


NightTrace

Version 5.1.1 Release Notes (PowerMAX)

July 2002

0890398-5.1.1



Copyright

Copyright 2002 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 5.1.1 | 6 |
| 5.1 Changes in This Release | 6 |
| 5.1.1 Modification of Shared Memory Segment | 6 |
| 5.2 Enhancements | 6 |
| 5.2.1 NightTrace Session Manager | 6 |
| 5.2.2 Streaming | 6 |
| 5.2.3 Dynamic daemon coupling with user applications | 6 |
| 5.2.4 trace_begin() | 6 |
| 5.2.5 Dynamic enabling and disabling of trace events | 7 |
| 5.2.6 Reduced overhead for inactive trace_event() calls | 7 |
| 5.2.7 Kernel display page improvement | 7 |
| 5.2.8 Removing data-sets from the display | 7 |
| 5.2.9 Additional spin lock protection mechanism | 8 |
| 5.3 Known Deficiencies | 9 |
| 6.0 Cautions | 11 |
| 6.1 TCP High-Water Mark Settings (PowerMAX OS) | 11 |
| 7.0 Direct Software Support | 12 |

1.0. Introduction

NightTrace™ is part of the NightStar™ tool set used for developing, debugging and analyzing multiprocess and multiprocessor real-time applications on PowerMAX OS™ systems.

The NightTrace toolset consists of an interactive debugging and performance analysis tool, trace data collection daemons, and an Application Programming Interface (API) allowing user applications to log data values. Trace applications can be written in C, Fortran, or Ada. NightTrace allows you to graphically display information about important events in both your application and the kernel - for example, event occurrences, timings, and data values.

NightTrace consists of the following parts:

| | |
|-----------------|---|
| NightTrace API | libraries and include files for use in user applications that log trace events to shared memory |
| ntrace | a graphical tool that controls daemon sessions and displays user and kernel trace events in trace event file(s) |
| ntraceud | a daemon program that copies user applications' trace events from shared memory to trace event file(s) |

NightTrace is flexible. As a user, you control:

- selection of user tracing of your application and/or kernel tracing
- selection of timestamp source
- trace-point placement within your application
- the source language of the trace application
- the number of processes and CPUs you gather data on
- the amounts and types of information you display
- trace event searches and summaries

2.0. Documentation

Table 2-1 lists the NightTrace 5.1.1 documentation available from Concurrent.

Table 2-1. NightTrace Version 5.1.1 Documentation

| Manual Name | Pub. Number |
|--|---------------|
| <i>NightTrace Manual</i> | 0890398-080 |
| <i>NightTrace Version 5.1.1 Release Notes (PowerMAX)</i> | 0890398-5.1.1 |

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 X Window System utility, **nhelp**
- on the Concurrent Computer Corporation web site at www.ccur.com

3.0. Prerequisites

Prerequisites for NightTrace Version 5.1.1 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 the NightTrace Session Manager (see “NightTrace Session Manager” on page 6). 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
 - 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**). See “System Installation” on page 4 for more information about these packages.

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.

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 NightTrace Session Manager (see “NightTrace Session Manager” on page 6), the following software package must be installed on the NightTrace host and on any target system:

| NightStar Tool | Package | Version |
|-------------------------|----------------|----------------|
| NightStar Daemon Server | nstar | 1.1 |

This package is shipped with the NightTrace product.

5.0. Overview of NightTrace 5.1.1

5.1. Changes in This Release

5.1.1. Modification of Shared Memory Segment

The internal structure of the shared memory segment used to communicate between user applications and **ntraceud** has changed. If your user applications include the statically linked version of **libntrace** then you must relink your applications after installing the new version of NightTrace. Otherwise, the `trace_start()` API call will fail with `NTBADVERSION` and no trace events will be logged. Checking the return status from `trace_start()` is highly advised to avoid confusion in this situation.

5.2. Enhancements

5.2.1. NightTrace Session Manager

The NightTrace Session Manager allows users to manage user and kernel NightTrace daemons. It provides users with the ability to define a session consisting of one or more daemon definitions which can be saved for future use. These definitions include daemon collection modes and settings, daemon priorities and CPU bindings, and data output formats, as well as the trace event types that are logged by that particular daemon.

Using the NightTrace Session Manager, users can manage multiple daemons simultaneously on multiple target systems from a central location.

The NightTrace Session Manager offers the user the ability to start, stop, pause, and resume execution of any of the daemons under its management. The user may also view statistics as trace data is being gathered as well as dynamically enable and disable events while a particular daemon is executing.

5.2.2. Streaming

In addition to sending trace output to a file for later analysis, the NightTrace Session Manager also offers a streaming output method. When streaming, trace output is sent directly to the NightTrace display buffer for immediate analysis even while additional trace data is being collected.

5.2.3. Dynamic daemon coupling with user applications

NightTrace 5.1.1 now allows trace daemons to join user applications on-the-fly. Previously, users were required to first start the **ntraceud** daemon, then run their applications to completion, then stop **ntraceud**. NightTrace 5.1.1 allows you to start and stop **ntraceud** daemons at will, or use the NightTrace Session Manager (see “NightTrace Session Manager” on page 6) to “peek” at trace data sets generated from your application when so desired.

5.2.4. `trace_begin()`

A new API routine, `trace_begin()`, has been provided which replaces the `trace_start()` routine. (`trace_start()` is still provided for backward compatibility but is now considered obsolete.)

In addition to the actions that `trace_start()` performed, `trace_begin()` allows the user to specify configuration settings that were previously only available when invoking `ntraceud` from the command line. This is especially important for applications that may run for long periods of time before a user daemon joins the application.

The `trace_begin()` routine performs the following operations:

- verifies that the version of the NightTrace library linked with the application is compatible with the version used by the user daemon (if already running)
- verifies that the supplied configuration settings are not in conflict with a pre-existing daemon or, if the user daemon does not yet exist, defines the configuration with these settings
- verifies that the RCIM synchronized tick clock (if selected as the timestamp source) is counting
- attaches the shared memory buffer (after creating it if necessary)
- initializes the preemption control mechanism
- locks critical NightTrace library routine pages in memory
- initializes tracing in the current process

5.2.5. Dynamic enabling and disabling of trace events

The NightTrace Session Manager allows users to dynamically enable or disable selected trace event types while a particular daemon is running.

5.2.6. Reduced overhead for inactive `trace_event()` calls

NightTrace 5.1.1 significantly reduces the overhead for applications which make `trace_event()` API calls when no daemon is present. Most users will find the overhead so insignificant that `trace_event()` calls can be present in their applications during the entire development stage or even in final production.

5.2.7. Kernel display page improvement

In previous versions of NightTrace, kernel events for specific CPUs were discarded initially or after event loss until a context switch was seen so that the appropriate process ID/name could be displayed and so syscall, interrupt, and exception state graphs were accurate.

NightTrace 5.1.1 takes a different approach. For the initial state, and after any pause/resume or event loss, all kernel events are shown, even if no context switch has occurred. All syscall, interrupt, and exception states are considered inactive in the display until appropriate events which would start these states are seen. The process ID associated with events will indicate "unknown" until a context switch is seen.

This is especially useful for applications which consume an entire CPU and rarely context switch. Previously, no kernel events associated with these CPUs would be displayed. In NightTrace 5.1.1, all kernel events are shown.

5.2.8. Removing data-sets from the display

NightTrace 5.1.1 allows you to remove a specific data-set from the display buffer so that the associated events no longer appear in display pages.

This can be especially useful when occasionally streaming kernel data (taking several second snapshots over a period of time) as the data-set can drastically increase in size over time.

Use the **Close Trace File...** menu item from the main NightTrace window under the **NightTrace** menu bar.

NOTE

The filename associated with streaming data-sets is included in the text pane of the main NightTrace window.

5.2.9. Additional spin lock protection mechanism

The NightTrace API and associated daemons use high-performance, low-intrusion spin locks to protect critical sections involved in logging and consuming trace events.

Instead of the system priority level (SPL) spin lock protection mechanism, rescheduling variables can now be used. This is sufficient for applications that do not log trace points from user-level interrupt handlers. See **resched_cnt1(2)** for more information. In addition, see the section titled “Locking Policies” in the “NightTrace Session Manager” chapter of the *NightTrace Manual* (0890398).

5.3. Known Deficiencies

NightTrace 5.1.1 contains the following deficiencies:

- The NightTrace Session Manager offers a streaming output method. When streaming, trace data is sent directly to the NightTrace display buffer for immediate analysis even while additional trace data is being collected.

In the current implementation, streaming relies heavily on temporary external files. These files are written to `/tmp` (or to the location specified by the `TMPDIR` environment variable) and are named:

`stream_identifier`

where *identifier* is a randomly-generated alphanumeric string.

The name of this file is listed in the NightTrace Global Window when streaming is activated. For example, when streaming user trace data, a line similar to the following would appear in the NightTrace Global Window:

```
User trace event log file: /tmp/stream_OdOISS
```

When streaming kernel data, a vectors file of the form:

`stream_identifier.vec`

is generated in addition to the kernel trace event log file.

Data captured while streaming can be viewed on subsequent invocations of NightTrace by specifying the `stream_identifier` file as an argument.

Note that for large amounts of data (or for streams left active for long periods of time), NightTrace may appear to be sluggish and utilize large amounts of disk space. A more efficient and flexible implementation is underway and will be available in a future release of NightTrace.

- When invoking NightTrace with extremely large data files (or when displaying large data files with the `Display` button in the NightTrace Session Manager), NightTrace may be unresponsive to GUI activities. In reality, even though no hourglass appears, NightTrace will become responsive again when it finishes initializing the trace data.
- Due to required kernel support, kernel bufferwrap mode is currently unsupported.

Kernel filewrap mode is supported, however, and can be used when extended trace periods are desired. It is anticipated that the required kernel support for bufferwrap will be available in a forthcoming PowerMAX OS patch at which time support will be activated within NightTrace.

- On PowerMAX OS 4.3 systems, some processes in kernel display pages will only have PID values and not symbolic names associated with them. This will be true for processes that start executing after kernel tracing has begun but exit before kernel tracing has terminated. This is especially common when using the streaming feature while capturing kernel data.

This is not an issue on PowerMAX OS 5.1 systems.

- The dialog presented by selecting the `Close Trace File...` item from the NightTrace menu in the main NightTrace window allows you to remove a data-set from the current display buffer. This is especially useful when streaming kernel data in an on-and-off manner, as the kernel data-set can become extraordinarily large. While this is a handy capability, the dialog should allow you to choose from a pre-populated list of current files representing each data-set. Currently, however, the

dialog is simply a file-selection dialog. Look in the text pane of the main window to locate the file-name associated with the data-set you wish to remove.

6.0. Cautions

6.1. TCP High-Water Mark Settings (PowerMAX OS)

If you experience kernel event loss while running kernel daemons with the Session Manager 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.

