


NightTrace

Version 5.1 Release Notes (Linux)

July 2002

0898398-5.1



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

Night Hawk is a registered trademark of Concurrent Computer Corporation.

Motorola is a registered trademark of Motorola, Inc.

PowerStack is a trademark of Motorola, Inc.

Linux is a registered trademark of Linus Torvalds.

Red Hat is a registered trademark of Red Hat, Inc.

Intel is a registered trademark of Intel Corporation.

X Window System is a trademark of The Open Group.

Contents

1.0 Introduction	1
2.0 Documentation	2
3.0 Prerequisites	3
3.1 Host System	3
3.1.1 Software	3
3.1.2 Hardware	4
3.2 Target System	4
3.2.1 Software	4
3.2.2 Hardware	4
4.0 System Installation	5
4.1 Separate Host Installation	5
4.2 Target Installation	7
5.0 Overview of NightTrace 5.1	9
5.1 Changes in This Release	9
5.1.1 Modification of Shared Memory Segment	9
5.2 Enhancements	9
5.2.1 NightTrace Session Manager	9
5.2.2 Streaming	9
5.2.3 Dynamic daemon coupling with user applications	9
5.2.4 trace_begin()	9
5.2.5 Dynamic enabling and disabling of trace events	10
5.2.6 Reduced overhead for inactive trace_event() calls	10
5.2.7 Kernel display page improvement	10
5.2.8 Removing data-sets from the display	10
5.2.9 Additional spin lock protection mechanism	11
5.3 Known Deficiencies	12
6.0 Cautions	14
6.1 TCP High-Water Mark Settings (PowerMAX OS)	14
7.0 Direct Software Support	15

1.0. Introduction

NightTrace™ is part of the PowerWorks™ Linux Development Environment (PLDE) and is used for developing, debugging and analyzing multiprocess and multiprocessor real-time applications.

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. A NightTrace daemon on a target system running PowerMAX OS™ writes trace event records to user-specified files. These trace event records are generated by the execution of user application programs written in C, Fortran, or Ada. The generated trace files are input to the **ntrace** program on the host system running Linux®. The **ntrace** interactive trace-event analysis tool provides summary statistics as well as a graphical view of the run-time behavior of those application programs as well as the PowerMAX OS system on which they were run.

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 documentation available from Concurrent.

Table 2-1. NightTrace Version 5.1 Documentation

Manual Name	Pub. Number
<i>NightTrace Manual</i>	0890398-080
<i>NightTrace Version 5.1 Release Notes (Linux)</i>	0898398-5.1

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 PowerWorks Linux Development Environment utility, **nhelp**
- in PDF format in the **documentation** directory of the PLDE Installation CD
- on the Concurrent Computer Corporation web site at www.ccur.com

3.0. Prerequisites

Prerequisites for NightTrace Version 5.1 for both the host system and target system are as follows:

3.1. Host System

3.1.1. Software

- Red Hat® Linux*
- Required capabilities

NOTE

The following capabilities are normally installed as part of the standard installation of Red Hat Linux and the PowerWorks Linux Development Environment. During installation of the PLDE, the user will be notified if required capabilities do not exist on the Linux system.

- PowerWorks Linux Development Environment

Capabilities	RPMs providing these capabilities
ccur-HyperHelp ccur-HyperHelp-scripts	ccur-x11progs-6.4.2-003 ccur-HyperHelp-scripts-6.4.2-002

- Red Hat Linux

Capabilities	RPMs providing these capabilities
/bin/sh ld-linux.so.2 libX11.so.6 libXext.so.6 libc.so.6 libm.so.6 libnsl.so.1 libresolv.so.2 libstdc++-libc6.1-1.so.2 libc.so.6(GLIBC_2.0) libc.so.6(GLIBC_2.1) libm.so.6(GLIBC_2.0) libm.so.6(GLIBC_2.1)	<i>Red Hat 7.1:</i> bash-2.04-21 compat-glibc-6.2-2.1.3.2 compat-libstdc++-6.2-2.9.0.14 glibc-2.2.4-19.3 XFree86-libs-4.1.0-15 <i>Red Hat 7.2:</i> bash-2.05-8 compat-glibc-6.2-2.1.3.2 compat-libstdc++-6.2-2.9.0.16 glibc-2.2.4-13 XFree86-libs-4.1.0-3

* This product has been extensively tested on Red Hat Linux 7.1 and 7.2. However, this product has not been tested with versions of Linux supplied by other vendors.

3.1.2. Hardware

- an Intel®-based PC - 300Mhz or higher (recommended minimum configuration)
- 64MB physical memory (recommended minimum configuration)

3.2. Target System

3.2.1. Software

- PowerMAX OS 4.3 or later
- **ntracelog** 5.1 or later (see “Target Installation” on page 7 for more information)

3.2.2. 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

4.0. System Installation

Installation of the host portion of NightTrace is normally done as part of the general installation of the PowerWorks Linux Development Environment software suite. A single command installs (or uninstalls) all software components of the PLDE, as described in the *PowerWorks Linux Development Environment Release Notes* (0898000).

The following section describes how to install (or uninstall) NightTrace separately from the PLDE suite for those rare cases when this is necessary.

In addition, NightTrace requires that certain packages be installed on the PowerMAX OS target systems. See “Target Installation” on page 7 for more information.

4.1. Separate Host Installation

In rare cases, it may be necessary to install (or uninstall) NightTrace independent of the installation of the PowerWorks Linux Development Environment software suite. This may be done using the standard Linux product installation mechanism, **rpm** (see **rpm(8)**).

The rpm name associated with NightTrace 5.1 is:

ccur-ntrace

and the file associated with this rpm is:

ccur-ntrace-5.1-000.i386.rpm

which can be found in the **linux-i386** directory on the PowerWorks Linux Development Environment Installation CD.

NOTE

The user must be root in order to use the **rpm** product installation mechanism on the Linux system.

To install the NightTrace rpm, issue the following commands on your Linux system:

1. Insert the PowerWorks Linux Development Environment Installation CD in the CD-ROM drive
2. Mount the CD-ROM drive (assuming the standard mount entry for the CD-ROM device exists in **/etc/fstab**)

```
mount /mnt/cdrom
```

3. Change the current working directory to the directory containing the NightTrace rpm

```
cd /mnt/cdrom/linux-i386
```

4. Install the rpm

```
rpm -i ccur-ntrace-5.1-000.i386.rpm
```

By default, the product is installed in `/usr/opt`.

5. Change the current working directory outside the `/mnt/cdrom` hierarchy

```
cd /
```

6. Unmount the CD-ROM drive (otherwise, you will be unable to remove the PowerWorks Linux Development Environment Installation CD from the CD-ROM drive)

```
umount /mnt/cdrom
```

To uninstall the NightTrace rpm, use the following command:

```
rpm -e ccur-ntrace
```

4.2. Target Installation

NightTrace requires that the following software package is installed on the PowerMAX OS target systems:

NightStar Tool	Package	Version
NightTrace	ntracelog	5.1 or later

This package is normally installed during the installation of the corresponding NightStar tool on the PowerMAX OS system. However, the package is also included on the PowerWorks Linux Development Environment Installation CD and may be installed by following the procedures below if the above package has not been installed on the target system, or if the version installed on the target system is earlier than that specified above.

NOTE

The version of the required software package installed on the PowerMAX OS target systems must be the same as or later than that listed above. If this is not the case, follow the procedures below.

If your PowerMAX OS system has a CD-ROM device:

```
mkdir /mnt/tmp
mount -F cdfs /dev/dev_name /mnt/tmp
cd /mnt/tmp/powermax-ppc604
```

If your PowerMAX OS system does not have a CD-ROM device:

On the Linux system, insert the PowerWorks Linux Development Environment Installation CD in the CD-ROM drive, mount the CD-ROM drive, and use NFS to export it to the PowerMAX OS system:

```
mount /mnt/cdrom
exportfs powermax_os_system:/mnt/cdrom
```

On the PowerMAX OS system, mount the CD-ROM as an NFS filesystem:

```
mkdir /tmp/cdrom
mount -F nfs linux_system:/mnt/cdrom /tmp/cdrom
cd /tmp/cdrom/powermax-ppc604
```

On the PowerMAX OS system, install the target portions of the NightStar tools via the following commands:

```
cp ./nightstar.Z /tmp/nightstar.Z
zcat /tmp/nightstar.Z | pkgadd -d -
```

If your PowerMAX OS system has a CD-ROM device:

```
cd /  
umount /mnt/tmp
```

If your PowerMAX OS system does not have a CD-ROM device:

On the PowerMAX OS system, unmount the NFS filesystem:

```
cd /  
umount /tmp/cdrom
```

On the Linux system, unexport the NFS filesystem, and then unmount the CD-ROM:

```
cd /  
exportfs -u powermax_os_system:/mnt/cdrom  
umount /mnt/cdrom
```

5.0. Overview of NightTrace 5.1

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 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 allows you to start and stop **ntraceud** daemons at will, or use the NightTrace Session Manager (see “NightTrace Session Manager” on page 9) 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 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 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, all kernel events are shown.

5.2.8. Removing data-sets from the display

NightTrace 5.1 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 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.

