is similar to the information that the Show Cluster utility (SHOW CLUSTER) displays for the CIRCUITS, CONNECTIONS, and COUNTERS classes. However, in Availability Manager Version 1.3 and earlier, SHOW CLUSTER displays show connections and virtual circuits to other OpenVMS nodes only; they do not show connections and virtual circuits to the DIGITAL Storage Architecture (DSA) devices, such as FDDI or DSSI disk controllers. This limitation will be fixed in a future release of the Availability Manager software.

You can display virtual circuit data by double-clicking the SCS name of a node. Figure 3–19 is an example of virtual circuit data on a Cluster Summary page. (Because it is very wide, the illustration splits the screen into two sections; the bottom screen is a continuation of the top screen.)

Figure 3–19 Virtual Circuit Data Page



(continued on next page)

Getting Information about Nodes

3.2 Node Data Pages

Figure 3–19 (Cont.) Virtual Circuit Data Page

The screenshot shows the 'Cluster Summary' window for node ENQ. It includes a summary box with the following data:

Formed:	10-May-2000 12:15	Members In:	18
Last Trans:	12-May-2000 14:57	Members Out:	2
Votes:	12	Quorum:	7
Expected Votes:	13	QD Votes:	65835
Failover Step:	58	Failover ID:	740

Below the summary are two tables of virtual circuit data:

#	SCS Name	BCSID/CSID	Votes	Expect	Quorum	LckDir/A	Status	Transition Time
00		299	6	197	11	1720	0	0
54		1	0	0	0	0	0	0
61		2192	6	1024	11	9062	0	0
13		2225	6	1109	11	10216	0	0

MsgRcvd	KB Mapped	BlkDataSent	BlkDataRead	BlkTranSent	BlkTranRead	DOBent	DOFRcvd	Credit/A	BOT/A
0.00	7	6	1	9	2	0	0	0	0
0.00	0	0	0	0	0	0	0	0	0
0.00	0	0	0	0	0	0	0	0	0
0.52	2218	0	1108	2	10212	0	0	34	0

#	#Chars	TimeOpened	TimeClosed
00			
44	630	6	302
25	0	0	0
10	40	6	0
43	574	6	266
49	2210	25	1062

Table 3–7 describes virtual circuit data. Each line shows either a summary of all system applications (SYSAPs) that use virtual circuit communication or the communication on the connection between a local and a remote SYSAP.

Table 3–7 Virtual Circuit Data

Data	Description
NodeName	SCS name of the remotely connected node.
LPortName	The device name of the local port associated with the circuit.
RPortType	The type of remote port associated with the circuit.
State	The state of the virtual circuit connection.
MsgSent	Rate that messages are sent on the virtual circuit.
MsgRcvd	Rate that messages are received on the virtual circuit.
KB Mapped	Number of kilobytes mapped for block data transfer using the virtual circuit.
BlkDataSent	Number of kilobytes sent on the virtual circuit by this node using send block data transfers.
BlkDataReqd	Number of kilobytes requested on the virtual circuit by this node using request block data transfers.

(continued on next page)

Table 3–7 (Cont.) Virtual Circuit Data

Data	Description
BlkTransSent	Number of send block data transfers by this node on the virtual circuit.
BlkTransReqd	Number of request block data transfers by this node on the virtual circuit.
DGSent	Number of datagrams sent by this node on the virtual circuit.
DGRcvd	Number of datagrams received from the remote system on the virtual circuit.
CreditWt	Number of times any connection on the virtual circuit had to wait for a send credit.
BDTWt	Number of times any connection on the virtual circuit had to wait for a buffer descriptor.

3.2.7.4 SCS Connections Data

You can display SCS connections data by double-clicking the SCS Connections icon displayed under a virtual circuit node name.

See Figure 3–19 for SCS connections data on a Cluster Summary page. Table 3–8 describes the SCS connections data.

Table 3–8 SCS Connections Data

Data	Description
Local SYSAP	Name of the local system application using the SCS connection.
Remote SYSAP	Name of the remote system application communicated with using the SCS connection.
State	The state of the SCS connection.
MsgSent	Rate that messages are sent to the remote SYSAP using the SCS connection.
MsgRcvd	Rate that messages are received from the remote SYSAP using the SCS connection.
KB Mapped	Number of kilobytes mapped for block data transfer using the SCS connection. Note: This field is available only in raw data format.
BlkDataSent	Number of kilobytes sent on the SCS connection by this node using send block data transfers.
BlkDataReqd	Number of kilobytes the data requested on the SCS connection by this node using request block data transfers.
BlkTransSent	Number of send block data transfers by this node on the SCS connection.
BlkTransReqd	Number of request block data transfers by this node on the SCS connection.
DGSent	Number of datagrams sent by this node on the SCS connection.
DGRcvd	Number of datagrams received from the remote system on the SCS connection.
CreditWt	Number of times the connection had to wait for a send credit.
BDTWt	Number of times the connection had to wait for a buffer descriptor.

Getting Information about Nodes

3.2 Node Data Pages

3.2.7.5 LAN Virtual Circuit Details Data

You can display LAN virtual circuit details by double-clicking the icon preceding that heading on the Cluster Summary page. The data items displayed depend of the type of virtual circuit. Currently, this feature is available only for LAN virtual circuits.

See Figure 3–19 for a LAN virtual circuit data display. Table 3–9 describes the LAN virtual circuit details data.

Table 3–9 LAN Virtual Circuit Details Data

Data	Description
BufSz	Maximum data buffer size in use by this LAN virtual circuit.
ReXMTRatio	Ratio of retransmitted to transmitted packets during the most recent data collection interval.
ReXMTTimOut (µs)	Retransmission timeout (in microseconds). The amount of time the virtual circuit will wait for acknowledgment of a packet's receipt before retransmitting the packet.
WinSz	Current value of the transmit window (or pipe quota). Maximum number of packets that will be sent before stopping to await an acknowledgment. After a timeout, the transmit window is reset to 1 to decrease congestion and is allowed to increase as acknowledgments are received.
WinSzMax	Maximum transmit window size currently allowed for the virtual circuit.
#Chans	Number of channels ¹ available to the virtual circuit.
TimeOpened	Most recent time the virtual circuit was opened. Not implemented in this release.
TimeClosed	Most recent time the virtual circuit was closed. Not implemented in this release.

¹A channel is a logical communication path between two LAN adapters located on different nodes. Channels between nodes are determined by the pairs of LAN adapters and the connecting network. For example, two nodes, each having two LAN adapters, could establish four channels. The data carried by a particular virtual circuit can be sent over any channel connecting the two nodes.

For more detailed explanations of these data, refer to the NISCA troubleshooting information in the *OpenVMS Cluster Systems* manual.

3.2.8 OpenVMS Network Interconnect for the System Communication Architecture (NISCA) Pages

The Network Interconnect for the System Communication Architecture (NISCA) is the transport protocol responsible for carrying messages such as disk I/Os and lock messages across Ethernet and FDDI LANs to other nodes in the cluster. More detailed information about the protocol is in the *OpenVMS Cluster Systems* manual.

The NISCA pages show detailed information about the LAN (Ethernet or FDDI) connection between two nodes. The Availability Manager displays one window for each LAN virtual circuit.

The pages are intended primarily as real-time aids to diagnosing LAN-related problems. The *OpenVMS Cluster Systems* manual describes the parameters shown in these pages and tells how to diagnose LAN-related cluster problems. The pages provide the same information as the OpenVMS System Dump Analyzer (SDA) command `SHOW PORTS/VC=VC_node-name`. (VC refers to

Getting Information about Nodes 3.2 Node Data Pages

a virtual circuit; *node-name* refers to a node in the cluster. The system defines *VC-node-name* after a SHOW PORTS command is issued from SDA.)

To display NISCA details, double-click the buffer size data item (BufSz) under the LAN Virtual Circuit Details heading on the Cluster Summary page (Figure 3–20).

Figure 3–20 LAN Virtual Circuit Details



The system then displays a Transmit Data page (Figure 3–24) that contains five additional tabs you can select. These pages are described in the following sections.

Notes about the Display

Following are notes about displaying data in this window:

- You cannot sort or filter the data displayed.
- You can change collection intervals.

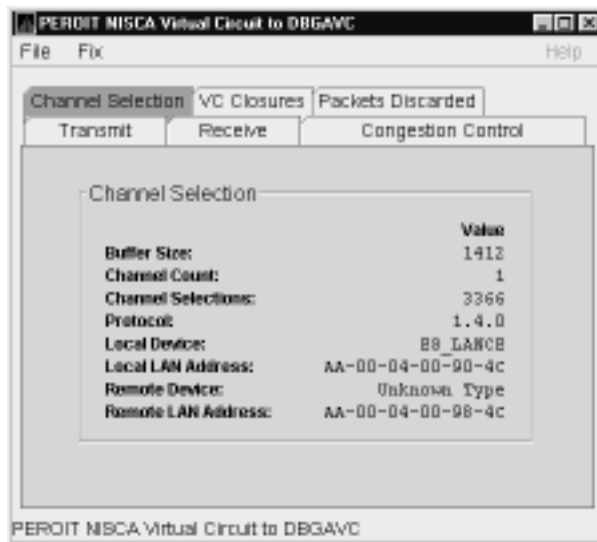
3.2.8.1 Channel Selection Data

Channel selection data includes information about the selection of virtual circuit channels. Figure 3–21 is an example of a Channel Selection Data page.

Getting Information about Nodes

3.2 Node Data Pages

Figure 3–21 Channel Selection Data Page



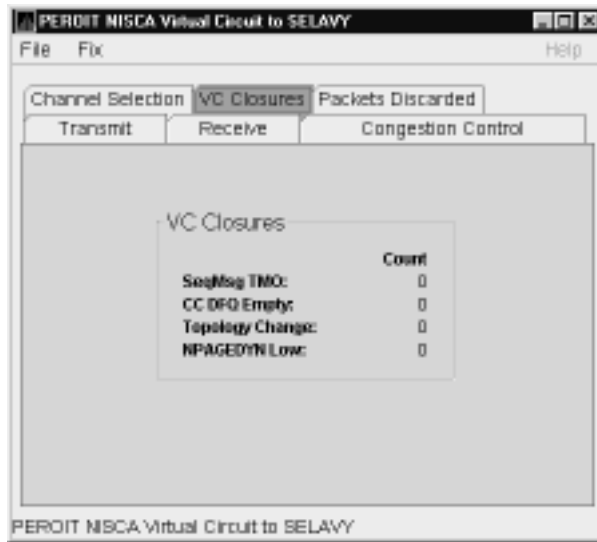
The following table describes the data displayed:

Data	Description
Buffer Size	Maximum data buffer size for this virtual circuit.
Channel Count	Number of channels available for use by this virtual circuit.
Channel Selections	Number of channel selections performed.
Protocol	NISCA protocol version.
Local Device	Name of the local LAN device that the channel uses to send and receive packets.
Local LAN Address	Address of the local LAN device that performs sends and receives.
Remote Device	Name of the remote LAN device that the channel uses to send and receive packets.
Remote LAN Address	Address of the remote LAN device performing the sends and receives.

3.2.8.2 Virtual Circuit (VC) Closures Data

VC closures data is information about the number of times a virtual circuit has closed for a particular reason. Figure 3–22 is an example of a VC Closures Data page.

Figure 3–22 Virtual Connect (VC) Closures Data Page



The following table describes the displayed data:

Data	Description
SeqMsg TMO	Number of times the VC was closed because of sequenced transmit timeouts.
CC DFQ Empty	Number of times the VC was closed because the channel control data-free queue (DFQ) was empty.
Topology Change	Number of times the VC was closed because PEDRIVER performed a failover from a LAN path with a large packet size to a LAN path with a smaller packet size, necessitating the closing and reopening of the virtual circuit.
NPAGEDYN Low	Number of times the virtual circuit was lost because of a pool allocation failure on the local node.


3.2.8.3 Packets Discarded Data

Packets discarded data is information about the number of times packets were discarded for a particular reason. Figure 3–23 is an example of a Packets Discarded Data page.

Getting Information about Nodes

3.2 Node Data Pages

Figure 3–23 Packets Discarded Data Page




The following table describes the displayed data:

Data	Description
No Xmt Chan	Number of times there was no transmit channel.
Rcv Short Msg	Number of times an undersized transport message was received.
Ill Seq Msg	Number of times an illegal sequenced message was received.
Bad Checksum	Number of times there was a checksum failure on a received packet.
TR DFQ Empty	Number of times the transmit data-free queue (DFQ) was empty.
TR MFQ Empty	Number of times the transmit message-free queue (MFQ) was empty.
CC MFQ Empty	Number of times the channel control MFQ was empty.
Cache Miss	Number of packets that could not be placed in the virtual circuit's receive cache because the cache was full.

3.2.8.4 Transmit Data

Transmit data is information about data packet transmission. Figure 3–24 is an example of a Transmit Data page.

Figure 3–24 Transmit Data Page



The following table describes the displayed data:

Data	Description
Packets	(Raw) count and rate of packets transmitted through the virtual circuit to the remote node, including both sequenced and unsequenced (channel control) messages and lone acknowledgments.
Unsequenced (DG)	(Raw) count and rate of the number of unsequenced packets transmitted.
Sequenced	(Raw) count and rate of sequenced packets transmitted. Sequenced packets are guaranteed to be delivered.
Lone ACK	(Raw) count and rate of packets sent solely for the purpose of acknowledging receipt of one or more packets.
ReXMT Count	Number of packets retransmitted. Retransmission occurs when the local node does not receive an acknowledgment for a transmitted packet within a predetermined timeout interval.
ReXMT Timeout	Number of retransmission timeouts that have occurred.
ReXMT Ratio	Ratio of current and past retransmission count to the current and past number of sequenced messages sent.
Bytes	(Raw) count and rate of bytes transmitted through the virtual circuit.


3.2.8.5 Receive Data

Receive data is information about the receipt of data packets. Figure 3–25 is an example of a Receive Data page.

Getting Information about Nodes

3.2 Node Data Pages

Figure 3–25 Receive Data Page



The following table describes the displayed data:


Data	Description
Packets	(Raw) count and rate of packets received on the virtual circuit from the remote node, including both sequenced and unsequenced (channel control) messages and lone acknowledgments.
Unsequenced (DG)	(Raw) count and rate of unsequenced packets received.
Sequenced	(Raw) count and rate of sequenced packets received.
Lone ACK	(Raw) count and rate of lone acknowledgments received.
Duplicate	Number of redundant packets received by this system. Duplicates occur when the sending node retransmits a packet and both the original and retransmitted packets are received.
Out of Order	Number of packets received out of order by this system.
Illegal ACK	Number of illegal acknowledgments received.
Bytes	(Raw) count and rate of bytes received through the virtual circuit.

3.2.8.6 Congestion Control Data

Congestion control data is information about transmit congestion. The values indicate the number of messages that can be sent to the remote node before receiving an acknowledgment and the retransmission timeout.

Figure 3–26 is an example of a Congestion Control Data page.

Figure 3–26 Congestion Control Data Page



The following table describes the displayed data:

Data	Description
Transmit Window Current	Current value of the transmit window (or pipe quota). After a timeout, the pipe quota is reset to 1 to decrease congestion and is allowed to increase as quickly as acknowledgments are received.
Transmit Window Grow	The slow growth threshold. The size at which the window's increase rate is slowed to avoid congestion on the network again.
Transmit Window Max	Maximum transmit window size currently allowed for the virtual circuit based on channel limitations.
Transmit Window Reached	Number of times the entire transmit window was full. If this number is small compared with the number of sequenced messages transmitted, then either the local node is not sending large bursts of data to the remote node or acknowledging packets are being received so promptly that the window limit is never reached.
Roundtrip µs	Average round-trip time (in microseconds) for a packet to be sent and acknowledged.
Roundtrip Deviation µs	Average deviation (in microseconds) of the round-trip time.
Retransmit Timeout µs	Value (in microseconds) used to determine packet retransmission timeout. If a packet does not receive either an acknowledging or a responding packet, the packet is assumed to be lost and will be resent.
UnAcked Messages	Number of unacknowledged messages.
CMD Queue Length	Current length of all command queues.
CMD Queue Max	Maximum number of commands in queues so far.

Getting Information about Nodes

3.2 Node Data Pages

3.2.9 OpenVMS Single Process

When you double-click a PID on the lower part of an OpenVMS CPU Process (Figure 3–7), Memory (Figure 3–9), or I/O (Figure 3–10) page, the Availability Manager displays the first of several OpenVMS Single Process pages. These pages allow you to click tabs to display specific data about one process.

The following pages display data about a single process and are described in subsequent sections:

- Process Information
- Working Set
- Execution Rates
- Process I/O Quotas
- Wait States
- Job Quotas

3.2.9.1 Process Information

When you click the **Process Information** tab, the Availability Manager displays the OpenVMS Process Information page, as shown in Figure 3–27. The data on this page is displayed at the default intervals shown for Single Process Data on the Data Collection customizations.

Figure 3–27 Process Information Page



The following table describes the data on this page:

Data	Description
Process name	Name of the process.
Username	User name of the user who owns the process.
Account	Account string that the system manager assigns to the user.

Getting Information about Nodes

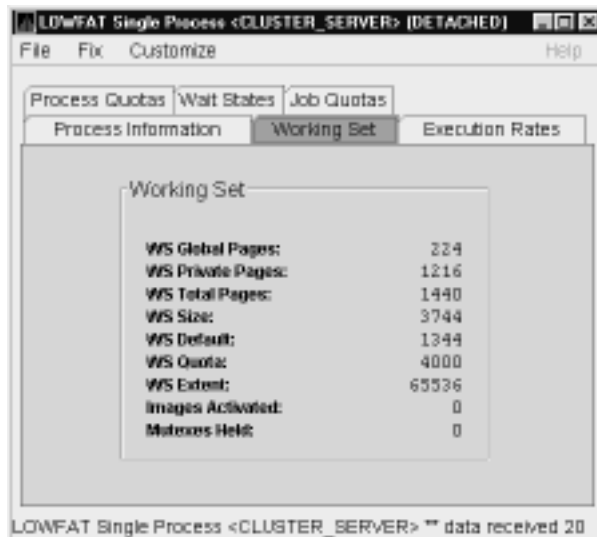
3.2 Node Data Pages

Data	Description
UIC	User identification code (UIC). A pair of numbers or character strings that designate the group and user.
PID	Process identifier. A 32-bit value that uniquely identifies a process.
Owner ID	Process identifier of the process that created the process displayed on the page. If the PID is 0, then the process is a parent process.
PC	Program counter. On OpenVMS Alpha systems, this value is displayed as 0 because the data is not readily available to the Data Collector node.
PSL	Processor status longword (PSL). This value is displayed on VAX systems only.
Priority	Computable and base priority of the process. Priority is an integer between 0 and 31. Processes with higher priority are given more CPU time.
State	One of the process states listed in Appendix A.

3.2.9.2 Working Set

When you click the **Working Set** tab, the Availability Manager displays the OpenVMS Working Set page (Figure 3–28).

Figure 3–28 Working Set Page



The following table describes the data on this page:

Data	Description
WS Global Pages	Shared data or code between processes, listed in pages (measured in pagelets).
WS Private Pages	Amount of accessible memory, listed in pages (measured in pagelets).
WS Total Pages	Sum of global and private pages (measured in pagelets).

Getting Information about Nodes

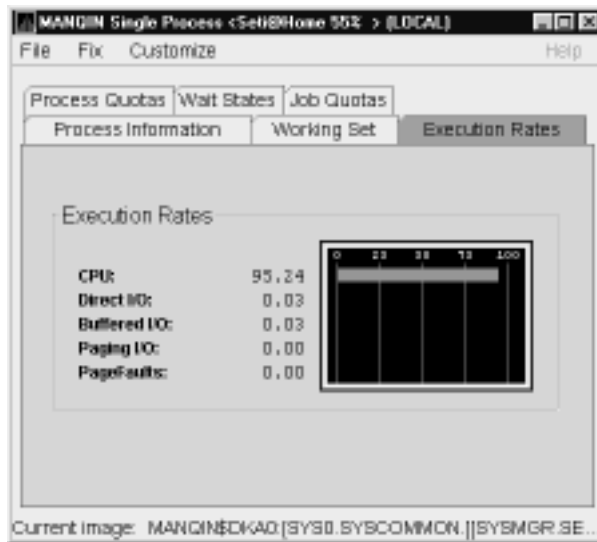
3.2 Node Data Pages

Data	Description
WS Size	Working set size. The number of pages (measured in pagelets) of memory the process is allowed to use. This value is periodically adjusted by the operating system based on analysis of page faults relative to CPU time used. Increases in large units indicates that a process is taking many page faults, and its memory allocation is increasing.
WS Default	Working set default. The initial limit of the number of physical pages (measured in pagelets) of memory the process can use. This parameter is listed in the user authorization file (UAF); discrepancies between the UAF value and the displayed value are due to page/longword boundary rounding or other adjustments made by the operating system.
WS Quota	Working set quota. The maximum amount of physical pages (measured in pagelets) of memory the process can lock into its working set. This parameter is listed in the UAF; discrepancies between the UAF value and the displayed value are due to page/longword boundary rounding or other adjustments made by the operating system.
WS Extent	Working set extent. The maximum number of physical pages (measured in pagelets) of memory the system will allocate for the process. The system provides memory to a process beyond its quota only when it has an excess of free pages and can be recalled if necessary. This parameter is listed in the UAF; any discrepancies between the UAF value and the displayed value are due to page/longword boundary rounding or other adjustments made by the operating system.
Images Activated	Number of times an image is activated.
Mutexes Held	Number of mutual exclusions (mutexes) held. Persistent values other than zero (0) require analysis. A mutex is similar to a lock but is restricted to one CPU. When a process holds a mutex, its priority is temporarily increased to 16.

3.2.9.3 Execution Rates

When you click the **Execution Rates** tab, the Availability Manager displays the OpenVMS Execution Rates page, as shown in Figure 3–29.

Figure 3–29 Execution Rates Page



The following table describes the data on this page:

Data	Description
CPU	Percent of CPU time used by this process. The ratio of CPU time to elapsed time.
Direct I/O	Rate at which I/O transfers take place from the pages or pagelets containing the process buffer that the system locks in physical memory to the system devices.
Buffered I/O	Rate at which I/O transfers take place for the process buffer from an intermediate buffer from the system buffer pool.
Paging I/O	Rate of read attempts necessary to satisfy page faults. This is also known as page read I/O or the hard fault rate.
Page Faults	Page faults per second for the process.


3.2.9.4 Process I/O Quotas

When you click **Process Quotas**, the Availability Manager displays the OpenVMS Process Quotas page, as shown in Figure 3–30.

Getting Information about Nodes

3.2 Node Data Pages

Figure 3–30 Process Quotas Page




The following table describes the data on this page. Note that when you display the SWAPPER process, no values are listed in this section. The SWAPPER process does not have quotas defined in the same way as other system and user processes do.

Data	Description
Direct I/O	The current number of direct I/Os used compared with the limit possible.
Buffered I/O	The current number of buffered I/Os used compared with the possible limit.
ASTs	Asynchronous system traps. The current number of ASTs used compared with the possible limit.
CPU Time	Amount of time used compared with the possible limit (in ms). "No Limit" is displayed if the limit is 0.

3.2.9.5 Wait States

When you click the **Wait States** tab, the Availability Manager displays the OpenVMS Wait States page, as shown in Figure 3–31.

Figure 3–31 Wait States Page



The following table describes the data on this page. The wait state specifies why a process cannot execute, based on calculations made on collected data. Each value is calculated over a data collection period of approximately 2 minutes.

“Current” refers to the percentage of elapsed time each process spends in one of the computed wait states. If a process spends all its time waiting in one state, the total approaches 100%.

Appendix A contains descriptions of process wait states.

Data	Description
Compute	Relative value indicating that the process is waiting for CPU time. Possible states are COM, COMO, or RWCAP.
Memory	Relative value indicating that the process is waiting for a page fault that requires data to be read from disk; this is common during image activation. Possible states are PFW, COLPG, FPG, RWPAG, RWNPG, RWMPE, or RWMPB.
Direct I/O	Relative value indicating that the process is waiting for data to be read from or written to a disk. The possible state is DIO.
Buffered I/O	Relative value indicating that the process is waiting for data to be read from or written to a slower device such as a terminal, line printer, or mailbox. The possible state is BIO.
Control	Relative value indicating that the process is waiting for another process to release control of some resource. Possible states are CEF, MWAIT, LEF, LEFO, RWAST, RWMBX, RWSCS, RWCLU, RWCSV, RWUNK, or LEF waiting for an ENQ.
Quotas	Relative value indicating that the process is waiting because the process has exceeded some quota. Possible states are QUOTA or RWAST_QUOTA.
Explicit	Relative value indicating that the process is waiting because the process asked to wait, such as a hibernate system service. Possible states are HIB, HIBO, SUSP, SUSPO, or LEF waiting for a TQE.

Getting Information about Nodes

3.2 Node Data Pages

3.2.9.6 Job Quotas

When you click the **Job Quotas** tab, the Availability Manager displays the OpenVMS Job Quotas page, as shown in Figure 3–32.

Figure 3–32 Job Quotas Page



The following table describes the data on this page:

Data	Description
File Count	Current number of open files compared with the possible limit.
Page File Count	Current number of disk blocks in the page file that the process can use compared with the possible limit.
Enqueue Count	Current number of resources (lock blocks) queued compared with the possible limit.
TQE Count	Current number of timer queue entry (TQE) requests compared with the possible limit.
Process Count	Current number of subprocesses created compared with the possible limit.
Byte Count	Current number of bytes used for buffered I/O transfers compared with the possible limit.

Getting Information about Events

Note

Before you start this chapter, be sure to read the explanation of data collection, events, thresholds, and occurrences in Chapter 1.

The Availability Manager indicates resource availability problems in the Events pane of the main Application window (Figure 4-1).

Figure 4-1 Application Window

Node Name	CPU	MEM	BIO	DIO	CPU Qs	OS Version	HW Model
AFF551	2	11	6	0	0	Open VMS V7.1-1H1	AlphaServer 2100 4/275
EBJB35	0	0	0	0	0	Open VMS V621.P2E	AlphaServer DS20 500 MHz
HMS	15	64	12	2	0	Open VMS V7.2-1	AlphaServer 400 4233
MANQIN	100	79	58	4	2	Open VMS V7.2	AlphaServer 400 4233
MCAREY	9	73	0	14	10	Open VMS V7.2	VAXstation 4000-90
MSCP	0	0	0	0	0	Open VMS V7.2	VAX 6000-620
PRF34A	1	36	0	0	0	Open VMS V6V4.FT1	DEC 3000 Model 400
PRF34B	0	37	0	0	0	Open VMS V6V4.FT1	DEC 3000 Model 400
XDELTA	0	0	0	0	0	Open VMS V7.2-1	AlphaStation 255/300

Node	Group	Date & Time	Severity	Event	Description
MSCP	DECAMDS	06-Mar-2000 16:12:22.04	100	PTHLS1	MSCP path lost, uptime was 0 01:44:44.93
MCAREY	DECAMDS	06-Mar-2000 16:13:17.27	99	FITINO	MCAREY has timed out for node FID to File...
BBJB35	DECAMDS	06-Mar-2000 16:16:25.83	90	DPGERR	EBJB35 error executing driver program, pr...
MCAREY	DECAMDS	06-Mar-2000 16:12:37.40	80	DSEKRM	MCAREY (2034DKA0) (2034DKA0) disk FMAIT co...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2044DKA400) (2044DKA400) disk FMAI...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2054DKA300) (2054DKA300) disk FMAI...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2034DKA300) (2034DKA300) disk FMAI...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2084DKA100) (2084DKA100) disk FMAI...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2084DKA400) (2084DKA400) disk FMAI...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2034DKA0) (2034DKA0) disk FMAIT co...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2074DKA200) (2074DKA200) disk FMAI...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2014DKA300) (2014DKA300) disk FMAI...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2014DKA100) (2014DKA100) disk FMAI...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2044DKA100) (WIL02 BSE) disk FMAIT ...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2094DKA0) (2094DKA0) disk FMAIT co...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2034DKA100) (2034DKA100) disk FMAI...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2084DKA300) (2084DKA300) disk FMAI...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2054DKA200) (2054DKA200) disk FMAI...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2044DKA300) (2044DKA300) disk FMAI...
MCAREY	DECAMDS	06-Mar-2000 16:12:38.70	80	DSEKRM	MCAREY (2044DKA0) (QTV4TOOLS) disk FMAIT co...

Collection [WDM_GROUP] has 2 nodes 16:16:28

The Events pane, which occupies the lower section of the window, allows you to identify a system problem. In many cases, you can apply fixes to correct these problems as well (see Chapter 5). The Availability Manager displays a warning message in the Events pane whenever it detects a resource availability problem.

Getting Information about Events

4.1 Displaying Event Information

4.1 Displaying Event Information

The Availability Manager automatically displays events for all nodes currently in communication with the Data Analyzer. When an event of a specific severity occurs, the Availability Manager adds the event to a list in the Events pane.

By default, the Events pane displays more serious events in red and less serious events in yellow; more serious events are listed first.

The length of time an event is displayed depends on the severity of the event. Less severe events are displayed for a short period of time (30 seconds); more severe events are displayed until you explicitly remove the event from the Events pane (explained in Section 4.1.2).

4.1.1 Data in the Events Pane

Table 4–1 identifies the data items displayed in the Events pane.

Table 4–1 Events Pane Data

Data Item	Description
Node	Name of the node causing the event
Group	Group of the node causing the event
Date	Date the event occurred
Time	Time that an event was detected
Sev	Severity: a value from 0 to 100
Event	Alphanumeric identifier of the type of event
Description	Short description of the resource availability problem

Appendix B contains tables of events that are displayed in the Event pane. In addition, these tables contain an explanation of each event and the recommended remedial action.

4.1.2 Events Pane Menu Options

When you right-click a node name or data item in the Events pane, the Availability Manager displays a popup menu with the following options:

Menu Option	Description
Display	Displays the Node Summary page associated with that event.
Fix	Displays a list of fix options (see Chapter 5.)
Remove	Removes an event from the display.
Freeze/Unfreeze	Freezes a value in the display until you “unfreeze” it; a snowflake icon is displayed to the left of an event that is frozen.
Customize	Allows you to customize events.

4.2 Criteria Signaling and Displaying an Event

The Availability Manager uses the following criteria to determine whether to signal an event and display it in the Events pane:

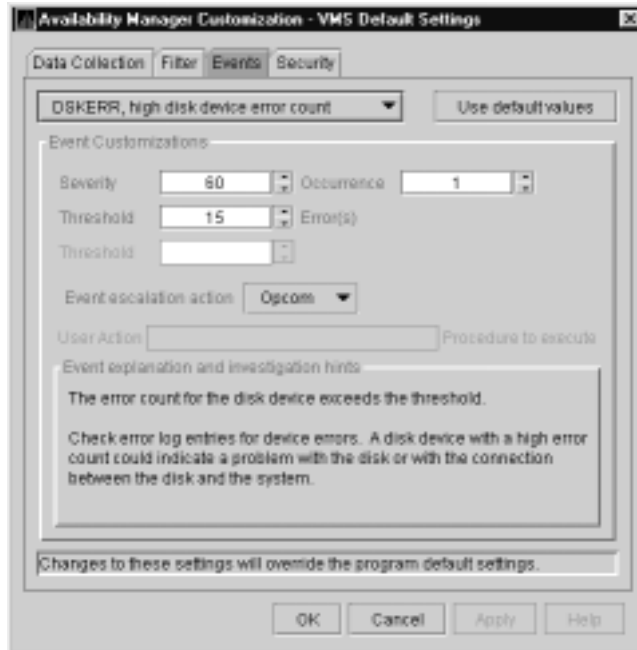
- Data collection signals an event if the appearance of that event exceeds the number of data collections on the Event customization page corresponding to

Getting Information about Events

4.2 Criteria Signaling and Displaying an Event

that event. For example, Figure 4–2 is a sample Event Customization page in which with an Occurrence value of 1.

Figure 4–2 Sample Event Customization Page



With an Occurrence value of 1, if an event exceeds its threshold at any data collection interval, it is signaled in the Events pane.

- When an event is signaled, data is collected at the **Event** interval shown on the Data Collection page (Figure 4–3).

Getting Information about Events

4.2 Criteria Signaling and Displaying an Event

Figure 4-3 OpenVMS Data Collection Page



- The data value displayed in the Node pane that is associated with the event turns red when an event is signaled (see Figure 4-4).

Figure 4-4 OpenVMS Node Pane

Node Name	CPU	MEM	BIO	DIO	CPU Qs	OS Version	HW Model
CMOVEQ	0	36	0	0	0	Open VMS V7.2-1	AlphaStation 200 4/166
DBGAVC	3	19	23	76	0	Open VMS V7.1-1H2	AlphaServer 2100A 5/300
DRINKS	0	44	1	0	0	Open VMS V7.1	VAXstation 4000-90
LOWFAT	0	38	0	0	0	Open VMS V6.2-1H3	AlphaServer 1000 4/200
MAWK	14	81	2	166	0	Open VMS V7.1	VAXstation 4000-60
PEROIT	1	31	1	0	0	Open VMS V7.1	VAXstation 4000-60
REDSOL	0	33	0	0	0	Open VMS V7.1-1H2	Digital Personal WorkStation
SELAVY	1	23	0	0	0	Open VMS V6.2	VAXstation 4000-60
SSPADE	6	51	3	0	0	Open VMS V7.1	VAXstation 4000-90

- An event, when signaled, is added to the Events Log File:

- On OpenVMS systems, the event log file is:

AMDS\$AM_LOG:AVAILMANEVENTS.LOG

A new version of this file is created each time you access the Availability Manager.

- On Windows NT systems, the events log is:

EVENTS.LOG

This file, which is in the installation directory, is overwritten each time you access the Availability Manager.

Getting Information about Events

4.2 Criteria Signaling and Displaying an Event

The following example shows a partial event log file:

```
VAXJET 01-22-1999 11:24:50.67 0 CFGDON VAXJET configuration done
DBGAVC 01-22-1999 11:25:12.41 0 CFGDON DBGAVC configuration done
AFFS5 01-22-1999 11:25:13.23 0 CFGDON AFFS5 configuration done
DBGAVC 01-22-1999 11:25:18.31 80 LCKCNT DBGAVC possible contention for resource REG$MASTER_LOCK
VAXJET 01-22-1999 11:25:27.47 40 LOBIOQ VAXJET LESSACP_V30 has used most of its BIOLM process quota
PEROIT 01-22-1999 11:25:27.16 0 CFGDON PEROIT configuration done
KOINE 01-22-1999 11:25:33.05 99 NOSWFL KOINE has no swap file
MAWK 01-22-1999 11:26:20.15 99 FXTIMO MAWK Fix timeout for FID to Filename Fix
MAWK 01-22-1999 11:26:24.48 60 HIDIOR MAWK direct I/O rate is high
REDSQL 01-22-1999 11:26:30.61 10 PRPGFL REDSQL _FTA2: high page fault rate
REDSQL 01-22-1999 11:26:31.18 60 PRPIOR REDSQL _FTA7: paging I/O rate is high
MAWK 01-22-1999 11:26:24.48 60 HIDIOR MAWK direct I/O rate is high
AFFS52 01-22-1999 11:25:33.64 60 DSKMNV AFFS52 $4$DUA320(OMTV4) disk mount verify in progress
VAXJET 01-22-1999 11:38:46.23 90 DPGERR VAXJET error executing driver program, ...
REDSQL 01-22-1999 11:39:18.73 60 PRCPWT REDSQL _FTA2: waiting in PWAIT
REDSQL 01-22-1999 11:44:37.19 75 PRCCUR REDSQL _FTA7: has a high CPU rate
```

4.3 Displaying Additional Event Information

For more detailed information about a specific event, double-click any event data item in the Events pane. The Availability Manager first displays a data page that most closely corresponds to the cause of the event. You can choose other tabs for additional detailed information.

For a description of data pages and the information they contain, see Chapter 3.

Performing Fixes on OpenVMS Nodes

You can perform **fixes** on OpenVMS nodes to resolve resource availability problems and improve system availability.

This chapter discusses the following topics:

- Understanding fixes
- Performing fixes

Caution

Performing certain fixes can have serious repercussions, including possible system failure. Therefore, only experienced system managers should perform fixes.

5.1 Understanding Fixes

When you suspect or detect a resource availability problem, in many cases you can use the Availability Manager to analyze the problem and to perform a fix to improve the situation.

Availability Manager fixes fall into two categories:

- Node fixes
- Process fixes

You can access fixes, by category, from the pages listed in Table 5–1.

Table 5–1 Accessing Availability Manager Fixes

Fix Category and Name	Available from This Page
Node fixes:	Node Summary
Crash node	CPU
Adjust cluster quorum	Memory
	I/O

(continued on next page)

Performing Fixes on OpenVMS Nodes

5.1 Understanding Fixes

Table 5–1 (Cont.) Accessing Availability Manager Fixes

Fix Category and Name	Available from This Page
Process fixes:	All of the process fixes are available from the following pages:
<ul style="list-style-type: none"> • General process fixes: <ul style="list-style-type: none"> Delete a process Exit an image Suspend a process Resume a process Change a process priority • Process memory fixes: <ul style="list-style-type: none"> Purge working set Adjust working set • Process limits fixes: <ul style="list-style-type: none"> Direct I/O Buffered I/O AST Open file Lock Timer I/O Byte 	<ul style="list-style-type: none"> • Memory • I/O • CPU Process • Single Process

Table 5–2 summarizes various problems, recommended fixes, and the expected results of fixes.

Table 5–2 Summary of Problems and Matching Fixes

Problem	Fix	Result
Node resource hanging cluster	Crash Node	Node fails with operator-requested shutdown.
Cluster hung	Adjust Quorum	Quorum for cluster is adjusted.
Process looping, intruder	Delete Process	Process no longer exists.
Endless process loop in same PC range	Exit Image	Exit from current image.
Runaway process, unwelcome intruder	Suspend Process	Process is suspended from execution.
Process previously suspended	Resume Process	Process starts from point it was suspended.
Runaway process or process that is overconsuming	Change Process Priority	Base priority changes to selected setting.
Low node memory	Purge Working Set	Frees memory on node; page faulting might occur for process affected.
Working set too high or low	Adjust Working Set	Removes unused pages from working set; page faulting might occur.
Process quota has reached its limit and has entered RWAIT state	Adjust Process Limits	Process receives greater limit, which in many cases frees the process to continue execution.

Most process fixes correspond to an OpenVMS system service call, as shown in the following table:

Performing Fixes on OpenVMS Nodes

5.1 Understanding Fixes

Process Fix	System Service Call
Delete a process	\$DELPRC
Exit an image	\$FORCEX
Suspend a process	\$SUSPND
Resume a process	\$RESUME
Change a process priority	\$SETPRI
Purge working set	\$PURGWS
Adjust working set	\$ADJWSL
Adjust process limits of the following:	None
Direct I/O (DIO)	
Buffered I/O (BIO)	
Asynchronous system trap (AST)	
Open file (FIL)	
Lock queue (ENQ)	
Timer queue entry (TQE)	
Subprocess (PRC)	
I/O byte (BYT)	

Note

Each fix that uses a system service call requires that the process execute the system service. A hung process will have the fix queued to it, where the fix will remain until the process is operational again.

Be aware of the following facts before you perform a fix:

- Fixes are optional.
- You must have write access to perform a fix.
- You cannot undo many fixes. For example, after using the crash node fix, the node must be rebooted (either by the node if the node reboots automatically, or by a person performing a manual boot).

5.2 Performing Fixes

Standard OpenVMS privileges restrict users' write access. When you run the Data Analyzer, you must have the CMKRNL privilege to send a write (fix) instruction to a node with a problem.

The following options are displayed at the bottom of all fix pages:

Option	Description
OK	Applies the fix and then exits the page. Any message associated with the fix is displayed in the Event pane.
Cancel	Cancels the fix.
Apply	Applies the fix and does not exit the page. Any message associated with the fix is displayed in the Return Status section of the page and in the Event pane.

The following sections explain how to perform nodes fixes and process fixes and describe specific fixes you can make.

Performing Fixes on OpenVMS Nodes

5.2 Performing Fixes

5.2.1 Node Fixes

The Availability Manager node fixes allow you to deliberately fail (crash) a node or to adjust cluster quorum.

To perform a node fix, follow these steps:

1. On the Node Summary, CPU, Memory, or I/O page, click the **Fix** menu.
2. Click **Fix Options**.

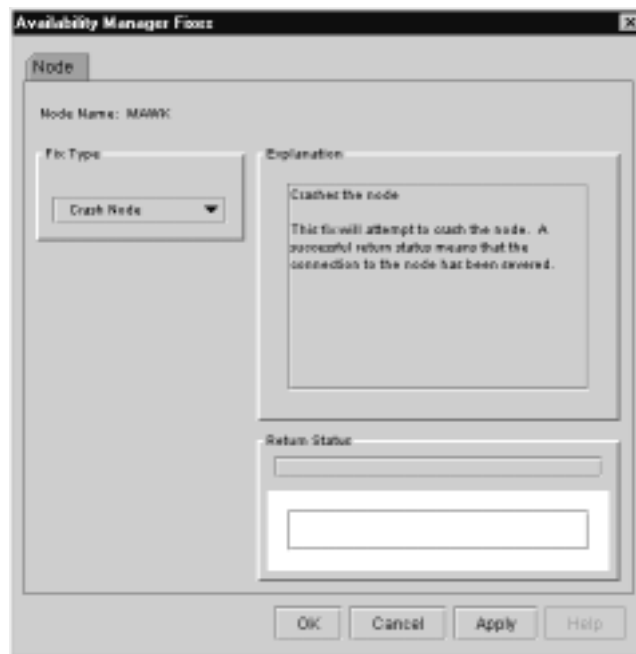
5.2.1.1 Crash Node

Caution

The crash node fix is an operator-requested bugcheck from the driver. It takes place as soon as you click **OK** in the Crash Node page. After you perform this fix, the node cannot be restored to its previous state. After a crash, the node must be rebooted.

When you select the Crash Node option, the Availability Manager displays the Crash Node page, shown in Figure 5–1.

Figure 5–1 Crash Node Page



Note

Because the node cannot report a confirmation when a node crash fix is successful, the crash success message is displayed after the timeout period for the fix confirmation has expired.

5.2.1.2 Adjust Quorum

The Adjust Quorum fix forces the node to refigure the quorum value. This fix is the equivalent of the Interrupt Priority C (IPC) mechanism used at system consoles for the same purpose. The fix forces the adjustment for the entire cluster so that each node in the cluster will have the same new quorum value.

The Adjust Quorum fix is useful when the number of votes in a cluster falls below the quorum set for that cluster. This fix allows you to readjust the quorum so that it corresponds to the current number of votes in the cluster.

When you select the Adjust Quorum option, the Availability Manager displays the page shown in Figure 5–2.

Figure 5–2 Adjust Quorum Page



5.2.2 Performing Process Fixes

To perform a process fix, follow these steps:

1. On the Memory or I/O page, right-click a process name.
2. Click **Fix Options**.

The Availability Manager displays three Process tabs:

- Process General
- Process Memory
- Process Limits

3. Click one of these tabs to bring it to the front.
4. Click the down arrow to display the process fixes in this group.
5. Select one process fix (for example, Change Process Priority, as shown in Figure 5–3,) to display a fix page.

Performing Fixes on OpenVMS Nodes

5.2 Performing Fixes

Figure 5–3 Change Process Priority Page



Some of the fixes, like Change Process Priority, require you to use a slider to change the default value. When you have finished setting a new process priority, click one of the options at the bottom of the page.

5.2.3 General Process Fixes

The following sections describe Availability Manager general process fixes.

5.2.3.1 Delete Process

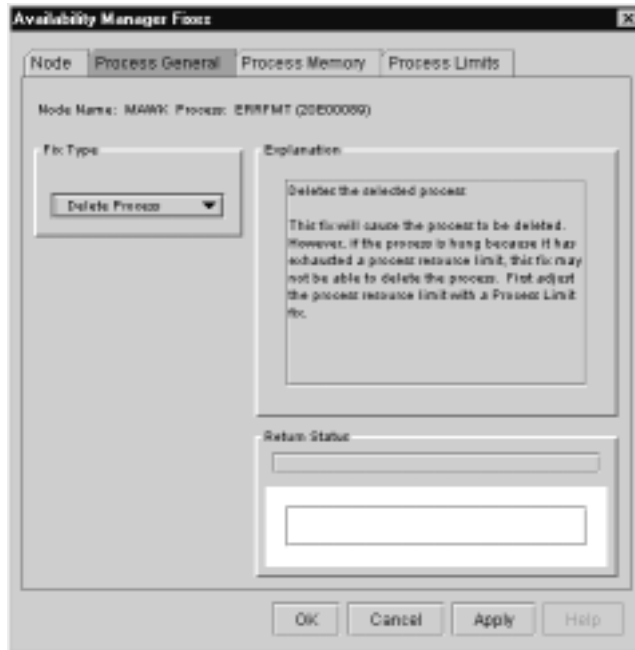
In most cases, a Delete Process fix deletes a process. However, if a process is waiting for disk I/O or is in a resource wait state (RWAST), this fix might not delete the process. In this situation, it is useless to repeat the fix. Instead, depending on the resource the process is waiting for, a Process Limit fix might free the process. As a last resort, reboot the node to delete the process.

Caution

Deleting a system process on a system process could cause the system to hang or become unstable.

When you select the Delete Process option, the Availability Manager displays the page shown in Figure 5–4.

Figure 5–4 Delete Process Page



After reading the explanation, select one of the options displayed at the bottom of the page. A message displayed on the page indicates that the fix has been successful.

5.2.3.2 Exit Image

Exiting an image on a node can stop an application that a user requires. Check the Single Process page first to determine which image is running on the node.

Caution

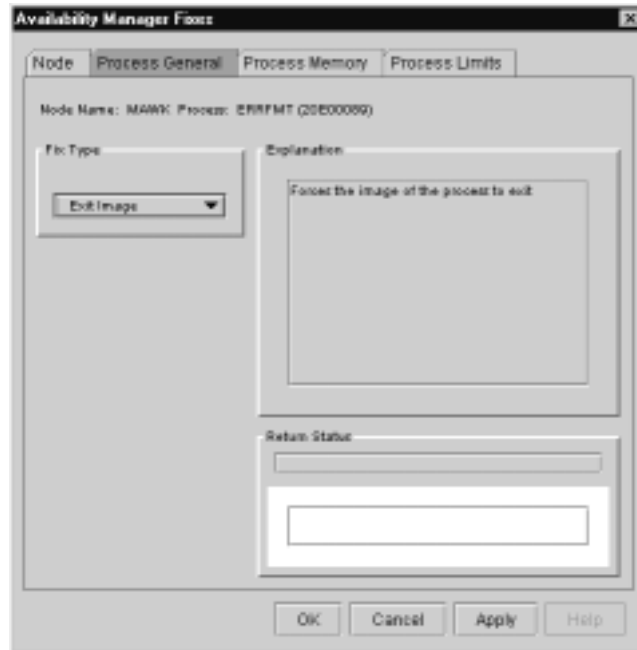
Exiting an image on a system process could cause the system to hang or become unstable.

Performing Fixes on OpenVMS Nodes

5.2 Performing Fixes

When you select the Exit Image option, the Availability Manager displays the page shown in Figure 5–5.

Figure 5–5 Exit Image Page



After reading the explanation in the page, select one of the options displayed at the bottom of the page. A message displayed on the page indicates that the fix has been successful.

5.2.3.3 Suspend Process

Suspending a process that is consuming excess CPU time can improve perceived CPU performance on the node by freeing the CPU for other processes to use. (Conversely, resuming a process that was using excess CPU time while running might reduce perceived CPU performance on the node.)

Caution

Do not suspend system processes, especially JOB_CONTROL, because this might make your system unusable. (See the *OpenVMS Programming Concepts Manual, Volume I* for more information.)

When you select the Suspend Process option, the Availability Manager displays the page shown in Figure 5–6.

Figure 5–6 Suspend Process Page



After reading the explanation, select one of the options displayed at the bottom of the page. A message displayed on the page indicates that the fix has been successful.

5.2.3.4 Resume Process

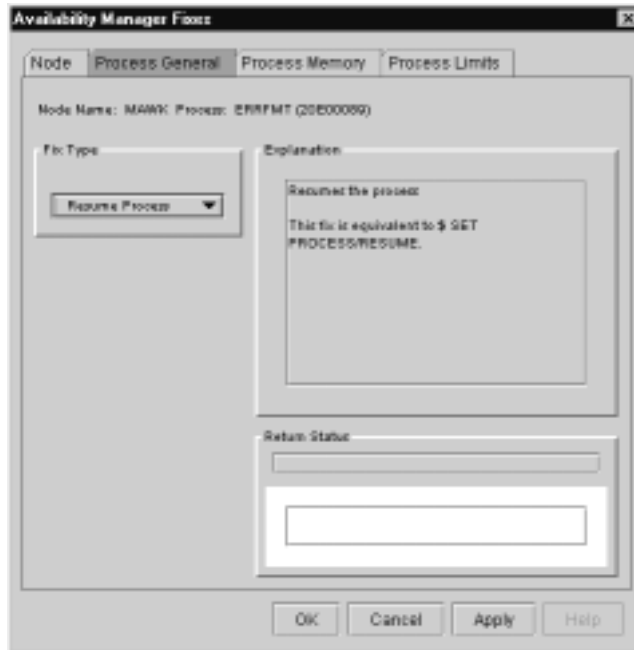
Resuming a process that was using excess CPU time while running might reduce perceived CPU performance on the node. (Conversely, suspending a process that is consuming excess CPU time can improve perceived CPU performance by freeing the CPU for other processes to use.)

Performing Fixes on OpenVMS Nodes

5.2 Performing Fixes

When you select the Resume Process option, the Availability Manager displays the page shown in Figure 5-7.

Figure 5-7 Resume Process Page



After reading the explanation, select one of the options displayed at the bottom of the page. A message displayed on the page indicates that the fix has been successful.

5.2.3.5 Change Process Priority

If the priority of a compute-bound process is too high, the process can consume all the CPU cycles on the node, affecting performance dramatically. On the other hand, if the priority of a process is too low, the process might not obtain enough CPU cycles to do its job, also affecting performance.

When you select the Process Priority option, the Availability Manager displays the page shown in Figure 5–8.

Figure 5–8 Change Process Priority Page



To change the base priority for a process, drag the slider on the scale to the number you want. The current priority number is displayed in a small box above the slider. You can also click the line above or below the slider to adjust the number by one.

When you are satisfied with the new base priority, select one of the options displayed at the bottom of the page. A message displayed on the page indicates that the fix has been successful.

5.2.4 Process Memory Fixes

The following sections describe the Availability Manager fixes you can use to correct process memory problems.

5.2.4.1 Purge Working Set

This fix purges the working set to a minimal size. You can use this fix to reclaim a process's pages that are not in active use. If the process is in a wait state, the working set remains at a minimal size, and the purged pages become available for other uses. If the process becomes active, pages the process needs are page-faulted back into memory, and the unneeded pages are available for other uses.

Be careful not to repeat this fix too often: a process that continually reclaims needed pages can cause excessive page faulting, which can affect system performance.

Performing Fixes on OpenVMS Nodes

5.2 Performing Fixes

When you select the Purge Working Set option, the Availability Manager displays the page shown in Figure 5–9.

Figure 5–9 Purge Working Set Page



After reading the explanation on the page, select one of the options displayed at the bottom of the page. A message displayed on the page indicates that the fix has been successful.

5.2.4.2 Adjust Working Set

Adjusting the working set of a process might prove to be useful in situations similar to the following ones:

- If a process is page-faulting because of insufficient memory, you can reclaim unused memory from other processes by decreasing the working set of one or more of them.
- If a process is page-faulting too frequently because its working set is too small, you can increase its working set.

When you select the Adjust Working Set fix, the Availability Manager displays the page shown in Figure 5–10.

Figure 5–10 Adjust Working Set Page



To perform this fix, use the slider to adjust the working set to the limit you want. You can also click the line above or below the slider to adjust the number by one.

When you are satisfied with the new working set limit, select one of the options displayed at the bottom of the page. A message displayed on the page indicates that the fix has been successful.

The following sections describe Availability Manager process limits fixes.

5.2.5 Process Limits Fixes

If a process is waiting for a resource, you can use a Process Limits fix to increase the resource limit so that the process can continue. The increased limit is in effect only for the life of the process, however; any new process is assigned the quota that was set in the UAF.

When you click the **Process Limits** tab, you can select any of the options described in the following sections.

Performing Fixes on OpenVMS Nodes

5.2 Performing Fixes

5.2.5.1 Direct I/O Count Limit

You can use this fix to adjust the direct I/O count limit of a process. When you select the Direct I/O option, the Availability Manager displays the page shown in Figure 5–11.

Figure 5–11 Direct I/O Count Limit Page



To perform this fix, use the slider to adjust the direct I/O count to the limit you want. You can also click the line above or below the slider to adjust the number by one.

When you are satisfied with the new direct I/O count limit, select one of the options displayed at the bottom of the page. A message displayed on the page indicates that the fix has been successful.

5.2.5.2 Buffered I/O Count Limit

You can use this fix to adjust the buffered I/O count limit of a process. When you select the Buffered I/O option, the Availability Manager displays the page shown in Figure 5–12.

Figure 5–12 Buffered I/O Count Limit Page



To perform this fix, use the slider to adjust the buffered I/O count to the limit you want. You can also click the line above or below the slider to adjust the number by one.

When you are satisfied with the new buffered I/O count limit, select one of the options displayed at the bottom of the page. A message displayed on the page indicates that the fix has been successful.

Performing Fixes on OpenVMS Nodes

5.2 Performing Fixes

5.2.5.3 AST Queue Limit

You can use this fix to adjust the AST queue limit of a process. When you select the AST option, the Availability Manager displays the page shown in Figure 5–13.

Figure 5–13 AST Queue Limit Page



To perform this fix, use the slider to adjust the AST queue limit to the number you want. You can also click the line above or below the slider to adjust the number by one.

When you are satisfied with the new AST queue limit, select one of the options displayed at the bottom of the page. A message displayed on the page indicates that the fix has been successful.

5.2.5.4 Open File Limit

You can use this fix to adjust the open file limit of a process. When you select the Open File option, the Availability Manager displays the page shown in Figure 5–14.

Figure 5–14 Open File Limit Page



To perform this fix, use the slider to adjust the open file limit to the number you want. You can also click the line above or below the slider to adjust the number by one.

When you are satisfied with the new open file limit, select one of the options displayed at the bottom of the page. A message displayed on the page indicates that the fix has been successful.


Performing Fixes on OpenVMS Nodes

5.2 Performing Fixes

5.2.5.5 Lock Queue Limit

You can use this fix to adjust the lock queue limit of a process. When you select the Lock option, the Availability Manager displays the page shown in Figure 5–15.

Figure 5–15 Lock Queue Limit Page



To perform this fix, use the slider to adjust the lock queue limit to the number you want. You can also click the line above or below the slider to adjust the number by one.

When you are satisfied with the new lock queue limit, select one of the options displayed at the bottom of the page. A message displayed on the page indicates that the fix has been successful.

5.2.5.6 Time Queue Entry Limit

You can use this fix to adjust the time queue entry limit of a process. When you select the Timer option, the Availability Manager displays the page shown in Figure 5-16.

Figure 5-16 Time Queue Entry Limit Page



To perform this fix, use the slider to adjust the timer queue entry limit to the number you want. You can also click the line above or below the slider to adjust the number by one.

When you are satisfied with the new timer queue entry limit, select one of the options displayed at the bottom of the page. A message displayed on the page indicates that the fix has been successful.

5.2.5.7 Subprocess Creation Limit

You can use this fix to adjust the creation limit of the subprocess of a process. When you select the Subprocess option, the Availability Manager displays the page shown in Figure 5-17.

Performing Fixes on OpenVMS Nodes

5.2 Performing Fixes

Figure 5–17 Subprocess Creation Limit Page



To perform this fix, use the slider to adjust a process's subprocess creation limit to the number you want. You can also click the line above or below the slider to adjust the number by one.

When you are satisfied with the new subprocess creation limit, select one of the options displayed at the bottom of the page. A message displayed on the page indicates that the fix has been successful.

5.2.5.8 I/O Byte

You can use this fix to adjust the I/O byte limit of a process. When you select the I/O Byte option, the Availability Manager displays the page shown in Figure 5–18.

Figure 5–18 I/O Byte Page



To perform this fix, use the slider to adjust the I/O byte limit to the number you want. You can also click the line above or below the slider to adjust the number by one.

When you are satisfied with the new I/O byte limit, select one of the options displayed at the bottom of the page. A message displayed on the page indicates that the fix has been successful.

Customizing the Availability Manager

This chapter explains how to customize the following features of the Availability Manager:

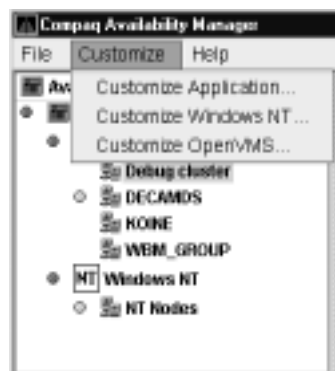
Feature	Description
Nodes or groups of nodes	You can select one or more groups or individual nodes to monitor.
Group membership	You can change a node's default group membership.
Types of data collection	For OpenVMS nodes, you can select the types of data you want to collect as well as several types of collection intervals. (On Windows NT nodes, specific types of data are collected by default.)
Data filters	For OpenVMS nodes, you can specify a number of parameters and values that limit the amount of data that is collected.
Event filters	You can specify the severity of events that are displayed as well as several other filtering settings.
Security features	On Data Analyzer and Data Collector nodes, you can change passwords. On OpenVMS Data Collector nodes, you can edit a file that contains security triplets.

How to Specify Which Nodes to Customize

In the Application window, you can select the Customize option in either of these places:

- On the Application window menu bar, click the **Customize** menu. The Availability Manager displays the menu shown in Figure 6–1.

Figure 6–1 Customize Menu in Application Window



- Right-click a node name in the Node pane of the Application window Figure 1–3. The Availability Manager displays the popup menu shown in Figure 6–2.

Customizing the Availability Manager

Figure 6–2 Customize Menu in Node Pane

Node Name	CPU	MEM	BIO	DIO	CPU Qs	OS Version	HW Model
CMOVEQ						Open VMS V7.2-1	AlphaStation 280 4/166
DBG						Open VMS V7.2-1	AlphaServer 2100A 5/300
DFRM						Open VMS V7.1	VAXstation 4000-90
LOW						Open VMS V6.2-1H3	AlphaServer 1000 4/200
MM						Open VMS V7.1	VAXstation 4000-60
PERM						Open VMS V7.1	VAXstation 4000-60
RED						Open VMS V7.2-1	Digital Personal WorkStation
SAL						Open VMS V7.2-1	DEC 4000 Model 6-90
SEL						Open VMS V6.2	VAXstation 4000-60
SSPADE	5	50	0			Open VMS V7.1	VAXstation 4000-90

Depending on which Customize menu you use and your choice of menu items, your customizations can affect one or more nodes, as indicated in the following table.

Nodes Affected	Action
All nodes	Click Customize Application... on the menu shown in Figure 6–1.
All Windows NT nodes	Click Customize Windows NT... on the menu shown in Figure 6–1.
All OpenVMS nodes	Click Customize OpenVMS... on the menu shown in Figure 6–1.
One node	Click Customize... on the popup menu shown in Figure 6–2. The customization options you choose will affect only the node you have selected.

6.1 Specifying Groups or Individual Nodes to Monitor

When you select **Customize Application** on the Application window Customize menu (Figure 6–1), the Availability Manager displays a **Group/Node Lists** tab (Figure 6–3).

Customizing the Availability Manager

6.1 Specifying Groups or Individual Nodes to Monitor

Figure 6–3 Group/Node Lists Customization Page



You can select either of the following lists of nodes to monitor:

- **Use Group List**

Using the correct case, enter the names of the groups of nodes you want to monitor.

- **Use Node List**

Using the correct case, enter the names of individual nodes you want to monitor.

If you decide to return to the default (monitor all nodes) or to enter names again, click **Use default values**.

After you enter a list of groups or individual nodes, click one of the following choices:

Option	Description
OK	Accepts the choice of names you have entered and exits the page.
Cancel	Cancels the choice of names and does not exit the page.
Apply	Accepts the choice of names you have entered but does not exit the page.

To put the list into effect, exit the Availability Manager and restart it.

Customizing the Availability Manager

6.2 Changing the Group Membership of a Node

6.2 Changing the Group Membership of a Node

Each Availability Manager Data Collector node is assigned to the DECamDS group by default. The following sections explain how to change the group membership of nodes.

Note

You need to place nodes that are in a cluster in the same group. If nodes within a cluster are in different groups, some of the data collected might be misleading.

6.2.1 Changing the Group of an OpenVMS Node

You need to edit a logical on each Data Collector node to change the group for that node. To do this, follow these steps:

1. Assign a unique name of up to 15 alphanumeric characters to the AMDSSGROUP_NAME logical name in the AMDSSAM_SYSTEM:AMDSSAM_LOGICALS.COM file. For example:

```
$ AMDSSDEF AMDSSGROUP_NAME FINANCE ! Group FINANCE; OpenVMS Cluster alias
```

2. Apply the logical name by restarting the Data Collector:

```
$ @SYS$STARTUP:AMDSSAM_STARTUP RESTART
```

Note

The configuration files for DECamds and the Availability Manager are separate; only one set is used, depending on which startup command procedure you use to start the driver.

Refer to the *Compaq Availability Manager Version 1.4 Installation Instructions for OpenVMS Alpha Systems* for further explanation of the configuration files set up for both DECamds and the Availability Manager.

6.2.2 Changing the Group of a Windows NT Node

You need to edit the Registry to change the group of a Windows NT node. To edit the Registry, follow these steps:

1. Click the Windows NT **Start** button. On the lists displayed, first click **Programs** and then **Command Prompt**.

2. Type REGEDIT after the angle prompt (>).

The system displays a screen for the Registry Editor, with a list of entries under My Computer.

3. On the list displayed, click **HKEY_LOCAL_MACHINE**.
4. Click **SYSTEM**.
5. Click **CurrentControlSet**.
6. Click **damdrv**.
7. Click **Parameters**.

Customizing the Availability Manager

6.2 Changing the Group Membership of a Node

8. Double-click **Group Name**. Then type a new group name of 15 alphanumeric characters or fewer, and click **OK** to make the change.
9. On the Control Panel, click **Services**, and then click **Stop** for “PerfServ.”
10. Again on the Control Panel, click **Devices**, and then click **Stop** for “damdrv.”
11. First restart **damdrv** under “Devices,” and then restart **PerfServ** under “Services.”

This step completes the change of groups for this node.

6.3 Customizing OpenVMS Data Collection

Note

Before you start this section, be sure to read the explanation of data collection, events, thresholds, and occurrences in Chapter 1.

When you click the **Customize OpenVMS...** menu option in the Application window and then click the **Data Collection** tab, the Availability Manager displays the OpenVMS Data Collection page, where you can choose types of data you want to collect. You can also change the default Availability Manager intervals at which data is collected or updated.

Figure 6–4 shows the OpenVMS Data Collection page.

Figure 6–4 OpenVMS Data Collection Page



Customizing the Availability Manager

6.3 Customizing OpenVMS Data Collection

Table 6–1 identifies the page where each type of data collected appears and indicates whether or not that type of data collection is a default.

Table 6–1 Data Collection Choices

Type of Data Collection	Default	Page That Data Is Collected For
Cluster summary data	No	Cluster Summary page
CPU mode data	No	CPU Modes Summary page
CPU summary data	No	CPU Process States page
Disk status data	No	Disk Status Summary page
Disk volume data	No	Disk Volume Summary page
I/O data	No	I/O Summary page
Lock contention data	No	Lock Contention page
Memory data	No	Memory Summary page
Node summary data	Yes	Node pane, Node Summary page, and the top pane of the CPU, Memory, and I/O pages
Page/Swap file data	No	I/O Page Faults page
Single disk data	Yes ¹	Single Disk Summary page
Single process data	Yes ²	Data collection for the Process Information page

¹Data is collected by default when you open a Single Disk Summary page.

²Data is collected by default when you open a Single Process page.

You can choose additional types of data collection by selecting the Collect check box on the Data Collection page of the **Customize OpenVMS...** menu (Figure 6–4). A check mark indicates that data will be collected at the intervals described in Table 6–2.

Table 6–2 Data Collection Intervals

Interval (in seconds)	Description
Display	How often the data display is refreshed.
Event	If events occur related to the screen, how often data is collected.
NoEvent	If no events occur related to the screen, how often data is collected.

You can enter a different collection interval by double-clicking the value; then delete the old value and enter a new one. If you change your mind and decide to return to the default collection interval, select one or more rows of data items; then click **Use default values**. The system displays the default values for all the collection intervals.

When you finish customizing your data collection, select one of the following options at the bottom of the page:

Option	Description
OK	To confirm any changes you have made and exit the page.

Option	Description
Cancel	To cancel any changes you have made and exit the page.
Apply	To confirm and apply any changes you have made and not exit the page.

6.4 Customizing OpenVMS Data Filters

When you select the **Customize OpenVMS...** menu option and click the **Filter** tab, the Availability Manager displays a page that allows you to display specific kinds of data by selecting check boxes. The types of data filters available are as follows:

- CPU Filters
- Disk Status
- Disk Volume
- I/O
- Lock Contention
- Memory Filters
- Page/Swap File

Filters can vary depending on the type of data collected. For example, filters might be process states or a variety of rates and counts. The following sections describe data filters that are available for various types of data collection.

Each filter page contains the following note: “Changes to these settings will override the program default settings.” This means that settings that are changed on the node level override settings made at the application and program (or GUI) level.

You can display the default values for the data items on the page at any time. To do this, select the item and click **Use default values**.

When you finish modifying a filters page, click one of the following options at the bottom of the page:

Option	Description
OK	To confirm any changes you have made and exit the page.
Cancel	To cancel any changes you have made and exit the page.
Apply	To confirm and apply any changes you have made and continue to display the page.

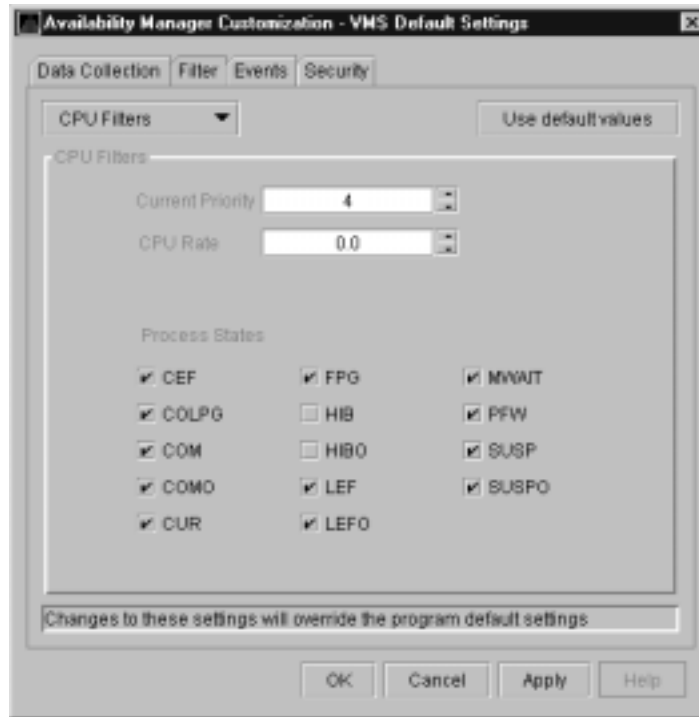
6.4.1 OpenVMS CPU Filters

When you click **CPU Filters** on the Filter page, the Availability Manager displays the OpenVMS CPU Filters page (Figure 6–5).

Customizing the Availability Manager

6.4 Customizing OpenVMS Data Filters

Figure 6–5 OpenVMS CPU Filters Page



This page allows you to change and select values that are displayed on the OpenVMS CPU Process States page (Figure 3–7).

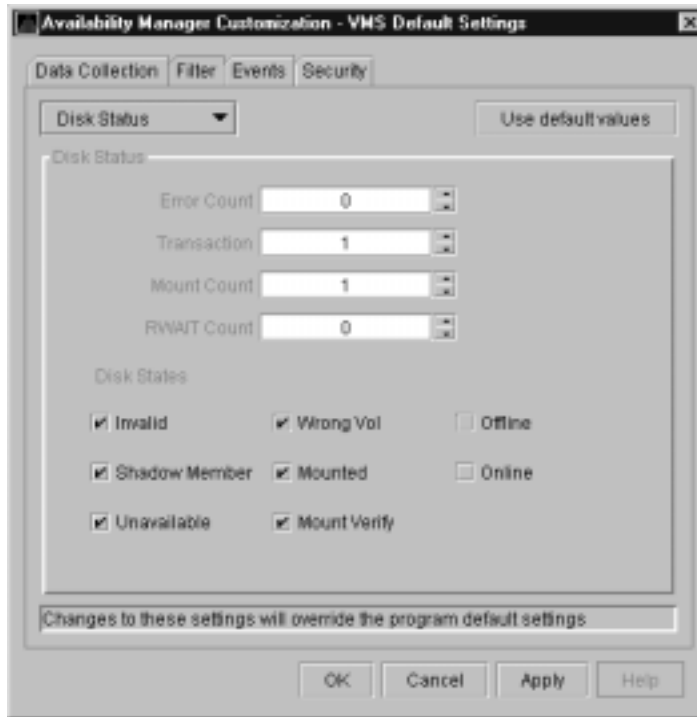
You can change the current priority and rate of a process. By default, a process is displayed only if it has a Current Priority of 4 or more. Click the up or down arrow to increase or decrease the priority value by one. The default CPU rate is 0.0, which means that processes with any CPU rate used will be displayed. To limit the number of processes displayed, you can click the up or down arrow to increase or decrease the CPU rate by .5 each time you click.

The OpenVMS CPU Filters page also allows you to select the states of the processes that you want to display on the CPU Process States page. Select the check box for each state you want to display. (Process states are described in Appendix B.)

6.4.2 OpenVMS Disk Status Filters

When you click **Disk Status** on the Filter page, the Availability Manager displays the OpenVMS Disk Status Filters page (Figure 6–6).

Figure 6–6 OpenVMS Disk Status Filters Page



The OpenVMS Disk Status Summary page (Figure 3–12) displays the values you set on this filters page.

This filters page lets you change the following default values:

Data	Description
Error Count	The number of errors generated by the disk (a quick indicator of device problems).
Transaction	The number of in-progress file system operations for the disk.
Mount Count	The number of nodes that have the specified disk mounted.
RWAIT Count	An indicator that a system I/O operation is stalled, usually during normal connection failure recovery or volume processing of host-based shadowing.

This filters page also lets you check the states of the disks you want to display, as described in the following table:

Disk State	Description
Invalid	Disk is in an invalid state (Mount Verify Timeout is likely).
Shadow Member	Disk is a member of a shadow set.
Unavailable	Disk is set to unavailable.
Wrong Vol	Disk was mounted with the wrong volume name.
Mounted	Disk is logically mounted by a MOUNT command or a service call.

Customizing the Availability Manager

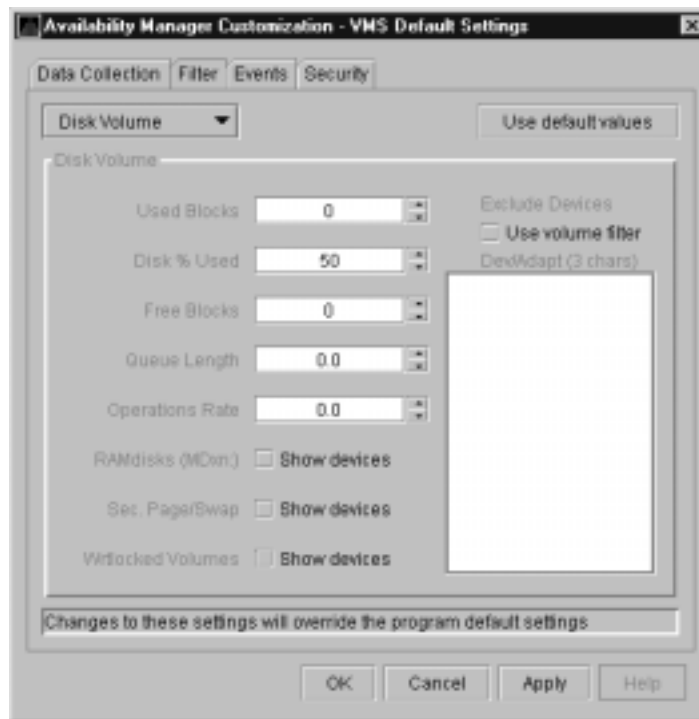
6.4 Customizing OpenVMS Data Filters

Disk State	Description
Mount Verify	Disk is waiting for a mount verification.
Offline	Disk is no longer physically mounted in device drive.
Online	Disk is physically mounted in device drive.

6.4.3 OpenVMS Disk Volume Filters

When you click **Disk Volume** on the Filter page, the Availability Manager displays the OpenVMS Disk Volume Filters page (Figure 6–7).

Figure 6–7 OpenVMS Disk Volume Filters Page



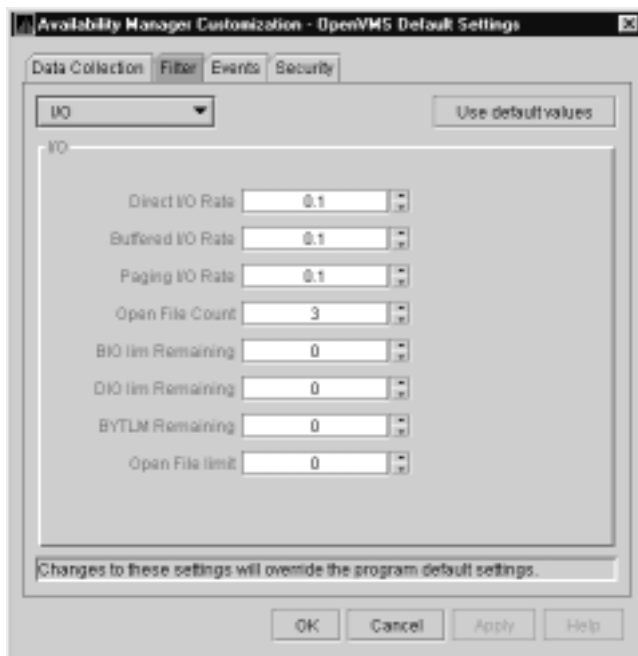
The OpenVMS Disk Volume Filters page allows you to change the values for the following data:

Data	Description
Used Blocks	The number of volume blocks in use.
Disk % Used	The percentage of the number of volume blocks in use in relation to the total volume blocks available.
Free Blocks	The number of blocks of volume space available for new data.
Queue Length	Current length of all queues.
Operations Rate	The rate at which the operations count to the volume has changed since the last sampling. The rate measures the amount of activity on a volume. The optimal load is device specific.

6.4.4 OpenVMS I/O Filters

When you click **I/O** on the Filter page, the Availability Manager displays the OpenVMS I/O Filters page (Figure 6–8).

Figure 6–8 OpenVMS I/O Filters Page



The OpenVMS I/O Summary page (Figure 3–11) displays the values you set on this filters page.

This filters page allows you to change values for the following data:

Data	Description
Direct I/O Rate	The rate at which I/O transfers occur between the system devices and the pages or pagelets that contain the process buffer that the system locks in physical memory.
Buffered I/O Rate	The rate at which I/O transfers occur between the process buffer and an intermediate buffer from the system buffer pool.
Paging I/O Rate	The rate of read attempts necessary to satisfy page faults (also known as Page Read I/O or the Hard Fault Rate).
Open File Count	The number of open files.
BIO lim Remaining	The number of remaining buffered I/O operations available before the process reaches its quota. BIOLM quota is the maximum number of buffered I/O operations a process can have outstanding at one time.
DIO lim Remaining	The number of remaining direct I/O limit operations available before the process reaches its quota. DIOLM quota is the maximum number of direct I/O operations a process can have outstanding at one time.

Customizing the Availability Manager

6.4 Customizing OpenVMS Data Filters


Data	Description
BYTLM Remaining	The number of buffered I/O bytes available before the process reaches its quota. BYTLM is the maximum number of bytes of nonpaged system dynamic memory that a process can claim at one time.
Open File limit	The number of additional files the process can open before reaching its quota. FILLM quota is the maximum number of files that can be opened simultaneously by the process, including active network logical links.

6.4.5 OpenVMS Lock Contention Filters

The OpenVMS Lock Contention Filters page allows you to remove (filter out) resource names from the Lock Contention page (Figure 3–17).

When you click **Lock Contention** on the Filter page, the Availability Manager displays the OpenVMS Lock Contention Filters page (Figure 6–9).

Figure 6–9 OpenVMS Lock Contention Filters Page




Each entry on the Lock Contention Filters page is a resource name or part of a resource name that you want to filter out. For example, the STRIPES entry filters out any value that starts with the characters STRIPES\$. To redisplay original entries, click **Use default values**.

6.4.6 OpenVMS Memory Filters

When you click **Memory Filters** on the Filter page, the Availability Manager displays the OpenVMS Memory Filters page (Figure 6–10).

Figure 6–10 OpenVMS Memory Filters Page



The OpenVMS Memory page (Figure 3–9) displays the values on this filters page.

The OpenVMS Memory Filters page allows you to change values for the following data:

Data	Description
Working Set Count	The number of physical pages or pagelets of memory that the process is using.
Working Set Size	The number of pages or pagelets of memory the process is allowed to use. The operating system periodically adjusts this value based on an analysis of page faults relative to CPU time used. An increase in this value in large units indicates a process is receiving a lot of page faults and its memory allocation is increasing.
Working Set Extent	The number of pages or pagelets of memory in the process's WSEXTENT quota as defined in the user authorization file (UAF). The number of pages or pagelets will not exceed the value of the system parameter WSMAX.
Page Fault Rate	The number of page faults per second for the process.
Page I/O Rate	The rate of read attempts necessary to satisfy page faults (also known as page read I/O or the hard fault rate).

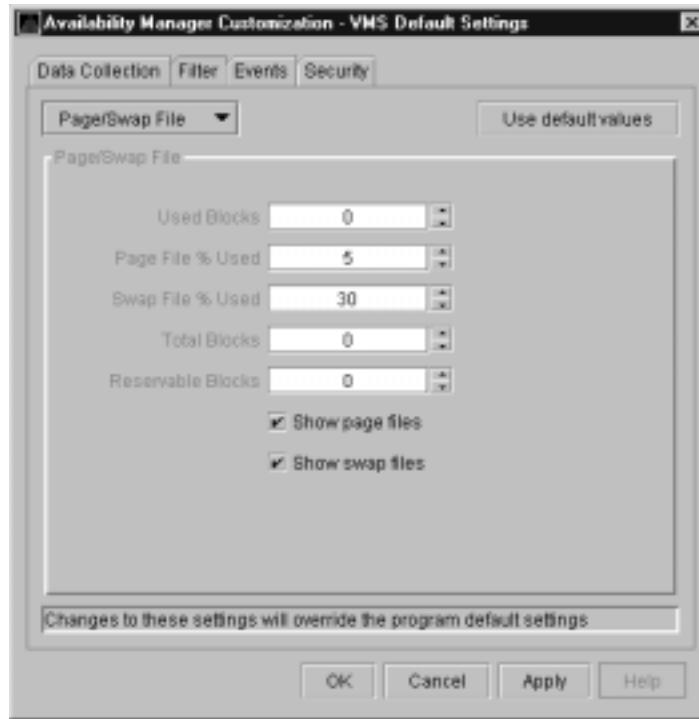
Customizing the Availability Manager

6.4 Customizing OpenVMS Data Filters

6.4.7 OpenVMS Page/Swap File Filters

When you click **Page/Swap File** on the Filter page, the Availability Manager displays the OpenVMS Page/Swap File Filters page (Figure 6–11).

Figure 6–11 OpenVMS Page/Swap File Filters Page



The OpenVMS I/O Page Faults page (Figure 3–11) displays the values that you set on this filters page.

This filters page allows you to change values for the following data:

Data	Description
Used Blocks	The number of used blocks within the file.
Page File % Used	The percentage of the blocks from the page file that have been used.
Swap File % Used	The percentage of the blocks from the swap file that have been used.
Total Blocks	The total number of blocks in paging and swapping files.
Reservable Blocks	Number of reservable blocks in each paging and swapping file currently installed. Reservable blocks can be logically claimed by a process for a future physical allocation. A negative value indicates that the file might be overcommitted. Note that a negative value is not an immediate concern but indicates that the file might become overcommitted if physical memory becomes scarce.

You can also select (turn on) or clear (turn off) the following options:

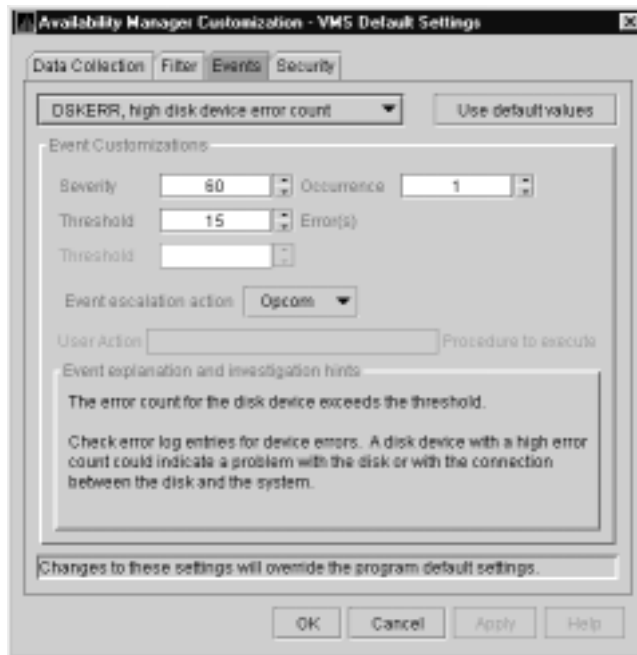
- Show page files
- Show swap files

6.5 Customizing Events

You can customize a number of characteristics of the events that are displayed in the Events pane of the Application window (Figure 1–3).

When you select the **Customize OpenVMS...** or **Customize Windows NT...** option from the Application window's **Customize** menu, the Availability Manager displays a tabbed page similar to the one shown in Figure 6–12.

Figure 6–12 Event Customization Page



You can change the values for any data that is available (does not appear dimmed) on this page. The following table describes the data you can change:

Data	Description
Severity	Controls the severity level at which events are displayed in the Event pane. By default, all events are displayed. Increasing this value reduces the number of event messages in the Event pane and can improve perceived response time.

Customizing the Availability Manager

6.5 Customizing Events

Data	Description
Occurrence	<p>Each Availability Manager event is assigned an occurrence value, that is, the number of consecutive data samples that must exceed the event threshold before the event is signaled. By default, events have low occurrence values. However, you might find that a certain event indicates a problem only when it occurs repeatedly for an extended period. You can change the occurrence value assigned to that event so that the Availability Manager signals it only when necessary.</p> <p>For example, suppose page fault spikes are common in your environment, and the Availability Manager frequently signals intermittent <i>HITTLP, total page fault rate is high</i> events. You could change the event's occurrence value to 3, so that the total page fault rate must exceed the threshold for three consecutive collection intervals before being signaled to the event log.</p> <p>To avoid displaying insignificant events, you can customize an event so that the Availability Manager signals it only when it occurs continuously.</p>
Threshold	<p>Most events are checked against only one threshold; however, some events have dual thresholds, where the event is triggered if either one is true. For example, for the <i>LOVLSP, node disk volume free space is low</i> event, the Availability Manager checks both of the following thresholds:</p> <ul style="list-style-type: none">• Number of blocks remaining (<i>LowDiskFreeSpace.BlkRem</i>)• Percentage of total blocks remaining (<i>LowDiskFreeSpace.Percent</i>)
Escalation	Not yet implemented.
User Action	Not yet implemented.

The section of the page called “Event explanation and investigation hints” includes a description of the event and a hint about how to correct any problems that the event signals.

6.6 Customizing Security Features

The following sections explain how to change these security features:

- Data Analyzer passwords for OpenVMS and Windows NT Data Collector nodes
- OpenVMS Data Collector security triplets
- A Windows NT Data Collector password

Note

OpenVMS Data Collector nodes can have more than one password: each password is part of a security triplet. (Windows NT nodes allow you to have only one password per node.)

6.6.1 Changing Data Analyzer Passwords

You can change the passwords that the Windows NT Data Analyzer uses for OpenVMS Data Collector nodes and for Windows NT Data Collector nodes. The following sections explain how to perform both actions.

Customizing the Availability Manager

6.6 Customizing Security Features

6.6.1.1 Changing a Data Analyzer Password for an OpenVMS Data Collector Node

When you click **Customize OpenVMS...** on the **Customize** menu of the Application window, the Availability Manager displays a **Security** page (Figure 6–13).

Figure 6–13 OpenVMS Security Customization Page



To change the default password for the Data Analyzer to use to access OpenVMS Data Collector nodes, enter a password of exactly 8 uppercase alphanumeric characters. The Availability Manager will use this password to access OpenVMS Data Collector nodes. This password must match the password that is part of the OpenVMS Data Collector security triplet (see Section 1.4).

When you are satisfied with your password, click **OK**. Exit the Availability Manager, and restart the application for the password to take effect.

6.6.1.2 Changing a Data Analyzer Password for a Windows NT Data Collector Node

When you click **Customize Windows NT...** on the **Customize** menu of the Application window, the Availability Manager displays a **Security** page (Figure 6–14).

Customizing the Availability Manager

6.6 Customizing Security Features

Figure 6–14 Windows NT Security Customization Page



To change the default password for the Data Analyzer to use to access Windows NT Data Collector nodes, enter a password of exactly 8 alphanumeric characters. Note that this password is case sensitive; any time you type it, you must use the original capitalization.

This password must also match the password for the Windows NT Data Collector node that you want to access. (See Section 6.6.3 for instructions for changing that password.)

When you are satisfied with your password, click **OK**. Exit and restart the Availability Manager for the password to take effect.

6.6.2 Changing Security Triplets on OpenVMS Data Collector Nodes

To change security triplets on an OpenVMS Data Collector node, you must edit the `AMDS$DRIVER_ACCESS.DAT` file, which is installed on all Data Collector nodes. The following sections explain what a security triplet is, how the Availability Manager uses it, and how to change it.

6.6.2.1 Understanding OpenVMS Security Triplets

A security triplet determines which nodes can access system data from an OpenVMS Data Collector node. The `AMDS$DRIVER_ACCESS.DAT` file on OpenVMS Data Collector nodes lists security triplets.

On OpenVMS Data Collector nodes, the `AMDS$AM_CONFIG` logical translates to the location of the default security file, `AMDS$DRIVER_ACCESS.DAT`. This file is installed on all OpenVMS Data Collector nodes.

A security triplet is a three-part record whose fields are separated by backslashes (`\`). A triplet consists of the following fields:

- A network address (hardware address or wildcard character)

Customizing the Availability Manager

6.6 Customizing Security Features

- An 8-character alphanumeric password
The password is not case sensitive (so the passwords “testtest” and “TESTTEST” are considered to be the same).
- A read or write (R or W) access verification code

The exclamation point (!) is a comment delimiter; any characters to the right of the comment delimiter are ignored.

Example

All Data Collector nodes in group FINANCE have the following AMDS\$DRIVER_ACCESS.DAT file:

```
*\FINGROUP\R ! Let anyone with FINGROUP password read
!
2.1\DEVGROUP\W ! Let only DECnet node 2.1 with
! DEVGROUP password perform fixes (writes)
```

6.6.2.2 How to Change a Security Triplet

Note

The configuration files for DECams and the Availability Manager are separate; only one set is used, depending on which startup command procedure you use to start the driver.

See *Installing the Availability Manager Version 1.4 on OpenVMS Alpha Systems and Running DECams and the Availability Manager Concurrently* for a further explanation of the configuration file setup for both DECams and the Availability Manager.

On each Data Collector node on which you want to change security, you must edit the AMDS\$DRIVER_ACCESS.DAT file. The data in the AMDS\$DRIVER_ACCESS.DAT file is set up as follows:

```
Network address\password\access
```

Use a backslash character (\) to separate the three fields.

To edit the AMDS\$DRIVER_ACCESS.DAT file, follow these steps:

1. Edit the network address.

The network address can be either of the following:

- Hardware address

The hardware address field is the physical hardware address in the LAN adapter chip. It is used if you have multiple LAN adapters or are running the Compaq DECnet-Plus for OpenVMS networking software on the system (not the Compaq DECnet Phase IV for OpenVMS networking software).

For adapters provided by Compaq, the hardware address is in the form 08-00-2B-xx-xx-xx, where the 08-00-2B portion is Compaq's valid range of LAN addresses as defined by the IEEE 802 standards, and the xx-xx-xx portion is chip specific.

To determine the value of the hardware address on a node, use the OpenVMS System Dump Analyzer (SDA) as follows:

```
$ ANALYZE/SYSTEM
SDA> SHOW LAN
```

Customizing the Availability Manager

6.6 Customizing Security Features

These commands display a list of available devices. Choose the template device of the LAN adapter you will be using, and then enter the following command:

```
SDA> SHOW LAN/DEVICE=xxA0
```

- **Wildcard address**

The wildcard character (*) allows any incoming triplet with a matching password field to access the Data Collector node. Use the wildcard character to allow read access and to run the console application from any node in your network.

Because the Data Analyzer does not use this field, you should use the wildcard character in this field in the AMDS\$CONSOLE_ACCESS.DAT file.

Caution: Use of the wildcard character for write-access security triplets enables any person using that node to perform system-altering fixes.

2. Edit the password field.

The password field **must be** an 8-byte alphanumeric field.

The password field gives you a second level of protection when you want to use the wildcard address denotation to allow multiple modes of access to your monitored system.

3. Enter either R or W as an access code.

R means READONLY allowance for the Data Analyzer. W means READ/WRITE allowance for the Data Analyzer. (WRITE implies READ.)

OpenVMS Data Collector nodes accept more than one password. Therefore, you might have several security triplets in an AMDS\$DRIVER_ACCESS.DAT file for one Data Collector node. For example:

```
*\1DECAMDS\R  
*\KOINECLS\R  
*\KOINEFIX\W
```

In this example, Data Analyzer nodes with the passwords 1DECAMDS and KOINECLS would be able to see the Data Collector data, but only the Data Analyzer node with the KOINEFIX password would be able to write or change information, including performing fixes, on the Data Collector node.

If you want, you can set up your AMDS\$DRIVER_ACCESS.DAT file to allow anyone in the world to read from your system but allow only a certain node or nodes to write or change information on your system.

Note

After editing the AMDS\$DRIVER_ACCESS.DAT file, you must stop and then restart the Data Collector. This action loads the new data into the driver.

6.6.2.3 How the Availability Manager Ensures Security

The Availability Manager performs these steps when using security triplets to ensure security among Data Analyzer and Data Collector nodes:

1. A message is broadcast at regular intervals to all nodes within the LAN indicating the availability of a Data Collector node to communicate with a Data Analyzer node.
2. The node running the Data Analyzer receives the availability message and returns a security triplet that identifies it to the Data Collector, and requests system data from the Data Collector.
3. The Data Collector examines the security triplet to determine whether the Data Analyzer is listed in the AMDSS\$DRIVER_ACCESS.DAT file to permit access to the system.
 - If the AMDSS\$DRIVER_ACCESS.DAT file lists Data Analyzer access information, then the Data Provider and the Data Analyzer can exchange information.
 - If the Data Analyzer is not listed in the AMDSS\$DRIVER_ACCESS.DAT file or does not have appropriate access information, then access is denied and a message is logged to OPCOM. The Data Analyzer receives a message stating that access to that node is not permitted.

Table 6–3 describes how the Data Collector node interprets a security triplet match.

Table 6–3 Security Triplet Verification

Security Triplet	Interpretation
08-00-2B-12-34-56\HOMETOWN\W	The Data Analyzer has write access to the node only when the Data Analyzer is run from a node with this hardware address (multiadapter or DECnet-Plus system) and with the password HOMETOWN.
2.1\HOMETOWN\R	The Data Analyzer has read access to the node when run from a node with DECnet for OpenVMS Phase IV address 2.1 and the password HOMETOWN.
*\HOMETOWN\R	Any Data Analyzer with the password HOMETOWN has read access to the node.

6.6.3 Changing a Password on a Windows NT Data Collector

To change the Data Collector password in the Registry, follow these steps:

1. Click the Windows NT **Start** button. First click **Programs** and then **Command Prompt**.
2. Type `regedit` after the angle prompt (>).
The system displays a screen for the Registry Editor, with a list of entries under My Computer.
3. On the list displayed, click **HKEY_LOCAL_MACHINE**.
4. Click **SYSTEM**.
5. Click **CurrentControlSet**.

Customizing the Availability Manager

6.6 Customizing Security Features

6. Click **Services**.
7. Click **damdrv**.
8. Click **Parameters**.
9. Double-click **Read Password**. Then type a new 8-character alphanumeric password, and click **OK** to make the change.
10. To store the new password, click **Exit** under File on the main menu bar.
11. On the Control Panel, click **Services** and then **Stop** for “PerfServ.”
12. Again on the Control Panel, click **Devices** and then **Stop** for “damdrv.”
13. First restart **damdrv** under “Devices” and then restart **PerfServ** under “Services.”

This step completes the change of your Data Collector password.

CPU Process States

The CPU process states shown in the following table are displayed in the OpenVMS CPU Process States page (see Figure 3-7) and in the OpenVMS Process Information page (see Figure 3-27).

Table A-1 CPU Process States

Process State	Description
CEF	Common Event Flag, waiting for a common event flag
COLPG	Collided Page Wait, involuntary wait state; likely to indicate a memory shortage, waiting for hard page faults
COM	Computable; ready to execute
COMO	Computable Outswapped, COM, but swapped out
CUR	Current, currently executing in a CPU
FPW	Free Page Wait, involuntary wait state; most likely indicates a memory shortage
LEF	Local Event Flag, waiting for a Local Event Flag
LEFO	Local Event Flag Outswapped; LEF, but outswapped
HIB	Hibernate, voluntary wait state requested by the process; it is inactive
HIBO	Hibernate Outswapped, hibernating but swapped out

(continued on next page)

CPU Process States

Table A-1 (Cont.) CPU Process States

Process State	Description																																				
MWAIT	Miscellaneous Resource Wait, involuntary wait state, possibly caused by a shortage of a systemwide resource, such as no page or swap file capacity or no synchronizations for single-threaded code. Types of MWAIT states are shown in the following table:																																				
	<table border="1"> <thead> <tr> <th>MWAIT State</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>BYTLM Wait</td> <td>Process waiting for buffered I/O byte count quota.</td> </tr> <tr> <td>JIB Wait</td> <td>Process in either BYTLM Wait or TQELM Wait state.</td> </tr> <tr> <td>TQELM Wait</td> <td>Process waiting for timer queue entry quota.</td> </tr> <tr> <td>EXH</td> <td>Kernel thread in exit handler (not currently used).</td> </tr> <tr> <td>INNER_MODE</td> <td>Kernel thread waiting to acquire inner-mode semaphore.</td> </tr> <tr> <td>PSXFR</td> <td>Process waiting during a POSIX fork operation.</td> </tr> <tr> <td>RWAST</td> <td>Process waiting for system or special kernel mode AST.</td> </tr> <tr> <td>RWMBX</td> <td>Process waiting because mailbox is full.</td> </tr> <tr> <td>RWNBX</td> <td>Process waiting for nonpaged dynamic memory.</td> </tr> <tr> <td>RWPFF</td> <td>Process waiting because page file is full.</td> </tr> <tr> <td>RWPAG</td> <td>Process waiting for paged dynamic memory.</td> </tr> <tr> <td>RWMPE</td> <td>Process waiting because modified page list is empty.</td> </tr> <tr> <td>RWMPB</td> <td>Process waiting because modified page writer is busy.</td> </tr> <tr> <td>RWSCS</td> <td>Process waiting for distributed lock manager.</td> </tr> <tr> <td>RWCLU</td> <td>Process waiting because OpenVMS Cluster is in transition.</td> </tr> <tr> <td>RWCAP</td> <td>Process waiting for CPU that has its capability set.</td> </tr> <tr> <td>RWCSV</td> <td>Kernel thread waiting for request completion by OpenVMS Cluster server process.</td> </tr> </tbody> </table>	MWAIT State	Definition	BYTLM Wait	Process waiting for buffered I/O byte count quota.	JIB Wait	Process in either BYTLM Wait or TQELM Wait state.	TQELM Wait	Process waiting for timer queue entry quota.	EXH	Kernel thread in exit handler (not currently used).	INNER_MODE	Kernel thread waiting to acquire inner-mode semaphore.	PSXFR	Process waiting during a POSIX fork operation.	RWAST	Process waiting for system or special kernel mode AST.	RWMBX	Process waiting because mailbox is full.	RWNBX	Process waiting for nonpaged dynamic memory.	RWPFF	Process waiting because page file is full.	RWPAG	Process waiting for paged dynamic memory.	RWMPE	Process waiting because modified page list is empty.	RWMPB	Process waiting because modified page writer is busy.	RWSCS	Process waiting for distributed lock manager.	RWCLU	Process waiting because OpenVMS Cluster is in transition.	RWCAP	Process waiting for CPU that has its capability set.	RWCSV	Kernel thread waiting for request completion by OpenVMS Cluster server process.
MWAIT State	Definition																																				
BYTLM Wait	Process waiting for buffered I/O byte count quota.																																				
JIB Wait	Process in either BYTLM Wait or TQELM Wait state.																																				
TQELM Wait	Process waiting for timer queue entry quota.																																				
EXH	Kernel thread in exit handler (not currently used).																																				
INNER_MODE	Kernel thread waiting to acquire inner-mode semaphore.																																				
PSXFR	Process waiting during a POSIX fork operation.																																				
RWAST	Process waiting for system or special kernel mode AST.																																				
RWMBX	Process waiting because mailbox is full.																																				
RWNBX	Process waiting for nonpaged dynamic memory.																																				
RWPFF	Process waiting because page file is full.																																				
RWPAG	Process waiting for paged dynamic memory.																																				
RWMPE	Process waiting because modified page list is empty.																																				
RWMPB	Process waiting because modified page writer is busy.																																				
RWSCS	Process waiting for distributed lock manager.																																				
RWCLU	Process waiting because OpenVMS Cluster is in transition.																																				
RWCAP	Process waiting for CPU that has its capability set.																																				
RWCSV	Kernel thread waiting for request completion by OpenVMS Cluster server process.																																				
PFW	Page Fault Wait, involuntary wait state; possibly indicates a memory shortage, waiting for hard page faults.																																				
RWAST	Resource Wait State, waiting for delivery of an asynchronous system trap (AST) that signals a resource availability; usually an I/O is outstanding or a process quota is exhausted.																																				
RWBRK	Resource Wait for BROADCAST to finish																																				
RWCAP	Resource Wait for CPU Capability																																				
RWCLU	Resource Wait for Cluster Transition																																				
RWCSV	Resource Wait for Cluster Server Process																																				

(continued on next page)

Table A-1 (Cont.) CPU Process States

Process State	Description
RWIMG	Resource Wait for Image Activation Lock
RWLCK	Resource Wait for Lock ID data base
RWMBX	Resource Wait on MailBox, either waiting for data in mailbox (to read) or waiting to place data (write) into a full mailbox (some other process has not read from it; mailbox is full so this process cannot write).
RWMPB	Resource Wait for Modified Page writer Busy
RWMPE	Resource Wait for Modified Page list Empty
RWNPG	Resource Wait for Non Paged Pool
RWPAG	Resource Wait for Paged Pool
RWPFF	Resource Wait for Page File Full
RWQUO	Resource Wait for Pooled Quota
RWSCS	Resource Wait for System Communications Services
RWSWP	Resource Wait for Swap File space
SUSP	Suspended, wait state process placed into suspension; it can be resumed at the request of an external process
SUSPO	Suspended Outswapped, suspended but swapped out

Tables of Events

This appendix contains the following tables of events:

- OpenVMS Events
- Windows NT Events

Each table provides the following information:

- Alphabetical list of the events that the Availability Manager signals in the Events pane of the Application window (see Figure 1–3)
- Abbreviation and brief description of each event (also displayed in the Events pane)
- Explanation of the event and a suggestion for remedial action, if applicable

Table B–1 OpenVMS Events

Event	Description	Explanation	Recommended Action
CFGDON	Configuration done	The server application has made a connection to the node and will start collecting the data according to the Customize Data Collection options.	This informational event indicates that the node is recognized. No further investigation is required.
DPGERR	Error executing driver program	The Data Collector has detected a program error while executing the data collection program.	This event can occur if you have a bad driver program library, or there is a bug in the driver program. Make sure you have the program library that shipped with the kit; if it is correct, contact your customer support representative with the full text of the event.
DSKERR	High disk error count	The error count for the disk device exceeds the threshold.	Check error log entries for device errors. A disk device with a high error count could indicate a problem with the disk or with the connection between the disk and the system.
DSKINV	Disk is invalid	The valid bit in the disk device status field is not set. The disk device is not considered valid by the operating system.	Make sure that the disk device is valid and is known to the operating system.

(continued on next page)

Tables of Events

Table B-1 (Cont.) OpenVMS Events

Event	Description	Explanation	Recommended Action
DSKMNV	Disk in mount verify state	The disk device is performing a mount verification.	<p>The system is performing a mount verification for the disk device. This could be caused by:</p> <ul style="list-style-type: none"> • A removable disk on a local or remote node was removed. • A disk on a local or remote node has gone offline due to errors. • The node that serves the disk is down. • The connection to a remote disk is down.
DSKOFF	Disk device is off line	The disk device has been placed in the off line state.	Check whether the disk device should be off line. This event is also signalled when the same device name is used for two different physical disks. The volume name in the event is the second node to use the same device name.
DSKQLN	High disk queue length	The average number of pending I/Os to the disk device exceeds the threshold.	More I/O requests are being queued to the disk device than the device can service. Reasons include a slow disk or too much work being done on the disk.
DSKRWT	High disk RWAIT count	The RWAIT count on the disk device exceeds the threshold.	RWAIT is an indicator that an I/O operation has stalled, usually during normal connection failure recovery or volume processing of host-based shadowing. A node has probably failed and shadowing is recovering data.
DSKUNA	Disk device is unavailable	The disk device has been placed in the Unavailable state.	The disk device state has been set to /NOAVAILABLE. See DCL help for the SET DEVICE/AVAILABLE command.
DSKWRV	Wrong volume mounted	The disk device has been mounted with the wrong volume label.	Set the correct volume name by entering the DCL command SET VOLUME/LABEL on the node.
ELIBCR	Bad CRC for exportable program library	The CRC calculation for the exportable program library does not match the CRC value in the library.	The exportable program library may be corrupt. Restore the exportable program library from its original source.
ELIBNP	No privilege to access exportable program library	Unable to access the exportable program library.	Check to make sure that the Availability Manager has the proper security access to the exportable program library file.
ELIBUR	Unable to read exportable program library	Unable to read the exportable program library for the combination of hardware architecture and OpenVMS version.	The exportable program library may be corrupt. Restore the exportable program library from its original source.

(continued on next page)

Table B–1 (Cont.) OpenVMS Events

Event	Description	Explanation	Recommended Action
FXCPKT	Received a corrupt fix response packet from node	The Availability Manager tried to perform a fix, but the fix acknowledgment from the node was corrupt.	This event could occur if there is network congestion or some problem with the node. Confirm the connection to the node, and reapply the fix if necessary.
FXCRSH	Crash node fix	The Availability Manager has successfully performed a Crash Node fix on the node.	This informational message indicates a successful fix. Expect to see a Path Lost event for the node.
FXDCPR	Decrement process priority fix	The Availability Manager has successfully performed a Decrement Process Priority fix on the process.	This informational message indicates a successful fix. Setting a process priority too low takes CPU time away from the process.
FXDCWS	Decrement process working set size fix	The Availability Manager has successfully decreased the working set size of the process on the node by performing an Adjust Working Set fix.	This informational message indicates a successful fix. This fix disables the automatic working set adjustment for the process.
FXDLPR	Delete process fix	The Availability Manager has successfully performed a Delete Process fix on the process.	This informational message indicates a successful fix. If the process is in RWAST state, this fix does not work. This fix also does not work on processes created with the no delete option.
FXEXIT	Exit image fix	The Availability Manager has successfully performed an Exit Image fix on the process.	This informational message indicates a successful fix. Forcing a system process to exit its current image can corrupt the kernel.
FXINPR	Increment process priority fix	The Availability Manager has successfully performed an Increment Process Priority fix on the process.	This informational message indicates a successful fix. Setting a process priority too high takes CPU time away from other processes. Set the priority above 15 only for “real-time” processing.
FXINQU	Increment process quota limits fix	The Availability Manager has successfully increased the quota limit of the process on the node by placing a new limit value in the limit field of the quota.	This informational message indicates a successful fix. This fix is only for the life of the process. If the problem continues, change the limit for the account in the UAF file.
FXINWS	Increment process working set size fix	The Availability Manager has successfully increased the working set size of the process on the node by performing an Adjust Working Set fix.	This informational message indicates a successful fix. This fix disables the automatic working set adjustment for the process. The adjusted working set value cannot exceed WSQUOTA for the process or WSMAX for the system.
FXNOPR	No-change process priority fix	The Availability Manager has successfully performed a Process Priority fix on the process that resulted in no change to the process priority.	This informational message indicates a successful fix. The Fix Value slider was set to the current priority of the process.

(continued on next page)

Tables of Events

Table B–1 (Cont.) OpenVMS Events

Event	Description	Explanation	Recommended Action
FXNOQU	No-change process quota limits fix	The Availability Manager has successfully performed a quota limit fix for the process that resulted in no change to the quota limit.	This informational message indicates a successful fix. The Fix Value slider was set to the current quota of the process.
FXNOWS	No-change process working set size fix	The Availability Manager has successfully performed Adjust Working Set fix on the process.	This informational message indicates a successful fix. The Fix Value slider was set to the current working set size of the process.
FXPGWS	Purge working set fix	The Availability Manager has successfully performed a Purge Working Set fix on the process.	This informational message indicates a successful fix. The purged process might page fault to retrieve memory it needs for current processing.
FXPRIV	No privilege to attempt fix	The Availability Manager cannot perform a fix on the node due either to no CMKRNL privilege or to unmatched security triplets.	See Chapter 6 for details about setting up security.
FXQUOR	Adjust quorum fix	The Availability Manager has successfully performed an Adjust Quorum fix on the node.	This informational message indicates a successful fix. Use this fix when you find many processes in RWCAP state on a cluster node.
FXRESM	Resume process fix	The Availability Manager has successfully performed a Resume Process fix on the process.	This informational message indicates a successful fix. If the process goes back into suspend state, check the AUDIT_SERVER process for problems.
FXSUSP	Suspend process fix	The Availability Manager has successfully performed a Suspend Process fix on the process.	This informational message indicates a successful fix. Do not suspend system processes.
FXTIMO	Fix timeout	The Availability Manager tried to perform a fix, but no acknowledgment for the fix was received from the node within the timeout period.	This event can occur if there is network congestion, if some problem is causing the node not to respond, or if the fix request failed to reach the node. Confirm the connection to the node, and reapply the fix if necessary.
FXUERR	Unknown error code for fix	The Availability Manager tried to perform a fix, but the fix failed for an unexpected reason.	Please contact your Compaq customer support representative with the text of this event. The event text is also recorded in the event log.
HIBIOR	High buffered I/O rate	The node's average buffered I/O rate exceeds the threshold.	A high buffered I/O rate can cause high system overhead. If this is affecting overall system performance, use the I/O Summary to determine the high buffered I/O processes, and adjust their priorities or suspend them as needed.
HICOMQ	Many processes waiting in COM or COMO	The average number of processes on the node in the COM or COMO queues exceeds the threshold.	Use the CPU Mode Summary to determine which processes are competing for CPU resources. Possible adjustments include changing process priorities and suspending processes.

(continued on next page)

Table B-1 (Cont.) OpenVMS Events

Event	Description	Explanation	Recommended Action
HIDIOR	High direct I/O rate	The average direct I/O rate on the node exceeds the threshold.	A high direct I/O rate can cause high system overhead. If this is affecting overall system performance, use the I/O Summary to determine the high direct I/O processes, and adjust their priorities or suspend them as needed.
HIHRDP	High hard page fault rate	The average hard page fault rate on the node exceeds the threshold.	A high hard page fault indicates that the free or modified page list is too small. Check Chapter 6 for possible actions.
HIMWTQ	Many processes waiting in MWAIT	The average number of processes on the node in the Miscellaneous Resource Wait (MWAIT) queues exceeds the threshold.	Use the CPU and Single Process pages to determine which resource is awaited. See Chapter 6 for more information about wait states.
HINTER	High interrupt mode time	The average percentage of time the node spends in interrupt mode exceeds the threshold.	Consistently high interrupt time prohibits processes from obtaining CPU time. Determine which device or devices are overusing this mode.
HIPWIO	High paging write I/O rate	The average paging write I/O rate on the node exceeds the threshold.	Use the Process I/O and Memory Summary pages to determine which processes are writing to the page file excessively, and decide whether their working sets need adjustment.
HIPWTQ	Many processes waiting in COLPG, PFW, or FPG	The average number of processes on the node that are waiting for page file space exceeds the threshold.	Use the CPU Process States and Memory Summary to determine which processes are in the COLPG, PFW, or FPG state. COLPG and PFW processes might be constrained by too little physical memory, too restrictive working set quotas, or lack of available page file space. FPG processes indicate too little physical memory is available.
HISYSP	High system page fault rate	The node's average page fault rate for pageable system areas exceeds the threshold.	These are page faults from pageable sections in loadable executive images, page pool, and the global page table. The system parameter SYSMWCNT might be set too low. Use AUTOGEN to adjust this parameter.
HITTLP	High total page fault rate	The average total page fault rate on the node exceeds the threshold.	Use the Memory Summary to find the page faulting processes, and make sure that their working sets are set properly.
HMPSYN	High multiprocessor (MP) synchronization mode time	The average percentage of time the node handles multiprocessor (MP) synchronization exceeds the threshold.	High synchronization time prevents other devices and processes from obtaining CPU time. Determine which device is overusing this mode.
KTHIMD	Kernel thread waiting for inner-mode semaphore	The average percentage of time that the kernel thread waits for the inner-mode semaphore exceeds the threshold.	Use SDA to determine which kernel thread of the process has the semaphore.

(continued on next page)

Tables of Events

Table B–1 (Cont.) OpenVMS Events

Event	Description	Explanation	Recommended Action
LCKBLK	Lock blocking	The process holds the highest priority lock in the resource's granted lock queue. This lock is blocking all other locks from gaining access to the resource.	Use the Single Process Windows to determine what the process is doing. If the process is in an RW _{xxx} state, try exiting the image or deleting the process. If this fails, crashing the blocking node might be the only other fix option.
LCKCNT	Lock contention	The resource has a contention situation, with multiple locks competing for the same resource. The competing locks are the currently granted lock and those that are waiting in the conversion queue or in the waiting queue.	Use Lock Contention to investigate a potential lock contention situation. Locks for the same resource might have the NODLCKWT wait flag enabled and be on every member of the cluster. Usually this is not a lock contention situation, and these locks can be filtered out.
LCKWAT	Lock waiting	The process that has access to the resource is blocking the process that is waiting for it. Once the blocking process releases its access, the next highest lock request acquires the blocking lock.	If the blocking process holds the resource too long, check to see whether the process is working correctly; if not, one of the fixes might solve the problem.
LOASTQ	Process has used most of ASTLM quota	Either the remaining number of asynchronous system traps (ASTs) the process can request is below the threshold, or the percentage of ASTs used compared to the allowed quota is above the threshold.	If the amount used reaches the quota, the process enters RWAST state. If the process requires a higher quota, you can increase the ASTLM quota for the process in the UAF file. ASTLM is only a count; system resources are not compromised by increasing this count.
LOBIOQ	Process has used most of BIOLM quota	Either the remaining number of Buffered I/Os (BIO) the process can request is below the threshold, or the percentage of BIOs used is above the threshold.	If the amount used reaches the quota, the process enters RWAST state. If the process requires a higher quota, you can increasing the BIOLM quota for the process in the UAF file. BIOLM is only a count; system resources are not compromised by increasing this count.
LOBYTQ	Process has used most of BYTLM quota	Either the remaining number of bytes for the buffered I/O byte count (BYTCNT) that the process can request is below the threshold, or the percentage of bytes used is above the threshold.	If the amount used reaches the quota, the process enters RWAST state. If the process requires a higher quota, you can raise the BYTLM quota for the process in the UAF file. BYTLM is the number of bytes in nonpaged pool used for buffered I/O.
LODIOQ	Process has used most of DIOLM quota	Either the remaining number of Direct I/Os (DIOs) the process can request is below the threshold, or the percentage of DIOs used is above the threshold.	If the amount used reaches the quota, the process enters RWAST state. If the process requires a higher quota, you can increase the DIOLM quota for the process in the UAF file. DIOLM is only a count; system resources are not compromised by increasing this count.
LOENQU	Process has used most of ENQLM quota	Either the remaining number of lock enqueues (ENQ) the process can request is below the threshold, or the percentage of ENQs used is above the threshold.	If the limit reaches the quota, the process is not able to make further lock queue requests. If the process requires a higher quota, you can increase the ENQLM quota for the process in the UAF file.

(continued on next page)

Table B-1 (Cont.) OpenVMS Events

Event	Description	Explanation	Recommended Action
LOFILQ	Process has used most of FILLM quota	Either the remaining number of files the process can open is below the threshold, or the percentage of files open is above the threshold.	If the amount used reaches the quota, the process must first close some files before being allowed to open new ones. If the process requires a higher quota, you can increase the FILLM quota for the process in the UAF file.
LOMEMY	Free memory is low	For the node, the percentage of free memory compared to total memory is below the threshold.	Use the automatic Purge Working Set fix, or use the Memory and CPU Summary to select processes that are either not currently executing or not page faulting, and purge their working sets.
LOPGFQ	Process has used most of PGFLQUOTA quota	Either the remaining number of pages the process can allocate from the system page file is below the threshold, or the percentage of pages allocated is above the threshold.	If the process requires a higher quota, you can raise the PGFLQUOTA quota for the process in the UAF file. This value limits the number of pages in the system page file that the account's processes can use.
LOPGSP	Low page file space	Either the remaining number of pages in the system page file is below the threshold, or the percentage of page file space remaining is below the threshold.	Either extend the size of this page file or create a new page file to allow new processes to use the new page file.
LOPRCQ	Process has used most of PRCLM quota	Either the remaining number of subprocesses the current process is allowed to create is below the threshold, or the percentage of created subprocesses is above the threshold.	If the amount used reaches the quota, the process is not allowed to create more subprocesses. If the process requires a higher quota, you can increase the PRCLM quota for the process in the UAF file.
LOSTVC	Lost virtual circuit to node	The virtual circuit between the listed nodes has been lost.	Check to see whether the second node listed has failed or whether the connection between the nodes is broken. The VC name listed in parentheses is the communication link between the nodes.
LOSWSP	Low swap file space	Either the remaining number of pages in the system page file is below the threshold, or the percentage of page file space remaining is below the threshold.	Either increase the size of this page file, or create a new page file to allow new processes to use the new page file.
LOTQEQ	Process has used most of TQELM quota	Either the remaining number of Time Queue Entries (TQEs) the process can request is below the threshold, or the percentage of TQEs used to the allowed quota is above the threshold.	If the amount used reaches the quota, the process enters RWAST state. If the process requires a higher quota, you can raise the TQELM quota for the process in the UAF file. TQELM is only a count; system resources are not compromised by raising it.
LOVLSP	Low disk volume free space	Either the remaining number of blocks on the volume is below the threshold, or the percentage of free blocks remaining on the volume is below the threshold.	You must free up some disk volume space. If part of the purpose of the volume is to be filled, such as a page/swap device, then you can filter the volume from the display.

(continued on next page)

Tables of Events

Table B–1 (Cont.) OpenVMS Events

Event	Description	Explanation	Recommended Action
LOVOTE	Low cluster votes	The difference between the number of VOTES and the QUORUM in the cluster is below the threshold.	Check to see whether voting members have failed. To avoid the hang that results if VOTES goes below QUORUM, use the Adjust Quorum fix.
LOWEXT	Low process working set extent	The process page fault rate exceeds the threshold, and the percentage of working set size compared to working set extent exceeds the threshold.	This event indicates that the WSEXTENT value in the UAF file might be too low. The process needs more physical memory but cannot obtain it; therefore, the process page faults excessively.
LOWSQU	Low process working set quota	The process page fault rate exceeds the threshold, and the percentage of working set size exceeds the threshold.	This event indicates the process needs more memory but might not be able to obtain it because one of the following is true: <ul style="list-style-type: none"> • The WSQUOTA value in the UAF file is set too low for the size of memory allocation requests or • The system is memory constrained.
LRGHSH	Remote lock hash table too large to collect data on	The Availability Manager cannot investigate the node's resource hash table (RESHASHTBL). It is either too sparse or too dense to investigate efficiently.	This event indicates that the Availability Manager will take too many collection iterations to analyze lock contention situations efficiently. Make sure that the SYSGEN parameter RESHASHTBL is set properly for the node.
NOPGFL	No page file	The Availability Manager cannot find a page file on the node.	Use SYSGEN to create and connect a page file on the node.
NOPLIB	No program library	The program library for the combination of hardware architecture and OpenVMS version was not found.	Check to see that all the program library files exist in the program library directory.
NOPRIV	Not allowed to monitor node	The Availability Manager cannot monitor the node due to unmatched security triplets.	See Chapter 6 for details on setting up security.
NOPROC	Specific process not found	The Availability Manager cannot find the process name selected in the Process Name Search dialog box on the Node Summary page.	This event can occur because the listed process no longer exists, or the process name is listed incorrectly in the dialog box.
NOSWFL	No swap file	The Availability Manager cannot find a swap file on the node.	If you do not use swap files, you can ignore this event. Otherwise, use SYSGEN to create and connect a swap file for the node.
PKTFER	Packet format error	The data packet sent to the remote node was not in the correct format for the remote node to process.	Please contact your Compaq customer support representative with the full text of the event, the version of the Availability Manager, the configuration of the node running the Availability Manager, and the configuration of the nodes being monitored.

(continued on next page)

Table B–1 (Cont.) OpenVMS Events

Event	Description	Explanation	Recommended Action
PLIBNP	No privilege to access program library	Unable to access the program library.	Check to see that the Availability Manager has the proper security access to the program library file.
PLIBUR	Unable to read program library	Unable to read the program library for the combination of hardware architecture and OpenVMS version.	The program library is either corrupt or from a different version of the Availability Manager. Restore the program library from the last installation.
PRBIOR	High process buffered I/O rate	The average buffered I/O rate of the process exceeds the threshold.	If the buffered I/O rate is affecting overall system performance, lowering the process priority or suspending the process would allow other processes to obtain access to the CPU.
PRBIOW	Process waiting for buffered I/O	The average percentage of time the process is waiting for a buffered I/O to complete exceeds the threshold.	Use SDA on the node to ensure that the device to which the process is performing buffered I/Os is still available and is not being overused.
PRCCOM	Process waiting in COM or COMO	The average number of processes on the node in the COM or COMO queues exceeds the threshold.	Use the CPU Summary to determine which processes should be given more CPU time, and adjust process priorities and states accordingly.
PRCCUR	Process has a high CPU rate	The average percentage of time the process is currently executing in the CPU exceeds the threshold.	Make sure that the listed process is not looping or preventing other processes from gaining access to the CPU. Adjust process priority or state as needed.
PRCMUT	Process waiting for a mutex	The average percentage of time the process is waiting for a particular system mutex exceeds the threshold.	Use SDA to help determine which mutex the process is waiting for and to help determine the owner of the mutex.
PRCMWT	Process waiting in MWAIT	The average percentage of time the process is in a Miscellaneous Resource Wait (MWAIT) state exceeds the threshold.	Various resource wait states are part of the collective wait state called MWAIT. See Appendix A for a list of these states. The CPU Process page and the Single Process page display which state the process is in. Check the Single Process page to determine which resource the process is waiting for and whether the resource is still available for the process.
PRCPSX	Process waiting in PSXFR	The average percentage of time the process waits during a POSIX fork operation exceeds the threshold.	
PRCPUL	Most of CPULIM process quota used	The remaining CPU time available for the process is below the threshold.	Make sure the CPU time allowed for the process is sufficient for its processing needs. If not, increase the CPU quota in the UAF file of the node.
PRCPWT	Process waiting in COLPG, PFW or FPG	The average percentage of time the process is waiting to access the system page file database exceeds the threshold.	Check to make sure the system page file is large enough for all the resource requests being made.

(continued on next page)

Tables of Events

Table B–1 (Cont.) OpenVMS Events

Event	Description	Explanation	Recommended Action
PRCQUO	Process waiting for a quota	The average percentage of time the process is waiting for a particular quota exceeds the threshold.	Use the Single Process pages to determine which quota is too low. Then adjust the quotas of the account in the UAF file.
PRCRWA	Process waiting in RWAST	The average percentage of time the process is waiting in the RWAST state exceeds the threshold. RWAST indicates the process is waiting for an asynchronous system trap to complete.	Use the Single Process pages to determine if RWAST is due to the process quota being set too low. If not, use SDA to determine if RWAST is due to a problem between the process and a physical device.
PRCRWC	Process waiting in RWCAP	The average percentage of time the process is waiting in the RWCAP state exceeds the threshold. RWCAP indicates that the process is waiting for CPU capability.	When many processes are in this state, the system might be hung because not enough nodes are running in the cluster to maintain the cluster quorum. Use the Adjust Quorum fix to correct the problem.
PRCRWM	Process waiting in RWMBX	The average percentage of time the process is waiting in the RWMBX state exceeds the threshold. RWMBX indicates the process is waiting for a full mailbox to be empty.	Use SDA to help determine which mailbox the process is waiting for.
PRCRWP	Process waiting in RWPAG, RWNPG, RWMPE, or RWMPB	The average percentage of time the process is waiting in the RWPAG, RWNPG, RWMPE, or RWMPB state exceeds the threshold. RWPAG and RWNPG are for paged or nonpaged pool; RWMPE and RWMPB are for the modified page list.	Processes in the RWPAG or RWNPG state can indicate you need to increase the size of paged or nonpaged pool, respectively. Processes in the RWMPB state indicate that the modified page writer cannot handle all the modified pages being generated. Refer to Chapter 6 for suggestions.
PRCRWS	Process waiting in RWSCS, RWCLU, or RWCSV	The average percentage of time the process is waiting in the RWSCS, RWCLU, or RWCSV state exceeds the threshold. RWCSV is for the cluster server; RWCLU is for the cluster transition; RWSCS is for cluster communications. The process is waiting for a cluster event to complete.	Use the Show Cluster utility to help investigate.
PRCUNK	Process waiting for a system resource	The average percentage of time the process is waiting for an undetermined system resource exceeds the threshold.	The state in which the process is waiting is unknown to the Availability Manager.
PRDIOR	High process direct I/O rate	The average direct I/O rate of the process exceeds the threshold.	If the I/O rate is affecting overall system performance, lowering the process priority might allow other processes to obtain access to the CPU.
PRDIOW	Process waiting for direct I/O	The average percentage of time the process is waiting for a direct I/O to complete exceeds the threshold.	Use SDA on the node to ensure that the device to which the process is performing direct I/Os is still available and is not being overused.

(continued on next page)

Table B-1 (Cont.) OpenVMS Events

Event	Description	Explanation	Recommended Action
PRLCKW	Process waiting for a lock	The average percentage of time the process is waiting in the control wait state exceeds the threshold.	The control wait state indicates that a process is waiting for a lock. Although no locks might appear in Lock Contention, the awaited lock might be filtered out of the display.
PRPGFL	High process page fault rate	The average page fault rate of the process exceeds the threshold.	The process is memory constrained; it needs an increased number of pages to perform well. Make sure that the working set quotas and extents are set correctly. To increase the working set quota temporarily, use the Adjust Working Set fix.
PRPRIOR	High process paging I/O rate	The average page read I/O rate of the process exceeds the threshold.	The process needs an increased number of pages to perform well. Make sure that the working set quotas and extents are set correctly. To increase the working set quota temporarily, use the Adjust Working Set fix.
PTHLST	Path lost	The connection between the server and collection node has been lost.	Check to see whether the node failed or whether the LAN segment to the node is having problems. This event occurs when the server no longer receives data from the node on which data is being collected.
RESDNS	Resource hash table dense	The percentage of occupied entries in the hash table exceeds the threshold.	A densely populated table can result in a performance degradation. Use the system parameter RESHASHTBL to adjust the total number of entries.
RESPRS	Resource hash table sparse	The percentage of occupied entries in the hash table is less than the threshold.	A sparsely populated table wastes memory resources. Use the system parameter RESHASHTBL to adjust the total number of entries.
UEXPLB	Using OpenVMS program export library	The program library for the combination of hardware architecture and OpenVMS version was not found.	Check to see that all the program library files exist in the program library directory.
UNSUPP	Unsupported node	The Availability Manager does not support this combination of hardware architecture and OpenVMS version.	Check the product SPD for supported system configurations.

Tables of Events

Table B–2 Windows NT Events

Event	Description	Explanation	Recommended Action
CFGDON	Configuration done	The server application has made a connection to the node and will start collecting the data according to the Customize Data Collection options.	An informational event to indicate that the node is recognized. No further investigation is required.
NODATA	Unable to collect performance data	The Availability Manager is unable to collect performance data from the node.	The performance data is collected by the PerfServ service on the remote node. Check to see that the service is up and running properly.
NOPRIV	Not allowed to monitor node	The Availability Manager cannot monitor the node due to a password mismatch between the Data Collector and the Data Analyzer.	See Chapter 6 for details on setting up security.
PTHLST	Path lost	The connection between the Data Analyzer and the Data Collector has been lost.	Check if the node crashed or if the LAN segment to the node is having problems. This event occurs when the server no longer receives data from the node on which data is being collected.
PVRMIS	Packet version mismatch	This version of the Availability Manager is unable to collect performance data from the node because of a data packet version mismatch.	The version of the Availability Manager Data Collector is more recent than the Data Analyzer. To process data from the node, upgrade the Data Analyzer to correspond to the Data Collector.

OpenVMS Events by Type of Data Collection

This appendix shows the events that can be signaled for each type of OpenVMS data collected. The events have been categorized as follows:

- Threshold events: Table C-1
- Nonthreshold events: Table C-2

Appendix B describes these events in detail and provides recommended actions.

Table C-1 OpenVMS Threshold Events

Type of Data Collected	Event	Description
Disk status	DSKERR	High disk error count
	DSKINV	Disk is invalid
	DSKMNV	Disk in mount verify state
	DSKOFF	Disk device is off line
	DSKRWT	High disk RWAIT count
	DSKUNA	Disk device is unavailable
	DSKWRV	Wrong volume mounted
Disk volume	DSKQLN	High disk queue length
	LOVLSP	Low disk volume free space
Node summary	HIBIOR	High buffered I/O rate
	HICOMQ	Many processes waiting in COM or COMO
	HIDIOR	High direct I/O rate
	HIHRDP	High hard page fault rate
	HIMWTQ	Many processes waiting in MWAIT
	HINTER	High interrupt mode time
	HIPWIO	High paging write I/O rate
	HIPWTQ	Many processes waiting in COLPG, PFW, or FPG
	HISYSP	High system page fault rate
	HITTLF	High total page fault rate
	HMPSYN	High multiprocessor (MP) synchronization mode time
	LOMEMY	Free memory is low

(continued on next page)

OpenVMS Events by Type of Data Collection

Table C-1 (Cont.) OpenVMS Threshold Events

Type of Data Collected	Event	Description
Lock contention	LCKCNT	Lock contention
	LRGSHS	Remote lock hash table too large to collect data
	RESDNS	Resource hash table dense
	RESPRS	Resource hash table sparse
Single lock	LCKBLK	Lock blocking
	LCKWAT	Lock waiting
Single process	KTHIMD	Kernel thread waiting for inner-mode semaphore
	LOASTQ	Process has used most of ASTLM quota
	LOBIOQ	Process has used most of BIOLM quota
	LOBYTQ	Process has used most of BYTLM quota
	LODIOQ	Process has used most of DIOLM quota
	LOENQU	Process has used most of ENQLM quota
	LOFILQ	Process has used most of FILLM quota
	LOPGFQ	Process has used most of PGFLQUOTA quota
	LOPRCQ	Process has used most of PRCLM quota
	LOTQEQ	Process has used most of TQELM quota
	LOWEXT	Low process working set extent
	LOWSQU	Low process working set quota
	PRBIOR	High process buffered I/O rate
	PRBIOW	Process waiting for buffered I/O
	PRCCOM	Process waiting in COM or COMO
	PRCCUR	Process has a high CPU rate
	PRCMUT	Process waiting for a mutex
	PRCPSX	POSIX process waiting during fork operation
	PRCPUL	Most of CPULIM process quota used
	PRCPWT	Process waiting in COLPG, PFW, or FPG
	PRCQUO	Process waiting for a quota
	PRCRWA	Process waiting in RWAST
	PRCRWC	Process waiting in RWCAP
	PRCRWM	Process waiting in RWMBX
	PRCRWP	Process waiting in RWPAG, RWNPG, RWMPE, or RWMPB
	PRCRWS	Process waiting in RWSCS, RWCLU, or RWCSV
	PRCUNK	Process waiting for a system resource
PRDIOR	High process direct I/O rate	

(continued on next page)

OpenVMS Events by Type of Data Collection

Table C-1 (Cont.) OpenVMS Threshold Events

Type of Data Collected	Event	Description
	PRDIOW	Process waiting for direct I/O
	PRLCKW	Process waiting for a lock
	PRPGFL	High process page fault rate
	PRPIOR	High process paging I/O rate
Process I/O	LOBIOQ	Process has used most of BIOLM quota
	LOBYTQ	Process has used most of BYTLM quota
	LODIOQ	Process has used most of DIOLM quota
	LOFILQ	Process has used most of FILLM quota
	PRBIOR	High process buffered I/O rate
	PRDIOR	High process direct I/O rate
	PRPIOR	High process paging I/O rate
Page/swap file	LOPGSP	Low page file space
	LOSWSP	Low swap file space
	NOPGFL	No page file
	NOSWFL	No swap file
Cluster summary	LOVOTE	Low cluster votes
Memory	LOWEXT	Low process working set extent
	LOWSQU	Low process working set quota
	PRPGFL	High process page fault rate
	PRPIOR	High process paging I/O rate
CPU process	PRCCOM	Process waiting in COM or COMO
	PRCCUR	Process has a high CPU rate
	PRCMWT	Process waiting in MWAIT (See Appendix A for a breakdown of MWAIT state.)
	PRCPWT	Process waiting in COLPG, PFW, or FPG

OpenVMS Events by Type of Data Collection

Table C-2 OpenVMS Nonthreshold Events

Type of Data Collected	Event	Description
Node-level event	CFGDON	Configuration done
	DPGERR	Error executing driver program
	NOPRIV	Not allowed to monitor node
	PKTFER	Packet format error
Program library error	ELIBCR	Bad CRC for exportable program library
	ELIBNP	No privilege to access exportable program library
	ELIBUR	Unable to read exportable program library
	NOPLIB	No program library
	PLIBNP	No privilege to access program library
	PLIBUR	Unable to read program library
	UEXPLB	Using exportable program library
	UNSUPP	Unsupported node
Events generated by fixes	FXCPKT	Received a corrupt fix response packet from node
	FXCRSH	Crash node fix
	FXDCPR	Decrement process priority fix
	FXDCWS	Decrement process working set size fix
	FXDLPR	Delete process fix
	FXEXIT	Exit image fix
	FXINPR	Increment process priority fix
	FXINQU	Increment process quota limits fix
	FXINWS	Increment process working set size fix
	FXPGWS	Purge working set fix
	FXPRIV	No privilege to attempt fix
	FXQUOR	Adjust quorum fix
	FXRESM	Resume process fix
	FXSUSP	Suspend process fix
	FXTIMO	Fix timeout
	FXUERR	Unknown error code for fix

A

- Access control lists (ACLs), 1–11
- Adjust queue limits fix, 5–16
- Adjust quorum fix, 5–5
- Adjust working set fix, 5–12
- AMDSSAM_CONFIG logical name, 6–18
- AMDSSAM_LOGICALS.COM file, 6–4
- AMDSSCONSOLE_ACCESS.DAT file, 6–20
- AMDSSDRIVER_ACCESS.DAT file, 6–18
- AMDSEVENT_LOG.LOG file, 4–4
- AMDSSGROUP_NAME logical name, 6–4
- AMDSSSYSTARTUP_VMS.COM file, 2–1
- APCs (asynchronous procedure calls), 3–6
- Application window
 - components, 2–4
 - how to display data, 2–4
 - menu bar, 2–4
 - panes in, 1–1
- AST fix, 5–16
- ASTLM (AST limit) quota, B–6
- Asynchronous procedure calls (APCs), 3–6
- Asynchronous system traps (ASTs)
 - adjusting queue limits fix, 5–16
- Automatic data collection, 1–6
- Availability messages
 - sent to Data Analyzer, 6–21
- AVAIL command, 2–2

B

- Background data collection, 1–5
- Blocks
 - in use, remaining, and so on, 6–10
- Buffered I/O
 - byte limit (BYTLM), 3–13
 - limit, 3–13
 - rate, 3–12, 6–11, B–4
- Buffered I/O (BIO) fix, 5–15
- Byte limit remaining for process I/O, 3–13

C

- Change process priority fix, 5–2, 5–10
- Channels
 - LAN virtual circuit, 3–29
- Clusters
 - See OpenVMS Clusters
- Collecting data
 - background and foreground, 1–5
 - specifying types, 2–7
- Collection intervals, 1–4, 3–1
- Configuration, 1–2
- Congestion control
 - transmitting data, 3–34
- CPU modes
 - OpenVMS, 3–6, 3–7
 - Windows NT, 3–5
- CPU process states, A–1
- CPUs (central processing units)
 - improving performance by suspending, 5–9
 - modes, 3–4
 - summary information, 3–4
 - number active on a node, 3–4
 - process states, A–1
 - process summary, 3–8
 - setting process priorities, 5–11
 - usage, 3–4
 - wait state, 3–8
- Crash node fix, 5–4
- Customization
 - events, 1–9
- Customizing
 - access codes, 6–20
 - node selection, 6–1
- OpenVMS
 - data collection, 6–5
 - data filters, 6–7
 - events, 6–15
 - group membership, 6–4
 - security features, 6–16
 - specifying nodes to monitor, 6–2
 - security features, 6–16
 - security triplets, 6–19
- Windows NT
 - events, 6–15
 - group membership, 6–4

- Customizing
 - Windows NT (cont'd)
 - security features, 6-16
 - specifying nodes to monitor, 6-2

D

- Data Analyzer
 - description, 1-2
 - event log file, 4-4
 - nodes, 1-2
 - passwords, 1-10
 - security, 1-10
- Data collection
 - automatic, 1-6
 - background, 1-5
 - changing collection intervals, 6-6
 - customizing settings, 6-5
 - foreground, 1-6
 - frequency of, 1-7
 - intervals, 1-7
 - selecting items, 2-7
- Data Collector
 - nodes, 1-2
 - restarting, 6-4
 - security
 - private LAN transport, 1-11
 - read and write access, 1-11
 - security triplets, 1-10
- Data filters
 - changing values, 6-7
- Data packets
 - receipt, 3-33
 - transmission, 3-32
- DECams, 1-1
 - changes and enhancements
 - no installation of server, 2-1
- Deferred procedure calls (DPCs), 3-5
- Delete process fix, 5-6
- DIOLM (Direct I/O limit), 3-13
- Direct I/O fix, 5-14
- Direct I/O rate, 3-12
- Disks
 - OpenVMS
 - single disk summary, 3-16
 - status summary, 3-15
 - summaries, 3-15
 - volume summary, 3-17
 - Windows NT
 - logical summary, 3-18, 3-19
 - physical summary, 3-20
- Disk status
 - filtering data, 6-8
- Display data collection interval, 1-7

E

- Event data collection interval, 1-7
- Event log file, 1-1, 4-4
- Events
 - displaying information, 4-2, 4-5
 - escalation of, 6-16
 - occurrence value, 6-16
 - OpenVMS, B-1
 - posting, 1-8
 - severity, 4-2, 6-15
 - signaling performance problems, 1-4
 - testing for, 1-9
 - thresholds, 6-16
 - thresholds for signaling, 1-8
 - Windows NT, B-12
- EVENTS.LOG file on Windows NT systems, 4-4
- Events pane, 4-1
- Exit image fix, 5-7
- Extended LAN
 - communication between Data Collector and Data Analyzer, 1-2

F

- File protection
 - for security, 1-11
- Filtering data
 - methods, 6-1
- Filters
 - OpenVMS CPU, 6-7
 - OpenVMS disk status, 6-8
 - OpenVMS disk volume, 6-10
 - OpenVMS I/O, 6-11
 - OpenVMS lock contention, 6-12
 - OpenVMS memory, 6-13
 - OpenVMS page/swap file, 6-14
 - specifying types of data to collect, 6-7
- Fixes
 - adjust quorum, 5-5
 - adjust resource limits, 5-13
 - changing process priority, 5-10
 - changing working set size, 5-12
 - cluster hung, 5-2
 - CMKRNL privilege required, 5-3
 - crashing a node, 5-4
 - deleting a process, 5-6
 - description of, 5-1
 - exiting an image, 5-7
 - intruder, 5-2
 - list of available, 5-1
 - memory too low, 5-2
 - memory usage, 5-2, 5-3
 - node, 5-1, 5-4
 - process, 5-1
 - process looping, 5-2
 - purging a working set, 5-11

Fixes (cont'd)

- recommended, 5-3
- results, 5-3
- resuming a process, 5-9
- runaway process, 5-2
- suspending a process, 5-8
- system service calls associated with, 5-1, 5-2
- understanding, 5-1
- working set too high or too low, 5-2

Foreground data collection, 1-6

G

Graphical user interface

See GUI

Group pane

See also Groups, Panes

selecting a group, 2-5

Groups

See also Group pane, Panes

changing, 6-4

displaying nodes in, 2-5

selecting for monitoring, 2-5, 6-2

user-definable, 2-5

GUI (graphical user interface), 1-1

H

Hardware

security triplet address, 6-19

Hardware model, 3-4

Help

getting, 2-9

Hung cluster fix, 5-2

I

I/O (input/output)

adjusting limits fix, 5-14

current, threshold, and peak values, 3-12

default data collection, 6-6

page/swap files, 3-14

page fault rate, 3-11

process quotas, 3-39

rates per process, 3-13

summaries, 3-11

I/O byte fix, 5-21

IEEE 802.3 Extended Packet format protocol, 1-2

Increasing resource limits fix, 5-13

Interrupts per second, 3-6

Intruder fix, 5-2

IPID (internal PID), 3-7

J

Java GUI, 1-1

Job quotas in use

single process, 3-42

JOB_CONTROL process, 5-12

L

LAN

virtual circuits, 3-22

detailed data, 3-28

Local area network (LAN), 1-1

Lock contention

OpenVMS, 3-20, 6-6

Lock queue limit fix, 5-18

Locks

contention for, 3-20

Logical disks

Windows NT, 3-19

Low node memory fix, 5-2

M

Memory, 3-4

count, 3-11

data, 6-13

default data collection, 6-6

fixes, 5-2

low memory fix, 5-2

OpenVMS summary, 3-10

summaries, 3-8

total for a node, 3-4

Windows NT summary, 3-8

Memory usage

displaying, 3-10

fixes, 5-3

Menu bar

in Application window, 2-4

Modes

See CPU modes

Monitoring nodes

selecting nodes, 6-2

Mutexes

held, 3-38

number in node, 3-3

MWAIT state

resource wait table, A-2

N

Network address

security triplet, 6-19

Network Interconnect for the System

Communications Architecture

See NISCA

Network protocol, 1-1
Networks
 interconnect data, 3-28
NISCA
 data, 3-28
Node data
 OpenVMS
 table, 3-2
 summary, 3-4
 Windows NT, 3-2
Node pane, 2-6
 See also Nodes
 data in OpenVMS display, 3-2
 data in Windows NT display, 3-2
 Windows NT, 3-2
Nodes
 See also Node pane
 adjusting quorum, 5-5
 crash fix, 5-4
 displaying data, 3-1
 displaying in Node pane, 2-6
 fixes, 5-1, 5-4
 memory usage, 3-10
 OpenVMS, 3-1
 selecting data to collect, 2-7
 summary information, 3-4
NoEvent data collection interval, 1-7
Nonpaged pool
 displaying size, 3-10

O

Occurrences
 criterion for posting an event, 1-8
 event, 1-8
 low values, 6-16
Open file limit fix, 5-17
OpenVMS Clusters
 hung, 5-2
 interconnect summary, 3-22
 members data, 3-24
 running Availability Manager in, 1-2
 summary data, 3-24
OS (operating system)
 version, 3-4

P

Packets discarded
 LAN virtual circuit, 3-31
Page/swap files, 3-14
Paged pool
 displaying size, 3-10
Page faults, 3-12
 adjust working set fix, 5-12
 purging working sets fix, 5-12
 rate, 3-11

Page files
 data collection, 6-6
Paging write I/O rate, 3-12
Panels
 Application window, 1-1, 2-3
 changing sizes, 2-4
 Event, 2-4
 Group, 2-3
 Node, 2-3
Passwords, 1-10
 changing, 6-16, 6-19, 6-21
 default, 1-10
 security triplets, 1-10
Performance
 identifying problems, 1-4
Physical disks
 Windows NT, 3-20
PID (process identifier), 3-7, 3-11, 3-13
PIO (paging I/O)
 rate, 3-13
Pipe quota
 for transmitted data, 3-35
PRCLM process limit, B-7
Private LAN transport security, 1-11
Processes
 adjusting limits, 5-2
 data filtering, 6-7
 execution rates, 3-38
 fixes, 5-1
 information about, 3-36
 job quotas, 3-42
 looping process fix, 5-2
 privileges, 1-11
 Process priority fix, 5-10
 quotas
 displaying, 3-39
 Resume process fix, 5-9
 Runaway process fix, 5-2
 single process data, 3-36
 Suspend process fix, 5-8
 wait states, 3-40
 working sets, 3-37
Process limit fixes, 5-13
Process looping fix, 5-2
Process priority
 changing, 5-3
Process quotas
 adjusting, 5-2
 displaying data, 3-39
Purge working set fix, 5-11, 5-12

Q

Quotas
 adjusting, 5-2
 job, 3-42
 process I/O, 3-39
 working set, B-5

R

- Receiving information, 1-3
- Requesting information, 1-3
- Resource availability
 - displaying, 4-1
 - fixes, 5-1
- Restarting the Data Collector, 6-4
- Resume process fix, 5-9
- Runaway process
 - Change process priority fix, 5-2
 - Suspend process fix, 5-2

S

- SCA (System Communications Architecture)
 - virtual circuit data, 3-25
- SCS (System Communications Services)
 - connections data, 3-22, 3-27
- Security, 1-10
 - access control lists (ACLs), 1-11
 - changing groups, 6-4
 - changing passwords, 6-16, 6-19, 6-21
 - data transfer, 1-11
 - file protection, 1-11
 - private LAN transport, 1-11
 - process privileges, 1-11
 - triplets, 1-10, 6-18
- Security triplets
 - access verification code, 6-19
 - changing, 6-19
 - description, 6-18
 - files, 1-11
 - format, 6-19
 - hardware address, 6-19
 - network address, 6-18
 - operation, 6-21
 - password, 6-18
 - verifying, 6-21
 - wildcard address, 6-20
- Semaphores, 3-3
- Single disks, 3-16
- Single process
 - data, 3-36
 - OpenVMS
 - execution rates, 3-38
 - job quotas, 3-42
 - process I/O quotas, 3-39
 - process information, 3-36
 - wait states, 3-40
 - working set, 3-37
- SMP (symmetric multiprocessing), 3-6
- Sorting data, 2-8
- Starting the Availability Manager, 2-2
- Status bar
 - in Application window, 2-4

Subprocess

- adjusting creation limit, 5-19
- Subprocess creation limit fix, 5-19
- Suspend process fix, 5-2, 5-8
- Swap files
 - data collection, 6-6
- SWAPPER process
 - displaying, 3-40
- SYSSSTARTUP directory, 2-1
- SYSAPs (system applications)
 - See System applications (SYSAPs)
- System applications (SYSAPs), 3-26
- System cache
 - displaying size in use, 3-9
- System service calls
 - associated with fixes, 5-2

T

- Threads, 3-2
- Thresholds
 - criteria for posting an event, 1-8
 - events, 6-16
- Time queue entry fix, 5-19
- Title bar
 - in Application window, 2-4

U

- Uptime, 3-4

V

- Virtual circuits
 - LAN
 - channel selection data, 3-29
 - closures data, 3-30
 - congestion control data, 3-34
 - detailed data, 3-28, 3-29
 - for individual nodes, 3-22
 - packets discarded data, 3-31
 - receive data, 3-33
 - transmit data packets, 3-32
 - SCS connections data, 3-27
 - System Communications Architecture (SCA)
 - data, 3-25
- Virtual memory
 - displaying size, 3-9
- Volume
 - default data collection, 6-6

W

- Wait states
 - CPU, 3-8
 - process, 3-40

- Wildcard address
 - security triplet, 6-20
- Working set extent, 3-11
- Working sets
 - data, 3-37
 - pages, 3-37
 - purging, 5-11
 - size, 3-11
 - size fix, 5-12
 - too high or too low, 5-2
- Working set size fix, 5-12