


NightTrace

Version 5.4 Release Notes (Linux)

April 2004

0898398-5.4



Copyright

Copyright 2003 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, PowerWorks, PowerMAXION, PowerMAX OS, Power Hawk, RedHawk, and TurboHawk 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	3
3.2 Target Systems	4
3.2.1 RedHawk Systems	4
3.2.1.1 Software	4
3.2.1.2 Hardware	4
3.2.2 PowerMAX Systems	4
3.2.2.1 Software	4
3.2.2.2 Hardware	4
4.0 System Installation	5
4.1 Separate Host Installation	5
4.2 Target Installation	7
4.2.1 RedHawk Target Installation	7
4.2.2 PowerMAX Target Installation	7
5.0 Overview of NightTrace 5.4	9
5.1 Changes in This Release	9
5.1.1 Analysis API	9
5.1.2 Tag Annotation	12
5.1.3 User-Defined Summary Scripts	12
5.1.4 Changes to Textual Description of Events	14
5.1.5 New Daemon Options	14
5.1.6 Enhanced Capabilities for Threads Programs	14
5.1.7 Accelerator Key Review	15
5.2 Common Problems and Limitations	17
5.2.1 Privileges and Capabilities	17
5.2.2 Multithreaded Programs and the NightTrace API	17
5.2.3 NightTrace User's Guide	19
6.0 Direct Software Support	20

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 (Linux)</i>	0898398-5.4

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 for both the host system and target systems are as follows:

3.1. Host System

3.1.1. Software

- RedHawk™ Linux *or* Red Hat® Linux*
- Required capabilities**

NOTE

The following capabilities are normally installed by the installation script on the CD containing the NightStar tools. The user will be notified if required capabilities do not exist on the Linux system.

- NightStar™ Tools

Capabilities	RPMs providing these capabilities
<code>ccur-HyperHelp</code> <code>ccur-HyperHelp-scripts</code> <code>ccur-elanlm</code>	<code>ccur-HyperHelp-scripts-6.4.2-002</code> <code>ccur-elanlm-5.0-8</code> <code>ccur-x11progs-6.4.2-006</code>

3.1.2. Hardware

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

* This product has been extensively tested on RedHawk Linux 1.4, 2.1, Red Hat Linux 8.0, 9.0 and Enterprise Linux 3.0. However, this product has not been tested with versions of Linux supplied by other vendors.

** The “Capabilities” listed may be found in those versions of the RPMs listed under “RPMs providing these capabilities” or in later versions.

3.2. Target Systems

3.2.1. RedHawk Systems

3.2.1.1. Software

- RedHawk Linux 1.4 or later
- Required RedHawk Linux RPMs (see “RedHawk Target Installation” on page 7 for more information)

3.2.1.2. Hardware

- any iHawk Series 860 system

3.2.2. PowerMAX Systems

3.2.2.1. Software

- PowerMAX OS 4.3 or later
- Required PowerMAX OS packages (see “PowerMAX Target Installation” on page 7 for more information)

3.2.2.2. 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

4.0. System Installation

Installation of the host portion of NightTrace is normally done as part of the general installation of either the RedHawk NightStar Tools or the PowerWorks Linux Development Environment. A single command installs (or uninstalls) all required software components. See the *RedHawk NightStar Tools Release Notes* (0898008) or *PowerWorks Linux Development Environment Release Notes* (0898000) for more information.

The following section describes how to install (or uninstall) NightTrace independently for those rare cases when this is necessary.

NOTE

NightTrace requires that certain packages are installed on the target systems. See “Target Installation” on page 7 for more information.

4.1. Separate Host Installation

At times, it may be necessary to install (or uninstall) NightTrace independent of the installation of the software suite in which it is normally distributed. This may be done using the standard Linux product installation mechanism, **rpm** (see **rpm (8)**).

The RPM name associated with NightTrace 5.4 is:

ccur-ntrace

and the file associated with this rpm is:

ccur-ntrace-5.4-000.i386.rpm

which can be found in the **RPM** directory on the 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 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/RPM

4. Install the rpm

```
rpm -Uvh ccur-ntrace-5.4-000.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 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

4.2.1. RedHawk Target Installation

When targeting a RedHawk Linux system, NightTrace requires that the following RPMs are installed on that target system:

NightStar Tool	RPM
NightTrace server	ccur-ntracelog-5.4-000.i386.rpm
NightTrace Analysis API	ccur-ntraceapi-5.4-000.i386.rpm
NightStar daemon	ccur-nstar-1.2-007.i386.rpm
Élan License Manager	ccur-elanlm-5.0-8.i386.rpm

These RPMs may be installed on the target system by installing the RedHawk NightStar Tools on that system. However, the individual RPMs can be found in the **RPM** subdirectory on the installation CD and may be installed separately. See “Separate Host Installation” on page 5 for the procedure to install an individual RPM.

4.2.2. PowerMAX Target Installation

When targeting a PowerMAX system, NightTrace requires that the following software packages are installed on that target system:

NightStar Tool	Package	Version
NightTrace	ntracelog	5.4 or later
NightTrace Analysis API	ntraceapi	5.4 or later
NightStar daemon	nstar	1.2-007

These packages are normally installed during the installation of the corresponding NightStar tool on the PowerMAX OS system. However, the packages are also included on the PowerWorks Linux Development Environment Installation CD and may be installed by following the procedures below if the packages listed above have not been installed on the target system, or if the versions installed on the target system are not the same as those specified above.

NOTE

The versions of the required software packages installed on the PowerMAX OS target systems must be the same as those 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 5.4

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 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
```

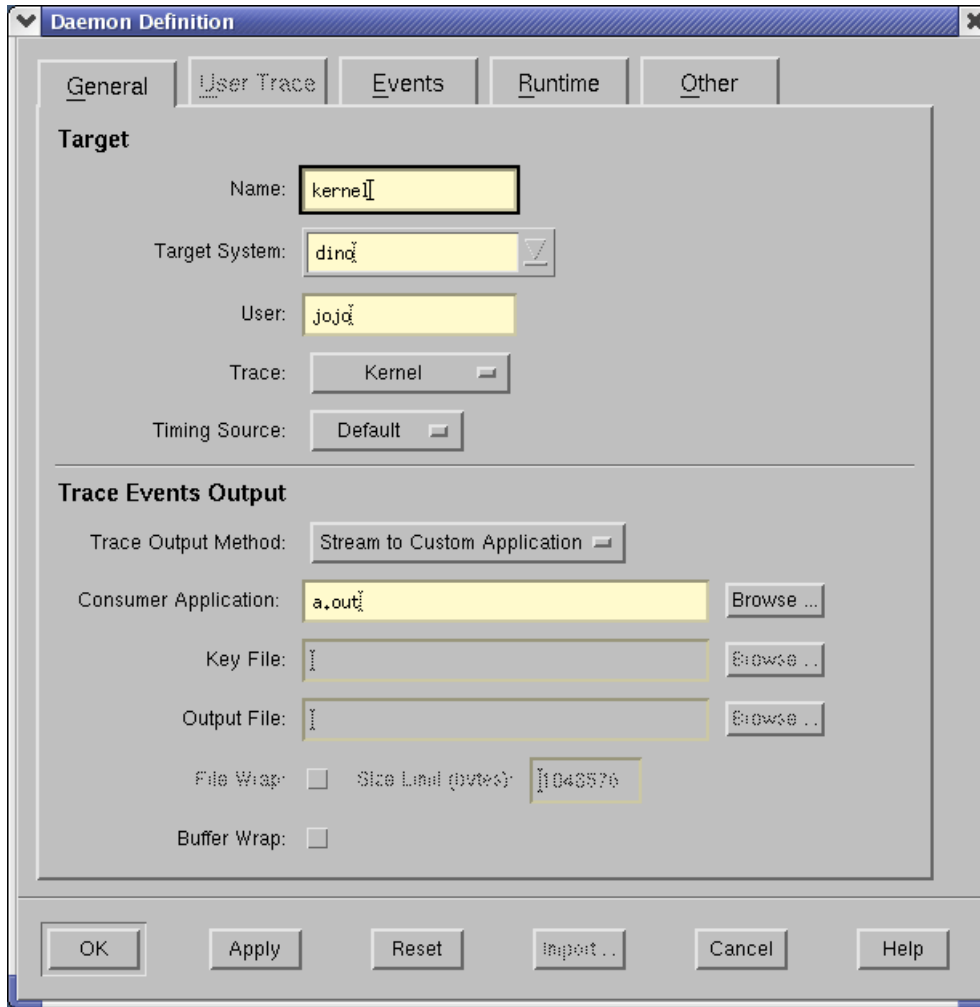
See the **README** and **Makefile** in that directory for information on the sample programs.

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 example 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 Events Output** section of the **General** tab.

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 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.

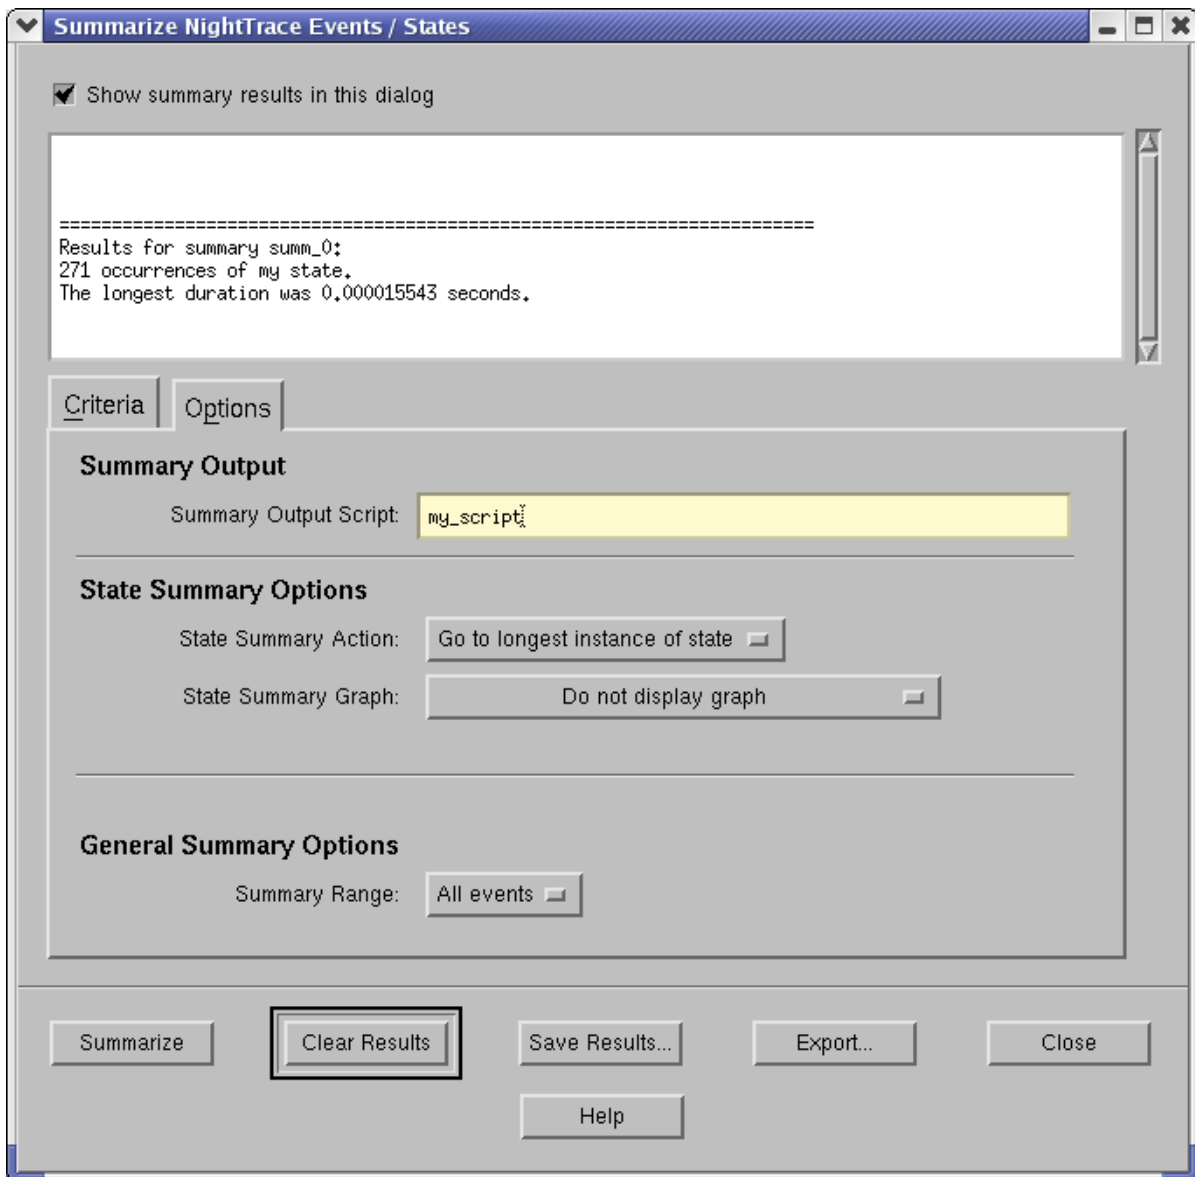
5.1.2. Tag Annotation

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 in 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 of 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** 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 ntracekd	--stream	Causes binary trace data to be output to <code>stdout</code> ; this option is intended to provide streaming input to NightTrace Analysis API applications. See “Analysis API” on page 9.

5.1.6. Enhanced Capabilities for Threads Programs

Many of the limitations held in previous releases for programs using threads and the NightTrace API have been reduced or lifted. Trace events for individual, named threads can now be easily identified. See “Multithreaded Programs and the NightTrace API” on page 17 for details.

5.1.7. 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 is provided for reference and summarizes 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

Some of the locking policies recommended for use with NightTrace daemons and the NightTrace API require special privileges.

A common problem when using the NightTrace API is that the `trace_begin()` (or `trace_start()`) call may fail if the user lacks sufficient privileges to lock pages or utilize appropriate spin lock protection. If these calls fail due to lack of privilege, subsequent API calls will fail to log events.

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.

Either root access or user registration in a privileged capabilities role is required when page locking or rescheduling variables are used (the `--lock` and `--resched` options to `ntraceud` and the `ntc_lock_pages` and `ntc_use_resched` fields of an `ntconfig_t` structure passed to `trace_begin()`).

Have your system administrator register you as a NightTrace user according to the following instructions:

1. Add the following line to the `/etc/pam.d/rsh` and `/etc/pam.d/login` files:

```
session required /lib/security/pam_capability.so
```

NOTE

For those users that log into their system directly from the Gnome or KDE graphical desktop environment, it is necessary to add the above line to `/etc/pam.d/gdm` or `/etc/pam.d/kde`, respectively. In addition, you must restart your X server or reboot your system before these changes will take effect.

2. Add the following line to the ROLES section of the `/etc/security/capability.conf` file:

```
role traceuser cap_sys_nice cap_sys_rawio cap_sys_lock
```

3. Add the following line to the bottom of the `/etc/security/capability.conf` file:

```
user user traceuser
```

where *user* is the login name of the desired user.

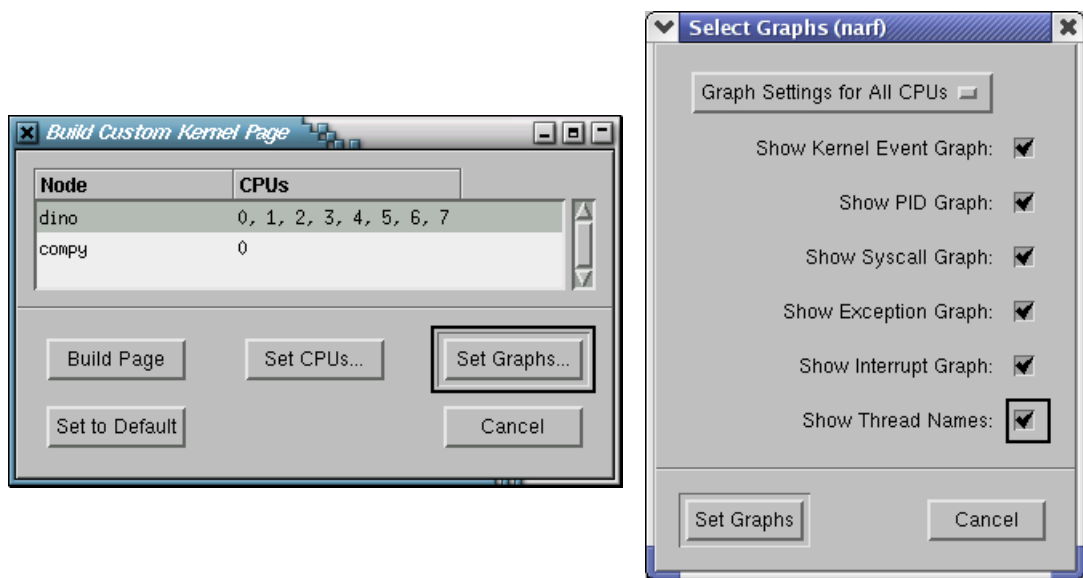
After these activities are complete, you must log off and log back onto the RedHawk system

5.2.2. Multithreaded Programs and the NightTrace API

The NightTrace API has been changed to provide enhanced functionality with multi-threaded user programs.

The following improvements have been made in this release, subject to limitations and API changes as described below:

- *Rescheduling Variables* can be selected as the *Locking Policy*, ensuring efficient and correct execution of critical sections in API calls.
- Multiple threads can make `trace_open_thread()` calls in the same application. This allows individual threads to be identified by name during data analysis.
- The thread associated with kernel trace events can be identified by name during data analysis if user and kernel data is included in the same analysis session.
- Kernel display pages have been enhanced to optionally display both the program name and thread name. From the **Pages** menu of the main window and click on **Custom Kernel Page...**



Click on **Set Graphs...** to bring up the **Select Graphs** dialog and click on **Show Thread Names**. Thread names on kernel events are only available for threads which have user trace data in the same NightTrace session.

- Textual descriptions from **ntrace** have been 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.

In order to effectively use these enhancements, the user must link with an alternative NightTrace API library and must replace POSIX thread creation calls:

- Link with `-lntrace_thr` instead of `-lntrace`
- Replace `pthread_create()` thread creation calls in your user application with `Pthread_create()`; this is a simple string change operation -- the function profile and POSIX threads semantics are otherwise identical. Threads created using `pthread_create()` will fail all subsequent NightTrace API calls.
- Include `<ntrace_thr.h>` instead of `<ntrace.h>` (provides declaration of `Pthread_create()`)

In order to properly and efficiently protect critical sections within the NightTrace API, the Locking Policy should be set to use rescheduling variables using the following mechanisms:

- Select the **Rescheduling Variables** option from the **Locking Policies** section of the **User Trace** tab in the **Daemon Definition** dialog when launching daemons from the graphical user program.
- Select the **--resched** option when launching daemons using the command-line utility **ntraceud**
- Set `ntc_use_resched` to `TRUE` when specifying `ntconfig_t` records as passed to the `trace_begin` API function. This is not necessary if the user application will always be launched after a user daemon has been initiated, however, it is highly recommended to ensure consistent results (in case the user daemon is accidentally launched after the application).

Limitations if **-lntrace_thr** is not used:

- Use of rescheduling variables is prohibited. Use of the NightTrace API in threaded programs without specifying rescheduling variables as the locking policy can cause API calls to spin indefinitely and have a dramatic performance impact on the system.
- Individual threads cannot be reliably located by thread name
- Kernel events cannot be reliably associated with thread names
- Only one call to `trace_open_thread()` is allowed per program

In either case, when multithreaded applications and the NightTrace API are used with RedHawk 1.4, the process ID logged with threaded programs may produce unexpected results, since threads do not share the same process ID under RedHawk 1.4. Under RedHawk 2.1, the PID value is shared between all threads in a program.

5.2.3. 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. 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.

