



NightView User's Guide

Version 8.1

(RedHawk™ Linux®)

0898395-420

July 2025

Copyright 2011-2025 by Concurrent Real-Time. All rights reserved. This publication or any part thereof is intended for use with Concurrent Real-Time products by Concurrent Real-Time personnel, customers, and end-users. It may not be reproduced in any form without the written permission of the publisher.

The information contained in this document is believed to be correct at the time of publication. It is subject to change without notice. Concurrent Real-Time makes no warranties, expressed or implied, concerning the information contained in this document.

To report an error or comment on a specific portion of the manual, photocopy the page in question and mark the correction or comment on the copy. Mail the copy (and any additional comments) to Concurrent Real-Time, 2881 Gateway Drive, Pompano Beach, FL 33069-4324. Mark the envelope “**Attention: Publications Department.**” This publication may not be reproduced for any other reason in any form without written permission of the publisher.

Concurrent Computer Corporation and its logo are registered trademarks of Concurrent Computer Corporation. All other Concurrent product names are trademarks of Concurrent while all other product names are trademarks or registered trademarks of their respective owners.

Linux[®] is used pursuant to a sublicense from the Linux Mark Institute.

NightStar’s integrated help system is based on Assistant, a Qt[®] utility. Qt is a registered trademark of Digia Plc and/or its subsidiaries.

NVIDIA[®] CUDA[™] is a trademark of NVIDIA Corporation.

Preface

NightView is a general purpose source-level program debugger. Some of the features make it useful for debugging systems of real-time programs, but it can also be used to debug a single ordinary program.

NightView can debug programs written in multiple languages. C, C++ and Fortran are supported.

NightView can debug multiple processes on the local system or on different hosts.

NightView has been designed to be as flexible as possible. The NightView command interpreter includes macro processing so that you can write your own NightView commands.

You communicate with NightView with one of three user interfaces. The command-line interface is useful when no advanced terminal capabilities are present. A simple full-screen interface is available for ASCII terminals. The graphical user interface provides the most functionality.

Scope of Manual

This document is the user manual for the NightView debugger. It is intended for anyone using NightView, regardless of their previous level of experience with debuggers. This manual describes how to use NightView, by way of tutorial and reference guide. There is also material for system administrators.

Structure of Manual

The manual begins with the short tutorials, Chapter 1 [A Quick Start] on page 1-1 and Chapter 2 [A Quick Start - GUI] on page 2-1, giving you just enough information to get you started. For more complete tutorial, see Chapter 4 [Tutorials] on page 4-1.

The next section describes the major concepts you will need to understand in order to get the best use out of NightView. See Chapter 3 [Concepts] on page 3-1.

More detailed information about the NightView commands is found in Chapter 6 [Command-Line Interface] on page 6-1.

The next chapter describes a simple full-screen interface to NightView. See Chapter 7 [Simple Full-Screen Interface] on page 7-1.

The next chapter describes the graphical user interface for NightView. See Chapter 8 [Graphical User Interface] on page 8-1.

Syntax Notation

The following notation is used throughout this guide:

italic

Books, reference cards, and items that the user must specify appear in *italic* type. Special terms and comments in code may also appear in *italic*.

list bold

User input appears in **list bold** type and must be entered exactly as shown. Names of directories, files, commands, options and man page references also appear in **list bold** type.

`list`

Operating system and program output such as prompts and messages and listings of files and programs appears in `list` type. Keywords also appear in `list` type.

`window`

Keyboard sequences and window features such as push buttons, radio buttons, menu items, labels, and titles appear in `window` type.

[]

Brackets enclose command options and arguments that are optional. You do not type the brackets if you choose to specify such option or arguments. Mutually exclusive choices are separated by the pipe (|) character.

{ }

Braces enclose mutually exclusive choices separated by the pipe (|) character, where one choice must be selected. You do not type the braces or the pipe character with the choice.

...

An ellipsis follows an item that can be repeated.

::=

This symbol means *is defined as* in Backus-Naur Form (BNF).

Related Publications

The following publications are referenced in this document:

0890514	<i>NightBench™ User's Guide</i>
0898004	<i>RedHawk Linux User's Guide</i>
0898008	<i>NightStar Installation Guide</i>

0898008	<i>NightStar Tutorial</i>
0898398	<i>NightTrace™ User's Guide</i>
0898465	<i>NightProbe™ User's Guide</i>
0898480	<i>NightSim™ User's Guide</i>
0898515	<i>NightTune™ User's Guide</i>

Reporting Bugs

Please report any bugs you encounter. Software support is available from the Concurrent Software Support Center at our toll free number 1-800-245-6453. For calls outside the United States, the number is 1-954-283-1822. The Software Support Center operates Monday through Friday from 8 a.m. to 5 p.m. Eastern Standard Time. You may also submit a bug report or a request for assistance at any time by using the Concurrent Real-Time web site at <http://www.concurrent-rt.com/support> or by sending an email to support@concurrent-rt.com.

Contents

Chapter 1 A Quick Start

Sample Program	1-1
Starting Up	1-2
Getting Help	1-3
Setting a Breakpoint	1-4
Finishing up	1-5

Chapter 2 A Quick Start - GUI

Sample Program - GUI	2-1
Starting Up - GUI	2-2
Getting Help - GUI	2-3
Setting a Breakpoint - GUI	2-5
Finishing up - GUI	2-5

Chapter 3 Concepts

Debugging	3-2
Accessing Files	3-2
Programs and Processes	3-3
Multiple Processes	3-3
Families	3-4
Attaching	3-4
Detaching	3-4
Core Files	3-5
Qualifiers	3-5
Dialogues	3-6
Dialogue I/O	3-7
Real-Time Debugging	3-7
Remote Dialogues	3-7
Remote File Access	3-9
ReadyToDebug	3-10
Finding Your Program	3-10
Controlling Your Program	3-11
Eventpoints	3-11
Interactions Between Conditions, Ignore Counts, etc.	3-14
Breakpoints	3-14
Monitorpoints	3-14
Patchpoints	3-15
Tracepoints	3-15
Heappoints	3-15
Watchpoints	3-16

Syscallpoints	3-18
Timing	3-18
Timing Info	3-19
underflow	3-19
maximum recursion depth	3-20
current recursion depth	3-20
Signals	3-20
Exceptions	3-21
Restarting a Program	3-22
Restart Mechanism	3-23
Restart Information	3-23
Restart Macros	3-24
Exited and Terminated Processes	3-24
Process States	3-25
Operations While the Process Is Executing	3-25
Examining Your Program	3-26
Expression Evaluation	3-27
Ada Expressions	3-28
C Expressions	3-29
C++ Expressions	3-30
Fortran Expressions	3-30
Function Calls	3-31
Overloading	3-31
Program Counter	3-32
Context	3-32
Scope	3-32
Stack	3-33
Current Frame	3-33
Registers	3-34
Inline Subprograms	3-34
Interesting Subprograms	3-35
Monitor Window	3-36
Debugging the Heap	3-37
Levels and Common Errors	3-38
Fences	3-39
Hardware Overrun Protection	3-40
Retained Free Blocks	3-41
Heap Check	3-41
Leak Detection	3-42
Branch Tracking	3-43
Errors	3-43
Command Streams	3-44
Interrupting the Debugger	3-45
Macros	3-45
Convenience Variables	3-46
Smart Printing	3-46
Logging	3-46
Value History	3-47
Command History	3-47
Initialization Files	3-47
Optimization	3-47
Debugging Ada Programs	3-48
Packages	3-48
Exception Handling	3-49

Multithreaded Programs	3-49
Stopping	3-49
Executing	3-50
Thread Tags	3-50
Protected thread tag	3-51
CUDA Debugging	3-51
Limitations and Warnings	3-53
DWARF Debug Information	3-53
Setuid Programs	3-54
Attach Permissions	3-54
Architecture Interoperability	3-55
Frequency-Based Scheduler	3-55
NightTrace Daemon	3-56
Memory Mapped I/O	3-56
Blocking Interrupts	3-56
Debugging with Shared Libraries	3-56

Chapter 4 Tutorials

General Graphical Tutorial	4-1
Topical Tutorials	4-29
Thread Tags Tutorial	4-29
Tracing Tutorial	4-35

Chapter 5 Invoking NightView

nview	5-1
nview-save-core-file	5-3

Chapter 6 Command-Line Interface

Command Syntax	6-1
Selecting Overloaded Entities	6-3
Special Expression Syntax	6-5
Predefined Convenience Variables	6-7
IA-32 Registers	6-8
AMD64 Registers	6-11
CUDA Registers	6-15
Casts of Vector Registers	6-16
Special Ininsics	6-16
__eventpoint_data__	6-16
Location Specifiers	6-17
Qualifier Specifiers	6-19
Eventpoint Specifiers	6-20
Regular Expressions	6-21
Wildcard Patterns	6-23
Repeating Commands	6-24
Replying to Debugger Questions	6-25
Controlling the Debugger	6-25
Quitting NightView	6-26
quit.	6-26
Managing Dialogues	6-27
login.	6-27

debug	6-29
nodebug	6-30
set-debug-file-directory	6-31
translate-object-file	6-31
logout	6-33
on dialogue	6-33
apply on dialogue.	6-35
Dialogue Input and Output	6-36
!	6-36
set-show	6-37
show	6-38
Managing Processes	6-40
run	6-40
rerun	6-41
set-notify	6-42
notify	6-43
attach	6-43
detach	6-44
kill	6-45
symbol-file	6-45
core-file	6-46
save-core-file	6-47
exec-file	6-48
on program	6-50
apply on program	6-52
on restart	6-52
checkpoint	6-53
family	6-54
set-children	6-55
set-exit	6-56
set-shared-lib-update	6-56
wait	6-57
mreserve	6-57
Heap Debugging	6-58
heapdebug	6-59
Setting Modes	6-65
set-echo	6-66
set-log	6-66
set-language	6-67
set-qualifier	6-68
set-history	6-68
set-limits	6-69
set-prompt	6-69
set-format-hex	6-71
set-terminator	6-71
set-safety	6-71
set-restart	6-72
set-local	6-72
set-patch-area-size	6-73
interest	6-74
set-auto-frame	6-77
set-overload	6-77
set-search	6-77
set-editor	6-78

set-preallocate	6-78
set-resume	6-79
set-download	6-80
set-disassembly	6-81
set-branch-tracking	6-82
set-futurepoints	6-82
set-cuda	6-83
set-cuda-software-preemption	6-83
set-cuda-memcheck	6-83
Debugger Environment Control	6-85
cd	6-85
pwd	6-85
Source Files	6-86
Viewing and Editing Source Files	6-86
list	6-86
directory	6-88
edit	6-89
Searching	6-90
forward-search	6-90
reverse-search	6-90
Source Line Decorations	6-92
Examining and Modifying	6-95
backtrace	6-95
print	6-96
set	6-99
x	6-99
output	6-103
echo	6-104
data-display	6-104
memory-display	6-105
display	6-106
undisplay	6-107
redisplay	6-108
printf	6-108
load	6-109
branch-history	6-110
Manipulating Eventpoints	6-110
Eventpoint Modifiers	6-113
name	6-113
breakpoint	6-114
patchpoint	6-116
set-trace	6-119
tracepoint	6-119
monitorpoint	6-121
heappoint	6-123
mcontrol	6-124
clear	6-125
commands	6-126
condition	6-127
delete	6-128
disable	6-128
enable	6-129
ignore	6-130
modify-eventpoint	6-131

tbreak	6-131
tpatch.	6-132
watchpoint.	6-133
syscallpoint	6-135
Controlling Execution	6-137
set-run-mode	6-138
continue.	6-138
resume	6-139
step	6-141
next	6-142
stepi.	6-144
nexti	6-145
finish	6-146
stop	6-147
jump	6-148
signal.	6-149
handle	6-150
handle Signals	6-152
handle Ada Exceptions	6-153
handle C++ Exceptions	6-154
Selecting Context.	6-154
frame	6-154
up	6-156
down	6-156
select-context	6-157
Miscellaneous Commands	6-159
help	6-159
refresh	6-160
shell.	6-161
source	6-161
declare-thread-tag	6-162
set-tag	6-163
set-thread-name.	6-163
delay	6-164
Info Commands	6-165
Status Information.	6-166
info log.	6-167
info eventpoint.	6-167
info breakpoint.	6-168
info tracepoint	6-169
info patchpoint.	6-170
info monitorpoint.	6-171
info heappoint	6-172
info watchpoint	6-173
info syscallpoint.	6-174
info frame	6-175
info directories.	6-176
info convenience	6-176
info display	6-176
info history.	6-176
info limits.	6-177
info registers	6-177
info signal	6-178
info process	6-178

info memory	6-179
info dialogue	6-182
info family	6-183
info name	6-183
info on dialogue	6-184
info on program	6-184
info on restart	6-184
info exception	6-185
info threads	6-186
info timing	6-187
heapcheck	6-187
Symbol Table Information	6-188
info args	6-189
info locals	6-189
info variables	6-189
info address	6-190
info sources	6-190
info functions	6-190
info types	6-191
info whatis	6-191
info representation	6-191
info declaration	6-192
info files	6-192
info line	6-192
Defining and Using Macros	6-193
define	6-193
Referencing Macros	6-196
info macros	6-198
Smart Printing	6-199
smart-print	6-200
replace Smart Printers	6-201
struct Smart Printers	6-201
container Smart Printers	6-202
Smart Printing Limitations	6-203
Predefined Smart Printers	6-203
Intrinsics for Smart Printing	6-204
Timing	6-205
timing	6-205

Chapter 7 Simple Full-Screen Interface

Using the Simple Full-Screen Interface	7-1
Editing Commands in the Simple Full-Screen Interface	7-2
Monitor Window - Simple Full-Screen	7-2

Chapter 8 Graphical User Interface

NightView GUI Concepts	8-1
GUI Online Help	8-2
Context-Sensitive Help	8-2
Help Buttons	8-3
Help Command	8-3
Dialogues and Dialog Boxes	8-3

Context Menu	8-3
Current Process	8-4
GUI Configuration	8-4
Main Window.	8-4
Menu bar.	8-5
File Menu.	8-5
View Menu.	8-6
Shell Menu.	8-10
Process Menu.	8-11
Source Menu	8-12
Eventpoint Menu	8-15
Data Menu	8-17
Tools Menu	8-20
Help Menu	8-21
Toolbars.	8-23
Command Toolbar	8-23
Process Toolbar	8-24
Run Mode Toolbar	8-26
Timing Toolbar	8-26
Eventpoint Toolbar	8-27
Value Toolbar.	8-29
Source Display Toolbar	8-29
Status Bar	8-29
List of Shortcuts	8-31
Main Window Dialog Boxes.	8-32
Run Program in Shell Dialog Box.	8-33
Attach Dialog Box	8-33
Source Selection Dialog Box.	8-34
File Selection Dialog Box	8-35
Eventpoint Dialog Boxes.	8-35
System Call Selection Dialog	8-41
Debug Heap Dialog Box	8-41
Remote Login Dialog Box.	8-43
Remote Login General Page.	8-43
Remote Login Advanced Page	8-44
Remote Login Action Buttons	8-45
Preferences Dialog Box.	8-46
Preferences General Page.	8-47
Safety.	8-47
Automatically Resume On	8-47
Searching	8-47
Data Panel	8-48
Eventpoint Panel	8-48
CUDA	8-48
Display Limits.	8-48
Source Panel Keystrokes.	8-48
Source File Size.	8-48
Preferences Source and Assembly Page.	8-49
Source Display	8-49
Disassembly	8-50
Source Editor.	8-50
Preferences Fonts Page.	8-50
Global NightStar Fonts	8-52
My NightView Fonts.	8-53

Effective NightView Fonts	8-53
NightStar Global Fonts Dialog	8-54
Changes Fonts For...	8-55
Apply Fonts To...	8-55
Set Default Fonts	
Set Panel Fonts	8-55
Save & Close	8-55
Save.	8-55
Cancel	8-56
Help.	8-56
Preferences Terminal Page	8-56
Preferences Advanced Page	8-57
Remote Object File Cache	8-58
Memory Preallocation.	8-58
Restart.	8-58
Value History	8-58
Expression Evaluation Automatic Overloading	8-58
Future Eventpoints	8-58
Restore Defaults	8-58
Preferences Binary Viewer Page	8-59
Size and Format	8-59
Search	8-59
Background Colors	8-59
Process Settings Dialog Box	8-59
Process Settings General Page	8-60
Debug Children.	8-60
Set Run Mode	8-60
Branch Tracking	8-60
Stop Before Exiting	8-60
Expression Language	8-60
Refresh debug info when shared libs change	8-61
Program.	8-61
Process Settings Interest Page	8-61
Process Settings Signals Page	8-61
Process Settings Ada Exceptions Page.	8-61
Process Settings C++ Exceptions Page	8-61
Browse Ada Exceptions Dialog Box.	8-62
Select a Type Dialog Box	8-62
Rename Page Dialog Box	8-62
Print Dialog Box	8-62
List Location Dialog Box	8-62
Eventpoint Panel Update Interval Dialog Box	8-63
Set New Timing Start/End Dialog Box	8-63
Set Timing for Function Dialog Box.	8-63
Panels.	8-64
Find Bar	8-65
Source Panel	8-65
Source Panel Target Line.	8-66
Source Panel Expression Tooltip.	8-66
Source Panel Context Menu	8-67
Source Panel Tracking.	8-71
Source Panel Keystrokes.	8-72
Shell Panel	8-73
Message Panel	8-74

Eventpoint Panel	8-74
Context Panel	8-77
Locals Panel	8-77
Monitor Panel	8-77
Registers Panel	8-77
Data Panel	8-78
Monitor Bar	8-78
Data Items	8-78
Expression Data Item	8-81
Local Variables Data Item	8-82
Registers Data Item	8-82
Register Data Item	8-82
Stack Data Item	8-84
Branch History Data Item	8-85
Threads Data Item	8-86
Processes Data Item	8-87
Shells Data Item	8-87
Heap Information Data Item	8-88
Heap Errors Data Item	8-89
Leak Sets / Still Allocated Sets Data Items	8-89
Block Data Item	8-89
Monitorpoint Values Data Item	8-90
Data Panel Context Menu	8-90
Data Panel Dialog Boxes	8-99
Data Panel Item Dialog Box	8-100
Data Panel Add Expression	8-100
Data Panel Add Heap Errors	8-101
Data Panel Add Heap Leaks	8-101
Data Panel Add Still Allocated Blocks	8-101
Data Panel Call Stack Frames	8-102
Data Panel Edit Expression	8-102
Data Panel Expand Tree	8-102
Data Panel Describe	8-102
Data Panel Load Layout	8-102
Data Panel Pointer Array Dimension	8-103
Data Panel Save Layout	8-103
Data Panel Save Snapshot	8-103
Data Panel Subscript Array	8-104
Data Panel Subscript Enum Array	8-104
Data Panel Linked List Expression Dialog	8-104
Data Panel Condition Filter Expression Dialog	8-105
Monitorpoint Update Interval Dialog Box	8-106
CUDA Coordinates Panel	8-107
CUDA Lanes Panel	8-107
CUDA Lanes Context Menu	8-108
CUDA Warp Locals Panel	8-109
CUDA Warp Locals Panel Context Menu	8-109
Memory Segments Panel	8-110
Memory Segments Panel Context Menu	8-111
Binary Viewer Panel	8-112
Binary Viewer Searching	8-114
Binary Viewer Editing	8-115
Binary Viewer Panel Context Menu	8-115
Timing Regions Panel	8-117

Timing Regions Context Menu	8-117
Timing Results Panel	8-118
Timing Results Context Menu	8-119
Help Window	8-122

Appendix A Summary of Commands

Appendix B Quick Reference Guide

Invoking NightView	B-1
Controlling the Debugger	B-1
Quitting NightView	B-2
Managing Dialogues	B-2
Dialogue Input and Output	B-2
Managing Processes	B-3
Heap Debugging	B-4
Setting Modes	B-4
Debugger Environment Control	B-6
Source Files	B-6
Viewing and Editing Source Files	B-6
Searching	B-7
Examining and Modifying	B-7
Manipulating Eventpoints	B-8
Controlling Execution	B-10
Selecting Context	B-10
Miscellaneous Commands	B-11
Info Commands	B-11
Status Information	B-12
Symbol Table Information	B-13
Defining and Using Macros	B-14
Smart Printing	B-14

Appendix C Implementation Overview

Reporting Bugs	C-2

Appendix D Tutorial Files

C Files	D-1
msg.h	D-1
main.c	D-1
parent.c	D-2
child.c	D-2
Fortran Files	D-3
main.f	D-3
parent.f	D-4
child.f	D-4
ftint.c	D-5
Ada Files	D-6
main.a	D-6
parent.a	D-6
child.a	D-7

Glossary

Index

Tables

Table 3-1. Eventpoint Summary	3-13
Table 6-1. Special '\$' Constructs	6-5
Table 6-2. Predefined Convenience Variables	6-7
Table 6-3. IA-32 Registers (iHawk Series 860)	6-8
Table 6-4. AMD64 Registers (iHawk Series 870)	6-12
Table 6-5. Regular Expressions.	6-21
Table 6-6. Wildcard Patterns	6-23
Table 6-7. Source Line Decorations	6-92
Table 6-8. Eventpoint Commands.	6-111

A Quick Start

This chapter is for people who want to start using the command-line version of the debugger before reading the whole manual. You may also be interested in the graphical-user-interface (GUI) version of this chapter in Chapter 2 [A Quick Start - GUI] on page 2-1. There is a more thorough tutorial in Chapter 4 [Tutorials] on page 4-1.

If you are familiar with the GNU debugger, `gdb`TM, you should have very few problems with NightView. The commands are almost all identical. The biggest difference between NightView and other debuggers is how you tell NightView what program to debug and how you start that program.

If you get any errors, the error message tells which section of the manual can help you determine what went wrong. At any time, you can ask the debugger to display help on an error message by mentioning that section's name as the argument to the **help** command (see “help” on page 6-154).

The rest of this chapter goes through a sample debug session on a small program. Feel free to dive right into the debugger. If you get into trouble, use the **help** command to get out of it.

Sample Program

This section lists the program used as an example through the remainder of the chapter. The program does not have any bugs in it; it will be used to show how to run a program, set breakpoints, look at variables, etc. You can copy this file from `/usr/lib/NightView/fact.c` into your own directory. The following program is in the file `fact.c`:

```

1  #include <stdio.h>
2
3  static int factorial(x)
4      int x;
5  {
6      if (x <= 1) {
7          return 1;
8      } else {
9          return x * factorial(x-1);
10     }
11 }
12
13 void
14 main(argc, argv)
15     int argc;
16     char ** argv;
17 {
18     int i, errors;
19     for (i = 1; i < argc; ++i) {
20         long xl;
21         int x;
22         int answer;
23         char * ends = NULL;
24         xl = strtol(argv[i], &ends, 10);
25         x = (int)xl;
26         answer = factorial(x);
27         printf("factorial(%d) == %d\n", x, answer);
28     }
29     exit(0);
30 }

```

The remainder of this chapter assumes that you compiled **fact.c** and put the resulting executable in **fact**:

```
cc -g -o fact fact.c
```

Starting Up

You can start NightView with or without a program name. If you start it with a program name, NightView runs the program in a dialogue shell (see “Dialogues” on page 3-6). If you start NightView without a program name or you want to debug another program, you must execute the program with the **run** command (see “run” on page 6-39) in a dialogue shell.

Below is an example of starting up the debugger with a program name and a program argument. Note that throughout the quick start, the version and the link time might not match exactly for your version of NightView. Also, some of the shell output and other messages may not come out exactly as shown. Some messages might not appear, or additional messages might appear, depending on your environment.


```
$ nview -nogui ./fact 7
NightView 7.7, linked Fri Jun 21 17:13:19 EDT 2018

In case of confusion, type "help"
```

Note that you invoked NightView with a program name argument `./fact` and program argument `7`. NightView responded with information about the debugger.

```
New process: local:2347 parent pid: 2340
Process local:2347 is executing /users/bob/fact.
Reading symbols from /users/bob/fact...done
Executable file set to /users/bob/fact
/usr/lib/NightView/ReadyToDebug
$ /usr/lib/NightView /ReadyToDebug
$ ./fact 7
(local)
```

NightView always runs a special program, `/usr/lib/NightView/ReadyToDebug`. This program helps NightView synchronize with the shell. That's why you see that line in the output. You might see only one echo of `/usr/lib/NightView/ReadyToDebug`, depending on how quickly the dialogue shell starts. The dollar signs (`"$"`) are prompts from the shell.

NightView automatically created a dialogue named `local`; it also displayed the string `local` as the prompt, showing that by default, commands apply to that dialogue (or the processes running in that dialogue).

The debugger waited for the new program to get started. Because sending input to a dialogue is just like typing commands to a shell (the dialogue is really running the same shell program you normally use), this caused the `fact` program to be executed with the single argument `7`.

If the `fact` program had required input, you would have used the `!` command to send the input to the program. See `"!"` on page 6-35.

When the dialogue executed the program, NightView got control and informed you that a new process was just started in dialogue `local` and told you that the process id was 2347.

Because this is the only program running in dialogue `local`, you do not have to do anything special to cause any commands you type to refer to this process; the default qualifier is already set to `local`, so commands will automatically apply to the one process running there.

Getting Help

Next you will enter a bogus command. Note that throughout this section, the help text and display size may not exactly match your NightView session.

```
(local) foo
Error: Unrecognized command "foo". [E-command_proc003]
```

NightView responded to the bogus command with an error message and an error code (`[E-command_proc003]`).

Now get NightView to tell you more about the error message.

```
(local) help
E-command_proc003:
Unrecognized command "string".
```

STRING is not a valid NightView command. See "Summary of Commands".

You typed **help** without any arguments to see more information about the error message. NightView showed the extended error information.

In the command-line and simple screen interfaces, online help is available only for error messages. Consult a printed manual or view the online help with NightView's graphical user interface.

If you are familiar with **gdb**, the remainder of this chapter will be fairly boring because (once you get the program started) NightView and **gdb** look very much alike (at least for all the commands demonstrated in this simple example).

Setting a Breakpoint

You will now use the **list** command to look at the source.

```
(local) l 1
1      | #include <stdio.h>
2      |
3      | static int factorial(x)
4      |     int x;
5      | {
6 *    |     if (x <= 1) {
7 *    |         return 1;
8      |     } else {
9 *    |         return x * factorial(x-1);
10     |     }
(local)
```

You told the **list** command (abbreviated to **l** in this example) to list at line 1.

You now decide where you want to set a breakpoint. An interesting spot in this program is the `return` statement in the recursive routine `factorial` where it is about to start backing out of the recursive calls.

```
(local) b 7
local:2347 Breakpoint 1 set at fact.c:7
(local)
```

The `return` was on line 7, so you used the **breakpoint** command (abbreviated to **b**) to set a breakpoint on line 7.

Complete descriptions of the commands you used here appear in “list” on page 6-86 and “breakpoint” on page 6-114.

Finishing up

Now run the program until it reaches the breakpoint.

```
(local) c
local:2347: at Breakpoint 1, 0x100026fc in factorial(int
x = 1) at fact.c line 7
7 B=|          return 1;
(local)
```

You used the **continue** command (abbreviated to **c**) without any arguments. This told the program to start running. It ran until it hit the breakpoint that you had set on line 7. Note that your process ID and addresses will differ.

Now look at the call stack.

```
(local) bt
#0  0x100026fc  in factorial(int x = 1) at fact.c line 7
#1  0x1000271c  in factorial(int x = 2) at fact.c line 9
#2  0x1000271c  in factorial(int x = 3) at fact.c line 9
#3  0x1000271c  in factorial(int x = 4) at fact.c line 9
#4  0x1000271c  in factorial(int x = 5) at fact.c line 9
#5  0x1000271c  in factorial(int x = 6) at fact.c line 9
#6  0x1000271c  in factorial(int x = 7) at fact.c line 9
#7  0x10002784  in main(int argc = 2,
char **argv = 0x2ff7eaec)
                        at fact.c line 26
(local)
```

You used the **bt (backtrace)** command to display the call stack. You saw all the expected recursive calls (see “backtrace” on page 6-92).

Now look at the value of the variable `x`.

```
(local) p x
$1: x = 1
(local)
```

You used the **p (print)** command to print the variable `x`, verifying that it was equal to 1.

Now finish running the program.

NOTE

If your system has debug information installed for system libraries, the process may appear to be stopped in the `_exit()` library routine after the command below. If so, enter the command **up** until the debugger reports that the process is in `main`.

```
(local) c
factorial(7) == 5040
Process local:2347 is about to exit normally
#0  0x100027ac  in main(int argc = 2,
unsigned char **argv = 0x2ff7eaec)
        at fact.c line 29
29 <>|      exit(0);
(local)
```

You used the **c** (**continue**) command to allow the process to run to completion.

Exit from NightView.

```
(local) q
Kill all processes being debugged? y
You are now leaving NightView...
Process local:2347 exited normally
Dialogue local has exited.
$
```

Finally you typed **q** (**quit**) to leave the debugger. The **fact** program had not fully exited, so NightView prompted, asking if the program should be killed. You responded with **y**, and the sample session ended. The commands used in this section appear in “continue” on page 6-138, “backtrace” on page 6-95, “print” on page 6-96, and “quit” on page 6-26.

2

A Quick Start - GUI

This chapter is for people who want to start using the graphical-user-interface (GUI) version of the debugger before reading the whole manual. You may also be interested in the command-line version of this chapter in Chapter 1 [A Quick Start] on page 1-1. There is a more thorough tutorial in Chapter 4 [Tutorials] on page 4-1.

In this manual, the words click, drag, press, and select always refer to mouse button 1.

This entire manual is available through the online help system built into the debugger. If you get any errors, the error message tells which section of the manual can help you determine what went wrong. At any time, you can ask the debugger to display any section of the manual by clicking on the **Help** menu or using the **H** mnemonic. See “Help Menu” on page 8-20. Click on the **NightView User's Guide** menu item or use the **U** mnemonic. NightView puts up a Help Window that displays the table of contents for the manual. See “Help Window” on page 8-118. You can read this manual section by clicking on **A Quick Start - GUI**.

The rest of this chapter goes through a sample debug session on a small program. Feel free to dive right into the debugger. If you get into trouble, use the **Help** menu to get out of it.

Sample Program - GUI

This section lists the program used as an example through the remainder of the chapter. The program does not have any bugs in it; it will be used to show how to run a program, set breakpoints, look at variables, etc. You can copy this file from **/usr/lib/NightView/fact.c** into your own directory. The following program is in the file **fact.c**:

```

1  #include <stdio.h>
2
3  static int factorial(x)
4      int x;
5  {
6      if (x <= 1) {
7          return 1;
8      } else {
9          return x * factorial(x-1);
10     }
11 }
12
13 void
14 main(argc, argv)
15     int argc;
16     char ** argv;
17 {
18     int i, errors;
19     for (i = 1; i < argc; ++i) {
20         long xl;
21         int x;
22         int answer;
23         char * ends = NULL;
24         xl = strtol(argv[i], &ends, 10);
25         x = (int)xl;
26         answer = factorial(x);
27         printf("factorial(%d) == %d\n", x, answer);
28     }
29     exit(0);
30 }

```

The remainder of this chapter assumes that you compiled **fact.c** and put the resulting executable in **fact**:

```
cc -g -o fact fact.c
```

Starting Up - GUI

You can start NightView with or without a program name and arguments. If you start it with a program name, NightView begins debugging the program immediately. If you start NightView without a program name, or you want to debug another program, you may run the program with the Run menu item in the **Process** menu, or by typing in the shell in a shell panel. See “Shell Panel” on page 8-70. In either case, the program is run in a dialogue shell (see “Dialogues” on page 3-4).

Below is an example of starting up the debugger with a program name and a program argument. Note that throughout the quick start, the version and the link time might not match exactly for your version of NightView. Also, some of the messages might not come out exactly as shown. Some messages might not appear, or additional messages might appear, depending on your environment.

```
$ nview ./fact 7
```

NightView displays the main window. See “Main Window” on page 8-4.

Starting the debugger with the program name `./fact` and argument `7` sent the line `./fact 7` to the `local` dialogue and caused the debugger to wait for the new program to get started. Because sending input to a dialogue is just like typing commands to a shell (the dialogue is really running the same shell program you normally use), this caused the `fact` program to be executed with the single argument `7`.

If the `fact` program had required input, you would have typed the input into a shell panel. See “Shell Panel” on page 8-70.

The message panel (see “Message Panel” on page 8-71) contains a message like the following:

```
New process: local:2347      parent pid: 2340
Process local:2347 is executing /users/bob/fact.
Reading symbols from /users/bob/fact...done
Executable file set to
/users/bob/fact
```

When the dialogue executed the program, NightView got control and informed you that a new process was just started in dialogue `local` and told you that the process id was 2347.

The status bar at the bottom of the window displays the program name, `fact`, the dialogue name and PID, `local:2347`, and the state, `Stopped for exec`. See “Status Bar” on page 8-28. The source panel title bar displays the program name, the dialogue name and PID, and the name of the source file, `fact.c`. The source code from file `fact.c` appears in the source panel, centered around `main`. See “Source Panel” on page 8-63.

Getting Help - GUI

Next you will enter a bogus command. Note that throughout this section, the help text and display size may not exactly match your NightView session.

The command toolbar is labeled **Command:**. Click in the combo box of the command toolbox (see “Command Toolbar” on page 8-22) and issue the following command:

```
foo
```

Press **Return** to enter the command.

NightView responded to the bogus command with the following message and error code:

```
Error: Unrecognized command "foo". [E-command_proc003]
```

Now get NightView to tell you more about the error message. Click on the **Help** menu or use the **H** mnemonic. See “Help Menu” on page 8-20. Click on the **On Last Error** menu item or use the **E** mnemonic. NightView puts up a Help Window that displays the following extended error information:

E-command_proc003

MESSAGE

ERROR: Unrecognized command "*string*".

EXPLANATION

string is not a valid NightView command. See Summary of Commands.

Next, dismiss the Help Window by selecting **Exit** from the **File** menu. See “Help Window” on page 8-118.

Next you will read about the **list** command. Click on the **Help** menu or use the **H** mnemonic. See “Help Menu” on page 8-20. Click on the **On Commands** menu item or use the **m** mnemonic. NightView puts up the following Help Window with a menu of NightView commands.

Summary of Commands

This section gives a summary of all the commands in NightView. The table is organized alphabetically by command. The abbreviations for the commands are included with the corresponding commands, rather than alphabetically.

Also, remember that you can abbreviate commands by using a unique prefix.

!

Pass input to a dialogue.

apply on dialogue

Execute on dialogue commands for existing dialogues.

(etc.)

Most of the information would not fit on your display. The Help Window showed this by having only a small thumb or slider on the vertical scroll bar. Scroll down to the **list** command by moving the thumb or by clicking on the arrow heads of the vertical scroll bar. Click on the **list** command. NightView displayed the following Help Window with

information about the **list** command.

list

List a source file. This command has many forms, which are summarized below.

list *where-spec*

List ten lines centered on the line specified by *where-spec*.

list *where-spec1*, *where-spec2*

List the lines beginning with *where-spec1* up to and including the *where-spec2* line.

(etc.)

To see more about the **list** command, you could move the thumb or click on the arrow heads of the vertical scroll bar. However, rather than reading more, you make the Help Window go away by selecting **Exit** from the **File** menu.

Setting a Breakpoint - GUI

You now decide where you want to set a breakpoint. An interesting spot in this program is the `return` statement in the recursive routine `factorial` where it is about to start backing out of the recursive calls.

Right-click on the line with the `return` statement (line 7) in the source panel. The line becomes highlighted and a context menu appears. See “Source Panel Keystrokes” on page 8-69. Select the first item, **Set Simple Breakpoint**.

The source line decoration beside line 7 is now a stop sign  to indicate a breakpoint.

See “breakpoint” on page 6-108 and “Source Line Decorations” on page 6-88. The eventpoint panel now has an entry for the breakpoint.

The message panel shows:

```
local:2347 Breakpoint 1 set at fact.c:7
```

Finishing up - GUI

Now you want to run the program until it reaches the breakpoint. Click on the **Resume** button in the process toolbar.  See “Process Toolbar” on page 8-22.

Clicking on **Resume** told the program to start running. It ran until it hit the breakpoint that you had set on line 7. The source line decoration beside line 7 is now a stop sign overlaid with a triangle pointing to the right  to indicate where execution will resume.

NightView responds with:

```
local:2347: at Breakpoint 1, 0x100026fc in factorial(int
x = 1) at fact.c line 7
```

Note that your process ID and addresses will differ. The status bar indicates the process is **Stopped at breakpoint 1**. Now look at the call stack. The context panel and the locals panel are in the same area with tabs below them. Click on the **Context** tab. The context panel has an entry for each frame on the stack, displayed in tree form. See “Context Panel” on page 8-74. You see all the expected recursive calls. Scroll to the bottom of the panel. One of the icons is an arrowhead pointing down.  Click that icon to show more stack frames, until you see the call to `main`. Then scroll to the top again and click on the first frame.

Now look at the local variables. Click on the **Locals** tab. You see the local variables displayed in tree form. In this case, there is only one local variable, `x`. The locals panel tracks the current context, which you set when you clicked in the context panel. The value of `x` in this frame is 1. See “Locals Panel” on page 8-74.

Now finish running the program. Click on the **Resume** button. See “Process Toolbar” on page 8-22.

This allowed the process to run to completion. The program printed a message, which appeared in the message panel:

```
factorial(7) == 5040
```

NOTE

If your system has debug information installed for system libraries, the process may appear to be stopped in the `_exit()` library routine. If so, click the Up button  until the debugger reports that the process is in `main`.

NightView showed the call to `exit(0)` in the source panel and displayed the following message in the message panel.

```
Process local:2347 is about to exit normally
```

Exit from NightView by selecting the **File** menu. See “File Menu” on page 8-4. Click on **File** or use the **F** mnemonic. Click on the **Exit NightView** menu item or use the **X** mnemonic.

NightView responds with a warning dialog box. The warning dialog box says:

```
Kill all processes being debugged?
```

Finally you click on the **OK** button to leave the debugger. The **fact** program had not fully exited, so NightView prompted, asking if the program should be killed. You responded by clicking **OK**, and the sample session ended.

3

Concepts

This section describes concepts you will need to understand in order to use the debugger effectively.

Many of the concepts described in this section are also defined in the glossary. The glossary is an alphabetical list of the concepts — the description here is organized hierarchically.

Debugging

The term *debugger* is actually a misnomer. A debugger does not remove bugs from your program. Instead, it is a tool to help you monitor and examine your program so that you can find the bugs and remove them yourself.

A debugger primarily lets you do two things:

1. start and stop the execution of your program; and,
2. examine and alter the contents of the program's memory.

There are many ways to do these things, so there are lots of debugger commands. Also, some of the commands control the debugger itself.

NightView is a symbolic debugger. That means that you can talk about your program using the same high-level language constructs that you use when you write programs. You can refer to variables, expressions and procedures as they appear in your program source. You can also refer to source files and line numbers within those files. For example, you can tell your program to stop at a particular line. In order to use the symbolic capabilities of the debugger, you must compile and link your program with options that tell the compiler and linker to save the symbolic information along with your program.

Sometimes, you want to be able to debug at a lower level, referring to machine language instructions and registers. NightView lets you do that, too.

Accessing Files

During the course of debugging, NightView will likely have to access a number of files: executable files for programs being debugged, source files for those programs, and possibly object and library files. Those files must all reside, or be accessible from, the system on which NightView is executing.

If you are debugging processes running on some other system, you will probably want to have some of that system's files mounted via NFS™ on the system running NightView. Furthermore, your debugging will probably go much easier if the pathnames to those files (especially the executables) are the same on both systems. This will allow NightView to find the executable files automatically most of the time. See “Finding Your Program” on page 3-9. If the pathnames of the executable files are different, you can use the **translate-object-file** command to tell how to translate the names. See “translate-object-file” on page 6-29. In addition, remote files can be specified by using the form *user@host:/path*. See “Remote File Access” on page 3-7.

Programs and Processes

It is necessary to distinguish between a *program* and a *process*. A *program* is something that you write, compile and link to form a program file. A *process* is an instance of execution of a program. There may be several processes running the same program.

Multiple Processes

The most typical use for NightView is debugging a single program running as a single process, but NightView can also be used to debug an *application* consisting of multiple processes, so the debugger has ways to describe multiple processes. If you come to a section of the manual that describes multiple processes, and you are only debugging one process, you can usually just ignore the parts about multiple processes.

You may inadvertently create multiple processes, even though you only want to debug one. This may happen if your program *forks*. For example, your program may call **system**. This call works by using the **fork** service to create another process, which then runs a shell. A process created this way is called a *child process*. Because NightView has the capability of debugging child processes, you are notified when this happens. If you don't want to debug the child process, then you should **detach** from it, which allows it to run without further interference from the debugger. See “detach” on page 6-41. If you know in advance that you don't want to debug any child processes, you can use the **set-children** command to specify this. See “set-children” on page 6-52.

If you use pipelines in the dialogue shell, or invoke shell scripts which call many other programs, you are likely to get multiple processes which you are not interested in debugging. (Dialogues are described in a later section, see “Dialogues” on page 3-4.) Again, if you don't want to debug those other processes, you should detach from them.

Another way to determine which processes are debugged is to use **debug** and **nodebug**, which let you describe which processes you want to debug by their program names. See “nodebug” on page 6-28.

Families

One of the handy things NightView lets you do is group processes together into *families*.

You do this by giving the family a name and telling the debugger what processes are in that family. For example, you might have several processes executing the same program, and you might want to set a breakpoint at the same source line in all of them. You could define a family containing all of the processes and then use that family name with the **breakpoint** command. See “family” on page 6-51.

Attaching

Sometimes you want to debug a process that is already running, rather than starting up a new process running the same program. You can do this with the **attach** command (see “attach” on page 6-40) or with the Attach Dialog Box (see “Attach Dialog Box” on page 8-31.)

In order to attach to a process, you must know its process identifier (or PID). You can get a list of running processes and their PIDs by clicking on the **Attach** menu item in the **Process** menu (see “Process Menu” on page 8-9) to bring up the Attach Dialog Box.

As an alternative, you can run the **ps (1)** program. You can use the **shell** command (see “shell” on page 6-152) to run **ps (1)**. If you want to attach to a process running on another machine, you may have to use the remote shell command (`/usr/bin/rsh`) to run **ps (1)** on the right machine.

Once you have attached to a process, you can debug it in the same way you would debug a process started normally from a dialogue.

For the security restrictions on **attach**, see “Attach Permissions” on page 3-49.

If the process to which you attach is stopped (<CONTROL Z> stops a foreground process in most shells), then the attach will not take effect until the process is continued from the shell.

Detaching

Detaching a process is the inverse of attaching one. When you detach a process it starts running independently of the debugger. Nothing it does will get the debugger's attention. Any children it forks will also be ignored by the debugger. You have to explicitly attach to the process again to make the debugger notice it.

Detaching from an exited or terminated process completely removes the process from the system. See “Exited and Terminated Processes” on page 3-22. Detaching from or killing a pseudo-process associated with a core file (see “Core Files” on page 3-4) is the only way to make that pseudo-process go away.

Detaching from a process causes NightView to forget all the eventpoint settings and other information it remembers about the process.

When detaching from a process under the RedHawk kernel, any patches installed in the program for patchpoints, etc., will be left in the process and will continue to apply. When detaching from a process with NightStar LX on a non-RedHawk kernel, any patches

installed will be removed before the detach occurs; they will not continue to apply after the patch.

Attaching to a process from which you have detached is not supported on Linux. Avoid detaching from processes unless you are sure you will not want to debug them further.

Core Files

A core file is a snapshot image of a process created by the system when the process aborts (typical reasons for creating a core file include referencing an address outside the memory allocated to the process, dividing by zero, floating-point exceptions, etc.). NightView allows you to debug core files as well as processes (see “core-file” on page 6-43). Since a core file is not actually a running process, all you can do is look at it. None of the commands which require a running process will work on core files (for example, you cannot **continue** a core file and you cannot evaluate any expression containing a function call).

If a core file is from a process that used dynamic linking, the core file must be debugged on the same system where the process was running, otherwise information from the libraries may not match the core file.

Qualifiers

If you are not debugging multiple processes, you will probably never need to worry about command qualifiers, but for multiprocess debugging, they are essential. A qualifier is used to restrict a command so it operates only on specific processes. There is always a default qualifier in effect, but any command may be given an explicit qualifier.

Most qualified commands act as though the command was specified once for each process (for instance, the **breakpoint** command sets a separate breakpoint in each of the processes specified in its qualifier).

Some commands treat the qualifier in special ways, and other commands ignore the qualifier. Any special treatment is described in the section on each command.

Qualifiers are specified as a prefix on the command. The complete description may be found in “Command Syntax” on page 6-1 and “Qualifier Specifiers” on page 6-17.

Dialogues

Dialogues are one of the most important (and unique) concepts in NightView. Essentially, a *dialogue* is just an ordinary shell where you run commands as you would normally run them in the shell (in fact, you are running your normal shell), but in a dialogue, you have the opportunity to debug any or all of the programs you run in the dialogue shell. Most debuggers have special commands to tell the debugger which program to debug and what arguments to give it. In NightView, the way to debug a

program is to run it within a dialogue shell. This means you can debug a program that is a member of a pipe, or is invoked by some other program, and you can run the program in the debugger using the exact same invocation you would normally use outside the debugger. For instance, if your programs run under the control of the Frequency-Based Scheduler, you could invoke **rtcp** or NightSim™ from your dialogue.

The environment variable `NIGHTVIEW_ENV` is set to 1 within a dialogue shell. This allows you to alter the behavior of programs and scripts running in the dialogue shell. For example, you may wish to avoid running some programs in a shell initialization file when the shell is a dialogue shell.

NightView sets the `TERM` environment variable to `dumb` in the dialogue shell, to avoid problems with some shell programs.

Once the shell is started, you can change directory, set environment variables, or set **ulimit(1)** parameters just like a normal shell. Any processes you start in the dialogue will automatically be debugged, except for programs in the standard directories such as `/bin`. You may alter this default behavior using the **debug** and **nodebug** commands. See “debug” on page 6-27 and “nodebug” on page 6-28.

When you start a program in a dialogue shell, the debugger prints a message describing the new process that just started in the dialogue. The information printed includes the program name, the arguments it received on startup and the process identifier (PID). This new process is stopped immediately prior to executing any code. At this point you can decide what to do with the process (set breakpoints, etc.) and tell it to continue, or detach from it and let it run without being debugged.

At startup, NightView provides an initial dialogue named `local`. This initial dialogue shell inherits the current working directory and environment variables in existence at the time you started the debugger. However, the `LD_LIBRARY_PATH` environment variable is not inherited from the environment that started the debugger; instead, it is unset.

You may create additional dialogues at any time (see “login” on page 6-25). Multiple dialogues allow you to debug distributed systems of processes running on different computers. Each dialogue has a name. Unless you specify otherwise, the name of a dialogue is the host name of the system to which it is connected. You may use dialogue names in command qualifiers to tell NightView to which system you wish to talk, such as, when you want to run a command in a particular dialogue.

Dialogue I/O

You send input to a dialogue shell or to a program you are debugging in the dialogue by using the **!** command (see “!” on page 6-34) or the **run** command (see “run” on page 6-38). The qualifier on the command determines which dialogue receives the input data. In the graphical user interface, you can send input to a dialogue with a shell panel (see “Shell Panel” on page 8-70) for that dialogue.

Since each dialogue is a separate shell, the programs running in separate dialogues may generate output at any time. In the command-line interface, it would be confusing to have these print at any time. Instead, all the output generated by each dialogue shell and the programs running in it is logged by NightView. You can control this log using the **set-show** command (see “set-show” on page 6-35), and you can review the log with the **show** command (see “show” on page 6-36). In the graphical user interface, dialogue

output goes to the dialogue I/O area for that dialogue.

Real-Time Debugging

By running NightView on a development system and starting a dialogue on a real-time system you are debugging, you can minimize the impact of the debugger on the real-time system. Most of the debugger runs on the development system, and only a NightView control program and the dialogue shell run on the real-time system. You can also control the CPU, memory, and other resource allocations of debugger processes to help minimize the impact of the debugger on critical resources. See "Remote Dialogues" on page 3-6.

Monitorpoints provide a means of monitoring the value of variables in your program without stopping it. See "Monitorpoints" on page 3-12.

NightTrace™ is another tool you may find useful in debugging real-time programs. It allows you to gather performance information and record limited amounts of data with minimal overhead. NightView provides facilities for using NightTrace from within the debugger; see "Tracepoints" on page 3-13.

Remote Dialogues

A *remote dialogue* is a shell, controlled by NightView, running on a system other than the one on which NightView was initially invoked. We refer to the system where NightView was invoked as the "local system", while the system where the remote dialogue shell is running is referred to as the "target" or "remote system".

You may need to use a remote dialogue if:

- you need to debug programs running on multiple target systems simultaneously;
- your application uses most of the system's CPU or memory resources, leaving insufficient resources for NightView;
- the source files for your programs are not accessible on the target system;
- you do not wish to install all of NightView on the target system, perhaps to conserve disk space on the target;
- you need to reduce network traffic on the target system by eliminating NightView's GUI overhead;
- you need to reduce disk loading on the target system by eliminating NightView's reading of source and object files.

When you use a remote dialogue, a process with architecture-specific knowledge named **NightView.p** runs on the local system, while another process named **NightView.pp** runs on whichever system you are debugging on to access and control the processes you are debugging.

The following activities are performed on the local system:

- all user interaction, including command input/output and window manipulation and updating;
- reading source and object files, including reading and interpreting debug information in your program;
- evaluation of expressions in commands such as **print** and **x**, except that retrieving data from a debugged process (such as variable values) is performed on the remote system.

The activities performed on the remote system are limited to storing and retrieving data to and from a debugged process, controlling execution of a debugged process, and supplying target-dependent information to the local system portion of NightView. Additionally, NightView sometimes runs the C compiler on the target system to generate code for eventpoints. See “Eventpoints” on page 3-9.

You may wish to control how the target system allocates resources to NightView.p and the dialogue shell, both to prevent them from interfering with your application and to ensure that they get sufficient resources to give adequate response in NightView. You can control the allocation of CPU and memory resources as well as the scheduling policy and priority through either the **login** command or the remote login dialog. See “login” on page 6-25. See “Remote Login Dialog Box” on page 8-41.

Note that the parameters you specify for the remote dialogue will be inherited by processes you execute within that dialogue shell. You may wish to use the **run (1)** shell command when you run your application in the dialogue shell.

There are some things you need to be aware of when you use a remote dialogue. Because source files and debug information are read on the local system, those files (or copies of them) need to be accessible on the local system. This is particularly true of the executable program file, because that is where the debug information resides. When a debugged process **execs** a new program, NightView attempts to determine the location of the executable program file. See “Finding Your Program” on page 3-9. With a remote dialogue, NightView assumes that the pathname of the executable program file is the same (or locates identical files) on both systems. If this is not true, then NightView is not able to read debug information for that program until you specify the correct pathname with the **symbol-file** command or use object filename translations. See “symbol-file” on page 6-42. Also, see “translate-object-file” on page 6-29.

Creating a new dialogue involves logging into a system (see “login” on page 6-25) via **ssh (1)**. You may login again as yourself, or as another user (subject to a password check). When a dialogue is created, it executes your login shell (or, more accurately, the login shell of the user whom you logged in as). For convenience with logging in, you might want to investigate **ssh-agent (1)**.

Logging in runs your **.profile** or other initialization file appropriate to your normal login shell. Your **.profile** should avoid reading from the standard input if **NIGHTVIEW_ENV** has a non-empty value.

Remote File Access

Referencing remote files can be useful either during remote debugging or when the files of interest reside on another host.

For remote debugging, in most situations NightView can find files automatically and you don't need to worry about it. However, sometimes you need to provide more information. In those situations, you need to know the rules used by NightView to find files.

If NightView cannot find an object file on the local host, including using object translations, and the download mode is not off, it attempts to download it from the target system. See "translate-object-file" on page 6-29 and "set-download" on page 6-76.

If NightView cannot find a source file on the local host and the download mode is not off, it tries to download it from the target system. See "directory" on page 6-84.

NightView always interprets **exec-file**, **core-file**, and **load** filenames relative to the target system. See "exec-file" on page 6-46, "core-file" on page 6-43, and "load" on page 6-104.

Filenames on other commands are interpreted relative to the local host by default; however, you may explicitly refer to files on other hosts using the form *user@host:/path*. If you omit the *user@* portion, your current user is used. When you use this form, NightView transfers the file from the *host* system onto the local host system. (The file is downloaded into a file cache. You can control the behavior of the file cache with the **set-download** command (see "set-download" on page 6-76) or with the Preferences Advanced Page in the graphical user interface [see "Preferences Advanced Page" on page 8-55].)

As a special case, if **exec-file** sees the *user@host:/path* form, the **exec-file** command is treated as a **symbol-file** command (see "symbol-file" on page 6-42).

If you need to refer to remote files but do not want NightView to transfer them, don't use the *user@host:/path* form. Instead, set up another way for NightView to see the files, such as an NFS mount. NightView automatically sets up object file translations for NFS mounts. See "translate-object-file" on page 6-29.

The following commands interpret filenames relative to the host by default, but can take the *user@host:/path* form:

- **symbol-file** (see "symbol-file" on page 6-42)
- **translate-object-file** (see "translate-object-file" on page 6-29)
- **directory** (see "directory" on page 6-84)
- **source** (see "source" on page 6-153)
- **list** (see "list" on page 6-82)

Example:

Assuming you have a remote dialogue to system *fred* (see "Remote Dialogues" on page 3-6), and the program and its source are on that system, all you need to do is ensure that the download mode is set (see "set-download" on page 6-76) and then run the program.

```
set-download temporary
```

Example:

Suppose that the program is in **/usr/biff** on system *fred*, but it has been stripped of debugging information. The version with the debugging information is in **/usr/joe** on

system barney. The source is in **/usr/bob** on system betty. Use the following commands:

```
set-download temporary
translate-object-file /usr/biff/ barney:/usr/joe/
directory betty:/usr/bob
```

ReadyToDebug

The program **/usr/lib/NightView-release/ReadyToDebug** is a special program used by NightView to synchronize with the dialogue shell (*release* is the NightView release level). You will probably see this program name echoed when a dialogue shell starts up. When NightView sees this program run, it knows that the shell is through with any initialization. NightView then considers any new processes that run in the shell to be candidates for debugging. This allows the initialization to take place without debugging the programs that might run during that time.

Finding Your Program

When a program is started up from a dialogue, NightView is notified that a new program is executing, but there is currently no way for NightView to find out exactly what program is running.

NightView tries to guess the name of your program by looking at the arguments, the current working directory, and the **PATH** environment variable of the program. Usually, these correctly identify the program name, but not always. Then NightView can't tell what the program name is. Also, sometimes NightView may guess wrong.

NightView prints a message with the name of the program when the program starts up. If this name is wrong, then you will need to tell NightView the name of the program by using the **exec-file** command. See “exec-file” on page 6-46.

Most shells already do this correctly, so you will rarely need to worry about it. The problem sometimes occurs in programs that run other programs.

Controlling Your Program

NightView provides many ways to control the execution of a program you are debugging.

Eventpoints

An eventpoint is a generic term which includes breakpoints, patchpoints, monitorpoints, tracepoints, watchpoints, heappoints, and syscallpoints. All of these are different ways to

debug or modify the behavior of your program, and all of them are assigned unique numbers by the debugger when you create them. These numbers are unique across all processes. For example, if you use a command qualifier to set a breakpoint in many processes at once, each breakpoint in each process is assigned a unique eventpoint number.

Breakpoints, monitorpoints, patchpoints, tracepoints, and heappoints are *inserted eventpoints*. They are implemented by inserting code or traps into your process at a user-specified `pc` location. Syscallpoints and watchpoints are not inserted eventpoints. This difference is mostly transparent to the user, but it does cause some minor differences in behavior. Those differences are noted where appropriate.

NightView allows you to set conditions on eventpoints, so the action associated with the eventpoint is taken only if the condition is satisfied. For inserted eventpoints, the condition is an arbitrary expression in the language of the routine where the eventpoint is set (in other words, if you set a conditional eventpoint in a Fortran subroutine, you would write the conditional expression in Fortran). Whenever possible, NightView actually compiles the conditional expressions and patches them into the program, so evaluating the condition does not require the debugger to take control. This means that setting a conditional eventpoint only adds the overhead required to evaluate the condition and the program will run at almost full speed until the condition is satisfied. See “condition” on page 6-121.

With RedHawk 7.3.2 or newer kernels, a condition on a syscallpoint or watchpoint also is compiled and patched into the program. Control is transferred to that expression by the RedHawk kernel immediately after the hardware trap, and does not require the debugger to take control. So, the overhead required to evaluate the condition is minimal, and the program will run at almost full speed until the condition is satisfied. However, for older kernels or non-RedHawk kernels, a condition on a syscallpoint or a watchpoint is evaluated in the debugger.

For syscallpoints and watchpoints, the language of the expression is determined by your language setting. See “set-language” on page 6-63. Because syscallpoint and watchpoint conditions are always evaluated in the global scope, if your language setting is `auto`, NightView evaluates the condition in the language of the main program.

You can also specify an ignore count for an eventpoint. This means you must execute past the eventpoint a certain number of times before it might be taken. The ignore count is checked prior to the condition, so if you have both ignore counts and conditions, the condition will not be checked until the ignore count is down to zero. See “ignore” on page 6-124. Like conditions, the code to implement ignore counts is patched into the program for inserted eventpoints whenever possible, so the program will execute at nearly full speed until the ignore count reaches zero. Likewise, the code to implement ignore counts is patched into the program for watchpoints and syscallpoints with RedHawk 7.3.2 or newer. However, it is performed by the debugger for older or non-RedHawk kernels.

The ignore count and other attributes of inserted eventpoints can be accessed in conditions, etc. with the `__eventpoint_data__` intrinsic (see “`__eventpoint_data__`” on page 6-14).

CUDA does not support patching of debugger-defined code into a CUDA kernel. Because of this, conditions and ignore counts must be evaluated by NightView and cannot be evaluated by a patch. See “CUDA Debugging” on page 3-46 for more information.

There are several commands to manipulate eventpoints, but not every type of manipulation makes sense for every type of eventpoint. Deleting, disabling, enabling, and attaching ignore counts and conditions works for all types of eventpoints. See “Manipulating Eventpoints” on page 6-106.

Table 3-1. Eventpoint Summary

	Action	Code is inserted	May have commands
breakpoint	stop the process when the breakpoint is reached	X	X
heappoint	check the heap or configure the heap debugger	X	
monitorpoint	display the value of expressions in the monitorpoint window	X	X
patchpoint	execute an expression or modify the flow of the program	X	
tracepoint	record an event when execution reaches the tracepoint	X	
syscallpoint	stop the process when a system call is entered or existed (optionally restarting it)		X
watchpoint	stop the process when the process reads or writes a variable in memory		X

Inserted eventpoints evaluate their conditions and ignore counts at full program speed, and may be manipulated while the process is running. However, this is not the case for inserted eventpoints in CUDA code, because patching is prohibited there. Watchpoint conditions and ignore counts are evaluated in the debugger. Watchpoints may be enabled and disabled only while the process is stopped.

Inserted eventpoints always use a user-specified `pc` location to determine where in the application they are inserted. These are specified with a location specifier (see “Location Specifiers” on page 6-15). If the location specified is not a valid location, particularly for those with line numbers, they may be interpreted as approximations. That is, if an eventpoint is set at line 12 and there is no code at line 12, the actual breakpoint may be shifted to line 14 where there is code. This often is convenient for code-compile-debug cycles where the lines are moving around slightly. (See “Eventpoint Modifiers” on page 6-107.) Furthermore, NightView will accept inserted eventpoints for completely nonexistent locations in case the user’s intention was to set an eventpoint that will appear later during the execution of the application. (See “Future Eventpoints” on page 8-56 or “set-future-points” on page 6-78.)

NightView supports detection of shared libraries and CUDA code that appear after inserted eventpoints already have been set. By default, NightView will consider any inserted eventpoints in the context of those new shared libraries and CUDA code and will move them to locations within those shared libraries or CUDA code, if those locations are more appropriate. For instance, if an eventpoint was specified at line 12, but originally there was no code at line 12 so it was inserted at line 20, then later a shared library or CUDA kernel was loaded that did have code at line 12, the eventpoint would be moved from line 20 back to line 12.

Interactions Between Conditions, Ignore Counts, etc.

The following pseudo-code describes the control flow regarding eventpoint conditions, ignore counts, etc. It sometimes is useful to know how they interact if using multiple features on the same eventpoint. The field names used here are as described for the `__eventpoint_data__` intrinsic (see “`__eventpoint_data__`” on page 6-14):

```
if (!thread->protected || Protected) {
    If (EnabledFlag) {
        CrossingCount++;
        if (IgnoreCount != 0) {
            IgnoreCount--;
        } else {
            if (Condition) {
                HitCount++;
                if (EnabledOnce) {
                    EnabledFlag = false;
                }
            }
        }
    }
}
```

Eventpoint Actions

Breakpoints

A breakpoint is one of the most frequently used features of a debugger. You can set a breakpoint at any place in a program you are debugging, and when execution reaches that point, the program will stop. You may then use the debugger to examine the current values of variables, set additional breakpoints, etc. See “breakpoint” on page 6-108.

You may also specify an arbitrary set of debugger commands to execute each time a breakpoint is hit (if it is a conditional breakpoint, that means only when the condition is satisfied). See “commands” on page 6-120.

Monitorpoints

If you are debugging a real-time program, you may wish to monitor the value of one or more variables without interrupting the execution of your program. Monitorpoints allow you to do this. A monitorpoint is code inserted at a specified location by the debugger that will save the value of one or more expressions, which you specify. Because the expressions are evaluated by the program within a specific context, the expressions may

reference local stack variables and machine registers and may call functions in your program. The saved values are then periodically displayed by NightView in a Monitor Window (see “Monitor Window” on page 3-32). You can set a monitorpoint using the **monitorpoint** command. See “monitorpoint” on page 6-115.

Note that the expressions you specify are evaluated *every time* execution passes the location of the monitorpoint (unless the monitorpoint is disabled or has a condition or an ignore count). However, NightView may not display every value saved by the monitorpoint. If the monitorpoint location is executed more frequently than NightView can update the Monitor Window, you will miss seeing some of the values evaluated by the monitorpoint.

Note that there may be some delay between the time that NightView reads the values saved by a monitorpoint and the time the values appear on your display. Therefore, values sampled by *different* monitorpoints cannot reliably be related in time. However, you may be sure that all the values sampled by a *single* monitorpoint were all evaluated at the same time.

Patchpoints

During the course of debugging, you may find a small error you would like to fix, but you would also like to continue debugging the program without recompiling and relinking. The **patchpoint** command (see “patchpoint” on page 6-110) allows you to patch in a change to the memory image of the process and continue running. (Note that it does *not* change the disk copy of the program file; recompiling and relinking is the only way to make a permanent change.)

A patchpoint can cause an expression (including function calls) to be evaluated, modify a variable, or cause the program to branch to a new location.

The **load** command (see “load” on page 6-104) provides the ability to make larger scale changes by loading in whole object files. This feature may be used to replace defective routines, or to load custom designed debugging routines that can do things like verify complex data structures, or search through linked lists.

Tracepoints

The manual for the NightTrace tool describes a library that may be used to generate trace records by calling trace routines in your program. If you didn't initially build a program with trace calls, (or you did, but decided later additional trace calls were necessary) the **tracepoint** command (see “tracepoint” on page 6-113) may be used to patch in tracepoints. The values traced may then be examined with the **ntrace** tool. For more information on NightTrace, see **ntrace (1)**.

Because the program runs at full speed through a tracepoint, you can use tracepoints in real-time applications where breakpoints are unacceptable.

One significant difference between a tracepoint and a monitorpoint is that values recorded by a tracepoint are all available for later analysis; values will not be “lost” because of delays in displaying, as they may with a monitorpoint. Another difference is that tracepoints provide a reliable means of relating values of expressions at different points of execution to the times those values were evaluated. Monitorpoints do not.

Heappoints

Heappoints can be used to narrow the search for a memory bug. A heappoint either checks the process's heap or changes the configuration of heap debugging, depending on which options you specify. Heappoints may be set only after configuring the heap debugger with the **heapdebug** command (see “heapdebug” on page 6-56) or from the **Process** menu (see “Process Menu” on page 8-9). Heappoints are set with the **heappoint** command (see “heappoint” on page 6-117).

Watchpoints

A watchpoint stops your program when a particular area of memory is read or written. This is most useful in determining when a variable (or other program element) is being changed to a "bad" value during execution. You could set a watchpoint on the variable, and then the program would stop whenever the variable is modified. Watchpoints are set with the **watchpoint** command. See “watchpoint” on page 6-127.

Often you know what the bad value is. If so, you can set a condition on the watchpoint so that the program will stop only when the variable is changed to the bad value. The condition is evaluated *after* the instruction that triggers the watchpoint has completed. NightView provides a process-local convenience variable, `$is`, that is useful in watchpoint conditional expressions. See “Convenience Variables” on page 3-41. `$is` contains the value of the variable (or other program element) after the instruction that triggers the watchpoint has completed.

NOTE

In some cases it would be nice to have a `$was` (the value of the variable before the instruction began). NightView has no way to know what the value was immediately before the instruction executed. You can approximate `$was` by adding some code to your conditional expression. However, note that this picks up the old value only when the watchpoint is triggered and that there are circumstances, described below, that can change the value without triggering the watchpoint.

Example:

This example code assumes the current language is C. Suppose you want the process to stop if the old value was 1. Set the condition on the watchpoint to this expression:

```
$was=$prev, $prev=$is, $was==1
```

Set these convenience variables before resuming the process:

```
set $was=0
set $prev=0
set-local $was
set-local $prev
```


A watchpoint condition is evaluated relative to the global scope of your program. The language of the condition is controlled by your current language setting. If the setting is `auto`, then the condition is evaluated in the language of the main program.

Unlike other eventpoints, a watchpoint is not associated with a code location. A watchpoint is not an *inserted eventpoint*. See “Eventpoints” on page 3-9.

You can have many watchpoints per process, but there is a limit on the number of watchpoints that can be *enabled* at the same time. On an IA-32 or AMD64 system, at most 4 watchpoints can be enabled at one time.

A watchpoint can be set only on a program element in memory, not in a register. You should be careful about setting a watchpoint on a variable on the stack, because the watchpoint probably will not be meaningful once the routine that contains the variable returns.

For watchpoint restart information, NightView always uses the same address that it calculates when you originally create the watchpoint. Note that the specific address may or may not be interesting in another run of your program, depending on exactly what your program does. For example, a variable on the heap may always be allocated in the same place each time your program runs, or it may be allocated at a different address depending on when it is allocated, what other allocations are done, timing of external events, etc. You may need to delete a watchpoint that was created by restarting and create a different watchpoint. See “Restarting a Program” on page 3-20.

When you have a watchpoint set, your process does not incur any performance penalty until it references the addresses being watched. When that happens, NightView gets control.

NOTE

If the target system is an IA-32 system, then the mechanism NightView uses for watchpoints watches 1, 2 or 4 bytes. If the variable you are watching is 3 bytes long, then you may get some extraneous triggers on the next byte. If the variable you are watching is longer than 4 bytes, only the smallest address bytes of the variable are watched. If you need to watch more bytes of the variable, you can use multiple watchpoints, specifying addresses and lengths.

If the target system is an AMD64 system, then the mechanism NightView uses for watchpoints watches 1, 2, 4 or 8 bytes. If the variable you are watching is 3, 5, 6 or 7 bytes long, then you may get some extraneous triggers on the trailing bytes. If the variable you are watching is longer than 8 bytes, only the smallest address bytes of the variable are watched. If you need to watch more bytes of the variable, you can use multiple watchpoints, specifying addresses and lengths.

The performance of watchpoints depends on the underlying operating system. With RedHawk 7.3.2 and later kernels, watchpoints are evaluated at close to full program speed. With older RedHawk kernels, or with non-RedHawk kernels, the debugger evaluates any ignore count and condition, so the ignore count and condition are not evaluated at full program speed. See “Eventpoints” on page 3-9.

A watchpoint is not triggered if the variable is accessed by other processes through shared memory (unless they are also being debugged and have watchpoints set) or if the variable is accessed through I/O using direct memory access (DMA), such as a low-level read from disk.

Syscallpoints

A syscallpoint stops your program when a system call is entered or exited. It prints the name of the system call and whether the program is about to enter it, or is exiting from it. Options are available to resume the process automatically after printing such information and to control whether entry, exit, or both are trapped.

You can specify a specific set of system calls to trace or trace all of them.

Syscallpoints are set with the **syscallpoint** command (see “syscallpoint” on page 6-129 or the Syscallpoint Dialog (see “Eventpoint Dialog Boxes” on page 8-33).

Timing

NightView provides a mechanism to quickly and easily time regions of code from within the debugger. This is done by identifying start and end locations within executable code and associating them with a *region* name.

These locations are patched with NightView code to record timestamps and calculate the time spent within those regions. Using this information, for each region, it keeps track of the following:

- Total time spent in the region over the whole execution
- Number of crossings of the region
- Minimum time spent in the region for any single crossing
- Maximum time spent in the region for any single crossing
- Average time spent in the region for all crossings
- Additional info: warnings and/or errors detected

Each timing region supports an arbitrary number of *start* and *end* locations so that regions of code that can be entered or exited via many paths can be supported. This is particularly useful for timing whole functions, which often have multiple exits; and for Fortran programs with **ENTRY** statements to support multiple entries (although the implementation of **ENTRY** statements varies from compiler to compiler).

In addition, the patches have support for *subevents*. If two regions are defined such that one region encloses another, the enclosed region is considered a subevent of the enclosing region. In this description, “enclosing” is in terms of the dynamic execution, and not necessarily in terms of the source code. For instance, a region in a calling function would enclose a region within a called function. In general, an enclosing region could enclose an

arbitrary number of enclosed regions. In these cases, the timing patches keep track of both:

- complete time of the enclosing region, including the subevents
- the time of the enclosing region, excluding all subevents

For multi-threaded applications, all data are collected on a per-thread basis, so the results can be displayed for each thread separately, or for all threads combined.

A mechanism also is provided for *crossthread* timing. This mechanism completely disregards the threads, so that a timing region can start in one thread and end in another thread. This is useful for timing communication mechanisms between threads. Note that the overhead for crossthread timing is somewhat higher, because mutexes are required.

To support timing of whole functions, support is provided for determining the set of entry locations (usually only one, aside from Fortran ENTRY statements), and exit locations automatically. Each entry location is set as a start location, and each exit location is set as an end location.

NOTE

The timing of functions relies on DWARF debug information, and heuristics involving analysis of instructions. It is possible that not enough information may exist to do this correctly.

A further caveat relates to optimized code using tail branches and tail recursion. These are optimizations where a calling function arranges to call another function such that it never will return to the calling function. Because of this, such a tail branch must be considered an exit, so any time spent in the "called" function cannot be included in the function of the "caller".

Timing regions, start, and end locations can be configured, and timing results can be displayed using **timing** (see “Timing” on page 6-196). Timing start and end locations can be displayed using **info timing** (see “info timing” on page 6-178).

Timing Info

In addition to the timing data, sometimes unusual circumstances will be detected, which can indicate errors in where timing start or end locations are specified. These are indicated in timing results reports as **Info**. The conditions that can be detected are:

underflow

Underflows indicate end location crossings without corresponding start location crossings. This can indicate extraneous end locations or missing start locations.

However, if timing locations were added to a process during its execution, this could just reflect that execution already was within a region when a timing location was added. (Usually this results in only a single underflow, unless the process is multi-threaded or the region is executing recursively.)

maximum recursion depth

The recursion depth is the number of start location crossings without corresponding end location crossings. Values of 2 or higher indicate actual recursion.

The maximum recursion depth is the maximum observed so far. NightView supports independent timing of only 2 levels of recursion. For recursion beyond that limit, further nested regions are ignored.

current recursion depth

The recursion depth is the number of start location crossings without corresponding end location crossings. Values of 2 or higher indicate actual recursion.

The current recursion depth is the depth observed most recently. A value of 1 is normal if execution of the region is in progress. A value of 2 or more may be normal if the region also is known to be executed recursively.

If neither of these is the case, it may indicate extraneous start locations, or missing end locations.

Signals

Usually, your process is stopped and the debugger gets control if the process receives a signal. Signals may be generated by error conditions (such as dividing by zero or trying to write to a write-protected location). Other signals may be generated under program control (the program can request the system to send it a SIGALRM periodically, or another program may explicitly send a signal with the **kill (2)** system service).

Several ways in which to handle a signal are described in the **handle** command (see “handle” on page 6-144).

In addition, you may use the debugger to explicitly send a signal to a process (see “signal” on page 6-142). This is useful if you need to test the signal handler code in a program (however, the debugger itself uses SIGTRAP, so it should not be used in any of your code).

If you specify **nostop**, **noprint**, and **pass** for a signal, then the system will deliver the signal to the process normally and bypass the debugger. This avoids any performance penalty to your program if it makes frequent use of signals.

Signals may cause somewhat different behavior when you are single-stepping your program (see “Controlling Execution” on page 6-131). If a signal occurs while you are single-stepping, NightView's reaction depends on whether you specified **stop** or **nostop** and **pass** or **nopass** in the **handle** command (see “handle” on page 6-144). The four possible combinations are explained below.

`nostop, pass`

The single-step operation continues, but the signal will be passed to the program. If you have a signal handler in your program, it will be executed *without* single-stepping. When the handler finishes executing, single-stepping will be resumed until it is complete or another signal occurs.

`nostop, nopass`

The signal has no effect (other than temporarily interrupting execution). The single-step operation continues until it is completed or another signal occurs.

`stop, pass`

The single-step operation is terminated and the process is stopped. If you issue another single-step command or a **continue** command, or a **resume** command with no argument, the signal is passed on to the process when it resumes execution.

`stop, nopass`

The single-step operation is terminated and the process is stopped. The signal is discarded.

Some signals can have additional information passed to the signal handler via **siginfo(5)**. However, NightView has no mechanism for the user to specify this information, so signals sent to the process using the **signal** or **resume** commands will have no associated **siginfo(5)** information.

If a process stops with a signal that has associated **siginfo(5)** information, that information is preserved by NightView whenever possible. If you specified `pass` for that signal and you continue execution using the **continue** command or the **resume** command with no argument, the **siginfo(5)** information will be delivered to the process along with the signal. However, no **siginfo(5)** information is ever delivered if you explicitly specify a signal number on the **signal** or **resume** commands.

Exceptions

When an Ada program raises an exception, or a C++ program throws an exception, the respective runtime libraries propagate this exception through calling frames until an exception handler or catch block is found. The behavior is the same when debugging, with certain special cases:

Stepping

When stepping over a function call using commands like **next**, **nexti**, or possibly **finish**, if the function raises or throws an exception such that it would propagate back to the frame where the stepping was done, the stepping command is interrupted.

Debugger function calls

When a user calls a function from within NightView (directly, and not via a patch-point), if that function raises or throws an exception such that it propagates back to

the point where NightView initiated the call, the exception is captured and the function call is interrupted.

Exception handling:

The debugger can be configured to stop and/or print information about caught exceptions using the **handle** command (see “handle” on page 6-144).

In each of these cases, for a C++ exception, the **\$ex** convenience variable is created, and contains a reference to the thrown exception object.

Restarting a Program

Restarting execution of a program under NightView is different than in many other debuggers, because instead of being executed directly by the debugger, programs are executed from a dialogue shell, or by other programs. The typical way you restart a program is to invoke it again in the dialogue shell. See “run” on page 6-38.

When NightView recognizes that a program is being run again, it automatically applies the same eventpoints, and other information, to the new instance of the program. NightView considers two programs to be the same if they have the same full pathname.

This method of restarting programs was chosen because of NightView's multi-process nature. You may actually want to debug multiple copies of the same program, and in that case you may or may not want to have the same eventpoints set in each copy. However, if you are debugging just one instance of one program, you can easily restart its execution without having to manually duplicate your eventpoint settings.

Occasionally you may wish to run a program again and again without stopping when it **execs** or when it exits. For instance, if a program sometimes dies with a signal, you could run it repeatedly until the signal occurs and then examine where it occurred. To avoid having the process stop when it **execs**, put a **resume** command (see “resume” on page 6-133) inside an **on program** command (see “on program” on page 6-47), like this:

```
on program yourprogram do
    resume
end on program
```

The **resume** command will not actually take effect until after the process has been initialized, so **on program** and **on restart** commands that set eventpoints and otherwise modify the process work as expected. Note that the process does actually stop when it **execs**, but the **resume** command tells it to start running again as soon as NightView is finished initializing it.

To avoid having the process stop when it exits, use the **set-exit** command. See “set-exit” on page 6-53. These two mechanisms, in combination, allow you to run a program repeatedly and only stop it if it hits a breakpoint or a watchpoint or gets a signal.

The following sections describe the details of how restarting works. Most users will not need to know these details. The normal automatic mechanism handles most situations.

Restart Mechanism

At certain times in the execution of a program, NightView takes a *checkpoint* on that program. A checkpoint saves information about the eventpoints, signal disposition, etc. This information is called the *restart information*. Each checkpoint replaces the previous restart information.

The restart information is stored as a sequence of commands associated with your program name via an **on restart** command. See “on restart” on page 6-49. The commands restore the eventpoints and other information in the new program.

Each time you execute a program, NightView checks to see if an **on restart** command matches your program. If one matches, NightView executes the sequence of commands associated with your program.

Unlike other command streams, execution of an **on restart** command stream is not terminated by an error. See “Command Streams” on page 3-39.

NightView takes a checkpoint on a process when:

- It is about to exit, terminate with a signal, or be killed by NightView.
- It is about to **exec** a new program.
- You enter a **checkpoint** command. See “checkpoint” on page 6-50.

It is not possible to turn off checkpoints. However, you can control whether restart information is applied. See “set-restart” on page 6-68.

Note that if you have a program that has not yet taken a checkpoint and you start a new instance of that program, then no restart information is applied to the new instance because there is none for that program.

You can save restart information to a file. See “info on restart” on page 6-175. This allows you to save the information across debug sessions. Or, you can edit the file to change the restart information. In either case, you would then **source** the file to restore the restart information. See “source” on page 6-153.

Restart Information

This section describes the restart information saved during a checkpoint.

- Any memory reservations made with the **mreserve** command. See “mreserve” on page 6-54.
- Eventpoints, including any names, conditions, ignore counts and commands associated with each eventpoint. See “Eventpoints” on page 3-9.
- Directory search path. See “directory” on page 6-84.
- Child disposition. See “set-children” on page 6-52.
- Signal disposition. See “handle” on page 6-144.
- Display list. See “display” on page 6-101.
- Symbol file. See “symbol-file” on page 6-42.

- Default language. See “set-language” on page 6-63.
- Whether or not the process will stop before exiting. See “set-exit” on page 6-53.
- The interest level threshold, the interest level for `inline`, `justlines`, and `nodebug`, and any explicit interest levels for subprograms. See “interest” on page 6-70.
- Information to reproduce the items in the data panel. See “Data Panel” on page 8-75. See “data-display” on page 6-100.

Restart Macros

If an **on restart** command is created by a checkpoint, then in addition to commands to restore eventpoints and other program information, there are two macros: `restart_begin_hook`, at the beginning of the commands, and `restart_end_hook` at the end of the commands. Both macros are called with the name of the program being restarted as an argument.

These macros let you customize restart processing. The initial definition of these macros is

```
define restart_begin_hook(program_name) apply on program
define restart_end_hook(program_name) echo
```

This means that **on program** commands will be applied before any restart processing, and nothing will be done afterwards. (`restart_end_hook` is defined as **echo** because there is no way to make an empty macro.)

You can define these macros to be anything you wish. See “Defining and Using Macros” on page 6-184. For example, you could define `restart_begin_hook` to be **echo** to disable the **on program** processing. See “on program” on page 6-47.

Exited and Terminated Processes

When a process terminates normally, it flushes its I/O buffers, closes any open files, then calls the exit service. By default, NightView automatically arranges for a process to stop when it calls the **exit** system service. (You may alter this behavior with the **set-exit** command. See “set-exit” on page 6-53.) When a process terminates abnormally, it receives a signal, which causes the process to stop and NightView to get control. Thus, you may always examine a program that is about to exit or terminate abnormally. The process will still exist, so you can examine memory and registers.

If you continue execution of a process in one of these states, the process will cease to exist and NightView will forget about all the eventpoints set in that process. The PID for that process will be removed from all families (see “Families” on page 3-2) in which it appears. Detaching from such a process has the same effect (see “Detaching” on page 3-3).

Process States

A process is normally in one of two states; it is either *running*, or it is *stopped*. A process is said to be stopped when it gets a signal (and it is being debugged) or it hits a breakpoint or watchpoint (meaning that the point of execution reached the breakpoint or the watchpoint was triggered, and all the conditions on the breakpoint or watchpoint were satisfied). When it is stopped, the debugger has control. The debugger may continue to execute commands attached to that breakpoint or watchpoint, but once the debugger initially gets control, the process is considered to be stopped. (This is not the same type of stop as job control in the C shell or the Korn shell.)

Some debugger commands require the process to be stopped. It is meaningful to examine or modify stack locations or variables only if the process is stopped. Monitorpoints and tracepoints provide ways to examine variables without stopping a process. See “Monitorpoints” on page 3-12. See “Tracepoints” on page 3-13. The first *inserted eventpoint* in a process must be set while the process is stopped, unless eventpoint memory preallocation is on. See “Eventpoints” on page 3-9. See “set-preallocate” on page 6-74. A watchpoint may be enabled or disabled only when the process is stopped. See “Watchpoints” on page 3-14.

In addition to being stopped or running, a process may be exiting or terminated, or it may be a pseudo-process associated with a core file. A pseudo-process cannot be continued. Continuing an exiting or terminated process causes the process to cease existence.

Operations While the Process Is Executing

This section lists what you can do while the process is executing.

- Examine and modify statically-allocated variables. This includes `static` and global variables in C, and `COMMON` variables and variables with the `SAVE` attribute in Fortran. It does not include variables allocated to registers or the stack.
- Examine and modify absolute memory locations. This includes accessing memory referenced by a pointer variable, if the pointer variable is accessible as noted above.
- Evaluate expressions involving the above items. See “Expression Evaluation” on page 3-25. Note that a function call is not allowed.

For the purposes of establishing the scope and meaning of variable names, and also the language for the expression, NightView uses the location where the process was last stopped to determine the context of the expression (see “Context” on page 3-28). You can use the special forms NightView provides to change this context, if you want to access variables local to a procedure, for instance. See “Special Expression Syntax” on page 6-3. However, note that the forms that refer to specific stack frames are not allowed while the process is running, because the state of the stack is indeterminate.

- Examine, modify, and disassemble executable code.
- Create, manipulate, and destroy *inserted eventpoints* (breakpoints, monitorpoints, patchpoints, tracepoints, and heappoints). See “Eventpoints” on page 3-9. These types of eventpoints may be enabled and disabled, have conditions added or removed, and have ignore counts modified. You may modify the commands attached to breakpoints, monitorpoints and watchpoints. You may also get information about any type of eventpoint. See “Manipulating Eventpoints” on page 6-106.

Enabling or disabling watchpoints requires the process to be stopped. Any of the other operations may be performed on watchpoints while the process is executing. However, since, by default, watchpoints are enabled when created, and disabled when destroyed, you cannot ordinarily create or destroy a watchpoint while the process is executing. See “Watchpoints” on page 3-14.

If preallocation is on (default), the debugger performs special processing to support eventpoints and monitorpoints as soon as the program starts. See “set-preallocate” on page 6-74. If preallocation is off, the following restrictions apply to setting the first eventpoint and the first monitorpoint:

- The *first inserted eventpoint* within a particular text region must be set while the process is stopped. A text region is either your program or the dynamic libraries it references.
- The first *monitorpoint* must be set while the process is stopped, regardless of whether other eventpoints have been set in that region. See “Monitorpoints” on page 3-12.

This is necessary because NightView needs to do special processing when the first eventpoint is created within a text region, or when the first monitorpoint is created. That special processing requires the process to be stopped.

- Create and delete timing regions, and their start and end locations. See “Timing” on page 3-16. The start and end locations are patched, so the pre-allocation restrictions that apply to eventpoints also apply to timing start and end locations.

While the process is executing, you may not use forms of commands that depend on knowing the program counter or the value of any machine register. See “Predefined Convenience Variables” on page 6-5.

Note that monitorpoints and tracepoints also provide ways of monitoring your program without stopping it. See “Real-Time Debugging” on page 3-6.

Examining Your Program

If you specify running processes in the qualifier of a command which requires stopped processes, you get a warning message about each running process, but the command executes normally on any of the stopped processes in the qualifier.

Expression Evaluation

Because NightView is a symbolic debugger supporting multiple languages, you are allowed to evaluate expressions written in different languages, but this does not mean you have access to all the features of each language. (Specific language syntax is not described here; consult the reference manuals for the language for that information.)

One important point to note is that the debugger may not always precisely follow the language semantics when evaluating an expression. In particular, the results of a floating-point expression evaluated by the debugger may not be bit for bit identical to the results the same expression would give if it were compiled and executed in your program. See “Special Expression Syntax” on page 6-3.

A program written in multiple languages may define identical names for different global objects. NightView looks first for the name as defined in the language of the current context (see “Context” on page 3-28). If there is no current context, it uses the current language setting to determine which symbols to look at first (see “set-language” on page 6-63).

The debugger can evaluate arithmetic or logical expressions (essentially anything that may appear on the right hand side of an assignment). The debugger cannot declare new variables.

In general, the debugger cannot execute statements, it can only evaluate expressions. For Fortran, the concept of an expression is extended to assignment. (Assignment is an expression in C and C++.)

In some ways the debugger is more flexible than the compiler. The debugger usually allows you to evaluate expressions or assign new values to variables without the type checking done by the compiler. Unless the expression simply makes no sense, the debugger will evaluate it.

-
- Function Calls - How function calls are handled in the debugged executable.

C Expressions

All C expressions are supported.

The debugger supports array slices in expressions using the following syntax:

array_name[*l..u*]

where *l* is the lower bound and *u* is the upper bound. The *array_name* may be any expression that denotes either an array object or a pointer. The type of an array slice is an array whose bounds are the values of *l* and *u*, respectively.

In addition, to support querying multiple CUDA memory segments, additional keywords are allowed in type specifiers, particularly for use in type casts. They are:

```
__code__  
__constant__  
__generic__  
__global__  
__local__  
__parameter__  
__register__  
__shared__  
__texture__
```

If a pointer type references a type qualified by one of these keywords, then it implies that the pointer points to the specified CUDA memory segment. For instance, to obtain the content of CUDA shared memory at address 0x1000, you may use an expression such as `*(__shared__ float*) 0x1000`.

C++ Expressions

Most C++ expressions are supported, with a few exceptions noted below.

The debugger supports array slices in C++. See also “C Expressions” on page 3-25.

In function calls and assignments, the debugger copies an object by copying the bytes of the object. No copy constructor or user-defined assignment operator is called.

These C++ features are not supported:

- Exceptions.
- Templates.

Operator and function overloading is supported with additional input from the user used to select the desired function. See “Overloading” on page 3-27.

A special case form of the `dynamic_cast<>` function is supported. You may use `dynamic_cast<>`, spelled exactly this way (with no type name given as a template argument inside the `<>`). This form of dynamic casting will cast an object or a pointer to the actual type of that object as determined by run time type information provided by the compiler.

The additional CUDA segment keywords supported for C also are supported for C++. See “C Expressions” on page 3-25.

Fortran Expressions

All Fortran expressions are supported.

Fortran subroutines are treated as if they were functions with no return value. Fortran assignments are supported except for Concurrent Fortran array assignments.

The debugger cannot execute statements of any kind (except assignments and procedure calls), including Fortran I/O statements.

Function Calls

An expression may contain calls to functions in the debugged executable. NightView implements this by preparing the process to execute a function as if the program itself had called it, and then running the process until the completion of that function. NightView then regains control of the process.

When a user calls a function from within NightView (directly, and not via a patchpoint), if that function throws an exception and it propagates back to the point where NightView initiated the call, the exception is captured and the function call is terminated.

If the debugged process has multiple threads, for the duration of the called function, *all* threads are allowed to proceed by default. This prevents deadlock that could result if the called function performs a blocking operation, like a mutex lock. If this behavior is not desired, it can be controlled by changing the current run mode. See “Set Run Mode” on page 8-58.

This is implemented for both Ada and C++ exceptions. If the exception is a C++ exception, then a \$ex convenience variable is created, and contains a reference to the throw exception.

Overloading

Overloading means that more than one entity with the same name is visible at the same point in the program. Overloading is allowed for location specifiers and for expressions. In C++ language mode, overloading of functions and operators is allowed. See “set-language” on page 6-63. NightView refers to the appropriate entity if it has enough context to determine that there is only one choice. Otherwise, you need to provide NightView with additional information in the form of special syntax added to the expression or location specifier where the overloaded name is used.

This is typically a two step process. You run the command once and get an error which displays the possible choices. Then you run the command again with additional syntax to request the specific candidate number from that list.

The special syntax used to request candidates from the list is described in “Selecting Overloaded Entities” on page 6-2. Overloaded names are supported in language expressions (see “Expression Evaluation” on page 3-25) and location specifiers (see “Location Specifiers” on page 6-15), and the same syntax is used for both.

The **set-overload** command (see “set-overload” on page 6-73) may also be used to make NightView automatically generate overload candidate lists by turning on either of the two separate overload modes for routine names and language operators. This automates the first step of the two step process. The special syntax may be used to request overload candidate information for a single function or operator even when the corresponding overload mode is off.

If overloading is on, NightView interprets overloaded entities according to the current language. If overloading is off, NightView uses the built-in meaning of all operators, if possible, and interprets all function and procedure calls as referring to one function or procedure it arbitrarily picks from the list of candidates. If operator overloading is off and the built-in operator does not make sense in the context in which it is used, NightView

gives an error.

If overloading is on, but a unique meaning for an overloaded operator or routine cannot be determined, NightView gives an error that includes the list of the possible overload candidates. You may then run the command again, adding the syntax to select the correct candidate.

The numbers assigned to the choices are unique for the specific context (see “Context” on page 3-28) where the expression or location specifier appears. If, for example the 5th item in a list of choices refers to a particular instance of the overloaded function `funcname` when you are stopped at one point in your program, you may not assume the 5th item will refer to that same instance when you are stopped at a different location.

The one number you can rely on is 1 for overloaded operators. The built in language operator is always number 1, and any user or library defined operators have numbers greater than 1.

Program Counter

When a process is stopped, it has stopped at one specific place in the program, which is the address of the next instruction to be executed. This place is where the program counter points. Different machines have different sets of registers, but the program counter is always referred to as `$pc`.

If the currently selected frame is not the most recently called frame, then the `$cpc` register points to the instruction that made the call and the `$pc` register points to the place where execution will return after the call. In the most recently called frame, `$cpc` and `$pc` point to the same place.

Context

The location pointed to by `$cpc` implies a specific context for evaluating expressions. `$cpc` is located in some procedure (or routine, or function — the terms are used interchangeably throughout this document). This procedure was coded in some language (C, C++, Fortran, or assembler). By default, the language of the routine containing the `$cpc` is the language used to evaluate any expressions.

Another component of the context is the current stack frame (see “Current Frame” on page 3-29). It establishes which instance of a given local variable you are actually referring to in an expression. NightView provides special syntax (see “Special Expression Syntax” on page 6-3) for referencing variables in other contexts besides the current one.

Scope

Most languages have scoping rules, with local variables visible only in inner blocks and more widely visible variables in outer blocks. Often the same name is used for different variables in different scopes. Just as the `$cpc` is located in a particular routine, it is also located in a particular block of the routine. The variables that are directly visible to the

debugger are determined by the language rules and current block nesting structure of the program at that point.

When debugging, you may need to look at other variables which would normally not be visible by the strict language rules. NightView makes every effort to make any additional variables visible for use in expressions (as long as the names do not conflict). If you cannot reference a variable due to a naming conflict, NightView provides special syntax (see “Special Expression Syntax” on page 6-3) for referencing variables visible in other scopes.

Stack

When a process stops, it not only stops at a particular program counter, but it also has a current stack. The stack is used to hold local variables and return address information for each routine. As a routine calls another routine, new entries (called *frames*) are made on the stack. The stack can be examined to show the routines which were called to get to the current routine using the **backtrace** command (see “backtrace” on page 6-91).

The debugger assigns numbers to each frame. The most recent frame is always frame zero.

In a program with multiple threads, each thread has its own stack. See “select-context” on page 6-149.

Frames corresponding to uninteresting subprograms are not numbered and they are not shown in a backtrace. See “Interesting Subprograms” on page 3-31.

Current Frame

When a process stops, the current frame is initially the stack frame associated with the most recently called routine (where `$cpc` points). This frame contains the local variables for that routine, and these variables may be referenced in expressions you evaluate. Each frame also contains the return address indicating the specific point in the older routine where the `$pc` will be located when the current frame returns.

You may wish to examine the variables in one of the routines that called the current routine. To do that, you may use the **up** command (“up” on page 6-148) or the **frame** command (“frame” on page 6-147) to change the current frame. As you move up the stack (towards older routines, or in the same direction a return will go), the new stack frame becomes the *current frame*. Any variables referenced are now evaluated in the context of this new frame and new `$cpc` indicated by the called frame.

NightView also provides special syntax in expressions as an alternative to using the **up** or **frame** commands. See “Special Expression Syntax” on page 6-3.

Registers

Each stack frame also contains locations where registers are saved while in one routine so they can be restored when returning to the calling routine. As the current frame is moved, the debugger notices which registers will be saved and restored. If you look at registers using the **info registers** command, or examine local variables which are being kept in registers, you see the values as they will be restored when the process finally returns to that frame. Referencing a specific register using the predefined convenience variable also refers to the register relative to the current frame.

When examining a variable allocated to a register, you must be aware that the variable may exist in that register for only a short time. Therefore, the contents of the register may not accurately reflect the value of the variable. See “Optimization” on page 3-43 for more information.

Inline Subprograms

C++ programs can have inline subprograms. The code for these subprograms is expanded directly into the calling program rather than being called with a transfer of control. There is usually a time savings, sometimes at a cost in the size of the code.

NightView generally treats inline subprogram calls the same as non-inline calls. Although an inline call does not create a stack frame, NightView creates a frame for it to match the semantics of the language and to simplify the model of debugging. You can use the usual commands to move up and down the stack frames and view variables within each frame. See “Current Frame” on page 3-29.

You can use single step commands to step into inline subprograms, to step over them, or to finish them. See “step” on page 6-135, “next” on page 6-136, and “finish” on page 6-140.

NOTE

If you step to a source line, and the instructions corresponding to that line begin with an inline call, NightView positions you at the beginning of the inline subprogram, rather than on the line with the call.

If you set an eventpoint within an inline subprogram, NightView modifies each instance of the subprogram. If there are a lot of calls to the subprogram, this may take a long time. If execution is stopped in an inline subprogram and you set an eventpoint using the default location specifier (which corresponds to `$pc`), the location specifier refers only to *that particular instance* of the inline subprogram as opposed to all instances. See “Location Specifiers” on page 6-15.

You can set an interest level for individual inline subprograms. The interest level applies to all instances of an inline. You can also set an interest level to avoid seeing any inline subprograms. See “Interesting Subprograms” on page 3-31. This may be desirable depending on how your program uses inline subprograms.

You may not call an inline subprogram in an expression, unless the compiler has created an out-of-line instance of the subprogram. See “Expression Evaluation” on page 3-25.

Interesting Subprograms

NightView considers some subprograms to be *interesting* and the rest to be *uninteresting*. NightView avoids showing you uninteresting subprograms. Single-step commands do not normally stop in an uninteresting subprogram. See “step” on page 6-135. A stack walk-back does not display frames corresponding to uninteresting subprograms. See “Stack” on page 3-29.

In general, subprograms compiled with debug information are usually interesting and the rest are usually uninteresting. NightView gives you control over which subprograms are considered interesting by using the **interest** command. See “interest” on page 6-70.

Each process has a current *interest level threshold*. The default threshold is 0. NightView uses rules to decide on the interest level of a subprogram. If the interest level of the subprogram is greater than or equal to the interest level threshold, then the subprogram is considered to be interesting.

NightView uses these rules, in order, to determine the interest level for a subprogram:

1. The interest level may be specified for that subprogram with the **interest** command.
2. If the subprogram is an inline subprogram, the value of the `inline` interest level is compared to the interest level threshold. If the `inline` interest level is less than the interest level threshold, then the interest level for the subprogram is the minimum value. Otherwise, continue with the next rule.
3. The interest level may be recorded in the debug information for that subprogram by the compiler. Some compilers have a way of designating an interest level in the source.
4. If the subprogram has debug information, but no explicit interest level, the interest level is 0.
5. If the subprogram has line number information, but no other debug information, the interest level is the value of the `justlines` interest level for that process.
6. If the subprogram has no debug information at all, the interest level is the value of the `nodebug` interest level for that process.

In some situations there may be no interesting subprograms on the stack. In that case, the most recently called subprogram is considered interesting.

You can make all subprograms interesting by setting the interest level threshold to the minimum value.

Monitor Window

The Monitor Window shows the values of expressions being monitored by monitorpoints (see “Monitorpoints” on page 3-12). When you set a monitorpoint (see “monitorpoint” on page 6-115), the Monitor Window is created if it does not already exist, and the expressions associated with that monitorpoint are automatically displayed in the Monitor Window. The values in the window are updated approximately once a second to show the values computed the last time each monitorpoint was executed.

The **mcontrol** command (see “mcontrol” on page 6-118) controls the monitorpoint display. You can remove monitorpoint items from the display window (and add them back in later). You can change the rate at which the window updates take place, and you can stop updates completely, then start them again later. You can also turn the Monitor Window off to remove it from your screen, then restore it later.

Note that interrupting the debugger implicitly causes the Monitor Window to stop updating. See “Interrupting the Debugger” on page 3-40.

The Monitor Window is not available in the command-line interface of the debugger. You must use either the simple full-screen interface (see Chapter 7 [Simple Full-Screen Interface] on page 7-1) or the graphical user interface (see Chapter 8 [Graphical User Interface] on page 8-1) in order to take advantage of monitorpoints.

The monitored items are displayed in the Monitor Window using built-in information about the precision of the data type to decide how many columns to use for the value. You have some control over this by using the format codes on the print command.

You also have some control over the layout of the items in the window. New items are added across a line, from left to right, until there is not enough space remaining on the line to add the current item. Then a new line is started. If you remove some items (by using **mcontrol nodisplay** or by removing the monitorpoints), the remaining items are shifted left and up to pack the display. If you then add the items back, they are added at the end of the display (*not* in their original positions).

By default, each item is displayed with an identification string, a *stale data indicator*, then the value itself laid out left to right. The stale data indicator can be turned on and off via **mcontrol**. There are 3 possible states that this indicator can denote:

Updated

The monitorpoint location was executed and values were saved since the last time NightView updated the display. Note that the location may have been executed many times in between successive display updates. The displayed value represents the value as it existed the last time the monitorpoint location was executed.

Not executed

Execution has not reached the monitorpoint location since the last time NightView updated the display. This may happen if that location is executed infrequently, if the process gets suspended for some reason, or if the process is stopped by a signal or breakpoint. The displayed value still represents the value as it existed the last time the monitorpoint location was executed.

Executed but not sampled

Execution reached the monitorpoint location, but no values were saved because of an ignore count or unsatisfied condition. In this case, the displayed value is not necessarily the same as the value of the expression the last time the monitorpoint location was executed.

The actual form of the stale data indicator depends on the interface being used. See “Monitor Window - Simple Full-Screen” on page 7-2. See “Monitorpoint Values Data Item” on page 8-86.

Debugging the Heap

NightView has features to help debug problems with a program's heap (i.e. memory obtained via `malloc`, `calloc`, `realloc`, etc.). Common problems with a program's heap include buffer overruns, reads or writes of memory through "dangling" pointers which reference freed memory, and memory leaks. The debugger can also provide information about memory usage, such as the number of blocks that are allocated.

The heap debugger can be enabled and configured with the **heapdebug** command (see “heapdebug” on page 6-56), or with the **Debug Heap...** item in the **Process** menu. See “Process Menu” on page 8-9.

Once enabled, the heap debugger intercepts calls to the following heap routines:

- `calloc`
- `free`
- `malloc`
- `memalign`
- `posix_memalign`
- `pvalloc`
- `realloc`
- `valloc`

Before allowing the allocator to perform the requested operation, the heap debugger performs some checks for each such call. In addition, it is usually configured to perform a "heap check" with a specified frequency. Also, depending on the configuration, it may allocate extra memory for each block or may fill certain regions with fill bytes. See “heapdebug” on page 6-56 for details on configuration.

In addition, heappoints may be inserted at user-specified locations in the program. They can perform heap checks or change the configuration of heap debugging dynamically. See “Heappoints” on page 3-14.

The heap debugger can remember a walkback list for each allocator operation. The walkback list has the program counter for the caller of the heap routine, and the caller of that routine, and so on. If you encounter a heap error or a memory leak, this tells you which part of your program allocated the block.

The heap debugger can be used to provoke bugs by filling allocated blocks with trash to reveal problems with uninitialized fields, or by filling freed blocks with trash to reveal problems with dangling pointer references. You can also discover how your program behaves when it runs out of memory by restricting the amount of memory the heap debugger will allocate.

You can also hide the effects of some bugs. This is not intended as a remedy, but rather as information about what might be wrong with your program. For example, you can allocate extra memory for each block, which helps determine if your program is not allocating blocks of the right size, or have `malloc` zero-fill each block, which helps determine if your program is not initializing fields before using them.

Levels and Common Errors

If you have no interest in the details of heap debugging, you may want to use heap debugging levels. The levels are just a convenient way to configure heap debugging.

Level 0 (Disable) sets the controls so that the heap debugger does as little as possible, but can still issue errors for invalid heap operations.

Level 1 (Low) sets the heap debugger to do some heap checking with relatively small overhead, and is the default setting. Level 1 also enables all the features that do not change the behavior of the system allocator. That is, if you run your program without the heap debugger and then you run your program with the heap debugger at level 1, you will get the same pattern of block addresses and block sizes (however, note that some versions of the operating system vary the address space layout randomly for each process). Enabling more features may change the pattern.

Level 2 (Medium) sets the heap debugger to do more checking with greater overhead in memory used and execution time. It may also cause subtle changes in the behavior of the system allocator which can hide or expose different bugs.

Level 3 (High) does a very high level of checking at the cost of extreme overhead in memory used and execution time. In particular, the heap is checked before every heap operation and all freed blocks are retained. It may also cause subtle changes in the behavior of the system allocator.

Similarly, using the `common_errors` keyword in the **heapdebug** command (see “heapdebug” on page 6-56), or clicking one of the **Common Errors** buttons in the graphical user interface, provides a convenient way to configure heap debugging for particular kinds of errors.

Entering

```
heapdebug common_errors=block_overrun
```

or clicking the **Block Overrun** button in the graphical user interface, configures the heap debugger to detect if the program references past the end of an allocated block.

Entering

```
heapdebug common_errors=dangling_pointer
```

or clicking the **Dangling Pointer** button in the graphical user interface, configures the heap debugger to detect if the program references a block after it has been freed.

Entering

```
heapdebug common_errors=uninitialized_field
```


or clicking the **Uninitialized Field** button in the graphical user interface, configures the heap debugger to detect if the program reads from a block that it has failed to initialize.

Once you have selected a level or a common error, you may then make more detailed customizations if you wish. The heap debugging level does not affect error control, fill byte values, heap size, internal checks or slop.

Fences

The heap debugger can set "fences" on either end of each block. The fence before the beginning of the block is called the pre-fence. The fence after the end of the block is called the post-fence. The fence bytes are filled with a specified fill pattern. During a heap check, or when the block is freed, the heap debugger checks that the fence bytes have not been altered. If the fence bytes have been altered, it is an indication that your program is writing outside the block, and the heap debugger stops your process with an error status.

As a special case, blocks allocated with a size of zero have no pre-fence, no post-fence, and no slop.



This figure shows the layout of heap debugger overhead in a block. The pre-fence, slop, and post-fence are optional. See “slop=n” on page 6-61.

Hardware Overrun Protection

The heap debugger can also set up hardware overrun protection. When enabled, each block is placed as close as possible to the end of a page, and the following page is protected from reads and writes. Then if your program tries to read or write past the end of the block, it gets a signal (SIGSEGV). Note that this is not reported as a heap error in the same


way that errors appear in the **info memory** report (see “info memory” on page 6-170) or in a data panel.

The advantage of hardware overrun protection is that it catches stray references immediately, which makes it easier to find bugs, and it catches both reads and writes. The main disadvantage of hardware overrun protection is that it uses a great deal of overhead memory in order to position each block. (The protected page takes up virtual memory in your address space, but does not use any physical memory or system swap space. However, the overhead to position the block does take system memory.) It is possible for your program to exhaust the system's memory. You may need to adjust your "ulimit -v" setting or talk to your system administrator about increasing the system configuration variables `vm.max_map_count` and `vm.overcommit_ratio`.

This option is useful for small programs, and for large programs that overrun blocks before they have allocated a lot of memory. This option is also useful if you enable it during only part of the process's execution so that only particular blocks are protected. See “Heappoints” on page 3-14.

Note also that the heap debugger cannot place every block such that it ends right at the end of a page. The beginning of each block is aligned to an 8-byte boundary on IA-32 and a 16-byte boundary on AMD64. (For some allocation calls, your program may specify a larger alignment.) If the block's size is not a multiple of this alignment, then there will be a gap of a few bytes before the protected page. In this case, the program will not get the SIGSEGV unless it strays outside the block further than the gap.

When hardware overrun protection is turned on, the heap debugger automatically fills the gap with post-fence fill bytes, to help catch stray references into the gap. The gap is never made wider than it has to be to accommodate the alignment and size restrictions. For hardware overrun protection, the number of post-fence bytes is 8 on IA-32 (16 on AMD64), or the number you specify, whichever is larger, but no more than will fit in the gap.



This figure shows the layout of heap debugger overhead in a block with hardware overrun protection. The pre-fence and slop are optional. See “slop=n” on page 6-61. The gap may be zero sized; otherwise it typically is filled with the post-fence fill byte.

Retained Free Blocks

You can configure the heap debugger to retain some number of free blocks, or even all free blocks. A retained free block is not available for reuse. Once the desired number of retained free blocks is reached, blocks are made available for reuse in the same order they were freed by your program.

Retaining free blocks can help you find dangling pointer bugs. This is particularly effective when hardware overrun protection also is used, because when your program frees a block it is protected from reads and writes. However, note that retaining a large number of free blocks can use a lot more memory than your program would use normally, especially when you also use hardware overrun protection.

Heap Check

Heap checks are scans of allocated blocks, and possibly some free blocks (see “Retained Free Blocks” on page 3-37), looking for errors that could not be detected immediately as they occurred. The possible errors are:

- free-fill modified
- post-fence modified
- pre-fence modified

A heap check will be performed in the following circumstances:

- repeated automatically after a specified number of heap operations
- when the **heapcheck** command with no expression is issued (see “heap-check” on page 6-178)
- when selecting the **Heap Errors...** item from the Debug Display Menu, and selecting either:
 - Check Heap For New Errors First
 - Check Heap For All Errors First
- when selecting one of the following menu items from the Data Item Popup Menu:
 - Check Heap and Report New Errors
 - Check Heap and Report All Errors
- at a user-specified location with a **heappoint check** (see “heappoint” on page 6-117)

In addition, a limited heap check, pertaining only to a single block, is performed in the following circumstances:

- automatically when `free` or `realloc` is called on a block
- when the **heapcheck** command with an expression is issued

- when selecting the **Update Block Errors** menu item from the Data Item Popup Menu

Leak Detection

Programs sometimes "leak" blocks, which wastes memory. If NightView's heap debugging functionality is turned on with the **heapdebug** command (see "heapdebug" on page 6-56) or from the **Process** menu (see "Process Menu" on page 8-9), then it is possible to obtain a report of leaked heap blocks. See "info memory" on page 6-170 or "Leak Sets / Still Allocated Sets Data Items" on page 8-85.

An allocated block is considered leaked if no pointer in your program references it. The means of detecting leaks generally is conservative. It does not distinguish pointers from other data types. So if, for example, an integer is encountered which happens to have a value equivalent to that of a pointer to a leaked block, then leak detection will believe that it is a pointer to the leaked block, and therefore will not consider the block leaked.

Despite this conservatism, there are ways to disguise pointers which will cause leak detection to believe that a heap block is leaked when it really is not. Typically this arises when a pointer to a heap block is stored in some non-standard format. This could happen if a pointer is stored in big endian format on a little endian machine, or vice versa, or if a pointer is marshalled or pickled, and the original unmarshalled or unpickled pointer is destroyed. In practice, these situations occur rarely.

Branch Tracking

NightView and the RedHawk kernel contains support for tracking all branches in an application. There is significant time overhead with this feature, so it is disabled by default. But the user can enable it any time it seems likely that the information might be useful. It is particularly useful when an application bug causes it to branch in totally unexpected ways, such as calling a function pointer whose value is garbage, or returning from a function when the return address on the stack has been overwritten with garbage.

The **set-branch-tracking** command (see "set-branch-tracking" on page 6-78) or the **Process Settings** dialogue (see "Branch Tracking" on page 8-58) can be used to enable or disable this functionality.

For each branch instruction tracked, two addresses are remembered: the address containing the branch, and the target of the branch, where execution was transferred. Once branch tracking is enabled, and the process has been allowed to execute, this information can be displayed with the **branch-history** command (see "branch-history" on page 6-105) or in the Data panel (see "Branch History Data Item" on page 8-81).

Branch tracking is supported only for host threads, and not for CUDA contexts.

Branch tracking is only supported for RedHawk Linux version 6.0 or later and only on newer Intel chips.

Errors

NightView error messages always have this format:

severity: text [error-message-id]

The *severity* can be one of:

Caution

Usually just an informational message. It is not serious.

Warning

A little more serious, but NightView tries to finish the current command as you requested.

Error

A serious error. This level of error terminates the current command. It also terminates a command stream. See “Command Streams” on page 3-39.

Abort

So serious that NightView cannot continue running. This does not usually indicate that you have done something wrong; either there is a system problem or there is a bug in NightView.

The *text* is a brief explanation of the problem.

The *error-message-id* is a section name you can use with the **help** command to find out more about the error and possibly how to fix it. An *error-message-id* begins with E-.

NOTE

Some libraries used by NightView, such as the X Window System™, issue their own error messages in certain circumstances. These error messages *do not* follow the format described above. You can recognize these messages because they do not have the *[error-message-id]* appended to the message.

Command Streams

A command stream is a set of commands that the debugger executes sequentially. There are three kinds of command streams:

- Interactive command streams. These are commands entered directly by the user.

- A file of commands being read by the **source** command is also a command stream. Execution of the **source** command suspends execution of the command stream it appears in and creates a new one that endures until the file is exhausted.
- Event-driven command streams. For example, commands attached to a breakpoint are an event-driven command stream. Each instance of hitting a breakpoint creates a new command stream; the stream terminates when the commands attached to the breakpoint are finished. These non-interactive command streams always operate with safety level set to **unsafe** (see “set-safety” on page 6-67).

The debugger may interleave the execution of two or more command streams. For instance, it may execute some of the commands attached to one breakpoint, then execute some of the commands attached to a different breakpoint (on behalf of a different process), then execute more of the commands attached to the first breakpoint.

The debugger stops executing a command stream if it encounters a serious error (such as an unknown command, or a badly formed command). A less severe error (such as a warning about a process not being stopped) simply generates an error message, but the debugger continues to execute the remaining commands. If a serious error terminates a command stream, and that command stream was created by another command stream, then the older command stream is also terminated. This goes on until the interactive command stream is reached. The interactive command stream is not terminated.

on restart command streams are an exception to this rule. They continue to execute even if the commands get errors. See “on restart” on page 6-49.

Interrupting the Debugger

The shell interrupt character (normally <CONTROL C>) does not terminate NightView. Instead, it terminates whatever command is currently executing, if any. You may wish to use it if you accidentally ask NightView to print a large quantity of information you don't want. To type <CONTROL C>, press the **c** key while holding down the control key.

In the graphical user interface, you can interrupt the debugger by clicking the **Interrupt** button in the process toolbar. See Chapter 8 [Graphical User Interface] on page 8-1. See “Process Toolbar” on page 8-22.

If you interrupt the debugger, all command streams except the standard input stream are terminated. The standard input stream is interrupted, but not terminated, so it will prompt for the next command immediately.

Furthermore, any output from debugged processes is temporarily halted (it is still buffered, but not displayed) until after you enter the next debugger command. This gives you a chance to type a command without interference from the debugger or the debugged processes. See “Dialogue I/O” on page 3-5 for more information about controlling the output from debugged processes.

Interrupting the debugger stops the Monitor Window from updating. See “Monitor Window” on page 3-32.

Macros

A *macro* is a named set of text, possibly with arguments, that can be substituted later in any NightView command. When you define a macro, you specify its name, the names of the formal arguments, and the text to be substituted. The text to be substituted is called the *body* of the macro.

When you reference the macro in a NightView command, you again specify its name, along with the actual arguments. Actual arguments are the text you want substituted for the references to the formal arguments in the macro body. See “Defining and Using Macros” on page 6-184 for details on how to define and reference macros.

Macro expansion, the process of replacing the reference to a macro with its body, is simply a textual substitution. Very little analysis is performed on the substituted text, so macros can be a very powerful facility. Furthermore, a macro reference is expanded only when it is needed.

Macros provide a way for you to extend the set of NightView commands. They also provide a way to define shortcuts for things frequently used in commands or expressions.

Convenience Variables

NightView provides an unlimited number of convenience variables. These are variables you can assign values and reference in expressions, but they are managed by the debugger, not stored in your program. You don't have to declare these variables, just assign to them. They remember the data type and value last assigned to them.

There are two kinds of convenience variables — global and process local. Variables are global by default, but by using the **set-local** command (“set-local” on page 6-68) you can make a variable local to a process. Once you declare a variable name process local, each process maintains a separate copy of that convenience variable (a variable cannot be local in one process, but shared among all other processes). It is possible to imagine other types of scoping for convenience variables (such as breakpoint local or dialogue local), but process local and global are the only kinds currently implemented.

Because conditions on *inserted eventpoints* and the expressions associated with monitor-points, patchpoints, and tracepoints are compiled code executed in the process being debugged, references to convenience variables in these expressions always treat the convenience variable as a constant, using the value the variable had at the time the expression was defined. On the other hand, the *commands* associated with a breakpoint or watch-point, and conditions attached to *watchpoints*, are always executed by the debugger, so a convenience variable referenced in a command gets the value at the time the command or condition is evaluated.

Smart Printing

NightView supports *smart printing*, which is the capability of recognizing certain complex data types and presenting them in a simpler conceptual form that hides the details of their implementation. This is particularly useful for classes like `std::string`, which logically contains a string, but where the actual implementation is more complex and makes determination of the logical string complicated. It also is useful for container classes like `std::list` that have a simple logical organization, but where the implementation is considerably more complicated.

The **smart-print** command (see “smart-print” on page 6-191) allows definition of *smart printers*. A smart printer recognizes a type if its name matches a user-defined pattern. If the type matches, any display of an object of the type is replaced with a simpler logical description of the object.

Smart printing can be enabled or disabled for a whole process (see “smart-print” on page 6-191), for a single **print** or **output** command (see “print” on page 6-91), or for a single item in the data panel (see “Data Panel Context Menu” on page 8-86).

Logging

Each dialogue retains a buffer showing the output generated by the programs run in that dialogue shell. This output may also be logged to a file (see “set-show” on page 6-35).

In addition to the output log for each dialogue, you may log the commands you type, or the entire debug session (see “set-log” on page 6-62).

Value History

NightView keeps the results of the **print** command (see “print” on page 6-91) on a value history list. There is only one list for all the processes, and all printed values go on this list regardless of the process. You can review this history (see “info history” on page 6-167), or use previous history values in new expressions (see “Special Expression Syntax” on page 6-3).

Command History

NightView keeps a record of the commands you enter during a debugging session. There are mechanisms in the simple full-screen interface and in the graphical user interface to retrieve any of these commands, edit them, and re-enter them if desired. See “Editing Commands in the Simple Full-Screen Interface” on page 7-2. See “Command Toolbar” on page 8-22.

NightView does not add a command to the command history if it is the same as the previous command. Empty lines are never added. Commands are added only from interactive command streams. See “Command Streams” on page 3-39.

Initialization Files

When the debugger starts up, it looks for a file named `.NightViewrc` in the current working directory. If it can not find one there, it looks for `$home/.NightViewrc`. The file, if found, is then automatically executed as though it appeared as an argument to the **source** command (see “source” on page 6-153).

You can specify other initialization files, and you may disable the automatic execution of the default initialization files, using options on the NightView command line. See Chapter 5 [Invoking NightView] on page 5-1.

Optimization

There are some problems associated with debugging optimized code. These are the most common problems, but there are others:

- Machine language code may be moved around so that it does not correspond line for line to the source code in your program.
- Variables may not have the values you expect. The most common reason for this is that the value of the variable is not needed at the current location in your program and the register storing the value of the variable has been reused for another value.

Concurrent compilers generate additional debug information that indicates where variables are (i.e., register or stack) at different locations in your program. NightView uses this information to access the variables when their location(s) contain accurate values, and to prevent you from accessing them when no location contains an accurate value.

In general, you must be alert to the possibility that the compiler has changed things in your program. It may be easier to debug if you temporarily compile your program without optimization, provided your bug is still reproducible in that case.

Compilers generate debugging information at high optimization levels because it is more useful than to have nothing; however, the debug information is often inadequate to describe an optimized program. (Future compilers may generate more accurate debug information.) So, be careful and consult the appropriate manual for details.

Multithreaded Programs

NightView gives you facilities for debugging threads. On Linux, threads are implemented with separate “processes” that share resources, including memory, file descriptors, etc. In this manual, those processes are referred to as thread processes, or more simply, threads. In addition, if CUDA is in use, there may be CUDA threads present (see “CUDA Debugging” on page 3-46). When this manual refers to a process, that means all the threads together. See **pthread_create(3)**. Generally, use of the term thread in this manual includes ordinary threads, CUDA threads.

NightView provides you some control over the execution of a thread independently of the others in that process. There are two mechanisms for this. For controlling which threads stop when any of them stop, the `protected thread` tag is used (see “Protected thread tag” on page 3-46 and “Stopping” on page 3-44). For controlling which threads are executed when a command or GUI action is performed, the `run mode` is used (see “Executing” on page 3-44).

Stopping

When a thread stops, by default NightView stops the entire process (i.e. all other threads that are still executing are stopped). NightView shows you the thread that caused the process to stop, known as the *current thread*. To see other threads use the **select-context** command (see “select-context” on page 6-149), which changes the current thread. When using the graphical interface, all threads are shown in a tree display in a data panel. The current thread is shown in green.

However, it is possible to mark some threads as protected using the predefined `protected thread` tag (see “Protected thread tag” on page 3-46), defined for every thread (excluding CUDA threads). If, for a given thread, this tag is set to `true` (1), actions that cause the process to stop (e.g. hitting a breakpoint) will not stop the thread.

Executing

When you resume execution (e.g. via “resume” on page 6-133) of the process, you can choose whether all threads are allowed to execute, or just a single thread. The behavior depends on the `run mode` (see “set-run-mode” on page 6-131) and the options provided to the command. The `run mode` controls whether a single thread resumes execution or all threads resume execution. The mode is consulted when you use the graphical interface to cause execution to resume, or when you use any commands which lack explicit options to override the mode.

Similarly, if you issue a single-step command (see “step” on page 6-135), the selected thread will be stepped according to the command, but the other threads may also execute freely during that time — depending on the `run mode` or the options you use with the single-step command. The next command works similarly (see “next” on page 6-136).

All the commands that cause execution to resume have the following options: **/all** and **/one**. The former indicates that all threads should execute, regardless of the `run mode`.

The **/one** option indicates that only the current thread will execute, regardless of the run mode.

It is worth noting that the command **step /all** does **not** single step all threads. It single steps the current thread and all other threads are allowed to execute freely during the single step and they are stopped when the current thread's single step completes. The command **next /all** operates in the same manner -- it does **not** imply that all threads step over one line of code.

IMPORTANT

The **/one** option should be used with care. It is not uncommon for a single thread to block waiting on operations to be completed by other threads (e.g. by **pthread_cond_wait(3)** and many other services that hold resources). If those other threads are prevented from executing, the one thread that was resumed may block indefinitely. In such circumstances, stop the thread and resume all threads using **/all**.

Thread Tags

NightView provides the ability to associate "tag variables" with individual threads in your program. An artificial record with the reserved convenience variable name `$thr` is created for any threads in which you set thread specific tag values. The tag values are members in the `$thr` struct and each thread gets a unique `$thr` struct.

When new tags are created, the `$thr` structs are automatically reallocated to hold the new fields, as required.

The tag values can be created simply by referencing them as members in the `$thr` struct (e.g. `set $thr.tag = expr`) or setting them with the **set-tag** command (see "set-tag" on page 6-155). Tags also may be created with the **tags** syntax of the **patchpoint** command (see "patchpoint" on page 6-110 or "Patchpoint Action" on page 8-37). Any thread which executes through the patch code gets the new tag values set in that thread's `$thr` struct.

Any tags created without an explicit declaration are always one bit boolean flags with a default value of `false (0)`. The **declare-thread-tag** command may be used to declare tags of other types. Not all types are allowed: Tag types have to be fixed size types with no dynamic components. The default value of tags is always all zero bits.

Eventpoints can test values in the `$thr` struct in an eventpoint condition, so thread tags provide a way to make eventpoints conditional on which threads execute them.

Tag values are shown in the data panel thread descriptions as well as the **info threads** command output, so setting unique tag values where unique kinds of threads are created in your program serves as a way to mark threads to easily locate different threads in your program. If a program uses general purpose "worker bee" threads that are dispatched on different tasks as there are things to do, you can set a **tags** patchpoint in the thread dispatching code to mark the thread with a tag that says what it is doing while it is at work,

then clear that tag when it goes back into the pool of worker bees. To keep the output uncluttered, only tag values with non-zero values are printed, and if you declare tags with complex aggregate types, those tags are not printed in the thread descriptions. You can always use the **print \$thr** command to see the entire set of tags for the current thread.

This support is implemented by calling various library routines such as `pthread_getspecific` and `pthread_setspecific` as well as `malloc` and `realloc`, so tags will only function in programs linked with the normal **pthread** and **libc** libraries and using the pthread threading model. This also means that tags do cause a small amount of time and memory overhead to be added to your program (but not much more than any other kind of eventpoint adds). Of course, if you declare lots of tags, or use large types for tag values, the memory usage can be arbitrarily large. Also note that you should never count on the address of `$thr` to remain fixed. If you define new tags, NightView may need to realloc the space for the `$thr` struct and copy the values to a new location.

Protected thread tag

There is a predefined `protected` member present for every thread (excluding CUDA threads). It is used by NightView to determine whether or not each thread is protected from being stopped when the process as a whole is stopped. See “Stopping” on page 3-44 for more details.

There is a convenience checkbox in the data panel context menu to allow changing this tag's value easily (see “Thread Protection” on page 8-38).

CUDA Debugging

NightView supports debugging NVIDIA CUDA code. CUDA code executes on a separate CUDA device which is a wholly different architecture than the host. CUDA code is arranged into contexts, which are presented as siblings of host threads. Within these contexts are individual threads, possibly numbering in the thousands. Logically, they are organized into blocks of threads, and the blocks may be contained within multiple grids. If multiple grids are present, each may be executing a distinct CUDA kernel. Physically, the threads are organized onto the computational resources of the CUDA devices, which are broken down into devices, symmetric multiprocessors (SMs), warps, and lanes. Because the number of CUDA threads potentially can be very high, they are presented as parts of either a physical or logical hierarchy to ease understanding.

In many ways, debugging of CUDA threads is similar to debugging host threads, but there are some differences.

Similar to the way host threads function by default, when an application with CUDA code stops, all threads and CUDA devices are stopped. When resuming a CUDA thread, all CUDA threads are resumed. The **/one** option is supported, but only insofar as it will avoid starting host threads. The CUDA threads will all start running. This is a technical limitation of the CUDA architecture.

Stepping CUDA code behaves differently depending on circumstances. It does not resume any host threads because of isolation of the CUDA contexts from host threads. It is desirable to step only a single warp, but this is not always possible. NightView will step a single warp if the following conditions apply:

- The current frame is the innermost frame (i.e. frame 0 if there are no non-interesting frames).
- You do not step over a `__syncthreads()` operation.
- You do not step over any called procedures.

There are a number of features which either do not make sense for CUDA code, or which are not supported, usually because of technical limitations of the CUDA architecture or its driver. Notably, these include:

- Eventpoints other than breakpoints
- Heap debugging
- Thread tags

CUDA code is not loaded or available immediately at the start of the host application. It is loaded later by the `libcuda.so` shared library. Because of this, when setting breakpoints in CUDA code, the source decorations often are presented on other lines until the CUDA code actually is loaded. Those breakpoint lines will be changed to appropriate locations in the CUDA code once that code is loaded. Similarly, attempts to disassemble or otherwise interact with the CUDA code before it is loaded may result in information about stubs in the host code.

Eventpoints are implemented differently in CUDA code than in host code. It is not possible to compile conditions and execute them in the CUDA code, nor to perform ignore count processing in the CUDA code. So these operations are handled by NightView directly after hitting a trap in CUDA code. This is a technical limitation of the CUDA driver.

WARNING

When debugging an application with CUDA code that executes on a CUDA device, problems can arise if that same device also is used for an X11 display. Doing so can cause the device to hang. If the device was displaying the NightView windows, it may be difficult to recover from this situation.

On CUDA 5.5 or higher, and with a device with CUDA capability 3.5 or higher, CUDA software preemption may be used (see “set-cuda-software-preemption” on page 6-79 or “CUDA” on page 8-46).

On older CUDA versions or hardware, either another display device or a remote X11 display should be used.

Limitations and Warnings

DWARF Debug Information

All modern compilers can generate DWARF debug information, as described in the "DWARF Debugging Information Format" document. Typically, the option to enable this is the **-g** option.

With the gcc (and g++, gfortran, etc.) compilers versions 3.x through 4.7, the **-g** option generates DWARF version 2. Starting with gcc 4.8, the **-g** option began generating DWARF version 4, including features from DWARF 3. Starting with gcc 11, the **-g** option began producing DWARF version 5. All these versions are supported by NightView.

The gcc compilers have some additional options which generate styles of DWARF 5 information that are **not** supported by NightView, including:

- **-gdwarf64**, which causes the 64-bit DWARF format to be used. Note that this is **not** necessary to support 64-bit applications. It is only necessary if the DWARF information itself exceeds the capacity of 32-bit offsets.
- **-gsplit-dwarf**, which separates debug information into separate `.dwo` objects.

Regardless of the DWARF version, some descriptive features of DWARF are not utilized by NightView, including:

- Implicit pointer support, which allows descriptions of pointer variables which have been eliminated by optimization, but where their dereferenced objects are still accessible.
- Tail branch indications, which could allow a heuristic deduction of eliminated stack frames in walkbacks, based on the visible caller and callee.
- DWARF Expression and Location Descriptions using the typed DWARF stack.

Setuid Programs

Setuid and setgid programs can be run in a dialogue shell. If you are the superuser or the owner of the setuid program, you may also debug the program. Otherwise, NightView issues a warning message telling you that it has automatically detached from the process and the program runs without being debugged. In this case, you also cannot debug any child processes of such a program.

Note that programs run using the **shell** command (see "shell" on page 6-152) are not controlled by the debugger and so may run setuid.

Attach Permissions

You are only allowed to attach to processes running as the same user and group as the dialogue in which the **attach** command was issued, or, if a qualifier was specified, as the dialogue in the qualifier. More precisely, the dialogue's effective UID must be the same as the real and saved UID of the process you want to attach, and the dialogue's effective GID must be the same as the real and saved GID of the process you want to attach. However, the root user can attach to any process.

If attaching using a name instead of a PID, only processes owned by the same UID as the dialogue will be considered when searching for a process with the name. If associated with the root user, all processes will be considered.

Architecture Interoperability

By default, if debugging on a 64-bit system, if a 64-bit application running in a dialogue execs a 32-bit application, that application cannot be debugged by NightView and it will be detached automatically.

But NightView does have capabilities for debugging 32-bit applications on 64-bit systems. If the debugger is started with the **--arch=i386** option (see “**--arch=i386**” on page 5-1), or if a Remote Shell is started using the **Debug 32-bit applications on the x86_64 target** checkbox (see “**Debug i386 programs on x86_64 target**” on page 8-42), then either the initial dialogue or the new remote dialogue, respectively, will start a 32-bit shell instead of the default 64-bit shell. Any 32-bit applications run under that shell, directly or indirectly, can be debugged by NightView. However, if that 32-bit shell or any applications run under it exec a 64-bit application, it cannot be debugged and will be detached automatically.

Attaching to a 32-bit process is possible from a 32-bit dialogue. However, attaching to 64-bit processes is prohibited from a 32-bit dialogue.

It is possible to debug both 64-bit and 32-bit applications from a single NightView session by using a Remote Shell, but all 32-bit programs must be launched or attached from the 32-bit dialogue and its 32-bit shell, while all 64-bit programs must be launched or attached from the 64-bit dialogue and its normal 64-bit shell.

The 32-bit shell used by NightView resides in:

```
/usr/lib/NightView-version-i386/nviewsh
```

This shell is a 32-bit version of **pdsh** built specially to run on 64-bit systems. It is possible for the user to override this shell if desired. NightView first searches the **/usr/lib/NightView-version-i386/** directory for a shell with the same name as the user's desired login shell as specified in **/etc/passwd**. Only if no such shell is found will it use the **nviewsh** mentioned above. If the user wishes to build a 32-bit shell capable of running on their 64-bit system and installs it in that path with their login shell name, it will be executed instead.

Frequency-Based Scheduler

When a process running under control of the Frequency-Based Scheduler (FBS) hits a

breakpoint or watchpoint, or receives a signal that is handled by the debugger, the FBS stops running. This means that other processes under control of the same FBS will no longer be scheduled. Any other processes that are currently running will continue to run, but once they do an **fbswait(2)** call, they will not start running again until the FBS is restarted (it is as if the clock running the scheduler was stopped).

If you continue the stopped process, it will resume running, but once it executes an **fbswait(2)** call, it will also go to sleep and not wake up until the scheduler is restarted.

It is your responsibility to start the scheduler running again. This can be done via the **resume** command of the **rtcp(1)** program (perhaps using NightView's **shell** command) or by clicking **Resume** in NightSim.

NightTrace Daemon

The **tracepoint** command (see “tracepoint” on page 6-113) can be used to trace variables in a process. Tracing only works if a NightTrace daemon has been started prior to adding tracepoints to the process.

It is the responsibility of the user to start a NightTrace daemon, using either the command line tool **ntraceud(1)** or the **Daemons** dialog in the tool **ntrace(1)**. See the *NightTrace User's Guide*.

Memory Mapped I/O

Special purpose programs often attach to regions of memory mapped to I/O space. This memory is sometimes very sensitive to the size of reads and writes (often requiring an 8-bit or 16-bit reference). The debugger may access memory using 8-bit, 16-bit, or 32-bit references. This means you should probably avoid referencing I/O mapped memory unless the size of access does not matter.

Be especially careful of printing pointers to strings (e.g., variables declared to be (char *) in C or C++), because the debugger automatically dereferences these variables to print the referenced string.

Note that accesses made by tracepoints, monitorpoints, and patchpoints will be made according to the natural data type of the variable accessed, so those accesses should normally work correctly.

Blocking Interrupts

If you are debugging a program containing sections of code that block interrupts, you can easily get a CPU hung or crash the system by attempting to single step through this code (or by hitting a breakpoint or watchpoint in a section of code which executes with blocked interrupts).

Debugging with Shared Libraries

NightView provides the ability to debug programs that reference shared libraries, but there are a few things you need to know to use this effectively. This section describes how NightView interacts with shared libraries.

Shared libraries are a mechanism that allows many programs to share libraries of common code without duplicating that code in each executable file. The executable files for those programs contain the names of the shared-library files referenced by that program. These references must be *resolved* before the program can reference data or functions in the libraries. When the program first starts executing, a routine called the *dynamic linker* gets control and resolves references to shared libraries.

However, NightView gets control of a process **before** the dynamic linker executes. This is useful for NightView, but not very useful for you the user, because until the dynamic linker runs, you cannot reference any of the data or functions in the shared libraries. For instance, you could not set a breakpoint in a function residing in a shared library.

Therefore, when NightView detects that the process references shared libraries, it lets the dynamic linker execute before giving you control of the process. This allows you to debug the entire program, without needing to know which parts reside in which shared library.

One consequence of this action, however, concerns signals. If your process should receive a signal while the dynamic linker is running, NightView will detect it and give you an error message. You will not be able to reference the shared-library parts of your program, and most likely the process will not be able to continue executing properly. One source of such a signal is the dynamic linker itself. If it cannot find one or more of the shared-library files referenced by the program, it will abort the process with a signal.

Some programs require more flexibility in their use of shared libraries. These programs call the **dlopen(3)** service to load a shared library when it is needed. Because this happens after the program has initialized, NightView is unaware that a new shared library has been brought into the program's address space.

However, it is easy to make NightView aware of any dynamically loaded libraries at any time. Once your program has loaded a library or libraries using **dlopen**, you can use the **exec-file** command to force NightView to reexamine the list of shared libraries referenced by the program. See “exec-file” on page 6-46. After your program has called **dlopen**, enter the following command:

exec-file *program-name*

where *program-name* is the name of the program you are running (the one that calls **dlopen**). NightView updates its database of shared libraries, and you can then reference data and procedures in the dynamically loaded libraries.

You can issue this **exec-file** command as often as you wish. If your program loads several libraries at various points during its execution, you may want to issue the **exec-file** command several times.

Vfork

Linux supports the **vfork(2)** system call. This is a special form of **fork** where the parent and child processes shared the same virtual memory. The intention is that the child will **exec(2)** promptly, or possibly just exit. Following the **vfork** and before the **exec(2)** or **exit**, the parent is prevented from executing. NightView fully supports **vfork(2)** when the **set-children** setting is either `none` or `exec` (see “set-children” on page 6-52). However, if the setting is `all`, support for debugging the child is extremely limited. This is to prevent changes made to the child from impacting the behavior of the parent after the child has done its **exec(2)** or **exit**.

4

Tutorials

This chapter is divided into two sections.

The first is a general graphical tutorial which introduces the very basic functions of NightView and symbolic program debugging in general.

The second portion is devoted to short tutorials that concentrate on a single feature or concept. These currently include the following.

- The “Thread Tags Tutorial” on page 4-27 discusses the advantages of using Thread Tags when debugging multi-threaded processes.
- The “Tracing Tutorial” on page 4-33 discusses inserting NightTrace trace events into user code for timing analysis.

General Graphical Tutorial

This is the general tutorial for the graphical user interface (GUI) version of NightView. NightView’s graphical user interface runs only on X servers. For more information about the graphical user interface, see Chapter 8 [Graphical User Interface] on page 8-1. There is a much shorter general tutorial in Chapter 2 [A Quick Start - GUI] on page 2-1. Additionally, there are short topical tutorials in the second half of this chapter, beginning on page 27.

About the Tutorial

This tutorial shows only the most common debugger commands and features. It expects you to know the basics about window system concepts, processes, and signals, but you do not need to know about NightView and debugging concepts.

The supplied tutorial program spawns a child process. The parent writes a message to std-out, sleeps, sends signal SIGUSR1 to the child, and loops. The child writes a message to std-out when it receives the signal.

The source files used in this tutorial are found under the `/usr/lib/NightView/Tutorial` directory in tar files, and include the following:

C	Fortran	
<code>msg.h</code>	–	Defines constants
<code>main.c</code>	<code>main.f</code>	Forks a child and calls other routines

parent.c	parent.f	Sends signals to the child
child.c	child.f	Receives signals from the parent
-	ftint.c	Provides Fortran interfaces to system services

This tutorial takes significant time to complete and since each section must be performed in order, it's not easy (or wise) to skip to later sections. If you want a short general tutorial, you may want to see Chapter 2 [A Quick Start - GUI] on page 2-1.

RECOMMENDATION

Perform each step as indicated, or your results may differ from those provided in later steps of the tutorial.

Since this is a live debugging session and NightView is supported on many Linux distributions, the process IDs and hexadecimal program addresses will likely differ from those shown in the tutorial. Additionally, the line breaks in your output may differ from those shown because the lengths of displayed data items may vary.

Some messages might not appear, or additional messages might appear, depending on your environment.

Some of the commands that appear in this tutorial are an abbreviated form of the corresponding canonical NightView command. You may abbreviate NightView commands and some keywords to the shortest unambiguous prefix. For more information, see "Command Syntax" on page 6-1. You cannot abbreviate file names, symbolic names, or NightView construct names.

You could run this entire tutorial with commands and operations from the keyboard. However, use the mouse whenever possible during this tutorial.

Creating the Program

- Create a directory named `nview` where you can create files for this tutorial, and move into that directory.

```
mkdir nview
cd nview
```

Decide what language program you want to debug.

Make the `msg` program contain debug information. For the Fortran program, you should also build the `ftint.c` interface, but, for this tutorial, do *not* build it with debug information.

-

- For C, you should enter the following:

```
tar xvf /usr/lib/NightView/Tutorial/C.tar
cc -g -o msg *.c
```


- For Fortran, you should enter:

```
tar xvf /usr/lib/NightView/Tutorial/Fortran.tar
cc -c ftint.c
cf77 -g -o msg ftint.o *.f
```

You should now have a `msg` program with debug information in your `nview` directory.

Starting NightView

- Enter the following command.

```
nview
```

Note that in this tutorial `msg` does not appear on the `nview` invocation line, although NightView does accept program invocations and program arguments on the command line.

Shells in NightView

NightView communicates with a system through a dialogue which contains a shell where you run shell commands and debug running programs. For information about dialogues, see “Dialogues” on page 3-4.

The shell is always present but isn’t shown by default. If you’re just debugging a single program and don’t need to provide terminal-oriented input to it, then you don’t really need to have a shell panel visible -- since all output the program generates will also be shown in the NightView Message panel.

In most debug sessions, you simply put the program to debug and any arguments it needs on the `nview` command line, or, run the process using the **Process** menu’s **Run...** option. A shell panel is useful when you need to debug multiple programs or perhaps need to redirect program output or input.

In this tutorial, for demonstration purposes, we will use a shell panel. We will create a shell panel on a separate page.

- Click on the **View** menu and select **Add Page**.

A new tabbed page appears with the name **Page 2**. The old page now shows a tab with the name **Main**.

- Click on the **View** menu and select **Show Shell Panel**.

We’ll talk about what is displayed in the shell panel in a minute, but it is important to understand the distinction between a shell panel and a shell. The shell always exists, whereas a shell panel is simply a view of the existing shell that allows you to interact with it, just as if you were interacting with a shell in a terminal session.

In a single NightView session, you can create additional shells; normally you do this when debugging a process on a remote system. When multiple shells exist, NightView always has you decide which one you mean when you choose to create a shell panel for it, or attach to a process from it.

Each shell has a name; the default shell is `local`. The page now has a panel with the title **local shell**. For more information about shell panels, see “Shell Panel” on page 8-70.

The shell panel we created in the step above displays any output generated from your shell's initialization routine (e.g. `.bash_init` or `.kshrc`) as well as the following text:

```
/usr/lib/NightView-release/ReadyToDebug
$ /usr/lib/NightView-release/ReadyToDebug
```

NightView runs the **ReadyToDebug** program automatically as part of its initialization. It does this so it knows when your shell's initialization is complete and when it should start considering processes spawned by the shell as processes it should debug.

You might see only one echo of `/usr/lib/NightView-release/ReadyToDebug`, depending on how quickly the dialogue shell starts (*release* is the NightView release level). For information about **ReadyToDebug**, see "ReadyToDebug" on page 3-9. Note that in this tutorial the dialogue shell prompt is "\$ ". Yours may differ.

Getting Help

NightView has an integrated help system. You can get help on a widget, a command, or get help on the last error that occurred. Use the **Help** menu to access help, or type the following command:

```
help
```

Starting Your Program

Most NightView features operate on existing processes in a running program. Because you did not specify a program when you started the debugger, there haven't been any processes to debug. You must start `msg` now to debug it and to use most of the rest of the NightView features in this tutorial.

- If the current tabbed page shown is not **Page 2**, switch to it by clicking on its tab.
- Click in the shell panel to give it the keyboard focus.
- In the shell panel, you should enter the following command and press **Return**:

```
./msg
```

- Switch back to the other page by clicking on the tab labeled **Main**.

NightView should look similar to the following figure:




Figure 4-1. General Tutorial - NightView after program launch

The status bar at the bottom of the window shows that `msg` is the executable program the process is running.

The message panel shows:

```
New Process: local: 31518 parent pid: 17882
Process local:31518 is executing /raptor/fang/jojo/nview/
msg.
Reading symbols from /raptor/fang/jojo/nview/msg...done
Executable file set to
/raptor/fang/jojo/nview/msg
```

If `msg` was dynamically linked, NightView also displays the following messages:

```
Program was dynamically linked.
Dynamic linking completed.
```

NightView shows the process ID (PID) of the new process and the path where your executable exists. Your PID and the path where your executable exists will probably differ from

those in this tutorial (in fact, if they don't, you should go by a lottery ticket immediately). For information about processes, see "Programs and Processes" on page 3-2.

The source panel title bar shows the program being debugged, `msg`, the qualifier, `local:31518`, and the name of the source file that is being displayed in the source panel, `main.c`, `main.f`, or `main.a`.

In the source panel, NightView displays numbered source lines. Executable lines have a small blue diamond  source line decoration beside the line numbers.

For more information about source line decorations, see "Source Line Decorations" on page 6-88. The vertical and horizontal scroll bars in the source panel let you examine the rest of the source file.

The status bar shows the status **Stopped for exec**. This means that the process has just **exec (3)**'ed a new program image.

The context panel has an entry for this process. The header shows the qualifier, `local:31518`, and the name of the program this process is running, `msg`. The context panel entry is for the current stack frame, which is in a start-up routine that gets control before `main`. Later, we will see process entries here, but for now there is only one process, and the context panel shows process entries only if there is more than once process. See "Context Panel" on page 8-74.

Note that by appending an ampersand (&) to the `./msg`, you could have started your program in the background of the shell. This is generally a good idea because it gives you the flexibility to debug multiple programs in one NightView session; however, in this tutorial, you will be supplying the program with input, so the program needs to be running in the foreground.

Note also that although this tutorial does not ask you to do so, you can rerun a program by invoking it again in the shell panel, or by clicking on the **Rerun** button in the process toolbar.

Debugging All Child Processes

By default, NightView debugs child processes only when they have called **exec (3)**. In the `msg` program, the child process never calls `exec`. To be able to debug this child process, you must tell NightView to debug children *before* `msg` forks the child process. Also, you have tell NightView to debug children *after* invoking `./msg` so this setting can be applied to existing processes. See "Multiple Processes" on page 3-2.

- Click on the **Process** menu and select the **Process Settings...** entry.




Figure 4-2. General Tutorial - Process Settings General Page

- In the **General** page, in the **Debug Children** area, click on the combo box and select **always**.
- Click on the **OK** button to apply the change and dismiss the dialog box.

NightView echoes a **set-children** command in the message panel.

Handling Signals

By default, signals stop execution under the debugger. In the `msg` program, the parent process sends signal `SIGUSR1` to the child process. It then sleeps as a crude way of synchronizing the sending and receiving of signals. Having execution stop because of this signal is not desirable in this case.

- Click on the **Process** menu and select the **Process Settings...** entry, and then click on the **Signals** tab.




Figure 4-3. General Tutorial - Process Settings Signals Page

- Scroll down to the entry for SIGUSR1. Turn off the checkbox for **Stop**, but leave the checkboxes set for **Print** and **Pass**.
- Click on the **OK** button to apply the change and dismiss the dialog box.

NightView echoes a **handle** command in the message panel.

Note: you had to change the signal settings *after* invoking `./msg` so they could be applied to existing processes.

Setting the First Breakpoints

A breakpoint is set on the executable statement where you want program execution suspended. The program stops at the breakpoint *before* it executes the instruction where the breakpoint is set.

- Right click on the line 18 line in the source panel.

Right clicking brings up the context menu for the source panel. Clicking on the particular line identifies it as the target for the subsequent action. This is indicated by changing the background color of the line.

- Select **Set simple breakpoint** in the context menu. Repeat this for the other breakpoints as indicated below.

For the C program, the lines are 18, 25, and 30. NightView displays the following information in the message panel.

```
local:31518 Breakpoint 1 set at main.c:18
local:31518 Breakpoint 2 set at main.c:25
local:31518 Breakpoint 3 set at main.c:30
```

For the Fortran program, the lines are 18, 26, and 28. NightView displays the following information in the message panel.

```
local:31518 Breakpoint 1 set at main.f:18
local:31518 Breakpoint 2 set at main.f:26
local:31518 Breakpoint 3 set at main.f:28
```

An *eventpoint* is a generic term which includes breakpoints, patchpoints, monitorpoints, tracepoints, heappoints, watchpoints, and syscallpoints. NightView gives each eventpoint an ordinal identification number beginning at 1.

Note that you can put breakpoints only on executable statements. NightView did not give you an error for attempting to put a breakpoint on a comment line. Instead, it put the breakpoint on the executable statement that immediately follows the comment line. However, the message in the message panel has the number of the line you clicked on.

The Eventpoints panel will look similar to the following figure:

Type	ID ▲	Location	PID	Enabled	Ignore	Hits	Crossings	Cmds	Condition
● Break	1	main.c line 18 in main()	31518	✗	0	0	0	No	
● Break	2	main.c line 25 in main()	31518	✗	0	0	0	No	
● Break	3	main.c line 31 in main()	31518	✗	0	0	0	No	

Figure 4-4. General Tutorial - Eventpoint Panel

NightView changes the source panel when you set a breakpoint. Note that each line with a breakpoint on it now has a stop sign  source line decoration.

Continuing Execution

To make use of the breakpoints you set, you must allow the `msg` program to execute up to the statement with the breakpoint.

- Click on the **Resume** button. 

The status bar shows the status **Stopped at breakpoint 1**. This means that the process hit breakpoint number 1.

NightView changes the source line decoration on the statement with the breakpoint to a stop sign overlaid with a green triangle pointing to the right . The stop sign still indicates a breakpoint, and the triangle indicates that execution is stopped there.

For the C program, NightView displays the following in the message panel:

```
local:31518: at Breakpoint 1, 0x10002818 in main() at
main.c line 18
```

For the Fortran program, NightView displays the following in the message panel:

```
local:31518: at Breakpoint 1, 0x10003878 in main() at
main.f line 18
```

Not Entering Functions

Execution is stopped at the line that prompts for the number of signals to send. You don't want to enter the code for the output statement (or function) because it is part of the library, not part of your program.

- Click on the **Next** button. 

In the source panel, NightView changes the source line decoration of the next line to a green triangle pointing to the right,  which shows that execution is stopped there.

The status bar shows the status **Stopped after step**. This means that the process has finished a stepping command.

- Switch to **Page 2**.

The `msg` program writes the prompt "How many signals should the parent send the child?" in the shell panel.

Entering Input

You must respond to the `msg` program prompt "How many signals should the parent send the child?".

Remember that you may need to click in the shell panel to put the keyboard focus there.

- In the shell panel, you should enter:

10

and press **Return**.

- Switch back to the first page by clicking on the tab labeled **Main**.

Continuing Execution Again

Before you can examine aspects of `parent_routine` and `child_routine`, you must get NightView to stop at the calls to these routines.

- Click on the **Resume** button. 

The status bar shows the status **Stopped at breakpoint 3**. This means that the process hit breakpoint number 3.

For the C program, NightView displays the following in the message panel:

```
local:31518: at Breakpoint 3, 0x1000284c in main() at
main.c line 31
```

For the Fortran program, NightView displays the following in the message panel:

```
local:31518: at Breakpoint 3, 0x10003904 in main() at
main.f line 29
```

The source line decoration is now a stop sign overlaid with a triangle pointing to the right. The stop sign still indicates a breakpoint, and the triangle indicates that execution is stopped there. 

The context panel has a new entry for the child process, as shown in the following figure.




Figure 4-5. General Tutorial - Context Panel

The child process is the one with the status **New Process**. The parent process, 31518, is the *current process*, shown with green text. Buttons, menus and commands generally apply to the current process.

You would like to view the child process as the current process.

- In the context panel, click on the entry for the child process.

Now the source panel is displaying the child process.

The status bar still shows that `msg` is the executable program the current process is running. (The child is executing the same program as the parent.) The qualifier in the status bar now shows the qualifier of the child process.

NOTE

If your system has debug information installed for system libraries, the process may appear to be stopped in the `fork()` library routine. If so, click the Up button  until the debugger reports that the process is in `main`.

For the C program, the message panel shows:

```
New process: local:31575      parent pid: 31518
#0  0x10002838  in main() at main.c line 20
```

For the Fortran program, the message panel shows:

```
New process: local:31575      parent pid: 31518
#0  0x100038e4  in main() at main.f line 22
```

In this example, the child process has process ID 31575, and the parent process has process ID 31518. Note that your process IDs will differ. Note also that after the `fork`, only the parent process continued execution; the child process is still at the `fork`.

The source panel shows the main program because execution is stopped in a routine (**`fork(2)`**) which is hidden because it is uninteresting. NightView usually does not show you system library routines. See “Interesting Subprograms” on page 3-32. The source line

decoration, a gray (rather than green) triangle pointing to the right, indicates that this line made a subprogram call which has not yet returned. 

The status bar shows the status **New process**. This means that the process has just been created by a **fork (2)** call in the parent process. The process is stopped. See “Multiple Processes” on page 3-2.

The status bar shows the qualifier, `local:31575`.

The context panel lists entries for processes 31518 and 31575.

Catching up the Child Process

We want to get the **child** process to continue execution up to the breakpoint on the call to `child_routine` (line 25 in `main.c`, line 26 in `main.f`).

- With the *child* as the current process, click on the **Resume** button.



For the C program, NightView displays in the message panel:

```
local:31575: at Breakpoint 5, 0x10002840 in main() at main.c line 25
```

For the Fortran program, NightView displays in the message panel:

```
local:31575: at Breakpoint 4, 0x100038fc in main() at main.f line 26
```

The debug source file name is `main.c` or `main.f`.

NightView puts a source line decoration of a stop sign overlaid by a green triangle pointing to the right  in the source panel on line 25 for the C program and line 26 for the Fortran program.

The status bar shows the status **Stopped at breakpoint 5**. This means that the process hit breakpoint number 5. Breakpoint 5 in the child corresponds to breakpoint 2 in the parent. Inherited eventpoints get new identifiers, but the order of the eventpoint identifiers is unpredictable, so your breakpoint may have a different number.

Verifying Data Values

We want to look at the value of variables in the `msg` program.

- In the source panel, start at one side of any instance of the `total_sig` variable, hold down mouse button 1, drag it across the entire variable name, and release. (Alternatively, you could double click on the variable name where it appears surrounded by spaces).

Only the variable name should be highlighted.

- Click on the **Print** button. 

NightView displays in the message panel:

```
$1: total_sig = 10
```

The **Print** button always prints integers in decimal. NightView keeps a history of printed values. The \$1 means that this is the first value in this history. For more information about the printed value history, see “Value History” on page 3-44.

Note that if you had looked at the `total_sig` variable *after* its last use, you might have seen gibberish. This happens when the location holding a value gets overwritten. For more information, see “Optimization” on page 3-44. In the Fortran program, `total_sig` was put in `COMMON` so you could consistently see its value in the tutorial.

Listing the Source

We want to look at the source code for `child_routine`.

- Switch to the **parent** process by clicking on the parent process’s entry in the context panel. (The parent has the status **Stopped at breakpoint 3**.)
- Click on the **Source** menu, and select **List Function/Unit...**

After clicking on the parent process, the status bar shows **Stopped at breakpoint 3**. The source panel shows that execution is stopped at the call to `parent_routine`.

After clicking in the **Source** menu, NightView puts up the **Select a Function/Unit** dialog box.

- In the **Select a Function/Unit** dialog box, you should enter `child_routine` as the regular expression, and click on the **Filter** button. (For more information about regular expressions, see “Regular Expressions” on page 6-19.)

NightView finds the `child_routine` function and puts it in the list.

- In the **Select a Function/Unit** dialog box, you should click on the **Select** button.

NightView closes the **Select a Function/Unit** dialog box.

The title bar of the source panel changes the file name to `child.c` or `child.f`, and the source panel shows the source code.

Entering Functions

At this point, the parent process is about to run `parent_routine`, and the child process is about to run `child_routine`.

- In the command toolbar, enter the following command:

(all) step

(Remember, you may need to click in the command toolbar to get the keyboard focus to be there.)

Note that if you had wanted to enter a routine in only one process, you could have qualified the **step** command with the process ID, or you could have made the process the current process before entering the command.

Because you used the (all) qualifier, the **step** command causes both processes to step.

For the C program, NightView displays in the message panel:

```
#0  0x10002884  in child_routine(int total_sig = 10) at c
hild.c line 14
#0  0x10002944  in parent_routine(pid_t child_pid = 31575
, int total_sig = 10)
                        at parent.c line 11
```

For the Fortran program, NightView displays in the message panel:

```
#0  0x1000393c  in child_routine() at child.f line 17
#0  0x10003a48  in parent_routine(INTEGER child_pid /
31575 / )
                        at parent.f line 16
```

NightView tells you when a **step** command takes you into (or out of) a subprogram call. The lines that begin with #0 announce that you have entered `child_routine` in the child process and `parent_routine` in the parent process.

Note that the order of the lines displayed may vary.

Both the process entries in the context panel show the status **Stopped after step**. This means that the processes have finished a stepping command. The status bar shows the same status for the parent process.

The source file name in the title bar of the source panel changes to `parent.c` or `parent.f`, and the source panel shows the source code. Alternatively, the current context could be that of the child, in which case `child.c` or `child.f` would be shown.

Line 11 of `parent.c` or line 16 of `parent.f` in the source panel has the source line decoration of a green triangle pointing to the right,  which indicates that execution is stopped there.

Examining the Stack Frames

Observe the context panel. You may need to scroll and to expand the frames in the child process by clicking on the box with the + sign: There are entries for the two processes. Under each process are entries for the stack frames.

For the C program, NightView displays in the context panel:

```
31518 msg  Stopped after single step
#0  0x10002944  in parent_routine(pid_t child_pid =
31575, int total_sig = 10) at parent.c line 11
#1  0x10002854  in main() at main.c line 31
#2  0xb7dd5879  in __libc_start_main(...)
31575 msg  Stopped after single step
#0  0x10002884  in child_routine(int total_sig =
10) at
child.c line 14
#1  0x10002848  in main() at main.c line 25
#2  0xb7dd5879  in __libc_start_main(...
```

For the Fortran program, NightView displays in the context panel:

```
31518 msg Stopped after single step
#0 0x10003a48 in parent_routine(INTEGER child_pid /
31575 / ) at parent.f line 16
#1 0x10003910 in main() at main.f line 29
31575 msg Stopped after single step
#0 0x1000393c in child_routine() at child.f line 17
#1 0x10003900 in main() at main.f line 26
```

On lines labeled #0, NightView shows its location within the current routine. On lines labeled #1, NightView shows the location of the call to the current routine within the calling routine.

Moving in the Stack Frames

You may want to move among the stack frames to examine and modify variables, run functions, etc., in other frames. For example, suppose that you want to examine the value of local variable `tracefile` in `main`.

- Click on the entry for `parent_routine` in the parent process in the context panel.
- In the process toolbar, you should click on the Up button. 

The file name in the source panel title bar changes to `main.c`, `main.f`, or `main.a`, and the source panel shows the source code.

For the C program, NightView displays in the message panel:

```
Output for process local:31518
#1 0x10002854 in main() at main.c line 31
```

For the Fortran program, NightView displays in the message panel:

```
Output for process local:31518
#1 0x10003910 in main() at main.f line 29
```

The source panel for the C program is shown in the following figure:

```

1 #include "msg.h"
2
3 /* This program spawns a child process and sends
4  * signals from the parent to the child.
5  *
6  */
7
8 main()
9 {
10     int total_sig;
11     pid_t pid;
12     char *tracefile = "msg_file";
13     extern void parent_routine();
14     extern void child_routine();
15     extern void signal_handler();
16
17     signal( SIGUSR1, signal_handler );
18     printf( "How many signals should the parent send the child?\n" );
19     scanf( "%d", &total_sig );
20     pid = fork();
21
22     if( pid == 0 )
23     {
24         /* It's the child */
25         child_routine( total_sig );
26     }
27
28     else
29     {
30         /* It's the parent */
31         parent_routine( pid, total_sig );
32     }
33

```

Figure 4-6. General Tutorial - Source Panel for C Program

The source line decoration in the source panel is a gray triangle pointing to the right, which indicates that execution will resume there when the called routine returns. This source line decoration appears on line 34 of `main.c` and line 29 of `main.f`. The source line decoration may appear on different lines depending on which compiler you used.

Verifying Data Values in Other Stack Frames

From `main`, you can examine local variables, run functions, etc.

The locals panel may be sharing screen space with the context panel. In that case there will be tabs at the bottom of those panels. Click the one labeled **Locals**.

The Locals panel for the C program is shown in the following figure:

Item	Value
pid	31575
total_sig	10
tracefile	0x80487a0 "msg_file"

Figure 4-7. General Tutorial - Locals Panel for C Program

For the C program, NightView displays in the locals panel:

```
pid 31575
total_sig 10
tracefile 0x80487a0 "msg_file"
```

For the Fortran program, NightView displays in the locals panel:

```
pid 31575
sigusr1 10
total_sig 10
tracefile "msg_file"
```

Returning to a Stack Frame

We want to return to `parent_routine`.

- Click on the entry for `parent_routine` under the parent process in the Context panel. This frame becomes the current stack frame.

For the C program, NightView displays in the message panel:

```
Output for process local:31518
#0  0x10002944 in parent_routine(pid_t child_pid =
31575, int total_sig = 10)
    at parent.c line 11
```

For the Fortran program, NightView displays in the message panel:

```
Output for process local:31518
#0  0x10003a48 in parent_routine(INTEGER child_pid /
31575 / )
    at parent.f line 16
```

The file name in the title bar of the source panel changes to `parent.c` or `parent.f`, and the source panel shows the source code.

The green triangle  source line decoration in the source panel indicates that execution stopped there. This source line decoration appears on line 11 of `parent.c` and line 15 of `parent.f`.

Resuming Execution

We want to continue the execution of the child process so that it will get signals as soon as they are sent by the parent process.

- Switch to the **child** process by clicking on the child process's entry. Then you should click on the **Resume** button. 

After clicking the child process's entry, the file name in the source panel title bar is `child.c` or `child.f`.

After pressing **Resume**, NightView disables (dims) most of the buttons in the process toolbar.



Figure 4-8. General Tutorial - Process Toolbar

The status bar status bar shows the status **Running**. This means that the process is currently executing.

Removing a Breakpoint

Breakpoint 1 (set in “Setting the First Breakpoints” on page 4-8) is no longer needed.

- Right-click on the entry for breakpoint 1 in the eventpoint panel. The context menu appears. Select **Delete**.

NightView deletes the breakpoint from the eventpoint panel.

Setting Conditional Breakpoints

It is often useful to suspend execution conditionally.

We will set a breakpoint on the line that displays how long the parent is sleeping in `parent_routine`; the breakpoint should suspend execution when the value of `isec` equals the value of `total_sig`.

- In the context panel, click on the parent process.
- In the source panel, click on the line containing the print code. For `parent.c`, it is line 16. For `parent.f`, it is line 17. You should click on the **Eventpoint** menu. Select **Set Breakpoint...**

NightView displays the breakpoint dialog box.




Figure 4-9. General Tutorial - Set New Breakpoint Dialog

Do *not* press Return after you enter the following text.

For the C program, you should enter in the Condition If: text input area:

```
isec == total_sig
```

For the Fortran program, you should enter in the Condition If: text input area:

```
isec .eq. total_sig
```

You are ready to finish setting the conditional breakpoint.

- Click on the OK button.

NightView closes the breakpoint dialog box.

For the C program, NightView displays in the message panel:

```
local:31518 Breakpoint 7 set at parent.c:16
```

For the Fortran program, NightView displays in the message panel:

```
local:31518 Breakpoint 7 set at parent.f:17
```

The indicated line gets a stop sign source line decoration in the source panel.

Attaching an Ignore Count to a Breakpoint

Sometimes you won't want to monitor each iteration of a loop. For example, assume that a loop runs many times, and somewhere during the loop an error occurs. You could ignore the first half of the loop values to determine in which half of the iterations the error occurred.

We will set a breakpoint on the line that displays how long the parent is sleeping in `parent_routine`, ignoring the next five iterations.

- In the source panel, click on the line containing the print code. For `parent.c`, it is line 16. For `parent.f`, it is line 17. You should click on the **Eventpoint** menu. Select **Set Breakpoint...**

NightView displays the breakpoint dialog box.

Enter **5** in the **Ignore Count**: spin box. Do *not* press **Return**.

You are ready to finish attaching an ignore count to a breakpoint.

- Click on the **OK** button.

NightView closes the breakpoint dialog box.

For the C program, NightView displays in the message panel:

```
local:31518 Breakpoint 8 set at parent.c:16
```

For the Fortran program, NightView displays in the message panel:

```
local:31518 Breakpoint 8 set at parent.f:17
```

Attaching Commands to a Breakpoint

You can attach arbitrary NightView commands to a breakpoint. They run when that particular breakpoint is hit.

We will add a command stream that prints out the value of `total_sig` only when you hit the breakpoint you set in the previous step (set in “Attaching an Ignore Count to a Breakpoint” on page 4-20).

- In the eventpoint panel, you should right-click on the entry for breakpoint 8. The context menu appears. Select **Edit...**

NightView displays the breakpoint dialog box.

Note that **5** is in the **Ignore Count**: text input area from “Attaching an Ignore Count to a Breakpoint” on page 4-20.

Do *not* press **Return** after you enter the following text.

- In the **Commands**: text input area, enter the following command:

```
print total_sig
```

- Click on the **OK** button.

NightView closes the breakpoint dialog box.

Automatically Printing Variables

You can create a list of one or more expressions to be displayed each time execution stops.

We will tell NightView to display the value of the `sig_ct` variable.

- In the source panel, start at one side of any instance of the `sig_ct` variable, hold down mouse button 1, drag it across the entire variable name, and release. (Alternatively, you could double click on the variable name where it appears surrounded by spaces.) Only the variable name should be highlighted. Click on the **Data Display** button. 

A data panel appears, with a line for `sig_ct`. The value displayed is meaningless, because `sig_ct` has not yet been initialized by the program.

Note that the data panel is updated every time execution stops, and the **print** command from “Attaching Commands to a Breakpoint” on page 4-20 runs only when execution stops at a specific breakpoint.

Watching Inter-Process Communication

You already resumed the execution of the child process, so NightView did not wait for the child process.

- Click on the **Resume** button to continue execute of the **parent** process.



In the shell I/O area, NightView responds with something like the following:

```

1. Parent sleeping for 2 seconds
2. Parent sleeping for 2 seconds
Process local:31575 received SIGUSR1
Child got ordinal signal #1
3. Parent sleeping for 2 seconds
Process local:31575 received SIGUSR1
Child got ordinal signal #2
4. Parent sleeping for 2 seconds
Process local:31575 received SIGUSR1
Child got ordinal signal #3
5. Parent sleeping for 2 seconds
Process local:31575 received SIGUSR1
Child got ordinal signal #4
Process local:31575 received SIGUSR1
Child got ordinal signal #5
```

Because of the ignore count on breakpoint 8, the parent process sent only five out of ten signals to the child process before the breakpoint was hit. The source code is written so that the lines that begin with a number come from the parent process, and the lines that begin with the word "Child" come from the child process. The lines that mention signal SIGUSR1 appear because the signal settings are implicitly set to `print` and explicitly set to `nostop`.

The status bar status bar shows the status **Stopped at breakpoint 8**. This means that the process hit breakpoint number 8.

For the C program, NightView displays something like the following in the message panel:

```
local:31518: at Breakpoint 8, 0x10002950 in
parent_routine(
    pid_t child_pid = 31575, int total_sig
    = 10)
    at parent.c line 16
$3: total_sig = 10
```

For the Fortran program, NightView displays something like the following in the message panel:

```
local:31518: at Breakpoint 8, 0x105d0 in parent_routine(
    INTEGER child_pid / 31575 / ) at
parent.f line 17
$3: total_sig = 10
```

Initial lines show where execution stopped. One line shows the value of `total_sig` from the **print** command attached to breakpoint 8.

Note that the order of the displayed lines may vary.

The data panel shows the value of `sig_ct` as 6.

Patching Your Program

You just watched the parent process sleep for 2 seconds between sending signals to the child process. Look at how this is done in the source.

You will notice that the variable `isec` always has the value 2. Instead, you could vary the sleep interval `isec` by assigning it a value from 1 through 3, based on the signal count `sig_ct`.

- Click on the line that displays how long the parent is sleeping, then click on the **Eventpoint** menu. Select **Set Patchpoint...**

NightView displays the patchpoint dialog box.




Figure 4-10. General Tutorial - Set New Patchpoint Dialog

Do *not* press Return after you enter the following text.

For the C program, you should enter in the **Evaluate:** text input area:

```
i sec = sig_ct % 3 + 1
```

For the Fortran program, you should enter in the **Evaluate:** text input area:

```
i sec = mod( sig_ct, 3 ) + 1
```

You are ready to finish patching your program.

- Click on the OK button.

NightView closes the patchpoint dialog box.

Note that the line in the source panel with a patchpoint on it now has the multiple event-point  source line decoration, because it now has multiple *kinds* of eventpoints, breakpoint and patchpoint. This is overlaid with the program counter decoration. 

For the C program, NightView displays in the message panel:

```
local:31518 Patchpoint 9 set at parent.c:16
```

For the Fortran program, NightView displays in the message panel:

```
local:31518 Patchpoint 9 set at parent.f:17
```

Disabling a Breakpoint

We want to run `msg` to completion without stopping at breakpoint 8.

- In the eventpoint panel, click the checkbox in the **Enabled** field of the entry for breakpoint 8 to clear the checkbox.

Examining Eventpoints

We want to examine the types, locations, and statuses of the eventpoints we have set in `msg`.

Observe the eventpoint panel.

Type	ID ▲	Location	PID	Enabled	Ignore	Hits	Crossings	Cmds	Condition
 Break	2	main.c line 25 in main()	31518	☒	0	0	0	No	
 Break	3	main.c line 31 in main()	31518	☒	0	1	1	No	
 Break	4	main.c line 18 in main()	31575	☒	0	1	1	No	
 Break	5	main.c line 25 in main()	31575	☒	0	1	1	No	
 Break	6	main.c line 31 in main()	31575	☒	0	0	0	No	
 Break	7	parent.c line 16 in par...	31518	☒	0	1	6	No	<code>isec == total_sig</code>
 Break	8	parent.c line 16 in par...	31518	☐	0	1	6	Yes	
 Patch	9	parent.c line 16 in par...	31518	☒	0	0	0	No	

Figure 4-11. General Tutorial - Reviewing Eventpoint Panel for C Program

NightView displays all eventpoints for process `local:31518` followed by the eventpoints for process `local:31575`.

Breakpoints 1, 2, and 3 were set in “Setting the First Breakpoints” on page 4-8. Breakpoint 1 has no entry because it was deleted in “Removing a Breakpoint” on page 4-18. Breakpoints 2 and 3 are still enabled.

When the child process was forked, it inherited the parent process’s breakpoints. The child’s breakpoints 4, 5, and 6 correspond to the parent’s breakpoints 1, 2, and 3. The order of the eventpoint numbers for inherited eventpoints is not necessarily the same as in the parent.

Breakpoint 7 was set in “Setting Conditional Breakpoints” on page 4-18 and is still enabled.

Breakpoint 8 was set in “Attaching an Ignore Count to a Breakpoint” on page 4-20 and was disabled in “Disabling a Breakpoint” on page 4-24.

Patchpoint 9 was set in “Patching Your Program” on page 4-22 and is still enabled.

Continuing to Completion

There’s nothing else to look at, so let’s run `msg` to completion.

- If the parent process is not the currently displayed process, switch to it.

- Click on the **Resume** button.

NightView displays in the message panel:

```

6. Parent sleeping for 1 seconds
7. Parent sleeping for 2 seconds
Process local:31575 received SIGUSR1
Child got ordinal signal #6
8. Parent sleeping for 3 seconds
Process local:31575 received SIGUSR1
Child got ordinal signal #7
9. Parent sleeping for 1 seconds
Process local:31575 received SIGUSR1
Child got ordinal signal #8
10. Parent sleeping for 2 seconds
Process local:31575 received SIGUSR1
Child got ordinal signal #9
Process local:31575 received SIGUSR1
Child got ordinal signal #10

```

The source code is written so that the lines that begin with a number come from the parent process, and the lines that begin with the word "Child" come from the child process. Note that the sleep interval varies from 1 through 3 because of the patch you made in “Patching Your Program” on page 4-22. The lines that mention signal SIGUSR1 appear because the signal settings are implicitly set to `print` and explicitly set to `nostop`.

Note the order of the displayed lines may vary.

NOTE

If your system has debug information installed for system libraries, the process may appear to be stopped in the `_exit()` library routine. If so, click the **Up** button  until the debugger reports that the process is in `main`.

The source panel shows the main program, at the call to `exit`.

The status bar status bar shows the status **About to exit**. This means that the process has called the exit system service. See “Exited and Terminated Processes” on page 3-22.

NightView displays in the message panel:

```
Process local:31518 is about to exit normally
```

The data panel shows that the `sig_ct` variable is not visible at this point in the parent process.

```
Identifier "sig_ct" is not visible in the given context.
```

Depending on which compiler you used, the value may still be visible.

Leaving the Debugger

- Click on the **File** menu. Select **Exit NightView**.

Neither process has completely exited, so NightView puts up a warning dialog box, asking the following question:

`Kill all processes being debugged?`

- In the warning dialog box, click on the **OK** button.

The main window is removed.

This concludes the general tutorial.

Topical Tutorials

This section contains short tutorials which concentrate on a single topic or feature. A general working knowledge of debugging with NightView is assumed.

Thread Tags Tutorial

This tutorial demonstrates how to use thread tags to enhance the experience of debugging a multi-threaded application.

- Compile and link the `thread.c` file which is located in `/usr/lib/NightView/Tutorial`, with commands similar to the following:

```
cat /usr/lib/NightView/Tutorial/threads.c > threads.c
cc -g threads.c -lpthread -lm -lrt
```

NOTE

On some RedHawk Linux systems, if the above command fails to link the program, you may need to specify `-lccur_rt` instead of `-lrt`.

The example program is written in C and creates a `director` thread which does out work to a set of three `worker_bee` threads.

- Start the program under NightView, using the following command:

```
nvview ./a.out
```

- Once NightView starts, add a **Data** panel with a **Threads** display, by using the **Data** menu and selecting the **Threads** option.
- Now let the process run for a bit by issuing the following two debugger commands, separated by a few seconds:

```
resume
stop
```

The program has executed for a few seconds and then stopped. The Threads display in the Data panel will look similar to the following figure:

Item	Value
 Threads	local:11566 a.out
 11566	C thread 0x401df6c0
 11570	C thread 0x403e0b90 (director)
 11571	C thread 0x405e1b90 (worker_bee)
 11572	C thread 0x407e2b90 (worker_bee)
 11573	C thread 0x409e3b90 (worker_bee)

Figure 4-12. Thread Tags Tutorial - Threads Automatically Named

For each thread (other than the main program thread), NightView has determined the name of the start routine for the thread and displays that in parentheses.

The name is derived from the start routine address passed to `pthread_create(3)`, which is the first user function that executes when a new thread is created.

In our example program, there is a single `director` thread and three instances of a `worker_bee` thread. Using worker threads is a fairly common programming paradigm; the activities of those threads are dynamically assigned by the application as it executes.

In our application, the `director` thread creates a workload message and queues the message waiting for a `worker_bee` thread to service the message.

Here's an excerpt from the source of the `worker_bee` routine (some lines were deleted for brevity).

```

62: for (;;) {
64:     mq_receive(q, (char*)&msg, sizeof(msg_t), &priority);
71:     switch (msg.job) {
72:         case _EXIT:
73:             break;
74:         default:
75:             calc = (*msg.work)(msg.arg);
76:     }
81: }
```

The `worker_bee` thread receives a message (when one is available) and then starts executing the body of work that the message indicates. In this case, it calls a `work` function supplied by the message.

It would be convenient to know what job each thread is running whenever we look at the thread display in NightView.

To achieve that, we'll create a thread tag which represents the type of job being executed. Then we'll insert a patchpoint into the `worker_bee` routine to assign the current job to the current thread whenever it is dispatched.

- Issue the following command in NightView:

```
declare-thread-tag job jobs
```

Here we have declared that a thread tag called `job` to be of type `jobs`, which in this case is defined as an enumerated type in the program on line 14.

```
13
14 typedef enum { _Nuthin, _Sine, _Cosine, _Pi, _EXIT } jobs;
15
```

- Click on line 71 in the source panel and then right-click and select the **Set eventpoint** menu option and the **Set Patchpoint** option from that sub-menu.
- The Patchpoint dialog appears.




Figure 4-13. Thread Tags Tutorial - Patchpoint Dialog

- Click the **Set thread local tags values** radio button
- Type the following in the **Thread Tags** field

```
job = msg.job
```

- Click the OK button to close the dialog
- Resume the process for about five seconds, then stop it again.

When the process stops, the **Threads** display in the **Data** panel now includes the value of the thread tag `job` for all threads where its value is non-zero.

Item	Value
 Threads	local:11566 a.out
  11566	C thread 0x401df6c0
  11570	C thread 0x403e0b90 (director)
  11571	C thread 0x405e1b90 (worker_bee), job=_Pi (3)
  11572	C thread 0x407e2b90 (worker_bee), job=_Sine (1)
  11573	C thread 0x409e3b90 (worker_bee), job=_Pi (3)

Figure 4-14. Thread Tags Tutorial - Threads display with job tag

NOTE

Since the `directory` threads assigns jobs randomly and we are randomly stopping the process, the jobs assigned to each thread will likely differ from the figure above, as will the current thread, which is shown in green.

NOTE

If there is no value shown next to the word `job` in the **Threads** display, then perhaps you forgot the **declare-thread-tag** step described above (if you forgot, then the default type of `job` is `boolean`; so it only shows up because its value is non-zero -- i.e. `true`). If so, execute the **declare-thread-tag** command now, remove the patchpoint and then set it again and proceed.

The `director` thread assigns a priority to each message that it sends; it might be nice to see which threads are executing the most urgent messages. Here we can set a `boolean` thread tag to indicate this.

Since thread tags are of type `boolean` by default, we don't need to declare the type of this new tag we're going to create, we can just start using it.

- Issue the following command in NightView:

```
patch threads.c:65 tags urgent = priority > 2
```

- Resume the process and let it execute for a few seconds, then stop it again.

Item	Value
 Threads	local:11566 a.out
 11566	C thread 0x401df6c0
 11570	C thread 0x403e0b90 (director)
 11571	C thread 0x405e1b90 (worker_bee), job=_Pi (3)
 11572	C thread 0x407e2b90 (worker_bee), job=_Sine (1)
 11573	C thread 0x409e3b90 (worker_bee), job=_Sine (1) urgent

Figure 4-15. Thread Tags Tutorial - Urgent Processing Shown

The Threads display in the Data panel will now include the word `urgent` for any thread that is executing an urgent message. Tags whose values are all zero are simply omitted from the list of tags in the Threads display.

NOTE

It is possible you stopped the process when no urgent messages were being serviced. If you don't see the `urgent` tag in the Threads display, repeatedly resume and stop the process until you do.

- Click on a thread in the Threads display in the Data panel which does not have the `urgent` tag displayed and then enter the following command:

p \$thr

This command prints the values of all thread tags for the current thread, even if their value is zero:

```
$1: $thr = {
    job = _Pi (3);
    urgent = FALSE
}
```

You can also declare thread tags to have descriptive text. In our example, the message received by the `worker_bee` threads includes a string.

- Issue the following command in NightView:

declare-thread-tag handle char[20]

We've declared that the thread tag `handle` is an array of 20 characters. However, assigning to strings in C requires us to use a syntax other than a simple assignment (`=`) operator. As such, an alternative form of patchpoint can be used.

- Issue the following command in NightView:

```
patch threads.c:65 eval strcpy($thr.handle,msg.handle)
```

In this case, we used the standard `eval` form of a patchpoint, and directly referenced the handle tag from the convenience variable `$thr`, which contains all tags.

- Resume the process in NightView, wait a few seconds, and then stop it.

Item	Value
 Threads	local:11571 a.out
 11566	C thread 0x401df6c0
 11570	C thread 0x403e0b90 (director)
 11571	C thread 0x405e1b90 (worker_bee), job= Cosine (2) handle="cosine"
 11572	C thread 0x407e2b90 (worker_bee), job=_Pi (3) urgent handle="apple"
 11573	C thread 0x409e3b90 (worker_bee), job=_Pi (3) handle="apple"

Figure 4-16. Thread Tags Tutorial - String Tag Shown

The list of non-zero thread tags now includes the handle for the `worker_bee` threads.

Thread tags are not only useful for providing feedback to you on thread activities and state, you can use thread tags in eventpoint conditions.

- Issue the following command in NightView:

```
break do_math if $thr.job==_Cosine
```

Here we've referenced a thread tag in a breakpoint, which is a convenient way to set thread-specific conditions.

- Resume the process and let it hit the breakpoint we just defined.

Within a few seconds, the process should stop and NightView will show that we've reached the breakpoint on line 115 and that the current thread is one executing the `_Cosine` job, as indicated in the following screen shot.




Figure 4-17. Thread Tags Tutorial - Breakpoint Reached

In the figure above, we hovered the mouse cursor over the breakpoint decoration on line 115; this caused NightView to pop-up a tool-tip which summarized the breakpoint, its condition, and the hit and crossing counts. We can see that the breakpoint was crossed 4 times; 3 times the condition was false, so processing for a message other than `_Cosine` must have occurred during that time.

- Terminate the NightView session.

This concludes the Thread Tags tutorial.

Tracing Tutorial

This tutorial describes how to use NightTrace tracepoints while debugging an application with NightView.

NightTrace is part of the NightStar tool suite and allows you to instrument your application with function calls which log integer trace identifier values, along with any type of argument. Each logged trace event is timestamped. You can analyze the data graphically using the `ntrace(1)` command, or obtain ASCII listings of the data in chronological order, or write programs to consume the logged data using the NightTrace Analysis API.

This tutorial does not teach you how to use NightTrace, but will show you the minimal commands you need to use to collect and report on NightTrace data. The NightStar Tutorial is highly recommended as an introduction to using NightTrace, but you don't need to stop and read it to follow along in this tutorial.

In our example, we will insert trace events into an application that does not already use the NightTrace API.

- Compile the example program using commands similar to the following:

```
cat /usr/lib/NightView/Tutorial/tracing.c > tracing.c
cc -g tracing.c -lpthread -lrt -lm
```

NOTE

On some RedHawk Linux systems you may need to replace `-lrt` with `-lccur_rt` if the `cc` command above does not successfully link the program.

The example program is written in C and creates a `director` thread which does out work to a set of three `worker_bee` threads.

- Start the process under NightView with the following command:

```
nvview ./a.out &
```

NOTE

Be sure to background the NightView invocation, as we will be using this terminal session while NightView is still executing.

- Once NightView starts, add a **Data** panel with a **Threads** display, by using the **Data** menu and selecting the **Threads** option.
- Now let the process run for a bit by issuing the following two debugger commands, separated by a few seconds:

```
resume
stop
```

The program has executed for a few seconds and then stopped. The **Threads** display in the **Data** panel will look similar to the following figure:

Item	Value
 Threads	local:27761 a.out
  27761	C thread 0x401df6c0
  27762	C thread 0x403e0b90 (director)
  27763	C thread 0x405e1b90 (worker_bee)
  27764	C thread 0x407e2b90 (worker_bee)
  27765	C thread 0x409e3b90 (worker_bee)

Figure 4-18. Trace Tutorial - director and worker_bee threads

In our example program, there is a single `director` thread and three instances of a `worker_bee` thread. The `director` thread creates a workload message with a priority between 1 and 3, and queues the message waiting for a `worker_bee` thread to service the message.

Each `worker_bee` thread waits for a message, and then executes the work function specified in the message.

We might be interested in how long it takes to send a message, have it received, and then serviced.

We can do this by inserting NightTrace tracepoints into the application and use NightTrace to report on their instances and the times associated with them.

The `director` thread sends messages on line 48.

```

48     if (mq_send(q, (const char*)&msg, sizeof(msg), priority)) {
49         fprintf (stderr, "warning: mq_send() fails: %s\n", strerror(errno));
50         sleep(10);
51     }

```

Figure 4-19. Trace Tutorial - snippet of director source code

- Click on line 48 in the source panel and then right-click and select the **Set eventpoint** menu option and the **Set Tracepoint** option from that sub-menu.
- The Tracepoint dialog appears.




Figure 4-20. Thread Tags Tutorial - Tracepoint Dialog

- Ensure the **Location** field says `tracing.c:48`; if not, make it so.
- Enter the value `10` in the **Event ID** field.
- Click the **OK** button to close the dialog.

A tracepoint decoration now appears near the line number in the source display.

```

48  if (mq_send(q, (const char*)&msg, sizeof(msg), priority)) {
49      fprintf (stderr, "warning: mq_send() fails: %s\n", strerror(errno));
50      sleep(10);
51  }
```

Figure 4-21. Trace Tutorial - Source Panel w/ Tracepoint Decoration

Now we'll add another tracepoint at the point where messages are finished being serviced.

- Enter the following NightView command:

```
trace 11 at tracing.c:78
```

Now we're almost ready to start collecting trace data. However, in order to use tracing, the application must make a call to the NightTrace Logging API to initialize tracing. Our application doesn't use the NightTrace Logging API, so we can have NightView take care of this for us.

- Enter the following NightView command:

```
set-trace tracefile=tracing.data
```

The NightView message panel will respond with text similar to the following:

```
set-trace tracefile=tracing.data
tracing initialized, but no trace daemon running
```

NightView has linked in the required modules from `/usr/lib/libntrace_thr.so` and inserted a call to **trace_begin(3)** on the applications's behalf, passing "tracing.data" as the name of the trace file. NightView also detected that it could not locate a running NightTrace daemon associated with the `tracing.data` file.

NightTrace requires that a daemon is run to collect the trace data and either send it to a file, or send it to a live **ntrace(1)** analysis session. You can run your application without a daemon being active, but no trace records will be collected, until a daemon is subsequently launched.

- Resume process execution in NightView using the **resume** command.

The process is now executing and it is attempting to log tracepoints at lines 48 and 78, but since no daemon is active, they are harmlessly discarded.

- In the terminal session used to launch NightView, issue the following three commands:

```
ntraceud --join tracing.data
sleep 7
ntraceud --quit-now tracing.data
```

These commands initiated a NightTrace daemon to collect the trace events, waited for seven seconds, and then terminated the daemon.

We can now use NightTrace to summarize the collected events.

IMPORTANT

Remember that this tutorial does not teach general use of NightTrace and as such will be only using the most primitive of NightTrace features so as not to confuse new users.

- Invoke the following command in the terminal session used in the previous step:

```
ntrace --summary=st:10-11 tracing.data
```

We have told NightTrace to summarize the data in file `tracing.data` and to report on a state which represents sending a message and subsequently servicing it.

The start of the state is defined by event 10, which we set at line 48 in the `director` thread at the call to `mq_send`. The end of the state is defined by event 11, which we set at line 78 in the `worker_bee` thread after processing the workload defined by the message.

NightTrace will generate output similar to the following:

```
=====
Summary: States starting with event 10, ending with event
11.

State Summary Results
=====

Number of states found:          3028

Maximum state duration:         0.003980141 at offset:
10201
Minimum state duration:         0.000006065 at offset: 933
Average state duration:         0.001202657
Total of state durations:       3.641646082

Number of state gaps found:     3028

Maximum state gap:              0.006318099 at offset:
10222
Minimum state gap:              0.000004960 at offset: 5043
Average state gap:              0.000668079
Total of state gaps:            2.022944709
```

Of most significance to us is the average state duration. We can see that on average, it took 1.2 milliseconds to send and subsequently receive and process a message.

We will now refine our tracepoints to get some more detailed data.

- In NightTrace, double-click on the tracepoint event in the Eventpoint panel to launch the Edit Tracepoint dialog.
- Set the Condition text field to the following:
priority==1
- Press the OK button to close the dialog.
- Do the same for the second tracepoint shown in the Eventpoint panel.
- Now issue the following NightView commands; you can copy and paste these commands from this text into the command field of NightView if you wish:

```
trace 20 at tracing.c:48 if priority==2
trace 21 at tracing.c:78 if priority==2
trace 30 at tracing.c:48 if priority==3
trace 31 at tracing.c:78 if priority==3
```

The Eventpoint panel should now look like the following figure:

Type	ID ▲	Location	PID	Enabled	Ignore	Hits	Crossings	Cmds	Condition
 Trace	1	tracing.c line 48 in director()	27761		0	0	0	No	priority==1
 Trace	2	tracing.c line 78 in worker_...	27761		0	0	0	No	priority==1
 Trace	3	tracing.c line 48 in director()	27761		0	0	0	No	priority==2
 Trace	4	tracing.c line 78 in worker_...	27761		0	0	0	No	priority==2
 Trace	5	tracing.c line 48 in director()	27761		0	0	0	No	priority==3
 Trace	6	tracing.c line 78 in worker_...	27761		0	0	0	No	priority==3

Figure 4-22. Tracing Tutorial - 6 Tracepoints

Additionally, the source decorations for lines 48 and 78 should look like the following figure:

```

 48      if (mq_send(q, (const char*)&msg, sizeof(msg), priority)) {
  ♦ 49          fprintf (stderr, "warning: mq_send() fails: %s\n", strerror(errno));
  ♦ 50          sleep(10);
    51      }
  ♦ 77          calc = (*msg.work)(msg.arg);
 78      if (min == 0.0 || calc < min) min = calc;
  ♦ 79      if (max == 0.0 || calc > max) max = calc;
```

Figure 4-23. Tracing Tutorial - Conditional Tracepoint Decorations

The salmon-colored background shown on the tracing decoration indicates that a condition has been applied to the tracepoint at that line. The background color is obscured on line 78 because of the blue background used to indicate line selection.

- Issue the following commands in the terminal session we've been using in this tutorial:

```
ntraceud --join tracing.data
```

```
sleep 7
ntraceud --quit-now tracing.data
```

The commands above collected trace data for seven seconds.

- Issue the following command to summarize the trace event data and print the average state durations.

```
ntrace --summary=st:10-11,st:20-21,st:30-31 \
fgrep "Average state duration"
```

That instructed NightTrace to summarize the events collected, in terms of three states; each state corresponds to messages of priority 1, 2, and 3; since we applied those conditions on the tracepoints in NightView.

The output of the last command should be similar to the following:

```
Average state duration:      0.001176966
Average state duration:      0.000314065
Average state duration:      0.000309296
```

Here we see that the average time required to send and subsequently receive and service messages seems to be affected by the message's priority.

To give you a glimpse into additional NightTrace capabilities, execute the following command:

```
ntrace /usr/lib/NightView/Tutorial/tracing.session \
tracing.data
```

Here we have passed a NightTrace session file to **ntrace(1)** which includes some customization of graphs and event and state names. A window containing figures similar to the following should appear:




Figure 4-24. Tracing Tutorial - NightTrace Graph of Message Servicing

Events							
Offset	Event	CPU	Process	Thread	Time (sec)	Tag	Description
6791	msg_prio3_sent		a.out	25888	3.462986899		
6792	msg_prio3_serviced		a.out	25888	3.463274884		
6793	msg_prio3_sent		a.out	25888	3.463651885		
6794	msg_prio3_serviced		a.out	25888	3.464490607		
6795	msg_prio2_send		a.out	25888	3.464703035		
6796	msg_prio2_serviced		a.out	25888	3.465209941		
6797	msg_prio2_send		a.out	25888	3.466898824		
6798	msg_prio2_serviced		a.out	25888	3.467183829		
6799	msg_prio3_sent		a.out	25888	3.467710834		
6800	msg_prio3_serviced		a.out	25888	3.467995250		

Figure 4-25. Tracing Tutorial - NightTrace Event List

The panel with the three graphs is showing the individual durations of each instance of a state; the red graph corresponds to priority 1 messages, blue to priority 2, and green to priority 3.

The data boxes to the left of the graphs show the individual service times of the message most recently serviced to the left of the "current timeline". The current timeline is defined

by wherever you click the mouse in the graph. The graphs can be zoomed in and out and panned left and right.

The textual events panel shows a simple text listing of every event, including the time at which it was logged. The current timeline in the data graphs is linked to the current row in the events panel; they always refer to the same event, regardless of which panel you interact with.

- Exit NightTrace.
- Exit NightView.

You can learn more about using NightTrace and its powerful features in the NightStar Tutorial. You can view the tutorial from the **Help** menu of any NightStar tool.

This concludes the Tracing tutorial.

5

Invoking NightView

This chapter describes how to start a NightView session.

nview

You can start NightView without any arguments at all, if you wish. The arguments available on the NightView command line control the initial state of the debugger, and optionally allow you to specify the first program to be debugged. The command line to invoke NightView looks like this:

```
nview [--arch=i386] [-attach pid|name]
      [-config config-file] [-core core-file]
      [-help]
      [-nogui] [-nolocal] [-nx] [-prompt string]
      [-safety safe-mode] [-singlescreen] [-version]
      [-Xoption ...] [-x command-file]
      [program-name [program-argument ...]]
```

--arch=i386

The **--arch=i386** option instructs NightView to start with the local dialogue debugging 32-bit applications instead of 64-bit applications. It is meant for use on a 64-bit system when you wish to debug 32-bit applications there. See “Architecture Interoperability” on page 3-55.

-attach pid|name

Attach to the process specified by *pid* or *name* in the local dialogue. This is similar to using the **attach** command. See “attach” on page 6-41. This option is not meaningful with **-nolocal**.

-config config-file

Load the configuration contained in *config-file*. This is similar to using the Load Config... item in the File menu. See “File Menu” on page 8-4. This option is valid only in the graphical user interface.

-core corefile-name

When you supply a **-core** option, NightView starts out by creating a pseudo-process for the given core file. See “Core Files” on page 3-5 and “core-file” on page 6-46.

-help

Causes NightView to print its command line syntax followed by a brief description of each option and then exit with code 0.

-nogui

Prevents NightView from automatically invoking the graphical user interface. See Chapter 8 [Graphical User Interface] on page 8-1.

-nolocal

Prevents NightView from starting a dialogue on the local system. See “Dialogues” on page 3-6. In the graphical user interface, if **-nolocal** is used, NightView pops up a Remote Login Dialog Box (see “Remote Login Dialog Box” on page 8-43).

-nx

Prevents NightView from reading commands from the default initialization file. See “Initialization Files” on page 3-47.

-prompt *string*

Sets NightView's initial prompt string to *string*. See “set-prompt” on page 6-69.

-safety *safe-mode*

Sets the initial safety level to *safe-mode*, which can be *forbid*, *verify*, or *unsafe*. The default level is *verify*. This controls the debugger's response to dangerous commands. See “set-safety” on page 6-71.

-singlescreen

Directs NightView to use a simple full-screen interface. This option implies **-nogui**. See Chapter 7 [Simple Full-Screen Interface] on page 7-1.

-version

Causes NightView to display its current version and then exit with code 0.

-Xoption

NightView accepts a subset of the standard X Toolkit command line options (see **X(7x)**). These options are allowed only when using the graphical user interface. See below for a list of the options accepted.

-x *command-file*

Directs NightView to read commands from *command-file* before reading commands from the default initialization file or from standard input. You may supply more than one **-x** option if you like; the files are read in the order of their appearance on the command line.

program-name [*program-argument* . . .]

If *program-name* is specified, NightView begins debugging that program.

All options may be abbreviated to unique prefixes. For example,

nview -si

invokes NightView with the simple full-screen interface.

If the environment variable `DISPLAY` is set, or the standard X Toolkit command line option **-display** is used, then NightView communicates through a graphical user interface. In this case, a subset of other standard X Toolkit command line options are also allowed, e.g., **-geometry** *geometry-string*. See Chapter 8 [Graphical User Interface] on page 8-1.

When using the graphical user interface, the X Toolkit options accepted include:

-display *display*

-geometry *geometry*

-fn *font* or **-font** *font*

-bg *color* or **-background** *color*

-fg *color* or **-foreground** *color*

-btn *color* or **-button** *color*

-name *name*

-title *title*

NightView uses the NightStar License Manager (NSLM) to control access to the NightStar tools. See “Licensing...” on page 8-22 for more information.

All NightView command line options are case-insensitive. However, note that X Toolkit options are case-sensitive.

When NightView starts execution, it first attempts to read commands from any files specified in **-x** options. It then looks for any initialization files to read (see “Initialization Files” on page 3-47), unless the **-nx** option was specified. When those files have all been processed, NightView reads commands from standard input until it encounters the end of the file or the **quit** command is executed (see “quit” on page 6-25).

nview-save-core-file

When transporting a core file to another system for analysis, it is important to take all the shared libraries related to its execution. Locating them can be a cumbersome task that NightView can do for you more easily.

This command is provided as a convenience. It invokes NightView and saves all files associated with a core file in a compressed tar file for subsequent analysis and then exits.

The options are similar to the **save-core-file** command (see “save-core-file” on page 6-47).

nview-save-core-file [Options] *core-file* [*core-args*]

--nozip

Do not compress the tar archive used to save the core file (the default is to generate a compressed file).

--nodebuginfo

Do not include the debuginfo files (if there are any). The default is to include them. Debuginfo files are special files that are optionally installed on systems which provide debug information for standard libraries. NightView knows how to locate these files, and by default, automatically consults them when they are present.

--replace

If the output file already exists, do not terminate the save operation, but instead to try to overwrite it.

--keep

If there are errors during the writing of the output file, keep the possibly broken file rather than removing it.

--include=file

Add *file* to the directory of information being packaged with the core file. The specified file can be of any type -- perhaps some data files used as input to your program or maybe just some notes you make so you can remember the circumstances relating to the core file. You can use multiple *include* options to include multiple files.

--note=string

Add the *string* as a note in the generated script that will be used to debug the core file subsequently. This note will be echoed at the end of the execution of the script that this script creates and places in the tar file.

--output=file

Defines the name of the output file; the final file will be named *file.tar.gz* or *file.tar*, depending on whether **--nozip** is specified. If this option is omitted, *file* defaults to *corefile*.

--quiet

Hide all output from the NightView session that is invoked as part of this script; that session is used to actually implement this script.

core-file [*core-args*]

Specifies the name of the core-file to be saved and passes *core-args* as arguments to the NightView **core-file** command which is invoked as part of the operation of this script.

This command will take a single core file and will generate the (optionally) compressed tar file. This can be useful if there is limited access time available on the current system or you wish to send the core file off to another person.

The tar file can then be transferred to another system for subsequent analysis. The system must be of the same architecture as the original system (i.e. 32-bit or 64-bit).

The tar file will contain all required files as well as a script called **debug-core-file**. Simply executing that script while positioned in the extracted directory will allow you to debug the core file.

Command-Line Interface

This chapter describes how to interact with NightView through commands.

In some cases, this may be your only means of directing the debugger's actions. If you are using the graphical user interface (see Chapter 8 [Graphical User Interface] on page 8-1), however, commands are only one of several ways to control the debugger and your programs.

Command Syntax

This section describes the general syntax and conventions of NightView commands. Most commands have three parts. An optional qualifier appears first (in parentheses) and is used to restrict the command to a certain set of processes or dialogues. Next comes the keyword indicating which command is to be executed. The command arguments follow as the third part. In general, you must separate syntactic items (like keywords and argument values) with white space, unless they are separated by punctuation characters. White space consists of one or more space characters. These rules may be different within expressions, where the rules of the programming language apply.

Some commands apply to individual processes; others apply to dialogues. The *qualifier* is a prefix that determines the dialogues and/or processes to which the following command applies. A qualifier is simply a list of dialogues and/or processes enclosed in parentheses. If a command applies only to dialogues, and the qualifier includes specific processes, the command applies to the dialogues containing the processes. If a command applies only to processes, but the qualifier includes dialogues, the command applies to all processes in those dialogues. If a command affects neither dialogues nor processes, the qualifier is ignored. You can set a default qualifier that will be applied when you don't provide one. For more information on the syntax and operation of qualifiers, see “Qualifier Specifiers” on page 6-17.

On startup, NightView provides you with a dialogue, `local`, for debugging on the local machine. The initial default qualifier is set to `all` to indicate all dialogues and processes.

After the qualifier, if any, all commands start with a *keyword*, which may be abbreviated to the shortest unambiguous prefix. Many frequently used commands also have special abbreviations. Most commands have one or more *arguments*; some arguments are also keywords, while others are information you supply. A keyword argument can usually be abbreviated if it is unambiguous; any exceptions to this rule are noted in the section describing the command. Both command and argument keywords are case-insensitive; they can be entered in either upper or lower case. You cannot abbreviate file names, symbolic names, or NightView construct names.

Commands are terminated by the end of the input line.

If you enter a line interactively consisting solely of a newline, NightView will usually

repeat the previous command. This is explained more fully later; see “Repeating Commands” on page 6-21.

You can include comment lines with your commands. A comment line starts with the # character, which must be the first non-blank character on the line, and terminates at the end of the input line. Comments are most useful when you write debugger source files or macros (see “Defining and Using Macros” on page 6-184 and “source” on page 6-153).

NightView prompts you for input. The format of the prompt may be controlled by the **set-prompt** command (see “set-prompt” on page 6-65). The default prompt includes the names of all the dialogues in the default qualifier and looks like this:

```
(local)
```

Some NightView commands require multiple lines of input. For these commands, the command-line and simple full-screen interfaces change the prompt to > to remind you that you are entering a multi-line command.

```
>
```

To terminate NightView, use the **quit** command, which can be abbreviated **q** (see “quit” on page 6-24).

The subsections below explain some common syntactic constructs that are used in a variety of NightView commands.

•

Selecting Overloaded Entities

For general information about function and operator overloading, see “Overloading” on page 3-27.

The special overloading syntax used in both expressions and location specifiers is always introduced by a number sign character (#) used as a suffix directly following the entity (an operator in an expression or a function or procedure name). The # is followed by additional information indicating the specific kind of overload request. There are three forms of # syntax:

1. #?

A number sign followed by a question mark is a query. It always makes the command it appears in fail, but the error message shows all the possible choices for overloading the name or operator (even if there is only 1 choice). The choices will be numbered starting at 1, and the number may be used to select the specific function.

2. ##

Two number signs in a row act just as if **set-overload** were on for that one name. If there is only one possible choice, it is used; if there are multiple choices, the command fails and the error message shows the list.

3. #<digits>

A number sign followed by a number is the syntax used to pick a specific overloaded function or operator from the list printed in the error message.

In C++, the function call and subscript operators don't appear in a single location, but are "spread out" with arguments or subscripts between the parenthesis or brackets. In these cases the final bracket or parenthesis is the character which should be suffixed with the #. For example:

```
function#5(12,3)
```

This picks the 5th instance of the name `function` from a list of overloaded functions.

```
object(12,3)#5
```

This, on the other hand, picks the 5th version of an overloaded `operator()` function call operator applied to the `object` variable.

Special Expression Syntax

For general information about expression evaluation, see "Expression Evaluation" on page 3-25. In addition to the standard language syntax, NightView offers a special syntax for referencing convenience variables and variables from other scopes or stack frames.

The special constructs all start with '\$' as shown in the following table.

Table 6-1. Special '\$' Constructs

\$

A simple '\$' by itself is a special convenience variable which always refers to the last value history entry (see "print" on page 6-91). See "Value History" on page 3-42.

\$\$

The name '\$\$' refers to the value history entry immediately prior to '\$'. See "Value History" on page 3-42.

\$*number*

A '\$' followed by a number refers to that number entry in the value history. See "Value History" on page 3-42.

\${-*number*}

A '\$' followed by a negative number enclosed in braces refers to value history entries prior to the most recent one. '\${-0}' is a complicated way to refer to the same thing as '\$', and '\${-1}' is the same as '\$\$'. This syntax is useful when you want to reference values farther back than -1. See "Value History" on page 3-42.

\$identifier

This is the standard syntax for convenience variables. Many names are predefined (for instance, all the machine registers may be referenced using predefined convenience variables). See “Convenience Variables” on page 3-41, and “Predefined Convenience Variables” on page 6-5.

\${file:line expression}

This syntax is used to evaluate the expression in the context specified by the given file and line number. This is most useful for referencing static variables which are not visible in the current context. If you reference a local stack or register variable from some other context, the results are not defined.

\${+number:routine expression}

This syntax is used to go up the stack (see “up” on page 6-148) until you see *number* previous occurrences of *routine* relative to the current frame. (It does not matter what the current routine name is, this construct always backs up the frame first, then starts looking for frames associated with the given routine.) The given *expression* is then evaluated in that context. For example, *'\${+1:fred x}'* refers to the variable named *'x'* in the first routine named *fred* above the current routine.

\${+number expression}

This syntax simply refers to previous stack frames, regardless of the routine name. The immediately previous frame is *'+1'*.

\${-number:routine expression}

This syntax is useful only if you have changed your current frame with the **up** command. This allows you to refer to frames down the stack and is analogous to the version above which uses the *'+'* syntax.

\${-number expression}

This is also analogous to the corresponding *'+'* syntax, but refers to frames down, rather than up the stack.

\${=number expression}

This syntax evaluates the expression in the context of the given absolute frame number, regardless of the current frame. You can determine absolute frame numbers by using the **backtrace** command (see “backtrace” on page 6-91).

*\${*frame-addr expression}*

This syntax uses *frame-addr*, which must be a numeric constant, as an absolute frame address. It evaluates *expression* in the context of this frame address, regardless of the current frame. If there is no frame on the current stack with this address, the results are undefined.

You may wish to use this form in **display** expressions (see “display” on page 6-101) to refer to a specific stack frame regardless of where it appears relative to the current frame. You can use the **info frame** command (see “info frame” on page 6-166) to get the frame address for any stack frame.

The above constructs may be used freely in any language expression. This means they may be nested (in case you want to do something like back up the stack frame, then shift to a different local scope in that routine). Because different frames may be associated with routines in different languages, the expressions evaluated in any given context may be expressions in different languages. This might not always make sense because different languages support different data types. If NightView cannot figure out how to evaluate a mixed language expression, it returns an error.

If you use any of these constructs in a conditional expression for an *inserted eventpoint* (breakpoint, monitorpoint, patchpoint, tracepoint, or heappoint), or in a monitorpoint, patchpoint or tracepoint expression, they are evaluated at the time you establish the expression, not when the expression is evaluated within the eventpoint. This is because the eventpoint expressions are compiled into your program by the debugger, and these constructs must be evaluated at compile time.

In the rare case of a user program which contains variables that have a '\$' in their name, the user program variable is always referenced in preference to the convenience variable.

Predefined Convenience Variables

You may create any number of convenience variables simply by assigning values to new names, but some variables are predefined and have special values. The '\$' and '\$\$' variables have already been documented (see “Special Expression Syntax” on page 6-3). The following special variables are all automatically defined on a per process basis.

Table 6-2. Predefined Convenience Variables

\$_

This variable holds the address of the last item dumped with the **x** command (see “x” on page 6-95). It is also set by the eventpoint status commands to the address of the last eventpoint listed, and the **info line** command to the address of the first executable instruction in the line. If you were dumping words, it holds the address of the last word. If you were dumping bytes, it holds the address of the last byte, etc. See “x” on page 6-95, “info eventpoint” on page 6-158, “info breakpoint” on page 6-159, “info tracepoint” on page 6-160, “info patchpoint” on page 6-161, “info monitorpoint” on page 6-162, “info heappoint” on page 6-163, and “info line” on page 6-183.

\$__

This variable holds the contents of the last item printed by the **x** command. If you were dumping words, it holds the last word. If you were dumping bytes, it holds the last byte, etc.

\$pc

This variable provides access to the program counter. This is a machine register, but every machine has a \$pc, so this name is common to all machines. When a program is stopped, \$pc is the location where it stopped. On any given machine, \$pc may not map directly onto a specific machine register (RISC machines often have multiple program counters), but it always represents the address at which the program stopped. See “Program Counter” on page 3-28.

`$cpc`

`$cpc` is similar to `$pc`. In frame 0, if there are no hidden frames below frame 0 (because of uninteresting subprograms), `$cpc` has the same value as `$pc`. See “Interesting Subprograms” on page 3-31. In other frames (including frame 0 if there are hidden frames below it), `$cpc` is the address of the instruction that is currently executing. In most cases, this is the call instruction that caused the frame immediately below the current frame to be created. For the frame immediately above a signal-handler stack frame, `$cpc` is the address of the instruction that was executing when the signal occurred.

`$sp`

Most machines have a stack pointer. The stack pointer is always called `$sp`.

`$fp`

Most machines either have a frame pointer, or have an implicit frame pointer derived from information in the program. The `$fp` variable always represents the frame address (even if it is not a specific hardware register), and local variables are always described with some offset from the frame pointer (see “info address” on page 6-180 and “info frame” on page 6-166).

`$cfa`

`$cfa` is the *canonical frame address*. This is how the debug information describes the locations of the return address and the low-level registers saved in a frame. This may or may not be the same as `$fp`. See “info frame” on page 6-166.

`$is`

`$is` is defined when a watchpoint is triggered. See “Watchpoints” on page 3-14. `$is` is the value of the variable being watched after the instruction that causes the trigger has completed.

`$ex`

`$ex` is defined when a C++ exception is handled by NightView (see “handle” on page 6-144), or captured as part of a stepping command or a function call from the debugger. `$ex` is a reference to the thrown exception object.

IA-32 Registers

IA-32 machines are based on the Intel IA-32 architecture (see *IA-32 Intel Architecture Software Developer's Manual* for architectural details). See “info registers” on page 6-168.

In addition to the common register definitions for stack pointer (`$sp`), frame pointer (`$fp`), and program counter (`$pc`), the IA-32 machines support the registers shown in the following table.

Table 6-3. IA-32 Registers (iHawk Series 860)

`$eax, $ebx, $ecx, $edx, $esi, $edi, $ebp, $esp`

These names map onto the 8 general purpose registers. Note that `$sp` is the same as `$esp`, and `$fp` may be the same as `$ebp`, depending on how the compiler generates code.

`$ax, $bx, $cx, $dx, $si, $di, $bp, $sp16`

These names map onto the lower 16 bits of each of the 8 general purpose registers mentioned above, respectively. Note that the lower 16 bits of the ESP register is more commonly known as simply SP. But the name `$sp` is reserved for an architecture-independent stack pointer in NightView. So the name `$sp16` is used for the lower 16 bits of the ESP register, instead.

`$al, $bl, $cl, $dl`

These names map onto the low order 8 bits of each of the AX, BX, CX, and DX registers, respectively. In other words, they map onto bits 0-7 of each of the EAX, EBX, ECX, and EDX registers, respectively.

`$ah, $bh, $ch, $dh`

These names map onto the high order 8 bits of each of the AX, BX, CX, and DX registers, respectively. In other words, they map onto bits 8-15 of each of the EAX, EBX, ECX, and EDX registers, respectively.

`$eflags`

The program status and control register. NightView and the kernel use the TF flag of this register to implement stepping. See “step” on page 6-135, “stepi” on page 6-138, “next” on page 6-136, and “nexti” on page 6-139. Users should not modify the TF field of the `$eflags` register. Other flags in this register are used by the kernel. Care should be taken if modifying this register.

`$eip`

The instruction pointer register. This is the same as the `$pc` register.

`$cs`

The code segment register. The IA-32 architecture uses this register to determine the location of the executable code in memory. Users should not modify this register. Modification of this register in patchpoints and eventpoint conditions is prohibited.

`$ss`

The stack segment register. The IA-32 architecture uses this register to determine the location of the process stack. Users should not modify this register.

`$ds, $es, $fs, $gs`

The data segment registers. The IA-32 architecture uses the `$ds` register to determine the location of the process data. Users should not modify that register. Care should be taken if modifying `$es`, `$fs`, or `$gs`.

`$csbase, $dsbase, $esbase, $fsbase, $gsbase, $ssbase`

These names map onto internal processor and kernel LDT data structures which hold the base addresses associated with the `$cs`, `$ds`, `$es`, `$fs`, `$gs`, and `$ss` registers, respectively. They are useful particularly for determining the location of thread-specific data. For instance, if a disassembly address mode references memory with `%fs:8` or `%gs:(%eax)`, then the location can be determined in NightView with `$fsbase+8` or `$gsbase+$eax`, respectively.

`$st0 through $st7`

These names map onto the 8 floating point registers. The floating point registers on the IA-32 always hold 80 bit double extended precision (i.e. long double) values. Note that the architecture defines these registers as a stack. Also note that these registers are aliases of the registers `$mm0` through `$mm7`.

`$r0 through $r7`

These names map onto `$st0` through `$st7`, but always referenced as though the floating point stack pointer were zero.

`$cwd and $fctrl`

These names map onto the floating point control register. They are synonyms.

`$swd or $fstat`

These names map onto the floating point status register. They are synonyms.

`$twd or $ftag`

These names map onto the floating point tag word register. They are synonyms. These names may be used in the **info registers** command or in expressions in the **set** and **print** commands, but not in patchpoints or eventpoint conditions. This register may be read, but not modified. See `$fxtag`.

`$fxtag`

This name maps to the floating point tag word register, but in a different form from `$ftag`. The form of this register is one byte with each bit corresponding to a floating point register. This name does not have the restrictions of `$ftag`.

`$fip or $fioff`

These names map onto the lower 32 bits of the floating point last instruction pointer register. They are synonyms.

`$fcs or $fiseg`

These names map onto the upper 16 bits of the floating point last instruction pointer register. They are synonyms.

`$fop`

This name maps onto the floating point opcode register.

`$foo` or `$fooff`

These names map onto the lower 32 bits of the floating point last data (operand) pointer register. They are synonyms.

`$fos` or `$foseg`

These names map onto the upper 16 bits of the floating point last data (operand) pointer register. They are synonyms.

`$dr0` through `$dr3`

These names map onto the debug address registers. NightView uses these registers to implement watchpoints. See “Watchpoints” on page 3-14. Users should not modify these registers.

`$dr6`

This name maps onto the debug status register. NightView uses this register to implement watchpoints and single step. See “Watchpoints” on page 3-14, “step” on page 6-135, “stepi” on page 6-138, “next” on page 6-136, and “nexti” on page 6-139. Users should not modify this register.

`$dr7`

This name maps onto the debug control register. NightView uses this register to implement watchpoints. See “Watchpoints” on page 3-14. Users should not modify this register.

`$mm0` through `$mm7`

These names map onto the 64 bit vector registers available with Intel MMX Technology. Note that these registers are aliases of the registers `$st0` through `$st7`. However, the `$mm0` through `$mm7` registers are not defined as a stack.

`$xmm0` through `$xmm7`

These names map onto the 128 bit vector registers available with the Streaming SIMD Extensions (SSE).

`$mxcsr`

This name maps onto the SSE MXCSR control and status register.

Note that the floating point, debug, MMX, and SSE registers are not normally displayed by the **info registers** command (see “info registers” on page 6-168). If you want to display those registers, you can do so with the following commands:

info registers st.*	<i>displays floating point registers</i>
info registers dr.*	<i>displays debug registers</i>
info registers mm.*	<i>displays MMX registers</i>
info registers xmm.*	<i>displays SSE registers</i>
info registers .*	<i>displays all registers</i>

AMD64 Registers

AMD64 machines are based on the AMD AMD64 architecture (see *AMD64 Architecture Programmer's Manual* for architectural details). See “info registers” on page 6-168.

In addition to the common register definitions for stack pointer (\$sp), frame pointer (\$fp), and program counter (\$pc), the AMD64 machines support the registers shown in the following table.

Table 6-4. AMD64 Registers (iHawk Series 870)

\$rax, \$rbx, \$rcx, \$rdx, \$rsi, \$rdi, \$rbp, \$rsp, \$r8 through \$r15

These names map onto the 16 general purpose registers. Note that \$sp is the same as \$rsp, and \$fp may be the same as \$rbp, depending on how the compiler generates code.

\$eax, \$ebx, \$ecx, \$edx, \$esi, \$edi, \$ebp, \$esp, \$r8d through \$r15d

These names map onto the lower 32 bits of each of the 16 general purpose registers mentioned above, respectively.

\$ax, \$bx, \$cx, \$dx, \$si, \$di, \$bp, \$sp16, \$r8w through \$r15w

These names map onto the lower 16 bits of each of the 16 general purpose registers mentioned above, respectively. Note that the lower 16 bits of the RSP register is more commonly known as simply SP. But the name \$sp is reserved for an architecture-independent stack pointer in NightView. So the name \$sp16 is used for the lower 16 bits of the RSP register, instead.

\$al, \$bl, \$cl, \$dl, \$sil, \$dil, \$bpl, \$spl, \$r8b through \$r15b

These names map onto the low order 8 bits of each of the AX, BX, CX, DX, SI, DI, R8W, R9W, R10W, R11W, R12W, R13W, R14W, and R15W registers, respectively. In other words, they map onto bits 0-7 of each of the RAX, RBX, RCX, RDX, RSI, RDI, RBP, RSP, R8, R9, R10, R11, R12, R13, R14 and R15 registers, respectively.

\$ah, \$bh, \$ch, \$dh

These names map onto the high order 8 bits of each of the AX, BX, CX and DX registers, respectively. In other words, they map onto bits 8-15 of each of the RAX, RBX, RCX, and RDX registers, respectively.

\$eflags

The program status and control register. NightView and the kernel use the TF flag of this register to implement stepping. See “step” on page 6-135, “stepi” on page 6-138, “next” on page 6-136, and “nexti” on page 6-139. Users should not modify the TF field of the \$eflags register. Other flags in this register are used by the kernel. Care should be taken if modifying this register.

\$rip, \$eip

\$rip is the instruction pointer register. This is the same as the \$pc register. \$eip is the lower 32 bits of the \$rip.

`$fs, $gs`

The data segment registers. Care should be taken if modifying `$fs` or `$gs`.

`$fsbase, $gsbase`

These names map onto the `FS.base` and `GS.base` model-specific registers or internal processor and kernel LDT data structures which hold the base addresses associated with the `$fs` and `$gs` registers, respectively. They are useful particularly for determining the location of thread-specific data. For instance, if a disassembly address mode references memory with `%fs:8` or `%gs:(%eax)`, then the location can be determined in NightView with `$fsbase+8` or `$gsbase+$eax`, respectively.

`$st0 through $st7`

These names map onto the 8 floating point registers. The floating point registers on the AMD64 always hold 80 bit double extended precision (i.e. long double) values. Note that the architecture defines these registers as a stack. Also note that these registers are aliases of the registers `$mm0` through `$mm7`.

`$fpr0 through $fpr7`

These names map onto `$st0` through `$st7`, but always referenced as though the floating point stack pointer were zero.

`$cwd and $fctrl`

These names map onto the floating point control register. They are synonyms.

`$swd or $fstat`

These names map onto the floating point status register. They are synonyms.

`$twd or $ftag`

These names map onto the floating point tag word register. They are synonyms. These names may be used in the **info registers** command or in expressions in the **set** and **print** commands, but not in patchpoints or eventpoint conditions. This register may be read, but not modified. See `$fxtag`.

`$fxtag`

This name maps to the floating point tag word register, but in a different form from `$ftag`. The form of this register is one byte with each bit corresponding to a floating point register. This name does not have the restrictions of `$ftag`.

`$frip`

This name maps onto the 64-bit floating-point last instruction pointer register.

`$fioff`

This name maps onto the lower 32 bits of the floating-point last instruction pointer register.

\$frdp

This name maps onto the 64-bit floating-point last data (operand) pointer register.

\$fcs or \$fiseg

These names map onto the upper 16 bits of the floating point last instruction pointer register. They are synonyms.

\$fop

This name maps onto the floating point opcode register.

\$foo or \$fooff

These names map onto the lower 32 bits of the floating point last data (operand) pointer register. They are synonyms.

\$fos or \$foseg

These names map onto the upper 16 bits of the floating point last data (operand) pointer register. They are synonyms.

\$dr0 through \$dr3

These names map onto the debug address registers. NightView uses these registers to implement watchpoints. See “Watchpoints” on page 3-14. Users should not modify these registers.

\$dr6

This name maps onto the debug status register. NightView uses this register to implement watchpoints and single step. See “Watchpoints” on page 3-14, “step” on page 6-135, “stepi” on page 6-138, “next” on page 6-136, and “nexti” on page 6-139. Users should not modify this register.

\$dr7

This name maps onto the debug control register. NightView uses this register to implement watchpoints. See “Watchpoints” on page 3-14. Users should not modify this register.

\$mm0 through \$mm7

These names map onto the 64 bit vector registers available with Intel MMX Technology. Note that these registers are aliases of the registers \$st0 through \$st7. However, the \$mm0 through \$mm7 registers are not defined as a stack.

\$xmm0 through \$xmm15

These names map onto the 128 bit vector registers available with the Streaming SIMD Extensions (SSE).

\$xmm0gf through \$xmm15gf

These names map onto the single low-order 8-byte floating-point value for each of the \$xmm0 through \$xmm15 registers, respectively.

`$xmm0wf` through `$xmm15wf`

These names map onto the single low-order 4-byte floating-point value for each of the `$xmm0` through `$xmm15` registers, respectively.

`$mxcsr`

This name maps onto the SSE MXCSR control and status register.

`$ymm0` through `$ymm15`

These names map onto the 256 bit vector registers available with the Advanced Vector Extensions (AVX). The lower halves of these registers are aliases for the `$xmm0` through `$xmm15` registers.

`$ymm0h` through `$ymm15h`

These names map onto upper 128 bits of the 256 bit vector registers available with the Advanced Vector Extensions (AVX). The corresponding 128 bit lower halves of these registers are available as `$xmm0` through `$xmm15`. The whole 256 bit registers are available as `$ymm0` through `$ymm15`.

Note that the floating point, debug, MMX, SSE, and AVX registers are not normally displayed by the **info registers** command (see “info registers” on page 6-168). If you want to display those registers, you can do so with the following commands:

info registers st.*	<i>displays floating point registers</i>
info registers dr.*	<i>displays debug registers</i>
info registers mm.*	<i>displays MMX registers</i>
info registers xmm.*	<i>displays SSE registers</i>
info registers ymm.*	<i>displays AVX registers</i>
info registers .*	<i>displays all registers</i>

CUDA Registers

The CUDA architecture is used by CUDA code running on CUDA devices under the control of a host application. In addition to the common register definitions for stack pointer (`$sp`), frame pointer (`$fp`), and program counter (`$pc`), the CUDA architecture supports a uniform set of registers named `$Rn`. Note that case is significant in these register names. The number of such registers is determined by the particular CUDA architecture. For instance, Tesla (1.x) and Fermi (2.x) architectures support `$R0` through `$R63`, whereas Kepler (3.x) architectures support `$R0` through `$R255`.

CUDA registers are only available when the process is stopped and the current context is a CUDA thread. When the current context is a CUDA thread, registers for the host architecture are unavailable.

See “CUDA Debugging” on page 3-46 for a description of NightView’s CUDA support.

Casts of Vector Registers

Some IA-32 and AMD64 registers are vector registers and are capable of holding vectors of differing configurations (e.g. two `double` elements, eight `short` elements, four `int` elements, etc.). If a scalar type cast is applied to such a vector register, the register is interpreted as an array of the scalar type, up to the size of the register. For example, consider the following expressions to display register `$xmm0`:

```
print (double)$xmm0
$0: (double)$xmm0 = {[0] = 1.2345, [1] = 6.7890}
print (int)$xmm0
$1: (int)$xmm0 = {[0] = 309237645, [1] = 1072939139,
[2] = -1649267442, [3] = 1075521519}
```

Special Intrinsics

For general information about expression evaluation, see “Expression Evaluation” on page 3-25. NightView provides an additional intrinsic function to allow access to attributes of eventpoints.

`__eventpoint_data__`

`__eventpoint_data__(eventpoint)`

The `__eventpoint_data__` intrinsic takes a single *eventpoint* parameter, which identifies the eventpoint. The parameter may be any of the following:

- An integer expression which evaluates to the eventpoint number. If `__eventpoint_data__` is used within a patchpoint expression or eventpoint condition, then *eventpoint* must be a compile-time constant (usually just the literal eventpoint number).
- The name of an eventpoint (see “name” on page 6-107 or the **name** parameter to any eventpoint command). The name must be specified as an identifier and not as a quoted string, and must resolve to exactly one eventpoint. It is resolved to the associated eventpoint when the command is executed, so subsequent changes to eventpoint names will have no effect on it.
- The special `.` (period) syntax. This may be used only if `__eventpoint_data__` is used within an expression associated with an eventpoint (patchpoint expression, eventpoint condition, eventpoint command, etc.). This syntax indicates the associated eventpoint.

No matter how the expression is specified, the indicated eventpoint must be an inserted eventpoint (see “Eventpoints” on page 3-9), or the intrinsic will fail. Notably, this means that it cannot indicate a watchpoint or syscallpoint.

The result of the `__eventpoint_data__` intrinsic is a *pointer* to the following C struct:

```
struct eventpoint_data {
    int    IgnoreCount;
```

```

int    CrossCount;
int    HitCount;
short EnabledFlag;
short EnabledOnce;
};

```

The meanings of the members are as follows:

- `IgnoreCount` is the remaining number of crossings to be ignored (“ignore” on page 6-124).
- `CrossCount` is the number of times program execution has crossed the eventpoint since the program started execution (even if the ignore count or condition was not satisfied, this count is incremented).
- `HitCount` is the number of times the eventpoint has been *hit* since the program started execution (this count is incremented only if the condition and ignore count are satisfied).
- `EnabledFlag` is a boolean indicating whether or not the eventpoint is enabled (see “enable” on page 6-123 and “disable” on page 6-122).
- `EnabledOnce` is a boolean indicating whether or not the eventpoint is marked to be automatically disabled after it was hit once (see **enable /once** under “enable” on page 6-123, “tbreak” on page 6-125, and “tpatch” on page 6-126).

This intrinsic can be used to create unusual condition checks on eventpoints. For instance:

```
(local) break location if __eventpoint_data__(.)->CrossCount % 12 == 11
```

For more information on the precise interactions of these members, see “Interactions Between Conditions, Ignore Counts, etc.” on page 3-12.

Location Specifiers

A *location-spec* is used in various commands to specify a location in the executable program. It can be any of the following:

function_name or *unit_name*

specifies the beginning of the named function

file_name:line_number

specifies the first instruction generated for the given line in the given file

line_number:function_name or *line_number:unit_name*

specifies the first instruction for the given line in the file containing the given function or unit

file_name: function_name

specifies the beginning of the specified function declared in the given file (this is required for `static` functions that are not globally visible)

line_number

specifies the first instruction generated for the given line in the current file

**expression*

specifies the address given by *expression*

If a location specifier is omitted, then the location used is the next instruction to be executed in the current stack frame. See “Current Frame” on page 3-29.

Function names always refer to the location of the first instruction following any prologue code (the *prologue* is code that allocates local stack space, saves the return address, etc.). To refer to the actual entry point of a function, use the **expression* form and write an expression that evaluates to the function entry point address (in C language mode, this would look like `&function`).

NOTE

A location specifier may sometimes designate multiple locations; for instance, a line number within a procedure that has been expanded inline several times will designate every location where that procedure was expanded. If such a location specifier is used to set an eventpoint (see “Manipulating Eventpoints” on page 6-106), NightView will set the eventpoint at each of the corresponding locations. An eventpoint set at multiple locations is still considered to be a single eventpoint. If you wish to set an eventpoint at some subset of the locations that are implied by a particular location specifier, the **info line** command (see “info line” on page 6-183) may be used to determine the locations corresponding to the particular location specifier. The **expression* form of location specifier may then be used to designate the proper location.

Wherever a *file_name* appears, it may be enclosed in double quotes. This is necessary if the *file_name* contains special characters.

Wherever a function name appears in a location specifier, it may also appear with an overloading suffix to distinguish between multiple functions with the same name (see “Selecting Overloaded Entities” on page 6-2). The names of operator functions in C++ may also be used as function names. In C++ the keyword `operator` should be followed by the operator name (the same syntax used to declare operator functions in the language). Because the function name form of operator functions is always used in location specifiers, the only **set-overload** mode which affects location specifiers is the *routine* mode (see “set-overload” on page 6-73).

All commands that accept a *location-spec* argument allow the keyword **at** to precede the *location-spec*. In most cases, the **at** keyword is optional, but a few commands require it

to be present. The syntax of each command indicates whether the keyword is required or optional.

Qualifier Specifiers

Qualifiers are used to apply NightView commands to specific processes or dialogues. A qualifier is simply a list of *qualifier specifiers*, each specifier representing one or more processes or dialogues. You can supply a qualifier explicitly, in parentheses as a prefix to the command, or implicitly, by using the **set-qualifier** command (see “set-qualifier” on page 6-64). In a prefix qualifier, the list of specifiers is separated by either blanks or tabs.

Each *qualifier specifier* can be any one of the following items:

family-name

A name given by you to a set of processes, called a *family*. See “family” on page 6-51.

dialogue-name

The name of a dialogue in your NightView session. This is usually the name of the system on which the dialogue is running, but you may also specify a different name (see “login” on page 6-25). In contexts where the qualifier is being used to specify a set of processes, a *dialogue-name* refers to all the processes being debugged in that dialogue.

PID

The numeric value of the process ID of one of the processes being debugged by NightView. You can use this form only if the process ID is unique among all the processes being debugged. This may not be true if you have multiple dialogues, but it is always true if you have only one dialogue.

If your process is threaded, the threads are implemented as multiple processes sharing resources (thread processes). See “See “list” on page 6-82. See “step” on page 6-135. Multithreaded Programs” on page 3-44. You may use the PID of any of those thread processes in a qualifier. NightView considers them all to refer to the same process.

dialogue-name:PID

This is how you specify a particular process when processes in different dialogues have the same process ID.

all

This keyword designates all processes or dialogues known to NightView.

auto

This keyword designates the one process that is currently stopped and has been stopped for the longest time. You may want to specify `auto` as your default qualifier if you want to work on only one process at a time (see “set-qualifier” on page

6-64). NightView gives you an error message if you use `auto` when there are no processes stopped.

Note that, because a qualifier specifier can be either a family name or a dialogue name, you cannot have a dialogue with the same name as a process family.

In general, the specifiers in a qualifier are not *evaluated* until a qualified command requests the information. A qualifier is evaluated when a command qualified by it needs the information; that is, when the command is applied to the processes or dialogues in the qualifier. Certain NightView commands ignore their qualifier, so they do not request evaluation of the specifiers in the qualifier. This has several effects on you:

- A *family-name* appearing in a qualifier may remain undefined until a command requires evaluation of the qualifier. You may also change the definition of a *family-name* currently in use in a qualifier; such a change will affect the next command that evaluates that qualifier.
- Evaluating a *dialogue-name* yields all the processes in the dialogue at the time of the evaluation. Since evaluation is generally delayed until the last possible moment, using a *dialogue-name* is usually a good way to reference all the currently-existing processes in a dialogue.
- The specifiers `all` and `auto` are evaluated at the time a command is actually executed.

Eventpoint Specifiers

Eventpoints may be grouped together and assigned a name (see “name” on page 6-107). In addition, the name ``.`` is a reserved name that always refers to the set of eventpoints most recently created by a single command. (If an eventpoint creation command fails, the definition of ``.`` is cleared.) Eventpoint numbers and eventpoint names are two types of *eventpoint specifiers*.

Another kind of eventpoint specifier is a location-spec. The location-spec must begin with the keyword `at`. See “Location Specifiers” on page 6-15. A location-spec eventpoint specifier with a line number refers only to eventpoints set at the beginning of that source line, not to any eventpoints that may be set on addresses within the line. Note also that a location-spec eventpoint specifier may refer to multiple locations, such as when a breakpoint is set in an inline function that is expanded multiple times.

Some commands expect more arguments after the eventpoint specifier. These commands do not accept a location-spec as an eventpoint specifier, because a location-spec eventpoint specifier must be the last argument.

Eventpoint specifiers that refer only to breakpoints may also be called *breakpoint specifiers* (*tracepoint specifiers*, *patchpoint specifiers*, *monitorpoint specifiers*, *heappoint specifiers*, *watchpoint specifiers* and *syscallpoint specifiers* are similarly defined).

Regular Expressions

A *regex* is used by many of the commands to specify a pattern used to match against a set of names (like variable names or register names in the **info** commands). Regular expressions may be case-sensitive or case-insensitive depending on the **set-search** command (see “set-search” on page 6-73).

Regular expressions are similar to wildcard patterns, but are more flexible. Regular expressions and wildcard patterns are used for different things in the debugger (see “Wildcard Patterns” on page 6-20). The descriptions of the commands tell if they take a regular expression or a wildcard pattern.

The regular expression syntax recognized is similar to that recognized by many other common tools, but the details (as always) vary somewhat.

Table 6-5. Regular Expressions

.	A dot matches any character except a newline.
*	A star matches zero or more occurrences of the preceding regular expression. For example, <code>.*</code> matches zero or more of any character except a newline.
+	A plus matches one or more occurrences of the preceding regular expression.
{ <i>m</i> }	Matches exactly <i>m</i> occurrences of the preceding regular expression.
{ <i>m</i> , }	Matches <i>m</i> or more occurrences of the preceding regular expression.
{ <i>m</i> , <i>n</i> }	Matches from <i>m</i> to <i>n</i> occurrences of the preceding regular expression.
^	A caret matches at the beginning of a string.
\$	A dollar sign matches at the end of a string.
()	Parentheses are used to group regular expressions.

[]

Brackets define a set of characters, any one of which will match. Within the brackets, additional special characters are recognized:

^

If the first character inside the brackets is a caret, then the set of characters matched will be the inverse of the set specified by the remaining characters in the brackets.

-

A range of characters may be matched by specifying the starting and ending characters in the range separated by a dash.

To define a set that includes a - character, specify the dash as the first or last character in the set.

Any other character matches itself.

To literally match one of the special characters defined above, use a backslash (\) character in front of it (to literally match a backslash, use two of them (\)).

The *m* and *n* match counts above must be positive integers less than 256.

Most commands that use regular expressions do not require the use of '^' and '\$' because they implicitly assume that an *anchored* match is called for. Other commands (such as the **forward-search** and **reverse-search** commands) assume that only a partial match is called for (and does not imply an *anchored* match). The description of each command that uses regular expressions specifies whether or not it implicitly assumes its regular expressions are to be anchored.

If you do not need the full expressive power of regular expressions, you can just use a normal string.

Examples:

r[1-5]

This example matches the strings 'r1', 'r2', 'r3', 'r4', and 'r5'. This might be a good expression to match register names.

child_pid

This example matches only the string 'child_pid'. This might be a good expression to match a program variable name.

Wildcard Patterns

Wildcard patterns are used by the commands **debug**, **nodebug** and **on program**. See “debug” on page 6-27, “nodebug” on page 6-28, and “on program” on page 6-47.

Wildcard patterns are similar to regular expressions, but are usually more convenient for

representing file names. See “Regular Expressions” on page 6-19.

If the wildcard pattern starts with a `/`, it is assumed to be a pattern that must match a complete absolute path name. Otherwise the pattern is matched against the rightmost (trailing) components of the program name. Patterns are always matched to component boundaries. Spaces and tabs are not allowed in wildcard patterns.

Wildcards are similar to wildcards in **sh**.

Table 6-6. Wildcard Patterns

Matches zero or more characters (but does not match a `/`).

{ [chars] }

Matches any of the characters in the set. A dash (`-`) can be used to separate a range of characters and a leading bang (`!`) matches any characters except the ones in the set (but not a `/`).

?

Matches any single character (except a `/`).

Any other character matches itself.

Unlike **sh**, a ***** matches a leading dot (`.`) in a file name.

If you do not need the full expressive power of wildcards, you can just use the file name.

Examples:

/bin/*

This matches any file in the directory **/bin**.

test*

This matches any file that begins with the letters **test**, in any directory.

***.c**

This matches any source file that ends with **.c**, in any directory. This might be a good expression to match file names.

/usr/bob/myprog

This matches only the file **/usr/bob/myprog**.

Repeating Commands

A line typed from an interactive terminal consisting solely of a newline (no other characters, including blanks) generally causes NightView to repeat the previous

command. Note that the blank line must come from an interactive device; a blank line in a macro or in a disk file read by the **source** command does not cause repetition. The command that gets repeated may, however, come from a macro.

Not all commands can be repeated in this manner. In general, commands whose result would not be any different when repeated will not repeat. Typing a blank line after a non-repeating command has no effect; it acts the same as a comment. If the description of a command does not say it is repeatable, then it isn't.

A few commands, such as **list** or **x**, alter their behavior slightly when repeated: instead of exactly repeating the command, they typically repeat the action on a different set of data. These differences in behavior are documented in the description of the command.

In the following examples, assume all commands were entered interactively.

```
(local) list 20:func
(local)
(local)
```

In this example, lines 16-25 (approximately) of function **func** would be listed by the **list** command. Repeating this command lists the next set of 10 lines, lines 26-35. Note that **list** is one of the commands whose behavior changes when it is repeated.

```
(local) define mac(ln) as
> list @ln:func
> end define
(local) @mac(20)
(local)
(local)
```

This example is equivalent to the previous one. It demonstrates that the repeated command may come from a macro.

```
(local) define mac(vn) as
> x/20x @vn
> echo
> end define
(local) @mac(xstruct)
(local)
(local)
```

This example demonstrates how to write a macro that does not repeat at all. Since **echo** is a non-repeating command, entering a blank line after the **@mac(xstruct)** line does nothing.

Replying to Debugger Questions

This section describes how to respond when the debugger asks you a question.

Certain forms of some debugger commands are considered unsafe and will check the debugger's safety level (see "set-safety" on page 6-67) before executing. When the safety level is **verify**, these commands will ask a question of the user and wait for

verification. The possible responses to the question are always "yes" and "no" (case insensitive). These responses may be abbreviated to their first letter if desired. The response must be terminated by a carriage return.

A "yes" response indicates that the unsafe action is to be performed.

A "no" response indicates that the unsafe action is *not* to be performed.

In the graphical user interface, the debugger pops up a warning dialog box.

Controlling the Debugger

This section describes how to exit NightView, and the commands used to control debugged processes and your interaction with them.

Quitting NightView

quit

Stop everything. Exit the debugger.

quit

Abbreviation: **q**

This command terminates the debugger. If the safety level (see “set-safety” on page 6-67) is `forbid`, you will not be allowed to quit unless there are no processes being debugged. In other safety levels, any active processes will be killed when you quit. If the safety level is `verify`, you will be prompted for confirmation before quitting causes any debugged processes to be killed (see “Replying to Debugger Questions” on page 6-22).

The processes killed include all active processes started in any dialogue shell and not explicitly detached. NightView detaches from any processes that are being controlled but are not being debugged by you because of a **nodebug** command. See “Detaching” on page 3-3. See “nodebug” on page 6-28.

Processes started using the **shell** command are independent of the debugger, and are not affected by a **quit**.

Managing Dialogues

A *dialogue* is an interaction with a particular host system for the purpose of debugging one or more processes on that system under a particular user name. You may have as many dialogues as you wish; there can even be more than one dialogue with a particular host system. Dialogues are described in more detail in the Concepts chapter (see “Dialogues” on page 3-4).

login

Login to a new dialogue shell.

```
login [/conditional] [/popup] [name=dialogue name]  
[user=login name] [others ...] machine
```

NOTE

If present, the options `/conditional` and `/popup` must appear before the machine name and before any keywords.

The **login** command takes many keyword parameters. The most commonly used are:

`/conditional`

Ignore this **login** command if a dialogue with this name already exists. This is useful from macros (see “Defining and Using Macros” on page 6-184) and for other programs that communicate with NightView.

`/popup`

Pop up the Remote Login Dialog Box (see “Remote Login Dialog Box” on page 8-41) initialized with the machine name and the values of the `name=` and `user=` keywords. No other keywords are allowed with this option. This option is meaningful only in the graphical user interface.

`name=dialogue name`

Give this parameter to specify a name for the dialogue you are creating. If you leave it off, the dialogue name is the same as the name of the machine running the dialogue. To run multiple dialogue shells on the same machine you must give them unique names. No dialogue name may be the same as a family name (see “family” on page 6-51). A dialogue name must start with an alphabetic character and may be followed by any number of alphabetic, numeric or underscore characters.

`user=login name`

Login as this user. Normally your current user name is used, but you may login as any user.

machine

Specify the machine where the programs to be debugged are located and the dialogue shell will run. This is a required parameter. It may be a host name, with or without domain qualification, or it may be an IP address.

The following parameters are less frequently used, but are provided to allow you to control the execution environment of the remote dialogue.

nice=nice value

The dialogue normally runs with normal interactive priority. A positive nice value lowers the priority (makes other programs seem more important). You must have special privileges to specify a negative nice value.

cpu=cpu list

Set the CPU bias for the dialogue. *cpu list* is a comma-separated list of CPU IDs or CPU ID ranges. For example: "**0,2-4,6**". *cpu list* may also be **active** or **boot** to specify all active processors or the boot processor, respectively.

priority=value

Specify the priority of the remote dialogue processes. The scheduling policy determines what values may be specified for the priority. *value* must be an integer value that is valid for the current scheduling policy. Higher numerical values represent more favorable scheduling priorities.

scheduling=sched_keywords

Control the scheduling policy that will be used for the dialogue. The allowed keywords are: **sched_fifo**, **fifo**, **sched_rr**, **rr**, **sched_other**, and **other**.

quantum=time

Control the time slice quantum size for the process. A quantum value is meaningful only under the **sched_rr** and **sched_other** scheduling policies. *time* is specified either as a nice value or a millisecond value corresponding to a nice value. Nice values must be between -20 and 19 inclusive. By default, a quantum value of -20 results in a ~300ms slice, and a quantum value of 19 results in a ~10ms slice. Millisecond values are in the form *numberms* and must reflect the times defined in the system for the nice values. If a non-defined millisecond value is supplied, an "unsupported quantum" error message is returned.

The **cpu**, **scheduling**, **priority**, and **quantum** parameters all accept the same arguments as the corresponding options on the **run(1)** command — see the man page for details.

Any programs started in the dialogue shell will inherit all the above parameters. The **run(1)** command can control all these parameters, and may be used within the dialogue shell to debug programs and change the parameters.

When you use the **login** command you are asked for a password. See "Remote

Dialogues” on page 3-6 for a general discussion of how to use remote dialogues.

Example:

```
(afamily) login fred
To begin a remote debug session on 'fred', enter the
password for user 'wilma'.
Password: enter wilma's password
(afamily) login user=barney name=fredII fred
To begin a remote debug session on 'fred', enter the
password for user 'barney'.
Password: enter barney's password
(afamily)
```

The above example shows the creation of two new dialogues. The first **login** command starts a dialogue on a machine named *fred* and logs in as the current user (*wilma* in this example). This dialogue is named *fred*, because no explicit name was given.

The second creates a dialogue on machine *fred* named *fredII*. In this case the user logged into *fred* is *barney*.

The **login** command is creating a new dialogue, so the qualifier has no effect on this command.

debug

Specify names for programs you wish to debug.

debug *pattern* . . .

pattern

A wildcard pattern matching the name of a program to be debugged. Spaces and tabs are not allowed in *pattern*. See “Wildcard Patterns” on page 6-20.

This command and its inverse (see “nodebug” on page 6-28) allow you to control which programs get debugged. The list of programs applies to the individual dialogues specified in the **debug** command qualifier (different dialogues may have different lists of programs to be debugged).

The **debug** and **nodebug** commands work by remembering the list of **debug** and **nodebug** commands. When a new file needs to be checked to see if it should be debugged, the name is first compared to the pattern in the most recent command, then the pattern in the next most recent command, and so on.

The first pattern that matches the file name determines what to do with the associated process. If the matching pattern is on a **debug** command, then the process will be debugged. If it was on a **nodebug** command, then the process will not be debugged.

The pattern ***** matches everything, so the list of patterns is always reset when ***** appears as an argument. Since each dialogue always starts with either **debug *** or **nodebug *** in the list, it is impossible to pick a file name that does not match at some point in the list.

The default pattern list for a dialogue is:

```

nodebug /bin/* /sbin/* /usr/X11R6/bin/* /usr/ada/*/bin/*
        /usr/ada/bin/* /usr/bin/* /usr/bin/X11/* /usr/ccs/*/*
        /usr/ccs/*/*/*/* /usr/kerberos/bin/* /usr/kerberos/sbin/*
        /usr/lib/* /usr/lib/*/* /usr/lib/*/*/* /usr/lib/*/*/*/*
        /usr/lib/*/*/*/*/* /usr/lib/*/*/*/*/*/*
        /usr/lib/*/*/*/*/*/* /usr/local/bin/* /usr/local/sbin/*
        /usr/sbin/* /usr/ucb/* /usr/ucblib/*
debug *

```

To print the list of **debug** and **nodebug** patterns, see “info dialogue” on page 6-173.

nodebug

Specify names for programs you do not wish to debug.

nodebug *pattern* . . .

pattern

A wildcard pattern matching the name of a program to avoid debugging.

This command is typically used in combination with the **debug** command to control which programs are debugged in a dialogue. The complete syntax of wildcards and the algorithm used to match files is described in the **debug** command (see “debug” on page 6-27).

Example:

```

(afamily) nodebug *
(afamily) debug x*

```

This example uses **nodebug** ***** to turn off all debugging. It then uses **debug** to turn on debugging for any programs started where the basename begins with the letter **x**.

Note that even if one command is not debugged, its children may be debugged. To avoid debugging a command as well as any children, you must use the **detach** command (see “detach” on page 6-41).

To print the list of **debug** and **nodebug** patterns, see “info dialogue” on page 6-173.

set-debug-file-directory

Tell NightView where to look for .debug files.

set-debug-file-directory [*path*]

path

The name of the directory in which to find .debug files.

The **set-debug-file-directory** sets the directory to use when searching for .debug files associated with shared libraries, for each dialogue in the qualifier. With no argument, NightView prints the current setting.

The default path is **/usr/lib/debug** which is where .debuginfo rpms usually install their files.

If you don't care about debugging libraries, you can improve performance by setting this to a directory that does not exist.

translate-object-file

Translate object filenames for a remote dialogue.

translate-object-file [*from* [*to*]]

Abbreviation: **xl**

from

The filename or filename prefix as seen by the remote system.

to

The filename or filename prefix as seen by the local system.

A file residing on a system other than the target system can be specified using the form *user@host:/path*. NightView will download this file from the specified system to the host. See “Remote File Access” on page 3-7.

If both *from* and *to* are present, a translation is added. If only *from* is present, the translation exactly matching *from* is removed. If neither is present, all translations are removed.

NOTE

from and *to* are *not* wildcard patterns or regular expressions. See “Wildcard Patterns” on page 6-20. See “Regular Expressions” on page 6-19.

The **translate-object-file** command manages translations for object filenames for each dialogue in the qualifier. Translations are useful when:

- An object file is visible from both systems, but its position in the file system is different. For example, **/usr** on system *fred* may be mounted as **/fred/usr** on the local system.
- An object file is not visible from the local system, but you have a copy of the file. For example, you might have a development directory from which the image on the remote system is created.
- The object file on the remote system has been stripped, but you have a copy with debugging information.

Object filenames from **exec-file** and **load** commands are subject to object filename translation. See “exec-file” on page 6-46. See “load” on page 6-104. Dynamic library

names are also subject to object filename translation. See “Debugging with Shared Libraries” on page 3-51. Object filenames from **symbol-file** commands are *not* subject to object filename translation. See “symbol-file” on page 6-42.

NightView attempts to match translations to the initial characters of the filename. Filename component boundaries are not treated as a special case. If you want to match to component boundaries, include slashes in the strings. NightView tries all translations that match the strings, beginning with the longest matching translation, until it finds a translated filename with the same text segment contents as the executing program. If no file is found with the same text segment contents, NightView gives a warning and uses the first translation that matched the object filename.

NightView automatically supplies a default set of translations when a remote dialogue is created. The default set is made by inspecting the local system mount table and by considering the set of cross-development environments on the local system. In many cases, these translations are sufficient; additional translations are not necessary.

Translate-object-file commands take effect in existing processes as well as future ones.

Examples:

Suppose the object files that exist on the remote system under the directory **/wilma/pebbles** exist on the local system under the directory **pebbles** (relative to your current working directory).

```
(fred) xl /wilma/pebbles/ pebbles/
```

This command translates any object filename beginning with the string **/wilma/pebbles/** to the same filename with **/wilma/pebbles/** replaced by **pebbles/**. For example, **/wilma/pebbles/hair** becomes **pebbles/hair**. Note that **pebbles/hair** will be evaluated relative to NightView's current working directory. See “pwd” on page 6-81.

Suppose the object files that exist on the remote system under **/betty** exist on the local system under **/barney**. However, the files under **/betty** whose name begins with **bam** should be found under **/dino**.

```
(fred) xl /betty/ /barney/
(fred) xl /betty/bam /dino/bam
```

These commands translate any object filename beginning with the string **/betty/** to the same filename with **/betty/** replaced by **/barney/** and any object filename beginning with the string **/betty/bam** to the same filename with **/betty/bam** replaced by **/dino/bam**. NightView picks **/betty/bam** in preference to **/betty/** because **/betty/bam** is longer. For example,

/betty/dress	becomes	/barney/dress
/betty/bambam	becomes	/dino/bambam
/betty/bambino	becomes	/dino/bambino

A good place to put a **translate-object-file** command is in an **on dialogue** command in your **.NightViewrc** file. See “on dialogue” on page 6-31. Also, see “Initialization Files” on page 3-43.

Example:

```
(all) on dialogue fred.* do
>      xl /usr/ /fred/usr/
>      end on dialogue
```

This command translates the directory **/usr** on the remote system to the directory **/usr/fred** on the local system, for dialogues whose name begins with **fred**.

logout

Terminate a dialogue.

logout

The **logout** command terminates any dialogues named in the command qualifier. If your safety level is **unsafe** then *all* processes being debugged in the dialogues are killed (see “set-safety” on page 6-67). If your safety level is **verify** then you are prompted for confirmation before the logout causes any debugged processes to be killed (see “Replying to Debugger Questions” on page 6-22). If your safety level is **forbid**, then the logout does not occur. If you want any processes to continue running, you must **detach** them prior to using **logout** (see “detach” on page 6-41). NightView detaches from any processes that are being controlled but are not being debugged by you because of a **nodebug** command. See “Detaching” on page 3-3. Also, see “nodebug” on page 6-28.

If the dialogue shell is still running at **logout** time, it is killed (you may send an **exit** command to the shell to terminate it normally prior to logging out).

Example:

```
(adialogue) detach
(adialogue) !exit
(adialogue) logout
```

The example shows how to avoid having any processes killed. The **detach** command allows all processes in the dialogue to continue running independently of the debugger. The **!exit** command sends an exit command to the dialogue shell to terminate it normally, then the **logout** command terminates the debugger dialogue.

on dialogue

Specify debugger commands to be executed when a dialogue is created.

```
on dialogue [regexp]

on dialogue regexp command

on dialogue regexp do

    regexp
```

A regular expression to match against the names of newly created dialogues. See “Regular Expressions” on page 6-19.

command

A debugger command to be executed when a new dialogue whose name matches *regex* is created.

In the third form of the **on dialogue** command, the debugger commands to be executed must begin on the line following the **do** keyword. The list of debugger commands to execute is terminated when a line containing only the words **end on dialogue** is encountered.

The **on dialogue** command allows a user-specified sequence of one or more debugger commands to be executed immediately after creating a new dialogue within NightView. When a new dialogue is created, the list of all **on dialogue** regular expressions is checked to see if any of them match the name of the new dialogue. The most recently specified **on dialogue** command whose regular expression matches the dialogue name will have its commands executed.

In its first form (given only a regular expression), the **on dialogue** command will remove any commands that were associated with the given regular expression. If no regular expression is given, then *all* previously defined **on dialogue** commands are removed. If your safety level is set to *forbid*, you are not allowed to remove all **on dialogue** commands. If your safety level is set to *verify*, NightView requests verification before removing all **on dialogue** commands. See “set-safety” on page 6-67.

In its second and third forms, the **on dialogue** command will associate a sequence of one or more user-specified debugger commands with the given regular expression. Macro invocations are *not* expanded when reading the commands to associate with the regular expression.

If dialogue *local* is started up automatically by NightView, then it will exist *before* any commands in your *.NightViewrc* file are read. In this case, NightView automatically runs the **on dialogue** command after all the initialization files have been processed. See “apply on dialogue” on page 6-33. See “Initialization Files” on page 3-43.

The default qualifier for all commands associated with the given regular expression will be the newly created dialogue.

The commands specified by **on dialogue** are event-triggered commands: they have an implied safety level (which may be different from the safety level that was set using **set-safety**).

If you wish to list all **on dialogue** commands, or see which **on dialogue** commands would be executed for a particular dialogue name, you should use the **info on dialogue** command.

Example:

```
(local)on dialogue ben.*  nodebug /usr/bin/*
```

After issuing the above command, if we now create a new dialogue named *ben_hur*, then we will automatically set it up so that programs residing in the directory named */usr/bin* are not debugged by NightView.

Now suppose we do the following:

```
(local) on dialogue .*jerry do
>         nodebug /usr/remote/*
>         nodebug /usr/local/*
> end on dialogue
```

At this point, if we create another dialogue named `ben_n_jerry`, then this newly created dialogue will automatically be set up so that programs residing in the directories `/usr/remote` and `/usr/local` are not debugged by NightView. Note that even though the name `ben_n_jerry` also matches the regular expression `ben.*`, this dialogue *will* try to debug programs that reside in the directory `/usr/bin`. This is because **on dialogue** regular expressions are matched in reverse-chronological order (most recent first), and only the first match found is used.

```
(local) info on dialogue ben_n_jerry
on dialogue .*jerry do
    nodebug /usr/remote/*
    nodebug /usr/local/*
end on dialogue
```

If we were to now issue the command:

```
(local) on dialogue .*jerry
```

Then this would remove `.*jerry` (and its associated commands) from the debuggers **on dialogue** command list. Now, if we create yet another dialogue named `benny_and_jerry`, then this third dialogue will *not* automatically debug programs that reside in the directory `/usr/bin`, but it *will* debug programs that reside in `/usr/remote` and `/usr/local` (just like the first one did).

```
(local) info on dialogue benny_and_jerry
on dialogue ben.* do
    nodebug /usr/bin/*
end on dialogue
```

apply on dialogue

Execute **on dialogue** commands for existing dialogues.

apply on dialogue

The **apply on dialogue** command allows **on dialogue** commands to be executed for existing dialogues. See “on dialogue” on page 6-31. For each dialogue specified by the qualifier, the **on dialogue** commands which would match the name of the dialogue are immediately executed on behalf of the dialogue.

When the debugger automatically creates a local dialogue, it does an **on dialogue** command with a qualifier of `(local)` after processing all the initialization files. See “Initialization Files” on page 3-43. Because dialogue `local` exists *before* the customization commands in the user's `.NightViewrc` file are interpreted by the debugger, the **on dialogue** command by itself cannot initialize the environment for dialogue `local` (since it only applies to dialogues that will be created *after* the **apply on dialogue** command is issued). The automatic **on dialogue** executes any **on**

dialogue commands that refer to dialogue `local`.

Dialogue Input and Output

Because each dialogue is a separate shell, each dialogue has its own input and output streams. NightView has several options for sending input to dialogues and managing the output data generated by the dialogue shell and the programs being run within it.

!

Pass input to a dialogue.

! [*input line*]

input line

If *input line* is specified, it is passed to the dialogue (or dialogues) determined by the command qualifier.

If *input line* is not specified, then this command switches to a special dialogue input mode.

If the qualifier for this command specifies more than one dialogue, then the same input data is sent to all the dialogues. This can make sense if you are doing something like debugging two versions of the same program and you want to see where they diverge. It is up to you to insure that the input is sensible to all the dialogues (or that the command qualifier only refers to one dialogue).

When you use the **!** command without an *input line* argument to switch to dialogue input mode, everything you type goes to the specified dialogues. Nothing you type is treated as a debugger command until a special terminator string is recognized. The default terminator string is ```-.` (note that this is not the same as the ```~.` used by **rlogin(1)** or **cu(1)**). See “set-terminator” on page 6-67, for information on how to change the terminator string.

The **!** command without an *input line* argument cannot be used inside a macro (see “Defining and Using Macros” on page 6-184), nor can it be used in the graphical or full-screen user interfaces.

Macros are *not* expanded when reading the input (or arguments) to this command.

This command does not care if it is talking to the dialogue shell or to a program running in the shell. If you start a program that requests input, you can pass the input to it using this command.

See “Repeating Commands” on page 6-21.

Example:

```
(afamily) !pwd
(afamily) !
PATH=/extra/progs:$PATH
```



```

ulimit -m 200
ulimit -d 100
ulimit -s 100
- .
(afamily)

```

The first line just sends a **pwd** command to the dialogue. The second switches to dialogue input mode and then several lines of input are sent directly to the dialogue to set up environment variables and limits on the amount of memory subsequent processes will be allowed to use. The final "- ." switches back to normal command input mode.

Note that if you just want to send a program name to the shell and wait for that program to start, you may want to use the **run** command instead. See “run” on page 6-38.

set-show

Control where dialogue output goes.

```

set-show [silent | notify=mode | continuous=mode]
           [log[=filename]] [buffer=number]

```

silent

Just buffer the dialogue output, do not display it. The **show** command may be used to see what has accumulated (see “show” on page 6-36).

notify=mode

Do not display the dialogue output, but do print a notice when output first becomes available.

continuous=mode

Display dialogue output when it is generated.

The **notify** and **continuous** modes both accept one of the following keyword arguments:

immediate

In immediate mode the notification or actual output is displayed as soon as output becomes available.

atprompt

In the *atprompt* mode, the output is displayed only when the debugger is not requesting input. This is typically immediately prior to printing a new prompt to request additional commands, but it also prints output when the debugger is waiting for some event and has not yet prompted for new input.

Additional parameters on the **set-show** command control logging to a file and the size of the internal buffer.

`log[=filename]`

The `log` parameter without the `=filename` option turns off logging to a file and resumes buffering a limited amount of output in memory. When a file name is specified, the output from the dialogue is logged to that file until the `log` parameter is changed.

`buffer=number`

The `buffer` parameter is used to set the size of the buffer holding all the most recent output from the dialogue. The default size is 10240 (10K bytes). When the buffer fills up, the oldest output is discarded. When logging to a file, this parameter does not have any effect — a log file may grow until disk space is exhausted.

This command only logs the output from dialogues. It does not log debugger commands, nor does it directly log the input to a dialogue; however, the input will normally be echoed by the system, so it will be logged as output from the dialogue.

To log the entire debug session, see “set-log” on page 6-62.

Each dialogue starts off in the default mode:

```
(all) set-show buffer=10240 continuous=atprompt
```

show

Control dialogue output.

show [*number* | `all` | `none`] [| *shell-command*]

number

The number of old output lines you wish to see again.

`all`

Specifying `all` instead of a number means show all the buffered output from the dialogue shell.

`none`

The `none` keyword is used to tell the debugger you are not interested in any of the buffered output. It pretends you have already seen any data currently in the buffer.

| *shell-command*

You may use a vertical bar (shell pipe operator) to request the output be sent to an arbitrary shell command, rather than being displayed. You may use this to run the output through a pager or filter of some kind.

The debugger always internally buffers output generated by dialogues. The **show** command displays any buffered output from a dialogue which you have not yet seen. The *number* or `all` arguments tell the debugger to display that many lines of previous output in

addition to the new output (so the total number of lines displayed may be greater than *number*). The **set-show** command is used to control when dialogue output is printed without a specific request via the **show** command (see “set-show” on page 6-35).

Managing Processes

run

Run a program in a dialogue and wait for NightView to start debugging it.

run *input line*

input line

The shell command that will start a program (or programs) to debug.

This command is very similar to the **!** command (see “!” on page 6-34): it sends the specified *input line* to the dialogue shell (or shells) specified by the qualifier. The difference between **run** and **!** is that **run** waits for a new process to be debugged in one of the dialogues specified by the qualifier.

NOTE

Even if the qualifier specifies multiple dialogues, the **run** command terminates as soon as one new process has started.

The **run** command does not check the given *input line* for validity; it simply passes it unchanged to the dialogue shell, just like the **!** command. If it does not start a new process to be debugged, then **run** will just continue waiting forever (or until you type <CONTROL C>). If you issue a **run** command that starts more than one program, **run** will only wait until one of them starts up and is noticed by NightView. The other programs will start up and be debugged, but you probably won't know about them until after you have entered the next command.

If you just want to send input to a program that is reading from the shell's input terminal, or you want to start up a program or programs without waiting for them, just use the **!** command.

If you want to run the same program again, use the **run** command again or use the **rerun** command. See “Restarting a Program” on page 3-20. If you want multiple programs to run concurrently, end the shell commands with & (ampersand). (You can't do this if your program expects input from you.)

rerun

Run a program again.

rerun

The qualifier must evaluate to exactly one process or exactly one dialogue and no process. This command takes no arguments.

Whenever a process starts up, NightView remembers the most recent dialogue input line and associates it with the new process.

If the qualifier contains a process, NightView kills the process and sends the associated dialogue input line again.

If there is no process, NightView sends the dialogue input line associated with the process that terminated most recently in the specified dialogue.

The method of remembering recent dialogue input lines works for nearly all situations, but there may be situations of complex process start-up where NightView cannot send an appropriate dialogue input line and this command should not be used.

set-notify

Control how you are notified of events.

set-notify [*silent* | *continuous=mode*]

silent

Only report events when explicitly requested.

continuous=mode

Display events when they happen.

The *continuous* mode accepts one of the following keyword arguments:

immediate

In *immediate* mode the notification is displayed as soon as the event happens.

atprompt

In the *atprompt* mode, the notification is displayed only when the debugger is not requesting input. This is typically immediately prior to printing a new prompt to request additional commands, but it also prints notifications when the debugger is waiting for some event and has not yet prompted for new input.

This command controls how the debugger tells you what is happening to the processes you are debugging. Individual processes may be set to notify you in different ways (using the command qualifier).

Events that might cause notification include hitting a breakpoint or watchpoint, getting a signal (but see “handle” on page 6-144), or `'exec'`ing a new program. New processes to be debugged also cause notification, but this notification is controlled by the notification setting of the parent of the new process. Processes created directly by the dialogue shell always cause notification in the default notify mode. When a process exits, you will be notified by the process' dialogue (but see “show” on page 6-36 and “set-show” on page 6-35).

The output generated by any commands attached to a breakpoint (or watchpoint) or any automatic display expressions is also controlled by **set-notify**. If you set notify mode to *silent* for a process, all debugger output associated with that process will be buffered up and saved until you ask to see it.

Any change to the notify mode of a process takes place immediately, so changing the mode from `silent` to `continuous` may also result in large amounts of accumulated event notifications and other buffered output being generated.

The **notify** command (see “notify” on page 6-40) can be used to explicitly request notification of any events that have been saved up (this is the only way to find out about events that have happened in a process where the notify mode is `silent`).

If no arguments are given to the `notify` command, then the current notify mode of each process in the qualifier is printed.

The default notify mode is:

```
(all) set-notify continuous=atprompt
```

notify

Ask about pending event notifications.

notify

If you have been suppressing event notification on certain processes (see “set-notify” on page 6-39), the `notify` command may be used to request any notifications that have not yet been printed. It only tells you about pending events in the processes specified by the command qualifier.

attach

Attach the debugger to a process that is already running.

```
attach [{/resume | /stop}] { PID | name }
```

```
/resume
```

Resume the process when the attach is complete.

```
/stop
```

Keep the process stopped when the attach is complete.

```
/nodebug
```

Attach to the process, but treat it as nodebug. This means that the process will not be debugged, but NightView can gain control of any children it forks. (See “set-children” on page 6-52.)

PID

The process ID of the running process.

name

The name of a running process.

This command allows a program to be debugged even if it was not started from a debugger dialogue shell (see “Attaching” on page 3-3). The qualifier on this command must specify a single dialogue indicating which machine is running the specified PID. An error is reported if the qualifier implies multiple dialogues. It is also an error to attempt to attach to a program already being debugged, or to attach any of the processes required to run the debugger.

If the argument is a number, it is treated as a *PID*, and an attachment to the specified PID is attempted. Otherwise, the argument is treated as a *name*. If treated as a name, it is interpreted as a *regular expression* which is matched against the command name of all processes on the system owned by the dialogue’s user (or, if the dialogue’s user is **root**, all processes). If a *name* matches multiple processes, an attach to the first match is attempted and all others are ignored.

Since the program to which you are attaching is already running independently of the debugger, you will not be able to send it input through the normal dialogue input mechanism (see “!” on page 6-34) or see the output it generates (the input and output for the process remain connected to the same streams they were connected to prior to the **attach**).

Once you attach to a process, any future children it forks will also be debugged. See “set-children” on page 6-52. Children created prior to the attach must be explicitly attached if you want to debug them.

See “Attach Permissions” on page 3-49 for a description of what processes you are allowed to attach.

Once the **attach** is complete, the process will stay stopped or will be resumed depending on the setting from the **set-resume** command (see “set-resume” on page 6-75). You can override that setting by explicitly giving a */resume* or */stop* option.

If the **/nodebug** option is specified, NightView will have minimal control of the specified process, but it will not be debugged, and will automatically resume. However, it can gain control of and debug any forked processes. (See “set-children” on page 6-52.) A possible use of this feature is to **attach/nodebug** to a shell, not to debug the shell, but rather to debug any processes started by that shell.

detach

Stop debugging a list of processes.

detach

The **detach** command terminates the debugger’s connection to all the processes named in the command qualifier. Any breakpoints, monitorpoints, heappoints, watchpoints, or syscallpoints set in those processes are removed, but patchpoints and tracepoints remain if they are enabled when you execute the **detach** command. See “breakpoint” on page 6-108, “patchpoint” on page 6-110, “monitorpoint” on page 6-115, “heappoint” on page 6-117, “tracepoint” on page 6-113, “watchpoint” on page 6-127, and “syscallpoint” on page 6-129.

The processes are allowed to continue running normally and the debugger will not be notified of any subsequent events that occur in those processes. If any of the processes fork or exec new programs, the debugger will not see them.

When the safety level is `unsafe` (see “set-safety” on page 6-67), detaching a process that was stopped while evaluating a debugger expression containing a function call aborts any expression evaluation in progress. This returns the process to the state it was in when you asked to evaluate the expression. At `verify` safety level, it asks first, and at safety level `forbid`, it refuses to let you detach the process.

For another way of avoiding debugging certain processes, see “nodebug” on page 6-28. Also, see “set-children” on page 6-52.

kill

Terminate a list of processes.

kill

The **kill** command terminates all the processes named in the command qualifier.

In the graphical user interface, if you use a ‘Kill’ button (as opposed to manually typing the **kill** command) the debugger will check your safety level (see “set-safety” on page 6-67) before permitting you to kill the desired processes. If your safety level is `forbid` then you will *not* be permitted to kill the selected processes. If your safety level is `verify` then you will be prompted for verification. If your safety level is `unsafe` then the processes are terminated with no questions asked.

symbol-file

Establish the file containing symbolic information for a program.

symbol-file *program-name*

program-name

This must be the name of an executable file corresponding to the programs running in the specified processes. It should contain symbolic debug information for the program.

If *program-name* is a relative pathname, it is interpreted relative to NightView's current working directory.

If the program is on a remote system other than the specified target system, use the form *user@host:/path*. NightView will download this file from the remote system to read the debug information. See “Remote File Access” on page 3-7.

program-name is *not* subject to object filename translations. See “translate-object-file” on page 6-29.

A *symbol file* is an executable file from which NightView obtains information about symbols in a program being debugged. Normally, the symbol file is the same as the program's executable file, but it may be different if, for example, you are debugging a stripped program (see **strip(1)**). In this case, you need to specify an unstripped version of the program in the **symbol-file** command, if you want to access information symbolically.

The **symbol-file** command is applied to each process in the qualifier. You should make sure that each of those processes is running the same program; otherwise, you may get unpredictable results from the debugger when you examine variables or memory.

Note: If you have not specified a symbol file for a process, NightView attempts to obtain the information from the executable file (see “exec-file” on page 6-46).

In some situations, such as when debug information is needed from shared libraries, an object filename translation is more appropriate than a **symbol-file** command. See “translate-object-file” on page 6-29.

core-file

Create a pseudo-process for debugging an aborted program's core image file.

```
core-file corefile-name [exec-file=program-name] [with-translations]
                        [interpreter-base=address]
```

corefile-name

The name of a core file. When used in a remote dialogue, this file must reside on the target system.

If *corefile-name* is a relative pathname, it is interpreted relative to NightView's current working directory.

exec-file=program-name

Specifies the name of the executable program that created the given core file. When used in a remote dialogue, this file must reside on the target system.

If *program-name* is a relative pathname, it is interpreted relative to NightView's current working directory.

with-translations

Indicates that the lines following the **core-file** command are library translations.

interpreter-base=address

Specify the address of the dynamic loader. This may be useful for a core file generated on a system different from the target system.

A *core file* is a copy of a process's memory made when a process is terminated abnormally. You can examine these core files using NightView by specifying the core file name in the **core-file** command. NightView responds with a process ID (PID) corresponding to a newly-created *pseudo-process*. This is not a real executing process; a pseudo-process is merely a mechanism for dealing with core files in NightView. The PID NightView assigns does not correspond to any running process, but you can use it in qualifiers, and you can also include it in process families using the **family** command. See “family” on page 6-51.

The qualifier for the **core-file** command is used only to determine with which dialogue the pseudo-process should be associated. (Among other things, this determines the type of machine that created the core file. The core file must have been created on the system the dialogue is running on.) Thus, the qualifier should specify exactly one dialogue; otherwise, NightView issues an error message and refuses to honor the command.

If you specify the `exec-file=program-name` option, it is equivalent to executing an **exec-file** command (see “exec-file” on page 6-46) on the pseudo-process created by the **core-file** command. This is seldom required, since NightView attempts to determine the location of the executable program from information saved in the core file (see “Finding Your Program” on page 3-9). If NightView is unable to correctly determine the executable program, you will need to specify the `exec-file=program-name` option or use the **exec-file** command to specify the name of the executable program.

When debugging a core file, NightView uses the executable program file for two purposes. NightView uses this file to obtain symbolic information about variables and procedures in your program, just as it does when debugging normal processes. For core files, NightView also must use this file to obtain the contents of read-only memory, including the machine instructions of the program. If NightView is unable to locate the executable program, then you will only be able to examine writable memory by absolute address. You can specify the file, or files, NightView should use by specifying the `exec-file=program-name` option or by using the **exec-file** and **symbol-file** commands (see “exec-file” on page 6-46 and “symbol-file” on page 6-42).

If you specify `with-translations`, then the lines following the **core-file** command are library translations of the form:

from-string to-string

End the translations with a line that contains only:

end translations

This allows debugging core files from dynamically-linked programs on systems where the installed libraries do not match the libraries that were being used when the core file was generated. This is not necessary for most users. The translations are similar to the object file translations in the **translate-object-file** command, but they refer to dynamic libraries and are applied only to this process (see “translate-object-file” on page 6-29). For remote debugging, the translations are applied on the target system, not the host system.

Note that, unlike other debuggers, NightView allows you to examine the core file of a process at the same time you are executing the program that produced the core file. This allows you to try executing your program again to try to find the problem, while still accessing information from the core file. For instance, you may find from the core file that a certain global variable has an incorrect value. You could then run the program again, stopping it at interesting points to check the value of that global variable. By using an appropriate qualifier, you can easily print out the values of variables in both the running program and the core file for easy comparison.

save-core-file

Packages up all files required for subsequent core file analysis into a compressed file.

When transporting a core file to another system for analysis, it is important to take all the shared libraries related to its execution. Locating them can be a cumbersome task that NightView can do for you more easily.

save-core-file [/nozip] [/nodebuginfo] [/replace] [/keep]
[include=*file*] [note=*string*] *savename*

/nozip

Do not compress the tar archive used to save the core file (the default is to generate a compressed file).

/nodebuginfo

Do not include the debuginfo files (if there are any). The default is to include them. Debuginfo files are special files that are optionally installed on systems which provide debug information for standard libraries. NightView knows how to locate these files, and by default, automatically consults them when they are present.

/replace

If the output file already exists, do not terminate the save operation, but instead to try to overwrite it.

/keep

If there are errors during the writing of the output file, keep the possibly broken file rather than removing it.

include=*file*

Add *file* to the directory of information being packaged with the core file. The specified file can be of any type -- perhaps some data files used as input to your program or maybe just some notes you make so you can remember the circumstances relating to the core file. You can use multiple include options to include multiple files.

note=*string*

Add the *string* as a note in the generated script that will be used to debug the core file subsequently. This note will be echoed at the end of the script.

savename

The name of the file where the (optionally) compressed tar file will be written. The actual file name will be *savename.tar.gz* or *savename.tar*.

This command will take a single core file process and will generate the *savename* tar file. This can be useful if there is limited access time available on the current system or if you wish to send the core file off to another person.

The tar file can then be transferred to another system for subsequent analysis. The system must be of the same architecture as the original system (i.e. 32-bit or 64-bit).

If you are not already debugging the core file of interest, you can use the command **nview-save-core-file** to save the core file. See “nview-save-core-file” on page 5-3 for information on that script.

exec-file

Specify the location of the executable file corresponding to a process.

exec-file *program-name*

program-name

Specifies the file containing the executable program corresponding to the specified processes.

If *program-name* is a relative pathname, it is interpreted relative to NightView's current working directory.

If *program-name* is on a remote system other than the specified target system, use the form *user@host:/path*. NightView will download this file from the specified system to read the debug information. In this case, **exec-file** is treated as though you had used **symbol-file**.

program-name is subject to object filename translations. See “translate-object-file” on page 6-29.

This command tells NightView where to find the executable file corresponding to the processes specified by the qualifier. Obviously, you should ensure that all those processes are, in fact, running the same program; otherwise, you may get strange behavior. (NOTE: NightView does not do this verification for you because the processes may be executing different copies of the same program on several different systems. NightView would not be able to tell that these were the same program.)

You usually use this command in conjunction with the **core-file** command (see “core-file” on page 6-43). You may also need to use it if NightView is unable to determine the executable file corresponding to a new process being debugged. See “Finding Your Program” on page 3-9.

If you do not explicitly specify a symbol file for a process (see “symbol-file” on page 6-42), NightView uses the executable file. Since the symbolic information is usually contained in the executable file anyway, this is most often what you want. You can specify the executable file and symbol file in any order for a given process.

When a new executable file is specified, any **on program** commands that match the new file name are executed. See “on program” on page 6-47.

Examples:

```
(local) core-file ./mycore
New process: local:65536
/users/bob/mycore
was last modified on Wed Nov 18 17:48:38 1992
Core file indicates the executable file is
/users/bob/myprog
Executable file set to
```

```

/users/bob/myprog
Pseudo-process assigned PID 65536
Process 65536 terminated with SIGQUIT
(local) family mycore 65536
(local) (mycore) exec-file ./stripped_prog
(local) (mycore) symbol-file ./full_prog

```

The first command creates a new pseudo-process for the file **mycore** in NightView's current directory. NightView assigns this pseudo-process PID number 65536. The **family** command then gives the name **mycore** to this pseudo-process. The **exec-file** command then establishes the file **stripped_prog** as the executable file for that process, while the **symbol-file** command establishes **full_prog** as the name of the symbol file.

When dealing with shared libraries, an object filename translation is more appropriate than a **exec-file** command. See “translate-object-file” on page 6-29.

on program

Specify debugger commands to be executed when a program is 'exec'ed.

on program [*pattern*]

on program *pattern command*

on program *pattern do*

pattern

A wildcard pattern to match against the executable file names of newly 'exec'ed programs. See “Wildcard Patterns” on page 6-20.

command

A debugger command to be executed when a new program whose executable file name matches *pattern* is 'exec'ed.

In the third form of the **on program** command, the debugger commands to be executed must begin on the line following the **do** keyword. The list of debugger commands to execute is terminated when a line containing only the words **end on program** is encountered.

The **on program** command allows a user-specified sequence of one or more debugger commands to be executed immediately after 'exec'ing a program that is being debugged by NightView. When a debugged process performs an 'exec' (or the **exec-file** command is used to change the location of the executable file name), the list of **on program** patterns for that process's controlling dialogue is checked to see if any of the patterns match the executable file name of the program that was just 'exec'ed. The most recently specified **on program** command whose pattern matches the executable file name of the newly 'exec'ed program will have its commands executed.

on program processing is related to **on restart** processing. When a program execs (or the **exec-file** command is used), NightView first checks the **on restart** patterns. See “on restart” on page 6-49. If a match is found, then the commands

associated with the matching pattern are executed. In this case, no **on program** patterns are checked. However, **on restart** commands created by a checkpoint always begin with a call to the macro `restart_begin_hook`. The initial definition of this macro invokes the **apply on program** command. So, by default, **on program** patterns are checked and matching commands are run *before* the **on restart** commands are run. See “Restarting a Program” on page 3-20.

If no match is found in the **on restart** patterns, then NightView checks the **on program** patterns.

In its first form (given only a pattern), the **on program** command will remove any commands that were associated with the given pattern for each dialogue specified in the qualifier. If no pattern is given, then *all* previously defined **on program** commands are removed from each dialogue specified in the qualifier. If your safety level is set to `forbid`, you are not allowed to remove all **on program** commands. If your safety level is set to `verify`, NightView requests verification before removing all **on program** commands. See “set-safety” on page 6-67.

In its second and third forms, the **on program** command will associate a sequence of one or more user-specified debugger commands with the given pattern for each dialogue specified by the qualifier. Macro invocations are *not* expanded when reading the commands to associate with the pattern.

The default qualifier for all commands associated with the given pattern will be the process performing the ‘exec’.

The commands specified by **on program** are event-triggered commands: they have an implied safety level (which may be different from the safety level that was set using **set-safety**), and may be terminated automatically if they resume execution of the ‘exec’ing process. See “Command Streams” on page 3-39.

If you wish to list all **on program** commands, or see which **on program** commands would be executed for a particular program name, you should use the **info on program** command.

Example:

```
(local)on program ren* break main.c:24
```

After issuing the above command, if we now run a program in dialogue `local` named `ren_n_stimpy`, then we will automatically set a breakpoint in it at line 24 of the file `main.c`.

Now suppose we do the following:

```
(local)on program *stimpy do
> handle 5 noprint nostop
> handle 6 noprint nopass
> end on program
```

At this point, if we run `ren_n_stimpy` again, then this newly ‘exec’ed program will handle signals 5 and 6 in the specified manner. Note that even though the name `ren_n_stimpy` also matches the pattern `ren*` that a breakpoint will *not* automatically be set at line 24 of `main.c` in this new invocation of `ren_n_stimpy`. This is because **on program** patterns are matched in reverse-chronological order (most recent first), and only the first match found is used.

```
(local) info on program ren_n_stimpy
on program *stimpy do
    handle 5 noprint nostop
    handle 6 noprint nopass
end on program
```

If we were to now issue the command:

```
(local) on program *stimpy
```

Then this would remove `*stimpy` (and its associated commands) from the `on program` list for dialogue `local`. Now, if we run `ren_n_stimpy` a third time, then this third invocation will automatically have a breakpoint set at line 24 of `main.c` (just like the first one did).

```
(local) info on program ren_n_stimpy
on program ren* do
break main.c:24
end on program
```

apply on program

Execute `on program` commands for existing processes.

apply on program

The `apply on program` command allows `on program` commands to be executed for existing processes. (See “on program” on page 6-47). For each process specified by the qualifier, the `on program` commands which would match the executable file name of the process are immediately executed on behalf of the process.

Example:

Suppose I want to set a breakpoint at the subroutine named `main` in all programs both new and old that are debugged in dialogue `local`. Using the `on program` and `apply on program` commands, this could be accomplished as follows:

```
(local) on program *      b main
(local) apply on program
```

on restart

Specify debugger commands to be executed when a program is restarted.

```
on restart [pattern]
```

```
on restart pattern command
```

```
on restart pattern do
```

pattern

A wildcard pattern to match against the executable file names of newly executed programs. See “Wildcard Patterns” on page 6-20.

command

A debugger command to be executed when a new program whose executable file name matches *pattern* is executed.

In the third form of the **on restart** command, the debugger commands to be executed must begin on the line following the **do** keyword. The list of debugger commands to execute is terminated when a line containing only the words **end on restart** is encountered.

The **on restart** command is primarily intended to be used internally by the debugger as part of the restart processing. See “Restarting a Program” on page 3-20. You may use **on restart** explicitly, if desired, but you should be wary of conflicts with the debugger's use. The debugger creates **on restart** commands as a result of a checkpoint.

on restart is virtually identical to **on program** in form and function. See “on program” on page 6-47 for a description of the parameters and functionality of these commands. That section also describes the interaction of these two commands.

If you wish to list all **on restart** commands, or see which **on restart** commands would be executed for a particular program name, use the **info on restart** command. See “info on restart” on page 6-175.

checkpoint

Take a restart checkpoint now.

checkpoint

The **checkpoint** command saves restart information for the program running in each process in the qualifier.

In most cases, you do not need to use the **checkpoint** command, because checkpoints are taken automatically at certain times. See “Restarting a Program” on page 3-20. **checkpoint** gives you a way to explicitly take a checkpoint at a time you choose. Note that any later checkpoints (either explicit or automatic) will replace the restart information.

Example:

In this example, you are debugging a complex program. You know some good places to set breakpoints, and you know that you need some more to find the bug, but are not sure yet where they should be. You set your known breakpoints, take a checkpoint, and save the restart information to a file. Then you experiment with some different breakpoints.

```
(local) # set known good breakpoints
(local) breakpoint fred.c:123
set other known breakpoints ...
```



```
(local) checkpoint
(local) info on restart output=restart_info

(local) # now try experimental breakpoints
(local) breakpoint pebbles.c:456
set other experimental breakpoints ...
```

You decide to start the program again and want only the known breakpoints. You kill your process, which takes a checkpoint, including the experimental breakpoints. Then you **source** the file containing the restart information. The restart information is replaced with only the known breakpoints. When you restart your program, only the known breakpoints are restored.

```
(local) kill
(local) source restart_info
restart program
```

family

Give a name to a family of one or more processes.

family *family-name* *[[-] qualifier-spec]* ...

family-name

The family name to be defined. This must not be the same as the name of any dialogue you currently have. The family-name must consist only of alphanumeric characters and underscores and must begin with an alphabetic character. The family-name may be of arbitrary length.

qualifier-spec

Identifies one or more processes to be included or excluded in the family named by *family-name*. See “Qualifier Specifiers” on page 6-17.

The total set of processes is accumulated by scanning the *qualifier-spec* arguments left to right. An argument is added to the set unless it is preceded by a '-', in which case it is subtracted from the set accumulated so far.

If no *qualifier-spec* is included, then this command removes any previous definition of the *family-name*. If your safety level is set to **forbid**, you are not allowed to remove the definition of a *family-name* that is present in the default qualifier. If your safety level is set to **verify**, NightView requests verification before removing such a definition. See “set-safety” on page 6-67.

If one or more *qualifier-spec* arguments are supplied, they are immediately evaluated (see “Qualifier Specifiers” on page 6-17) and the *family-name* is defined as the list of processes indicated by those arguments. Evaluation of the arguments has the following implications:

- Any *family-name* appearing in the argument list must be defined. Subsequent changes made to the definition of that *family-name* will have no

effect on the processes implied by the *family-name* being defined in the **family** command.

- The processes denoted by any *dialogue-name* appearing in the argument list are just those that exist at the time the **family** command is executed.
- The argument **all** denotes only those processes that exist at the time the **family** command is executed.
- The argument **auto** denotes the process that has been stopped the longest at the time the **family** command is executed.

Any qualifier applied to this command has no effect.

Note that you may use a *family-name* in a qualifier before it is actually defined, but you must define the *family-name* before executing any command that needs to know what the *family-name* refers to.

Examples:

```
(local) family fam1 12 25 18
(local) family fam2 fam1 99
(local) family fam1 fam1 16
```

The first command gives the name *fam1* to the processes identified by PIDs 12, 18, and 25. The second command gives the name *fam2* to the three processes in *fam1* plus process 99. The third command extends the definition of *fam1* to include process 16; thus *fam1* is a synonym for four processes: 12, 16, 18, and 25. Note that extending *fam1* has no effect on *fam2*, which still consists of processes 12, 18, 25, and 99.

Using the families defined in the previous examples, the use of a minus sign on arguments can be illustrated by the following examples:

```
(local) family fam3 fam1 fam2 -12
(local) family fam3 fam1 -12 fam2
```

The first command defines *fam3* to be the processes 16, 18, 25, and 99. In contrast, the second command defines *fam3* to be the processes 12, 16, 18, 25, and 99. In this case, the argument **-12** removed process 12 from the set accumulated from *fam1*, but the **fam2** argument adds that process back in. In general, it is a good idea to put all the subtracted arguments at the end of the list.

set-children

Control whether children should be debugged.

```
set-children { all [ resume ] | exec | none }
```

```
all
```

Debug all children. If the optional keyword **resume** is specified, then a child process is resumed automatically after NightView has prepared it for debugging. This is useful if your program creates many child processes that you want to debug, but all you need to do is inherit the eventpoints and debug settings from the parent process. See “Multiple Processes” on page 3-2.

`exec`

Debug children only when they have called **exec(3)** (that is, when they are running a different program). The program name is checked against the debug/nodebug list for the controlling dialogue to see if the program should be debugged. See “debug” on page 6-27. This is the default setting for direct children of the dialogue shell and processes debugged with the **attach** command. See “attach” on page 6-40.

`none`

Ignore all children.

Sometimes you are not interested in the child processes of the process you are debugging. For example, your program may make many calls to **system(3)** which you are not interested in debugging. The **set-children** command gives you a way of controlling which children will be debugged without having to detach from each one individually. See “detach” on page 6-41.

The **set-children** command applies to future children of the processes specified by the qualifier. Existing children are not affected.

This mode is inherited by future children.

set-exit

Control whether a process stops before exiting.

set-exit [stop | nostop]

`stop`

The process will stop if the **exit** system service is called.

`nostop`

The process will *not* stop before exiting.

The **set-exit** command controls whether the processes specified by the qualifier will stop before exiting. The default state for a process is to stop before exiting. See “Exited and Terminated Processes” on page 3-22.

If no arguments are specified to the command, the command prints the current state for each process in the qualifier. If an argument is specified, the command changes the state of each process in the qualifier accordingly and then prints the new state.

Note that the initial **set-exit** mode for each process comes from the global **set-resume** mode. See “set-resume” on page 6-75. Note also that the mode persists for the entire life of the process, even across an **exec** system call, until modified by another **set-exit** command. In the case of an **exec**, an **on program** or **on restart** command might specify a **set-exit** command that changes the mode. See “on program” on page 6-47 and “on restart” on page 6-49. See also “Restarting a Program” on page 3-20.

If you also want a process to automatically resume execution after an `exec`, use the **set-resume** command, or put a **resume** command in an **on program** specification.

See “set-resume” on page 6-75, “resume” on page 6-133 and “on program” on page 6-47.

set-shared-lib-update

Control whether NightView re-reads debug information each time a new shared object is loaded.

set-shared-lib-update [on | off]

If this setting is `on`, then each time a shared object is loaded (see **dlopen(3)**), debug information is re-read, equivalent to an **exec-file** command, to discover new information in the shared object. The default is `off`, because re-reading the debug information can be time-consuming in large applications, especially if they perform multiple `dlopen` calls.

This functionality can be done manually at any time with **exec-file** (see “exec-file” on page 6-46).

wait

Wait for processes to stop.

wait [{all | any} [new]]

all

Wait for all processes in the qualifier to stop.

any

Wait for any process in the qualifier to stop.

new

Implicitly add any new processes that show up to the qualifier.

The **wait** command waits for processes in the qualifier to stop. See “Process States” on page 23. That is, no more commands are read from this command stream until the specified processes stop. See “Command Streams” on page 39. See “Interrupting the Debugger” on page 40.

If no arguments are specified, the default behavior is **wait** any new.

mreserve

Reserve a region of memory in a process.

mreserve start=address {length=bytes | end=address}

`start=address`

Specify the start address of the region.

`length=bytes`

Specify the length of the region in bytes.

`end=address`

Specify the end address of the region.

The `start=address` parameter is required. You must specify either a `length` or an `end` address.

The **mreserve** command reserves a region of memory for each process specified by the qualifier. This means that NightView will not allocate space for patch areas in that region. See Appendix C [Implementation Overview] on page C-1.

This command does not directly affect the process. It is only an indication to NightView to avoid placing patch areas in the specified region, presumably because your program will be using that region later in its execution.

mreserve only affects *future* allocations. You should reserve memory before using any commands that allocate space in the process, including eventpoint commands, the **load** command, or any command with an expression that involves a function call. See “Eventpoints” on page 3-9. See “load” on page 6-104. See “Expression Evaluation” on page 3-25.

You should exercise some caution with this command. It is possible to reserve memory in such a way that NightView cannot function.

For convenience, you are allowed to specify reservations that overlap or contain existing regions in your process.

Memory reservations are printed as part of the **info memory** command. See “info memory” on page 6-170.

Memory reservations are remembered as part of the restart information. See “Restart Information” on page 3-21. During restart, memory reservations are applied before any commands that would allocate space in the process.

You cannot reserve a region of CUDA memory. This is a technical limitation of the CUDA driver. All addresses are treated as host memory addresses by the **mreserve** command.

Heap Debugging

heapdebug

Specify parameters for heap debugging.

```
heapdebug [check_free_fill={0|1}]
           [common_errors={block_overrun |
                           dangling_pointer |
                           uninitialized_field}]
           [do_free_fill={0|1}]
           [do_malloc_fill={0|1}]
           [error-name [{noprint | nostop | print | stop} ...]]
           [free_fill_byte={n | trash}]
           [frequency=n[{k|m}]]
           [heap_size={n[{k|m}] | unlimited}]
           [internal_checks={0|1}]
           [level={0|1|2|3}]
           [malloc_fill_byte={n | trash}]
           [off]
           [on]
           [post_fence_size=n]
           [post_fill_byte={n | trash}]
           [pre_fence_size=n]
           [pre_fill_byte={n | trash}]
           [protected={0|1}]
           [retain_free_blocks={n[{k|m}] | unlimited}]
           [slop=n]
           [walkback=n]
```

Abbreviation: **hd**

```
check_free_fill={0|1}
```

During heap checks, check that the free fill has not been disturbed in retained free blocks. Setting this to 0 (turning it off) improves performance, but does not detect as many errors. The default value is 1 (check free fill).

```
common_errors={block_overrun | dangling_pointer |
               uninitialized_field}
```

This is a convenient way to set parameters to detect common program errors.

```
block_overrun
```

detect program writing past the end of a block

```
dangling_pointer
```

detect program referencing a freed block

```
uninitialized_field
```

detect program failing to initialize fields in a block

```
do_free_fill={0|1}
```

When a block is freed, fill it with the `free_fill_byte`. Free fill applies to free blocks that are retained and also to free blocks that are immediately available for reuse. Setting this to 0 (turning it off) will disable free filling. The default value is 1 (fill free blocks).

```
do_malloc_fill={0|1}
```

When a block is allocated, fill it with `malloc_fill_byte`. Setting this to 0 (turning it off) disables malloc filling. The default value is 1 (fill allocated blocks).

```
error-name [{noprnt | nostop | print | stop} ...]
```

Specify how the debugger responds when an error condition is detected.

```
stop
```

stop the process when the error occurs; implies `print`

```
nostop
```

let the process continue when the error occurs

```
print
```

print a message when the error occurs

```
noprnt
```

do not print a message when the error occurs; implies `nostop`

error-name can be any of the following:

```
free_fill_modified
free_not_at_beginning
free_unallocated
internal_error
malloc_zero
memalign_not_power_2
out_of_memory
post_fence_modified
pre_fence_modified
realloc_not_at_beginning
realloc_unallocated
```

The default for all the errors is `stop print`, except for `malloc_zero`, `memalign_not_power_2`, and `out_of_memory`, which are not normally considered to be heap errors; the defaults for those errors is `nostop noprnt`.

Blocks are checked for errors during a heap check (see “Heap Check” on page 3-37) and when they are freed or `realloc`'ed. Other errors are detected during heap operations.

`free_fill_modified`

The free fill pattern in a retained free block has been modified.

`post_fence_modified`

The post-fence fill pattern in a block has been modified.

`pre_fence_modified`

The pre-fence fill pattern in a block has been modified.

`free_not_at_beginning`

`free` was called with an address that is within an allocated block, not at the beginning of a block.

`free_unallocated`

`free` was called with an address that does not correspond to any currently allocated block.

`internal_error`

An inconsistency was found in the internal data structures. There is a bug in the heap debugger, or the process has modified the heap debugger's internal data structures.

`malloc_zero`

The program asked for a block of size 0 bytes. This is not normally considered to be an error. The default disposition for `malloc_zero` is `nostop, noprint`.

`memalign_not_power_2`

The program called `memalign` with an alignment that is not a power of 2. This is not normally considered to be an error. The default disposition for `memalign_not_power_2` is `nostop, noprint`.

`out_of_memory`

The process ran out of memory, either because the system could not satisfy the request or because of the setting of the `heapsize` parameter. This is not normally considered to be an error. The default disposition for `out_of_memory` is `nostop, noprint`.

`realloc_not_at_beginning`

`realloc` was called with an address that is within an allocated block, not at the beginning of a block.

`realloc_unallocated`

`realloc` was called with an address that does not correspond to any currently allocated block.


```
free_fill_byte={n | trash}
```

The value to put in each byte of each block when it is freed if `do_free_fill` is 1. The default is `trash`, which, for `free_fill_byte`, is `0xc3`.

```
frequency=n[{k|m}]
```

The heap is checked every *n* heap operations (mallocs, frees, etc.). You may append *k* to multiply *n* by 1024 or *m* to multiply by 1048576. If *n* is zero, the heap is checked only by a **heappoint** (see “heappoint” on page 6-117) or a **heapcheck** command (see “heapcheck” on page 6-178). The default value is 10000.

```
heap_size={n[{k|m}] | unlimited}
```

The program is not allowed to allocate more than *n* total bytes. The default value is `unlimited`. You may append *k* to multiply *n* by 1024 or *m* to multiply by 1048576.

```
internal_checks={0|1}
```

If set to 1 (turned on), then during a heap check, check internal data structures for integrity. This adds a large overhead to each heap check. The default value is 0 (do not check internal data structures).

```
level={0|1|2|3}
```

This is a convenient way to set many of the other parameters.

0

disable checking

1

minimal checking

2

a medium level of checking

3

extreme checking

See “Levels and Common Errors” on page 3-34 for a discussion of heap debugging levels.

```
malloc_fill_byte={n | trash}
```

The value to put in each byte of a block when it is allocated, if `do_malloc_fill` is 1. The default is `trash`, which, for `malloc_fill_byte`, is `0xc5`.

off

Turn heap debugging off.

If heap debugging is off when the process makes its first allocation, the heap debugger adds little or no overhead. If heap debugging is turned off after the first allocation, the heap debugger still adds overhead, but it no longer checks for errors.

on

Turn heap debugging on.

Heap debugging may be turned on before the program makes its first allocation. After the program makes its first allocation, heap debugging may be turned on only if it was on when the program made its first allocation.

post_fence_size=*n*

Add *n* bytes after the end of a block when it is allocated, fill them with `post_fill_byte`, and check them during a heap check. The default is zero (no fence).

In hardware overrun protection mode, there may be a gap between the end of the block and the protected page, due to alignment requirements and the size of the block. At most *n* bytes of the gap are filled and checked. See “Hardware Overrun Protection” on page 3-35.

post_fill_byte={*n* | trash}

The value to put in each post-fence byte of a block when it is allocated. The default is `trash`, which, for `post_fill_byte`, is `0xaf`.

pre_fence_size=*n*

Add *n* bytes before the beginning of a block when it is allocated, fill them with `pre_fill_byte`, and check them during a heap check. The default is zero (no fence).

pre_fill_byte={*n* | trash}

The value to put in each pre-fence byte of a block when it is allocated. The default is `trash`, which, for `pre_fill_byte`, is `0xbf`.

protected={0|1}

If set to 1, turn on hardware overrun protection. Each block is allocated such that the end of the block is as near as possible to the end of a page. The following page is protected from reads and writes. See “Hardware Overrun Protection” on page 3-35.

The default value is 0 (no hardware overrun protection).

```
retain_free_blocks={n [{k|m}] | unlimited}
```

The number of recently-freed blocks to retain. You may append *k* to multiply *n* by 1024 or *m* to multiply by 1048576. Retained free blocks are not immediately available for reuse. See “Retained Free Blocks” on page 3-37.

The default value is 0 (no retained blocks).

```
slop=n
```

Add *n* bytes to the size of each allocation. For example, if *n* is 4 and the program calls `malloc(8)`, the allocation proceeds as though the program had called `malloc(12)`. The default value is 0 (no slop).

```
walkback=n
```

The maximum number of walkback entries to keep for each heap operation (`malloc`, `free`, etc.). More walkback entries may help you identify which routines are causing heap problems. The default value is 8 entries. This count refers to physical walkback entries. The number of walkback frames may differ from this number when displayed in NightView. The number of frames displayed may include extra inline frames, as they are not physical frames. The number of frames displayed may be fewer if certain frames are deemed uninteresting (see “interest” on page 6-70). See “Debugging the Heap” on page 3-33.

The **heapdebug** command configures the heap debugger in each of the processes in the qualifier. See “Debugging the Heap” on page 3-33. Another way to configure the heap in the graphical user interface is with the **Debug Heap...** item in the **Process** menu (see “Process Menu” on page 8-9).

All arguments may be abbreviated to the shortest unambiguous prefix.

The heap debugger remembers its settings when turned off. This way, it can be turned back on at a later time and will retain all of its former settings.

NOTE

Heap debugging is not supported on the CUDA architecture.

Setting Modes

-
- **set-cuda-software-preemption** - Allows a device to be used by other applications when an application running on the device is stopped for debugging.
- **set-echo** - Controls whether or not each subsequent user command is echoed back.

set-echo

Controls whether or not each subsequent user command is echoed back.

set-echo [{off | on}]

{off | on }

Turn echoing of commands off or on as specified. If no argument is specified, the state of echoing is toggled.

By default, echoing is `off`.

set-log

Log session to file.

set-log *keyword filename*

keyword

The *keyword* parameter must be one of the following:

`all`

Log entire session (commands as well as the output generated by commands).

`commands`

Log just commands typed.

`close`

Close a log file.

filename

Name of the log file.

This command starts logging the debugger session to a file. If the file already exists, the log information is appended to it. You may log just the commands (by using the `commands` keyword) or the entire session (`all` keyword) to a file (if the named file is

already an open log file, specifying a different keyword simply changes the mode of the log). You may open multiple log files (although more than one of each type of log would be rather redundant).

The `close` keyword is used to close the log associated with the file. (See “info log” on page 6-158).

The qualifier does not have any effect on this command. Any logs are global to the debug session.

Note that this command logs everything that happens during the debug session (essentially, everything you see on your terminal). The **set-show** command may be used to log output from a single dialogue (see “set-show” on page 6-35).

set-language

Establish a default language context for variables and expressions.

set-language {auto | c | c++ | fortran}

auto

Indicates that the default language should be determined automatically.

c

Indicates that the default language should be C.

c++

Indicates that the default language should be C++.

fortran

Indicates that the default language should be Fortran.

The arguments to this command can be in any mixture of upper and lower case.

For each process specified by the qualifier, **set-language** sets the default language used to interpret expressions and variables in commands. If a default language has not been established, or if the default has been set to `auto`, NightView decides the language in one of two ways. If the object file contains DWARF, then it contains the language information. Otherwise, NightView infers the language from the extension (the last few characters) of the source file name associated with the frame selected when the expression or variable is mentioned. The following extensions are recognized:

.c

The language is assumed to be C.

.C

The language is assumed to be C++.

.f

The language is assumed to be Fortran.

.s

Although this indicates an assembler source file, NightView uses the C language for such files. C expressions include nearly all the operators allowed by the assembler, plus much more.

The language determines the meaning of operators and constants in expressions; determines the syntax of some kinds of expressions (e.g., C type casts); controls the visibility of variable names; and controls the significance of case (upper versus lower) in variable names. The language also controls the formatting of output from the **print** command (see “print” on page 6-91), especially the way the type of an expression is indicated.

set-qualifier

Specify the default list of processes or dialogues that will be affected by subsequent commands which accept qualifiers.

set-qualifier [*qualifier-spec* ...]

qualifier-spec

Specifies a process or dialogue to be included in the default qualifier list (see “Qualifier Specifiers” on page 6-17). Any family names in the *qualifier-spec* are evaluated at the time of each command, not at the time of **set-qualifier**.

If no argument is specified, the default qualifier is set to null, meaning that a qualifier must be supplied to subsequent commands that require qualification.

set-history

Specify the number of items to be kept in the value history list.

set-history *count*

count

The number of items to be kept in the value history.

The qualifier is ignored on this command. The default history list size is 1000. If more history items than that are created, the oldest ones are discarded. No matter how many items are in the list, each new history item gets the next highest number.

set-limits

Specify limits on the number of array elements, string characters, or program addresses printed when examining program data.

```
set-limits {array=number | string=number | addresses=number |
             source=number} ...
```

```
array=number
```

The `array` keyword parameter specifies the maximum number of array elements to be printed. If you want unlimited output, specify zero as the limit.

```
string=number
```

The `string` keyword parameter specifies the maximum number of characters of a string to be printed. If you want unlimited output, specify zero as the limit.

```
addresses=number
```

The `addresses` keyword parameter specifies the maximum number of addresses to be printed for a particular location (See “Location Specifiers” on page 6-15). If you want unlimited output, specify zero as the limit.

```
source=number
```

The `source` keyword parameter specifies the maximum number of bytes a source file can have to be displayed in a source panel. This is useful for extremely large source files which overwhelm NightView due to the overhead involved in building individual widgets associated with each line. The number is in units of 1000 bytes. The default value is 4000, which indicates ~4MB. The source for files that exceed the limit are not displayed, but the assembly associated with the function associated with the current stack frame within the file is displayed.

The `array`, `string`, `addresses`, and `source` keywords may be specified in any order.

The qualifier is ignored on this command. The limits set by **set-limits** apply to all output of variables or expressions or program locations. If a printed value is truncated because of these limits, the value will be followed by ellipses.

Note that the limitation on array elements applies to each dimension of a multi-dimensional array. If you print a 50 x 20 two-dimensional array, and you have the array limit set to 5, then you will see the first 5 elements of the each of the first 5 rows (or columns, for Fortran).

The default limits are 100 array elements, 100 characters, and 10 addresses. To find out what the current limits are, use the **info limits** command (See “info limits” on page 6-168).

set-prompt

Set the string used to prompt for command input.

```
set-prompt string
```

string

Specify the string the debugger uses to prompt for command input. The string must be enclosed in double quotes. If you include any of the following substrings in the prompt, they will be expanded by the debugger immediately prior to printing the prompt.

`%q`

Expands to the current default qualifier. This prints out the same way the qualifier was defined. If you used a family name, it shows the family name (not the individual PIDs), etc. If the default qualifier is `auto`, it prints the current automatically selected PID.

`%p`

Expands to the complete list of PIDs implied by the current default qualifier.

`%d`

Expands to the complete list of dialogues implied by the current default qualifier.

`%a`

Expands to the complete list of dialogues, if the current default qualifier is `all`. Otherwise, this expands to the current default qualifier.

`%%`

Expands to the single character `%`.

The *string* argument may also include the escape sequences recognized in C language strings, such as `'\n'` to indicate a newline.

The string ``` (%a) "` is the default prompt.

The qualifier on the **set-prompt** command is ignored.

Examples:

```
(afamily) set-prompt "%p> "
local:2047,2048>
```

The above example shows what happens when the default qualifier is a process family named `afamily` assumed to contain two PIDs (2047 and 2048), both in dialogue `local`. The initial prompt is `" (%q) "` and the **set-prompt** command changes it to expand to a list of PIDs.

```
(afamily) set-prompt "Dialogues: %d\nProcesses: %p>"
Dialogues: mach1,mach2
Processes: mach1:15 mach2:15,549,2047,2048>
```

The above example prints two lines as a prompt, the first containing a list of dialogues and the second containing a list of processes.

set-format-hex

Control whether or not many values are displayed in hexadecimal in addition to their natural format. This setting affects commands such as `print` (see “`print`” on page 6-91) and the data panel (see “Data Panel” on page 8-75).

set-format-hex [{on | off}]

on

Enable hexadecimal display.

off

Disable hexadecimal display.

set-terminator

Set the string used to recognize end of dialogue input mode.

set-terminator *string*

string

Define the *string* used to terminate dialogue input mode (see “!” on page 6-34).

When the **!** command is used to switch all input to a dialogue, the terminator string is recognized to switch input back to the debugger. The terminator string must appear on a line by itself to be recognized. The default string is “-.” (different from **rlogin** and **cu**).

Unlike normal debugger commands, this string must be typed exactly as specified in the **set-terminator** command. The case of the letters must match, and the full string must be typed.

Only one terminator string is defined. The qualifier on this command is ignored.

Leading and trailing whitespace in the specified terminator string is ignored. Macros are *not* expanded when reading the new terminator string.

If no terminator string is given, then the current terminator string is printed, otherwise the new terminator string is printed.

set-safety

Control debugger response to dangerous commands.

set-safety [forbid | verify | unsafe]

forbid

In `forbid` mode, the debugger simply refuses to execute a dangerous command and explains why it will not execute. (You may have tried to `quit` while processes were still running, etc.).

verify

In `verify` mode, the debugger tells you what dangerous thing you are about to do and asks if you really meant that (see “Replying to Debugger Questions” on page 6-22). If you answer `yes`, it goes ahead and does it. This is the default safety level of the debugger.

unsafe

In `unsafe` mode, the debugger simply tells you what it did. It assumes you meant what you said and does not try to stop you.

If no mode is specified then the **set-safety** command prints the current safety level.

The qualifier on the **set-safety** command is ignored.

set-restart

Control whether restart information is applied.

set-restart [`always` | `never` | `verify`]

always

Restart information is unconditionally applied when a program starts. This is the default mode.

never

Restart information is never applied when a program starts.

verify

When a program starts, you are asked whether to apply restart information to it.

If no keyword is specified then the **set-restart** command prints the current restart mode.

The restart mode is a global mode, not a per-process or per-dialogue mode. The qualifier on the **set-restart** command is ignored.

See “Restarting a Program” on page 3-20.

set-local

Define process local convenience variables.

set-local *identifier* ...

identifier

The name of a convenience variable (the leading '\$' on each identifier, normally used to reference convenience variables, is optional).

Each named identifier is defined to be a process local convenience variable.

A process local variable always has a unique value in each process. If the variable was already defined as a global at the time it appears in a **set-local** command, then each process gets a separate copy of the current global value, but future changes will be unique for each process.

The command qualifier does not have any effect on this command. It is not possible to define a variable to be local for only one process, but globally shared among other processes.

set-patch-area-size

Control the size of patch areas created in your process.

set-patch-area-size {data=*data-size* | eventpoint=*eventpoint-size* |
monitor=*monitor-size* | text=*text-size* |
timing=*timing-size*} ...

data=*data-size*

The data keyword parameter specifies the size of the data area in kilobytes.

monitor=*monitor-size*

The monitor keyword parameter specifies the size of the shared memory region used by all monitorpoints in this dialogue, in kilobytes.

text=*text-size*

The text keyword parameter specifies the size of the text area in kilobytes.

eventpoint=*eventpoint-size*

The eventpoint keyword parameter specifies the size of the eventpoint areas in kilobytes.

timing=*timing-size*

The timing keyword parameter specifies the size of the shared memory region used for all timing data in this process, in kilobytes.

The data, monitor, text, eventpoint and timing keywords may be abbreviated and may be specified in any order.

NightView creates some regions in your process, and uses these regions to store text and data. There is usually one data region, one text region, one or more eventpoint regions. If there are any monitorpoints in the process, there is one shared memory region for their data. There is one shared memory region for timing data.

These regions are called *patch areas*. See Appendix C [Implementation Overview] on page C-1.

You can adjust the sizes of the patch areas with this command. For example, if you have a lot of conditional eventpoints, then you may need to make the size of the eventpoint and text regions larger so that NightView has room to allocate all the code necessary for those eventpoints. Similarly, if you have a lot of monitorpoints, then you may need to make the size of the monitorpoint shared memory region larger. On the other hand, if system memory resources are scarce, then you may need to make some of these regions smaller.

The patch area size values are associated with each dialogue and apply to all processes within the dialogue. This command sets the values for each dialogue specified in the qualifier.

Note that these values only apply to patch areas created in the future. Existing regions are not changed. Therefore, if you want to debug a program and use a large text or data area, you need to specify that before you run your program (i.e., before the process calls `exec`). (For `fork`, the child process inherits its regions from the parent, so the regions are the same size in the child and the parent.)

Each process has its own data, eventpoint, text, and timing areas, but the monitorpoint shared memory region is shared by all the processes that have monitorpoints in the dialogue, and by the dialogue itself. Therefore, if you want to change the size of the monitorpoint shared memory region, you need to do so before creating any monitorpoints in the dialogue. See “Monitorpoints” on page 3-12.

The initial values of the patch area sizes are 512 kilobytes each for the data and text patch areas, 256 kilobytes for the eventpoint areas, 32 kilobytes for the monitorpoint shared memory region, and 32 kilobytes for the timing shared memory region. This is adequate for most applications.

Use **info dialogue** to see the current patch area size values. (see “info dialogue” on page 6-173).

You can see information about the patch areas in an existing process with the **info memory** command (see “info memory” on page 6-170).

interest

Control which subprograms are interesting.

interest [*level*] [[*at*] [*location-spec*]]

Set or query the interest level for a subprogram.

interest inline[=*level*]

interest justlines[=*level*]

interest nodebug[=*level*]

interest cuda_syscall[=*level*]

interest threshold[=*level*]

Set or query the interest keyword values.

level

Specify a level for the subprogram defined by *location-spec*, or a value for the specified keyword. *level* is a signed integer or the keywords `minimum` or `maximum`. If this argument is not present, then this command queries the level of the subprogram or the specified keyword.

[*at*] *location-spec*

Set or query the interest level for the subprogram specified by *location-spec*. See “Location Specifiers” on page 6-15. If no *location-spec* is present, it defaults to `*$cpc`. If the `at` keyword is present, it must be followed by a *location-spec*. If no *level* is specified, then the `at` keyword is required to distinguish some forms of location specifiers from a *level*.

inline

Set or query the inline interest level. If this level is less than the interest level threshold, then all inline subprograms have the `minimum` interest level unless their interest level has been explicitly set with **interest** *level location-spec*. The initial value of this level is 0.

justlines

Set or query the interest level for subprograms with line number information but no other debug information. The initial value is -2.

nodebug

Set or query the interest level for subprograms with no debug information (e.g., system library routines). Without debug information, the interest level cannot be specified for individual subprograms, so NightView uses the value specified by this form. The initial value is -4.

cuda_syscall

Set or query the interest level for artificial subprograms defined by the CUDA runtime for execution of its internal syscalls. The initial value is -4.

threshold

Set or query the interest level threshold NightView uses to decide whether a subprogram is interesting. The initial value is 0.

The **interest** command sets or queries the information NightView uses to decide which subprograms are interesting for each process in the qualifier. See “Interesting Subprograms” on page 3-31.

The `minimum` keyword specifies the lowest possible interest level. The `maximum` keyword specifies the highest possible interest level.

A query prints the interest information requested. If an interest level is being set, the command prints the new interest level.

Some compilers provide a means to specify the interest level of a subprogram through the debug information. If the subprogram has debug information, but it does not specify an interest level, the default level is 0. The **interest** command overrides an interest level set at compile time.

The interest levels and the interest level threshold are remembered as part of the restart information. See “Restart Information” on page 3-21. For a way to see all the interest levels that have been explicitly set, see “info on restart” on page 6-175.

If an interest level or the interest level threshold is changed, then NightView checks the current frame to see if it has become uninteresting. See “Current Frame” on page 3-29. If it has, then the current frame is reset to frame 0 of the current context and frame information is printed. See “select-context” on page 6-149. Even if the current frame does not have to be reset, it gets a different frame number if frames below it have become hidden or unhidden.

Examples:

```
(local) run fact 7
...process startup information...
(local) interest
local:6729: Interest level is -4 (uninteresting) for 0x100024d0
(nodebug)
```

You query the interest level, using the default location specifier of `*$cpc`. The program begins in the C runtime startup routine, which has no debug information, so it is uninteresting.

```
(local) breakpoint 26
local:6729 Breakpoint 1 set at fact.c:26
(local) continue
local:6729: at Breakpoint 1, 0x10002780 in main(int argc = 2,
unsigned char ** argv = 0x2ff7eae4) at fact.c line 26
26 B=|          answer = factorial(x);
(local) step
#0 0x100026f4 in factorial(int x = 7) at fact.c line 6
6 = |    if (x <= 1) {
(local) interest -1
local:6729: Interest level set to -1 (uninteresting) for
factorial
#0 0x10002780 in main(int argc = 2, unsigned char ** argv =
0x2ff7eae4) at fact.c line 26S
26 B<>|          answer = factorial(x);
```

You step into the `factorial` function, then decide that it is not interesting. You mark `factorial` uninteresting, using the default location specifier. Your current frame becomes uninteresting, so it is reset to frame 0. Frame 0 is now the frame for `main`, because `factorial` is not interesting. The source decorations for line 26 show that `$pc` and `$cpc` are within that line. See “Source Line Decorations” on page 6-88.

```
(local) interest threshold=-1
local:6729: threshold interest level set to -1
(local) frame
Output for process local:6729
#1 0x10002780 in main(int argc = 2, unsigned char ** argv =
0x2ff7eae4) at fact.c line 26
26 B<>|          answer = factorial(x);
```

You change the interest level threshold, which makes `factorial` interesting again. Your current frame is still interesting, so it is not reset to frame 0. The **frame** command shows that your current frame is still the frame for `main`, but now that frame is frame number 1.

set-auto-frame

Control the positioning of the stack when a process stops.

set-auto-frame *args* ...

The functionality of this command has been subsumed by the **interest** command. See “interest” on page 6-70. This command has been retained for compatibility, but it might be removed in some future release.

set-overload

Control how NightView treats overloaded operators and routines in expressions.

set-overload [operator={on | off}] [routine={on | off}]

operator={on | off}

Turn operator overloading on or off.

routine={on | off}

Turn routine overloading on or off.

The **set-overload** command determines how NightView treats overloaded operators, functions, and procedures in expressions. See “Expression Evaluation” on page 3-25. This behavior can be controlled for operators separately from functions and procedures using the keywords on the command. The specified settings apply to all expressions evaluated by NightView. The qualifier is ignored by the **set-overload** command. The `routine` mode also controls overloading of function names which appear in location specifiers.

After setting the specified overloading modes, the **set-overload** command prints the new settings. If no arguments are specified, the command simply prints the existing overloading modes.

For a discussion of how overloading works in NightView see “Overloading” on page 3-27. For the details of the syntax used to specify overloading in expressions and location specifiers see “Selecting Overloaded Entities” on page 6-2.

When NightView starts, the overloading modes are initially:

set-overload operator=off routine=on

set-search

Control case sensitivity of regular expressions in NightView.

set-search [sensitive | insensitive]

sensitive

Make regular expressions case sensitive (this is the default setting).

insensitive

Make regular expressions case insensitive.

The **set-search** command controls case sensitivity for the regular expressions (see “Regular Expressions” on page 6-19) used by several commands as well as some dialog boxes in the graphical interface.

When the **set-search** command is run with no argument, it reports (but does not change) the current mode setting.

When the *sensitive* argument is specified, regular expressions become case sensitive. The case of alphabetic characters must match exactly as written in the regular expression. This is the default **set-search** mode.

When the *insensitive* argument is specified, regular expressions become case insensitive. Either the upper case or the lower case form of an alphabetic character will match both the upper and lower case form of that same character.

set-editor

Set the mode for editing commands in the simple full-screen interface.

set-editor *mode*

mode

One of *emacs*, *gmacs*, *vi*, or *dumb*.

Determine which kind of keystroke commands are available to edit commands in the simple full-screen interface.

See “Editing Commands in the Simple Full-Screen Interface” on page 7-2.

set-preallocate

Control how NightView preallocates memory for eventpoints, monitorpoint buffers, and timing buffers.

set-preallocate [/eventpoint] [/monitorpoint]
[/timing] [{off | on}]

/eventpoint

Indicates the eventpoint mode should be set or queried.

/monitorpoint

Indicates the monitorpoint buffer mode should be set or queried.


```
/timing
```

Indicates the timing buffer mode should be set or queried.

```
{off | on}
```

Turn preallocation off or on for eventpoints, monitorpoint buffers, or timing buffers as specified.

The default is for NightView to preallocate space in the user process for eventpoints, monitorpoint buffers, and timing buffers. If these modes are off, then NightView allocates space only when needed. However, space can be allocated only when the process is stopped. With preallocation, you do not need to worry about whether the process is stopped when you set an eventpoint or define a timing region. See “Process States” on page 3-23. See “Operations While the Process Is Executing” on page 3-23.

Monitorpoints cannot be used in the command-line interface, so in that interface monitorpoint buffers are never preallocated.

If NightView is prevented from preallocating space for eventpoints, then support for CUDA will be disabled. This is because a hidden patch must be installed at program start time to detect the presence of CUDA code later in the application execution.

This is a global mode. The qualifier is ignored. The setting affects only future processes, not currently running processes.

With no arguments, **set-preallocate** prints the current settings.

set-resume

Control NightView’s behavior on events

```
set-resume [/attach] [/exec] [/exit] [/fork]
             [/cuda] [/cuda_system]
             [{off | on}]
```

```
/attach
```

Indicates the attach mode should be set or queried.

```
/exec
```

Indicates the exec mode should be set or queried.

```
/exit
```

Indicates the exit mode should be set or queried.

```
/fork
```

Indicates the fork mode should be set or queried. NightView pays attention to this for a child only if the parent process has **set-children all**. See “set-children” on page 6-52.

`/cuda`

Indicates the `cuda` mode should be set or queried. This mode determines whether or not the process is stopped when a CUDA application kernel is launched.

`/cuda_system`

Indicates the `cuda_system` mode should be set or queried. This mode determines whether or not the process is stopped when a CUDA system kernel (e.g. the kernel for `cudaMemset`) is launched.

`{off | on}`

Turn automatic resume off or on for the modes specified.

A process is normally stopped when NightView attaches to it, when it execs, when it is about to exit, and when it is created (i.e., its parent forks). A process normally is not stopped when NightView detects any CUDA kernel launches. This command allows you to control that for each case.

This is a global mode. The qualifier is ignored.

If `off` or `on` is not specified, **set-resume** displays the current settings.

set-download

Control how NightView downloads files from remote targets.

set-download `[[off | permanent | temporary]]`
`[directory=path-to-cache]`

`off`

This disables the download feature.

`permanent`

This enables downloading files to the cache. The files are left in the cache for future use.

`temporary`

This enables downloading files to the cache. When NightView exits, it removes all the files in the cache that were downloaded under temporary mode. This is the default setting.

`directory=path-to-cache`

This specifies the directory to use to build the cache. The default directory name is `~/NightViewCache`. The downloads will be faster if this directory is local to the system where you are running NightView.

If NightView cannot find an object file on the local system, and this process is in a remote dialogue, then NightView searches on the target system for the file and copies it to the cache directory. See “Remote File Access” on page 3-7.

This is a global mode. Any qualifier is ignored.

To manipulate this mode in the graphical user interface, see “Preferences Advanced Page” on page 8-55.

With no arguments, **set-download** prints the current download information.

set-disassembly

Control how NightView displays disassembled instructions.

```
set-disassembly [flavor={att | intel}] [symbols={off | on}]
                  [offset={decimal | hexadecimal}]
                  [comment_level=number]
```

```
flavor={att | intel}
```

Set the flavor of disassembly to the `att` style or the `intel` style.

The default flavor is `att`. The `att` flavor is the one used by the standard compilation tools. The `intel` flavor is the one described in the *IA-32 Intel Architecture Software Developer's Manual* or the *AMD64 Architecture Programmer's Manual*. For a discussion of the differences in the two flavors, enter this command in a shell outside of NightView:

```
info 'gas' 'Machine Dependencies' i386-Dependent
```

```
symbols={off | on}
```

Indicate whether each line of the disassembly should include the name of the routine being disassembled. If this is set to `on`, `<routine+offset>` is appended to each address, where *routine* is the name of the routine being disassembled and *offset* is the offset (in bytes) from the beginning of *routine*. If this is set to `off`, then the routine name appears once at the beginning of the disassembly, and `<+offset>` is appended to each address. References to addresses outside the routine being disassembled include the name in either mode. The default is `off`, which makes the disassembly listing more compact.

```
offset={decimal | hexadecimal}
```

Indicate whether symbol offsets in either of the above forms should be presented in decimal or hexadecimal.

```
comment_level=number
```

Indicate the kind of comments the disassembler should provide. Comments provide more information about the instructions.

number is one of:

0

never print comments

1

print operation information for certain instructions

2

print "aka" ("also known as") opcode aliases in addition to the level 1 comments

This command sets the disassembly modes for the debug session. Any qualifier is ignored. See “x” on page 6-95. See “Source Menu” on page 8-11.

If no arguments are specified, **set-disassembly** prints the current disassembly information.

set-branch-tracking

Control whether or not NightView and the RedHawk kernel are tracking branch instructions. See “Branch Tracking” on page 3-38.

set-branch-tracking [{on | off}]

on

Enable branch tracking.

off

Disable branch tracking.

There is considerable time overhead while tracking branches, so this mode should not be enabled without consideration. But it is very useful when attempting to understand bugs where control of the program has been transferred to totally unexpected places, as might happen when calling a function pointer whose value has been overwritten with garbage, or when returning from a function when then return address on the stack has been overwritten with garbage.

set-futurepoints

Control whether or not NightView accepts location specifiers for locations which do not exist yet.

set-futurepoints [{create | ask | error}]

create

Always will accept location specifiers for locations which do not exist yet.

ask

Ask before accepting location specifiers for locations which do not exist yet.

error

Never accept location specifiers for locations which do not exist yet.

To support setting eventpoints in shared libraries or CUDA object which are not loaded at program start time and may not be loaded for some time thereafter, NightView can accept inserted eventpoints at locations which do not exist yet. When a shared library or CUDA kernel later is loaded, if the eventpoint location is meaningful within the context of that new code, it will be inserted there.

But the acceptance of locations which do not exist yet can be confusing if a location was simply mis-specified as a typo. So NightView provides the above three behaviors as options to the user.

set-cuda

Enable or disable CUDA support.

```
set-cuda [{off | on}]
```

```
{off | on}
```

Turn CUDA support off or on as specified.

By default, CUDA support is enabled. If you wish to debug an application with CUDA code, but do not wish to debug that code, you may disable support for it by turning this off. In that case, the CUDA hardware will be treated only as a device.

Note that CUDA support also can be disabled with the **set-preallocate** command (see “set-preallocate” on page 6-74).

This is a global mode. The qualifier is ignored. The setting affects only future processes, not currently running processes, even if currently running processes have not yet executed any CUDA code.

set-cuda-software-preemption

Enable or disable CUDA software preemption.

```
set-cuda-software-preemption [{off | on}]
```

```
{off | on}
```

Turn CUDA software preemption off or on as specified.

CUDA software preemption allows a device to be used by other applications when an application running on the device is stopped for debugging. It requires CUDA 5.5 or higher, and a device with CUDA capability 3.5 or higher.

NOTE

This feature often is used to allow debugging a CUDA application using the same device as an X11 display.

set-cuda-memcheck

Enable or disable CUDA memcheck support for more precise memory exceptions.

```
set-cuda-memcheck [{off | on}]
```

```
{off | on}
```

Turn CUDA memcheck support off or on as specified.

By default, CUDA memory exceptions are imprecise. If one is detected, the device may stop at a `pc` which is a bit removed from the place that performed the incorrect memory load or store, possibly even in a different thread, and the description of the exception will be imprecise.

CUDA memcheck support changes memory faults into a more precise form of exception. They will be reported at the place which caused the exception, and they will contain the exact address that caused the exception.

The use of CUDA memcheck substantially increases the overhead of the CUDA code, possibly even to the point of changing the behavior of the CUDA code. It should be used with caution. But it can be useful in determining the exact thread and `pc` where a memory exception occurred.

CUDA memcheck can only be enabled before any CUDA code has executed in the process. Usually, this means it must be set as soon as the program starts. Any attempt to turn CUDA memcheck on after CUDA code has executed will be rejected.

CUDA memcheck support was removed in CUDA 12.0 and later releases.

Debugger Environment Control

cd

Set the debugger's default working directory.

cd *dirname*

dirname

The name of the directory.

The **cd** command changes the working directory of NightView to the specified directory. You usually use this command to control the search for source files, core files, and program files. It affects the behavior of the following commands:

- **shell** (see “shell” on page 6-152)
- **list** (see “list” on page 6-82)
- **directory** (see “directory” on page 6-84)
- **symbol-file** (see “symbol-file” on page 6-42)
- **core-file** (see “core-file” on page 6-43)
- **exec-file** (see “exec-file” on page 6-46)

The **cd** command does not affect commands executed in dialogue shells (see “login” on page 6-25). Also, the qualifier does not have any effect on this command.

You can use the **pwd** command to find out what NightView's current working directory is. See “pwd” on page 6-81.

pwd

Print NightView's current working directory.

pwd

This command prints the current working directory of the debugger. Note that this directory may not be the same as the current working directory of your dialogue shells, nor need it be the same as the current working directory of any program you are debugging.

You can use the **cd** command to set the current working directory. (see “cd” on page 6-81).

The qualifier does not have any effect on this command.

Source Files

This section describes commands to view and edit source files and to search for text in source files.

Viewing and Editing Source Files

list

List a source file. This command has many forms, which are summarized below.

list *where-spec*

List ten lines centered on the line specified by *where-spec*.

list *where-spec1*, *where-spec2*

List the lines beginning with *where-spec1* up to and including the *where-spec2* line.

list ,*where-spec*

List ten lines ending at the line specified by *where-spec*.

list *where-spec*,

List the ten lines starting at *where-spec*. Note the comma.

list +

List the ten lines just after the lines last listed.

list -

List the ten lines immediately preceding the lines last listed.

list =

List the last set of lines listed. If the previous command was a search command, list the ten lines around the line found by the search.

list

If a list command has not been given since the current source file was last established (see below), this form lists the ten lines centered around the line where execution is stopped in the current source file. Otherwise, this form lists the ten lines just after the last lines listed.

Abbreviation: **l**

Each *where-spec* argument can be any one of the following forms.

[at] *location-spec*

Specifies a location in the program or a source file (See “Location Specifiers” on page 6-15). No matter which form of *location-spec* you use, it is always translated into a source line specification for this command. If you give two arguments on the **list** command, they cannot specify different source files.

[at] *file_name*

Specifies the first line of the file. The *file_name* may be a quoted or unquoted string, but be aware that an unquoted string may be ambiguous. A string without quotes will be interpreted first as a function name; if no such function unit exists, the string will then be interpreted as a file name.

A file name on a remote system can be specified using the form *user@host:/path*. See “Remote File Access” on page 3-7.

+*n*

Specifies the line that is *n* lines *after* the last line in the last group listed (see below). If this is the second *where-spec*, it specifies the line *n* lines after the first argument.

-*n*

Like +*n*, except it specifies the line *n* lines *before* the last line in the last group listed (see below). If this is the second *where-spec*, it specifies the line *n* lines before the first argument.

The **list** command is applied to each process in the qualifier. If the qualifier specifies more than one process, you get one listing for each process; each listing is preceded by a notation indicating which process the listing is for. The specified source file is found using the directory search path you established using the **directory** command (see “directory” on page 6-84). Note that each program has its own directory search path.

NightView maintains, for each process, a current source file. The current source file is usually the most recent file listed or searched. However, when the process stops execution, the current source file is automatically set to the file where execution stopped. The context selection commands (see “Selecting Context” on page 6-147) also set the current source file to the one associated with the selected stack frame. When a process first starts execution, the current source file is the one containing the main program. If the first argument to the **list** command does not explicitly specify a source file, then the current source file is used.

When you list one or more lines in a source file, NightView remembers the first and last line of that group. If you subsequently give a **list** command that uses a relative *where-spec* or contains just a + or - argument, those arguments are interpreted relative to the lines in the last group listed. Arguments containing a + are relative to the last line in the group, and arguments containing a - are relative to the first line in the group. This also affects the **forward-search** and **reverse-search** commands. See “forward-search” on page 6-86 and “reverse-search” on page 6-86.

Repeating the **list** command by entering a blank line behaves differently depending on the form of **list** you used last. In most cases, repeating the command lists the next ten lines following the last line in the last group. However, if you used the **list -** form last, then repetition lists the ten lines preceding the first line in the last group.

The listed source lines are preceded by *source decorations*. (see “Source Line Decorations” on page 6-88).

You can use the **info line** command to determine the location in your program of the code for a particular source line. (see “info line” on page 6-183).

directory

Set the directory search path.

directory [*dirname* ...]

dirname

The name of a directory to include in the search path. If this is not an absolute pathname, it is interpreted relative to NightView's current working directory and transformed into an absolute pathname. Thus, if you later change NightView's working directory, the search path will not be affected. See “cd” on page 6-81 and “pwd” on page 6-81.

The search can be performed on remote systems by specifying *dirname* in the form *user@host:/dirname*. See “Remote File Access” on page 3-7.

The **directory** command sets the directory search path for the program in each process in the qualifier. The arguments are used in order as the elements of the directory search path. Subsequent **directory** commands contribute directories to the head of the current search path.

The directory search path is used for displaying source files. When you list a source file (see “list” on page 6-82), NightView looks for the source file in each of the directories in the search path, starting at the beginning of the search path each time.

If no **directory** command has been specified for the program, the search path implicitly contains the path to the executable file and NightView's current working directory. Once a **directory** command is specified for the program, these directories are no longer implicit in the search path.

If you enter a **directory** command with no arguments, the search path is reset to its initial state.

The directory search path is associated with a program, not with a process. If you debug multiple instances of a program, the directory search path is the same for each instance. If your process calls **exec(3)**, the directory search path is implicitly set for the new program.

Use the **info directories** command to display the directory search path for a program. See “info directories” on page 6-167.

For ELF programs, the debugging information contains absolute pathnames to source files, so the directory search path may not be needed. It is still sometimes useful to indicate that a source tree is not where the debugging information indicates.

Examples:

Suppose your ELF program was compiled from two source files:

`/usr/bob/src/main/main.c` and `/usr/bob/src/doit/doit.c`. You want to debug your program, but you have moved the source files to `/usr/joe/main/main.c` and `/usr/joe/doit/doit.c`. Enter a **directory** command to indicate the new root of the source tree:

```
(local) directory /usr/joe
```

Similarly, if the source files are now on system "oursys", use this command:

```
(local) directory oursys:/usr/joe
```

edit

Edit the current source file.

edit

This command invokes a text editor on the source file currently displayed in the source panel.

This command can be used only from the graphical user interface. See "Source Panel" on page 8-63.

Searching

forward-search

Search forward through the current source file for a specified regular expression.

forward-search [*regex*]

Abbreviation: **fo**

regex

The regular expression to search for. *No* anchored match is implied. (see “Regular Expressions” on page 6-19). If *regex* is omitted, the previous *regex* is used.

The search command is applied to the current source file of each process specified by the qualifier.

The search starts at the first line displayed by the last **list** command, the last place the process stopped, or the last place a search was satisfied, whichever was most recent, and proceeds forward through the file to the end. In the graphical user interface, the search position is not affected by scrolling the source panel. If the regular expression is found, the containing source line is listed. This will affect subsequent **list** commands that specify relative arguments.

If the end of the file is encountered without finding the regular expression, a message is printed indicating the search was unsuccessful. For a definition of current source file, see “list” on page 6-82.

The regular expression case sensitivity depends on the current search mode (see “set-search” on page 6-73).

reverse-search

Search backwards through the current source file for a specified regular expression.

reverse-search [*regex*]

regex

The regular expression to search for. *No* anchored match is implied. (see “Regular Expressions” on page 6-19). If *regex* is omitted, the previous *regex* is used.

The search command is applied to the current source file of each process specified by the qualifier. The search starts at the last line displayed by the last **list** command, the last place the process stopped, or the last place a search was satisfied, whichever was most recent, and proceeds backward through the file to the beginning. In the graphical user interface, the search position is not affected by scrolling the source panel. If the regular expression is found, the containing source line is listed. This will affect subsequent **list** commands that specify relative arguments.

If the beginning of the file is encountered without finding the regular expression, a message is printed indicating the search was unsuccessful. For a definition of current source file, see “list” on page 6-82.

The regular expression case sensitivity depends on the current search mode (see “set-search” on page 6-73).

Source Line Decorations

When NightView lists source lines in your program or displays the disassembled instructions of your program, it precedes each line with decorations providing information about that line. Every source line gets a *line number*, which is relative to the beginning of that file. Each instruction displayed is preceded by the line number of the source line that generated it (see “x” on page 6-95).

In the serial interface, line numbers precede the decorations. If multiple decorations are needed, they are concatenated, such as BP=.

In the graphical user interface, line numbers follow the decorations. If multiple eventpoints are present, that is represented by a separate icon. If the program is stopped on the line, the program counter icon is overlaid on any other icon. See “Graphical User Interface” on page 8-1. If you hover the mouse pointer over the icon, a tooltip shows the eventpoint information for the line.

Also, in the graphical user interface, the icon may indicate additional information. Here are the rules NightView uses for choosing which icon is displayed:

- If there are different kinds of eventpoints on the line, a generic eventpoint icon is shown. No other information is represented with this icon. 
- If all the eventpoints on the line are disabled, the icon has reduced color. 
- If the eventpoint has a non-zero ignore count, or the eventpoint has a condition, there are multiple eventpoints (of the same kind) on the line, or the location of the eventpoint is not yet known (see “set-futurepoints” on page 6-78), the icon has a pink background. This indicates that more information is available. 
- Otherwise, this is a single simple eventpoint and the regular icon is used. 

The following table lists the source line decorations. The decoration for the serial interface is listed first, followed by the icon for the graphical user interface.

Table 6-7. Source Line Decorations

'B' 

Indicates that one or more breakpoints, possibly disabled, are set somewhere within this source line. When displaying instructions, this indicates that one or more breakpoints are set on this instruction. (See “breakpoint” on page 6-108).

'H' 

Indicates that one or more heappoints, possibly disabled, are set somewhere within this source line. When displaying instructions, this indicates that one or more heappoints are set on this instruction. (See “heappoint” on page 6-117).

'M' 

Indicates that one or more monitorpoints, possibly disabled, are set somewhere within this source line. When displaying instructions, this indicates that one or more monitorpoints are set on this instruction. (See “monitorpoint” on page 6-115).

'P' 

Indicates that one or more patchpoints, possibly disabled, have been inserted somewhere within this source line. (See “patchpoint” on page 6-110). When displaying instructions, this indicates the instruction where the patchpoint was inserted, and the patched expressions are displayed elsewhere.

'T' 

Indicates that one or more tracepoints, possibly disabled, are set within this source line. When displaying instructions, this indicates a tracepoint immediately preceding this instruction. (See “tracepoint” on page 6-113).



Indicates that multiple kinds of eventpoints, possibly disabled, are set within this source line. When displaying instructions, this indicates multiple kinds of eventpoints immediately preceding this instruction. In the serial interface, eventpoint characters are concatenated.

'=' 

Indicates that execution is stopped somewhere within or at the beginning of this line. When displaying instructions, this indicates the instruction at which execution is stopped (the one that will next be executed). In the graphical user interface, this icon is overlaid over any other icon.

'>'

Indicates the line (or instruction) in the current frame (see “frame” on page 6-147), where execution will resume when the called routine returns. This is not represented in the graphical user interface.

This decoration is not displayed if the current frame is frame #0 (with no hidden frames below frame 0); in this case the '=' decoration will appear in its place.

'<' 

Indicates the line (or instruction) in the current frame (see “frame” on page 6-147), which was executing when the called frame was created, i.e., `$CPC`. See “Program Counter” on page 3-28.

This decoration is not displayed if the current frame is frame #0 (with no hidden frames below frame 0); in this case the '=' decoration will appear in its place.

'*' 

Indicates that this source line corresponds to executable code. A line that appears executable may still not have executable code associated with it because of optimization or conditional compilation. Not used when displaying instructions.

This decoration is not displayed if there are any other indicators also on that line, since the other indicators imply there is executable code for the line.

'@' ♦

Used only when displaying instructions, this decoration indicates that the associated instruction is the first for the corresponding source line.

In the serial interface, NightView reserves enough columns for displaying a 3-digit line number, 2 decoration characters, and a 2-character separator. If the line number and decorations fit within this space, the source text displayed lines up in columns just as it does in the source file. If more space is needed for line number or decorations, the line is shifted over accordingly. In the graphical user interface, the line numbers are expanded as necessary and the source decorations always take the same amount of space.

In the serial interface source listing, the 2-character separator is a vertical bar followed by a space. This helps distinguish decorations from source characters. In the serial interface disassembly listing, the 2-character separator consists of 2 spaces.

Example source listing, in the serial interface:

```

20  | void
21 * | main(argc, argv)
22  |     int argc;
23  |     char ** argv;
24  | {
25  |     int i, errors;
26 * |     errors = 0;
27 * |     for (i = 1; i < argc; ++i) {
28  |         long xl;
29  |         int x;
30  |         int answer;
31 * |         char * ends = NULL;
32 T |         xl = strtol(argv[i], &ends, 10);
33 B=|         x = (int)xl;
34 B |         answer = factorial(x);
35 P |         printf("factorial(%d) == %d\n", x, answer);
36  |     }
37 * |     exit(errors);
38  | }
```

In this example, line 32 has a tracepoint set on it; line 33 has a breakpoint set somewhere within the line, and execution is stopped on the line (but not necessarily at the breakpoint). Line 34 has a breakpoint set somewhere within the line (perhaps on the return from `factorial`). Line 35 has a patchpoint inserted somewhere within it. Apart from these lines, the other lines with asterisks on them have executable code associated with them.

Example instruction listing:

```

31 @ 0x10002788 <main+52>:    li r6,0
31  0x1000278c <main+56>:    stw r6,0x40(r1)
32 @T 0x10002790 <main+60>:   slwi r5,r16,2
32   0x10002794 <main+64>:   lwzx r3,r17,r5
32   0x10002798 <main+68>:   addi r4,r1,64
32   0x1000279c <main+72>:   li r5,10
32   0x100027a0 <main+76>:   bl 0x100010e0 <strtoul>
```



```

33 @B= 0x100027a4 <main+80>: mr r20,r3
34 @ 0x100027a8 <main+84>: bl 0x10002700 <factorial>
34 B 0x100027ac <main+88>: mr r5,r3
35 @P 0x100027b0 <main+92>: lis r3,12288
35 0x100027b4 <main+96>: addi r3,r3,12528
35 0x100027b8 <main+100>: mr r4,r20
35 0x100027bc <main+104>: bl 0x10001100 <printf>

```

This is a partial disassembly listing for the preceding example source listing.

Examining and Modifying

backtrace

Print an ordered list of the currently active stack frames.

backtrace [*number-of-frames*]

Abbreviation: **bt**

number-of-frames

Number of stack frames to print, starting with the currently executing frame.

The **backtrace** command prints, for each process specified in the qualifier, a summary of the active stack frames, starting with the currently executing frame. Each subsequent entry corresponds to the caller of the frame which precedes it in the listing. All active frames are indicated, unless a value for *number-of-frames* is given, in which case, the given number of frames is printed.

Each entry in the **backtrace** listing includes the frame number (the first frame is numbered 0), the program counter, the subprogram name (if known), the arguments of the subprogram (if known), the source file name (if known), and the line number (if known).

For information on changing the current stack frame, see “frame” on page 6-147, “up” on page 6-148, or “down” on page 6-149.

Frames corresponding to uninteresting subprograms are not shown in the listing. See “Interesting Subprograms” on page 3-31.

print

Print the value of a language expression.

print [/print-format-code] *expression*

Abbreviation: **p**

print-format-code

One of the following codes specifying the format in which to print each component value of the expression:

a

Print the value of the expression in hexadecimal and as an address relative to a program symbol.

c

Treat the rightmost (least significant) eight bits of the value as a character constant and print the constant.

cx

Treat the rightmost (least significant) eight bits of the value as a character constant and print the constant. Also print the eight bits as a hexadecimal value.

d

Print the bit representation of the value in signed decimal.

dx

Print the bit representation of the value in both signed decimal and hexadecimal.

f

Print the bit representation of the value as a single floating-point number using floating-point syntax.

If the expression is of a floating-point type, then that type is used to interpret the bit representation. If the expression is not of a floating-point type, then the size implies how it should be interpreted:

- 4 bytes: C float (IEEE 754 binary32)
- 8 bytes: C double (IEEE 754 binary64)
- 10 or 12 bytes:
IA-32 or AMD64: C long double
(80-bit double extended precision)
- 16 bytes:
IA-32 or AMD64: C __float128 (IEEE 754 binary128)
ARM64: C long double (IEEE 754 binary128)

fx

Print the bit representation of the value as a floating-point value (as above for f) and also as a hexadecimal value.

n

Disable smart printing (see “Smart Printing” on page 3-42) for this command, overriding the smart-printing mode..

o

Print the bit representation of the value in octal.

ox

Print the bit representation of the value in both octal and hexadecimal.

s

Print the data as a character string. Arrays of characters will print as one character string (terminated with a zero byte if the language is C or C++); scalar types will print using their default format plus the bytes of the value will be printed as a string. (You might want to use this in Fortran if you put Hollerith data in INTEGER variables.)

See note below about limits on the length of printed strings.

u

Print the bit representation of the value in unsigned decimal.

ux

Print the bit representation of the value in both unsigned decimal and hexadecimal.

x

Print the bit representation of the value in hexadecimal.

y

Enable smart printing (see “Smart Printing” on page 3-42) for this command, overriding the smart-printing mode.

expression

A language expression (see “Expression Evaluation” on page 3-25).

print displays the value of a language expression in each process specified by the qualifier. When the expression is an aggregate item, such as an array, record, or union, each component value of the expression is printed, along with the appropriate subscript, record field name, etc.

The space between **print** and / may be omitted. If no *print-format-code* is given, *expression* is printed in a format corresponding to the data type of the expression in the currently defined language. For many types, (e.g. integer-like types), values will be displayed in both their natural format and in hexadecimal. This can be controlled with the set-format-hex command (see “set-format-hex” on page 6-67).

The printed value is given a value history number (see “Value History” on page 3-42),

indicated in the output by \$ followed by the history number.

If the value printed contains an array or a character string, the number of array elements and characters will be limited to the values set by the **set-limits** command (see “set-limits” on page 6-64).

NOTE

For ease in debugging C and C++ programs, the **print** command treats expressions of type 'char *' specially. Whenever **print** prints the value of a 'char *' pointer, it also prints the string it points to, inside double-quote marks; **print** assumes the string is terminated by a null byte.

Most other commands that print expressions or variables also treat 'char *' pointers in this manner.

Examples:

```
(local) (12) p/x var_name*4
(local) (12) p array_name
```

The first example prints, in hexadecimal, a number equal to four times the value of `var_name`, for process 12. The second example prints the value of each member of the array `array_name` in a format based on the data type of `array_name`, for process 12.

set

Evaluate a language expression without printing its value.

set *expression*

expression

A language expression (see “Expression Evaluation” on page 3-25).

This command is similar to the **print** command (see “print” on page 6-91), in that it evaluates a language expression for each process specified in the qualifier. However, **set** does not accept a format specifier, print the value of the expression, or place the value of the expression in the value history. It is useful for doing assignments to language objects (e.g., memory addresses preceded by the C language cast syntax, variables, and array elements) and convenience variables, as well as for performing calls to subprograms whose return value is unimportant.

Examples:

```
(local) set $i = 98
(local) (27) set vector[5] = x * 2.5
(local) set *(int *)0x1234 = 0xabcd0123
(local) set routine(3,4)
```

The first example assigns the value 98 to the convenience variable `$i`. The second example assigns the value of `x * 2.5` to element five of array `vector`, in process 27. The third example assigns the hexadecimal value `abcd0123` to the hexadecimal absolute memory location 1234. The final example performs a call to the subprogram routine.

X

Print the contents of memory beginning at a given address.

x `[ / [repeat-count] [size-letter] [x-format-code] ] [addr-expression]`

repeat-count

Decimal number of consecutive memory units to print, where a unit is defined by the *size-letter* and the *x-format-code*.

size-letter

One of the following codes specifying the size of each memory unit:

b

Each memory unit is one byte (8 bits) long.

h

Each memory unit is one halfword (two bytes) long.

w

Each memory unit is one word (four bytes) long.

g

Each memory unit is one giant word (eight bytes) long.

p

Each memory unit is the size of a pointer on the target system. On an IA-32 system, this is 4 bytes. On an AMD64 system, this is 8 bytes.

t

Each memory unit is the size of a C `long double` on the target system. On an IA-32 system, this is 12 bytes. On an AMD64 or ARM64 system, this is 16 bytes. A `t` memory unit cannot be printed as decimal `d` or `u`.

A `t` memory unit also applies to C `__float128`, on an AMD64 or ARM64 system. See `f` and `e` format codes below.

The *size-letter* may appear either before or after the *x-format-code*.

x-format-code

One of the following codes specifying the format in which to print the contents of memory:

a

Print as an integer in hexadecimal and as an address relative to a program symbol. This format ignores *size-letter* and always uses p.

c

Print as character constants. This format ignores *size-letter* and always uses b.

cx

Print values as both character constants and hexadecimal. This format ignores *size-letter* and always uses b.

d

Print as signed integers in decimal format.

dx

Print values as both signed integer decimal and hexadecimal.

f

Print as floating-point values (favoring 80-bit double extended precision for 16-byte objects).

If the size is 16 bytes and the architecture is x86_64 (AMD64), then there are two possible formats: 80-bit double extended precision (C long double), and binary128 (C __float128). f selects the 80-bit double extended precision format. For other sizes or architectures, there is no ambiguity.

fx

Print values as both floating-point (as above for f), and hexadecimal.

e

Print as floating-point values (favoring binary128 for 16-byte objects).

If the size is 16 bytes and the architecture is x86_64 (AMD64), then there are two possible formats: 80-bit double extended precision (C long float), and binary128 (C __float128). e selects the binary128 format. For other sizes or architectures, there is no ambiguity.

ex

Print values as both floating-point (as above for e), and hexadecimal.

i

Print as machine instructions in assembler syntax, using the length of each instruction as the unit size. A *repeat-count* given with this format indicates how many instructions to print.

See “set-disassembly” on page 6-77 to control the form of the disassembly.

You can also view disassembly in a source panel (see “Source Menu” on page 8-11).

o

Print as unsigned integers in octal format.

ox

Print values as both unsigned integer octal and hexadecimal.

s

Print as a null-terminated string, using the length of the string (including the null byte) as the specified unit size; the *size-letter*, if any, is ignored. A *repeat-count* given with this format indicates how many strings to print.

If the string to be printed is longer than the string limit set by the **set-limits** command, the initial characters of the string are printed, with an ellipsis following the closing quote. (see “set-limits” on page 6-64).

u

Print as unsigned integers in decimal format.

ux

Print values as both unsigned integer decimal and hexadecimal.

x

Print as unsigned integers in hexadecimal format.

z

Print as unsigned integers in hexadecimal format with a display of the corresponding ASCII characters.

addr-expression

An expression yielding a memory address (see “Expression Evaluation” on page 3-25).

The **x** command prints the contents of memory beginning at the address specified by *addr-expression* in each process specified by the qualifier. If an *addr-expression* is not given, the address corresponds to the byte following the end of the memory contents printed in the last **x** command.

If an *addr-expression* is specified and its type is a pointer type, and the current context is that of a CUDA thread, then any CUDA segment information in that type will be used to determine the segment of the address. For instance, an expression like `(__shared__ float*)0x1000` will cause the **x** command to display that address in CUDA shared memory.

The space between **x** and / may be omitted. If *repeat-count* is omitted, one memory unit is printed. If either *size-letter* or *x-format-code* is omitted, the default is the last value used in an **x** command (beginning defaults are **p** and **d**, respectively).

If the **x** command is repeated, memory contents are printed using the same *repeat-count*, *size-letter*, and *x-format-code* as in the previous **x** command, and the beginning address corresponds to the byte following the end of the memory contents printed in the previous command.

A 0 precedes octal numbers. A 0x precedes hexadecimal numbers. Thus decimal 64 would appear in hexadecimal as 0x40 and in octal as 0100.

The *x-format-code* **z** produces a hexadecimal display *without* the leading 0x prefix. The character display shows non-printable characters replaced by . (period). Here, *printable* is determined by the current locale. The display of characters is framed in | and |.

After an **x** command, the convenience variables `$_` and `$__` are set and ready to use in expressions (see “Predefined Convenience Variables” on page 6-5). The convenience variable `$_` is set to the address of the last memory unit examined. The convenience variable `$__` is set to the contents and type of the last memory unit examined.

Examples:

```
(local) (14544) x/4i $pc
7  @B= 0x1000271c <factorial+28>: li r3,1
7      0x10002720 <factorial+32>: lwz r16,0x40(r1)
7      0x10002724 <factorial+36>: lwz r13,0x58(r1)
7      0x10002728 <factorial+40>: mtlr r13
```

For the process with process id 14544, print memory as four machine instructions starting with the address of the current program counter. See “Source Line Decorations” on page 6-88 for a description of the characters at the beginning of each line of this format.

```
(local) x /4wx 0x40a188
0x0040a188: 0x77767574 0x73727170 0x6f6e6d6c 0x6b6a6968
(local) x /8bz 4235656
0x0040a188: 77 76 75 74 73 72 71 70 |wvutsrqpl|
(local)
0x0040a190: 6f 6e 6d 6c 6b 6a 69 68 |onmlkjih|
(local) p $_ - 4235656
17: $_ - 4235656 = 0xf
(local) p $__
$18: $__ = 104 'h'
```

Print memory as four words (four-byte memory units) starting at hexadecimal address 0x0040a188 as unsigned integers in hexadecimal format with 0x prefixes.

Print memory as eight bytes (one-byte memory units) starting at the same address expressed in decimal (4235656) as unsigned integers in hexadecimal format with a display of the printable characters.

Print in the same format and repeat count starting at the next address (0x0040a190).

Print an expression `$_ - 4235656` to show the relative difference between the address of the last memory unit printed `$_ - 4235656` and address of the first memory unit two commands ago `4235656`.

Print expression `$__` to show the value of the last memory unit printed.

output

Print the value of a language expression with minimal output.

output [*/print-format-code*] *expression*

print-format-code

A code specifying the format in which to print the expression, as described in the **print** command (see “print” on page 6-91).

expression

A language expression (see “Expression Evaluation” on page 3-25).

output prints the value of a language expression for each process specified by the qualifier in the same manner as the **print** command, except that a newline is not printed, the value is not entered in the value history, and the “*\$history-number* = ” string does not prefix the output.

The space between **output** and / may be omitted. If no *print-format-code* is given, *expression* is printed in a format corresponding to the data type of the expression.

echo

Print arbitrary text.

echo *text*

text

Arbitrary text to be printed, up to the end of the line. Non-printing characters may be represented with C language escape sequences, such as ‘\n’ for new-line.

This command prints the given text. It is intended as an adjunct to the other commands which print information about the program, so that the output can be customized to whatever is desired.

A backslash (‘\’) may be used to correctly print leading and trailing spaces. In other words, a backslash may be used at the beginning of *text* to print leading spaces appearing after the backslash, and one may be used at the end of *text* to print the spaces appearing before the backslash. The backslash characters themselves are not printed.

Note that a newline is not printed unless the newline sequence (`'\n'`) is included.

Examples:

```
(local) echo \  Text with two leading spaces and a newline\n
(local) echo      A backslash (\\) and the number three (\063)
```

The first example prints " Text with two leading spaces and a newline", followed by a newline. The second example prints "A backslash (\) and the number three (3)", but does not print a newline.

data-display

Control items in a data panel.

data-display [/window="*window name*"] {/kind=*value* | *expression*}

/window="window name"

Determines which data panel is affected by this command.

The default is Data.

/kind=value

value indicates which kind of item to placed in the data panel. *value* is one of locals, registers, callstack or threads.

expression

An expression to place in a data window. There should not be a */kind* keyword in this form of the command.

The **data-display** command is not intended to be used directly by users. Its main use is in restart information. See “Restart Information” on page 3-21. A description of all the forms of this command is beyond the scope of this document. However, users may sometimes have a use for the simplest forms of the **data-display** command described here.

memory-display

Control memory displayed in a binary viewer panel.

memory-display [/title="*window name*"] [/group=*group size*]
 [/format=*format*] [/start=*offset*]
 [/context=*CUDA context*|process]
 /bytes=size[b|k|m|g] *expression*

Abbreviation: **md**

/title="window name"

Sets or changes the name of the binary viewer panel.

The default is Binary Data.

/group=group size

group size indicates the data element size to be used by the panel. It must be one of these values: 1, 2, 4, 8.

/format=format

format indicates the display format to be used by the panel. Possible values are: hex, signed decimal, unsigned decimal, octal, hex plus ascii, IEEE float, and INTEL float. Also permitted is any unambiguous prefix of those strings (e.g. s for signed decimal). Beyond that, the prefixes h and he are permitted and designate hex, despite being ambiguous. The meanings of the values are described in further detail in “Binary Viewer Panel” on page 8-107

/start=offset

offset indicates the initial offset within the memory range to be displayed on the first line.

/context=CUDA context|process

If the memory to be displayed is within the process memory, then the `process` should be specified. This also is the default. If the memory to be displayed is within a CUDA context, then *CUDA context* represents its context value.

/bytes=size [b|k|m|g]

size indicates the size of the range to be displayed. By default, it is interpreted as a number of bytes. But a suffix character can be added to change the interpretation: b for bytes, k for kilobytes, m for megabytes, or g for gigabytes.

expression

An expression which is evaluated as the starting address of the range to be displayed.

The **memory-display** command is not intended to be used directly by users. Its main use is in restart information. See “Restart Information” on page 3-21. A description of all the forms of this command is beyond the scope of this document. However, users may sometimes have a use for the simplest forms of the **memory-display** command described here.

display

Add to the list of expressions to be printed each time the process stops.

display [[/print-format-code] *expression*]
display / [repeat-count] [size-letter] [x-format-code] *addr-expression*

print-format-code

A code specifying the format in which to print the expression, as in the **print** command (see “print” on page 6-91).

expression

A language expression (see “Expression Evaluation” on page 3-25).

repeat-count

Decimal number of consecutive memory units to print, where a unit is defined by the *size-letter* and the *x-format-code*.

size-letter

A letter specifying the size of each memory unit, as described in the **x** command (see “x” on page 6-95). The *size-letter* may appear either before or after the *x-format-code*.

x-format-code

A code specifying the format in which to print the contents of memory, as described in the **x** command (see “x” on page 6-95).

addr-expression

An expression yielding a memory address (see “Expression Evaluation” on page 3-25).

The display item list contains language and memory address expressions which will be used to print expression values or contents of memory, respectively, each time one of the specified processes in the qualifier stops (hits a breakpoint, receives a signal, etc.). **display** adds a language or memory address expression to the list.

In order to determine whether the given expression is a language or address expression, the parameters before the expression are first examined. If a *repeat-count* or *size-letter* is given, or if either of the *x-format-codes* 's' or 'i' is given, then the expression is treated as an *addr-expression*. Otherwise, the expression is treated as a language *expression*.

If an *addr-expression* is specified and its type is a pointer type, and the current context is that of a CUDA thread, then any CUDA segment information in that type will be used to determine the segment of the address. For instance, an expression like (`__shared__ float*`) 0x1000 will cause the **x** command to display that address in CUDA shared memory.

When one of the processes specified by the qualifier stops, each enabled item in the display item list is evaluated. The indicated expression value or memory location is displayed, each item beginning on a new line. Each display item has an item number, followed by the text of the expression and then the expression's value or the contents of memory. If a language expression for an item cannot be evaluated in the currently defined language, output will not appear for that item; however, a summary of the unevaluated items will appear at the end of the **display** output.

The space between **display** and / may be omitted. If no *print-format-code* is given for a language expression, *expression* is printed in a format corresponding to the data type of the expression at the time the process stops. If *repeat-count* is omitted, one memory unit will be printed. If *size-letter* or *x-format-code* is omitted, the defaults are *w* and *d*, respectively. If no *x-format-code* is specified, values will be displayed in both their natural format and in hexadecimal. This can be controlled with the *set-format-hex* command (see “set-format-hex” on page 6-67).

If **display** is entered on a line by itself, the current values of the expressions or contents of memory for each item on the display list are printed. To simply see the expressions themselves, use the **info display** command (see “info display” on page 6-167).

Examples:

```
(local) (12) display/x var_name
(local) (12) display/4d 0x1234
```

If these commands are entered, then each time process 12 stops, the value of *var_name* will be printed in hexadecimal on one line, and four words of memory starting at hexadecimal address 1234 will be printed on the next line.

undisplay

Disable an item from the display expression list.

undisplay *item_number* ...

item_number

An item number of an item to be disabled in the list of expressions to be printed each time the program stops, as specified in previous **display** commands (see “display” on page 6-101).

The **undisplay** command disables the given items in each of the processes specified by the qualifier. The associated expressions or memory locations cease to be displayed when the corresponding process stops, until you enable them again using the **redisplay** command (see “redisplay” on page 6-103). The effect of the qualifier on this command is to limit the items to be disabled to only those that occur in the specified processes.

Item numbers prefix each displayed language expression and memory section. The item numbers also may be viewed by entering the **info display** command (see “info display” on page 6-167).

redisplay

Enable a display item.

redisplay *item_number* ...

item_number

An item number of an item to be enabled in the list of expressions to be printed each time the program stops, as specified in previous **display** commands (see “display” on page 6-101).

The **redisplay** command enables the specified display items so that they once again print data when the corresponding process stops. The **redisplay** command reverses the effect of the **undisplay** command. The effect of the qualifier on this command is to limit the items to be enabled to only those that occur in the specified processes.

Item numbers prefix each displayed language expression and memory section. The item numbers also may be viewed by entering the **info display** command (see “info display” on page 6-167).

printf

Print the values of language expressions using a format string.

printf *format-string* [, *expression* ...]

format_string

A string within quotes containing text to be printed and print formats for expressions to be printed.

expression

A language expression (see “Expression Evaluation” on page 3-25).

printf prints user-specified text plus, optionally, values of language expressions evaluated in the currently defined language, for each process specified in the qualifier. This command acts the same as the C language library routine **printf(3)**, with the exception of the ‘%n’ format descriptor. As in that routine, each print format (i.e., substring beginning with ‘%’ and or width specifier ‘*’) in the *format-string* corresponds to one language expression in the specified list. The number of language expressions entered must match the number of print formats.

If a ‘%n’ format descriptor is present in the format string, it is considered a syntax error and the **printf** command is aborted.

Example:

```
(local) (27) printf "The value of var_name = %d.\n", var_name
```

This example prints "The value of var_name = " followed by the decimal value of var_name and a newline, for the process with PID 27.

load

Dynamically load an object file, possibly replacing existing routines.

load *object*

object

The name of an object file to be loaded into the program.

object is subject to object filename translations (see “translate-object-file” on page 6-29).

This command dynamically loads the designated object file into the address space of the running program. If the loaded file contains any routines which are already defined in the program, the entry points of the existing routines are patched to jump directly to the new routines just loaded. If there are any active stack frames for old routines, the return addresses in the stack still point to the old code. New calls made following the **load** will call the new routines.

If you had any breakpoints or other eventpoints set in the old routine, you may need to set equivalent ones again in the new routine (the old ones are still there, but since the old routine will never be called again, you will probably never hit any of them).

The primary purpose of this command is to allow you to replace an existing routine with a new version, avoiding the overhead of forcing you to stop debugging the program, relink it, and rerun to get back to the point of interest.

This command must be used with care. If the new object file contains any global data definitions, you are very likely to wind up with an erroneous program in which old routines refer to the original data locations and new routines refer to the newly loaded data definitions. Patching the old routine entry points to jump to the new routine definitions is simple, but it is not possible to locate all the places that might refer to data items defined in the object file, so loading object files that define static data items is likely to generate unexpected results.

If the object file refers to other routines or external data items that are not already defined in the program file, you are told about the undefined symbols, and the object file is not loaded. If you load an object file that defines new symbols, they are added to the symbol table for the program, so subsequent loads may refer to the new names.

This command checks for obvious problems with the new object file and warns you of anything that is likely to be a mistake, but it loads the new object anyway.

branch-history

Display the branch history if branch tracking has been enabled. See “Branch Tracking” on page 3-38.

branch-history [*number-of-branches*]

Abbreviation: **bh**

number-of-branches

The maximum number of branches to display.

The **branch-history** command displays, for each process specified in the qualifier, a pair of addresses for each branch instruction tracked up to the *number-of-branches* maximum. If no maximum was specified, it defaults to 5 branches. The description for each branch instruction is of the form:

From: *address*
To: *address*

The *address* description always contains a hex address. If available, it will contain function name and source file,line information. If that is not available but symbols are, it will contain *<symbol+offset>* information.

Branch tracking is only supported for RedHawk Linux version 6.0 or later and only on newer Intel chips.

Manipulating Eventpoints

This subsection describes the various commands that are used to set and modify eventpoints.

Some of the commands which operate on breakpoints also operate on patchpoints, tracepoints, monitorpoints, heappoints, watchpoints and syscallpoints as well. The following table indicates which types of eventpoints may be affected by which commands:

Table 6-8. Eventpoint Commands

Command Name	What the Command May Apply To						
	Breakpoints	Patchpoints	Tracepoints	Monitorpoints	Watchpoints	Heappoints	Syscallpoints
name	X	X	X	X	X	X	X
clear	X	X	X	X		X	
commands	X			X	X		X
condition	X	X	X	X	X	X	X
delete	X	X	X	X	X	X	X
disable	X	X	X	X	X	X	X
enable	X	X	X	X	X	X	X
ignore	X	X	X	X	X	X	X
modify-eventpoint	X	X	X	X	X	X	X
tbreak	X						
tpatch		X					

Eventpoint Modifiers

An *eventpoint modifier* modifies the setting of eventpoints in a program.

The modifiers come after the eventpoint commands as follows:

command [*modifier ...*]

The eventpoint modifiers are:

/delete

Causes the eventpoint to be deleted after the first hit. This eventpoint modifier is valid only with the **enable** command (see “enable” on page 6-123).

/disabled

Causes the eventpoint to be created in a disabled state. You must use the **enable** command to activate the eventpoint (see “enable” on page 6-123).

/f

Causes any line numbers appearing in a location specifier for this eventpoint to be interpreted as *fixed*. The eventpoint will not be inserted at a later point if the specified line has no code. For example, by default, if an eventpoint was set at line 12, but no code was present at line 12, the eventpoint might be inserted at line 14 if there was code there. This is prevented by use of the */f* modifier.

name

Give a name to a group of eventpoints.

name [/add] *name* [[-] *eventpoint-spec*] . . .

/add

Add the eventpoints to the named set, rather than redefining the set.

name

The name of the set of eventpoints to be defined. This must not be the same as the name of any dialogue you currently have, or of any process family that is currently defined. The name must consist only of alphanumeric characters and underscores and must begin with an alphabetic character. The name may be of arbitrary length.

eventpoint-spec

An eventpoint specifier. See “Eventpoint Specifiers” on page 6-18.

The total set of eventpoints is accumulated by scanning the *eventpoint-spec* arguments left to right. An argument is added to the set unless it is preceded by a ‘-’, in which case it is subtracted from the set accumulated so far.

If no *eventpoint-spec* is given, then this command removes any previous definition of *name*.

Any qualifier applied to this command has the effect of restricting the set of eventpoints named to those which exist in the processes specified by the qualifier.

Examples:

```
(local) name evpt1 12 25 18
(local) name evpt2 evpt1 99
(local) name evpt1 evpt1 16
```

The first command gives the name *evpt1* to three eventpoints identified by eventpoints 12, 18, and 25. The second command gives the name *evpt2* to the three eventpoints in *evpt1* plus eventpoint 99. The third command extends the definition of *evpt1* to include eventpoint 16; thus *evpt1* is a synonym for four eventpoints: 12, 16, 18, and 25. Note that extending *evpt1* has no effect on *evpt2*, which still consists of eventpoints 12, 18, 25, and 99.

Using the names defined in the previous examples, the use of a minus sign on arguments can be illustrated by the following examples:

```
(local) name evpt3 evpt1 evpt2 -12
(local) name evpt3 evpt1 -12 evpt2
```

The first command defines *evpt3* to be the eventpoints 16, 18, 25, and 99. In contrast, the second command defines *evpt3* to be the eventpoints 12, 16, 18, 25, and 99. In this case, the argument *-12* removed eventpoint 12 from the set accumulated from *evpt1*, but the *evpt2* argument adds that eventpoint back in.

breakpoint

Set a breakpoint.

```
breakpoint [eventpoint-modifier] [/cuda|/process] [/protected]
              [name=breakpoint-name] [[at] location-spec]
              [if conditional-expression]
```

Abbreviation: **b**

eventpoint-modifier

Specifies the breakpoint modifier. See “Eventpoint Modifiers” on page 6-107.

/cuda or /process

The use of */cuda* insists that the breakpoint will apply only to CUDA code. Similarly, the use of */process* insists that the breakpoint will apply only to host (non-CUDA) code. If neither is specified, then the breakpoint will apply to CUDA code if there is CUDA code at the specified location; otherwise, it will apply to host code.

/protected

By default, if a breakpoint is hit, in addition to the thread which hit the breakpoint, NightView will stop all non-protected threads (see “Stopping” on page 3-44). This option overrides this behavior to include all threads, regardless of whether or not they are protected.

name=breakpoint-name

Gives a name to the breakpoint for later reference. (see “name” on page 6-107). If *breakpoint-name* is already defined, then this command adds the newly created breakpoints to the list of eventpoints associated with the name.

location-spec

Specifies the breakpoint location. (see “Location Specifiers” on page 6-15).

If *location-spec* is omitted, then the location used is the next instruction to be executed in the current stack frame. See “Current Frame” on page 3-29.

if conditional-expression

Specifies a breakpoint condition. The language and scope of the expression is determined by the location at which the breakpoint is set (see “Scope” on page 3-28 and “Context” on page 3-28). See also “Expression Evaluation” on page 3-25.

NOTE

The `at`, `if`, and `name` keywords may not be abbreviated in this command.

breakpoint sets a breakpoint in each of the processes specified by the qualifier. This causes the program to suspend execution at the breakpoint location. An optional condition may be applied to the breakpoint which causes execution to be suspended only if the condition evaluates to TRUE. The conditional expression is evaluated in the user program when the breakpoint location is reached (unless the breakpoint is currently being ignored, see “ignore” on page 6-124).

If more than one breakpoint is set (through the use of more than one process in the qualifier) then each breakpoint in each process is assigned a unique breakpoint number.

You can specify debugger commands to be executed when a breakpoint is hit. See “commands” on page 6-120.

It is possible (and sometimes useful) to set more than one breakpoint at the same location in a process. Perhaps you have two breakpoints set at the same place and each has its own set of commands. By enabling only one of the two breakpoints at a time, you can effectively toggle the set of commands that gets executed when the process reaches that location.

If more than one breakpoint is set at the same location in a given process, then the oldest breakpoint with an ignore count of zero and a condition that evaluates to TRUE will be the first breakpoint responsible for stopping the process. After this breakpoint has stopped

the process, before continuing on to the next instruction, NightView will check for any remaining breakpoints at that location which may stop the process. If there are any, then the process will stop at least once more (at the same location) before continuing on to the next instruction.

Example:

```
(local) (441 115) break name=loop sort.c:42
```

This example sets two breakpoints at line 42 of the file named **sort.c** and associates both breakpoints with the name 'loop'. One of the breakpoints is set in process 441 and the other breakpoint is set in process 115. Each of the two breakpoints is assigned a unique breakpoint number.

patchpoint

Patchpoints can be used to:

- Install a small patch to a routine.
- Insert an expression in the program which modifies the value of a variable or register or invokes a function.
- Insert a branch in the program.
- Set a thread-specific tag value

```
patchpoint [eventpoint-modifier] [name=patchpoint-name]
              [[at] location-spec] eval expression
```

```
patchpoint [eventpoint-modifier] [name=patchpoint-name]
              [[at] location-spec] goto location-spec
```

```
patchpoint [eventpoint-modifier] [name=patchpoint-name]
              [[at] location-spec] tag tag-assignments
```

eventpoint-modifier

Specifies the patchpoint modifier. See “Eventpoint Modifiers” on page 6-107.

name=patchpoint-name

Patchpoints are assigned event numbers, and the **name=** syntax as well as the **name** command (see “name” on page 6-107) may be used to give them names. See “Manipulating Eventpoints” on page 6-106.

at location-spec

Specify the exact point in the program to execute the patchpoint (see “Location Specifiers” on page 6-15). The patchpoint is executed immediately prior to any existing code at this location.

If *location-spec* is omitted, then the location used is the next instruction to be executed in the current stack frame. See “Current Frame” on page 3-29.

`eval expression`

This variant of the **patchpoint** command specifies an expression to insert in the program at the designated *location-spec*. C and C++ programmers should note that this is an expression and not a statement; therefore, it does *not* end with a semicolon. (The concept of *expression* is extended to include assignments and procedure calls in Fortran.) See “Expression Evaluation” on page 3-25.

`goto location-spec`

This variant of the **patchpoint** command specifies a location to branch to when the program reaches the point of the patchpoint. The instruction originally at the patchpoint location will not be executed.

Note that if an expression is used as a *location-spec*, the expression is evaluated only once for each process in the qualifier. For example, if the *location-spec* is `*$lr`, the value of register `lr` in the *current context* is used as the location to branch to.

`tags tag-assignment [, tag-assignment...]`

This variant of the **patchpoint** command creates a patchpoint that sets one or more tag values in `$thr` for any thread that executes through the patchpoint (see “Thread Tags” on page 3-45 for more information on `$thr`).

Following the **tags** keyword, one or more comma separated *tag-assignments* should appear. A *tag-assignment* has the following syntax:

tagname assign-op expression

- where, *tagname* is a simple identifier that names the tag value to be set. It names the member within the `$thr` struct.
- where *assign-op* may be one of `=`, `:=`, `-=`, `+=`, `*=`, `/=`. The `=` and `:=` operators are identical and represent simple assignment. The `-=`, `+=`, `*=`, and `/=` operators are like the C operators (but in a **tags** patch will work even for Ada and Fortran expression).
- where *expression* is the expression to evaluate and use in the assignment to the tag. Unless the *tagname* has been previously declared using the **declare-thread-tag** command, the expression must be a 1-bit boolean value (`true`, `false`, `1`, or `0`).

Additionally, a *tag-assignment* will function as expected if the *tagname* is declared as an array of `char` (see “declare-thread-tag” on page 6-154), the *assign-op* is a simple assignment (`=` or `:=`), and the *expression* is of type `char*` or something similar: it assigns the character contents of the *expression* to *tagname*.

NOTE

The keywords `name`, `at`, `eval`, and `goto` and `tags` may not be abbreviated in this command.

Once an `eval` patchpoint is installed, the language expression will be executed each time control reaches *location-spec* in the program. After the patchpoint is executed, the original instruction will also execute.

Once a `goto` patchpoint is installed, the branch will be executed before the patched instruction each time execution reaches *location-spec* in the program. It is important to note that the original instruction is not executed if the patchpoint is *hit* (that is, depending on the enabled status, the ignore count and any eventpoint condition on the patchpoint). If the patchpoint is not hit, the original instruction is executed normally.

When patching in a `goto`, you should be aware that the compiler has probably generated code which expects certain register contents and altering the flow of control in your program can very easily send it to a new location with unexpected values in registers, so the `goto` patchpoint should be used only when you are sure you know all the consequences.

You may attach a condition or ignore count to both kinds of patchpoints, using the **condition** (see “condition” on page 6-121) or **ignore** (see “ignore” on page 6-124) commands. This suppresses execution of the patched expression unless the ignore count is zero and the conditional expression evaluates to TRUE.

Patchpoints are implemented by modifying the executable code for the program, so they will remain in effect until the program exits, even if you **detach** the debugger from the program, unless the patchpoint was disabled when you detached (see “detach” on page 6-41 and “disable” on page 6-122). Note that the disk copy of the program is not modified; you must edit your source, recompile and relink to make a permanent modification to the program.

If multiple patchpoints are made at the same point in the program, they will all be executed in the order they were applied. This is especially important to note for **goto** patchpoints, because once a **goto** is executed, any subsequent patchpoints (or other kinds of eventpoints, such as breakpoints and tracepoints) at that same location will not be executed. If a **goto** patchpoint is not hit (because it was disabled, or the ignore count or condition caused it to be skipped), then the branch will not be taken and subsequent patchpoints will be executed, as well as the original patched instruction.

Example:

```
(local) patchpoint file.c:12 eval i=0
```

This C example patches the code to initialize the variable `i` to zero immediately prior to executing line 12 in the file `file.c`. Note that no semicolon appears in this example.

Example:

```
(local) patchpoint file.c:12 tags working=true
```

This C example patches the code to define the `$tags` field *working* the value *true* for each thread that executes line 12 of `file.c`.

Patchpoints are not supported in CUDA code. This is a technical limitation of the CUDA

driver.

set-trace

The **set-trace** command is used to specify information that may be useful before any tracepoints are set in a process (see “tracepoint” on page 6-113).

set-trace [*eventmap=event-map-file*] [*tracefile=key-filename*]

eventmap=event-map-file

Names the file that contains the mapping between symbolic trace-event tags and numeric trace-event IDs. This should be the same as the event-map file passed to **ntrace(1)**.

If you want to use symbolic trace-event tags rather than numeric trace-event IDs as the *event-id* parameter of the **tracepoint** command, then you must specify an event-map file. You may specify multiple event-map files by repeating the *eventmap* parameter. As long as the files do not contain conflicting definitions for tags, all the tags will be defined for use as trace-event identifiers

tracefile=key-filename

This argument causes NightView to initiate application tracing by issuing the NightTrace API `trace_begin` call. **tracepoints** aren't useful if the application has not initialized the NightTrace API subsystem. If your application already calls `trace_begin` you should not use this argument.

The *key-filename* is passed as the first parameter to `trace_begin` (in the form of a string constant). The *key-filename* identifies the daemon session which collects the NightTrace events that are logged by the application and by NightView **tracepoint** eventpoints. See **ntraceud(1)** and **trace_begin(3)** for more information on NightTrace.

To specify a file on a remote system, use the form *user@host:key-filename*. See “Remote File Access” on page 3-7.

NightView will automatically linking in the required NightTrace API library modules if they don't already exist in the program.

tracepoint

Set a tracepoint.

tracepoint [*eventpoint-modifier*] *event-id* [*name=tracepoint-name*]
 [[*at*] *location-spec*]
 [*value=logged-expression* [, *logged-expression*...]]
 [*if conditional-expression*]

eventpoint-modifier

Specifies the tracepoint modifier. See “Eventpoint Modifiers” on page 6-107.

event-id

An identifier for the trace event to be traced by NightTrace. This is either a numeric trace-event ID or a symbolic trace-event tag obtained from the event-map file specified by the *eventmap* parameter of the **set-trace** command (see “set-trace” on page 6-113).

name=tracepoint-name

Gives a name to the tracepoint for later reference. See “name” on page 6-107. If *tracepoint-name* is already defined, then this command adds the newly created tracepoints to the list of eventpoints associated with the name.

location-spec

Specifies the tracepoint location. See “Location Specifiers” on page 6-15.

If *location-spec* is omitted, then the location used is the next instruction to be executed in the current stack frame.

value=logged-expression [, logged-expression...]

Specifies that the value of each *logged-expression* should be recorded with the trace event. The expressions are separated by commas. (To include a comma in an expression, surround the expression by parentheses.)

The number of expressions and the type of the expressions must match a *trace_event* routine defined by the **ntrace** library. See the section “*trace_event()* and Its Variants” in the “Using the NightTrace Logging API” chapter of the *NightTrace User's Guide*.

The expressions are evaluated in the user program, so they obey the same rules that conditional and patchpoint expressions do. See “Expression Evaluation” on page 3-25.

if conditional-expression

Specifies a tracepoint condition. The language and scope of the expression is determined by the location at which the tracepoint is set (see “Scope” on page 3-28 and “Context” on page 3-28). See also “Expression Evaluation” on page 3-25.

NOTE

The *name*, *value*, and *if* keywords may not be abbreviated in this command.

The **tracepoint** command sets a tracepoint in each of the processes specified by the qualifier. This causes the program to emit special tracing output at the tracepoint location. An optional condition may be applied to the tracepoint which causes tracing to be

performed only if the condition evaluates to TRUE. The conditional expression *conditional-expression* is evaluated in the user program when the tracepoint location is reached (unless the tracepoint is currently being ignored, see “ignore” on page 6-124).

Tracepoints set in a process remain set even if you **detach** the debugger from the program, unless the tracepoint was disabled at the time you detached (See “detach” on page 6-41 and “disable” on page 6-122).

NOTE

The debugger does *not* start a NightTrace daemon, which is required in order to actually collect the events logged with tracepoints. You must do that using **ntraceud(1)** or the Daemons panel in **ntrace(1)** (see “NightTrace Daemon” on page 3-50).

If more than one tracepoint is set (through the use of more than one process in the qualifier) then each tracepoint in each process is assigned a unique tracepoint eventpoint number, but all instances will share the same *event-id* you specified on the command.

It is possible (and sometimes useful) to set more than one tracepoint at the same location in a process. Perhaps there is more than one noteworthy event that takes place at the same location in your program. If more than one tracepoint is set at the same location in a given process, then the tracepoints at that location are recorded in the order they were defined.

Example:

```
(local) (441 115) tracepoint 27 name=loop_trace sort.c:42
```

This example sets two tracepoints at line 42 of the file named **sort.c** and associates both tracepoints with the name 'loop_trace'. One of the tracepoints is set in process 441 and the other tracepoint is set in process 115. Each of the two tracepoints is assigned a unique tracepoint number. The ID of the trace event to trace is given by the number 27.

Tracepoints are not supported in CUDA code. This is a technical limitation of the CUDA driver.

monitorpoint

Monitor the values of one or more expressions at a given location.

```
monitorpoint [eventpoint-modifier] [name=monitorpoint-name]
               [[at] location-spec]
```

eventpoint-modifier

Specifies the monitorpoint modifier. See “Eventpoint Modifiers” on page 6-107.

name=monitorpoint-name

Gives a name to the monitorpoint for later reference. See “name” on page 6-107. If *monitorpoint-name* is already defined then this command adds the newly created monitorpoints to the list of eventpoints associated with the name.

location-spec

Specifies the monitorpoint location. See “Location Specifiers” on page 6-15.

If *location-spec* is omitted, then the location used is the next instruction to be executed in the current stack frame.

The **monitorpoint** command sets a monitorpoint in each of the processes specified by the qualifier. Each line following the **monitorpoint** command must be a special form (described later) of **print** command; each **print** command specifies an expression to be evaluated and monitored at the location of the monitorpoint. To end the list of **print** commands, type **end monitor** on a line by itself.

In the command-line and simple full-screen interfaces, the prompt changes to > while you are entering the attached **print** commands. See “Command Syntax” on page 6-1.

When the monitorpoint is executed, the expressions specified in the attached commands will be evaluated and their values saved in a location reserved by NightView. The monitored values are displayed periodically in a monitor display area; see “Monitor Window” on page 3-32. For a more detailed description of monitorpoints, see “Monitorpoints” on page 3-12.

The syntax of the commands attached to a monitorpoint is:

```
print [/print-format-code] [id="string"] expression
```

This syntax is identical to the **print** NightView command (see “print” on page 6-91), with the addition of the optional *id="string"* argument. The *string*, if specified, is used to identify the monitored expression in the monitor display area. If you do not specify the *id=* parameter, the text of the expression itself is used as the identifying string. Note that you may not abbreviate the *id=* keyword to anything shorter (like “i”).

Once you have created a monitorpoint, you can change the set of commands attached to it (and thus the expressions being monitored) using the **commands** command. See “commands” on page 6-120.

Example:

```
(local) monitorpoint file.c:12
> print variable1
> print id="Velocity (ft/sec)" variable2
> end monitor
```

In this example, two variables will be monitored at line 12 of **file.c**. The first variable, *variable1*, will be displayed using its name as the identifying string. The second variable, *variable2*, will be displayed with the string *Velocity (ft/sec)*.

Monitorpoints are not supported in CUDA code. This is a technical limitation of the CUDA driver.

heappoint

Check the heap for errors, or change the heap debugger settings, at a given location.

heappoint [*eventpoint-modifier*] [*name=heappoint-name*] [[*at*] *location-spec*]
 [{*check* | *debug* *parameters*}] [*if conditional-expression*]

eventpoint-modifier

Specifies the heappoint modifier. See “Eventpoint Modifiers” on page 6-107.

name=heappoint-name

Gives a name to the heappoint for later reference. See “name” on page 6-107. If *heappoint-name* is already defined then this command adds the newly created heappoints to the list of eventpoints associated with the name.

location-spec

Specifies the heappoint location. See “Location Specifiers” on page 6-15.

If *location-spec* is omitted, then the location used is the next instruction to be executed in the current stack frame.

check

Specifies that the heappoint does a heap check. This is the default if neither *check* nor *debug* is specified.

debug parameters

Specifies that the heappoint changes the heap debugger settings. *parameters* are the same as the arguments to the **heapdebug** command (see “heapdebug” on page 6-56).

if conditional-expression

Specifies a heappoint condition. The language and scope of the expression is determined by the location at which the heappoint is set (see “Scope” on page 3-28 and “Context” on page 3-28). See also “Expression Evaluation” on page 3-25.

The **heappoint** command sets a heappoint in each of the processes specified by the qualifier. See “Heappoints” on page 3-14.

When the heappoint is executed, the process does a heap check if *check* was specified, or changes the heap debugger settings if *debug* was specified. The *check* and *debug* parameters are mutually exclusive.

Putting *check* heappoints at various places in your program can help you narrow down where heap problems are occurring.

Changing the heap debugger settings dynamically within your program can help you get reasonable performance while still getting strong heap checking. For example, if you have a suspicious section of code, you could set a heappoint at the beginning of the section to set automatic heap checks to occur before every heap operation, and set another heappoint

at the end of the section to set automatic heap checks to occur only every 10,000 heap operations.

Heappoints are not supported in CUDA code. This is a technical limitation of the CUDA driver.

mcontrol

Control the monitor display window.

```
mcontrol {display | nodisplay} [monitorpoint-spec ...]
```

Turn on or off the display of individual monitorpoints in the monitor window.

```
mcontrol delay milliseconds
```

Set the milliseconds to delay between monitor window updates.

```
mcontrol {off | on | stale | nostale | hold | release}
```

Toggle a monitoring parameter.

Abbreviation: **hold**

This is an abbreviation for **mcontrol hold**.

Abbreviation: **release**

This is an abbreviation for **mcontrol release**.

```
display nodisplay
```

These keywords are used to enable or disable the display of specific monitorpoints in the monitor window. The monitorpoints appearing in the argument and in the processes specified by the qualifier are either added to or removed from the monitor window display area. This does not affect the monitorpoint itself, it simply determines which monitorpoints are shown in the window. See “monitorpoint” on page 6-115.

```
on off
```

These keywords turn the monitor window on or off. You may wish to turn off the monitor window to reclaim screen space, then turn it back on later. Turning off the window also does a **hold**, but turning the window on does not implicitly do a **release**.

```
stale nostale
```

The monitor window normally displays a stale data indication next to each value. The **nostale** keyword causes the monitor window to display blank space rather than one of the stale data indicators. The indicators may be turned back on with the **stale** keyword.

`hold release`

The `hold` and `release` keywords are used to hold or release updates of the monitor window. When the window is held, the values displayed in the monitor window will no longer change (the processes containing the values are not affected, they continue to run). The `release` keyword allows the monitor window to start updating the values again.

Interrupting the debugger implicitly causes the Monitor Window to stop updating. See “Interrupting the Debugger” on page 3-40.

`delay`

The monitor window normally waits one second (1000 milliseconds) between updates. A different number of milliseconds may be specified following the `delay` keyword. If you tell it to wait zero milliseconds, it updates the monitor window as fast as it possibly can.

All of the **mcontrol** parameters allow you to control various aspects of the monitor display window (see “Monitor Window” on page 3-32).

You may not combine parameters on the **mcontrol** command. Only one keyword may be used in one invocation of the command. The command qualifier is only used when the `display` or `nodisplay` keywords are used to specify a list of monitorpoints.

clear

Clear all eventpoints at a given location.

clear `[ [at] location-spec ]`

location-spec

Specifies the location from which all eventpoints are to be removed. See “Location Specifiers” on page 6-15.

If *location-spec* is omitted, then the location used is the next instruction to be executed in the current stack frame. See “Current Frame” on page 3-29.

clear removes all eventpoints at the specified location in each process. Once an eventpoint has served its purpose, the eventpoint may be removed by using the **clear** or **delete** commands (see “delete” on page 6-122). Both commands remove an eventpoint. **clear** removes eventpoints based on where they are in the process. **delete** removes eventpoints specified by name or by eventpoint-number.

NOTE

A location specifier may sometimes designate multiple locations (see “Location Specifiers” on page 6-15). Hence, it is possible for a single eventpoint to be set at multiple locations. If any of the locations at which an eventpoint is set match any of the locations implied by the location specifier for the **clear** command, then that eventpoint will be removed (from *all* of its corresponding locations).

It is unnecessary to clear a breakpoint in order to continue execution after the breakpoint has stopped the program.

Example:

```
(local) clear sort.c:42
```

This example removes all eventpoints set at line 42 of the file named **sort.c** in each of the processes specified by the default qualifier.

commands

Attach commands to a breakpoint, monitorpoint, or watchpoint.

commands *eventpoint-spec*

eventpoint-spec

The breakpoints, monitorpoints, or watchpoints to which the given commands are attached. See “Eventpoint Specifiers” on page 6-18.

The **commands** command attaches the given list of commands to the given breakpoints, monitorpoints, or watchpoints in processes specified by the qualifier. Each line following the **commands** command-line should be a command to associate with the eventpoints. To end the list of commands, type 'end' on a line by itself.

Each of the commands given is implicitly qualified with the PID of the process associated with the eventpoint.

In the command-line and simple full-screen interfaces, the prompt changes to > while you are entering this command. See “Command Syntax” on page 6-1.

If the first line given is 'silent', then the usual message that is printed when a breakpoint or watchpoint stops the process will be suppressed. Furthermore, the 'silent' command will also prevent the current source line from being listed, and will prevent any displays from being updated. The 'silent' command is valid only when attached to a breakpoint or watchpoint and is useful for breakpoints or watchpoints that are intended only to print a specific message and then resume execution.

Certain commands (such as **continue**, **resume**, and **signal**), once executed, will automatically terminate the command stream associated with a set of commands that were attached to a breakpoint or watchpoint using the **commands** command. See

“continue” on page 6-132, “resume” on page 6-133, and “signal” on page 6-142.

Although you can use the **commands** command to attach commands to breakpoints, monitorpoints, or watchpoints, the eventpoints specified on the command line must be all of the same type. Also note that the commands allowed for monitorpoints are restricted to **print** commands. See “monitorpoint” on page 6-115.

condition

Attach a condition to an eventpoint.

condition *eventpoint-spec* [*conditional-expression*]

eventpoint-spec

The eventpoints associated with the condition. See “Eventpoint Specifiers” on page 6-18.

conditional-expression

The condition to be associated with the eventpoints. See “Expression Evaluation” on page 3-25.

The simplest type of breakpoint is one which stops the program each time it is encountered (an *unconditional breakpoint*). Often however, you may wish to stop the program at a given location only after a certain event has occurred or when a specified condition has been met (a *conditional breakpoint*). The **condition** command may be used to attach a condition to a breakpoint.

In a similar manner, conditions may also be attached to tracepoints, monitorpoints, heappoints, patchpoints, watchpoints, and syscallpoints, causing the associated action to take effect only when the attached condition evaluates to TRUE.

The **condition** command attaches the condition *conditional-expression* to one or more eventpoints in the processes specified by the qualifier. If *conditional-expression* is omitted, then any condition attached to the specified eventpoint is removed in each of the processes specified by the qualifier, and the eventpoint becomes an unconditional one. If the specified eventpoint already has a condition attached to it, the existing condition is replaced with *conditional-expression*.

Examples:

```
(local) breakpoint name=loop at foo.c:12
(local) condition loop (index == 0)
(local) condition loop
```

The first **condition** command attaches a condition to the breakpoint named 'loop' so that it only stops the program when the variable 'index' is zero. The second **condition** command removes any condition associated with the breakpoint named 'loop' (thus making it an unconditional breakpoint).

```
(local) trace MyEvent name=tracel at foo.c:12
(local) condition tracel (x>12)
```

In this example, a tracepoint named 'trace1' is set, and the condition 'x>12' is attached to the tracepoint. Therefore, the event will be traced only when 'x' is greater than 12.

delete

Delete an eventpoint.

delete [*eventpoint-spec* ...]

Abbreviation: **d**

eventpoint-spec

The eventpoints to be deleted. See “Eventpoint Specifiers” on page 6-18.

delete removes the specified eventpoints in each of the processes specified by the qualifier. Both **delete** and **clear** may be used to delete eventpoints (see “clear” on page 6-119). The difference is that **delete** removes eventpoints specified by name or by eventpoint-number and **clear** removes eventpoints specified by location.

If *eventpoint-spec* is omitted and your safety level is `unsafe` then *all* eventpoints in the processes specified by the qualifier are removed (see “set-safety” on page 6-67). If *eventpoint-spec* is omitted and your safety level is `verify`, then you are prompted for confirmation before the eventpoints are removed (see “Replying to Debugger Questions” on page 6-22). If *eventpoint-spec* is omitted and your safety level is `forbid` then no eventpoints are removed.

The effect of the qualifier on this command is to limit the eventpoints deleted to be only those that occur in the processes specified by the qualifier.

Examples:

```
(local) d loop
(local) d 2 5
```

The first example removes all eventpoints associated with the name 'loop'. The second example removes eventpoints 2 and 5.

disable

Disable an eventpoint.

disable [*eventpoint-spec* ...]

eventpoint-spec

The eventpoints to be disabled. See “Eventpoint Specifiers” on page 6-18.

The **disable** command disables the given eventpoints in each of the processes specified by the qualifier. Disabling an eventpoint is not quite the same as removing an eventpoint. When an eventpoint is removed, it is made inoperative and all the information associated

with the eventpoint is removed. When an eventpoint is disabled, it is simply made inoperative. It may still be seen, however, if you use the **info eventpoint** command (see “info eventpoint” on page 6-158). All information associated with the eventpoint is still retained so that the eventpoint may later be reactivated using the **enable** command (see “enable” on page 6-123).

If *eventpoint-spec* is omitted and your safety level is `unsafe` then *all* eventpoints in the processes specified by the qualifier are disabled (see “set-safety” on page 6-67). If *eventpoint-spec* is omitted and your safety level is `verify`, then you are prompted for confirmation before the eventpoints are disabled (see “Replying to Debugger Questions” on page 6-22). If *eventpoint-spec* is omitted and your safety level is `forbid` then no eventpoints are disabled.

The effect of the qualifier on this command is to limit the eventpoints disabled to be only those that occur in the processes specified by the qualifier.

Example:

```
(local) disable 4
(local) (115 441) disable calvin
(local) (549) disable 8 hobbes 12 14
```

The first example disables eventpoint number 4 in the processes specified by the default qualifier. The second example disables the eventpoints associated with the name ‘calvin’ in process 115 and in process 441. The third example disables the eventpoints associated with the name ‘hobbes’ and disables eventpoints numbered 8, 12, and 14 in process 549.

enable

Enable an eventpoint for a specified duration.

enable [/once|/delete] [*eventpoint-spec* ...]

/once

Specify whether the given eventpoints are to be enabled once only and then immediately disabled after the next time they are hit. There need not be a space between the command name and the ‘/’.

/delete

Valid only for breakpoints and watchpoints. Specify whether the given breakpoints and watchpoints are to be enabled once only and then immediately deleted after the next time they are executed. There need not be a space between the command name and the ‘/’.

eventpoint-spec

The eventpoints to be enabled. See “Eventpoint Specifiers” on page 6-18.

The **enable** command enables for the specified duration each of the eventpoints in the processes specified by the qualifier. If neither `/once` nor `/delete` is specified, then the given eventpoints are simply enabled. If `/once` is specified, then the given eventpoints

are temporarily enabled. The eventpoints will be disabled again after the next time they are hit. If `/delete` is specified, then for each process in the qualifier, the given breakpoints and watchpoints are enabled and also marked for deletion. The breakpoints and watchpoints will be deleted after the next time they are hit.

If *eventpoint-spec* is omitted and your safety level is `unsafe` then *all* eventpoints in the processes specified by the qualifier are enabled (see “set-safety” on page 6-67). If *eventpoint-spec* is omitted and your safety level is `verify`, then you are prompted for confirmation before the eventpoints are enabled (see “Replying to Debugger Questions” on page 6-22). If *eventpoint-spec* is omitted and your safety level is `forbid` then no eventpoints are enabled.

The effect of the qualifier on this command is to limit the eventpoints enabled to be only those that occur in the processes specified by the qualifier.

Examples:

```
(local) enable calvin
(local) enable /once 4 6 23
(local) enable /delete 8 hobbes
```

The first example enables all eventpoints associated with the name 'calvin' in the default qualifier. The second example enables eventpoints number 4, 6, and 23 for once-only execution (the eventpoints will be disabled after the next time they are hit). The third example enables breakpoint number 8, and the breakpoints and watchpoints associated with the name 'hobbes' for deletion (these breakpoints and watchpoints will be deleted after the next time they are hit).

ignore

Attach an ignore-count to an eventpoint.

ignore *eventpoint-spec* *count*

eventpoint-spec

The eventpoints to be ignored. See “Eventpoints” on page 3-9.

count

The number of times to ignore the eventpoint. Specifying an ignore-count of zero has the effect of causing the eventpoints to no longer be ignored. The ignore-count is evaluated in the user's process.

The **ignore** command causes the specified eventpoints to be skipped the next *count* times execution reaches them (even if the eventpoint is a conditional eventpoint). This is accomplished by attaching an ignore-count to the given eventpoints. In the case of a breakpoint, any NightView commands associated with the breakpoint will not be executed until the breakpoint is hit.

Example:

```
(local) ignore calvin 4
```

This example causes the eventpoints associated with the name 'calvin' to be ignored 4 times before they may be hit again.

modify-eventpoint

Change the hit count, crossing count, or protected flag for an eventpoint.

```
modify-eventpoint eventpoint-spec [hits=hits] [crossings=crossings]
                  [protected={on|off}]
```

eventpoint-spec

The eventpoints to be modified. See “Eventpoints” on page 3-9.

hits=*hits*

Specifies that the hit count of the specified eventpoint should be changed to *hits*. See “Interactions Between Conditions, Ignore Counts, etc.” on page 3-12 for details.

crossings=*crossings*

Specifies that the crossing count of the specified eventpoint should be changed to *crossings*. See “Interactions Between Conditions, Ignore Counts, etc.” on page 3-12 for details.

protected={on|off}

Specifies that the protected flag for the given eventpoint should be set to `on` or `off`. This parameter can be specified only for breakpoints. See “Protected thread tag” on page 3-46 for details.

Note that the **modify-eventpoint** command cannot be applied to a futurepoint before its location has been discovered and the eventpoint has been inserted into the process. (see “Eventpoints” on page 3-9).

tbreak

Set a temporary breakpoint.

```
tbreak [name=breakpoint-name] [[at] location-spec]
        [if conditional-expression]
```

name=*breakpoint-name*

Gives a name to the breakpoint for later reference. See “name” on page 6-107. If *breakpoint-name* is already defined then this command adds the newly created breakpoints to the list of eventpoints associated with the name.

location-spec

Specifies the breakpoint location. See “Location Specifiers” on page 6-15.

if conditional-expression

Specifies an eventpoint condition. The language and scope of the expression is determined by the location at which the breakpoint is set (see “Scope” on page 3-28 and “Context” on page 3-28). See “Expression Evaluation” on page 3-25.

Note: The `at`, `if`, and `name` keywords may not be abbreviated in this command.

Like the **breakpoint** command (see “breakpoint” on page 6-108), the **tbreak** command sets a breakpoint. The difference between the two is that **tbreak** sets a one-time-only breakpoint in each of the processes specified by the qualifier. The breakpoint will be disabled after being hit once.

Example:

```
(local) (115) tbreak sort.c:48
```

This example sets a temporary breakpoint in process 115 at line 48 of the source file **sort.c**.

tpatch

Set a patchpoint that will execute only once.

tpatch [*name=patchpoint-name*] [[*at*] *location-spec*] *eval expression*

Insert an expression in the program that will be executed the next time the patchpoint is hit, then never executed again unless explicitly enabled. See “enable” on page 6-123.

tpatch [*name=patchpoint-name*] [[*at*] *location-spec*] *goto location-spec*

Overwrite an instruction in the program with a branch that will only be taken once. Subsequent execution will ignore the patchpoint and execute the original instruction.

name= patchpoint-name

Patchpoints are assigned event numbers, and the *name=* syntax as well as the **name** command (see “name” on page 6-107) may be used to give them names. See “Manipulating Eventpoints” on page 6-106.

at location-spec

Specify the exact point in the program to execute the patchpoint. See “Location Specifiers” on page 6-15. The patchpoint is executed immediately prior to any existing code at this location.

If *location-spec* is omitted, then the location used is the next instruction to be executed in the current stack frame. See “Current Frame” on page 3-29.

eval expression

This variant of the **patchpoint** command specifies an expression to insert in the program at the designated *location-spec*. C and C++ programmers

should note that this is an expression and not a statement; therefore, it does *not* end with a semicolon. (The concept of *expression* is extended to include assignments and procedure calls in Fortran.) See “Expression Evaluation” on page 3-25.

`goto location-spec`

This variant of the **patchpoint** command specifies a location to branch to when the program reaches the point of the patchpoint. The instruction originally at the patchpoint location will not be executed.

NOTE

The keywords `name`, `at`, `eval`, and `goto` may not be abbreviated in this command.

The **tpatch** command is a variant of the **patchpoint** command. See “patchpoint” on page 6-110. It works exactly like the patchpoint command, but a temporary patchpoint will automatically disable itself after executing one time. A temporary patchpoint may be enabled later, in which case it will act exactly like a normal patchpoint. See “enable” on page 6-123.

A temporary patchpoint may be useful for patching in initialization code which should only execute once.

watchpoint

Set a watchpoint.

```
watchpoint [eventpoint-modifier] [/once] [/read] [/write]
             [name=watchpoint-name] [at] lvalue
             [if conditional-expression]
```

```
watchpoint [eventpoint-modifier] [/once] [/read] [/write] /address
             [name=watchpoint-name] [at] address-expression
             {size size-expression | type expression}
             [if conditional-expression]
```

eventpoint-modifier

Specifies the watchpoint modifier. See “Eventpoint Modifiers” on page 6-107.

`/once`

The watchpoint is enabled only until the first time it is hit.

`/read`

Watchpoint processing occurs for a read (i.e., a “load”) of the specified address. Either or both of `/read` and `/write` may be specified.

`/write`

Watchpoint processing occurs for a write (i.e., a "store") of the specified address. Either or both of `/read` and `/write` may be specified. If neither is specified, the default is `/write`.

Watchpoint processing always occurs for a write, even if `/write` is omitted, because it is not possible to create a read-only watchpoint on an IA-32 or AMD64.

`/address`

Indicates this is the *address-expression* form of the command.

`name=watchpoint-name`

Gives a name to the watchpoint for later reference. (see "name" on page 6-107). If *watchpoint-name* is already defined, then this command adds the newly created watchpoints to the list of eventpoints associated with the name.

lvalue

An expression that yields an addressable item to watch. For example, *lvalue* may be a variable name or an array element.

address-expression

An expression that yields an address to watch.

`size size-expression`

The size of the item to watch, in bytes.

`type expression`

An expression whose type indicates the size of the item to watch. `type` is used only in restart information.

`if conditional-expression`

Sets a condition on the watchpoint. The watchpoint is considered to be hit only if *conditional-expression* evaluates to TRUE. The *conditional-expression* is always evaluated in the global scope. *conditional-expression* is evaluated *after* the process has executed the instruction causing the trap.

conditional-expression may refer to the process-local convenience variable `$is`. `$is` is the value of the watched item after the process has executed the instruction causing the trap. See "Watchpoints" on page 3-14.

NOTE

The `at`, `if`, `name`, `size` and `type` keywords may not be abbreviated in this command.

watchpoint sets a watchpoint in each of the processes specified by the qualifier. This causes the process to stop when it accesses the *lvalue* or *address-expression*. See “Watchpoints” on page 3-14.

You can specify commands to be executed when the watchpoint is hit. See “commands” on page 6-120.

Watchpoints are not supported on CUDA memory locations. This is a technical limitation of the CUDA architecture.

syscallpoint

Print system call information on entry and exit from system calls.

syscallpoint [*eventpoint-modifiers*] [*name=eventpoint-name*] [*syscall-list*]
[*if conditional-expression*]

eventpoint-modifiers:

/disabled
/delete
/once

These modifiers have their standard meaning, as described in “Event-point Modifiers” on page 6-107.

/before
/after

The **/before** and **/after** modifiers control when the **syscallpoint** applies (before the syscall is executed, after, or both). If both options are omitted, then the **syscallpoint** applies to both entry and exit of the service call.

/nostop

The **/nostop** modifier causes the debugger to automatically resume execution of the process after it prints the system call information; thus manual intervention is not required. Regardless, the process **is** stopped while the debugger gets control and prints a message (and resumes the process unless **/nostop** is specified). Note that all threads are stopped during this time.

/except

The **/except** modifier changes the meaning of the *syscall-list*. Instead of matching those system calls, it matches all except those system calls. The *syscall-list* must be non-empty if you use **/except** (because an empty *syscall-list* implies all system calls -- and using **/except** in such a circumstance would be useless).

name=*eventpoint-name*

Gives a name to the syscallpoint for later reference. See “name” on page 6-107. If *eventpoint-name* is already defined then this command adds the newly created syscallpoints to the list of eventpoints associated with the name.

syscall-list

Specifies the list of system calls that you are interested in tracing. If *syscall-list* is omitted, it implies all system calls. Otherwise, *syscall-list* should be one or more valid system call names, separated by commas or spaces.

The list of valid system call names can be found using the graphical interface in the **Syscallpoint Dialog** by pressing the **Select...** button (see “System Call Selection Dialog” on page 8-39).

It is possible while debugging multiple processes from different target systems at the same time in the same NightView session that the list of valid system calls can differ between processes; although, the majority of system call names are common between Linux kernel versions.

if *conditional-expression*

Specifies a syscallpoint condition. If specified, the expression is evaluated to see if the syscallpoint applies or not. This can test more complex conditions than merely the service call number. Since NightView cannot predict ahead of time where a syscallpoint might stop, the condition is always evaluated in the global scope.

The **commands** command (see “commands” on page 6-120) can be used to add commands to be executed when the process stops on an eventpoint. No commands are allowed if the **/nostop** modifier is specified.

Special Restrictions

The combination of **/before** and **/after** modifiers as well as the list of system calls must be unique across the entire process. You are not allowed to define multiple syscallpoints which have overlapping system calls. The **/except** option may be useful to define a catch-all syscallpoint that does not apply to any of the other syscallpoints already defined.

Due to the way the Linux kernel implements system call tracing, NightView is limited when stopped at a syscallpoint. Anything that requires NightView to execute the process will not work at a syscallpoint. This means that you cannot call functions in expressions (which may mean you can't do some operations internally implemented as functions that aren't obvious).

These restrictions also apply to the syscallpoint condition and any commands you may want to execute when the syscallpoint stops. For the same reason, you can't do things like modify the program counter in an expression and prevent the service call from completing.

Due to Linux internals, the system call number on exit from a system call is sometimes reported as -1 rather than the actual service number. This behavior is typically limited to special service calls used to do things like return from signal handlers.

Controlling Execution

This section describes commands used to control the execution of a process.

Most of the commands described in this section cause the processes specified in the qualifier to resume execution and then wait for something to happen. (This is what you usually want when you are debugging a single process.) Only **resume** resumes execution and then returns immediately for another command.

Some of the commands continue until something special happens. For example, **step** continues until control crosses a source line boundary. However, you should be aware that another event, such as a signal or hitting a breakpoint, may cause the process to stop sooner.

If the process stopped because of a signal, then it will receive that signal when the process resumes, subject to the setting of the **handle** command, see “handle” on page 6-144. If you want the process to receive a different signal, or no signal at all, then use the **signal** command. See “signal” on page 6-142.

NOTE

On Linux, there is no way to pass SIGSTOP to a process being debugged.

If you ask to continue execution of a process with any of the commands here, and that process is already executing, then you get a warning message. Any other processes specified by the qualifier are continued.

If a process is stopped at a breakpoint or watchpoint, it is *not* necessary to remove the breakpoint or watchpoint before continuing.

set-run-mode

Controls whether all threads in a multi-threaded process resume execution when a single thread is told to resume execution.

set-run-mode *mode*

one

In this mode, commands or actions that cause a single thread to resume execution have no effect on other threads in the process. Just the current thread executes.

all

In this mode, commands or actions that cause a thread to resume execution cause all other threads in the process to resume execution as well.

The run mode is consulted when a thread's execution is resumed by an action in the graphical interface (e.g. pressing the resume icon, using a keyboard shortcut in a source panel, etc.) or when a command is typed that does not supply overriding run mode options (e.g. `/one` or `/all`). See "Executing" on page 3-44 for more detail.

When using the *one* mode on a CUDA context, host threads will remain stopped but all CUDA code will resume execution. This is a technical limitation of the CUDA driver.

IMPORTANT

The *one* mode should be used with care. It is not uncommon for a single thread to block waiting on operations to be completed by other threads (e.g. by `pthread_cond_wait(3)` and many other services that hold resources). If those other threads are not permitted to resume, the one thread that was resumed may block indefinitely. In such circumstances, stop the blocked thread and resume all threads using `/all` on your command or change the run mode to *all*.

continue

Continue execution and wait for something to happen.

continue [/one | /all] [count]

Abbreviation: **c**

count

If the *count* argument is specified, the processes will not stop at the current breakpoint or watchpoint again until they have hit it *count* times. This argument is ignored for any processes that are not stopped at breakpoints or watchpoints.

/one

This option causes only the current thread to resume execution. It has no effect on any other threads. If both this option and `/all` are omitted, the current run mode controls which threads execute (see "set-run-mode" on page 6-131). This option is redundant for single-threaded processes. When applied to a CUDA context, all CUDA code will execute. See "set-run-mode" on page 6-131.

/all

This option causes all threads in the process to resume execution. If both this option and `/one` are omitted, the current run mode controls which threads execute (see "set-run-mode" on page 6-131). This option is redundant for single-threaded processes.

continue causes the processes specified by the qualifier to resume execution at the point where they last stopped. Processes run concurrently. Each process will execute until some event, such as hitting a breakpoint, causes it to stop.

If this command is entered interactively, the debugger does not prompt for any more commands until one of the processes specified by the qualifier stops executing for some reason. Note that only one of the specified processes has to stop for the **continue** command to complete; it does not wait for *all* of the processes to stop. Note also that a process is considered to be stopped the moment it hits a breakpoint or watchpoint; if the breakpoint or watchpoint has commands attached to it, they probably will not execute before you receive a prompt for another command.

If a **continue** command in a breakpoint (or watchpoint) command stream continues execution of the process stopped at that breakpoint or watchpoint, the command stream is terminated; no further commands are executed from that stream. If a **continue** command continues execution of a process that is currently executing another breakpoint (or watchpoint) command stream, the **continue** command does not take effect until that command stream has completed execution. See “Command Streams” on page 3-39.

If a **continue** command continues execution of a process that is currently executing an **on program** or **on restart** command stream, the **continue** command does not take effect until the affected process has been completely initialized by NightView and is ready to be debugged.

continue is similar to **resume**. See “resume” on page 6-133.

Example:

```
(local) c 5
```

The processes specified by the default qualifier are resumed and will not stop again at the current breakpoint or watchpoint until it has been hit 5 times.

resume

Continue execution.

```
resume [/one | /all] [sigid]
```

sigid

The processes receive the specified signal when they resume execution. *sigid* is a signal name or number. You may specify a signal name with or without the SIG prefix; the name is case-insensitive. If *sigid* is 0, then the processes receive no signal when they resume execution. See “signal” on page 6-142.

If this argument is not present, then the processes are resumed with the signal that caused them to stop, similar to **continue**.

NOTE

On Linux, there is no way to pass SIGSTOP to a process being debugged.

`/one`

This option causes only the current thread to resume execution. It has no effect on any other threads. If both this option and `/all` are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131). This option is redundant for single-threaded processes. When applied to a CUDA context, all CUDA code will execute. See “set-run-mode” on page 6-131.

`/all`

This option causes all threads in the process to resume execution. If both this option and `/one` are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131). This option is redundant for single-threaded processes.

resume causes the processes or threads specified by the qualifier to resume execution at the point where they last stopped. The processes or threads run concurrently. Each process will execute until some event, such as hitting a breakpoint or watchpoint, causes it to stop.

If a **resume** command in a breakpoint (or watchpoint) command stream continues execution of the process stopped at that breakpoint or watchpoint, the command stream is terminated; no further commands are executed from that stream. If a **resume** command continues execution of a process that is currently executing another breakpoint (or watchpoint) command stream, the **resume** command does not take effect until that command stream has completed execution. See “Command Streams” on page 3-39.

If a **resume** command continues execution of a process that is currently executing an **on program** or **on restart** command stream, the **resume** command does not take effect until the affected process has been completely initialized by NightView and is ready to be debugged.

The difference between **resume** and **continue** is that **resume** does not wait for the processes to stop. The debugger continues to read and process commands. See “continue” on page 6-132.

Example:

```
(local) resume 0
```

The processes specified by the default qualifier are resumed with no signal.

Example:

```
(local) resume 2
```

The processes specified by the default qualifier are resumed with signal number 2.

step

Execute one line, stepping into procedures.

step [/one | /all] [*repeat*]

Abbreviation: **s**

repeat

The *repeat* argument specifies the number of lines to single step. The default is one line.

/one

This option causes only the current thread to resume execution. It has no effect on any other threads. When applied to a CUDA context, multiple threads will execute, possibly including all CUDA code. See “set-run-mode” on page 6-131 and “CUDA Debugging” on page 3-46. If both this option and /all are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131). This option is redundant for single-threaded processes.

/all

This option causes all threads in the process to resume execution. If both this option and /one are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131). This option is redundant for single-threaded processes.

It is worth noting that the /all option does **not** imply that all threads will single step. All threads will execute while the current thread steps and then all threads will be stopped when the step operation completes. The other threads may execute many source lines or only a few instructions.

step causes the processes specified by the qualifier to continue execution until they have crossed a source line boundary. With a repeat count, this happens *repeat* times.

step follows execution into called procedures. That is, if the current line is a procedure call, and you **step**, then the process will execute until it is in that new procedure and then stop. If you want to step over the procedure, use **next**. See “next” on page 6-136.

If a **step** command causes execution to enter or leave a called procedure, then the output includes the equivalent of a **frame 0** command to show this. See “frame” on page 6-147.

This command completes only when all of the processes specified by the qualifier have completed the single step or stopped for some other reason (like receiving a signal). See “Signals” on page 3-18 for a discussion of the interactions between single-stepping and signals.

step is interpreted relative to the current frame. See “Current Frame” on page 3-29. That is, any lower frames are automatically finished before stepping.

There are also commands to single step individual instructions. See “stepi” on page 6-138

and “nexti” on page 6-139.

When the program has just started, **step** steps to the beginning of the procedure that calls static initializers or library-level elaboration procedures, if any. If there are none, **step** steps to the beginning of the main procedure.

Because of optimization and other considerations, a process may appear to stop multiple times in the same line or not at all in some lines. The decorations that appear when you list the source can help you decide which lines are executable (see “Source Line Decorations” on page 6-88). Also, disassembly can help you determine the flow of control through your program (see “x” on page 6-95 and “Source Menu” on page 8-11).

If the **step** command causes execution to enter a procedure which is uninteresting, the **step** acts like **next**. See “Interesting Subprograms” on page 3-31. See “next” on page 6-136.

NOTE

If you step to a source line, and the instructions corresponding to that line begin with an inline call, NightView positions you at the beginning of the inline subprogram, rather than on the line with the call.

When stepping in CUDA code, special rules apply with regard to how much of the CUDA device state may change during the step operation. You will step only a single warp if the current frame is the innermost frame (i.e. frame 0 if there are no non-interesting frames), and you do not step over a `__syncthreads()` operation. Otherwise, the step operation will resume execution of all CUDA devices until the step completes.

next

Execute one line, stepping over procedures.

next [/one | /all] [*repeat*]

Abbreviation: **n**

repeat

The *repeat* argument specifies the number of lines to single step. The default is one line.

/one

This option causes only the current thread to resume execution. It has no effect on any other threads. When applied to a CUDA context, multiple threads will execute, possibly including all CUDA code. See “set-run-mode” on page 6-131 and “CUDA Debugging” on page 3-46. If both this option and /all are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131). This option is redundant for single-threaded processes.

`/all`

This option causes all threads in the process to resume execution. If both this option and `/one` are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131). This option is redundant for single-threaded processes.

It is worth noting that the `/all` option does **not** imply that all threads will step over a single line. All threads will execute while the current thread steps over its current line and then all threads will be stopped when the step operation completes. The other threads may execute many source lines or only a few instructions.

next causes the processes specified by the qualifier to continue execution until they have crossed a source line boundary. With a repeat count, this happens *repeat* times.

next steps over called procedures, including “inline” procedures. See “Inline Subprograms” on page 3-30. That is, if the current line is a procedure call, and you single step with **next**, then the process will execute until that new procedure has returned. If you want to follow execution into the procedure, use **step**. See “step” on page 6-135.

If a **next** command causes execution to leave a called procedure, then the output includes the equivalent of a **frame 0** command to show this. See “frame” on page 6-147.

This command completes only when all of the processes specified by the qualifier have completed the single step or stopped for some other reason (like receiving a signal). See “Signals” on page 3-18 for a discussion of the interactions between single-stepping and signals.

next is interpreted relative to the current frame. See “Current Frame” on page 3-29. That is, any lower frames are automatically finished before stepping.

There are also commands to single step individual instructions. See “nexti” on page 6-139 and “stepi” on page 6-138.

When the program has just started, **next** steps to the beginning of the main procedure.

Because of optimization and other considerations, each process may appear to stop multiple times in the same line or not at all in some lines. The decorations that appear when you list the source can help you decide which lines are executable (see “Source Line Decorations” on page 6-88). Also, disassembly can help you determine the flow of control through your program (see “x” on page 6-95 and “Source Menu” on page 8-11).

NOTE

If you step to a source line, and the instructions corresponding to that line begin with an inline call, NightView positions you at the beginning of the inline subprogram, rather than on the line with the call.

When stepping in CUDA code, special rules apply with regard to how much of the CUDA device state may change during the step operation. You will step only a single warp if the

current frame is the innermost frame (i.e. frame 0 if there are no non-interesting frames), you do not step over a `__syncthreads()` operation, and you do not step over any called procedures. Otherwise, the step operation will resume execution of all CUDA devices until the step completes.

stepi

Execute one instruction, stepping into procedures.

stepi [/one | /all] [repeat]

Abbreviation: **si**

repeat

The *repeat* argument specifies the number of instructions to single step. The default is one instruction.

/one

This option causes only the current thread to resume execution. It has no effect on any other threads. When applied to a CUDA context, multiple threads will execute, possibly including all CUDA code. See “set-run-mode” on page 6-131 and “CUDA Debugging” on page 3-46. If both this option and /all are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131). This option is redundant for single-threaded processes.

/all

This option causes all threads in the process to resume execution. If both this option and /one are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131). This option is redundant for single-threaded processes.

It is worth noting that the /all option does **not** imply that all threads will single step an instruction. All threads will execute while the current thread steps and then all threads will be stopped when the step operation completes. The other threads may execute many source lines or only a few instructions.

stepi executes a single machine instruction in each of the processes specified by the qualifier.

This is very similar to **step**, except that **step** executes lines and **stepi** executes individual instructions. See “step” on page 6-135.

This command completes only when all of the processes specified by the qualifier have completed the single step or stopped for some other reason (like receiving a signal). See “Signals” on page 3-18 for a discussion of the interactions between single-stepping and signals.

stepi is interpreted relative to the current frame. See “Current Frame” on page 3-29. That is, any lower frames are automatically finished before stepping.

Sometimes, when stepping by instructions, it is useful to set up a **display** command to show the instruction that is just about to be executed each time the process stops. To do that, say

```
(local) display/i $pc
```

See “display” on page 6-101.

If the **stepi** command causes execution to enter a procedure which is uninteresting, the **stepi** acts like **nexti**. See “Interesting Subprograms” on page 3-31. See “nexti” on page 6-139.

When stepping in CUDA code, special rules apply with regard to how much of the CUDA device state may change during the step operation. You will step only a single warp if the current frame is the innermost frame (i.e. frame 0 if there are no non-interesting frames), and you do not step over a `__syncthreads()` operation. Otherwise, the step operation will resume execution of all CUDA devices until the step completes.

nexti

Execute one instruction, stepping over procedures.

```
nexti [/one | /all] [repeat]
```

Abbreviation: **ni**

repeat

The *repeat* argument specifies the number of instructions to single step. The default is one instruction.

/one

This option causes only the current thread to resume execution. It has no effect on any other threads. When applied to a CUDA context, multiple threads will execute, possibly including all CUDA code. See “set-run-mode” on page 6-131 and “CUDA Debugging” on page 3-46. If both this option and */all* are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131). This option is redundant for single-threaded processes.

/all

This option causes all threads in the process to resume execution. If both this option and */one* are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131). This option is redundant for single-threaded processes.

It is worth noting that the */all* option does **not** imply that all threads will step over a single instruction. All threads will execute while the current thread steps and then all threads will be stopped when the step operation completes. The other threads may execute many source lines or only a few instructions.

nexti executes a single machine instruction in each of the processes specified by the qualifier, except that **nexti** steps over procedure calls and inlined procedures. See “Inline Subprograms” on page 3-30.

This is very similar to **next**, except that **next** executes lines and **nexti** executes individual instructions. See “next” on page 6-136.

This command completes only when all of the processes specified by the qualifier have completed the single step or stopped for some other reason (like receiving a signal). See “Signals” on page 3-18 for a discussion of the interactions between single-stepping and signals.

nexti is interpreted relative to the current frame. See “Current Frame” on page 3-29. That is, any lower frames are automatically finished before stepping.

When stepping in CUDA code, special rules apply with regard to how much of the CUDA device state may change during the step operation. You will step only a single warp if the current frame is the innermost frame (i.e. frame 0 if there are no non-interesting frames), you do not step over a `__syncthreads()` operation, and you do not step over any called procedures. Otherwise, the step operation will resume execution of all CUDA devices until the step completes.

finish

Continue execution until the current function finishes.

finish [/one | /all]

/one

This option causes only the current thread to resume execution. It has no effect on any other threads. When applied to a CUDA context, all CUDA code will execute. See “set-run-mode” on page 6-131. If both this option and /all are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131). This option is redundant for single-threaded processes.

/all

This option causes all threads in the process to resume execution. If both this option and /one are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131 and “CUDA Debugging” on page 3-46). This option is redundant for single-threaded processes.

It is worth noting that the /all option does **not** imply that all threads will finish out of their current routines. All threads will execute while the current thread finishes executing the current routine and then all threads will be stopped. The other threads may execute only a few instructions, many source lines, or many functions.

finish causes a process to continue execution until the current frame returns. This happens in each process specified by the qualifier.

Note that this may cause the process to finish multiple procedures, depending on which frame is the current frame. See “frame” on page 6-147. If the current frame is in the context of a thread or thread process chosen by the **select-context** command, execution continues until that thread or thread process completes execution of that procedure, or until the process stops for some other reason. See “list” on page 6-82. See “step” on page 6-135. Multithreaded Programs” on page 3-44.

In general, the exact action of this command is dependent on the language being debugged.

The **finish** command causes execution to leave a called procedure, so the output includes the equivalent of a **frame 0** command to show this.

This command completes only when all of the processes specified by the qualifier have completed the function execution or stopped for some other reason (like receiving a signal). The discussion in “Signals” on page 3-18 concerning interactions between single-stepping and signals also applies to the **finish** command.

When stepping in CUDA code, special rules apply with regard to how much of the CUDA device state may change during the step operation. You will step only a single warp if the current frame is the innermost frame (i.e. frame 0 if there are no non-interesting frames), you do not step over a `__syncthreads()` operation, and you do not step over any called procedures. Otherwise, the step operation will resume execution of all CUDA devices until the step completes.

stop

Stop a process.

```
stop [/protected]
```

```
    /protected
```

This option causes the command also to stop any protected threads (see “Stopping” on page 3-44).

The **stop** command stops each of the processes specified by the qualifier. In many cases (such as setting breakpoints), NightView requires a process to be stopped before a command may be applied to the process.

The **stop** command does not complete until all of the specified processes have been stopped. If a specified process is already stopped, this command silently ignores that process.

In general, when the **stop** command is stopping a process with multiple threads, it will stop all non-protected threads within that process (see “Stopping” on page 3-44). This behavior can be overridden with the **/protected** option, which causes all threads to be stopped regardless of whether or not they are protected.

Example:

```
(local) (addams) stop
```

This example stops each of the processes in the process family named ‘addams’.

jump

Continue execution at a specific location.

jump [/one | /all] [at] *location-spec*

location-spec

The *location-spec* specifies where to continue execution. See “Location Specifiers” on page 6-15.

/one

This option causes only the current thread to resume execution. It has no effect on any other threads. If both this option and /all are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131). This option is redundant for single-threaded processes.

/all

This option causes all threads in the process to resume execution. If both this option and /one are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131). This option is redundant for single-threaded processes.

It is worth noting that the /all option does **not** imply that all threads will begin execution at the specified location. Only the current thread will resume execution at the specified location; all other threads will resume execution from their current PC.

jump causes execution to continue at the specified location. This happens for each process specified in the qualifier.

jump does not modify the stack frames or registers, it just modifies the program counter and continues execution. Unless you are sure the registers have the right contents for the new location, you are cautioned to avoid using this command.

You must be in frame 0, with no hidden frames below frame zero, to use **jump**. See “Interesting Subprograms” on page 3-31.

The **jump** command is not supported in CUDA code.

signal

Continue execution with a signal.

signal [/one | /all] *sigid*

sigid

Specifies the name or number of the signal with which to continue. If *sigid* is 0, then the processes are continued without a signal. You may specify a signal name with or without the SIG prefix; the name is case-insensitive.

/one

This option causes only the current thread to resume execution. It has no effect on any other threads. If both this option and /all are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131). This option is redundant for single-threaded processes.

/all

This option causes all threads in the process to resume execution. If both this option and /one are omitted, the current run mode controls which threads execute (see “set-run-mode” on page 6-131). This option is redundant for single-threaded processes.

signal resumes execution of the processes specified in the qualifier, passing them a signal.

signal is useful if a process has received a signal (causing it to stop and be recognized by the debugger), but you don't want it to see the signal. Then, rather than using **continue** to continue the process, use **signal 0**.

Or, perhaps you want the process to receive a different signal. **signal** can resume your process with any signal.

If a **signal** command in a breakpoint (or watchpoint) command stream continues execution of the process stopped at that breakpoint or watchpoint, the command stream is terminated; no further commands are executed from that stream. If a **signal** command continues execution of a process that is currently executing another breakpoint (or watchpoint) command stream, the **signal** command does not take effect until that command stream has completed execution. See “Command Streams” on page 3-39.

If a **signal** command continues execution of a process that is currently executing an **on program** or **on restart** command stream, the **signal** command does not take effect until the affected process has been completely initialized by NightView and is ready to be debugged.

If a **signal** command continues execution of a CUDA context, then the signal number or name is ignored, because signals are not a meaningful concept for CUDA contexts. The CUDA devices are resumed, however.

For a way to have the debugger deal with signals automatically, see “handle” on page 6-144. **signal** overrides the **pass** setting of **handle**.

NOTE

On Linux, there is no way to pass SIGSTOP to a process being debugged.

Type **info signal** to get a list of all of the signals on your system. See “info signal” on page 6-169.

Example:

```
(local) signal 2
```

The processes resume with signal number 2.

handle

Specify how to handle signals and in the user process.

```
handle [/signal] sigid keyword ...

handle /exception unit-name keyword ...
handle /exception all keyword ...
handle /cpp_exception type-name keyword
handle /cpp_exception * keyword
handle /cpp_exception type-name clear
handle /cpp_exception clear
handle /unhandled_exception keyword
```

/signal

Specifies handling of a signal. This is the default.

sigid

Specifies the name or number of a signal to handle. You may specify a signal name with or without the SIG prefix; the name is case-insensitive.

/exception

Specifies handling of an Ada exception. The **/exception** option is a synonym for **/ada_exception**, provided for backward compatibility with older versions of NightView.*keyword*

keyword is one of **stop**, **nostop**, **print**, **noprint**, **pass** or **nopass**. Multiple keywords may be specified.

clear

The **clear** keyword removes the handle specification for Ada or C++ exceptions. If an exception-name, unit-name or type-name is specified, specifications for the given exception(s) are removed. If no name is specified, all specifications for either Ada or C++ exceptions are removed. However, neither the **all** nor ***** specifications are changed.

handle tells the debugger how to deal with signals sent to the user program.

Here are the meanings of the keywords:

stop

The process stops when it gets this signal. **print** is implied with this keyword.

nostop

The process continues executing automatically after the signal. You may still use **print** to tell you when the signal has occurred.

print

NightView notifies you that the signal has occurred. In the command-line interface, a message is printed to your terminal. In the graphical user interface, a message is printed in the Debug Message Area. See Chapter 8 [Graphical User Interface] on page 8-1. See “Message Panel” on page 8-71.

noprint

You do not receive notification when the signal occurs. **nostop** is implied with this keyword.

pass

The signal will be passed to your process the next time it executes. This keyword is not applicable to exceptions.

NOTE

There is no way to pass SIGSTOP to a process being debugged.

nopass

The signal is discarded, after stopping and printing if that's appropriate. This keyword is not applicable to exceptions.

handle Signals

In most cases, a signal sent to a debugged program will cause that program to be stopped and NightView to be notified of the signal. NightView's normal action for most signals is to notify you of the signal and save it to be passed to the process the next time it is continued. For example, the default setting for SIGQUIT would be described as:

```
(local) handle sigquit stop print pass
```

This default behavior can be altered by the **handle** command. Some settings allow the system to avoid stopping your process and notifying NightView of the signal. See “Signals” on page 3-18 for more information about this.

The default action for a few signals is different than the behavior described above. Consider SIGALRM, which is not usually an error; it is used in the normal functioning of the program. You usually don't want to know when your program gets a SIGALRM (but your program does) so the default setting for SIGALRM is:

```
(local) handle sigalrm nostop noprint pass
```

This says that if NightView discovers that your process has been sent a SIGALRM, it will automatically resume execution and pass the signal to the process without notifying you.

(NightView may not even be aware of the signal with these settings of the **handle** command. See “Signals” on page 3-18.)

SIGINT is handled a little differently; when the process receives a SIGINT, the process stops and NightView notifies you, but the signal is discarded, so that the process never sees it. The normal setting for SIGINT is:

```
(local) handle sigint stop print nopass
```

For a way to deal with signals one at a time, see “signal” on page 6-142.

To find out the current settings for all the signals, see “info signal” on page 6-169.

If two conflicting keywords are specified, they are both applied, in the order they appear. For example, if the initial setting for signal number 1 is stop, print, pass, and you say:

```
(local) handle 1 noprint print
```

then the new setting is nostop, print, pass, because noprint implies nostop.

handle applies to all the processes specified in the qualifier.

handle Ada Exceptions

When determining the applicable behavior for an Ada exception, a specification with the given exception name is checked first. (This also includes specifications done using a unit-name.) If none is found, the specification for `all` is used. If the user has not specified it, then its default settings are `nostop, noprint`.

If the settings are set to stop, then execution is stopped in the Ada run-time routine that routes exceptions to the proper handler. This routine usually is uninteresting, so the current frame is set to the code that caused the exception. See “Interesting Subprograms” on page 3-31.

If the settings are set to print, the user is informed of the name of the exception and the *Ada Reference Manual* references.

To find out how one or more Ada exceptions will be handled, you may use the `info exception/ada` command. See “info exception” on page 6-176.

See “info exception” on page 6-176.

handle C++ Exceptions

When determining applicable behavior for a C++ exception, a specification with the given type name is checked first. If none is found, and the type is a C++ class type, its superclasses are checked. For single inheritance, this simply searches through the chain of superclasses in order. For multiple inheritance, this is done as a breadth-first search through the superclass hierarchy, first checking each specifically named superclass, then searching each of their superclasses in turn. If that also reveals no specification, then the specification for `*` is checked. If the user has not specified it, then its default settings are `nostop, noprint`.

If the settings are set to stop, then execution generally is stopped in compiler-generated code in the catching function, but often in inline dispatching code that will branch to the correct catch block. The exact organization of this code can vary from compiler to compiler, and even compiler version to version.

In addition, if the settings are set to stop, the `$ex` convenience variable is created, and contains a reference to the thrown exception object.

If the settings are set to print, the user is informed of the name of the exception.

To find out how one or more C++ exceptions will be handled, you may use the `info exception/cpp` command. See “info exception” on page 6-176.

Selecting Context

frame

Select a new stack frame or print a description of the current stack frame.

frame [*frame-number*]

frame **expression* [*at location-spec*]

Abbreviation: **f**

frame-number

Frame number selected as the new current stack frame. Frame number zero corresponds to the currently executing frame. Frame numbers for all the currently available stack frames may be obtained with the **backtrace** command (see “backtrace” on page 6-91).

**expression*

Expression which yields an address at which the stack frame should start. This is the value that `$cfa` would have, not the value of `$sp`.

location-spec

Specifies a location in the program to use to interpret the stack frame at the address given by **expression*. See “Location Specifiers” on page 6-15. If you do not supply this argument, the default is the current value of `$cpc`.

NOTE

The `at` keyword may not be abbreviated in this command.

If no argument is given, a brief description of the current stack frame is printed. If multiple processes are specified in the command qualifier, each of them is described

separately. For a more complete description of a frame, see “info frame” on page 6-166.

If a *frame-number* is given, the chosen stack frame is selected as the current frame (see “Current Frame” on page 3-29).

The **expression* form of this command is provided for those occasions in which the stack is in an inconsistent state, or you wish to examine some memory whose contents look like stack frames. You should be very careful when using this form, observing the following cautions.

- A stack frame cannot be interpreted except in the context of some program-counter value. Therefore, you must be sure that the *location-spec* you give (or the value of `$pc`) is consistent with the stack frame you are examining.
- The values of the machine registers are not altered by this form of the **frame** command. This means that variables that reside in registers cannot be reliably examined.
- The **up**, **down**, and **backtrace** commands are executed relative to the given frame address and program-counter value. However, the register contents for calling frames may still be incorrect, since only the registers saved in the stack can be restored by NightView.
- Modifying a register (or a variable stored in a register) may alter the current value of a machine register, or it may alter the value of that register stored on the stack. You must be very careful when doing this.
- Unless you have modified `$pc` or other machine registers, resuming execution of the process will resume with the state the process was in before the **frame** command was issued.

Once you have issued a **frame** command with a **expression* argument, you can restore the previous view of the stack by issuing a **frame** command with a *frame-number* argument. This restores NightView's view of the stack to what it was before you issued the **frame** **expression* command.

We recommend that, while you have the frame set using the **expression* form, you should restrict yourself to just using the **up**, **down**, **backtrace**, and **print** commands, and that you print only global variables or variables stored on the stack.

up

Move one or more stack frames toward the caller of the current stack frame.

up [*number-of-frames*]

number-of-frames

Number of stack frames to advance toward the oldest calling frame. The number zero may be used to restore the current source position in the current frame (see “Current Frame” on page 3-29). If a negative number is specified, then frames are advanced toward the newest stack frame (see “down” on page 6-149).

If *number-of-frames* is not given, the number defaults to one, corresponding to the caller of the current frame.

This command is applied to each process in the qualifier.

down

Move one or more stack frames toward frames called by the current stack frame.

down [*number-of-frames*]

number-of-frames

Number of stack frames to advance toward the currently executing (newest) stack frame. The number zero may be used to restore the current source position in the current frame (see “Current Frame” on page 3-29). If a negative number is specified, then frames are advanced toward the oldest stack frame (see “up” on page 6-148).

If *number-of-frames* is not given, the number defaults to one, corresponding to the frame called by the current frame.

This command is applied to each process in the qualifier.

select-context

Select the context of a thread or a thread process.

select-context default

select-context thread=*expression*

select-context pid=*pid*

select-context name=*name*

select-context stackaddress=*expression*

select-context cuda [*context context*] [*sm sm*]
[*warp warp*] [*lane lane*]

select-context cuda [*context context*] [*grid grid*]
[*block x[,y[,z]]*] [*thread x[,y[,z]]*]

default

This keyword selects the stack frame for the context where the process has stopped. If the process has threads, the default context is the thread process that stopped the process. See “See “list” on page 6-82. See “step” on page 6-135. Multithreaded Programs” on page 3-44.

`thread=expression`

The `thread=` keyword selects the context of a thread created by `thr_create(3thread)`. The *expression* must be the `thread_t` value returned by `thr_create` for a currently active thread.

`pid=pid`

The `pid=` keyword selects the context of a specific thread process. The *pid* is the ID of the thread process whose context is selected.

`name=name`

The `name=` keyword selects the context of a specific thread by matching the supplied name; *name* must identify a single thread in the current process. See “set-thread-name” on page 6-155 for more information on thread names.

`stackaddress=address`

The `stackaddress=` keyword selects the context of a specific thread by matching the supplied address; *address* must be a process address value inside some thread's stack.

This can be useful if you know the `pthread_self()` values reported by threads, but you are looking at a core file which has lost the thread associations. For all but the main thread, the `pthread_self()` value is normally contained within the stack segment for the associated thread.

`cuda`

The `cuda` keyword indicates that subsequent keywords all specify a change to a CUDA thread. Those keywords must be either all physical or all logical. The physical keywords are `sm`, `warp`, and `lane`. The logical keywords are `grid`, `block`, and `thread`. The `context` keyword may be specified in either case

It is not necessary to specify all of the keywords of the selected type. If any are left unspecified, the command will attempt to leave those coordinate values the same as in the current context. If the result would not be a valid set of coordinates, it will select the closest values that it can find which are valid. It is not possible to change context to a CUDA thread that is not currently executing (i.e. has not yet been dispatched, has completed, or currently is evicted from the device).

`context context`

The `context` keyword indicates the specified CUDA *context*.

`sm sm`

The `sm` keyword indicates the specified symmetric multiprocessor (*sm*).

`warp warp`

The `warp` keyword indicates the specified *warp* within the symmetric multiprocessor (SM).

lane *lane*

The `lane` keyword indicates the specified *lane* within the warp.

grid *grid*

The `grid` keyword indicates the specified *grid*.

block *x[,y[,z]]*

The `block` keyword indicates the specified *x*; (*x,y*); or (*x,y,z*) block coordinates within the grid.

thread *x[,y[,z]]*

The `thread` keyword indicates the specified *x*; (*x,y*); or (*x,y,z*) thread coordinates within the block.

When a process that contains multiple threads, thread processes, or CUDA contexts stops, the *current thread* becomes the specific thread thread process, or CUDA thread that caused it to stop. The *current context* becomes the innermost interesting frame of that thread. (For a discussion of how this choice is made, see “See “list” on page 6-82. See “step” on page 6-135. Multithreaded Programs” on page 3-44.) The **select-context** command allows you to change the current context, so that subsequent commands interact with the specified context (see “Examining Your Program” on page 3-24 and “See “list” on page 6-82. See “step” on page 6-135. Multithreaded Programs” on page 3-44). In all cases except when using the `stackaddress=` keyword, the current thread also is changed.

Once a context has been selected, all **frame**, **up**, **down**, and **backtrace** commands apply to that context. All expressions and references to registers also refer to that context.

To get a list of the threads and thread processes in a process, see “info threads” on page 6-177.

Miscellaneous Commands

help

Access the online help system.

help [*section*]

section

The name of a section in this manual (anything in the table of contents).

You can read any section in this document by giving the section name (or a unique prefix of the section name) as an argument to the **help** command.

If you type **help** without arguments, the help system displays the document section most

relevant to the last error you received. Type **help** again to see help on the previous error you received, and so on.

Error message identifiers are section names, so you can get help for a specific error by giving the **help** command with the error message identifier. An error message identifier, beginning with E-, is printed with each error message. See “Errors” on page 3-39.

In the non-graphical user interfaces, **help** prints to the terminal. In the graphical user interface, **help** uses another program to display the documentation in a separate window. See “GUI Online Help” on page 8-1.

NOTE

In the non-graphical user interfaces, help is available only for error messages.

The **help** command ignores the command qualifier.

Examples:

```
(local) help Summary of Commands
```

The above example displays the section of the document that contains a brief description of each command.

```
(local) help backtrace
```

Display the description of the **backtrace** command.

```
(local) help E-command_proc003
```

Display help for the error with error message identifier E-command_proc003.

refresh

Re-read source files and refresh the terminal screen.

refresh

NightView normally notices when source files have changed when it switches the current source display. If you want to update the display without switching files, or if you suspect that a system problem has fooled the automatic mechanism, you can force source files to be read again with the **refresh** command.

In the simple full-screen interface, the **refresh** command also clears the terminal screen and redraws it. This is helpful when the screen becomes garbled, such as with a modem and noisy phone lines. See Chapter 7 [Simple Full-Screen Interface] on page 7-1.

shell

Run an arbitrary shell command.

shell [*shell-command*]

The **shell** command is used to execute a single line in a subshell. This command has nothing to do with debugging and the qualifier is ignored. It is simply provided because it is sometimes convenient to have a way to execute a shell command without having to suspend or exit the debugger.

If you just type **shell** without arguments, the debugger puts you in a shell where you can execute arbitrary commands until you exit the shell, at which time the debugger will get control again. You cannot use this form of the **shell** command inside a macro (See “Defining and Using Macros” on page 6-184).

The programs run by this command run on the local system only (the same one you are running NightView on) and inherit the current working directory of the debugger (see “cd” on page 6-81).

If you start background processes via **shell**, they will continue to run normally even if you quit out of the debugger.

The shell used is determined by looking for the SHELL environment variable, and if that is not found, by using your login shell.

In the simple full-screen interface, NightView does not have control over the terminal while you are executing a **shell** command, so after the command has completed you are asked to press return. This gives you a chance to view the command output before NightView redraws the screen. See Chapter 7 [Simple Full-Screen Interface] on page 7-1.

source

Input commands from a source file.

source *command-file*

command-file

The file to read.

To specify a file on a remote system, use the form *user@host:/path*. See “Remote File Access” on page 3-7.

This command reads the designated file and treats each line in the file as though it were a command you typed in. After reading all the commands in the file, the debugger returns to reading commands from the keyboard again. (If **source** commands are nested, ending one file returns to reading from the previous file.)

If NightView encounters any serious error, it stops reading from a **source** file. See “Command Streams” on page 3-39.

The qualifier on the **source** command has no effect. The default qualifier is applied to

any commands in the source file which do not have explicit qualifiers.

declare-thread-tag

Declare the type of a thread tag.

declare-thread-tag *tagname typespec*

A thread tag that is used without being declared is always a 1 bit boolean flag value. Use **declare-thread-tag** to give the tag a different type before assigning to the *tagname* via a command like **set-tag** *tagname expr* (see “set-tag” on page 6-155) or **set \$thr.tagname = expr**, or with a **patchpoint** command (see “patchpoint” on page 6-110).

The *typespec* consists of one or two parts. The first part names the type, and the second (optional) part declares it to be a simple one dimensional array with a fixed size.

The name of the type can be specified two different ways:

It can be a type from the program, as long as it is a type that can be referenced with a name (where compound names are allowed). The compound name syntax will accept either `.` or `::` to separate the name components. The leading components should name scopes to look in for the next component. The last component should be the name of a type.

It can also be a simple C-like type composed of some (sensible) combination of the keywords: `bool`, `char`, `short`, `int`, `long`, `float`, `double`, `signed`, `unsigned`.

The optional array declaration consists of a C-like open bracket `[` followed by a positive integer constant, followed by a close bracket `]`.

For example,

```
declare-thread-tag foo int
declare-thread-tag bar char[20]
declare-thread-tag zort TypesClass::VectorType
declare-thread-tag color MyPackage.Types.Colors
```

The latter two invocations referring to types declared in a source code present in your program; for example:

```
// C++ Example
class TypesClass {
public:
    typedef float VectorType[10];
};

-- Ada Example
package MyPackage is
    type Colors is (blue, green, chartreuse, vermilion);
end MyPackage;
```

See “Thread Tags” on page 3-45 for more information.

set-tag

Modify the value of a thread tag for the current thread.

set-tag *tagname expression*

The **set-tag** command modifies the value of the specified thread *tag* to the result of the *expression*. It is similar to the command **set \$thr.tagname = expression**. However, the **set-tag** command is able to modify a tag value even while a thread is executing. Internally, this is implemented by remembering the change to be applied the next time the *tag* is referenced by the executing code or by NightView. This implementation is transparent to the user.

Additionally, the **set-tag** command will function as expected if the *tagname* is declared as an array of `char` (see “declare-thread-tag” on page 6-154), and the *expression* is of type `char*` or something similar: it assigns the character contents of the *expression* to *tagname*.

set-thread-name

Set the name of a thread.

set-thread-name [*modifier*] "*name*"

Where *modifier*, if supplied, is one of the following:

name=*"current_name"*

The *name*= keyword selects the thread that matches the supplied *current_name*. This command will fail if the specified *current_name* is not unique among the current list of threads. In this form of the command, two names are specified: the *current_name* which is used to identify the thread; and *name*, the thread's new name. Both names must be specified as quoted strings.

task=*task_name_or_id*

The *task*= keyword selects the thread by its Ada task name or its Ada task ID, which is a pointer value assigned by the Ada runtime.

thread=*thread_id*

The *thread*= keyword selects the thread by its thread ID value, as returned by **pthread_self(3)** and **pthread_create(3)**.

pid=*pid*

The *pid*= keyword selects the thread by its **gettid(2)** value; displayed as its *pid* in NightView. Under Linux, the main thread's **getpid(2)** and **gettid(2)** values are identical; all other threads in the process share the same **getpid(2)** value, but have unique **gettid(2)** values.

If *modifier* is omitted, then the command applies to the current thread.

The **set-thread-name** command allows you to change the thread name used by NightView to further identify the thread in **info-thread** commands and in the Context Panel and in Context, Process, and Thread displays in Data Panels.

By default, NightView automatically assigns names to threads using the name of the start routine whose address is passed to **pthread_create(3)** when the thread is created.

This command overrides that default name.

User defined names are not preserved on restart, so any names assigned to threads using the **set-thread-name** command are not automatically reassigned when rerunning the process. For a way to identify threads across restart, see “Thread Tags” on page 3-45.

This command does not select a thread to be the "current thread" in NightView; it merely changes the specified thread's name.

The **set-thread-name** command may not be used on CUDA contexts or threads.

delay

Delay NightView command execution for a specified time.

delay [*milliseconds*]

milliseconds

The number of milliseconds to delay command execution. If not specified, the default is 1.

This command delays the execution of NightView commands for at least the specified time period, expressed in milliseconds. The actual delay may be longer than the specified period. The command following a **delay** command in the same command stream will not execute until at least the specified time has elapsed.

The primary use of the **delay** command is in command scripts, when you may want to prevent a command from executing immediately after the preceding one. For instance, you may wish to allow time for your program to execute for some length of time between the execution of two NightView commands.

The qualifier on the **delay** command has no effect.

Info Commands

The info commands all start with the word **info**, which may always be abbreviated to the single character **i**. The keyword following **info** identifies one of the many topics for which info is available. Each info command may also have additional arguments specific to the individual command.

The info commands can be broadly divided into two basic categories:

- Status queries, returning information about the current state of the debugger and the processes being debugged.
- Symbol table queries, returning information about program variables and type definitions.

Status Information

The status info commands allow you to query various information about the current state of the debugger (e.g., what breakpoints are set, how many dialogues are active, etc.).

info log

Describe any open log files.

info log

Describes any open log files currently in use by the debugger. The log files may be created by **set-log** (see “set-log” on page 6-62) or by **set-show** (see “set-show” on page 6-35).

info eventpoint

Describe current state of breakpoints, tracepoints, patchpoints, monitorpoints, heappoints, watchpoints, and syscallpoints.

info eventpoint [/verbose] [*eventpoint-spec*] ...

/verbose

Specify that the locations of all eventpoints displayed will be in verbose format. Verbose location format includes the program counter address (or addresses) of the eventpoint and, where possible, the corresponding function name, file name, and line number. The number of PC addresses printed is subject to the limit on printing addresses (see “set-limits” on page 6-64). The *verbose* keyword may be abbreviated.

eventpoint-spec

An eventpoint specifier, which is an eventpoint name, an eventpoint number, or the keyword **at** followed by a location specifier. See “Eventpoint Specifiers” on page 6-18.

This command describes eventpoints associated with the processes in the command qualifier. An eventpoint is any of a breakpoint, tracepoint, patchpoint, monitorpoint, heappoint, watchpoint, or syscallpoint. See “breakpoint” on page 6-108, “tracepoint” on page 6-113, “patchpoint” on page 6-110, “monitorpoint” on page 6-115, “watchpoint” on page 6-127, “syscallpoint” on page 6-129, and “heappoint” on page 6-117.

The information printed includes:

- The eventpoint ID.
- The eventpoint type.
- Current state of eventpoint (enabled, disabled, temporary).
- The eventpoint location. If */verbose* was specified, then the location will be printed in verbose format. Otherwise it will be printed in the format in

which it was specified when the eventpoint was created. For watchpoints, information is printed about the address being watched.

- The number of times program execution has crossed the eventpoint since the program started execution (even if the ignore count or condition was not satisfied, this count is incremented).
- The number of times the eventpoint has been *hit* since the program started execution (this count is incremented only if the condition and ignore count are satisfied).
- Any conditions attached to the eventpoint.
- The current ignore count.
- Any commands attached to the eventpoint (if it is a breakpoint, monitorpoint, watchpoint, or syscallpoint).
- For heappoints, the word *check* if the heappoint does a heap check, or the word *debug* followed by the new settings if the heappoint changes the heap debugger settings.

For more information on the precise interactions of many of these items, see “Interactions Between Conditions, Ignore Counts, etc.” on page 3-12.

This command sets the default **x** command dump address as well as the `$_` predefined convenience variable to the address of the last eventpoint listed. See “x” on page 6-95 and “Predefined Convenience Variables” on page 6-5.

info breakpoint

Describe current state of breakpoints.

info breakpoint [/verbose] [*eventpoint-spec*] ...

Abbreviation: **i b**

/verbose

Specify that the locations of all breakpoints displayed will be in verbose format. Verbose location format includes the program counter address (or addresses) of the breakpoint and, where possible, the corresponding function name, file name, and line number. The number of PC addresses printed is subject to the limit on printing addresses (see “set-limits” on page 6-64). The *verbose* keyword may be abbreviated.

eventpoint-spec

An eventpoint specifier, which is an eventpoint name, an eventpoint number, or the keyword **at** followed by a location specifier. See “Eventpoint Specifiers” on page 6-18.

This command normally describes all breakpoints associated with the processes indicated by the command qualifier. If you specify a list of eventpoint names or numbers, only those events are described. If any of the specified eventpoints are not breakpoints, they are ignored. Breakpoints are created with the **breakpoint** command. See “breakpoint”

on page 6-108.

The information printed includes:

- The breakpoint ID.
- Current state of breakpoint (enabled, disabled, temporary).
- The breakpoint location. If `/verbose` was specified, then the location will be printed in verbose format. Otherwise it will be printed in the format in which it was specified when the breakpoint was created.
- The number of times program execution has crossed the breakpoint since the program started execution (even if the ignore count or condition was not satisfied, this count is incremented).
- The number of times the breakpoint has been *hit* since the program started execution (this count is incremented only if the condition and ignore count are satisfied).
- Any conditions attached to the breakpoint.
- The current ignore count.
- Any commands attached to the breakpoint.

This command sets the default **x** command dump address as well as the `$_` predefined convenience variable to the address of the last breakpoint listed. See “x” on page 6-95 and “Predefined Convenience Variables” on page 6-5.

info tracepoint

Describe current state of tracepoints.

info tracepoint [`/verbose`] [*eventpoint-spec*] ...

/verbose

Specify that the locations of all tracepoints displayed will be in verbose format. Verbose location format includes the program counter address (or addresses) of the tracepoint and, where possible, the corresponding function name, file name, and line number. The number of PC addresses printed is subject to the limit on printing addresses (see “set-limits” on page 6-64). The `verbose` keyword may be abbreviated.

eventpoint-spec

An eventpoint specifier, which is an eventpoint name, an eventpoint number, or the keyword **at** followed by a location specifier. See “Eventpoint Specifiers” on page 6-18.

This command describes tracepoints in the processes indicated by the qualifier. Normally all tracepoints are described, but if an argument is given, only those named are described. Any eventpoints specified in the argument list which are not tracepoints are ignored. Tracepoints are created with the **tracepoint** command. See “tracepoint” on page 6-113.

The information printed includes:

- The tracepoint ID.
- Current state of tracepoint (enabled, disabled, temporary).
- The tracepoint location. If `/verbose` was specified, then the location will be printed in verbose format. Otherwise it will be printed in the format in which it was specified when the tracepoint was created.
- The tracepoint event ID.
- The number of times program execution has crossed the tracepoint since the program started execution (even if the ignore count or condition was not satisfied, this count is incremented).
- The number of times the tracepoint has been *hit* since the program started execution (this count is incremented only if the condition and ignore count are satisfied).
- Any conditions attached to tracepoint.
- The current ignore count.
- The expression being recorded at the tracepoint.

This command sets the default **x** command dump address as well as the `$_` predefined convenience variable to the address of the last tracepoint listed. See “x” on page 6-95 and “Predefined Convenience Variables” on page 6-5.

info patchpoint

Describe current state of patchpoints.

info patchpoint [/verbose] [*eventpoint-spec*] ...

/verbose

Specify that the locations of all patchpoints displayed will be in verbose format. Verbose location format includes the program counter address (or addresses) of the patchpoint and, where possible, the corresponding function name, file name, and line number. The number of PC addresses printed is subject to the limit on printing addresses (see “set-limits” on page 6-64). The *verbose* keyword may be abbreviated.

eventpoint-spec

An eventpoint specifier, which is an eventpoint name, an eventpoint number, or the keyword **at** followed by a location specifier. See “Eventpoint Specifiers” on page 6-18.

This command describes patchpoints in the processes indicated by the qualifier. Normally all patchpoints are described, but if an argument is given, only those named are described. Any eventpoints specified in the argument list which are not patchpoints are ignored. Patchpoints are created using the **patchpoint** command. See “patchpoint” on page 6-110.

The information printed includes:

- The patchpoint ID.
- Current state of patchpoint (enabled, disabled, temporary).
- The patchpoint location. If `/verbose` was specified, then the location will be printed in verbose format. Otherwise it will be printed in the format in which it was specified when the patchpoint was created.
- The number of times program execution has crossed the patchpoint since the program started execution (even if the ignore count or condition was not satisfied, this count is incremented).
- The number of times the patchpoint has been *hit* since the program started execution (this count is incremented only if the condition and ignore count are satisfied).
- Any conditions attached to patchpoint.
- The current ignore count.
- The expression patched in at that point, or a description of where the program will branch.

This command sets the default **x** command dump address as well as the `$_` predefined convenience variable to the address of the last patchpoint listed. See “x” on page 6-95 and “Predefined Convenience Variables” on page 6-5.

info monitorpoint

Describe current state of monitorpoints.

info monitorpoint [`/verbose`] [*eventpoint-spec*] ...

/verbose

Specify that the locations of all monitorpoints displayed will be in verbose format. Verbose location format includes the program counter address (or addresses) of the monitorpoint and, where possible, the corresponding function name, file name, and line number. The number of PC addresses printed is subject to the limit on printing addresses (see “set-limits” on page 6-64). The `verbose` keyword may be abbreviated.

eventpoint-spec

An eventpoint specifier, which is an eventpoint name, an eventpoint number, or the keyword **at** followed by a location specifier. See “Eventpoint Specifiers” on page 6-18.

This command describes monitorpoints in the processes indicated by the qualifier. Normally all monitorpoints are described, but if an argument is given, only those named are described. Any eventpoints specified in the argument list which are not monitorpoints are ignored. Monitorpoints are created with the **monitorpoint** command. See “monitorpoint” on page 6-115.

The information printed includes:

- The monitorpoint ID.
- Current state of monitorpoint (enabled, disabled, temporary).
- The monitorpoint location. If `/verbose` was specified, then the location will be printed in verbose format. Otherwise it will be printed in the format in which it was specified when the monitorpoint was created.
- The number of times program execution has crossed the monitorpoint since the program started execution (even if the ignore count or condition was not satisfied, this count is incremented).
- The number of times the monitorpoint has been *hit* since the program started execution (this count is incremented only if the condition and ignore count are satisfied).
- Any conditions attached to monitorpoint.
- The current ignore count.
- The commands attached to the monitorpoint.

This command sets the default **x** command dump address as well as the `$_` predefined convenience variable to the address of the last monitorpoint listed. See “x” on page 6-95 and “Predefined Convenience Variables” on page 6-5.

info heappoint

Describe the current state of heappoints.

info heappoint [`/verbose`] [*eventpoint-spec*] ...

/verbose

Specify that the locations of all heappoints displayed will be in verbose format. Verbose location format includes the program counter address (or addresses) of the heappoint and, where possible, the corresponding function name, file name, and line number. The number of PC addresses printed is subject to the limit on printing addresses (see “set-limits” on page 6-64). The `verbose` keyword may be abbreviated.

eventpoint-spec

An eventpoint specifier, which is an eventpoint name, an eventpoint number, or the keyword **at** followed by a location specifier. See “Eventpoint Specifiers” on page 6-18.

This command describes heappoints in the processes indicated by the qualifier. Normally all heappoints are described, but if an argument is given, only those named are described. Any eventpoints specified in the argument list which are not heappoints are ignored. Heappoints are created with the **heappoint** command. See “heappoint” on page 6-117.

The information printed includes:

- The heappoint ID.

- Current state of heappoint (enabled, disabled, temporary).
- The heappoint location. If `/verbose` was specified, then the location will be printed in verbose format. Otherwise it will be printed in the format in which it was specified when the heappoint was created.
- The number of times program execution has crossed the heappoint since the program started execution (even if the ignore count or condition was not satisfied, this count is incremented).
- The number of times the heappoint has been *hit* since the program started execution (this count is incremented only if the condition and ignore count are satisfied).
- Any conditions attached to heappoint.
- The current ignore count.
- The word `check` if this heappoint does a heap check, or the word `debug` followed by the new settings if this heappoint changes the heap debugger settings.

This command sets the default **x** command dump address as well as the `$_` predefined convenience variable to the address of the last heappoint listed. See “x” on page 6-95 and “Predefined Convenience Variables” on page 6-5.

info watchpoint

Describe current state of watchpoints.

info watchpoint [`/verbose`] [*eventpoint-spec*] ...

/verbose

The `verbose` keyword is accepted for compatibility with other watchpoints, but is ignored. The `verbose` keyword may be abbreviated.

eventpoint-spec

An eventpoint specifier, which is an eventpoint name, an eventpoint number, or the keyword **at** followed by a location specifier. See “Eventpoint Specifiers” on page 6-18.

This command describes watchpoints in the processes indicated by the qualifier. Normally all watchpoints are described, but if an argument is given, only those named are described. Any eventpoints specified in the argument list which are not watchpoints are ignored. Watchpoints are created with the **watchpoint** command. See “watchpoint” on page 6-127.

The information printed includes:

- The watchpoint ID.
- Current state of the watchpoint (enabled, disabled, temporary).
- The address being watched.

- The number of times the process accessed the address being watched since the program started execution. This count is incremented even if the ignore count or condition was not satisfied. This number is displayed as `#crossings` (for consistency with other eventpoint types).
- The number of times the watchpoint has been *hit* since the program started execution (this count is incremented only if the condition and ignore count are satisfied).
- Any conditions attached to the watchpoint.
- The current ignore count.
- Any commands attached to the watchpoint.

info syscallpoint

Describe current state of watchpoints.

info syscallpoint [*eventpoint-spec*] . . .

eventpoint-spec

An eventpoint specifier, which is an eventpoint name or an eventpoint number. See “Eventpoint Specifiers” on page 6-18.

This command describes syscallpoints in the processes indicated by the qualifier. Normally all syscallpoints are described, but if an argument is given, only those named are described. The name or number identifies the syscallpoint and not the system calls themselves.

Any eventpoints specified in the argument list which are not syscallpoints are ignored. Syscallpoints are created with the **syscallpoint** command. See “syscallpoint” on page 6-129.

The information printed includes:

- The syscallpoint ID.
- Current state of the syscallpoint (enabled, disabled, temporary).
- The system call(s) being traced.
- The number of times the process encountered a matching system call since the program started execution. This count is incremented even if the ignore count or condition was not satisfied. This number is displayed as `#crossings` (for consistency with other eventpoint types).
- The number of times the syscallpoint has been *hit* since the program started execution (this count is incremented only if the condition and ignore count are satisfied).
- Any conditions attached to the syscallpoint.
- The current ignore count.
- Any commands attached to the syscallpoint.

info frame

Describe a stack frame.

info frame [/v] [**expression* [at *location-spec*]]

/v

If this option is supplied, NightView prints detailed, machine-specific, information about the requested stack frame. You are seldom likely to be interested in this information; it is provided primarily for detecting problems with the generated debugging information.

**expression*

The address of a stack frame. This is the value that `$cfa` would have, not `$sp`.

location-spec

Specifies a location in the program to use to interpret the stack frame at the address given by **expression*. See “Location Specifiers” on page 6-15. If you do not supply this argument, the default is the current value of `$cpc`.

NOTE

The `at` keyword may not be abbreviated in this command.

This command describes all available information about the current stack frame for a process (see “Current Frame” on page 3-29). See also “frame” on page 6-147.

If multiple processes are specified in the command qualifier, each of them is described separately. An error message is printed if any of the processes are running.

If the optional **expression* is given, then the frame at that address is described (but the current frame is not changed). If you supply the *location-spec* argument, the frame is interpreted as a frame for the routine at the resulting address. If you omit this argument, the current value of `$cpc` is used in decoding the frame.

If **expression* does not evaluate to a valid frame address, or the frame at that address does not correspond to the given program location, the information printed will probably be nonsense.

The information printed about a frame includes:

- The address of the frame.
- The addresses of the adjacent frames (if any).
- The frame size.
- The saved return address and its location on the stack (or in a register).
- Any saved registers and their locations on the stack.

- Which registers are currently in use as stack and/or frame pointers and their relation to the current frame.
- The name of the subroutine associated with the frame along with the source line and file name (if known).
- How `$fp` is computed for the frame.

info directories

Print the search path used to locate source files.

info directories

Print the search path used to locate source files. If multiple processes are given in the qualifier, print the list of directories for each process. See “directory” on page 6-84, for the command used to set the search path.

info convenience

Describe convenience variables.

info convenience

This command describes all the convenience variables that have been defined. Convenience variables may be global or process local (see “set-local” on page 6-68). This command first describes the global variables, then (for each process specified by the command qualifier) describes the process local variables. The name, data type, and value of each variable is listed.

The convenience variables that correspond to the process registers are not described by this command (see “info registers” on page 6-168).

info display

Describe expressions that are automatically displayed.

info display

This command describes the set of expressions that are automatically displayed each time a program stops (see “display” on page 6-101).

info history

Print value history information.

info history [*number*]

number

Specifies an item in the value history list (each value has a unique sequence number). The default value is the most recent history list entry.

This command prints ten history-list values centered around the specified entry. It also prints information about how many history items currently exist. See “set-history” on page 6-64.

info limits

Print information about limits on expression and location output.

info limits

The command prints the limits on array elements and character-string elements printed by expression output commands, and the limits on program locations printed by other **info** commands. See “set-limits” on page 6-64.

The qualifier is ignored by this command.

info registers

Print information about registers.

info registers [*regex*]

regex

A regular expression matching register names. An anchored match *is* implied. See “Regular Expressions” on page 6-19.

If the *regex* argument is not given, this command prints all the normally accessible registers that are of general interest to most programmers (such as accumulators, program counter, stack pointer, etc.). If you give a regular expression argument, any register with a name matching that regular expression is printed. To print *all* the registers, you must specify the regular expression `.*` as an argument (this includes all the obscure control registers and any other registers not normally of interest to a programmer). See “Predefined Convenience Variables” on page 6-5.

The regular expression case sensitivity depends on the current search mode (see “set-search” on page 6-73).

Registers are printed relative to the current frame (see “Current Frame” on page 3-29). This means that any register saving is logically unwound as you change frames (the register contents are not actually modified). You see the value the register would have if you returned to the current frame. (If the current frame is also the most recent frame at the end of the stack, the current machine register contents are the correct contents relative to frame zero.)

If the current frame is not frame zero, but you want to see the current active contents of the machine registers, you have to move to frame zero before running the **info registers** command (see “frame” on page 6-147).

If the command qualifier names multiple processes, the registers from each process are printed separately. If any of the processes are running, an error is printed.

Since this command operates only on register names, the dollar sign (\$) normally used to refer to registers is optional for this command.

Some registers are defined by the architecture to be composed of various fields. **info registers** expands those fields symbolically. If a field is a single bit, NightView prints an abbreviation for that field only if the value of the field is 1. See the architecture manual for descriptions of the fields and a list of the abbreviations for each register.

info signal

Print information about signals.

info signal [*signal* ...]

signal

A signal number or signal name.

This command describes how signals will be handled by the process receiving them. If the command qualifier specifies multiple processes, then the signal information is listed separately for each process. The information printed includes:

- The signal name.
- The signal number.
- The way the debugger will handle this signal. (see “handle” on page 6-144).

If no *signals* are specified, then information for all signals is printed.

info process

Describe processes being debugged.

info process

This command lists information about all the processes specified in the command qualifier (qualify with *all*) to list all of them). The information includes:

- The process ID (PID).
- The controlling dialogue for the process.
- The arguments passed to the program on startup (*argv* array).
- The current process state (running, running one thread, stopped).
- When the process state is stopped, list where and why it stopped.
- The current language setting. See “set-language” on page 6-63.

- The disposition of child processes; that is, under what circumstances a child process will be debugged. See “set-children” on page 6-52.
- If the current source file resides on a remote system, its name is preceded by *user@host*. Local source files will not have that designation. If the current source file is different than that recorded in the program’s debug information, both source paths are listed.
- The run-mode of the process (see “set-run-mode” on page 6-131).

info memory

Print information about memory, which may include information about the virtual address space, or the heap.

```
info memory [/ranges] [/heap] [/leaks] [/allocated]
               [/all] [/append=filename] [/output=filename]
               [/verbose] [expression]
```

/ranges

If this option is specified, the command prints information about the virtual address space. If an expression is specified, then it should evaluate to an address, and only information about the region that contains that address is displayed. If no expression is specified, then all regions in each process are displayed. For each region of memory displayed, this command displays the following information:

- The beginning address and ending address of the region.
- The size, in bytes, of the region.
- The filename, if the region is backed by a file (e.g. a shared library).
- Whether the region is readable, writable, executable, shared, or locked in physical memory.
- Indications if the region is part of the process' stack, memory heap, fast system call library (**vdso** and/or **vsyscall**), kernel **vvar** variables, etc.
- If the region was created by NightView, what the region is for and how much space is left in the region. See “Implementation Overview” on page C-1. If the `/verbose` option is specified, NightView prints information about the individual blocks allocated in the region.

The list also includes any regions reserved by the user with the **mreserve** command. See “mreserve” on page 6-54.

/heap

If this option is specified, the command prints information about the heap. This option is supported only for processes which have turned on heap debugging.

If an expression is specified, then it should evaluate to an address, and the command displays the following heap information about the memory block that contains the address, if any:

- its state, which will be one of:
 - allocated
 - freed, but retained
 - freed or never allocated, but owned by heap
 - not owned by heap

In the case of the latter two, no further information is displayed.

- the address range of its memory block
- its size in bytes
- descriptions of any errors detected pertaining to it
- information pertaining to each heap operation (allocation, most recent `realloc`, and `free`) that has happened for the block, including this information for each:
 - number and address range of post-fence bytes and the post-fence fill byte
 - number and address range of pre-fence bytes and the pre-fence fill byte
 - number of slop bytes
 - whether free filling was enabled and the free fill byte
 - whether malloc filling was enabled and the malloc fill byte
 - whether hardware overrun protection was enabled
 - walkback of stack frames at the time of that operation, as restricted by the walkback setting at that time

If no expression is specified, the command displays the following global heap information:

- totals, including:
 - number of blocks ever allocated
 - number of bytes ever allocated
 - number of additional bytes of debugger overhead ever allocated
 - number of blocks ever freed
 - number of bytes ever freed
 - number of additional bytes of debugger overhead ever freed
 - number of blocks currently allocated

- number of bytes currently allocated
 - number of additional bytes of debugger overhead currently allocated
 - number of blocks currently freed but still retained
 - number of bytes currently freed but still retained
 - number of additional bytes of debugger overhead currently freed but still retained
- whether heap debugging is on or off
 - number of post-fence bytes and the post-fence fill byte
 - number of pre-fence bytes and the pre-fence fill byte
 - number of slop bytes
 - whether free filling is enabled and the free fill byte
 - whether malloc filling is enabled and the malloc fill byte
 - whether hardware overrun protection is enabled
 - frequency of automatic heap checks (i.e. the number of heap operations between automatic heap checks)
 - maximum heap size, which may be "unlimited"
 - maximum number of retained free blocks, which may be "unlimited"
 - maximum number of walkback frames per heap operation
 - whether or not to check fill bytes of free blocks

`/leaks`

If this option is specified, the command prints information about heap blocks which very likely have leaked. See “Leak Detection” on page 3-38 for accuracy limitations on leak detection. An expression may not be specified with this option. This option is supported only for processes which have turned on heap debugging.

If the `/all` option is specified, then all leaks will be displayed. Otherwise, only new leaks since the last leak report will displayed.

If the `/verbose` option is omitted, then heap blocks are reported as sets. A set contains all heap blocks with identical sizes and walkbacks at the time of their allocations (or most recent `reallocs`), regardless of other characteristics. For each set, the following information is reported:

- number of blocks in the set
- block size of the blocks in the set
- walkback of stack frames, as restricted by the walkback setting at that time of the allocation (or most recent `realloc`) operation, of the blocks in the set

If the `/verbose` option is specified, then each heap block is reported individually. For each block, the following information is reported:

- its beginning address
- its size in bytes
- walkback of stack frames at the time of the allocation (or most recent `realloc`) operation, as restricted by the walkback setting at that time

`/allocated`

If this option is specified, the command prints information about heap blocks which are still allocated. An expression may not be specified with this option. This option is supported only for processes which have turned on heap debugging.

If the `/all` option is specified, then all blocks still allocated will be displayed. Otherwise, only blocks still allocated and allocated since the last still allocated blocks report will displayed.

The format of the output is identical to that for the `/leaks` option. The only difference is that all allocated blocks are reported, instead of only those which are determined to be leaks.

`/all`

This option is meaningful only with the `/leaks` or `/allocated` options. See the descriptions of those options for its effect.

`/append=filename`

Write the output of the command to the specified *filename*, appending to any existing contents.

`/output=filename`

Write the output of the command to the specified filename, replacing any existing contents.

`/verbose`

Indicates that extra information should be printed.

This command displays information about the virtual address space, or about the heap, for each process specified in the command qualifier. If no options are specified, the default behavior is like that for the `/ranges` and `/heap` options.

info dialogue

Print information about active dialogues.

info dialogue

This command lists information about all the dialogues specified in the command

qualifier (qualify with `(all)` to list all of them). The information includes:

- The machine running the dialogue.
- The sizes that will be used for patch areas created in the future. See “set-patch-area-size” on page 6-69.
- The list of **debug** and **nodebug** patterns for this dialogue. See “debug” on page 6-27.
- The processes being debugged under control of the dialogue.
- The user running the dialogue.
- The status of any dialogue output (see “set-show” on page 6-35).
- The list of object filename translations for this dialogue. See “translate-object-file” on page 6-29.

info family

Print information about an existing process family.

info family [*regex*]

regex

A regular expression matching family names. An anchored match *is* implied. See “Regular Expressions” on page 6-19.

For each family name that matches *regex* this command lists each process that is a member of that family (see “family” on page 6-51). If *regex* is omitted, then the contents of all process families are printed.

The regular expression case sensitivity depends on the current search mode (see “set-search” on page 6-73).

info name

Print information about an existing eventpoint-name.

info name [*regex*]

regex

A regular expression matching eventpoint-names. An anchored match *is* implied. See “Regular Expressions” on page 6-19.

For each eventpoint-name that matches *regex*, this command lists each eventpoint that is a member of that eventpoint-name (see “name” on page 6-107). If *regex* is omitted, then the contents of all eventpoint-names are printed.

The regular expression case sensitivity depends on the current search mode (see “set-search” on page 6-73).

Each eventpoint is identified by a dialogue-name, a process-id (PID), and an eventpoint-id that is unique for that process.

info on dialogue

Print **on dialogue** commands.

info on dialogue [*name*]

name

The name of a prospective dialogue.

If no arguments are given, then all existing **on dialogue** commands are printed. If a dialogue name is given, then only the **on dialogue** commands that would be executed if a dialogue named *name* were to be created are printed. See “on dialogue” on page 6-31

info on program

Print **on program** commands.

info on program [*program*]

program

The path name of a prospective executable file.

If no arguments are given, then **info on program** prints all existing **on program** commands for each dialogue specified by the qualifier. If a program path is given, then **info on program** prints the **on program** commands that would be executed if *program* were run in each dialogue specified by the qualifier. See “on program” on page 6-47.

info on restart

Print **on restart** commands.

info on restart [output=*outname* | append=*outname*] [*program*]

output=*outname*

Write the information to *outname*.

append=*outname*

Append the information to *outname*.

program

The path name of a prospective executable file.

If no *program* is given, then **info on restart** prints all existing **on restart** commands for each dialogue specified by the qualifier. If a *program* path is given, then

info on restart prints the **on restart** commands that would be executed if *program* were run in each dialogue specified by the qualifier. See “on restart” on page 6-49.

If no *outname* is specified, then the output is to the terminal or to the GUI message area.

info on restart may be used to preserve restart information in a file for use in a later debug session. See “source” on page 6-153. See “Restarting a Program” on page 3-20. For an example, see “checkpoint” on page 6-50.

info exception

Print information about exception handling.

info exception/ada

info exception/ada *exception-name...*

info exception/ada *unit-name*

info exception/cpp

info exception/cpp *type-name...*

info exception *exception-name...*

info exception *unit-name...*

info exception

Abbreviation: **exception**

exception-name

Specifies the name of a particular Ada exception.

unit-name

Specifies all Ada exceptions defined in the specified unit.

type-name

Specifies the name of a C++ type used as an exception.

This command describes the current Ada or C++ exception handling settings for the processes specified by the qualifier. See “handle” on page 6-144. In either case, the command will list:

- The exception name, or the keyword `all` for Ada, or `*` for C++, denoting the default.
- The exception handling settings.

If the `/ada` option is specified, Ada exceptions are displayed. With no arguments, the current default handling of exceptions (`all`) is displayed along with the handling of any specific Ada exceptions to which the default is not applicable. If an argument is given,

the handling of only those specific Ada exceptions is displayed.

If the `/cpp` option is specified, C++ exceptions are displayed. With no arguments, the current default handling of exceptions (*) is displayed along with the handling of any specific C++ exceptions to which the default is not applicable. If an argument is given, the handling of those specific C++ exceptions is displayed.

If neither option is specified, the behavior varies. With no arguments, both Ada and C++ exceptions are displayed. If an argument is given, it is interpreted as an Ada exception name, for backward compatibility with older versions of NightView.

info threads

Describe C threads and thread processes.

```
info threads [/verbose]
                [/cuda] [/partition { physical | logical |
                                     hierarchy_logical | pc }]
```

```
/verbose
```

Show all thread tags for each thread, using a multi-line display. By default, only values for non-zero simple thread tag types are shown.

```
/cuda
```

Show all threads within any CUDA contexts.

```
/partition physical | logical | hierarchy_logical | pc
```

If showing all threads within a CUDA context, specify the partitioning method used to group the threads. The `physical` method organizes the threads by physical characteristics of the CUDA device: SM, warp, and lane. The `logical` method organizes the threads by grid, then by block coordinates and then by thread coordinates. The `hierarchy_logical` method is similar to `logical`, but also organizes the grids into a parent-child hierarchy. The `pc` method organizes the threads into groups that have common values of the program counter (\$`pc`).

This command describes the C threads and thread processes for the processes specified by the qualifier. If CUDA code is present in the application, it also describes CUDA contexts and optionally the threads therein. The identifiers listed for each thread type may be used with the **select-context** command to switch to that thread. See “select-context” on page 6-149.

The current thread is marked with the leading characters `=>`.

If only one thread is running (see “set-run-mode” on page 6-131) then the leading characters `=>` are followed by the word `Running` for the running thread.

For each host thread, the following information is printed:

- The process ID (PID) for the thread. This is the value that would be returned by the `gettid()` system call (which is not available as a library call). It is also the value you would see for the PID in a **ps (1)** listing if

you asked to see thread information. For Linux threads, all threads share the same value as returned by `getpid(2)`, but each thread has its own `gettid()` value.

- The `pthread_t` value for the thread, which is assigned by **`pthread_create(3)`** and returned by **`pthread_self(3)`**.
- The name of the start routine of the thread, if NightView is able to locate it. This is the name of the user's function whose address is passed to **`pthread_create(3)`**.
- A list of all simple **Thread Tags** which have non-zero values. See "Thread Tags" on page 3-45 for information on setting user-defined, thread-specific values. Only tags with simple types and those which are similar to a `C char[]` are shown by default; to see all thread tag values, use the `/verbose` option.

In the graphical user interface, you can also see this thread information from the Display menu. See "Data Menu" on page 8-16.

info timing

Display timing regions and their start and end locations.

info timing [`/verbose`]

`/verbose`

Specify that the start and end locations displayed will be in verbose format. Verbose location format includes the program counter address (or addresses) of the location and, where possible, the corresponding function name, file name, and line number. The number of PC addresses printed is subject to the limit on printing addresses (see "set-limits" on page 6-64). The verbose keyword may be abbreviated.

This command displays each named timing region. Indented below each timing region are the set of start and end locations defined for that region.

heapcheck

Check the heap for errors.

heapcheck [`/all`] [`/append=filename`] [`/output=filename`] [`expression`]

`/all`

Report all existing heap errors. Without this option, the only errors shown are the ones that have occurred since the most recent **heapcheck** command, automatic heap check performed by the process during heap operations (see "Heap Check" on page 3-37), or `heappoint` check (see "heappoint" on page 6-117).

/append=filename

Write the output of the command to the specified filename, appending to any existing contents.

/output=filename

Write the output of the command to the specified filename, replacing any existing contents.

The **heapcheck** command checks the heap for errors, such as overwritten fences, for each process in the qualifier. See “Fences” on page 3-35. If an expression is specified, then it should evaluate to a heap address, and only the block that contains that address is checked. If no expression is specified, then all heap blocks are checked. Heap debugging must have been turned on already, via the **heapdebug** command (see “heapdebug” on page 6-56) or the **Debug Heap...** item in the **Process** menu (see “Process Menu” on page 8-9).

The output is identical to error reporting when heap errors are discovered by automatic heap checks, or by a heappoint check. Possible errors are:

- post-fence modified in allocated block (value=address)
- pre-fence modified in allocated block (value=address)
- free-fill modified in free block (value=address)

Symbol Table Information

The info commands in this section are used to lookup and report on information recorded in the debug tables of program files. This includes the names and declarations of variables, the address of generated code for source lines, etc.

info args

Print description of current routine arguments.

info args

This command prints a description of each argument of the subroutine associated with the current frame (see “Current Frame” on page 3-29).

info locals

Print information about local variables.

info locals [*regex*]

regex

A regular expression matched against local variable names. An anchored match *is* implied. See “Regular Expressions” on page 6-19.

Print a description of every local variable visible in the current context. If the *regex* argument is given, print only the variables with names matching the regular expression.

The regular expression case sensitivity depends on the current search mode (see “set-search” on page 6-73).

The term *local variables* is defined to include all variables with any sort of restricted scope. External variables visible throughout the program are never listed by this command.

The information listed for each variable includes:

- The name of the variable.
- The type of the variable.
- The current value of the variable.
- The location of the variable.
- The scope of the variable (directly visible, inherited from an outer block, etc.).

info variables

Print global variable information.

info variables [*regex*]

regex

A regular expression matched against global variable names. An anchored match *is* implied. See “Regular Expressions” on page 6-19.

This command prints information about global variables. When the *regex* argument is given, it prints only variable names matching the regular expression.

The regular expression case sensitivity depends on the current search mode (see “set-search” on page 6-73).

info address

Determine the location of a variable.

info address *identifier*

identifier

The name of the variable to be described.

Print out information about where the given variable (visible in the current context) is located. If the variable is in a register, it prints the register name. If it is on the stack, it prints the stack frame offset. If it is in static memory, it prints the absolute location.

To determine the absolute address of a particular instance of a stack variable you must

use the **print** command to evaluate an expression which returns the address (for the C language, this would be something like **print &name**, see “print” on page 6-91).

info sources

List names of source files.

info sources [/v] [*pattern*]

pattern

Wildcard pattern to match against source file names. See “Wildcard Patterns” on page 6-20.

This command lists the names of the source files recorded in the debug tables. If a wildcard pattern is given, it lists only file names matching the wildcard pattern.

If the /v option is supplied, it lists the full pathnames of the files as recorded in the debug tables.

If multiple processes are specified in the command qualifier, the source files for each process are listed separately.

info functions

List names of functions, subroutines.

info functions [*regex*]

regex

A regular expression to match against function names. An anchored match is implied. See “Regular Expressions” on page 6-19.

This command lists the names of functions or subroutines recorded in the debug tables. If a regular expression is given, it lists only names matching the regular expression.

The regular expression case sensitivity depends on the current search mode (see “set-search” on page 6-73).

info types

Print type definition information.

info types [*regex*]

regex

A regular expression to match against type names. An anchored match is implied. See “Regular Expressions” on page 6-19.

This command prints information about type definitions. When the *regex* argument is given, it prints only type names matching the regular expression; otherwise, it prints all

the types defined in the program.

The regular expression case sensitivity depends on the current search mode (see “set-search” on page 6-73).

info whatis

Describe the result type of an expression visible in the current context.

info whatis *expression*

Abbreviation: **whatis**

expression

An expression for which the data type is to be determined. See “Expression Evaluation” on page 3-25.

Describe the result type of the expression. The expression is not normally evaluated, but operations which require run time type determination may require portions of the expression to be evaluated. If the expression includes the C++ `dynamic_cast<>` function, the operands must be evaluated in order to determine the actual type of the result.

info representation

Describe the storage representation of an expression.

info representation *expression*

Abbreviation: **representation**

expression

An expression for which the data type is to be determined. See “Expression Evaluation” on page 3-25.

Describe the storage representation of the result type of the expression. The expression is not evaluated.

info declaration

Print the declaration of variables or types.

info declaration *regex*

Abbreviation: **p_type**

regex

A regular expression to match against type names and variable names. An anchored match *is* implied. See “Regular Expressions” on page 6-19.

The *regex* parameter may specify type or variable names visible in the current context.

This command prints the complete declaration of all matching names.

The regular expression case sensitivity depends on the current search mode (see “set-search” on page 6-73).

info files

Print the names of the executable, symbol table and core files.

info files

For each process specified in the command qualifier, print the names of the executable file, symbol table file, and core file associated with the process (the executable and symbol table files are usually the same). If the file resides on a remote system, its name is preceded by *user@host*. Local files will not have that designation.

info line

Describe location of a source line.

info line [*at*] *location-spec*

location-spec

Query the source line number associated with this location.

Describe the location of the source line implied by the *location-spec* argument (see “Location Specifiers” on page 6-15). The information printed includes:

- The address of the *location-spec*.
- The ranges of addresses occupied by the generated code for the line. The number of address ranges printed is subject to the current limit on addresses (see “set-limits” on page 6-64). If this is the first range for an instance of the line, the address range is preceded by *@*.
- The source file and line number.
- The function containing the line.

This command sets the default **x** command dump address as well as the `$_` predefined convenience variable to the address of the first instruction in the line. See “x” on page 6-95 and “Predefined Convenience Variables” on page 6-5.

info debuginfo

info debuginfo

For the executable being debugged, and for each shared object used by that executable, it describes whether or not it contains debug information. If the executable or shared object has separate debug information, it provides the pathname for the file that contains it.

Note that this does not guarantee that the debug information for these files is complete. It is possible that an executable or shared object was linked with a mixture of object files, some containing debug information, and some not.

Defining and Using Macros

NightView provides a macro facility so you can augment the NightView commands with your own features. Macros can either be used as part of another command, or as a new command.

A *macro* is a named set of text, possibly with arguments, that can be substituted later in any NightView command. The arguments allow macros to expand to different text in different circumstances. Macros are useful in extending the command set available in NightView; they can also serve as shortcuts for frequently used constructs in commands or expressions.

define

Define a NightView macro.

```
define macro-name [ (arg-name [, arg-name] ...) ] [text]
```

```
define macro-name [ (arg-name [, arg-name] ...) ] as
```

macro-name

This is the name of the macro. Macro names follow the usual rules for identifiers in most languages: they must begin with an alphabetic character, followed by zero or more alphanumeric characters or underscore. There is no limit to the length of a macro name.

A macro name can be the same as a NightView command name, but this may render the command unusable. See “Referencing Macros” on page 6-187 for more information.

arg-name

A *formal argument* name. These names follow the same rules as *macro-name*.

text

The text to be substituted when the macro is invoked. In this form, the substituted text will not contain any newline characters, so the *text* becomes part of whatever command the macro invocation appears in.

NOTE

There must not be any blanks separating the *macro-name* from the left parenthesis that introduces the formal arguments.

In the second form of the **define** command, the text of the macro begins on the line following the **define** command and extends until a line containing only the words **end define** is encountered. Except for the newline character immediately following the **as** keyword and the newline immediately preceding the **end define** command, the newline characters within the body of the macro will be retained in the substituted text. Thus, each line of text in the macro body must normally be a complete NightView command.

Comments appearing in the body of the macro become part of the body. Thus, they appear in the text that is substituted for a reference to the macro. You should avoid having a comment as the last line of a macro, because it may cause any text following the macro invocation to be ignored.

In the command-line and simple full-screen interfaces, the prompt changes to **>** while you are entering the second form of the **define** command. (See “Command Syntax” on page 6-1.)

The **define** command associates a body of text with the given *macro-name*. When the macro is invoked (see “Referencing Macros” on page 6-187), the macro name and its actual arguments are replaced by the associated text. The text of the macro, called the *macro body*, may contain references to other macros (in particular, they will want to reference their formal arguments). A macro may not reference itself, either directly or indirectly; that is, macros cannot be recursive.

Within the body of a macro, each *arg-name* becomes a macro without arguments that expands to its corresponding actual argument. “Referencing Macros” on page 6-187 describes the syntax of macro invocations and actual arguments.

A macro body should not contain another **define** command.

The **define** command ignores any qualifier supplied for it.

If the given *macro-name* was previously defined as a macro, the new definition replaces the old one. If you omit the *text* in a one-line definition, or the **end define** command appears on the line immediately following the **define...as** command, any prior definition of *macro-name* is removed.

Examples:

```
(local) define printhex(str,x) printf "The value of %s is 0x%x\n",
@str, @x
```

The above example defines a macro that prints a descriptive string and the value of an arbitrary variable, using the **printf** command.

```
(local) define advance(p) as
>     set @p = @p->next
>     print *@p
>     end define
```

The preceding example defines a macro that advances a pointer to the next item in a linked list, then prints the item. Note that this macro requires the language context to be C or C++, but the type of the argument pointer can be a pointer to any structure that contains an appropriately-typed field named "next".

```
(local) define short (VERY_LONG_NAME(INDEX*2,INDEX-1)*SOME_CONSTANT)
```

This example simply defines a shorthand for a long Fortran expression. Note that it does not have any arguments; the parentheses surround the substituted text to make sure that precedence of operators is preserved when the macro is invoked.

Referencing Macros

Macros are usually referenced by preceding the macro name with the `@` character, and following the macro name with a parenthesized list of arguments, if the macro was defined with arguments. If you wish, you may enclose the macro name inside of `'{'` and `'}'` (but any argument list must appear *outside* of the braces). The number of arguments you supply must be the same as the number of formal arguments (i.e., the *arg-names*) specified in the **define** command; otherwise, NightView issues an error. Arguments are matched with each formal argument name by position.

A reference to a macro without any arguments consists solely of the `@` character followed (without intervening blanks) by the macro name. A reference to a macro with one or more arguments consists of the `@` character, the macro name, and a list of actual arguments. The actual arguments begin with a left parenthesis and end with a matching right parenthesis. If more than one argument is given, a comma must separate them. If an actual argument contains a left parenthesis, then the argument extends until a matching right parenthesis is encountered, irrespective of any other characters, including commas, in the intervening text. Note that an unmatched right parenthesis appearing in an actual argument prematurely ends the list of actual arguments; this may cause an error, or it may produce unexpected results.

An actual argument may contain an invocation of another macro; that invocation is expanded immediately when the actual argument is read during the processing of the enclosing macro invocation. This can lead to some surprising results, because NightView expands these actual arguments without regard to the context in which they will ultimately appear.

For example:

```
(local) define abc xyz
(local) define printit(x) printf "The value is %s\n", @x
(local) printf "The value is %s\n", "@abc"
(local) @printit("@abc")
```

The **print** command will print "The value is `@abc`", because macros are not normally expanded within string literals. However, the `@printit` command will print "The value is xyz", because NightView expands the macro `@abc` when it is processing the invocation of macro `@printit`. At that time, it does not know that the double quotes imply a string literal.

String literals as actual arguments can cause other problems as well. For example:

```
(local) define mymac body_does_not_matter
(local) # Illegal reference:
(local) @mymac("This has a left-parenthesis(")
(local) # Okay:
(local) @mymac("This has two parentheses(")
```

The first invocation of **mymac** is invalid because the actual argument contains an unmatched left parenthesis. Since NightView attempts to balance parentheses without regard to any other text (including quotes), the right parenthesis matches the left parenthesis in the argument, leaving the argument list without a closing right parenthesis.

If a macro invocation appears where a command keyword is expected, then you can leave off the `@` prefix character (but the macro name may *not* be enclosed between `'{'` and `'}'`). This allows macros to be used conveniently as command shortcuts. However, if the macro requires arguments, these must still be placed within parentheses after the macro name.

Macros take precedence over commands when the macro name appears in place of a command keyword. This means that if you name a macro the same as a built-in NightView command, you may not be able to reference the built-in command anymore. However, you cannot abbreviate the macro name in an invocation, so you may be able to use an abbreviation for the built-in command. If you name a macro the same as a built-in command abbreviation, you won't be able to use that particular abbreviation for the built-in command later, but you can still use the full form, or a different abbreviation. If you accidentally name a macro the same as a built-in command, you can remove the definition by entering

```
(local) # Note, no text given in definition.
(local) define macro-name
```

You may want to refer to the Summary of Commands (see Appendix A [Summary of Commands] on page A-1) for a complete list of the NightView commands, so you can avoid these kinds of conflicts.

Macro references can generally appear anywhere within a NightView command, but you should be aware of the following rules:

- NightView never expands macros that appear within command comments.
- NightView usually does not expand macros that appear within string literals. However, if the literal appears as an actual argument in another macro invocation, macros within the string literal may be expanded.
- Macros are not expanded in the *format-string* argument to the **printf** command. See “printf” on page 6-104.
- Macros appearing in an **echo** command are expanded. See “echo” on page 6-99.
- Macros appearing in a **!** (see “!” on page 6-34), **run** (see “run” on page 6-38), or **shell** (see “shell” on page 6-152) command are not expanded.
- A macro referenced within a language expression must expand to text that makes sense as part of that expression.
- A macro can be used to form part of a syntactic item, or token, in a NightView command. For example, you could form a variable name in an expression from the results of two macro invocations. However, you cannot use this technique to construct the name of a macro to be invoked.
- If the `@` character is preceded by a non-whitespace character, the macro is expanded only if the macro name is enclosed between `'{'` and `'}'`.

Examples:

```
(local) define short (VERY_LONG_NAME (INDEX*2, INDEX-1) *SOME_CONSTANT)
(local) set $x=i + @{short}*10
```

The above example uses a macro in an expression.

```
(local) define printhex(str,x) printf "The value of %s is 0x%x\n", @str, @x
(local) printhex("ptr1", ptr1)
(local) printhex("ptr1->next", (ptr1=ptr1->next, ptr1))
```

This example invokes the macro 'printhex' twice. The second invocation demonstrates how an expression containing a comma can be included as a formal argument.

```
(local) directory user@{host}:/mydir # a macro invocation
(local) directory user@host:/mydir # not a macro invocation
```

In this example, the @ character is preceded by "r", so you must use '{' and '}' to have the macro expanded.

The following C fragment defines some data types for use in the next example:

```
struct list_element {
    struct list_element * next ;
    struct data          * the_data ;
};
extern struct list_element * hd ;
```

Example NightView commands:

```
(local) define printdata(p) as
>         printf "The data is:\n"
>         print *(@p)->the_data
>         end define
(local) define next(p) as
>         set @p = (@p)->next
>         end define
```

info macros

Print a description of one or more NightView macros.

info macros [*regex*]

regex

A regular expression matching macro names. An anchored match *is* implied. See "Regular Expressions" on page 6-19.

If the *regex* argument is not given, the **info macros** command prints a description of every macro you have defined. If you give a *regex* argument, a description of every macro whose name matches the regular expression is printed.

The regular expression case sensitivity depends on the current search mode (see "set-search" on page 6-73).

The description of each macro includes:

- The name of the macro.
- The formal argument names, if any, of the macro.

- The macro body text, exactly as it will appear when substituted, except that the last line of the macro will be followed by a newline.

Smart Printing

NightView supports *smart printing*, which is the capability of recognizing certain complex data types and presenting them in a simpler conceptual form that hides the details of their implementation. See “Smart Printing” on page 3-42.

A *smart printer* is a definition that matches types based on a pattern (which also can simply be a string). For any type that it matches, attempts to display an object of that type will be based on the content of the smart printer. There are several types of smart printers:

- **replace**, which replaces the entire content of the object with the result of a user-defined expression (most likely based on the original object).
- **struct**, which is used for struct types and which allows some fields to be hidden, others to be displayed with particular formats, and allows control over which fields are automatically expanded or collapsed in the data panel.
- **container**, which usually is used for container types (e.g. lists, vectors, sets). It defines a number of convenience variables which instruct NightView how to iterate through elements of the container. Objects of these types are displayed as abstract containers with each of the elements determined from the convenience variables.

Smart printers are associated with a single process. Different processes or applications with totally different types that happen to have the same names can have completely different smart printers for those types.

Predefined smart printers are provided with NightView for many C++ STL container types, and for many Qt types (see “Predefined Smart Printers” on page 6-194).

smart-print

Define, undefined, view, enable, or disable smart printers.

```
smart-print info [pattern]

smart-print { [ on | off ] }

smart-print replace pattern
    replace-definition
end-smart-print

smart-print struct pattern
    struct-definition
end-smart-print

smart-print container pattern
    container-definition
end-smart-print

smart-print undef pattern

smart-print reload

    pattern
```

This is a pattern used to identify types to which this smart printer applies. It may be a simple string which matches only a type with the same exact name, or it may contain the * wildcard character which matches zero or more characters in the type name. Because smart printers are used so often on C++ template types, the * wildcard is required to have balanced < and > characters in the matching substring (where no < or > character at all also is considered balanced).

struct-definition

replace-definition

container-definition

These definitions are explained in subsequent sections.

With the `info` keyword, the command displays the pattern string for each smart printer defined for the process. If a *pattern* is specified, it displays only the smart printer that matches that pattern.

With the `on` or `off` keyword, all smart printing can be enabled or disabled for the process. This is a convenient way to disable smart printing temporarily to actually look at the complicated details of a type that has a smart printer. While `off`, all smart printers remain defined but are disabled. When `on`, those smart printers function normally. The default mode is `on`.

With the `struct`, `replace`, or `container` keywords, new smart printers are defined. They match the *pattern* as specified in the command. The syntax and meanings of the definitions are explained in subsequent sections.

With the `undef` keyword, an existing smart printer with the given pattern is removed from the process.

With the `reload` keyword, NightView re-loads all the predefined smart printers into the current process. It has no effect on smart printers defined by the user, so long as their patterns do not conflict with the reloaded patterns.

replace Smart Printers

A replace smart printer can be used to modify the display of any object. It replaces the entirety of the object with the result of a single expression. The *replace-definition* contains at most one of each of the following lines:

```
replacement expression
format format-code
```

The `replacement` line is required and provides the expression which replaces the object. That expression is evaluated and its result will be displayed in lieu of the original object. Within that expression, the original object can be referenced using the special `$val` convenience variable.

The `format` line is optional and can be used to specify the format code that should be used to display the replacement expression.

struct Smart Printers

A struct smart printer is used to modify the display of a struct object. The *struct-definition* in such a smart printer is a sequence of lines of the form:

```
self action
field name action
default action
```

The set of possible actions is:

```
hide
show
format format-code
expand
collapse
```

The `self` keyword describes the smart printer modification applied to the whole struct. Likely the most common action here is `expand`, so that it always is expanded by default in the data panel. The `self` keyword should be used at most once in a smart printer definition.

The `field` keyword describes the smart printer modification applied to the field *name*.

The `default` keyword describes the smart printer modification applied to every field in the struct that is not otherwise explicitly named by a `field` definition.

The `hide` action causes the given object or field to be hidden from display. The `show` action causes the given object or field to be shown.

The `format` action causes the given object or field's format to be changed from the default to the specified *format-code*.

The `expand` action causes the given object or field to be expanded automatically in the data panel. The `collapse` action causes it to be collapsed automatically in the data panel.

container Smart Printers

A container smart printer is used to describe container types. Objects with these smart printers are described as abstract containers with a sequence of elements. The smart printer describes a number of convenience variables which instruct NightView how to iterate through the elements of the container.

The *container-definition* in such a smart printer is in two parts. Its syntax is:

```
initialize-line
...
iterate
iterate-line
...
```

The first part is a sequence of *initialize-lines*. Each line declares and initializes a user-defined convenience variable which describes the first element in the container, and which will be reused in the second part. The second part is a sequence of *iterate-lines* which instruct NightView how to iterate from one element in the container to the next element in the container. The names of the convenience variables are entirely user-defined and can be anything the user wishes. There is one convenience variable that is predefined: `$val`, which contains the original object being smart printed.

An *initialize-line* is one of the following:

```
element $variable = expression
```

Declare the *\$variable* convenience variable. It is interpreted by NightView as the value of the first element. In the `iterate` section, it is interpreted as the value of each subsequent element.

```
index $variable = expression
```

Declare the *\$variable* convenience variable. It is interpreted by NightView as the index for the first element. In the `iterate` section, it is interpreted as the index of each subsequent element. For containers which have no index concept (e.g. `std::set`), this should still be defined, but it may be arbitrary such as initializing its value to 0 in the `initialize` section and incrementing it by 1 in the `iterate` section.

```
endflag $variable = expression
```

There should be exactly one line with the `endflag` keyword. It declares the *\$variable* convenience variable, which is interpreted as a boolean value. If it evaluates to `true` (or any non-zero value if it isn't strictly boolean), this tells NightView that the iterator has reached the end of the container and that there are no more elements. In the `initialize` section, if it evaluates to `true`, this tells NightView that the container is empty. If it evaluates to `true` in the `iterate` section, this tells NightView that it

has run out of elements. In any event, no further computations will be performed once this evaluates to `true`.

```
temp $variable = expression
```

Declare the *\$variable* convenience variable with the initial value *expression*. This convenience variable is never used directly by NightView, but it may be used indirectly in any of the *expressions* for the other keywords.

The values of the convenience variables above determine the display of the first element (or lack thereof in the case of an empty container). The display of each subsequent element is determined by evaluating the full set of iterate-lines. An iterate-line is of the form:

```
$variable = expression
```

The *\$variable* convenience variable should be one defined in an *initialize-line*. Their meanings to Nightview already are specified by the appropriate keyword in the *initialize-line* that declared them. That is, the value of the convenience variable declared with `index` will be displayed as the subsequent index. The value of the convenience variable declared with `element` will be displayed as the subsequent value. The value of the convenience variable declared with `endflag` will be used to determine if the end of the container has been reached.

Smart Printing Limitations

Function calls are not allowed in any of the *expressions* in a smart printer definition. However, intrinsic operations which happen to look like function calls are permitted (see “Intrinsics for Smart Printing” on page 6-195).

For container smart printers, attempting to display high-numbered elements requires iterating through all previous elements. Naturally, if the expressions are complex, this can be quite slow, so expect delays. The debugger can be interrupted if this is taking longer than expected.

Predefined Smart Printers

NightView provides predefined smart printers for types from the C++ STL and for types from the Qt toolkit. These smart printers are defined in the following files:

```
/usr/lib/NightView/lib/stl.print  
/usr/lib/NightView/lib/qt.print
```

NightView loads these files automatically. However, it will load any other file whose name is of the form **/usr/lib/NightView/lib/anything.print**. If the user has other smart printer definitions that they want accessible universally on the system, they can be installed in a new file in that same directory.

The predefined set of smart printers includes support for the following STL types:

```
std::auto_ptr<*>
std::basic_string<char,*,*>
std::deque<*,*>
std::list<*,*>
std::map<*,*,*,*>
std::multimap<*,*,*,*>
std::multiset<*,*,*>
std::queue<*,*>
std::set<*,*,*>
std::stack<*,*>
std::unique_ptr<*>
std::vector<*,*>
```

The predefined set of smart printers includes support for the following Qt types:

```
QAtomicPointer<*>
QCache<*,*>
QHash<*,*>
QLinkedList<*,*>
QList<*>
QMap<*,*>
QMultiHash<*,*>
QMultiMap<*,*>
QPointer<*>
QQueue<*>
QSet<*>
QStack<*>
QString
QVector<*>
```

Intrinsics for Smart Printing

A number of intrinsics are implemented in NightView. They actually are available for general-purpose use, but their primary purpose is for use in smart printing expressions:

`__typeof__ (expression)`

Because smart printers can match many different types because of pattern matching and templates, this allows the determination of the type of any expression. It emulates the GNU C `__typeof__` intrinsic.

`__alignof__ (expression-or-type)`

This returns the default alignment for the expression or type. It does not necessarily know about special options that affect alignment passed to the compiler, so it could return a surprising result if any such options were used.

`__nview_iconv__ (buffer, size, fromcode[, tocode])`

This uses the **iconv** library routine to convert a block of data pointed to by *buffer* and of the specified *size* from the character size specified by *fromcode* to the character set specified by *tocode*. If *tocode* is not specified, it converts to the native locale.

The result of this intrinsic is a char array containing the converted bytes. The primary purpose of this intrinsic is to convert data stored in a program in a form inconvenient to read (e.g. **UTF-16**) to a human-readable character array.

```
__nview_vl_rb_tree_increment__ (nodeptr)
```

This is a very specialized intrinsic designed to advance from one element to the next in the red-black trees used by STL map and set containers. It expects that *nodeptr* will be a pointer to a struct with fields `_M_left`, `_M_right`, and `_M_parent` and knows how to advance to the next element.

```
__nview_ne__(array, numelt, eltsize, match)
```

This searches the given *array* and returns the index of the first element that is not equal to *match*. The number of elements in the array is *numelt*, and the size of each element in bytes is *eltsize* (which must be 1, 2, 4, or 8). If every element in the array matches, then it returns -1. This is useful for searching hash tables for non-empty cells.

Timing

timing

Manipulate timing regions, set or delete timing start and end locations, and display timing results.

```
timing create [/crossthread] name
timing delete name
timing rename newname oldname
timing start [/crossthread] name [at] location-spec
timing end   [/crossthread] name [at] location-spec
timing delstart name [at] location-spec
timing delend   name [at] location-spec
timing function name function-name
timing function name
timing clear name
timing [/column ... ]
```

```
/crossthread
```

Used to indicate that a newly-defined timing region is crossthread, which disregards threads. See “Timing” on page 3-16. This option can only be used when first defining a timing region with keyword create, start, or end.

```
name, newname, oldname
```

Specifies a timing region name. See “Timing” on page 3-16.

```
location-spec
```

Specifies the location. See “Location Specifiers” on page 6-15.

function-name

Specifies a function. This is a restricted form of “Location Specifiers” on page 6-15), allowing only forms that include **function_name**.

With the **create** keyword, a new timing region name is defined. Note that it is not necessary to use this form of the command. Any command that mentions a previously-undefined timing region *name* will define it.

With the **delete** keyword, an existing timing region *name* is deleted.

With the **rename** keyword, an existing timing region *oldname* is renamed to *newname*.

With the **start** keyword, a timing start location for the *name* region is defined at the specified location (see “Location Specifiers” on page 6-15).

With the **end** keyword, a timing end location for the *name* region is defined at the specified location (see “Location Specifiers” on page 6-15).

With the **delstart** keyword, an existing timing start location for the *name* region at the specified location is deleted (see “Location Specifiers” on page 6-15).

With the **delend** keyword, an existing timing end location for the *name* region at the specified location is deleted (see “Location Specifiers” on page 6-15).

With the **function** keyword, the *function-name* is analyzed and all entry and exit instructions are set as start or end locations, respectively. If the *function-name* is omitted, the *name* is used instead, doubling as both the region name and the function name.

With the **clear** keyword, all collected times and crossing counts for the *name* region are cleared, as are any errors, warnings or other diagnostics.

If no keyword is specified, the command provides a report of the currently collected timing results. Any number of column options may be specified. The allowed column options are:

/count	The number of times the region was crossed (start to end)
/total	The total time spent in the region
/average	The average time spent in the region
/min	The minimum time for any single crossing of the region
/max	The maximum time for any single crossing of the region
/subevents	Show times including subevents
/nosubevents	Show times excluding subevents
/info	Show errors, warnings & other diagnostics. See “Timing Info” on page 3-17.
/threads	Show data for each thread separately (the default is to combine data from all threads). This has no effect on crosstthread regions, which disregard threads.

When manipulating the start or end locations of a timing region via **start**, **end**, **delstart**, **delend**, or **function**, all collected times for region, if any, are cleared automatically.

Simple Full-Screen Interface

NightView is designed to be able to debug multiple processes asynchronously. That means your processes may be running and producing output or hitting breakpoints, all at the same time. You might be entering NightView commands at the same time as well.

This can be a little confusing. It would be especially confusing if NightView were to write to your terminal at the same time you are trying to enter a command. For this reason, NightView doesn't usually show you output or event notifications while it is reading your commands (It will do that if you want it to, though. See “set-show” on page 6-37.)

This means that NightView may have output or event notifications to show you, but it will not show them to you because it is waiting for you to type a command. You can press carriage return a few times to see output you are expecting, but that can be annoying.

A full-screen interface gives NightView a way to show you output and event notifications as soon as they are available without interfering with your typing.

The simple full-screen interface has the same basic functionality as the command-line interface. All the commands are the same. In fact, the simple full-screen interface looks a lot like the command-line interface. The main difference is that NightView has control over the entire screen, so it can print output to you while you are “at a prompt”.

Using the Simple Full-Screen Interface

To use the simple full-screen interface, you should have your `TERM` environment variable set to the type of your terminal. If you are using a full-screen editor, such as `vi(1)`, you probably have already taken care of this.

Invoke NightView with the `-simplercreen` option:

```
nview -simplercreen
```

NightView clears the screen before it writes its welcome message. Then the prompt is written to the bottom line and you can type a command.

NightView does not have control over the terminal while you are executing a `shell` command, so after the command has completed you are asked to press return. This gives you a chance to view the command output before NightView redraws the screen. See “shell” on page 6-161.

The simple full-screen interface creates a special window when you use monitorpoints. See “Monitor Window - Simple Full-Screen” on page 7-2 for more information about this window.

Editing Commands in the Simple Full-Screen Interface

You can use special key sequences to edit your commands. The key sequences are based on the line editing modes of **ksh (1)**. NightView implements the `emacs`, `gmacs` and `vi` modes of **ksh**. In particular, you can use the various key sequences to retrieve previously entered commands.

The initial editor mode is set from your `VISUAL` or `EDITOR` environment variables. If NightView cannot determine the mode from those variables, then the default mode is `emacs`. You can explicitly set the editor mode with the `set-editor` command. See “set-editor” on page 6-75.

Monitor Window - Simple Full-Screen

The Monitor Window is created when you use monitorpoints while running NightView with the simple full-screen interface. See “Monitor Window” on page 3-36.

In the simple full-screen interface, the Monitor Window appears at the top of the screen and takes up as many lines as it needs for the number of items displayed, plus one status line, while leaving at least ten lines for other debugger operations at the bottom of the screen.

Only the items that fit in the space available at the top of the screen are displayed. Any further items are left in the same state they would be in following an `mcontrol nodisplay` command (See “mcontrol” on page 6-124)

The stale data indicators used in the simple full-screen Monitor Window are simple characters used to indicate each state. A space () is used to indicate updated values. A period (.) is used for monitorpoints that have not been executed. An exclamation point (!) is used for monitorpoints which have executed but not taken a sample. For more information about stale data indicators, see “Monitor Window” on page 3-36.

A status line at the bottom of the simple full-screen Monitor Window divides it from the remainder of the screen. The status line indicates the state of the Monitor Window (`held` or `running`) and shows the current delay time in milliseconds between updates of the window.

Graphical User Interface

This chapter describes the graphical user interface (GUI) for NightView. The GUI provides more flexibility and functionality than either the command-line interface or the simple full-screen interface.

The graphical user interface for NightView is based on the Qt toolkit. NightView runs in the environment of the X Window System™ Version 11, Release 6 (or later).

This chapter assumes that you have a basic understanding of window system concepts such as selecting objects by clicking with the mouse and working with dialog boxes and standard controls. Use mouse button 1 when you are told to click, drag, press, and select.

Sample debug sessions showing how to use the NightView graphical user interface are available. See Chapter 2 [A Quick Start - GUI] on page 2-1. See Chapter 4 [Tutorials] on page 4-1.

NightView GUI Concepts

This section explains concepts that you need to understand so that you can use the NightView graphical user interface to its fullest advantage.

GUI Online Help

The graphical user interface provides several ways of providing help on particular topics.

- Context-sensitive help is available in the main window and all the dialogs. See “Context-Sensitive Help” on page 8-2.
- The main window has a **Help** menu. See “Help Menu” on page 8-21.
- Pressing the **F1** function key displays help for the part of the window that has the current focus.
- The dialog boxes have help buttons that pop up help for the particular dialog box.
- You can use the **help** command from the command-line interface. See “help” on page 6-154.

Help information is displayed in a help window. NightView uses a separate program to display the help window. Once a help window is displayed, you can move around in the help system in a variety of ways. You can keep the help window on your screen, or dismiss it. You can also iconify it, and it redisplay itself the next time you ask for help. See “Help Window” on page 8-124.

Context-Sensitive Help

Context-sensitive help is available through the **Help** menu. See “Help Menu” on page 8-21. In addition, the **F1** function key displays help information for the currently selected window component.

Generally, help is not provided on individual graphical items, such as individual buttons. Instead, you are given help for the region you have selected. For example, if you select help on the Kill button, the help window displays information about the process toolbar. See “Process Toolbar” on page 8-24.

To get context-sensitive help using the **Help** menu, select the **On Context...** menu item. The pointer changes to a question mark with an arrow. Place the point of the arrow over the graphical region for which you want help and click mouse button 1. The help window is displayed with information about that region. The pointer changes back to its original shape.

To get context-sensitive help using the **F1** key, select a window component that you have a question about. Press the **F1** key. A help window is displayed with information about that region.

Help Buttons

Dialog boxes include a **Help** button in the lower right corner. You can click on this button to receive help on the dialog box. See “Dialogues and Dialog Boxes” on page 8-3.

Help Command

You can type the **help** command, followed by the topic you want help on, into the command toolbar to obtain online help. See “help” on page 6-159. A help window is displayed that contains information about the requested topic. See “Help Window” on page 8-124. See “Command Toolbar” on page 8-23.

If a help window does not exist, NightView displays one for you. Otherwise, the text of the existing help window changes to show you the information that you requested.

If NightView cannot find the information you requested, a warning dialog box and a help window are displayed.

Dialogues and Dialog Boxes

NightView has a concept called a *dialogue*, which is a way of communicating with an ordinary command shell. See “Dialogues” on page 3-6. Note that this kind of dialogue is spelled with a “ue” at the end.

The graphical user interface uses another term: *dialog box*. This is not related to the NightView concept of a *dialogue*. *Dialog box* refers to a particular type of window that may appear during your session. A dialog box usually appears only briefly and typically allows you to specify a particular item, such as a file name.

These two concepts are distinct and unrelated, even though they sound alike.

Context Menu

Each panel has a context menu with entries appropriate for that panel type. You show the context menu by right-clicking in the panel.

In the source panel, where you click may affect the operation of the menu items by setting the source panel target line. See “Source Panel Context Menu” on page 8-68 and “Source Panel Target Line” on page 8-67.

In the data panel and related panels, different context menu entries appear depending on where you click. In some cases the menu entry’s operation is directed to the data item on which you click. See “Data Panel Context Menu” on page 8-91.

In the eventpoint panel, you may select one or more eventpoints (rows) and then right-click. If you right-click on a row that is not selected, the selection is cleared and the row you clicked on becomes selected. If you right-click on a row that is selected, the selection does not change. The context menu’s entries are enabled or disabled based on which rows are selected.

Current Process

In the graphical user interface, NightView has the concept of a current process. When you click on toolbar buttons, or enter commands, the operation is performed on the current process. (If you are debugging only one process, that process is the current process.)

The status bar shows the status of the current process. See “Status Bar” on page 8-29. A locals panel shows variables in the current frame of the current process. See “Locals Panel” on page 8-78. If you have more than one process, a context panel shows the current process with green underlined text. See “Context Panel” on page 8-78.

You can switch to a different process by clicking on the other process in a context panel or by clicking on a source panel displaying source for the other process.

GUI Configuration

NightView can save your current GUI configuration. The configuration includes the geometry of the main window, the positions of the toolbars, which pages are present, which panels are present on each page, and the geometry and other information about each panel.

When NightView starts up, it looks for a configuration in the following places, in order.

- filename supplied with the **-config** option
- **.NightView_config** in the current directory
- **\$HOME/.NightView_config**
- **/usr/lib/NightView/lib/config**

You can explicitly load or save a configuration from the **File** menu. See “File Menu” on page 8-5.

Main Window

The main window has a menu bar, toolbars and a status bar. The remaining space is for docking various panels. See “Panels” on page 8-65. NightView can remember the arrangement of the panels. See “GUI Configuration” on page 8-4.

Menu bar

From the menu bar you can perform global NightView actions, perform actions on a shell or a process, choose source to display or edit, manipulate eventpoints, change the way you view the window, select items to display, invoke other NightStar tools, and obtain online help.

File Menu

Mnemonic: **F**

The **File** menu has the following items.

Load Config File...

Mnemonic: **O**

This brings up a file selection dialog to load a configuration. See “GUI Configuration” on page 8-4.

Load System Default Config

Mnemonic: **D**

This loads the default configuration in **/usr/lib/NightView/lib/config**. This is useful if you have modified your configuration and want to return to the default state. See “GUI Configuration” on page 8-4.

Save Config File

Mnemonic: **S**

Accelerator: **Ctrl+S**

If the configuration was loaded from the default, **/usr/lib/NightView/lib/config**, then this saves the configuration to **\$HOME/.NightView_config**. Otherwise, this saves the configuration to the configuration file previously loaded or saved. See “GUI Configuration” on page 8-4.

Save Config File As...

Mnemonic: A

This brings up a file selection dialog to save the configuration. See “GUI Configuration” on page 8-4.

Preferences...

Mnemonic: F

This brings up a dialog that lets you set preferences, such as fonts. See “Preferences Dialog Box” on page 8-46.

Save Preferences

Mnemonic: V

This saves the current preference settings to disk.

Print Window...

Mnemonic: P

This brings up a dialog that lets you print an image of the main window to a printer.

Exit NightView (quit NightView)

Mnemonic: X

Accelerator: Ctrl+Q

Selecting this menu item causes NightView to exit. This has the same effect as the **quit** command. See “quit” on page 6-25.

Depending on the safety level (see “set-safety” on page 6-71), NightView displays a warning dialog box when you click the **Exit** menu item if there are any active processes.

View Menu

Mnemonic: V

The **View** menu lets you create new panels, modify pages, change the size of text in panels, or modify which toolbars are shown.

New Context Panel

Mnemonic: C

Selecting this menu item creates a new context panel. See “Context Panel” on page 8-78.

New Locals Panel

Mnemonic: L

Selecting this menu item creates a new locals panel. See “Locals Panel” on page 8-78.

New Source Panel

Mnemonic: S

Selecting this menu item creates a new source panel. See “Source Panel” on page 8-66.

New Registers Panel

Mnemonic: R

Selecting this menu item creates a new registers panel. See “Registers Panel” on page 8-78.

New Data Panel

Mnemonic: D

Selecting this menu item creates a new data panel. See “Data Panel” on page 8-79.

New Monitor Panel

Mnemonic: M

Selecting this menu item creates a new monitor panel. See “Monitor Panel” on page 8-78.

New Shell Panel

Mnemonic: H

Selecting this item creates a new shell panel. This does not create a new shell. It merely creates a new panel for an existing shell. See “Shell Panel” on page 8-74.

If there is more than one shell, then this item opens a sub-menu from which you can choose which shell the new panel represents.

New Message Panel

Mnemonic: G

Selecting this item creates a new message panel. See “Message Panel” on page 8-75.

New Eventpoint Panel

Mnemonic: E

Selecting this item creates a new eventpoint panel for the current shell. See “Event-point Panel” on page 8-75.

New CUDA Panel

Mnemonic: U

Selecting this item opens a sub-menu that lets you select a particular kind of CUDA-specific panel.

Coordinates

Mnemonic: C

Selecting this item creates a new CUDA Coordinates panel. See “CUDA Coordinates Panel” on page 8-108.

Lanes

Mnemonic: L

Selecting this item creates a new CUDA Lanes panel. See “CUDA Lanes Panel” on page 8-108.

Warp Locals

Mnemonic: W

Selecting this item creates a new CUDA Warp Locals panel. See “CUDA Warp Locals Panel” on page 8-110.

New Memory Segments Panel

Mnemonic: O

Selecting this menu item creates a new memory segments panel. See “Memory Segments Panel” on page 8-111.

New Binary Viewer Panel

Mnemonic: V

Selecting this menu item creates a new binary viewer panel. See “Binary Viewer Panel” on page 8-113.

Add Page

Mnemonic: A

Accelerator: Ctrl+A

Selecting this item creates a new tabbed page.

Rename Current Page...

Mnemonic: R

Selecting this item brings up a dialog box that lets you give the current page a new name. See “Select a Type Dialog Box” on page 8-62. You can create a mnemonic that switches to the page: put an ampersand (&) in front of one of the characters of the name. The page’s tab will have an underscore under that character. To get a real &, use &&.

Delete Current Page

Mnemonic: T

Selecting this item deletes the current page.

Text Size

Mnemonic: Z

Selecting this item opens a sub-menu that lets you select the size of text in the panels.

Increase

Mnemonic: I

Accelerator: Ctrl++

Selecting this item increases the size of the text in the panels. Panels that use a fixed-width font and panels that use a variable-width font are adjusted separately.

Decrease

Mnemonic: D

Accelerator: Ctrl+-

Selecting this item decreases the size of the text in the panels. Panels that use a fixed-width font and panels that use a variable-width font are adjusted separately.

Normal

Mnemonic: N

Accelerator: Ctrl+0 (zero)

Selecting this item resets the size of the text in the panels to normal.

Toolbars

Mnemonic: B

Selecting this item opens a sub-menu that lets you choose which toolbars are shown. See “Toolbars” on page 8-23.

Shell Menu

Mnemonic: L

The Shell menu lets you start a remote shell, terminate a shell, or create a new shell panel.

Start Remote Shell...

Mnemonic: R

Selecting this menu item allows you to create a remote dialogue on a target system of your choice. A dialog box is displayed that allows you to choose parameters for the remote dialogue. See “Remote Login Dialog Box” on page 8-43.

Terminate shell

Mnemonic: T

Selecting this item terminates the dialogue. This is similar to using the **logout** command. See “logout” on page 6-32.

If there is more than one shell, then this item opens a sub-menu from which you can choose the shell to terminate.

Depending on the safety level (see “set-safety” on page 6-71) and whether there are any active processes, NightView may display a warning dialog box when you use the **Terminate shell** menu item.

Show Shell Panel

Mnemonic: H

Selecting this item creates a new shell panel for the current shell. It does not create a new shell. See “Source Panel” on page 8-66.

If there is more than one shell, then this item opens a sub-menu from which you can choose which shell the new panel represents.

New Terminal

Mnemonic: E

Selecting this item executes a new terminal emulation program. This program does not appear within the NightView window, but rather runs in its own distinct window. However, NightView has control of this terminal emulation program and automatically will debug any children that it forks. Because of this, it can be used similarly to a Shell panel. Because it actually is a terminal emulation program, it will respond exactly as expected to ANSI escape sequences, curses, etc. But, I/O in this window will not be reflected in the **Global** panel. Also, the **Rerun Process** button will not function for processes started in this terminal. The user may specify alternate terminal emulation programs (see the “Preferences Terminal Page” on page 8-56).

If there is more than one dialogue shell active, then this item opens a sub-menu from which you can choose which dialogue shell the new panel represents.

Process Menu

Mnemonic: P

This menu is used to perform actions on processes.

Run...

Mnemonic: R

Use this item to run a program. Selecting this item pops up a dialog box for you to enter a shell command line. See “Run Program in Shell Dialog Box” on page 8-33. If your program takes no input and you want to debug more than one program at a time, end the command with `&`.

If there is more than one shell, then this item opens a sub-menu from which you can choose in which shell to run the program.

Another way to run a program is to type a command into a shell panel. See “Shell Panel” on page 8-74.

Attach...

Mnemonic: A

Selecting this item pops up a dialog box you can use to view the processes on the system and attach to one of them. See “Attach Dialog Box” on page 8-33.

If there is more than one shell, then this item opens a sub-menu from which you can choose in which shell to attach.

Detach

Mnemonic: D

Selecting this item causes NightView to detach from the current process. See “Current Process” on page 8-4.

This is similar to using the **detach** command. See “detach” on page 6-42.

Depending on the safety level (see “set-safety” on page 6-71), NightView may display a warning dialog box when you use the **Detach** menu item.

Kill

Mnemonic: K

Selecting this item causes NightView to terminate the current process. See “Current Process” on page 8-4.

This is similar to using the **kill** command. See “kill” on page 6-43.

Depending on the safety level (see “set-safety” on page 6-71), NightView may display a warning dialog box when you use the **Kill** menu item.

Set Timing Start or End Location...

Mnemonic: T

Selecting this menu item pops up a dialog box to allow setting a start or end location for a timing region. See “Set New Timing Start/End Dialog Box” on page 8-63.

Set Timing for a Function

Mnemonic: F

Clicking this button pops up a dialog box that allows you to set timing locations for a function. See “Set Timing for a Function Dialog Box” on page 8-64.

Debug Heap...

Mnemonic: H

Selecting this item pops up the Debug Heap dialog box, which allows you to turn on and adjust heap debugging for the current process. See “Debug Heap Dialog Box” on page 8-41. See “Debugging the Heap” on page 3-37. See “Current Process” on page 8-4.

Process Settings...

Mnemonic: S

Selecting this item pops up the Process Settings dialog box, which allows you to change how NightView treats the current process. See “Process Settings Dialog Box” on page 8-59. See “Current Process” on page 8-4.

Refresh Shared Libs

Mnemonic: L

Selecting this item causes NightView to re-read the list of shared libraries in use by the process and to refresh any debug information that may have changed or is new. This is not normally necessary, however, if your process dynamically loads libraries via **dlopen (2)**, NightView needs to be informed.

You can also set a mode to have NightView automatically detect this; there is some overhead involved however when that mode is set to automatic. See “Debugging with Shared Libraries” on page 3-56 and “set-shared-lib-update” on page 6-56 for more information.

Source Menu

Mnemonic: S

This menu provides ways of changing the program code displayed in source panels and editing source files that are listed. See “Source Panel” on page 8-66.

The items that change which source file is displayed act on all source panels displaying the current process. The items that select source or disassembly are enabled only if there is exactly one source panel, or if there is a source panel that has a target line. See “Source Panel Target Line” on page 8-67.

List Function/Unit...

Mnemonic: U

Selecting this menu item pops up a dialog box that allows you to list the program code of a function in the debug source display. See “Source Panel” on page 8-66.

This dialog box is titled **Select a Function/Unit**. The title bar also displays the process's qualifier specifier. See “Qualifier Specifiers” on page 6-19. It allows you

to optionally enter a regular expression that is used to search for function names that NightView knows about. (An anchored match is *not* implied.) See “Regular Expressions” on page 6-21. For example, enter `set$` to search for function names ending with 'set'. A list of functions is displayed, and one function can be selected for display in the debug source display. For C++, the regular expression is only applied to the final component of a name.

The regular expression case sensitivity depends on the current search mode (see “set-search” on page 6-77).

The **Select a Function/Unit** dialog box is one variation of the debug source selection dialog box, which is also used by the **List Source File...** menu item. See “Source Selection Dialog Box” on page 8-34.

List Source File...

Mnemonic: S

Selecting this menu item pops up a dialog box that allows you to list a source file in the source panels. See “Source Panel” on page 8-66.

This dialog box is titled **Select a Source File**. The title bar also displays the process's qualifier specifier. See “Qualifier Specifiers” on page 6-19. It allows you to optionally enter a wildcard pattern which is used to search for source file names that NightView knows about. See “Wildcard Patterns” on page 6-23. For example, enter `mod*.c` to search for source file names that start with 'mod' followed by any number of characters and ending with '.c'. A list of source files is displayed, and one source file can be selected for display in the debug source display.

The **Select a Source File** dialog box is one variation of the debug source selection dialog box, which is also used by the **List Function/Unit...** menu item. See “Source Selection Dialog Box” on page 8-34.

List Any File...

Mnemonic: A

Selecting this menu item pops up a file selection dialog box that allows you to choose any file you wish and list it in the source panels displaying the current process. See “Source Panel” on page 8-66.

This dialog box is titled **Select a File**.

List Location...

Selecting this menu item pops up a dialog box in which you can enter arguments for a **list** command, to apply to any source panels displaying the current process. See “List Location Dialog Box” on page 8-62 and “list” on page 6-86.

Explain Source...

Selecting this menu item pops up a dialog box that explains the file system locations that NightView searched for the current source file. This dialog box is titled **Explain Source File Search**.

Edit

Mnemonic: I

Selecting this item lets you edit the current process's current source. See "Source Panel" on page 8-66.

Note that once you have edited the source file, NightView displays the *new* contents, but the debugging information still refers to the *old* contents. For this reason, the source decorations may no longer match. Also, you might get confusing results from using the special keys in the debug source display or from entering commands based on the new contents.

Source Search List...

Selecting this menu item opens the **Process Settings** dialog. See "Process Settings Dialog Box" on page 8-59

Show Source

Mnemonic: O

Accelerator: **Ctrl+O**

Selecting this menu item causes the source panel to list source, if possible. If the debugger tries to show a position, such as a library routine, that does not have a corresponding source file, then the source panel shows disassembly instead.

When switching between display modes, NightView uses the position of the text cursor to determine the line or address to show in the new mode.

This item is enabled only if there is exactly one source panel, or if a source panel has a target line. See "Source Panel Target Line" on page 8-67. Each source panel has its own display mode, so, for example, you can show source in one source panel and disassembly in another source panel.

Show Mixed Source and Disassembly

Mnemonic: M

Accelerator: **Ctrl+M**

In this mode the debugger shows a line of source followed by the instructions that correspond to that line. Source lines that do not produce code are not shown. Only one source line is shown for each group of instructions, so statements that span lines are only partially shown. Note that because of inlining and optimization, not all the instructions that follow a line are generated by that line. Also note that lines from multiple files may be shown in this mode.

See the description of the **Show Source** menu item for more information.

Show Disassembly

Mnemonic: D

Accelerator: **Ctrl+D**

Selecting this menu item causes the source panels to show assembly instructions.

The range of instructions displayed usually corresponds to a single subprogram.

See the description of the **Show Source** menu item for more information.

File Coding

Specify the type of file coding being used. Select either:

- Best Guess
- Latin1
- UTF-8

Eventpoint Menu

Mnemonic: E

This menu provides ways to set eventpoints and to see a summary of eventpoints. See “Eventpoints” on page 3-11.

Before selecting one of the eventpoint menu items, select the line of interest in a source panel. See “Source Panel” on page 8-66. NightView uses this line to initialize the location specifier for you. See “Location Specifiers” on page 6-17.

Set Breakpoint...

Mnemonic: B

Accelerator: **Ctrl+B**

Selecting this menu item pops up a breakpoint dialog box that allows you to set a new breakpoint at a given location and apply eventpoint attributes to it. See “Manipulating Eventpoints” on page 6-110. See “breakpoint” on page 6-114.

For information on using the breakpoint dialog box, see “Eventpoint Dialog Boxes” on page 8-35.

Set Monitorpoint...

Mnemonic: M

Selecting this menu item pops up a monitorpoint dialog box that allows you to set a new monitorpoint at a given location and apply eventpoint attributes to it. See “Manipulating Eventpoints” on page 6-110. See “monitorpoint” on page 6-121.

For information on using the monitorpoint dialog box, see “Eventpoint Dialog Boxes” on page 8-35.

Set Patchpoint...

Mnemonic: P

Accelerator: **Ctrl+P**

Selecting this menu item pops up a patchpoint dialog box that allows you to set a new patchpoint at a given location and apply eventpoint attributes to it. See “Manipulating Eventpoints” on page 6-110. See “patchpoint” on page 6-116.

For information on using the patchpoint dialog box, see “Eventpoint Dialog Boxes” on page 8-35.

Set Tracepoint...

Mnemonic: T

Selecting this menu item pops up a tracepoint dialog box that allows you to set a new tracepoint at a given location and apply eventpoint attributes to it. See “Manipulating Eventpoints” on page 6-110. See “tracepoint” on page 6-119.

For information on using the tracepoint dialog box, see “Eventpoint Dialog Boxes” on page 8-35.

Set Heappoint...

Mnemonic: H

Selecting this menu item pops up a heappoint dialog box that allows you to set a new heappoint at a given location and apply eventpoint attributes to it. See “Manipulating Eventpoints” on page 6-110. See “heappoint” on page 6-123.

For information on using the heappoint dialog box, see “Eventpoint Dialog Boxes” on page 8-35.

Set Watchpoint...

Mnemonic: W

Selecting this menu item pops up a watchpoint dialog box that allows you to set a new watchpoint at a given location and apply eventpoint attributes to it. See “Manipulating Eventpoints” on page 6-110. See “watchpoint” on page 6-133.

For more information on using the watchpoint dialog box, see “Eventpoint Dialog Boxes” on page 8-35.

Set Syscallpoint...

Mnemonic: S

Selecting this menu item pops up a syscallpoint dialog box that allows you to set a new syscallpoint and apply eventpoint attributes to it. See “Manipulating Eventpoints” on page 6-110. See “syscallpoint” on page 6-135.

For more information on using the syscallpoint dialog box, see “Eventpoint Dialog Boxes” on page 8-35.

New Eventpoint Panel

Mnemonic: E

Selecting this menu item creates a new eventpoint panel. See “Eventpoint Panel” on page 8-75.

Eventpoint Panels Update Interval...

Mnemonic: U

Clicking this menu item brings up a dialog box that lets you set the interval between eventpoint panel automatic updates. See “Eventpoint Panel Update Interval Dialog Box” on page 8-63.

Data Menu

Mnemonic: D

Use this menu to select a data item to place in a data panel or to load or save a data panel layout.

When you add a data item, if there is no data panel, then one is created. If there is one data panel, then the data item goes to that panel. If there is more than one data panel, then NightView pops up a dialog box to ask you which data panel to add the new item to. For **Expression...**, the dialog box is always popped up and also has a field for entering the expression. See “Data Panel” on page 8-79. The **Save Snapshot** button lets you save the current contents of the panel. The **Save Layout...** and **Load Layout...** buttons let you save the current data panel layout or restore an old layout.

Most data items are added for the current process. See “Current Process” on page 8-4. Shells items, processes items and monitorpoint items are global, not just for one process.

The items in this menu are similar to some of the items in the data panel context menu. See “Data Panel Context Menu” on page 8-91.

The menu items are:

Expression...

Mnemonic: E

Accelerator: **Ctrl+E**

The dialog box allows you to enter an expression. A data item for that expression is placed in the data panel. See “Data Panel Add Expression” on page 8-101. See “Expression Data Item” on page 8-82.

Local Variables...

Mnemonic: L

A local-variables data item is placed in the data panel. See “Local Variables Data Item” on page 8-83.

Registers...

Mnemonic: R

A registers data item is placed in the data panel. See “Registers Data Item” on page 8-83.

Stack...

Mnemonic: S

A stack data item is placed in the data panel. See “Stack Data Item” on page 8-85.

Threads...

Mnemonic: T

A threads data item is placed in the data panel. See “Threads Data Item” on page 8-87.

Processes...

Mnemonic: C

A processes data item is placed in the data panel. See “Processes Data Item” on page 8-88.

Shells...

Mnemonic: H

A shells data item is placed in the data panel. See “Shells Data Item” on page 8-88.

Heap Information...

Mnemonic: H

A heap information data item is placed in the data pane. See “Heap Information Data Item” on page 8-89.

Heap Errors...

Mnemonic: P

A heap errors data item is placed in the data panel. See “Data Panel Add Heap Errors” on page 8-102. See “Heap Errors Data Item” on page 8-90.

Heap Leaks...

Mnemonic: K

A heap leaks data item is placed in the data panel. See “Data Panel Add Heap Leaks” on page 8-102. See “Leak Sets / Still Allocated Sets Data Items” on page 8-90.

Still Allocated Blocks

Mnemonic: A

A still allocated blocks data item is placed in the data panel. See “Data Panel Add Still Allocated Blocks” on page 8-102. See “Leak Sets / Still Allocated Sets Data Items” on page 8-90.

Monitorpoint Values

A monitorpoint values data item is placed in the data panel. See “Monitorpoint Values Data Item” on page 8-91.

Save Snapshot...

Mnemonic: V

This menu item lets you save the current contents of the data panel to a text file. Clicking on this button brings up a dialog box that lets you specify the name of a file in which to save the data. You can also record a comment in the file to describe the data that is being saved. See “Data Panel Save Snapshot” on page 8-104.

Save Layout...

Mnemonic: A

Selecting this menu item pops up a dialog box that lets you save the layout of all the data items for a particular process in all the data panels. The information saved includes the type and format of each data item, and to which data panel the item belongs. See “Data Panel Save Layout” on page 8-104.

Load Layout...

Mnemonic: O

Selecting this menu item pops up a dialog box that lets you load a saved layout for one or more processes. Any data panels mentioned in the layout are created if they do not exist. See “Data Panel Load Layout” on page 8-103.

Set Stack Frames...

Clicking on this button pops up a dialog box that lets you set the number of stack frames displayed in a context panel or data panel. See “Data Panel Call Stack Frames” on page 8-103.

Set Pointer as Array Indices...

Clicking on this button pops up a dialog box that lets you set the number of elements to show when displaying a pointer as an array. See “Data Panel Pointer Array Dimension” on page 8-104.

Tools Menu

Mnemonic: T

The **Tools** menu may be used to invoke other NightStar tools. The tools are invoked on the same display as NightView.

NightView makes a menu entry only for those tools that are installed on your system.

NightProbe Monitor

Mnemonic: P

Opens the NightProbe Data Monitoring application.

NightProbe™ is a real-time graphical tool for monitoring, recording, and altering program data within one or more executing programs without significant intrusion. NightProbe can be used in a development environment as a tool for debugging, or in a production environment to create a “control panel” for program input and output.

See the *NightProbe User's Guide* for more information.

NightSim Scheduler

Mnemonic: S

Opens the NightSim Application Scheduler.

NightSim™ is a tool for scheduling and monitoring real-time applications which require predictable, repetitive process execution. With NightSim, application builders can control and dynamically adjust the periodic execution of multiple coordinated processes, their priorities, and their CPU assignments.

See the *NightSim User's Guide* for more information.

NightTrace

NightTrace™ is a graphical tool for analyzing the dynamic behavior of single and multiprocessor applications. NightTrace can log application data events from simultaneous processes executing on multiple CPUs or even multiple systems and can combine these application events with kernel events to present a synchronized view of the entire system. NightTrace allows users to zoom, search, filter, summarize, and analyze events in a wide variety of ways. In addition, NightTrace allows users to manage user and kernel NightTrace daemons, providing the user with the ability to start, stop, pause, and resume execution of any of the daemons under its management.

See the *NightTrace User's Guide* for more information.

From the NightTrace menu option, select either NightTrace Analyzer or NightTrace Illuminator:

NightTrace Analyzer

Mnemonic: T

Opens a new session of NightTrace.

NightTrace Illuminator

Mnemonic: I

Opens a new session of NightLight.

NightTune Tuner

Mnemonic: U

Opens the NightTune Tuner.

NightTune™ is a graphical tool that can be used to tune an application or to monitor various aspects of the system. NightTune monitors CPU utilization, context switching levels, virtual memory paging activity, disk I/O activity, interrupt activity, and network traffic levels. In addition, NightTune can monitor processes and threads, and can change their scheduling parameters and CPU binding. NightTune can also set CPU shielding and change interrupt CPU affinities.

See the *NightTune User's Guide* for more information.

Help Menu

Mnemonic: H

The **Help** menu has the following items.

On Context

Mnemonic: C

This item provides help about a particular graphical region of a window. See “Context-Sensitive Help” on page 8-2.

On Last Error

Mnemonic: E

If NightView just displayed an error message, you can get help on that error by selecting this menu item.

Selecting this item is similar to using the **help** command with no argument. See “help” on page 6-154.

On Commands

Mnemonic: M

This item gives a summary of NightView commands.

On Keys

Mnemonic: K

This item gives help about using special keys in NightView. See “List of Shortcuts” on page 8-31.

A Quick Start

Mnemonic: Q

This item takes you to the beginning of the GUI quick start chapter. See Chapter 2 [A Quick Start - GUI] on page 2-1.

NightView Tutorial

Mnemonic: R

This item takes you to the beginning of the GUI tutorial chapter. See Chapter 4 [Tutorials] on page 4-1.

NightView User's Guide

Mnemonic: U

The item opens the online version of the *NightView User's Guide* in the help window.

NightStar Tutorial

Mnemonic: T

This item opens the online version of the *NightStar Tutorial* in the help window. This tutorial incorporates the use of NightSim, NightProbe, NightView, NightTrace, and NightTune in one complete example.

Licensing...

Mnemonic: L

Opens a dialog box with information about how many licenses are in use and how to get licenses.

Report a Problem...

Mnemonic: R

Opens a dialog box with information on how to receive software support via website, email, or phone.

On Version

Mnemonic: V

This item pops up an information dialog box that describes which version of NightView you are running.

Check for Updates...

Mnemonic: U

Launches NUU (Network Update Utility) enabling you to update your system. This requires network access to Concurrent's Updates website. Updates require a login

and user ID issued by Concurrent. Refer to <https://redhawk.concurrent-rt.com/home> for complete information.

This menu item is not shown on Debian-based systems. Use the standard Debian update tools to obtain NightSim updates. See the URL shown in the paragraph above for more information.

Toolbars

Most of the controls in the toolbars apply to the current process. See “Current Process” on page 8-4. Some buttons may be disabled (dimmed) under certain circumstances. For example, when the process is running, the **Resume** button is disabled.

A toolbar may be moved to different areas around the main window by dragging on the drag handle.

Toolbars may be hidden or shown by clicking on an area of a toolbar that has no controls. This brings up a menu of the toolbars. Choose which toolbars you want shown. Another way to choose which toolbars are shown is to select the **Toolbars...** entry in the **View** menu.

-

Command Toolbar

The command toolbar has a label **Command:** and a combo box, and is initially near the bottom of the main window. The combo box is used to enter NightView commands. All the command-line interface commands, except for **shell**, can be entered in the command toolbar.

Input to this area is similar to using the command-line interface. For example, you can enter an explicit qualifier followed by a command. If you do not specify a qualifier, the command is implicitly qualified by the current process. See “Current Process” on page 8-4.

The combo box has entries for older commands. To retrieve older commands, press the down arrow key. To see the whole list, click on the downward-pointing triangle.

Process Toolbar

Resume process 

Clicking on this button is similar to using the **resume** command with no argument, or the **r** keystroke in the Source Panel. See “Source Panel Keystrokes” on page 8-70. See “resume” on page 6-134.

Stop process 

Clicking on this button is similar to using the **stop** command, except that the button also interrupts any pending commands for the current process. See “stop” on page 6-142.

Next over

Clicking on this button is similar to using the **next** command with no argument, or the **n** keystroke in the Source Panel. See “Source Panel Keystrokes” on page 8-70. See “next” on page 6-137.

Step into

Clicking on this button is similar to using the **step** command with no argument, or the **s** keystroke in the Source Panel. See “Source Panel Keystrokes” on page 8-70. See “step” on page 6-136.

Finish routine

Clicking on this button is similar to using the **finish** command, or the **f** keystroke in the Source Panel. See “Source Panel Keystrokes” on page 8-70. See “finish” on page 6-141.

Run to Here

Run the process until it reaches the target line in the source panel. See “Source Panel Target Line” on page 8-67 and see “Source Panel” on page 8-66. This allows you to use the **Run to Here** button or the **h** keystroke in the Source Panel to quickly skip past chunks of code without single-stepping through each line.

Clicking on this button combines the actions of three commands: First, it sets a **breakpoint** at the target line. Next, it runs **enable/delete** on that breakpoint (which will cause it to be deleted when it is hit). Finally, it **resumes** the process. See “breakpoint” on page 6-114. See “enable” on page 6-129. See “Source Panel Keystrokes” on page 8-70. See “resume” on page 6-134.

When you press the button, you will see the source line decoration for the breakpoint appear and the message area will print a message about the new breakpoint. When the process finally stops at that breakpoint, the breakpoint will be deleted, and the decoration will disappear. See “Message Panel” on page 8-75.

Next over one instruction

Clicking on this button is similar to using the **nexti** command with no argument or the **N** keystroke in the Source Panel. See “Source Panel Keystrokes” on page 8-70. See “nexti” on page 6-140.

Step into one instruction

Clicking on this button is similar to using the **stepi** command with no argument or the **S** keystroke in the Source Panel. See “Source Panel Keystrokes” on page 8-70. See “stepi” on page 6-139.

Up frame

Clicking on this button advances one stack frame toward the oldest calling frame. This action is similar to using the **up** command with no argument or the **u** keystroke in the Source Panel. See “Source Panel Keystrokes” on page 8-70. See “up” on page 6-151.

Down frame

Clicking on this button advances one stack frame toward the currently executing (newest) stack frame. This action is similar to using the **down** command with no argument or the **d** keystroke in the Source Panel. See “Source Panel Keystrokes” on page 8-70. See “down” on page 6-151.

Kill process

Clicking this button causes NightView to terminate the current process. See “Current Process” on page 8-4.

This is similar to using the **kill** command. See “kill” on page 6-43.

Depending on the safety level (see “set-safety” on page 6-71), NightView may display a warning dialog box when you use the Kill menu item.

Rerun process

Clicking this button causes NightView to terminate the current process, if there is one, and issue the shell command associated with this process. If there is no current process, then NightView issues the shell command associated with the process that terminated most recently. This action is similar to using the **rerun** command with no argument. See “down” on page 6-151.

Interrupt debugger

Clicking on this button interrupts whatever the debugger is doing. This is similar to using the shell interrupt character in the command-line interface. See “Interrupting the Debugger” on page 3-45.

Display Expression

Create a data item in the **Data** panel using the current selection from the **Source** panel as the expression for that data item. Clicking this button is similar to the data-display command where the expression is the current selection from the **Source** panel. See “data-display” on page 6-104.

Print Expression

Print the result of evaluating the current selection from the **Source** panel. Clicking this button is similar to the **print** command where the expression is the current selection from the **Source** panel, or the **p** keystroke in the **Source** panel. See “Source Panel Keystrokes” on page 8-70. See “Print” on page 8-29.

Show Source

Causes the **Source** panel to list source code, if possible. Clicking this button is similar to the Show Source menu item in the **Source** panel Context Menu. See “Source Panel Context Menu” on page 8-68.


Show Mixed Source and Disassembly

Causes the **Source** panel to show lines of source followed by the instructions that correspond to each line. Clicking this button is similar to the Show Mixed Source and Disassembly menu item in the **Source** panel Context Menu. See “Source Panel Context Menu” on page 8-68.

Show Disassembly

Causes the **Source** panel to show assembly instructions. Clicking this button is similar to the Show Disassembly menu item in the **Source** panel Context Menu. See “Source Panel Context Menu” on page 8-68.

Run Mode Toolbar



This option list reflects the current run mode and allows you to change it. See “set-run-mode” on page 6-138 and “Executing” on page 3-50 for descriptions of run mode.

Timing Toolbar

The items in the timing toolbar control timing regions, and their start and end locations.

Region

The **Region** combo box is used to specify the region that other buttons in the timing toolbar interact with. One can enter new region names, or select already existing region names from a drop down menu. If this combo box is left empty, it defaults to the name **default**.

Set Timing Start Location

Clicking this button sets a timing start location at the currently selected line in the source panel for the timing region specified in the **Region** combo box.

Set Timing End Location

Clicking this button sets a timing end location at the currently selected line in the source panel for the timing region specified in the **Region** combo box.

Clear timing results for the specified region

Clicking this button clears timing results for the specified region, if any.

Set Timing Start or End Region

Clicking this button pops up a dialog box that allows you to set start or end locations for any timing region. See “Set New Timing Start/End Dialog Box” on page 8-63.

Set Timing for a Function

Clicking this button pops up a dialog box that allows you to set timing locations for a function. See “Set Timing for a Function Dialog Box” on page 8-64.

Eventpoint Toolbar

Before clicking on a button in the eventpoint toolbar, you may want to click on a target line in a source panel. The **Location:** field of the dialog box is initialized from the target line, if there is one. See “Source Panel Target Line” on page 8-67.

Set Breakpoint

Clicking on this button pops up a breakpoint dialog box that allows you to set a new breakpoint at a given location and apply eventpoint attributes to it. See “Manipulating Eventpoints” on page 6-110. See “breakpoint” on page 6-114.

For information on using the breakpoint dialog box, see “Eventpoint Dialog Boxes” on page 8-35.

Set Patchpoint

Clicking on this button pops up a patchpoint dialog box that allows you to set a new patchpoint at a given location and apply eventpoint attributes to it. See “Manipulating Eventpoints” on page 6-110. See “patchpoint” on page 6-116.

For information on using the patchpoint dialog box, see “Eventpoint Dialog Boxes” on page 8-35.

Set Monitorpoint

Clicking on this button pops up a monitorpoint dialog box that allows you to set a new monitorpoint at a given location and apply eventpoint attributes to it. See “Manipulating Eventpoints” on page 6-110. See “monitorpoint” on page 6-121.

For information on using the monitorpoint dialog box, see “Eventpoint Dialog Boxes” on page 8-35.

Set Tracepoint

Clicking on this button pops up a tracepoint dialog box that allows you to set a new tracepoint at a given location and apply eventpoint attributes to it. See “Manipulating Eventpoints” on page 6-110. See “tracepoint” on page 6-119.

For information on using the tracepoint dialog box, see “Eventpoint Dialog Boxes” on page 8-35.

Set Heappoint

Clicking on this button pops up a heappoint dialog box that allows you to set a new heappoint at a given location and apply eventpoint attributes to it. See “Manipulating Eventpoints” on page 6-110. See “heappoint” on page 6-123.

For information on using the heappoint dialog box, see “Eventpoint Dialog Boxes” on page 8-35.

Set Watchpoint

Clicking on this button pops up a watchpoint dialog box that allows you to set a new watchpoint at a given location and apply eventpoint attributes to it. See “Manipulating Eventpoints” on page 6-110. See “watchpoint” on page 6-133.

For more information on using the watchpoint dialog box, see “Eventpoint Dialog Boxes” on page 8-35.

Set Syscallpoint

Clicking on this button pops up a syscallpoint dialog box that allows you to set a new syscallpoint and apply eventpoint attributes to it. See “Manipulating Eventpoints” on page 6-110. See “syscallpoint” on page 6-135.

For more information on using the syscallpoint dialog box, see “Eventpoint Dialog Boxes” on page 8-35.

Clear

Clicking on this button is similar to using the **clear** command for the target line. See “clear” on page 6-120. You must have set the target line to the source line where you want to clear eventpoints. See “Source Panel Target Line” on page 8-67. When you press this button, any eventpoints that are set at the first instruction of this line are removed. (If you have eventpoints set at instructions within the line, they will not be cleared.) You see the source line decoration change and a message is displayed in the message panel. See “Message Panel” on page 8-75.

New Eventpoint Panel

Create a new panel to show eventpoint information.

Eventpoint Panels Update Interval

Change the interval at which eventpoint panels refresh their data.

Value Toolbar

Data Display

Clicking on this button is similar to using the **Expression...** button in the **Data** menu. See “Data Menu” on page 8-17. You must have selected an expression in a source panel (or another panel) before pressing this button. See “Source Panel” on page 8-66. When you press the button, the selected expression is added to the default data panel. See “Data Panel” on page 8-79.

Print

Clicking on this button is similar to using the **print** command. See “print” on page 6-92. You must have selected an expression in a source panel (or another panel) before pressing this button. See “Source Panel” on page 8-66. When you press the button, the value of the selected expression is printed using the default format for the type of the expression.

Source Display Toolbar

The buttons in this toolbar let you change the source display between source, mixed source and disassembly and just disassembly. These act the same as the corresponding menu items in the **Source** menu. See “Source Menu” on page 8-12.

Status Bar

This area shows the name of the executable program that the current process is running, the qualifier specifier for the current process, and the current process's current status. See “Current Process” on page 8-4. See “Qualifier Specifiers” on page 6-19. For threaded programs, the process ID in the qualifier indicates which of the thread processes last stopped. See “Multithreaded Programs” on page 3-49.

A progress bar is shown here when the debugger is doing some lengthy operations, such as processing debugging information, or during commands that wait, such as the run command.

Here are the values that may appear as the process status:

About to exit

The process called the `_exit(2)` system service. See “Exited and Terminated Processes” on page 3-24.

Calling function

The process is executing to evaluate a function call.

Exited

The process has exited. See “Exited and Terminated Processes” on page 3-24. This status does not normally appear, because the process is removed from the window when the process exits.

Finish frame

The process is executing until a designated instance of a subprogram returns to its caller. See “finish” on page 6-146.

New process

This process has just been created by a `fork()` call in the parent process. The process is stopped. See “Multiple Processes” on page 3-3.

Running

The process is currently executing.

Stepping

The process is executing because of a stepping command. See “step” on page 6-141.

Stopped after finish

The process has completed a **finish** command. See “finish” on page 6-141.

Stopped after step

The process has finished a stepping command. See “step” on page 6-141.

Stopped at breakpoint *number*

The process hit breakpoint number *number*. See “Breakpoints” on page 3-14.

Stopped at watchpoint *number*

The process stopped because of watchpoint *number*. See “Watchpoints” on page 3-16.

Stopped for watchpoint error

The process stopped because of an error during watchpoint processing. An error message in the message panel should explain the problem. See “Watchpoints” on page 3-16. See “Message Panel” on page 8-75.

Stopped after unexpected trap

The process stopped due to an `int3` instruction at a location that NightView was not expecting. Either the user program has an `int3` instruction in it, which would be unusual, or else there is an internal error in NightView.

Stopped by attach

The process has just been attached by the debugger. See “Attaching” on page 3-4.

Stopped by user

The process stopped because of a **stop** command. See “stop” on page 6-142.

Stopped for exec

The process has just `exec()`ed a new program image. See “Programs and Processes” on page 3-3.

Stopped due to CUDA exception

The process has stopped because a CUDA exception was raised. Most CUDA exceptions are fatal.

Stopped due to CUDA error

The process has stopped because of some problem with the CUDA device, CUDA driver, or CUDA intermediary process.

Stopped with *signal*

The process stopped with signal *signal*. See “Signals” on page 3-20.

Terminated with *signal*

The process terminated with signal *signal*. See “Exited and Terminated Processes” on page 3-24. This status appears only for core files. See “Core Files” on page 3-5.

List of Shortcuts

These shortcut keys work when the focus is in any panel except a shell panel. In a shell panel, control keys are passed to the shell. See “Shell Panel” on page 8-74. There are extra keys you can type into a source panel. See “Source Panel Keystrokes” on page 8-73.

Ctrl+A	New page. See “View Menu” on page 8-6.
Ctrl+B	Breakpoint dialog. See “Eventpoint Menu” on page 8-15.
Ctrl+C	Copy in a text field.
Ctrl+D	Show disassembly. See “Source Menu” on page 8-12.
Ctrl+E	Add expression to data panel. See “Data Menu” on page 8-17.
Ctrl+F	Find. See “Find Bar” on page 8-66.
Ctrl+G	Find again. See “Find Bar” on page 8-66.
Ctrl+M	Show mixed source and disassembly. See “Source Menu” on page 8-12.
Ctrl+O	Show source. See “Source Menu” on page 8-12.

Ctrl+P	Patchpoint dialog. See “Eventpoint Menu” on page 8-15.
Ctrl+Q	Exit the debugger. See “File Menu” on page 8-5.
Ctrl+S	Save config. See “File Menu” on page 8-5.
Ctrl+V	Paste in a text field.
Ctrl+X	Cut in a text field.
Ctrl+Z	Undo in a text field.
Ctrl++	Increase font size. See “View Menu” on page 8-6.
Ctrl+-	Decrease font size. See “View Menu” on page 8-6.
Ctrl+0 (zero)	Restore default font size. See “View Menu” on page 8-6.

Main Window Dialog Boxes

This section describes the dialog boxes you might use while debugging, except for the dialog boxes related to the data panel. See “Data Panel Dialog Boxes” on page 8-100.

-

Run Program in Shell Dialog Box

This dialog box pops up when you click **Run...** in the **Process** menu. See “Process Menu” on page 8-11. This dialog box lets you enter a program to debug. You may enter the program name yourself or click **Browse...** to find the program in a file browser. Then enter any arguments to the program and click **OK**. The command is sent to the dialogue shell. See “Dialogue I/O” on page 3-7.

Attach Dialog Box

This dialog box pops up when you select the **Attach...** item in the **Process** menu. See “Process Menu” on page 8-11. This dialog box lets you view the processes on the system and select one or more to attach to.

The dialog box provides two checkboxes to enable optional behavior:

Automatically resume process after attach

Check this box to indicate that the processes should continue to run after the attach. The initial value for this is from the global mode, set with the **set-resume** command or the preferences dialog box. See “set-resume” on page 6-79 and “Preferences Dialog Box” on page 8-46.

Attach, but debug only new children

Check this box to indicate that NightView should gain minimal control of the specified processes, but they will not be debugged, and will automatically resume. Gen-

erally, the purpose of this is so NightView can gain control of and debug any forked processes. (See “set-children” on page 6-55.) A possible use of this feature is to attach to a shell, not to debug the shell, but rather to debug any processes started by that shell. This is equivalent to the **attach /nodebug** option (see “attach” on page 6-43).

Select which processes you want to see with the filter fields, then click on **Refresh**. Select the processes you want to attach to by clicking on them, then click on **Attach**.

The filter has three text input areas. The text input areas each take a regular expression. When you click on the **Refresh** button, the **Processes** list is filled in with the processes that match all three regular expressions: one for the process identifier (PID), one for the **User** and one for the **Program** name. See “Regular Expressions” on page 6-21. Each regular expression must match the entire corresponding string (that is, each one uses an anchored match). The initial value of the PID and **Program** regular expressions is “. *”, which match all processes. The initial value of the **User** regular expression is the name of the user logged in to the dialogue.

The **Processes** list indents the program names to show the parent/child relationship. Each process appears below its parent process and indented relative to the parent process.

The **Attach** button closes the dialog box and attaches to the selected processes.

Click on **Cancel** to dismiss the dialog box without attaching to any processes. Click on **Help** to get help for this dialog box.

Source Selection Dialog Box

This dialog box pops up when you ask to list a function or ask to list a source file from the **Source** menu. See “Source Menu” on page 8-12. It allows you to change the program code that is listed in the source panel by selecting a function or source file name from a list. You can interact with other NightView windows while this dialog box is displayed.

This dialog box is titled **Select a Function/Unit** or **Select a Source File**, depending on which menu item you selected, and displays the qualifier of the current process.

Enter search criteria.

Enter the regular expression (if you are searching for functions) or wildcard pattern (if you are searching for source files) you want to search for, then either press **Return** or click on **Search**. (For a regular expression, an anchored match is *not* implied.) See “Regular Expressions” on page 6-21. See “Wildcard Patterns” on page 6-23.

If you do not want to enter a regular expression or wildcard pattern, you can simply press **Return** or click on **Search** and all functions or files are displayed.

For C++, the regular expression is only applied to the final component of a name.

The next time you use this dialog box, this text is redisplayed.

Select a list item.

If NightView finds any functions or source files, their names are displayed in the list area. If no functions or files are found, a message is displayed in the message panel. See “Message Panel” on page 8-75.

Select an item in the list. If you double-click on an item in the list, the OK button is activated.

Choose an action button.

Click on **OK** to list that function or source file in the source display area. See “Source Panel” on page 8-66. This button is disabled (dimmed) if the list is empty.

You can cancel the listing of the selected function or source file by clicking on **Cancel**.

You can get help for this dialog box by clicking on **Help**.

File Selection Dialog Box

This dialog box pops up when you select **List Any File...** from the **Source** menu. It allows you to list a file of your choice in a source panel.

Select a file name.

Select the file you want to list. If you double-click on a file name in the **Files** list, the **OK** button is activated.

Choose an action button.

If you are satisfied with the file you selected, click on **OK**.

Clicking on **Cancel** cancels the action and closes this dialog box.

You can get help for this dialog box by clicking on **Help**.

Eventpoint Dialog Boxes

NightView provides a dialog box for each type of eventpoint. See “Eventpoints” on page 3-11. These dialog boxes pop up when you use the **Eventpoint** menu, the source panel context menu, or the eventpoint panel to set or change an eventpoint. See “Eventpoint Menu” on page 8-15.

All types of eventpoints share common traits; some eventpoints have additional optional or required information.

The eventpoint dialog boxes generally present the common eventpoint information first, followed by any data that is specific to a given eventpoint. The watchpoint dialog box first presents information specific to watchpoints, followed by the common eventpoint information. Similarly, the syscallpoint dialog first presents widgets allowing you to specify the system calls of interest, followed by the common eventpoint information.

For *inserted eventpoints*, NightView provides default settings for new eventpoints, including a default location specifier. See “Location Specifiers” on page 6-17. In addition, you can enter other information to define the eventpoint. Required data that

must be provided by you before NightView can set the eventpoint is visually emphasized.

Depending on whether you are setting a new eventpoint, or changing an existing eventpoint, NightView allows or disallows access to certain fields in the eventpoint dialog boxes.

Define the eventpoint.

Description (display only)

The title bar of each eventpoint dialog box indicates which kind of eventpoint the dialog box deals with and whether the dialog box allows you to set a new eventpoint or to change an existing eventpoint.

Location

This field is displayed only for inserted eventpoint dialog boxes, not for watchpoint dialog boxes.

When the dialog box appears, the **Location** field contains a location specifier.

When setting a new eventpoint, NightView determines this value from the target line in the source panel. See “Source Panel Target Line” on page 8-67. You can edit this text input area.

When changing an existing eventpoint, NightView displays the location specifier associated with this eventpoint. You cannot change this location.

Line numbers in location must match exactly

This may be checked to indicate that any line number specified in the location will be interpreted as *fixed*. That is, if the specified line contains no code, a subsequent line that does contain code will not be selected instead. See /f under “Eventpoint Modifiers” on page 6-113.)

Watchpoint options (watchpoint dialog box only)

These controls let you indicate whether you want to specify an L-value (e.g., a variable name) or an explicit program address and size. You can also control whether you want the watchpoint to be for memory reads, memory writes, or both. If the target is an IA-32 or AMD64, watchpoints always trap on memory writes, but you can control whether they also trap on memory reads.

When changing an existing watchpoint, these controls cannot be changed.

Watchpoint target (watchpoint dialog box only)

This text input area lets you enter an L-value or an explicit program address, depending on the setting of the controls in the watchpoint options area.

When changing an existing watchpoint, this field cannot be changed.

Watchpoint size (watchpoint dialog box only)

This combo box lets you select the size, in bytes, of the watchpoint target if you have selected **Watch address and size** in the watchpoint options area. If you have not selected **Watch address and size**, then this area is not enabled.

When changing an existing watchpoint, this field cannot be changed.

Syscallpoint selection (syscallpoint dialog box only)

This text input area lets you enter one or more system call names, separated by commas or spaces. Leaving the list empty, or the word **All**, refers to all system calls.

To the right of the text input area is a **Select...** button; pressing it launches a dialog which aids you in selecting system call names. See “System Call Selection Dialog” on page 8-41 for information on this dialog.

Below the text input area is a checkbox, which, when checked, negates the specified list of system calls. Thus you can easily use the selection area to match all system calls but a select few. This can be especially useful in situations where you already have a syscallpoint that matches one system call, but want to take different actions for all other system calls with a new syscallpoint event.

When changing an existing syscallpoint, these fields cannot be changed.

Syscallpoint actions and modifiers (syscallpoint dialog box only)

The action to be taken when a syscallpoint is hit is specified by selection of one of the two radio buttons:

- Print only; do not stop
- Stop and print

Further, the modifiers to the right allow you to match a system call on entry, exit, or both.

When changing an existing syscallpoint, these fields cannot be changed.

Eventpoint Number (display only)

When changing an existing eventpoint, NightView displays the eventpoint number.

Enable Options

When setting a new eventpoint, you can choose from several enable options. By default, the eventpoint is created enabled. This is similar to using the **enable** or **disable** commands. See “enable” on page 6-124. See “disable” on page 6-123.

When changing an existing eventpoint, NightView displays the eventpoint's enabled state. You can select a different enable option by clicking on one of the choices.

Enable

This is the default choice when setting a new eventpoint. The eventpoint is enabled.

Enable, disable after next hit

You can have the eventpoint be disabled automatically after the next hit.

For breakpoints, this is similar to using the **tbreak** command, or the **enable/once** command. See “tbreak” on page 6-126.

For patchpoints, this is similar to using the **tpatch** command, or the **enable/once** command. See “tbreak” on page 6-126.

For other eventpoint types, this is similar to using the **enable/once** command.

Enable, delete after next hit

Valid for breakpoints and watchpoints only. You can have the eventpoint be deleted automatically after the next hit. This is similar to using the **enable/delete** command.

Disable

You can disable the eventpoint.

Condition

You can attach a condition to this eventpoint, or change an existing condition, by editing this text input field. This is similar to using the **condition** command. See “condition” on page 6-122.

If you delete an existing condition, the eventpoint becomes unconditional.

Ignore Count

You can attach an ignore count to this eventpoint, or change an existing ignore count, by entering a number in this text input area. This is similar to using the **ignore** command. See “ignore” on page 6-125.

The default ignore count is zero and is represented by a blank field.

Name

When setting a new eventpoint, you can assign a name to it by entering text in this text input area. The name must consist only of alphanumeric characters and underscores and must begin with an alphabetic character. The name may be of arbitrary length. This is similar to using the **name** command. See “name” on page 6-108.

You cannot change an existing eventpoint's name using the dialog box. Use the **name** command to change eventpoint names.

Commands

Valid for breakpoints, syscallpoints, and watchpoints only. You can attach commands to this eventpoint, or change existing commands, by entering one command per line in this multi-line text input area. This is similar to using the **commands** command. See “commands” on page 6-121.

There are additional restrictions for commands when used with watchpoints and syscallpoints; see “watchpoint” on page 6-133 and “syscallpoint” on page 6-135 for more information.

Monitorpoint Expressions

This area is shown for monitorpoints only. One or more expressions are *required* to set a monitorpoint. Enter an expression in the first column.

The default format to print the is determined by the type of the expression. If you want a different format, click on the format field. The field changes to a combo box. Select the desired format.

Enter a label if desired. If no label is entered, then the expression string is used as the label.

Click on **New** to make another entry in the expression table. You can also use the **Tab** key in the label field to make a new entry.

Patchpoint Action

Valid for patchpoints only; you are *required* to enter either an expression, a location specifier, or a *Thread Tags* assignment to set a patchpoint. Select the appropriate choice by clicking on it. The radio button appears filled for your selection, and the label for the text input area changes to either **Evaluate**, **Go to**, or **Tag**. Enter the expression or location specifier in the text input area.

Insert an expression at this location

This field represents the *eval* argument of one variant of the **patchpoint** command. See “patchpoint” on page 6-111. This is the default choice.

Branch to a different location

This field represents the *goto* argument of one variant of the **patchpoint** command.

Set thread local tag values

This field must contain the tag assignment(s) to be executed. See the *tag assignment* description of the **patchpoint** command for more information.

NightTrace Event (for tracepoints only)

ID

You are *required* to supply an ID number to set a tracepoint. This field represents the *event-id* argument of the **tracepoint** command. You must enter a trace-event number or symbolic name. See “tracepoint” on page 6-114.

Value

This optional field represents the `value=` argument of the **tracepoint** command. You can enter an expression whose value should be logged with the trace event.

Once set, these fields cannot be changed.

Heappoint Action -- Check - Debug Settings

This area is shown for heappoints only. Select whether you want to perform a heap check or change the heap debugger settings by clicking on one of the radio buttons. If you want to change the heap debugger settings, enter the arguments in the **Settings:** text input area. Valid arguments are the same as for the **heapdebug** command. See “heapdebug” on page 6-57.

Show/Hide Advanced Controls

This button controls whether or not additional *advanced* controls are displayed. If they are displayed, a few more options are available.

Thread Protection

This option controls whether or not a breakpoint hit causes protected threads to be stopped along with all non-protected threads. Normally, protected threads would not be stopped unless the breakpoint was actually hit within one of them.

Address Space

You may restrict the address space to which the given eventpoint applies:

Cuda only

The eventpoint may only be inserted into a CUDA address space.

Process only

The eventpoint may only be inserted in the host address space.

Cuda or Process

The eventpoint may be inserted into either the host address space or a CUDA address space. If a CUDA location applies, it will be selected over a host location.

Dialog Buttons

Click on **OK** to set or change the eventpoint. The dialog box is dismissed.

Click on **Delete** to delete this eventpoint. The dialog box is dismissed. This button is present only for an existing eventpoint.

Clicking on **Cancel** cancels the action and closes this dialog box.

You can get help for this dialog box by clicking on **Help**. The dialog box is not dismissed.

If you are setting a new eventpoint or deleting an existing one, you see the source line decoration change and the eventpoint panel change. NightView displays a message in the message panel to tell you if the eventpoint was set.

If you make an error while entering data, NightView may display an error dialog box and allow you to re-enter the data. Other warnings or errors associated with setting or changing this eventpoint are displayed in the message panel. See “Message Panel” on page 8-75.

You can use the **info eventpoint** command or the eventpoint panel to check the eventpoint settings. See “info eventpoint” on page 6-160. See “Eventpoint Panel” on page 8-75.

System Call Selection Dialog

This dialog aids in selecting one or more system calls for the Syscallpoint Dialog. You can use the search facility to quickly locate a specific call, or you can scroll and select individual items. This dialog supports multiple selection; use **Ctrl+<click>** to select multiple non-contiguous system calls, or **Shift+Ctrl+<click>** to select a range of system calls.

When focus is in the list of displayed system calls, keystrokes entered take you to the next system call that starts with the letter associated with your keystroke.

Debug Heap Dialog Box

This dialog box pops up when you use the **Debug Heap...** item in the **Process** menu (see “Process Menu” on page 8-9). This dialog box allows you to turn on heap debugging and control the heap debugger. The effect of using this dialog box is similar to using the **heapdebug** command (see “heapdebug” on page 6-57).

Enable Heap Debugging

The most important control is the checkbox to turn heap debugging on and off. If you are a casual user of heap debugging, you may want to restrict your attention to this control and the level buttons and common error buttons. See “heapdebug” on page 6-57 for information about turning heap debugging on and off.

Debugging Level buttons

The level buttons provide a convenient way of setting some of the other controls. See “Levels and Common Errors” on page 3-38.

Common Errors Detection buttons

The common error buttons provide a convenient way of setting some of the other controls to configure the heap debugger to detect a particular kind of program error. See “Levels and Common Errors” on page 3-38.

General Settings

If **Hardware Overrun Protection** is checked, each block is placed at the end of a page and the following page is protected from reads and writes. See “Hardware Overrun Protection” on page 3-40.

Check **Specify check heap freq** to enable automatic heap checking. In the text field, set the number of heap operations between heap checks. Uncheck the box to turn off automatic heap checks.

Check **Specify retained free blocks** to give the number of free blocks that should be retained. Uncheck the box retain all free blocks. See “Retained Free Blocks” on page 3-41.

Check **Specify heap size** to limit the size of the heap to the specified number of bytes. Uncheck the box to indicate that the total size of the heap is limited only by system resources.

Enter the number of extra bytes to add to each allocation size in **Slop size**.

Enter the number of walkback entries to keep for each heap operation in **Walkback Entries Per Block**. This number refers to physical walkback entries. The number of walkback frames may differ from this number when displayed in NightView. The number of frames displayed may include extra inline frames, as they are not physical frames. The number of frames displayed may be fewer if certain frames are deemed uninteresting (see “interest” on page 6-74).

Pre-fence size is the number of bytes to fill and check before each block. **Post-fence size** is the number of bytes to fill and check after each block. See “Fences” on page 3-39.

NOTE

When specifying a number of heap operations, blocks, or bytes, you may append the letter **k** to multiply the number by 1024, or the letter **m** to multiply by 1048576.

Fill Settings

If **Fill malloc space** is checked, blocks are filled with the **Malloc fill byte** when they are allocated.

If **Fill free space** is checked, blocks are filled with the **Free fill byte** when they are freed.

If **Check free fill** is checked, retained free blocks are checked for the **Free fill byte** during a heap check.

Pre-fence fill byte is the value used to fill the pre-fence. **Post-fence fill byte** is the value used to fill the post-fence.

Error Control

These controls specify how the debugger responds when an error condition is detected. Each error has a **Stop** checkbox and a **Print** checkbox. If **Stop** is checked, the process will stop when it gets that error. If **Print** is checked, the debugger prints a message when the process gets that error.

Action buttons

Click the **OK** button to configure heap debugging with these settings and dismiss the dialog box. Click the **Reset** button to restore the settings to be the same as when the dialog box popped up. Click the **Cancel** button to dismiss the dialog box without making any changes to the heap debugging configuration. Click the **Help** button to get help about the dialog box.

Remote Login Dialog Box

This dialog box pops up when you use the **Shell** menu's **Start Remote Shell...** item. See “Shell Menu” on page 8-10. This dialog box allows you to specify the parameters for creating a remote NightView session. See “Remote Dialogues” on page 3-7. Some of these parameters are required, but most are optional.

The parameters specified in this dialog apply to the NightView processes that execute on the remote system. These processes include a NightView target program, a dialogue shell, and (unless you specify otherwise using the **run (1)** shell command) all the processes started by that dialogue shell.

Remote Login General Page

Target

This is the name or address of the remote system on which you want a remote dialogue. This field is required information.

Login name

This specifies the user name to use to log into the remote system. This field is required, but it defaults to the user running NightView.

After you click on **Login**, you will be prompted for the passphrase for this user in another dialog box. If you leave the passphrase empty, then you will be prompted for the password for this user. For security, the passphrase or password you type is not echoed in the window; instead, an asterisk (*) replaces each character.

Shell Name

This field specifies the name to give to the dialogue. See “Qualifier Specifiers” on page 6-19. If you leave this field empty, the name of the dialogue will default to be the same as the **Target** field. If the remote system name is not a valid dialogue name, an error dialog will appear. A common reason for the

remote system to be an invalid dialogue name is that the remote system name contains period (.) characters (e.g., it includes domain names), or it is an IP address instead of a name.

Show the remote shell in a new shell panel

Indicate with the checkbox whether you want a new shell panel to be created for this shell. See “Shell Panel” on page 8-74.

Debug i386 programs on x86_64 target

Indicate with the checkbox if you want to start a 32-bit shell instead of the default shell. Use of this checkbox allows you to debug 32-bit applications in this shell, but prohibits debugging of 64-bit applications. See “Architecture Interoperability” on page 3-55.

Remote Login Advanced Page

This area allows you to set scheduling attributes which will be applied to the NightView processes which will run on the specified target system.

Scheduling Class

This combo box allows you to select the POSIX scheduling class for the NightView processes:

- Other

This class corresponds to the SCHED_OTHER scheduling policy which provides for general process scheduling with urgency less favorable than the other two classes. Processes in this class have their priority adjusted by the operating system based on CPU usage.

- Round Robin

This class corresponds to the SCHED_RR scheduling policy which provides real-time process scheduling using a time-slicing algorithm to share CPU resources with other SCHED_RR processes of the same priority.

- First In First Out

This class corresponds to the SCHED_FIFO scheduling policy which provides the strictest real-time process scheduling. Processes are not time-sliced with other SCHED_RR or SCHED_FIFO processes of the same priority.

See **`sched_setscheduler(2)`** for more information on these scheduling classes.

Class Attributes

Depending on the scheduling class chosen, the following attributes can be selected:

- Nice Value

You can use this spin-box to set the initial nice value to be associated with the NightView processes. A nice value provides a bias to the default priority of a process, thereby affecting the effective priority. Positive values correspond to less favorable scheduling urgency. This attribute is only available when the **Other** class is selected. See **nice(1)** for more information on the effect of nice values.

- Real-time Priority

You can use this spin-box to set the real-time priority of the NightView processes. Priority values are constrained to be between 1..99. Selecting a priority value exceeding 90 is not recommended as it may interfere with kernel daemon processing. A priority of one is sufficient to give the process more urgency than any process in the **SCHED_OTHER** class. Higher priority numbers correspond to more favorable scheduling urgency.

- Time Quantum

You can use this spin-box to set the duration of the time-slice for processes using the **SCHED_RR** class. This attribute is not applicable to any other scheduling class.

CPU List

NightView server process execution will be constrained to the CPUs listed here. Additionally, the remote shell being created and all programs you run in it will be similarly constricted, unless you specifically change their interrupt affinity programmatically or with the **run(1)** command.

By default, the list is **all**, which means the server process can run on any CPU on the target system which isn't shielded from process execution (consult the **shield(1)** man page for more information on shielding).

The list can either be the word **all**, or a comma-separated list of CPU numbers or ranges of CPU numbers; for example: **0, 2-3**.

To the right of the text field a description of the resultant CPU mask is shown. Some system interfaces require CPU affinity to be specified as a mask, with each bit in the mask representing a CPU. The mask is shown to remind you that the numbers you enter into the text field here are logical CPU numbers, not hexadecimal characters in a CPU mask.

If you enter something invalid into the text field, the description to the right changes to the word **invalid**, shown in red. Ultimately, syntactically-invalid CPU lists are automatically replaced with a list indicating **all**.

NOTE

Selection of the Round Robin or First In First Out scheduling class or a non-default CPU bias requires privileged access.

Remote Login Action Buttons

Login

When you press the **Login** button, the remote dialogue is created and the remote login dialog box is dismissed. If the remote dialogue cannot be created, either an error dialog box will appear or the remote login dialog disappears and a message is displayed in the message panel. See “Message Panel” on page 8-75.

Reset

Set all the fields back to their default values.

Cancel

Pressing the **Cancel** button dismisses the dialog box without creating a remote dialogue.

Help

Pressing the **Help** button brings up the online help with information about the remote login dialog box.

Preferences Dialog Box

This dialog box pops up when you click on **Preferences...** in the **File** menu. See “File Menu” on page 8-5. With it you can adjust things such as fonts and how you interact with the debugger. Preferences you set here apply only to this debug session unless you save them to disk. You can save your preferences to disk by clicking the **Save** button in this dialog box or with the **Save Preferences** item in the **File** menu.

Many of these preferences may be set by commands. This dialog box and the commands refer to the same preferences, so, for example, if you change a preference with a command, you can see the same value when you bring up this dialog box.

For a way to change settings that relate to a process, see “Process Settings Dialog Box” on page 8-59.

Click **OK** to apply the preferences and dismiss the dialog.

Click **Apply** to apply the preferences and leave the dialog up.

Click **Reset** to set the preferences back the way they were when the dialog popped up (or when **Apply** the button was clicked).

Click **Save** to apply the preferences, save them to disk and dismiss the dialog.

Click **Cancel** to dismiss the dialog without applying any changes.

Click **Help** to get help about the dialog.

Preferences General Page

General | Source And Assembly | Fonts | Terminal | Advanced

Safety
 Safety level: **verify** ☐ Ask before killing processes when exiting.

Automatically Resume On
☐ attach ☐ exec ☐ exit ☐ fork ☒ cuda ☒ cuda (system)

Searching
☐ case-sensitive

Data Panel
 Stack frames: **5** Pointer As Array Indices: **2**

Eventpoint Panel
☒ Auto Update Update Interval: **0.3 Secs**

CUDA
☒ enabled ☐ software preemption

Display Limits
 array elements: **2048** string characters: **256** addresses: **10**

Source Panel Keystrokes
☒ enabled

Source File Size
 max source file limit: **4000 kB**

OK Apply Reset Save Cancel Help

Safety

Select a safety level. This is similar to using **set-safety**. See “set-safety” on page 6-71. If the level is **verify**, you can say whether you want to be warned if you ask to exit when you are still debugging processes.

Automatically Resume On

The debugger normally stops a process when it is first attached, when it execs, when it is about to exit, or when it forks. It normally does not stop a process when NightView detects any CUDA kernel launches. Check the conditions for which you want the process to continue running. See “set-resume” on page 6-79.

Searching

Indicate whether you want searching to be case-sensitive. This affects the forward-search and reverse-search commands, and is the initial setting for find bars. See “set-search” on page 6-77.

Data Panel

Select the default number of stack frames to display in a data panel (assuming at least that many stack frames exist). Select the number of indices to show when treating a pointer as an array. Additional cases are checked for pointers and may result in "invalid" and a red foreground:

- Invalid run-time type information (RTTI)
- RTTI which indicates that the object pointed to is incorrect, e.g. `T*` pointer which points to an object not of type `T` or of a class descended from type `T`.

Hover over the value to display a tooltip with details of the problem.

See “Data Panel Call Stack Frames” on page 8-103 and “Data Panel Pointer Array Dimension” on page 8-104.

Eventpoint Panel

Specify whether or not the Eventpoint panel should update automatically. If so, also specify the update interval in seconds.

CUDA

Specify whether or not CUDA support should be enabled. If disabled, interactions with the CUDA hardware will be treated like any other device. See “set-cuda” on page 6-83.

Also specify whether or not CUDA software preemption should be enabled. CUDA software preemption allows a device to be used by other applications when an application running on the device is stopped for debugging. It requires CUDA 5.5 or higher, and a device with CUDA capability 3.5 or higher.

NOTE

CUDA software preemption often is used to allow debugging a CUDA application using the same device as an X11 display.

Display Limits

Select the display limits for array elements, string characters and addresses in location specifier listings. These limits are to prevent large amounts of output from overwhelming the display. See “set-limits” on page 6-69.

Source Panel Keystrokes


You can perform many actions by pressing keys when the keyboard focus is in the source panel. Use this control to turn source panel keystrokes on or off. See “Source Panel Keystrokes” on page 8-73.

Source File Size

Select the maximum number of bytes a source file can have to be displayed in a source panel. This is useful for extremely large source files which overwhelm NightView due to the overhead involved in building individual widgets associated with each line. The num-

ber is in units of 1000 bytes. The default value is 4000, which indicates ~4MB. The source for files that exceed the limit are not displayed, but the assembly associated with the function associated with the current stack frame within the file is displayed. See “set-limits” on page 6-69.

Preferences Source and Assembly Page



Source Display

Choose whether you want the source to be colored based on syntax and whether you want to see line numbers in the source panel. See “Source Panel” on page 8-66.

In addition, choose the **Source File Coding**. This is the character encoding used by source files. NightView generally expects files to be encoded with either Latin 1 (ISO/IEC 8859-1) or UTF-8 (which permits the entirety of Unicode). Its default is **Best Guess**, which uses heuristics to determine the correct encoding between these two. Alternately, either **Latin1** or **UTF-8** may be chosen explicitly to disable the guessing. Finally, an **Add/Edit...** item may be used to specify additional character encodings. The names of encodings are as supported by the **iconv** utility. To get a listing of all possible encodings, execute **iconv -l** from a command-line shell. For example, to use the GBK (Chinese) encoding, select **Add/Edit...**, enter GBK into the text field, press **Add Encod-**

ing and then press OK. The file coding also can be set from the Source Panel's context menu (see "Source Panel Context Menu" on page 8-68).

Disassembly

Choose how you want to see disassembled code. See "set-disassembly" on page 6-81.

Source Editor

This section determines how to run the editor when you select the **Edit...** item in the **Source** menu or use the **e** key in the source panel. Enter the command you want to use to edit files. The combo box is already loaded with two popular editor commands, and one is selected based on your `EDITOR` environment variable. You can pass information about the file and the current position with `%` specifiers.

`%`

Replaced by `%`. That is, to get a `%`, use `%%`.

`s`

Replaced by the name of the source file.

`l`

Replaced by the line number of the current position.

`p`

Replaced by the offset, in characters, of the current position from the beginning of the file.

`c`

Replaced by the column of the current position.

A `%` followed by any other character is ignored.

The **Run this command in a terminal program** checkbox determines whether or not the source editor should be run under a terminal emulator (e.g. **xterm**). If clear, it is assumed that the editor can communicate with the X Window System display directly. In that case, the editor runs on the same display as NightView. If checked, the editor is run under a terminal emulator. The terminal emulator can be specified on the **Terminal** page (see "Preferences Terminal Page" on page 8-56).

Preferences Fonts Page

NightView uses multiple fonts to present text in the most effective manner throughout the various display areas of the tool.

Variable-width fonts are most commonly used; these fonts most closely resemble how people write or print words.

Fixed-width fonts require that all characters and numbers have the same width (visual footprint). Fixed-width fonts are of benefit when source code is being displayed or manipulated or when columns of numbers are viewed.

NightView further divides the use of fonts into the following categories; default and panel.

Default fonts are used for text associated with operational description and control, including: menus, buttons, selection devices, labels, tool tips, status bar messages, and generally descriptive verbiage.

Panel fonts are used in NightView panels, which display the data of highest importance.

Fonts are selected by querying font preferences from the following sources until a preference is found:

- Your NightView preference
- Your NightStar-wide preference
- The system's NightView preference
- The system's NightStar-wide preference
- NightView's ultimate default



Figure 8-1. Font Preferences Page

This page is divided into three sections.

Global NightStar Fonts

The **Change...** button in this area launches the **NightStar Global Fonts** dialog which allows you to set your Nightstar-wide preferences, your preferences for another specific NightStar tool, or the system's tool or NightStar-wide preferences.

Note:

Setting a NightStar preference for the system typically requires root access.

Changes saved in the **NightStar Global Fonts** dialog are always saved to disk and apply to the current and subsequent NightView invocations.

See “NightStar Global Fonts Dialog” on page 8-52 for more information.

My NightView Fonts

This area allows you to set or clear your user’s preferences for NightView.

Selection of the checkboxes for the individual font categories control whether or not your preferences are to be consulted. Clearing a checkbox effectively removes your user preference for that category. Setting a checkbox allows you to select specific fonts within the category.

Changes to any of the settings in this area, including individual fonts or category checkboxes, are immediately reflected in the **Effective NightView Fonts** area at the bottom of the page so you can see the ultimate effect a change will have.

To change a specific font, ensure that the corresponding category’s checkbox is checked and then press the **Change...** button. This will launch a standard font selection dialog. When you select a font from the dialog and press **OK**, the name of the font family is displayed to the left of the **Change...** button and is displayed in the selected font as well.

Effective NightView Fonts

This area shows you the effective fonts that will be used based on your user settings and consultation of global settings which aren’t shown in the page.

The values in this area immediately change to reflect the effective font whenever any change is made within the page.

Your changes in the **My NightView Fonts** area are applied to the current invocation of NightView when you press the **OK** button. However, your changes are not saved to disk and will not affect subsequent invocations of NightView unless you press the **Save** button.

Separation of apply and **Save** operations make it easy to experiment with fonts in the current invocation without affecting long-term usage.

Note:

Changes to font preferences in the NightStar Global Fonts dialog are always saved to disk and apply to the current and subsequent NightView invocations; i.e. there is no way to experiment with a global font preference without affecting subsequent NightView invocations.

NightStar Global Fonts Dialog

The NightStar Global Fonts dialog allows you to set your Nightstar-wide preferences, your preferences for another specific NightStar tool, or the system's tool or NightStar-wide preferences.



Figure 8-2. NightStar Global Fonts Dialog

Keep in mind that fonts are selected by querying font preferences from the following sources until a preference is found:

- Your NightView preference
- Your NightStar-wide preference
- The system's NightView preference
- The system's NightStar-wide preference
- NightView's ultimate default

This dialog has two control areas which define the scope of font preference application.

Changes Fonts For...

By default, the dialog is set up to apply font preferences to your user account. Select the **Entire System** button if you wish to set the system's preferences.

Note:

Changing font preference for the system typically requires `root` access.

Apply Fonts To...

This area additionally controls the scope of font preference application. You can change a preference for a specific NightStar tool or change the NightStar-wide preference.

If you wish to change the font for more than one tool from this dialog, but not change the NightStar-wide preference, select the first tool of interest, make your preference change in the areas below, and then press the **Save** button. Then select the second tool of interest and repeat.

Set Default Fonts Set Panel Fonts

These areas contain the variable and fixed-width font preferences for each of the font categories, identified by the label next to each checkbox.

To remove the preferences in a category, clear its checkbox.

To change a specific font, ensure that the category's checkbox is checked and then press the **Change...** button. This will launch a standard font selection dialog. When you select a font from the dialog and press **OK**, the name of the font family is displayed to the left of the **Change...** button and is displayed in the selected font as well.

The buttons at the bottom of the page control the application of your changes.

Save & Close

Saves any changes made in this dialog to disk, thus affecting subsequent tool invocations, and closes the dialog.

These changes may affect the effective font preferences for the current invocation of NightView. When the dialog is closed, the fonts shown in the **Effective NightView Fonts** section of the **Preferences** dialog are updated. If you apply the changes in that dialog, they will take effect in the current invocation of NightView.

Save

Applies the preferences from the dialog to the current invocation of NightView, saves the preferences to disk thereby affecting subsequent NightView invocations.

These changes may affect the effective font preferences for the current invocation of NightView. When this dialog is subsequently closed, the fonts shown in the **Effective**

NightView Fonts section of the Preferences dialog are updated. If you apply the changes in that dialog, they will take effect in the current invocation of NightView.


Cancel

Cancels any unsaved changes and closes the dialog.

Help

Opens the help system to display this section.

Preferences Terminal Page



This page allows selection of the preferred system terminal emulation tool (e.g. **xterm**). The user may specify a number of tools in priority order. NightView will use the first tool in the list that can be found on the user's **\$PATH**. The terminal emulator will be executed

using any arguments specified. A few replacement strings are provided. If they appear in the arguments, they will be replaced as specified:

Table 8-1. Terminal replacement strings

%C	Command to be run as multiple arguments
%Q	Command to be run as properly quoted single string argument
%%	The % character

The list of terminal emulators can be manipulated with the buttons to the right. The **New** button allows addition of a new terminal emulator. The **Delete** button removes the selected terminal emulator from the list. The **Up** and **Down** buttons move the selected terminal emulator up and down in the list, respectively, changing their position in the priority order.

Preferences Advanced Page

The screenshot shows the 'Advanced' tab of a preferences dialog. The 'Remote Object File Cache' section has a 'Cache mode' dropdown set to 'permanent' and a 'Cache directory' text field containing '/rat1/users/todd/.NightViewCache' with a 'Browse...' button. The 'Eventpoint Memory Preallocation' section has two checked checkboxes: 'eventpoints' and 'monitorpoint buffers'. The 'Restart' section has an 'Apply restart information:' dropdown set to 'always'. The 'Value History' section has a 'Keep' spinner set to '1000' and the text 'entries.'. The 'Expression Evaluation Automatic Overloading' section has two checkboxes: 'names' (checked) and 'operators' (unchecked). The 'Future Eventpoints' section contains explanatory text and three radio buttons: 'Always ask about undefined eventpoint locations' (selected), 'Always treat undefined eventpoint locations as an error', and 'Always treat undefined eventpoint locations as futurepoints'. At the bottom of the dialog is a 'Restore all preferences on every tab to their default values:' label and a 'Restore defaults' button. The bottom of the window features a row of buttons: 'OK', 'Apply', 'Reset' (highlighted with a dashed border), 'Save', 'Cancel', and 'Help'.

Remote Object File Cache

Control how NightView downloads files from remote targets. See “set-download” on page 6-80.

Memory Preallocation

Control how NightView preallocates memory for eventpoints, monitorpoint, and timing buffers. See “set-preallocate” on page 6-78.

Restart

Control whether restart information is applied. See “set-restart” on page 6-72.

Value History

Specify the number of items to be kept in the value history list. See “set-history” on page 6-68.

Expression Evaluation Automatic Overloading

Control how NightView treats overloaded operators and routines in expressions. See “set-overload” on page 6-77.


Future Eventpoints

Control how NightView treats location specifiers in eventpoint commands when the locations specified do not exist yet. See “set-futurepoints” on page 6-82.

Restore Defaults

Click this button to give all the preferences their default values. The preferences on all the tabs of the preferences dialog are affected.

Preferences Binary Viewer Page



Size and Format

These are the default values used for the data element size and format when a new binary viewer panel is opened.

Search

This is the default **Search all memory** setting when a new binary viewer panel is opened.

Background Colors

These are the colors used for highlighting various data elements in a binary viewer panel:

Process Settings Dialog Box

This dialog box pops up when you click on **Process Settings...** in the **Process** menu. See “Process Menu” on page 8-11. With it you can adjust things such as whether children are debugged and how signals are handled.

Many of these settings may be set by commands. This dialog box and the commands refer to the same settings, so, for example, if you change a setting with a command, you can see the same value when you bring up this dialog box.

For a way to change user preferences that do not relate to individual processes, see “Preferences Dialog Box” on page 8-46.

Click **OK** to apply the settings and dismiss the dialog.

Click **Apply** to apply the settings and leave the dialog up.

Click **Reset** to put the settings back the way they were when the dialog popped up (or when **Apply** the button was clicked).

Click **Cancel** to dismiss the dialog without applying any changes.

Click **Help** to get help about the dialog.

Process Settings General Page

Debug Children

Control whether children of this process should be debugged. See “set-children” on page 6-55.

Set Run Mode

Controls the execution of threads in a multi-threaded process when resuming a thread -- does one thread run or do all threads run. See “set-run-mode” on page 6-138 for a full description.

Branch Tracking

Control whether NightView and the RedHawk kernel are tracking branch instructions in this process. See “Branch Tracking” on page 3-43.

Stop Before Exiting

Control whether this process stops before exiting. See “set-exit” on page 6-56.

Expression Language

Establish a default language context for variables and expressions. See “set-language” on page 6-67.

Refresh debug info when shared libs change

When this option is checked, NightView automatically detects when shared libraries are loaded (e.g. use of the **dlopen(2)** service) after the process starts. There is some overhead in using this option. Alternatively, you can use the **Refresh Shared Libs** option from the **Process** menu manually when you wish NightView to re-read the shared library list.

Program

Type the name of the executable program or use the **Browse...** button to find the executable program with a file browser. See “exec-file” on page 6-48.

Process Settings Interest Page

Control which subprograms are interesting. See “interest” on page 6-74.

Process Settings Signals Page

Specify how to handle signals in the user process. For each signal, specify whether you want the process to stop when it receives the signal, whether you want NightView to print a message when the process receives the signal, and whether you want the signal to be passed to the process when it resumes. Note that **stop** implies **print**. See “handle” on page 6-150.

Process Settings C++ Exceptions Page

This tab appears only for a C++ program.

Select whether you want the process to stop and to print for C++ exceptions in general. If you want to specify different actions for specific C++ types, use the **Add...** button to add types to the table. The **Add...** button pops up a browser for types. Select whether you want the process to stop and to print for the individual types in the table. If the type is a class type, the specification also will apply to subclasses.

In all cases stop implies print.

See “Exceptions” on page 3-21 and “handle” on page 6-150

Select a Type Dialog Box

This dialog box provides two methods of adding types to the Process Settings C++ Exceptions Page.

- Enter the name of the type into the **Name** field and press the **OK** button to the right of this field. This is the only way to add POD types such as **int**, **float**, **char***, etc. It also is useful for types with short names that you already know.
- Performing a filtered search for user-defined types such as **class**, **structs**, and **typedefs**. This is done by entering a partial type name in the **Filter** field and pressing the **Filter** button to the right of this field. This will cause the application to be searched for all types that match the filter. Then, one or more types can be selected from the list, and the **Select** button can be pressed to add those types.

Rename Page Dialog Box

This dialog box pops up when you click **Rename Current Page...** in the **View** menu. See “View Menu” on page 8-6.

Type the new name for the page. You can create a mnemonic for the page by preceding one of the characters with an ampersand (&). The page's tab will have an underscore under that character. To get a real &, use &&.

Print Dialog Box

This dialog box pops up when you click on **Print Window...** in the **File** menu. Click **Print** to print the main window.

List Location Dialog Box

This dialog box pops up when you click on **List Location...** in the **Source** menu or in the source panel context menu.

Enter a where-spec as you would for the list command. See “list” on page 6-86. If you clicked on **List Location...** in the **Source** menu, any source panels displaying the current process are affected (see “Source Menu” on page 8-12). If you clicked on **List Location...** in the source panel context menu, the source panel in which you clicked is the only one affected (see “Source Panel Context Menu” on page 8-68).

If desired, you can specify a file on another host with the form *user@host:/path*. See “Remote File Access” on page 3-9.

Eventpoint Panel Update Interval Dialog Box

This dialog box pops up when you click on **Eventpoint Panels Refresh Rate...** in the **Eventpoint** menu.

NightView can discover some of the eventpoint information only by querying the process. To reduce overhead, this is done only every few seconds.

Select the number of seconds between updates and whether you want automatic updates. You can use the eventpoint context panel to do updates manually. The eventpoint panel is updated when you make any eventpoint change regardless of whether you have automatic updates on.

The eventpoint panel update interval is not related to the monitorpoint update interval. See “Monitorpoint Update Interval Dialog Box” on page 8-107.

Set New Timing Start/End Dialog Box

This dialog box is activated from the timing toolbar's **Set Timing Start or End Region** button (see “Set Timing Start or End Region” on page 8-27), or from the **Process Menu**. It allows setting start or end locations for a new or existing timing region.

Location

Specify a program location for the start or end location. This defaults to the target line in the source panel (see “Source Panel Target Line” on page 8-67).

Region Name

Specify the timing region with which to associate the new start or end location. This defaults to the region name in the timing toolbar's **Region** combo box (see “Region” on page 8-26).

Cross-thread region

If checked, it indicates that the region should be a crossthread timing region (see “Timing” on page 3-18). This option only is available only for new timing regions.

Position

This indicates whether the location is to be a start or end location.

Set Timing for a Function Dialog Box

This dialog box is activated from the timing toolbar's **Set Timing for a Function** button (see “Set Timing for a Function” on page 8-12), or from the **Process Menu**. It allows setting timing locations for a function.

Region Name

Specify the timing region with which to associate the new timing locations. If the name is left empty, then the name of whatever function is entered or selected below also is used as the region name.

This dialog box provides two methods of selecting a function name:

- Entering the function name into the **Function Name** field and pressing the **OK** button to the right of this field.
- Performing a filtered search for the function. This is done by entering a partial function name in the **Function Filter** field and pressing the **Filter** button to the right of this field. This will cause the application to be searched for all functions that match the filter. Then, the function can be selected from the list, and the **Select** button can be pressed to add the timing locations for that function.

Panels

NightView shows different kinds of information in different panels. The panels are docked within the main window or may be undocked, separate windows. See “Main Window” on page 8-4.

Find Bar

In most panels you can use the find bar to search for text in that panel. The find bar shows up at the bottom of the panel when you use the **Find...** item in the panel's context menu. See “Context Menu” on page 8-3. You can also make the find bar appear by typing **Ctrl+F** when the focus is in the panel. (However, typing **Ctrl+F** in a shell panel does not show the find bar. Instead, the character is sent to the shell.)

The find bar has a button with an **X** you can use to close the find bar. Next is a text field where you can enter a search string. The search string is a regular expression. See “Regular Expressions” on page 6-21. As you type in more characters, the search progresses. If a matching line is found, the line is highlighted. If no match is found, the background color of the text entry field changes. A label at the end of the find bar shows the status of the search. You can press the **Escape** key here to close the find bar. You can press **Enter** to move to the next matching text.

The find bar has a button to find the next matching text and one to search backwards for the previous matching text. There is a check box to indicate whether you want the search to pay attention to case. The initial state of the check box comes from the **Searching** section of the preferences dialog box. See “Preferences Dialog Box” on page 8-46.

Once you have a search text in a panel, you can use the **Find Again** item in the panel's context menu, or you can type **Ctrl+G** when the focus is in the panel, to find the next matching text. (However, typing **Ctrl+G** in a shell panel does not search. Instead, the character is sent to the shell.)

Source Panel

The source panel lists the program source code or the disassembled instructions corresponding to the current frame in the current process. See “Current Frame” on page 3-33. See “Current Process” on page 8-4. You can select various display modes to display source or disassembly with the **Source** menu. See “Source Menu” on page 8-12. See “list” on page 6-86, for information on how the current source file is determined.

The text in this area includes the program source or disassembled instructions along with line numbers and source decorations. See “Source Line Decorations” on page 6-92. You can turn off the line numbers with the preferences dialog box. See “File Menu” on page 8-5. If you hover the mouse pointer over the source line decoration, a tooltip shows any eventpoint information for the line.

The text in this area changes if you use the **Source** menu to list other functions or files, if you use the **list** command (see “list” on page 6-83), and when the process stops or you change the current context.

The titlebar of the source panel shows the name of the program, the qualifier of the process associated with this source panel, and the name of the source file displayed in the panel.

If the source panel is in disassembly display mode, the title bar shows information about the region displayed, such as `Function main`.

Source Panel Target Line

When you click on a line in a source panel, that line becomes the *target line*. The source panel shows the target line with a different background. (This is independent of text selection, which has a different highlighting background.) At most one line in one source panel can be the target line.

Buttons and keystrokes that deal with eventpoints use the target line to identify the source file and line number to operate on. If the button brings up a dialog box, then the target line is used to initialize the location field. Otherwise, the target line is acted on immediately. See “Source Panel Keystrokes” on page 8-73. See “Eventpoint Menu” on page 8-15. See “Source Panel Context Menu” on page 8-68.

If there is more than one source panel, the target line also indicates which source panel to operate on in some items in the **Source** menu. See “Source Menu” on page 8-12.

In addition to setting the target line, clicking in a source panel also sets the current process to be the process currently associated with the source panel. See “Current Process” on page 8-4. If the current process changes, the target line is cleared.

Source Panel Expression Tooltip

If you select the text of an expression in the source panel, and then hover the mouse pointer over the selection, NightView displays a tooltip showing the value of the expression. (This is text selection, not the same as clicking to set the source panel target line. See “Source Panel Target Line” on page 8-67.) If there is some error, such as that the selected text is not a valid expression, the tooltip shows the error instead. If the process is stopped,

then NightView evaluates the expression with respect to the current frame. See “Current Frame” on page 3-33. If the process is running, the evaluation is done in the global scope.

The surrounding text is ignored, so, for example, if you have a line that includes `a.b.c` and you select and hover over only the `c`, NightView tries to find a variable named `c`, not the member named `c` from `a.b`.

The expression is not evaluated if it would modify the process or cause the process to run, so you will get an error if you select and hover over `"i = 1"` or `"factorial(x)"`.

To select text, either drag the mouse pointer over it, or move the text cursor to one end of the expression and then hold down **Shift** and press the **Right** or **Left** arrow key until the whole expression is selected.

Source Panel Context Menu

Right-click in the source panel to bring up the context menu. This menu provides ways of changing the program code displayed in the panel, manipulating eventpoints, and editing source files that are listed. See “Source Panel” on page 8-66.

Note that the right-click also sets the target line. See “Source Panel Target Line” on page 8-67.

The items that change which source file is displayed, and the items that select source or disassembly act on this panel only.

Set simple breakpoint

Mnemonic: **B**

Sets a simple breakpoint on the target line.

Clear eventpoint

Mnemonic: **L**

Clears all the eventpoints on the target line.

Run to Here

Mnemonic: **H**

Runs the process until it reaches the target line (by setting a breakpoint and deleting the breakpoint when it is hit).

Set eventpoint

Mnemonic: **E**

Brings up a sub-menu of eventpoint types. Click on an entry in the sub-menu to bring up a dialog to set an eventpoint on the target line.

Set Breakpoint...

Mnemonic: **B**

Accelerator: Ctrl+B

Set Monitorpoint...

Mnemonic: M

Set Patchpoint...

Mnemonic: P

Accelerator: Ctrl+P

Set Tracepoint...

Mnemonic: T

Set Heappoint...

Mnemonic: H

Set Watchpoint...

Mnemonic: W

Set Syscallpoint...

Mnemonic: S

Edit eventpoint

If there is a single eventpoint on the target line, this item appears as **Edit eventpoint *n***, where *n* is the eventpoint number. If there are multiple eventpoints on the target line, this item brings up a sub-menu that lists the eventpoints. Selecting any of these items brings up a dialog box allowing you to edit the eventpoint attributes.

List Function/Unit...

Mnemonic: U

Selecting this menu item pops up a dialog box that allows you to list the program code of a function in the source panel. See “Source Panel” on page 8-66.

This dialog box is titled **Select a Function/Unit**. The title bar also displays the process's qualifier specifier. See “Qualifier Specifiers” on page 6-19. It allows you to optionally enter a regular expression that is used to search for function names that NightView knows about. (An anchored match is *not* implied.) See “Regular Expressions” on page 6-21. For example, enter `set$` to search for function names ending with 'set'. A list of functions is displayed, and one function can be selected for display in the source panel. For C++, the regular expression is only applied to the final component of a name.

The regular expression case sensitivity depends on the current search mode (see “set-search” on page 6-77).

The **Select a Function/Unit** dialog box is one variation of the source selection dialog box, which is also used by the **List Source File...** menu item. See “Source Selection Dialog Box” on page 8-34.

List Source File...

Mnemonic: S

Selecting this menu item pops up a dialog box that allows you to list a source file in the source panels. See “Source Panel” on page 8-66.

This dialog box is titled **Select a Source File**. The title bar also displays the process's qualifier specifier. See “Qualifier Specifiers” on page 6-19. It allows you to optionally enter a wildcard pattern which is used to search for source file names that NightView knows about. See “Wildcard Patterns” on page 6-23. For example, enter `mod*.c` to search for source file names that start with 'mod' followed by any number of characters and ending with '.c'. A list of source files is displayed, and one source file can be selected for display in the source panel.

The **Select a Source File** dialog box is one variation of the source selection dialog box, which is also used by the **List Function/Unit...** menu item. See “Source Selection Dialog Box” on page 8-34.

List Any File...

Mnemonic: A

Selecting this menu item pops up a file selection dialog box that allows you to choose any file you wish and list it in the source panel. See “Source Panel” on page 8-66.

This dialog box is titled **Select a File**.

List Location...

Brings up a dialog box in which you can enter arguments for a **list** command, to apply only to this source panel. See “List Location Dialog Box” on page 8-62 and “list” on page 6-86.

Explain Source...

Selecting this menu item pops up a dialog which explains how the currently displayed source file was determined. It can be useful if multiple different versions of a source file exist in different locations. It explains the decision tree that led to a particular source file being displayed.

List History

Mnemonic: R

Selecting this menu item brings up a sub-menu with entries for source files you have viewed recently. The most recently viewed files are listed first. (If you were viewing disassembly, then the entry describes the address region you were viewing.) Click on an entry in the sub-menu to view that file again.

If you view different parts of the same file, separate entries are kept in the list so you can easily switch back and forth.

A separate list is kept for each dialogue and program name. If you debug the same program again, the list is still available.

Find...

Mnemonic: **F**

Accelerator: **Ctrl+F**

Brings up the find bar and sets the keyboard focus there. See “Find Bar” on page 8-66.

Find again

Mnemonic: **G**

Accelerator: **Ctrl+G**

Finds the search string in the find bar again. The search begins from the text cursor position. At the end of each search the text cursor is left at the end of the found string.

Edit

Mnemonic: **I**

Selecting this item lets you edit the source file that is currently displayed in the source panel. See “Source Panel” on page 8-66. This item is disabled (dimmed) if the source panel is displaying disassembly.

Note that once you have edited the source file, NightView displays the *new* contents, but the debugging information still refers to the *old* contents. For this reason, the source decorations may no longer match. Also, you might get confusing results from using the special keys in the source panel or from entering commands based on the new contents.

Show Source

Mnemonic: **O**

Accelerator: **Ctrl+O**

Selecting this menu item causes the source panel to list source, if possible. If the debugger tries to show a position, such as a library routine, that does not have a corresponding source file, then the source panel shows disassembly instead.

When switching between display modes, NightView uses the position of the text cursor to determine the line or address to show in the new mode.

Each source panel has its own display mode, so, for example, you can show source in one source panel and disassembly in another source panel.

Show Mixed Source and Disassembly

Mnemonic: **M**

Accelerator: **Ctrl+M**

In this mode the debugger shows a line of source followed by the instructions that correspond to that line. Source lines that do not produce code are not shown. Only

one source line is shown for each group of instructions, so statements that span lines are only partially shown. Note that because of inlining and optimization, not all the instructions that follow a line are generated by that line. Also note that lines from multiple files may be shown in this mode.

Show Disassembly

Mnemonic: D

Accelerator: Ctrl+D

Selecting this menu item causes the source panel to show assembly instructions.

The range of instructions displayed usually corresponds to a single subprogram.

See the description of the **Show Source** menu item for more information.

File Coding

Selecting this menu item brings up a selection of possible character encodings for the source file. By default, this list includes the following options:

- Best Guess: Uses heuristics to choose between Latin 1 (ISO/IEC 8859-1) or UTF-8 character encoding.
- Latin1: Uses the Latin 1 (ISO/IEC 8859-1) character encoding.
- UTF-8: Uses the UTF-8 character encoding (which permits the entirety of Unicode).

In addition, the list of choices includes any other encodings that may have been added using the Preferences dialog. See “Source Display” on page 8-49.

This is a global setting which affects all source files.

Track any process

Mnemonic: Y

The panel changes whenever another process becomes the current process. See “Source Panel Tracking” on page 8-72.

Track one process

Mnemonic: N

The panel ignores changes for other processes. See “Source Panel Tracking” on page 8-72.

Panel locked

Mnemonic: K

The panel ignores changes for all processes. See “Source Panel Tracking” on page 8-72.

Source Panel Tracking

A source panel usually shows source for the current process and changes when the process stops or you select a different frame. When another process becomes the current process, the source panel also changes to show the new process. We say the source panel tracks any process. See “Current Process” on page 8-4.

You can restrict a source panel to track a particular process. To do this, select the process as the current process in the context panel, right-click in the source panel to get the source panel’s context menu, then select **Track One Process**. See “Context Panel” on page 8-78. See “Source Panel Context Menu” on page 8-68. In this mode the source panel responds only to events for that process, such as clicking on frames for that process in the context panel, or the process stopping.

You can also lock a source panel. Right-click in the source panel to get the context menu, then select **Panel Locked**. In this mode the panel does not respond to the process stopping or changing frames.

In any mode, the source panel still responds to things you explicitly tell it to do, such as displaying a different file by using the source panel’s context menu.

If a source panel is tracking one process or locked, and the associated process terminates, the source panel goes back to tracking any process.

Source Panel Keystrokes

There are several special keys that may be used when the keyboard focus is in a source panel. The function of most keys is independent of the target line. Some keys, like **b** and **h**, do depend on the target line so that NightView can determine the source line of interest. See “Source Panel Target Line” on page 8-67. Note that the meaning of these keys does not change between source display mode and disassembly display mode. For example, **s** means step one line in any display mode.

In addition to these keys that work only in a source panel, there are shortcuts that work in any panel (except for a shell panel). See “List of Shortcuts” on page 8-31.

b	This key sets a breakpoint. It performs the same action as the Breakpoint button (see “Process Toolbar” on page 8-24).
d	This key changes the current frame to the callee. This is similar to using the down command with no argument. See “down” on page 6-151.
Enter	Sets the line with the text cursor to be the target line. See “Source Panel Target Line” on page 8-67.
e	This key is similar to selecting the Edit item in the Source menu. See “Source Menu” on page 8-12. This key is disabled in disassembly display mode.
f	This key runs the process until it returns from the current frame. It is similar to using the finish command. See “finish” on page 6-141.
h	Run the process until it reaches the target line. This key is identical to the Run to Here button. See “Process Toolbar” on page 8-24. It combines the actions of breakpoint , enable/delete , and resume .

N	Step one instruction without entering called routines. This key is similar to using the nexti command with no argument. See “nexti” on page 6-140.
n	Step one line without entering called routines. This key is similar to using the next command with no argument. See “next” on page 6-137.
p	Print the result of evaluating the current selection. This key performs the same action as the Print button in the process toolbar. See “Process Toolbar” on page 8-24).
r	Resume the process. This key is similar to using the resume command with no argument. See “resume” on page 6-134.
S	Step one instruction, entering a called routine. This key is similar to using the stepi command with no argument. See “stepi” on page 6-139.
s	Step one line, entering any called routine. This key is similar to using the step command with no argument. See “step” on page 6-136.
u	This key changes the current frame to the caller. This is similar to using the up command with no argument. See “up” on page 6-151.
=	Move to the newest stack frame. This key is similar to using the frame 0 command. See “an Ada task, ,frame” on page 6-149.
>	Print information about the current frame. This key is similar to using the frame command with no arguments. See “an Ada task, ,frame” on page 6-149.
[	This key sets a timing start location for the timing region indicated in the Timing Toolbar's Region. See “Timing Toolbar” on page 8-26.
]	This key sets a timing end location for the timing region indicated in the Timing Toolbar's Region. See “Timing Toolbar” on page 8-26.

Shell Panel

This area allows you to interact with the dialogue shell and with your programs. See “Dialogues” on page 3-6. You can run your program here, just as you would normally run it, providing any arguments that it needs. Shell and program output is displayed here. You can also enter input to the shell and to your programs. This panel acts something like a little terminal. If your shell lets you do command-line editing, then you can do that in this panel, too. In addition to passing escape sequences and control keys that you type, NightView also passes standard escape sequences for the arrow keys (**Right**, **Left**, **Up** and **Down**) and the **Home** and **End** keys, for shells that recognize those sequences.

NightView shortcuts cannot be used here because the keystrokes are passed to the shell.

Any programs that you run in the shell can be debugged and manipulated by NightView.

Alternatives to using the shell panel are the **Run...** item in the **Process** menu, the **run** command, and giving the program and program arguments as arguments when invoking NightView. See “Process Menu” on page 8-11. See “run” on page 6-40.

Right-click in the panel to bring up the context menu. The shell panel's context menu contains these entries:

Find...

Mnemonic: F

Brings up the find bar and sets the keyboard focus there. See “Find Bar” on page 8-66.

Find again

Mnemonic: G

Finds the search string in the find bar again. The search begins from the text cursor position. At the end of each search the text cursor is left at the end of the found string.

Message Panel

This panel displays messages including process status messages, error messages, output from commands, and output from processes and the shell. All message panels have the same contents.

Right-click in the panel to bring up the context menu. The message panel’s context menu contains these entries:

Find...

Mnemonic: F

Accelerator: Ctrl+F

Brings up the find bar and sets the keyboard focus there. See “Find Bar” on page 8-66.

Find again

Mnemonic: G

Accelerator: Ctrl+G

Finds the search string in the find bar again. The search begins from the text cursor position. At the end of each search the text cursor is left at the end of the found string.

Eventpoint Panel

The eventpoint panel shows you a table of existing eventpoints for all processes. It also provides ways for you to change eventpoints. See “Eventpoints” on page 3-11.

The table has a row for each eventpoint. The columns show:

- the type of the eventpoint (breakpoint, heappoint, monitorpoint, patchpoint, tracepoint, syscallpoint, or watchpoint), with an icon for the type.

- the eventpoint ID number. Each eventpoint has a unique ID.
- the location of the eventpoint in the program. Watchpoints and syscallpoints do not have a location in the program: for watchpoints, this area describes the location being watched; for syscallpoints, this area describes the system calls being traced.
- the ID of the process this eventpoint is in.
- whether the eventpoint is enabled.
- the ignore count.
- the hit count.
- the crossing count.
- whether the eventpoint has commands (only breakpoints, monitorpoints, syscallpoints, and watchpoints can have commands).
- any condition on the eventpoint.

For more information on the precise interactions of many of these items, see “Interactions Between Conditions, Ignore Counts, etc.” on page 3-14.

If the program is stopped at a breakpoint, syscallpoint, or watchpoint, the eventpoint icon is overlaid with a green triangle pointing to the right,  as it is in the source panel. If more than one breakpoint was hit, the triangle indicates the last breakpoint hit by the current thread.

You can sort on the various columns in the table by clicking on the headers. You can rearrange the columns by dragging the headers.

One way to change eventpoints in this panel is to edit the fields in the table directly. The **Enabled**, **Ignore**, **Hits**, **Crossings**, and **Condition** fields can be edited by clicking on them. The field changes to a control appropriate for changing that field.

NOTE

It is important to click somewhere else after changing the fields in the table directly. The eventpoint will not be changed until you click somewhere else.

Right-click in the panel to bring up the context menu. You may select one or more eventpoints (rows) and then right-click. If you right-click on a row that is not selected, the selection is cleared and the row you clicked on becomes selected. If you right-click on a row that is selected, the selection does not change. The context menu's entries are enabled or disabled based on which rows are selected.

The eventpoint panel's context menu contains these entries:

Edit...

Mnemonic: I

Brings up a dialog box to change the attributes of the selected eventpoint. See “Eventpoint Dialog Boxes” on page 8-35.

Enable

Mnemonic: E

Enable the selected eventpoints. This is similar to using the **enable** command. See “enable” on page 6-129.

Disable

Mnemonic: D

Disable the selected eventpoints. This is similar to using the **disable** command. See “disable” on page 6-128.

Delete

Mnemonic: L

Delete the selected eventpoints. This is similar to using the **delete** command. See “delete” on page 6-123.

Once deleted, you cannot refer to these eventpoints again. If you think you may want to "turn off" an eventpoint temporarily, then use it again later, you should disable the eventpoint and enable it when you are ready to use it.

Clear Ignore Count

Mnemonic: G

Set the ignore count to zero for the selected eventpoints.

Clear Hit Count

Mnemonic: I

Set the hit count to zero for the selected eventpoints.

Clear Crossing Count

Mnemonic: R

Set the crossing count to zero for the selected eventpoints.

Clear Commands

Mnemonic: M

Remove any commands on the selected eventpoints.

Clear Condition

Mnemonic: N

Remove any condition on the selected eventpoints.

List Source

Mnemonic: S

Show the source corresponding to the location of the selected eventpoint.

Resize Columns to Data

Mnemonic: Z

Adjust the width of the columns of the table to fit the data.

Update Now

Refresh the data in the table.

Warnings or errors associated with changing eventpoints are displayed in the message panel. See “Message Panel” on page 8-75.

You can also use the **info eventpoint** command to check eventpoint settings. See “info eventpoint” on page 6-160.

Context Panel

The context panel is a special data panel that lets you browse within the processes you are debugging and their stack frames. See “Data Panel” on page 8-79.

A context panel is essentially a specialized data panel that shows you shells, processes, threads, CUDA contexts, stack frames, and local variables within stack frames. However, it hides some of these panel components if they are uninteresting; shells are shown only if you have multiple shells, process entries appear only if you have multiple processes, etc.

The current stack frame is shown in green underlined text. If threads and CUDA contexts are shown, the current thread or CUDA context is shown in green underlined text. If processes are shown, the current process is shown in green underlined text. If shells are shown, the shell that contains the current process is shown in green underlined text.

To change the current context, click on a stack frame, process or shell. The source panels and the status bar are updated for that context and the items related to that context become green underlined text.

Locals Panel

A locals panel is a special data panel that has a single local variables data item. (The root of the data item is hidden.) This panel always shows the local variables for the current frame for the current process. See “Data Panel” on page 8-79. See “Local Variables Data Item” on page 8-83. See “Current Frame” on page 3-33. See “Current Process” on page 8-4.

Monitor Panel

A monitor panel is a special data panel that has a single monitorpoint values data item. (The root of the data item is hidden.) This panel always shows monitorpoint values. See “Data Panel” on page 8-79. See “Monitorpoint Values Data Item” on page 8-91. See “Current Frame” on page 3-33. See “Current Process” on page 8-4.

Registers Panel

A registers panel is a special data panel that has a single local registers data item. (The root of the data item is hidden.) This panel always shows the registers for the current frame for the current process. See “Data Panel” on page 8-79. See “Registers Data Item” on page 8-83. See “Current Frame” on page 3-33. See “Current Process” on page 8-4.

Data Panel

A data panel displays various information about your debug session. Each data panel has a name. If there is more than one data panel, then you are prompted for the name of a data panel when you place an item in the panel. If no panel of that name exists, then one is created with that name. There is no limit on the number of data panels.

Items are placed into a data panel by using the **Data** menu, or by using the **Data Display** button, or by using the data panel’s context menu, or by invoking the **data-display** command. See “Data Menu” on page 8-17. See “Value Toolbar” on page 8-29. See “Data Panel Context Menu” on page 8-91. See “data-display” on page 6-101.

Monitor Bar

The monitor bar appears when you select **Show Monitor Bar** in the context menu for a monitorpoint values item or in a monitor panel. See “Data Panel Context Menu” on page 8-91.

This area contains a button to hide the monitor bar, the **Hold/Release** button, and a spinbox with the update interval in milliseconds.

Use the **Hold/Release** button to hold or release monitorpoint updates. The button is labeled **Hold** when updates are running, and **Release** when updates are held. You can also hold and release updates with the **mcontrol** command or its aliases **hold** and **release**. See “mcontrol” on page 6-119.

If you modify the update interval spinbox, you need to click somewhere else to have the value take effect.

Data Items

Each data item shows one piece of data from your process. The data items are arranged in a tree. A data item has a label and a value field. The format of the value field depends on the kind of data item.

If the data item has sub-items, then it appears with a small button to the left. If the button is shown with a +, then any sub-items the data item has are currently collapsed (not currently displayed). You can expand the sub-items by clicking on the + button.

If the button is shown with a -, then any sub-items the data item has are currently expanded (displayed). You can collapse the sub-items by clicking on the - button.

If there is no button with + or -, then the data item has no sub-items.

Some data items can be modified. In general, this is allowed for any item which is of a scalar type and which represents an lvalue (e.g. variable, register, indirection through a pointer, etc.). Modification can be begun by double clicking in the value cell, clicking an already-selected value cell, or simply by beginning to type when a value cell is selected and in focus. This transforms the value cell into an editor. After **Enter** is pressed in the editor, the text entered will be interpreted as an expression and assigned to the item. So, it is permitted for the new value to be a literal, another variable, or a complex expression, so long as it results in a value which is meaningful for the type of the data item. If the item has known values (e.g. a variable of an enumeration type), the editor will be a combo box with suggestions for possible values for that item. However, it still is possible to disregard them and use any meaningful expression.

You can right-click on the data item to pop up a context menu for the data item. See “Data Panel Context Menu” on page 8-91.

You can change the size of scroll regions in this panel with the **Data** menu or the data panel context menu. See “Data Menu” on page 8-17.

For expression data items showing arrays, and for stack data items, you can extend the number of items shown by clicking on the arrowhead items .  See “Expression Data Item” on page 8-82 and “Stack Data Item” on page 8-85.

NightView may present the value field with the term "(invalid)" appended and using a red foreground. This happens with expression data items for pointer expressions, and for block data items and for some of the sub-items therein. For a pointer data item, this indicates that the pointer references nonexistent, freed, or never allocated memory (although NULL always is considered valid). For block data items, which are available if heap debugging is turned on, this indicates a heap block for which an error has been detected. For sub-items therein, it indicates the nature of the error.

Top-level data items may be moved by dragging them with the mouse. They may be moved to a different top-level position in the same data panel or to another data panel. Items in a context, locals or monitor panel may not be moved.

Expression Data Item

An expression data item displays the value of an expression, such as a variable.

The expression is re-evaluated whenever the process stops.

The expression is re-evaluated in the context that was current at the time the data item was created, or in the context that is current at the time of the re-evaluation, depending on the setting when you created the data item. See “Data Panel Add Expression” on page 8-101.

Item	Value
life_and_the_universe	42
str	0x804b008 "Oh Brave New World that has such debuggers in it!"
vector	0x804b180
*	34.00000000000000
(heap info)	
state	allocated
range	0x0804b180 .. 0x0804b1cf
size	80 bytes
errors	none detected (as of last heap check)
allocation information	0x0804858d in main() at data-items.c line 36
configuration	
walkback	0x0804858d in main() at data-items.c line 36
critter	hamster (4)
linked_list	0x804b040 struct node * {linked list: \$p->next [0:1]}
-> [0th]	0x804b040 struct node *
-> [1st]	0x804b060 struct node *
value	777
ptr	0x804b070
next	0x804b080 struct node *
(heap info)	
(heap info)	

Figure 8-3. Expression Data Items

In the figure above, several different kinds of expressions have been added to the data panel.

If the value is a C struct, then the sub-items are the members of the struct.

If the value is an array, the sub-items are the elements of the array. A limited number of elements is shown.

If the value is a pointer, the sub-item is the result of indirecting through the pointer.

Pointers can be treated as arrays or linked lists. When treated as a linked list, sub-items represent nodes in the linked list. See “Data Panel Linked List Expression Dialog” on page 8-105 on how linked lists are interpreted.

If the value is an array, or is being treated as an array, there are special arrowhead items on either end. ▲ ▼ Click on the arrowhead items to reveal more array elements. (NightView lets you reference elements beyond the ends of an array.) Linked lists can be manipulated the same way, although such lists terminate when the link is NULL.

If heap debugging is turned on, and the value is a pointer which references heap memory, then it will have an additional sub-item named "(heap info)", which is a block data item describing the heap block containing the memory referenced by the pointer. See "Block Data Item" on page 8-90.

Local Variables Data Item

A local variables data item has sub-items for all the local variables visible in the current scope, including subprogram arguments. For C++ member functions, this is also included in the local variables.

This data item is updated whenever the process stops or you change the current context, e.g., by clicking on Up or Down in the process toolbar. See "Process Toolbar" on page 8-24.

Registers Data Item

A registers data item has sub-items for all of the registers.

This data item is updated whenever the process stops or you change the current context, e.g., by clicking on Up or Down in the process toolbar. See "Process Toolbar" on page 8-24.

Register Data Item

A register data item displays the value for a particular register.

The best format for a given register may depend on circumstance, but best guesses for the default format are made based on typical usage for each register. For example, on the x86_64 architecture, `rax` and `rdx` often hold arithmetic values, and so are displayed as a signed integer. However, `rsi` and `rbp` often hold pointers, and so are displayed as unsigned integers.

For vector registers (e.g. x86_64 `xmmn` registers), the format varies substantially depending on circumstance. So, for these registers, sub-items are provided for each possible hardware format. For example, this might include "64-bit doubles" and "32-bit floats". Furthermore, sub-items for each of those formats is provided to display the individual values of the vector for whichever format is expanded. So, if "64-bit doubles" is expanded, it would show each of the 64-bit double elements of an x86_64 `xmmn` register.

If it is known that all registers of a given register file (e.g. all `xmmn` registers of the XMM register file) will use the same format, then a specific Vector Register View can be specified for that register file. See "Vector Register Views" on page 8-100.

For special registers with known bit fields (e.g. x86_64 `eflags` or aarch64 `pstate`), the register has sub-items for each defined field. For some registers, interpreting some of the fields may be complex, and so additional sub-items are provided which synthesize information from multiple fields. For example, the x86_64 `eflags` registers contains both

"signed interpretation" and "unsigned interpretation" synthetic fields to interpret the two different interpretations after a `cmp` instruction.

Stack Data Item

A stack data item has sub-items for each frame on the stack. Initially, the number of frames shown is limited to the default number of stack frames. See “Data Menu” on page 8-17. The highest numbered stack frame is followed by an item with an arrowhead ▼. Click on the arrowhead item to reveal more stack frames. Once you have reached the end of the stack, the arrowhead item goes away.

Expanding a frame shows the local variables in that frame.

The current frame is shown in green underlined text. See “Current Frame” on page 3-33.

If branch tracking has been enabled (see “Branch Tracking” on page 3-43), a stack data item has an additional Branch History sub-item to display the branch history. This information is independent of stack frames and may cover branches within the current frame, any previous frames, and even frames which have completed and returned.

Item	Value
	<u>local:28519</u>
	#0 0x08048491 in calc(int v = 12) at data-items.c line 17
	#1 0x080484aa in work(int x = 3) at data-items.c line 19
	<u>#2 0x08048625 in main() at data-items.c line 44</u>
	i
	10
	n
	0x804b160 struct node *
	value
	172
	ptr
	0x804b170
	*
	736
	(heap info)
	0 struct node *
	next
	0 struct node *
	(heap info)

Figure 8-4. Stack Data Item

Branch History Data Item

A branch history data item appears only as a sub-item of a Stack data item if branch tracking has been enabled (see “Branch Tracking” on page 3-43). It has sub-items for each branch that was tracked. Initially, the number of branches shown is limited to 5. The highest numbered branch shown is followed by an item with an arrowhead ▼. Click on the arrowhead item to reveal more branches. Once you have reached the total number of branches tracked, the arrowhead item goes away.

The most recent branch is displayed first, with each subsequent branch being the next most recent.

Each branch is displayed with a **From** *address* location and a **To** *address* location. The *address* description always contains a hex address. If available, it will contain function name and source file, line information. If that is not available but symbols are, it will contain *<symbol+offset>* information.

Threads Data Item

A threads data item has sub-items to describe C threads, and CUDA contexts. Expanding a sub-item shows the stack for that thread. If the program is not threaded, then it is considered to have one thread.

Item	Value
 Threads	<u>local:28131 threst</u>
 28127	C thread 0x402e0b30
 28130	C thread 0x404e2b70 (one)
 28131	C thread 0x406e3b70 (two)
 28132	C thread 0x408e4b70 (thr)

Figure 8-5. Threads Data Item

When only a single thread is running (see “Multithreaded Programs” on page 3-49 for a description of thread control), the thread’s icon will be green with a small arrow wrapped around it, as shown in the figure above.

When the process is stopped, the current thread is shown in green underlined text.

Item	Value
 Threads	<u>local:28131 threst</u>
 28127	C thread 0x402e0b30
 #0 0x08048854 in main() at threst.c line 61	
 act	struct sigaction
 t	1083067248
 zero	0.0000000000000000
 28130	C thread 0x404e2b70 (one)
 #0 0x0804869c in one(void * unused = 0) at threst.c line 16	
 unused	0
 <u>28131</u>	<u>C thread 0x406e3b70 (two)</u>
 #0 0x080486e5 in two(void * context = 0xbffff400) at threst.c line 28	
 angle	70.6858347057698
 context	0xbffff400
 d	1.0000000000000000
 pi	3.14159265358979
 28132	C thread 0x408e4b70 (thr)
 #0 0x0804873e in thr(void * unused = 0) at threst.c line 38	
 unused	0

Figure 8-6. Threads Data Item -- Expanded

Expanding a thread will show its stack frames -- expanding a stack frame will show the local variables associated with the stack frame.

Expanding a CUDA context will show the threads in that context partitioned into a hierarchy, which the user can specify with the context menu. By default, it shows the physical hierarchy.




Figure 8-7. Threads Data Item with CUDA context and threads

You can switch context between threads by simply clicking on them in the panel. Switching context means that the current thread changes. The source panel will react to the change. Subsequent actions that cause the thread to run will cause the current thread to run, by itself when the run mode is **one**, or in tandem with all threads if the run mode is **all**. See “Run Mode Toolbar” on page 8-26 and “resume” on page 6-139.

Processes Data Item

A processes data item has sub-items to describe the processes you are debugging. The sub-items are the program’s threads. If the program is not threaded, then it is considered to have one thread. The current process is shown in green underlined text.

Shells Data Item

A shells data item has sub-items to describe the shells you have. The sub-items are the processes you are debugging in each shell. The shell that has the current process is shown in green underlined text.

Normally you only have one shell active, but at times multiple shells are useful. You can open up a remote shell and debug processes on the remote system at the same time you debug processes on the local system. This is useful when debugging client and server processes that communicate across a network.

Note that you always have at least one shell, but by default, the local shell panel isn’t displayed. You can view that shell panel by selecting **Show Shell Panel** from the View menu.

Shell panels are used to invoke programs, provide input to programs that read from `stdin`, and to see output generated by such programs.

Shell data items simply represent the active shells in your NightView session in a data panel, and show you the current processes in each shell in tree form.

Heap Information Data Item

A heap information data item has the following sub-items:

- Totals
- Configuration

The `Totals` data item has the following sub-items which show totals for the heap:

- number of blocks ever allocated
- number of bytes ever allocated
- number of additional bytes of debugger overhead ever allocated
- number of blocks ever freed
- number of bytes ever freed
- number of additional bytes of debugger overhead ever freed
- number of blocks currently allocated
- number of bytes currently allocated
- number of additional bytes of debugger overhead currently allocated
- number of blocks currently freed but still retained
- number of bytes currently freed but still retained
- number of additional bytes of debugger overhead currently freed but still retained

The `Configuration` data item has the following sub-items which show the current configuration of heap debugging:

- whether heap debugging is on or off
- number of post-fence bytes and the post-fence fill byte
- number of pre-fence bytes and the pre-fence fill byte
- number of slop bytes
- whether free filling is enabled and the free fill byte
- whether malloc filling is enabled and the malloc fill byte
- whether hardware overrun protection is enabled
- frequency of automatic heap checks (i.e. the number of heap operations between automatic heap checks)

- maximum heap size, if any
- maximum number of retained free blocks
- maximum number of walkback frames per heap operation
- whether or not to check fill bytes of free blocks

Heap Errors Data Item

A heap errors data item describes the set of heap errors most recently reported (see “Heap Check” on page 3-41). The set of errors is updated whenever new errors are discovered by a heap operation, by an automatic heap check, or by a heappoint check. Additionally, the set of errors is updated whenever a heap check is requested explicitly by the user, regardless of whether or not it detects any errors.

The heap errors data item has one sub-item for each block with any errors in the last report. Each block item has sub-items for each separate error in the block.

Leak Sets / Still Allocated Sets Data Items

A leak sets data item contains a snapshot of any allocated blocks which likely leaked at the time that the data item is created. See “Leak Detection” on page 3-42 for accuracy limitations on leak detection.

A still allocated sets data item contains a snapshot of all allocated blocks at the time that the data item is created.

Note that these data items do not update automatically. In most programs, the determination of these sets can take a considerable amount of time, and this would impose an unacceptable delay if the sets were recomputed automatically.

For both of these data items, there is a sub-item for every set of heap blocks. A set contains all heap blocks with identical sizes and walkbacks at the time of their allocations (or most recent `reallocs`), regardless of other characteristics.

For each set data item, the following sub-items exist:

- walkback
- blocks

The walkback data item has a sub-item for each stack frame in the walkback at the time of the allocation (or most recent `realloc`) operation.

The blocks data item has a block sub-item for each individual block in the set.

Block Data Item

A block data item has the following sub-items:

- its state, which will be one of:
 - allocated
 - freed, but retained

- freed or never allocated, but owned by heap
- not owned by heap

In the case of the latter two, no further sub-items are available.

- its address range
- its size in bytes
- an errors item
- an item for each heap operation (allocation, most recent `realloc`, and `free`) that has happened for the block

An errors data item contains a sub-item for each distinct error detected in the block, if any.

A heap operation (allocation, most recent `realloc`, or `free`) data item contains the following sub-items:

- configuration
- walkback

A heap operation's configuration data item contains the following sub-items:

- number and address range of post-fence bytes and the post-fence fill byte
- number and address range of pre-fence bytes and the pre-fence fill byte
- number of slop bytes
- whether free filling was enabled and the free fill byte
- whether malloc filling was enabled and the malloc fill byte
- whether hardware overrun protection was enabled

A heap operation's walkback data item has a sub-item for each stack frame in the walkback at the time of its allocation (or most recent `realloc`) operation.

Monitorpoint Values Data Item

A monitorpoint values data item has a sub-item for each monitorpoint value. See “Monitorpoints” on page 3-14.

Data Panel Context Menu

This menu pops up when you right-click in a data panel (or a context, locals or monitor panel). The menu lets you operate on the data item or its sub-items. In a data panel, the lower portion of the menu is the same as the **Data** menu and lets you add new data items or do other operations on the panel. See “Data Menu” on page 8-17. In a context, locals or monitor panel, you cannot add new items, so those menu entries do not appear.

Where you click determines which menu items appear. The menu items below appear only if they are appropriate for the type of the data item.

Collapse

Mnemonic: L

Expand

Mnemonic: X

Either **Collapse** or **Expand** is shown depending on whether the sub-items are currently expanded. Clicking on this button is the same as clicking on the + or - button.

Treat as Pointer to Single Item

Mnemonic: P

Click this radio button to consider the pointer as pointing to a single program element. The sub-item is the result of indirecting through the pointer.

Treat as Pointer to Array

Mnemonic: A

Click this radio button to consider the pointer as pointing to an array of elements, so that the sub-items are the elements of the array. The value column changes to show {array}.

Treat as Pointer to Linked List...

Mnemonic: K

Click this radio button to consider the pointer as pointing to a linked list of elements, so that the sub-items are the elements of the linked list. Clicking this radio button brings up a dialog so you can enter an expression that the debugger can use to get to successive elements. See “Data Panel Linked List Expression Dialog” on page 8-105.

The initial number of elements shown is the same as the number of indices to show when treating a pointer as an array. See “Data Panel Context Menu” on page 8-91. See “Preferences General Page” on page 8-47. As with an array, you can click on the arrowhead item ▼ to show the next element. You can also use the **Show nth Element...** item in the data panel context menu to get to a particular element. The value column changes to show {linked list: *next-expression*}.

The list is considered to terminate if a next pointer is NULL or points to an item already displayed, or causes some error.

Overriding Segment

Use this to open a sub-menu to allow you to override the default segment for a CUDA pointer. CUDA memory logically has multiple address spaces. By default, an appropriate segment will be deduced from pointer type information. But if this is not available, you may override the segment to reference a different address space. The available segments are:

- **None:** Use the default segment as determined from the pointer type.
- **Code:** Treat this as a pointer into `__code__` memory.

- **Constant:** Treat this as a pointer into `__constant__` memory.
- **Generic:** Treat this as a pointer into `__generic__` memory.
- **Global:** Treat this as a pointer into `__global__` memory.
- **Local:** Treat this as a pointer into `__local__` memory.
- **Parameter:** Treat this as a pointer into `__parameter__` memory.
- **Shared:** Treat this as a pointer into `__shared__` memory.
- **Texture:** Treat this as a pointer into `__texture__` memory.

Expand Tree...

Mnemonic: T

Expand all the sub-items of this data item and their sub-items, etc. Clicking on this button pops up a dialog box to ask you how many levels of sub-items to expand. See “Data Panel Expand Tree” on page 8-103.

Collapse Tree

Mnemonic: R

All the sub-items of this data item and their sub-items, etc., are collapsed.

Re-evaluate

Mnemonic: V

The data item is re-evaluated. The new value is displayed in the value field.

Describe...

Pop up a dialog box with additional information about this data item. See “Data Panel Describe” on page 8-103.

Show

Subscript...

Show nth Element...

Mnemonic: B (Show Subscript...)

Mnemonic: H (Show nth Element...)

This button is meaningful for arrays and linked lists. If the data item is a pointer, it is treated as an array. This menu item becomes **Show nth Element...** when the selected item is a linked list or when an array or linked list is filtered. You are prompted for a subscript or an element number. When you click on **OK**, the range of sub-items displayed is increased to include the subscript and the display is scrolled to make that sub-item visible. See “Data Panel Subscript Array” on page 8-105.

Set Watchpoint

Mnemonic: W

A new watchpoint is created for the given data item, if it an object in memory. See “Watchpoints” on page 3-16.

View Memory

Mnemonic: M

The binary viewer panel is configured to view the given data item, if it is an object in memory. If no binary viewer panel exists, one will be created. See “Binary Viewer Panel” on page 8-113.

Copy line to clipboard

Copies the selected line to the clipboard. The line is then available to be pasted into a text entry field, such as the command toolbar, with **Ctrl-V**. You may also click on either the label field or the value field and then use **Ctrl-C** to copy just that field to the clipboard.

Filter Elements with a Condition...

Change Condition Filter...

Use this to search for elements in an array or a linked list. Selecting this menu item brings up a dialog allowing you to enter a condition expression. When you click on **OK**, only elements matching the condition are displayed. See “Data Panel Condition Filter Expression Dialog” on page 8-106. If you already have a condition filter, this menu entry changes to **Change Condition Filter...**

CUDA Partitioning

Mnemonic: C

Use this to open a sub-menu with radio buttons for each of the supported CUDA thread partitioning methods. They are:

- **Physical:** Partition according to physical compute units. Specifically, this arranges threads into SM (symmetric multiprocessor), Warp, and Lane.
- **Logical:** Partition according to grids, blocks, and threads within those blocks.
- **Hierarchical Logical:** Partition according to grids, blocks, and threads within those blocks. If multiple grids are present, they will be organized by parent/child relationships.
- **PC:** Partition according to threads which are stopped at a common PC.

Select Frame

Mnemonic: S

This button is meaningful only for a local variables data item (including sub-items of stack items). The frame becomes the current frame. See “Current Frame” on page 3-33.

Show Source

Mnemonic: H

This button is meaningful only for a local variables data item (including sub-items of stack items), a heap operation data item, a walkback data item or one of its sub-items. The location indicated by the data item is listed. See “list” on page 6-86.

Check Heap and Report New Errors

Mnemonic: N

This button is meaningful only for heap errors data items. It causes a heap check to be performed, searching for new errors since the last heap check was performed. See “Heap Check” on page 3-41.

Check Heap and Report All Errors

Mnemonic: A

This button is meaningful only for heap errors data items. It causes a heap check to be performed, searching for all errors. See “Heap Check” on page 3-41.

Update Block Errors

Mnemonic: U

This button is meaningful only for a block data item or block errors data item. It causes a heap check to be performed searching for all errors in the corresponding single block. See “Heap Check” on page 3-41.

Edit...

Mnemonic: E

This button pops up a dialog box that lets you modify the expression in this data item. See “Data Panel Edit Expression” on page 8-103.

Delete

Mnemonic: D

The data item is removed from the data panel.

Overriding Format

Mnemonic: F

This button pops up a sub-menu to allow you to control how the data item is printed in the value field. The **Overriding Format** applies to the data item and to each sub-item which does not override its own format.

None
None (use *format* from **parent**)
None (use *format* from monitorpoint)

Mnemonic: N

The value is normally displayed according to its type. See “print” on page 6-96. However, the format may be overridden by the item, by a parent item, or for a monitorpoint item, by the monitorpoint.

The precedence for choosing the format to display an item is:

1. the overriding format on the item
2. the overriding format on a parent item
3. the overriding format on the monitorpoint expression, for monitorpoint items only
4. the natural format for the type of the item

The label for this menu item varies depending on whether the format is currently overridden by a parent (or by the monitorpoint for a monitorpoint item).

Address + Offset

Mnemonic: A

The value is displayed in hexadecimal and as an address relative to a program symbol. See “print” on page 6-96.

Character

Mnemonic: C

The value is displayed in character format. See “print” on page 6-96.

Decimal

Mnemonic: D

The value is displayed in decimal format. See “print” on page 6-96.

Decimal & Hexadecimal

The value is displayed in both decimal and hexadecimal format. See “print” on page 6-96.

Float

Mnemonic: F

The value is displayed in float format. See “print” on page 6-96.

Float & Hexadecimal

The value is displayed in both float and hexadecimal format. See “print” on page 6-96.

Hexadecimal

Mnemonic: H

The value is displayed in hexadecimal format. See “print” on page 6-96.

Octal

Mnemonic: O

The value is displayed in octal format. See “print” on page 6-96.

String

Mnemonic: S

The value is displayed in string format. See “print” on page 6-96.

Unsigned Decimal

Mnemonic: U

The value is displayed in unsigned decimal format. See “print” on page 6-96.

Unsigned Decimal & Hexadecimal

The value is displayed in both unsigned decimal and hexadecimal format. See “print” on page 6-96.

Default Smart Print

This item does not specify whether smart printing is enabled or disabled. It will be determined by any override setting for any parent item, or by the process-wide smart printing mode (see “smart-print” on page 6-200).

Never Smart Print

This item disables smart printing, overriding any override setting for the parent and overriding the smart printing mode (see “smart-print” on page 6-200).

Always Smart Print

This item enables smart printing, overriding any override setting for the parent and overriding the smart printing mode (see “smart-print” on page 6-200).

Find...

Mnemonic: F

Accelerator: **Ctrl+F**

Brings up the find bar and sets the keyboard focus there. See “Find Bar” on page 8-66.

Find again

Mnemonic: G

Accelerator: **Ctrl+G**

Finds the search string in the find bar again. The search begins from the selected data item. At the end of each search item found is selected.

The following menu items are similar to the items in the **Data** menu. See “Data Menu” on page 8-17.

New Expression...

The dialog box allows you to enter an expression. A data item for that expression is placed in the data panel. See “Data Panel Add Expression” on page 8-101. See “Expression Data Item” on page 8-82.

New Local Variables...

A local-variables data item is placed in the data panel. See “Local Variables Data Item” on page 8-83.

New Registers...

A registers data item is placed in the data panel. See “Registers Data Item” on page 8-83.

New Stack...

A stack data item is placed in the data panel. See “Stack Data Item” on page 8-85.

New Threads...

A threads data item is placed in the data panel. See “Threads Data Item” on page 8-87.

New Processes...

A processes data item is placed in the data panel. See “Processes Data Item” on page 8-88.

New Shells...

A shells data item is placed in the data panel. See “Shells Data Item” on page 8-88.

New Heap Information...

A heap information data item is placed in the data panel. See “Heap Information Data Item” on page 8-89.

New Heap Errors...

A heap errors data item is placed in the data panel. See “Data Panel Add Heap Errors” on page 8-102. See “Heap Errors Data Item” on page 8-90.

New Heap Leaks...

A heap leaks data item is placed in the data panel. See “Data Panel Add Heap Leaks” on page 8-102. See “Leak Sets / Still Allocated Sets Data Items” on page 8-90.

New Still Allocated Blocks

A still allocated blocks data item is placed in the data panel. See “Data Panel Add Still Allocated Blocks” on page 8-102. See “Leak Sets / Still Allocated Sets Data Items” on page 8-90.

New Monitorpoint Values

A monitorpoint values data item is placed in the data panel. See “Monitorpoint Values Data Item” on page 8-91.

Save Snapshot...

This menu item lets you save the current contents of the data panel to a text file. Clicking on this button brings up a dialog box that lets you specify the name of a file in which to save the data. You can also record a comment in the file to describe the data that is being saved. See “Data Panel Save Snapshot” on page 8-104.

Save Layout...

Selecting this menu item pops up a dialog box that lets you save the layout of all the data items for a particular process in all the data panels. The information saved includes the type and format of each data item, and to which data panel the item belongs. See “Data Panel Save Layout” on page 8-104.

Load Layout...

Selecting this menu item pops up a dialog box that lets you load a saved layout for one or more processes. Any data panels mentioned in the layout are created if they do not exist. See “Data Panel Load Layout” on page 8-103.

Set Stack Frames...

Clicking on this button pops up a dialog box that lets you set the number of stack frames displayed for items in this panel. See “Data Panel Call Stack Frames” on page 8-103.

Protected Thread

This checkbox is present for thread entries. If checked, it causes the thread to be protected. See “Protected thread tag” on page 3-51 and “Stopping” on page 3-49 for more information.

Set Thread Name...

Selecting this menu item pops up a dialog box which allows you to set the thread's name. See “set-thread-name” on page 6-163 for more information. This option is only shown for context menus associated with thread items.

Set Pointer as Array Indices...

Clicking on this button pops up a dialog box that lets you set the number of elements to show when displaying a pointer as an array. See “Data Panel Pointer Array Dimension” on page 8-104.

The remaining menu items appear only in a monitor panel or in a data panel when clicking on a monitorpoint value item.

Hold Monitor Updates

Clicking on this button causes monitorpoint items to stop updating. See “mcontrol” on page 6-124. When monitorpoint updates are held, this button reads **Release Monitor Updates** and clicking it causes monitorpoint items to resume updating.

Change Update Interval...

Mnemonic: I

Selecting this menu item pops up a dialog box that lets you change the interval between monitorpoint updates. See “Monitorpoint Update Interval Dialog Box” on page 8-107.

Show Monitor Bar

Clicking on this button causes the monitor bar to be shown. See “Monitor Bar” on page 8-79.

Vector Register Views

This button pops up a sub-menu to allow you to control how data for all registers within a specified vector register file (e.g. all xmm registers within the XMM register file) are displayed. By default, all possible formats for the register file are displayed. But if it is known that only a specific format is needed for every register within a given register file, it allows selection of the given format.

When a given format is specified here, for each register in the register file, instead of having sub-items for each format with further sub-items for each format showing the individual field values, the sub-items will show individual field values directly without the intermediate format layer.

Data Panel Dialog Boxes

This section describes the dialog boxes related to the data panel. For other dialog boxes, see “Main Window Dialog Boxes” on page 8-32.

Data Panel Item Dialog Box

This section describes common information for several of the data panel dialog boxes. These dialog boxes all contain controls to set the data panel name, plus OK, Cancel and Help buttons.

Data Panel Name

Enter the name of the data panel to receive the data item, or select a name from the list by clicking on the arrow.

If no data panel exists with this name, one is created. The default name is "Data".

OK

Click on this button to perform the operation and dismiss the window.

Cancel

Click on this button to dismiss the window without performing any operation.

Help

Click on this button to get help about the specific dialog box.

Data Panel Add Expression

This dialog box pops up when you click on **Expression...** in the **Data** menu. See “Data Menu” on page 8-17.

Enter an expression and click on **OK** to add the expression to the data panel.

Radio buttons let you select the context for later re-evaluation.

When the expression is re-evaluated, it can be evaluated in the current context at the time of the re-evaluation ("**Always evaluate in context where process stops**"), or it can be evaluated with the context saved when the expression data item is created ("**Always evaluate in context saved with expression.**").

This is a Data Panel Item Dialog Box. See “Data Panel Item Dialog Box” on page 8-101.

Data Panel Add Heap Errors

This dialog box pops up when you click on **Heap Errors...** in the **Data** menu. See “Data Menu” on page 8-17.

Click on **OK** to add a heap errors data item to the data panel.

This dialog box contains 3 mutually exclusive buttons which may be used to perform a heap check before displaying the heap errors data item. If **Don't Check Heap First**, the default, is selected, then no heap check is performed and the last reported heap errors are displayed. If **Check Heap for New Errors First** is selected, then a heap check is performed looking for new errors since the last heap check, and those errors are displayed. If **Check Heap for All Errors First** is selected, then a heap check is performed looking for all errors, and those errors are displayed. See “Heap Errors Data Item” on page 8-90.

This is a Data Panel Item Dialog Box. See “Data Panel Item Dialog Box” on page 8-101.

Data Panel Add Heap Leaks

This dialog box pops up when you click on **Heap Leaks...** in the **Data** menu or in the data panel context menu. See “Data Menu” on page 8-17.

Click on **OK** to add a leak sets data item to the data panel.

This dialog box contains 2 mutually exclusive buttons. If **New** is selected, then only new leaks since the last leak report will displayed. If **All** is selected, then all leaks will be displayed. See “Leak Sets / Still Allocated Sets Data Items” on page 8-90.

This is a Data Panel Item Dialog Box. See “Data Panel Item Dialog Box” on page 8-101.

Data Panel Add Still Allocated Blocks

This dialog box pops up when you click on **Still Allocated Blocks...** in the **Data** menu or in the data panel context menu. See “Data Menu” on page 8-17.

Click on **OK** to add a still allocated sets data item to the data panel.

This dialog box contains 2 mutually exclusive buttons. If **New** is selected, then only blocks still allocated and allocated since the last still allocated blocks report will dis-

played. If All is selected, then all blocks still allocated will be displayed. See “Leak Sets / Still Allocated Sets Data Items” on page 8-90.

This is a Data Panel Item Dialog Box. See “Data Panel Item Dialog Box” on page 8-101.

Data Panel Call Stack Frames

This dialog box pops up when you click on **Call Stack Frames...** in the **Data** menu or in the data panel context menu. See “Data Menu” on page 8-17.

Use the spin box to enter the number of frames you would like to see for items in this data panel.

Click on the check button if you want to change the number of frames displayed for all the existing items in this data panel. Otherwise, only future data items use the new number of frames.

Click on OK to complete the operation.

Data Panel Edit Expression

This dialog box pops up when you click on **Edit...** in the data panel context menu. See “Data Panel Context Menu” on page 8-91.

Change the expression as desired.

The controls are the same as the Data Panel Add Expression dialog box. See “Data Panel Add Expression” on page 8-101.

Click on OK to complete the operation.

Data Panel Expand Tree

This dialog box pops up when you click on **Expand...** in the data panel context menu. See “Data Panel Context Menu” on page 8-91.

Use the spin box to enter how many levels of sub-items you want expanded.

Click on OK to complete the operation.

Data Panel Describe

This dialog box pops up when you click on **Describe...** in the data panel context menu. See “Data Panel Context Menu” on page 8-91.

The dialog box contains additional information about the data item.

Click on OK to complete the operation.

Data Panel Load Layout

This dialog box pops up when you select **Load Layout...** in the **Data** menu or the data panel context menu. See “Data Menu” on page 8-17.

It allows you to select a file from which to load data item layout information for the current process.

This is a file selection dialog box.

Select a filename.

Use the file selection controls to select a file. If you double-click on a filename in the **Files** section, the **OK** button is activated.

Choose an action button.

Click on **Open** to load the data items and close the window.

Clicking on **Cancel** cancels the action and closes this dialog box.

Data Panel Pointer Array Dimension

This dialog box pops up when you click on **Pointer Array Dimension...** in the **Data** menu or the data panel context menu. See “Data Menu” on page 8-17.

Use the spin box to enter the number of array elements you would like to see for pointers being treated as arrays in this data panel.

Click on the check button if you want to change the number of elements displayed for all the existing items in this data panel. Otherwise, only future data items use the new number of elements.

Click on **OK** to complete the operation.

Data Panel Save Layout

This dialog box pops up when you select **Save Layout...** in the **Data** menu or the data panel context menu. See “Data Menu” on page 8-17.

It allows you to select a file in which to save the layout information for all data items for a particular process, for use in a future debug session.

This is a file selection dialog box with an additional list for selecting a process.

Select a filename.

Use the file selection controls to select a file. If you double-click on a filename in the **Files** section, the **OK** button is activated.

Select a process from the list.

Choose an action button.

Click on **OK** to save the layout and close the window.

Clicking on **Cancel** cancels the action and closes this dialog box.

Click on **Filter** to update the files displayed based on the current filter string.

You can get help for this dialog box by clicking on **Help**.

Data Panel Save Snapshot

This dialog box pops up when you select **Save Snapshot...** in the **Data** menu or the data panel context menu. See “Data Menu” on page 8-17. It allows you to select a file in which to save a snapshot of the data panel.

The items are saved at their current level of expansion. For example, if a struct is shown in the data panel, but it is not expanded (i.e., the members of the struct are not shown), then the struct is saved in the snapshot, but the members of the struct are not.

This is a file selection dialog box.

Select a filename.

Use the file selection controls to select a file. If you double-click on a filename in the **Files** section, the **OK** button is activated.

Enter comments.

Enter some comments to save with this file. For example, "Stopped in blarg just before error occurs." The comments are saved at the beginning of the snapshot, followed by a timestamp.

Select append or overwrite.

Use the radio buttons to indicate whether the snapshot should overwrite the file or append to it if it already exists.

Choose an action button.

Click on **OK** to save the snapshot and close the window.

Clicking on **Cancel** cancels the action and closes this dialog box.

Data Panel Subscript Array

This dialog box pops up when you click on **Show Subscript...** or **Show nth Element...** in the data panel context menu. See “Data Panel Context Menu” on page 8-87.

Use the spin box to enter the subscript of the array element, linked list element or filter element you want to see.

If the array or linked list is filtered, then the subscript you specify is applied to the filter, not to the underlying array or linked list. Enter **7** to see the 8th item that matches the filter.

Click on **OK** to show the element.

Data Panel Linked List Expression Dialog

This dialog pops up when you click on the **Treat as Pointer to Linked List...** radio button in data panel context menu (see “Data Panel Context Menu” on page 8-91).

Enter an expression in the combo box. The debugger uses this expression to get to the next element in the linked list. Note that any side effects in the expression, such as assignment, will affect your process. When the expression is evaluated, a temporary convenience variable (“Convenience Variables” on page 3-46), `$p`, holds the address of the current element. For example, if you have a linked list made of structures like this:

```
struct foo {  
    struct foo * next;  
    int value;  
};
```

you would enter this as the expression:

```
$p->next
```

NightView is often able to supply a suitable expression. If there are multiple pointers in the structure, there may be multiple entries in the combo box.

Click on **OK** to complete the operation.

Clicking on **Cancel** cancels the action and closes this dialog box.

Data Panel Condition Filter Expression Dialog

This dialog pops up when you click on **Filter Elements with a Condition...** or **Change Condition Filter...** in the data panel context menu.

Enter an expression in the combo box. The debugger evaluates this expression for each element of the array or linked list. The debugger displays only the elements for which the expression value is true.

When the expression is evaluated, some temporary convenience variables are available to be used in the expression. See “Convenience Variables” on page 3-46.

```
$i  is      the      index      of      the      element  
$p  is the address of the element. This is not available for some element types.  
$v  is the element.
```

Setting a condition filter does not increase the number of elements displayed for the underlying array or linked list. If you want to see more filtered elements after you set the filter, you can click on the arrowhead items   if they are available, or use **Show nth Element...** in the context menu (see “Data Panel Context Menu” on page 8-91). If no elements match, then none are displayed.

The combo box holds expressions you have entered previously.

Note that your process can be affected by side effects in the condition expression.

Example:

Stop in a main program that uses `argv`. In a local variables panel, right-click on `argv` to bring up the data panel context menu. Click on **Treat as Pointer to Array**. Then right-click on `argv` again and then click on **Filter Elements with a Condition...** Enter this condition:

```
$v != 0 && strstr($v, "HOME") != 0
```

You probably see only the up and down arrows. Click on the down arrow. You should see your `HOME` environment variable. (When your program starts, the environment variables are stored just beyond the program arguments.)

If you are having trouble getting your condition right, try an expression like this (note the comma after the `printf` call):

```
printf("%p\n", $v), $v != 0 && strstr($v, "HOME") != 0
```

This causes the program to print each pointer as the condition is evaluated.

Example:

If you have an array of structures like this:

```
struct                                foo                                {
    int                                value;
};
```

you might use this condition:

```
$v.value > 5
```

Example:

To see every fifth element, use this condition:

```
($i % 5) == 0
```

There is a limit on the number of array (or linked-list) elements NightView will search to find the next element to display. The limit is imposed to prevent NightView from using excessive resources and time in a fruitless search. You can adjust the limit here. You may also use the **Interrupt** button to stop a search. See “Process Toolbar” on page 8-24.

Click on **OK** to complete the operation. The value column for the array or linked list changes to show {filtered: *condition*}. Filtered elements are displayed in blue. The label field shows “...” at the end if the following element is not displayed due to not matching the condition.

Click on **Cancel** to cancel the action and close the dialog box.

Monitorpoint Update Interval Dialog Box

This dialog box pops up when you select **Change Update Interval...** in the context menu of the monitor panel or the data panel. See “Monitor Panel” on page 8-78. See “Data Panel” on page 8-79.

Enter the number of milliseconds to delay between updates.

Click on **OK** to change the interval and close the dialog box.

Click on **Cancel** to cancel the action and close the dialog box.

The update interval can also be changed with the **mcontrol** command. See “mcontrol” on page 6-119.

The monitorpoint update interval is not related to the eventpoint panel update interval. See “Eventpoint Panel Update Interval Dialog Box” on page 8-63.

CUDA Coordinates Panel

A CUDA coordinates panel displays the physical and logical coordinates for the current thread when stopped in a CUDA context. It is greyed out when stopped in a host thread. The physical coordinates are the SM, Warp, and Lane on the CUDA device that are executing the given thread. The logical coordinates are the grid, and block (x,y,z) , and thread (x,y,z) coordinates as declared by the application. If the application did not declare y or z values for either the block or thread, they will be greyed out.




Figure 8-8. CUDA Coordinates Panel

Most of these values are spin boxes allowing the user to enter new values or adjust the current values with arrow buttons. The logical grid is a combo box containing all the known grids executing or evicted from the CUDA device.

If the user changes the values in the physical section, the values in the logical section will change to reflect the corresponding thread executing on that physical unit. If the physical coordinates specify a compute unit that is not executing any thread at the moment, the logical section will be greyed out. Similarly, if the user changes the values in the logical section, the values in the physical section will change to reflect the corresponding compute unit executing that thread. If no compute unit is executing that thread, the physical section will be greyed out.

Each of the physical and logical sections has a **Find** button. If the current coordinates are not valid, it attempts to find the closest physical or logical coordinates, respectively, that are valid.

Changing the physical or logical coordinate values in this panel does not immediately change the current thread. But if you do wish to change it, you may use the **Apply** button to do so. If the currently displayed coordinates specify a valid thread executing on a compute unit, then the current thread will be changed to that thread. The **Reset** button may be used to reset the values in this panel back to those for the current thread.

CUDA Lanes Panel

The CUDA lanes panel allows you to quickly see which threads are being executing by the current warp. The display is broken into cells, with one cell for each lane in the warp. It displays the lane number and the logical thread coordinates for the thread being executed on that lane. Necessarily, all threads in the warp are executing from the same block, so this information is not reiterated in each block. It is available in the panel's title, though.




Figure 8-9. CUDA Lanes Panel with Divergent and Unused Lanes

If a lane is executing the current thread, the text of its cell will be displayed in green underlined text.

If a lane is executing a thread which has diverged from the current thread, then it is displayed with a different background color. In case there are multiple levels of divergence, lanes are grouped such that those which share a common background color are executing with each other, even though that subset of lanes is divergent with the current thread.

If a lane is not executing any thread, then its background color will be gray.

This panel allows changing the current thread by clicking on a valid lane.

CUDA Lanes Context Menu

The CUDA lanes panel has a context menu which allows you specify the configuration of lanes in the panel. The options are:

Arrange horizontally

Mnemonic: H

Arrange the lanes in one horizontal line. It is very likely that this will exceed the width of the display, necessitating scrolling.

Arrange as wide grid

Mnemonic: W

Arrange the lanes in a grid with 8 cells along the horizontal and 4 cells along the vertical. This usually is the most space-efficient approach that shows all lanes at once, so it is the default.

Arrange as tall grid

Mnemonic: T

Arrange the lanes in a grid with 4 cells along the horizontal and 8 cells along the vertical. This may be desirable in some user-defined panel configurations.


Arrange vertically

Mnemonic: V

Arrange the lanes in one vertical line.

CUDA Warp Locals Panel

The CUDA warp locals panel shows the values of local variables for all lanes in the current warp in one panel. It is arranged as a table with one column for each local variable, and one row for each lane in the warp. The special value threadIdx is included as the first local variable to give some context for the thread executing there.



	threadIdx	i	myFormal	result
0	divergent	divergent	divergent	divergent
1	(1,0,0)	1	0x500100000	2.00000
2	divergent	divergent	divergent	divergent
3	(3,0,0)	3	0x500100000	2.00000
4	divergent	divergent	divergent	divergent
5	(5,0,0)	5	0x500100000	2.00000
6	divergent	divergent	divergent	divergent
7	(7,0,0)	7	0x500100000	2.00000
8	divergent	divergent	divergent	divergent
9	(9,0,0)	9	0x500100000	2.00000
10	divergent	divergent	divergent	divergent
11	(11,0,0)	11	0x500100000	2.00000
12	divergent	divergent	divergent	divergent
13	(13,0,0)	13	0x500100000	2.00000

Figure 8-10. CUDA Warp Locals Panel

For any lane within the warp which is divergent, an attempt is made to determine the corresponding value in that lane, even though it is executing at a different program location. It will be displayed if the variable is defined at that other program location (i.e. the variable's scope includes that other location), and if it has a storage location (e.g. it is in a valid register). This is determined irrespective of visibility rules, so if the other program location has a definition which hides the variable being displayed, the value displayed would be that for the original variable and not the one hiding it.

This panel allows changing the current thread by clicking on the lane numbers on the left side of the table.

CUDA Warp Locals Panel Context Menu

The CUDA warp locals panel has a context menu with this item:


Configure Locals

Mnemonic: C

This opens a dialogue box. Displayed in the dialog box is a checkbox for each local variable in the routine. If you wish to prune the set of local variables because you only wish to view a few of them, you may uncheck the undesired local variables. These changes will be applied when you press the OK button. Pressing the Cancel button cancels the operation and there is no effect on the panel.

Memory Segments Panel

The memory segments panel displays memory segments in the debugged application. By default, it shows all Process memory segments in the entire application, although the set can be restricted in various ways. A memory segment is a contiguous region of an address space which was mapped by a single operation during program startup (e.g. application code, shared libraries), by later `dlopen` calls, or by user `mmap`.



Clicking on any row in the memory segments panel will cause that memory segment to be displayed in a binary viewer panel. A new binary viewer panel will be opened if one does not already exist.

A row of filters is present above the list of memory segments. They may be used to restrict the set of memory segments displayed. The filters are:

Memory Kind

The two supported states are **Process**, which shows all Process memory segments; and **Nview Internal**, which shows internal NightView patch blocks within patch areas. Users generally will be interested only in Process memory. It is not possible to show CUDA memory segments, because the CUDA driver and library do not provide this information.

Range

The memory segments can be filtered to include only those segments within a range of addresses, specified by **Start** address and number of **Bytes**.

Flags

The memory segments can be filtered based on their **Write** and **Execute** flags. For each of these fields, if the checkbox is checked, then the filter is disabled. If the checkbox is cleared, however, then any memory segments which have the flag will be excluded.

Description Regexp

The memory segments can be filtered based on an arbitrary regular expression specified in this field. A memory segment will be displayed only if its description column matches this regular expression.

The columns within the memory segment list have the following meanings:

Start

This is the start address of the memory segment.

Bytes

This is the number of bytes in the memory segment.

Flags

This is the set of flags for the memory segment. There are 3 possible flags, each represented by a letter: **r** for readable, **w** for writable, and **x** for executable. If the flag is set, that letter will appear. If the flag is not set, then the letter will be replaced with a **-** character.

Description

This is a description of the memory segment. If it is backed by a file (e.g. the static part of an application or a shared object), then it will contain the name of that file. Otherwise, it will contain a descriptive string, if any description is known. For arbitrary regions mapped by `mmap`, it may be empty.

Memory Segments Panel Context Menu

The memory segments panel has a context menu with these items:

Refresh

Mnemonic: R

Selecting this menu item causes the panel to scan the process for any changes to the memory segments since it was created, or since its last refresh.

Show only this Process

Mnemonic: P

By default, if the user changes the current context to that of a different process, the memory segments panel will change to show that new process. However, if this menu item is checked, it will be locked to whichever process was current at that time.


Reset Filters

Mnemonic: F

Selecting this menu item causes all the filter fields above the memory segments list to be reset to their default values. The default values show all Process memory segments.

Binary Viewer Panel

The binary viewer panel displays raw memory for a range of addresses in a variety of formats.



A pair of combo boxes at the bottom left control the format of the display. The first indicates the size of each data element in bytes. The second indicates the format used to display each data element, selected from the following values:

hex

Display in hexadecimal.

signed decimal

Interpret as a signed integer and display in decimal.

unsigned decimal

Interpret as an unsigned integer and display in decimal.

octal

Interpret as an unsigned integer and display in octal.

hex plus ascii

Display in hexadecimal, but also display a separate section to the right showing each byte as an ASCII character. The size is irrelevant for the ASCII portion of this display. For bytes which do not correspond to printable ASCII characters, a . (period) is displayed instead.

IEEE float

Interpret as an IEEE 754 floating-point value and display with enough decimal digits to show the least significant mantissa bits.

If the size is 16 bytes and the architecture is x86_64 (AMD64), then there are two possible formats: 80-bit double extended precision, and binary128. IEEE float selects the binary128 format. For other sizes or architectures, there is no ambiguity.

INTEL float

Interpret as an IEEE 754 floating-point value and display with enough decimal digits to show the least significant mantissa bits.

If the size is 16 bytes and the architecture is x86_64 (AMD64), then there are two possible formats: 80-bit double extended precision, and binary128. INTEL float selects the 80-bit double extended precision format. This format is supported only for 12 byte or 16 byte sizes. For the 12 byte size, it is the only possible interpretation, and so there is no ambiguity.

There are some combinations which are not permissible, such as 1 byte IEEE float, or 12 byte IEEE float.

If the range of memory being displayed does not contain an integer multiple of the data element size, then the final element will be displayed highlighted with a purple background. This indicates that it contains additional 0 byte padding which does not reflect memory beyond the end of that range (which may or may not even exist). This color is the default, but is configurable from the Preferences dialog (see “Preferences Binary Viewer Page” on page 8-59).

Binary Viewer Searching

The bottom right of the binary viewer panel contains controls for searching memory. One or more values may be entered into the **Search** field. If multiple values are specified, they must be separated by spaces or commas. Those values are interpreted as if they are of a type and display format consistent with the size and format fields. The memory is searched for data elements with those same values. If more than one value is specified, then a match will be found only if they appear in the specified order in memory.

NOTE

If values are being displayed in hexadecimal then, in the **Search** field, a value like 10 will be interpreted as 0x10 (16), even though it does not have an 0x prefix. Similarly, if values are displayed in octal, then a value like 10 will be interpreted as 010 (8), even though it lacks a 0 prefix.

Additionally, an * (asterisk) wildcard may be specified in place of any number and will match any value. This allows searches for a sequence like 1 * 3, which would find an item containing 1 followed by an item with any value and then an item containing 3.

If the data element size is 1, then the Search field may contain quoted string values. This corresponds to a sequence of bytes which correspond to the characters contained in the string. This does not include a trailing NUL character as in C strings. For example, a search for the sequence 7f "ELF" 2 would match the bytes 7f 45 4c 46 02.

NOTE

Searches are performed only on data elements that would be displayed in the panel. That is, if the start address is aligned to the data element size, then only aligned data elements are compared. If the start address is misaligned to the data element size, then only similarly misaligned data elements are compared.

By default, searches are restricted to the memory that the binary viewer panel was set to display (although the memory need not be visible on the screen). However, the **Search all memory** checkbox can be checked, which allows searching the entirety of the application's address space. This works by loading successive memory segments until a match is found (or the search wraps around to its starting point, which indicates failure).

After performing a successful search, a match will be highlighted with a green background. If a search matches only a portion of a data element, then it is highlight with a yellow background. Usually, this is possible only after changing the data element size after a previous successful search. These colors are the defaults, but are configurable from the Preferences dialog (see "Preferences Binary Viewer Page" on page 8-59).

Binary Viewer Editing

The binary viewer panel supports direct edits to memory locations. Clicking within a data item allows editing the value there. In addition, arrow keys can be used to navigate within the panel. If a value is modified, it will be displayed highlighted to indicate that an edit is in progress. If the value currently is valid, it will be highlighted with a blue background. If the value currently is invalid, it will be highlighted with a pink background. Values are interpreted as if they are of a type and display format consistent with the size and format fields, including special floating-point values like `nan` or `-inf`. These colors are the defaults, but are configurable from the Preferences dialog (see “Preferences Binary Viewer Page” on page 8-59).

If the format is **hex plus ascii**, then it also is possible to perform edits within the ASCII area on the right.

While an edit is in progress, new **Apply** and **Discard** buttons appear. The memory location is modified only when the **Apply** button is pressed. Alternatively, an edit can be rejected by pressing the **Discard** button, in which case the original value will be displayed again.

It is permissible to change multiple data elements before pressing **Apply** or **Discard**.

While performing an edit, it is not possible to change the data element size. However, it is possible to change the data element format. If so, if the modified values are valid, they will be converted into appropriate text for the new format.

Binary Viewer Panel Context Menu

The binary viewer panel has a context menu with these items:

Modify Region...

Mnemonic: M

Selecting this menu item pops up a dialog allowing modification of the region being displayed. See “Binary Viewer Modify Binary Region Dialog Box” on page 8-118.

Refresh

Mnemonic: R

Selecting this menu item causes the panel to scan the process for any changes to memory since it was created, or since its last refresh.

Clear

Mnemonic: C

Selecting this menu item clears the **Search** pattern field, and clears any highlighted search matches.

Query address of value

Mnemonic: Q

Selecting this menu item pops up a dialog window with the address of the data item over which the context menu was opened.

Show Address...

Mnemonic: A

Selecting this menu item pops up a dialogue requesting a new address within the valid ranges of addresses for the panel. Once such an address is entered, it will make that address visible in the panel.

Values per line...

Mnemonic: V

Selecting this menu item pops up a dialogue which can be used to set the number of items appearing on each line. The two values are:

Draw as many values as will fit on each line

This option will show as many data items as will fit within the horizontal dimension of the panel.

Fixed number of values per line

This option will show a specified number of data items regardless of the dimensions of the panel.

Show only this Process

Mnemonic: P

By default, if a request is made to view memory in a different process, and an existing binary viewer panel is present, its current process and range are changed to conform to the new request. However, if this menu item is checked for a binary viewer panel, then requests to view memory in a different process will not modify that binary viewer. Instead, they will create a new binary viewer panel for the request.

Show only this Address Range

Mnemonic: T

By default, if a request is made to view a different range of memory or memory in a different process, and an existing binary viewer panel is present, its current process and range are changed to conform to the new request. However, if this menu item is checked for a binary viewer panel, then requests to view either a different range of memory or memory in a different process will not modify that binary viewer. Instead, they will create a new binary viewer panel for the request.

Show Joystick

Mnemonic: J

If this menu item is checked, then an additional navigation item called the Joystick is presented. It appears as a vertical slider bar to the right of the ordinary scrollbar. When not interacted with, the slider snaps to its central position. If it is dragged to a position above its center, then the addresses shown are scrolled down. Similarly, if it is dragged to a position below its center, then the address shown are scrolled up. The distance that the slider is dragged determines the speed of scrolling in either direction.

This mechanism exists to give fine control over the memory addresses when the binary viewer is displaying a very large region of memory, such that even a minute movement of the scrollbar results in very large adjustments to the addresses shown.

Edit

If an edit is in progress, then an Edit sub-menu is available. It has the following items:

Apply all Edits

Apply all the current edits in progress. This is identical to pressing the **Apply** button.

Apply this Edit

If the context menu was opened by clicking on an individual data element being edited, then this item will apply just that single edit.

Discard all Edits

Discard all the current edits in progress. This is identical to pressing the **Discard** button.

Discard this Edit

If the context menu was opened by clicking on an individual data element being edited, then this item will discard just that single edit.

Binary Viewer Modify Binary Region Dialog Box

This dialog box pops up when you click on **Modify Region...** in the **Binary Viewer Panel Context Menu**. See “Binary Viewer Panel Context Menu” on page 8-116. It allows modification of the following properties of the region being viewed:

- **Process:** selection of a different process being debugged
- **Title:** renames the title of the panel
- **Start:** start address of the region being viewed
- **Size:** size, in bytes, of the region being viewed.

Timing Regions Panel

The **Timing Regions** panel displays timing regions and their start and end locations. The items are displayed in a tree such that each item is a:

- Timing region - This includes the associated PID, and an indication of whether the region is cross-thread, or per-thread (the default).
- Start or end location, as a child of a timing region - This includes the associated PID, and the program location.

Timing Regions Context Menu

The **Timing Regions** panel has a context menu which allows manipulation of the timing regions or start/end locations. The options are:

Delete

Mnemonic: L

Each selected item is deleted. If a timing region is selected, all its start and end locations necessarily also are deleted.

Clear Region

Mnemonic: C

This clears the timing results for each selected region. If a start or end location is selected, this applies to the timing region with which it is associated.

List Source

Mnemonic: S

Source the source corresponding to the selected start or end location. This option is not available if a timing region is selected.

Resize Columns to Data

Mnemonic: Z

Adjust the width of the columns of the table to fit the data.

Update Now

Mnemonic: U

Refresh the data in the table.

Timing Results Panel

The Timing results panel provides a table of the currently collected timing results. Each line in the table corresponds to a single timing region. The columns indicate the name of the region and individual timing data for that region. The set of columns is configurable from the context menu, and can be rearranged by clicking and dragging columns.

The columns that represent time durations default to displaying in seconds with fractional parts down to the nanosecond, and with underscores separating those fractions into milliseconds, microseconds, and nanosecond parts; but this behavior is configurable from the context menu.

By default, this panel updates only upon request, but this behavior is configurable from the context menu.

The columns are:

Name	Timing region name
Count	Number of crossings of the timing region
Minimum no subevents	The minimum time for any single crossing of the region, excluding time spent in subevents.
Minimum with subevents	The minimum time for any single crossing of the region, including time spent in subevents.
Average no subevents	The average time spent in the region, excluding time spent in subevents.
Average with subevents	The average time spent in the region, including time spent in subevents.
Maximum no subevents	The maximum time for any single crossing of the region, excluding time spent in subevents.
Maximum with subevents	The maximum time for any single crossing of the region, including time spent in subevents.
Total no subevents	The total time spent in the region, excluding time spent in subevents.
Total with subevents	The total time spent in the region, including time spent in subevents.
Info	Show errors, warnings & other diagnostics. See “Timing Info” on page 3-19.

If the application being debugged is multi-threaded, an additional combo box appears at the bottom of the panel. It can be used to control which thread(s) are displayed in the panel. It defaults to "All Threads", which combines the results from all the threads together. But an individual thread can be selected instead, in which case the results are restricted to just the selected thread.

Timing Results Context Menu

The Timing results panel has a context menu which allows control of the update behavior of the panel, the set of columns visible, and the formatting of time durations. The options are:

Delete

Mnemonic: L

Each selected timing region is deleted. All its start and end locations necessarily also are deleted.

Clear Timing Results

Mnemonic: C

This clears the timing results for each selected region.

Hold Panel

Mnemonic: H

If this item is checked (the default), the panel only updates upon request, using the Refresh Panel item below. If this item is unchecked, this panel updates automatically once per update interval, configurable using **Change Update Interval** below.

Change Update Interval...

Mnemonic: I

This item opens a dialog which allows changing the update interval for the panel. This value is meaningful only if the **Hold Panel** item above is unchecked.

Refresh Panel

Mnemonic: R

This item is used to manually request an update to the results in the panel. It can be used regardless of whether or the **Hold Panel** item above is checked.

Resize Columns

Mnemonic: Z

Adjust the width of the columns of the table to fit the data.

Save Snapshot...

Mnemonic: S

This menu item lets you save the current contents of the data panel to a text file. Clicking on this button brings up a dialog box that lets you specify the name of a file in which to save the data. You can also record a comment in the file to describe the data that is being saved.

Display Columns

Mnemonic: C

Selecting this item opens a sub-menu that lets you control which columns are displayed in the panel. Each item controls a column directly with a checkbox:

Column	Mnemonic
Count	C
Minimum (no subevents)	M
Minimum (with subevents)	
Average (no subevents)	A
Average (with subevents)	
Maximum (no subevents)	X

Column	Mnemonic
Maximum (with subevents)	
Total (no subevents)	T
Total (with subevents)	
Info	I

Additional sub-menu items control several columns at once:

Show All

Mnemonic: S

This item makes all columns visible.

Hide All

Mnemonic: H

This item makes all columns invisible (other than **Name**).

Show with subevents

This item makes all columns that include "(with subevents)" visible.

Hide with subevents

This item makes all columns that include "(with subevents)" invisible.

Show no subevents

This item makes all columns that include "(no subevents)" visible.

Hide no subevents

This item makes all columns that include "(no subevents)" invisible.

Display Formats

Mnemonic: F

Selecting this item opens a sub-menu that lets you control how columns that display time durations are formatted.

s: Seconds

Mnemonic: S

Change the units to seconds, with fractional parts.

ms: Milliseconds

Mnemonic: M

Change the units to milliseconds, with fractional parts.

us: Microseconds

Mnemonic: U

Change the units to microseconds, possibly with fractional parts, if Include nanoseconds in time fractions below is checked.

Display _ underscores in time fractions

If checked, when displaying fractional parts, use underscores to separate 3-digit triples, effectively separating milliseconds, microseconds, and nanoseconds. This is only meaningful if displaying more than 3 digits of the fraction (when displaying seconds, or when displaying milliseconds and including nanoseconds).

Include nanoseconds in time fractions

Mnemonic: N

If checked, include the fractional part down to nanoseconds. Often, this level of detail is very noisy, so disabling it can make the data easier to interpret. But the data is available if desired.

Help Window

NightView displays online help in the help window. The help window allows you to display any section of the *NightView User's Guide* and provides different methods to allow you to navigate from one section to another.

For a general discussion of NightView's online help, see “GUI Online Help” on page 8-2.

A

Summary of Commands

This section gives a summary of all the commands in NightView. The table is organized alphabetically by command. The abbreviations for the commands are included with the corresponding commands, rather than alphabetically.

Also, remember that you can abbreviate commands by using a unique prefix.

!

Pass input to a dialogue. See “!” on page 6-35 for more information.

apply on dialogue

Execute **on dialogue** commands for existing dialogues. See “apply on dialogue” on page 6-34 for more information.

apply on program

Execute **on program** commands for existing processes. See “apply on program” on page 6-50 for more information.

attach

Attach the debugger to a process that is already running. See “attach” on page 6-41 for more information.

backtrace

bt

Print an ordered list of the currently active stack frames. See “backtrace” on page 6-92 for more information.

branch-history

bh

Display the branch history if branch tracking has been enabled. See “branch-history” on page 6-106 for more information.

breakpoint

b

Set a breakpoint. See “breakpoint” on page 6-109 for more information.

cd

Set the debugger’s default working directory. See “cd” on page 6-82 for more information.

checkpoint

Take a restart checkpoint now. See “checkpoint” on page 6-51 for more information.

clear

Clear all eventpoints at a given location. See “clear” on page 6-120 for more information.

commands

Attach commands to a breakpoint or monitorpoint. See “commands” on page 6-121 for more information.

condition

Attach a condition to an eventpoint. See “condition” on page 6-122 for more information.

continue

c

Continue execution and wait for something to happen. See “continue” on page 6-133 for more information.

core-file

Create a pseudo-process for debugging an aborted program's core image file. See “core-file” on page 6-44 for more information.

data-display

Control items in a data panel. See “data-display” on page 6-101 for more information.

debug

Specify names for programs you wish to debug. See “debug” on page 6-28 for more information.

declare-thread-tag

Declare the type of a thread tag. See “declare-thread-tag” on page 6-156 for more information.

define

Define a NightView macro. See “define” on page 6-186 for more information.

delay

Delay NightView command execution for a specified time. See “delay” on page 6-159 for more information.

delete

d

Delete an eventpoint. See “delete” on page 6-123 for more information.

detach

Stop debugging a list of processes. See “detach” on page 6-42 for more information.

directory

Set the directory search path. See “directory” on page 6-85 for more information.

disable

Disable an eventpoint. See “disable” on page 6-123 for more information.

display

Add to the list of expressions to be printed each time the process stops. See “display” on page 6-102 for more information.

down

Move one or more stack frames toward frames called by the current stack frame. See “down” on page 6-151 for more information.

echo

Print arbitrary text. See “echo” on page 6-100 for more information.

edit

Edit the current source file. See “edit” on page 6-86 for more information.

enable

Enable an eventpoint for a specified duration. See “enable” on page 6-124 for more information.

exec-file

Specify the location of the executable file corresponding to a process. See “exec-file” on page 6-47 for more information.

family

Give a name to a family of one or more processes. See “family” on page 6-52 for more information.

finish

Continue execution until the current function finishes. See “finish” on page 6-141 for more information.

forward-search

fo

Search forward through the current source file for a specified regular expression. See “forward-search” on page 6-87 for more information.

frame

f

Select a new stack frame or print a description of the current stack frame. See “an Ada task, ,frame” on page 6-149 for more information.

handle

Specify how to handle signals in the user process. See “handle” on page 6-145 for more information.

heapcheck

Check the heap for errors. See “heapcheck” on page 6-181 for more information.

heapdebug

hd

Specify parameters for heap debugging. See “heapdebug” on page 6-57 for more information.

heappoint

Check the heap for errors, or change the heap debugger settings, at a given location. See “heappoint” on page 6-118 for more information.

help

Access the online help system. See “help” on page 6-154 for more information.

ignore

Attach an ignore-count to an eventpoint. See “ignore” on page 6-125 for more information.

info address

Determine the location of a variable. See “info address” on page 6-183 for more information.

info args

Print description of current routine arguments. See “, or Ada unit namesinfo args” on page 6-182 for more information.

info breakpoint

i b

Describe current state of breakpoints. See “info breakpoint” on page 6-161 for more information.

info convenience

Describe convenience variables. See “info convenience” on page 6-169 for more information.

info debuginfo

Describe debug information for each executable and shared object.

info declaration

ptype

Print the declaration of variables or types. See “info declaration” on page 6-185 for more information.

info dialogue

Print information about active dialogues. See “info dialogue” on page 6-176 for more information.

info directories

Print the search path used to locate source files. See “info directories” on page 6-169 for more information.

info display

Describe expressions that are automatically displayed. See “info display” on page 6-169 for more information.

info eventpoint

Describe current state of breakpoints, tracepoints, patchpoints, monitorpoints, heap-points, and watchpoints. See “info eventpoint” on page 6-160 for more information.

See “info exception” on page 6-185 for more information.**info family**

Print information about an existing process family. See “info family” on page 6-176 for more information.

info files

Print the names of the executable, symbol table and core files. See “info files” on page 6-185 for more information.

info frame

Describe a stack frame. See “info frame” on page 6-168 for more information.

info functions

List names of functions or subroutines. See “info functions” on page 6-183 for more information.

info heappoint

Describe the current state of heappoints. See “info heappoint” on page 6-165 for more information.

info history

Print value history information. See “info history” on page 6-169 for more information.

info limits

Print information about limits on expression and location output. See “info limits” on page 6-170 for more information.

info line

Describe location of a source line. See “info line” on page 6-185 for more information.

info locals

Print information about local variables. See “info locals” on page 6-182 for more information.

info log

Describe any open log files. See “Ada tasks,,info log” on page 6-160 for more information.

info macros

Print a description of one or more NightView macros. See “info macros” on page 6-191 for more information.

info memory

Print information about memory, which may include information about the virtual address space, or the heap. See “info memory” on page 6-172 for more information.

info monitorpoint

Describe current state of monitorpoints. See “info monitorpoint” on page 6-164 for more information.

info name

Print information about an existing eventpoint-name. See “info name” on page 6-176 for more information.

info on dialogue

Print **on dialogue** commands. See “info on dialogue” on page 6-177 for more information.

info on program

Print **on program** commands. See “info on program” on page 6-177 for more information.

info on restart

Print **on restart** commands. See “info on restart” on page 6-177 for more information.

info patchpoint

Describe current state of patchpoints. See “info patchpoint” on page 6-163 for more information.

info process

Describe processes being debugged. See “info process” on page 6-171 for more information.

info registers

Print information about registers. See “info registers” on page 6-170 for more information.

**info representation
representation**

Describe the storage representation of an expression. See “info representation” on page 6-184 for more information.

info signal

Print information about signals. See “info signal” on page 6-171 for more information.

info sources

List names of source files. See “info sources” on page 6-183 for more information.

info syscallpoint

List information about syscallpoints. See “info syscallpoint” on page 6-167 for more information.

info threads

Describe C threads and thread processes. See “info threads” on page 6-179 for more information.

info tracepoint

Describe current state of tracepoints. See “info tracepoint” on page 6-162 for more information.

info types

Print type definition information. See “info types” on page 6-184 for more information.

info variables

Print global variable information. See “info variables” on page 6-182 for more information.

info watchpoint

Describe current state of watchpoints. See “info watchpoint” on page 6-166 for more information.

info whatis

whatis

Describe the result type of an expression visible in the current context. See “info
whatis” on page 6-184 for more information.

interest

Control which subprograms are interesting. See “interest” on page 6-71 for more
information.

jump

Continue execution at a specific location. See “jump” on page 6-143 for more infor-
mation.

kill

Terminate a list of processes. See “kill” on page 6-43 for more information.

list

l

List a source file. See “list” on page 6-83 for more information.

load

Dynamically load an object file, possibly replacing existing routines. See “load” on
page 6-105 for more information.

login

Login to a new dialogue shell. See “login” on page 6-26 for more information.

logout

Terminate a dialogue. See “logout” on page 6-32 for more information.

mcontrol

hold

release

Control the monitorpoint value display. See “mcontrol” on page 6-119 for more
information.

memory-display

md

Control memory displayed in a binary viewer panel. See “memory-display” on page
6-101 for more information.

modify-eventpoint

Modify eventpoint hits, crossing counts, and protected flag. See “modify-event-
point” on page 6-126 for more information.

monitorpoint

Monitor the values of one or more expressions at a given location. See “monitorpoint” on page 6-116 for more information.

mreserve

Reserve a region of memory in a process. See “mreserve” on page 6-55 for more information.

name

Give a name to a group of eventpoints. See “name” on page 6-108 for more information.

next

n

Execute one line, stepping over procedures. See “next” on page 6-137 for more information.

nexti

ni

Execute one instruction, stepping over procedures. See “nexti” on page 6-140 for more information.

nodebug

Specify names for programs you do not wish to debug. See “nodebug” on page 6-29 for more information.

notify

Ask about pending event notifications. See “notify” on page 6-41 for more information.

on dialogue

Specify debugger commands to be executed when a dialogue is created. See “on dialogue” on page 6-32 for more information.

on program

Specify debugger commands to be executed when a program is `execed`. See “on program” on page 6-48 for more information.

on restart

Specify debugger commands to be executed when a program is restarted. See “on restart” on page 6-50 for more information.

output

Print the value of a language expression with minimal output. See “output” on page 6-100 for more information.

patchpoint

Install a small patch to a routine. See “patchpoint” on page 6-111 for more information.

print

p

Print the value of a language expression. See “print” on page 6-92 for more information.

printf

Print the values of language expressions using a format string. See “printf” on page 6-105 for more information.

pwd

Print NightView's current working directory. See “pwd” on page 6-82 for more information.

quit

q

Stop everything. Exit the debugger. See “quit” on page 6-25 for more information.

redisplay

Enable a display item. See “redisplay” on page 6-104 for more information.

refresh

Re-read source files and refresh the terminal screen. See “refresh” on page 6-155 for more information.

rerun

Run a program again. See “rerun” on page 6-39 for more information.

resume

Continue execution. See “resume” on page 6-134 for more information.

reverse-search

Search backwards through the current source file for a specified regular expression. See “reverse-search” on page 6-87 for more information.

run

Run a program in a dialogue and wait for NightView to start debugging it. See “run” on page 6-39 for more information.

save-core-file

Saves the core file and any required shared library files needed for subsequent core file analysis in a compressed file. See “save-core-file” on page 6-45 for more information.

select-context

Select the context of a thread a thread process, or a CUDA context. See “select-context” on page 6-152 for more information.

set

Evaluate a language expression without printing its value. See “set” on page 6-95 for more information.

set-auto-frame

Control the positioning of the stack when a process stops. See “set-auto-frame” on page 6-74 for more information.

set-branch-tracking

Control whether or not NightView and the RedHawk kernel are tracking branch instructions. See “set-branch-tracking” on page 6-79 for more information.

set-children

Control whether children should be debugged. See “set-children” on page 6-53 for more information.

set-cuda

Control whether or not CUDA support is enabled in this session. See “set-cuda” on page 6-80 for more information.

set-cuda-memcheck

Enable or disable CUDA memcheck, which makes CUDA exceptions more precise at the expensive of very high overhead. See “set-cuda-memcheck” on page 6-84 for more information.

set-cuda-software-preemption

Allow a CUDA device to be used by other applications when an application running on the device is stopped for debugging. See “set-cuda-software-preemption” on page 6-80 for more information.

set-debug-file-directory

Tell NightView where to look for .debug files. See “set-debug-file-directory” on page 6-29 for more information.

set-disassembly

Control how NightView displays disassembled instructions. See “set-disassembly” on page 6-78 for more information.

set-download

Control how NightView downloads files from remote targets. See “set-download” on page 6-77 for more information.

set-echo

Controls whether or not each subsequent user command is echoed back. See “set-echo” on page 6-63 for more information.

set-editor

Set the mode for editing commands in the simple full-screen interface. See “set-editor” on page 6-75 for more information.

set-exit

Control whether a process stops before exiting. See “set-exit” on page 6-54 for more information.

set-format-hex

Control whether or not many values are displayed in hexadecimal in addition to their natural format.

See “set-format-hex” on page 6-68 for more information.

set-futurepoints

Control whether or not NightView accepts location specifiers for locations which do not exist yet. See “set-futurepoints” on page 6-79 for more information.

set-history

Specify the number of items to be kept in the value history list. See “set-history” on page 6-65 for more information.

set-language

Establish a default language context for variables and expressions. See “set-language” on page 6-64 for more information.

set-limits

Specify limits on the number of array elements, string characters, or program addresses printed when examining program data. See “set-limits” on page 6-66 for more information.

set-local

Define process local convenience variables. See “set-local” on page 6-70 for more information.

set-log

Log session to file. See “set-log” on page 6-63 for more information.

set-notify

Control how you are notified of events. See “set-notify” on page 6-40 for more information.

set-overload

Control how NightView treats overloaded operators and routines in expressions. See “set-overload” on page 6-74 for more information.

set-patch-area-size

Control the size of patch areas created in your process. See “set-patch-area-size” on page 6-70 for more information.

set-preallocate

Control how NightView preallocates memory for eventpoints and monitorpoint buffers. See “set-preallocate” on page 6-75 for more information.

set-prompt

Set the string used to prompt for command input. See “set-prompt” on page 6-67 for more information.

set-qualifier

Specify the default list of processes or dialogues that will be affected by subsequent commands which accept qualifiers. See “set-qualifier” on page 6-65 for more information.

set-restart

Control whether restart information is applied. See “set-restart” on page 6-69 for more information.

set-resume

Control NightView’s behavior on events for which you might want the process to stop automatically. See “set-resume” on page 6-76 for more information.

set-run-mode

Control whether all threads in a multi-threaded process resume execution when a single thread is told to resume execution.

set-safety

Control debugger response to dangerous commands. See “set-safety” on page 6-69 for more information.

set-search

Control case sensitivity of regular expressions in NightView. See “set-search” on page 6-75 for more information.

set-shared-lib-update

Control whether NightView re-reads debug information each time a new shared object is loaded.

set-show

Control where dialogue output goes. See “set-show” on page 6-36 for more information.

set-tag

Modify the value of a thread tag. See “set-tag” on page 6-157 for more information.

set-terminator

Set the string used to recognize end of dialogue input mode. See “set-terminator” on page 6-68 for more information.

set-thread-name

Set the name of a thread. See “set-thread-name” on page 6-158 for more information.

set-trace

Establish tracing parameters. See “set-trace” on page 6-114 for more information.

shell

Run an arbitrary shell command. See “shell” on page 6-155 for more information.

show

Control dialogue output. See “show” on page 6-37 for more information.

signal

Continue execution with a signal. See “signal” on page 6-144 for more information.

smart-print

Define, undefined, view, enable, or disable smart printers. See “smart-print” on page 6-193 for more information.

source

Input commands from a source file. See “source” on page 6-156 for more information.

step

s

Execute one line, stepping into procedures. See “step” on page 6-136 for more information.

stepi

si

Execute one instruction, stepping into procedures. See “stepi” on page 6-139 for more information.

stop

Stop a process. See “stop” on page 6-142 for more information.

symbol-file

Establish the file containing symbolic information for a program. See “symbol-file” on page 6-43 for more information.

syscallpoint

Set an eventpoint to trace one or more system calls. See “syscallpoint” on page 6-130 for more information.

tbreak

Set a temporary breakpoint. See “tbreak” on page 6-126 for more information.

tpatch

Set a patchpoint that will execute only once. See “tpatch” on page 6-127 for more information.

tracepoint

Set a tracepoint. See “tracepoint” on page 6-114 for more information.

translate-object-file

x1

Translate object filenames for a remote dialogue. See “translate-object-file” on page 6-30 for more information.

undisplay

Disable an item from the display expression list. See “undisplay” on page 6-104 for more information.

up

Move one or more stack frames toward the caller of the current stack frame. See “up” on page 6-151 for more information.

wait

Wait for processes to stop. See “wait” on page 6-55 for more information.

watchpoint

Set a watchpoint. See “watchpoint” on page 6-128 for more information.

x

Print the contents of memory beginning at a given address. See “x” on page 6-96 for more information.

B**Quick Reference Guide****Invoking NightView**

```
nview [-attach pid|name] [-config config-file] [-core core-file]
      [-help] [-nogui] [-nolocal] [-nx] [-prompt string]
      [-safety safe-mode] [-simplescreen] [-version]
      [-Xoption ...] [-x command-file]
      [program-name [program-argument ...]]
```

Controlling the Debugger**Quitting NightView**

```
quit
```

Abbreviation: **q**

Managing Dialogues

```
login [/conditional] [/popup] [name=dialogue name]
      [user=login name] [others ...] machine
```

```
debug pattern ...
```

```
nodebug pattern ...
```

```
set-debug-file-directory [path]
```

```
translate-object-file [from [to]]
```

Abbreviation: **xl**

```
logout
```

```
on dialogue [regexp]
```

```
on dialogue regexp command

on dialogue regexp do

apply on dialogue
```

Dialogue Input and Output

```
! [input line]

set-show [silent | notify=mode | continuous=mode]
         [log[=filename]] [buffer=number]

show [number | all | none] [| shell-command]
```

Managing Processes

```
run input line

rerun

set-notify [silent | continuous=mode]

notify

attach [{/resume | /stop}] { pid | name }

detach

kill

symbol-file program-name

core-file corefile-name [exec-file=program-name]
                       [with-translations]

save-core-file [/nozip] [/nodebuginfo] [/replace] [/keep]
               [include=file] [note=string] savename

exec-file program-name

on program [pattern]

on program pattern command

on program pattern do

apply on program
```

```

on restart [pattern]

on restart pattern command

on restart pattern do

checkpoint

family family-name [[-] qualifier-spec ] ...

set-children { all [ resume ] | exec | none }

set-exit [stop | nostop]

wait [{all | any} [new]]

mreserve start=address {length=bytes | end=address}

```

Heap Debugging

```

heapdebug [check_free_fill={0|1}]
            [common_errors={block_overrun |
                           dangling_pointer |
                           uninitialized_field}]
            [do_free_fill={0|1}]
            [do_malloc_fill={0|1}]
            [error-name [{noprint |
                           nostop |
                           print |
                           stop } ...]]
            [free_fill_byte={n | trash}]
            [frequency=n[{k|m}]]
            [heap_size={n[{k|m}] | unlimited}]
            [internal_checks={0|1}]
            [level={0|1|2|3}]
            [malloc_fill_byte={n | trash}]
            [off]
            [on]
            [post_fence_size=n]
            [post_fill_byte={n | trash}]
            [pre_fence_size=n]
            [pre_fill_byte={n | trash}]
            [protected={0|1}]
            [retain_free_blocks={n[{k|m}] | unlimited}]
            [slop=n]
            [walkback=n]

```

error-name can be any of the following:

```
free_fill_modified
free_not_at_beginning
free_unallocated
internal_error
malloc_zero
memalign_not_power_2
out_of_memory
post_fence_modified
pre_fence_modified
realloc_not_at_beginning
realloc_unallocated
```

Abbreviation: **hd**

Setting Modes

set-log *keyword filename*

set-language {auto | c | c++ | fortran}

set-qualifier [*qualifier-spec ...*]

set-history *count*

set-limits {array=*number* | string=*number* |
addresses=*number*} ...

set-format-hex [{on | off}]

set-prompt *string*

set-terminator *string*

set-safety [forbid | verify | unsafe]

set-restart [always | never | verify]

set-local *identifier ...*

set-patch-area-size {data=*data-size* |
eventpoint=*eventpoint-size* |
monitor=*monitor-size* |
text=*text-size*} ...

interest [*level*] [[at] [*location-spec*]]

interest inline[=*level*]

interest justlines[=*level*]

interest nodebug[=*level*]


```

interest cuda_syscall[=level]

interest threshold[=level]

set-auto-frame args...

set-overload [ operator={on | off} ]
               [ routine={on | off} ]

set-search [ sensitive | insensitive ]

set-editor mode

set-preallocate [/eventpoint] [/monitorpoint] [{off | on}]

set-resume [/attach] [/exec] [/exit] [/fork]
            [/cuda] [/cuda_system]
            [{off | on}]

set-download [{off | permanent | temporary}]
              [directory=path-to-cache]

set-disassembly [flavor={att | intel}]
                 [symbols={off | on}]
                 [comment_level=number]

set-branch-tracking [{on | off}]

set-futurepoints [{create | ask | error}]

set-cuda [{off | on}]

set-cuda-memcheck [{off | on}]

set-cuda-software-preemption [{off | on}]

set-echo [{off | on}]

```

Debugger Environment Control

```

cd dirname

pwd

```

Source Files

Viewing and Editing Source Files

```
list where-spec  
list where-spec1, where-spec2  
list ,where-spec  
list where-spec,  
list +  
list -  
list =  
list  
Abbreviation: l
```

```
directory [dirname ...]  
edit
```

Searching

```
forward-search [regex]  
Abbreviation: fo  
reverse-search [regex]
```

Examining and Modifying

```
backtrace [number-of-frames]  
Abbreviation: bt  
print [/print-format-code] expression  
Abbreviation: p
```

set *expression*

x [/repeat-count] [size-letter] [x-format-code]] [*addr-expression*]

output [/print-format-code] *expression*

echo *text*

data-display [/window="window name"]
 {/kind=value | *expression*}

memory-display [/title="window name"] [/group=group size]
 [/format=format] [/start=offset]
 [/context=CUDA context|process]
 /bytes=size[b|k|m|g] *expression*

Abbreviation: **md**

display [/print-format-code] *expression*

display /repeat-count] [size-letter] [x-format-code] *addr-expression*

undisplay *item-number* ...

redisplay *item-number* ...

printf *format-string* [, *expression* ...]

load *object*

branch-history [number-of-branches]

Manipulating Eventpoints

name [/add] *name* [[-] *eventpoint-spec*] ...

breakpoint [*eventpoint-modifier*] [/cuda | /process]
 [name=breakpoint-name] [[at] *location-spec*]
 [if *conditional-expression*]

Abbreviation: **b**

patchpoint [*eventpoint-modifier*] [name=patchpoint-name]
 [[at] *location-spec*] eval *expression*

patchpoint [*eventpoint-modifier*] [name=patchpoint-name]
 [[at] *location-spec*] goto *location-spec*

```

set-trace [eventmap=event-map-file]

tracepoint [eventpoint-modifier] event-id [name=tracepoint-name]
           [[at] location-spec]
           [value=logged-expression [, logged-expression...]]
           [if conditional-expression]

monitorpoint [eventpoint-modifier] [name=monitorpoint-name]
           [[at] location-spec]

heappoint [eventpoint-modifier] [name=heappoint-name]
           [[at] location-spec]
           [{check | debug parameters}]
           [if conditional-expression]

mcontrol {display | nodisplay} [monitorpoint-spec ...]

mcontrol delay milliseconds

mcontrol {off | on | stale | nostale | hold | release}

           Abbreviation: hold

           Abbreviation: release

clear [[at] location-spec]

commands eventpoint-spec

condition eventpoint-spec [conditional-expression]

delete [eventpoint-spec ...]

           Abbreviation: d

disable [eventpoint-spec ...]

enable [/once|/delete] [eventpoint-spec ...]

ignore eventpoint-spec count

modify-eventpoint eventpoint-spec [hits=hits] [crossings=crossings]
           [protected={on|off}]

tbreak [name=breakpoint-name] [[at] location-spec]
           [if conditional-expression]

tpatch [name=patchpoint-name]
           [[at] location-spec] eval expression

tpatch [name=patchpoint-name]
           [[at] location-spec] goto location-spec

```

watchpoint [*eventpoint-modifier*] [/once] [/read] [/write]
 [name=*watchpoint-name*] [at] *lvalue*
 [if *conditional-expression*]

watchpoint [*eventpoint-modifier*] [/once] [/read] [/write]
 /address
 [name=*watchpoint-name*]
 [at] *address-expression*
 {size *size-expression* | type *expression*}
 [if *conditional-expression*]

syscallpoint [*eventpoint-modifier*] [/nostop] [/before] [/after]
 [name=*patchpoint-name*] [*syscall-list*]
 [if *conditional-expression*]

Controlling Execution

continue [*count*]

Abbreviation: **c**

resume [*sigid*]

step [*repeat*]

Abbreviation: **s**

next [*repeat*]

Abbreviation: **n**

stepi [*repeat*]

Abbreviation: **si**

nexti [*repeat*]

Abbreviation: **ni**

finish

stop

jump [at] *location-spec*

signal *sigid*

handle [/signal] *sigid keyword ...*

Selecting Context

frame [*frame-number*]

frame **expression* [at *location-spec*]

Abbreviation: **f**

up [*number-of-frames*]

down [*number-of-frames*]

select-context default

select-context thread=*expression*

select-context pid=*pid*

select-context cuda context *context*

select-context cuda sm *sm*

select-context cuda warp *warp*

select-context cuda lane *lane*

select-context cuda block *x[,y[,z]]*

select-context cuda thread *x[,y[,z]]*

Miscellaneous Commands

help [*section*]

refresh

shell [*shell-command*]

source *command-file*

delay [*milliseconds*]

Info Commands

Status Information

info log

info eventpoint [/verbose] [*eventpoint-spec*] ...

info breakpoint [/verbose] [*eventpoint-spec*] ...

Abbreviation: **i b**

info tracepoint [/verbose] [*eventpoint-spec*] ...

info patchpoint [/verbose] [*eventpoint-spec*] ...

info monitorpoint [/verbose] [*eventpoint-spec*] ...

info heappoint [/verbose] [*eventpoint-spec*] ...

info syscallpoint [*eventpoint-spec*] ...

info watchpoint [/verbose] [*eventpoint-spec*] ...

info frame [/v] [**expression* [at *location-spec*]]

info directories

info convenience

info display

info history [*number*]

info limits

info registers [*regex*]

info signal [*signal* ...]

info process

info memory [/ranges] [/heap] [/leaks] [/allocated]
 [/all] [/append=*filename*] [/output=*filename*]
 [/verbose] [*expression*]

info dialogue

info family [*regex*]

```
info name [regex]  
  
info on dialogue [name]  
  
info on program [program]  
  
info on restart [output=outname|append=outname] [program]  
  
info threads [/verbose] [/cuda]  
                [/partition { physical | logical | pc }]  
  
heapcheck [/all] [/append=filename] [/output=filename] [expression]
```

Symbol Table Information

```
info args  
  
info locals [regex]  
  
info variables [regex]  
  
info address identifier  
  
info sources [pattern]  
  
info functions [regex]  
  
info types [regex]  
  
info whatis expression  
    Abbreviation: whatis  
  
info representation expression  
    Abbreviation: representation  
  
info declaration regex  
    Abbreviation: ptype  
  
info files  
  
info line [at] location-spec
```


Defining and Using Macros

```
define macro-name [ (arg-name [, arg-name] ...) ] [text]  
define macro-name [ (arg-name [, arg-name] ...) ] as  
info macros [regex]
```

Smart Printing

```
smart-print info [pattern]  
smart-print { [ on | off ] }  
smart-print replace pattern  
          replace-definition  
end-smart-print  
smart-print struct pattern  
          struct-definition  
end-smart-print  
smart-print container pattern  
          container-definition  
end-smart-print  
smart-print undef pattern  
smart-print reload
```


C

Implementation Overview

This section gives a very high-level description of how the debugger is implemented.

The user invokes **nview**. **nview** is a script that runs either **snview** or **xnview**. **snview** implemen

ts the command-line and simple full-screen interfaces. **xnview** implements the graphical user interface. (Users are discouraged from invoking **snview** or **xnview** directly.) The user interface programs deal with all aspects of the user interface and with managing the symbolic debugging information from executable files. See Chapter 5 [Invoking NightView] on page 5-1.

NightView runs an instance of `NightView.p` for each dialogue. This process is run on the local system, regardless of whether the dialogue is on the local system or on a remote system. The `NightView.p` process is aware of many low-level architectural details, which may differ between the local system if the dialogue is for a remote system.

In addition, NightView runs `NightView.pp` for each dialogue. If the dialogue is for the local system, then NightView runs `NightView.pp` on the local system, and communicates with it via a shared memory region. There is one such shared memory region per invocation of NightView. See “Dialogues” on page 3-6. For remote dialogues, NightView runs `NightView.pp` on the remote system, and establishes a socket connection with it.

`NightView.pp` is responsible for controlling the user processes by a combination of the `/proc` file system and the **ptrace** system service.

Monitorpoints communicate with `NightView.pp` via a shared memory region created in your process. There is one shared memory region for each dialogue using monitorpoints. See “Monitorpoints” on page 3-14. The shared memory region is placed in your process at a preferred address if that address is available. Otherwise, it is placed anywhere NightView can find space. On IA-32, the preferred address is `0xafe78000`. On AMD64, the preferred address is `0x15aaaa2000`. On ARM64, the preferred address is `0x3ffffe0000`.

Each dialogue runs a shell and controls it using `/proc` and **ptrace**. This is not to debug the shell, but so that the debugger is notified of the shell's children, which are the processes to be debugged. The shell interacts with a pseudo-terminal (pty) controlled by the debugger, so that the debugger can capture the program I/O.

Watchpoints are implemented by setting architecture-specific hardware registers intended for this purpose (`DRn` registers on IA-32 or AMD64 systems, or `DBGWCRn` & `DBGWVRn` on ARM64). Syscallpoints are implemented using the `ptrace` (`PTRACE_SYSCALL`) system call. Other eventpoints are implemented in architecture-specific ways: On IA-32 and AMD64 systems, the instruction at the target address is replaced by a trap instruction. When the debugged program hits the trap, the kernel translates this into a branch to a patch area. On ARM64, the instruction at the target address is replaced by a branch to the patch area. In either case, the patch area contains instructions to implement the particular eventpoint, emulate the replaced instruction, and return to the target address.

Space for a patch area is acquired by using `mmap` or by creating a shared memory region in the process's address space. The debugger usually creates one data patch area, one text patch area, and multiple eventpoint patch areas. The user can adjust the sizes of the patch areas. See “set-patch-area-size” on page 6-73. By default, NightView preallocates patch areas to avoid needing to allocate them later, which can impact real-time performance. But this behavior can be controlled. See “set-preallocate” on page 6-78 for the command.

On IA-32, NightView tries to put all the patch areas between 0xa0000000 and 0xb0000000.

On AMD64 and ARM64, an eventpoint patch area may need to be placed near the instruction being patched. On AMD64, NightView tries to put all the other patch areas between 0x15aaaab000 and 0x2000000000. On ARM64, NightView tries to put all the other patch areas between 0x4000000000 and 0x6000000000.

You can see where NightView has placed patch areas with the **info memory** command (see “info memory” on page 6-172).

The user process is sometimes forced to execute code on behalf of the debugger. This is how function calls work in evaluated expressions, and it is also used to do some of the housekeeping chores, e.g., creating memory regions. On AMD64, NightView may need to create additional memory areas to hold this code.

To implement heap debugging, NightView patches a heap debugger module into your program and arranges for calls to `malloc`, `free`, and other heap functions, to be intercepted by this module. The heap debugger modifies each allocation request to allow for slop and fences and then passes the new request to the system allocator. The heap debugger remembers where each block is and what attributes it has.

To implement CUDA support, NightView sets a hidden trap to detect the `dlopen()` of **libcudart.so**, if any should occur. If such a `dlopen()` does occur, NightView sets flags in the debugged process to inform the CUDA library that it is to be debugged. This causes it to create a clone process which acts as an intermediary between the debugged CUDA code running on the device and **NightView.pp**. Communication with this intermediary is established using named pipes and all interactions with the CUDA code take place through this intermediary.

Reporting Bugs

To report a problem or request software assistance, contact Concurrent Real-Time via the web, email, or by phone.

Support URL

<http://www.concurrent-rt.com/support>

Email

support@concurrent-rt.com

Phone

800.245.6453 (954.283.1822 for customers outside the United States).

D

Tutorial Files

The following sections show source listings for the files used in the tutorials. These files all reside under the `/usr/lib/NightView/Tutorial` directory.

-

C Files

msg.h

```

1  #include <stdlib.h>
2  #include <stdio.h>
2  #include <sys/types.h>
3  #include <signal.h>

```

main.c

```

1  #include "msg.h"
2
3  /* This program spawns a child process and sends
4   * signals from the parent to the child.
5   *
6   */
7
8  main()
9  {
10     int total_sig;
11     pid_t pid;
12     char *tracefile = "msg_file";
13     extern void parent_routine();
14     extern void child_routine();
15     extern void signal_handler();
16
17     signal( SIGUSR1, signal_handler );
18     printf( "How many signals should the parent send the child?\n" );
19     scanf( "%d", &total_sig );
20     pid = fork();
21
22     if( pid == 0 )
23     {
24         /* It's the child */
25         child_routine( total_sig );

```

```
26     }
27
28     else
29     {
30         /* It's the parent */
31         parent_routine( pid, total_sig );
32     }
33
34     exit( 0 );
35 }
```

parent.c

```
1  #include "msg.h"
2
3  /* Every time the parent sends the child a signal,
4   * the parent writes a message.
5   */
6
7  void parent_routine( child_pid, total_sig )
8  pid_t child_pid;
9  int    total_sig;
10 {
11     int isec = 2;
12     int sig_ct;
13
14     for( sig_ct = 1; sig_ct <= total_sig; ++sig_ct )
15     {
16         printf( "%d. Parent sleeping for %d seconds\n", sig_ct, isec
17 );
18         sleep( isec );
19         kill( child_pid, SIGUSR1 );
20     }
21 }
```

child.c

```
1  #include "msg.h"
2
3  /* Every time the child receives a signal from
4   * the parent, the child writes a message.
5   */
6
7  int sig_ct_child = 0;
8
9  void child_routine( total_sig )
10 int total_sig;
11 {
12     extern void signal_handler();
13
14     signal( SIGUSR1, signal_handler );
15
16     while( sig_ct_child < total_sig )
```



```

17     {
18         pause();
19         printf( "Child got ordinal signal %#d\n", sig_ct_child );
20     }
21 }
22
23
24
25 /* Count how many signals have been received */
26
27 void signal_handler(sig_num)
28 int sig_num;
29 {
30     signal( SIGUSR1, signal_handler );
31     ++sig_ct_child;
32 }

```

Fortran Files

main.f

```

1  C      This program spawns a child process and sends
2  C      signals from the parent to the child.
3  C
4      program main
5      common /msg_comm/ total_sig
6      common /usr1_comm/ sigusr1
7      integer total_sig
8      integer sigusr1
9      integer pid
10     integer ftfork
11     integer ftsigusr1
12     character *8 tracefile
13     external parent_routine
14     external child_routine
15
16     sigusr1 = ftsigusr1()
17     tracefile = "msg_file"
18     write( 6, FMT='(A)' )
19     X    "How many signals should the parent send the child?"
20
21     read( 5,* ) total_sig
22     pid = ftfork()
23
24     if( pid .eq. 0 ) then
25     C      It's the child
26         call child_routine()
27     else
28     C      It's the parent
29         call parent_routine( pid )
30     end if
31
32     call exit
33     end

```

parent.f

```
1 C      Every time the parent sends the child a signal,
2 C      the parent writes a message.
3
4      subroutine parent_routine( child_pid )
5      common      /msg_comm/ total_sig
6      common      /usr1_comm/ sigusr1
7      integer     child_pid
8      integer     total_sig
9      integer     sigusr1
10     integer     isec
11     integer     ireturn
12     integer     sig_ct
13     integer     kill
14     data        isec/2/
15
16     do 10 sig_ct = 1, total_sig
17         write( 6, FMT='(I3, A, I2, X, A)' ) sig_ct,
18 X      ". Parent sleeping for", isec, "seconds"
19         call sleep( isec )
20         ireturn = kill( child_pid, sigusr1 )
21 10    continue
22
23     return
24     end
```

child.f

```
1 C      Every time the child receives a signal from
2 C      the parent, the child writes a message
3
4      subroutine child_routine( )
5      common      /msg_comm/ total_sig
6      common      /sig_comm/ sig_ct_child
7      common      /usr1_comm/ sigusr1
8      integer     total_sig
9      integer     sig_ct_child
10     integer     sigusr1
11     integer     ireturn
12     integer     ftpause
13     integer     ftsignal
14     external     signal_handler
15     integer     signal_handler
16
17     ireturn = ftsignal( sigusr1, signal_handler )
18
19     do while ( sig_ct_child .lt. total_sig )
20         ireturn = ftpause()
21         ireturn = ftsignal( sigusr1, signal_handler )
22         write( 6, FMT='(A, I3)' ) "Child got ordinal signal #",
23 X      sig_ct_child
24     end do
25
26     return
27     end
```

```

28
29
30
31 C      Count how many signals have been received
32
33      integer function signal_handler()
34      common /sig_comm/ sig_ct_child
35      integer sig_ct_child
36      data    sig_ct_child /0/
37
38      sig_ct_child = sig_ct_child + 1
39      return
40      end

```

ftint.c

```

1  /*
2   * C routines to provide simple Fortran interfaces
3   * to some system services, so that the tutorial
4   * works the same way on different systems.
5   */
6
7  #include <signal.h>
8  #include <unistd.h>
9
10 int
11 ftfork_()
12 {
13     int status = fork();
14     return status;
15 }
16
17 int
18 ftpause_()
19 {
20     int status = pause();
21     return status;
22 }
23
24 typedef void (*sig_handler)(int);
25
26 int
27 ftsignal_(int * signum, sig_handler handler)
28 {
29     int status = (int)signal(*signum, handler);
30     return status;
31 }
32
33 int
34 ftsigusr1_()
35 {
36     return SIGUSR1;
37 }

```


Glossary

This glossary defines terms used in NightView. Terms in *italics* are defined here.

accelerator

A special key used to select a menu item quickly in the graphical user interface. See also *mnemonic*. See “List of Shortcuts” on page 8-31.

anchored match

The entire string must match the regular expression. Put another way, a $\wedge$ is implied at the beginning of the regular expression, and a $\$$ is implied at the end of the regular expression. See “Regular Expressions” on page 6-21.

application

A group of related processes. The processes may be running the same program or different programs.

attaching

Attaching to a process means that the debugger will have control over it. This is how you debug processes that already exist. See “attach” on page 6-43.

breakpoint

A breakpoint is a place in your program where execution will stop. You can set a breakpoint with the **breakpoint** command. See “Breakpoints” on page 3-14. Breakpoints may be conditional, see *conditional breakpoint*. Breakpoints may have debugger commands associated with them, see *breakpoint commands*.

breakpoint commands

A set of debugger commands to be executed when a breakpoint is hit. See *breakpoint*.

checkpoint

A checkpoint saves information about the eventpoints, signal disposition, and other information, for a program. This information is used when a program is *restarted*. See “Restarting a Program” on page 3-22.

child process

When a process *forks*, a new process is created that looks just like the old process. The new process is called a child process and the old process is called the parent process. A process may have many child processes, but only one parent process. You can control

whether the child process is debugged with the **set-children** command. See “set-children” on page 6-53.

command history

NightView keeps a history of all the commands you enter. You can retrieve commands, edit them, and re-enter them. See “Command History” on page 3-47.

command-line interface

A command-line interface deals with only one line at a time. This kind of interface can be used from a terminal or from other programs that expect simple behavior, such as a shell running in emacs. Contrast this with a *full-screen interface* and a *graphical user interface*. See Chapter 6 [Command-Line Interface] on page 6-1.

command stream

A command stream is a set of commands executed sequentially by NightView. The commands attached to a breakpoint form a command stream, as do the commands you type as input to NightView. Execution of commands in one command stream may be interleaved with the execution of commands from another command stream. See “Command Streams” on page 3-44.

conditional breakpoint

A breakpoint may have a language expression associated with it. The breakpoint is “hit” only if the expression evaluates to TRUE when the breakpoint is encountered. See *breakpoint*.

context

Context refers to the information the debugger uses to determine how to evaluate an expression. The main components of the context are the *program counter*, which determines the *scope*, and the *stack*. Context determines the language (i.e., C, C++ or Fortran) as well as the type and location of variables in the program. NightView allows you to specify the context to be used in interpreting an expression. See “Context” on page 3-32.

context menu

In the graphical user interface, a context menu pops up when you right-click in a window. “Context” here refers to the position of the mouse in the window and has nothing to do with the debugging concept of *context*. See “Context Menu” on page 8-3.

context panel

In the graphical user interface, a context panel is a convenient way to navigate in your process's *stack*, and also to navigate between processes and between shells. See “Context Panel” on page 8-78.

convenience variables

A convenience variable is a variable maintained by the debugger to hold the value of an expression. The type of a convenience variable is determined by the type of the expression assigned to it. See “Convenience Variables” on page 3-46.

core file

A core file is a snapshot of a process’s memory created by the operating system when the process is aborted. You can examine this process state using NightView. See “Core Files” on page 3-5.

crossing count

A crossing count is the number of times program execution has crossed an eventpoint since the program has started execution. This count is updated even if the ignore count or condition was not satisfied. The crossing count is not updated if the eventpoint is disabled.

CUDA

CUDA refers to CUDA contexts and threads running on a separate CUDA device. See “CUDA Debugging” on page 3-51.

current frame

The current frame is one of the frames on the stack of a stopped process. It is often the same as the *currently executing frame*, but other frames can be selected using the **up**, **down**, and **frame** commands. The current frame is used to determine the *context* for evaluating an expression. See “Current Frame” on page 3-33.

current process

In the graphical user interface, one process is the current process. Button clicks and commands apply to that process. If there is more than one process, the *context panel* indicates the current process with green text. See “Current Process” on page 8-4.

currently executing frame

The currently executing frame is the stack frame associated with the most recently called routine in a stopped process. Contrast this with *current frame*.

data item

Each data item shows one piece of data from your process in the *data panel*. A data item can show an expression, local variables, registers, the stack, or the process threads. See “Data Items” on page 8-79.

data panel

In the graphical user interface, a data panel allows you to view data items in your process. See “Data Panel” on page 8-79.

data panel layout

The organization of *data items* in the *data panel* in your debug session. This information can be saved and reused in future debug sessions. See “Data Menu” on page 8-17.

debugger

A debugger is a tool to help you debug programs. A debugger lets you control the execution of your program and look at your program's memory.

debug session

A debug session is one invocation of NightView; it lasts until you exit from the debugger. See Chapter 5 [Invoking NightView] on page 5-1. See “Quitting NightView” on page 6-26.

default font

The default font is specified by the Motif `fontList` resource and applies only to the graphical user interface. See “Preferences Dialog Box” on page 8-46.

detaching

Detaching from a process means that the debugger no longer has control over that process and any future children that are created by that process. The debugger still has control over previously created children. See “detach” on page 6-44.

dialogue

NightView provides dialogues as a means of starting processes, via a shell, and communicating with those processes. See “Dialogues” on page 3-6. See also *remote dialogue*.

disassembly

A symbolic representation of the raw machine language that makes up your program. To disassemble part of your program, use the **x** command with the **i** format. See “x” on page 6-96. In the graphical user interface, you can view disassembly in the source panels by using the **Source** menu. See “Source Menu” on page 8-12.

display item

A display item is an expression or memory location whose value or contents are to be printed out whenever the associated process stops. NightView assigns a unique number to each display item in each process. See “display” on page 6-106 and “info display” on page 6-176.

DWARF

DWARF is the standard format for symbolic debugging information used with ELF files. See *ELF*.

ELF

Executable and Linking Format. This is a standard for the format and contents of an executable file. It also determines the form and content of information about your program available to the debugger.

event-map file

An event-map file lets you associate or map symbolic *trace-event tags* and numeric *trace-event IDs*. This file appears on the **ntrace** invocation line when performing *NightTrace* tracing. See *trace*.

eventpoint

An eventpoint is a generic name given to the various kinds of modifications NightView can insert at a particular location of a process. The different kinds of eventpoints are: *breakpoint*, *monitorpoint*, *heappoint*, *tracepoint*, *patchpoint*, and *watchpoint*. See “Eventpoints” on page 3-11.

eventpoint modifier

An eventpoint modifier modifies the meaning of an eventpoint command. The only eventpoint modifiers are */delete* and */disabled*. The modifier */delete* is valid only for breakpoints and watchpoints, and tells NightView to delete the eventpoint after the next time it is hit. The modifier */disabled* tells NightView to create the eventpoint, but leave it disabled initially. See “Eventpoint Modifiers” on page 6-113.

eventpoint panel

In the graphical user interface, an eventpoint panel is a convenient way to view and modify information about *eventpoints*. See “Eventpoint Panel” on page 8-75.

family

A group of related processes. See “family” on page 6-54.

fence

A region of memory preceding or following a block, when heap debugging is turned on. It can be used to detect block overrun bugs. See “Fences” on page 3-39.

find bar

In the graphical user interface, you can use the *context menu* to bring up a find bar in a panel to search for text in that panel. See “Find Bar” on page 8-66.

focus

See *keyboard focus*.

fork

Create a new process. The debugger informs you when your process forks. See *child process*.

frame

See *stack frame*.

full-screen interface

A full-screen interface uses the capabilities of a terminal to control the display of information on the entire screen, rather than just writing to the terminal one line at a time. Contrast this with a *command-line interface* and a *graphical user interface*. See Chapter 7 [Simple Full-Screen Interface] on page 7-1.

graphical user interface

A graphical user interface may be used on a graphics display. This kind of display allows much more flexibility and functionality than a text display. Contrast this with a *command-line interface* and a *full-screen interface*. See Chapter 8 [Graphical User Interface] on page 8-1.

GUI

A *graphical user interface*.

heap

Memory allocated and deallocated dynamically via calls to `malloc` and `free`. NightView can help you debug problems with heap usage. See “Debugging the Heap” on page 3-37.

heap check

A complete check of all allocated and retained free blocks in the heap, when heap debugging is turned on. It can detect numerous heap-related bugs. See “Heap Check” on page 3-41.

heap debugger

A module that NightView loads into your process to monitor memory allocation and deallocation. See “Debugging the Heap” on page 3-37.

heappoint

An *eventpoint* that checks the *heap* or changes the heap debugging parameters. See “Heappoints” on page 3-15.

Help Window

In the graphical user interface, the Help Window displays NightView's online help information. You can choose to look at any part of the *NightView User's Guide*. See also *online help system*. See "Help Window" on page 8-124.

hit a breakpoint

A breakpoint is hit when execution reaches the breakpoint location and the ignore count and conditions, if any, are satisfied. Thus, hitting a breakpoint stops the process. See "Breakpoints" on page 3-14.

hit an eventpoint

An *inserted eventpoint* is hit when execution reaches the eventpoint location and the ignore count and conditions, if any, are satisfied. A watchpoint is hit when the specified addresses are referenced, and the ignore count and conditions are satisfied. Thus, hitting an eventpoint causes that eventpoint to perform its specified action; e.g., a breakpoint stops the process, a monitorpoint evaluates its expressions and saves their values, a tracepoint logs a trace event, and so on. See *eventpoint*, *breakpoint*, *monitorpoint*, *tracepoint*, *heappoint*, and *watchpoint*.

ignore count

An ignore count causes NightView to skip an eventpoint the next *count* times that execution reaches the eventpoint. You use the **ignore** command to attach an ignore count to an eventpoint. See "ignore" on page 6-125.

initialization file

An initialization file is a file containing NightView commands that are executed before NightView reads commands from standard input. NightView has a default initialization file, and you can specify others on the NightView invocation line. See "Initialization Files" on page 3-47.

inline subprogram

A subprogram that is expanded directly into the calling program. See "Inline Subprograms" on page 3-34.

inline interest level

The level that determines if any inline subprograms are interesting. You may set an *interest level* for individual inline subprograms to override this level. See "Inline Subprograms" on page 3-34. You can change or query this level with the **interest** command. See "interest" on page 6-71.

inserted eventpoint

An *eventpoint* that is associated with a location in your program. Inserted eventpoints are implemented by inserting code into your process. See "Eventpoints" on page 3-11.

interest level

Each subprogram has an associated interest level. NightView compares the interest level to the *interest level threshold* to determine if the subprogram is interesting. NightView generally avoids showing you uninteresting subprograms. See “Interesting Subprograms” on page 3-35. You can change or query the interest level with the **interest** command. See “interest” on page 6-71.

interest level threshold

Each process has an interest level threshold. If the *interest level* of a subprogram is less than the interest level threshold, the subprogram is considered to be uninteresting. See “Interesting Subprograms” on page 3-35.

keyboard focus

The keyboard focus determines which field receives keyboard input in the graphical user interface.

lane

A lane is a physical component of a CUDA device which executes a CUDA thread. See “CUDA Debugging” on page 3-51.

leaked block

A block which was allocated from the heap, which was never freed, and which the program no longer references. See “Leak Detection” on page 3-42.

locals panel

In the graphical user interface, a locals panel displays the local variables in the *current frame* of the *current process*. See “Locals Panel” on page 8-78.

macro

A macro is a named set of text, possibly with arguments, that can be substituted in a NightView command by referencing the name. This is a means of extending the facilities provided by NightView. See “Defining and Using Macros” on page 6-193.

message panel

In the graphical user interface, a message panel displays process status, command output, error messages and process and shell output. See “Message Panel” on page 8-75.

mnemonic

A mnemonic is a way of selecting a menu or a menu item quickly in the graphical user interface. See also *accelerator*. See “List of Shortcuts” on page 8-31.

monitorpoint

A monitorpoint is a location in a debugged process where one or more expressions are evaluated and the values saved. The saved values are displayed periodically by NightView. Monitorpoints thus provide a means of viewing program data while the program is executing. See “Monitorpoints” on page 3-14 and “monitorpoint” on page 6-121.

monitor bar

In the graphical user interface, in the *monitor panel*, you can use the *context menu* to bring up a monitor bar to control the display of *monitorpoint* values. See “Monitor Bar” on page 8-79.

monitor panel

In the graphical user interface, a monitor panel displays the values of *monitorpoints*. See “Monitor Panel” on page 8-78.

NightTrace

An interactive debugging and performance analysis tool that lets you examine trace events logged by user applications and the kernel. See *trace*. See the *NightTrace Manual* for details.

NightView

A pretty good debugger.

online help system

All of the *NightView User's Guide* is available to you, online, through NightView's online help system. In the graphical user interface, help information is displayed in the Help Window. See also *Help Window*. See “help” on page 6-159. See “GUI Online Help” on page 8-2.

overloading

Overloading means that more than one entity with the same name is visible at some point in the program. See “Overloading” on page 3-31.

patch

A patch is an expression (or a branch) inserted into a debugged process to alter its behavior (usually to fix a bug). See *patchpoint*. See “Patchpoints” on page 3-15.

patch area

NightView creates regions, known as patch areas, in your process. This is where NightView puts code and data that is inserted into your process. See Appendix C [Implementation Overview] on page C-1. See “set-patch-area-size” on page 6-73.

patchpoint

A patchpoint is a location in a debugged process where a patch is inserted. See *patch*. See “Patchpoints” on page 3-15.

pattern

A pattern is used in the **debug** and **nodebug** commands to control which programs will be debugged in a particular *dialogue*. Patterns are similar to shell wildcard patterns. See “debug” on page 6-28.

PID

A process identifier. This is an integer from 1 to 30000 which uniquely identifies a process on a particular system. In some situations, NightView may create false PIDs, outside the normal range, to identify false processes, e.g., core files.

procedure

See *routine*.

process

The execution of a program. Many processes may be executing the same program. See “Programs and Processes” on page 3-3.

process state

A process state describes whether the process is actively executing and what you can do with the process using NightView. The two most common process states are *running* and *stopped*. See “Process States” on page 3-25.

program

A file containing instructions and data. A program is usually created with the **ld(1)** program. An executing program represents a *process*. See “Programs and Processes” on page 3-3.

program counter

The program counter is a register that locates the instruction that is to be executed next. See “Program Counter” on page 3-32.

qualifier

A qualifier specifies the set of processes or dialogues that a command affects. See “Qualifiers” on page 3-5.

registers

Registers are special storage locations in the CPU for holding frequently accessed data. In NightView, you can access most of these registers using specially-named convenience variables. See “Predefined Convenience Variables” on page 6-7.

remote dialogue

A remote dialogue is a *dialogue* started on a system other than the one on which Night-View was invoked. See “Remote Dialogues” on page 3-7.

restarted

When a program is run again in the same debug session, it is considered to be *restarted*. Information from the most recent *checkpoint* is applied to the process. See “Restarting a Program” on page 3-22.

retained free block

A block which was freed by the user, but which the heap debugger has not yet made available for reuse. It can be used to detect dangling pointer bugs. See “Retained Free Blocks” on page 3-41.

routine

Routine is a generic term denoting a function or subroutine in a program. Different languages use different terms for this concept; other similar terms are subprogram and procedure.

scope

A scope is a section of your program where a particular set of variables can be referenced. Scope forms a part of the *context*. See “Scope” on page 3-32.

shell

The shell is the program the system normally executes when you log in. There are several varieties of shell: Bourne shell, C shell, and Korn shell are some examples. In NightView, each *dialogue* you create executes an instance of your login shell.

shell panel

In the graphical user interface, a shell panel provides a way to interact with your *dialogue* shell. See “Shell Panel” on page 8-74.

signal

A signal is a notification of some event to your process. This event may be external to your process, or it may be the result of an erroneous action by the process itself. Night-View allows you to control how signals are delivered to your process. See “Signals” on page 3-20.

SM

A SM, or symmetric multiprocessor, is a physical component of a CUDA device which executes a number of CUDA warps. See “CUDA Debugging” on page 3-51.

smart printer

A smart printer is a definition which recognizes types by their names and replaces descriptions of their objects with a user-defined form. See “Smart Printing” on page 3-46.

source panel

In the graphical user interface, a source panel displays program source or *disassembly*. You can also interact with the process through the source panel. See “Source Panel” on page 8-66.

source panel target line

In the graphical user interface, When you click on a line in a source panel, that line becomes the source panel target line. Some buttons and keystrokes use the target line to identify the source file and line number to operate on. See “Source Panel Target Line” on page 8-67.

stack

An area of memory used to hold local variables and return information for each active routine. The stack consists of a sequence of *stack frames*. Calling a routine pushes a new frame onto the stack; returning from the routine removes that frame from the stack. See “Stack” on page 3-33.

stack frame

A stack frame is a contiguous set of locations in the process' stack that corresponds to the execution of an active routine. The stack frame holds the local automatic variables of the routine, and it also holds information needed to return to the calling routine. See “Stack” on page 3-33.

stale data indicator

A stale data indicator is a character or icon displayed with a monitored value to indicate the validity and reliability of that value. See *monitorpoint*.

symbol file

An executable file containing symbolic debug information. Normally, the symbol file is the same as the program's executable file, but it may be different if, for example, you are debugging a stripped program. See “symbol-file” on page 6-45.

syscallpoint

An *eventpoint* that informs you of system call entry and exit. See “Heappoints” on page 3-15.

thread

Each instance of execution of a program contains one or more threads of execution. Some programs have a single thread. See “Multithreaded Programs” on page 3-49.

thread processes

Threaded programs are implemented with multiple processes that share resources, including memory. This manual refers to these processes as *thread processes*.

trace

The collection of data produced by executing *tracepoints* in a process is called a trace. See *NightTrace*.

trace-event ID

An integer that identifies a *NightTrace* trace event. User trace event IDs are in the range 0 through 4095, inclusive. See *event-map file* and *trace-event tag*.

trace-event tag

A symbolic name that identifies a *NightTrace* trace event. It is mapped to a numeric *trace-event ID* in an *event-map file*.

tracepoint

A tracepoint is a call to one of the **ntrace(3X)** library routines for recording the time when execution reached the tracepoint. You can insert a tracepoint in your source, or you can use NightView to insert them after starting your process. See “Tracepoints” on page 3-15.

value history

The value history is a list of values you have printed in your NightView session. You can view this list, and you can reference the values in other expressions. See “Value History” on page 3-47.

warp

A warp is a physical component of a CUDA device which executes a number of CUDA threads. See “CUDA Debugging” on page 3-51.

watchpoint

A watchpoint stops the process when the process reads or writes a variable in memory. See “Watchpoints” on page 3-16.

Index

Symbols

__eventpoint_data 6-16
 ! 1-3, 3-7, 6-36, 6-41, 6-71, 6-197
 !exit 6-33
 . 6-20
 -. (input terminator) 6-36, 6-71
 .NightViewrc file 3-47, 6-33, 6-34, 6-36
 .profile file 3-9
 'body 6-17, 6-18
 'pos 8-104
 'specification 6-17, 6-18
 (local) prompt 1-3
 @ (macro invocation) 6-196
 @ (source line decoration) 6-94, 6-193
 * (source line decoration) 6-93
 /disabled eventpoint modifier 6-113, Glossary-5
 /f eventpoint modifier 6-113
 /proc 3-4, 3-56, C-1
 /usr/lib/NightView-release/ReadyToDebug
 1-3, 3-10, 4-5
 /usr/ucb/rsh 3-4
 & 3-29, 4-7, 6-18, 6-41, 6-190
 # (comment) 6-2
 + (list argument) 6-86, 6-88
 <= 8-86
 = (list argument) 6-87
 = (source line decoration) 6-93
 = key 8-73
 > prompt 6-2, 6-122, 6-126, 6-194
 > (source line decoration) 6-93
 | (show argument) 6-39
 \$ 6-5, 6-73, 6-98
 \$ prompt 1-3, 4-5
 \$ _ 6-7
 \$ _ 6-7
 \$\$ 6-5
 \$cfa 6-8, 6-155, 6-175
 \$cpc 3-32, 3-33, 6-7, 6-75, 6-93, 6-155, 6-175
 \$fp 6-8, 6-176
 \$is 6-8
 \$pc 3-26, 3-32, 3-33, 6-7, 6-103, 6-156
 \$sp 6-8
 \$was 3-16

A

Abbreviations

b (breakpoint) 6-114, B-8
bh (branch-history) 6-110
bt (backtrace) 6-95, B-7
c (continue) 6-138, B-10
 command 6-2, 6-197
d (delete) 6-128, B-9
exception (info exception) 6-185
f (frame) 6-155, B-11
fo (forward-search) 6-90, B-7
hd (heapdebug) 6-59, B-4
hold (mcontrol hold) 6-124, B-9
i b (info breakpoint) 6-168, B-12
l (list) 6-87, B-7
md (memory-display) 6-105, B-7
n (next) 6-142, B-10
ni (nexti) 6-145, B-10
p (print) 6-96, B-7
ptype (info declaration) 6-192, B-13
q (quit) 6-26, B-2
release (mcontrol release) 6-124, B-9
representation (info representation)
 6-191, B-13
s (step) 6-141, B-10
si (stepi) 6-144, B-10
whatis (info whatis) 6-191, B-13
xl (translate-object-file) 6-31, B-2

Abnormal termination 6-46

Abort 3-43

Accelerator Glossary-1

Ctrl+B 8-15

Ctrl+P 8-16

Ctrl+Q 8-6

Ctrl+S 8-5

Accessing files 3-2, 3-9

Actual argument

macro 6-194, 6-196, 6-197

Ada i-iii, 3-28, 3-32, 3-48, 4-4, 4-9, 4-10, 4-12, 4-13,
 4-15, 4-16, 4-17, 4-19, 4-20, 4-21, 4-22, 6-67,
 6-117, 6-132, 6-153, 8-13, 8-34, 8-68,
 Glossary-2

Ada elaboration 3-49

Ada exception 3-29, 6-151, 6-185, Glossary-5

Ada exception handling 3-49, 6-150, 6-153, 6-185

Ada expressions 3-28

Ada packages 3-48

Ada task 3-49, 6-147, 6-157, 6-186, 8-1, 8-86,
 Glossary-1, Glossary-12

Ada unit 6-17, 6-18, 6-87, 6-151, 6-185, 6-190

Address

printing 6-190

Addresses limits 6-177
 printing 6-69
addr-expression 6-102, 6-106
Aggregate item 6-98, 6-177
AMD64 registers (iHawk Series 870) 6-11
Anchored match 6-22, 6-90, 6-177, 6-183, 6-189, 6-190,
 6-191, 6-192, 6-198, 8-13, 8-33, 8-34, 8-68,
 Glossary-1
Application 3-3, Glossary-1
apply on dialog 6-34, 6-35
apply on program 6-52
--arch option 5-1
Argument
 actual 6-194
 command 6-2
 macro 6-194, 6-196
 printing 6-178, 6-189
Array 6-98, 6-177, 8-81
 printing 6-69
Array slices 3-29, 3-30
Assembly listing 8-14, 8-15, 8-70
Assignment 3-27, 3-28, 6-117
attach 3-4, 3-54, 5-1, 6-43, 6-55
Attach button 8-33
Attach Dialog Box 3-4, 8-33
-**attach** option 5-1
Attaching 3-4, 3-54, 5-1, 6-43, 8-11, 8-31, 8-33,
 Glossary-1
Attaching commands to a dialogue 6-34
Attaching commands to a program 6-50, 6-52

B

b (breakpoint) 1-4, 2-5, 6-114, B-8
B (source line decoration) 6-92
b key 8-72
Background process 6-161
backtrace 1-5, 2-6, 3-33, 6-6, 6-95, 6-155, 6-156,
 6-159, 6-164
bh (branch-history) 6-110
Binary Viewer Panel 6-105, 8-112
Blank line 6-24, 6-88
Body
 macro 3-45, 6-194, 6-199
Branch instruction 6-116, 6-132
branch-history 6-110
Breakpoint 3-11, 3-14, 3-24, 3-55, 6-92, 6-94, 6-111,
 6-118, 6-140, 6-149, 8-27, 8-30, 8-72,
 Glossary-1, Glossary-5
 changing 8-35
 clearing 6-125, 6-126
 commands on 3-23, 3-44, 3-46, 4-23, 6-126, 6-168,

 6-169, 8-38, Glossary-1
 condition on 3-23, 4-21, 6-115, 6-127, 6-168,
 6-169, 8-38, Glossary-2
 deleting 4-20, 6-44, 6-128, 8-40
 disabling 4-26, 6-128, 8-37
 displaying 4-26, 6-167, 6-168
 enabling 6-129, 8-37
 hitting 3-25, 6-42, 6-107, 6-137, 6-139, 6-169, 8-37,
 Glossary-7
 ignoring 3-23, 4-22, 6-115, 6-130, 6-168, 6-169,
 8-38, Glossary-7
 named 3-23, 6-114, 6-115, 6-131, 8-38
 protected 6-115
 saving 3-23
 setting 1-4, 2-5, 4-9, 4-21, 4-22, 6-109, 6-114,
 6-131, 8-27, 8-35, 8-72
 state 6-169, 8-37
 temporary 6-131, 8-37
breakpoint 1-4, 2-5, 3-4, 3-5, 6-114, 6-132, 6-168,
 8-24, 8-27, 8-72
Breakpoint button 8-27, 8-72
Breakpoint crossing count Glossary-3
Breakpoint Dialog Box 4-21, 4-22, 4-23, 8-15, 8-27,
 8-35
Browse selection policy 8-34
bt 1-5, 2-6
bt (backtrace) 6-95, B-7
Buffered output 3-45
Building a program 1-2, 2-2
Button
 Attach 8-33
 Breakpoint 8-27, 8-72
 Cancel 8-34, 8-35, 8-40, 8-100, 8-103, 8-104,
 8-105, 8-106
 Clear 8-28
 Data Display 8-29
 Delete 8-40
 dimmed 4-20, 8-34, 8-40
 disabled 4-20, 8-34, 8-40
 Dismiss 8-3
 Down 4-19
 Down 8-25
 Filter 8-33
 Finish 8-24
 Help 8-3, 8-34, 8-35, 8-40, 8-100, 8-103
 Hold 8-78
 Interrupt 8-25
 Kill 6-45, 8-25
 Next 8-24
 Nexti 8-24
 OK 2-6, 8-34, 8-35, 8-40, 8-100, 8-101, 8-102,
 8-103, 8-104, 8-105, 8-106
 Print 2-6, 8-29, 8-72
 radio 8-39, 8-101

- Release 8-78
 - Rerun 8-25
 - Resume 2-5, 2-6, 8-24
 - Run to Here 8-24, 8-72
 - Search 8-34
 - Step 8-24
 - Stepi 8-24
 - Stop 8-24
 - toggle 8-39
 - Up 4-17
 - Up 8-25
- C**
- C 4-4
 - c (continue)** 1-5, 6-138, B-10
 - C language i-iii, 3-29, 3-32, 3-48, 3-56, 6-67, 6-70, 6-97, 6-99, 6-104, 6-108, 6-117, 6-132, 6-190, 6-195, Glossary-2
 - C string 3-56, 6-98
 - C thread 6-186, 8-86
 - C++ i-iii, 3-30, 3-32, 3-48, 3-56, 6-67, 6-97, 6-117, 6-132, 6-195, 8-13, 8-34, 8-68, Glossary-2
 - Calling macros 6-196
 - Cancel button 8-34, 8-35, 8-40, 8-100, 8-103, 8-104, 8-105, 8-106
 - Canonical frame address 6-8, 6-155, 6-175
 - Cautions 3-43
 - cc option
 - g 1-2, 2-2
 - cd** 6-85
 - Changing a breakpoint 8-35
 - Changing a heappoint 8-35
 - Changing a monitorpoint 8-35
 - Changing a patchpoint 8-35
 - Changing a tracepoint 8-35
 - Changing a watchpoint 8-35
 - Changing an eventpoint 8-35
 - Changing eventpoints 8-35, 8-74
 - Character string 6-98, 6-177
 - printing 6-69
 - Checkpoint 3-23, 6-50, 6-53, Glossary-1, Glossary-11
 - checkpoint** 3-23, 6-53
 - Child process 3-3, 3-4, 4-2, 4-7, 4-13, 6-30, 6-44, 6-55, Glossary-2
 - clear** 6-111, 6-125, 6-128, 8-28
 - Clear button 8-28
 - Clearing breakpoints 6-125
 - Clearing eventpoints 6-125, 8-28
 - Clearing heappoints 6-125
 - Clearing monitorpoints 6-125
 - Clearing patchpoints 6-125
 - Clearing terminal 6-160
 - Clearing tracepoints 6-125
 - Clicking
 - double 8-34, 8-35, 8-103, 8-104
 - Clicking on objects 4-3, 4-5
 - Command abbreviations 6-2, 6-197
 - b (breakpoint)** 6-114, B-8
 - bh (branch-history)** 6-110
 - bt (backtrace)** 6-95, B-7
 - c (continue)** 6-138, B-10
 - d (delete)** 6-128, B-9
 - exception (info exception)** 6-185
 - f (frame)** 6-155, B-11
 - fo (forward-search)** 6-90, B-7
 - hd (heapdebug)** 6-59, B-4
 - hold (mcontrol hold)** 6-124, B-9
 - i b (info breakpoint)** 6-168, B-12
 - l (list)** 6-87, B-7
 - md (memory-display)** 6-105, B-7
 - n (next)** 6-142, B-10
 - ni (nexti)** 6-145, B-10
 - p (print)** 6-96, B-7
 - ptype (info declaration)** 6-192, B-13
 - q (quit)** 6-26, B-2
 - release (mcontrol release)** 6-124, B-9
 - representation (info representation)** 6-191, B-13
 - s (step)** 6-141, B-10
 - si (stepi)** 6-144, B-10
 - whatis (info whatis)** 6-191, B-13
 - x1 (translate-object-file)** 6-31, B-2
 - Command arguments 6-2
 - Command case 6-2
 - Command execution
 - delaying 6-164
 - Command file 6-161
 - Command history 3-47
 - Command input 6-161, 6-164
 - Command qualifier 3-5, 6-1, 6-19, 6-68, Glossary-10
 - Command repetition 6-2, 6-24, 6-88
 - Command replacement 6-197
 - Command stream 3-44, 3-45, 3-47, 6-139, 6-140, 6-149, 6-161, 6-164, Glossary-2
 - event-driven 3-44
 - Command summary 8-21, A-1
 - Command syntax 6-1
 - Command toolbar 8-23
 - Command-line interface i-iii, 1-1, 3-36, 4-2, 6-2, 6-122, 6-126, 6-152, 6-194, 7-1, 8-1, 8-23, Glossary-2
 - Command-line user interface 6-160
 - Commands
 - ! 6-36
 - apply on dialogue** 6-35
 - apply on program** 6-52

attach 5-1, 6-43
backtrace 6-95
branch-history 6-110
breakpoint 6-114
cd 6-85
checkpoint 6-53
clear 6-125
commands 6-126
condition 6-127
continue 6-138
core-file 6-46, 6-47
data-display 3-24, 6-104
debug 6-29
declare-thread-tag 6-162
define 6-193
delay 6-164
delete 6-128
detach 6-44
directory 6-88
disable 6-128
display 6-106
down 6-156
echo 6-104
edit 6-89
enable 6-129
exec-file 6-48
family 6-54
finish 6-146
forward-search 6-90
frame 6-154
handle 6-150
heapcheck 3-42, 6-62, 6-187
heapdebug 3-15, 3-37, 3-39, 3-42, 6-59, 6-123, 6-188, 8-40, 8-41
heappoint 3-15, 3-42, 6-62, 6-123, 6-172
help 6-159
ignore 6-130
info address 6-190
info args 6-189
info breakpoint 6-168
info convenience 6-176
info declaration 6-192
info dialogue 6-182
info directories 6-176
info display 6-176
info eventpoint 6-167
info exception 6-185
info family 6-183
info files 6-192
info frame 6-175
info functions 6-190
info heappoint 6-172
info history 6-176
info limits 6-177
info line 6-192
info locals 6-189
info log 6-167
info macros 6-198
info memory 6-179
info monitorpoint 6-171
info name 6-183
info on dialogue 6-184
info on program 6-184
info on restart 6-184
info patchpoint 6-170
info process 6-178
info registers 6-177
info representation 6-191
info signal 6-178
info sources 6-190
info tracepoint 6-169
info types 6-191
info variables 6-189
info watchpoint 6-173, 6-174
info whatis 6-191
interest 6-74
jump 6-148
kill 6-45
list 6-86
load 6-109
login 6-27
logout 6-33
mcontrol 6-124
memory-display 6-105
modify-eventpoint 6-131
monitorpoint 6-121
mreserve 6-57
name 6-113
next 6-142
nexti 6-145
nodebug 6-30
notify 6-43
on dialogue 6-33
on program 6-50
on restart 6-52
output 6-103
patchpoint 6-116
print 6-96, 6-108
pwd 6-85
quit 6-26
redisplay 6-108
refresh 6-160
rerun 6-41
resume 6-139
reverse-search 6-90
run 6-40
select-context 6-157, 6-163
set 6-99

- set-auto-frame** 6-77
- set-children** 6-55
- set-cuda** 6-83
- set-cuda-memcheck** 6-83
- set-debug-file-directory** 6-31
- set-disassembly** 6-81
- set-download** 6-79, 6-80
- set-editor** 6-78
- set-exit** 6-56
- set-format-hex** 6-71
- set-futurepoints** 6-82
- set-history** 6-68
- set-language** 6-67
- set-limits** 6-69
- set-local** 6-72
- set-log** 6-66
- set-notify** 6-42
- set-overload** 6-77
- set-patch-area-size** 6-73
- set-preallocate** 6-78
- set-prompt** 6-69
- set-qualifier** 6-68
- set-restart** 6-72
- set-safety** 6-71
- set-search** 6-77
- set-show** 6-37
- set-tag** 6-163
- set-terminator** 6-71
- set-trace** 6-119
- shell** 6-161
- show** 6-38
- signal** 6-149
- smart-print** 6-200
- source** 6-161
- step** 6-141
- stepi** 6-144
- stop** 6-147
- symbol-file** 6-45
- syscallpoint** 6-135
- tbreak** 6-131
- tpatch** 6-132
- tracepoint** 6-119
- translate-object-file** 6-31
- undisplay** 6-107
- up** 6-156
- wait** 6-57
- watchpoint** 6-133
- x** 6-99
- commands** 3-14, 6-111, 6-126
 - Commands attached to a dialogue 6-34
 - Commands attached to a program 6-50, 6-52
 - Commands on breakpoint 3-23, 4-23, 6-114, 6-115, 6-126, 8-38, Glossary-1
 - Commands on eventpoint 8-38, 8-39
 - Commands on monitorpoint 3-23, 6-121, 6-126, 8-38, 8-39
 - Commands on watchpoint 3-23
 - Comments 6-2, 6-194
 - Compilation
 - by debugger 3-11, 3-46
 - Compiling 1-2, 2-2, 3-15, 3-48
 - Condition
 - breakpoint 3-23, 4-21, 6-168, 6-169, 8-38, Glossary-2
 - eventpoint 3-11, 3-23, 6-127, 6-130, 6-132, 6-168, 8-38
 - heappoint 3-23, 6-168, 6-173, 8-38
 - monitorpoint 3-23, 6-168, 6-172, 8-38
 - patchpoint 3-23, 6-168, 6-171, 8-38
 - syscallpoint 6-174
 - tracepoint 3-23, 6-168, 6-170, 8-38
 - watchpoint 3-16, 3-23, 6-168, 6-174, 8-38
 - condition** 6-111, 6-118, 6-127
 - Condition removal 6-127, 8-38
 - conditional-expression* 3-16, 3-29, 6-114, 6-120, 6-127, 6-131, 6-133
 - config** option 5-1
 - config** option 5-1
 - Configuration 5-1, 8-4, 8-5, 8-6
 - Context 3-25, 3-32, 3-48, 3-49, 6-5, 6-67, 6-87, 6-147, 6-178, 6-190, 6-191, 6-192, Glossary-2, Glossary-5, Glossary-8
 - Context Menu 8-3
 - context panel 4-7, 4-12, 4-13
 - context, current 6-159
 - Context-sensitive help 8-2, 8-21
 - continue** 1-5, 3-5, 3-21, 6-138, 6-140, 6-149
 - Continuing execution 1-5, 2-5, 2-6, 3-50, 4-10, 4-20, 4-27, 6-138, 6-139, 6-146, 6-148, 6-149, 8-73
 - Convenience variable 3-46, 6-5, 6-7, 6-99, Glossary-3
 - global 3-46, 6-73, 6-176
 - predefined 3-34, 6-5, 6-7, 6-168, 6-169, 6-170, 6-171, 6-172, 6-173, 6-177, 6-193
 - process local 3-46, 6-7, 6-73, 6-176
 - Core file 3-4, 3-5, 3-25, 5-1, 6-46, 6-192, 8-31, Glossary-3
 - core** option 5-1
 - core-file** 6-46, 6-47, 6-85
 - CPU bias 6-28
 - CPU hang 3-56
 - Crossing count Glossary-3
 - CUDA 3-51, 6-83, 8-48, Glossary-3
 - CUDA and X11 Display 6-83, 8-48
 - CUDA Coordinates Panel 8-107
 - CUDA Lanes Panel 8-107
 - CUDA Partitioning 6-186, 8-93
 - CUDA registers 6-15
 - CUDA software preemption 6-83, 8-48

- CUDA Warp Locals Panel 8-109
- Cuda_syscall interest level 6-75
- current context 6-159
- Current frame 3-33, 6-142, 6-143, 6-145, 6-146, 8-25
- Current Process 8-4
- Current process Glossary-3
- Current source file 6-87, 6-90, 6-91
- Current stack frame 3-32, 3-33, 4-17, 4-19, 4-27, 6-93, 6-115, 6-116, 6-120, 6-122, 6-123, 6-125, 6-132, 6-147, 6-154, 6-156, 6-175, 6-177, Glossary-3
- current thread 3-49, 6-159
- Current working directory 6-85

D

- d (delete)** 6-128, B-9
- d key 8-72
- Data definitions
 - global 6-109
 - static 6-109
- Data Display button 8-29
- Data item 8-78, Glossary-3
- Data menu 8-17
- Data Panel
 - saving contents 8-19, 8-98
- Data panel 8-17
- Data Panel Add Heap Errors Dialog Box 8-101
- Data Panel Add Heap Leaks Dialog Box 8-101
- Data Panel Add Still Allocated Blocks Dialog Box 8-101
- Data Panel Condition Filter Expression Dialog Box 8-105
- Data Panel Describe Dialog Box 8-102
- data panel layout Glossary-4
- Data Panel Linked List Expression Dialog Box 8-104
- Data type
 - printing 6-191
- Data Window 3-24, 4-23, 4-25, 4-28, 6-104, 8-29, 8-78
 - loading layout 8-19, 8-98
 - saving layout 8-19, 8-98
- Data Window Add Expression Dialog Box 8-100
- Data Window Call Stack Frames Dialog Box 8-102
- Data Window Edit Expression Dialog Box 8-102
- Data Window Expand Tree Dialog Box 8-102
- Data Window Item Dialog Box 8-100
- Data Window layout 8-19, 8-98, 8-102, 8-103
- Data Window Load Layout Dialog Box 8-102
- Data Window Pointer Array Dimension Dialog Box 8-103
- Data Window Save Layout Dialog Box 8-103
- Data Window Save Snapshot Dialog Box 8-103

- Data Window Subscript Array Dialog Box 8-104
- Data Window Subscript Enum Array Dialog Box 8-104
- data-display** 3-24, 6-104
- debug** 3-4, 3-6, 6-23, 6-29, 6-30, 6-183
- Debug command area 2-3, 2-6
- Debug Eventpoint Dialog Boxes 8-35
- Debug Eventpoint menu 8-35
- Debug file directory 6-31
- Debug File Selection Dialog Box 8-35
- Debug group area 4-11, 4-15, 4-20, 4-24, 4-28
- Debug identification area 2-3, 4-12
- Debug information 4-4, 6-45, 6-80
- Debug Interrupt button 8-25
- Debug menu bar 8-5
- Debug message area 2-3, 2-6, 4-8, 4-9, 4-10, 4-11, 4-12, 4-13, 4-14, 4-15, 4-19, 4-22, 4-24, 4-26, 8-34
- Debug NightView menu 2-6
- Debug session Glossary-4
- Debug source display 2-3, 2-5, 2-6, 4-10, 4-15, 4-17, 4-20, 4-22, 4-26, 8-13, 8-14, 8-15, 8-29, 8-34, 8-36, 8-65, 8-68, 8-69, 8-70
- Debug source file name 2-3, 4-13, 4-17, 4-20, 4-28
- Debug Source menu 8-12, 8-34, 8-65, 8-72
- Debug Source Selection Dialog Box 8-13, 8-34, 8-68
- Debug status area 2-3, 4-11, 4-15, 4-20, 4-24, 4-28
- Debug status message 8-74
- Debug table 6-188, 6-190
- Debug View menu 4-27
- Debug Window 2-3, 4-15, 8-5, 8-30, 8-34, 8-35
- Debugger 3-2, Glossary-4
 - gdb 1-1, 1-4
 - NightView i-iii, 3-2
 - symbolic 3-2
- Debugging
 - multiple processes 3-3
 - real-time 3-7
 - single process 3-3
- Declaration
 - printing 6-192
- declare-thread-tag** 6-162
- Default font Glossary-4
- define** 6-193, 6-196
- Defining a macro 6-194
- delay** 6-164
- delete** 6-111, 6-126, 6-128
- Delete button 8-40
- Deleting breakpoints 4-20, 6-128, 8-40
- Deleting eventpoints 6-128, 8-40
- Deleting heappoints 6-128, 8-40
- Deleting monitorpoints 6-128, 8-40
- Deleting patchpoints 6-128, 8-40
- Deleting tracepoints 6-128, 8-40
- Deleting watchpoints 6-128, 8-40
- detach** 6-30, 6-33, 6-44, 8-11

- Detaching 3-3, 3-24, 6-26, 6-33, 6-44, 6-120, 8-11, Glossary-4
- Dialog Box 2-3, 2-6, 8-3
 - Attach 3-4, 8-33
 - Breakpoint 4-21, 4-22, 4-23, 8-15, 8-27, 8-35
 - Data Panel Add Heap Errors 8-101
 - Data Panel Add Heap Leaks 8-101
 - Data Panel Add Processes 8-18, 8-97
 - Data Panel Add Shells 8-18, 8-97
 - Data Panel Add Still Allocated Blocks 8-101
 - Data Panel Condition Filter Expression 8-105
 - Data Panel Describe 8-102
 - Data Panel Linked List Expression 8-104
 - Data Window Add Expression 8-17, 8-96, 8-100
 - Data Window Add Heap Errors 8-18, 8-97
 - Data Window Add Heap Information 8-18, 8-97
 - Data Window Add Heap Leaks 8-19, 8-97
 - Data Window Add Local Variables 8-18, 8-97
 - Data Window Add Registers 8-18, 8-97
 - Data Window Add Stack 8-18, 8-97
 - Data Window Add Still Allocated Blocks 8-19, 8-97
 - Data Window Add Threads 8-18, 8-19, 8-97, 8-98
 - Data Window Call Stack Frames 8-102
 - Data Window Edit Expression 8-102
 - Data Window Expand Tree 8-102
 - Data Window Item 8-100
 - Data Window Load Layout 8-102
 - Data Window Pointer Array Dimension 8-103
 - Data Window Save Layout 8-103
 - Data Window Save Snapshot 8-103
 - Data Window Subscript Array 8-104
 - Data Window Subscript Enum Array 8-104
- Debug Eventpoint 8-35
- Debug File Selection 8-35
- Debug Heap 8-41
- Debug Source Selection 8-13, 8-34, 8-68
- Eventpoint 8-35
- File Selection 8-13, 8-69
- Heappoint 8-16, 8-28, 8-35
- Monitorpoint 8-16, 8-28, 8-35
- Monitorpoint Update Interval 8-106
- Patchpoint 4-25, 8-16, 8-27, 8-35
- Remote Login 8-43
- Select a Function/Unit 4-14, 8-13, 8-34, 8-68
- Select a Source File 8-13, 8-34, 8-68, 8-69
- Tracepoint 8-16, 8-28, 8-35
- Warning 4-28, 8-3, 8-6, 8-10, 8-11, 8-25
- Watchpoint 8-16, 8-17, 8-28, 8-35
- Dialogue 3-6, 3-7, 4-4, 4-11, 4-13, 6-19, 6-27, 6-29, 6-43, 6-183, 8-3, 8-10, 8-43, C-1, Glossary-4
 - commands on 6-34
 - local 3-6, 5-1, 6-2
 - local - with on dialogue 6-34, 6-36
 - printing 6-182
 - remote 3-7, 6-29, 8-10, 8-43
 - starting 6-27
 - terminating 6-33
- Dialogue I/O area 2-2, 2-3, 3-7, 4-23, 4-27
- Dialogue input 6-36, 6-41, 6-71
- Dialogue name 6-29
- Dialogue NightView menu 2-6
- Dialogue output 3-7, 6-36, 6-38
- Dialogue prompt 5-2
- Dialogue shell 1-2, 2-2, 3-6, 3-7, 8-73, C-1
 - login 6-27
 - logout 6-33
- Dialogue Window 2-3
- Dimmed button 4-20, 8-34, 8-40
- Dimmed label 8-37
- Directory
 - current 6-85
 - debug file 6-31
- directory** 3-23, 6-85, 6-87, 6-88
- Directory searching 3-23, 6-87, 6-88, 6-176
- disable** 6-111, 6-128
- Disabled button 4-20, 8-34, 8-40
- Disabling a breakpoint 4-22, 4-26, 6-128, 6-130, 8-37
- Disabling a heappoint 6-130, 8-37
- Disabling a monitorpoint 6-130, 8-37
- Disabling a patchpoint 6-130, 8-37
- Disabling a tracepoint 6-130, 8-37
- Disabling a watchpoint 6-130, 8-37
- Disabling an eventpoint 6-128, 6-130, 8-37
- Disassembly 6-81, 6-101, 6-142, 6-143, 6-145, 8-65, Glossary-4
- Disassembly display mode 8-14, 8-65, 8-70, 8-72
- Dismiss button 8-3
- display** 3-23, 6-6, 6-106, 6-108, 6-145, 6-176
- Display Ada exception handling 6-185
- Display addresses limits 6-69, 6-177
- Display arguments 6-178, 6-189
- Display array 6-69
- Display breakpoint 4-26, 6-167, 6-168
- Display checkpoint information 6-184
- Display convenience variables 6-176
- Display declaration 6-192
- Display dialogue information 6-182
- Display display variables 6-176
- DISPLAY environment variable 5-3
- Display eventpoint 4-26, 6-167, 6-183
- Display expression 6-176, 6-191, 8-17, 8-96
- Display expression limits 6-177
- Display family information 6-183
- Display file names 6-192
- Display function names 6-190
- Display global variable 6-189
- Display heap errors 8-18, 8-97

- Display heap information 8-18, 8-97
- Display heap leaks 8-19, 8-97
- Display heappoint 6-167, 6-172
- Display item Glossary-4
- Display line number 6-192
- Display local variables 6-189, 8-18, 8-82, 8-97
- Display log file information 6-167
- Display macro 6-198
- Display mode
 - disassembly 8-14, 8-65, 8-70, 8-72
 - source 8-65
- Display monitorpoint 6-167, 6-171
- Display monitorpoint values 8-19, 8-98
- Display on program commands 6-184
- Display on restart commands 6-184
- Display patchpoint 4-26, 6-167, 6-170
- Display process information 6-178, 8-18, 8-97
- Display registers 8-82
- Display search path 6-176
- Display shell information 8-18, 8-97
- Display source file 6-86, 6-160
- Display source file names 6-190
- Display stack frame
 - all 6-95, 6-104, 8-18, 8-97
 - one 6-175
- Display still allocated blocks 8-19, 8-97
- Display string limits 6-69
- Display thread information 8-18, 8-97
- Display tracepoint 6-167, 6-169
- Display type 6-192
- Display type information 6-191
- Display value history 6-176
- Display variable 1-5, 2-6, 3-23, 4-14, 4-18, 4-23, 6-96, 6-192
- Display variable address 6-190
- Display watchpoint 6-167, 6-173, 6-174
- Documentation
 - online 1-1, 1-3, 2-1, 2-3, 6-159, 8-2, Glossary-7, Glossary-9
- Double clicking 8-34, 8-35, 8-103, 8-104
- down** 6-155, 6-156, 6-159, 8-25, 8-72
- Down button 4-19
- Down button 8-25
- DWARF 3-48, 6-67, Glossary-4
- dynamic library translations 6-47
- Dynamic linker 3-56
- Dynamically loaded library 3-5, 3-26, 3-56, 3-57, 6-32

E

- e** key 8-72
- echo** 6-24, 6-104, 6-197

- edit** 6-89
- Edit source file 6-89, 8-14, 8-70, 8-72
- Editor
 - emacs 6-78, 7-2
 - gmacs 6-78, 7-2
 - vi 6-78, 7-2
- EDITOR environment variable 7-2
- Elaboration 3-49
- ELF 3-48, 6-89, Glossary-5
- emacs editor 6-78, 7-2
- enable** 6-111, 6-113, 6-128, 6-129, 8-24, 8-72
- Enabling a breakpoint 6-129, 8-37
- Enabling a heappoint 6-129, 8-37
- Enabling a monitorpoint 6-129, 8-37
- Enabling a patchpoint 6-129, 8-37
- Enabling a tracepoint 6-129, 8-37
- Enabling a watchpoint 3-17, 6-129, 8-37
- Enabling an eventpoint 6-129, 8-37
- end define** 6-194
- end on dialogue** 6-34
- end on program** 6-50
- end on restart** 6-52
- Enter key 8-72
- Entry point 6-109
- Environment variable
 - DISPLAY 5-3
 - EDITOR 7-2
 - NIGHTVIEW_ENV 3-6, 3-9
 - PATH 3-10
 - SHELL 6-161
 - TERM 3-6, 7-1
 - VISUAL 7-2
- Error
 - abort 3-43
 - caution 3-43
 - warning 3-43
- Error message 6-175, 8-74
- Errors 1-3, 2-3, 3-43, 3-49, 8-21, Glossary-5
- Evaluation of expressions 3-25, 6-99, 6-104
- Event notification 6-42, 6-43
- Event-driven command streams 3-44
- Event-map file 6-119, 6-120, Glossary-5
- Eventpoint 3-11, 3-22, 3-23, 3-24, 6-20, 6-109, 6-111, 6-118, 6-128, 6-169, 6-170, 6-171, 6-172, 8-15, 8-35, 8-74, Glossary-5
 - changing 8-35
 - clearing 6-125, 8-28
 - commands on 8-38, 8-39
 - condition on 3-11, 3-23, 3-46, 6-127, 6-130, 6-132, 6-168, 8-38
 - deleting 6-128, 8-40
 - disabling 6-128, 8-37
 - displaying 4-26, 6-167
 - enabling 6-129, 8-37

- hitting 6-17, 6-168, 8-37, Glossary-7
 - ignoring 3-12, 3-23, 6-130, 6-168, 8-38
 - inserted 3-11, 3-13, 3-17, 3-25, 3-26, 3-46, 6-6, Glossary-7, Glossary-8
 - named 3-23, 8-38
 - naming 6-20, 6-114
 - printing 6-183
 - removing 6-126
 - saving 3-23
 - setting 3-5, 8-35
 - state 6-167, 8-37
 - Eventpoint crossing count Glossary-3
 - Eventpoint Dialog Boxes 8-35
 - Eventpoint ID 8-37
 - Eventpoint menu** 8-15, 8-35
 - Debug 8-35
 - Eventpoint modifier 6-113, Glossary-5
 - /delete 6-113
 - /disabled 6-113, Glossary-5
 - /f 6-113
 - Eventpoint number 3-11, 8-37
 - Eventpoint state 8-74
 - Eventpoint summary 8-35, 8-74
 - Eventpoints
 - changing 8-35, 8-74
 - Event-triggered commands 6-34, 6-51
 - Exception 6-142, 6-143, 6-145, 6-146, 6-147
 - Ada 3-29, 6-151, 6-185, Glossary-5
 - exception(info exception)** 6-185
 - Exception handling 3-49, 6-150, 6-153, 6-185
 - saving 3-23
 - exec 3-23, 4-7, 6-44, 6-50, 6-52, 6-55, 6-56, 8-31
 - exec-file** 3-10, 3-57, 6-32, 6-46, 6-48, 6-50, 6-85
 - Executable
 - stripped 6-32, 6-45
 - Executable and linking format Glossary-5
 - Executable file 3-2, 6-45, 6-46, 6-49, 6-188, 6-192
 - Execution
 - continuing 1-5, 2-5, 2-6, 3-50, 4-10, 4-20, 4-27, 6-138, 6-139, 6-146, 6-148, 6-149, 8-73
 - restarting 3-22, 3-23, 6-50, 6-53, 6-56, Glossary-1, Glossary-11
 - resuming 1-5, 2-5, 2-6, 4-10, 4-20, 4-27, 6-137, 6-138, 8-24
 - starting 1-2, 2-2, 3-22
 - stopping 1-4, 2-5, 3-49, 4-9, 4-21, 4-22, 6-114, 6-133
 - Exit messages 4-28
 - Exiting 1-6, 2-6, 3-23, 6-26, 6-56
 - Expression 3-32, 6-68
 - Ada 3-28
 - conditional 3-29, 4-21, 6-115, 6-120
 - displaying 6-104, 6-176, 8-17, 8-29, 8-96
 - evaluation 3-25, 3-27, 6-99, 6-104
 - floating-point 3-27
 - insertion 6-116, 6-132
 - intrinsics 6-16
 - language 6-106, 6-107, 6-108, 6-118
 - logging 6-120
 - memory address 6-107
 - patchpoint 6-120, 8-39
 - printing 1-5, 4-14, 4-18, 4-23, 6-96, 6-107, 6-171, 6-191, 8-29
 - regular 6-21, 6-34, 6-77, 6-90, 6-177, 6-183, 6-189, 6-190, 6-191, 6-192, 6-198, 8-13, 8-34, 8-68
 - regular examples 6-22
 - syntax 6-5
 - Expression Evaluation 3-32
 - Expression limits 6-177
 - Expressions
 - monitoring 6-122
 - External data definitions 6-109
 - External variable 6-156
 - printing 6-189
- ## F
- f(frame)** 6-155, B-11
 - f key 8-72
 - F1 key 8-2
 - fact program 1-1, 2-1
 - Family 3-4, 3-24, 6-19, 6-28, 6-54, 6-68, 6-113, 6-148, 6-183, Glossary-5
 - printing 6-183
 - family** 6-46, 6-49, 6-54
 - FBS 3-55
 - fbswait 3-55
 - File
 - .NightViewrc 3-47, 6-33, 6-34, 6-36
 - .profile 3-9
 - commands 6-161
 - core 3-4, 3-5, 3-25, 5-1, 6-46, 6-192, 8-31, Glossary-3
 - event-map 6-119, 6-120, Glossary-5
 - executable 3-2, 6-45, 6-46, 6-49, 6-188, 6-192
 - initialization 5-2, 5-3, 6-161, Glossary-7
 - library 3-2
 - log 6-66, 6-167
 - object 3-2, 3-15, 6-109
 - ReadyToDebug 1-3, 3-10, 4-5
 - source 3-2, 4-3, 4-7, 4-10, 4-13, 4-15, 4-17, 4-20, 4-22, 4-26, 6-86, 6-87, 6-88, 6-89, 6-161, 6-176, 6-190, 6-193, 8-12, 8-13, 8-65, 8-68
 - symbol 6-45, 6-49
 - trace event-map 6-119, 6-120, Glossary-5

- File access 3-2, 3-9
- File Coding 8-71
- File menu 8-5
- File name
 - printing 6-192
- File Selection Dialog Box 8-13, 8-69
- Filter
 - PID 8-33
 - Program 8-33
 - User 8-33
- Filter button 8-33
- finish** 6-146, 8-24, 8-30, 8-72
- Finish button 8-24
- Floating-point expressions 3-27
- fo (forward-search)** 6-90, B-7
- Focus
 - keyboard 8-2, Glossary-6, Glossary-8
- Font
 - default Glossary-4
- forbid safety level 5-2, 6-26, 6-33, 6-45, 6-54, 6-72, 6-128, 6-129
- Forking 3-3, 3-4, 4-7, 4-13, 6-44, 8-30, Glossary-6
- Formal argument
 - macro 6-194, 6-196
- Fortran i-iii, 3-11, 3-30, 3-32, 3-48, 4-3, 4-4, 4-9, 4-12, 4-15, 4-16, 4-17, 4-19, 4-20, 4-21, 4-22, 6-67, 6-97, 6-98, 6-117, 6-132, 6-195, Glossary-2
- forward-search** 6-22, 6-88, 6-90
- Frame
 - displaying 6-175
 - stack 3-32, 3-33, 4-17, 4-19, 4-27, 6-6, 6-87, 6-93, 6-109, 6-115, 6-116, 6-120, 6-122, 6-123, 6-125, 6-132, 6-147, 6-154, 6-156, 6-175, 6-177, 6-190, Glossary-3, Glossary-6, Glossary-12
 - stack - printing 6-95, 8-84
- frame** 3-33, 6-141, 6-147, 6-154, 6-159, 8-73
- Frame address 6-8, 6-155, 6-175
- Frame pointer 6-8, 6-176
- Frame zero 6-7, 6-76, 6-93, 6-149, 6-155, 6-156, 6-177
- Frames
 - hidden 6-7, 6-76, 6-93, 6-149
- Frequency-Based Scheduler 3-55
- Full-screen interface i-iii, 1-1, 3-36, 5-2, 6-2, 6-122, 6-126, 6-161, 6-194, 7-1, 7-2, 8-1, Glossary-6
- Full-screen user interface 6-160
- Function 4-10, 4-15, 6-141, 6-142, 6-144, 6-145, 6-146, 6-176
 - static - location of 6-17
- Function arguments
 - printing 6-189
- Function name
 - list 6-190

G

- g option 1-2, 2-2
- gdb 1-1, 1-4
- GID 3-54
- Global data definitions 6-109
- Global variable 3-25, 6-156
 - printing 6-189
- gmacs editor 6-78, 7-2
- Graphical user interface i-iii, 2-1, 3-36, 3-45, 5-2, 5-3, 6-1, 6-45, 6-152, 6-160, 8-1, Glossary-6
- Group ID 3-54
- GUI 2-1, 3-36, 3-45, 5-2, 5-3, 6-1, 6-45, 6-152, 8-1, Glossary-6
- Guide
 - command summary A-1

H

- H (source line decoration) 6-92
- h key 8-72
- handle** 3-20, 3-23, 3-49, 6-137, 6-150, 6-178
- hd (heapdebug)** 6-59, B-4
- Heap 6-179
- heapcheck** 3-42, 6-62, 6-187
- heapdebug** 3-15, 3-37, 3-39, 3-42, 6-59, 6-123, 6-188, 8-40, 8-41
- Heappoint 3-11, 6-44, 6-92, 6-111, Glossary-5
 - changing 8-35
 - clearing 6-125
 - condition on 3-23, 6-127, 6-168, 6-173, 8-38
 - deleting 6-128, 8-40
 - disabling 6-128, 8-37
 - displaying 6-167, 6-172
 - enabling 6-129, 8-37
 - hitting 6-173, 8-37
 - ignoring 3-23, 6-130, 6-168, 6-173, 8-38, Glossary-7
 - named 3-23, 6-114, 6-123, 8-38
 - saving 3-23
 - setting 6-109, 8-35
 - state 6-172, 8-37
 - temporary 8-37
- heappoint** 3-15, 3-42, 6-62, 6-123, 6-172
- Heappoint crossing count Glossary-3
- Heappoint Dialog Box 8-16, 8-28, 8-35
- Help
 - context-sensitive 8-2, 8-21
- help** 1-1, 1-3, 3-44, 6-159, 8-2, 8-3, 8-21, Glossary-7, Glossary-9
- Help button 8-3, 8-34, 8-35, 8-40, 8-100, 8-103

- Help menu 2-1, 2-3, 2-4, 4-5, 8-2, 8-21
 - help option 5-1
 - Help system
 - movement 8-2, 8-122
 - Help Window 2-1, 2-3, 2-4, 8-2, 8-3, 8-122, Glossary-7
 - Help window
 - exiting 2-4, 2-5
 - Hidden frames 6-7, 6-76, 6-93, 6-149
 - History
 - command 3-47
 - value 3-47, 4-14, 6-5, 6-68, 6-98, 6-99, 6-103, 6-176, Glossary-13
 - Hit a breakpoint 8-37
 - Hit a heappoint 8-37
 - Hit a monitorpoint 8-37
 - Hit a patchpoint 8-37
 - Hit a tracepoint 8-37
 - Hit a watchpoint 8-37
 - Hit an eventpoint 8-37
 - hold(mcontrol hold)** 6-124
 - hold(mcontrol-hold)** B-9
 - Hold button 8-78
 - Hollerith data 6-97
- I**
- i b(info breakpoint)** 6-168, B-12
 - I/O 3-7
 - Iconifying windows 8-2
 - ID
 - group 3-54
 - process 3-4, 3-6, 3-24, 4-7, 4-13, 6-19, 8-33, Glossary-10
 - trace-event 6-119, 6-120, 8-39, Glossary-13
 - user 3-54
 - ignore** 6-111, 6-118, 6-130, Glossary-7
 - Ignore count Glossary-7
 - Ignoring breakpoints 3-23, 4-22, 6-115, 6-130, 6-138, 6-168, 6-169, 8-38, Glossary-7
 - Ignoring eventpoints 3-12, 3-23, 6-130, 6-168, 8-38
 - Ignoring heappoints 3-23, 6-130, 6-168, 6-173, 8-38, Glossary-7
 - Ignoring monitorpoints 3-23, 6-130, 6-172, 8-38, Glossary-7
 - Ignoring patchpoints 3-23, 6-118, 6-130, 6-168, 6-171, 8-38, Glossary-7
 - Ignoring syscallpoints 6-174
 - Ignoring tracepoints 3-23, 6-120, 6-130, 6-168, 6-170, 8-38, Glossary-7
 - Ignoring watchpoints 3-23, 6-130, 6-168, 6-174, 8-38, Glossary-7
 - iHawk Series 860 3-17, 6-8, 6-100, 6-134, 8-30, 8-36, C-1, C-2
 - iHawk Series 870 3-17, 6-11, 6-100, 6-134, 8-30, 8-36, C-1, C-2
 - info address** 6-190
 - info args** 6-189
 - info breakpoint** 6-167, 6-168
 - info convenience** 6-176
 - info declaration** 6-192
 - info dialogue** 6-30, 6-74, 6-182
 - info directories** 6-89, 6-176
 - info display** 6-107, 6-108, 6-176
 - info eventpoint** 6-128, 6-167
 - info exception** 6-154, 6-185
 - info family** 6-183
 - info files** 6-192
 - info frame** 6-6, 6-175
 - info functions** 3-48, 6-190
 - info heappoint** 6-172
 - info history** 6-176
 - info limits** 6-69, 6-177
 - info line** 6-88, 6-192
 - info locals** 6-189
 - info log** 6-167
 - info macros** 6-198
 - info memory** 6-74, 6-179
 - info monitorpoint** 6-167, 6-171
 - info name** 6-183
 - info on dialogue** 6-34, 6-184
 - info on program** 6-51, 6-184
 - info on restart** 3-23, 6-53, 6-76, 6-184
 - info patchpoint** 6-167, 6-170
 - info process** 6-178
 - info registers** 3-34, 6-11, 6-15, 6-176, 6-177
 - info representation** 6-191
 - info signal** 6-150, 6-178
 - info sources** 6-190
 - info tracepoint** 6-167, 6-169
 - info types** 6-191
 - info variables** 6-189
 - info watchpoint** 6-173, 6-174
 - info whatis** 6-191
 - Initialization file 5-2, 5-3, 6-161, Glossary-7
 - Initialize tracing 6-119
 - Inline interest level 6-75
 - Inline subprograms 3-34, 6-143, 6-146
 - Input
 - dialogue 6-36, 6-41, 6-71
 - program 1-3, 2-3, 3-7, 4-11, 6-37, 6-71
 - shell 8-73
 - Input area 4-21, 4-22, 4-23, 4-25
 - Input command 6-161
 - Input terminator 6-36, 6-71
 - Inserted eventpoints 3-11, 3-13, 3-17, 3-25, 3-26, 3-46, 6-6, Glossary-7, Glossary-8

Instruction

branch 6-116, 6-132

Intel IA-32 registers (iHawk Series 860) 6-8

interest 3-35, 6-74, 6-77, Glossary-8

Interest level

cuda_syscall 6-75

inline 6-75

justlines 3-24, 3-35, 6-75

nodebug 3-24, 3-35, 6-75

subprogram 3-24, 3-35, 6-74, Glossary-8

Interest level threshold 3-24, 3-35, 6-75, Glossary-8

Interesting subprograms 3-24, 3-33, 3-35, 4-13, 6-7, 6-75, 6-142, 6-145, Glossary-8

Interface

command-line i-iii, 1-1, 3-36, 4-2, 6-2, 6-122, 6-126, 6-152, 6-194, 7-1, 8-1, 8-23, Glossary-2

full-screen i-iii, 1-1, 3-36, 5-2, 6-2, 6-122, 6-126, 6-161, 6-194, 7-1, 7-2, 8-1, Glossary-6

graphical user i-iii, 2-1, 3-36, 3-45, 5-2, 5-3, 6-1, 6-45, 6-152, 8-1, Glossary-6

Interrupt button 8-25

Interrupt the debugger 8-25

Interrupting the debugger 3-36, 3-45, 6-125, 8-25

Interrupts 3-56

Intrinsic

__eventpoint_data__ 6-16

Intrinsics 6-16

Invoking NightStar Tools 8-20

Invoking the debugger 1-2, 2-2, 3-47, 5-1

IPL register 3-56

J

Job control 3-25

jump 6-148

Justlines interest level 3-24, 3-35, 6-75

K

Key

= 8-73

> 8-73

b 8-72

d 8-72

e 8-72

Enter 8-72

f 8-72

F1 8-2

h 8-72

N 8-72

n 8-72

p 8-72

r 8-72

Return 4-21, 4-22, 4-23, 4-25, 6-24, 6-25, 6-88, 7-1, 8-34

S 8-72

s 8-73

u 8-73

Keyboard focus 8-2, Glossary-8

Kill 8-25

kill 3-20

kill 6-45, 8-11, 8-25

Kill button 6-45, 8-25

Killing processes 3-23, 6-26, 6-45, 8-11, 8-25

L

l (list) 1-4, 6-87, B-7

Label

dimmed 8-37

lane Glossary-8

Language 6-120, 6-147

machine 3-2, 3-48

Language expression 6-106, 6-107, 6-108, 6-118

Language support i-iii, 3-27, 3-32, 4-3, 6-67, 6-178

Library

dynamically loaded 3-5, 3-26, 3-56, 3-57, 6-32

shared 3-5, 3-26, 3-56, 3-57, 6-32

Library file 3-2

licences 5-3

License manager 5-3

Limits

addresses 6-177

expression 6-177

Line decorations 2-5, 4-7, 4-10, 4-11, 4-13, 4-16, 4-18, 4-20, 4-22, 4-26, 6-88, 6-92, 6-142, 6-143, 8-14, 8-24, 8-27, 8-28, 8-65, 8-70

Line number

printing 6-192

Linking 1-2, 2-2, 6-121

dynamic Glossary-5

Linux 3-49

list 1-4, 3-48, 6-24, 6-85, 6-86, 6-88, 6-90

List assembly code 8-14, 8-15, 8-70

List function names 6-190

List mode 8-65

List source file 6-86

List source file names 6-190

load 3-15, 6-32, 6-109

Loading Data Window layout 8-19, 8-98

local dialogue 1-3, 3-6, 5-1, 6-2, 6-51

- with on dialogue 6-34, 6-36
- Local system 3-7
- Local variable 3-25, 6-5, 6-104, 8-18, 8-82, 8-97
 - printing 6-189
- Location
 - in executable program 6-17
 - printing 6-69, 6-190
- Location specifier 3-32, 6-17, 6-20, 8-15, 8-24, 8-28, 8-35, 8-36, 8-39
- Log
 - dialogue 6-38
- Log file 6-66, 6-167
- Logging
 - session 6-66
- login** 3-8, 6-27
- logout** 6-33, 8-10

M

- M (source line decoration) 6-93
- Machine language 3-2, 3-48
- Macro 3-24, 3-45, 6-24, 6-37, 6-193, Glossary-8
 - actual arguments 6-196
 - argument 6-194, 6-196
 - definition 6-194
 - example 6-24
 - formal arguments 6-194, 6-196
 - printing 6-198
 - recursion 6-194
 - referencing 6-196
 - replacing 6-194
 - restart_begin_hook 3-24, 6-50
 - restart_end_hook 3-24
 - string 6-196
- Macro body 3-45, 6-194, 6-199
- Macro expansion 6-34, 6-51
- Manual
 - online 1-1, 1-3, 2-1, 2-3, 6-159, 8-2, Glossary-7, Glossary-9
- Manual section 6-159
- mcontrol** 3-36, 6-124, 7-2
- md (memory-display)** 6-105, B-7
- Memory 6-155
 - output 6-102, 6-106
 - shared 6-73, C-1, C-2
 - static 6-190
- Memory address expression 6-107
- Memory layout 6-179
- Memory mapped I/O 3-56
- Memory Segments Panel 8-110
- memory-display** 6-105
- Menu

- context 8-3
- Data** 8-17
- Debug Eventpoint 8-35
- Debug NightView 2-6
- Debug Source 8-12, 8-34, 8-65, 8-72
- Debug View 4-27
- Dialogue NightView 2-6
- Eventpoint 8-15, 8-35
- File 8-5
- Help 2-1, 2-3, 2-4, 4-5, 8-2, 8-21
- Process 8-11, 8-12
- Shell 8-10
- Source 8-12, 8-34, 8-65, 8-72
- Tools 8-20
- View 8-6
- Menu bar
 - Debug 8-5
- Message
 - error 6-175, 8-74
 - exit 4-28
 - process status 8-74
- Message panel 8-74
- message panel 4-6
- mmap C-2
- Mnemonic Glossary-8
 - A 8-6, 8-8, 8-11, 8-13, 8-19, 8-69, 8-91, 8-94, 8-95, 8-116
 - B 8-9, 8-15, 8-67, 8-92
 - C 8-6, 8-8, 8-18, 8-21, 8-95, 8-110, 8-115
 - D 8-5, 8-7, 8-9, 8-11, 8-15, 8-17, 8-70, 8-75, 8-94, 8-95
 - E 2-3, 8-8, 8-10, 8-15, 8-17, 8-21, 8-67, 8-75, 8-94
 - F 8-5, 8-6, 8-69, 8-73, 8-74, 8-94, 8-95, 8-96, 8-112
 - G 8-7, 8-69, 8-73, 8-74, 8-76, 8-96
 - H 2-1, 2-3, 2-4, 8-7, 8-10, 8-12, 8-16, 8-18, 8-21, 8-67, 8-68, 8-93, 8-95, 8-108
 - I 8-9, 8-14, 8-70, 8-75, 8-76, 8-99
 - J 8-116
 - K 8-11, 8-12, 8-19, 8-22, 8-91
 - k 8-71
 - L 8-7, 8-8, 8-10, 8-18, 8-22, 8-67, 8-76, 8-91
 - M 8-7, 8-14, 8-15, 8-21, 8-67, 8-70, 8-76, 8-93
 - m 2-4
 - N 2-6, 8-9, 8-71, 8-76, 8-94
 - n 2-1
 - O 8-5, 8-8, 8-14, 8-70, 8-95
 - P 8-6, 8-11, 8-12, 8-16, 8-18, 8-20, 8-67, 8-112, 8-116
 - Q 8-22, 8-115
 - R 8-7, 8-8, 8-10, 8-11, 8-18, 8-69, 8-76, 8-92, 8-112, 8-115
 - S 8-5, 8-7, 8-12, 8-13, 8-16, 8-18, 8-20, 8-68, 8-76, 8-93, 8-96
 - T 8-9, 8-10, 8-16, 8-18, 8-19, 8-20, 8-22, 8-67,

- 8-91, 8-92, 8-108, 8-116
- U 8-7, 8-13, 8-17, 8-21, 8-22, 8-68, 8-94, 8-96
- V 8-6, 8-8, 8-19, 8-22, 8-92, 8-109, 8-116
- W 8-8, 8-16, 8-68, 8-92, 8-108
- X 2-6, 8-6, 8-91
- Y 8-71
- Z 8-9, 8-76
- Mode
 - list 8-65
- modify-eventpoint** 6-131
- Monitor refresh rate 6-125
- Monitor Window 3-36, 6-124, 7-2
 - simple full-screen 7-2
- Monitoring expressions 6-122
- Monitorpoint 3-11, 3-14, 3-15, 3-25, 3-29, 3-36, 3-56, 6-44, 6-73, 6-93, 6-111, 6-122, 6-124, Glossary-5, Glossary-9
 - changing 8-35
 - clearing 6-125
 - commands on 3-23, 6-126, 6-168, 6-172, 8-38, 8-39
 - condition on 3-23, 6-127, 6-168, 6-172, 8-38
 - deleting 6-128, 8-40
 - disabling 6-128, 8-37
 - displaying 6-167, 6-171
 - enabling 6-129, 8-37
 - hitting 6-172, 8-37
 - ignoring 3-23, 6-130, 6-168, 6-172, 8-38, Glossary-7
 - named 3-23, 6-114, 6-121, 8-38
 - saving 3-23
 - setting 6-109, 6-121, 8-35
 - state 6-171, 8-37
 - temporary 8-37
- monitorpoint** 3-14, 3-36, 6-121, 6-171, Glossary-9
- Monitorpoint crossing count Glossary-3
- Monitorpoint Dialog Box 8-16, 8-28, 8-35
- Monitorpoint Update Interval Dialog Box 8-106
- Motif 8-1
- Mouse button 1 2-1, 4-3, 8-1, 8-2, 8-78
- Mouse button 3 8-79
- mreserve** 6-57
- msg program 4-4, 4-7, 4-10, 4-20
- Multiple processes 3-3, 3-5, 7-1

N

- n (next)** 6-142, B-10
- N key 8-72
- n key 8-72
- name** 6-111, 6-113, 6-116, 6-132
- Named breakpoint 3-23, 6-114, 6-115, 6-131, 8-38
- Named eventpoint 3-23, 6-20, 6-114, 8-38

- Named heappoint 3-23, 6-123, 6-136, 8-38
- Named monitorpoint 3-23, 6-121, 8-38
- Named patchpoint 3-23, 6-116, 6-132, 8-38
- Named tracepoint 3-23, 6-120, 8-38
- Named watchpoint 3-23, 6-134
- Newline 6-24, 6-88
- next** 6-141, 6-142, 6-146, 8-24, 8-72
- Next** button 8-24
- nexti** 6-142, 6-143, 6-145, 8-24, 8-72
- Nexti** button 8-24
- NFS 3-3
- ni (nexti)** 6-145, B-10
- nice value 6-28
- NightProbe 8-20
- NightSim 3-55, 8-20
- NightStar Licence Manager 5-3
- NightStar tool set 8-20
 - NightProbe 8-20
 - NightSim 8-20
 - NightTrace 8-20
 - NightTune 8-21
- NightTrace 3-7, 3-15, 6-120, 6-121, 8-20, Glossary-9
- NightTune 8-21
- NightView i-iii, 3-2, Glossary-9
- NightView version 1-2, 5-2, 8-22
- NIGHTVIEW_ENV environment variable 3-6, 3-9
- NLSM 5-3
- nodebug** 3-4, 3-6, 6-23, 6-26, 6-30, 6-33, 6-35
- Nodebug interest level 3-24, 3-35, 6-75
- nogui option 1-2, 5-2
- nolocal option 5-2
- Notification of events 6-37, 6-42, 6-43, 6-152, 7-1
- notify** 6-42, 6-43
- ntrace 3-15, 6-119
- ntraceud 3-56, 6-121
- nview
 - invoking 5-1
- nview** 1-2, 2-2, 7-1
- nview
 - exiting 1-6, 2-6, 6-26
 - invoking 1-2, 2-2, 3-47
- nview option
 - arch 5-1
 - config 5-1
 - core 5-1
 - nogui 1-2, 5-2
 - nolocal 5-2
 - simplescreen 7-1
- nview options 5-1

O

Object activation 4-3
 Object file 3-2, 6-109
 Object filename translations 6-31, 6-45, 6-49, 6-109, 6-183
 Object selection 4-3, 4-5
 OK button 2-6, 8-34, 8-35, 8-40, 8-100, 8-101, 8-102, 8-103, 8-104, 8-105, 8-106
on dialogue 6-33, 6-35, 6-184
 on dialogue
 with local dialogue 6-34, 6-36
on program 3-24, 6-23, 6-49, 6-50, 6-52, 6-53, 6-184
on restart 3-23, 3-24, 6-50, 6-52, 6-184
 Online documentation 1-1, 1-3, 2-1, 2-3, 6-159, 8-2, Glossary-7, Glossary-9
 Online help system Glossary-9
 Optimization 3-47, 6-142, 6-143
 Option
 --arch 5-1
 -attach 5-1
 -config 5-1
 -core 5-1
 -g 1-2, 2-2
 -nogui 1-2, 5-2
 -nolocal 5-1, 5-2
 -simplescreen 7-1
 -x 5-3
 -xrm 5-3
 Options
 nview 5-1
 Output 3-7
 buffered 3-45
 dialogue 3-7, 6-38
 logging 3-46
 memory 6-102, 6-106
 program 8-73
 session 6-66
 shell 8-73
 suppressed 7-1
 text 6-104
output 6-103
 Output addresses limits 6-69
 Output array 6-69
 Output string limits 6-69
 Output variable 1-5, 2-6, 4-14, 4-18, 4-23, 6-96
 Overloading 3-32, Glossary-9

P

p (print) 1-5, 6-96, B-7

P (source line decoration) 6-93
 p key 8-72
 Packages
 Ada 3-48
 Panel
 data 8-17
 message 8-74
 panel
 context 4-7, 4-12, 4-13
 shell 4-11
 Parent process 4-13, 8-34
 Patch Glossary-9
 patch area Glossary-9
 Patchpoint 3-11, 3-29, 3-56, 6-44, 6-93, 6-94, 6-111, Glossary-5, Glossary-10
 changing 8-35
 clearing 6-125
 condition on 3-23, 6-118, 6-127, 6-168, 6-171, 8-38
 deleting 6-128, 8-40
 disabling 6-128, 8-37
 displaying 4-26, 6-167, 6-170
 enabling 6-129, 8-37
 hitting 6-171, 8-37
 ignoring 3-23, 6-118, 6-130, 6-168, 6-171, 8-38, Glossary-7
 named 3-23, 6-114, 6-132, 8-38
 saving 3-23
 setting 4-25, 6-109, 6-116, 6-132, 8-35
 state 6-170, 8-37
 temporary 6-132, 8-37
patchpoint 3-15, 6-116, 6-170
 Patchpoint crossing count Glossary-3
 Patchpoint Dialog Box 4-25, 8-16, 8-27, 8-35
 Patchpoints named 6-116
 PATH environment variable 3-10
 Pattern
 wildcard 6-23, 6-30, 6-50, 6-52, 6-190, 8-13, 8-34, 8-68
 wildcard examples 6-23
 Pattern matching 6-21, 6-29, 6-30, 6-78, 6-90, 6-183, Glossary-10
 Pattern matching examples 6-22
 PID 3-4, 3-6, 3-24, 4-7, 4-13, 6-19, 8-33, Glossary-10
 Pipelines 3-3, 3-6
 Pointer
 question mark 8-2
 Predefined convenience variable 3-34, 6-5, 6-7, 6-168, 6-169, 6-170, 6-171, 6-172, 6-173, 6-177, 6-193
 Preferences 8-6
print 1-5, 3-47, 6-68, 6-96, 6-99, 6-103, 6-106, 6-108, 6-122, 6-156, 6-190, 6-196, 8-29, 8-72
 print
 command attached to monitorpoint 6-122

- Print Ada exception handling 6-185
- Print addresses limits 6-69, 6-177
- Print arguments 6-178, 6-189
- Print array 6-69
- Print breakpoint 4-26, 6-167, 6-168
- Print button 2-6, 8-29, 8-72
- Print checkpoint information 6-184
- Print convenience variables 6-176
- Print current directory 6-85
- Print declaration 6-192
- Print dialogue information 6-182
- Print display variables 6-176
- Print eventpoint 4-26, 6-167
- Print eventpoint information 6-183
- Print expression 1-5, 4-14, 4-18, 4-23, 6-96, 6-107, 6-191
- Print expression limits 6-177
- Print family information 6-183
- Print file names 6-192
- Print function names 6-190
- Print global variable 6-189
- Print heappoint 6-167, 6-172
- Print line number 6-192
- Print local variables 6-189
- Print log file information 6-167
- Print macro 6-198
- Print monitorpoint 6-171
- Print on dialogue commands 6-184
- Print on program commands 6-184
- Print on restart commands 6-184
- Print patchpoint 4-26, 6-167, 6-170
- Print process information 6-178
- Print registers 6-177
- Print search path 6-176
- Print signal 6-178
- Print source file names 6-190
- Print stack frame
 - all 6-95
 - one 6-175
- Print string limits 6-69
- Print text 6-104
- Print tracepoint 6-167, 6-169
- Print type information 6-191, 6-192
- Print value history 6-176
- Print variable 6-192
- Print variable address 6-190
- Print watchpoint 6-167, 6-173, 6-174
- Print Window 8-6
- printf** 6-195, 6-197
- Procedure 4-10, 4-15, 6-141, 6-142, 6-144, 6-145, 6-146, 6-176, Glossary-10
- Procedure arguments
 - printing 6-189
- Procedure call 3-29, 6-117
- Procedure name
 - list 6-190
- Process 3-3, 8-11, 8-12, C-1, Glossary-10
 - abnormal termination 6-46
 - attaching to 5-1, 6-43, 8-11, 8-33
 - background 6-161
 - child 3-3, 3-4, 4-2, 4-7, 4-13, 6-30, 6-44, 6-55, Glossary-2
 - current 8-4
 - exiting 3-23, 3-25, 6-56
 - killing 3-23, 6-26
 - multiple 3-3, 3-5, 7-1
 - parent 4-13, 8-34
 - printing 6-178
 - pseudo 3-4, 3-25, 5-1, 6-46, 6-49
 - running 3-25, 3-26
 - single 3-3
 - stopped 3-25, 3-26, 3-32, 3-33
 - stopping 6-147, 6-148
 - stopping debugging 6-44
 - terminated 3-25
 - terminating 3-23, 3-24
- Process families 3-4
- Process ID 3-4, 3-6, 3-24, 4-7, 4-13, 6-19, 8-33, Glossary-10
- Process menu 8-11, 8-12
- Process state 3-25, 6-178, Glossary-10
- Program 3-3, Glossary-10
 - commands on 6-50
 - compiling 1-2, 2-2, 3-48
 - fact 1-1, 2-1
 - msg 4-4, 4-7, 4-10, 4-20
 - restarting 3-22, 3-23, 6-50, 6-53, 6-56, Glossary-1, Glossary-11
 - running 1-2, 6-37, 6-40
 - setuid 3-4
 - starting 3-22
- Program arguments 5-1
- Program counter 3-26, 3-32, 3-33, 6-7, 6-149, 6-177, Glossary-10
- Program I/O C-1
- Program input 1-3, 2-3, 3-7, 4-11, 6-37, 6-71
- Program location
 - specifying 6-17
- Program name 1-1, 2-1, 3-10, 4-4, 4-7, 5-2, 8-29
- Program output 3-7, 3-45, 3-46, 8-73
- Program, run 8-11
- Prologue 6-18
- Prompt 6-2, 6-70
 - (local) 1-3
 - > 6-122, 6-126, 6-194
 - \$ 1-3, 4-5
 - dialogue 5-2, 6-2
 - shell 1-3, 4-5

protected 3-49
 Protected thread 3-51
 protected thread 3-49
 protected thread tag 6-115
 ps 3-4
 Pseudo process 3-4, 3-25, 5-1, 6-46, 6-49
ptrace 3-4, C-1
ptype(info declaration) 6-192, B-13
pwd 6-85

Q

q(quit) 1-6, 6-2, 6-26, B-2
 Qualifier 3-5, 6-1, 6-19, 6-68, Glossary-10
 Quick command summary A-1
 quit 5-3
quit 1-6, 6-2, 6-26, 8-6
 Quitting 1-6, 2-6, 6-26

R

r key 8-72
 Radio button 8-39, 8-101
raise 3-49
 ReadyToDebug 1-3, 3-10, 4-5
 Real-time debugging 3-7
 Recursion
 macro 6-194
redisplay 6-108
 Referencing macros 6-196
refresh 6-160
 Refreshing terminal 6-160
regex 6-21, 6-90, 6-177, 6-183, 6-189, 6-191, 6-192, 6-198
 Register
 IPL 3-56
 Register variable 6-5
 Registers 3-2, 3-26, 3-32, 3-34, 6-5, 6-7, 6-104, 6-149, 6-155, 6-176, 6-190, 8-18, 8-82, 8-97, Glossary-10
 AMD64 (iHawk Series 870) 6-11
 CUDA 6-15
 display 8-82
 IA-32 (iHawk Series 860) 6-8
 printing 6-177
 vector 6-16
 Regular expression 6-21, 6-34, 6-77, 6-90, 6-177, 6-183, 6-189, 6-190, 6-191, 6-192, 6-198, 8-13, 8-33, 8-34, 8-68
 Regular expression examples 6-22

release(mcontrol release) 6-124, B-9
 Release button 8-78
 Remote dialogue 3-7, 6-29, 8-10, 8-43, Glossary-11
 Remote file access 3-9
 Remote Login Dialog Box 8-43
 Remote system 3-7
 Repeating commands 6-2, 6-24, 6-88
 Replacing commands 6-197
representation(info representation)
 6-191, B-13
rerun 6-41, 8-25
 Rerun button 8-25
 Rerunning a program 3-22, 3-23, 6-50, 6-53, 6-56
 Restart
 commands on 6-52
 restart_begin_hook macro 3-24, 6-50
 restart_end_hook macro 3-24
 Restarting a program 3-23, 6-50, 6-56
 Restarting execution 3-22, 3-23, 6-50, 6-53, 6-56, Glossary-1, Glossary-11
resume 3-21, 3-50, 3-55, 6-126, 6-137, 6-139, 8-24, 8-72
 Resume button 2-5, 2-6, 8-24
 Resuming execution 1-5, 2-5, 2-6, 4-10, 4-20, 4-27, 6-137, 6-138, 8-24
 Return key 4-21, 4-22, 4-23, 4-25, 6-24, 6-25, 6-88, 7-1, 8-34
reverse-search 6-22, 6-88, 6-90
 rlogin 6-29
 Routine 4-10, 4-15, 6-141, 6-142, 6-144, 6-145, 6-146, 6-176, Glossary-11
 trace_open_thread 6-121
 trace_start 6-121
 Routine arguments
 printing 6-189
 Routine name
 list 6-190
 Routine replacement 6-109
 rtcp 3-55
 rtutil 3-55
 run 3-8, 6-29, 8-43
run 1-2, 3-7, 6-37, 6-40, 6-197
 Run a program 1-2, 3-22, 6-37, 6-40, 6-41, 8-11, 8-25
 Run Mode 8-26
 Run to Here button 8-24, 8-72

S

s(step) 6-141, B-10
 S key 8-72
 s key 8-73
 Safety level

- forbid 5-2, 6-26, 6-33, 6-45, 6-54, 6-72, 6-128, 6-129
- unsafe 3-44, 5-2, 6-25, 6-33, 6-44, 6-45, 6-72, 6-128, 6-129
- verify 5-2, 6-25, 6-26, 6-33, 6-45, 6-54, 6-72, 6-128, 6-129
- Saving breakpoints 3-23
- Saving Data Panel contents 8-19, 8-98
- Saving Data Window layout 8-19, 8-98
- Saving eventpoints 3-23
- Saving exception handling 3-23
- Saving heappoints 3-23
- Saving monitorpoints 3-23
- Saving patchpoints 3-23
- Saving tracepoints 3-23
- Saving watchpoints 3-23
- Scheduler
 - frequency-based 3-55
- Scope 3-32, 3-46, 6-115, 6-120, 6-123, 6-136, 6-189, Glossary-11
- Script
 - debugger 6-161, 6-164
- Scroll bar 2-4, 2-5, 4-7, 8-34
- Search button 8-34
- Searching
 - function 8-13, 8-68
 - path 6-87, 6-88, 6-176
 - regular expression 6-90, 8-34
 - wildcard pattern 8-13, 8-34, 8-68
- Section
 - manual 6-159
- Select a Function/Unit Dialog Box 4-14, 8-13, 8-34, 8-68
- Select a Source File Dialog Box 8-13, 8-34, 8-68, 8-69
- select-context** 3-49, 6-147, 6-157, 6-163
- Selection
 - object 4-3, 4-5
- Selection policy
 - Browse 8-34
- Semicolon 6-117, 6-132
- Session
 - debug Glossary-4
- Session logging 6-66
- set** 6-99
- set-auto-frame** 6-77
- set-children** 3-3, 3-23, 4-7, 6-55
- set-cuda** 6-83
- set-cuda-memcheck** 6-83
- set-debug-file-directory** 6-31
- set-disassembly** 6-81
- set-download** 6-79, 6-80
- set-editor** 6-78, 7-2
- set-exit** 6-56
- set-format-hex** 6-71
- set-futurepoints** 6-82
- set-history** 6-68
- set-language** 3-24, 6-67, 6-178
- set-limits** 6-69, 6-98, 6-101, 6-167, 6-168, 6-169, 6-170, 6-171, 6-172
- set-local** 3-46, 6-72
- set-log** 6-66, 6-167
- set-notify** 6-42
- set-overload** 3-32, 6-77
- set-patch-area-size** 6-73, 6-183, C-2, Glossary-9
- set-preallocate** 6-78
- set-prompt** 6-2, 6-69
- set-qualifier** 6-19, 6-68
- set-restart** 3-23, 6-72
- set-safety** 6-34, 6-51, 6-71
- set-search** 6-77
- set-show** 3-7, 6-37, 6-39, 6-67, 6-167
- set-tag** 6-163
- set-terminator** 6-71
- Setting a breakpoint 1-4, 2-5, 4-9, 4-21, 4-22, 6-114, 6-131, 8-27, 8-35, 8-72
- Setting a conditional breakpoint 3-11, 4-21, 6-168, 6-169, 8-38, Glossary-2
- Setting a conditional eventpoint 6-127, 6-130, 6-132, 6-168, 8-38
- Setting a conditional heappoint 6-168, 6-173, 8-38
- Setting a conditional monitorpoint 6-168, 6-172, 8-38
- Setting a conditional patchpoint 6-168, 6-171, 8-38
- Setting a conditional tracepoint 6-168, 6-170, 8-38
- Setting a conditional watchpoint 6-168, 6-174, 8-38
- Setting a heappoint 8-35
- Setting a monitorpoint 6-121, 8-35
- Setting a patchpoint 4-25, 6-116, 6-132, 8-35
- Setting a tracepoint 6-120, 8-35
- Setting a watchpoint 6-133, 8-35
- Setting an eventpoint 8-35
- set-trace** 6-119, 6-120
- Setuid programs 3-4
- Shared library 3-5, 3-26, 3-56, 3-57, 6-32
- Shared memory 6-73, C-1, C-2
- Shell Glossary-11
 - dialogue 3-6, 3-7, 8-73, C-1
- shell** 3-4, 3-55, 3-56, 6-26, 6-85, 6-161, 6-197, 7-1, 8-23
- SHELL environment variable 6-161
- Shell I/O 8-73
- Shell menu 8-10
- shell panel 4-11
- Shell prompt 1-3, 4-5
- show** 3-7, 6-37, 6-38
- si (stepi)** 6-144, B-10
- SIGABRT 6-153
- SIGADA 6-153

- SIGALRM 6-152
- SIGFPE 6-153
- siginfo 3-21
- SIGINT 6-153
- signal** 3-20, 6-126, 6-137, 6-149
- Signals 3-20, 3-23, 4-2, 6-42, 6-51, 6-107, 6-137, 6-140, 6-141, 6-143, 6-145, 6-146, 6-147, 6-149, 6-151, Glossary-11
 - printing 6-178
- SIGQUIT 6-152
- SIGSEGV 6-153
- SIGSTOP 6-137, 6-140, 6-150, 6-152
- SIGTRAP 3-20
- SIGUSR1 4-2
- Simple full-screen interface i-iii, 1-1, 3-36, 5-2, 6-2, 6-122, 6-126, 6-161, 6-194, 7-1, 7-2, 8-1, Glossary-6
 - editing commands 7-2
- simplescreen option 7-1
- Single process 3-3
- Single stepping 3-20, 3-50, 3-56, 4-10, 4-15, 6-141, 6-142, 6-144, 6-145, 6-147, 8-24, 8-30, 8-72, 8-73
- SM Glossary-11
- smart printer Glossary-12
- Smart Printing 6-199
- smart-print** 6-200
- source** 3-23, 3-44, 3-47, 6-24, 6-161, 6-185
- Source display area 8-12, 8-34
- Source display mode 8-65
- Source file 3-2, 4-3, 4-7, 4-10, 4-13, 4-15, 4-17, 4-20, 4-22, 4-26, 6-161, 6-193, 8-12, 8-13, 8-65, 8-68
 - current 6-87, 6-90, 6-91
 - displaying 6-86, 6-160
 - editing 6-89, 8-14, 8-70, 8-72
 - list 6-190
 - search path for 6-87, 6-88, 6-176
- Source line decorations 2-5, 4-7, 4-10, 4-11, 4-13, 4-16, 4-18, 4-20, 4-22, 4-26, 6-88, 6-92, 6-142, 6-143, 8-14, 8-24, 8-27, 8-28, 8-65, 8-70
- Source listing 1-4, 2-3, 6-24, 6-86
- Source** menu 8-12, 8-34, 8-65, 8-72
 - Debug 8-12, 8-34, 8-65, 8-72
- source panel 4-7, 4-13
- Stack Glossary-12
- Stack examination 1-5, 2-6, 6-95
- Stack frame 6-6, 6-87, 6-109, 6-190, Glossary-6, Glossary-12
 - current 3-32, 3-33, 4-17, 4-19, 4-27, 6-93, 6-115, 6-116, 6-120, 6-122, 6-123, 6-125, 6-132, 6-147, 6-154, 6-156, 6-175, 6-177, Glossary-3
 - displaying 6-104, 6-175, 8-18, 8-97
 - printing 6-95
- Stack pointer 6-8, 6-177
- Stack variable 3-25
- Stale data indicator 3-36, 7-2, Glossary-12
- Starting execution 1-2, 2-2, 3-22
- Starting the debugger 1-2, 2-2, 3-47, 5-1
- Starting tracing 6-119
- State
 - breakpoint 6-169, 8-37
 - eventpoint 6-167, 8-37
 - heappoint 6-172, 8-37
 - monitorpoint 6-171, 8-37
 - patchpoint 6-170, 8-37
 - process 3-25, 6-178, Glossary-10
 - tracepoint 6-169, 8-37
 - watchpoint 6-173, 6-174, 8-37
- Static data definitions 6-109
- Static function
 - specifying location of 6-17
- Static memory 6-190
- Static variable 3-25, 6-5
- Status bar 4-6
- status bar 4-7, 4-13
- Status information 6-165
- step** 3-35, 3-49, 3-50, 4-15, 6-137, 6-141, 6-143, 6-145, 8-24, 8-73
- Step** button 8-24
- stepi** 6-142, 6-143, 6-144, 8-24, 8-72
- Stepi** button 8-24
- stop** 6-147, 8-24, 8-31
- Stop** button 8-24
- Stopping a process 6-147, 6-148
- Stopping execution 1-4, 2-5, 3-49, 4-9, 4-21, 4-22, 6-114, 6-133
- Stream
 - command Glossary-2
- String
 - C 3-56, 6-98
 - character 6-98, 6-177
 - macro 6-196
- String limits
 - printing 6-69
- strip 6-45
- Stripped executable 6-32, 6-45
- stty 7-1
- Subprogram 4-10, 4-15, 6-141, 6-142, 6-144, 6-145, 6-146, 6-176
- Subprogram arguments
 - printing 6-189
- Subprogram interest level 3-24, 3-35, 6-74, Glossary-8
- Subprogram name
 - list 6-190
- Subprograms
 - inline 3-34
 - interesting 3-24, 3-33, 3-35, 4-13, 6-7, 6-75, 6-142,

- 6-145, Glossary-8
- uninteresting 3-24, 3-33, 3-35, 4-13, 6-7, 6-75, 6-142, 6-145
- Subroutine 4-10, 4-15, 6-141, 6-142, 6-144, 6-145, 6-146, 6-176
- Subroutine arguments
 - printing 6-189
- Subroutine call 3-30, 6-117
- Subroutine name
 - list 6-190
- Substitution
 - text 6-193
- Summary of commands 8-21, A-1
- Summary of eventpoints 8-35, 8-74
- Symbol file 3-23, 6-45, 6-49, Glossary-12
- Symbol table 6-109, 6-165, 6-188, 6-190, 6-192
- symbol-file** 3-8, 3-23, 6-32, 6-45, 6-47, 6-85
- Symbolic debug information 6-45
- Symbolic debugger i-iii, 3-2
- Symbols
 - undefined 6-109
- Syntax
 - command 6-1
 - expression 6-5
 - intrinsic 6-16
 - qualifier 6-2
- Syscallpoint
 - displaying 6-174
 - hitting 6-174
 - named 6-136
 - state 6-174
 - syscall-list 6-136
- syscallpoint** 6-135
- Syscallpoint Dialog Box 8-28
- System
 - local 3-7
 - remote 3-7
- system 3-3, 6-55
- System crash 3-56

T

- T (source line decoration) 6-93
- Tag
 - trace-event 6-119, 6-120, Glossary-13
- tag 3-50, 6-117, 6-162, 6-163
- Task 3-49, 6-147, 6-157
 - Ada 8-1, 8-86, Glossary-1, Glossary-12
- tbreak** 6-111, 6-131
- Temporary breakpoint 6-131, 8-37
- Temporary heappoint 8-37
- Temporary monitorpoint 8-37

- Temporary patchpoint 6-132, 8-37
- Temporary tracepoint 8-37
- Temporary watchpoint 8-37
- TERM environment variable 3-6, 7-1
- Terminal 8-56
- Terminal emulator 8-56
- Terminal refresh 6-160
- Terminating a process 3-23, 3-24, 6-45, 8-11, 8-25
- Termination
 - abnormal 6-46
- Terminator
 - input 6-36, 6-71
- Text
 - printing 6-104
- Text cursor 8-24, 8-28
- Text input area 4-21, 4-22, 4-23, 4-25, 8-36, 8-38
- Text substitution 6-193
- Thread 3-49, 6-147, 6-157, 8-86, Glossary-12
- Thread process 3-49, 6-19, 6-147, 6-157, 6-163, 6-186, 8-86, Glossary-13
- Thread tag 3-50
- thread tag 6-117, 6-162, 6-163
- thread, current 3-49, 6-159
- Threshold
 - interest level 3-24, 3-35, 6-75
- Toggle button 8-39
- Tools menu 8-20
- tpatch** 6-111, 6-132
- Trace Glossary-13
- Trace initialization 6-119
- trace_open_thread routine 6-121
- trace_start routine 6-121
- Trace-event ID 6-119, 6-120, 8-39, Glossary-13
- Trace-event map file 6-119, 6-120, Glossary-5
- Trace-event tag 6-119, 6-120, Glossary-13
- Trace-event Value 8-39
- Tracepoint 3-11, 3-15, 3-25, 3-29, 3-56, 6-44, 6-93, 6-94, 6-111, 6-118, Glossary-5, Glossary-13
 - changing 8-35
 - clearing 6-125
 - condition on 3-23, 6-127, 6-168, 6-170, 8-38
 - deleting 6-128, 8-40
 - disabling 6-128, 8-37
 - displaying 6-167, 6-169
 - enabling 6-129, 8-37
 - hitting 6-170, 8-37
 - ignoring 3-23, 6-120, 6-130, 6-168, 6-170, 8-38, Glossary-7
 - named 3-23, 6-114, 6-120, 8-38
 - saving 3-23
 - setting 6-109, 6-120, 8-35
 - state 6-169, 8-37
 - temporary 8-37
- tracepoint** 3-15, 3-56, 6-119, 6-169

Tracepoint crossing count Glossary-3
 Tracepoint Dialog Box 8-16, 8-28, 8-35
 Tracing 3-7, 3-15, 3-56
translate-object-file 3-8, 6-31, 6-45, 6-47,
 6-50
 Translations
 dynamic library 6-47
 object filename 6-31, 6-45, 6-49, 6-109, 6-183
 Tutorial
 command-line 4-2
 Type definition
 printing 6-191, 6-192

U

u key 8-73
 UID 3-54
 ulimit 3-6
 Undefined symbols 6-109
undisplay 6-107, 6-108
 Uninteresting subprograms 3-24, 3-33, 3-35, 4-13, 6-7,
 6-75, 6-142, 6-145
 unsafe safety level 3-44, 5-2, 6-25, 6-33, 6-44, 6-45,
 6-72, 6-128, 6-129
up 3-33, 6-6, 6-155, 6-156, 6-159, 8-25, 8-73
 Up button 4-17
 Up button 8-25
 User 6-183
 User ID 3-54
 User interface
 command-line i-iii, 1-1, 3-36, 4-2, 6-2, 6-122,
 6-126, 6-152, 6-160, 6-194, 7-1, 8-1, 8-23,
 Glossary-2
 full-screen i-iii, 1-1, 3-36, 5-2, 6-2, 6-122, 6-126,
 6-160, 6-161, 6-194, 7-1, 7-2, 8-1,
 Glossary-6
 graphical i-iii, 2-1, 3-36, 3-45, 5-2, 5-3, 6-1, 6-45,
 6-152, 6-160, 8-1, Glossary-6

V

Value history 3-47, 4-14, 6-5, 6-68, 6-98, 6-99, 6-103,
 6-176, Glossary-13
 Variable
 assignment 3-27, 3-28
 convenience 3-46, 6-5, 6-7, 6-73, 6-99, 6-176,
 Glossary-3
 declaration 3-27
 global 3-25, 6-156
 local 3-25, 3-32, 3-33, 3-34, 6-5, 6-104, 8-18, 8-82,

8-97
 predefined convenience 3-34, 6-5, 6-7, 6-168,
 6-169, 6-170, 6-171, 6-172, 6-173, 6-177,
 6-193
 printing 6-189, 6-192
 register 3-25, 6-5
 static 3-25, 6-5
 Vector registers 6-16
 verify safety level 5-2, 6-25, 6-26, 6-33, 6-45, 6-54,
 6-72, 6-128, 6-129
 Version
 NightView 1-2, 5-2, 8-22
 vi editor 6-78, 7-2
 View menu 8-6
 Virtual address space 6-179
 Virtual memory 6-179
 VISUAL environment variable 7-2

W

wait 6-57
 Warning Dialog Box 4-28, 8-3, 8-6, 8-10, 8-11, 8-25
 Warnings 3-43
 warp Glossary-13
 Watchpoint
 changing 8-35
 commands on 3-23
 condition on 3-16, 3-23, 6-168, 6-174, 8-38
 deleting 6-128, 8-40
 disabling 8-37
 displaying 6-167, 6-173
 enabling 3-17, 6-129, 8-37
 hitting 6-174, 8-37
 ignoring 3-23, 6-130, 6-168, 6-174, 8-38,
 Glossary-7
 named 3-23, 6-134
 saving 3-23
 setting 6-133, 8-35
 state 6-173, 8-37
 temporary 8-37
watchpoint 6-133
 Watchpoint Dialog Box 8-16, 8-17, 8-28, 8-35
what is (info what is) 6-191, B-13
 Wildcard pattern 6-23, 6-30, 6-50, 6-52, 6-190, 8-13,
 8-34, 8-68
 Wildcard pattern examples 6-23
 wildcard_pattern 6-190
 Window
 Binary Viewer 6-105
 Data 3-24, 4-23, 4-25, 4-28, 6-104, 8-29, 8-78
 Debug 2-3, 4-15, 8-5, 8-30, 8-34, 8-35
 Dialogue 2-3

Help 2-1, 2-3, 2-4, 8-2, 8-3, 8-122, Glossary-7
iconifying 8-2
Monitor 3-36, 6-124, 7-2
print 8-6

X

X 8-1
x 6-7, 6-24, 6-99, 6-106, 6-168, 6-169, 6-170, 6-171,
6-172, 6-173, 6-193
-x option 5-3
X Window System 3-44, 8-1
X11 Display and CUDA 6-83, 8-48
xl (translate-object-file) 6-31, B-2
-xrm option 5-3