System Event Analyzer

User Guide

System Event Analyzer (SEA) is a rules-based hardware fault management diagnostic tool that provides error event analysis and translation. The multi-event correlation analysis feature of SEA provides the capability to analyze events stored in the system's binary event log file and events from other sources.

The *System Event Analyzer User Guide* provides information about the features of SEA and explains how to operate the software.

Rev. 12/12/03–A

Operating System: Microsoft® Windows® 2000 and XP
HP Tru64 UNIX® versions 4.0F, 4.0G, 5.0A or higher
HP-UX version 11.0 or higher
Red Hat Linux versions 7.3 and 8.0
OpenVMS Alpha versions 7.2-2 or higher

Software Version: SEA 4.3.1

Hewlett-Packard Company
Technical Publications
305 Rockrimmon Boulevard South
Colorado Springs, Colorado 80919 • U.S.A.

December 2003

Change Summary

The following table summarizes changes to this document:

Revision	Description
12/12/03–A	Initial 4.3.1 copy

Hewlett-Packard Company
Technical Publications
305 Rockrimmon Boulevard South
Colorado Springs, Colorado 80919 • U.S.A.

Contents

Contents

List of Figures

List of Figures

List of Tables

List of Tables

Preface

System Event Analyzer (<System Name>) is a rules-based hardware fault management diagnostic tool that provides error event analysis and translation. The multi-event correlation analysis feature of <System Name> allows you to analyze events from a variety of sources, including those stored in the system's binary event log file.

Overview

The *System Event Analyzer User Guide* describes the features of <System Name> and explains how to use the application.

The book is divided into eight chapters, each covering a specific topic related to SEA.

User Guide Contents

Chapter	Contents
Chapter 1	Introduces <System Name> and WEBES.
Chapter 2	Describes how to interact with the WEBES Director.
Chapter 3	Provides information about the Command Line Interface.
Chapter 4	Provides detailed information about the Web User Interface.
Chapter 5	Describes the translation of system events and the analysis of error logs.
Chapter 6	Explains the analysis rules used by <System Name>.
Chapter 7	Discusses the <System Name> configuration settings.
Chapter 8	Describes how to configure automatic notification.
Appendix A	Shows sample output files.
Appendix B	Contains information about optimizing the performance of <System Name>.

User Guide Contents (continued)

Chapter	Contents
Appendix C	Details how to configure and use your browser with the web interface.
Appendix D	Describes <System Name> messages.
Appendix E	Explains how to use the CLI old common syntax, DECevent UNIX syntax, and DECevent VMS syntax.

Intended Audience

The *System Event Analyzer User Guide* is intended for system managers and service personnel who use the <System Name> software.

Documentation Conventions

The following conventions are used in this manual:

Bold	is used to indicate user entries and GUI tasks. These instructions include information that should be entered exactly as it appears in the document.
Italic	is used to indicate variables and other information that will vary depending on your system and user profile.
`Fixed-width font`	is used for system output, code examples, CLI commands, file names, directories, process names, and URLs.
CAPITALIZATION	is used for keyboard commands.

Note

Notes provide additional important information related to the text.

Further Information

<System Name> is a member of the Web-Based Enterprise Services (WEBES) suite of products. For more information on the other WEBES applications, visit the support web site at the following URL:

http://h18000.www1.hp.com/support/svctools/

For information about the supported products and limitations of the current release, refer to the *System Event Analyzer Release Notes*.

Information about the supported operating systems is contained in the *WEBES Install Guide* along with detailed installation instructions for each operating system.

1

Introduction

This chapter describes SEA, the supported platforms, the post-installation setup procedures, the WEBES and SEA processes, the procedures used to start and stop the Director, the locations of WEBES Director log files, and the nomenclature differences.

1.1 Description of SEA

SEA is a fault analysis utility designed to provide analysis for single error/fault events, as well as multiple events and complex analysis. In addition to the traditional binary error log, SEA provides system analysis capabilities that use other error/fault data sources.

SEA provides background automatic analysis by monitoring the active binary error log and processing events as they occur. The events in the binary error log file are checked against the analysis rules, and if one or more of the events in the binary error log file meets the conditions specified in the rules, the analysis engine collects the error data and creates a problem report containing a description of the problem and any corrective actions required. Once the problem report is created, it is distributed in accordance with the customer's notification preferences.

SEA supplies a web-based user interface that connects to the Director and can perform a variety of tasks from a remotely connected web browser. In addition, a set of command-line tools enable diagnosis of binary event logs without connecting to the Director.

1.2 Hewlett-Packard Service Tools

Hewlett-Packard has implemented a common Application Programming Interface (API) for many of its service tools called Web-Based Enterprise Services (WEBES). The tools included in the current WEBES release are:

- System Event Analyzer (SEA)
- Computer Crash Analysis Tool (CCAT)

SEA utilizes the common components of WEBES and adds it own functionality. The other WEBES service tools can be installed along with SEA and utilize the same common components.

1.3 Supported Products

The following list includes the products SEA supports.

This list also is available in the *System Event Analyzer Release Notes*. In the event of any discrepancy between this list and the *System Event Analyzer Release Notes*, the release notes take precedence.

Note

Do not confuse the supported products with the systems where WEBES can be installed. Installation requirements are given in the *WEBES Installation Guide*.

- Platforms: Analysis and Bit-To-Text Translation
    - HP AlphaServer DS10/DS10L/DS15/DS20/DS20E/DS25 (Tru64 UNIX and OpenVMS)
    - HP AlphaServer ES40/ES45 (Tru64 UNIX and OpenVMS)
    - HP AlphaServer GS80/GS160/GS320 (Tru64 UNIX and OpenVMS)
    - HP AlphaServer TS80/ES47/ES80/GS1280/GS1280 M64 (Tru64 UNIX and OpenVMS)
    - HP AlphaServer TS20/TS40 (Tru64 UNIX and OpenVMS)
    - HP AlphaServer TS202C (Tru64 UNIX and OpenVMS)
    - Memory Channel II (Tru64 UNIX and OpenVMS)
- Platforms: Bit-To-Text Translation only
    - HP AlphaServer DS20L (Tru64 UNIX and OpenVMS)
- I/O Devices: Analysis and Bit-To-Text Translation
    - Disk Storage based on SCSI specification (Tru64 UNIX, OpenVMS, and Windows)
    - EZ4X/EZ6X (Tru64 UNIX and OpenVMS)
    - EZ5X/EZ7X (Tru64 UNIX and OpenVMS)
    - HSG60/HSG80/HSZXX (Tru64 UNIX and OpenVMS)
    - HSG60/HSG80 (Windows)
    - KGPSA-CA/KGPSA-BC/KGPSA-BY/KGPSA-CB/KGPSA-CX/KGPSA-CY FCA2384/FCA2354/FCA2404/FCA2406 (Tru64 UNIX)
    - Smart Array 5304 Controller (Tru64 UNIX and OpenVMS)
    - Modular SAN Array 1000 (Tru64 UNIX and OpenVMS)
    - EMA16000, MA8000/EMA12000, MA6000, RA8000/ESA12000
- I/O Devices: Bit-To-Text Translation only
    - RA3000
    - KZPSC/KZPAC/KZPBA/KZPCM/KZPSA/KZPCC/KSPEA
    - KGPSA-CA/KGPSA-BC/KGPSA-BY/KGPSA-CB/KGPSA-CX/KGPSA-CY FCA2384/FCA2354/FCA2404/FCA2406 (OpenVMS)
    - CCMAB-AA
    - CIPCA-BA
- Storage Systems: Analysis and Bit-To-Text Translation
    - EVA 3000/5000 on VCS V2.0x and V3.0x for HSV100 and HSV110 controllers
    - MSA1000
- Storage System Components: Analysis and Bit-To-Text Translation
    - StorageWorks SAN 1 Gbps Switches:
        - DSGGA-AA 8 port, StorageWorks Fibre Channel switch
        - DSGGA-AB 16 port, StorageWorks Fibre Channel switch
        - DSGGB-AA 8 port, StorageWorks SAN switch 8
        - DSGGB-AB 16 port, StorageWorks SAN switch 16

- DSGGC-AA 8 port, SAN Switch 8-EL
- DSGGC-AB 16 port, SAN Switch 16-EL
- DSGGS SAN Switch Integrated /32 and /64 ports

- StorageWorks SAN 2 Gbps Switches:
  - DS-DSGGD-AA 16 port, SAN Switch 2/16
  - DS-DSGGD-AB 32 port, SAN Switch 2/32
  - DS-DSGGD-AC 8 port, SAN Switch 2/8-EL
  - DS-DSSGD-AD 16 port, SAN Switch 2/16-EL
  - DS-DSGGD-BB 32 port, SAN Switch 2/32
  - DS-DSGGD-DB 32 port, SAN Switch 2/32
  - DS-DSGGE-*xx* 64 port, Core Switch 2/64

1.4 Supported Operating Systems

Hewlett-Packard Sustaining Engineering maintains a schedule of support for the Tru64 UNIX, HP-UX, and OpenVMS operating systems. Hewlett-Packard does not commit to supporting WEBES when installed on an operating system version that has exceeded its end-of-support date. See the following URL:

```
http://www.compaq.com/services/software/ss_pvs_se_amap.html
```

SEA supports Windows 2000 and XP. Windows NT Alpha is not supported.

1.5 Security and Required Permissions

In order to enhance security, only privileged users can access the WEBES directory tree and run SEA commands. The requirements for each operating system are given here.

Tru64 UNIX

The following actions are restricted to privileged users:

- Running any WEBES or SEA commands (`desta` or `wsea` commands from the command prompt).
- Viewing the WEBES directory tree on a system.

Only the “root” user can perform these actions. The `/usr/opt/hp/svctools` directory is owned by root, and has rwx (read, write, and execute) permissions for root (owner), and no permissions for any other user (group or world).

HP-UX

The following actions are restricted to privileged users:

- Running any WEBES or SEA commands (`desta` or `wsea` commands from the command prompt).
- Viewing the WEBES directory tree on a system.

Only the "root" user can perform these actions. The `/opt/hp/svctools` directory is owned by root, and has rwx (read, write, and execute) permissions for root (owner), and no permissions for any other user (group or world).

Linux

The following actions are restricted to privileged users:

- Running any WEBES or SEA commands (`desta` or `wsea` commands from the command prompt).
- Viewing the WEBES directory tree on a system.

Only the "root" user can perform these actions. The `/usr/opt/hp/svctools` directory is owned by root, and has rwx (read, write, and execute) permissions for root (owner), and no permissions for any other user (group or world).

OpenVMS

Commands—To execute any SEA commands (DESTA or WSEA commands), the user needs all of the following OpenVMS privileges. Note that these are a subset of the privileges required to install, upgrade, or uninstall WEBES as described in the *WEBES Installation Guide*:

ALTPRI	DIAGNOSE	SYSPRV
BUGCHK	IMPERSONATE	TMPMBX
CMKRNL	NETMBX	

Files—File access is restricted in the WEBES installed directory tree pointed to by the SVCTOOLS_HOME logical (SYS$COMMON:[HP] by default). To view these files, you must be a member of the System group, your user ID must have all privileges, or you must issue the SET PROCESS /PRIV=ALL command.

All directories and files in the SVCTOOLS_HOME tree are owned by the System user, and have System, Owner, and Group permissions of RWED (Read, Write, Execute, and Delete). There are no permissions for World.

Windows

The following actions are restricted to privileged users:

- Running any of the WEBES programs from the Start menu (**Start>Programs>Hewlett-Packard Service Tools**).
- Running any WEBES or SEA command (`desta` or `wsea` commands from the command prompt).

- Accessing any files within the WEBES installed directory tree (`C:\Program Files\hp\svctools` by default).

To perform restricted actions, your user ID must be either:

- A member of the Administrators group on that machine.
- A member of a group that is a member of the Administrators group on that machine. For example, if your user ID is a Domain Admin, and you have added Domain Admins to the Administrators group on the local machine, you will have the necessary permissions. (The *WEBES Install Guide* describes how to add a group to the Administrators group.)

1.6 WEBES and SEA Processes

Each WEBES-based service tool adds functionality to the Director, a process (or set of processes) that executes continuously. SEA provides the Director with the capability to capture and interpret hardware events. Event analysis can be performed automatically or at the request of an outside process.

SEA includes a web browser interface that enables you to interact with the Director. Although only one Director can run on a machine at any time, many web browser connections can be active simultaneously, all connected to the single Director.

Note

WEBES (Web-Based Enterprise Services) and DESTA (Distributed Enterprise Service Tools Architecture) refer to the same common components.

1.6.1 Director

The Director manages its own host (either a machine, or a node in a cluster) and communicates to Directors on other hosts through TCP/IP sockets.

The Director captures, translates, and analyzes events as well as routing messages for the SEA system. The Director is idle except for the following circumstances:

- Events are received for processing
- Messages arrive from other WEBES processes on the same machine
- Messages arrive from a Director on another machine
- Another WEBES tool within the Director, performs any task

The Director is automatically started along with the machine and should not require any additional action.

See Sections 1.7 and 1.8 for more information regarding starting and stopping the Director.

1.6.2 Command Line Interface (CLI)

Many SEA operations can be performed from the command prompt by issuing commands beginning with `desta` and `wsea`. For example:

```
wsea analyze input myBinary.errlog
```

Each CLI command starts a process. Some CLI processes connect to the Director on the same machine to perform tasks. However, other CLI processes perform all tasks themselves without connecting to a Director.

CLI commands typically support many options, enabling you to specify input and output files, and filtering criteria.

Chapters 2 and 3 of this guide describes how to use the command line interface.

1.6.3 Web Interface

Using a web browser, such as Netscape Communicator or Internet Explorer, you can connect:

- directly to the URL of the Director on the same machine as the browser
- directly to the URL of the Director on a remote machine
- indirectly to a remote Director through a direct connection to the Director on the local or a remote machine.

The web interface can monitor multiple nodes by communicating with the Directors on other machines. You can establish a direct connection to the Director on any machine reachable by its TCP/IP socket port, and, through that connection, view the SEA processes on other nodes (via Director-to-Director communication). You do not need to have WEBES installed or running on the web browser's machine to connect directly to the Director on a remote machine.

Chapter 4 of this guide describes how to use the web interface and Appendix C discusses the browser requirements.

1.6.4 Director/Interface Interaction

Figure 1–1 shows an example of two machines running SEA processes.

Figure 1–1 Interaction Between Two Systems Running SEA

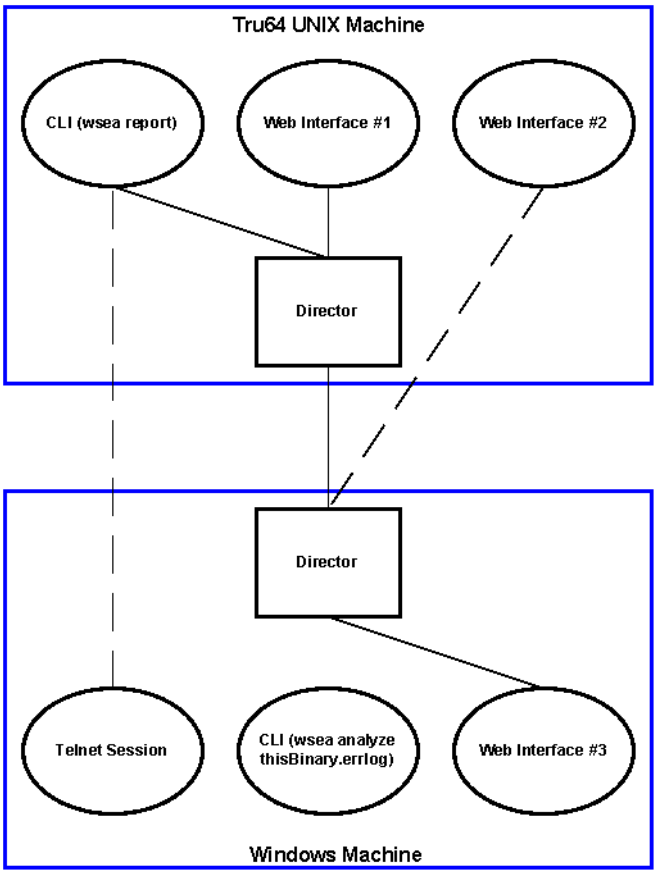

In this example, a UNIX machine and a Windows machine, each running a single Director, communicate with each other over a network.

Web interface #1 is a web browser running on the UNIX machine, directly connected to the local Director on the same machine (`http://localhost:7902`). It can also communicate with the Director on the Windows machine through the UNIX Director. This enables you to view the output produced by either machine (such as analysis results) using the same web interface.

Web interface #2 is also a web browser running on the UNIX machine, but it has directly connected to the Director on the Windows machine (`http://thatPC.mydomain.com:7902`). Using this web interface you can, if desired, connect back to the UNIX Director as well, but the UNIX Director need not be running at all.

A telnet session initiated from the Windows machine has logged on to the UNIX machine, and the user has issued the SEA CLI command `wsea report` (to view the results of automatic analysis). The CLI process connects to the UNIX machine's Director, which returns the current report data to the CLI process. The report text is then displayed to the user. Note that it is not necessary to have the Director running on the Windows machine for this type of remote connection.

The CLI command `wsea analyze input thisBinary.errlog` (which performs manual analysis on the binary log file named `thisBinary.errlog`) is issued from a command prompt on the Windows machine. In the example, the file `thisBinary.errlog` is assumed to be a log file from a UNIX machine that was transferred to the Windows machine using FTP. The Director is not required to perform manual analysis and therefore, the local Director is not required.

Finally, web interface #3 is a web browser running on the Windows machine. This interface is directly connected to the local Director on the same machine (`http://localhost:7902`), the same way that web interface #1 connects to its local UNIX Director.

1.7 Starting the Director

The Director is automatically started during system startup. Under normal operation, you should not need to manually start the Director. However, if circumstances require it, you can manually start the Director by following the instructions for your operating system.

Note

After a “desta stop” or “net stop desta_service” command completes, the operating system sometimes requires a few more seconds to stop all WEBES-related processes and release their resources (such as sockets). On rare occasions, restarting the Director too soon after stopping it can can result in errors in the Director log file, and the Director also may fail to restart.

To avoid this issue, wait 10 more seconds before restarting the Director, once the “desta stop” or “net stop desta_service” command completes.

Tru64 UNIX

At a shell prompt, enter:

```
/usr/sbin/desta start
```

On TruClusters, you can run the `/usr/sbin/webes_install_update` program and choose the Start WEBES Director option to start the Director on either all the nodes in the cluster or a selected group of nodes that you choose.

HP-UX

At a shell prompt, enter:

```
/usr/sbin/desta start
```

Linux

At a shell prompt, enter:

```
/usr/sbin/desta start
```

OpenVMS

At the OpenVMS command line prompt, enter:

```
desta start
```

On OpenVMS clusters, you can uses the SYSMAN utility to issue the command `do desta start` on either all the nodes in the cluster or a specific group of nodes that you choose.

Windows

To start the WEBES Director, start the `DESTA_Service` using one of the following methods:

- Select **Programs > Hewlett-Packard Service Tools > Web-Based Enterprise Service > Start Director** from the Start menu.
- In a Command Prompt window, enter:
  ```
  net start DESTA_Service
  ```
- Start `DESTA_Service` using the Services utility in the Control Panel.

Using the `desta start` command on a Windows system is unsupported and not recommended. While the command may start the Director, it will also generate error messages. Starting the director this way is not recommended because:

- Closing the command prompt window used to issue the command or logging out of the Windows session, will forcibly but incompletely kill the Director, leaving running processes behind (see the *WEBES Release Notes* if this situation occurs). In addition, open files may not be saved correctly, resulting in data corruption.
- Text log output from the Director process will only be displayed on the screen and will eventually scroll past the buffer.

On Windows, the desta start/stop functionality is only intended to be used as a tool for investigating WEBES operational problems. If the Director is started with `desta start`, it must be stopped with `desta stop`.

1.8 Stopping the Director

Under normal operation, you should not need to stop the Director. However, if circumstances require you to stop the Director, follow the instructions for your operating system.

Note

After a “desta stop” or “net stop desta_service” command completes, the operating system sometimes requires a few more seconds to stop all WEBES-related processes and release their resources (such as sockets). On rare occasions, restarting the Director too soon after stopping it can can result in errors in the Director log file, and the Director also may fail to restart.

To avoid this issue, wait 10 more seconds before restarting the Director, once the “desta stop” or “net stop desta_service” command completes.

Tru64 UNIX

At a shell prompt, enter:

```
/usr/sbin/desta stop
```

On TruClusters, you can run the `/usr/sbin/webes_install_update` program and choose the Stop WEBES Director option to stop the Director on either all the nodes in the cluster or a selected group of nodes that you choose.

HP-UX

At a shell prompt, enter:

```
/usr/sbin/desta stop
```

Linux

At a shell prompt, enter:

```
/usr/sbin/desta stop
```

OpenVMS

At a prompt, enter:

```
desta stop
```

On OpenVMS clusters, you can uses the SYSMAN utility to issue the command `do desta stop` on either all the nodes in the cluster or a specific group of nodes that you choose.

Windows

Stop the Director by stopping the `DESTA_Service` Windows service. You can use any of the following methods:

- Select **Programs** > **Hewlett-Packard Service Tools** > **Web-Based Enterprise Service** > **Stop Director** from the Start menu.

A Stop Director icon appears in the Task Bar, then disappears when the Director's shutdown has completed.

- At a command prompt, enter:
  `net stop DESTA_Service`
- Use the Services utility in the Control Panel.

Using the `desta stop` command on Windows systems is unsupported. Stopping the director this way is not recommended because:

- The Director may not stop completely, leaving running processes behind.
- Error messages may be displayed in either the logs for the Director process or in the `desta stop` output.
- The Director may take longer to stop than it normally would using one of the recommended methods, and may continue to run for a time even after the `desta stop` process has finished.

On Windows, the desta start/stop functionality is only intended to be used as a tool for investigating WEBES operational problems. If the Director is started with `desta start`, it must be stopped with `desta stop`.

1.9 Starting the WCCProxy

The WCCProxy is a set of processes, and a Service on Windows (see Section 1.11 Monitoring WEBES Processes for details). These processes are used by the WEBES Director to perform various system-level tasks. The Director will not perform correctly without the WCCProxy.

The WCCProxy is automatically started during system startup. Under normal operation, you should not need to manually start the WCCProxy. However, if circumstances require it, you can manually start the WCCProxy by following the instructions for your operating system.

Tru64 UNIX

At a shell prompt, enter:

```
# /usr/sbin/wccproxy start
```

HP-UX

At a shell prompt, enter:

```
# /usr/sbin/wccproxy start
```

Linux

At a shell prompt, enter:

```
# /usr/sbin/wccproxy start
```

OpenVMS

At the OpenVMS command line prompt, enter:

```
$ wccproxy start
```

On OpenVMS clusters, you can uses the SYSMAN utility to issue the command `do wccproxy start` on either all the nodes in the cluster or a specific group of nodes that you choose.

Windows

To start the WCCProxy, start the WCCProxy service using one of the following methods:

- In a Command Prompt window, enter either equivalent command:

  ```
  C:\>net start wccproxy
  ```

  ```
  C:\>wccproxy start
  ```

- Start the WCCProxy service using the Services utility in the Control Panel.

1.10 Stopping the WCCProxy

Under normal operation, you should not need to stop the WCCProxy. However, if circumstances require you to stop the WCCProxy, follow the instructions for your operating system.

Tru64 UNIX

At a shell prompt, enter:

```
# /usr/sbin/wccproxy stop
```

HP-UX

At a shell prompt, enter:

```
# /usr/sbin/wccproxy stop
```

Linux

At a shell prompt, enter:

```
# /usr/sbin/wccproxy stop
```

OpenVMS

At a prompt, enter:

```
$ wccproxy stop
```

On OpenVMS clusters, you can uses the SYSMAN utility to issue the command `do wccproxy stop` on either all the nodes in the cluster or a specific group of nodes that you choose.

Windows

Stop the WCCProxy by stopping the WCCProxy Windows service. You can use any of the following methods:

- In a Command Prompt window, enter either equivalent command:

  ```
  C:\>net stop wccproxy
  ```

  ```
  C:\>wccproxy stop
  ```

- Stop the WCCProxy service using the Services utility in the Control Panel.

If any of the processes associated with WCCProxy (see Section 1.11 Monitoring WEBES Processes) do not stop using any of the methods listed above, you can kill them with the following command:

```
wccproxy kill
```

1.11 Monitoring WEBES Processes

Monitor WEBES processes as follows.

Director Status

You can monitor the WEBES Director process with the following command:

```
desta status
```

This command generates a return value and state as shown in Table 1–1.

Table 1–1 Director Status Codes

Value	State
1	The Director is not running.
3	The Director is running.
5	The Director is starting up.

Table 1–1 Director Status Codes

Value	State
7	The Director is shutting down.
9	The Director's status file indicates it is running, but the process ID was not found, so the Director process in fact is *not* running and has terminated abnormally.
99	The Director's status could not be determined.

If the status is undetermined, or you want more detailed information about subprocesses, you may want to use the monitoring procedures specific to your operating system.

WCCProxy Status

You can monitor the WCCProxy (a set of processes used by the Director) using the following command:

```
wccproxy status
```

The command generates a value and state that indicate the current status of the WCCProxy processes as shown in Table 1–2.

Table 1–2 WCCProxy Status Codes

Value [a]	State
0 or 4	The WCCProxy status could not be determined.
1	The WCCProxy is running.
2	The WCCProxy is not running.
3	The WCCProxy service is not installed.

a. In WEBES 4.3.1, the values apply only to Windows. The values will be changed in a future release to values more like the desta status values shown in Table 1–1.

On all UNIX platforms, the return code is always zero. This will be corrected in a future release.

On OpenVMS, the return code is always %X10010001 (hexadecimal). This will be corrected in a future release.

If the status is undetermined, or you want more detailed information about subprocesses, you may want to use the monitoring procedures specific to your operating system.

Note

With UNIX and Windows, some WEBES processes are listed using the name `java`. Other applications running on your system may also be listed using the name `java`. Therefore, when using the procedures described here, it is important to make sure that the processes in question are WEBES processes.

Tru64 UNIX

Some WEBES processes are Java™-based, using the Java Runtime Environment bundled with WEBES. These WEBES processes run under the java executable. Other processes are C++ based and run under their own image name. The processes currently running can be displayed with the command:

```
ps ugxww | grep /usr/opt/hp/svctools | grep -v grep
```

This searches for the path containing all WEBES executable image names, including the "java" image in the Java Runtime Environment embedded in WEBES.

Example output is shown here:

```
root 146989 0.0 0.1 2.95M 552K pts/0 S N 16:31:43 0:00.05 /usr/opt/hp/svctools/common/wccproxy/
share/WCCProxy
root 147095 0.0 3.4 22.7M 17M pts/0 S N 16:31:49 0:05.24 /usr/opt/hp/svctools/common/jre/bin/../
bin/alpha/native_threads/java -classic -noverify -DSvctools.Home=/usr/opt/hp/svctools -DSwcc.Home=/
var/adm -Xmx99m com.compaq.svctools.desta.core.DESTAController
root 147114 0.0 2.9 20.8M 15M pts/0 S N 16:31:53 0:03.01 /usr/opt/hp/svctools/common/jre/bin/../
bin/alpha/native_threads/java -classic -noverify -DSvctools.Home=/usr/opt/hp/svctools -DSwcc.Home=/
var/adm -Xmx136m com.compaq.svctools.desta.util.DESTAProcessWrapper
root 147145 0.0 0.1 3.16M 640K pts/0 S N 16:31:57 0:00.78 /usr/opt/hp/svctools/common/wccproxy/
share/CAAgents -s 19 -p 3273 -1 -g
root 147148 0.0 0.1 2.92M 480K pts/0 S N 16:32:06 0:00.38 /usr/opt/hp/svctools/common/wccproxy/
share/WCCAgents -s 20 -p 2877 -1 -g
jones 147172 0.0 0.1 2.30M 344K pts/1 S + 16:33:01 0:00.03 wsea analyze /usr/opt/hp/svctools/common/
ca/examples/gs320_uce_ivp.errlog
jones 147180 0.1 0.1 2.49M 520K pts/1 S + 16:33:01 0:00.08 /usr/opt/hp/svctools/common/bin/desta
exec com.compaq.svctools.ca.cli.CACLIInterpreter analyze /usr/opt/hp/svctools/common/ca/examples/
gs320_uce_ivp.errlog
jones 147207 81.4 5.3 32.7M 27M pts/1 R + 16:33:01 0:05.60 /usr/opt/hp/svctools/common/jre/bin/../
bin/alpha/native_threads/java -classic -noverify -DSvctools.Home=/usr/opt/hp/svctools -DSwcc.Home=/
var/adm com.compaq.svctools.ca.cli.CACLIInterpreter analyze /usr/opt/hp/svctools/common/ca/
examples/gs320_uce_ivp.errlog
```

In this example:

- Process 146989 is the WCCProxy process, a C++ based launcher for WEBES processes, that communicates to the main Director process.
- Process 147095 is the main Java-based Director process, started with the DESTAController Java class.
- Process 147114 is a subprocess of the Director (subprocesses start with the DESTAProcessWrapper Java class), which only runs when needed.
- Process 147145 is a CAAgents process, a SEA C++-based process launched by WCCProxy to read the native binary event log and send events to the main Director process. There may be more than one CAAgents process running at a time, or none.

- Process 147148 is a WCCAgents process, a C++-based process launched by WCCProxy to send notifications. There may be more than one WCCAgents process running at a time, or none.
- Process 147172, its child process 147180, and its child process 147207 are all running a CLI command issued by the “jones” user, analyzing an example event log.

OpenVMS

Use the following command to show the processes running on an OpenVMS machine:

```
show system
```

Example output is shown here:

```
OpenVMS V7.2-2 on node THIS 15-OCT-2002 15:03:52.59 Uptime 39 05:37:42
Pid      Process Name    State  Pri I/O     CPU           Page flts  Pages
...
0000F68D WCCProxy        LEF    6   353     0 00:00:00.07  504        201
0000F68E DESTA Director  HIB    5   198456  0 00:01:10.09  154670     12301 M
0000F68F SMITH_2         HIB    6   23027   0 00:02:31.40  25089      6285 MS
0000F691 CA.A.19.54240   HIB    6   341     0 00:00:00.16  422        286
0000F695 CA.A.20.54249   LEF    6   248     0 00:00:00.11  465        239
0000F698 WCC.A.1200.8989 LEF    6   201     0 00:00:00.10  382        220
0000F89C JONES_1         HIB    6   291     0 00:00:00.05  316        133 S
0000F69E JONES_2         COM    4   2656    0 00:00:07.57  73623      7357 MS
0000F342 RCM           HIB    7   0      0 00:00:00.00  23        30
```

In the above example, the DESTA Director parent process is shown. That process has also spawned a subprocess named SMITH_2, which only runs when needed, so named because the user SMITH started the Director, but the relation is not apparent from the output. Other WEBES processes, such as SEA Command Line Interface commands, appear named after the user that started them, such as JONES_1 and its subprocess JONES_2 in this example, although it is not apparent that they are WEBES processes. The WCCProxy process is a C++ based launcher for WEBES processes that communicates to the main Director process. The CA.A.*nn.nnnn* and WCC.A.*nn.nnnn* processes are C++ based processes launched by WCCProxy to send notifications, read the native binary event log, and send events to the main Director process. There may be more than one CA.A.*nn.nnnn* or WCC.A.*nn.nnnn* process running at a time, or none.

Windows

Use the Windows Task Manager to monitor processes in Windows.

1. Press CTRL+ALT+DEL to start the Task Manager.
2. Click the Task Manager button.

   The Task Manager window appears.

3. Click the Processes tab to view the running processes.

WEBES Director processes consist of the following image names:

- DESTAService.ex (on Windows 2000) or DESTAService.exe (on XP)
- java.exe

- WCCProxy.exe
- CAAgents.exe
- WCCAgents.exe

The main parent Java-based Director process is the DESTAService process, which runs as a Windows service. It spawns a subprocess when needed, which runs under the process name java.exe. The WCCProxy process is a launcher for C++ based WEBES processes that communicates to the main Director process. The CAAgents.exe and WCCAgents.exe processes are C++ based processes launched by WCCProxy to send notifications, read the native binary event log, and send events to the main Director process. There may be more than one CAAgent or WCCAgent process running at any time, or none.

All CLI commands run under the process name `java.exe`. However, not all `java.exe` processes are guaranteed to be WEBES processes. Java-based applications other than WEBES may also appear as `java.exe`. You may be able to distinguish the Director set of processes from other WEBES and non-WEBES Java processes by looking at the Base Priority of the `java.exe` processes. The Director processes always run at Low priority, while all other WEBES processes run at Normal priority. However, other Java processes, not associated with WEBES, may also run at Low priority.

If the Base Priority column is not shown in the Task Manager list:

1. Choose Select Columns from the View pull-down menu.
2. Click Base Priority
3. Click OK

1.12 Log Files

SEA processes warnings and informational messages from the Director in log files.

Note

These warning and informational message files are different from binary event log files. See Section 1.15 for more information about the different log files.

If SEA appears to execute incorrectly, or does not respond as expected, check the Director log files for messages that may help you restart or recover. The files can be copied to new file names so that they are not overwritten later, and can be sent to your service provider for review.

All WEBES processes log their messages either to files or to the terminal window. For common messages you may encounter, refer to the *System Event Analyzer Release Notes* or Appendix D.

1.12.1 Location

The format of the log file messages is the same for all platforms, however, the file locations are different for each operating system.

Tru64 UNIX

The Director and web interface log standard output and error messages to:

```
/usr/opt/hp/svctools/specific/webes/logs/desta_dir.log
```

The Director appends to this log file each time it is started.

WEBES and the WEBES installer write additional log files containing information that might be useful to WEBES product support personnel when diagnosing a problem with WEBES or any of its component tools. These log files are stored in the following directories:

```
/usr/opt/hp/svctools/specific/ca/logs
/usr/opt/hp/svctools/specific/wccproxy/logs
/usr/opt/hp/svctools/specific/webes/logs
```

HP-UX

The Director and web interface log standard output and error messages to:

```
/opt/hp/svctools/specific/webes/logs/desta_dir.log
```

The Director appends to this log file each time it is started.

WEBES and the WEBES installer write additional log files containing information that might be useful to WEBES product support personnel when diagnosing a problem with WEBES or any of its component tools. These log files are stored in the following directories:

```
/opt/hp/svctools/specific/ca/logs
/opt/hp/svctools/specific/wccproxy/logs
/opt/hp/svctools/specific/webes/logs
```

Linux

The Director and web interface log standard output and error messages to:

```
/usr/opt/hp/svctools/specific/webes/logs/desta_dir.log
```

The Director appends to this log file each time it is started.

WEBES and the WEBES installer write additional log files containing information that might be useful to WEBES product support personnel when diagnosing a problem with WEBES or any of its component tools. These log files are stored in the following directories:

```
/usr/opt/hp/svctools/specific/ca/logs
/usr/opt/hp/svctools/specific/wccproxy/logs
/usr/opt/hp/svctools/specific/webes/logs
```

OpenVMS

The Director and web interface log standard output and error messages to:

```
SVCTOOLS_HOME:[SPECIFIC.WEBES.LOGS]DESTA_DIR.LOG
```

The Director creates a new log file each time it is started. The previous log file is saved as:

```
DESTA_DIR.LOG;n
```

Where *n* is the previous version number of the VMS filename.

WEBES and the WEBES installer write additional log files containing information that might be useful to WEBES product support personnel when diagnosing a problem with WEBES or any of its component tools. These log files are stored in the following directories:

```
SVCTOOLS_HOME:[SPECIFIC.CA.LOGS]
SVCTOOLS_HOME:[SPECIFIC.WCCPROXY.LOGS]
SVCTOOLS_HOME:[SPECIFIC.WEBES.LOGS]
```

Windows

The locations given here assume that SEA was installed in the default directory; if this is not the case, the location path will match the chosen installation directory.

The Director (and web interface) logs its standard output messages to:

```
C:\Program Files\hp\svctools\specific\webes\logs\director_out.txt
```

The Director's standard error messages are logged to:

```
C:\Program Files\hp\svctools\specific\webes\logs\director_err.txt
```

The Director creates new log files each time it is started. The previous log files are renamed to `director_err.txt.bck` and `director_out.txt.bck`, overwriting any previous versions of those files.

WEBES and the WEBES installer write additional log files containing information that might be useful to WEBES product support personnel when diagnosing a problem with WEBES or any of its component tools. These log files are stored in the following directories:

```
C:\Program Files\hp\svctools\specific\ca\logs
C:\Program Files\hp\svctools\specific\desta\logs
C:\Program Files\hp\svctools\specific\wccproxy\logs
C:\Program Files\hp\svctools\specific\webes\logs
```

1.12.2 Logging Level

The messages logged by WEBES processes are stored in the Director log files described in Section 1.12.1. The minimum severity level, or logging level, indicates the lowest priority message that will be written to the files. Only messages that meet or exceed the minimum severity level are written to the Director log files.

1.13 Service Obligations

A service obligation specifies your service provider, service agreement information, and the duration of your agreement. During the WEBES installation process, you will be prompted to enter the service obligation information. This information is included with the results of analysis.

Although SEA continues to function without a valid service obligation, local notification and reporting are disabled. In addition, the web interface will no longer operate after your service obligation has expired.

Refer to Sections 2.6 and 4.11 for information on viewing service obligations.

1.14 Environment Setup

For more information on automatic notification and the SEA configuration settings refer to the following sections:

- To set up Simple Mail Transfer Protocol (SMTP) E-mail notification of problem reports, refer to Section 8.2.
- To set up notification of problem reports using Automated Call Handling Service (ACHS), also known as System-Initiated Call Logging (SICL), which uses DSNlink services installed on the system, refer to Section 8.4.1.
- To set up notification of problem reports to a Customer Service Gateway (CSG), formerly known as Qualified Service Access Point (QSAP), for use with Proactive Remote Service (PRS), refer to Section 8.4.2.
- If you wish to change how the SEA components operate, you can change the configuration using the web interface. Refer to Chapter 7 for more information about configuration.

You can modify the SEA environment at any time.

1.15 Nomenclature

The term configuration is used in two different contexts in SEA:

- Hardware Configuration – identifying the Field Replaceable Unit (FRU) or hardware components currently installed in a machine.
- System Configuration – identifying the current software settings of the SEA system and each of the services it contains. Most of the settings can be changed using the SEA interfaces.

Log file is also used in two different contexts:

- A log file containing text errors or information written by a SEA or WEBES process, such as `/usr/opt/hp/svctools/specific/webes/logs/desta_dir.log` on Tru64 UNIX
- An error or event log file containing binary events written by the system event logger, such as `/var/adm/binary.errlog`, written by the binlogd daemon on Tru64 UNIX and translated and analyzed by SEA

2

WEBES Director

This chapter describes the WEBES Director and the commands associated with it.

2.1 Overview

You interact with the WEBES Director by sending Director commands from your system's command prompt. These commands impact the WEBES common components and are not limited to SEA. Using Director commands you can configure port settings, activate automatic notification, and view your service obligation.

2.1.1 Security and the CLI

In order to run any of the commands described in this chapter, you must be a privileged user. Refer to Section 1.5 for information on the permissions required for your operating system.

2.1.2 Clusters and the CLI

Even if SEA is installed on a cluster, commands only impact the local node. If you want to modify an entire cluster you will need to perform the desired operation on each node.

2.2 DESTA Commands

Director commands are formed using the following convention:

`desta` *`command_verb`*

Where *command_verb* indicates the action you want to perform.

Table 2–1 describes the command verbs used with `desta`.

Table 2–1 Command Verbs—desta

Verb	Description
`dri`	Controls DESTA registry entries including the amount of memory used by the Director process and subprocesses. Refer to Section 7.7.4 for more information.
`msg`	Changes the SEA port configuration settings. See Section 2.3 for more details on port settings.
`priority`	Changes the priority of the Director process. This command is only supported on UNIX and VMS. See Section 2.5 for more details.
`qsap`	Toggles on or off the SEA Qualified Service Access Point (QSAP) feature, which automatically log calls with Hewlett-Packard Services. See Section 2.4.2 for syntax information and Section 8.4.2 for more details on QSAP.
`servob`	Displays your SEA service obligation. See Section 2.6 for more details.

Table 2–1 Command Verbs—desta (continued)

Verb	Description
`sicl`	Toggles on or off the SEA System Initiated Call Logging (SICL) feature, which automatically log calls with Hewlett-Packard Services if DSNLink is installed on the system. See Section 2.4.1 for syntax information and Section 8.4.1 for more details on SICL.
`start`	Starts the Director if it has been stopped. See Section 1.7 for more details on starting the Director.
`status`	Shows the current status of the Director. See Section 1.11 for more details on the Director's status.
`stop`	Manually stops the Director. See Section 1.8 for more details on stopping the Director.

2.3 Configuration

You can configure the socket ports used by WEBES with the following command:

`desta msg -chgport` *nnn*

Changes the socket ports (see Section 7.4.2 for more information).

Note

There are more configuration settings that can be changed using the web interface. Refer to Chapter 7 for more information on configuration.

2.4 Notification

SEA can automatically send problem reports to Hewlett-Packard for faster problem resolution. With notifications, the results of SEA analysis are automatically sent to your service provider as they occur.

2.4.1 SICL Notifications

Note

The SICL command syntax was formerly `wsea sicl [on|off]`.

Begin using the new `desta sicl` command, and update any scripts that employ the `wsea sicl` command. The `wsea sicl` command will be completely phased out in a future release.

System-Initiated Call Logging (SICL), also known as Automatic Call Handling Services (ACHS), can send automatic notifications when DSNlink is installed on the same system where WEBES is installed. The following command turns SICL notification on or off:

```
desta sicl [on|off]
```

See Section 8.4.1 for more information.

2.4.2 PRS Notifications

Proactive Remote Service (PRS) can send automatic notifications without installing DSNlink or other software on the same system where WEBES is installed. PRS first sends problem reports to another system, the Customer Service Gateway (CSG), which then forwards the reports to the service provider. The CSG was formerly known as a Qualified Service Access Point (QSAP). The following command turns PRS notification on or off:

```
desta qsap [on|off]
```

See Section 8.4.2 for more information.

2.5 Priority

By default the Director process runs at low priority. On UNIX and VMS systems you can change the priority while the Director is running by entering the following command:

```
desta priority [compete | low]
```

Where `compete` assigns the Director a normal priority and `low` assigns the Director a low priority.

On OpenVMS systems, this command issues the `SET PROCESS /PRIORITY` command. The operating system may change the priority of any process at any time, and may not change the priority when the `SET PROCESS /PRIORITY` command is issued. Therefore, the `desta`

command may not change the priority of the DESTA Director process. It functions more like a suggestion to the operating system rather than a command.

2.6 Service Obligations

Your service obligation describes the details of your service agreement. You can view an existing service obligation from the command line. See Section 1.13 for more information about service obligations.

To view the service obligation for a machine, enter the following command:

```
desta servob show
```

This displays all the information associated with your service obligation. An example of this information is shown here:

```
WEBES Service Obligation Status
-------------------------------

   Service Obligation:            Valid
   Service Obligation Number:     50036123
   System Serial Number:          50036123
   Service Provider Company Name: Hewlett-Packard
```

2.7 Getting Help

You can access help from the CLI using the command for your operating system:

- Tru64 UNIX – `man desta` and `desta help`
- HP-UX – `man desta` and `desta help`
- Linux – `man desta` and `desta help`
- OpenVMS – `help desta` and `desta help`
- Windows – `desta help`

3

Command Line Interface

This chapter describes the Command Line Interface (CLI) for SEA including its usage and functionality.

3.1 CLI Overview

The command line interface (CLI) provides a text-based interface for SEA.

Security and the CLI

In order to run any of the commands described in this chapter, you must be a privileged user. Refer to Section 1.5 for information on the permissions required for your operating system.

Clusters and the CLI

Even if SEA is installed on a cluster, commands only impact the local node. Thus, if you want to modify an entire cluster you will need to perform the desired operation on each node.

Standalone CLI

The Director is not required to run all the CLI commands. The following CLI functions can be performed without the Director:

- Manual Analysis
- Translation
- Summary Report
- Create New Binary Log File
- List Registered Rule Sets
- Register/Unregister Rule Sets
- Change or View Syntax
- Reset the Automatic Analysis Database
- View the Status Information

Since these operations do not use the Director, messages that would otherwise be written to the Director's log files are included in the output for the command. The messages shown remain subject to the logging level. Refer to Section 1.12 for more information on log messages.

Conventions

Table 3–1 describes the conventions used to show CLI commands in this manual.

Table 3–1 Syntax Conventions

Convention	Meaning
Bold	Command text. Bold is used for information that must be typed as it appears. For example, command verbs are shown in bold.
Italic	Variables. Italics are used for information that varies depending on your requirements. For example, *inputfile* indicates that you should enter the name of the file you want to process.
[ ]	Optional Entries. Information shown in square brackets is not required. You may or may not include these optional modifiers. In most cases the optional entries pertain to input files, output files and filtering commands.
|	Mutually Exclusive Entries. The bar separates mutually exclusive entries.

3.2 Command Syntax

You interact with the CLI by issuing commands from the command prompt. Some SEA operations can be performed using several different commands, or syntaxes. The supported syntaxes are:

- New Common Syntax
- Old Common Syntax
- DECevent Emulation (only available on UNIX and VMS systems)

You can enter commands using any of the supported command formats. If desired, you can switch between the different syntaxes.

Note

This chapter only describes the New Common Syntax since it supports all the SEA functionality and is the default syntax at installation. If you want to use commands from another syntax, refer to Appendix E.

If you are using a command syntax other than the default, you must include a syntax designator in the command. Table 3–2 shows the syntax designators.

Table 3–2 Syntax Designators

Syntax Name	Syntax Designator	Command Preface
New Common Syntax	n	`wsea` or `wsea n`[a]
Old Common Syntax	x	`wsea x`
DECevent Emulator (UNIX)	u	`wsea u`
DECevent Emulator (VMS)	v	`wsea v`

a. Because the New Common Syntax is the default syntax, the syntax designator "n" is not required.

3.2.1 Setting the Default Syntax

When SEA is installed the new common syntax is the default for the CLI, however, you can change the default syntax. Any commands entered using the current default syntax do not require a syntax designator. To specify a default syntax, use the following command:

`wsea syntax` *`syntax_designator`*

Where *syntax_designator* refers to the letter corresponding to the desired default syntax (see Table 3–2 for the designator associated with each syntax).

For example, to set the DECevent UNIX syntax as the default syntax, use the following command:

```
wsea syntax u
```

Once the syntax is set, you can enter commands in your chosen syntax without specifying the syntax designator.

Note

Changing the default syntax affects all the users on the system. Thus, if someone else changes the default syntax, it will affect your SEA session. If there are multiple users logged into a system, you may want to include a syntax designator with all commands that support multiple syntaxes.

3.2.2 Showing the Default Syntax

To show the current default syntax, use the following command:

`wsea syntax`

3.3 Command Verbs

Some SEA commands are supported by multiple syntaxes, some are only supported by the new common syntax, and some are syntax independent.

- Commands that are syntax specific are formed using the following convention:
  wsea *syntax_designator command_verb*
- Commands that are syntax independent do not use a syntax designator.
  wsea *command_verb*

Note

If you enter the command `wsea` without any command verb or parameters, SEA defaults to translation. In this case, the system event log is translated and the output is sent to the screen.

Table 3–3 provides an overview of the `wsea` command verbs available in the new common syntax.

Table 3–3 Command Verbs—wsea (New Common Syntax)

Command Verb[a]	Description
`ana` (`analyze`)	Analyzes one or more binary event logs. See Section 3.5.1 for more details.
`aut` (`autoanalysis`)	Turns automatic analysis on or off. See Section 3.5.2 for more details.
`tra` (`translate`)	Translates one or more binary event logs, but does not analyze the events. See Section 3.6 for more details.
`sum` (`summarize`)	Returns a summary of all the events contained in a binary event log. See Section 3.7 for more details.
`bin` (`binary`)	Applies a filter to an existing binary event log and creates a new binary event log containing the subset of events returned after filtering. The bin command verb can also be used to merge existing binary event logs. See Section 3.8 for more details.
`lis` (`listrk`)	Lists the registered analysis rule sets. See Section 3.10 for syntax information and Chapter 6 for more details on rule sets.
`reg` (`regknw`)	Registers one or more analysis rule sets for use during automatic and manual event analysis. See Section 3.10 for syntax information and Chapter 6 for more details on rule sets.
`unr` (`unregknw`)	Unregisters one or more analysis rule sets so they are no longer considered during automatic and manual event analysis. See Section 3.10 for syntax information and Chapter 6 for more details on rule sets.

Table 3–3 Command Verbs—wsea (New Common Syntax) (continued)

Command Verb[a]	Description
`tes` (`test`)	Simulates automatic analysis. See Section 3.5.2.3 for syntax information and Chapter 5 for more details on analysis.
`res` (`reset`)	Resets the automatic analysis database. See Section 3.5.2.4 for syntax information and Chapter 5 for more details on analysis.
`sta` (`status`)	Displays system information such as the software version, obligation information, and notification status. See Section 3.11 for more information.
`help`	Displays a text-based help file. The text-file describes the syntaxes supported by your operating system.

a. The new common syntax allows abbreviations. You only need to enter the minimum number of characters required to uniquely identify the command to run the command (generally, the first three letters of a command verb). The full command verb is shown in parenthesis.

Table 3–4 describes the commands that are syntax independent.

Table 3–4 Command Verbs—wsea (Syntax Independent)

Command Verb	Description
`report`	Displays the active problem reports generated from automatic analysis. See Section 3.5.2.1 for more details.
`log`	Toggles the logging to a file of automatically generated problem reports on or off. See Section 3.5.2.2 for more details.
`sicl`	Toggles on or off the SEA System Initiated Call Logging (SICL) feature, which automatically log calls with Hewlett-Packard Services if DSNLink is installed on the system. See Section 2.4.1 for syntax information and Section 8.4.1 for more details on SICL. This command is being phased out and replaced by the `desta sicl` command.

3.4 Command Parameters

Note

With the new common syntax, command parameters can be abbreviated. You only need to enter the minimum number of characters required to uniquely identify the parameter. For example, `input` can be abbreviated as `inp` and `outhtml` can be abbreviated as `outh`.

Parameters are used to specify binary log files for processing, designate output files, and create filters. In most cases, SEA allows you to specify parameters in any order (the new common syntax sum command is an exception, see Section 3.7 for details). For example, the following commands using the new common syntax are equivalent:

```
wsea tra inp myinput.zpd out myoutput.txt index=start:10 brief
wsea brief index=start:10 out myoutput.txt inp myinput.zpd tra
```

Notice that even the placement of the command verb (`tra` in this case) may be changed.

3.5 Analysis

If the Director is installed, automatic analysis is initiated when you start your machine. This means that SEA automatically analyzes the default event log file and generates reports as necessary.

Manual analysis is the user-initiated process of selecting a log file for immediate processing using either the CLI or the web interface (see Chapter 4 for more information about the web interface).

For more information on analysis and the default analysis settings, refer to Chapter 5.

3.5.1 Manual Analysis

The analyze command allows you to perform manual analysis as well as filtered manual analysis on a binary event log. Binary event logs can include the system event log, a log obtained from the same system, or a log obtained from another system.

To manually analyze binary event logs using the new common syntax, use the following command format:

wsea ana [**input** *inputfile*] [**out** | **outhtml** *outputfile*]

Note

When performing filtered manual analysis, sort filters may cause invalid analysis results.

For more information on manual analysis operations and output, refer to Chapter 5.For information on performing manual analysis with another syntax, refer to Appendix E.

Input Files

By default, manual analysis processes the system event log. If you want to process a different binary log file, you must use the input keyword and specify the input file. Refer to Section 3.9.1 for more information on input files

Output Files

By default, output from manual analysis is displayed on the screen. To save output to a file, use either the out or the outhtml keyword and provide a file name. Refer to Section 3.9.2 for more information on output files.

3.5.2 Automatic Analysis

By default, when the Director is started SEA initiates automatic analysis on the binary system event log. Using the CLI, you can view the reports generated by automatic analysis or save them to a file.

To disable or enable automatic analysis, use the following command format:

```
wsea aut [on | off]
```

Automatic analysis is enabled by default, but you may want to disable it if SEA is running on a platform such as HP-UX or Linux, where the native error log is not currently analyzed.

For more information on automatic analysis and the problem reports generated by analysis, refer to Chapter 5. For information on using the command line interface to simulate automatic analysis, refer to Sections 3.5.2.3 and 5.7.

3.5.2.1 Viewing Automatic Analysis Reports

To view the active problem reports generated by automatic analysis, use the `report` command. Reports can be viewed in the command prompt window or saved to a file.

The syntax for the report command is shown here:

```
wsea report [outtext | outhtml outputfile]
```

If you do not include any optional parameters, the reports are shown on the screen. See Section 3.9.2 for more information about working with output files.

3.5.2.2 Logging Automatic Analysis Reports

SEA can automatically log generated problem reports in the `prob.log` file located in the `\specific\ca\logs` directory.

To turn automatic logging on, use the following command:

```
wsea log prob on
```

To turn automatic logging off, use the following command:

```
wsea log prob off
```

If the `prob.log` file already exists, the new data from subsequent logging operations is appended to the existing file. If you delete the `prob.log` file, it is automatically recreated during the next logging operation. Log output is flushed and the file is closed after each entry.

3.5.2.3 Simulate Automatic Analysis

The command line can also simulate automatic analysis with the following command (this command is only available for the new common syntax).

```
wsea tes [nosystem]
```

Tests automatic analysis and the system's error logging facilities. See Section 5.7 for more information on simulating automatic analysis.

3.5.2.4 Reset Automatic Analysis Results

Note

Resetting the automatic analysis results may significantly impact the results of future analysis. Refer to Section 5.3.2 for a complete description of the reset command.

The command line can clear the automatic analysis database (this command is only available for the new common syntax).

```
wsea res
```

Removes any currently active callouts and any stored analysis data (for example, thresholding information). Neither the FRU configuration data nor the marker of the most recently analyzed event are removed. See Section 5.3.2 for more detailed instructions on resetting the automatic analysis results and information about the impact that resetting automatic analysis results has on future analysis results.

3.5.3 Analysis Output

Chapter 5 describes the problem reports generated by analysis. Refer to Appendix A for an example of a report generated by analysis.

3.6 Translation

You can translate, or decompose, the events in a binary event log into a readable format using the translate command. Translation operates in manual mode, meaning you must enter the command every time you want to perform translation.

For more information about translation and its default settings, refer to Chapter 5.

To manually translate binary event logs using the new common syntax, use the following command format:

```
wsea tra [input inputfile] [out | outhtml outputfile] [filterstatement]
[brief | full]
```

For information on performing translation with another syntax, refer to Appendix E.

Input Files

By default, translation processes the system event log. If you want to process a different binary log file, you must use the input keyword and specify the input file. Refer to Section 3.9.1 for more information on input files

Output Files

By default, output from translation is displayed on the screen. To save output to a file, use either the out or the outhtml keyword and provide a file name. Refer to Section 3.9.2 for more information on output files.

Filtering Log Files

You can specify the events from a binary event log file that you want to translate by defining a filter. For more information on filtering refer to Section 3.9.3.

Report Type

You can specify either brief or full output for translation. Refer to Section 5.2.3 for more information on the report types.

Translation Output

Chapter 5 describes the results of translation. Refer to Appendix A for an example of a translated event and to see the difference between full and brief output.

3.7 Summary of Events

You can use the CLI to view a summary of the events contained in a binary log file. To do so with the new common syntax, use the following command format:

```
wsea sum [index] [input inputfile] [out | outhtml outputfile]
[filterstatement]
```

Note

Section 5.8 describes circumstances that can cause unexpected summary output.

For information on generating a summary with another syntax, refer to Appendix E.

Indexed Output

By default, a tallied list of all the events in the binary event log files is generated. However, you can generate an indexed list of all the events using the index modifier.

Be aware if you are using the new common syntax `sum` command and you want to generate indexed output, the `index` parameter must immediately follow the `sum` command verb. Otherwise, SEA will assume you are using the `index` filter keyword.

Input Files

By default, the summary command processes the system event log. If you want to process a different binary log file, you must use the input keyword and specify the input file. Refer to Section 3.9.1 for more information on input files

Output Files

By default, output from the summary command is displayed on the screen. To save output to a file, use either the out or the outhtml keyword and provide a file name. Refer to Section 3.9.2 for more information on output files.

Filtering Log Files

You can specify the events from a binary event log file that you want to view a summary report for by defining a filter. For more information on filtering refer to Section 3.9.3.

Example Output

The results of the summary command are displayed in the command prompt window.

An example of the standard, tallied output is shown here:

```
== /usr/opt/hp/svctools/common/ca/examples/gs320-unix-dir-620.errlog ==
  Qty   Type Description
----- ------ -------------------------------------------------------
    2    120 Correctable Error Throttling Notification Event Detected
    2    301 Tru64 UNIX Shutdown ASCII Message
    2    620 Correctable System Event
    2    300 Tru64 UNIX Start-up ASCII Message
    1    310 Tru64 UNIX Time Stamp Message
    2    199 Tru64 UNIX CAM Event
```

```
    3    110 Configuration Event
Total Entry Count: 14
First Entry Date: Mar 21, 2000 8:12:25 AM GMT-05:00
Last Entry Date: Mar 21, 2000 9:15:44 AM GMT-05:00
```

An example of the indexed output is shown here:

```
== /usr/opt/hp/svctools/common/ca/examples/gs320-unix-dir-620.errlog ==
Index Type Description                                 Date/Time
----- ---- ------------------------------------------ ------------------
    1  110 Configuration Event                        03/21/00 08:12:25 AM
    2  310 Tru64 UNIX Time Stamp Message              03/21/00 08:22:25 AM
    3  301 Tru64 UNIX Shutdown ASCII Message          03/21/00 08:31:21 AM
    4  110 Configuration Event                        03/21/00 09:07:15 AM
    5  300 Tru64 UNIX Start-up ASCII Message          03/21/00 09:07:16 AM
    6  120 Correctable Error Throttling Notification Event Detected 03/21/00
09:07:32 AM
    7  199 Tru64 UNIX CAM Event                       03/21/00 09:07:42 AM
    8  301 Tru64 UNIX Shutdown ASCII Message          03/21/00 09:08:41 AM
    9  110 Configuration Event                        03/21/00 09:11:16 AM
   10  300 Tru64 UNIX Start-up ASCII Message          03/21/00 09:11:17 AM
   11  120 Correctable Error Throttling Notification Event Detected 03/21/00
09:11:33 AM
   12  199 Tru64 UNIX CAM Event                       03/21/00 09:11:43 AM
   13  620 Correctable System Event                   03/21/00 09:15:41 AM
   14  620 Correctable System Event                   03/21/00 09:15:44 AM
```

3.8 Creating New Binary Event Log Files

You can filter the contents of existing binary event logs and create a new binary event log file containing a subset of the events from the originals. When you create a new binary log file, SEA checks the events in the original binary event log file against the filter statement. All the events that meet the criteria specified by the filter statement are added to the new binary event log file. The new binary event log file can then be used for analysis, translation, or any other SEA process. The syntax for creating new binary event log files using the new common syntax is shown here:

wsea bin [**input** *inputfile(s)*] **out** *outputfile* [*filterstatement*] [**skipconfig**]

For information on creating a new binary event log file with another syntax, refer to Appendix E.

Input Files

By default, the system event log is used as the input file. If you want to process a different binary log file or files, you must specify the input file location and name.

See Section 3.9.1 for more information on working with input files.

Note

You can specify multiple input files and merge them into a single binary log file (in this case, filtering occurs for each input file before events are written to the new file). Be aware that SEA does not remove duplicate events.

Output Files

You must specify a file name and location where the new binary output file will be saved. The output file parameter is mandatory when you are creating a new binary event log file.

Filtering Log Files

You can specify the events from a binary event log file that you want to include in the new log file by defining a filter. If you do not define a filter, the new log file will contain all the events in the existing log file. For more information on filtering refer to Section 3.9.3.

Skipping Configuration Entries

If you are using the new common syntax, you can keep configuration entries from being automatically inserted by adding the `skipconfig` modifier to your command. This modifier prevents configuration entries from the original log files that are needed for analysis from being inserted into the new log file if they would normally be filtered out.

3.9 Modifying Commands

By default, the analysis, translation, summary and new binary log file commands all process the system event log. The output from analysis, translation and summary commands is displayed on the screen. You can change these defaults in order to process other binary log files and save the processing results to a file. When you are performing translation, generating a summary, performing analysis, or creating a new binary log file you can further restrict the events that are processed by filtering the binary log file used for input. The following sections describe how to use these features.

3.9.1 Input Files

Many of the commands used in manual mode enable you to specify an input binary event log file. To specify an input file using the new common syntax, append the following modifier to the command:

`input` *`filename`*

Where *filename* indicates the path and name of the input file.

When you are specifying an input file, the following guidelines apply:

- Specifying an input file is optional. If you do not specify either a directory or a file, SEA processes the binary system event log. An example of a command without any input file criteria is shown here:

```
wsea ana
```

- You can use the relative directory structure to specify input files. Thus, if you were in the `C:\Program Files\hp\svctools\common\ca` directory and you wanted to analyze the `ds20.errlog` binary event log located in the `C:\Program Files\hp\svctools\common\ca\examples` directory, you could enter the following command:

```
wsea ana input examples\ds20.errlog
```

- If you specify a directory but no file name, SEA processes all the files with a `.errlog`, `.sys`, `.zpd`, or `.evt` extension located in the provided directory. An example of a command that only indicates a directory is shown here:

```
wsea analyze input examples
```

- Multiple filenames can be specified by separating them with a comma and space, as shown in the following example:

```
wsea ana input examples\ds20.errlog, examples\hscir1.zpd
```

- You can use wildcards to specify multiple files. In the example shown here, all the files located in the `examples` directory with a name that starts with `ds` and an `.errlog` extension are analyzed:

```
wsea ana input examples\ds*.errlog
```

3.9.2 Output Files

With many commands, you can save the results of processing to a file rather than viewing the output on the screen.

Note

These output file guidelines do not apply when you are creating a new binary event log. Refer to Section 3.8 for more details.

To send the output of an operation to a file, use one of the following modifiers:

out *filename*
outhtml *filename*

The out modifier creates a text output file and the outhtml modifier creates a HTML output file. The *filename* indicates the path and name where you want to save the output.

The following examples show commands that specify output files:

```
wsea ana out results.txt
wsea ana outhtml results.html
```

3.9.3 Filtering

The `tra`, `sum`, `bin`, and `ana` commands enable you to filter a binary event log file and only process a subset of the events. The following general rules apply when you use filters:

- You can include multiple filter statements by separating them with comma and a space.
- You can abbreviate the filter parameters. You only need to enter the minimum number of letters required to uniquely identify a parameter. For example, `index` could be abbreviated as `ind`.
- On Windows systems, any argument that includes a comma must be enclosed in quotation marks. This includes arguments that contain a date.

Table 3–5 describes the filtering statements available with the new common syntax.

Table 3–5 Filtering Statements (New Common Syntax)

Filter Statement	Description

`begin="`*`date`*`"`
`since="`*`date`*`"`
`end="`*`date`*`"`
	Filters based on the time the event occurred. No events that occurred before the given start time or after the given end time are processed. The date can be entered in any format supported by Java (for example, *dd-mmm-yyyy,hh:mm:ss*). You do not need to include the time (*hh:mm:ss*) with the date. Be aware of the following guidelines:
	• The begin and since statements are equivalent.
	• You can use the keywords YESTERDAY and TODAY.
	• With the begin and since keywords, you can enter a negative integer value to process based on a relative date. For example, entering -3 processes events from the last three days.

`include=`*`keyword`* `exclude=`*`keyword`*	Filters based on the numeric entry type. You must enter a keyword rather than the actual entry type. Refer to Table 3–6 for information on supported keywords.
`node=`*`name`*	Filters based on the node responsible for generating the event. The *name* argument is case sensitive.
`index=`*`nn`* `index="start:`*`nn`*`, end:`*`nn`*`"`	Filters based on the event's position in the event log. The first event in the file is event index 1.
`reverse`	Processes the events in reverse order according to the event index number.

Table 3–6 Event Type Keywords (New Common Syntax)

Keyword	Description
`mchk`	All machine check events.
`cam`	All SCSI entries logged by the CAM logger (199).
`configurations`	Configuration entries (110).
`control_entries`	System startup entries or new error log creation entries (32, 35, 300).
`environmental_entries`	Power entries (mchk-env).
`swxcr`	Entries logged by SWXCR (198).
`machine_checks` `mchks`	Events with machine checking information (mchk).
`operating_system=`*`value`* `os=`*`value`*	Events with a specific operating system type. The *value* parameter indicates the numeric code for the desired operating system.
`panic`	Crash re-start, system panic, or user panic entries (37, 302).
`software_informationals` `swi`	Events with lastfail, system startup, or system configuration information (volume mounts, volume dismounts, new error logs, timestamp entries) (32, 35, 37, 38, 39, 64, 65, 250, 300, 301, 310).
`osf_entry`	Events logged on a Tru64 UNIX operating system.
`mchk_sys`	All system machine check events.
`mchk_cpu`	All cpu machine check events.
`mchk_env`	All environmental machine check events.

Examples – New Common Syntax

The following examples show sample commands that use filtering.

Processes events from the system described by *ComputerName*:

```
wsea tra node=ComputerName
wsea sum node=ComputerName
wsea bin input inputfile.zpd out outputfile.bin node=ComputerName
```

Processes events that occurred before 8:33:57 PM on January 31, 2000:

```
wsea tra end="31-Jan-2000,20:33:57"
wsea sum end="31-Jan-2000,20:33:57"
wsea bin input inputfile.zpd out outputfile.bin end="31-Jan-
2000,20:33:57"
```

Processes all CPU machine check and system machine check events. The translation command presents the output in reverse chronological order:

```
wsea tra include="mchk_cpu, mchk_sys reverse"
wsea sum include="mchk_cpu, mchk_sys"
wsea bin input inputfile.zpd out outputfile.bin include="mchk_cpu,
mchk_sys"
```

Processes all the events after the fifteenth event in the log file:

```
wsea tra index=start:15
wsea sum index=start:15
wsea bin input inputfile.zpd out outputfile.bin index=start:15
```

3.10 Knowledge Rule Sets

Rule sets are used in conjunction with analysis. The events in a binary log file are compared with rule sets. Depending on the results of this comparison problem reports are generated. The following new common syntax commands are used to work with rule sets.

wsea lis

Lists the registered rule sets used by analysis (see Section 6.3.1 for more information).

wsea reg

Registers the rule sets used by analysis (see Section 6.3 for more information).

wsea unr

Unregisters the rule sets used by analysis (see Section 6.3 for more information).

Refer to Appendix E for information on the old common syntax commands used to work with rule sets.

3.11 Show Status Information

The new common syntax provides a command used to view information about SEA.

wsea sta

Shows the version, service obligation, and notification status.

An example of the output is shown here:

```
SEA for Tru64 UNIX V4.3.1 (Build 417)
Service Tools Home: /usr/opt/hp/svctools
Service Obligation Start Date: Fri Oct 04 00:00:00 MDT 2002
Service Obligation End Date: Sat Oct 04 00:00:00 MDT 2003
SICL/DSNlink notification: disabled.
CSG/QSAP notification: enabled.
```

3.12 Getting Help

You can access help from the CLI using the command for your operating system:

- Tru64 UNIX – `man wsea` and `wsea help`
- HP-UX – `man wsea` and `wsea help`
- Linux – `man wsea` and `wsea help`
- OpenVMS – `help wsea` and `wsea help`
- Windows – `wsea help`

Help is also available through the User Guide and the *System Event Analyzer Release Notes*. Documents are installed by the kit in PDF (Adobe Acrobat), text, and HTML formats, in the following directory:

- Tru64 UNIX – `/usr/opt/hp/svctools/common/ca/docs`
- HP-UX – `/opt/hp/svctools/common/ca/docs`
- Linux – `/usr/opt/hp/svctools/common/ca/docs`
- OpenVMS – `SVCTOOLS_HOME:[COMMON.CA.DOCS]`
- Windows – `C:\Program Files\hp\svctools\common\ca\docs`

The text versions do not include the graphics and formatting available with the other formats, and should only be used if the other formats cannot be easily viewed.

The HTML version of the *System Event Analyzer User Guide* is also available through the web interface, see Section 4.9 for details.

4

Web Interface

This chapter describes how to access and use the SEA web interface.

4.1 About the Web Interface

The web interface provides browser-based access to SEA. You can use the web interface to connect to the Director on your local machine or on remote machines and analyze and translate their binary event log files.

4.1.1 About Translation

Event information in the system event log is stored in binary format. Translation is the process of converting this binary data into readable text. The web interface does not automatically perform translation; each event that you want to translate must be manually selected.

- See Section 4.5 for more information on how the web interface presents translation information.
- Refer to Chapter 5 for more information on translation, interpreting translated events, and default translation settings.

4.1.2 About Analysis

Information from a binary event log file can be used to detect hardware failures on the system. The process of reading binary event log files, interpreting events, and creating problem reports with proposed resolutions is called analysis.

As the system writes events to the binary event log file, SEA processes each event according to the registered rule sets. The rule sets contain the information necessary to interpret events. Then, when an event matches the conditions described in the rule sets, SEA creates a problem report containing information about the event and proposed resolutions.

The web interface can perform both automatic and manual analysis.

- See Section 4.5 for more information on how the web interface presents analysis information.
- Refer to Chapter 5 for more information on analysis and its results.

4.1.2.1 Automatic Analysis

When the Director is started, SEA initiates automatic analysis. In automatic mode, SEA continuously monitors the binary system event log and processes events as they arrive. Problem reports are generated as necessary.

For more information about automatic analysis operations and output, refer to Chapter 5.

4.1.2.2 Manual Analysis

Manual analysis also compares the events from log files to the registered rule sets and generates problem reports. However, unlike automatic analysis, you must manually select each binary event log file you want to process.

For more information about manual analysis operations and output, refer to Chapter 5.

4.1.3 Notification

The results of automatic analysis can be sent to remote systems using SMTP, ACHS/SICL using DSNlink, or CSG/QSAP.

Refer to Chapter 8 for more information on notification.

4.1.4 Create New Binary Log File

You can filter the contents of existing binary event logs and create a new binary event log file containing a subset of the events from the originals. When you create a new binary log file, SEA checks the events in the original binary event log file (input file) against the filter statement. All the events that meet the criteria specified by the filter statement are added to the new binary event log file (output file). The new binary event log file can then be used for analysis, translation, or any other SEA process.

For more information on using the web interface to create a new binary event log file, refer to Section 4.6.

4.2 Starting the Web Interface

It is not necessary to have the Director running on your machine in order to use SEA. In fact, WEBES need not be installed on the browser's machine at all. However, WEBES must be installed and the Director must be running on the target machine in order to connect to its SEA system. Therefore, before using the web interface, you must ensure the Director is started on the target machine.

For additional information about supported browsers and configuring your browser for SEA, refer to Appendix C.

Accessing the Web Interface

1. Start the Director on the machine(s) you want to connect to (if they have not been started already). Refer to Section 1.7 for details.
2. Start your web browser.
3. Enter the URL of the target machine to connect to it.

- If you are connecting to a remote host, enter:
  `http://`*`hostname`*`.`*`domain`*`.com:7902`
- If you are connecting to the local machine, enter:
  `http://localhost:7902`

In some network configurations, the name `localhost` may not be recognized. Enter the machine's hostname or its IP address (such as `http://14.77.189.23:7902`) instead.

If you are using Internet Explorer, be sure to include the `http://`.

4. Enter the profile name you want to use in the Logon window (Figure 4–1) and click the Logon button or press Enter. See Section 7.5 for more information on profiles.

Figure 4–1 Logon Window

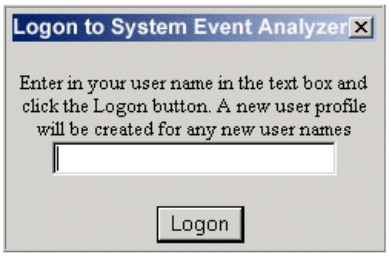

Although you must login to SEA, the logout process is automatic. Refer to Section 4.10 for a description of the automatic logout process.

4.3 Using The Web Interface

After you log on, the browser displays the web interface main screen (Figure 4–2).

Figure 4–2 Main Screen

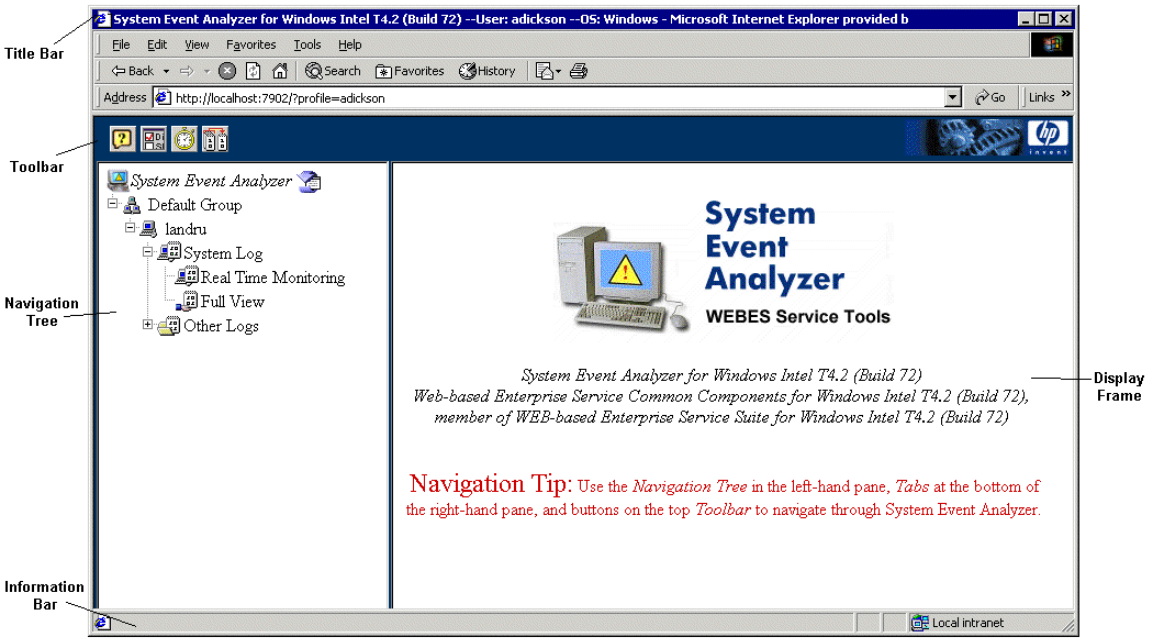

Note that the value of the URL field includes the *hostname* for the machine you logged into, as well as your *username*, indicating the current profile.

```
http://hostname:7902/?profile=username
```

Tip

If you need to change profiles while using SEA, you can edit your browser's URL field by replacing the current profile username with a different one.

The components of the web interface display are described in Table 4–1.

Table 4–1 Web Interface Components

Component	Description
Title Bar	Shows the software version, active profile, and operating system.
Toolbar	By default, provides access to the on-line help, system configuration, processing statistics, and new binary error log creation. The toolbar is dynamically updated, and additional features are available with some SEA screens. See Section 4.3.1 for more information.

Table 4–1 Web Interface Components (continued)

Component	Description
Navigation Tree	Lists the available groups, nodes, categories, and log files.
Display Frame	Displays interactive screens and system information. When SEA loads, the display frame shows product information.
Information Bar	Displays messages from the browser and usage tips. See Section 4.9.1 for more information on the web interface's usage tips.

Note

SEA allows you to work in multiple browser windows. If you are using the same profile, the navigation trees in all the windows will automatically synchronize.

4.3.1 Toolbar

Figure 4–3 shows the default web interface toolbar.

Figure 4–3 Toolbar

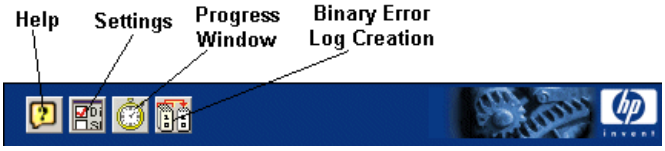

The toolbar buttons update dynamically depending on what you are doing. Table 4–2 describes the toolbar commands that are always available:

Table 4–2 Toolbar – Default Buttons

Component	Description
Help Button	Opens a new browser window containing the on-line user guide. See Section 4.9 for more information on getting help.
Settings Button	Opens the settings screen. See Section 4.8 for more information on changing the settings.
Progress Window Button	Opens a new browser window that reports the processing status of log files. See Section 4.5.2 for more information on processing status.
New Binary Log Button	Opens the New Binary Log screen in the display frame. See Section 4.6 for more information on creating a new binary log file.

The following buttons may also appear in the toolbar, depending on the feature being used:

Table 4–3 Toolbar – Dynamic Buttons

Component	Description
Clear	Available when viewing automatic analysis details. See Section for more information.
Refresh	Available when viewing manual analysis details. See Section for more information.
Analyze	Available when viewing manual translation details. See Section for more information.
Analyze Filtered Events	Available after processing a file with a filter applied. See Section for more information.

4.3.2 Navigation

Using SEA, it is possible to monitor the binary event log files generated by a wide variety of computers all from a single web interface. In order to simplify the process of monitoring these diverse information sources, the web interface uses a hierarchical navigation tree composed of groups, nodes, categories, and binary event log files.

4.3.2.1 Navigation Tree Hierarchy

The entries in the navigation tree are as follows:

Figure 4–4 Navigation Tree - Hierarchy

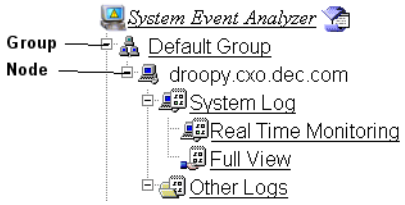

Table 4–4 Navigation Tree - Hierarchy

Folder	Description
Groups	Multiple computers that are logically associated. Groups contain one or more nodes.
Nodes	Individual computers. Each node contains two types of log files: System Log and Other Logs.
System Log	The binary system event log where the computer writes system information. By default, the System log contains Real Time Monitoring and Full View.
Real Time Monitoring	Automatic analysis results.
Full View	Manual analysis results for the system event log.
Other Logs	Any other binary event log files saved on the computer. These can include old files, files from other systems, and examples. Optionally, the other logs can be further divided by categories (See Section 4.8.1 for information on modifying SEA to use categories).

4.3.2.2 Features of the Navigation Tree

Figures 4–5 and 4–6 describe the features and functions of the navigation tree.

Figure 4–5 Navigation Tree - Collapsed

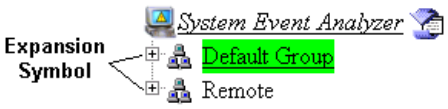

Figure 4–6 Navigation Tree - Expanded

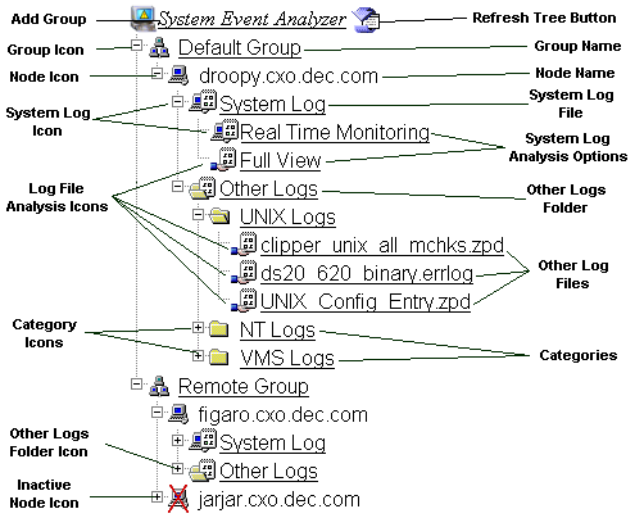

Table 4–5 Navigation Tree - Features

Feature	Description
Current Selection is Highlighted	In most browsers, the currently selected entry in the Navigation Tree is highlighted (Figure 4–5).
Collapsing Navigation	The tree structure can be collapsed to the group level (Figure 4–5).
Expanding Navigation	Click on the expansion symbol for an entry to view its contents. Once an entry is expanded, the expansion symbol changes to a collapse symbol. To hide the contents again, click the collapse symbol.
Icons	Each entry in the tree has a name and an icon that indicates its type. For example, in Figure 4–6 you can tell that the jarjar.cxo.dec.com node is inactive because of its icon.
Customizing the Navigation Tree	You can customize the navigation tree by adding and removing groups, nodes, categories, and binary event log files (see Section 4.4).
Viewing Results	You can view the results of automatic analysis and initiate manual analysis from the navigation tree (see Section 4.5).
Refreshing Navigation	If you modify the entries in the navigation tree, you may need to refresh the display so your changes appear. To refresh the navigation tree, click the Refresh Tree button.

4.4 Customizing the Navigation Tree

The first time you run the web interface using your profile, only one entry appears in the navigation tree: the node name for the Default Group. Ordinarily, this is the machine that you logged into.

You can customize the navigation tree display by creating new groups, adding nodes to groups, and selecting log files.

After you submit changes to the navigation tree, SEA refreshes the display. The refresh process may take a few seconds.

4.4.1 Groups

From the navigation tree, you can create new groups and remove existing groups.

4.4.1.1 Adding Groups

To add new groups follow these steps:

1. Click the “System Event Analyzer” link at the top of the navigation tree.

   The “Group Maintenance For System Event Analyzer” screen appears in the display frame (Figure 4–7). The Add Groups tab is already selected.

Figure 4–7 Add Group

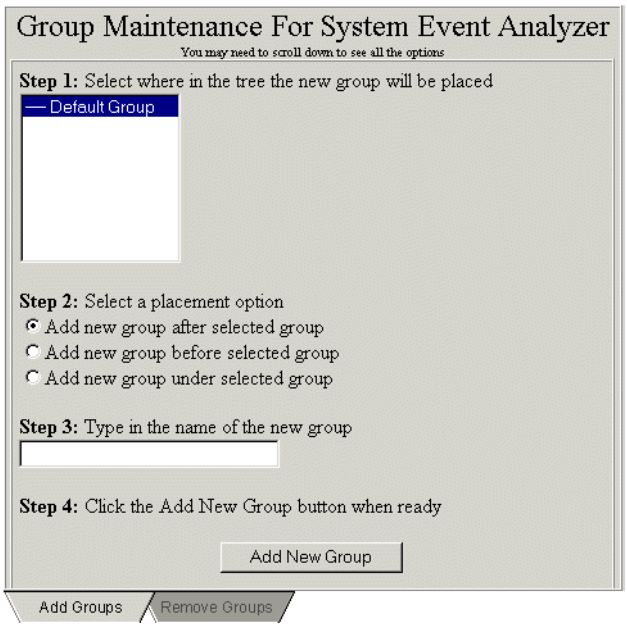

The location and placement options determine where you would like the new group to appear in the navigation tree relative to existing groups. By default, new groups are added after the selected group.

2. Select an existing group from the list.
3. Select a placement option from the radio buttons.
4. Enter the name for the new group in the text field. Be sure to follow these rules for naming groups:
    - Group names should be unique. If you enter a group name that is already in the navigation tree at the same level, SEA will not create the new group.
    - Group names should not use punctuation characters. These characters can cause JavaScript errors in the web interface.
    - Group names should be descriptive. If you leave this field blank, the group is named “newGroup” by default.
5. Click the Add New Group button.

   The new group appears in the navigation tree.

4.4.1.2 Removing Groups

Removing a group removes all the nodes and files contained in the group as well as all of the lower level groups nested under it.

To remove existing groups, follow these steps:

1. Click the System Event Analyzer link at the top to the navigation tree.
2. Select the Remove Groups tab at the bottom of the “Group Maintenance...” screen (Figure 4–8).

Figure 4–8 Remove Group

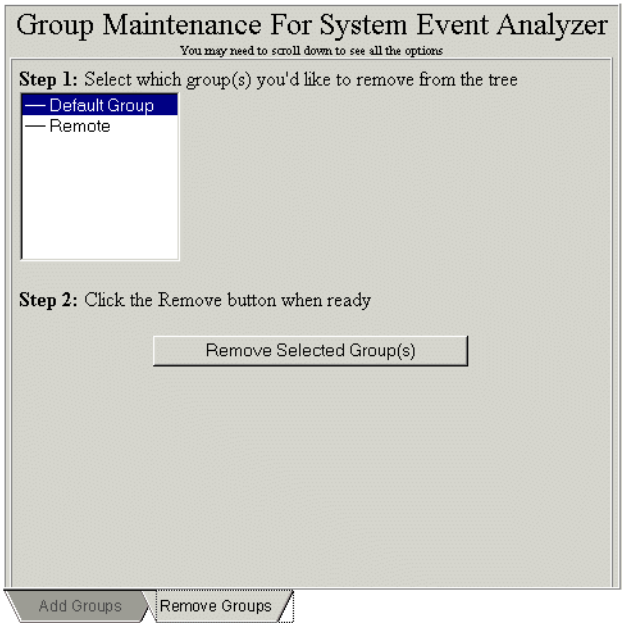

3. Select the group name or names you want to remove from the list.

   To select multiple groups, press CTRL and click on each group. If the groups are consecutive, press SHIFT and click on the first and last group names.

4. Click the Remove Selected Group(s) button.

   The groups disappear from the navigation tree.

4.4.2 Nodes

Expanding a group in the navigation tree displays the nodes contained in that group. Nodes can be expanded by clicking on the expansion symbol next to their name to reveal the log file types included in that node. You can add and remove nodes from the groups in the navigation tree.

4.4.2.1 Adding Nodes

Any computer where the Director is running can be added to your web interface navigation tree as a node. To add additional nodes follow these steps:

1. Select the group you want to add nodes to from the navigation tree.

The “Node Maintenance” screen appears in the display frame (Figure 4–9). The Add Nodes tab is already selected.

Figure 4–9 Add Node

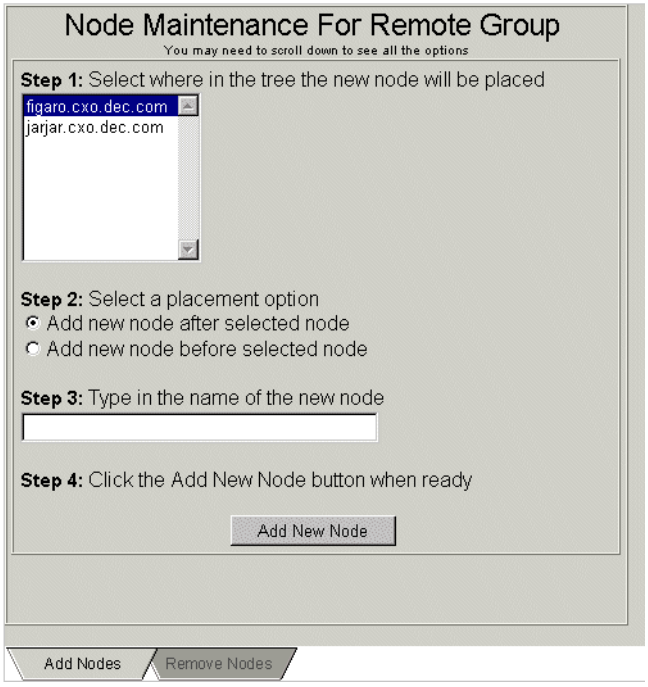

The location and placement options determine where you would like the new node to appear in the navigation tree relative to existing nodes. By default, new nodes are added after the selected node.

If no nodes currently exist for the group, skip steps 2 and 3.

2. Select an existing node from the list.
3. Select a placement option from the radio buttons.
4. Enter the name for the new node in text field. Be sure to follow these rules:
    - Node names should be valid hostnames or IP addresses. Hostnames must be accessible through the nameserver to be valid. For example, the hostname of a Windows machine using DHCP is not listed with the nameserver. In this instance, you would need to enter the IP address.
    - Node names should be unique. Entering the name of a node you are already connected to will overwrite the existing node and any Other Logs settings associated with it.
5. Click the Add New Node button.

The new node appears under its group in the navigation tree.

4.4.2.2 Removing Nodes

Removing a node removes all the additional binary event log files contained in the node from the navigation tree.

To remove existing nodes, follow these steps:

1. Select the group you want to remove nodes from in the navigation tree.
2. Select the Remove Nodes tab at the bottom of the screen (Figure 4–10).

Figure 4–10 Remove Node

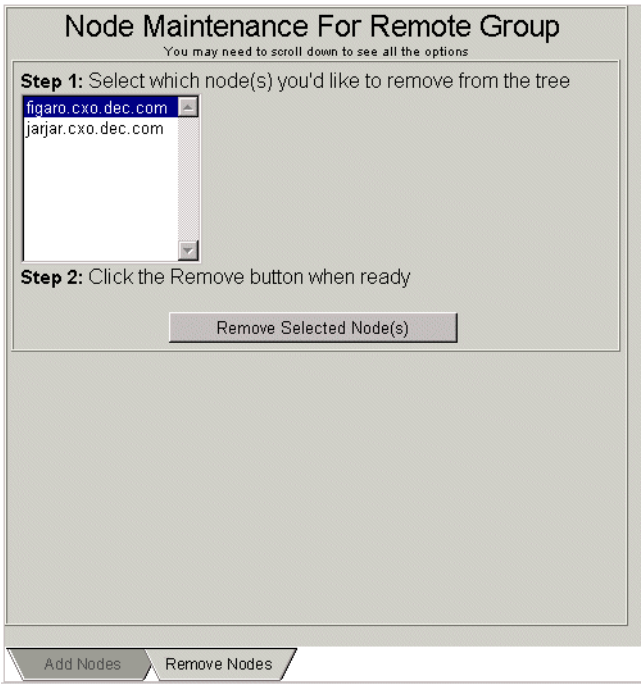

3. Select the node name or names from the list.

   To select multiple nodes, press CTRL and click on each node. If the nodes are consecutive, press SHIFT and click on the first and last node names.

4. Click the Remove Selected Node(s) button.

   The nodes disappear from the navigation tree. If the selected node is contained in multiple groups, removing it from one group will not affect its presence in other groups.

4.4.2.3 Activating Nodes

Nodes are either active or inactive and by default when you connect to a node or load a profile that connects to other nodes, all the nodes are active. A node is only classified as inactive if SEA cannot connect to it. Inactive nodes appear in the navigation tree with a red “X” through their icon.

If a node is inactive, you can try to connect to it manually. To connect to a inactive node, follow these steps:

1. Click the expansion icon for the node.

   The only available option is “Activate this node” (Figure 4–11).

Figure 4–11 Activate Node

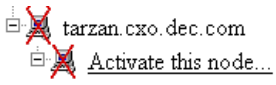

2. Click the “Activate this node” link.

   If the Director on the remote node is accessible, a message appears in the display frame (Figure 4–12) and the navigation tree is updated to show the new status.

Figure 4–12 Activating Node Message

   If the Director is not accessible, a message appears in the display frame (Figure 4–13) and the navigation tree is not changed.

Figure 4–13 Unable to Activate Node Message

4.4.3 Categories

Categories are an optional feature that is disabled by default. If you want to use categories, you must enable the feature using the User Settings tab on the Settings screen (see Section 4.8).

Categories provide a method for grouping the log files listed under the Other Logs folder. If you use categories, SEA provides another layer of folders under the Other Logs folder. This feature may be useful if you monitor numerous log files.

4.4.3.1 Adding Categories

Once you have enabled the categories feature, you can add categories to the navigation tree. To add categories, follow these steps:

1. Select the Other Logs folder for the node you want to have new categories.

   The Category Maintenance screen appears in the display frame (Figure 4–14). The Add Categories tab is already selected.

Figure 4–14 Add Category

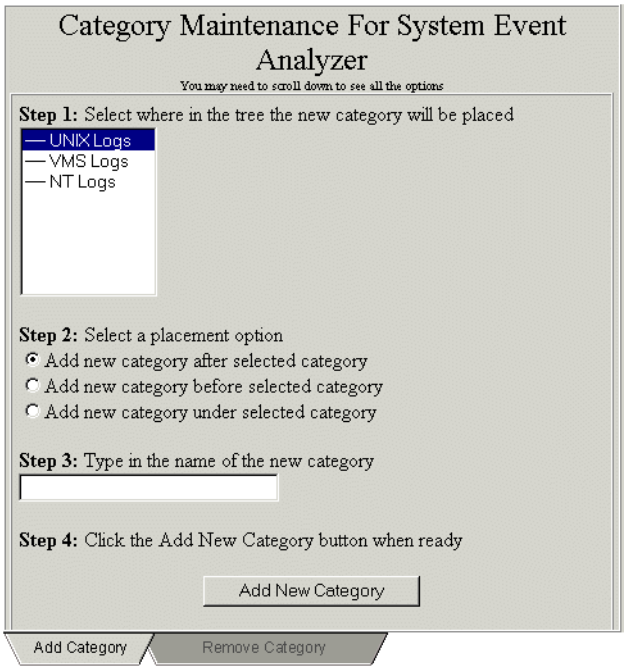

   The location and placement options determine where you would like the new category to appear in the navigation tree relative to existing categories. By default, new categories are added after the selected category.

If no categories currently exist for the group, skip steps 2 and 3.

2. Select an existing category from the list.
3. Select a placement option from the radio buttons.
4. Enter the name for the new category in the text field. Be sure to follow these rules for naming categories:
    - Category names should be unique. If you enter the name of an existing category, SEA will not create the new category.
    - Category names should not use punctuation characters. These characters can cause JavaScript errors in the web interface.
    - Category names should be descriptives. If you leave this field blank, the category is named “newCat” by default.
5. Click the Add New Category button.

   The new category appears in the navigation tree.

4.4.3.2 Removing Categories

Removing a category removes all the binary event log files contained in the category from the navigation tree.

To remove existing categories, follow these steps:

1. Select the Other Logs folder for the node you want to remove categories from in the navigation tree.
2. Select the Remove Category tab at the bottom of the screen (Figure 4–15).

Figure 4–15 Remove Category

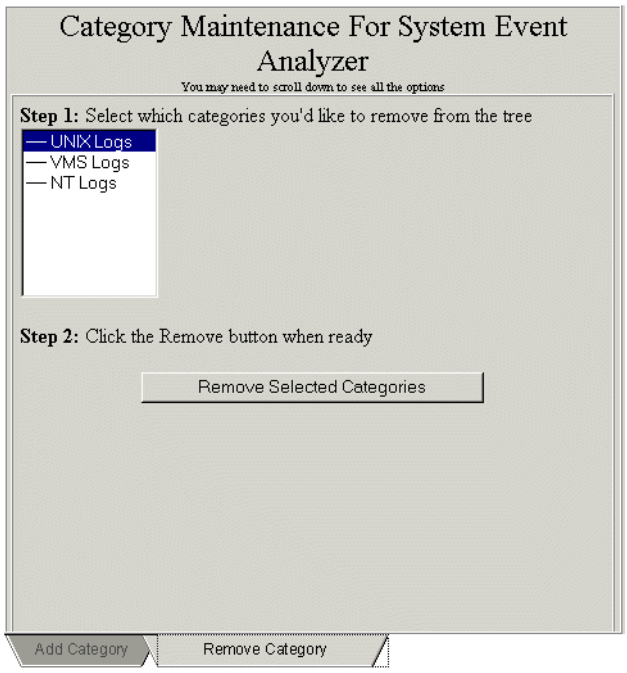

3. Select the category name or names you want to remove from the list.

   To select multiple categories, press CTRL and click on each category. If the categories are consecutive, press SHIFT and click on the first and last category names.

4. Click the Remove Selected Categories button.

   The categories disappear from the navigation tree. If a log file is contained in multiple categories, removing it from one of the categories will not affect its presence in the others.

4.4.4 Log Files

Each node contains binary event log files. Log files are separated into two different types: the binary system event log and all other binary event logs.

4.4.4.1 System Log

The system log is the binary event log file where system events are written. You cannot change this log file. Click the expansion symbol to view the analysis options for the system log in the navigation tree.

- Real Time Monitoring – shows the results of automatic analysis in the display frame.

- Full View – manually analyzes the system event log and processes all the events in the file.

See Sections 4.1.2 and 4.5 for more information on analysis.

4.4.4.2 Other Logs

The Other Logs folder in the navigation tree contains entries for binary event log files other than the system event log. These can include the example binary log files included with SEA, or any other binary event log file located on the node. Initially, there are no sub-entries under Other Logs in the navigation tree.

If you are using categories, the Other Logs entry contains the categories you have created and the category folders contain entries for binary event log files.

In order to add saved log files to the navigation tree, they must be viewable in the Add Log Files list. For a file to be viewable, it must meet both of these criteria:

- The log file must have a `.sys`, `.evt`, `.zpd`, or `.errlog` extension. If you wish to add a file with a different extension, you will need to rename the file so it uses an acceptable file extension.
- The log file must be saved in the `svctools` directory (created during installation), one of its subdirectories, or one of the directories you configured in the `CA.WUI.OLDirs` key in the DESTA registry. Files that are stored in these locations are automatically displayed in the list. For more information, see section 7.7.2.

The best place to store log files (as well as other user data) is in one of the userdata subdirectories:

```
svctools\specific\ca\userdata
svctools\common\ca\userdata
```

Files stored in these subdirectories are automatically backed up and saved if you uninstall, reinstall, or upgrade WEBES. For more information on storing user data, see the *WEBES Installation Guide*.

If you want to store files elsewhere, you can configure WEBES by adding a comma separated list of file paths to the `CA.WUI.OLDirs` key in the DESTA registry. For more information, see section 7.7.2.

You can also enable a text entry field for specific users. The text field allows users to add log files to the Other Logs list by entering the path and filename of an event log located anywhere in the file system. For more information, see section 7.7.3.

Adding Other Logs

Follow these steps to add other log files:

1. Open the Other Logs screen in the display frame.

If you are using categories, click on the category name for the node. Otherwise, click on the Other Logs link for the node.

The Other Logs screen opens in the display frame (Figure 4–16). The Add Files tab is already selected.

Figure 4–16 Add Log Files Tab

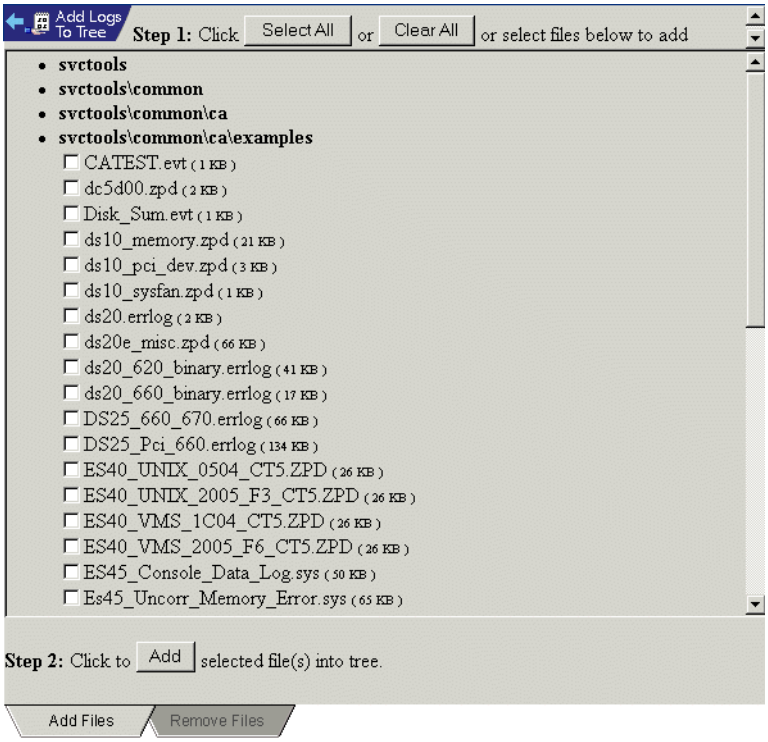

2. Select the desired binary event log files:
   - Click the Select All button to select all the listed log files.
   - Click the check box for each file. You can select multiple check boxes.
   - Click the Clear All button or uncheck a selected check box to deselect files.
3. (Optional) If enabled, enter the path and filename in the text field (see section 7.7.3 for more information).
4. Click the Add button.

   The binary event log file is added to the navigation tree under the Other Logs entry or appropriate category for the node.

Removing Other Logs

To remove binary event log files from the navigation tree, follow these steps:

1. Open the Other Logs screen in the display frame.

If you are using categories, click on the category name for the node. Otherwise, click on the Other Logs link for the node.

2. Select the Remove Files tab from the bottom of the screen (Figure 4–17).

Figure 4–17 Remove Log File Tab

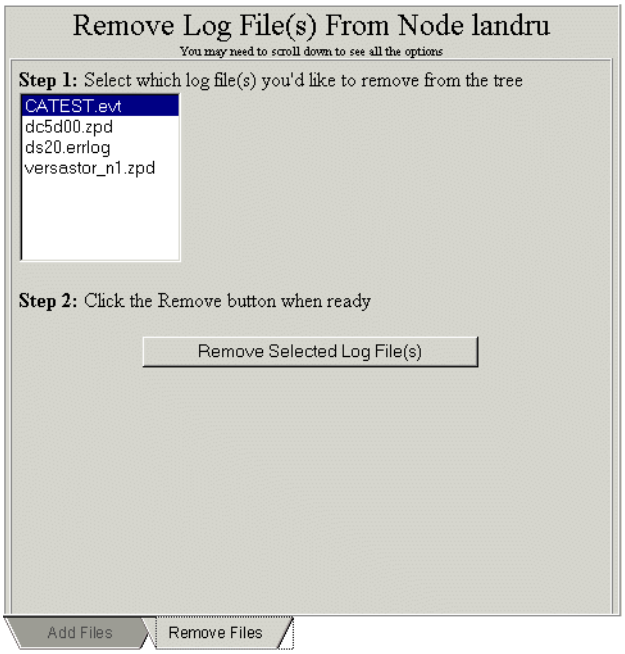

3. Select the log file name you want to remove from the list.

   To select multiple files, press CTRL and click on each file name. If the files are consecutive, press SHIFT and click on the first and last file names.

4. Click the Remove Selected Log File(s) button.

   The navigation tree is refreshed to reflect your changes.

4.5 Processing Log Files

You can process a log file, check its status, and view the results using any of the following methods:

- Selecting System Log or Real Time Monitoring runs automatic analysis on a node.
- Clicking Full View manually analyzes a node's system event log and display the results.
- Clicking a Log File name under Other Logs runs manual analysis on the file and displays the results.

Viewing Process Status

When analysis is successfully started, the log file's icon is animated. Once the file is processed, the icon in the toolbar changes to reflect the status of the log file (see Section 4.5.2) and the results of processing are shown in the display frame.

Viewing Results

Both automatic and manual analysis results are shown in the display frame. The information is organized under the following tabs:

- Problem Reports – results of analysis
- Summary – description of the contents of the log file (only available with manual analysis)
- Events – translation of the events contained in the log file

Note

If you have configured the User Settings so SEA only performs manual translation, rather than translation and analysis, the Problem Reports tab is empty. See Section 4.8.1 for more information on User Settings.

Analysis Failed

If the file cannot be processed for any reason, a message similar to the one in Figure 4–18 is shown.

Figure 4–18 Analysis Failed Message

4.5.1 Additional Toolbar Functions

SEA provides additional functionality depending on the type of processing you are performing.

Figure 4–19 Additional Toolbar Functions

Button	Name	When Does It Appear in the Toolbar?	Description
	Clear Results Button	When you are performing Automatic Analysis.	The Clear button removes all the entries (problem reports and events) from the display tabs.
	Reprocess File Button	When you are performing Manual Analysis.	The Reprocess button forces SEA to discard the previous analysis results and reprocess binary log files.
	Analyze File Button	When the User Settings are configured to perform manual translation.	Clicking the Analyze button will perform analysis for the current log file. Thus, if you need to perform analysis, it is not necessary to change the User Settings and reprocess the file.
	Analyze Filtered Events Button	When you use a filter for processing a log file,	Clicking the Analyze Filtered event button allows you to repeat the analysis using only the events that met the filter criteria.

4.5.2 Processing Status

With large log files, translation and analysis operations are not instantaneous. After you have started processing a log file there are several ways to check the operations progress. You can check the processing status from either the navigation tree or the Progress window.

4.5.2.1 Navigation Tree

You can quickly determine the status of automatic or manual analysis by looking at the icons in the navigation tree. Figure 4–20 shows the icons used to indicate analysis results.

Figure 4–20 Status Icons

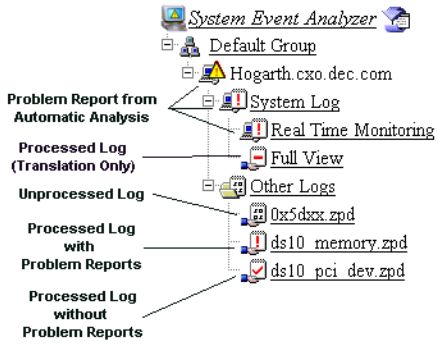

When automatic analysis generates a problem report, exclamation points are added to the icons for the node, system log, and real time monitoring. The icon remains changed until problem report details are viewed and the tree is refreshed. If another problem report is generated after the tree is refreshed, exclamation points are added to the icons again.

You can also determine the results of manual analysis on a binary event log file by checking the icons. SEA uses a animated yellow icon when a binary log file is being read and an animated green icon during analysis. If processing has completed and problem reports were generated, an exclamation point is added to the icon. Otherwise a check mark is added to the icon. Unlike the icon changes associated with automatic analysis, the manual analysis icon changes remain visible until you close the web interface session.

Note

If you are only performing translation, when processing completes the icon will appear with a dash. See Section 4.8.1 for more information on configuring the web interface to only perform translation.

4.5.2.2 Progress Window

You can open the Progress window by clicking on the Progress Window button in the toolbar (see Figure 4–3).

The Progress window opens in a new browser window (Figure 4–21).

Figure 4–21 Progress Window

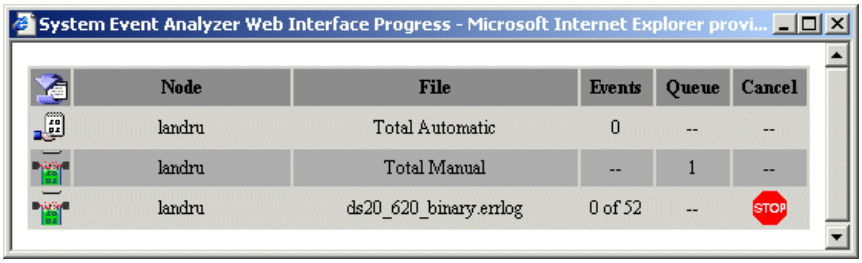

The Progress window provides statistics for all the log files that are currently being analyzed by SEA, including one automatic analysis job and multiple manual analysis jobs. The information in the Progress window includes:

- The node where the log file is located
- The name and location of the log file
- The number of events in the file

The position of each file in the queue is displayed, and information is dynamically updated as the processing changes. When a file finishes processing, it is removed from the window.

When monitoring the progress of a file, you can refresh the display manually by clicking the refresh icon in the upper left hand corner. To stop processing an active file, click on the Stop icon.

4.5.3 Working With Results

After processing, the results of analysis are shown on the tabs in the display frame (Figure 4–22).

Figure 4–22 Additional Entries Navigation

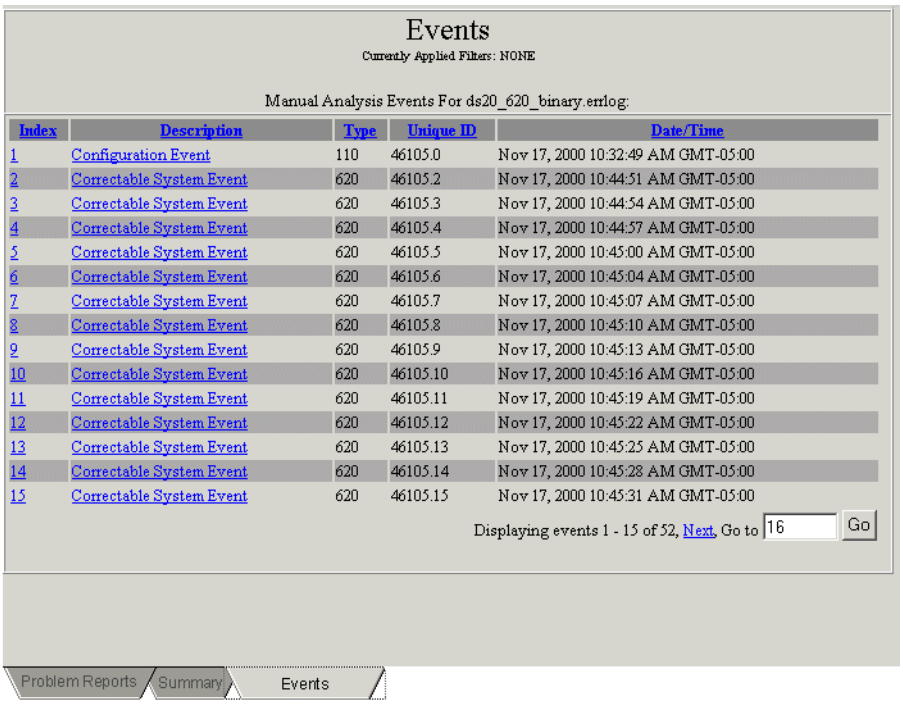

When there are many entries, you can use the navigation options to page through the results.

- Use the Previous and Next buttons to move between entry screens.
- Enter a number in the entry field and click Go to display a specific entry.

Note

You can control the number of entries shown in a tab with the options in the User Settings window. Refer to Section 4.8.1 for more details.

4.5.3.1 Problem Reports

The Problem Reports tab displays a list of the reports that were generated by analysis. An example of the problem report list is shown in Figure 4–23.

Figure 4–23 Problem Report Tab

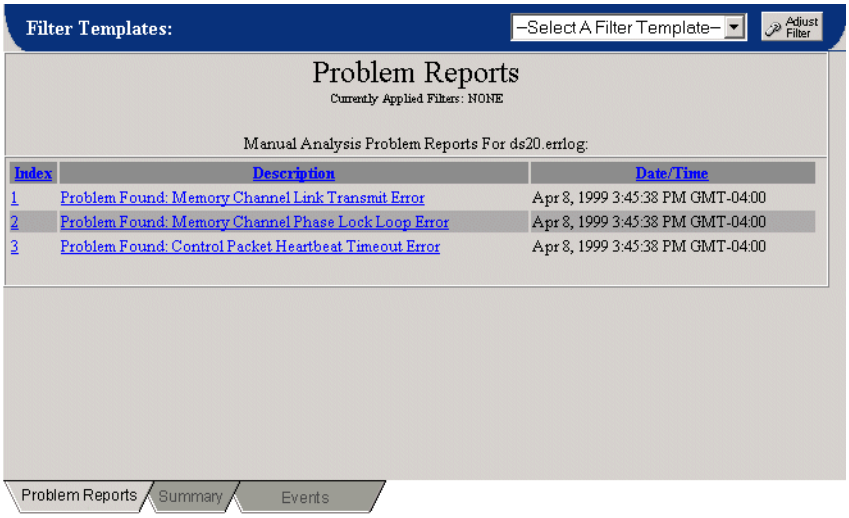

The filters used when generating the problem reports are listed at the top of the screen. However, the display only shows the filters that apply to problem reports and may not list all the filters you selected.

When working with problem reports, these options are available:

- To sort the entries in the report list select the column headers. See Section 4.5.3.4 for more details on sorting.
- To view the contents of a report, click on its entry in the list of available problem reports. See Section 4.5.3.5 for information on viewing reports.

The problem reports generated by the web interface are the same as those generated by the CLI.

- Refer to Chapter 5 for more information on analysis.
- Refer to Appendix A for an example of a problem report.

4.5.3.2 Summary

The Summary tab is only available when you perform manual analysis. If you select Real Time Monitoring from the Navigation Tree, for example, the Summary tab is not displayed.

When performing manual analysis, the Summary tab describes the event types contained in the binary event log file (Figure 4–24).

Figure 4–24 Summary Tab

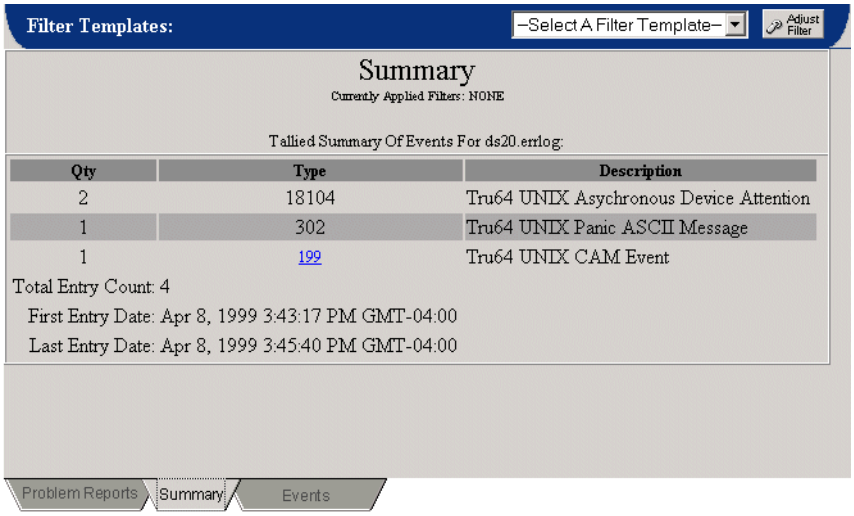

Each event type is listed along with the number of occurrences. The time stamps for the first and last events are listed under the summary information.

The filters that were applied are listed at the top of the screen. Be aware that the screen only shows the filters that apply to the summary report and may not list all the filters you selected.

Refer to Section 5.8 for details on the summary information presented.

4.5.3.3 Events

The Events tab shows a list of the events contained in the binary event log file. Depending on the filtering options that were applied during processing, all the events in the log file may or may not be shown (Figure 4–25).

Note

You can control the fields that are shown on the events tab from the User Settings window. Refer to Section 4.8.1 for more details.

Figure 4–25 Events Tab

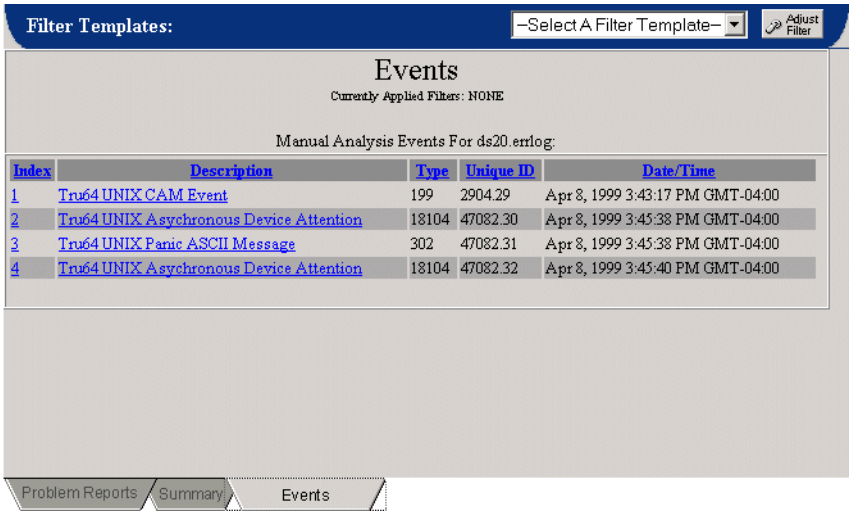

The filters that affected the output are listed at the top of the screen. Be aware that the screen only shows the filters that apply to events and may not list all the filters you applied.

When working with events, these options are available:

- To sort the events list, use the column headers. See Section 4.5.3.4 for more details on sorting.
- To view the translation of an event, click on its entry in the list. See Section 4.5.3.5 for information on viewing translation details.

The translated events shown by the web interface are the same as those shown by the CLI.

- Refer to Chapter 5 for more information on event translation
- Refer to Appendix A for an example of a translated event.

4.5.3.4 Sorting Results

You can sort the results of analysis using either the column headings on the tabs in the display frame, or by using a filter.

Sorting with Column Headings

- Sorting with the column headings only impacts the entries currently shown. Therefore, if there are too many entries to be listed on a tab, the column headings will only sort the entries that are displayed rather than all the output produced by processing the log file. In most cases, this limitation only impacts the Events tab.

- You can sort the results shown on any tab using the field names that appear in blue (i.e., as hypertext links). Simply click on the field name to sort based on that field. An arrow appears next to the field to indicate the direction of the sorting. The sorting options are applied to all the tabs, regardless of which tab was used to specify the sorting criteria.
- Entries can be sorted in either ascending or descending order. To change the sort order, click on the field name a second time. The arrow next to the field changes direction to indicate the new sort order. When the arrow is pointing up, it indicates an ascending sort. When the arrow is pointing down, it indicates a descending sort.
- If you are working in multiple windows, sorting only applies to the current window.

Sorting with a Filter

- Using a filter to sort entries impacts all the output generated by processing a log file, regardless of how many screens are required to show all the entries.
- For more information on using a filter to sort output, refer to the information on applying filters in Section 4.7.

4.5.3.5 Displaying Details

The Problem Reports tab lists the reports generated by analysis and the Events tab lists the events in the binary log. You can view the details of a problem report or the translated text of an event by clicking on an entry in the list. Depending on the User Settings selected (see Section 4.8.1), the details will either be shown in the display window or in a new browser window.

In order to make viewing details easier, navigation buttons are available at the top of each detailed entry. The navigation buttons for the Problem Reports tab and Events Tab are shown in Figures 4–26 and 4–27.

Figure 4–26 Navigation Buttons – Problem Reports

Figure 4–27 Navigation Buttons – Events

The buttons are used to move between entries in the list.

- You can view the details for other events in the list using the Previous and Next buttons.

  When paging between entries, the column heading sort order always reverts back to the Index column in ascending order. Filter sorts, however, still apply.

- Click the Index button to redisplay the list of entries.

  If you select “Put Event Details In A New Window” in your User Settings, the Index button is not available. Clicking the Previous and Next buttons displays all entries in the new window. See Section 4.8.1 for more information on user settings.

- The Event Details tab includes a drop down list that can be used to change the report type. Refer to Chapter 5 for more information on translation report types.

4.6 Creating New Log Files

To create a binary event log for use with SEA, follow these steps:

1. Click the New Binary Log File button in the toolbar (see Figure 4–3).

   The New Binary Log Screen appears in the display area (Figure 4–28).

Figure 4–28 New Binary Log Screen

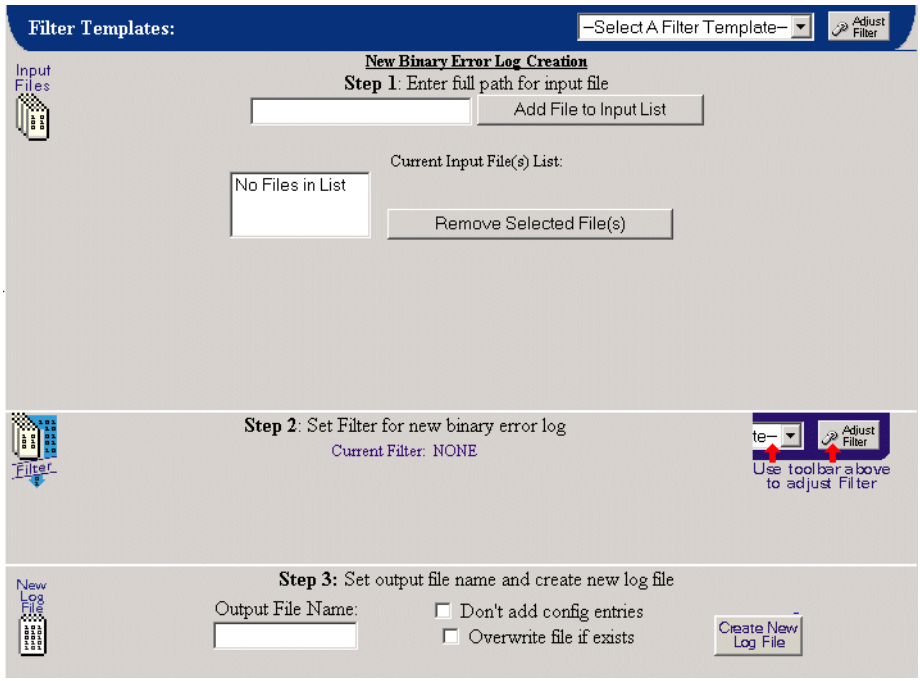

2. Enter the input file name, including its path, in the Input File text box.
3. Click the Add Input file Button.

   The file is added to the Currently Selected Input Files list.

4. Repeat steps 2 and 3 until all the desired input files are added.

Note

If you want to remove one of the input files you added, click on the filename in the Currently Selected Input Files list and click the Remove Selected Input Files button. You can select multiple files by holding the Ctrl key while you click on the filenames.

5. Specify the desired filtering options by either creating a new filter or applying an existing template.
   - To specify filtering criteria, click the Adjust Filter button at the top of the screen and use the Adjust Filter screen to select filtering options (see Section).
   - To apply an existing filter template, select the desired template from the drop down list at the top of the screen.

   For more information on filtering, refer to Sections 4.7 and 4.8.1.2.

6. Enter the output file name in the Output File text box.

Note

New binary log files are automatically stored in the `specific\ca\userdata` subdirectory located under the installation directory, hence it is not necessary to include a path with the Output filename. For more information on storing user data, see the *WEBES Installation Guide*.

7. If you have established a filter that excludes configuration entries and you want to preserve that filtering in the output file, select the “Don’t add config entries” check box.
8. If the output file name already exists and you want to replace the existing file, select the “Overwrite file if exists” check box.

   If you do not select this check box, and enter a filename that already exists, you will receive an error message.

9. Click the Create New Log File button to process the input files and create the new binary log file.

Note

It is possible to construct a filter that prevents any events from being added to the new log file. If this is the case, no log file will be created. However, even if this is the case, when the Overwrite option is selected any file with the same name as the output file will be lost.

4.7 Applying Filters

You can apply filters when processing existing log files and when creating new binary log files. You can also use filters to specify how problem reports and events are sorted. Specify the desired filter using the Filter Templates bar at the top of the screen (Figure 4–29).

Figure 4–29 Filter Templates Bar

If you have previously created filter templates, they will be listed in the drop-down list. You can either:

- Select an existing filter from the drop down list and if necessary modify it by clicking the Adjust Filter button and changing the filtering options.
- Click the Adjust Filter button and define a new filter.

Note

Modifying or defining a filter from the Filter Templates bar does not change an existing filter or save a new filter. Your changes are only used with the current operation. Use the Filters option under User Settings to create new templates (see Section 4.8.1.2).

When you use filters in conjunction with analysis and translation the filter description will be shown with the results. However, the filtering options you select are only applied to the appropriate output. Thus, if you select a filter that only affects event translation, rather than problem reports and translation, the filter will be listed with the event details but not with the problem reports details. Figure 4–30 depicts a filter description from the event details.

Figure 4–30 Filter Description

Refer to Section 4.8.1.2 for more information on creating and modifying filters.

4.8 Modifying Settings

The web interface settings enable you to control how the WEBES Director functions and modify the web interface to suit your preferences. To access the settings, click the settings button in the toolbar. This updates the web interface, replacing the normal navigation bar with the User Settings navigation bar. The display frame is updated to show the User Settings screen.

You can modify both User and Director settings.

4.8.1 User Settings

The user settings are used to modify the web interface, configure filtering information and determine what translation information is displayed. To access the User settings, click the Settings button in the toolbar and then select the User Settings tab.

Figure 4–31 User Settings

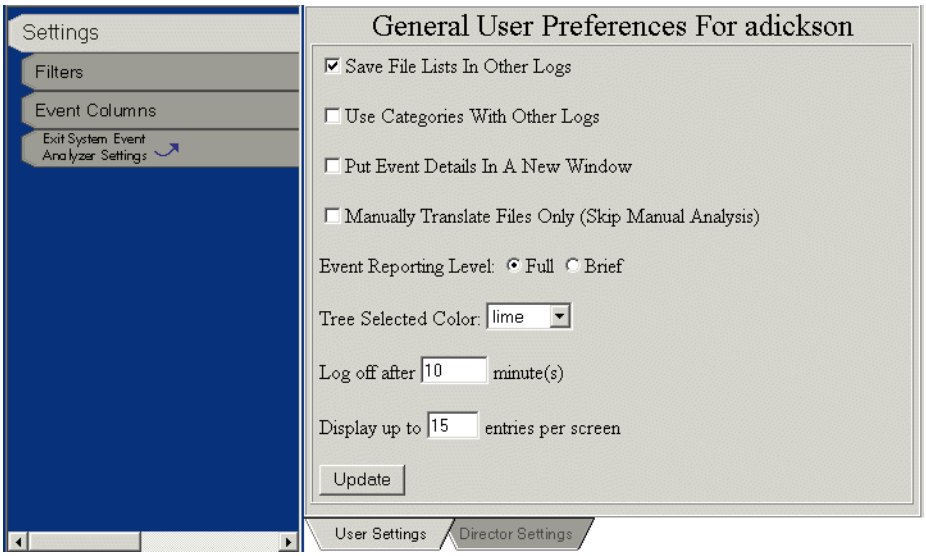

Use the tabs located at the left side of the screen to navigate the User settings (Figure 4–31).

Figure 4–32 User Settings Navigation

Option	Description
Settings	Displays the web interface general configuration options. Refer to Section 4.8.1.1 for more information.
Filters	Opens the Filter Preferences screen which is used to define filter templates and set a default filter. See Section 4.8.1.2 for more information.
Event Columns	Specifies the translation information you want to view. See Section 4.8.1.3 for more details.
Exit Settings	Closes the settings screen.

4.8.1.1 General Options

The general options screen is shown in Figure 4–31.The General User Settings screen presents the following options:

Table 4–6 General User Settings Options

Option	Description
Save File Lists in Other Logs	Select this option if you want the navigation tree to save a record of all the log files listed under Other Logs when you log off SEA. If this option is selected, the log files will remain in the navigation tree until you manually remove them. If this option is not selected, the Other Logs section of the tree will be empty when you logon.
Use Categories With Other Logs	Select this option to use categories with log files. Refer to Section 4.4.3 for more on categories.
Put Event Details In A New Window	Opens a new browser window for the details of a problem report or event selected from the list of entries. The list of entries will remain open in the original window.
Manually Translate Files Only (Skip Manual Analysis)	Prevents SEA from performing manual analysis for log files. This affects the output when you select an entry from the Other Logs area and when you perform manual analysis on the system event log.
Event Reporting Level	Specifies the default level of reporting for translated events. The available report types are brief and full. Refer to Section 5.2.3 for more information on report types.

Table 4–6 General User Settings Options

Option	Description
Tree Selected Color	Enables you to specify the color used to highlight selected entries in the navigation tree.
Entries per screen	Specifies the number of entries displayed at one time on the output tabs. Refer to Section 4.5.3 for more information.
Log Off Time	By default, SEA logs your profile off ten minutes after you close your connection with the Director. You can change the amount of time by entering a new value in the text box. All values are in minutes. See Section 4.10 for more information on logging off. (Setting the Log Off time to zero is not recommended. Refer to Section 4.10 for more details.)

Click the Update button to save your changes to the settings.

4.8.1.2 Filters

Filtering is used to reduce the number of events processed when you perform translation or create a new log file. With large log files, using only a subset of the events can improve processing time and enhance output by displaying only the most pertinent information.

Within the web interface, filtering is performed using templates. Templates contain pre-defined filtering functions that can be applied to SEA functions.

The Filtering Preferences screen allows you to create new filter templates, modify existing filter templates, or select default filtering options. To access the Filtering Preferences screen, click the Filters button in the User Settings navigation bar.

Figure 4–33 Filter Preferences

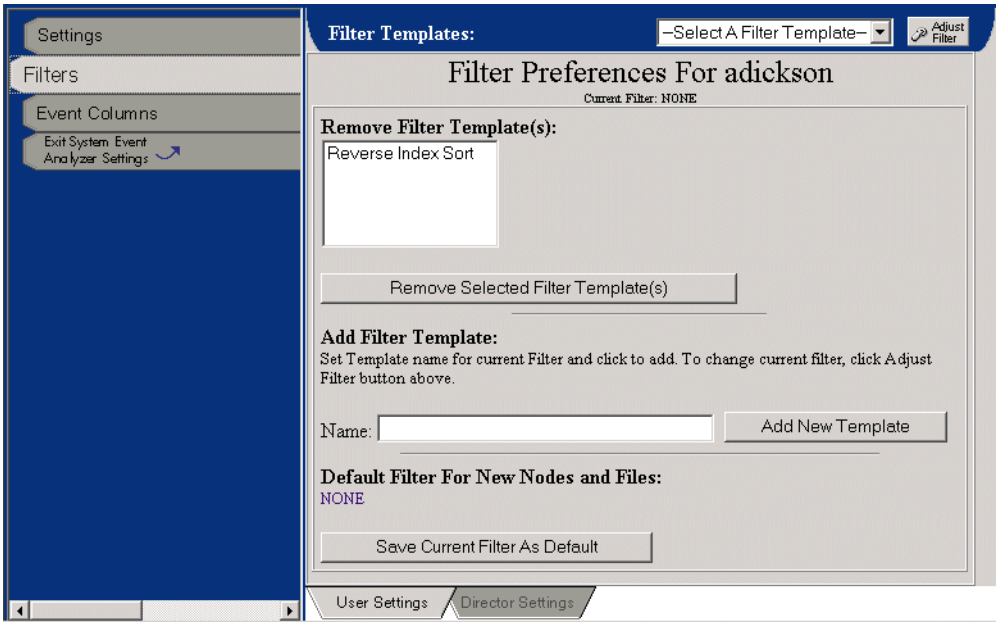

Creating and Modifying Filter Templates

To modify a filter template or create a new filter, use the following procedure:

1. Select the filter you want to modify from the drop-down list in the Filter Templates bar.

   If you want to create a new filter from scratch, you do not need to select an existing template.

2. Click the Adjust Filter button located in the Filter Templates bar.

   The Adjust Filter screen appears (Figure 4–34). If you are modifying an existing template, the contents of that filter are listed in the Currently Applied Filters list.

Figure 4–34 Adjust Filter

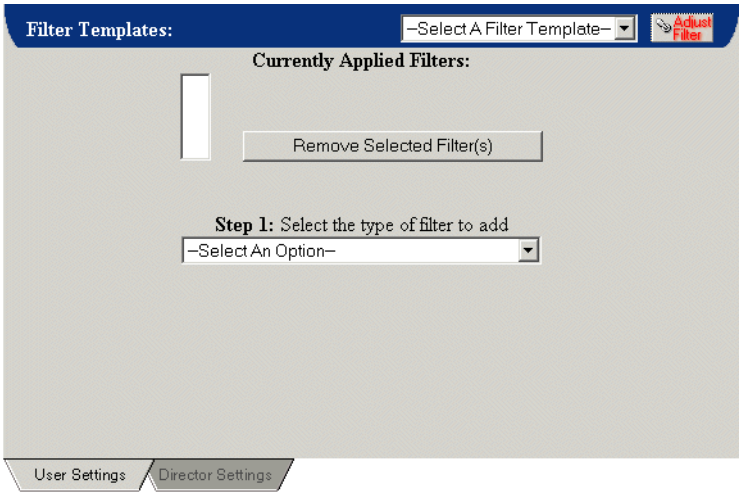

3. Ensure that all the filter information in the Currently Applied Filters list is correct.

   Initially, this field will display the contents of the filter template you selected. You can delete any filter by selecting it and clicking the Remove Selected Filters button. If you are creating a new filter the list is blank.

4. Choose any additional filtering criteria from the drop-down list.

   Once you have selected a filter type, the Filtering screen is dynamically updated to include the valid operators (Figure 4–35). Be aware that all the operators are not valid for all filter types.

5. Select the radio button that corresponds to the desired operator.
   - Not equal to (!=)
   - Equal to (=)
   - Greater than (>)
   - Less than (<)

Figure 4–35 Filtering Criteria

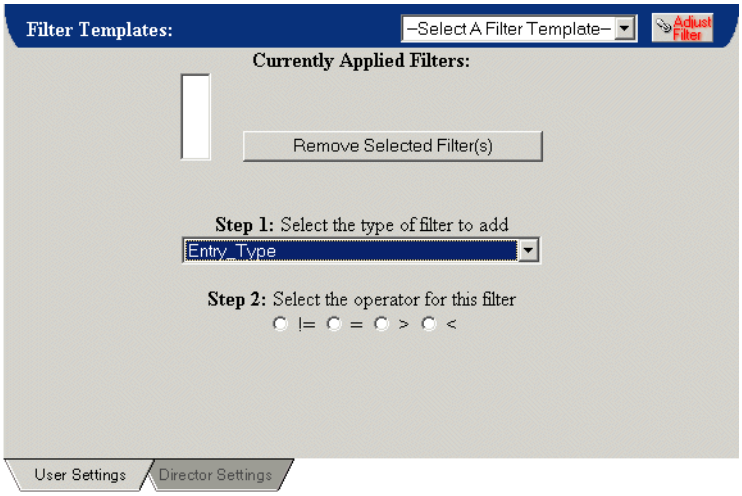

Once you have selected an operator, the screen is updated to include a drop-down list of values or a text entry field (Figure 4–36).

6. Select or enter the appropriate value.

Figure 4–36 Filtering Operators

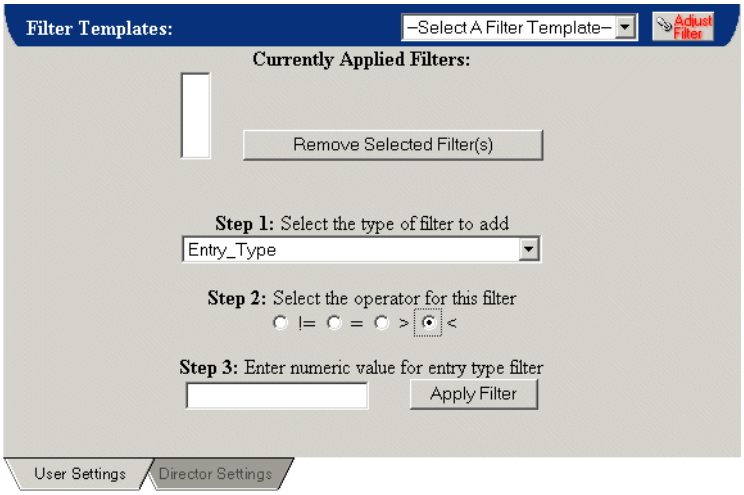

7. Click the Apply button.

   The filter is added to the list of Currently Applied Filters (Figure 4–37).

Figure 4–37 Applied Filter

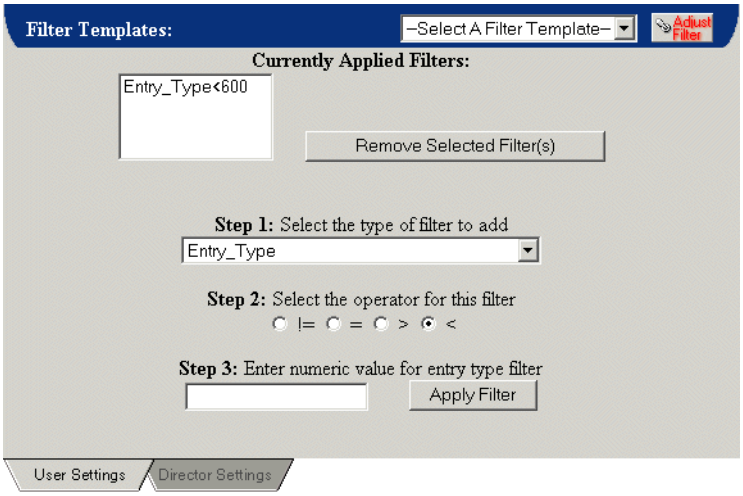

8. Repeat steps 3 to 7 until all the necessary filters have been added.
9. Click the Adjust Filter button again to close the Adjust Filter screen and return to the Filtering Preferences screen (Figure 4–33).

   The Filtering Preferences screen describes the contents of the new filter.

10. Save the new filter as a template by entering a filter name in the Name text box and click the Add New Template button.

    SEA will update the Filter Templates list and add the new filter to the drop-down list in the Filter Templates bar.

If you are creating a new filter from one of the details tabs rather than the User Settings window, the filter is saved for that file or automatic node, but not as a template that can be applied elsewhere. Otherwise the process is the same.

Default Filters

You can apply default filtering options to all the analysis and translation operations performed from the web interface using the Filter Preferences screen (Figure 4–33).

To set a default filter, use the following procedure:

1. Select the desired templates from the drop-down list in the Filter Templates bar.

   It is not necessary to select a template if you do not want to use an existing template.

2. If necessary, click the Adjust Filter button and modify the filter template or create a new template.
3. Click the Save Current Filter As Default button.

It is not necessary to save the default filter as a template. If you want to, you can use the Adjust Filter screen to create a filter and then save it as the default filter without saving it as a template.

Deleting Templates

You can delete a filter template from the Filter Preferences screen (Figure 4–33), using the following procedure.

1. Click on the name of the filter you want to delete in the Filter Templates list.

   You can select multiple filters by holding the CTRL key while you click the filter names.

2. Click the Remove Selected Filters button.

4.8.1.3 Event Columns

The Event Columns screen is used to specify the information displayed by translation on the Events tab (see Section 4.5.3.3 for more information on translation details).

Figure 4–38 Event Columns

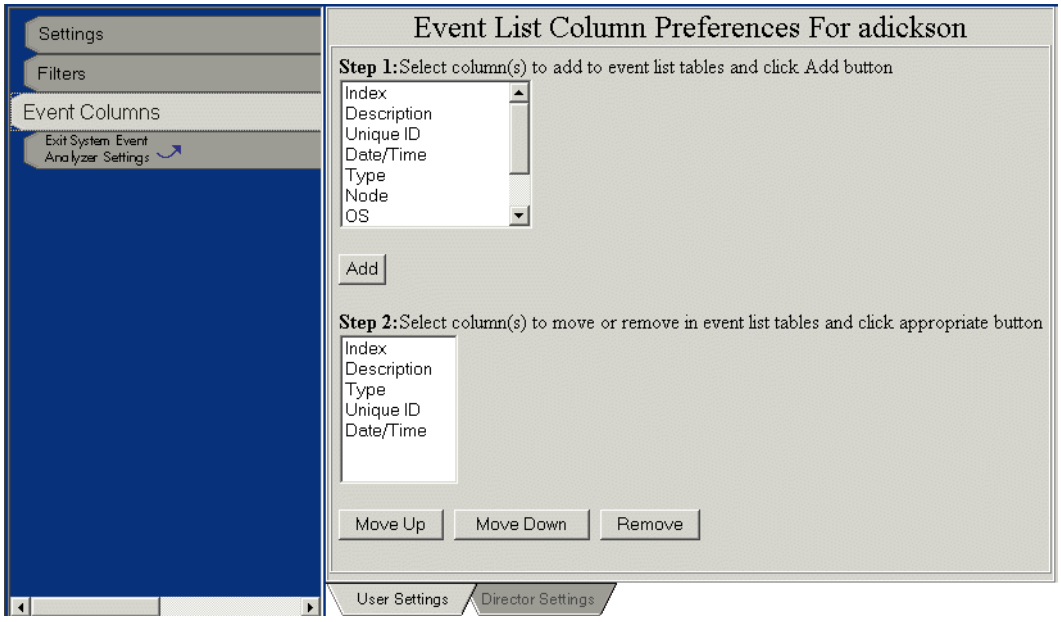

The Event Columns screen lists the field headings for event translation. You can designate which translation information is shown on the Events tab using the following procedures.

Adding Fields

To add fields, determine which additional translation fields need to be shown. The first list displays all the available translation fields and the second list indicates the fields that are currently shown.

1. Select the desired field from the first list by clicking on its name.

   You can select multiple entries by holding the Ctrl key while you select their names.

2. Click the Add button.

   The selected fields are added to the end of the second list and shown under the Events tab.

Rearranging Fields

The order of the fields in the second list indicates the order of the information on the Events tab. To rearrange the fields:

1. Select the field that needs to be moved by clicking on its name in the second list.
2. Move the field to its new location.
   - Click the Move Up button to move the field up in the list.
   - Click the Move Down button to move the field down in the list.

Removing Fields

To remove a field:

1. Select the field from the second list by clicking on its name.

   You can select multiple fields by holding the Ctrl key while you select their names.

2. Click the Remove button.

4.8.2 Director Settings

The Director settings are used to modify WEBES components, register rule sets. To access the Director settings, click the Settings button in the toolbar and then select the Director Settings tab.

Figure 4–39 Director Settings

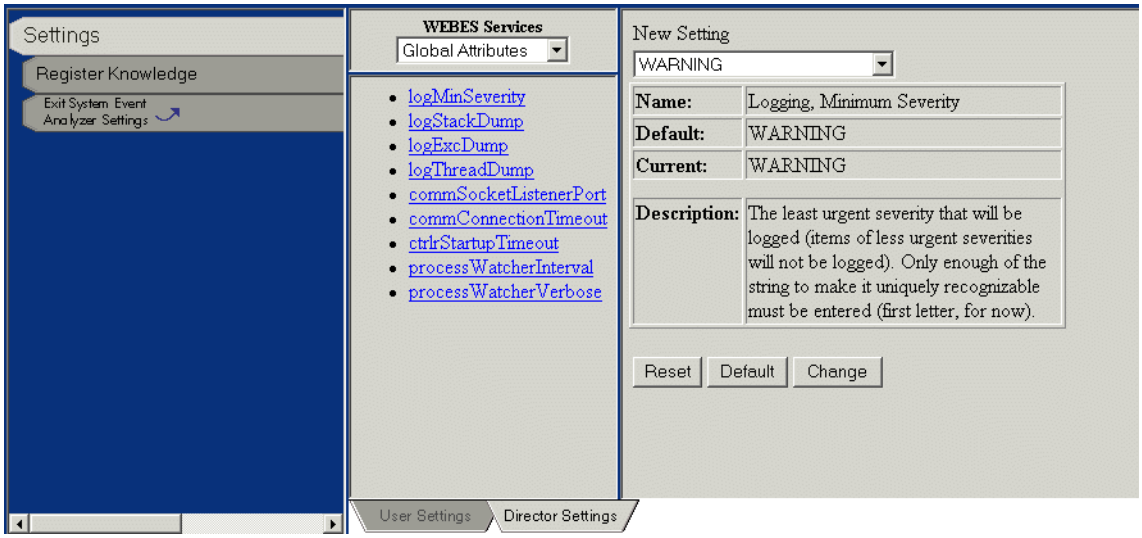

Use the buttons located at the right side of the window to navigate the Director settings.

Table 4–7 Director Settings Navigation

Option	Description
Settings	Displays the configuration settings for the Director. Refer to Chapter 7 for information on changing the Director settings.
Register Knowledge	Displays the knowledge rule sets that can be installed. Refer to Chapter 6 for more information on rule sets and analysis.
Exit System Event Analyzer Settings	Closes the settings screen.

4.9 Getting Help

The web interface provides usage tips and a link to the user guide.

4.9.1 Usage Tips

Position the cursor of your mouse over an element from the toolbar or navigation tree to view a brief description of the option in the information bar at the bottom of the browser window.

4.9.2 On-Line User Guide

Click on the Help button from the SEA toolbar to view an HTML version of the *System Event Analyzer User Guide*. The help opens in a new browser window.

4.10 Logging Off

It is not necessary to manually log off SEA. Once your connection to the Director is closed, SEA will automatically log off your profile after the log off time elapses. By default, the log off time is set to ten minutes, however, you can configure the time from the User Settings screen (see Section 4.8.1).

You can close your connection by exiting your browser or navigating to a web site outside of the SEA web interface. SEA continues to process requests and stores your data after you have closed your connection (as a result, SEA continues to consume memory resources). If you restore your connection to the Director before the log off time elapses, your data will remain intact. This enables you to browse other web sites without losing your SEA data. However, if the connection with the Director is closed, once the log off time elapses, all the data associated with your SEA session is lost and the memory resources used by SEA are released. Thus, if you return to SEA after the log off time has elapsed, you will not be able to view the results of processing.

For example, if your log off time is set to 120 minutes and you start processing a log file before switching to another web site, you have two hours during which SEA will continue to process the log file and maintain your data. If you return to SEA before the two hours elapses, processing will continue and all your data will be maintained. If you don’t return to SEA within the two hours, processing is terminated and your data will be lost as memory resources are cleaned up.

Note

If you set the log off time to zero, you will be logged off and lose your data if you click the refresh button in your browser or if you click a link that opens a page outside SEA.

The following list shows some possible log off times and the time frame they represent:

- 180 = 3 hours
- 1440 = 24 hours
- 10080 = 1 week
- 44640 = 31 days

Lost Connection

If your connection to the Director is lost for any reason, the message in Figure 4–40 appears in the toolbar.

Figure 4–40 Lost Connection Message

4.11 Service Obligation

You can view service obligation information by entering the following URL:

`http://`*`hostname`*`:7902/obligation`

Where *hostname* refers to the machine name or IP address.

An example of the service obligation information is shown here:

```
Service Obligation: Valid
Service Obligation Number: NI93202975
System Serial Number: NI93202975
Service Provider Company Name: Hewlett-Packard

Obligation Start Date: Sat May 13 00:00:00 MDT 2000
Obligation Ending Date: Sun May 13 00:00:00 MDT 2001
Time left on Obligation: 0 years, 355 days, 13 hours, 52 minutes, 57 seconds

History of changes:
  1. Sat May 13 15:46:22 MDT 2000: Installer (unknown) of Hewlett-Packard
Installation settings changed to start Sat May 13 00:00:00 MDT 2000 to Sun May
13 00:00:00 MDT 2001 (1 years, 0 days, 0 hours, 0 minutes, 0 seconds)
  2. Sat May 13 15:46:11 MDT 2000: WEBES (Web-based Enterprise Services
Common Components V3.0 (Build 12), member of WEBES V3.0 (Build 12)) of
Hewlett-Packard
Set initial obligation: 0 years, 5 days, 0 hours, 0 minutes, 0 seconds ending
Thu May 18 15:46:10 MDT 2000
Notifications to be sent
  1. 0 years, 60 days, 0 hours, 0 minutes, 0 seconds
  2. 0 years, 30 days, 0 hours, 0 minutes, 0 seconds
  3. 0 years, 15 days, 0 hours, 0 minutes, 0 seconds
  4. 0 years, 5 days, 0 hours, 0 minutes, 0 seconds
  5. 0 years, 4 days, 0 hours, 0 minutes, 0 seconds
  6. 0 years, 3 days, 0 hours, 0 minutes, 0 seconds
  7. 0 years, 2 days, 0 hours, 0 minutes, 0 seconds
  8. 0 years, 1 days, 0 hours, 0 minutes, 0 seconds
```

4.12 Disabling the Web Service

The following procedure describes how to turn off the SEA web service. If WEBES is installed on a cluster, you will need to repeat the procedure for every node where SEA is installed.

1. Stop the Director (see Section 1.8).
2. Edit the `ConfigDefaultsCA*.txt` file in the config directory.
   - Tru64 UNIX:

     `/usr/opt/hp/svctools/specific/desta/config/ConfigDefaultsCADUnix.txt`
   - OpenVMS:

     `svctools_home:[specific.desta.config]ConfigDefaultsCAOpenVMS.txt`
   - Windows:

     `c:\Program Files\hp\svctools\specific\desta\config\ConfigDefaultsCAWindows.txt`
3. Put a `#` in front of the line `com.compaq.svctools.ca.services.web.SEAWebService`.

   The contents of the file should look similar to this:

```
# ConfigDefaultsCAWindows.txt
#
# SEA Default Components, ** Windows Version **
#
# Default components of SEA, to enroll the first time the
# DESTA Director process is executed, as fully qualified Java class names.
# After DESTA runs the first time, the file Configuration.dat will be
# created, and it will be read on startup instead of ConfigDefaults*.txt.
#
com.compaq.svctools.ca.services.analysis.EvtAnalyzer
#com.compaq.svctools.ca.services.web.SEAWebService
com.compaq.svctools.ca.services.eventreaders.EvtMonitor
#
# Uncomment the next line if operation of the Unanalyzed Event Logging
service is desired
#com.compaq.svctools.ca.services.analysis.UnanalyzedEventLogger
```

4. Delete the `configuration.dat` file from the following directories (assuming you used the default install directory):
   - Tru64 UNIX – `/usr/opt/hp/svctools/specific/desta/config`
   - HP-UX – `/opt/hp/svctools/specific/desta/config`
   - Linux – `/usr/opt/hp/svctools/specific/desta/config`
   - OpenVMS – `svctools_home:[specific.desta.config]`
   - Windows – `C:\Program Files\hp\svctools\specific\desta\config`
5. Restart the Director (see Section 1.7).

5

Translation, Analysis, and Summary

This chapter describes event translation and explains how to view and interpret translation information. It also describes log file analysis, including automatic and manual analysis and how to view and interpret analysis information. Procedures for simulating automatic analysis are described as well. Exceptions that impact the results produced by summary operations are also detailed.

5.1 Translation, Analysis and Rules

The results produced by translation and analysis are dependent on rule sets. The rules sets are developed by Serviceability Engineers and registered with SEA. These rule sets determine what problem reports will be generated in response to the contents of a log file and determine what translated data is presented in SEA.

For more information on rule sets, refer to Chapter 6.

5.2 Manual Translation

SEA can translate the events in a binary event log and send the results to your computer. This activity is known as manual translation.

On supported platforms, SEA can read and translate error logs produced by any of the supported operating systems. For example, you can use the web interface running on your PC to connect to a Director running on a Tru64 UNIX machine to read, translate, and analyze an event file produced previously on an OpenVMS machine.

5.2.1 Translating Events

Translation information is available from the command line interface and the web interface. Refer to the following chapters for information on translating events:

- CLI – Chapter 3
- Web Interface – Chapter 4

5.2.2 Translation Defaults

By default some events are not processed. Under normal operation, correctable events are not translated. The events that are usually filtered include:

- Correctable System events (entry types 620 and 630)
- Correctable Error Throttling Notification events
- Miscellaneous events not used by analysis, such as:
    - Time Stamp events
    - Volume Mount/Dismount events
    - Cold Start (System Boot) and Shutdown events
    - Software-related events

5.2.3 Translation Report Type

When you translate an event, you can choose between brief and full output. The content differences between full and brief output are defined in the rule sets. Brief output generally only contains the most important data items from the event while full output generally includes most of the data items from the event. Since the exact contents of each report type are defined by the rules used to generate the report, the type of information contained in brief and full reports may vary for different events.

5.2.4 Interpreting Translation Information

Note

Translated events include a timestamp. For information on interpreting this information see Section 5.6.

A translated binary event consists of three layers of information: overall, frame, and field.

5.2.4.1 Overall

The overall binary event contains one or more translated frames of information. There are several types of binary events, each identified by its class name. In addition to the frames, some other information is stored at the overall layer, such as:

- The class name of the binary event (passed to Event Analysis but not displayed in the translated output in the CLI or web interface)
- The “match keys” for the event, a set of strings used in identifying analysis rules that may fire for this event (not displayed in the translated output in the CLI or web interface)

5.2.4.2 Frame

A frame within an event consists of one or more translated fields of information. There are many types of frames, each identified by its label. Each frame type contains a defined set of fields. In addition to the fields, some other information is stored at the frame layer, such as:

- The parent binary event of this frame
- The frame’s label, displayed at the beginning of each frame

5.2.4.3 Field

A field within a frame consists of the following:

- The parent frame of this field
- The field's label, both as an identifier (not shown) and as displayable text
- The field's value (of a type defined by the type of field) which is displayed in text form

5.2.4.4 Typical Frame of a Translated Binary Event

A typical frame of a translated binary event appears as follows:

```
HPM System Event Frame Subpacket - Version X
HPM_Elapsed_Time_Since_Srm_Boot 947                  Seconds Since Last
                                                     Console Boot
HPM_Event_Info_Block_1          x0040 AB81 0F0F 0010 H-Switch System Event
                                                     Information
   HPM_System_Event_Code[7:0]  x10                  HS Temperature in
                                                     Yellow Zone
   HPM_Supplementary_Code[15:8] x0                   Supplementary Code
   Gp0_Valid[16]                x1
   Gp1_Valid[17]                x1
   Gp2_Valid[18]                x1
   Gp3_Valid[19]                x1
   Hs_P0_Valid[24]              x1
   Hs_P1_Valid[25]              x1
   Hs_P2_Valid[26]              x1
   Hs_P3_Valid[27]              x1
   Csb_Master_Ena[32]           x1
   3_3_Dcok_2[42]              x0                    0 = NOT OK if
                                                     Regulator 2 is
                                                      Installed
   2_5_Dcok_2[44]              x0                    0 = NOT OK if
                                                     Regulator 2 is
                                                      Installed
   Pll_Dcok_2[46]              x0                    0 = NOT OK if
                                                     Regulator 2 is
                                                      Installed
   Csb_Address[55:48]           x40
```

This frame contains 17 fields. Each field has a single value, such as 947 (decimal) or x10 (hexadecimal, 16 decimal). Some fields are represented as both a Register (HPM_Event_Info_Block_1) containing the complete hexadecimal value, and again as a series of subfields such as HPM_System_Event_Code[7:0]. The [7:0] indicates that bits 0 through 7 of this register comprise this subfield, bit 0 being the least significant bit.

5.2.4.5 Unsupported Entries

Some of the events logged by a system or device are not used by SEA to diagnose hardware failures. The CLI translate command and the event listing in the web interface translate events with many different entry types, including some not used for analysis. However, there are some cases where SEA cannot translate an event:

- If the event type is not supported.

- If the system or device logged incorrect data for a supported entry type, causing it to be unrecognized.

If an event that is not supported or recognized is encountered during translation, an unsupported entry dump is shown in the output. The unsupported entry dump at the end of the event shows the entire event in hexadecimal format, from the first header byte to the last byte of the event.

Note

Each subsequent release of SEA supports the translation of new event types and incorporates better handling of incorrect input data. Events that currently result in a unsupported entry dump may be correctly translated in a future release.

The following example shows the translated output for an event that was logged incorrectly. The event should have been logged with major class 250 and minor class 0, which SEA would have correctly translated. However, the minor class was 18 and the event was unrecognized. As a result, an unsupported entry dump was generated.

```
Event: Unknown Combined Entry Type - UNSUPPORTED ENTRY - Major_Class: 250
Minor_Class: 18 occurred at Mon, 13 Aug 2001 18:30:36 +0200

COMMON EVENT HEADER (CEH) V2.0
OS_Type 1 -- Tru64 UNIX
Hardware_Arch 4 -- Alpha

CEH_Vendor_ID 3,564 -- Hewlett-Packard Company

Hdwr_Sys_Type 35 -- GS40/80/160/320 Series

Logging_CPU 0 -- CPU Logging this Event
CPUs_In_Active_Set 24
-- Unknown Combined Entry Type -
Entry_Type 18,250 UNSUPPORTED ENTRY -
Major_Class: 250
Minor_Class: 18

DSR_Msg_Num 1,969 -- AlphaServer GS320
Chip_Type 11 -- EV67 - 21264A
CEH_Device 255
CEH_Device_ID_0 x0000 0000
CEH_Device_ID_1 x0000 0000
CEH_Device_ID_2 x0000 0000
Unique_ID_Count 3
Unique_ID_Prefix 11,248

TLV Section of CEH
TLV_DSR_String AlphaServer GS320 6/731
TLV_OS_Version Tru64 UNIX V5.1 (Rev. 732)
TLV_Sys_Serial_Num QBB7.AJK01
TLV_Time_as_Local Mon, 13 Aug 2001 18:30:36 +0200
TLV_Computer_Name abcd101

com.compaq.svctools.desta.services.decomposers.DecompDataException:
EXCEPTION: Entry_Type_Support.java, DUNIX_Entry_Type(), No support for this
DUNIX Entry Type... Major_Class is: 250 Minor_Class is: 18
0000: FE FF FF FF 0C 01 00 00 ?yyy....
0008: 48 01 00 00 02 00 00 00 H.......
```

```
0010: 01 00 04 00 EC 0D 00 00 ....i...
0018: 23 00 00 00 00 00 00 00 ........
0020: 00 00 00 00 18 00 00 00 ........
0028: FA 00 12 00 B1 07 00 00 u.......
0030: FF 00 05 00 02 18 00 00 y.......
0038: 01 00 00 00 0B 00 00 00 ........
0040: 00 00 00 00 00 00 00 00 ........
0048: 00 00 00 00 03 00 F0 2B ......?.
0050: 00 00 00 00 00 00 00 00 ........
0058: 00 00 00 00 00 00 00 00 ........
0060: 00 00 00 00 00 00 00 00 ........
0068: 00 00 00 00 00 00 00 00 ........
0070: 00 00 00 00 00 00 00 00 ........
0078: 05 00 00 00 61 00 1F 00 ....a...
0080: 43 6F 6D 70 61 71 20 41 Compaq.A
0088: 6C 70 68 61 53 65 72 76 lphaServ
0090: 65 72 20 47 53 33 32 30 er.GS320
0098: 20 36 2F 37 33 31 00 00 .6.731..
00a0: 81 00 22 00 43 6F 6D 70 ....Comp
00a8: 61 71 20 54 72 75 36 34 aq.Tru64
00b0: 20 55 4E 49 58 20 56 35 .UNIX.V5
00b8: 2E 31 20 28 52 65 76 2E .1..Rev.
00c0: 20 37 33 32 29 00 00 00 .732....
00c8: C1 00 0B 00 51 42 42 37 A...QBB7
00d0: 2E 49 4F 52 30 31 00 00 .AJK01..
00d8: 41 00 18 00 32 30 30 31 A...2001
00e0: 30 38 31 33 31 38 33 30 08131830
00e8: 33 36 2C 30 30 30 30 32 36.00002
00f0: 30 30 00 00 21 01 14 00 00......
00f8: 68 73 31 31 30 31 61 00 abcd101.
0100: 00 00 00 00 00 00 00 00 ........
0108: 00 00 00 00 FA 00 00 00 ....u...
0110: 20 00 00 00 6D 63 68 61 ....mcha
0118: 6E 31 3A 20 20 6E 6F 64 n1...nod
0120: 65 20 32 20 68 61 73 20 e.2.has.
0128: 63 6F 6D 65 20 6F 6E 6C come.onl
0130: 69 6E 65 0A 00 00 00 00 ine.....
0138: 00 00 00 00 E8 00 00 00 ....e...
0140: 48 01 00 00 25 7E 3C 5E H.......
```

5.3 Automatic Analysis

Automatic analysis is the immediate analysis of an event that has been captured and decomposed by SEA as soon as the event is generated by the system (or shortly thereafter), regardless of any interfaces that may be running. No user intervention is required. Automatic analysis is always enabled while the Director is running. The Director is always running unless it is manually stopped or, during installation, you chose not to start the Director when the system is rebooted (Tru64 UNIX, HP-UX, Linux, or OpenVMS systems).

Problem reports resulting from automatic analysis are sent to all interfaces and to all recipients that are set up to be notified.

See Chapter 8 for information about setting up notification services.

5.3.1 Scavenge

Automatic analysis processes events as they occur. However, when the Director is stopped, SEA indicates the last event from the binary log file that was processed in the analysis database. When the system is restarted, SEA checks the database to see which events have been processed and processes all the events that occurred after that point. This operation is referred to as scavenging. The scavenge operation finds events that are still pending processing and ensures that no events are missed, even when the system is restarted. The first time scavenge occurs, it processes the entire event log. Once this is complete, new events are processed as they occur. The scavenge operation occurs four minutes after the Director is started. If the Director is started and stopped within four minutes, no scavenge occurs.

Initially, the entire system event log is read to find any events that can be analyzed. A filter is then applied to the analyzable events. All analyzable events that occurred within a week of the current time are processed.

If there are no analyzable events, the scavenge feature becomes dormant and a marker representing an unsupported system is stored in the automatic analysis database. As long as the unsupported system marker is present on the system, no scavenging occurs. If there is at least one recognized event, scavenging occurs every time the Director is stopped and started

Scavenging and the Web Interface

If you connect to the Web Interface before scavenging begins, events that arrive while the Web Interface is running will appear in the Real-Time Monitoring view. All the events that arrive before scavenging starts are processed once scavenging begins and any problem reports that result from scavenging also appear in the Real-Time Monitoring view. However, any events that were added to the event log before the Web Interface was started will not appear in the Real-Time Monitoring view.

5.3.2 Reset

Note

Resetting the automatic analysis results may significantly impact the results of future analysis.

You can use the command line interface to reset the automatic analysis database. Resetting the database erases the current callouts and any analysis data stored. Only the following information is retained:

- Configuration Information—specifies the computer hardware present.
- Scavenging Information—indicates the last event read from the system's binary event log.

To reset the database, stop the Director (see Section 1.8 for more information) and then use the reset command (only available in the new common syntax):

```
wsea reset
```

Once you have reset the database, restart the Director (see Section 1.7 for more information).

Impact of Resetting the Analysis Database

Be aware that resetting the analysis database clears both the active problem reports and all the storage units maintained by SEA. Storage units are records of past events that are used by some rules for thresholding and multiple event analysis. Hence, after a reset, analysis results may be significantly different than they otherwise would have been in the absence of a reset.

For example, if SEA had an accumulation of storage units that would count toward satisfaction of a threshold filter, resetting the database would erase these units. As a result, one or more problem reports that depend on thresholding filters could be delayed or completely suppressed. This type of issue most typically arises with correctable events. Whereas uncorrectable faults generally are reported immediately, correctable events such as intermittent disk read errors may be subject to threshold filtering. After a certain number of correctable events within a specified time-frame occur, a problem report can be generated to signal that a device is suspect, regardless the fact that it continues to function normally.

The best way to minimize the impact is to only use the reset command after carefully reviewing the recent events (reviewing the events from the last 24 hours is recommended as a minimum). During the review, look for recurring events, typically correctable errors, that involve any device that has not been called out in the problem reports. These events should help determine if any other devices on the system are suspect.

5.3.3 Disable

If necessary, automatic analysis can be disabled from the CLI as described in Chapter 3. You may want to disable automatic analysis if SEA is running on a platform such as HP-UX or Linux, where the native error log is not currently analyzed.

5.4 Manual Analysis

You can open a binary event log file and request that the events be translated and analyzed. This activity is known as manual analysis. Unlike automatic analysis, manual analysis relies on the time stamp information included with each event to determine when an event occurred.

Manual analysis can be performed from all the interfaces. Refer to the following chapters for information on manual analysis:

- CLI – Chapter 3
- Web Interface – Chapter 4

Regardless of the platform where it is installed, SEA can read and analyze binary event logs produced by any of the supported operating systems.

5.4.1 Resource Usage During Analysis

Whenever SEA starts, and when you run manual analysis, the program appears to use a lot of system resources and processor cycles. However, SEA uses only the capacity that is not being asked for by other programs.

SEA always relinquishes processor cycles to other programs whenever they need them. In other words, the program uses whatever resources are available.

At startup SEA needs the available capacity for the scavenge process. Depending on the system, and the size and content of the log, the initial startup pass can take many minutes or even hours to complete. After completing the scavenge process, SEA drops into idle mode, where resource usage hovers at only a few percent.

If you run SEA in manual mode, large amounts of system resources and processor cycles might also get used. As in the case of startup in automatic mode, the condition is directly related to the size and content of the log being processed. Once again by design, SEA uses as many resources as are available until processing is completed.

You can speed processing by managing the system error log so that it does not grow indefinitely. One way to accomplish this is to periodically archive and reset the current error log by following the guidelines in the *WEBES Installation Guide*. When you are using manual analysis, it may be beneficial to filter large log files in order to improve processing times.

5.5 Interpreting Analysis Information

Note

Problem reports generated by analysis include a timestamp. For information on interpreting this information see Section 5.6.

A report consists of a set of String and Value Pairs (SVP). A SVP can be short, for example:

```
Severity:
    2
```

A SVP also can be extensive, such as the Full Description or Evidence SVPs, which can contain many lines of information (see Appendix A for an output example). A problem report resulting from event analysis typically contains the following Strings, with Values describing the analysis results.

5.5.1 Problem Report Times

The Problem Report Times designator indicates the time when SEA generated the Problem Report, and is unrelated to the time of the event or events that caused the problem report.

5.5.2 Managed Entity

The Managed Entity designator provides service information regarding the system on which the problem was found. This includes the system host name (typically the computer name for networking purposes), and the type of computer system.

5.5.3 Service Obligation

The Service Obligation designator provides information about the service provider and the state of the service contract.

5.5.4 Brief Description

The Brief Description designator provides a high level description of the event. This typically includes whether the error event is related to the CPU, the system (PCI or Storage, for example), or the environmental subsystem within this managed entity.

5.5.5 Callout ID

The Callout ID designator provides information about the analysis rule set. Most characters within this designator are used for HP-specific reserved purposes.

5.5.6 Severity

The Severity designator provides the service relevance of the occurrence of the problem found. The current severity hierarchy is shown in Table 5–1.

Table 5–1 Problem Severity Levels

Severity Level	Service Relevance	Comments
1	Critical	This level is not currently used due to system operation required for SEA diagnosis.
2	Major	Fatal event that typically requires service if not already administered.
3	Minor	Non-Fatal or Redundant warning event that typically requires future service but system still operates normally.
4	Information	System service event such as enclosure PCI or Fan door is open and only requires system door closure.
5	Unknown	This level is not used currently.

5.5.7 Reporting Node

The Reporting Node designator is the node from which the error was reported. It is synonymous with the Managed Entity host name when SEA is used for system diagnosis for the system on which it is running. For future implementations, this may reflect a system server reporting about a client for which SEA is performing diagnosis within an enterprise computing environment.

5.5.8 Full Description

The full description designator provides detailed error information about the event. This can include the detected fault or error condition description, specific address or data bit where this fault or error occurred, and other service related information.

5.5.9 FRU List

The Field Replaceable Units (FRU) List designator lists the most probable defective FRUs. This list indicates that qualified service needs to be administered to one or more of these FRUs. This information typically provides the FRU probability, manufacturer, system device type, system physical location, part number, serial number, and firmware revision level (if applicable to the FRU).

5.5.10 Evidence

The Evidence designator provides the error event information that triggered the indictment. The evidence shown depends on the system that generated the error log and the registered rules. As a result the contents of the evidence field may vary.

Typically, the evidence includes the following:

- The time stamp of the event responsible for the callout.
- The event identifier, which is displayed differently depending on the responsible rule set. (In some cases, the event identifier uses new common event header Unique_ID_Prefix and Unique_ID_Count components. Where the Unique_ID_Prefix refers to an OS-specific identification for this event type and the Unique_ID_Count indicates the number of this event type that occurred.)
- The ruleset name and revision number may be included depending on the rule set.

5.5.11 Versions

The SEA Version and WCC Version designators provide the versions of SEA and WEBES that created the problem report.

5.6 Interpreting Time Stamps

If an event in a binary log includes a Storage Event Header (SEH) or Common Event Header (CEH), that information is used to provide the time stamp information for analysis and translation results. If the event only includes a Windows NT header, no time stamp is included with analysis results.

In addition, when you translate an event that includes a SEH or CEH header in addition to a Windows NT header, both time stamps are shown in the translation results. However, unless the machine responsible for logging the event is located in the GMT time zone, the time stamps will be different.

The event time is also displayed in the event description (located at the top of a translated event). Depending on the contents of the event and the SEA interface used to translate it, the translated output may include different information:

- If the event includes a SEH or CEH header, the time stamp information from that header is included in the event description. If the header has invalid date information the current date is shown along with an error message.
- If you are using the web interface and the event only has a Windows header, no date information is shown in the event description.

- If you are using the CLI to send the translation to the screen or a text file and the event only has a Windows header, the date information from the header is included in the event description.
- If you are using the CLI to send the translation to a HTML file and the event only has a Windows header, no date information is shown in the event description.

SEH and CEH Headers

SEH and CEH time stamps are stored as strings and reported in the `TLV_Time_as_Local` field of a translated event. This field has the following format:

```
Jan 11, 2002 3:06:09 AM GMT-0600
```

This indicates the time the event was logged, in the time zone where the machine responsible for logging the event is located. The time zone is shown as an offset, in hours, from GMT.

Windows Headers

The Windows NT header stores time stamp information as an integer indicating the number of seconds that have elapsed since epoch (January, 1 1970 00:00:00 AM GMT). These integers are translated into a date and time and reported in the `WNT_GMT_Time_Generated` and `WNT_GMT_Time_Written` fields of a translated event using the following format:

```
Jan 11, 2002 9:06:09 AM GMT
```

Since the Windows NT header does not include any information about the time zone where the logging machine is located, the GMT time zone is used. This does not mean the logging machine is located in the GMT time zone.

5.7 Simulation of Automatic Analysis

SEA can simulate the occurrence of events and their automatic analysis. The events are translated and analyzed as if they occurred on the local system and events and problem reports from analysis appear as automatic events do. Using the simulation, you can perform and end-to-end test of SEA.

Note

Problem reports created by simulated automatic analysis are identified as test callouts so that no action is taken by the customer service center. Translation results also indicate that the output was generated by the `test` command.

5.7.1 Sending A Test Event To The System Error Log

Use the following command to test SEA, from event detection to analysis and notification:

```
wsea test
```

This command sends an event with header fields but no further content to the system's error logging API. The action taken with this event is dependent on the system:

Tru64 UNIX, HP-UX, Linux, and OpenVMS

If the command was run on a supported platform, the system's error logging service takes the event content and wraps it with a Common Event Header (CEH). This is necessary because SEA only recognizes events with a CEH or a Storage Event Header (SEH). After the CEH is created and all its fields are populated, the event is written to the error log where it can be processed by automatic analysis, generate a problem report, and trigger notification.

Note

The event generated by the test command will be logged with a CEH on the following operating systems and platforms:
Tru64 UNIX 4.0E and above on all EV6 and above platforms
OpenVMS 7.1-2 and 7.2 and above on all platforms

Windows

The error logging service on Windows does not wrap event content with a CEH since that is usually done by the device drivers themselves. So, like a device driver, the `test` command creates a mock CEH which is used as the event content and passed to the system error logging API. The command does not provide values for all the fields in the mock CEH. Only the fields critical to translation, analysis, and human identification (including time, computer name, OS type and event ID) are given valid values. Most other fields are set to 0 or NULL values and do not affect translation or analysis. After Windows receives the event, it adds a NT header and the event is appended to the system error log. Once in the error log the event is processed by automatic analysis, generates a problem report, and triggers notification.

5.7.2 Bypassing The System Error Log

Use the following command to test SEA without sending an event through the system error log:

```
wsea test nosystem
```

The Director must be running in order to use the `test nosystem` command.

The `nosystem` option sends an event directly to the SEA event reader, bypassing the system altogether. This command is used to facilitate troubleshooting of a problem and determine if it is caused by SEA.

Regardless of the platform, the command creates a mock CEH for the event so that it can be recognized. Since SEA also requires an NT event header when running on Windows platforms, a mock NT header is also created when the command is executed on an Windows machine. Only the NT header fields necessary for translation, analysis, and human identification are populated with valid values. Fields set to 0 or NULL do not affect translation or analysis.

Since the event created by the `nosystem` option has a CEH (and for Windows, a NT header as well), it should always be recognized by SEA. However, since the event is never appended to the system error log, it cannot be seen when manually translating or analyzing the system error log. In addition, the problem report immediately expires and, as a result, it will not appear if you subsequently run the `wsea report` command. The only ways to view the problem report generated by analysis is by using the "Real Time Monitoring" view in the web interface, or the problem report logging functionality (see Section 3.5.2.2). The `wsea report` command will not show the problem report because it is designed to expire immediately.

Note

The `nosystem` option creates an event that can be translated and analyzed for all the supported operating systems, regardless of whether or not the hardware platform is supported.

5.8 Interpreting Summary Information

If a log file contains invalid data or lacks a recognizable (CEH or SEH) header, the results produced by the summary command will be affected.

- If the final event in a log file contains invalid data, SEA cannot determine the date information for the Last Entry Time field. In this case, the current date and time are shown in the Last Entry Time field.
- If an event does not include a recognized header, the event type is reported as `0`. In this case the summary command indicates that the event is `Unrecognized/Unsupported`. This applies to events that only contain a Windows header even if they are translated correctly.

6

Rule Sets

This chapter describes the rule sets and instance files used by SEA. Information on managing rule sets is also given.

6.1 Rule Sets

Binary events are analyzed by using an analysis engine to apply rules to them. Rules are designed to fire when a particular criteria, such as a threshold, is met. For example, if the number of events within a given time frame exceeds the threshold specified in a rule set, the rule fires.

Depending on the circumstances, a event may or may not fire any rules. Alternately, a single event can fire multiple rules. When a rule fires, it may or may not produce reports. In the case where reports are generated, a rule can create one or multiple reports. A report may be generated immediately, or may be generated after a gestation time period defined by the rule. Each report is stored in a instance file. After the report's expiration time period, as defined by the rules, the report is removed from the instance file.

Rules are also responsible for determining the output presented for a translated event.

Analysis rules are coded by Hewlett-Packard serviceability engineers or other domain knowledge specialists. These rule sets are stored in jar files located in the `svctools\common\jars` directory. Rule sets pertaining to the supported platforms are located in the jar files and can be installed, or "registered," for use with SEA. A rule set can later be "unregistered" if it is no longer applicable.

Note

It is possible to run SEA without any rule sets registered (if the rule sets have been unregistered or deleted). However, if there are no registered rule sets, analysis will not generate meaningful results. The problem report generated by analysis indicates if there are no registered rule sets or no applicable rule sets.

6.2 Analysis Data

SEA stores analysis data in the `svctools\common\ca\data` directory. These files contain information about:

- The rule set files to be used for analysis
- Input entry classes, derived from data in the binary events. Typically, the input classes are deleted after reports have been generated from them.
- Intermediate data such as complex storage classes, derived during analysis
- Output report classes (analysis results)

You can clear this state data using the `wsea reset` command described in Chapter 5.

6.3 Managing Rule Sets

SEA is installed with all rule sets pre-registered. You can manipulate the rule sets in the following ways:

- View the rule sets that are currently registered (see Section 6.3.1).
- If you receive or create new analysis rule files, you can register the new rule sets as needed (see Section 6.3.2).
- Unregister rule sets that are no longer needed (see Section 6.3.2).
- Re-register all the default rule sets (see Section 6.3.2).

Note

This section describes how to manage rule sets using the new common syntax. For the equivalent old common syntax commands, refer to Appendix E.

6.3.1 Viewing Registered Rules

Using the CLI or web interface, you can view the rulesets that are registered for use with SEA.

6.3.1.1 CLI

The new common syntax `lis` command provides a list of the paths and versions of the knowledge files registered with DeCOR. The syntax for the command is shown here:

```
wsea lis
```

Output

An example of the output is shown here:

```
      Ruleset                    Version
==================== ==================================
CATEST               Rules_v1_1
DS10                 Rev_030509
DS20                 Rev_030320
DS25                 Rev_030509
ES40                 Rev_030512
ES45                 Rev_030509
GS1280_EV7           V4_2
GS1280_IO7           V4_2
GS1280_RBOX          V4_2
GS1280_SM            V4_2
GS1280_ZBOX          V4_2
GS320_CE             V53_0953
GS320_SE             V53_0953
GS320_STARTUP        V53_0953
GS320_UCE            V53_0953
```

```
MCII                Rev_1
Storage             Rev_2.20
Storage_HSV_DRM     Rev_X1_00
Storage_HSV_EMU     Rev_X1_00
Storage_HSV_EXEC    Rev_X1_00
Storage_HSV_FCS     Rev_X1_00
Storage_HSV_FM      Rev_X1_00
Storage_HSV_SCMI    Rev_X1_00
TS202c              Rev_4_1_A0
Vstor               Rev_1.00
```

6.3.1.2 Web Interface

From the web interface:

1. Click the Settings button from the toolbar.
2. Select the Director Settings tab.
3. Click the Register Knowledge button in the navigation frame.

All the available rule sets are listed with a check box. Rule sets with a selected check box are registered.

6.3.2 Registering and Unregistering Rule Sets

You can register a set of rules using the CLI or the web interface.

6.3.2.1 CLI

The syntax for registering and unregistering rule sets is shown here (the first command shown is used to register rule sets and the second command is used to unregister rule sets).

Using the new common syntax:

```
wsea reg [ruleSet]
wsea unr [ruleSet]
```

Where *ruleSet* represents the name or names of the desired knowledge files. If you do not enter any rule set names, all the default rule sets are registered.

Wildcards can be used to specify multiple filenames, as shown in the following examples:

```
wsea reg ds*
wsea unr ds*
```

Note

If you are running analysis in the CLI, you will see the changes take effect immediately. However, to run analysis in the web interface, you must restart the Director first (see Sections 1.8 and 1.7).

6.3.2.2 Web Interface

To register or unregister a set of rules using the web interface, do the following:

1. Click on the Settings button in the toolbar.
2. Click the Director Settings tab at the bottom of the window.
3. Click the Register Knowledge button in the navigation frame (Figure 6–1).

Figure 6–1 Rules Files

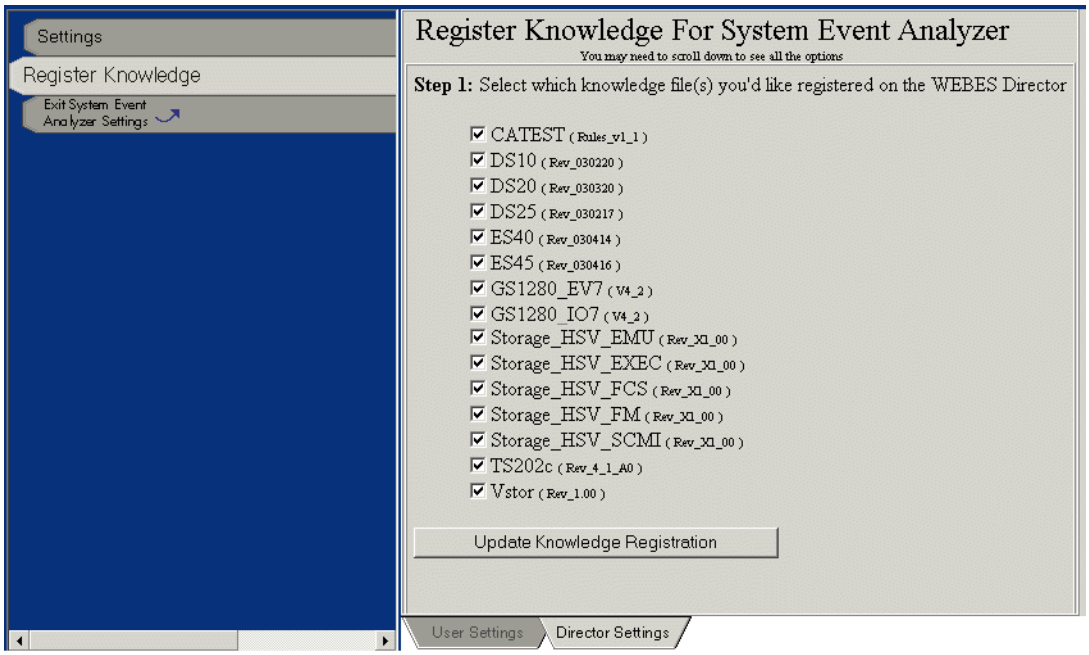

All the available rule sets are listed with a check box. If the check box is selected the rule set is registered, otherwise it is not registered.

4. Register or unregister the necessary rule sets.
   - To register a rule set that is not registered, select the check box next to its name.
   - To unregister a rule set that is currently registered, deselect the check box next to its name.
5. Click the Update Knowledge Registration button to save your changes.

Note

Changes will not take effect in the web interface for automatic analysis until the analyzer is restarted. This is done by stopping and restarting the Director. These changes will not affect manual analysis jobs already in progress.

6. Stop and Restart the Director to apply the changes (see Sections 1.8 and 1.7).

7

Configuration

This chapter describes configuration, including getting and changing the configuration, global and component configuration attributes, and creating and resetting the configuration.

7.1 Viewing the Configuration

You can view the configuration settings for your local Director from the web interface.

To view the configuration, use the following procedure:

1. Select the Settings button from the toolbar.
2. Select the Director Settings tab.

   The Director Settings window is shown in the display frame (see Figure 7–1). By default the Settings button is selected.

Figure 7–1 Settings

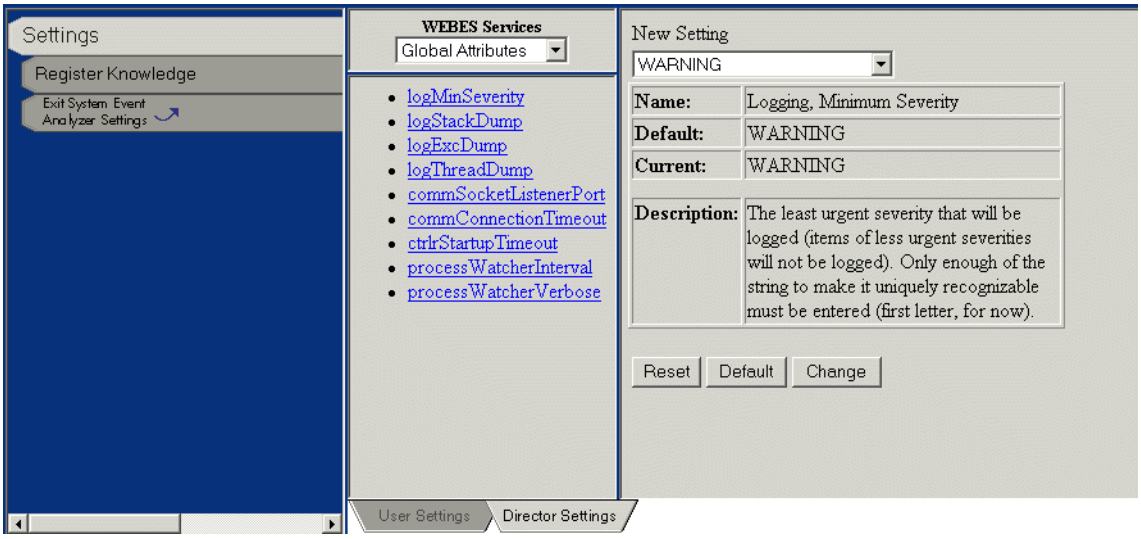

3. Select the service whose attributes you want to view from the drop-down list.

   By default, Global Attributes are shown; however, the drop-down list contains all the services currently enrolled in the system. The SEAWebService was selected in Figure 7–2.

Figure 7–2 Attribute Display

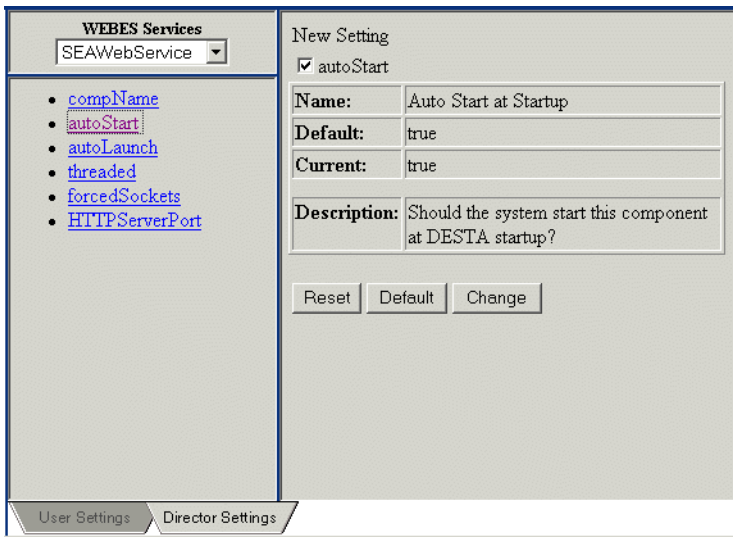

4. To view the current value of an attribute, click on its name on the left side of the window (see Figure 7–2).

   The attribute's full name and current and default values, are displayed on the right side of the window along with a description of the attribute.

   The automatic start attribute (autoStart) was selected in this example.

7.2 Component Configuration Attributes

Attributes for all components fall into two categories (indistinguishable in the web interface): common attributes and extended attributes.

Common Attributes

Attributes that each component contains by default are known as common attributes. They are still owned by their component, so the autoStart attribute for one component is independent from the autoStart attribute of another component.

Extended Attributes

Attributes specific to a particular component are known as extended attributes. For example, the "HTTPServerPort" attribute of the "SEAWebService" component does not exist in any other components, since it only applies to the web service.

7.3 Changing the Configuration

You can modify the attribute configuration settings from the web interface or make limited changes from the CLI.

Normally, it is not necessary to change the attribute settings. The following list describes the attributes that most often need changed and the location of the attribute in the web interface.

- commSocketListenerPort (Communications, Socket Listener Port Number) – under Global Attributes. Used to change the communications port number. Do not change the commSocket ListenerPort attribute from the web interface, see Section 7.4.2 for information on configuring ports.
  You may need to change the port number if there is another, conflicting application.
- commConnectionTimeout (Communications, Connection Handshake Timeout) – under Global Attributes. Used to change the amount of time that can elapse before the system times out.
  You may want to change the Timeout setting if your network is very slow and you want to allow more time for connections before timing out.
- autoMode (Automatic Mode) – under the EvtAnalyzer attribute. Used to enable or disable automatic processing of the binary system event log.
  You may want to change the autoMode setting if there are event entries for unsupported hardware in the event log.
- HTTPServerPort – under the SEAWebService attribute. Used to change the port used for http communications. See Section 7.4.2 for more information on configuring ports.
  You may need to change the port number if there is a usage conflict.

7.3.1 CLI

The CLI has limited configuration abilities.

Socket Ports

The socket ports can only be modified from the command line. Refer to Section 7.4.2 for details on changing the ports.

7.3.2 Web Interface

Using the web interface, you can change attributes from the Configuration Settings window (see Figure 7–1). Attributes that can be changed have a changeable field and three buttons in the System Configuration window. You must select an attribute to determine if it can be changed.

To change the value of an attribute, enter the new value in the New Setting field. Depending on the attribute that you want to change, you may be able to select the new attribute value from

a drop-down list or change a check-box setting. After changing attributes you have several choices.

- Click the Change button to apply the changes to the current attribute.
- Click the Reset button to change the values of the current attribute back to their last applied value.
- Click the Default button to change the values of the current attribute to their default values.

If you leave the Configuration Settings window without clicking the Change button, your modifications will be lost.

7.4 Global Configuration Attributes

The attributes listed under “Global Attributes” affect every component in the SEA system on the current machine, whether or not the component has been enrolled in the configuration.

7.4.1 Changing the Attributes

Changes to the Logging attributes (prefaced with “log”) take effect immediately.

Changes to the Communications and Controller attributes (prefaced with “comm” and “ctrlr,” respectively) take effect only when a new SEA process is started (such as the Director or another process that connects to the Director).

Be aware that changing a global configuration attribute affects both interfaces.

7.4.2 Changing Ports

Table 7–1 describes the ports used by SEA and indicates whether or not they can be configured.

Table 7–1 Ports

Port Number	Used For	Configurable
7901	Director-to-Director communications, and communicating with the Director on the local machine through the CLI.	Yes
7902	Director’s Web Interface listener port used by the web browser (e.g., http://machine.domain.com:7902)	Yes
7903	Communication between SEA’s applet (running inside the web browser) and the Director.	No

Table 7–1 Ports (continued)

Port Number	Used For	Configurable
7904	EVM connection to the Director. (Although EVM is a UNIX tool, the Director listens to this socket on all operating systems.)	No
7920	The WEBES WCCProxy process communicates with the Director on this port.	No.
1998	Service Cockpit	No
2069/8941	CSG/QSAP—the port number for CSG v4.5 and v5.0 is 2069. For v3.1 and v3.1B it is 8941. (See Section 8.4.2 for more details on CSG/QSAP.)	Yes
25	SMTP mail. This is the standard port used by TCP/IP systems for SMTP (see Section 8.2 for more details on configuring SMTP).	No

If a port is configurable, you can change the port number used. Most ports are configured using the web interface, however, the commSocketListenerPort, which is used for connections to the Director, can only be modified from the CLI.

Connections to the Director

The commSocketListenerPort defines the TCP/IP socket port used by the Director to communicate with other processes on the same machine or on other machines on the network (Port 7901, by default).

Note

Do not change the commSocketListenerPort attribute with the web interface. If you do, the Director cannot be stopped from that point on. After the socket port is changed, only a service that is already connected can stop the Director running on the old port.

To change the TCP/IP socket port attribute on all operating systems use the following command from the command prompt.

```
desta msg -chgport nnn
```

Where *nnn* is the new port number

This command changes the port number and then stops the Director and all connected processes. After the Director has finished shutting down, you can safely restart it on the new port.

Note

If the process hangs unexpectedly under Windows, kill the command and stop the Director manually. Press CTRL-C to exit the CLI command, and then enter **`net stop desta_service`**.

The Director can only communicate with Directors on other machines that have the same TCP/IP socket port number defined in their configuration. You can restrict access to your Director by changing the ports to nonstandard numbers and only disclosing the new port numbers to people who need access.

7.5 Profiles

When you are using the web interface, your changes to the configuration are saved in a profile. The profile for the current session is saved using the login name you entered (see Section 4.2). To restore your previous configuration settings when you restart the web interface, simply enter the same login name.

Your profile is saved on the machine where you logged on. If you logon to a different machine, then it will use the default settings. To customize the settings for the new machine, you will again need to create a new profile and change the configuration settings. This is true for each new machine you log onto.

Note

Profile names are case sensitive. Changing between upper case and lower case letters will create additional profiles. To access a profile, you must enter the profile name exactly as it was created.

7.6 Creating and Resetting the Configuration

The first time that SEA is started on a machine, a warning similar to the following is written to the Director log file. (See the *WEBES Installation Guide* and Section 1.12 of this guide for more information on log files.)

```
___.
WARNING on February 1, 2001 11:23:35 AM MST (0.023 sec elapsed)
      Configuration file /usr/opt/hp/svctools/desta/config/Configuration.dat
not found, creating it.
       Current Thread[main,5,main]
```

This warning is expected and correct. The `Configuration.dat` file is created based on the contents of the `ConfigDefaults*.txt` file in the `svctools/specific/desta/config`

directory. (The warning example shown is for a Tru64 UNIX system.) The classes named in those files will enroll themselves into the configuration, which is then saved as `Configuration.dat`, a binary file that should not be edited directly. Changes made from the web interface are saved in this file by the Director. This warning should not appear on subsequent starts of the Director.

If the configuration becomes damaged, or you wish to return to the default configuration state (the configuration when SEA was first started), make sure no SEA or WEBES processes are running (including the Director process), and delete the `Configuration.dat` file. When you restart SEA, the file will be recreated with the standard defaults, using `ConfigDefaults*.txt` the same way it was first time SEA was started.

7.7 Editing the Desta Registry

The Desta Registry contains information gathered about the user and the system during the installation process. Additionally, you can configure WEBES and SEA by making changes to the registry using the `desta dri` commands.

Note

In Windows, the WEBES registry is stored in the DESTA.REG file in the `svctools` installed directory tree, and should not be confused with the Windows Registry.

The `desta dri` commands allow you to add, view, edit, and remove registry keys.

Note

In OpenVMS, key names and parameters are always put in quotes in order to preserve mixed-case names and values. For example:
`desta dri get "KeyName"`

Adding a Registry Key

The `desta dri add` command creates the key within the registry. This command does not assign any values to the key, but you must create it before you can edit it. To add a key to the registry, enter the following:

```
desta dri add key_name
```

Viewing a Registry Key

The `desta dri get` command displays the current value assigned to a key. If the key returns a value of "null" (for example, `CA.WUI.OLMsgWait=null`) it does not exist, and you will need to add it before attempting to make any changes. To view a key, use the `get` command:

```
desta dri get key_name
```

Editing a Registry Key

The `desta dri set` command allows you to enter one or more values for an existing registry key. Multiple values can be assigned by entering a comma-separated list in quotation marks. To edit a key, use the `set` command:

```
desta dri set key_name parameter_value
```

When entering a comma-separated list:

```
desta dri set key_name "value1,value2,..."
```

Removing a Registry Key

The `desta dri del` command deletes all of the assigned values, and removes the key from the registry. To remove a key, use the `del` command:

```
desta dri del key_name
```

7.7.1 Configuring the Message Wait Timeout

The `CA.WUI.OLMsgWait` key allows you to set the message wait timeout value for the web interface. For example, you may be experiencing timeouts when loading the list of log files using the Other Logs link. By default, the value is 45 seconds. To reset the timeout to 90 seconds, add and set the key in the Desta Registry.

Windows, Tru64 UNIX, HP-UX, and Linux

1. Add the key to the registry if it does not already exist:
   `desta dri add CA.WUI.OLMsgWait`
2. Set the value of the key to 90 seconds:
   `desta dri set CA.WUI.OLMsgWait 90`
3. View the new value of the key:
   `desta dri get CA.WUI.OLMsgWait`

   The system displays the following:

   ```
   CA.WUI.OLMsgWait=90
   ```
4. To implement the change, restart the Director (See sections 1.7 and 1.8 for information on stopping and starting the Director).

OpenVMS

1. Add the key to the registry if it does not already exist:
   `desta dri add “CA.WUI.OLMsgWait”`
2. Set the value of the key to 90 seconds:
   `desta dri set “CA.WUI.OLMsgWait” 90`
3. View the new value of the key:
   `desta dri get “CA.WUI.OLMsgWait”`

   The system displays the following:

   ```
   CA.WUI.OLMsgWait=90
   ```
4. To implement the change, restart the Director (See sections 1.7 and 1.8 for information on stopping and starting the Director).

7.7.2 Configuring Additional Log File Directories

In order to add saved log files to the web interface’s navigation tree, files can be saved under the `svctools` directory, or in one or more directories you specify by editing the Desta registry.

To add log files which are saved in directories outside of the `svctools` path, you must first add the full path of each directory to the `CA.WUI.OLDirs` key. Multiple directories are added using a comma separated list.

For more information on Log Files, see Section 4.4.4.

Windows, Tru64 UNIX, HP-UX, and Linux

Follow these steps:

1. Add the key to the registry if it does not already exist:
   `desta dri add CA.WUI.OLDirs`
2. Set the new value for the key using the full path of each directory:
   `desta dri set CA.WUI.OLDirs “`*`directory1,directory2,...`*`”`**

   For example, in Windows you would enter:

   ```
   desta dri set CA.WUI.OLDirs “c:\morelogs,d:\evenmorelogs”
   ```
3. View the new values for the key:
   `desta dri get CA.WUI.OLDirs`

   In Windows, the system displays the following:

   ```
   CA.WUI.OLDirs=c:\morelogs,d:\evenmorelogs
   ```

4. To implement changes, restart the Director (See Sections 1.7 and 1.8 for information on stopping and starting the Director).

To delete the key and remove all directories from the search list, enter:

```
desta dri del CA.WUI.OLDirs
```

OpenVMS

Follow these steps:

1. Add the key to the registry if it does not already exist:
   ```
   desta dri add "CA.WUI.OLDirs"
   ```
2. Set the new value for the key using the full path of each directory:
   ```
   desta dri set "CA.WUI.OLDirs" "directory1,directory2,..."
   ```
3. View the new value for the key:
   ```
   desta dri get "CA.WUI.OLDirs"
   ```
4. To implement changes, restart the Director (See Sections 1.7 and 1.8 for information on stopping and starting the Director).

To delete the key and remove all directories from the search list, enter:

```
desta dri del "CA.WUI.OLDirs"
```

7.7.3 Enabling Text Entry in Other Logs Pane

When enabled, the text entry field in the Add Logs screen allows users to add log files by entering the path and filename for an event log located anywhere in the file system (Figure 7–3. For more information, see Section 4.4.4 and Figure 4–16).

Figure 7–3 Add Log Files Tab with Text Entry Field Enabled

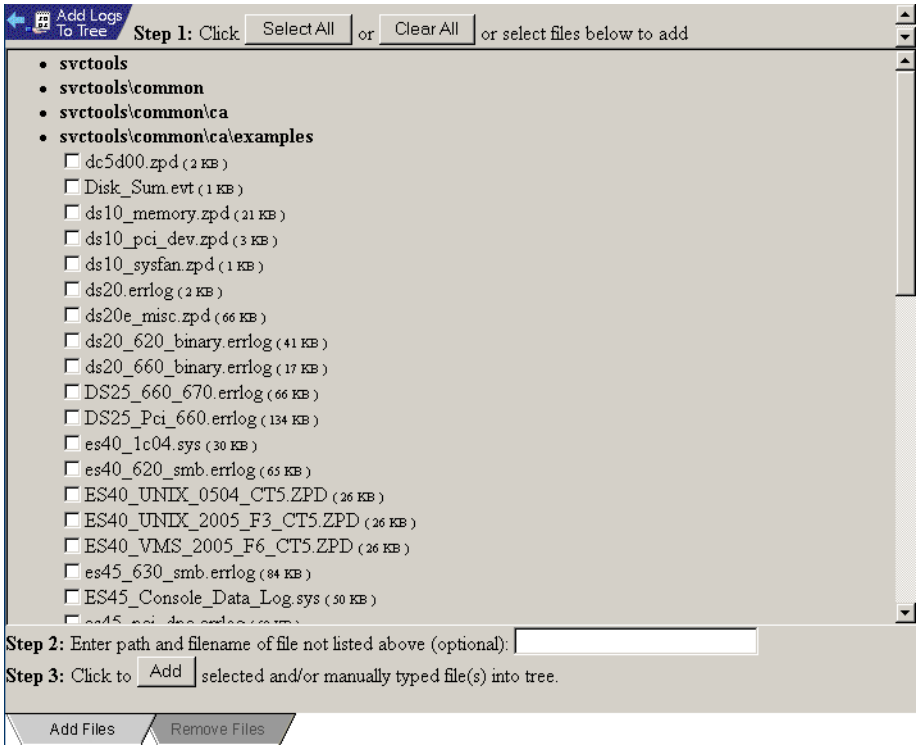

When entering a file name into the text entry field (Figure 7–4), the log file must have a `.sys`, `.evt`, `.zpd`, or `.errlog` extension. If you wish to add a file with a different extension, you will need to rename the file so it uses an acceptable file extension.

Figure 7–4 Text Entry Field

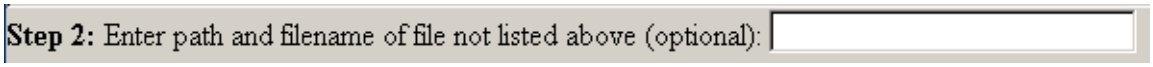

The text field can only be enabled for users you specify in the `CA.WUI.OLText` key. It cannot be enabled for all users unless you list each user individually.

Note

The list of usernames assigned to the `CA.WUI.OLText` key corresponds to the user profile entered by the user at the SEA Logon screen (see Section 4.2). SEA profiles and usernames are not related to the id a user enters to log onto a machine, and they are not authenticated by SEA during the logon process. It is therefore the responsibility of those with knowledge of text entry enabled user profiles to protect them from unauthorized use (i.e., not allowing open access to event logs anywhere on the system).

Windows, Tru64 UNIX, HP-UX, and Linux

Follow these steps:

1. Add the key to the registry if it does not already exist:
   `desta dri add CA.WUI.OLText`
2. Set the values for the key by entering a single username, or a comma-separated list of usernames:
   `desta dri set CA.WUI.OLText “`*`username1,username2,...`*`”`

   For example, in Windows you would enter:

   ```
   desta dri set CA.WUI.OLText “bill,ted”
   ```
3. View the new values for the key:
   `desta dri get CA.WUI.OLText`

   In Windows, the system displays the following:

   ```
   CA.WUI.OLText=bill,ted
   ```
4. To implement the change, restart the Director (See Sections 1.7 and 1.8 for information on stopping and starting the Director).

To delete the key and remove the text field for all users, enter the following:

`desta dri del CA.WUI.OLText`

OpenVMS

Follow these steps:

1. Add the key to the registry if it does not already exist:
   `desta dri add “CA.WUI.OLText”`
2. Set the values for the key by entering a single username, or a comma-separated list of usernames:
   `desta dri set “CA.WUI.OLText” “`*`username1,username2,...`*`”`

3. View the new values for the key:
   desta dri get “CA.WUI.OLText”
4. To implement the change, restart the Director (See Sections 1.7 and 1.8 for information on stopping and starting the Director).

To delete the key and remove the text field for all users, enter the following:

desta dri del “CA.WUI.OLText”

7.7.4 Controlling Memory Usage

The WEBES Director and Analyzer subprocesses run within a Java environment on all the supported operating systems. WEBES can override the default maximum amount of memory used by the Director process and any Java subprocesses that the Director spawns.

WEBES controls the memory usage by setting the following two DESTA registry entries:

- `desta.director.maxHeapSize` – Controls the memory used by the Director process.
- `desta.subprocess.maxHeapSize` – Controls the memory used by WEBES subprocesses.

In Java, the heap is the main block of memory that is allocated by the process. Setting the maximum size of the heap controls how much memory the process can allocate.

The following examples show the registry entries with values set:

- `desta.director.maxHeapSize=300m` – This registry limits the maximum memory for the Director process to 300 megabytes.
- `desta.subprocess.maxHeapSize=200m` – This registry entry limits the maximum memory for the Analyzer subprocess (and any other subprocess) to 200 megabytes.

Note

WEBES is installed with default heap settings. It is only necessary to adjust the values if you are having problems with out-of-memory errors.

7.7.4.1 Circumstances Requiring Memory Changes

If the Director hangs or terminates unexpectedly, check the Director log files (see Section 1.12 for more information on log files). If the log files contain errors mentioning “out of memory” conditions, one of the following conditions may apply:

- Your system has run out of memory or paging space.
- The Director process has reached its Java memory limits. These limits are set during WEBES installation, but may be overridden by setting the values on the registry entries described in this section.

If the Java memory limits are responsible for the problem, you can raise the memory limits applied to the Director process and its subprocesses. After the limits have been increased, you can restart the Director and perform the actions that caused the out of memory error. The limits can be set as high as necessary, and are only constrained by the memory and paging space available on the system.

To determine which registry entry to change, find the “out of memory” message in the Director log file. All messages from the subprocesses start with a “>” character at the beginning of the line. If the “out of memory” messages begin with “>” characters, as in the following example, then the subprocess heap limit needs to be raised.

```
> java.lang.OutOfMemoryError
> at sun.misc.Resource.getBytes(Resource.java, Compiled Code)
> at java.net.URLClassLoader.defineClass(URLClassLoader.java, Compiled Code)
...
```

The contents of the error message can vary widely. The important element is the OutOfMemoryError, which can be claimed by Java or other parts of the runtime system.

If the messages do not contain “>” characters at the beginning of the line, as in the following example, then the Director heap limit needs to be raised.

```
EXCEPTION java.lang.OutOfMemoryError
at com.compaq.svctools.ca.services.eventreaders.ReaderContext.readEvent
(ReaderContext.java, Compiled Code)
at com.compaq.svctools.ca.services.eventreaders.ReaderContext.getEvent
(ReaderContext.java, Compiled Code)
```

7.7.4.2 Changing Memory Settings

Before you begin changing the memory settings, check the current registry values to establish a baseline for your changes.

You can view the current values for the Director heap registry entry with the following commands:

- Windows, Tru64 UNIX, HP-UX, and Linux:

  desta dri get desta.director.maxHeapSize
- OpenVMS:

  desta dri get “desta.director.maxHeapSize”

You can view the current values for the subprocess heap registry entry with the following commands:

- Windows, Tru64 UNIX, HP-UX, and Linux:

  desta dri get desta.subprocess.maxHeapSize

- OpenVMS:

  ```
  desta dri get “desta.subprocess.maxHeapSize”
  ```

Once you have established a baseline value, you can modify the memory settings using the procedure for setting the heap size. The procedure varies slightly depending on your operating system.

Tru64 UNIX

To designate the maximum heap size for the Director set the value of the registry key:

1. Set the value of the registry key by entering the following command at the command prompt:

   ```
   # desta dri set desta.director.maxHeapSize XXm
   ```

   Where *XX* is the desired heap size in megabytes.

2. Restart the Director so the desta startup script will find the new heap setting and use it for the Director process. (See Sections 1.7 and 1.8 for information on stopping and starting the Director.)

To set the maximum heap size for subprocesses, use the following procedure:

1. Set the value of the registry key by entering the following command at the command prompt:

   ```
   # desta dri set desta.subprocess.maxHeapSize XXm
   ```

   Where *XX* is the desired heap size in megabytes.

2. Reset the subprocess command line in the desta registry by entering the following command at the command prompt:

   ```
   # desta setsub
   ```

3. Restart the Director so the startup script will find the new heap setting and use it for all subprocesses. (See Sections 1.7 and 1.8 for information on stopping and starting the Director.)

OpenVMS

Java on Windows and Tru64 UNIX uses more memory as needed up to the imposed limits. However, on VMS Java allocates the entire maximum heap size at startup for the lifetime of the process. Besides using the following commands to raise the heap sizes, you can also use them to reduce the heap sizes if the defaults are too resource-intensive for your system. Be aware that reducing the values limits the event processing that the Director can perform, and reducing them too much can cause the Director to fail during normal operation.

To designate the maximum heap size for the Director set the value of the registry key:

1. Set the value of the registry key by entering the following command at the command prompt:

   ```
   $ desta dri set “desta.director.maxHeapSize” “XXm”
   ```

Where *XX* is the desired heap size in megabytes.

2. Restart the Director so the desta startup script will find the new heap setting and use it for the Director process. (See Sections 1.7 and 1.8 for information on stopping and starting the Director.)

To set the maximum heap size for subprocesses, use the following procedure:

1. Set the value of the registry key by entering the following command at the command prompt:

   ```
   $ desta dri set “desta.subprocess.maxHeapSize” “XXm”
   ```

   Where *XX* is the desired heap size in megabytes.

2. Delete the subprocess command line registry key by entering the following command at the command prompt:

   ```
   $ desta dri del “desta.Subprocess.CommandLine”
   ```

3. Restart the Director so the startup script will find the new heap setting and use it for all subprocesses. (See Sections 1.7 and 1.8 for information on stopping and starting the Director.)

Windows

To set the maximum heap size for the Director process, adjust the value of the registry entry:

1. Set the value of the DESTA registry key with the following command:

   ```
   C:\> desta dri set desta.director.maxHeapSize XXm
   ```

   Where *XX* is the desired heap size in megabytes.

2. Restart the director so the desta service will find the new heap setting and use it for the director process. (See Sections 1.7 and 1.8 for information on stopping and starting the Director.)

To set the maximum heap size for subprocesses, use the following procedure:

1. Set the value of the registry key by entering the following command at the command prompt:

   ```
   C:\> desta dri set desta.subprocess.maxHeapSize XXm
   ```

   Where *XX* is the desired heap size in megabytes.

2. Restart the Director so the Java code will find the new heap setting and use it for all subprocesses. (See Sections 1.7 and 1.8 for information on stopping and starting the Director.)

7.8 Configuring Operating System-Specific Services

Some WEBES services are only appropriate for certain versions of the supported operating systems. This is usually because the earlier, older versions of the operating system do not provide the necessary support. Normally, WEBES determines which services are supported by the OS during installation and copies the necessary files. However, if you upgrade your system's OS version, you may want to add the version dependent services manually. The following sections describe the services that may need to be manually configured.

7.8.1 Drape

Drape is supported on systems running Tru64 UNIX v5.0 and newer. It provides event translation support for the Event Management (EVM) event viewer. The event viewer provides a graphical view of historical events through the common system management interface. The viewer can be launched through the SysMan Menu or through the SysMan Station. Refer to the `sysman(8)` reference page for more information.

If you upgrade your Tru64 UNIX system to a version that supports Drape, use the following procedure to configure the service:

1. Access the `/usr/opt/hp/svctools/common/ca/install` directory and locate the two enabling configuration files:
   - `DrapeConfigCA.txt`
   - `ConfigDefaultsDRAPE.txt`
2. Copy the configuration files to the `/usr/opt/hp/svctools/specific/desta/config` directory.
3. Execute the DESTA `ChangeEnrollments` command:

   ```
   /usr/sbin/desta exec
   com.compaq.svctools.desta.configuration.ChangeEnrollments -enroll
   ConfigDefaultsDRAPE.txt
   ```

The next time WEBES is started, the Drape service will be activated.

7.8.2 Indictment

The Indictment service is supported on both Tru64 UNIX and OpenVMS systems. It enables the system to automatically detect and shut down failing CPUs and certain PCI boards in order to avoid system crashes. Refer to the operating system documentation for more information on component indictment. The following sections describe how to configure Indictment on Tru64 UNIX and OpenVMS systems.

7.8.2.1 Tru64 UNIX

Indictment is supported on Tru64 UNIX v5.1 Rev 573 and newer.

Note

You can use the `sizer -v` command to check the operating system version and revision number.

In order to configure the Indictment service, use the following procedure:

1. Access the `/usr/opt/hp/svctools/common/ca/install` directory and locate the enabling configuration file `ConfigDefaultsIndictment.txt`.
2. Copy the enabling configuration file to the `/usr/opt/hp/svctools/specific/desta/config` directory.
3. Run the DESTA `ChangeEnrollments` command:

   ```
   /usr/sbin/desta exec
   com.compaq.svctools.desta.configuration.ChangeEnrollments -enroll
   ConfigDefaultsIndictment.txt
   ```

The next time WEBES is started, the Indictment service will be activated.

7.8.2.2 OpenVMS

Indictment is supported on OpenVMS V7.3-2 and newer. To configure Indictment on an upgraded system, use the following procedure:

1. Access the `SVCTOOLS_HOME:[common.ca.install]` directory and locate the enabling configuration file `ConfigDefaultsIndictment.txt`.
2. Copy the enabling configuration file to the `SVCTOOLS_HOME:[specific.desta.config]` directory.
3. Run DESTA `ChangeEnrollments` command:

   ```
   desta exec "com.compaq.svctools.desta.configuration.ChangeEnrollments" "-
   enroll" "ConfigDefaultsIndictment.txt"
   ```

The next time WEBES is started, the Indictment service will be activated.

8

Notification

This chapter describes how to configure Simple Mail Transfer Protocol (SMTP), System Initiated Call Logging (SICL), and Proactive Remote Service (PRS) for automatic notification as well as how to disable automatic notification.

8.1 Automatic Notification

Automatic notification enables you to distribute problem reports over e-mail without manual intervention. Be aware that problem reports generated by manual analysis are not sent out for notification. Only reports from automatic analysis are sent out.

The following sections describe how to configure automatic notification.

8.2 Configuring SMTP Mail Notification

Automatic notification provides the capability to send problem reports to recipients through the SMTP protocol.

Note

If you want to use SMTP (e-mail) automatic notification, your machine must either have connectivity to another SMTP server on the TCP/IP network, or it must have its own SMTP server. For further information on configuring your machine as a SMTP server, refer to your operating system documentation.

To set up SMTP (E-mail) notification of problem reports, you must edit the `NotifyCA.txt` file. You may have already entered the appropriate information during installation. If so, you will find the information stored in this file.

Use any text editor to open the file and specify what server to use for sending E-mail notification and the users to whom messages should be sent. The `NotifyCA.txt` file is in the following locations, depending on your operating system:

- Tru64 UNIX – `/usr/opt/hp/svctools/specific/desta/config`
- HP-UX – `/opt/hp/svctools/specific/desta/config`
- Linux – `/usr/opt/hp/svctools/specific/desta/config`
- OpenVMS – `svctools_home:[specific.desta.config]`
- Windows – *`install_directory`*`\specific\desta\config`

  Where *install_directory* indicates the directory where SEA was installed

The basic format of the text file is as follows:

```
SERVER=servername
FROM=username1@server.xxx.com
TO=username1@mailaddress1.com; username2@mailaddress2.com
CC=username3@mailaddress3.com
```

The servername must be either a machine currently running an SMTP server process, or localhost if the machine running SEA is also an SMTP server. The users you identify in the TO

and CC fields of the `NotifyCA.txt` file are automatically sent problem reports. When modifying the file, keep the following in mind:

- The domain name (the part of the address following the @ symbol) used in the FROM field must be valid. Some SMTP servers will not deliver the notification if the domain is invalid.
- Extraneous spaces are ignored and the semicolon can be used as a recipient separator in the TO and CC fields.
- The CC field is optional.
- For changes in the `NotifyCA.txt` file to take effect, you must stop the Director, and then restart it.

Note for UNIX

If your environment does not allow for SMTP forwarding using the normal protocol, you can add the following line to the `NotifyCA.txt` file:
`CMD=mailx -s '%s' %t`
The `mailx` command can be replaced with any other command for sending mail. The `%s` is substituted for the subject line of the problem report. The `%t` is substituted with a space-separated list of the mail addresses specified on the `TO=` lines of the `NotifyCA.txt` file.

Disabling and Enabling SMTP Notification

The SMTP Notification service is enabled by default, but will not perform any notification until the configuration procedures described in this section are performed (unless the necessary information was provided during installation). To disable any notification of problem reports, use the web interface to deselect the “autoStart” checkbox in the SMTP Notification service’s configuration attributes. The next time the Director is restarted, the Notification service will not be started, and no mail will be sent for problem reports. See Chapter 7 for more information regarding configuration.

To re-enable the service, select the “autoStart” checkbox and restart the Director.

8.3 Customer Profile File

You will need a customer profile file in order to automatically notify your Hewlett-Packard qualified service provider of problems detected by SEA. The profile file provides contact and system information used by your service provider.

Normally, the customer profile file is named `profile.txt` and depending on your operating system, the file’s location defaults to the following directory:

- Tru64 UNIX – `/usr/opt/hp/svctools/specific/desta/config`
- OpenVMS – `svctools_home:[specific.desta.config]`

- Windows – `\hp\svctools\specific\desta\config`

You can change the name and location of the profile file, however, you will need to modify the path to reflect those changes (see Section 8.3.2).

8.3.1 Profile File Contents

The installation process creates a profile file for you, however, if you need to change the file you can do so using a text editor. The file includes contact information, company information, and system information.

If you modify the profile file, you should maintain the format of the information and save your changes in the appropriate directory.

8.3.2 Path Setup

In order to use a profile file, you must specify the fully qualified path. Specify the path in the `\svctools\specific\desta\config\desta.reg` file. Add the following line to the `desta.reg` file:

`CA.ACHSProfile=`*`filename`*

Where *filename* is the path and name of the profile file. Be aware that backslash characters must be duplicated in order to be interpreted correctly.

For example, on a Windows system using the default file name and location, the path statement would appear as follows:

```
CA.ACHSProfile=C:\\Program Files\\hp\\svctools\\data\\profile.txt
```

8.4 Configuring Service Provider Notification

SEA can perform automatic notification to your service provider (such as Hewlett-Packard Services) with either SICL or PRS. Be aware that SICL and PRS notification are mutually exclusive.

8.4.1 Configuring SICL Notification

System Initiated Call Logging (SICL), also known as Automated Call Handling Services (ACHS), enables the SEA software to log service calls with a Hewlett-Packard Customer Support Center using DSNlink software. Before you enable SICL notification make sure that you have DSNlink installed.

Note

There are known DSNlink issues that may result in problems if you use DSNlink to provide SICL functionality and you connect to DSNlink via modem or X.25. For the best results, configure your DSNlink installation to use TCP/IP if your network supports it.

SICL notification is enabled and disabled from the command prompt. To enable SICL notification use one of the following commands:

```
wsea sicl on
desta sicl on
```

When you enable SICL, you will be prompted to specify an e-mail address where DSNlink will send a message indicating a call was logged. If SICL notification was previously enabled or if PRS notification is currently active, an error message displays.

To disable SICL notification, use the following command:

```
wsea sicl off
desta sicl off
```

Note

The enabling and disabling of SICL using DSNlink has changed from `wsea sicl [on|off]` to `desta sicl [on|off]`. Please use the `desta` syntax and update any scripts that refer to the `wsea sicl` command before this is completely phased out in a future release.

8.4.2 Configuring CSG Notification

The Proactive Remote Service (PRS) product contains support for a Customer Support Gateway (CSG), formerly known as a Qualified Service Access Point (QSAP). SEA can send notifications to a CSG to automatically log support calls with the Hewlett-Packard Customer Support Center. In order to use PRS, you will need to configure your site with a CSC and managed systems (refer to the PRS documentation for more information).

Before you can use PRS, you will need to enable communications with a CSG/QSAP node.

CSG/QSAP is enabled and disabled from the command prompt. Use the following command to enable it:

```
desta qsap on
```

When you enable CSG/QSAP, you will be prompted to specify the name of the Customer Service Gateway node (formerly the QSAP node) and the port QSAP listens on. The port number for CSG v4.5 and v5.0 is 2069; for v3.1 and v3.1B it is 8041. If PRS notification was previously enabled or if SICL notification is currently active, an error message displays.

To disable CSG/QSAP, use the following command:

```
desta qsap off
```

8.4.2.1 Event Log Settings

If your event log is completely full no more events can be logged and PRS will not be able to perform automatic notification. In order to ensure that the log does not fill, you should make sure the event log is set to automatically remove old events. Change the event log settings using the following procedure.

1. Open the Start menu and select **Programs > Administrative Tools (Common) > Event Viewer**.

   The Event Viewer opens.

2. Select Log Settings from the Log pull-down menu.

   The Event Log Settings dialog box opens (Figure 8–1).

Figure 8–1 Event Log Settings Dialog Box

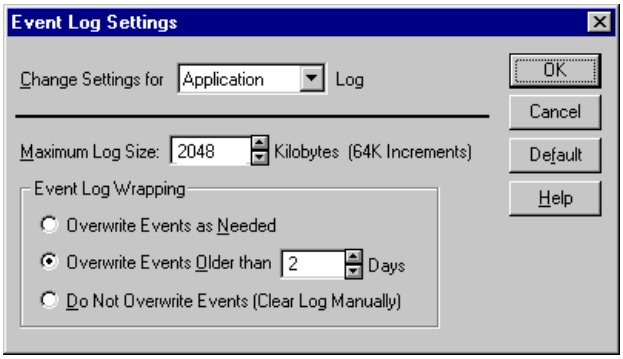

3. Select the Application Log using the drop-down list at the top of the dialog box.
4. Use the following settings for the Application Log:
   - Maximum Log Size – 2048 Kilobytes
   - Overwrite Events Older than 2 Days
5. Click the OK button to apply your changes.

A

Sample Outputs

This appendix provides examples of translated event output and analysis output.

A.1 Sample Analysis Output

```
-------- Problem Found: A Temperature Condition is being reported by the
Environmental Monitoring System


at Aug 25, 2001 10:14:09 AM GMT-04:00 ----------

Problem Report Times:
    Report Time:     Oct 18, 2002 10:58:37 AM GMT-06:00

Managed Entity:
 System Type :AlphaServer Marvel 7/800
 Computer Name :webshooter7
 System Serial Number


:MARVEL-001
 Operating System Version :Tru64 UNIX V5.1A (Rev. 1885)

Service Obligation Data:

   Service Obligation:            Valid
   Service Obligation Number:     A123456789
   System Serial Number:          A123456789
   Service Provider Company Name: Hewlett-Packard


Brief Description:
A Temperature Condition is being reported by the Environmental Monitoring
System

Callout ID:
x50FC85000007CD05

Severity:
2

Reporting Node:
webshooter7

Full Description:
 The Environmental Monitoring System has detected a Temperature change in the
System. The Condition is being reported as:

 The temperature measuring device indicates a GOOD reading.

FRU List:


 Probability      : High
 Fru Manufacturer : Not available
 Fru Model        : Not available
 Fru Part Number  : Not available
 Fru Serial Number: Not available
 Fru Firmware Rev : Not available
 Fru


Description  : The temperature sensor is on the MBM Module in the 8P Drawer
 Physical Location: Cabinet 2,


Drawer 3
 Fru Assembly     : MBM Module
 Fru Slot         : The MBM Module is accessed from the back of the
```

```
Cabinet.

Evidence:
Rule Set  : GS1280 SM Rule x1.2
Qualifiers: EFT-2
Event Id  : 54804 / 0
Event Time: Sat Aug 25 10:14:09 MDT


2001

SEA Version:
SEA for Windows Intel V4.3.1 (Build 302)


WCC Version:
Web-based Enterprise Service Common Components for
Windows Intel V4.3.1 (Build 303), member of WEB-based
Enterprise Service Suite for Windows Intel V4.3.1
(Build 302)
```

A.2 Sample Translated Event Output

The following samples show both full and brief translation output.

A.2.1 Full

```
Event:     2
Description: VMS Asychronous Device Attention  at Mar 1, 2001 9:59:34 AM GMT-
0500 from SABL15
File:    ./ca/examples/rx_data.zpd
================================================================

OS_Type              2             -- OpenVMS AXP
Hardware_Arch        4             -- Alpha
CEH_Vendor_ID        3,564         -- Hewlett-Packard Company
Hdwr_Sys_Type        22            -- Unrecognized System Type
Logging_CPU          0             -- CPU Logging this Event
CPUs_In_Active_Set   0
Entry_Type           128,098       -- VMS Asychronous Device Attention
DSR_Msg_Num          1,813         -- AlphaServer ES40
                                   .... CPU Slots: 1 (500Mhz)
                                   .... PCI Slots: 10
                                   .... MMB Slots: 8 (DIMMs)
Chip_Type            8             -- EV6  21264
CEH_Device           49
CEH_Device_ID_0     x0000 0000
CEH_Device_ID_1     x0000 0000
CEH_Device_ID_2     x0000 0000
Unique_ID_Count      93
Unique_ID_Prefix     2
TLV_DSR_String       AlphaServer ES40
TLV_DDR_String
TLV_Sys_Serial_Num   NI73702WH1
TLV_Time_as_Local    Mar 1, 2001 9:59:34 AM GMT-0500
TLV_OS_Version       X6O1-SSB
TLV_Computer_Name    SABL15
emb_ertcnt          x0000 0016
emb_class            128            Bus Class
emb_type             49             Memory Channel
emb_bcnt             0
```

```
emb_errcnt           1
emb_func             0
ucb_name_len         10
ucb_name             SABL15$MCA
ucb_dtname_len       0
ucb_dtname
Revision_Information   x0000 0001
Family_ID              x0000 0016
Member_MC_ID           x0000 0007
MC_PCI_Bus_Number      x0000 003D
MC_PCI_Slot_Number     x0000 0003
MC_PCI_Frame_Size      x0000 00A4
Vendor_ID              x1011
Device_ID_MC           x0018
Bus_Cmd                x0146
Bus_Status             x0400
Rev_ID                  176
RegProg                x00
Sub_Class              x80
Base_Class             x02
Cache_Line_Size        x00
Latency_Timer          x10
Header_Type            x00
BIST                   x00
Window_Cntl            x08
PCITbar                x78 0000
Base_Addr_1            x7800 0008
Base_Addr_2            x0000 0000
Base_Addr_3            x0000 0000
Base_Addr_4            x0000 0000
Base_Addr_5            x7800 0008
Cardbus_CIS            x0000 0000
Sys_Vendor_ID          x0000
Subsystem_ID           x0000
Expansion_ROM_Base_Addrx07C0 0000
Interrupt_Line          12
Interrupt_Pin           1
Min_Gnt                 0
Max_Lat                 0
PCT_Data               x0000 0000
MCLcsr                 x0000 C07A
   RPE[1]               x1
   Rx_Err_Ena[3]       x1
   Tx_Err_Ena[4]       x1
   MC_Int_Ena[5]       x1
   Port_Change_Ena[6]  x1
   Port_Change_Int[14] x1
   INT_Summary[15]     x1

PCIRbar                xF800 0000
MCError                x1202 0202
   Rx_Err_on_Data[1]   x1
   Cntl_Packet_History[9]x1
   Heartbeat_Ena[17]   x1
   Sum_Rx_Err[25]      x1
   Sum_Tx_Err[28]       x1

MCPort                 x5642 0000
   Line_Card_Slot[21:16]x02
   Hub_Type[24:22]     x1
   Rsvd_1[25]          x1
   Heartbeat_Timeout_Sel[26]x1
   Adapter_OK[28]      x1
   Hub_OK[30]          x1

Config                 x0000 001F
Port_Online            x0000 0000
Cluser_Status_Low      x0000 0002
Cluser_Status_High     x0000 0000
Node_0_Low             x0000 0000
```

```
Node_0_High             x0000 0000
Node_1_Low              x0000 0000
Node_1_High             x0000 0000
Node_2_Low              x0000 0009
Node_2_High             x0000 0000
Node_3_Low              x0000 0000
Node_3_High             x0000 0000
Node_4_Low              x0000 0009
Node_4_High             x0000 0000
Node_5_Low              x0000 0000
Node_5_High             x0000 0000
Node_6_Low              x0000 0000
Node_6_High             x0000 0000
Node_7_Low              x0000 0000
Node_7_High             x0000 0000
```

A.2.2 Brief

```
Event:     2
Description: VMS Asychronous Device Attention at Mon Mar 01 20:59:59 MST 2001
from SABL15
File:      ./ca/examples/rx_data.zpd
================================================================

OS_Type               2            -- OpenVMS AXP
Hardware_Arch         4            -- Alpha
CEH_Vendor_ID         3,564        -- Hewlett-Packard Company
Hdwr_Sys_Type         22           -- Unrecognized System Type
Logging_CPU           0            -- CPU Logging this Event
CPUs_In_Active_Set    0
Entry_Type            128,098      -- VMS Asychronous Device Attention
DSR_Msg_Num           1,813        -- AlphaServer ES40
                                   .... CPU Slots: 1 (500Mhz)
                                   .... PCI Slots: 10
                                   .... MMB Slots: 8 (DIMMs)
Chip_Type             8            -- EV6  21264
CEH_Device            49
CEH_Device_ID_0      x0000 0000
CEH_Device_ID_1      x0000 0000
CEH_Device_ID_2      x0000 0000
Unique_ID_Count       93
Unique_ID_Prefix      2
TLV_DSR_String        AlphaServer ES40
TLV_DDR_String
TLV_Sys_Serial_Num    NI73702WH1
TLV_Time_as_Local     Mar 1, 2001 9:59:34 AM GMT-0500
TLV_OS_Version        X6O1-SSB
TLV_Computer_Name     SABL15
emb_class             128            Bus Class
emb_type              49
```

A.3 Sample Configuration Entry

```
COMMON EVENT HEADER (CEH) V2.0
 OS_Type             1       -- Tru64 UNIX
 Hardware_Arch       4       -- Alpha
 CEH_Vendor_ID       3,564   -- Hewlett-Packard Company
 Hdwr_Sys_Type       35      -- GS40/80/160/320 Series
 Logging_CPU         0       -- CPU Logging this Event
 CPUs_In_Active_Set  1
 Entry_Type          110     -- Configuration Event
```

```
 DSR_Msg_Num        1,968   -- AlphaServer GS160
 Chip_Type          11      -- EV67 21264A
 CEH_Device         54
 CEH_Device_ID_0    x0000 03FF
 CEH_Device_ID_1    x0000 0007
 CEH_Device_ID_2    x0000 0007
 Unique_ID_Count    0
 Unique_ID_Prefix   32,640


TLV Section of CEH
 TLV_Time_as_Local       Mar 21, 2001 7:11:16 AM GMT-0500
 TLV_Computer_Name       wfsi21
 TLV_DSR_String          AlphaServer GS160 6/731
 TLV_OS_Version          Digital UNIX V4.0G (Rev. 1511)
 TLV_Sys_Serial_Num      PROTO-WF21


Configuration Entry
 NOTE
    -  CONFIGURATION ENTRY encountered in Event Log File.
    -  A Decomposed Configuration Tree Report is available
       for this event, and may be selected seperately for
       display in certain user modes.
```

B

Performance

This appendix describes the factors that may impact the performance of SEA and provides suggestions for optimizing it.

B.1 Performance and Resource Usage

Whenever SEA starts, and when you run manual analysis, the program appears to use a lot of system resources and processor cycles. However, SEA uses only the capacity that is not being asked for by other programs.

SEA always relinquishes processor cycles to other programs whenever they need them. In other words, the program uses whatever resources are available.

At startup SEA needs the available capacity for the scavenge process. Depending on the system, and the size and content of the log, the initial startup pass can take many minutes or even hours to complete. The initial analysis occurs only once, four minutes after the Director has been started. Subsequent restarts of the Director should not result in significant CPU usage except for the normal startup tasks, which may take from 10 to 30 seconds. After completing the scavenge process, SEA drops into idle mode, where resource usage hovers at only a few percent.

If you run SEA in manual mode, large amounts of system resources and processor cycles might also get used. As in the case of startup in automatic mode, the condition is directly related to the size and content of the log being processed. Once again, by design, SEA uses as many resources as are available until processing is completed.

For more information on controlling SEA's memory usage, refer to Section 7.7.4.

B.2 Performance Issues

The following symptoms are indications of a performance problem that may require your attention:

- Analysis aborts without completing.
- Translation does not produce output.
- Commands time-out.

  The DESTA Director process may be too busy scavenging to respond to other requests from the web interface or the CLI before their time-outs expire, thus, causing the request to fail. Manual translation or analysis of large binary event logs also may cause the Director to become too busy to respond to other requests in a timely manner.

- Memory errors occur.

  Processing may abort with an out-of-memory message, a communications error, or a streams error. If you are using the web interface, these errors are logged in the DESTA Director log. If you are using the command line interface, the errors will appear on the screen.

- Processing takes an excessive amount of time to complete.
- Director services fail to start up when the system is heavily loaded.

The Director will shutdown and record errors in the log. To correct for this problem, increase the ctrlrStartupTimeout value in the Director Settings (see the *SEA 4.2 Release Notes*).

B.3 Enhancing Performance

The following suggestions may improve performance and speed processing:

- In most cases, performance issues can be resolved by controlling the size of the error logs you process.

  Use filtering to create a smaller error log containing a subset of the events in the original log. Smaller error log files can speed processing and address performance issues associated with manual analysis and translation. Filtering may be performed using either the CLI or the web interface and information on filtering log files is available in Sections 3.8 and 4.6.

  Manage the system error log so that it does not grow indefinitely. One way to accomplish this is to periodically archive and reset the current error log by following the guidelines in the *WEBES Installation Guide*.

- Processing may be slowed by a fragmented disk. If processing is consistently slow, defragment your disk.
- If your system is performing a resource-intensive operation (such as scavenging), wait for the activity to complete and for the system to become idle again, then repeat the command or operation that failed.

B.3.1 Tru64 UNIX

If you have tried the above suggestions and still receive error messages (out-of-memory, communications error, or streams error) you may need to consider the following solutions:

Try increasing the total swap space allocation on your system. See the *WEBES Install Guide* for more information on swap space requirements.

On multiprocessor systems, if you have already tried creating a new log file and still receive processing errors, you may be able to eliminate those errors by forcing the DESTA Director to run on only one processor. When the DESTA Director runs on only one processor it is less susceptible to internal synchronization problems, and as a side benefit, it uses less memory. However, throughput is reduced.

To set DESTA Director to run on only one processor:

1. Stop the Director (`desta stop` from the command line).
2. Using any text editor, append the following line to the DESTA.REG file. (The default path for this file is `/usr/opt/hp/svctools/desta/config`.)

```
desta.CPUAffinity=t
```

3. Restart the Director (`desta start` from the command line).

Another workaround is to remove the swap space limitation that the Director imposes on itself to prevent it from using too much of the system's swap space. Normally, swap space usage is limited to half of the total swap space allocated by the system. Be aware that this workaround can potentially allow the Director to hang or crash the machine if it uses all the available system swap space. The Director process and the available swap space must be monitored during the time this workaround is in place (Refer Section 1.11 for details on monitoring the Director).

To remove the swap space restriction, use the following procedure:

1. Stop the Director (**`/usr/sbin/desta stop`**).
2. In the file `/usr/opt/hp/svctools/bin/desta`, change the following line:
   ```
   ulimit -v $ulimitvNEW
   ```
   To:
   ```
   ulimit -v $ulimitvOLD
   ```
3. Restart the Director (**`/usr/sbin/desta start`**).

The change only affects the Director process, not any other WEBES processes such as command-line analysis processes.

B.3.2 OpenVMS

If a VMS system continues to abort when you attempt to process a log file and other remedies have not solved the problem, copy the error log file to a platform running another operating system such as Windows NT or Tru64 UNIX, and analyze the OpenVMS error log from there instead.

C

Browsers And The Web Interface

This appendix describes how to configure your browser for SEA and provides troubleshooting tips for using browsers with the web interface.

C.1 Supported Web Browsers

Table C–1 lists the supported browser versions for SEA. Be aware that the appearance of the web interface may vary slightly when viewed with different browsers.

- Supported—fully tested
- As-is—not officially tested but may work reasonably well
- Unsupported—known not to work

Table C–1 SEA Browser Requirements

Category	Windows	UNIX variants	VMS
Supported	• Internet Explorer 6.0 • Netscape 7.x • Mozilla 1.3 or later	• Netscape 4.78 or 4.79 • Mozilla 1.4 or later	• HP Secure Web Browser (SWB) Version 1.2–1 or later (based on Mozilla)
As-Is	• Internet Explorer 5.5 • Mozilla earlier than 1.3	• Netscape earlier than 4.78 • Mozilla earlier than 1.4	• Mozilla, any HP version packaged separately from the SWB
Unsupported	• Internet Explorer earlier than 5.5 • Netscape earlier than 7.0	• Netscape 6.x	• Netscape, any version

Java Requirements

Web browsers can use different Java runtime environments, but the SEA web interface requires certain versions of Java for each web browser. The following affect all operating systems except OpenVMS which has special notes described later.

- Internet Explorer (IE) — either the Microsoft Java VM version 1.1.4, or a Sun JRE version 1.2 or higher.

  Internet Explorer on Windows 2000 includes its own Java VM 1.1.4, but no Java is included in IE on Windows XP, and Microsoft no longer supplies a Java VM. You must download and install a Sun JRE instead.

- Netscape — either the Netscape Java VM which is always included with Netscape, or a Sun JRE version 1.2 or higher.
- Mozilla — Sun JRE version 1.3.1 or higher.

  Mozilla does not include any Java VM. You must download and install a Sun JRE. You can check the version by selecting **Tools>Web Development>Java Console**. The Java version is given on the first line of the Java Console window.

Sun Java Runtime Environments can be downloaded from the following web site:

```
http://java.sun.com/j2se/downloads.html
```

You must have the desired web browser(s) installed before installing the Sun JRE. The JRE installation program will find and update any installed web browsers so they can use the Sun JRE.

Tru64 UNIX

Web browsers for Tru64 UNIX can be downloaded from the following web site:

```
http://h30097.www3.hp.com/internet/download.htm
```

Not all browsers on this site are supported by WEBES. Refer to the table above.

OpenVMS

HP now provides a fully supported Web browser for OpenVMS:

hp Secure Web Browser for OpenVMS™ Alpha™ (based on Mozilla) (SWB)

which can be downloaded from the following web site:

```
http://h71000.www7.hp.com/openvms/products/ips/cswb/cswb.html
```

Be sure to read the install documentation and release notes before using SWB for the SEA web interface.

Mozilla kits for VMS can be downloaded at:

```
h71000.www7.hp.com/openvms/products/ips/register_mozilla.html
```

Note

These are Mozilla builds later than the one upon which the Secure Web Browser (SWB) is based. They are offered on an “as-is” basis by HP, and are supported as-is by WEBES. The SWB is the preferred and fully supported browser for OpenVMS.

Be sure to read the install documentation and release notes before using Mozilla for the SEA web interface.

All web browsers for VMS require a Java runtime environment to use the SEA web interface or to access any web site that uses Java. You can either:

- Use the Java JRE embedded in WEBES (preferred when using the SEA web interface from an OpenVMS Web browser)

  Or

- Install and use the Software Development Kit (SDK) v 1.3.1-6 or later for OpenVMS, downloadable from the following web site:
  `http://h18012.www1.hp.com/java/alpha/`

Special notes apply depending on which option above you choose for accessing the SEA web interface:

To use the WEBES JRE:

1. Initialize Java in your terminal session by executing the script:
   ```
   $ @SVCTOOLS_HOME:[COMMON.JRE.LIB]JAVA$140_JRE_SETUP.COM
   ```
2. Launch the Web browser.

To use the SDK installed on the VMS system:

1. Initialize Java as described in the SDK Release Notes. For example, for the SDK v1.4.0, use either of the following two commands: (The command syntax will differ for different SDK versions.)
   ```
   $ @SYS$COMMON:[JAVA$140.COM]JAVA$140_SETUP FAST ! Use the Fast VM
   $ @SYS$COMMON:[JAVA$140.COM]JAVA$140_SETUP ! Use the Classic VM
   ```
2. Launch the Web browser.

   Java functionality within the Web browser should be identical for either initialization command above, but performance and memory usage may differ.

C.2 Browser Setup

The configuration requirements for the web interface are described here:

- Configure your browser to bypass your proxy server when you connect to the Director on any machine.
- Internet Explorer — The “Use HTTP 1.1” option must be enabled for the web interface to function properly.

  To enable the option, select Internet Options from the Tools menu. From the Options window, select the Advanced tab and make sure the check box next to “Use HTTP 1.1” is selected.

- Internet Explorer — The “Check for newer versions of stored pages” option should be set to “Every visit to the page”.

  To change the setting, select Internet Options from the Tools menu. On the General tab, click the “Settings...” button under “Temporary Internet files”. Select “Every visit to the page” and click OK.

- All Browsers – Java must be enabled for the web interface to function properly. To verify that Java is enabled, use the procedure for your browser:

Internet Explorer — select Internet Options from the Tools menu. Make sure that the check box next to Java Console Enabled is selected. Be aware that some versions of Windows XP do not include Java. If this is the case on your system, follow the instructions for installing the Sun JRE in Section C.1. (Microsoft no longer supports downloading the Microsoft VM.)

Netscape — select Preferences from the Edit menu. Click on the Advanced entry and make sure that the check box next to Java is selected.

C.3 Browser Usage

The following general operation notes apply when using the SEA web interface:

- If a screen does not automatically refresh itself, click the link that opened the screen again to manually refresh it.
- If the web interface is not functioning correctly, click the refresh button. This will reset the display and open the about screen in the display frame. (If you are using Mozilla re-login to the web interface; refer to Section C.4.3)
- Do not bookmark the web interface after logging on under a username. For example, bookmarking a URL such as to `http://`*`nodename`*`:7902/?profile=`*`user`* may result in errors. To bookmark the web interface, bookmark the logon screen (`http://`*`nodename`*`:7902`). This is true for all browsers.
- If you leave an active web interface session to visit a different web page and the logout time expires, clicking on the back button to return to your web interface session will result in multiple errors. In order to logon again, return to the root address of the node (`http://`*`nodename`*`:7902`) and repeat the logon procedure.
- Under normal operation, the color of hyper-text links changes after the link is visited. SEA presents dynamic data that is frequently updated, however, the links used to access the information do not change. As a result of this presentation, the color of links in the navigation tree may be erratic or incorrect. In most cases, the color of visited links will not change.
- Because the web pages that make up the interface are generated and refreshed dynamically, do not use the browser's back or forward buttons.

C.4 Browser Specific Limitations

Depending on the browser you use with the web interface, limitations may apply.

C.4.1 Internet Explorer

- When you access the web interface, you must preface the URL with `http://` (for example, enter **`http://16.23.132.145:7902/`** in the address line rather than **`16.23.132.145:7902/`**). If you do not enter the full URL, Internet Explorer will stop responding and the system may hang.
- Internet Explorer does not update the icons in the navigation frame quickly. Thus, if automatic analysis results in a problem report or manual analysis completes, the icon changes will not be visible immediately.
- If you are using SEA and open a new browser window, some of the icons in the first browser window may disappear. The icons can be restored by clicking the browser's Reload button.
- The progress bar at the bottom of the window indicates that loading is still occurring, even after a page is fully loaded.

  You can determine when loading has finished by watching the upper right corner of the web interface. The text "Loading New Page" appears while the page is loading and disappears once loading is completed.

C.4.2 Netscape Communicator

- If you are using Netscape 4.75 with SEA, you may notice excessive CPU usage. Some browser requests to SEA, may result in Netscape using 100% of the local system's CPU. This problem occurs if you are browsing with Netscape on the same system where SEA is running. When Netscape is using all of the CPU, SEA, which is a background process, does not respond in a reasonable amount of time. In most cases, this issue occurs in conjunction with requests such as adding files to Other Logs.

  If Netscape is using all of the CPU, the browser will appear to wait for SEA. Check your system's CPU usage and determine if Netscape is consuming the majority of the processing time.

  Wait twenty to thirty seconds and click the Stop button in the browser's toolbar. Any necessary updates are shown in the navigation tree, and you can continue to use SEA normally. If necessary, you can refresh the display frame by right-clicking on it and selecting Reload Frame from the pop-up menu. Do not use the Reload button located in the Netscape toolbar.

- Netscape may not display the contents of the navigation tree correctly. The entries in the tree may not collapse properly and as a result entries may appear to be overlapping and blank lines appear in the tree. To fix the navigation tree, click the Refresh Tree button in the navigation frame.

- Netscape for Windows inserts extra blank lines in saved problem reports. If you use the Save As option to save SEA problem reports in HTML format, the new HTML file will contain an extra blank line between every line of text. As a result, the new file appears double-spaced while the original appears single-spaced. When Netscape's Save As operation encounters the `<PRE>` tag in the original HTML file, it inserts extra lines into the source of the new file. Thus, regardless of the browser you use to open the new HTML file, the extra lines are present. Since this problem only affects text formatted with the `<PRE>` tag, it does not affect most translated events.

  To eliminate the extra spaces, right-click the Frame containing the HTML report and select View Frame Source from the pop-up menu. A text window containing the HTML source opens. In that window, press CTRL-A to select all the text and then press CTRL-C to copy it to the Clipboard. Paste the contents of the clipboard into an editor and save it to a file.

C.4.3 Mozilla and Netscape 7

- Mozilla 1.0 is the minimum version for the web interface
- The Refresh button on Netscape 7.0x does not function with the web interface. If you use the Refresh button, your current web session will stop functioning and you will need to login to the web interface again. To re-login, access the root web interface URL (`http://`*`hostname`*`.`*`domain`*`.com:7902`).

  Some Windows systems may not have this problem, but you should test your system before assuming that the Refresh button is safe to use.

  This problem does not apply to Netscape 7.1.

- Avoid opening the web interface in multiple windows using Netscape 7 and Mozilla. A frame update in one window can adversely affect the same named frame in another window. Instead, use tabs to run multiple sessions.

D

Known Messages in SEA

This appendix describes the return codes generated by CLI commands and known messages sent by SEA to its message logs (see Section 1.12 of this guide for more information on the message logs). Though the messages may appear to indicate problems, they are known and expected.

D.1 Return Codes

The following return codes are used with the SEA CLI commands.

All Commands

- 0 – No error

wsea log, wsea report, wsea sicl, wsea listrk, wsea regknw, wsea msg, desta msg, desta qsap, desta servob, desta sicl

- 386 – Insufficient arguments
- 10 – Too many arguments.
- 18 – Illegal number of arguments.
- 42 – No default krs files in default directory to process.
- 50 – Illegal arguments.
- 402 – Unknown option.
- 66 – DESTAException.
- 74 – Directory not found.
- 82 – krs files not found in directory.
- 354 – File I/O Error
- 106 – Service obligation expired.
- 114 – Bad user specified event log, or no default event logs in user specified
- 122 – Bad user specified krs file, or no default krs files in user specified
- 130 – No valid event log file(s) specified.
- 138 – No valid krs file(s) specified.
- 146 – Illegal output option argument.

wsea trans, wsea analyze, wsea filterlog, wsea fru, wsea summ

- 306 – Different argument expected
- 314 – Invalid command
- 322 – Invalid operator
- 330 – Numerical value expected
- 338 – Invalid keyword
- 346 – Invalid report type
- 354 – File I/O error
- 362 – Can not determine OS
- 370 – Invalid abbreviation
- 378 – Date value expected
- 386 – Insufficient arguments
- 394 – Command execution error

- 402 – Unknown option

desta status

- 1 – Director is not running
- 3 – Director is running
- 5 – Director is starting up
- 7 – Director is shutting down
- 9 – Director status file indicates that it is running, but the process ID was not found. As a result, the Director is assumed to be no longer running.
- 99 – Director is in an unknown state

Java VM Related Exit Codes

- 602 – VM error
- 610 – Unknown argument
- 618 – Unknown class
- 626 – Unknown method
- 634 – Missing environment
- 386 – Insufficient arguments

Installation Related Exit Codes

- 642 – The $SVCTOOLS_HOME directory does not exist.
- 650 – Could not find the Service Tools installed jar files.
- 658 – Could not find Java environment.
- 666 – Could not execute DESTA <DESTA program> executable.

Note

On VMS systems each error code has a severity of 2. Thus, an ON ERROR statement can be used in DCL scripts to trap for errors. For VMS, a bit-wise OR of the value 0x10000000 is performed on the published return code before the actual code is returned, which changes the value in $STATUS. Therefore, to determine the correct value, the leading 1 should be removed. For example, if an `Insufficient arguments` error is returned, an OR is performed with 0x10000000 and 0x00000182 (386 base 10) resulting in 0x10000182 or 268435842 base 10. Remove the leading 1 to obtain the correct decimal value.

D.2 Configuration File Created

```
___.
WARNING on February 1, 2001 11:23:35 AM MST (0.023 sec elapsed)
        Configuration file /usr/opt/hp/svctools/desta/config/
Configuration.dat not found, creating it.
```

```
Current Thread[main,5,main]
```

This warning is expected and correct the first time the WEBES Director is executed on a machine. See Chapter 7 of this guide for more information.

D.3 File Not Found

```
Could not find file: WCCApplet101BeanInfo.class
```

This message appears in the Director's log file the first time the web interface is activated. It does not affect proper operation of any part of SEA and can be ignored.

E

Other CLI Syntaxes

This appendix describes the old common syntax and DECevent emulator syntaxes available with some CLI commands.

E.1 Using Other Syntaxes

You can force a command to use a specific syntax using either of the following methods:

- Enter the syntax designator as part of the command.
- Change the default syntax.

Refer to Chapter 3 for more information on syntax designators and the default syntax.

The output generated by a command does not vary depending on syntax. Thus, manually analyzing a log file with the old common syntax will produce the same output as manually analyzing the same log file with the new common syntax.

Note

This appendix assumes that you have a working understanding of the SEA functionality. The other syntaxes described here provide the same output as their namesakes in the new common syntax. As a result, only command entry information is given here. For a more detailed description of a particular function refer to Chapter 3.

E.2 Conventions

Table E–1 describes the conventions used to show CLI commands in this manual.

Table E–1 Syntax Conventions

Convention	Meaning
Bold	Command text. Bold is used for information that must be typed as it appears. For example, command verbs are shown in bold.
Italic	Variables. Italics are used for information that varies depending on your requirements. For example, *inputfile* indicates that you should enter the name of the file you want to process.
[ ]	Optional Entries. Information shown in square brackets is not required. You may or may not include these optional modifiers. In most cases the optional entries pertain to input files, output files and filtering commands.
|	Mutually Exclusive Entries. The bar separates mutually exclusive entries.

E.3 Old Common Syntax

Old common syntax commands use the following format:

wsea x *command_verb*

Where *command_verb* indicates the action you want to perform.

Table E–2 describes the commands supported by the old common syntax:

Table E–2 Command Verbs—wsea (Old Common Syntax)

Command Verb	Description
`analyze`	Performs manual analysis one or more binary event logs. See Section E.3.1 for more details.
`trans`	Translates one or more binary event logs, but does not analyze the events. See Section E.3.2 for more details.
`summ`	Returns a summary of all the events contained in a binary event log. See Section E.3.3 for more details.
`filterlog`	Applies a filter to an existing binary event log and creates a new binary event log containing the subset of events returned after filtering. See Section E.3.4 for more details.
`listrk`	Lists the registered analysis rule sets. See Section E.3.6 for syntax information and Chapter 6 for more details on rule sets.
`regknw r`	Registers one or more analysis rule sets for use during automatic and manual event analysis. See Section E.3.6 for syntax information and Chapter 6 for more details on rule sets.
`regknw u`	Unregisters one or more analysis rule sets so they are no longer considered during automatic and manual event analysis. See Section E.3.6 for syntax information and Chapter 6 for more details on rule sets.
`help`	Displays a text-based help file. The text-file describes the new common syntax.

E.3.1 Manual Analysis

To perform manual analysis with the old common syntax, use the following command:

wsea x analyze [*inputfile*] [**outtext** | **outhtml** *outputfile*]

inputfile—enter the path and name of a binary log file. See Section E.3.5.1 for more details.

outputfile—enter the path and name where you want the output saved. See Section E.3.5.2 for more details.

E.3.2 Translation

To perform translation with the old common syntax, use the following command:

wsea x trans [*inputfile*] [**outtext** | **outhtml** *outputfile*] [**filter** *"filterstatement"*] [**brief** | **full**]

inputfile—specify the path and name of a binary log file. See Section E.3.5.1 for more details.

outputfile—specify the path and name where you want the output saved. See Section E.3.5.2 for more details.

filterstatement—enter a filterstatement to limit the events translated. See Section E.3.5.3 for more details.

Select the desired report type using the brief or full modifier.

E.3.3 Summary of Events

To view a summary of the events in a log file with the old common syntax, use the following command:

wsea x summ [**index**] [*inputfile*]

Create indexed output (instead of tallied output) by using the index modifier.

inputfile—provide the path and name of a binary log file. See Section E.3.5.1 for more details.

E.3.4 Creating New Binary Event Log Files

To create a new binary log file with the old common syntax, use the following command:

wsea x filterlog *inputfile outputfile* [*"filterstatement"*] [skipconfig]

inputfile—provide the path and name of the binary log file you want to filter to create a new log file. You must provide a input file, however, you cannot use multiple files. See Section E.3.5.1 for more details.

outputfile—provide the path and name of the new log file.

filterstatement—specify a filter to restrict the events added to the new log file. See Section E.3.5.3 for more information.

Skip the configuration entries in the input file by using the skipconfig keyword.

E.3.5 Modifying Commands

By default, the analysis, translation, summary and new binary log file commands all process the system event log. The output from analysis, translation and summary commands is displayed on the screen. You can change these defaults in order to process other binary log files and save the processing results to a file. With some of the commands you can further restrict the events that are processed by filtering the binary log file used for input. The following sections describe how to use these features.

E.3.5.1 Input Files

To change the binary log file used as input by a command, append the directory and file name of the desired file to the end of the command. For example:

```
wsea x analyze examples\ds20.errlog
```

When you are specifying an input file, the following guidelines apply:

- Specifying an input file is optional. If you do not specify either a directory or a file, SEA processes the binary system event log.

  The old common syntax `filterlog` command is the exception to this rule and requires an input file. Refer to Section E.3.4 for more information.

- You can use the relative directory structure to specify input files.
- If you specify a directory but no file name, SEA processes all the files with a `.errlog`, .sys, `.zpd`, or `.evt` extension located in the provided directory.
- Multiple filenames can be specified by separating them with spaces.
- You can use wildcards to specify multiple files.

E.3.5.2 Output Files

Note

These output file guidelines do not apply when you are creating a new binary event log. Refer to Section E.3.4 for more details.

To specify an output file, use one of the following modifiers:

```
outtext filename
outhtml filename
```

The outtext modifier creates a text output file and the outhtml modifier creates a HTML output file. The *filename* indicates the path and name where you want to save the output.

The following examples show commands that specify output files:

```
wsea x analyze outtext results.txt
wsea x analyze outhtml results.html
```

E.3.5.3 Filtering

The `trans` and `filterlog` commands enable you to filter a binary event log file and only process a subset of the events. The general rules that apply to filtering in the old common syntax are:

- Use the `filter` keyword before the filter statement when filtering with the `trans` command.
- Filter statements must be enclosed in quotation marks.
- You can join multiple filter statements by using an ampersand (&) between them.

Table E–3 describes the old common syntax filtering statements.

Table E–3 Filtering Statements (Old Common Syntax)

Filter Statement	Description
`dtb=`*`date`* (date_time_begin) `dte=`*`date`* (date_time_end)	Filters based on the time the event occurred. No events that occurred before the given start time or after the given end time are processed. The date can be entered in any format supported by Java (for example, *dd-mmm-yyyy,hh:mm:ss*). You do not need to include the time (*hh:mm:ss*) with the date.
`rtdb=`*`days`* (rel_time_days_begin) `rtde=`*`days`* (rel_time_days_end) `rthb=`*`hours`* (rel_time_hours_begin) `rthe=`*`hours`* (rel_time_hours_end)	Filters based on the time the event occurred relative to the time the first or last event in the log file occurred. Filtering based on days and hours is supported. For example, using the filter rtdb=3 will processes all the events that occurred within three days of the first event in the file.
`et=`*`nn`* `et!=`*`nn`* `et<`*`nn`* `et>`*`nn`* (entry_type)	Filters based on the numeric event type. Be aware of the following guidelines: • With the = and != operators you can enter multiple entry types by separating them with commas. • Instead of entering entry type numbers, you can use one of the supported keywords. Refer to Table E–4 for the supported keywords. •
`cn=`*`name`* `cn!=`*`name`* (computer_name)	Filters based on the node responsible for generating the event. • Using the = and != operators you can enter multiple entry types by separating them with commas. • The *name* argument is case sensitive.

Table E–3 Filtering Statements (Old Common Syntax) (continued)

Filter Statement	Description
`ost=`*n* `ost!=`*n* (os_type)	Filters based on the operating system type, using the numeric representation for each operating system. With the = and != operators you can enter multiple entry types by separating them with commas.
`idx=`*nn* `idx!=`*nn* `idx<`*nn* `idx>`*nn* (event_index)	Filters based on the event's position in the event log. The first event in the file is event index 1. With the = and != operators you can enter multiple entry types by separating them with commas.
`sort=`*keyword*	Used with a keyword to organize the output. The following keywords are supported: • entry – sorts based on entry type from highest entry type number to lowest • reventry – sorts based on entry type from lowest entry type number to highest • time – sorts based on entry time from most recent to oldest • revtime – sorts based on entry time from oldest to most recent • idx – sorts based on the entry index number from highest to lowest • revidx – sorts based on the entry index number from lowest to highest

Table E–4 Event Type Keywords (Old Common Syntax)

Keyword	Description
`mchk-all`	All machine check events.
`mchk`	All machine check events.
`mchk-sys`	All system machine check events.
`mchk-cpu`	All cpu machine check events.
`mchk-env`	All environmental machine check events.

Examples – Old Common Syntax

The following examples show sample commands that use filtering.

Processes events from the system described by *ComputerName*:

```
wsea x trans filter "computer_name=ComputerName"
wsea x filterlog inputfile.zpd outputfile.bin
"computer_name=ComputerName"
```

Processes events that did not occur on the system described by *ComputerName* that occurred after January 11, 2000:

```
wsea x trans filter “computer_name!=ComputerName & date_time_begin=11-
Jan-2000”
wsea x filterlog inputfile.zpd outputfile.bin
“computer_name!=ComputerName & date_time_begin=11-Jan-2000”
```

Processes events that occurred before 8:33:57 PM on January 31, 2000:

```
wsea x trans filter “date_time_end=31-Jan-2000,20:33:57”
wsea x filterlog inputfile.zpd outputfile.bin “date_time_end=31-Jan-
2000,20:33:57”
```

Processes events that occurred no more than four days after the first event in the log file:

```
wsea x trans filter “rel_time_days_begin=4”
wsea x filterlog inputfile.zpd outputfile.bin “rel_time_days_begin=4”
```

Processes events that occurred no more than 35 hours before the last event in the log file:

```
wsea x trans filter “rel_time_hours_end=35”
wsea x filterlog inputfile.zpd outputfile.bin “rel_time_hours_end=35”
```

Processes all CPU machine check events:

```
wsea x trans filter “entry_type=mchk-cpu”
wsea x filterlog inputfile.zpd outputfile.bin “entry_type=mchk-cpu”
```

Processes all events, except those of type 610, 620, and 630. Only the common syntax supports filtering based on specific entry types the other syntaxes must use keywords:

```
wsea x trans filter “entry_type!=610,620,630”
wsea x filterlog inputfile.zpd outputfile.bin “entry_type!=610,620,630”
```

Processes all events with a type greater than 600:

```
wsea x trans filter “entry_type>600”
wsea x filterlog inputfile.zpd outputfile.bin “entry_type>600”
```

Processes all events with a type less than 300 and an operating system of type 3:

```
wsea x trans filter “entry_type<300 & os_type=3”
wsea x filterlog inputfile.zpd outputfile.bin “entry_type<300 &
os_type=3”
```

Processes all events without an operating system type of 1 or 2. The translation command presents the output in reverse chronological order:

```
wsea x trans filter “os_type!=1,2 & sort=revtime”
wsea x filterlog inputfile.zpd outputfile.bin “os_type!=1,2”
```

Processes all the events after the fifteenth event in the log file:

```
wsea x trans filter “event_index>15”
wsea x filterlog inputfile.zpd outputfile.bin “event_index>15”
```

E.3.6 Knowledge Rule Sets

Rule sets are used in conjunction with analysis. The events in a binary log file are compared with rule sets. Depending on the results of this comparison problem reports are generated. The following old common syntax commands can be used to work with rule sets.

wsea x listrk

Lists the registered rule sets used by analysis (see Section 6.3.1 for more information).

wsea x regknw r [*ruleset*]

Registers the rule sets used by analysis (see Section 6.3 for more information).

wsea x regknw u [*ruleset*]

Unregisters the rule sets used by analysis (see Section 6.3 for more information).

E.4 DECevent UNIX Syntax

DECevent UNIX syntax commands use the following format:

wsea u *command_verb*

Where *command_verb* indicates the action you want to perform.

Table E–5 describes the commands supported by the DECevent UNIX syntax:

Table E–5 Command Verbs—wsea (DECevent UNIX syntax)

Command Verb	Description
ana	Performs manual analysis one or more binary event logs. See Section E.4.1 for more details.
-a	Translates one or more binary event logs, but does not analyze the events. See Section E.4.2 for more details.
-o sum	Returns a summary of all the events contained in a binary event log. See Section E.4.3 for more details.
-b	Applies a filter to an existing binary event log and creates a new binary event log containing the subset of events returned after filtering. See Section E.4.4 for more details.
hlp	Displays a text-based help file. The text-file describes the new common syntax.

E.4.1 Manual Analysis

To perform manual analysis with the DECevent UNIX syntax use the following command:

wsea u ana [**-f** *inputfile*] [**>** *outputfile*]

inputfile—enter the path and name of a binary log file. See Section E.4.5.1 for more details.

outputfile—enter the path and name where you want the output saved. See Section E.4.5.2 for more details.

E.4.2 Translation

To perform translation with the DECevent UNIX syntax use the following command:

wsea u -a [**-f** *inputfile*] [**brief** | **full**] [*filter flags*] [**>** *outputfile*]

inputfile—specify the path and name of a binary log file. See Section E.4.5.1 for more details.

Select the desired report type using the brief or full modifier.

filter flags—enter filter flags to limit the events translated. See Section E.4.5.3 for more details.

outputfile—specify the path and name where you want the output saved. See Section E.4.5.2 for more details.

E.4.3 Summary of Events

To view a summary of the events in a log file with the DECevent UNIX syntax use the following command:

wsea u -o sum [**-f** *inputfile*] [*filter flags*]

inputfile—provide the path and name of a binary log file. See Section E.4.5.1 for more details.

filter flags—enter filter flags to limit the events translated. See Section E.4.5.3 for more details.

E.4.4 Creating New Binary Event Log Files

To create a new binary event log file with the DECevent UNIX syntax use the following command:

wsea u -b *outputfile* [**-f** *inputfile*(*s*)] [*filter_flags*]

outputfile—provide the path and name of the new log file.

inputfile—provide the path and name of the binary log file you want to filter to create a new log file. See Section E.4.5.1 for more details.

filter flags—specify a filter to restrict the events added to the new log file. See Section E.4.5.3 for more information.

E.4.5 Modifying Commands

By default, the analysis, translation, summary and new binary log file commands all process the system event log. The output from analysis, translation and summary commands is displayed on the screen. You can change these defaults in order to process other binary log files and save the processing results to a file. With some of the commands you can further restrict the events that are processed by filtering the binary log file used for input. The following sections describe how to use these features.

E.4.5.1 Input Files

To change the input file used by a command, use the following modifier:

-f *filename*

Where *filename* indicates the path and name of the desired binary log file.

For example:

```
wsea u ana -f examples/ds20.errlog
```

When you are specifying an input file, the following guidelines apply:

- Specifying an input file is optional. If you do not specify either a directory or a file, SEA processes the binary system event log.
- You can use the relative directory structure to specify input files.
- If you specify a directory but no file name, SEA processes all the files with a `.errlog`, `.sys`, `.zpd`, or `.evt` extension located in the provided directory.
- Multiple filenames can be specified by separating them with spaces.
- You can use wildcards to specify multiple files.

E.4.5.2 Output Files

Note

These output file guidelines do not apply when you are creating a new binary event log. Refer to Section E.4.4 for more details.

To specify an output file, add the following modifier to the end of a command:

`> filename`

The modifier creates a text output file. The *filename* indicates the path and name where you want to save the output.

The following examples show commands that specify output files:

```
wsea u ana > results.txt
```

E.4.5.3 Filtering

The `-a`, `-o sum`, and `-b` commands enable you to filter a binary event log file and only process a subset of the events. You can include multiple filter statements by using more than one filtering flag in a command. In this case, separate each flag with a space.

Table E–6 describes the DECevent UNIX filtering statements.

Table E–6 Filtering Statements (DECevent UNIX syntax)

Filter Statement	Description
`-t "s:`*`date`* `e:`*`date`*`"`	Filters based on the time the event occurred. No events that occurred before the given start time or after the given end time are processed. The date can be entered in any format supported by Java (for example, *dd-mmm-yyyy,hh:mm:ss*). You do not need to include the time (*hh:mm:ss*) with the date. Be aware of the following guidelines: • The DECevent UNIX syntax combines the start and end times are in a single filter statement. • You can use the keywords YESTERDAY and TODAY.
`-i` *`keyword`* `-x` *`keyword`*	Filters based on the numeric entry type. You must enter a keyword rather than the actual entry type. Refer to Table E–7 for information on supported keywords.

Table E–6 Filtering Statements (DECevent UNIX syntax) (continued)

Filter Statement	Description
`-H` *name*	Filters based on the node responsible for generating the event. The *name* argument is case sensitive.
`-e s:`*nn* `e:`*nn*	Filters based on the event's position in the event log. The first event in the file is event index 1.
`-R`	Processes the events in reverse order according to the event index number.

Table E–7 Event Type Keywords (DECevent UNIX syntax)

Keyword	Description
`cam`	All SCSI entries logged by the CAM logger (199).
`configurations`	Configuration entries (110).
`control_entries`	System startup entries or new error log creation entries (32, 35, 300).
`cpus`	Machine check entries for AXP (mchk-cpu).
`environmental_entries`	Power entries (mchk-env).
`swxcr`	Entries logged by SWXCR (198).
`machine_checks` `mchks`	Events with machine checking information (mchk).
`operating_system=`*value* `os=`*value*	Events with a specific operating system type. The *value* parameter indicates the numeric code for the desired operating system.
`panic`	Crash re-start, system panic, or user panic entries (37, 302).
`software_informationals` `swi`	Events with lastfail, system startup, or system configuration information (volume mounts, volume dismounts, new error logs, timestamp entries) (32, 35, 37, 38, 39, 64, 65, 250, 300, 301, 310).
`osf_entry`	Events logged on a Tru64 UNIX operating system.

Examples – DECevent UNIX

The following examples show sample commands that use filtering.

Processes events from the system described by *ComputerName*:

```
wsea u -a -H ComputerName
wsea u -o sum -H ComputerName
wsea u -b outputfile.bin -f inputfile.zpd -H ComputerName
```

Processes events that occurred before 8:33:57 PM on January 31, 2000:

```
wsea u -a -t “e:31-Jan-2000,20:33:57”
wsea u -o sum -t “e:31-Jan-2000,20:33:57”
wsea u -b outputfile.bin -f inputfile.zpd -t “e:31-Jan-2000,20:33:57”
```

Processes all CPU machine check events:

```
wsea u -a -i cpu
wsea u -o sum -i cpu
wsea u -b outputfile.bin -f inputfile.zpd -i cpu
```

Processes all events without an operating system type of 1. The translation command presents the output in reverse chronological order:

```
wsea u -a -x operating_system=1 -R
wsea u -o sum -x operating_system=1
wsea u -b outputfile.bin -f inputfile.zpd -x operating_system=1
```

Processes all the events after the fifteenth event in the log file:

```
wsea u -a -e s:15
wsea u -o sum -e s:15
wsea u -b outputfile.bin -f inputfile.zpd -e s:15
```

E.5 DECevent VMS Syntax

DECevent VMS syntax commands use the following format:

wsea v *command_verb*

Where *command_verb* indicates the action you want to perform.

Table E–8 describes the commands supported by the DECevent VMS syntax:

Table E–8 Command Verbs—wsea (DECevent VMS syntax)

Command Verb	Description
/ana	Performs manual analysis one or more binary event logs. See Section E.5.1 for more details.
/tra	Translates one or more binary event logs, but does not analyze the events. See Section E.5.2 for more details.
/sum	Returns a summary of all the events contained in a binary event log. See Section E.5.3 for more details.
/bin	Applies a filter to an existing binary event log and creates a new binary event log containing the subset of events returned after filtering. See Section E.5.4 for more details.
/help	Displays a text-based help file. The text-file describes the new common syntax.

E.5.1 Manual Analysis

To perform manual analysis with the DECevent VMS syntax, use the following command:

wsea v /ana[**/out=***outputfile*] [*inputfile*]

outputfile—enter the path and name where you want the output saved. See Section E.5.5.2 for more details.

inputfile—enter the path and name of a binary log file. See Section E.5.5.1 for more details.

E.5.2 Translation

To perform translation with the DECevent VMS syntax, use the following command:

wsea v /tra[**/out=***outputfile*][**/brief** | **/full**][*filter flags*] [*inputfile*]

outputfile—specify the path and name where you want the output saved. See Section E.5.5.2 for more details.

Select the desired report type using the /brief or /full modifier.

filter flags—enter filter flags to limit the events translated. See Section E.5.5.3 for more details.

inputfile—specify the path and name of a binary log file. See Section E.5.5.1 for more details.

E.5.3 Summary of Events

To view a summary of the events in a log file with the DECevent VMS syntax use the following command:

wsea v /sum[*filter flags*] [*inputfile*]

filter flags—enter filter flags to limit the events translated. See Section E.5.5.3 for more details.

inputfile—provide the path and name of a binary log file. See Section E.5.5.1 for more details.

E.5.4 Creating New Binary Event Log Files

To create a new binary log file with the DECevent VMS syntax use the following command:

wsea v /bin=*outputfile*[*/filter_flags*] [*inputfile*(*s*)]

outputfile—provide the path and name of the new log file.

filter flags—specify a filter to restrict the events added to the new log file. See Section E.5.5.3 for more information.

inputfile—provide the path and name of the binary log file you want to filter to create a new log file. See Section E.5.5.1 for more details.

E.5.5 Modifying Commands

By default, the analysis, translation, summary and new binary log file commands all process the system event log. The output from analysis, translation and summary commands is displayed on the screen. You can change these defaults in order to process other binary log files and save the processing results to a file. With some of the commands you can further restrict the events that are processed by filtering the binary log file used for input. The following sections describe how to use these features.

E.5.5.1 Input Files

To change the input file used by a command, add the path and file name of the desired file to the end of the command.

For example:

```
wsea v /ana [.examples]ds20.errlog
```

When you are specifying an input file, the following guidelines apply:

- Specifying an input file is optional. If you do not specify either a directory or a file, SEA processes the binary system event log.
- You can use the relative directory structure to specify input files.
- If you specify a directory but no file name, SEA processes all the files with a `.errlog`, `.sys`, `.zpd`, or `.evt` extension located in the provided directory.
- Multiple filenames can be specified by separating them with spaces.
- You can use wildcards to specify multiple files.

E.5.5.2 Output Files

Note

These output file guidelines do not apply when you are creating a new binary event log. Refer to Section E.5.4 for more details.

To specify an output file, use the following modifier:

/out=*filename*

The modifier creates a text output file. The *filename* indicates the path and name where you want to save the output.

The following examples shows a command that specify output files:

```
wsea v /ana/out=results.txt
```

E.5.5.3 Filtering

The /tra, /sum, and /bin commands enable you to filter a binary event log file and only process a subset of the events. You can include multiple filter statements by using more than one filtering flag in a command.

Table E–9 describes the DECevent VMS filtering statements.

Table E–9 Filtering Statements (DECevent VMS syntax)

Filter Statement	Description
/SIN="*date*" /BEF="*date*"	Filters based on the time the event occurred. No events that occurred before the given start time or after the given end time are processed. The date can be entered in any format supported by Java (for example, *dd-mmm-yyyy,hh:mm:ss*). You do not need to include the time (*hh:mm:ss*) with the date. You can use the keywords YESTERDAY and TODAY.
/INC(*keyword*) /EXC(*keyword*)	Filters based on the numeric entry type. You must enter a keyword rather than the actual entry type. Refer to Table E–10 for information on supported keywords.
/NOD=*name*	Filters based on the node responsible for generating the event. The *name* argument is case sensitive.
/ENT=(S:*nn*,E:*nn*)	Filters based on the event's position in the event log. The first event in the file is event index 1.
/REV	Processes the events in reverse order according to the event index number.

Table E–10 Event Type Keywords (DECevent VMS syntax)

Keyword	Description
`cam`	All SCSI entries logged by the CAM logger (199).
`configurations`	Configuration entries (110).
`control_entries`	System startup entries or new error log creation entries (32, 35, 300).
`cpus`	Machine check entries for AXP (mchk-cpu).
`environmental_entries`	Power entries (mchk-env).
`swxcr`	Entries logged by SWXCR (198).
`machine_checks mchks`	Events with machine checking information (mchk).
`operating_system=`*`value`* `os=`*`value`*	Events with a specific operating system type. The *value* parameter indicates the numeric code for the desired operating system.
`panic`	Crash re-start, system panic, or user panic entries (37, 302).
`software_informationals swi`	Events with lastfail, system startup, or system configuration information (volume mounts, volume dismounts, new error logs, timestamp entries) (32, 35, 37, 38, 39, 64, 65, 250, 300, 301, 310).
`osf_entry`	Events logged on a Tru64 UNIX operating system.

Examples – DECevent VMS

The following examples show sample commands that use filtering.

Processes events from the system described by *ComputerName*:

```
wsea v /tra/nod=ComputerName
wsea v /sum/nod=ComputerName
wsea v /bin=outputfile.bin/nod=ComputerName inputfile.zpd
```

Processes events that occurred before 8:33:57 PM on January 31, 2000:

```
wsea v /tra/bef="31-Jan-2000,20:33:57"
wsea v /sum/bef="31-Jan-2000,20:33:57"
wsea v /bin/bef="31-Jan-2000,20:33:57"
```

Processes all CPU machine check events:

```
wsea v /tra/inc(cpu)
wsea v /sum/inc(cpu)
wsea v /bin=outputfile.bin/inc(cpu) inputfile.zpd
```

Processes all events without an operating system type of 1. The translation command presents the output in reverse chronological order:

```
wsea v /tra/EXC(operating_system=1)/rev
wsea v /sum/EXC(operating_system=1)
wsea v /bin=outputfile.bin/EXC(operating_system=1) inputfile.zpd
```

Processes all the events after the fifteenth event in the log file:

```
wsea v /tra/ent=(s:15)
wsea v /sum/ent=(s:15)
wsea v /bin=outputfile.bin/ent=(s:15) inputfile.zpd
```

Glossary

A

ACHS

Automatic Call Handling System. Within the service provider's customer service center, ACHS accepts incoming event analysis messages that were initiated by SICL.

analysis

The process of interpreting events from a binary event log and generating problem reports that describe any problems and possible corrective actions. There are two modes of analysis supported by <System Name>, automatic and manual.

attribute

A component of a service. Some attributes can be configured by the user to modify how <System Name> services operate.

Automated Call Handling Service

See ACHS.

automatic

One of the analysis modes supported by <System Name>. In automatic mode, <System Name> monitors the binary system event log, analyzes events, and generates reports without user intervention.

B

binary event log

A log file containing system data saved in binary format. Binary error logs are processed by <System Name> and the results of this analysis are presented in problem reports.

Bit To Text

See BTT.

BTT

Bit to Text. The process used to translate the events contained in a binary log file and produce text output. See also, translation.

C

CCAT

Computer Crash Analysis Tool. CCAT is a remote operating system failure analysis tool and is a WEBES component.

CEH

Common Event Header. The header format used for binary event logs on supported products. See the *System Event Analyzer Release Notes* for a list of the supported products.

CLI

Command Line Interface. The <System Name> interface that uses the command prompt to interact with the system. The CLI processes commands entered at the command prompt and returns information and results as text, either to the terminal window or to designated output file(s).

Command Line Interface

See CLI.

common attributes

Standard configuration settings available for all <System Name> services.

Common Event Header

See CEH.

Computer Crash Analysis Tool

See CCAT.

D

DESTA

Distributed Enterprise Service Tools Architecture. DESTA is Hewlett-Packard's high-availability system fault management architecture.

DHCP

Dynamic Host Configuration Protocol. DHCP is a protocol for automatic TCP/IP configuration that provides dynamic and static address allocation and management.

Director

The WEBES component responsible for managing a machine and communicating with other machines.

Distributed Enterprise Service Tools Architecture

See DESTA.

DSNlink

Automatic notification tool that sends the results of analysis to your service provider.

Dynamic Host Configuration Protocol

See DHCP.

E

event

System data written to the binary event log.

extended attributes

Configuration settings unique to a single <System Name> service.

F

field

Component of a frame containing a label and its corresponding value.

Field Replaceable Unit

See FRU.

frame

Part of an event consisting of one or more translated fields of information.

FRU

Field Replaceable Unit. A hardware component installed on a system.

G

global attribute

An attribute that affects all the <System Name> interfaces.

group

Multiple nodes associated in the navigation frame of the web interface.

J

Java

Platform-independent, object-oriented programming language.

L

log file

Either a binary file containing system events or a text file containing error and informational messages written by WEBES processes.

M

manual

One of the modes of operation supported by <System Name>. In manual mode, the binary log files and events to be analyzed must be specified by the user.

N

node

A remote system accessed through its Director.

notification

Procedure for relaying analysis information to the interested parties. <System Name> supports automatic notification via e-mail, SICL, or CSG/QSAP.

P

Proactive Remote Service

See PRS.

problem report

The output generated by analysis. Reports contain information about errors and suggested corrective actions.

profile

Configuration information that is associated with a log on name. The profile contains information about Director settings and navigation frame appearance that can be propagated to future sessions.

PRS

Proactive Remote Service. PRS is the next generation of SICL and is capable of operating effectively in a distributed environment.

Q

QSAP

Qualified Service Access Point. The QSAP acts as a gateway for PRS managed servers to connect with the outside world.

Qualified Service Access Point

See QSAP.

R

register

The process of installing or activating a knowledge rule set.

rule and rule set

Files that define what conditions must be met in order to trigger automatic analysis.

S

<System Name>

<System Name> is a remote system event monitoring tool and is a WEBES component.

service

A component responsible for providing a <System Name> function.

service obligation

An agreement with Hewlett-Packard for the use of the WEBES tools. The service obligation defines the terms of your support agreement with Hewlett-Packard.

SICL

System Initiated Call Logging. SICL refers to the concept of automatically sending fault and failure messages to the service provider’s customer service center. The messages are then received by ACHS, analyzed, and acted upon as appropriate.

Simple Mail Transfer Protocol

See SMTP.

SMTP

Simple Mail Transfer Protocol. SMTP is a TCP/IP protocol governing e-mail transmission and reception.

String and Value Pairs

See SVP.

SVP

String and Value Pairs. The format used to present information in generated reports. The string describes the type of information presented and the value indicates the system specific information.

system configuration

The software settings for <System Name>. The system configuration can be changed using any of the interfaces.

System Event Analyzer

See <System Name>.

System Initiated Call Logging

See SICL.

T

TCP/IP

Transmission Control Protocol/Internet Protocol. TCP/IP provides communication between computers across interconnected networks, even when the computers have different hardware architectures and operating systems.

translation

The process of converting binary event logs into readable output. See also BTT.

Transmission Control Protocol/Internet Protocol

See TCP/IP.

U

unregister

The process of removing or deactivating a knowledge rule set.

W

WBEM

Web-Based Enterprise Management. WBEM is distributed, web-based system management.

WCC

WEBES Common Components. The WCC are the portions of WEBES that allow the tool suite to function as an integrated installation. The WCC are separate from the individual tools in the WEBES suite (<System Name> and CCAT) and are transparent to the user.

Web-Based Enterprise Management

See WBEM.

Web-Based Enterprise Services

See WEBES.

WEBES

Web-Based Enterprise Services. WEBES is an integrated set of web-enabled service tools that include: <System Name> and Computer Crash Analysis Tool (CCAT). See also DESTA, WBEM.

WEBES Common Components

See WCC.

web interface

The <System Name> interface accessed through a web browser. The web interface uses graphical displays to present information and relies on a combination of mouse and keyboard actions to interact with the system.

Index

A

B

C

D

E

F

G

H

I

K

L

M

N

O

P

Q

R

S

T

U

V

W