


NightTrace

Version 5.4 Release Notes (PowerMAX)

April 2004

0890398-5.4



Copyright

Copyright 2004 by Concurrent Computer Corporation. All rights reserved. This publication or any part thereof is intended for use with Concurrent Computer Corporation products by Concurrent Computer Corporation personnel, customers, and end-users. It may not be reproduced in any form without the written permission of the publisher.

Disclaimer

The information contained in this document is subject to change without notice. Concurrent Computer Corporation has taken efforts to remove errors from this document, however, Concurrent Computer Corporation's only liability regarding errors that may still exist is to correct said errors upon their being made known to Concurrent Computer Corporation.

License

Duplication of this manual without the written consent of Concurrent Computer Corporation is prohibited. Any copy of this manual reproduced with permission must include the Concurrent Computer Corporation copyright notice.

Trademark Acknowledgments

NightTrace, NightStar, PowerMAXION, PowerMAX OS, TurboHawk, and Power Hawk are trademarks of Concurrent Computer Corporation.

Night Hawk is a registered trademark of Concurrent Computer Corporation.

Élan License Manager is a trademark of Élan Computer Group, Inc.

PowerStack is a trademark of Motorola, Inc.

OSF/Motif is a registered trademark of The Open Group.

X Window System and X are trademarks of The Open Group.

Contents

1.0 Introduction	1
2.0 Documentation	2
3.0 Prerequisites	3
3.1 Software	3
3.2 Hardware	3
4.0 System Installation	4
4.1 Target Installation	5
5.0 Overview of NightTrace 5.4	6
5.1 Changes in This Release	6
5.1.1 Analysis API	6
5.1.2 Tag Annotation	9
5.1.3 User-Defined Summary Scripts	9
5.1.4 Changes to Textual Description of Events	11
5.1.5 New Daemon Options	11
5.1.6 Accelerator Key Review	12
5.2 Common Problems and Limitations	14
5.2.1 Privileges and Capabilities	14
5.2.2 NightTrace User's Guide	14
6.0 Cautions	15
6.1 TCP High-Water Mark Settings	15
7.0 Direct Software Support	16

1.0. Introduction

NightTrace™ is used for debugging and analyzing applications and their interaction with the operating system. While useful for simple programs, NightTrace can be used to analyze complex multiprocess or multiprocessor programs with time-critical or real-time constraints.

The NightTrace toolset consists of an interactive debugging and performance analysis tool, trace data collection daemons, an Application Programming Interface (API) allowing user applications to log data values, and an API allowing user applications to subsequently analyze data that has been collected.

NightTrace APIs libraries and include files for use in user applications that log trace events to shared memory

libraries and include files for use in user applications that want to analyze data collected from user or kernel daemons

ntrace a graphical tool that controls daemon sessions and displays user and kernel trace events from trace event files

ntraceud a daemon program that copies user applications' trace events from shared memory to trace event files

ntracekd a daemon program that copies operating system trace events from kernel memory to trace event files

2.0. Documentation

Table 2-1 lists the NightTrace 5.4 documentation available from Concurrent.

Table 2-1. NightTrace Version 5.4 Documentation

Manual Name	Pub. Number
<i>NightTrace User's Guide</i>	0890398-110
<i>NightTrace Version 5.4 Release Notes (PowerMAX)</i>	0890398-5.4

Copies of the Concurrent documentation can be ordered by contacting the Concurrent Software Support Center. The toll-free number for calls within the continental United States is 1-800-245-6453. For calls outside the continental United States, the number is 1-954-283-1822 or 1-305-931-2408.

Additionally, the documentation listed above is available:

- online using the **nhelp** utility
- in PDF format in the **documentation** directory of the installation CD
- on the Concurrent Computer Corporation web site at www.ccur.com

3.0. Prerequisites

Prerequisites for NightTrace Version 5.4 are as follows:

3.1. Software

- PowerMAX OS 4.3 Patch 9 (P9) or later

or

PowerMAX OS 5.1 Service Release 3 (SR3) or later

NOTE

The aforementioned prerequisites are suggested in order to fully utilize the new kernel daemon capabilities provided by NightTrace. Otherwise, the following capabilities are unavailable and may cause unexpected behavior:

- **Pause/Resume** operations for kernel daemons will be unavailable
- **Flush** operation for kernel daemons will be unavailable
- statistics listed under the **Buffer** and **Lost** columns in the Daemon Details Area of the NightTrace Session Manager will be unavailable
- NightTrace Daemon Server package (**nstar**) (see “Target Installation” on page 5)
- X Window System (X11 Version 6.3 or later)¹
- Élan License Manager™ 5.0.1 or later

3.2. Hardware

- Computer Systems:
 - Power Hawk™ 620 and 640
 - Power Hawk 710, 720 and 740
 - Power Hawk 900 Series
 - PowerStack™ II and III
 - Night Hawk® Series 6000
 - TurboHawk™
 - PowerMAXION™
- Board-Level Products:
 - Motorola MVME2604
 - Motorola MVME4604
- An X Window System display device, such as an X™ terminal¹

¹ This is a prerequisite only for the NightTrace Event Analysis package (**ntrace**). It is not required by the NightTrace Event Logging package (**ntracelog**).

4.0. System Installation

The NightTrace product is installed as two standard PowerMAX OS software packages and utilizes the standard PowerMAX OS product installation mechanism, **pkgadd** (see **pkgadd (1)**).

The package names are **ntrace** and **ntracelog**. These names are case-sensitive.

ntrace	the NightTrace Event Analysis package. It consists of everything necessary to display trace event data.
ntracelog	the NightTrace Event Logging package. It consists of everything necessary to collect trace event data.
ntraceapi	the NightTrace API package. It consists of everything necessary for user applications to analyze data collected from user or kernel daemons.

The **ntrace** and **ntracelog** packages can be installed together or separately because there are no package dependencies between them. The **ntracelog** package was designed for installation on minimally configured systems that do not have the X11 packages installed and which will be used for event logging but not event analysis.

NOTE

In order to perform kernel tracing, the **trace** package, as shipped with PowerMAX OS, must be installed on the system. Once installed, the **trace** module needs to be activated via the **config(1)** command.

In addition, for proper execution of the NightTrace Session Manager, the PowerMAX OS package **nstar** (shipped with the NightTrace product) must be installed on the NightTrace host and any NightTrace target system (see “Target Installation” on page 5).

Please refer to the “Installing Add-on Software” chapter in the *System Administration Volume I* (0890429) manual and the *PowerMAX OS Release Notes* for instructions on software installation.

NightTrace may be installed in either the root directory or elsewhere. When you run **pkgadd (1M)** to install NightTrace, you are prompted to enter the name of the directory for installation. If you want to install in the root directory, just press the <return> key at the prompt. Otherwise, enter the name of the directory where you want NightTrace installed. If this directory does not exist, the installation procedures attempt to create it for you.

NightTrace *requires* the Élan License Manager. Follow the steps in the "Obtaining Licenses" section of the *Élan License Manager Release Notes* (0891055); the *feature alias* is NightTrace. If you are not already running the Élan License Manager, if you do not have a copy of the *Élan License Manager Release Notes*, or if you need a license key, contact Concurrent Software Distribution at 1-800-666-5405 (or 1-954-283-1836 outside the continental United States).

4.1. Target Installation

For proper operation of the session management features in NightTrace 5.4, the following software package must be installed on the NightTrace host and on any target system:

NightStar Tool	Package	Version
NightTrace	ntracelog	5.4
NightStar Daemon Server	nstar	1.2-007

This package is shipped with the NightTrace product.

5.0. Overview of NightTrace 5.4

5.1. Changes in This Release

5.1.1. Analysis API

NightTrace now includes an additional API which allows user applications to programatically analyze user or kernel trace data. The Analysis API provides users complete flexibility yet harnesses much of the power of the NightTrace graphical display. Complex conditions and state definitions can be exported from the graphical tool to aid in building analysis programs. Several example programs are included with NightTrace and can be found in the following directory:

```
/usr/lib/NightTrace/examples
```

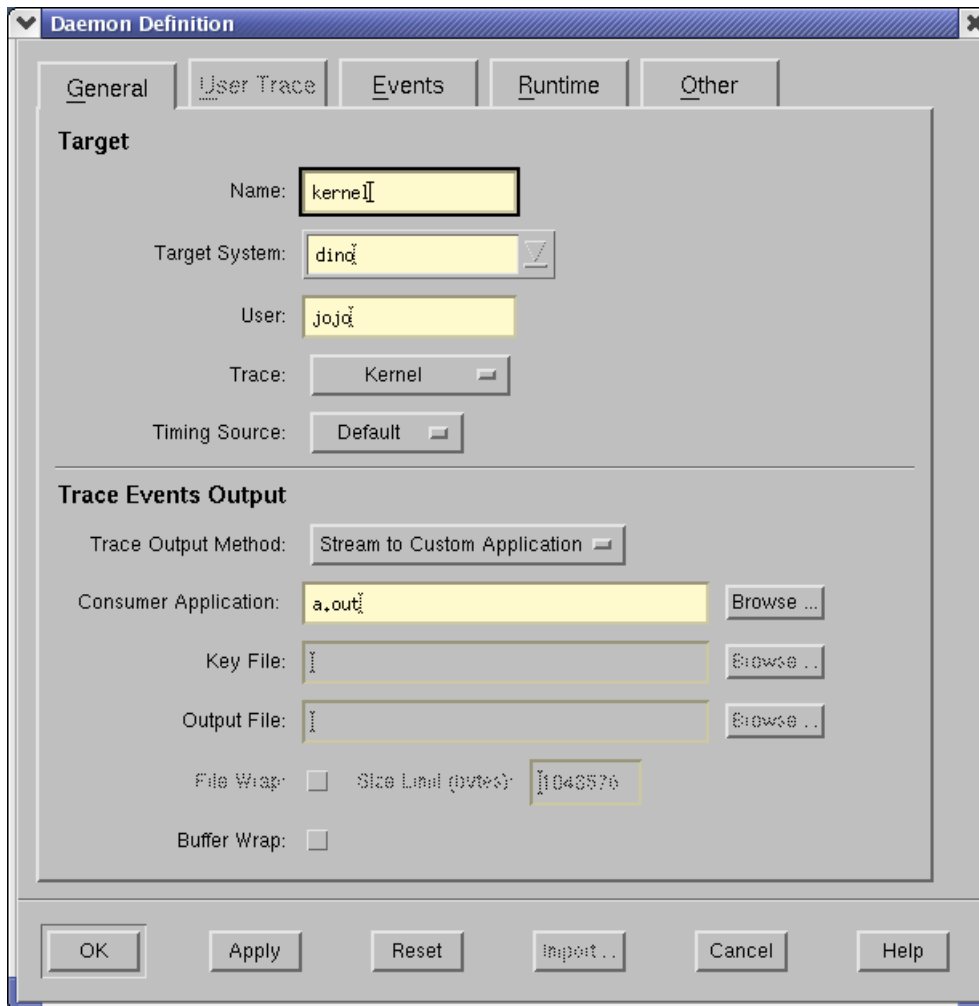
See the **README** and **Makefile** in that directory for information on the sample programs.

The Analysis API allows user programs to analyze data that was previously written to a file or can consume streaming data directly from a daemon.

The following commands will build one of the programs in the **examples** directory and execute it. It will describe context switches on CPU zero for 10 seconds before terminating:

```
cp /usr/lib/NightTrace/examples/watchdog.c .  
cc watchdog.c -lntrace_analysis  
ntracekd --stream --wait=10 /tmp/key | ./a.out 1
```

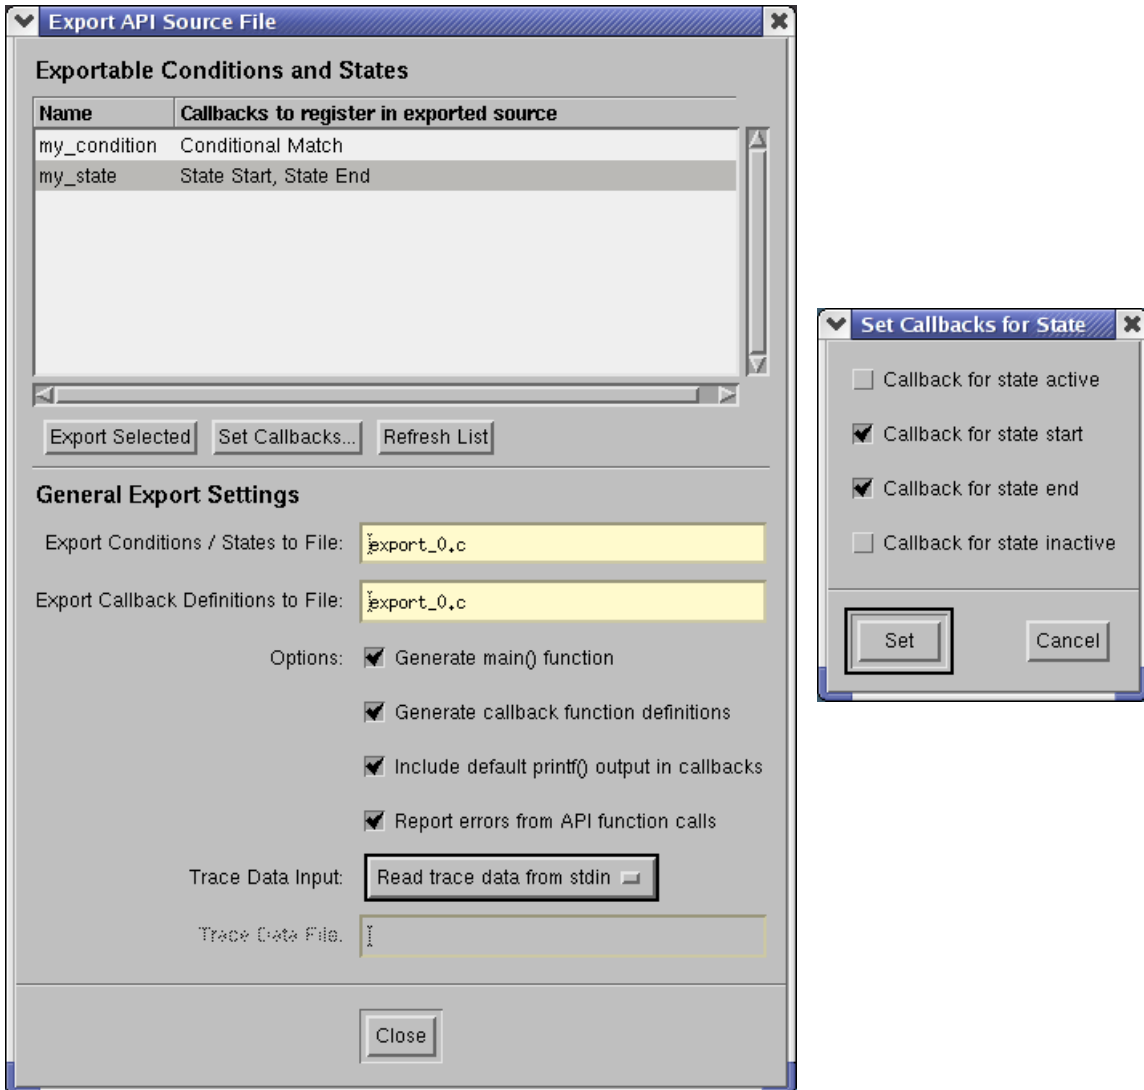
Alternatively, output from a daemon can be directed to a user application via the NightTrace graphical user interface Daemon Definition dialog.



Select **Stream to Custom Application** in the **Trace Output Method** section on the **General** page of the **Daemon Definition** dialog.

This allows user applications to monitor or analyze a system or program remotely, since the daemon may be launched on a remote system via the graphical user interface.

To aid in writing analysis programs, the NightTrace graphical utility has been enhanced with an Export API Source File dialog which is activated from the Export API Source... menu item under the File menu of the NightTrace Main window.



This dialog allows you to export specific conditions and states, generating callback routine stubs, and optionally, the main program. For each state or condition, individual callbacks may be selected.

The user need only fill in the body of the callback routine after exporting the source files and build the analysis program.

The list of exportable items includes any conditions or states that were formed during search and summary operations.

The NightTrace User's Guide has not yet been updated to include a detailed description of the NightTrace Analysis API. However, the include file, `/usr/include/ntrace_analysis.h`, includes detailed descriptions of each API call. Together with the sample programs mentioned above, users should be able to make effective use of the new API. The NightTrace User's Guide will updated in a forthcoming patch when the API documentation is complete.

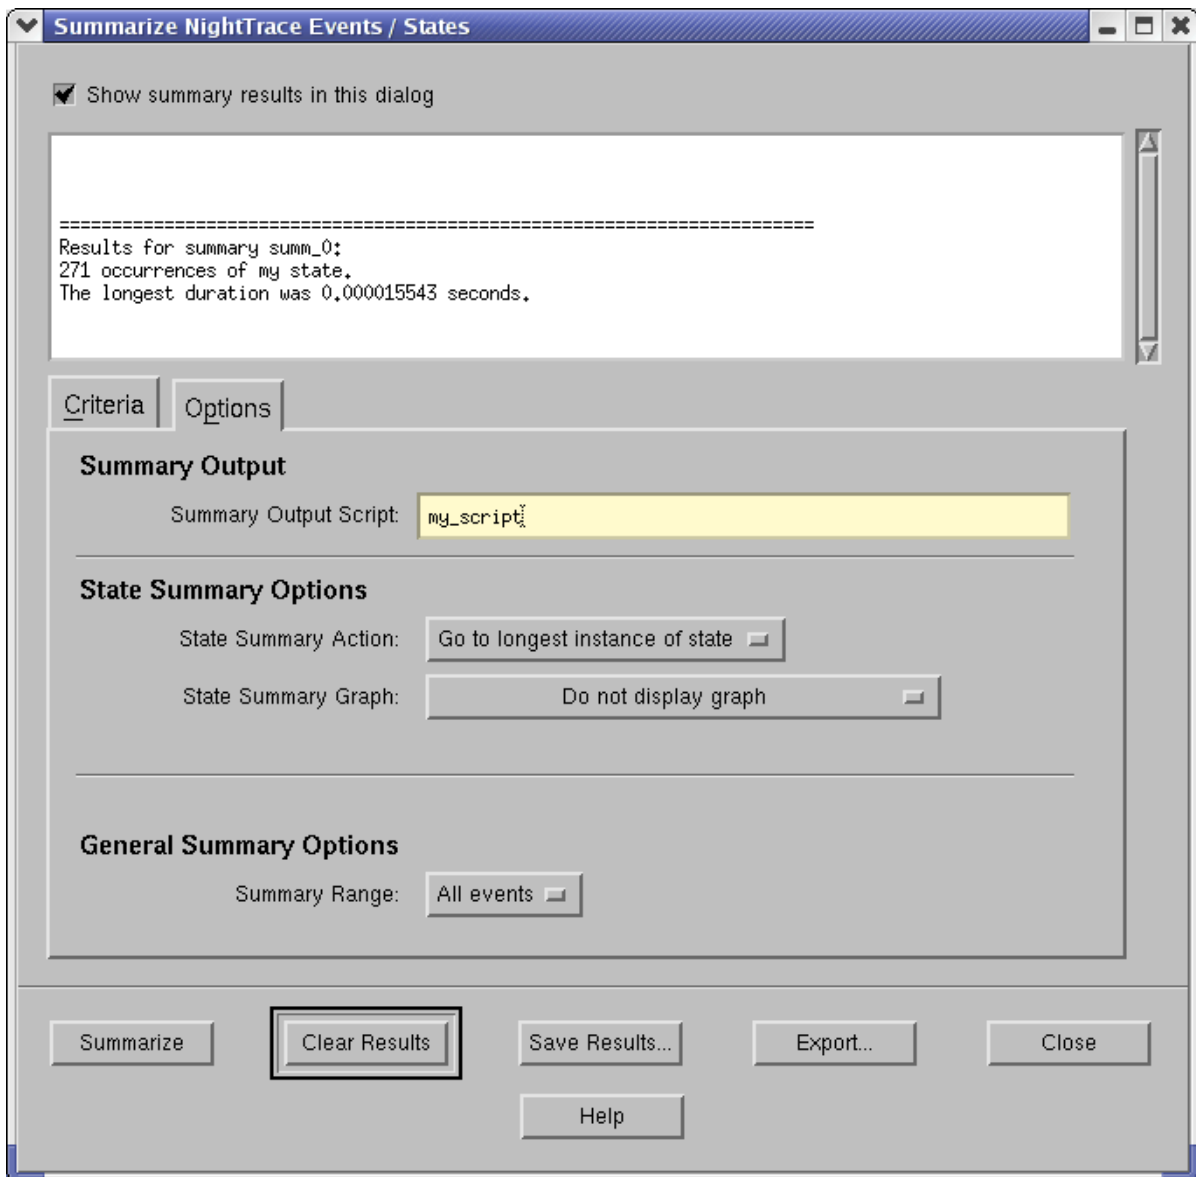
5.1.2. Tag Annotation

NightTrace now allows you to add textual comments and associate them with user-defined tags. The comments may be subsequently edited and are saved as part of a NightTrace session. Simply double-click on a tag icon on the ruler of a display page or select the **Note...** button on the **Tag Dialog** which is launched from the **Tags** menu item under the **Edit** menu on display pages.

By trimming the data set to just the set of events of interest and using the annotation feature of user-defined tags, expert analysis results can be easily exported and shared with other engineers.

5.1.3. User-Defined Summary Scripts

NightTrace now allows you to format summary information using scripts. The script is executed when the **Summarize** button is pushed and any output generated to **stdout** is displayed in the summary text area.



Scripts are associated with specific summaries in the Summary Output area under the Options tab of the Summarize NightTrace Events / States dialog. Summary script associations are saved as part of a NightTrace session. Subsequent use of named summaries on the command line will invoke the specified script.

Note that the dialog has been slightly reorganized in this release. The State Summary Options and General Summary Options were moved to the Options tab to help reduce the size of the dialog.

Summary information is passed to the script via environment variables. The following table defines the environment variables and their meaning:

Table 2-2. Summary Script Environment Variables

Variable	Meaning
NT_SUM_TYPE	Contains text describing the type of summary: Event Summary or State Summary.
NT_SUM_NUM	The number of occurrences of the state or event, expressed in decimal integer format.
NT_SUM_MIN_GAP	The minimum gap between occurrences of the state or event, expressed in seconds in decimal floating point format.
NT_SUM_MAX_GAP	The maximum gap between occurrences of the state or event, expressed in seconds in decimal floating point format.
NT_SUM_AVG_GAP	The average gap between occurrences of the state or event, expressed in seconds in decimal floating point format.
NT_SUM_TOTAL_GAP	The total time for all gaps between occurrences of the state or event, expressed in seconds in decimal floating point format.
NT_SUM_MIN_GAP_OFFSET	The offset at which the minimum gap between occurrences of the state or event occurred expressed in decimal integer format.
NT_SUM_MAX_GAP_OFFSET	The offset at which the maximum gap between occurrences of the state or event occurred expressed in decimal integer format.
NT_SUM_MIN_DURATION	For states, the minimum state duration expressed in seconds in decimal floating point format.
NT_SUM_MAX_DURATION	For states, the maximum state duration expressed in seconds in decimal floating point format.
NT_SUM_AVG_DURATION	For states, the average state duration expressed in seconds in decimal floating point format.

Table 2-2. Summary Script Environment Variables

Variable	Meaning
NT_SUM_TOTAL_DURATION	For states, the total of all state durations, expressed in seconds in decimal floating point format.
NT_SUM_MIN_DURATION_OFFSET	For states, the offset at which the minimum state duration occurred, expressed in decimal integer format.
NT_SUM_MAX_DURATION_OFFSET	For states, the offset at which the maximum state duration occurred, expressed in decimal integer format.

5.1.4. Changes to Textual Description of Events

Textual descriptions from **ntrace** have changed. The **tid=** description has been replaced by a **thr=** description, which is the name of the thread associated with the event or the value of the thread ID in decimal when a name cannot be located. On RedHawk 2.1 systems, the thread ID values are the those returned from the `gettid(2)` system call which uniquely identifies a thread system-wide. On RedHawk 1.4 systems, the thread ID values are those returned from the `getpid(2)` system call.

5.1.5. New Daemon Options

The following options were added to the **ntracekd** and **ntraceud** commands:

Daemon	Option	Meaning
ntraceud	--flush-now	Causes all events in the shared memory buffer to be immediately flushed to the output device. Most useful when the daemon is in buffer wrap mode or when streaming data to a NightTrace Analysis API application when the data rate is low.
ntracekd	--flush	Causes all events in kernel memory buffers to be immediately flushed to the output device. Most useful when the daemon is in buffer wrap mode.
ntraceud	--stream	Causes binary user trace data to be output to <code>stdout</code> . This option is intended to provide streaming user trace data as input to NightTrace Analysis API applications. See “Analysis API” on page 6.
ntracekd	--stream	Causes binary kernel trace data to be output to <code>stdout</code> . This option is intended to provide streaming kernel trace data as input to NightTrace Analysis API applications. See “Analysis API” on page 6.

5.1.6. Accelerator Key Review

Accelerator keys were added in the previous release for ease and convenience on the NightTrace Main Window as well as Display Pages. While no significant changes were made in this release, the following tables are provided for reference and summarize accelerator key usage.

Main Window	
ctrl-d	new daemon dialog
ctrl-e	edit daemon
ctrl-f	flush daemon buffer
ctrl-l	launch daemon (previously start)
ctrl-h	halt daemon (previously stop)
ctrl-p	pause daemon
ctrl-q	exit
ctrl-r	resume daemon
ctrl-s	save session
ctrl-y	copy session
ctrl-w	close the NightTrace Main Window
alt-q	exit immediately

Display Pages	
up-arrow	zoom out
down-arrow	zoom in
left-arrow	scroll left
right-arrow	scroll right
>	search forward with last criteria
<	search backward with last criteria
ctrl-a	select all
ctrl-c	copy select grid object
ctrl-d	launch Discard Events dialog
ctrl-e	launch qualified expressions dialog
ctrl-g	set Tag at current timeline
ctrl-i	zoom to region
ctrl-m	set the Mark to the current timeline
ctrl-n	center current timeline in interval
ctrl-o	toggle edit/view mode
ctrl-r	refresh interval display

Display Pages

ctrl-t	launch tags dialog
ctrl-v	paste new grid object
ctrl-x	delete selected grid object
ctrl-z	open summary dialog
ctrl-w	close current display page

NOTE

It is not necessary to press the Shift key to use the > or < accelerator.

5.2. Common Problems and Limitations

5.2.1. Privileges and Capabilities

A common problem when using the NightTrace API is that the `trace_begin()` (or `trace_start()`) call may fail if the lacks sufficient privileges to lock pages or utilize appropriate spin lock protection in which case no events are logged.

The application should be sure to check the return codes from `trace_begin()` (or `trace_start()`) and `trace_open_thread()` to ensure that subsequent `trace_event()` calls will actually log data.

5.2.2. NightTrace User's Guide

The *NightTrace User's Guide* has not yet been updated to reflect the changes in this release. Most notably, the NightTrace Analysis API is not described. A complete description of each API call is available in the include file `/usr/include/ntrace_analysis.h`. The *NightTrace User's Guide* will be updated in a forthcoming patch.

6.0. Cautions

6.1. TCP High-Water Mark Settings

If you experience kernel event loss while streaming, it may be due to insufficient TCP high-water mark settings on your PowerMAX OS system. By default, the TCP high-water marks are ~36K, which is an extremely low value, and can cause delays during I/O read and write operations of significant sizes.

We recommend adding the following commands to the `/etc/rc2.d/S69inet` script to raise the high-water mark settings.

```
$NDD -set /dev/tcp tcp_recv_hiwater_def 4194304
$NDD -set /dev/tcp tcp_xmit_hiwater_def 4194304
```

Add the above lines to the script *after* the existing lines of the form:

```
$NDD -set /dev/tcp tcp_old_urp_interpretation 1
$NDD -set /dev/ip ip_respond_to_timestamp 1
$NDD -set /dev/ip ip_respond_to_timestamp_broadcast 1
```

7.0. Direct Software Support

Software support is available from a central source. If you need assistance or information about your system, please contact the Concurrent Software Support Center at 1-800-245-6453. Our customers outside the continental United States can contact us directly at 1-954-283-1822 or 1-305-931-2408. The Software Support Center operates Monday through Friday from 8 a.m. to 7 p.m., Eastern Standard Time.

Calling the Software Support Center gives you immediate access to a broad range of skilled personnel and guarantees you a prompt response from the person most qualified to assist you. If you have a question requiring on-site assistance or consultation, the Software Support Center staff will arrange for a field analyst to return your call and schedule a visit.

