



NightView User's Guide

Version 7.1.1

(NightStar LX)

Copyright 2008 by Concurrent Computer Corporation. All rights reserved. This publication or any part thereof is intended for use with Concurrent products by Concurrent 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 Computer Corporation 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 Computer Corporation, 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 Qt's Assistant from Trolltech.

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 tutorials, see Chapter 4 [Tutorial] on page 4-1 and Chapter 5 [Tutorial - GUI] on page 5-1.

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 7 [Command-Line Interface] on page 7-1.

The next chapter describes a simple full-screen interface to NightView. See Chapter 8 [Simple Full-Screen Interface] on page 8-1.

The next chapter describes the graphical user interface for NightView. See Chapter 9 [Graphical User Interface] on page 9-1.

This manual also contains several appendixes that may not be of interest to all users, such as an implementation overview. A glossary of terms related to NightView and a quick reference guide are also provided.

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:

0897008	<i>NightStar LX Installation Guide</i>
0897008	<i>NightStar LX Tutorial</i>
0897398	<i>NightTrace™ User's Guide</i>
0897465	<i>NightProbe™ User's Guide</i>
0897515	<i>NightTune™ User's Guide</i>

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	3-1
Accessing Files	3-1
Programs and Processes	3-2
Multiple Processes	3-2
Families	3-2
Attaching	3-3
Detaching	3-3
Core Files	3-4
Qualifiers	3-4
Dialogues	3-4
Dialogue I/O	3-5
Real-Time Debugging	3-6
Remote Dialogues	3-6
Remote File Access	3-7
ReadyToDebug	3-9
Finding Your Program	3-9
Controlling Your Program	3-9
Eventpoints	3-9
Breakpoints	3-11
Monitorpoints	3-11
Patchpoints	3-12
Tracepoints	3-12
Heappoints	3-12
Watchpoints	3-13
Signals	3-14

Restarting a Program	3-16
Restart Mechanism	3-16
Restart Information	3-17
Restart Macros	3-18
Exited and Terminated Processes	3-18
Process States	3-18
Operations While the Process Is Executing	3-19
Examining Your Program	3-20
Expression Evaluation	3-20
C Expressions	3-21
C++ Expressions	3-21
Fortran Expressions	3-22
Overloading	3-22
Program Counter	3-23
Context	3-23
Scope	3-23
Stack	3-24
Current Frame	3-24
Registers	3-24
Inline Subprograms	3-25
Interesting Subprograms	3-25
Monitor Window	3-26
Debugging the Heap	3-28
Levels and Common Errors	3-29
Fences	3-30
Hardware Overrun Protection	3-30
Retained Free Blocks	3-31
Heap Check	3-32
Leak Detection	3-32
Errors	3-33
Command Streams	3-34
Interrupting the Debugger	3-34
Macros	3-35
Convenience Variables	3-35
Logging	3-36
Value History	3-36
Command History	3-36
Initialization Files	3-36
Optimization	3-37
Multithreaded Programs	3-37
Limitations and Warnings	3-38
Setuid Programs	3-38
Attach Permissions	3-38
NightTrace Analysis	3-38
Memory Mapped I/O	3-38
Blocking Interrupts	3-39
Debugging with Shared Libraries	3-39

Chapter 4 Tutorial

About the Tutorial	4-1
Creating a Program	4-2
Starting NightView	4-3

Getting General and Error Help	4-4
Starting Your Program	4-6
Debugging All Child Processes	4-7
Handling Signals	4-7
Listing the Source	4-8
Setting the First Breakpoints	4-9
Listing a Breakpoint	4-10
Continuing Execution	4-10
Not Entering Functions	4-10
Entering Input	4-11
Creating Families	4-12
Continuing Execution Again	4-13
Creating Families Again	4-14
Catching up the Child Process	4-14
Verifying Data Values	4-15
Entering Functions	4-16
Examining the Stack Frames	4-17
Moving in the Stack Frames	4-18
Verifying Data Values in Other Stack Frames	4-19
Returning to a Stack Frame	4-19
Resuming Execution	4-20
Setting the Default Qualifier	4-20
Removing a Breakpoint	4-21
Setting Conditional Breakpoints	4-21
Attaching an Ignore Count to a Breakpoint	4-22
Attaching Commands to a Breakpoint	4-23
Automatically Printing Variables	4-24
Watching Inter-Process Communication	4-24
Patching Your Program	4-25
Disabling a Breakpoint	4-26
Examining Eventpoints	4-27
Continuing to Completion	4-29
Leaving the Debugger	4-30

Chapter 5 Tutorial - GUI

About the Tutorial - GUI	5-1
Creating a Program - GUI	5-2
Starting NightView - GUI	5-3
Getting General and Error Help - GUI	5-4
Starting Your Program - GUI	5-6
Debugging All Child Processes - GUI	5-7
Handling Signals - GUI	5-8
Setting the First Breakpoints - GUI	5-9
Continuing Execution - GUI	5-10
Not Entering Functions - GUI	5-10
Entering Input - GUI	5-11
Continuing Execution Again - GUI	5-11
Catching up the Child Process - GUI	5-13
Verifying Data Values - GUI	5-14
Listing the Source - GUI	5-14
Entering Functions - GUI	5-15
Examining the Stack Frames - GUI	5-16

Moving in the Stack Frames - GUI	5-17
Verifying Data Values in Other Stack Frames - GUI	5-18
Returning to a Stack Frame - GUI	5-19
Resuming Execution - GUI	5-19
Removing a Breakpoint - GUI	5-20
Setting Conditional Breakpoints - GUI	5-20
Attaching an Ignore Count to a Breakpoint - GUI	5-21
Attaching Commands to a Breakpoint - GUI	5-22
Automatically Printing Variables - GUI	5-23
Watching Inter-Process Communication - GUI	5-23
Patching Your Program - GUI	5-25
Disabling a Breakpoint - GUI	5-26
Examining Eventpoints - GUI	5-26
Continuing to Completion - GUI	5-27
Leaving the Debugger - GUI	5-29

Chapter 6 Invoking NightView

Chapter 7 Command-Line Interface

Command Syntax	7-1
Selecting Overloaded Entities	7-2
Special Expression Syntax	7-3
Predefined Convenience Variables	7-5
IA-32 Registers	7-6
AMD64 Registers	7-9
Location Specifiers	7-13
Qualifier Specifiers	7-14
Eventpoint Specifiers	7-16
Regular Expressions	7-16
Wildcard Patterns	7-18
Repeating Commands	7-19
Replying to Debugger Questions	7-20
Controlling the Debugger	7-20
Quitting NightView	7-21
quit	7-21
Managing Dialogues	7-22
login	7-22
debug	7-24
nodebug	7-25
set-debug-file-directory	7-25
translate-object-file	7-26
logout	7-28
on dialogue	7-28
apply on dialogue	7-30
Dialogue Input and Output	7-30
!	7-31
set-show	7-32
show	7-33
Managing Processes	7-34
run	7-34
rerun	7-34
set-notify	7-35

notify	7-36
attach	7-36
detach	7-37
kill	7-37
symbol-file	7-38
core-file	7-38
exec-file	7-40
on program	7-41
apply on program	7-43
on restart	7-44
checkpoint	7-45
family	7-45
set-children	7-47
set-exit	7-47
wait	7-48
mreserve	7-49
Heap Debugging	7-50
heapdebug	7-50
Setting Modes	7-56
set-log	7-56
set-language	7-56
set-qualifier	7-57
set-history	7-58
set-limits	7-58
set-prompt	7-59
set-terminator	7-60
set-safety	7-60
set-restart	7-61
set-local	7-61
set-patch-area-size	7-62
interest	7-63
set-auto-frame	7-65
set-overload	7-66
set-search	7-66
set-editor	7-67
set-preallocate	7-67
set-resume	7-68
set-download	7-68
set-disassembly	7-69
Debugger Environment Control	7-71
cd	7-71
pwd	7-71
Source Files	7-72
Viewing and Editing Source Files	7-72
list	7-72
directory	7-74
edit	7-75
Searching	7-76
forward-search	7-76
reverse-search	7-76
Source Line Decorations	7-78
Examining and Modifying	7-81
backtrace	7-81
print	7-81

set	7-83
x	7-84
output	7-87
echo	7-87
data-display	7-88
display	7-88
undisplay	7-90
redisplay	7-90
printf	7-91
load	7-91
Manipulating Eventpoints	7-92
Eventpoint Modifiers	7-94
name	7-94
breakpoint	7-95
patchpoint	7-96
set-trace	7-98
tracepoint	7-99
monitorpoint	7-101
heappoint	7-102
mcontrol	7-103
clear	7-104
commands	7-105
condition	7-106
delete	7-107
disable	7-108
enable	7-108
ignore	7-109
tbreak	7-110
tpatch	7-111
watchpoint	7-112
Controlling Execution	7-113
continue	7-114
resume	7-115
step	7-116
next	7-117
stepi	7-118
nexti	7-119
finish	7-119
stop	7-120
jump	7-120
signal	7-121
handle	7-122
Selecting Context	7-124
frame	7-124
up	7-125
down	7-126
select-context	7-126
Miscellaneous Commands	7-127
help	7-127
refresh	7-128
shell	7-128
source	7-129
delay	7-129
Info Commands	7-130

Status Information	7-131
info log	7-131
info eventpoint.	7-131
info breakpoint	7-132
info tracepoint	7-133
info patchpoint.	7-134
info monitorpoint.	7-135
info heappoint	7-136
info watchpoint	7-137
info frame	7-138
info directories.	7-139
info convenience	7-139
info display	7-139
info history	7-140
info limits	7-140
info registers	7-140
info signal	7-141
info process	7-141
info memory	7-142
info dialogue	7-146
info family	7-146
info name.	7-146
info on dialogue.	7-147
info on program.	7-147
info on restart	7-147
info threads	7-148
heapcheck	7-148
Symbol Table Information	7-149
info args.	7-149
info locals	7-149
info variables.	7-150
info address	7-150
info sources	7-151
info functions.	7-151
info types.	7-151
info whatis.	7-152
info representation.	7-152
info declaration	7-152
info files.	7-153
info line	7-153
Defining and Using Macros.	7-153
define	7-154
Referencing Macros	7-156
info macros.	7-158

Chapter 8 Simple Full-Screen Interface

Using the Simple Full-Screen Interface.	8-1
Editing Commands in the Simple Full-Screen Interface	8-2
Monitor Window - Simple Full-Screen	8-2

Chapter 9 Graphical User Interface

NightView GUI Concepts	9-1
GUI Online Help	9-1
Context-Sensitive Help	9-2
Help Buttons	9-2
Help Command	9-2
Dialogues and Dialog Boxes.	9-2
Context Menu	9-3
Current Process	9-3
GUI Configuration	9-3
Main Window.	9-4
Menu bar	9-4
File Menu.	9-4
View Menu.	9-5
Shell Menu.	9-8
Process Menu.	9-8
Source Menu	9-10
Eventpoint Menu	9-12
Data Menu	9-14
Tools Menu	9-16
Help Menu	9-17
Toolbars.	9-19
Command Toolbar	9-19
Process Toolbar	9-19
Eventpoint Toolbar	9-21
Value Toolbar.	9-22
Source Display Toolbar	9-23
Status Bar	9-23
List of Shortcuts	9-25
Main Window Dialog Boxes.	9-25
Run Program in Shell Dialog Box	9-25
Attach Dialog Box	9-26
Source Selection Dialog Box.	9-26
File Selection Dialog Box	9-27
Eventpoint Dialog Boxes.	9-27
Debug Heap Dialog Box	9-32
Remote Login Dialog Box.	9-33
Remote Login General Page.	9-33
Remote Login Advanced Page	9-34
Remote Login Action Buttons	9-35
Preferences Dialog Box.	9-36
Preferences General Page.	9-36
Safety.	9-36
Automatically Resume On	9-37
Searching.	9-37
Data Panel	9-37
Display Limits.	9-37
Source Panel Keystrokes.	9-37
Preferences Appearance Page.	9-38
Source Display	9-38
Disassembly	9-38
Preferences Fonts Page.	9-39
Global NightStar Fonts	9-40

My NightView Fonts	9-41
Effective NightView Fonts	9-41
NightStar Global Fonts Dialog	9-42
Changes Fonts For...	9-43
Apply Fonts To...	9-43
Set Default Fonts	
Set Panel Fonts	9-43
Save & Close	9-43
Save.	9-43
Cancel	9-44
Help.	9-44
Preferences Advanced Page	9-44
Source Editor	9-45
Remote Object File Cache	9-45
Eventpoint Memory Preallocation	9-45
Restart	9-45
Value History	9-45
Expression Evaluation Automatic Overloading	9-46
Restore Defaults	9-46
Process Settings Dialog Box	9-46
Process Settings General Page	9-46
Debug Children.	9-46
Stop Before Exiting.	9-46
Expression Language	9-46
Program.	9-46
Process Settings Interest Page	9-47
Process Settings Signals Page	9-47
Rename Page Dialog Box	9-47
Print Dialog Box	9-47
List Location Dialog Box	9-47
Eventpoint Panel Update Interval Dialog Box	9-47
Panels.	9-48
Find Bar	9-48
Source Panel	9-48
Source Panel Target Line	9-49
Source Panel Expression Tooltip.	9-49
Source Panel Context Menu	9-50
Source Panel Tracking.	9-54
Source Panel Keystrokes.	9-54
Shell Panel	9-55
Message Panel	9-56
Eventpoint Panel.	9-57
Context Panel	9-59
Locals Panel	9-59
Monitor Panel.	9-59
Data Panel.	9-60
Monitor Bar.	9-60
Data Items	9-60
Expression Data Item.	9-61
Local Variables Data Item	9-61
Registers Data Item	9-62
Stack Data Item	9-62
Threads Data Item	9-62
Processes Data Item	9-62

Shells Data Item	9-62
Heap Information Data Item.	9-62
Heap Errors Data Item	9-63
Leak Sets / Still Allocated Sets Data Items	9-64
Block Data Item	9-64
Monitorpoint Values Data Item	9-65
Data Panel Context Menu	9-65
Data Panel Dialog Boxes.	9-72
Data Panel Item Dialog Box.	9-72
Data Panel Add Expression	9-72
Data Panel Add Heap Errors	9-73
Data Panel Add Heap Information.	9-73
Data Panel Add Heap Leaks.	9-73
Data Panel Add Local Variables.	9-73
Data Panel Add Monitorpoints.	9-74
Data Panel Add Processes	9-74
Data Panel Add Registers.	9-74
Data Panel Add Shells	9-74
Data Panel Add Still Allocated Blocks	9-74
Data Panel Add Stack.	9-75
Data Panel Add Threads.	9-75
Data Panel Call Stack Frames	9-75
Data Panel Edit Expression	9-75
Data Panel Expand Tree.	9-75
Data Panel Describe	9-76
Data Panel Load Layout.	9-76
Data Panel Pointer Array Dimension	9-76
Data Panel Save Layout	9-76
Data Panel Save Snapshot	9-77
Data Panel Subscript Array	9-77
Data Panel Linked List Expression Dialog	9-78
Data Panel Condition Filter Expression Dialog	9-78
Monitorpoint Update Interval Dialog Box.	9-80
Help Window	9-80

Appendix A NightStar LX Licensing

License Keys	A-1
License Requests	A-2
License Server	A-2
License Reports	A-3
Firewall Configuration for Floating Licenses	A-3
Serving Licenses with a Firewall	A-3
Running NightStar LX Tools with a Firewall.	A-5
License Support	A-6

Appendix B Kernel Dependencies

Advantages for NightView.	B-1
Advantages for NightTrace	B-2
Advantages for NightProbe	B-2
Advantages for NightTune.	B-3
Frequency Based Scheduler.	B-3

PCI Bar File System	B-3
---------------------------	-----

Appendix C Summary of Commands

Appendix D Quick Reference Guide

Invoking NightView	D-1
Controlling the Debugger	D-1
Quitting NightView	D-1
Managing Dialogues	D-1
Dialogue Input and Output	D-2
Managing Processes	D-2
Heap Debugging	D-3
Setting Modes	D-4
Debugger Environment Control	D-5
Source Files	D-5
Viewing and Editing Source Files	D-5
Searching	D-6
Examining and Modifying	D-6
Manipulating Eventpoints	D-7
Controlling Execution	D-8
Selecting Context	D-9
Miscellaneous Commands	D-9
Info Commands	D-10
Status Information	D-10
Symbol Table Information	D-11
Defining and Using Macros	D-11

Appendix E Implementation Overview

Appendix F Tutorial Files

C Files	F-1
msg.h	F-1
main.c	F-1
parent.c	F-2
child.c	F-2
Fortran Files	F-3
main.f	F-3
parent.f	F-4
child.f	F-4
ftint.c	F-5

Appendix G Reporting Bugs

Glossary

Index

Illustrations

Figure 9-1. Preferences Appearance Page	9-38
Figure 9-2. Preferences Fonts Page	9-40
Figure 9-3. NightStar Global Fonts Dialog	9-42
Figure 9-4. Preferences Advanced Page	9-44

Tables

Table 3-1. Eventpoint Summary	3-10
Table 7-1. Special '\$' Constructs	7-3
Table 7-2. Predefined Convenience Variables	7-5
Table 7-3. IA-32 Registers	7-6
Table 7-4. AMD64 Registers	7-10
Table 7-5. Regular Expressions	7-16
Table 7-6. Wildcard Patterns	7-18
Table 7-7. Source Line Decorations	7-78
Table 7-8. Eventpoint Commands	7-93

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 [Tutorial] 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 7-127).

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 3-4). 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 7-34) 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 debugger - Version 7.1, linked Fri Jun  8 10:24:51 EDT 2007
Copyright (C) 2007, Concurrent Computer Corporation

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 7-31.

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 7-72 and “breakpoint” on page 7-95.

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 7-81).

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 7-114, “backtrace” on page 7-81, “print” on page 7-81, and “quit” on page 7-21.

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 5 [Tutorial - GUI] on page 5-1.

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 9-17. 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 9-80. 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 9-55. 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 9-4.

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 9-55.

The message panel (see “Message Panel” on page 9-56) 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 9-23. 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 9-48.

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 9-19) 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 9-17. 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 9-80.

Next you will read about the **list** command. Click on the Help menu or use the H mnemonic. See "Help Menu" on page 9-17. 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 9-54. 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 7-95 and “Source Line Decorations” on page 7-78. 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 9-19.

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 9-59. 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 9-59.

Now finish running the program. Click on the **Resume** button. See “Process Toolbar” on page 9-19.

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 9-4. 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.

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 7-26. 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 7-37. 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 7-47.

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 7-25.

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 7-45.

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 7-36) or with the Attach Dialog Box (see “Attach Dialog Box” on page 9-26.)

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 9-8) 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 7-128) 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 3-38.

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 3-18. 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.

NightView typically uses some memory in the debugged process. If you detach and re-attach repeatedly, NightView will eventually be unable to find memory where it needs it in the process. See Appendix E [Implementation Overview] on page E-1. See also “set-patch-area-size” on page 7-62.

NOTE

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 7-38). 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 7-1 and “Qualifier Specifiers” on page 7-14.

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.

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 7-24 and “nodebug” on page 7-25.

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.

You may create additional dialogues at any time (see “login” on page 7-22). 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 7-31) or the `run` command (see “run” on page 7-34). 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 9-55) 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 7-32), and you can review the log with the `show` command (see “show” on page 7-33). 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 3-11.

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 3-12.

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, the NightView user interface runs on the local system, while another process, named NightView.p, runs on the remote system to access and control the processes you are debugging. The following activities are performed on the local system in this case:

- 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 7-22. See “Remote Login Dialog Box” on page 9-33.

Note that the parameters you specify for the remote dialogue will be inherited by processes you execute within that 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 7-38. Also, see “translate-object-file” on page 7-26.

Creating a new dialogue involves logging into a system (see “login” on page 7-22) 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 7-26 and “set-download” on page 7-68.

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 7-74.

NightView always interprets **exec-file** filenames and **core-file** filenames relative to the target system. See “exec-file” on page 7-40 and “core-file” on page 7-38.

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 7-68) or with the Preferences Advanced Page in the graphical user interface [see “Preferences Advanced Page” on page 9-44].)

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 7-38).

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 7-26.

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 7-38)
- **translate-object-file** (see “translate-object-file” on page 7-26)
- **directory** (see “directory” on page 7-74)
- **source** (see “source” on page 7-129)
- **load** (see “load” on page 7-91)
- **list** (see “list” on page 7-72)

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 7-68) 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 7-40.

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, and heappoints. 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 into your process. A watchpoint is not an inserted eventpoint. 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). 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 7-106. However, a condition on a watchpoint is evaluated in the debugger. For watchpoints, the language of the expression is determined by your language setting. See “set-language” on page 7-56. Because 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 7-109. Like conditions, the code to implement ignore counts is patched into the program for inserted eventpoints, so the program will execute at nearly full speed until the ignore count reaches zero. An ignore count on a watchpoint is evaluated in the debugger.

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 7-92.

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

Table 3-1. Eventpoint Summary

	Action	Code is inserted	May have commands
patchpoint	execute an expression or modify the flow of the program	X	
tracepoint	record an event when execution reaches the tracepoint	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. Watchpoint conditions and ignore counts are evaluated in the debugger. Watchpoints may be enabled and disabled only while the process is stopped.

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 7-95.

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 7-105.

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 3-26). You can set a monitorpoint using the **monitorpoint** command. See “monitorpoint” on page 7-101.

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 7-96) 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 7-91) 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 7-99) 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)** and the *NightTrace User's Guide*.

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 7-50) or from the **PROCESS** menu (see “Process Menu” on page 9-8). Heappoints are set with the **heappoint** command (see “heappoint” on page 7-102).

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 7-112.

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, `$is`, that is useful in watchpoint conditional expressions. See “Convenience Variables” on page 3-35. `$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 3-16.

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.

Because watchpoints are not inserted eventpoints, 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.

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 7-122).

In addition, you may use the debugger to explicitly send a signal to a process (see “signal” on page 7-121). 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 7-113). 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 7-122). 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.

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 7-34.

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 7-115) inside an **on program** command (see “on program” on page 7-41), 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 7-47. 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 7-44. 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 3-34.

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 7-45.

It is not possible to turn off checkpoints. However, you can control whether restart information is applied. See “set-restart” on page 7-61.

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 7-147. 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 7-129.

Restart Information

This section describes the restart information saved during a checkpoint.

- Any memory reservations made with the **mreserve** command. See “mreserve” on page 7-49.
- 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 7-74.
- Child disposition. See “set-children” on page 7-47.
- Signal disposition. See “handle” on page 7-122.
- Display list. See “display” on page 7-88.
- Symbol file. See “symbol-file” on page 7-38.
- Default language. See “set-language” on page 7-56.
- Whether or not the process will stop before exiting. See “set-exit” on page 7-47.
- The interest level threshold, the interest level for **inline**, **justlines**, and **nodebug**, and any explicit interest levels for subprograms. See “interest” on page 7-63.
- Information to reproduce the items in the data panel. See “Data Panel” on page 9-60. See “data-display” on page 7-88.

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 7-153. For example, you could define `restart_begin_hook` to be **echo** to disable the **on program** processing. See “on program” on page 7-41.

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 7-47.) 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 3-11. See “Tracepoints” on page 3-12. 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 7-67. A watchpoint may be enabled or disabled only when the process is stopped. See “Watchpoints” on page 3-13.

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 3-20. 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 3-23). 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 7-3. 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 7-92.

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 3-13.

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 7-67. 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 3-11.

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.

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 7-5.

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 7-3.

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 3-23). 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 7-56).

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.

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.

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 3-21.

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 3-22.

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.

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.

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 7-56. 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 7-2. Overloaded names are supported in language expressions (see “Expression Evaluation” on page 3-20) and location specifiers (see “Location Specifiers” on page 7-13), and the same syntax is used for both.

The **set-overload** command (see “set-overload” on page 7-66) 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 3-23) 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 3-24). 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 7-3) 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 7-3) 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 7-81).

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 7-126.

Frames corresponding to uninteresting subprograms are not numbered and they are not shown in a backtrace. See “Interesting Subprograms” on page 3-25.

Current Frame

When a process stops, the current frame is initially the stack frame associated with the most recently called routine (where `$pc` points). This frame contains the local variables for that routine, and these variables may be referenced in expressions you evaluate. Each frame also contains the return address indicating the specific point in the older routine where the `$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 7-125) or the **frame** command (“frame” on page 7-124) 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 `$pc` 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 7-3.

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 3-37 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 3-24.

You can use single step commands to step into inline subprograms, to step over them, or to finish them. See “step” on page 7-116, “next” on page 7-117, and “finish” on page 7-119.

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 7-13.

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 3-25. 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 3-20.

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 7-116. A stack walk-back does not display frames corresponding to uninteresting subprograms. See “Stack” on page 3-24.

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 7-63.

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 3-11). When you set a monitorpoint (see “monitorpoint” on page 7-101), 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 7-103) 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 3-34.

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 8 [Simple Full-Screen Interface] on page 8-1) or the graphical user interface (see Chapter 9 [Graphical User Interface] on page 9-1) 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 8-2. See “Monitorpoint Values Data Item” on page 9-65.

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 7-50), or with the **Debug Heap...** item in the **Process** menu. See "Process Menu" on page 9-8.

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 7-50 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 3-12.

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 7-50), 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_ouerrun
```

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
```

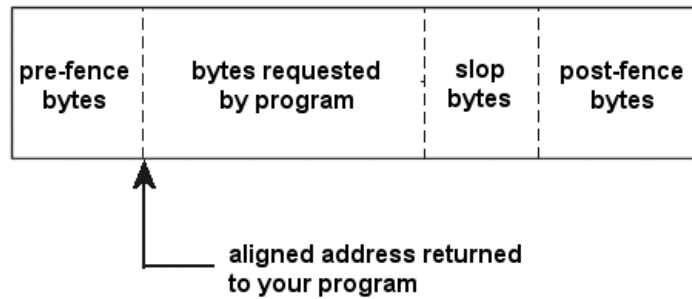
or clicking the **Uninitialized Field** button in the graphical user interface, configures the heap debugger to detect if the program reads from a block that it has failed to initialize.

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 7-55.

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 7-142) or in a data panel.

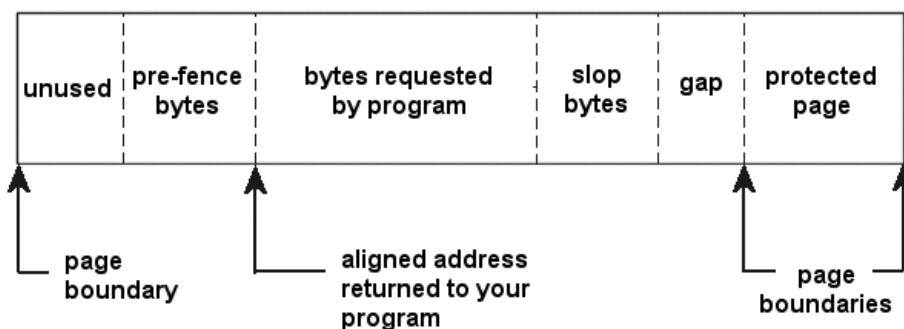
The advantage of hardware overrun protection is that it catches stray references immediately, which makes it easier to find bugs, and it catches both reads and writes. The main disadvantage of hardware overrun protection is that it uses a great deal of overhead memory in order to position each block. (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 dur-

ing only part of the process's execution so that only particular blocks are protected. See “Heappoints” on page 3-12.

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 7-55. 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 3-31), 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 “heapcheck” on page 7-148)
- 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 7-102)

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 7-50) or from the **Process** menu (see “Process Menu” on page 9-8), then it is possible to obtain a report of leaked heap blocks. See “info memory” on page 7-142 or “Leak Sets / Still Allocated Sets Data Items” on page 9-64.

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.

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 3-34.

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 7-60).

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 7-44.

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 9 [Graphical User Interface] on page 9-1. See “Process Toolbar” on page 9-19.

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 3-26.

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 7-153 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 7-61) 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 watchpoint, 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.

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 7-32).

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 7-56).

Value History

NightView keeps the results of the **print** command (see “print” on page 7-81) 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 7-140), or use previous history values in new expressions (see “Special Expression Syntax” on page 7-3).

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 8-2. See “Command Toolbar” on page 9-19.

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 3-34.

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 7-129).

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 6 [Invoking NightView] on page 6-1.

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. In this manual, those processes are referred to as thread processes. When this manual refers to a process, that means all the thread processes together. See `pthread_create(3)`.

When a process containing multiple threads stops, NightView shows you the thread process that caused the process to stop. To see other threads or thread processes, use the `select-context` command (see “select-context” on page 7-126).

It is important to note that NightView does **not** allow you to control the execution of a thread independently of the others in that process. When you resume execution (see “resume” on page 7-115), all thread processes are allowed to execute. If you issue a single-step command (see “step” on page 7-116), the selected thread will be stepped according to the command, but the other thread processes may also execute one or more

instructions — they are not restricted to stepping the current line or instruction.

Each time your process stops, NightView automatically sets the current context to the context of the thread that caused the process to stop. You may then use the **select-context** command to change the context.

Limitations and Warnings

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 7-128) 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.

NightTrace Analysis

The **tracepoint** command (see “tracepoint” on page 7-99) can be used to trace variables in a process. In order for tracepoints to actually log data, a NightTrace daemon must be running; any tracepoints logged before the daemon is started are lost.

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 7-40. 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.

This is the tutorial for the command-line version of NightView. NightView's command-line interface runs on all terminals. For more information about the command-line interface, see Chapter 7 [Command-Line Interface] on page 7-1. You may also be interested in the graphical-user-interface (GUI) version of this chapter in Chapter 5 [Tutorial - GUI] on page 5-1. There is a much shorter tutorial in Chapter 1 [A Quick Start] on page 1-1.

About the Tutorial

This tutorial shows only the most common debugger commands and features. It expects you to know the basics about 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 `stdout`, sleeps, sends signal `SIGUSR1` to the child, and loops. The child writes a message to `stdout` when it receives the signal.

Become familiar with the tutorial's source code; see Chapter F [Tutorial Files] on page F-1 or the files under the `/usr/lib/NightView/Tutorial` directory. The source files are:

C	Fortran	
<code>msg.h</code>	-	Defines constants
<code>main.c</code>	<code>main.f</code>	Forks a child and calls other routines
<code>parent.c</code>	<code>parent.f</code>	Sends signals to the child
<code>child.c</code>	<code>child.f</code>	Receives signals from the parent
-	<code>ftint.c</code>	Provides Fortran interfaces to system services

This tutorial takes at least two hours to do. Each section must be performed in order.

Exercises in this tutorial tell you to do things and ask you questions. Make the most of this tutorial and the manual; follow the steps below:

1. Look up the information.

2. Try to figure out the answer on your own.
3. Apply the provided solution. (**Warning:** Type the solutions exactly as they appear or your results may differ from those provided in later steps of the tutorial. Do not type anything until you see the words "you should enter" in the tutorial.)

You do not need to follow cross references in this tutorial **unless** you are explicitly told to read them.

This tutorial often displays process IDs. Your process IDs will probably differ from those shown. Also, the tutorial displays hexadecimal addresses. The addresses for your program may differ from those shown. Additionally, the line breaks in your output may differ from those shown because the lengths of displayed data items may vary.

The code produced when you create your program may vary slightly from the programs used to prepare this tutorial. In particular, the line shown as the return address from a sub-program may be different from what is shown here.

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

Creating a Program

NightView is mainly used with executables that contain debug information. To create such a program, compile source files with a particular option, and link edit them.

Exercise:

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

Solution:

You should enter:

```
$ mkdir nview
$ cd nview
```

Note: do not enter the `$`. It is part of the shell prompt.

Exercise:

Use the manual to find out what compiler option is necessary for debugging. (Hint: use the index.)

Solution:

From the index, **Compiling** has this information. The `-g` compiler option puts debug information into an executable.

Exercise:

Decide what language program you want to debug. Do *not* copy the source files from the `/usr/lib/NightView/Tutorial` directory, just compile and link these files. 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. What command or commands did you use?

Solution:

For C, you should enter:

```
$ cc -g -o msg /usr/lib/NightView/Tutorial/*.c
```

For Fortran, you should enter:

```
$ cc -c /usr/lib/NightView/Tutorial/ftint.c
$ f77 -g -o msg ftint.o \
    /usr/lib/NightView/Tutorial/*.f
```

You should now have a `msg` program with debug information in your `nview` directory. Note that for this tutorial, the source files should *not* be in this directory.

Starting NightView

You are ready to start up NightView *without* the graphical-user interface.

Exercise:

Read how to invoke the command-line interface of NightView. (You can find this information in the manual, on the `man` page, or by invoking `nview` with the `-help` option.) Start up the command-line interface of NightView.

Solution:

In the index, **Starting the debugger**, **Invoking the debugger**, and **nview, invoking** have this information. See Chapter 6 [Invoking NightView] on page 6-1. You should enter one of:

```
$ nview -nogui
$ nview -nog
```

Note that in this tutorial `msg` does not appear on the `nview` invocation line.

NightView responds with:

```
$ nview -nogui
NightView debugger - Version 5.1, linked Mon Jan 17
13:57:27 EST 2000
Copyright (C) 2000, Concurrent Computer Corporation
```

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

```
Use the 'run' command to run your program under the  
debugger  
(local)
```

These messages include NightView version information, copyright information, help information, and the command prompt, `(local)`. Your version number and date may differ. You will use online help later in this tutorial.

A dialogue contains a shell where you run shell commands and debug running programs. Each dialogue has a name; the default dialogue is `local`. The default qualifier is also `local`. The default command prompt is the qualifier in parentheses. For information about dialogues, see “Dialogues” on page 3-4. For information about qualifiers, see “Qualifiers” on page 3-4. For information about prompts, see “Command Syntax” on page 7-1.

In the command-line interface, NightView sometimes displays the command prompt before it completes its output display. You think NightView may have some undisplayed output.

Exercise:

To see the undisplayed output, wait a moment, press **Space**, then press **Return**.

NightView responds with:

```
/usr/lib/NightView-release/ReadyToDebug  
$ /usr/lib/NightView-release/ReadyToDebug  
$ (local)
```

NightView runs the `ReadyToDebug` program and your executable in the dialogue shell. 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 General and Error Help

This tutorial expects you to look up information in the NightView manual. For the command-line and simple screen interfaces, online help is available only for error messages. For general help you need to read the printed manual or consult the online help via NightView's graphical user interface. When this tutorial refers to another section of the manual, use one of those methods to read the section.

Exercise:

Try to use the non-existent "foo" command.

Solution:

You should enter:

```
(local) foo
```

Note: do not enter the `(local)`. It is part of the command prompt.

NightView responds with:

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

Exercise:

Now, invoke `help` without any arguments.

Solution:

You should enter one of:

```
(local) help
(local) he
```

NightView displays additional information about your most recent error and prints a new command prompt.

Note that `he` is not an official abbreviation for the `help` command; however, you may abbreviate NightView commands and some keywords to the shortest unambiguous prefix. For more information, see “Command Syntax” on page 7-1. You cannot abbreviate file names, symbolic names, or NightView construct names.

Exercise:

Once again, invoke `help` without any arguments.

Solution:

You should enter one of:

```
(local) help
(local) he
```

Note that NightView does not redisplay the extended error information; it assumes that you have already read that information. If there had been earlier errors, NightView would display help for the next most recent error now. However, there are no earlier errors, so NightView gives an error message indicating that.

NightView responds with the command prompt.

Starting Your Program

Most NightView commands 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 commands in this tutorial.

Exercise:

Read about the `run` command. Use it to start the `msg` program and have the program wait for debugging.

Solution:

You should enter one of:

```
(local) run ./msg
(local) ru ./msg
```

The preceding `./` is a safety precaution. Avoid accidentally debugging the wrong program by always providing some path information.

NightView responds with:

```
./msg
New process: local:15625    parent pid: 16428
Process local:15625 is executing /users/bob/nview/msg.
Reading symbols from /users/bob/nview/msg...done
Executable file set to
/users/bob/nview/msg
(local)
```

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

```
Program was dynamically linked.
Dynamic linking completed.
The file "file" does not contain symbolic debug
information,
only external symbols will be visible.
```

The long message *may* not appear.

NightView shows the process ID (PID) of the new process and its parent process, the path where your executable exists, and another local dialogue prompt. Your PIDs and the path where your executable exists will probably differ from those in this tutorial. For information about processes, see "Programs and Processes" on page 3-2.

Note that by appending an ampersand (&) to the `run` command, you could have started your program in the background of the dialogue 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 use the **run** command to rerun a program.

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 use the **set-children** command *before* `msg` forks the child process. Also, you have to issue the **set-children** command *after* the **run** command so the **set-children** command can be applied to existing processes.

Exercise:

Read about the **set-children** command. Issue the **set-children** command so that the child process in `msg` can be debugged.

Solution:

You should enter one of:

```
(local) set-children all
(local) set-c a
```

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.

Exercise:

Read about the **handle** command. Use it to adjust the default handling of the `SIGUSR1` signal so that the process does not stop.

Solution:

You should enter one of:

```
(local) handle SIGUSR1 nostop
(local) ha usr1 nos
```

NightView responds with:

```
Signal handling complete
(local)
```

Note: you had to issue the **handle** command *after* the **run** command so the **handle** command could be applied to existing processes.

Listing the Source

You probably want to look at the source files before debugging them.

Exercise:

Read about the **list** command. Notice all the syntax variations for this command, and use one of them to examine the source file where `main` is defined.

Solution:

You should enter one of:

```
(local) list main.c:1    (for the C program)
(local) l main.c:1      (for the C program)
(local) list main.f:1    (for the Fortran program)
(local) l main.f:1      (for the Fortran program)
        (local) list main
(local) l main
(local) list
(local) l
```

NightView responds by displaying ten numbered source lines. (You will see a different ten source lines depending on how you ran the **list** command.) Executable lines have an asterisk (*) source line decoration beside the line numbers. For more information about source line decorations, see “Source Line Decorations” on page 7-78.

The **list** command is repeatable. Press **Return**.

Now you see the next ten lines of the source file.

Keep pressing **Return** until you get an end of file message.

Exercise:

List the source file so the display is centered around line 16.

Solution:

You should enter one of:

```
(local) list main.c:16   (for the C program)
(local) l main.c:16     (for the C program)
(local) list main.f:16   (for the Fortran program)
(local) l main.f:16     (for the Fortran program)
        (local) list 16
(local) l 16
```

NightView responds by listing the lines.

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.

Exercise:

Read about the **breakpoint** command. Set a separate breakpoint to stop at each of the following places:

- The line that prompts for the number of signals to send
- The call to `child_routine`
- The comment before the call to `parent_routine`

Solution:

For the C program, this part of your debug session should look something like this:

```
(local) b 18
local:15625 Breakpoint 1 set at main.c:18
(local) b 25
local:15625 Breakpoint 2 set at main.c:25
(local) b 30
local:15625 Breakpoint 3 set at main.c:30
```

For the Fortran program, this part of your debug session should look something like this:

```
(local) b 19
local:15625 Breakpoint 1 set at main.f:19
(local) b 26
local:15625 Breakpoint 2 set at main.f:26
(local) b 28
local:15625 Breakpoint 3 set at main.f:28
```

Note that the preceding examples could have spelled out the **breakpoint** command. NightView gives each breakpoint an ordinal identification number beginning at 1. By default, breakpoints are set in the current list file, `main.c`, `main.f`, or `main.a` in this tutorial.

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.

Listing a Breakpoint

NightView changes the list display when you set a breakpoint.

Exercise:

Issue the **list** command that will relist the current lines.

Solution:

You should enter one of:

```
(local) list =  
(local) l =
```

NightView redisplay the ten lines you were viewing. Note that if you are displaying a line with a breakpoint on it, that line now has a B (for breakpoint) 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.

Exercise:

Read about the **continue** command. Use it to continue program execution up to the statement with the breakpoint.

Solution:

You should enter one of:

```
(local) continue  
(local) c
```

NightView displays the statement with the breakpoint. Note that the source line decoration is now a B=. The B still indicates a breakpoint, and the = indicates that execution is stopped there.

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.

Exercise:

Read about the `next` command. Use it to skip over the output statement (or function).

Solution:

You should enter one of:

```
(local) next
(local) n
```

The `msg` program writes the prompt "How many signals should the parent send the child?". NightView displays the next line. The `=` source line decoration shows that execution is stopped there.

Entering Input

You must respond to the `msg` program prompt "How many signals should the parent send the child?". By default, NightView interprets all input as debugger commands.

Exercise:

Assume that you want to send ten signals. See what problems arise when you simply enter the number ten.

Solution:

You should enter:

```
(local) 10
```

NightView responds with an error message.

Exercise:

Read about the `!` command. Use it to make NightView understand that the 10 is data for the `msg` program. (For information about Dialogue I/O, see "Dialogue I/O" on page 3-5 and "!" on page 7-31.)

Solution:

You should enter:

```
(local) !10
```

NightView responds with:

```
(local)
```

As described in “Starting NightView” on page 4-3, NightView sometimes has output that does not appear until you press **Return**.

Press **Space**, then press **Return** to see your input echoed.

NightView responds with:

```
10
(local)
```

Creating Families

Naming a process or process group has the following advantages over specifying PIDs.

- Mnemonic names are often easier to remember and type than numeric PIDs.
- You can group PIDs with a single name so that qualified NightView commands act only on the processes in the group.
- You can write generic NightView command files that use process names instead of specific PIDs.

In this tutorial, you will want to issue some NightView commands that pertain only to the parent process and others that pertain only to the child process.

Exercise:

Read about the **family** command. Use it to give the name `parent` to all processes that currently exist in your program. (There is only one process so far.)

Solution:

You should enter one of:

```
(local) family parent all
(local) fa parent all
```

Note that to name only the parent process, you had to issue this command *before* NightView executes the `fork` in the `msg` program. Note also that at this point, the `all` argument represents only one process, the parent process. Later you will see it represent multiple processes.

You will use the `parent` family name later in the tutorial.

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.

Exercise:

Continue executing the program so it stops at the next statement with a breakpoint.

Solution:

You should enter one of:

```
(local) continue
(local) c
```

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, the output below will not match what you see, but this is not significant to the remainder of the tutorial. Do not enter any extra commands.

For the C program, NightView responds with:

```
New process: local:13504      parent pid: 15625
#0 0x10002838 in main() at main.c line 20
20 < |      pid = fork();
local:15625: at Breakpoint 3, 0x1000284c in main() at
main.c line 31
31 B=|      parent_routine( pid, total_sig );
(local)
```

For the Fortran program, NightView responds with:

```
New process: local:13504      parent pid: 15625
#0 0x100038e4 in main() at main.f line 22
22 < |      pid = ftfork()
local:15625: at Breakpoint 3, 0x10003904 in main() at
main.f line 29
29 B=|      call parent_routine( pid )
(local)
```

The first few lines are from the child process. (In your output, the parent lines might appear first.) They show that you are currently calling `fork` (`ftfork` in the Fortran program). The `<` source line decoration indicates that this line made a subprogram call which has not yet returned. You may also see a `>` source line decoration, which indicates that control will return to, or within, this line. The subprogram that implements `fork` is hid-

den. NightView usually does not show you system library routines. See “Interesting Sub-programs” on page 3-25.

In this example, the child process has process ID 13504, and the parent process has process ID 15625. 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 line decoration in the parent process is now a B=. The B still indicates a breakpoint and the = indicates that execution is stopped there.

Creating Families Again

In this tutorial, you will want to issue some NightView commands that pertain only to the parent process and others that pertain only to the child process.

Exercise:

Use the `family` command to give the name `child` to only the newly forked child process.

Solution:

You should enter one of:

```
(local) family child all - parent
(local) fa child all - parent
```

At this time, the `all` argument consists of both the parent and child PIDs. In section “Creating Families” on page 4-12, you created the `parent` family so it consists of only the parent PID. Subtraction leaves only the child PID in the `child` family.

You will use the `child` family name later in the tutorial.

Note that to name only the child process, you had to issue this command *after* NightView executes the `fork` in the `msg` program.

Catching up the Child Process

To individually manipulate the parent and child processes, you must qualify your debugger commands.

Exercise:

Read about qualifiers. Get the child process to continue execution up to the breakpoint on the call to `child_routine` (line 25 in `main.c`, line 21 in `main.f`, and line 25 in `main.a`).

Solution:

You should enter one of:

```
(local) (child) continue
(local) (child) c
```

For the C program, NightView displays:

```
local:13504: at Breakpoint 5, 0x10002840 in main() at
main.c line 25
25 B=|          child_routine( total_sig );
(local)
```

For the Fortran program, NightView displays:

```
local:13504: at Breakpoint 4, 0x100038fc in main() at
main.f line 26
26 B=|          call child_routine()
(local)
```

This breakpoint in the child corresponds to breakpoint 2 in the parent. Inherited eventpoints get new identifiers. The order of eventpoint numbers in the child is unpredictable, so you might see a breakpoint number of 4, 5, or 6.

Note that you could have qualified the command with the child's process ID number instead of the child family name.

Verifying Data Values

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

Exercise:

Read about the `print` command. Use it to check that the `total_sig` variable has the value 10.

Solution:

You should enter one of:

```
(local) print total_sig
(local) p total_sig
```

NightView responds with:

```
Process local:15625:
$1: total_sig = 10
Process local:13504:
$2: total_sig = 10
```

By default, the 10 is printed 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 3-36.

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 3-37. In the Fortran program, `total_sig` was put in `COMMON` so you could consistently see its value in the tutorial.

NightView displays values for both processes because there are multiple processes in the default qualifier `local`.

Entering Functions

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

Exercise:

Read about the `step` command. Use it to simultaneously enter both routines.

Solution:

You should enter one of:

```
(local) step
(local) s
```

Note that if you had wanted to enter a routine in only one process, you would have had to qualify the `step` command. (For information about qualifiers, see "Qualifiers" on page 3-4.)

In all the following output descriptions, NightView displays the line you stepped to. The = source line decoration indicates that execution is stopped there.

For the C program, NightView displays:

```
#0 0x10002884 in child_routine(int total_sig = 10) at
child.c line 14
#0 0x10002944 in parent_routine(pid_t child_pid = 13504
, int total_sig = 10)
at parent.c line 11
14 = | signal( SIGUSR1, signal_handler );
11 = | int isec = 2;
(local)
```

Line 14 is from the child process. Line 11 is from the parent process.

For the Fortran program, NightView displays:

```
#0 0x1000393c in child_routine() at child.f line 17
```

```

17 = |          irecturn = ftsignal( sigusr1, signal_handler )
#0  0x10003a48  in parent_routine(INTEGER child_pid /
13504 / )
          at parent.f line 16
16 = |          do 10 sig_ct = 1, total_sig
(local)

```

Line 17 is from the child process. Line 16 is from the parent process.

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.

Examining the Stack Frames

It is often helpful to see how you got to a certain point in a program.

Exercise:

Read about the **backtrace** command. Use it to display the list of currently active stack frames.

Solution:

You should enter one of:

```

(local) backtrace
(local) bt

```

For the C program, NightView responds with:

```

Backtrace for process local:13504
#0  0x10002884  in child_routine(int total_sig = 10) at c
hild.c line 14
#1  0x10002848  in main() at main.c line 25
Backtrace for process local:15625
#0  0x10002944  in parent_routine(pid_t child_pid = 13504
, int total_sig = 10)
          at parent.c line 11
#1  0x10002854  in main() at main.c line 31
(local)

```

For the Fortran program, NightView responds:

```

Backtrace for process local:13504
#0  0x1000393c  in child_routine() at child.f line 17
#1  0x10003900  in main() at main.f line 26
Backtrace for process local:15625

```

```
#0  0x10003a48  in parent_routine(INTEGER child_pid /  
13504 / )  
                                at parent.f line 16  
#1  0x10003910  in main() at main.f line 29  
(local)
```

Note the order of the displayed lines may vary.

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`.

Exercise:

Read about the `up` command. Qualify the `up` command so the current stack frame of the parent process is `main`.

Solution:

You should enter:

```
(local) (parent) up
```

For the C program, NightView responds with:

```
Output for process local:15625  
#1  0x10002854  in main() at main.c line 31  
31 B<|          parent_routine( pid, total_sig );  
(local)
```

For the Fortran program, NightView responds with:

```
Output for process local:15625  
#1  0x10003910  in main() at main.f line 29  
29 B<|          call parent_routine( pid )  
(local)
```

The `<` source line decoration indicates that this line made a subprogram call which has not yet returned. You may also see a `>` source line decoration, which indicates that control will return to, or within, this line.

Note that you could have qualified the command with the parent's process ID number instead of the parent family name.

Verifying Data Values in Other Stack Frames

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

Exercise:

Qualify a **print** command so it displays the value of local variable `tracefile` in main only for the parent process.

Solution:

You should enter one of:

```
(local) (parent) print tracefile
(local) (parent) p tracefile
```

For the C program, NightView responds with:

```
$3: tracefile = 0x30003100 "msg_file"
(local)
```

For the Fortran program, NightView responds with:

```
$3: tracefile = "msg_file"
(local)
```

Note that you could have qualified the command with the parent's process ID number instead of the `parent` family name.

Returning to a Stack Frame

You want to return to `parent_routine`.

Exercise:

Read about the **down** command. Qualify the **down** command so the current stack frame of the parent process is `parent_routine`.

Solution:

You should enter one of:

```
(local) (parent) down
(local) (parent) do
```

For the C program, NightView responds with:

```
Output for process local:15625
```

```
#0 0x10002944 in parent_routine(pid_t child_pid =
13504, int total_sig = 10)
          at parent.c line 11
11 = |      int isec = 2;
(local)
```

For the Fortran program, NightView responds with:

```
Output for process local:15625
#0 0x10003a48 in parent_routine(INTEGER child_pid /
13504 / )
          at parent.f line 16
16 = |      do 10 sig_ct = 1, total_sig
(local)
```

Note: it is not meaningful to do a **down** without doing an **up** *first* (as you did in section “Moving in the Stack Frames” on page 4-18).

Resuming Execution

You 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. The **continue** command can do this, but it ties up the debugger's input mechanism while waiting for the process. You don't want to wait.

Exercise:

Read about the **resume** command. Qualify the **resume** command to resume execution of the *child* process without the waiting that occurs with the **continue** command.

Solution:

You should enter one of:

```
(local) (child) resume
(local) (child) res
```

Note that you could have qualified the command with the child's process ID number instead of the *child* family name.

Setting the Default Qualifier

As described in “Starting NightView” on page 4-3, the default qualifier is `local`, which means that unqualified commands affect all processes. It is cumbersome to keep qualifying commands that operate on a subset of these processes. The rest of the commands in this tutorial apply only to the parent process.

Exercise:

Read about the `set-qualifier` command. Use it to tell NightView that the default qualifier for the remaining commands is the family that consists of only the parent process.

Solution:

You should enter one of:

```
(local) set-qualifier parent
(local) set-q parent
```

NightView changes the prompt to your new qualifier, parent.

Removing a Breakpoint

Breakpoint 1 (set in “Setting the First Breakpoints” on page 4-9) is no longer needed.

Exercise:

Read about the `delete` command. Use it to remove breakpoint 1.

Solution:

You should enter one of:

```
(parent) delete 1
(parent) d 1
```

Setting Conditional Breakpoints

It is often useful to suspend execution conditionally.

Exercise:

Read about the `breakpoint` command. 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`.

Solution:

For the C program, you should enter one of:

```
(parent) breakpoint 16 if isec == total_sig
(parent) b 16 if isec == total_sig
```

For the Fortran program, you should enter one of:

```
(parent) breakpoint 17 if isec .eq. total_sig  
(parent) b 17 if isec .eq. total_sig
```

For the C program, NightView responds with:

```
local:15625 Breakpoint 7 set at parent.c:16
```

For the Fortran program, NightView responds with:

```
local:15625 Breakpoint 7 set at parent.f:17
```

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.

Exercise:

Read about the **ignore** command. Set a **breakpoint** command on the line that displays how long the parent is sleeping in `parent_routine`. NightView has a predefined name for the most-recently set breakpoint. For more information about this name, see "Eventpoint Specifiers" on page 7-16. Use this name on an **ignore** command on this line in `parent_routine`; ignore the next five iterations.

Solution:

For the C program, you should enter:

```
(parent) breakpoint 16  
(parent) ignore . 5
```

or

```
(parent) b 16  
(parent) ig . 5
```

For the Fortran program, you should enter:

```
(parent) breakpoint 17  
(parent) ignore . 5
```

or

```
(parent) b 17  
(parent) ig . 5
```

For the C program, NightView responds with:

```
local:15625 Breakpoint 8 set at parent.c:16
Will ignore next 5 crossings of Breakpoint #8 in
local:15625.
```

For the Fortran program, NightView responds with:

```
local:15625 Breakpoint 8 set at parent.f:17
Will ignore next 5 crossings of Breakpoint #8 in
local:15625.
```

Attaching Commands to a Breakpoint

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

Exercise:

Read about the **commands** command. Attach a command stream that prints out the value of `total_sig` only when you hit the breakpoint you set in the previous step. Note: use the NightView predefined name for the most-recently set breakpoint.

Solution:

You should enter one of:

```
(parent) commands .
(parent) com .
```

NightView responds with:

```
Type commands for when the breakpoints are hit, one per
line.
End with a line saying just "end".
>
```

You should enter:

```
> print total_sig
> end
```

or

```
> p total_sig
> end
```

Automatically Printing Variables

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

Exercise:

Read about the **display** command. Use a **display** command to display the value of the `sig_ct` variable.

Solution:

You should enter one of:

```
(parent) display sig_ct  
(parent) disp sig_ct
```

Note that this **display** command runs every time execution stops, and the **print** command from “Attaching Commands to a Breakpoint” on page 4-23 runs only when execution stops at a specific breakpoint.

Watching Inter-Process Communication

You already resumed the execution of the child process, so NightView gave you a prompt and did not wait for the child process.

Exercise:

Now continue execution for the parent process.

Solution:

You should enter one of:

```
(parent) continue  
(parent) c
```

NightView responds with something like the following:

```
1. Parent sleeping for 2 seconds  
2. Parent sleeping for 2 seconds  
Child got ordinal signal #1  
3. Parent sleeping for 2 seconds  
Process local:13504 received SIGUSR1  
Child got ordinal signal #2  
4. Parent sleeping for 2 seconds  
Process local:13504 received SIGUSR1  
Child got ordinal signal #3  
5. Parent sleeping for 2 seconds  
Process local:13504 received SIGUSR1  
Child got ordinal signal #4  
Process local:13504 received SIGUSR1  
Child got ordinal signal #5  
local:15625: at Breakpoint 8, 0x10002950 in parent_routine(
```

```

                                pid_t child_pid = 13504, int total_sig = 10)
                                at parent.c line 16
16 B=|          printf( "%d.  Parent sleeping for %d seconds\n", sig_ct, isec );
1: sig_ct = 6
(parent)

```

Note the order of the displayed lines may vary. For the Fortran program, NightView prints the argument or arguments to `parent_routine` differently.

Because of the **ignore** command 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 **handle** command is implicitly set to `print` and explicitly set to `nostop`. Two lines show where execution stopped; these lines will differ depending on your programming language. Another line shows the value of `sig_ct` because of the **display** command.

Note that the `print total_sig` output did not appear if NightView returned your prompt before the commands in the **commands** command stream completed their output.

Exercise:

To see the `print total_sig` output, enter a space *and* Return. (You do not need to do this if the output already appeared.)

WARNING

If you press Return *without* the space, you will repeat the **continue** command.)

NightView responds with the following:

```

$4: total_sig = 10
(parent)

```

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.

Exercise:

List the source file for the `parent_routine` so the display is centered around line 13.

Solution:

You should enter one of:

```
(parent) list parent.c:13    (for the C program)
```

```
(parent) 1 parent.c:13      (for the C program)
(parent) list parent.f:13   (for the Fortran program)
(parent) 1 parent.f:13     (for the Fortran program)
```

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`. Hint: In C the `%` operator, or in Fortran the `mod` function, may be useful.

Exercise:

Read about the `patchpoint` command. In the parent process, *on* the line that displays how long the parent is sleeping, patch in the assignment expression just described.

Solution:

For the C program, you should enter:

```
(parent) patchpoint at 16 eval isec = sig_ct % 3 + 1
```

For the Fortran program, you should enter:

```
(parent) patchpoint at 17 eval isec = mod( sig_ct, 3 ) + 1
```

For the C program, NightView responds with the following:

```
local:15625 Patchpoint 9 set at parent.c:16
```

For the Fortran program, NightView responds with the following:

```
local:15625 Patchpoint 9 set at parent.f:17
```

Disabling a Breakpoint

You want to run `msg` to completion without stopping at breakpoint 8.

Exercise:

Read about the `disable` command. Use it to disable breakpoint 8 (set in section “Attaching an Ignore Count to a Breakpoint” on page 4-22).

Solution:

You should enter one of:

```
(parent) disable 8
(parent) disa 8
```


Examining Eventpoints

An *eventpoint* is a generic term which includes breakpoints, patchpoints, monitorpoints, tracepoints, and heappoints. You want to examine the locations, associated commands, and statistics related to the eventpoints you have set in msg.

Exercise:

Read about the `info eventpoint` command. Use it to examine all eventpoints.

Solution:

You should enter one of:

```
(parent) (local) info eventpoint
(parent) (local) i ev
(parent) (all) info eventpoint
(parent) (all) i ev
```

For the C program, NightView responds with the following:

Eventpoints for process local:15625:

```
ID Typ Enb      Where
---
 2 B   Y  main.c:25
 3 B   Y  main.c:30
    #crossings=1 #hits=1
 7 B   Y  parent.c:16
    only if isec == total_sig
    #crossings=6
 8 B   N  parent.c:16
    #crossings=6 #hits=1
    commands:
      print total_sig
 9 P   Y  parent.c:16
    eval = isec = sig_ct % 3 + 1
```

Eventpoints for process local:13504:

```
ID Typ Enb      Where
---
 4 B   Y  main.c:18
    #crossings=1 #hits=1
 5 B   Y  main.c:25
    #crossings=1 #hits=1
 6 B   Y  main.c:30
(parent)
```

For the Fortran program, NightView responds with the following:

Eventpoints for process local:15625:

```

ID  Typ  Enb      Where
---  ---  ---  -----
  2  B    Y  main.f:26
  3  B    Y  main.f:28
      #crossings=1 #hits=1
  7  B    Y  parent.f:17
      only if isec .eq. total_sig
      #crossings=6
  8  B    N  parent.f:17
      #crossings=6 #hits=1
      commands:
          print total_sig
  9  P    Y  parent.f:17
      eval = isec = mod( sig_ct, 3 ) + 1

```

Eventpoints for process local:13504:

```

ID  Typ  Enb      Where
---  ---  ---  -----
  4  B    Y  main.f:26
      #crossings=1 #hits=1
  5  B    Y  main.f:28
  6  B    Y  main.f:19
      #crossings=1 #hits=1
(parent)

```

NightView displays all eventpoints for process local:15625 followed by the eventpoints for process local:13504.

Breakpoints 1, 2, and 3 were set in “Setting the First Breakpoints” on page 4-9. Breakpoint 1 has no entry because it was deleted in “Removing a Breakpoint” on page 4-21. Breakpoints 2 and 3 are still enabled. Breakpoint 3 has been crossed once and hit once. This means that its line has been executed once and stopped on once.

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 4-21 and is still enabled; note that NightView displays the condition on this breakpoint. This breakpoint has been crossed six times without being hit. This means that the line has been executed six times, but the condition has not been true yet.

Breakpoint 8 was set in “Attaching an Ignore Count to a Breakpoint” on page 4-22 and was disabled in “Disabling a Breakpoint” on page 4-26; note that NightView displays the commands (`print total_sig`) attached to this breakpoint. This breakpoint has been crossed six times and has been hit only once. This means that the line has been executed six times, but the **ignore** command has prevented the breakpoint from being hit more than once.

Patchpoint 9 was set in “Patching Your Program” on page 4-25 and is still enabled; note that NightView displays the expression associated with this patchpoint. This patchpoint has not been crossed or hit yet so these statistics are omitted from the display.

Continuing to Completion

There’s nothing else to look at, so you decide to run `msg` to completion.

Exercise:

Use the `continue` command to continue execution.

Solution:

You should enter one of:

```
(parent) continue
(parent) c
```

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, enter the command `up` until the debugger reports that the process is in `main`.

NightView responds with something like this:

```
6. Parent sleeping for 1 seconds
7. Parent sleeping for 2 seconds
Process local:13504 received SIGUSR1
Child got ordinal signal #6
8. Parent sleeping for 3 seconds
Process local:13504 received SIGUSR1
Child got ordinal signal #7
9. Parent sleeping for 1 seconds
Process local:13504 received SIGUSR1
Child got ordinal signal #8
10. Parent sleeping for 2 seconds
Process local:13504 received SIGUSR1
Child got ordinal signal #9
Process local:13504 received SIGUSR1
Child got ordinal signal #10
Process local:15625 is about to exit normally
Process local:13504 is about to exit normally
#1 0x1000285c in main() at main.c line 34
#1 0x1000285c in main() at main.c line 34
34 <>|      exit( 0 );
```

```
34 <>|      exit( 0 );  
--> Undisplayed items:  
   1: (print) sig_ct  
      (parent)
```

Note the order of the displayed lines may vary.

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. The lines that mention signal SIGUSR1 appear because the **handle** command is implicitly set to **print** and explicitly set to **nostop**.

The last two lines say that `sig_ct` is not displayed. This message appears because of the **display** command and because the `sig_ct` variable is not visible at this point in the parent process. For the Fortran program, depending on which compiler you used, `sig_ct` may still be available, so it may be displayed.

Leaving the Debugger

The tutorial is over.

Exercise:

Read about the **quit** command. Use it to leave the debugger.

Solution:

You should enter one of:

```
(parent) quit  
(parent) q
```

Neither process has completely exited, so NightView asks the following question:

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

Exercise:

Make the processes go away.

Solution:

You should respond:

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

NightView responds with:

```
You are now leaving NightView...
```

```
Process local:13504 exited normally  
Process local:15625 exited normally  
Dialogue local has exited.
```


This is the 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 9 [Graphical User Interface] on page 9-1. You may also be interested in the command-line version of this chapter in Chapter 4 [Tutorial] on page 4-1. There is a much shorter tutorial in Chapter 2 [A Quick Start - GUI] on page 2-1.

About the Tutorial - GUI

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 `stdout`, sleeps, sends signal `SIGUSR1` to the child, and loops. The child writes a message to `stdout` when it receives the signal.

Become familiar with the tutorial's source code; see Chapter F [Tutorial Files] on page F-1 or the files under the `/usr/lib/NightView/Tutorial` directory. The source files are:

C	Fortran	
<code>msg.h</code>	-	Defines constants
<code>main.c</code>	<code>main.f</code>	Forks a child and calls other routines
<code>parent.c</code>	<code>parent.f</code>	Sends signals to the child
<code>child.c</code>	<code>child.f</code>	Receives signals from the parent
-	<code>ftint.c</code>	Provides Fortran interfaces to system services

This tutorial takes at least two hours to do. Each section must be performed in order. If you do not have two hours, you may want to see Chapter 2 [A Quick Start - GUI] on page 2-1.

Exercises in this tutorial tell you to do things and ask you questions. Make the most of this tutorial and the manual; follow the steps below:

1. Look up the information.
2. Try to figure out the answer on your own.

3. Apply the provided solution. (**Warning:** Perform the solutions exactly as indicated, or your results may differ from those provided in later steps of the tutorial. Do not do anything until you see the words "you should" in the tutorial.)

You do not need to follow cross references in this tutorial **unless** you are explicitly told to read them.

Sometimes NightView displays a status so briefly that it seems to flash before being replaced by another status. This tutorial documents only the last status displayed.

This tutorial often displays process IDs. Your process IDs will probably differ from those shown. Also, the tutorial displays hexadecimal addresses. The addresses for your program may differ from those shown. Additionally, the line breaks in your output may differ from those shown because the lengths of displayed data items may vary.

The code produced when you create your program may vary slightly from the programs used to prepare this tutorial. In particular, the line shown as the return address from a sub-program may be different from what is shown here.

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

Some of the shortened commands that appear in this tutorial are not official abbreviations for NightView commands; however, you may abbreviate NightView commands and some keywords to the shortest unambiguous prefix. For more information, see "Command Syntax" on page 7-1. 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, to reduce confusion, use the mouse whenever possible during this tutorial. Use mouse button 1 when you are told to click, drag, and select.

Creating a Program - GUI

NightView is mainly used with executables that contain debug information. To create such a program, compile source files with a particular option, and link edit them.

Exercise:

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

Solution:

You should enter:

```
$ mkdir nview
$ cd nview
```

Note: do not enter the `$`. It is part of the shell prompt.

Exercise:

Use the manual to find out what compiler option is necessary for debugging. (Hint: use the index.)

Solution:

From the index, **Compiling** has this information. The **-g** compiler option puts debug information into an executable.

Exercise:

Decide what language program you want to debug. Do *not* copy the source files from the `/usr/lib/NightView/Tutorial` directory, just compile and link these files. 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. What command or commands did you use?

Solution:

For C, you should enter:

```
$ cc -g -o msg /usr/lib/NightView/Tutorial/*.c
```

For Fortran, you should enter:

```
$ cc -c /usr/lib/NightView/Tutorial/ftint.c
$ f77 -g -o msg ftint.o \
    /usr/lib/NightView/Tutorial/*.f
```

You should now have a `msg` program with debug information in your `nview` directory. Note that for this tutorial, the source files should *not* be in this directory.

Starting NightView - GUI

You are ready to start up NightView *with* the graphical user interface.

Exercise:

Read how to invoke the graphical user interface of NightView. (You can find this information in the manual, on the man page, or by invoking `nview` with the **-help** option.) Start up the graphical user interface of NightView.

Solution:

In the index, **Starting the debugger**, **Invoking the debugger**, and **nview, invoking** have this information. See Chapter 6 [Invoking NightView] on page 6-1. You should enter:

```
$ nview
```

Note that in this tutorial `msg` does not appear on the `nview` invocation line.

NightView responds by displaying the main window.

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

You communicate with a dialogue shell through a shell panel. A shell panel is used to type input to the shell and to see the shell output. For this tutorial we will create a shell panel on a separate page. Create a new 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**. Create a shell panel. Click on the **View** menu and select **New Shell Panel**. Each dialogue has a name; the default dialogue is `local`. The page now has a panel with the title `local shell`. For information about shell panels, see "Shell Panel" on page 9-55.

The shell panel displays:

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

NightView runs the **ReadyToDebug** program automatically as part of initialization. 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 General and Error Help - GUI

This tutorial expects you to look up information in the NightView manual. You may read the hard copy or the similar online manual. The online manual is accessible through the **Help** menu.

Exercise:

Try to use the non-existent "foo" command.

Solution:

Switch back to the first page by clicking on the tab labeled **Main**. The command toolbar is near the bottom of the window with a label **Command:** and a combo box where you can type commands. In the command toolbar, you should enter:

```
foo
```

and press **Return**.

NightView displays in the message panel:

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

Exercise:

Read about this error message.

Solution:

You should click on the **Help** menu and select **On Last Error**.

The help window displays additional information about your most recent error.

Read the information. Note that **Summary of Commands** appears highlighted.

Exercise:

Read about getting information about all NightView commands.

Solution:

In the help window, you should click on **Summary of Commands**.

The help window displays a list of NightView commands with each command highlighted. The vertical and horizontal scroll bars next to the help display let you examine the rest of the help text.

Exercise:

Read about the command toolbar.

Solution:

In the NightView window, you should click on the **Help** menu and select **On Context**.

NightView changes your pointer to a modified question mark.

Click on the menu bar.

NightView restores your pointer. The help window displays information about the menu bar.

When this tutorial asks you to read about buttons or panels, use this same procedure.

For now, you are finished using help.

Exercise:

Close the help window.

Solution:

In the help window, you should click on the **File** menu and select **Exit**. (The help window is running a separate program, so only that program will exit. NightView will still be running.)

The help window goes away.

This tutorial discusses the **Help** menu again in “Debugging All Child Processes - GUI” on page 5-7. For more information about help, see “GUI Online Help” on page 9-1.

Starting Your Program - GUI

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.

Exercise:

Start the `msg` program, and have it wait for debugging.

Solution:

You should switch back to the page with the shell panel by clicking on the tab labeled **Page 2**. You may need to click in the shell panel to give it the keyboard focus. In the shell panel, you should enter:

```
./msg
```

and press **Return**.

The preceding `./` is a safety precaution. Avoid accidentally debugging the wrong program by always providing some path information.

You should switch back to the other page by clicking on the tab labeled **Main**.

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: 15625 parent pid: 17882
Process local:15625 is executing /users/bob/nview/msg.
Reading symbols from /users/bob/nview/msg...done
Executable file set to
/users/bob/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. 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:15625`, 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 7-78. 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:15625`, 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 9-59.

Note that by appending an ampersand (&) to the `./msg`, you could have started your program in the background of the dialogue 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 - GUI

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.

Exercise:

Read about the process settings dialog.

Solution:

You should click on the **Help** menu and select **NightView User's Guide**. Open the contents tree to **NightView User's Guide**, then **Graphical User Interface**, then **Main Window**, then **Menu Bar**, then **Process Menu**. Scroll down to **Process Settings**, and then follow the link to **Process Settings Dialog Box**. Read the information that the help window displays about the process settings dialog box.

Exercise:

Use the File menu to close the help window.

Solution:

In the help window, you should click on the File menu and select Exit.

The help window goes away.

When this tutorial asks you to read about graphical components, use this same procedure.

Exercise:

Use the process settings dialog box to indicate that the child process in msg should be debugged.

Solution:

You should click on the Process menu and select the Process Settings... entry.

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 - GUI

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.

Exercise:

Use the process settings dialog box to adjust the default handling of the SIGUSR1 signal so that the process does not stop.

Solution:

You should click on the Process menu and select the Process Settings... entry.

In the 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 - GUI

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.

Exercise:

Read about the context menu in the source panel. Set a separate simple breakpoint to stop at each of the following places:

- The line that prompts for the number of signals to send
- The call to `child_routine`
- The comment before the call to `parent_routine`

Solution:

You should right click on the prospective breakpoint 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.

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


```
local:15625 Breakpoint 1 set at main.c:18
local:15625 Breakpoint 2 set at main.c:25
local:15625 Breakpoint 3 set at main.c:30
```

For the Fortran program, the lines are 19, 26, and 28. NightView displays the following information in the message panel.

```
local:15625 Breakpoint 1 set at main.f:19
local:15625 Breakpoint 2 set at main.f:26
local:15625 Breakpoint 3 set at main.f:28
```

An *eventpoint* is a generic term which includes breakpoints, patchpoints, monitorpoints, tracepoints, and heappoints. 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.

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 - GUI

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


Exercise:

Read about the **Resume** button in the process toolbar. Use it to continue program execution up to the statement with the breakpoint.

Solution:

You should 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:15625: at Breakpoint 1, 0x10002818 in main() at  
main.c line 18
```

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

```
local:15625: at Breakpoint 1, 0x10003878 in main() at  
main.f line 19
```


Not Entering Functions - GUI


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.

Exercise:

Read about the **Next** button. Use it to skip over the output statement (or function).

Solution:

You should 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 - GUI

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

Exercise:

Send ten signals.

Solution:

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 - GUI

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

Exercise:

Continue executing the program so it stops at the next statement with a breakpoint.

Solution:

You should 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:15625: at Breakpoint 3, 0x1000284c in main() at
main.c line 31
```

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

```
local:15625: 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. The child process is the one with the status **New Process**. The parent process, 15625, is the *current process*, shown with green underlined text. Buttons, menus and commands generally apply to the current process.

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

Exercise:

Read about the context panel. Switch to the child process.


Solution:

In the context panel, you should 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:13504    parent pid: 15625
#0 0x10002838 in main() at main.c line 20
```

For the Fortran program, the message panel shows:

```
New process: local:13504    parent pid: 15625
#0 0x100038e4 in main() at main.f line 22
```

In this example, the child process has process ID 13504, and the parent process has process ID 15625. 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 3-25. 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:13504`.


The context panel lists entries for processes 15625 and 13504.

Catching up the Child Process - GUI

Exercise:

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`).

Solution:

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


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

```
local:13504: at Breakpoint 5, 0x10002840 in main() at main.c line 25
```

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

```
local:13504: 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 - GUI

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

Exercise:

Read about the **Print** button in the value toolbar. Use it to check that the `total_sig` variable has the value 10.

Solution:

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 3-36.

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 3-37. In the Fortran program, `total_sig` was put in `COMMON` so you could consistently see its value in the tutorial.

Listing the Source - GUI

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

Exercise:

Read about the **Source** menu's **List Function/Unit...** item in “Source Menu” on page 9-10. With the **parent** as the current process, use this item to display the source code for `child_routine`.

Solution:

You should 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**.) Then you should 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.

Exercise:

Read about the **Filter** button in the **Select a Function/Unit** dialog box. Use it to search for `child_routine`.

Solution:

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 7-16.)

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

Exercise:

Read about the **Select** button in the **Select a Function/Unit** dialog box. Use it to change the source displayed.

Solution:

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 - GUI

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

When this tutorial asks you to read about commands, click on the **Help** menu and select **On Commands**. When the help window appears, scroll down to the desired command and click on the link for the command. Read the information about the command and then close the help window.

Exercise:

Read about the `all` qualifier. See “Qualifiers” on page 3-4. Read about the `step` command. Use it to simultaneously enter both routines.

Solution:

Remember, you may need to click in the command toolbar to get the keyboard focus to be there. In the command toolbar, you should enter one of:

```
(all) step
(all) s
```

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 = 13504
, 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 /
13504 / )
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.

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 - GUI

It is often helpful to see how you got to a certain point in a program.

Exercise:

Use the context panel to display the list of currently active stack frames for both processes.

Solution:

You should 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:

```
15625 msg  Stopped after single step
#0  0x10002944  in parent_routine(pid_t child_pid =
      13504, int total_sig = 10) at parent.c line 11
#1  0x10002854  in main() at main.c line 31
#2  0xb7dd5879  in __libc_start_main(...)
13504 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:

```
15625 msg  Stopped after single step
#0  0x10003a48  in parent_routine(INTEGER child_pid /
      13504 / ) at parent.f line 16
#1  0x10003910  in main() at main.f line 29
13504 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 - GUI

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`.

Exercise:

Make sure the **parent** process is the current process.


Solution:

You should click on the entry for the parent process in the context panel.

Exercise:

Read about the Up button in the process toolbar. Use the Up button to make the current stack frame of the **parent** process be main.

Solution:

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:15625
#1 0x10002854 in main() at main.c line 31
```

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

```
Output for process local:15625
#1 0x10003910 in main() at main.f line 29
```

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 - GUI

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

Exercise:

Use the locals panel to display the value of local variable `tracefile` in main for the parent process.

Solution:

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**.

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

```
tracefile 0x30003100 "msg_file"
```


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

```
tracefile "msg_file"
```

Returning to a Stack Frame - GUI

You want to return to `parent_routine`.

Exercise:

You could return to `parent_routine` with the **Down** button. Another way to move in the stack frames is to click on a stack frame in the context panel. Use the context panel to make the current stack frame of the **parent** process be `parent_routine`.

Solution:

You should 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:15625
#0  0x10002944 in parent_routine(pid_t child_pid =
13504, int total_sig = 10)
    at parent.c line 11
```

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

```
Output for process local:15625
#0  0x10003a48 in parent_routine(INTEGER child_pid /
13504 / )
    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 - GUI

You 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.

Exercise:

Use the **Resume** button in the process toolbar to resume execution of the **child** process.

Solution:

You should 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.

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

Removing a Breakpoint - GUI

Breakpoint 1 (set in "Setting the First Breakpoints - GUI" on page 5-9) is no longer needed.

Exercise:

Read about the eventpoint panel's context menu in "Eventpoint Panel" on page 9-57. Use the **Delete** item to remove breakpoint 1.

Solution:

You should 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 - GUI

It is often useful to suspend execution conditionally.

Exercise:

Read about the Eventpoint menu's **Set Breakpoint...** item in "Eventpoint Menu" on page 9-12. Use this feature to 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`.

Solution:

In the source panel, you should click on the line. 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.

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.

Exercise:

Save your changes and make the breakpoint dialog box go away.

Solution:

In the breakpoint dialog box, you should click on the **OK** button.

NightView closes the breakpoint dialog box.

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

```
local:15625 Breakpoint 7 set at parent.c:16
```

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

```
local:15625 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 - GUI

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.

Exercise:

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

Solution:

In the source panel, you should click on the line. 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.

Exercise:

Save your changes and make the breakpoint dialog box go away.

Solution:

In the breakpoint dialog box, you should click on the **OK** button.

NightView closes the breakpoint dialog box.

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

```
local:15625 Breakpoint 8 set at parent.c:16
```

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

```
local:15625 Breakpoint 8 set at parent.f:17
```

Attaching Commands to a Breakpoint - GUI

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

Exercise:

Attach 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 - GUI” on page 5-21).

Solution:

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 - GUI” on page 5-21.

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

In the Commands: text input area, you should enter one of:

```
print total_sig
p total_sig
```

Exercise:

In the breakpoint dialog box, save your changes and make the dialog box go away.

Solution:

In the breakpoint dialog box, you should click on the OK button.

NightView closes the breakpoint dialog box.


Automatically Printing Variables - GUI

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

Exercise:

Read about the Data Display button in the value toolbar. Use it to display the value of the `sig_ct` variable.

Solution:

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 - GUI” on page 5-22 runs only when execution stops at a specific breakpoint.

Watching Inter-Process Communication - GUI

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

Exercise:

Now continue execution for the **parent** process.

Solution:

You should click on the **Resume** button. 

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

```
1. Parent sleeping for 2 seconds
2. Parent sleeping for 2 seconds
Process local:13504 received SIGUSR1
Child got ordinal signal #1
3. Parent sleeping for 2 seconds
Process local:13504 received SIGUSR1
Child got ordinal signal #2
4. Parent sleeping for 2 seconds
Process local:13504 received SIGUSR1
Child got ordinal signal #3
5. Parent sleeping for 2 seconds
Process local:13504 received SIGUSR1
Child got ordinal signal #4
Process local:13504 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:15625: at Breakpoint 8, 0x10002950 in
parent_routine(
    pid_t child_pid = 13504, 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:15625: at Breakpoint 8, 0x105d0 in parent_routine(
    INTEGER child_pid / 13504 / ) 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 - GUI

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`. Hint: In C the `%` operator, or in Fortran the `mod` function, may be useful.

Exercise:

Read about the Eventpoint menu's `Set Patchpoint...` item in "Eventpoint Menu" on page 9-12. In the `parent` process, on the line that displays how long the parent is sleeping, patch in the assignment expression just described.

Solution:

You should click on the line in the source panel, then click on the Eventpoint menu. Select `Set Patchpoint...`

NightView displays the patchpoint dialog box.

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

For the C program, you should enter in the `Evaluate:` text input area:

```
isec = sig_ct % 3 + 1
```

For the Fortran program, you should enter in the `Evaluate:` text input area:

```
isec = mod( sig_ct, 3 ) + 1
```

You are ready to finish patching your program.



Exercise:

Save your changes and make the patchpoint dialog box go away.

Solution:

In the patchpoint dialog box, you should 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 eventpoint  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:15625 Patchpoint 9 set at parent.c:16
```

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

```
local:15625 Patchpoint 9 set at parent.f:17
```

Disabling a Breakpoint - GUI

You want to run `msg` to completion without stopping at breakpoint 8.

Exercise:

Disable breakpoint 8 (set in section “Attaching an Ignore Count to a Breakpoint - GUI” on page 5-21).

Solution:

In the eventpoint panel, you should click on the **Enabled** field of the entry for breakpoint 8, which currently has the value **Yes**. The enabled field becomes a combo box. Set the combo box to **No**. You must click somewhere else in the eventpoint panel to remove the combo box and update the value. For example, you might click on the **PID** field.

The eventpoint panel shows the new enabled status of breakpoint 8: **No**.

Examining Eventpoints - GUI

You want to examine the types, locations, and statuses of the eventpoints you have set in `msg`.

Exercise:

Examine all eventpoints.

Solution:

Refer to the eventpoint panel.

For the C program, NightView displays in the eventpoint list:

EvptID	Type	Enabled	Process - Address
-----	-----	-----	-----
2	B	Enabled	local:15625 at main.c:25
3	B	Enabled	local:15625 at main.c:30
7	B	Enabled	local:15625 at parent.c:16
8	B	Disabled	local:15625 at parent.c:16
9	P	Enabled	local:15625 at parent.c:16
4	B	Enabled	local:13504 at main.c:18
5	B	Enabled	local:13504 at main.c:25
6	B	Enabled	local:13504 at main.c:30

For the Fortran program, NightView displays in the eventpoint list:

EvptID	Type	Enabled	Process - Address
-----	-----	-----	-----
2	B	Enabled	local:15625 at main.f:26
3	B	Enabled	local:15625 at main.f:28
7	B	Enabled	local:15625 at parent.f:17
8	B	Disabled	local:15625 at parent.f:17
9	P	Enabled	local:15625 at parent.f:17
4	B	Enabled	local:13504 at main.f:26
5	B	Enabled	local:13504 at main.f:28
6	B	Enabled	local:13504 at main.f:19

NightView displays all eventpoints for process local:15625 followed by the eventpoints for process local:13504.

Breakpoints 1, 2, and 3 were set in “Setting the First Breakpoints - GUI” on page 5-9. Breakpoint 1 has no entry because it was deleted in “Removing a Breakpoint - GUI” on page 5-20. 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 - GUI” on page 5-20 and is still enabled.

Breakpoint 8 was set in “Attaching an Ignore Count to a Breakpoint - GUI” on page 5-21 and was disabled in “Disabling a Breakpoint - GUI” on page 5-26.

Patchpoint 9 was set in “Patching Your Program - GUI” on page 5-25 and is still enabled.

Continuing to Completion - GUI

There’s nothing else to look at, so you decide to run msg to completion.

Exercise:

Continue execution of msg.

Solution:

If the parent process is not the currently displayed process, you should switch to it. You should 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:13504 received SIGUSR1
Child got ordinal signal #6
8. Parent sleeping for 3 seconds
Process local:13504 received SIGUSR1
Child got ordinal signal #7
9. Parent sleeping for 1 seconds
Process local:13504 received SIGUSR1
Child got ordinal signal #8
10. Parent sleeping for 2 seconds
Process local:13504 received SIGUSR1
Child got ordinal signal #9
Process local:13504 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 - GUI" on page 5-25. 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 3-18.

NightView displays in the message panel:

```
Process local:15625 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 - GUI

The tutorial is over.

Exercise:

Read about the **File** menu. Use it to leave the debugger.

Solution:

You should 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?
```

Exercise:

Make the processes go away.

Solution:

In the warning dialog box, you should click on the **OK** button.

The main window is removed.

Invoking NightView

This section describes how to start a NightView session.

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 [-attach pid] [-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 ...]]
```

-attach *pid*

Attach to the process specified by *pid* in the local dialogue. This is similar to using the **attach** command. See “attach” on page 7-36. 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 9-4. 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 3-4 and “core-file” on page 7-38.

-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 9 [Graphical User Interface] on page 9-1.

-nolocal

Prevents NightView from starting a dialogue on the local system. See “Dialogues” on page 3-4. In the graphical user interface, if **-nolocal** is used, NightView pops up a Remote Login Dialog Box (see “Remote Login Dialog Box” on page 9-33).

-nx

Prevents NightView from reading commands from the default initialization file. See "Initialization Files" on page 3-36.

-prompt *string*

Sets NightView's initial prompt string to *string*. See "set-prompt" on page 7-59.

-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 7-60.

-simplscreen

Directs NightView to use a simple full-screen interface. This option implies **-nogui**. See Chapter 8 [Simple Full-Screen Interface] on page 8-1.

-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 9 [Graphical User Interface] on page 9-1.

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 “NightStar LX Licensing” on page A-1 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 3-36), 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 7-21).

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 9 [Graphical User Interface] on page 9-1), 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 blank or tab 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 7-14.

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 7-19.

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 7-153 and “source” on page 7-129).

NightView prompts you for input. The format of the prompt may be controlled by the **set-prompt** command (see “set-prompt” on page 7-59). 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 7-21).

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 3-22.

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 3-20. 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 7-1. Special '\$' Constructs

\$

A simple '\$' by itself is a special convenience variable which always refers to the last value history entry (see "print" on page 7-81). See "Value History" on page 3-36.

\$\$

The name '\$\$' refers to the value history entry immediately prior to '\$'. See "Value History" on page 3-36.

\$number

A '\$' followed by a number refers to that number entry in the value history. See "Value History" on page 3-36.

\${-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 3-36.

\$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

nience variables). See “Convenience Variables” on page 3-35, and “Predefined Convenience Variables” on page 7-5.

$\$\{file:line\ expression\}$

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 7-125) 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 7-81).

$\$\{*frame-addr\ expression\}$

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 7-88) 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 7-138) 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 7-3). The following special variables are all automatically defined on a per process basis.

Table 7-2. Predefined Convenience Variables

\$_

This variable holds the address of the last item dumped with the **x** command (see “x” on page 7-84). 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 7-84, “info eventpoint” on page 7-131, “info breakpoint” on page 7-132, “info tracepoint” on page 7-133, “info patchpoint” on page 7-134, “info monitorpoint” on page 7-135, “info heappoint” on page 7-136, and “info line” on page 7-153.

\$__

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 3-23.

`$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 3-25. 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 7-150 and “info frame” on page 7-138).

`$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 7-138.

`$is`

`$is` is defined when a watchpoint is triggered. See “Watchpoints” on page 3-13. `$is` is the value of the variable being watched after the instruction that causes the trigger has completed.

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 7-140.

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 7-3. IA-32 Registers

`$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 7-116, “stepi” on page 7-118, “next” on page 7-117, and “nexti” on page 7-119. 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.

`$fcs` or `$fiseq`

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.

`$fos` or `$foseq`

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 3-13. 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 3-13, “step” on page 7-116, “stepi” on page 7-118, “next” on page 7-117, and “nexti” on page 7-119. 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 3-13. 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 7-140). If you want to display those registers, you can do so with the following commands:

<code>info registers st.*</code>	<i>displays floating point registers</i>
<code>info registers dr.*</code>	<i>displays debug registers</i>
<code>info registers mm.*</code>	<i>displays MMX registers</i>
<code>info registers xmm.*</code>	<i>displays SSE registers</i>
<code>info registers .*</code>	<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 7-140.

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 7-4. AMD64 Registers

`$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 7-116, “stepi” on page 7-118, “next” on page 7-117, and “nexti” on page 7-119. 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 3-13. 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 3-13, “step” on page 7-116, “stepi” on page 7-118, “next” on page 7-117, and “nexti” on page 7-119. 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 3-13. 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.

Note that the floating point, debug, MMX, and SSE registers are not normally displayed by the `info registers` command (see “info registers” on page 7-140). If you want to display those registers, you can do so with the following commands:

<code>info registers st.*</code>	<i>displays floating point registers</i>
<code>info registers dr.*</code>	<i>displays debug registers</i>
<code>info registers mm.*</code>	<i>displays MMX registers</i>
<code>info registers xmm.*</code>	<i>displays SSE registers</i>
<code>info registers .*</code>	<i>displays all registers</i>

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 3-24.

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 7-92), 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 7-153) 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 7-2). 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 7-66).

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 7-57). 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 7-45.

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 7-22). 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 “Multithreaded Programs” on page 3-37. 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 7-57). 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 7-94). 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 7-13. 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*, and *watchpoint specifiers* are similarly defined).

Regular Expressions

A *regexp* 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 7-66).

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 7-18). 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 7-5. Regular Expressions

.

A dot matches any character except a newline.

*

A star matches zero or more occurrences of the preceding regular expression. For example, `.*` matches zero or more of any character except a newline.

+

A plus matches one or more occurrences of the preceding regular expression.

{*m*}

Matches exactly *m* occurrences of the preceding regular expression.

{*m*, }

Matches *m* or more occurrences of the preceding regular expression.

{*m*, *n*}

Matches from *m* to *n* 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 7-24, “nodebug” on page 7-25, and “on program” on page 7-41.

Wildcard patterns are similar to regular expressions, but are usually more convenient for representing file names. See “Regular Expressions” on page 7-16.

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 7-6. 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 7-60) 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 7-60) 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 7-20).

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 7-25.

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 7-153) and for other programs that communicate with NightView.

`/popup`

Pop up the Remote Login Dialog Box (see “Remote Login Dialog Box” on page 9-33) 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 7-45). 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.

Any programs started in the dialogue shell will inherit all the above 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 7-18.

This command and its inverse (see “`nodebug`” on page 7-25) 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 7-146.

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 7-24).

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 7-37).

To print the list of **debug** and **nodebug** patterns, see “info dialogue” on page 7-146.

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: **x1**

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 7-18. See "Regular Expressions" on page 7-16.

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 7-40. See "load" on page 7-91. Dynamic library names are also subject to object filename translation. See "Debugging with Shared Libraries" on page 3-39. Object filenames from **symbol-file** commands are *not* subject to object filename translation. See "symbol-file" on page 7-38.

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 7-71.

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 7-28. Also, see "Initialization Files" on page 3-36.

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 7-60). 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 7-20). 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 7-37). 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 7-25.

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 [regex]
```

```
on dialogue regex command
```

```
on dialogue regex do
```

regex

A regular expression to match against the names of newly created dialogues. See “Regular Expressions” on page 7-16.

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 7-60.

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 7-30. See “Initialization Files” on page 3-36.

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 7-28. 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 3-36. 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 7-60, 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 7-153), 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 7-19.

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 7-34.

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 7-33).

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 7-56.

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 7-32).

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 7-31): 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 3-16. 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 7-122), 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 7-33 and “set-show” on page 7-32).

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 7-36) 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 7-35), 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
```

```
/resume
```

Resume the process when the attach is complete.

```
/stop
```

Keep the process stopped when the attach is complete.

```
PID
```

The process ID of the 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.

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 7-31) 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 7-47. Children created prior to the attach must be explicitly attached if you want to debug them.

See “Attach Permissions” on page 3-38 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 7-68). You can override that setting by explicitly giving a `/resume` or `/stop` option.

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, or watchpoints 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 7-95, “patchpoint” on page 7-96, “monitorpoint” on page 7-101, “heappoint” on page 7-102, “tracepoint” on page 7-99, and “watchpoint” on page 7-112.

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. If the process has any patchpoints or tracepoints, then the debugger silently continues to monitor the process to control these eventpoints.

When the safety level is `unsafe` (see “set-safety” on page 7-60), 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 7-25. Also, see “set-children” on page 7-47.

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 7-60) 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 7-26.

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 7-40).

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 7-26.

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 7-45.

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 7-40) 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 7-40 and “symbol-file” on page 7-38).

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 7-26). 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.

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 7-26.

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 7-38). 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 7-38), 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 7-41.

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 7-26.

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 7-18.

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 7-44. 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 3-16.

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 7-60.

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 3-34.

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 7-41). 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 7-18.

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 3-16. 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 7-41 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 7-147.

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 3-16. **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 alphanu-

meric 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 7-14.

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 7-60.

If one or more *qualifier-spec* arguments are supplied, they are immediately evaluated (see “Qualifier Specifiers” on page 7-14) 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 7-24. This is the default setting for direct children of the dialogue shell and processes debugged with the `attach` command. See “attach” on page 7-36.

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 7-37.

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 3-18.

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 7-68. 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 7-41 and “on restart” on page 7-44. See also “Restarting a Program” on page 3-16.

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 7-68, “resume” on page 7-115 and “on program” on page 7-41.

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 18. That is, no more commands are read from this command stream until the specified processes stop. See “Command Streams” on page 34. See “Interrupting the Debugger” on page 34.

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 E [Implementation Overview] on page E-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 7-91. See “Expression Evaluation” on page 3-20.

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 7-142.

Memory reservations are remembered as part of the restart information. See “Restart Information” on page 3-17. During restart, memory reservations are applied before any commands that would allocate space in the process.

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 3-32) 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 7-102) or a **heapcheck** command (see “heapcheck” on page 7-148). 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 3-29 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 3-30.

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 3-30.

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 3-31.

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 7-63). See “Debugging the Heap” on page 3-28.

The **heapdebug** command configures the heap debugger in each of the processes in the qualifier. See “Debugging the Heap” on page 3-28. 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 9-8).

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.

Setting Modes

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 7-131).

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 7-32).

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 7-81), 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 7-14). 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*} ...

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 7-13). If you want unlimited output, specify zero as the limit.

The *array*, *string*, and *addresses* 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 7-140).

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-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 7-31).

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 7-20). 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 3-16.

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*} ...

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.

The `data`, `monitor`, `text`, and `eventpoint` 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, and, if there are any monitorpoints in the process, one shared memory region for the monitorpoints. These regions are called *patch areas*. See Appendix E [Implementation Overview] on page E-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 and text 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 3-11.

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, and 32 kilobytes for the monitorpoint 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 7-146).

You can see information about the patch areas in an existing process with the `info memory` command (see “info memory” on page 7-142).

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 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 7-13. 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.

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 3-25.

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 3-17. For a way to see all the interest levels that have been explicitly set, see “info on restart” on page 7-147.

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 3-24. 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 7-126. 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 7-78.

```
(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 7-63. 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 3-20. 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 3-22. For the details of the syntax used to specify overloading in expressions and location specifiers see “Selecting Overloaded Entities” on page 7-2.

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 7-16) 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` or `vi`.

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 8-2.

set-preallocate

Control how NightView preallocates memory for eventpoints and monitorpoint buffers.

set-preallocate [/eventpoint] [/monitorpoint] [{off | on}]

/eventpoint

Indicates the eventpoint mode should be set or queried.

/monitorpoint

Indicates the monitorpoint buffer mode should be set or queried.

{off | on}

Turn preallocation off or on for eventpoints or monitorpoint buffers as specified.

The default is for NightView to preallocate space in the user process for eventpoints and monitorpoint 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. See “Process States” on page 3-18. See “Operations While the Process Is Executing” on page 3-19.

Monitorpoints cannot be used in the command-line interface, so in that interface monitorpoint buffers are never preallocated.

This is a global mode. The qualifier is ignored.

With no arguments, **set-preallocate** prints the current settings.

set-resume

Control NightView's behavior on events that normally stop a process.

set-resume [/attach] [/exec] [/exit] [/fork] [{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 7-47.

{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). 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** prints 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 9-44.

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}]
                [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 GNU 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.

`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 7-84. See "Source Menu" on page 9-10.

If no arguments are specified, **set-disassembly** prints the current disassembly information.

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 7-128)
- **list** (see “list” on page 7-72)
- **directory** (see “directory” on page 7-74)
- **symbol-file** (see “symbol-file” on page 7-38)
- **core-file** (see “core-file” on page 7-38)
- **exec-file** (see “exec-file” on page 7-40)

The **cd** command does not affect commands executed in dialogue shells (see “login” on page 7-22). 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 7-71.

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 7-71).

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 7-13). 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 7-74). 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 7-124) 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 7-76 and “reverse-search” on page 7-76.

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 7-78).

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 7-153).

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 7-71 and “pwd” on page 7-71.

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 7-72), 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 7-139.

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/doi/doi.c`. You want to debug your program, but you have moved the source files to `/usr/joe/main/main.c` and `/usr/joe/doi/doi.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 9-48.

Searching

forward-search

Search forward through the current source file for a specified regular expression.

forward-search [*regexp*]

Abbreviation: **fo**

regexp

The regular expression to search for. *No* anchored match is implied. (see “Regular Expressions” on page 7-16). If *regexp* is omitted, the previous *regexp* 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 7-72.

The regular expression case sensitivity depends on the current search mode (see “set-search” on page 7-66).

reverse-search

Search backwards through the current source file for a specified regular expression.

reverse-search [*regexp*]

regexp

The regular expression to search for. *No* anchored match is implied. (see “Regular Expressions” on page 7-16). If *regexp* is omitted, the previous *regexp* 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 7-72.

The regular expression case sensitivity depends on the current search mode (see “set-search” on page 7-66).





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 7-84).

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 9-1. 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, or there are multiple eventpoints (of the same kind) on the line, 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 7-7. 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 7-95).
'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 7-102).
'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 7-101).

'P' 

Indicates that one or more patchpoints, possibly disabled, have been inserted somewhere within this source line. (See “patchpoint” on page 7-96). 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 7-99).



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 7-124), 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 7-124), which was executing when the called frame was created, i.e., `$cpc`. See “Program Counter” on page 3-23.

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 <strtol>
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 7-124, “up” on page 7-125, or “down” on page 7-126.

Frames corresponding to uninteresting subprograms are not shown in the listing. See “Interesting Subprograms” on page 3-25.

print

Print the value of a language expression.

print [*/print-format-letter*] *expression*

Abbreviation: **p**

print-format-letter

One of the following letters 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.

d

Print the bit representation of the value in signed decimal.

f

Print the bit representation of the value as a single precision floating-point number and print using floating-point syntax. If the data type of the language expression is double precision, however, then the bit representation is printed as a double precision floating-point number.

o

Print the bit representation of the value in octal.

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.

x

Print the bit representation of the value in hexadecimal.

expression

A language expression (see "Expression Evaluation" on page 3-20).

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-letter* is given, *expression* is printed in a format corresponding to the data type of the expression in the currently defined language.

The printed value is given a value history number (see “Value History” on page 3-36), 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 7-58).

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 3-20).

This command is similar to the **print** command (see “print” on page 7-81), 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-letter] ] [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-letter*.

size-letter

One of the following letters 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 system, this is 16 bytes. A `t` memory unit cannot be printed as decimal `d` or `u`.

The *size-letter* may appear either before or after the *x-format-letter*.

x-format-letter

One of the following letters 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.

d

Print as signed integers in decimal format.

f

Print as floating-point values.

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 7-69 to control the form of the disassembly.

You can also view disassembly in a source panel (see “Source Menu” on page 9-10).

o

Print as unsigned integers in octal format.

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 7-58).

u

Print as unsigned integers in decimal format.

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 3-20).

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.

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-letter* 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-letter* 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-letter* 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 7-5). 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 7-78 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 |wvutsrqp|
(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-letter] expression
```

print-format-letter

A letter specifying the format in which to print the expression, as described in the **print** command (see “print” on page 7-81).

expression

A language expression (see “Expression Evaluation” on page 3-20).

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-letter* 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 newline.

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 3-17. 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.

display

Add to the list of expressions to be printed each time the process stops.

display [[/print-format-letter] *expression*]
display / [repeat-count] [size-letter] [x-format-letter] *addr-expression*

print-format-letter

A letter specifying the format in which to print the expression, as in the **print** command (see “print” on page 7-81).

expression

A language expression (see “Expression Evaluation” on page 3-20).

repeat-count

Decimal number of consecutive memory units to print, where a unit is defined by the *size-letter* and the *x-format-letter*.

size-letter

A letter specifying the size of each memory unit, as described in the **x** command (see “x” on page 7-84). The *size-letter* may appear either before or after the *x-format-letter*.

x-format-letter

A letter specifying the format in which to print the contents of memory, as described in the **x** command (see “x” on page 7-84).

addr-expression

An expression yielding a memory address (see “Expression Evaluation” on page 3-20).

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-letters* ‘s’ or ‘i’ is given, then the expression is treated as an *addr-expression*. Otherwise, the expression is treated as a language *expression*.

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-letter* 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-letter* is omitted, the defaults are w and d, respectively.

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 7-139).

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 7-88).

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 7-90). 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 7-139).

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 7-88).

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 7-139).

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 3-20).

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.

To specify an object file on a remote system, use the form *user@host:/path*. NightView downloads the file from the remote system to the host system to be loaded into the program. See “Remote File Access” on page 3-7.

object is subject to object filename translations (see “translate-object-file” on page 7-26).

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.

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, and watchpoints as well. The following table indicates which types of eventpoints may be affected by which commands:

Table 7-8. Eventpoint Commands

Command Name	What the Command May Apply To					
	Breakpoints	Patchpoints	Tracepoints	Monitorpoints	Watchpoints	Heappoints
<code>name</code>	X	X	X	X	X	X
<code>clear</code>	X	X	X	X		X
<code>commands</code>	X			X	X	
<code>condition</code>	X	X	X	X	X	X
<code>delete</code>	X	X	X	X	X	X
<code>disable</code>	X	X	X	X	X	X
<code>enable</code>	X	X	X	X	X	X
<code>ignore</code>	X	X	X	X	X	X
<code>tbreak</code>	X					
<code>tpatch</code>		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 7-108).

/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 7-108).

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 7-16.

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] [name=breakpoint-name]
              [[at] location-spec] [if conditional-expression]
```

Abbreviation: **b**

eventpoint-modifier

Specifies the breakpoint modifier. See “Eventpoint Modifiers” on page 7-94.

name=breakpoint-name

Gives a name to the breakpoint for later reference. (see “name” on page 7-94). 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 7-13).

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 3-24.

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 3-23 and “Context” on page 3-23). See also “Expression Evaluation” on page 3-20.

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 7-109).

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 7-105.

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

Install a small patch to a routine.

```
patchpoint [eventpoint-modifier] [name=patchpoint-name]  
             [[at] location-spec] eval expression
```

Insert an expression in the program.

```
patchpoint [eventpoint-modifier] [name=patchpoint-name]  
             [[at] location-spec] goto location-spec
```


Insert a branch in the program.

eventpoint-modifier

Specifies the patchpoint modifier. See “Eventpoint Modifiers” on page 7-94.

name=patchpoint-name

Patchpoints are assigned event numbers, and the *name=* syntax as well as the **name** command (see “name” on page 7-94) may be used to give them names. See “Manipulating Eventpoints” on page 7-92.

at location-spec

Specify the exact point in the program to execute the patchpoint (see “Location Specifiers” on page 7-13). 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 3-24.

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 3-20.

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 **\$1r*, the value of register *1r* in the *current context* is used as the location to branch to.

NOTE

The keywords *name*, *at*, *eval*, and *goto* 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 7-106) or **ignore** (see “ignore” on page 7-109) 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 7-37 and “disable” on page 7-108). 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.

set-trace

Establish tracing parameters.

```
set-trace [eventmap=event-map-file] [tracefile=trace-file]
```

```
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)**.

```
tracefile=trace-file
```

Specifies the name of the file to which the user daemon logs trace events. This filename is passed to **trace_begin()** on behalf of the process.

The **set-trace** command is used to specify information that may be required before any tracepoints may be set in a process (see “tracepoint” on page 7-99).

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.

If your program does not initialize tracing with a call to **trace_begin()**, then you must use the **set-trace** command to pass the *trace-file* to the trace library before the process executes any tracepoints. The debugger causes the process to call **trace_begin()** with this filename.

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 7-94.

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 7-98).

name=tracepoint-name

Gives a name to the tracepoint for later reference. See “name” on page 7-94. 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 7-13.

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 3-20.

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 3-23 and "Context" on page 3-23). See also "Expression Evaluation" on page 3-20.

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 7-109).

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 7-37 and "disable" on page 7-108).

NOTE

If your program does not initialize tracing with a call to **trace_begin()**, then you must initialize tracing manually with the **set-trace** command before the process executes any tracepoints. Otherwise, the tracepoints will not save any trace events.

The debugger does *not* start the **ntraceud(1)** monitor process. You must do that manually (see "NightTrace Analysis" on page 3-38).

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 number.

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.

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 7-94.

name=monitorpoint-name

Gives a name to the monitorpoint for later reference. See “name” on page 7-94. 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 7-13.

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 7-1.

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 3-26. For a more detailed description of monitorpoints, see “Monitorpoints” on page 3-11.

The syntax of the commands attached to a monitorpoint is:

```
print [/print-format-letter] [id="string"] expression
```

This syntax is identical to the **print** NightView command (see “print” on page 7-81), 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 7-105.

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)`.

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 7-94.

name=heappoint-name

Gives a name to the heappoint for later reference. See “name” on page 7-94. 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 7-13.

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 7-50).

`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 3-23 and “Context” on page 3-23). See also “Expression Evaluation” on page 3-20.

The **heappoint** command sets a heappoint in each of the processes specified by the qualifier. See “Heappoints” on page 3-12.

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.

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 7-101.

`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 3-34.

`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 3-26).

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 7-13.

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 3-24.

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 7-107). 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 7-13). 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 7-16.

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 7-1.

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 7-114, “resume” on page 7-115, and “signal” on page 7-121.

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 7-101.

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 7-16.

conditional-expression

The condition to be associated with the eventpoints. See “Expression Evaluation” on page 3-20.

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, and watchpoints, 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=trace1 at foo.c:12
(local) condition trace1 (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 7-16.

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 7-104). 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 7-60). 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 7-20). 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 7-16.

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 7-131). 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 7-108).

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 7-60). 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 7-20). 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 7-16.

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 7-60). 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 7-20). 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.

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 7-94. 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 7-13.

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 3-23 and "Context" on page 3-23). See "Expression Evaluation" on page 3-20.

Note: The *at*, *if*, and *name* keywords may not be abbreviated in this command.

Like the **breakpoint** command (see "breakpoint" on page 7-95), 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 7-108.

```
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 7-94) may be used to give them names. See “Manipulating Eventpoints” on page 7-92.

at location-spec

Specify the exact point in the program to execute the patchpoint. See “Location Specifiers” on page 7-13. 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 3-24.

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 3-20.

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 7-96. 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 7-108.

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 7-94.

/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 7-94). 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 3-13.

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 3-13.

You can specify commands to be executed when the watchpoint is hit. See “commands” on page 7-105.

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 7-122. If you want the process to receive a different signal, or no signal at all, then use the **signal** command. See “signal” on page 7-121.

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.

continue

Continue execution and wait for something to happen.

continue [*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.

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 3-34.

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 7-115.

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 [*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 7-121.

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.

resume causes the processes specified by the qualifier to resume execution at the point where they last stopped. The processes 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 3-34.

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 7-114.

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 [repeat]
```

Abbreviation: **s**

repeat

The *repeat* argument specifies the number of lines to single step. The default is one line.

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 7-117.

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 7-124.

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 3-14 for a discussion of the interactions between single-stepping and signals.

step is interpreted relative to the current frame. See “Current Frame” on page 3-24. That is, any lower frames are automatically finished before stepping.

There are also commands to single step individual instructions. See “stepi” on page 7-118 and “nexti” on page 7-119.

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 7-78). Also, disassembly can help you determine the flow of control through your program (see “x” on page 7-84 and “Source Menu” on page 9-10).

If the **step** command causes execution to enter a procedure which is uninteresting, the **step** acts like **next**. See “Interesting Subprograms” on page 3-25. See “next” on page 7-117.

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.

next

Execute one line, stepping over procedures.

next [*repeat*]

Abbreviation: **n**

repeat

The *repeat* argument specifies the number of lines to single step. The default is one line.

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 3-25. 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 7-116.

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 7-124.

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 3-14 for a discussion of the interactions between single-stepping and signals.

next is interpreted relative to the current frame. See “Current Frame” on page 3-24. That is, any lower frames are automatically finished before stepping.

There are also commands to single step individual instructions. See “nexti” on page 7-119 and “stepi” on page 7-118.

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 7-78). Also, disassembly can help you determine the flow of control through your program (see “x” on page 7-84 and “Source Menu” on page 9-10).

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.

stepi

Execute one instruction, stepping into procedures.

stepi [*repeat*]

Abbreviation: **si**

repeat

The *repeat* argument specifies the number of instructions to single step. The default is one instruction.

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 7-116.

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 3-14 for a discussion of the interactions between single-stepping and signals.

stepi is interpreted relative to the current frame. See “Current Frame” on page 3-24. 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 7-88.

If the **stepi** command causes execution to enter a procedure which is uninteresting, the **stepi** acts like **nexti**. See “Interesting Subprograms” on page 3-25. See “nexti” on page 7-119.

nexti

Execute one instruction, stepping over procedures.

nexti [*repeat*]

Abbreviation: **ni**

repeat

The *repeat* argument specifies the number of instructions to single step. The default is one instruction.

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 3-25.

This is very similar to **next**, except that **next** executes lines and **nexti** executes individual instructions. See “next” on page 7-117.

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 3-14 for a discussion of the interactions between single-stepping and signals.

nexti is interpreted relative to the current frame. See “Current Frame” on page 3-24. That is, any lower frames are automatically finished before stepping.

finish

Continue execution until the current function finishes.

finish

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 7-124. 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 “Multithreaded Programs” on page 3-37.

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 3-14 concerning interactions between single-stepping and signals also applies to the **finish** command.

stop

Stop a process.

stop

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.

WARNING

It is possible, though unlikely, that the process will stop of its own accord (say by hitting a breakpoint) while NightView is trying to stop it. If that happens, your process may receive a spurious SIGTRAP signal the next time you resume its execution. This signal should be harmless; resuming your process after this signal occurs should get everything back to normal.

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 [at] location-spec
```


location-spec

The *location-spec* specifies where to continue execution. See “Location Specifiers” on page 7-13.

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 3-25.

signal

Continue execution with a signal.

signal *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.

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 3-34.

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.

For a way to have the debugger deal with signals automatically, see “handle” on page 7-122. **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 7-141.

Example:

```
(local) signal 2
```

The processes resume with signal number 2.

handle

Specify how to handle signals in the user process.

```
handle [/signal] sigid 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.

keyword

keyword is one of `stop`, `nostop`, `print`, `noprint`, `pass` or `nopass`. Multiple keywords may be specified.

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 9 [Graphical User Interface] on page 9-1. See “Message Panel” on page 9-56.

`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.

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.

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 3-14 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 3-14.)

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 7-121.

To find out the current settings for all the signals, see “info signal” on page 7-141.

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.

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 7-81).

**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 7-13. 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 7-138.

If a *frame-number* is given, the chosen stack frame is selected as the current frame (see “Current Frame” on page 3-24).

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 3-24). If a negative number is specified, then frames are advanced toward the newest stack frame (see “down” on page 7-126).

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 3-24). If a negative number is specified, then frames are advanced toward the oldest stack frame (see "up" on page 7-125).

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*

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 "Multithreaded Programs" on page 3-37.

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.

The **select-context** command allows you to examine the context (see “Examining Your Program” on page 3-20) of a thread or a thread process. Using **select-context**, you can get a backtrace (see “backtrace” on page 7-81) and examine registers and variables in the context of the selected thread or thread process.

When a process that contains multiple threads or thread processes stops, the current context becomes that of one specific thread or thread process. (For a discussion of how this choice is made, see “Multithreaded Programs” on page 3-37.) You can use the **select-context** command to temporarily change the context to that of some other thread or thread process.

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.

Note that execution control is on a process basis: if you resume execution, all threads are allowed to execute. If you enter a **finish**, **step**, **next**, **stepi**, or **nexti** command, the process executes until the selected thread or thread process completes the stepping operation, but other threads or thread processes may execute as well.

If you request evaluation of an expression containing a function call, the process is allowed to execute and all threads are allowed to run. If another thread hits a breakpoint, or stops for some other reason, the function call is terminated prematurely and an error message is issued.

To get a list of the threads and thread processes in a process, see “info threads” on page 7-148.

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 3-33.

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 9-1.

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 8 [Simple Full-Screen Interface] on page 8-1.

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 7-153).

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 7-71).

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 8 [Simple Full-Screen Interface] on page 8-1.

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 3-34.

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.

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 7-56) or by **set-show** (see “set-show” on page 7-32).

info eventpoint

Describe current state of breakpoints, tracepoints, patchpoints, monitorpoints, heappoints, and watchpoints.

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 7-58). 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 7-16.

This command describes eventpoints associated with the processes in the command qualifier. An eventpoint is any of a breakpoint, tracepoint, patchpoint, monitorpoint, heappoint, or watchpoint. See “breakpoint” on page 7-95, “tracepoint” on page 7-99, “patchpoint” on page 7-96, “monitorpoint” on page 7-101, “watchpoint” on page 7-112, and “heappoint” on page 7-102.

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, or watchpoint).
- 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.

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 7-84 and “Predefined Convenience Variables” on page 7-5.

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 7-58). 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 7-16.

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 7-95.

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 7-84 and “Predefined Convenience Variables” on page 7-5.

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 7-58). 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 7-16.

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 7-99.

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 7-84 and “Predefined Convenience Variables” on page 7-5.

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 7-58). 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 7-16.

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 7-96.

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 7-84 and “Predefined Convenience Variables” on page 7-5.

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 7-58). 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 7-16.

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 7-101.

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 7-84 and “Predefined Convenience Variables” on page 7-5.

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 7-58). 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 7-16.

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 7-102.

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 7-84 and “Predefined Convenience Variables” on page 7-5.

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 7-16.

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 7-112.

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 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 7-13. 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 3-24). See also “frame” on page 7-124.

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 7-74, 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 7-61). 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 7-140).

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 7-88).

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 7-58.

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 7-58.

The qualifier is ignored by this command.

info registers

Print information about registers.

info registers [*regexp*]

regexp

A regular expression matching register names. An anchored match *is* implied. See “Regular Expressions” on page 7-16.

If the *regexp* 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 7-5.

The regular expression case sensitivity depends on the current search mode (see “set-search” on page 7-66).

Registers are printed relative to the current frame (see “Current Frame” on page 3-24). 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 7-124).

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 7-122).

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 (a11) 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, stopped).
- When the process state is stopped, list where and why it stopped.
- The current language setting. See “set-language” on page 7-56.
- The disposition of child processes; that is, under what circumstances a child process will be debugged. See “set-children” on page 7-47.
- 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.

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.
- If the region is the first region associated with a shared library, the name of the library is printed.
- Whether the region is readable, writable, executable, shared, or locked in physical memory.
- Whether the region is being used as the process' stack or memory heap.
- 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 E-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 7-49.

/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 3-32 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 7-62.
- The list of **debug** and **nodebug** patterns for this dialogue. See “debug” on page 7-24.
- 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 7-32).
- The list of object filename translations for this dialogue. See “translate-object-file” on page 7-26.

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 7-16.

For each family name that matches *regex* this command lists each process that is a member of that family (see “family” on page 7-45). 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 7-66).

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 7-16.

For each eventpoint-name that matches *regex*, this command lists each eventpoint that is a member of that eventpoint-name (see “name” on page 7-94). 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 7-66).

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 7-28

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 7-41.

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 7-44.

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 7-129. See “Restarting a Program” on page 3-16. For an example, see “checkpoint” on page 7-45.

info threads

Describe C threads and thread processes.

info threads [/verbose]

`/verbose`

Show which thread is the current thread by placing `<=` next to it, for each thread type.

This command describes the C threads and thread processes for the processes specified by the qualifier. 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 7-126.

In the graphical user interface, you can also get thread information from the Display menu. See “Data Menu” on page 9-14.

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 3-32), or `heappoint` check (see “`heappoint`” on page 7-102).

`/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 3-30. 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 7-50) or the **Debug Heap...** item in the **PROCESS** menu (see “Process Menu” on page 9-8).

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 3-24).

info locals

Print information about local variables.

info locals [*regex*]

regexp

A regular expression matched against local variable names. An anchored match *is* implied. See “Regular Expressions” on page 7-16.

Print a description of every local variable visible in the current context. If the *regexp* 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 7-66).

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 [*regexp*]

regexp

A regular expression matched against global variable names. An anchored match *is* implied. See “Regular Expressions” on page 7-16.

This command prints information about global variables. When the *regexp* 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 7-66).

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 7-81).

info sources

List names of source files.

info sources [*pattern*]

pattern

Wildcard pattern to match against source file names. See “Wildcard Patterns” on page 7-18.

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 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 7-16.

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 7-66).

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 7-16.

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 7-66).

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 3-20.

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 3-20.

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: **ptype**

regex

A regular expression to match against type names and variable names. An anchored match *is* implied. See “Regular Expressions” on page 7-16.

The *regexp* 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 7-66).

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 7-13). 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 7-58). 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 7-84 and “Predefined Convenience Variables” on page 7-5.

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 7-156 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 7-1.)

The **define** command associates a body of text with the given *macro-name*. When the macro is invoked (see “Referencing Macros” on page 7-156), 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 7-156 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 C [Summary of Commands] on page C-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 7-91.
- Macros appearing in an **echo** command are expanded. See “echo” on page 7-87.
- Macros appearing in a **!** (see “!” on page 7-31), **run** (see “run” on page 7-34), or **shell** (see “shell” on page 7-128) 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 7-16.

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 7-66).

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.

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 7-32.)

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 `-simplscreen` option:

```
nview -simplscreen
```

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 7-128.

The simple full-screen interface creates a special window when you use monitorpoints. See "Monitor Window - Simple Full-Screen" on page 8-2 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 7-67.

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 3-26.

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 7-103)

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 3-26.

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 5 [Tutorial - GUI] on page 5-1.

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 9-2.
- The main window has a **Help** menu. See “Help Menu” on page 9-17.
- 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 7-127.

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 9-80.

Context-Sensitive Help

Context-sensitive help is available through the **Help** menu. See “Help Menu” on page 9-17. 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 9-19.

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 9-2.

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 7-127. A help window is displayed that contains information about the requested topic. See “Help Window” on page 9-80. See “Command Toolbar” on page 9-19.

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 3-4. 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 9-50 and “Source Panel Target Line” on page 9-49.

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 9-65.

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 9-23. A locals panel shows variables in the current frame of the current process. See “Locals Panel” on page 9-59. If you have more than one process, a context panel shows the current process with green underlined text. See “Context Panel” on page 9-59.

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 9-4.

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 9-48. NightView can remember the arrangement of the panels. See “GUI Configuration” on page 9-3.

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 9-3.

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 9-3.

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 9-3.

Save Config File As...

Mnemonic: A

This brings up a file selection dialog to save the configuration. See “GUI Configuration” on page 9-3.

Preferences...

Mnemonic: F

This brings up a dialog that lets you set preferences, such as fonts. See “Preferences Dialog Box” on page 9-36.

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 (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 7-21.

Depending on the safety level (see “set-safety” on page 7-60), 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 9-59.

New Locals Panel

Mnemonic: L

Selecting this menu item creates a new locals panel. See “Locals Panel” on page 9-59.

New Source Panel

Mnemonic: S

Selecting this menu item creates a new source panel. See “Source Panel” on page 9-48.

New Data Panel

Mnemonic: D

Selecting this menu item creates a new data panel. See “Data Panel” on page 9-60.

New Monitor Panel

Mnemonic: M

Selecting this menu item creates a new monitor panel. See “Monitor Panel” on page 9-59.

New Shell Panel

Mnemonic: H

Selecting this item creates a new shell panel. See “Source Panel” on page 9-48.

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 9-56.

New Eventpoint Panel

Mnemonic: E

Selecting this item creates a new eventpoint panel for the current shell. See “Eventpoint Panel” on page 9-57.

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 “Rename Page Dialog Box” on page 9-47. 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 9-19.

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 9-33.

Terminate shell

Mnemonic: T

Selecting this item terminates the dialogue. This is similar to using the **logout** command. See "logout" on page 7-28.

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 7-60) and whether there are any active processes, NightView may display a warning dialog box when you use the **Terminate shell** menu item.

New Shell Panel

Mnemonic: H

Selecting this item creates a new shell panel for the current shell. See "Source Panel" on page 9-48.

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.

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 9-25. 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 9-55.

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 9-26.

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 9-3.

This is similar to using the **detach** command. See “detach” on page 7-37.

Depending on the safety level (see “set-safety” on page 7-60), 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 9-3.

This is similar to using the **kill** command. See “kill” on page 7-37.

Depending on the safety level (see “set-safety” on page 7-60), NightView may display a warning dialog box when you use the Kill menu item.

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 9-32. See “Debugging the Heap” on page 3-28. See “Current Process” on page 9-3.

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 9-46. See “Current Process” on page 9-3.

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 9-48.

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 9-49.

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 9-48.

This dialog box is titled **Select a Function/Unit**. The title bar also displays the process's qualifier specifier. See “Qualifier Specifiers” on page 7-14. 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 7-16. 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 7-66).

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 9-26.

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 9-48.

This dialog box is titled **Select a Source File**. The title bar also displays the process's qualifier specifier. See “Qualifier Specifiers” on page 7-14. 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 7-18. 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 9-26.

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 9-48.

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 9-47 and “list” on page 7-72.

Edit

Mnemonic: I

Selecting this item lets you edit the current process’s current source. See “Source Panel” on page 9-48.

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.

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 9-49. 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 list 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.

Eventpoint Menu

Mnemonic: E

This menu provides ways to set eventpoints and to see a summary of eventpoints. See “Eventpoints” on page 3-9.

Before selecting one of the eventpoint menu items, select the line of interest in a source panel. See “Source Panel” on page 9-48. NightView uses this line to initialize the location specifier for you. See “Location Specifiers” on page 7-13.

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 7-92. See “breakpoint” on page 7-95.

For information on using the breakpoint dialog box, see “Eventpoint Dialog Boxes” on page 9-27.

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 7-92. See “monitorpoint” on page 7-101.

For information on using the monitorpoint dialog box, see “Eventpoint Dialog Boxes” on page 9-27.

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 7-92. See “patchpoint” on page 7-96.

For information on using the patchpoint dialog box, see “Eventpoint Dialog Boxes” on page 9-27.

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 7-92. See “tracepoint” on page 7-99.

For information on using the tracepoint dialog box, see “Eventpoint Dialog Boxes” on page 9-27.

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 7-92. See “heappoint” on page 7-102.

For information on using the heappoint dialog box, see “Eventpoint Dialog Boxes” on page 9-27.

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 7-92. See “watchpoint” on page 7-112.

For more information on using the watchpoint dialog box, see “Eventpoint Dialog Boxes” on page 9-27.

New Eventpoint Panel

Mnemonic: E

Selecting this menu item creates a new eventpoint panel. See “Eventpoint Panel” on page 9-57.

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 9-47.

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 9-60. 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 9-3. 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 9-65

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 9-72. See "Expression Data Item" on page 9-61.

Local Variables...

Mnemonic: L

A local-variables data item is placed in the data panel. See "Data Panel Add Local Variables" on page 9-73. See "Local Variables Data Item" on page 9-61.

Registers...

Mnemonic: R

A registers data item is placed in the data panel. See "Data Panel Add Registers" on page 9-74. See "Registers Data Item" on page 9-62.

Stack...

Mnemonic: S

A stack data item is placed in the data panel. See "Data Panel Add Stack" on page 9-75. See "Stack Data Item" on page 9-62.

Threads...

Mnemonic: T

A threads data item is placed in the data panel. See “Data Panel Add Threads” on page 9-75. See “Threads Data Item” on page 9-62.

Processes...

Mnemonic: C

A processes data item is placed in the data panel. See “Data Panel Add Processes” on page 9-74. See “Processes Data Item” on page 9-62.

Shells...

Mnemonic: H

A shells data item is placed in the data panel. See “Data Panel Add Shells” on page 9-74. See “Shells Data Item” on page 9-62.

Heap Information...

Mnemonic: H

A heap information data item is placed in the data panel. See “Data Panel Add Heap Information” on page 9-73. See “Heap Information Data Item” on page 9-62.

Heap Errors...

Mnemonic: P

A heap errors data item is placed in the data panel. See “Data Panel Add Heap Errors” on page 9-73. See “Heap Errors Data Item” on page 9-63.

Heap Leaks...

Mnemonic: K

A heap leaks data item is placed in the data panel. See “Data Panel Add Heap Leaks” on page 9-73. See “Leak Sets / Still Allocated Sets Data Items” on page 9-64.

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 9-74. See “Leak Sets / Still Allocated Sets Data Items” on page 9-64.

Monitorpoint Values

A monitorpoint values data item is placed in the data panel. See “Data Panel Add Monitorpoints” on page 9-74. See “Monitorpoint Values Data Item” on page 9-65.

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 9-77.

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 9-76.

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 9-76.

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 9-75.

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 9-76.

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.

NightTrace Analyzer

Mnemonic: T

Opens the NightTrace Analyzer.

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.

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 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 9-2.

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 7-127.

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 9-25.

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 5 [Tutorial - GUI] on page 5-1.

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 NightProbe, NightView, NightTrace, and NightTune in one complete example.

License Report...

Mnemonic: L

Opens a dialog box with information about how many licenses are in use and how to get licenses.

On Version

Mnemonic: V

This item pops up an information dialog box that describes which version of Night-View you are running.

Toolbars

Most of the controls in the toolbars apply to the current process. See “Current Process” on page 9-3. 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 9-3.

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 

Clicking on this button is similar to using the **resume** command with no argument. See “resume” on page 7-115.

Stop 


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 7-120.

Next 


Clicking on this button is similar to using the **next** command with no argument. See “next” on page 7-117.

Step 

Clicking on this button is similar to using the **step** command with no argument. See “step” on page 7-116.

Finish 

Clicking on this button is similar to using the **finish** command. See “finish” on page 7-119.

Run to Here 

Run the process until it reaches the target line in the source panel. See “Source Panel Target Line” on page 9-49. See “Source Panel” on page 9-48. This allows you to use the Run to Here button 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 7-95. See “enable” on page 7-108. See “resume” on page 7-115.

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 9-56.

Nexti 


Clicking on this button is similar to using the **nexti** command with no argument. See “nexti” on page 7-119.

Stepi 

Clicking on this button is similar to using the **stepi** command with no argument. See “stepi” on page 7-118.

Up 

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. See “up” on page 7-125.

Down 


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. See “down” on page 7-126.

Kill 


Clicking this button causes NightView to terminate the current process. See “Current Process” on page 9-3.

This is similar to using the **kill** command. See “kill” on page 7-37.

Depending on the safety level (see “set-safety” on page 7-60), NightView may display a warning dialog box when you use the Kill menu item.

Rerun 


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 7-126.

Interrupt 

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 3-34.


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 9-49.

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 7-92. See “breakpoint” on page 7-95.

For information on using the breakpoint dialog box, see “Eventpoint Dialog Boxes” on page 9-27.

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 7-92. See “patchpoint” on page 7-96.

For information on using the patchpoint dialog box, see “Eventpoint Dialog Boxes” on page 9-27.

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 7-92. See “monitorpoint” on page 7-101.

For information on using the monitorpoint dialog box, see “Eventpoint Dialog Boxes” on page 9-27.

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 7-92. See “tracepoint” on page 7-99.

For information on using the tracepoint dialog box, see “Eventpoint Dialog Boxes” on page 9-27.

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 7-92. See “heappoint” on page 7-102.

For information on using the heappoint dialog box, see “Eventpoint Dialog Boxes” on page 9-27.

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 7-92. See “watchpoint” on page 7-112.

For more information on using the watchpoint dialog box, see “Eventpoint Dialog Boxes” on page 9-27.

Clear

Clicking on this button is similar to using the **clear** command for the target line. See “clear” on page 7-104. You must have set the target line to the source line where you want to clear eventpoints. See “Source Panel Target Line” on page 9-49. 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 9-56.

Value Toolbar

Data Display

Clicking on this button is similar to using the **Expression...** button in the **Data** menu. See “Data Menu” on page 9-14. You must have selected an expression in a source panel (or another panel) before pressing this button. See “Source Panel” on page 9-48. When you press the button, the selected expression is added to the default data panel. See “Data Panel” on page 9-60.

Print

Clicking on this button is similar to using the **print** command. See “print” on page 7-81. You must have selected an expression in a source panel (or another panel) before pressing this button. See “Source Panel” on page 9-48. 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 9-10.

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 9-3. See “Qualifier Specifiers” on page 7-14. For threaded programs, the process ID in the qualifier indicates which of the thread processes last stopped. See “Multithreaded Programs” on page 3-37.

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 3-18.

Calling function

The process is executing to evaluate a function call.

Exited

The process has exited. See “Exited and Terminated Processes” on page 3-18. 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 7-119.

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 3-2.

Running

The process is currently executing.

Stepping

The process is executing because of a stepping command. See “step” on page 7-116.

Stopped after finish

The process has completed a **finish** command. See “finish” on page 7-119.

Stopped after step

The process has finished a stepping command. See “step” on page 7-116.

Stopped at breakpoint *number*

The process hit breakpoint number *number*. See “Breakpoints” on page 3-11.

Stopped at watchpoint *number*

The process stopped because of watchpoint *number*. See “Watchpoints” on page 3-13.

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 3-13. See “Message Panel” on page 9-56.

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 3-3.

Stopped by user

The process stopped because of a **stop** command. See “stop” on page 7-120.

Stopped for exec

The process has just `exec()`'ed a new program image. See “Programs and Processes” on page 3-2.

Stopped with *signal*

The process stopped with signal *signal*. See “Signals” on page 3-14.

Terminated with *signal*

The process terminated with signal *signal*. See “Exited and Terminated Processes” on page 3-18. This status appears only for core files. See “Core Files” on page 3-4.

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 9-55. There are extra keys you can type into a source panel. See “Source Panel Keystrokes” on page 9-54.

Ctrl+A	New page. See “View Menu” on page 9-5.
Ctrl+B	Breakpoint dialog. See “Eventpoint Menu” on page 9-12.
Ctrl+C	Copy in a text field.
Ctrl+D	Show disassembly. See “Source Menu” on page 9-10.
Ctrl+E	Add expression to data panel. See “Data Menu” on page 9-14.
Ctrl+F	Find. See “Find Bar” on page 9-48.
Ctrl+G	Find again. See “Find Bar” on page 9-48.
Ctrl+M	Show mixed source and disassembly. See “Source Menu” on page 9-10.
Ctrl+O	Show source. See “Source Menu” on page 9-10.
Ctrl+P	Patchpoint dialog. See “Eventpoint Menu” on page 9-12.
Ctrl+Q	Exit the debugger. See “File Menu” on page 9-4.
Ctrl+S	Save config. See “File Menu” on page 9-4.
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 9-5.
Ctrl+-	Decrease font size. See “View Menu” on page 9-5.
Ctrl+0 (zero)	Restore default font size. See “View Menu” on page 9-5.

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 9-72.

Run Program in Shell Dialog Box

This dialog box pops up when you click **Run...** in the **Process** menu. See “Process Menu” on page 9-8. 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 3-5.

Attach Dialog Box

This dialog box pops up when you select the **Attach...** item in the **Process** menu. See “Process Menu” on page 9-8. This dialog box lets you view the processes on the system and select one or more to attach to.

Check the box to indicate whether you want the processes to 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 7-68 and “Preferences Dialog Box” on page 9-36.

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 7-16. 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 9-10. 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 7-16. See “Wildcard Patterns” on page 7-18.

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 9-56.

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 9-48. 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 3-9. 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 9-12. You decide how you want the eventpoint set or changed, then select the OK button and NightView will set or modify it for you.

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.

For *inserted eventpoints*, NightView provides default settings for new eventpoints, including a default location specifier. See “Location Specifiers” on page 7-13. 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 9-49. 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.

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.

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 7-108. See “disable” on page 7-108.

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 7-110.

For patchpoints, this is similar to using the **tpatch** command, or the **enable/once** command. See “tpatch” on page 7-111.

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 7-106.

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 7-109.

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 7-94.

You cannot change an existing eventpoint's name using the dialog box. Use the **name** command to change eventpoint names.

Commands

Valid for breakpoints 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 7-105.

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.

Evaluate Expression - Go To Location

Valid for patchpoints only; you are *required* to enter either an expression or a location specifier to set a patchpoint. Select one of the two choices by clicking on it. The radio button appears filled for your selection, and the label for the text input area changes to either **Evaluate** or **Go to**. 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 7-96. This is the default choice.

Branch to a different location

This field represents the *goto* argument of one variant of the **patch-point** command.

Once set, this field cannot be changed.

Event ID

Valid for tracepoints only; *required* 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 7-99.

Once set, this field cannot be changed.

Value

Valid for tracepoints only. This 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, this field cannot be changed.

Check - Debug Settings

Valid 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 7-50.

Choose an action button.

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 9-56.

You can use the **info eventpoint** command or the eventpoint panel to check the eventpoint settings. See “info eventpoint” on page 7-131. See “Eventpoint Panel” on page 9-57.

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 9-8). 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 7-50).

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 7-50 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 3-29.

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 3-29.

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 3-30.

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 3-31.

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 7-63).

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 3-30.

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 9-8. This dialog box allows you to specify the parameters for creating a remote NightView session. See “Remote Dialogues” on page 3-6. 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 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 7-14. 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.

Indicate with the checkbox whether you want a new shell panel to be created for this shell. See “Shell Panel” on page 9-55.

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 Bias

These check-boxes allow you to set the CPU affinity of the NightProbe server process. Checking a box means that the process is allowed to run on the logical CPU associated with that number. At least one box must be checked. Leaving all boxes checked will allow the process to run on any available CPU.

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 disap-

pears and a message is displayed in the message panel. See “Message Panel” on page 9-56.

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 9-4. 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 9-46.

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

Safety

Select a safety level. This is similar to using **set-safety**. See “set-safety” on page 7-60. 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. Check the conditions for which you want the process to continue running. See “set-resume” on page 7-68.

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 7-66.

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. See “Data Panel Call Stack Frames” on page 9-75 and “Data Panel Pointer Array Dimension” on page 9-76.

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 7-58.

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 9-54.

Preferences Appearance Page

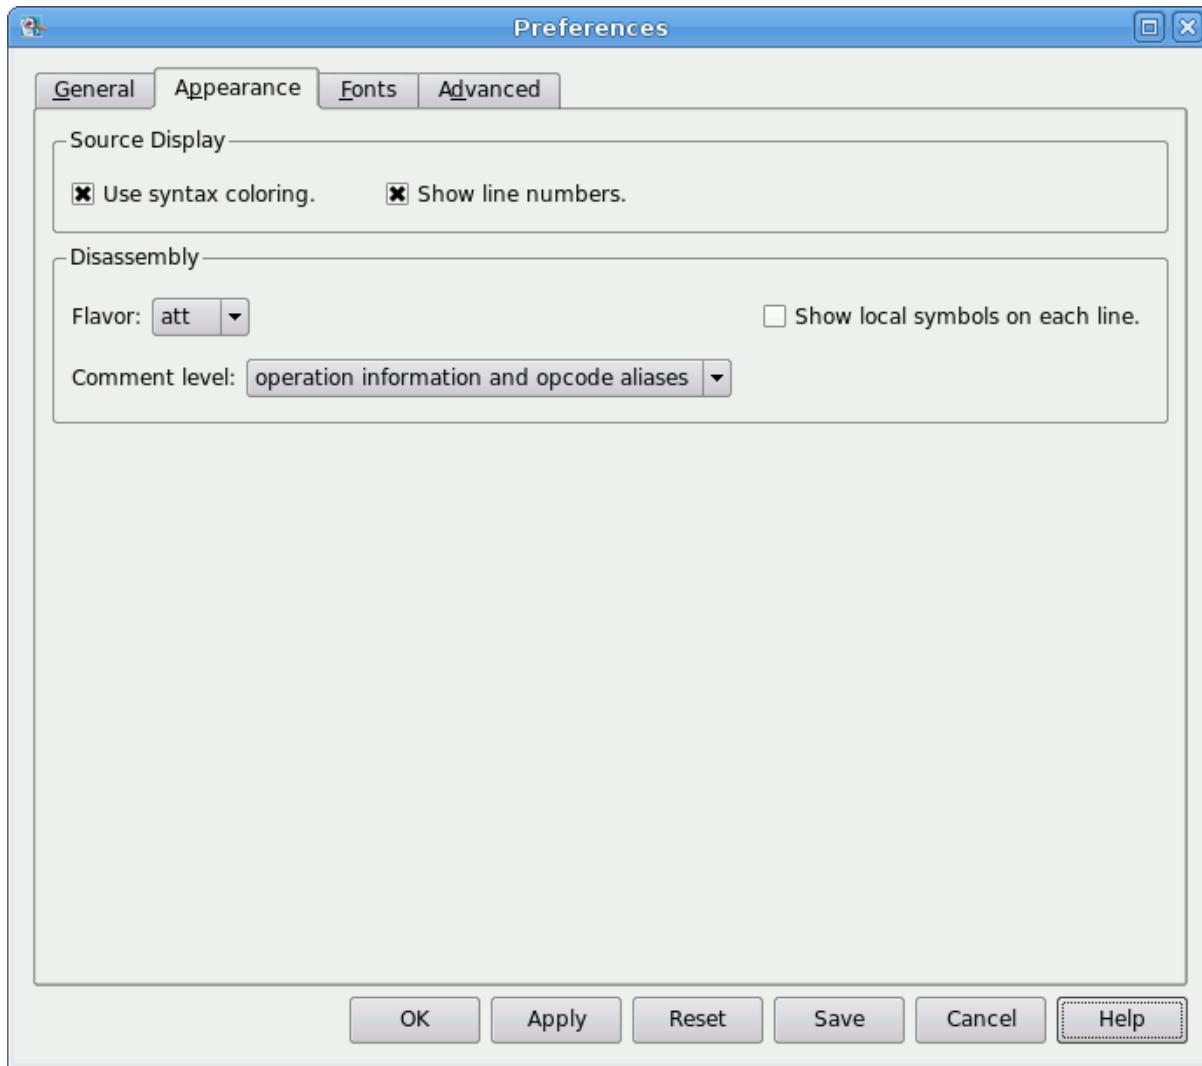


Figure 9-1. Preferences Appearance Page

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 9-48.

Disassembly

Choose how you want to see disassembled code. See “set-disassembly” on page 7-69.

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

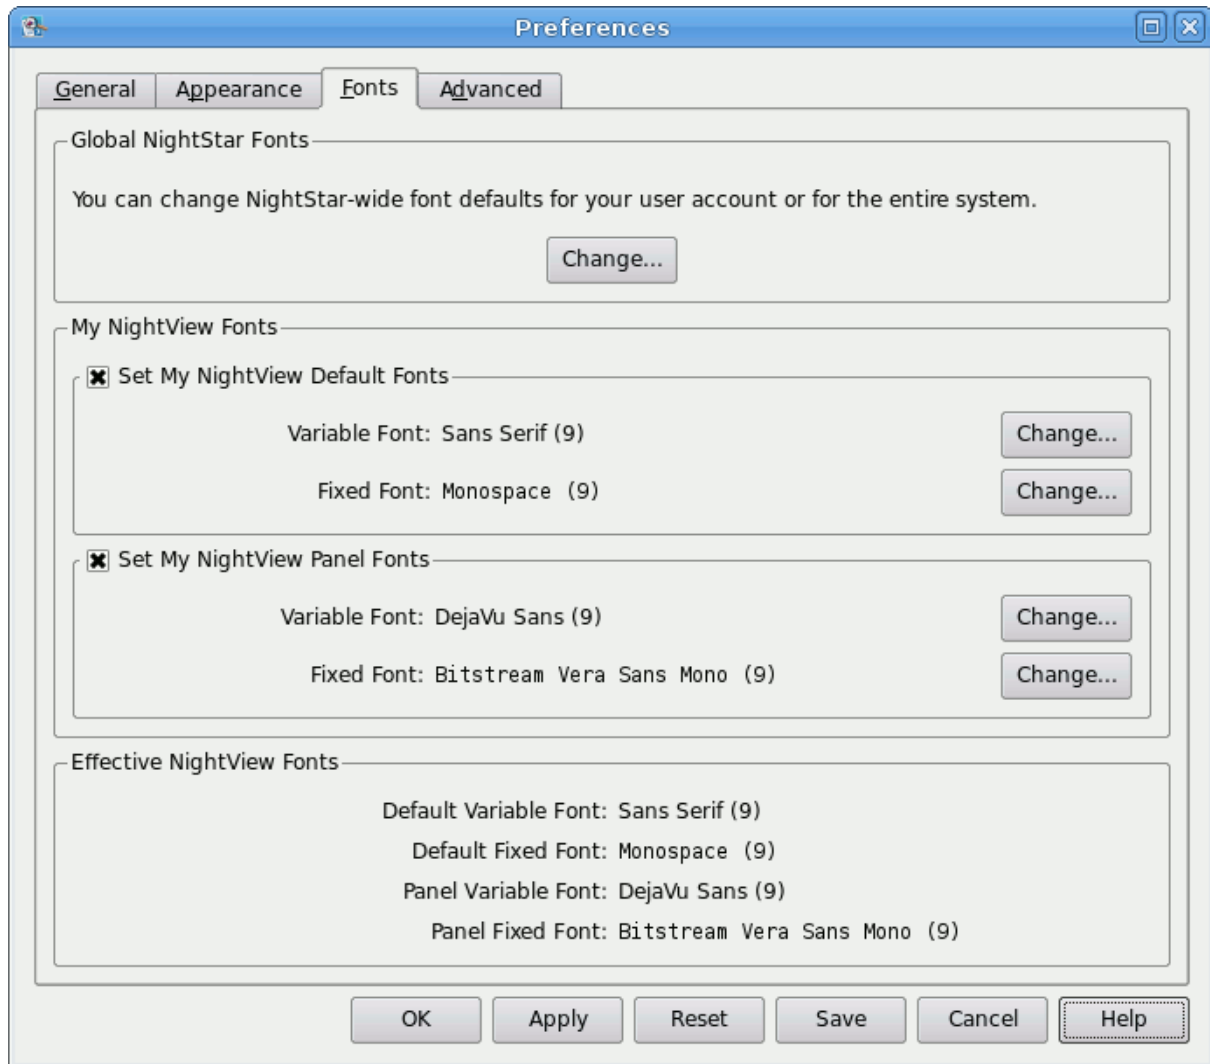


Figure 9-2. Preferences Fonts Page

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 9-42 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 9-3. 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 Advanced Page

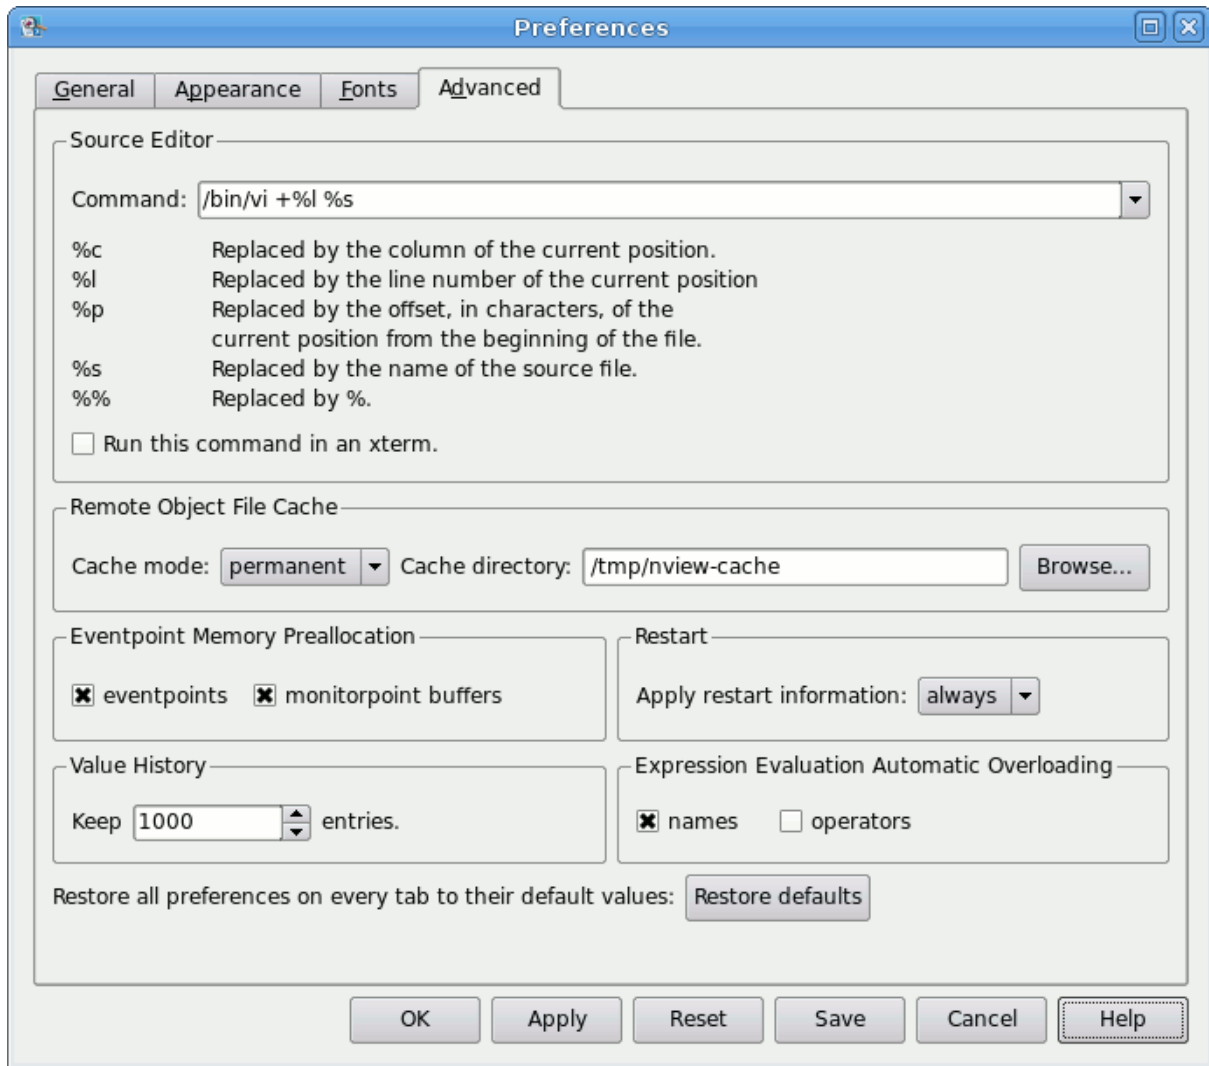


Figure 9-4. Preferences Advanced Page

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.

If your editor can communicate with the X Window System display directly, then you should clear the check box for **Run this command in an xterm**. The editor runs on the same display as NightView.

Remote Object File Cache

Control how NightView downloads files from remote targets. See “set-download” on page 7-68.

Eventpoint Memory Preallocation

Control how NightView preallocates memory for eventpoints and monitorpoint buffers. See “set-preallocate” on page 7-67.

Restart

Control whether restart information is applied. See “set-restart” on page 7-61.

Value History

Specify the number of items to be kept in the value history list. See “set-history” on page 7-58.

Expression Evaluation Automatic Overloading

Control how NightView treats overloaded operators and routines in expressions. See “set-overload” on page 7-66.

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.

Process Settings Dialog Box

This dialog box pops up when you click on **Process Settings...** in the **Process** menu. See “Process Menu” on page 9-8. 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 9-36.

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 7-47.

Stop Before Exiting

Control whether this process stops before exiting. See “set-exit” on page 7-47.

Expression Language

Establish a default language context for variables and expressions. See “set-language” on page 7-56.

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 7-40.

Process Settings Interest Page

Control which subprograms are interesting. See “interest” on page 7-63.

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 7-122.

Rename Page Dialog Box

This dialog box pops up when you click **Rename Current Page...** in the **View** menu. See “View Menu” on page 9-5.

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 7-72. If you clicked on **List Location...** in the **Source** menu, any source panels displaying the current process are affected (see “Source Menu” on page 9-10). 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 9-50).

If desired, you can specify a file on another host with the form *user@host:/path*. See “Remote File Access” on page 3-7.

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 9-80.

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 9-4.

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 9-3. 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 7-16. 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 9-36.

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 3-24. See “Current Process” on page 9-3. You can select various display modes to display source or disassembly with the **Source** menu. See “Source Menu” on page 9-10. See “list” on page 7-72, 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 7-78. You can turn off the line numbers with the preferences dialog box. See “File Menu” on page 9-4. 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 7-72), 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 9-54. See “Eventpoint Menu” on page 9-12. See “Source Panel Context Menu” on page 9-50.

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 9-10.

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 9-3. 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 9-49.) 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 3-24. 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 9-48.

Note that the right-click also sets the target line. See "Source Panel Target Line" on page 9-49.

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

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 9-48.

This dialog box is titled **Select a Function/Unit**. The title bar also displays the process's qualifier specifier. See “Qualifier Specifiers” on page 7-14. 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 7-16. 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 7-66).

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 9-26.

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 9-48.

This dialog box is titled **Select a Source File**. The title bar also displays the process's qualifier specifier. See “Qualifier Specifiers” on page 7-14. 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 7-18. For example, enter `mod* .c` to search for source file names that start with 'mod' followed by any num-

ber 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 9-26.

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 9-48.

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 9-47 and “list” on page 7-72.

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 9-48.

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 9-48. 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 list 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.

Track any process

Mnemonic: Y

The panel changes whenever another process becomes the current process. See “Source Panel Tracking” on page 9-54.

Track one process

Mnemonic: N

The panel ignores changes for other processes. See “Source Panel Tracking” on page 9-54.

Panel locked

Mnemonic: K

The panel ignores changes for all processes. See “Source Panel Tracking” on page 9-54.

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 9-3.

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 9-59. See “Source Panel Context Menu” on page 9-50. 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 9-49. 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 9-25.

b	This key sets a breakpoint. It performs the same action as the Breakpoint button (see “Process Toolbar” on page 9-19).
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 7-126.
Enter	Sets the line with the text cursor to be the target line. See “Source Panel Target Line” on page 9-49.
e	This key is similar to selecting the Edit item in the Source menu. See “Source Menu” on page 9-10. 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 7-119.
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 9-19. 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 7-119.
n	Step one line without entering called routines. This key is similar to using the next command with no argument. See “next” on page 7-117.
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 9-19).
r	Resume the process. This key is similar to using the resume command with no argument. See “resume” on page 7-115.
S	Step one instruction, entering a called routine. This key is similar to using the stepi command with no argument. See “stepi” on page 7-118.
s	Step one line, entering any called routine. This key is similar to using the step command with no argument. See “step” on page 7-116.
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 7-125.
=	Move to the newest stack frame. This key is similar to using the frame 0 command. See “frame” on page 7-124.
>	Print information about the current frame. This key is similar to using the frame command with no arguments. See “frame” on page 7-124.

Shell Panel

This area allows you to interact with the dialogue shell and with your programs. See “Dialogues” on page 3-4. 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 9-8. See "run" on page 7-34.

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 9-48.

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 9-48.

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 3-9.

The table has a row for each eventpoint. The columns show:

- the type of the eventpoint (breakpoint, heappoint, monitorpoint, patchpoint, tracepoint 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 do not have a location in the program: for watchpoints, this area describes the location being watched.
- 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 and watchpoints can have commands).
- any condition on the eventpoint.

If the program is stopped at a breakpoint or watchpoint, the eventpoint icon is overlaid with a green triangle pointing to the right,  as it is in the source panel.

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** 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 9-27.

Enable

Mnemonic: E

Enable the selected eventpoints. This is similar to using the **enable** command. See "enable" on page 7-108.

Disable

Mnemonic: D

Disable the selected eventpoints. This is similar to using the **disable** command. See "disable" on page 7-108.

Delete

Mnemonic: L

Delete the selected eventpoints. This is similar to using the **delete** command. See "delete" on page 7-107.

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 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 9-56.

You can also use the **info eventpoint** command to check eventpoint settings. See “info eventpoint” on page 7-131.

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 9-60. Processes appear in the context panel only if you are debugging more than one process. Similarly, shells appear in the context panel only if you have more than one shell.

The current stack frame 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 9-60. See “Local Variables Data Item” on page 9-61. See “Current Frame” on page 3-24. See “Current Process” on page 9-3.

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 9-60. See “Monitorpoint Values Data Item” on page 9-65. See “Current Frame” on page 3-24. See “Current Process” on page 9-3.

Data Panel

A data panel displays various information about your process. 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 9-14. See “Value Toolbar” on page 9-22. See “Data Panel Context Menu” on page 9-65. See “data-display” on page 7-88.

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 9-65.

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 7-103.

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.

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 9-65.

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 9-14.

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 9-61 and “Stack Data Item” on page 9-62.

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 9-72.

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.

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.)

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 9-64.

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 9-19.

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 9-19.

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 9-14. 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 3-24.

Threads Data Item

A threads data item has sub-items to describe C threads or thread processes. Expanding a sub-item shows the stack for that thread. If the program is not threaded, then it is considered to have one thread. The current thread is shown in green underlined text.

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.

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 3-32). 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 3-32 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 3-11.

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 9-14. 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 9-78.

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 9-65. See "Preferences General Page" on page 9-36. 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.

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 9-75.

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 9-76.

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 9-77.

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 9-78. If you already have a condition filter, this menu entry changes to **Change Condition Filter...**

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 3-24.

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 7-72.

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 3-32.

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 3-32.

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 3-32.

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 9-75.

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 7-81. 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 7-81.

Character

Mnemonic: C

The value is displayed in character format. See “print” on page 7-81.

Decimal

Mnemonic: D

The value is displayed in decimal format. See “print” on page 7-81.

Float

Mnemonic: F

The value is displayed in float format. See “print” on page 7-81.

Hexadecimal

Mnemonic: H

The value is displayed in hexadecimal format. See “print” on page 7-81.

Octal

Mnemonic: O

The value is displayed in octal format. See “print” on page 7-81.

String

Mnemonic: S

The value is displayed in string format. See “print” on page 7-81.

Unsigned Decimal

Mnemonic: U

The value is displayed in unsigned decimal format. See “print” on page 7-81.

Find...

Mnemonic: F

Accelerator: Ctrl+F

Brings up the find bar and sets the keyboard focus there. See “Find Bar” on page 9-48.

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 9-14.

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 9-72. See “Expression Data Item” on page 9-61.

New Local Variables...

A local-variables data item is placed in the data panel. See “Data Panel Add Local Variables” on page 9-73. See “Local Variables Data Item” on page 9-61.

New Registers...

A registers data item is placed in the data panel. See “Data Panel Add Registers” on page 9-74. See “Registers Data Item” on page 9-62.

New Stack...

A stack data item is placed in the data panel. See “Data Panel Add Stack” on page 9-75. See “Stack Data Item” on page 9-62.

New Threads...

A threads data item is placed in the data panel. See “Data Panel Add Threads” on page 9-75. See “Threads Data Item” on page 9-62.

New Processes...

A processes data item is placed in the data panel. See “Data Panel Add Processes” on page 9-74. See “Processes Data Item” on page 9-62.

New Shells...

A shells data item is placed in the data panel. See “Data Panel Add Shells” on page 9-74. See “Shells Data Item” on page 9-62.

New Heap Information...

A heap information data item is placed in the data panel. See “Data Panel Add Heap Information” on page 9-73. See “Heap Information Data Item” on page 9-62.

New Heap Errors...

A heap errors data item is placed in the data panel. See “Data Panel Add Heap Errors” on page 9-73. See “Heap Errors Data Item” on page 9-63.

New Heap Leaks...

A heap leaks data item is placed in the data panel. See “Data Panel Add Heap Leaks” on page 9-73. See “Leak Sets / Still Allocated Sets Data Items” on page 9-64.

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 9-74. See “Leak Sets / Still Allocated Sets Data Items” on page 9-64.

New Monitorpoint Values

A monitorpoint values data item is placed in the data panel. See “Data Panel Add Monitorpoints” on page 9-74. See “Monitorpoint Values Data Item” on page 9-65.

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 9-77.

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 9-76.

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 9-76.

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 9-75.

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 9-76.

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 7-103. 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 9-80.

Show Monitor Bar

Clicking on this button causes the monitor bar to be shown. See “Monitor Bar” on page 9-60.

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 9-25.

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 9-14.

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 9-72.

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 9-14.

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 9-63.

This is a Data Panel Item Dialog Box. See "Data Panel Item Dialog Box" on page 9-72.

Data Panel Add Heap Information

This dialog box pops up when you click on **Heap Information...** in the **Data** menu or in the data panel context menu. See "Data Menu" on page 9-14.

Click on **OK** to add a heap information data item to the data panel. See "Heap Information Data Item" on page 9-62.

This is a Data Panel Item Dialog Box. See "Data Panel Item Dialog Box" on page 9-72.

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 9-14.

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 be displayed. If **All** is selected, then all leaks will be displayed. See "Leak Sets / Still Allocated Sets Data Items" on page 9-64.

This is a Data Panel Item Dialog Box. See "Data Panel Item Dialog Box" on page 9-72.

Data Panel Add Local Variables

This dialog box pops up when you click on **Local Variables...** in the **Data** menu or in the data panel context menu if there is more than one data panel. See "Data Menu" on page 9-14.

Click on **OK** to add a local variables data item to the data panel.

This is a Data Panel Item Dialog Box. See “Data Panel Item Dialog Box” on page 9-72.

Data Panel Add Monitorpoints

This dialog box pops up when you click on **Monitorpoints...** in the **Data** menu or in the data panel context menu if there is more than one data panel. See “Data Menu” on page 9-14.

Click on **OK** to add a monitorpoints data item to the data panel.

This is a Data Panel Item Dialog Box. See “Data Panel Item Dialog Box” on page 9-72.

Data Panel Add Processes

This dialog box pops up when you click on **Processes...** in the **Data** menu or in the data panel context menu if there is more than one data panel. See “Data Menu” on page 9-14.

Click on **OK** to add a processes data item to the data panel.

This is a Data Panel Item Dialog Box. See “Data Panel Item Dialog Box” on page 9-72.

Data Panel Add Registers

This dialog box pops up when you click on **Registers...** in the **Data** menu or in the data panel context menu if there is more than one data panel. See “Data Menu” on page 9-14.

Click on **OK** to add a registers data item to the data panel.

This is a Data Panel Item Dialog Box. See “Data Panel Item Dialog Box” on page 9-72.

Data Panel Add Shells

This dialog box pops up when you click on **Shells...** in the **Data** menu or in the data panel context menu if there is more than one data panel. See “Data Menu” on page 9-14.

Click on **OK** to add a shells data item to the data panel.

This is a Data Panel Item Dialog Box. See “Data Panel Item Dialog Box” on page 9-72.

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 9-14.

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 be displayed. If **All** is selected, then all blocks still allocated will be displayed. See “Leak Sets / Still Allocated Sets Data Items” on page 9-64.

This is a Data Panel Item Dialog Box. See “Data Panel Item Dialog Box” on page 9-72.

Data Panel Add Stack

This dialog box pops up when you click on **Stack...** in the **Data** menu or in the data panel context menu if there is more than one data panel. See “Data Menu” on page 9-14.

Click on **OK** to add a stack data item to the data panel.

This is a Data Panel Item Dialog Box. See “Data Panel Item Dialog Box” on page 9-72.

Data Panel Add Threads

This dialog box pops up when you click on **Threads...** in the **Data** menu or in the data panel context menu if there is more than one data panel. See “Data Menu” on page 9-14.

Click on **OK** to add a threads data item to the data panel.

This is a Data Panel Item Dialog Box. See “Data Panel Item Dialog Box” on page 9-72.

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 9-14.

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 9-65.

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 9-72.

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 9-65.

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 9-65.

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 9-14.

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 9-14.

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 9-14.

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 9-14. 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 9-65.

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 9-65).

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 3-35), `$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 3-35.

```
$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 9-65). 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 9-19.

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 9-59. See “Data Panel” on page 9-60.

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 7-103.

The monitorpoint update interval is not related to the eventpoint panel update interval. See “Eventpoint Panel Update Interval Dialog Box” on page 9-47.

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 9-1.

NightStar LX Licensing

NightStar LX uses the NightStar License Manager (NSLM) to control access to the NightStar LX tools.

License installation requires a licence key provided by Concurrent (see “License Keys” on page A-1). The NightStar LX tools request a licence (see “License Requests” on page A-2) from a license server (see “License Server” on page A-2).

Two license modes are available, fixed and floating, depending on which product option you purchased. Fixed licenses can only be served to NightStar LX users from the local system. Floating licenses may be served to any NightStar LX user on any system on a network.

Tools are licensed per system, per concurrent user. A single license is shared among any or all of the NightStar LX tools for a particular user on a particular system. The intent is to allow n developers to fully utilize all the tools at the same time while only requiring n licenses. When operating the tools in remote mode, where a tool is launched on a local system but is interacting with a remote system, licenses are required only from the host system.

You can obtain a license report which lists all licenses installed on the local system, current usage, and expiration date for demo licenses (see “License Reports” on page A-3).

The default configuration includes a strict firewall which interferes with floating licenses. See “Firewall Configuration for Floating Licenses” on page A-3 for information on handling such configurations.

See “License Support” on page A-6 for information on contacting Concurrent for additional assistance with licensing issues.

License Keys

Licenses are granted to specific systems to be served to either local or remote clients, depending on the license model, fixed or floating.

License installation requires a license key provided by Concurrent. To obtain a license key, you must provide your system identification code. The system identification code is generated by the `nslm_admin` utility:

```
nslm_admin --code
```

System identification codes are dependent on system configurations. Reinstalling Linux on a system or replacing network devices may require you to obtain new license keys.

To obtain a license key, use the following URL and click on the [Licenses](#) link:

<http://www.ccur.com/NightStarKeys>

Provide the requested information, including the system identification code. Your license key will be immediately emailed to you.

Install the license key using the following command:

```
nslm_admin --install=xxx-xxx-xxx-xxx-xxx
```

where *xxx-xxx-xxx-xxx-xxx* is the key included in the license acknowledgment email.

License Requests

By default, the NightStar LX tools request a license from the local system. If no licenses are available, they broadcast a license request on the local subnet associated with the system's hostname.

You can control the license requests for an entire system using the `/etc/nslm.config` configuration file.

By default, the `/etc/nslm.config` file contains a line similar to the following:

```
:server @default
```

The argument `@default` may be changed to a colon-separated list of system names, system IP addresses, or broadcast IP addresses. Licenses will be requested from each of the entities found in the list, until a license is granted or all entries in the list are exhausted.

For example, the following setting prevents broadcast requests for licenses, by only specifying the local system:

```
:server localhost
```

The following setting requests a license from `server1`, then `server2`, and then a broadcast request if those fail to serve a license:

```
:server server1:server2:192.168.1.0
```

Similarly, you can control the license requests for individual invocations of the tools using the `NSLM_SERVER` environment variable. If set, it must contain a colon-separated list of system names, system IP addresses, or broadcast IP addresses as described above. Use of the `NSLM_SERVER` environment variable takes precedence over settings defined in `/etc/nslm.config`.

License Server

The NSLM license server is automatically installed and configured to run when you install NightStar LX.

The `nslm` service is automatically activated for run levels 2, 3, 4, and 5. You can check on these settings by issuing the following command:

```
/sbin/chkconfig --list nslm
```

In rare instances, you may need to restart the license server via the following command:

```
/sbin/service nslm restart
```

See `nslm(1)` for more information.

License Reports

A license report can be obtained using the `nslm_admin` utility.

```
nslm_admin --list
```

lists all licenses installed on the local system, current usage, and expiration date (for demo licenses). Use of the `--verbose` option also lists individual clients to which licenses are currently granted.

Adding the `--broadcast` option will list this information for all servers that respond to a broadcast request on the local subnet associated with the system's hostname.

See `nslm_admin(1)` for more options and information.

Firewall Configuration for Floating Licenses

The default operating system configuration includes a strict firewall which interferes with floating licenses.

If such a system is used to serve licenses, then at least one port must be opened in its firewall to allow server requests to pass. See "Serving Licenses with a Firewall" on page A-3 for more information.

Similarly, if such a system is host to the NightStar LX tools, then at least one port must be opened in its firewall so that it can receive licenses from the license server. If this is not done, a tool requesting a floating license will not receive it and will not function properly. See "Running NightStar LX Tools with a Firewall" on page A-5 for more information.

Serving Licenses with a Firewall

Following are a few approaches for allowing the NSLM license server to serve floating licences when the system on which it is running is configured with a firewall:

- disable the firewall on the system entirely

- allow NSLM license requests from a specific system (or one of several)
- allow NSLM license requests from any system on a particular subnet (or one of several)
- allow NSLM license requests from any system

NOTE

You must be root in order to modify the firewall configuration.

To disable the firewall entirely, execute:

```
service iptables stop
```

and then remove the `/etc/sysconfig/iptables` file:

```
rm -f /etc/sysconfig/iptables
```

This option may not be as dangerous as it seems. Often, whole networks are protected with a firewall so it is not necessary for individual systems on the network to be protected further. If unsure, check with your network administrator.

For the remaining cases, a simple modification should be made to the `/etc/sysconfig/iptables` file to allow license requests on UDP port 25517 and TCP port 25517. By default, that file should contain a line like the following:

```
-A RH-Firewall-1-INPUT -j REJECT --reject-with icmp-host-prohibited
```

To allow NSLM license requests from a specific system, insert the following lines before the REJECT line:

```
-A RH-Firewall-1-INPUT -p udp -m udp -s system --dport 25517 -j ACCEPT  
-A RH-Firewall-1-INPUT -p tcp -m tcp -s system --dport 25517 -j ACCEPT
```

Those lines can be repeated for multiple systems.

To allow NSLM license requests from any system on a particular subnet, insert the following lines before the REJECT line:

```
-A RH-Firewall-1-INPUT -p udp -m udp -s subnet/mask --dport 25517 -j ACCEPT  
-A RH-Firewall-1-INPUT -p tcp -m tcp -s subnet/mask --dport 25517 -j ACCEPT
```

The subnet might be of a form like `192.168.1.0` and the mask could be a traditional network mask like `255.255.255.0` or a single number like `24`, which indicates the number of bits from the left that are part of the mask. For example, `192.168.1.0/255.255.255.0` and `192.168.1.0/24` are equivalent.

Those lines can be repeated for multiple subnets.

To allow NSLM license requests from any system, insert the following lines before the REJECT line:

```
-A RH-Firewall-1-INPUT -p udp -m udp --dport 25517 -j ACCEPT
-A RH-Firewall-1-INPUT -p tcp -m tcp --dport 25517 -j ACCEPT
```

After modifying `/etc/sysconfig/iptables`, execute:

```
service iptables restart
```

Running NightStar LX Tools with a Firewall

Following are a few approaches for allowing a NightStar LX tool to receive floating licenses from a license server, when the system running the NightStar LX tool is configured with a firewall:

- disable the firewall on the requesting system entirely
- allow NSLM licenses from a specific license server (or one of several)
- allow NSLM licenses from any system on a particular subnet (or one of several)
- allow NSLM licenses from any system

NOTE

You must be root in order to modify the firewall configuration.

To disable the firewall entirely, execute:

```
service iptables stop
```

and then remove the `/etc/sysconfig/iptables` file:

```
rm -f /etc/sysconfig/iptables
```

This option may not be as dangerous as it seems. Often, whole networks are protected with a firewall so it is not necessary for individual systems on the network to be protected further. If unsure, check with your network administrator.

For the remaining cases, a simple modification should be made to the `/etc/sysconfig/iptables` file to allow license responses from UDP port 25517. By default, that file should contain a line like the following:

```
-A RH-Firewall-1-INPUT -j REJECT --reject-with icmp-host-prohibited
```

To allow NSLM licenses from a specific system running a license server, insert the following line before the REJECT line:

```
-A RH-Firewall-1-INPUT -p udp -m udp -s server --sport 25517 -j ACCEPT
```

That line can be repeated for multiple servers.

To allow NSLM licenses from any system running a license server on a particular subnet, insert the following before the REJECT line:

```
-A RH-Firewall-1-INPUT -p udp -m udp -s subnet/mask --sport 25517 -j ACCEPT
```

The subnet might be of a form like 192.168.1.0 and the mask could be a traditional network mask like 255.255.255.0 or a single number like 24, which indicates the number of bits from the left that are part of the mask. For example, 192.168.1.0/255.255.255.0 and 192.168.1.0/24 are equivalent.

That line can be repeated for multiple subnets.

To allow NSLM licenses from any system running a license server, insert the following line before the REJECT line:

```
-A RH-Firewall-1-INPUT -p udp -m udp --sport 25517 -j ACCEPT
```

After modifying `/etc/sysconfig/iptables`, execute:

```
service iptables restart
```

License Support

For additional aid with licensing issues, contact the Concurrent Software Support Center at our toll free number 1-800-245-6453. For calls outside the continental 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 request for assistance at any time by using the Concurrent Computer Corporation web site at http://www.ccur.com/isd_support_contact.asp or by sending an email to support@ccur.com.

Kernel Dependencies

Additional features and benefits are granted the NightStar tools when operating with real-time kernels from Concurrent Computer Corporation and Novell, Inc.

RedHawk Linux

The RedHawk Linux kernel from Concurrent Computer Corporation (<http://www.ccur.com>) provides real-time capabilities over and above the Red Hat kernel. RedHawk Linux operates on top of a Red Hat distribution.

SUSE Linux Enterprise Real-Time (SLERT)

The SLERT kernel from Novell, Inc. (<http://www.novell.com>) provides real-time capabilities over and above other SUSE Linux Enterprise kernels. SLERT is also available from Concurrent Computer Corporation (<http://www.ccur.com>).

The following sections describe the additional functionality and capabilities of the NightStar tools when running these real-time kernels.

Advantages for NightView

The following advantages are afforded NightView when a RedHawk or SLERT kernel is running:

- Application speed conditions

Provides “execution-speed” patches, conditions, and ignore counts.

- Signal handling

Allows NightView to pass signals directly to a particular process, avoiding context switching.

NOTE

NightView may not operate at all on older versions of SUSE or Red Hat without a real-time kernel.

Advantages for NightTrace

The following advantage is afforded NightTrace when a RedHawk or SLERT tracing kernel is running:

- Kernel tracing

Users of NightTrace gain the ability to obtain kernel trace data and combine that with user trace data. Kernel tracing is an incredibly powerful feature that not only provides insight into the operating system kernel but also provides useful information relating to the execution of user applications.

The real-time kernel is provided in three flavors:

- Tracing
- Debug
- Plain

The Tracing and Debug flavors provide the features required for NightTrace kernel tracing. These kernels can be selected at boot-time from the boot-loader menu.

Advantages for NightProbe

The following advantages are afforded NightProbe when a RedHawk or SLERT kernel is running:

- Minimal intrusion

Allows NightProbe to read and write variables without stopping the process for each sample or write operation.

- Sampling performance

Allows NightProbe to use direct memory fetches for data sampling (as opposed to programmed I/O) which is important for high-rate data acquisition.

- Concurrent debugging/probing

Allows NightProbe to probe programs already under the control of a debugger or another NightProbe session.

- PCI Device probing

Allows NightProbe to probe PCI device memory via the Base Address Register (BAR) file system.

Advantages for NightTune

The following advantage is afforded NightTune when a RedHawk or SLERT kernel is running:

- Context switch rate

Allows NightTune user to display the context switch counts per CPU instead of for the overall system.

- CPU shielding

Individual CPUs can be shielded from interrupts and processes allowing CPUs to be dedicated solely to specific interrupts and processes that are bound to the CPU.

- CPU sibling interference

Individual CPUs can be marked down to avoid interfering with hyperthreaded sibling CPUs and dual-core sibling CPUs. Hyperthreaded CPUs share all the resources of their sibling CPU. Dual-core CPUs share the CPU cache and a path to memory with their sibling CPU.

- Detailed memory information

Detailed process memory descriptions include the residency and lock state of any page in a process, and their association with physical memory pools for NUMA systems.

Frequency Based Scheduler

The Frequency Based Scheduler is only available on RedHawk systems from Concurrent Computer Corporation. It is required for all NightSim usage.

NightSim is only included in NightStar distributions intended for use on RedHawk systems.

PCI Bar File System

The PCI Bar File System is only available with the RedHawk kernel from Concurrent Computer Corporation and SLERT versions 1.0-1.6 kernel from Novell.

On other systems, PCI Device probing will be disabled within NightProbe.

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 7-31 for more information.

apply on dialogue

Execute **on dialogue** commands for existing dialogues. See “apply on dialogue” on page 7-30 for more information.

apply on program

Execute **on program** commands for existing processes. See “apply on program” on page 7-43 for more information.

attach

Attach the debugger to a process that is already running. See “attach” on page 7-36 for more information.

backtrace

bt

Print an ordered list of the currently active stack frames. See “backtrace” on page 7-81 for more information.

breakpoint

b

Set a breakpoint. See “breakpoint” on page 7-95 for more information.

cd

Set the debugger’s default working directory. See “cd” on page 7-71 for more information.

checkpoint

Take a restart checkpoint now. See “checkpoint” on page 7-45 for more information.

clear

Clear all eventpoints at a given location. See “clear” on page 7-104 for more information.

commands

Attach commands to a breakpoint or monitorpoint. See “commands” on page 7-105 for more information.

condition

Attach a condition to an eventpoint. See “condition” on page 7-106 for more information.

continue

c

Continue execution and wait for something to happen. See “continue” on page 7-114 for more information.

core-file

Create a pseudo-process for debugging an aborted program's core image file. See “core-file” on page 7-38 for more information.

data-display

Control items in a data panel. See “data-display” on page 7-88 for more information.

debug

Specify names for programs you wish to debug. See “debug” on page 7-24 for more information.

define

Define a NightView macro. See “define” on page 7-154 for more information.

delay

Delay NightView command execution for a specified time. See “delay” on page 7-129 for more information.

delete

d

Delete an eventpoint. See “delete” on page 7-107 for more information.

detach

Stop debugging a list of processes. See “detach” on page 7-37 for more information.

directory

Set the directory search path. See “directory” on page 7-74 for more information.

disable

Disable an eventpoint. See “disable” on page 7-108 for more information.

display

Add to the list of expressions to be printed each time the process stops. See “display” on page 7-88 for more information.

down

Move one or more stack frames toward frames called by the current stack frame. See “down” on page 7-126 for more information.

echo

Print arbitrary text. See “echo” on page 7-87 for more information.

edit

Edit the current source file. See “edit” on page 7-75 for more information.

enable

Enable an eventpoint for a specified duration. See “enable” on page 7-108 for more information.

exec-file

Specify the location of the executable file corresponding to a process. See “exec-file” on page 7-40 for more information.

family

Give a name to a family of one or more processes. See “family” on page 7-45 for more information.

finish

Continue execution until the current function finishes. See “finish” on page 7-119 for more information.

forward-search

fo

Search forward through the current source file for a specified regular expression. See “forward-search” on page 7-76 for more information.

frame

f

Select a new stack frame or print a description of the current stack frame. See “frame” on page 7-124 for more information.

handle

Specify how to handle signals in the user process. See “handle” on page 7-122 for more information.

heapcheck

Check the heap for errors. See “heapcheck” on page 7-148 for more information.

heapdebug

Specify parameters for heap debugging. See “heapdebug” on page 7-50 for more information.

heappoint

Check the heap for errors, or change the heap debugger settings, at a given location. See “heappoint” on page 7-102 for more information.

help

Access the online help system. See “help” on page 7-127 for more information.

ignore

Attach an ignore-count to an eventpoint. See “ignore” on page 7-109 for more information.

info address

Determine the location of a variable. See “info address” on page 7-150 for more information.

info args

Print description of current routine arguments. See “info args” on page 7-149 for more information.

info breakpoint

i b

Describe current state of breakpoints. See “info breakpoint” on page 7-132 for more information.

info convenience

Describe convenience variables. See “info convenience” on page 7-139 for more information.

info declaration

p type

Print the declaration of variables or types. See “info declaration” on page 7-152 for more information.

info dialogue

Print information about active dialogues. See “info dialogue” on page 7-146 for more information.

info directories

Print the search path used to locate source files. See “info directories” on page 7-139 for more information.

info display

Describe expressions that are automatically displayed. See “info display” on page 7-139 for more information.

info eventpoint

Describe current state of breakpoints, tracepoints, patchpoints, monitorpoints, heap-points, and watchpoints. See “info eventpoint” on page 7-131 for more information.

info family

Print information about an existing process family. See “info family” on page 7-146 for more information.

info files

Print the names of the executable, symbol table and core files. See “info files” on page 7-153 for more information.

info frame

Describe a stack frame. See “info frame” on page 7-138 for more information.

info functions

List names of functions or subroutines. See “info functions” on page 7-151 for more information.

info heappoint

Describe the current state of heappoints. See “info heappoint” on page 7-136 for more information.

info history

Print value history information. See “info history” on page 7-140 for more information.

info limits

Print information about limits on expression and location output. See “info limits” on page 7-140 for more information.

info line

Describe location of a source line. See “info line” on page 7-153 for more information.

info locals

Print information about local variables. See “info locals” on page 7-149 for more information.

info log

Describe any open log files. See “info log” on page 7-131 for more information.

info macros

Print a description of one or more NightView macros. See “info macros” on page 7-158 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 7-142 for more information.

info monitorpoint

Describe current state of monitorpoints. See “info monitorpoint” on page 7-135 for more information.

info name

Print information about an existing eventpoint-name. See “info name” on page 7-146 for more information.

info on dialogue

Print **on dialogue** commands. See “info on dialogue” on page 7-147 for more information.

info on program

Print **on program** commands. See “info on program” on page 7-147 for more information.

info on restart

Print **on restart** commands. See “info on restart” on page 7-147 for more information.

info patchpoint

Describe current state of patchpoints. See “info patchpoint” on page 7-134 for more information.

info process

Describe processes being debugged. See “info process” on page 7-141 for more information.

info registers

Print information about registers. See “info registers” on page 7-140 for more information.

**info representation
representation**

Describe the storage representation of an expression. See “info representation” on page 7-152 for more information.

info signal

Print information about signals. See “info signal” on page 7-141 for more information.

info sources

List names of source files. See “info sources” on page 7-151 for more information.

info threads

Describe C threads and thread processes. See “info threads” on page 7-148 for more information.

info tracepoint

Describe current state of tracepoints. See “info tracepoint” on page 7-133 for more information.

info types

Print type definition information. See “info types” on page 7-151 for more information.

info variables

Print global variable information. See “info variables” on page 7-150 for more information.

info watchpoint

Describe current state of watchpoints. See “info watchpoint” on page 7-137 for more information.

info whatis

whatis

Describe the result type of an expression visible in the current context. See “info whatis” on page 7-152 for more information.

interest

Control which subprograms are interesting. See “interest” on page 7-63 for more information.

jump

Continue execution at a specific location. See “jump” on page 7-120 for more information.

kill

Terminate a list of processes. See “kill” on page 7-37 for more information.

list

l

List a source file. See “list” on page 7-72 for more information.

load

Dynamically load an object file, possibly replacing existing routines. See “load” on page 7-91 for more information.

login

Login to a new dialogue shell. See “login” on page 7-22 for more information.

logout

Terminate a dialogue. See “logout” on page 7-28 for more information.

mcontrol

hold

release

Control the monitorpoint value display. See “mcontrol” on page 7-103 for more information.

monitorpoint

Monitor the values of one or more expressions at a given location. See “monitorpoint” on page 7-101 for more information.

mreserve

Reserve a region of memory in a process. See “mreserve” on page 7-49 for more information.

name

Give a name to a group of eventpoints. See “name” on page 7-94 for more information.

next

n

Execute one line, stepping over procedures. See “next” on page 7-117 for more information.

nexti

ni

Execute one instruction, stepping over procedures. See “nexti” on page 7-119 for more information.

nodebug

Specify names for programs you do not wish to debug. See “nodebug” on page 7-25 for more information.

notify

Ask about pending event notifications. See “notify” on page 7-36 for more information.

on dialogue

Specify debugger commands to be executed when a dialogue is created. See “on dialogue” on page 7-28 for more information.

on program

Specify debugger commands to be executed when a program is executed. See “on program” on page 7-41 for more information.

on restart

Specify debugger commands to be executed when a program is restarted. See “on restart” on page 7-44 for more information.

output

Print the value of a language expression with minimal output. See “output” on page 7-87 for more information.

patchpoint

Install a small patch to a routine. See “patchpoint” on page 7-96 for more information.

print

p

Print the value of a language expression. See “print” on page 7-81 for more information.

printf

Print the values of language expressions using a format string. See “printf” on page 7-91 for more information.

pwd

Print NightView’s current working directory. See “pwd” on page 7-71 for more information.

quit

q

Stop everything. Exit the debugger. See “quit” on page 7-21 for more information.

redisplay

Enable a display item. See “redisplay” on page 7-90 for more information.

refresh

Re-read source files and refresh the terminal screen. See “refresh” on page 7-128 for more information.

rerun

Run a program again. See “rerun” on page 7-34 for more information.

resume

Continue execution. See “resume” on page 7-115 for more information.

reverse-search

Search backwards through the current source file for a specified regular expression. See “reverse-search” on page 7-76 for more information.

run

Run a program in a dialogue and wait for NightView to start debugging it. See “run” on page 7-34 for more information.

select-context

Select the context of a thread or a thread process. See “select-context” on page 7-126 for more information.

set

Evaluate a language expression without printing its value. See “set” on page 7-83 for more information.

set-auto-frame

Control the positioning of the stack when a process stops. See “set-auto-frame” on page 7-65 for more information.

set-children

Control whether children should be debugged. See “set-children” on page 7-47 for more information.

set-debug-file-directory

Tell NightView where to look for .debug files. See “set-debug-file-directory” on page 7-25 for more information.

set-disassembly

Control how NightView displays disassembled instructions. See “set-disassembly” on page 7-69 for more information.

set-download

Control how NightView downloads files from remote targets. See “set-download” on page 7-68 for more information.

set-editor

Set the mode for editing commands in the simple full-screen interface. See “set-editor” on page 7-67 for more information.

set-exit

Control whether a process stops before exiting. See “set-exit” on page 7-47 for more information.

set-history

Specify the number of items to be kept in the value history list. See “set-history” on page 7-58 for more information.

set-language

Establish a default language context for variables and expressions. See “set-language” on page 7-56 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 7-58 for more information.

set-local

Define process local convenience variables. See “set-local” on page 7-61 for more information.

set-log

Log session to file. See “set-log” on page 7-56 for more information.

set-notify

Control how you are notified of events. See “set-notify” on page 7-35 for more information.

set-overload

Control how NightView treats overloaded operators and routines in expressions. See “set-overload” on page 7-66 for more information.

set-patch-area-size

Control the size of patch areas created in your process. See “set-patch-area-size” on page 7-62 for more information.

set-preallocate

Control how NightView preallocates memory for eventpoints and monitorpoint buffers. See “set-preallocate” on page 7-67 for more information.

set-prompt

Set the string used to prompt for command input. See “set-prompt” on page 7-59 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 7-57 for more information.

set-restart

Control whether restart information is applied. See “set-restart” on page 7-61 for more information.

set-resume

Control NightView's behavior on events that normally stop a process. See “set-resume” on page 7-68 for more information.

set-safety

Control debugger response to dangerous commands. See “set-safety” on page 7-60 for more information.

set-search

Control case sensitivity of regular expressions in NightView. See “set-search” on page 7-66 for more information.

set-show

Control where dialogue output goes. See “set-show” on page 7-32 for more information.

set-terminator

Set the string used to recognize end of dialogue input mode. See “set-terminator” on page 7-60 for more information.

set-trace

Establish tracing parameters. See “set-trace” on page 7-98 for more information.

shell

Run an arbitrary shell command. See “shell” on page 7-128 for more information.

show

Control dialogue output. See “show” on page 7-33 for more information.

signal

Continue execution with a signal. See “signal” on page 7-121 for more information.

source

Input commands from a source file. See “source” on page 7-129 for more information.

step

s

Execute one line, stepping into procedures. See “step” on page 7-116 for more information.

stepi

si

Execute one instruction, stepping into procedures. See “stepi” on page 7-118 for more information.

stop

Stop a process. See “stop” on page 7-120 for more information.

symbol-file

Establish the file containing symbolic information for a program. See “symbol-file” on page 7-38 for more information.

tbreak

Set a temporary breakpoint. See “tbreak” on page 7-110 for more information.

tpatch

Set a patchpoint that will execute only once. See “tpatch” on page 7-111 for more information.

tracepoint

Set a tracepoint. See “tracepoint” on page 7-99 for more information.

translate-object-file

x1

Translate object filenames for a remote dialogue. See “translate-object-file” on page 7-26 for more information.

undisplay

Disable an item from the display expression list. See “undisplay” on page 7-90 for more information.

up

Move one or more stack frames toward the caller of the current stack frame. See “up” on page 7-125 for more information.

wait

Wait for processes to stop. See “wait” on page 7-48 for more information.

watchpoint

Set a watchpoint. See “watchpoint” on page 7-112 for more information.

x

Print the contents of memory beginning at a given address. See “x” on page 7-84 for more information.

Invoking NightView

```
nview [-attach pid] [-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: **x1**

```
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
detach
kill
symbol-file program-name
core-file corefile-name [exec-file=program-name]
                    [with-translations]
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-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 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] [{off | on}]

set-download [{off | permanent | temporary}]
    [directory=path-to-cache]

set-disassembly [flavor={att | intel}]
    [symbols={off | on}]
    [comment_level=number]

```

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 [*regexp*]

Abbreviation: **fo**

reverse-search [*regexp*]

Examining and Modifying

backtrace [*number-of-frames*]

Abbreviation: **bt**

print [*/print-format-letter*] *expression*

Abbreviation: **p**

set *expression*

x [*/repeat-count*] [*size-letter*] [*x-format-letter*]] [*addr-expression*]

output [*/print-format-letter*] *expression*

echo *text*

data-display [*/window="window name"*]
 {*/kind=value* | *expression*}

display [*/print-format-letter*] *expression*

display [*/repeat-count*] [*size-letter*] [*x-format-letter*] *addr-expression*

undisplay *item-number* ...

redisplay *item-number* ...

printf *format-string* [, *expression* ...]

load *object*

Manipulating Eventpoints

name [/add] *name* [[-] *eventpoint-spec*] ...

breakpoint [*eventpoint-modifier*] [*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*] [*tracefile=trace-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

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]
```

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*

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 watchpoint [/verbose] [eventpoint-spec] ...

info frame [/v] [*expression [at location-spec]]

info directories

info convenience

info display

info history [number]

info limits

info registers [regexp]

info signal [signal ...]

info process

info memory [/ranges] [/heap] [/leaks] [/allocated]
           [/all] [/append=filename] [/output=filename]
           [/verbose] [expression]

info dialogue

info family [regexp]

info name [regexp]
```

```

info on dialogue [name]
info on program [program]
info on restart [output=outname|append=outname] [program]
info threads [/verbose]
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]
```

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** implements 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 6 [Invoking NightView] on page 6-1.

NightView runs `NightView.p` for each dialogue. If the dialogue is on the local machine, then NightView communicates with `NightView.p` via a shared memory region. There is one such shared memory region per invocation of NightView. See “Dialogues” on page 3-4. For remote dialogues, NightView establishes a socket connection with `NightView.p`.

`NightView.p` is responsible for controlling the user processes by a combination of the `/proc` file system and the **ptrace** system service.

Monitorpoints communicate with `NightView.p` via a shared memory region created in your process. There is one shared memory region for each dialogue using monitorpoints. See “Monitorpoints” on page 3-11. 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`.

Each dialogue runs a shell and controls it using `/proc` and **ptrace**. This is not to get control of the shell, but so that the debugger is notified of the shell's children, which are the processes to be debugged. The shell runs at a pseudo-terminal controlled by the debugger, so that the debugger can capture the program I/O.

Watchpoints are implemented by setting the `DRn` registers on IA-32 or AMD64 systems. Other eventpoints are implemented by replacing the instruction at the target address by a trap instruction. When your program hits the trap, the debugger translates this into a branch to a patch area. 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 one or two eventpoint patch areas. The user can adjust the sizes of the patch areas. See “set-patch-area-size” on page 7-62. Each region is only created in the process if necessary.

On IA-32, NightView tries to put all the patch areas between `0xa0000000` and `0xb0000000`.

On AMD64, an eventpoint patch area may need to be placed near the instruction being patched. NightView tries to put all the other patch areas between `0x15aaaab000` and `0x2000000000`.

You can see where NightView has placed patch areas with the **info memory** command (see “info memory” on page 7-142).

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.

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 <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 }
```

```

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
5      program main
6      common /msg_comm/ total_sig
7      common /usr1_comm/ sigusr1
8      integer total_sig
9      integer sigusr1
10     integer pid
11     integer ftfork
12     integer ftsigusr1
13     character *8 tracefile
14     external parent_routine
15     external child_routine
16
17     sigusr1 = ftsigusr1()
18     tracefile = "msg_file"
19     write( 6, FMT='(A)' )
20     X "How many signals should the parent send the child?"
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 ftsigsr1_()
35 {
36     return SIGUSR1;
37 }

```


Reporting Bugs

This section describes how to report bugs in NightView. It is important to report problems, otherwise we might never know about them. You can report a problem by calling the Concurrent Software Support Center. For more information, see the section *Direct Software Support* in the release notes for the current release.

It is also important to report a problem in a way that helps us understand and reproduce the problem.

Here is a list of things you should tell us in a problem report.

- What version of NightView are you using? You can get this by running `nview -v` or by activating **On Version** in the **Help** menu in the GUI.
- What type of machine are you running on, and what is the version of the operating system? Use the command `uname -a` to get this information.
- Are you using the GUI, the simple-screen interface, or the command line interface? Sometimes this is not obvious from the description of the problem.
- Try to be very explicit about what you see happen when the problem occurs. Do you get any error messages? Exactly what incorrect behavior do you see? How will we know when we have reproduced the problem?
- Try to isolate the problem to a small test program and a small series of actions in the debugger. This is not always possible, but to the extent that you can isolate the problem, that will help us reproduce and fix it.
- Be explicit about exactly how to reproduce the problem. Try not to leave out any facts, even if you think you know the cause of the problem.

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 9-25.

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 7-16.

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 7-36.

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 3-11. 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 3-16.

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 7-47.

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 3-36.

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 7 [Command-Line Interface] on page 7-1.

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 3-34.

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 3-23.

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 9-3.

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 9-59.

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 3-35.

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 3-4.

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.

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 3-24.

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 underlined text. See “Current Process” on page 9-3.

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 9-60

data panel

In the graphical user interface, a data panel allows you to view data items in your process. See “Data Panel” on page 9-60.

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 9-14.

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 6 [Invoking NightView] on page 6-1. See "Quitting NightView" on page 7-21.

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 9-36.

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 7-37.

dialogue

NightView provides dialogues as a means of starting processes, via a shell, and communicating with those processes. See "Dialogues" on page 3-4. 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 7-84. In the graphical user interface, you can view disassembly in the source panels by using the `Source` menu. See "Source Menu" on page 9-10.

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 7-88 and "info display" on page 7-139.

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 3-9.

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 7-94.

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 9-57.

family

A group of related processes. See “family” on page 7-45.

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 3-30.

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 9-48.

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 8 [Simple Full-Screen Interface] on page 8-1.

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 9 [Graphical User Interface] on page 9-1.

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 3-28.

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 3-32.

heap debugger

A module that NightView loads into your process to monitor memory allocation and deallocation. See “Debugging the Heap” on page 3-28.

heappoint

An *eventpoint* that checks the *heap* or changes the heap debugging parameters. See “Heappoints” on page 3-12.

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 9-80.

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 3-11.

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 7-109.

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 3-36.

inline subprogram

A subprogram that is expanded directly into the calling program. See “Inline Subprograms” on page 3-25.

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 3-25. You can change or query this level with the **interest** command. See “interest” on page 7-63.

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 3-9.

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 3-25. You can change or query the interest level with the **interest** command. See “interest” on page 7-63.

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 3-25.

keyboard focus

The keyboard focus determines which field receives keyboard input in the graphical user interface.

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 3-32.

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 9-59.

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 7-153.

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 9-56.

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 9-25.

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 3-11 and “monitorpoint” on page 7-101.

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 9-60.

monitor panel

In the graphical user interface, a monitor panel displays the values of *monitorpoints*. See “Monitor Panel” on page 9-59.

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 7-127. See "GUI Online Help" on page 9-1.

overloading

Overloading means that more than one entity with the same name is visible at some point in the program. See "Overloading" on page 3-22.

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 3-12.

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 E [Implementation Overview] on page E-1. See "set-patch-area-size" on page 7-62.

patchpoint

A patchpoint is a location in a debugged process where a patch is inserted. See *patch*. See "Patchpoints" on page 3-12.

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 7-24.

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 3-2.

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 3-18.

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 3-2.

program counter

The program counter is a register that locates the instruction that is to be executed next. See “Program Counter” on page 3-23.

qualifier

A qualifier specifies the set of processes or dialogues that a command affects. See “Qualifiers” on page 3-4.

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 7-5.

remote dialogue

A remote dialogue is a *dialogue* started on a system other than the one on which NightView was invoked. See “Remote Dialogues” on page 3-6.

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 3-16.

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 3-31.

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 3-23.

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 9-55.

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. NightView allows you to control how signals are delivered to your process. See “Signals” on page 3-14.

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 9-48.

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 9-49.

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 3-24.

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 3-24.

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 7-38.

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 3-37.

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 3-12.

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 3-36.

watchpoint

A watchpoint stops the process when the process reads or writes a variable in memory. See “Watchpoints” on page 3-13.

Symbols

! 1-3, 3-5, 4-11, 7-31, 7-34, 7-60, 7-157
!exit 7-28
(comment) 7-2
\$ 7-3, 7-62, 7-83
\$ prompt 1-3, 4-4, 5-4
\$\$ 7-3
\$_ 7-5
\$__ 7-5
\$cfa 7-6, 7-124, 7-138
\$cpc 3-23, 3-24, 7-6, 7-64, 7-79, 7-124, 7-125, 7-138
\$fp 7-6, 7-139
\$is 7-6
\$pc 3-20, 3-23, 3-24, 7-5, 7-86, 7-125
\$sp 7-6
\$was 3-13
& 4-6, 5-7, 7-13, 7-34, 7-151
(local) prompt 1-3, 4-4
* (source line decoration) 7-79
+ (list argument) 7-72, 7-73
. 7-16
- . (input terminator) 7-31, 7-60
.NightViewrc file 3-36, 7-27, 7-29, 7-30
.profile file 3-7
/disabled eventpoint modifier 7-94, Glossary-5
/proc 3-3, 3-39, E-1
/Tutorial 4-3
/usr/lib/NightView/Tutorial 5-3
/usr/lib/NightView-release/ReadyToDebug
1-3, 3-9, 4-4, 5-4
/usr/ucb/rsh 3-3
< (source line decoration) 7-79
<= 7-148, 9-62
= (list argument) 7-72
= (source line decoration) 7-79
= key 9-55
> prompt 7-2, 7-101, 7-106, 7-155
> (source line decoration) 7-79
@ (macro invocation) 7-156
@ (source line decoration) 7-80, 7-153
\ 7-88
\n 7-87
| (show argument) 7-33

' specification 7-13

A

Abbreviations

b (breakpoint) 7-95, D-7
bt (backtrace) 7-81, D-6
c (continue) 7-114, D-8
command 7-1, 7-157
d (delete) 7-107, D-7
f (frame) 7-124, D-9
fo (forward-search) 7-76, D-6
hd (heapdebug) 7-50, D-4
hold (mcontrol hold) 7-103, D-7
i b (info breakpoint) 7-132, D-10
l (list) 7-72, D-5
n (next) 7-117, D-8
ni (nexti) 7-119, D-9
p (print) 7-81, D-6
ptype (info declaration) 7-152, D-11
q (quit) 7-21, D-1
release (mcontrol release) 7-103, D-7
representation (info representation)
7-152, D-11
s (step) 7-116, D-8
si (stepi) 7-118, D-8
whatis (info whatis) 7-152, D-11
xl (translate-object-file) 7-26, D-1

Abnormal termination 7-39

Abort 3-33

Accelerator Glossary-1

Ctrl+B 9-12

Ctrl+P 9-12

Ctrl+Q 9-5

Ctrl+S 9-4

Accessing files 3-1, 3-7

Actual argument

macro 7-155, 7-156, 7-157

Address

printing 7-150

Addresses limits 7-140

printing 7-58

addr-expression 7-86, 7-89

Aggregate item 7-82, 7-140
AMD64 registers 7-9
Anchored match 7-17, 7-76, 7-140, 7-146, 7-149, 7-150,
7-151, 7-152, 7-158, 9-10, 9-26, 9-51,
Glossary-1
Application 3-2, Glossary-1
apply on dialogue 7-29, 7-30
apply on program 7-43
Argument
 actual 7-155
 command 7-1
 macro 7-154, 7-156
 printing 7-142, 7-149
Array 7-83, 7-140, 9-61
 printing 7-58
Array slices 3-21
Assembly listing 9-11, 9-12, 9-53
Assignment 3-21, 7-97
attach 3-3, 3-38, 6-1, 7-36, 7-47
Attach button 9-26
Attach Dialog Box 3-3, 9-26
-attach option 6-1
Attaching 3-3, 3-38, 6-1, 7-36, 9-9, 9-24, 9-26,
Glossary-1
Attaching commands to a dialogue 7-28
Attaching commands to a program 7-42, 7-44

B

b (breakpoint) 1-4, 2-5, 4-9, 7-95, D-7
B (source line decoration) 7-78
b key 9-55
Background process 7-129
backtrace 1-5, 2-6, 3-24, 4-17, 5-16, 7-4, 7-81,
7-124, 7-125, 7-127
Blank line 7-19, 7-73
Body
 macro 3-35, 7-154, 7-159
Branch instruction 7-97, 7-111
Breakpoint 3-9, 3-11, 3-18, 4-14, 5-13, 7-78, 7-80, 7-92,
7-98, 7-115, 7-121, 9-21, 9-24, 9-55,
Glossary-1, Glossary-5
 changing 9-28
 clearing 7-104, 7-105
 commands on 3-17, 3-34, 3-35, 4-23, 5-22, 7-105,
7-132, 7-133, 9-30, Glossary-1
 condition on 3-17, 4-21, 5-20, 7-96, 7-106, 7-132,
7-133, 9-29, Glossary-2
 deleting 4-21, 5-20, 7-37, 7-107, 9-31
 disabling 4-26, 5-26, 7-108, 9-29
 displaying 4-27, 5-26, 7-131, 7-132
 enabling 7-108, 9-29

 hitting 3-18, 7-35, 7-89, 7-113, 7-114, 7-133, 9-29,
Glossary-6
 ignoring 3-17, 4-22, 5-21, 7-96, 7-109, 7-132,
7-133, 9-30, Glossary-7
 named 3-17, 7-94, 7-95, 7-110, 9-30
 saving 3-17
 setting 1-4, 2-5, 4-9, 4-21, 4-22, 5-9, 5-20, 5-21,
7-92, 7-95, 7-110, 9-21, 9-28, 9-55
 state 7-133, 9-29
 temporary 7-110, 9-29
breakpoint 1-4, 2-5, 3-2, 3-4, 4-9, 4-22, 4-23, 7-95,
7-110, 7-132, 9-20, 9-21, 9-55
Breakpoint button 5-9, 9-21, 9-55
Breakpoint crossing count Glossary-3
Breakpoint Dialog Box 5-21, 5-22, 9-12, 9-21, 9-27
Browse selection policy 9-27
bt 1-5, 2-6
bt (backtrace) 7-81, D-6
Buffered output 3-35
Building a program 1-2, 2-2, 4-2, 5-3
Button
 Attach 9-26
 Breakpoint 5-9, 9-21, 9-55
 Cancel 9-26, 9-27, 9-31, 9-72, 9-76, 9-77, 9-78,
9-79, 9-80
 Clear 9-22
 Close 5-8
 Data Display 5-23, 9-22
 Delete 9-31
 dimmed 5-20, 9-27, 9-31
 disabled 5-20, 9-27, 9-31
 Dismiss 9-2
 Down 5-19
 Down 9-20
 Filter 9-26
 Finish 9-20
 Help 9-2, 9-26, 9-27, 9-31, 9-72, 9-77
 Hold 9-60
 Interrupt 9-21
 Kill 7-37, 9-20
 Next 5-10, 9-19
 Nexti 9-20
 OK 2-6, 5-15, 5-21, 5-22, 5-23, 5-29, 9-27, 9-31,
9-72, 9-73, 9-74, 9-75, 9-76, 9-77, 9-78,
9-80
 Print 2-6, 5-14, 5-18, 9-22, 9-55
 radio 9-30, 9-73
 Release 9-60
 Rerun 9-21
 Resume 2-5, 2-6, 5-10, 5-11, 5-20, 5-24, 5-28,
9-19
 Run to Here 9-20, 9-55
 Search 5-15, 9-26
 Step 9-20

- Stepi 9-20
 - Stop 9-19
 - Switch To 5-12
 - toggle 9-30
 - Up 5-17
 - Up 9-20
- C**
- c (continue)** 1-5, 7-114, D-8
 - C language i-iii, 3-21, 3-23, 3-38, 7-56, 7-59, 7-82, 7-83, 7-87, 7-91, 7-97, 7-111, 7-151, 7-155, Glossary-2
 - C string 3-38, 7-83
 - C thread 7-148, 9-62
 - C++ i-iii, 3-21, 3-23, 3-38, 7-56, 7-82, 7-97, 7-111, 7-155, 9-10, 9-27, 9-51, Glossary-2
 - Calling macros 7-156
 - Cancel button 9-26, 9-27, 9-31, 9-72, 9-76, 9-77, 9-78, 9-79, 9-80
 - Canonical frame address 7-6, 7-124, 7-138
 - Cautions 3-33
 - cc option
 - g 1-2, 2-2, 4-2, 5-3
 - cd** 7-71
 - Changing a breakpoint 9-28
 - Changing a heappoint 9-28
 - Changing a monitorpoint 9-28
 - Changing a patchpoint 9-28
 - Changing a tracepoint 9-28
 - Changing a watchpoint 9-28
 - Changing an eventpoint 9-28
 - Changing eventpoints 5-22, 5-26, 9-27, 9-57
 - Character string 7-83, 7-140
 - printing 7-58
 - Checkpoint 3-16, 3-17, 7-42, 7-44, 7-45, Glossary-1, Glossary-10
 - checkpoint** 3-17, 7-45
 - Child process 3-2, 3-3, 4-1, 4-7, 4-14, 5-1, 5-7, 5-13, 7-25, 7-37, 7-47, Glossary-1
 - clear** 7-92, 7-104, 7-107, 9-22
 - Clear** button 9-22
 - Clearing breakpoints 7-104
 - Clearing eventpoints 7-104, 9-22
 - Clearing heappoints 7-104
 - Clearing monitorpoints 7-104
 - Clearing patchpoints 7-104
 - Clearing terminal 7-128
 - Clearing tracepoints 7-104
 - Clicking
 - double 5-14, 5-18, 9-27, 9-76, 9-77
 - Clicking on objects 5-2, 5-5, 5-7, 5-14, 5-20, 5-21, 5-22, 5-25, 5-26, 5-29
 - Close button 5-8
 - Command abbreviations 7-1, 7-157
 - b (breakpoint)** 7-95, D-7
 - bt (backtrace)** 7-81, D-6
 - c (continue)** 7-114, D-8
 - d (delete)** 7-107, D-7
 - f (frame)** 7-124, D-9
 - fo (forward-search)** 7-76, D-6
 - hd (heapdebug)** 7-50, D-4
 - hold (mcontrol hold)** 7-103, D-7
 - i b (info breakpoint)** 7-132, D-10
 - l (list)** 7-72, D-5
 - n (next)** 7-117, D-8
 - ni (nexti)** 7-119, D-9
 - p (print)** 7-81, D-6
 - ptype (info declaration)** 7-152, D-11
 - q (quit)** 7-21, D-1
 - release (mcontrol release)** 7-103, D-7
 - representation (info representation)** 7-152, D-11
 - s (step)** 7-116, D-8
 - si (stepi)** 7-118, D-8
 - whatis (info whatis)** 7-152, D-11
 - x1 (translate-object-file)** 7-26, D-1
 - Command arguments 7-1
 - Command case 7-1
 - Command execution
 - delaying 7-130
 - Command file 7-129
 - Command history 3-36
 - Command input 7-129, 7-130
 - Command prompt 4-4, 4-11
 - Command qualifier 3-4, 4-20, 7-1, 7-14, 7-58, Glossary-10
 - Command repetition 4-14, 7-1, 7-19, 7-73
 - Command replacement 7-157
 - Command stream 3-34, 3-36, 7-115, 7-121, 7-129, 7-130, Glossary-2
 - event-driven 3-34
 - Command summary 5-5, 9-18, C-1
 - Command syntax 7-1
 - Command toolbar 9-19
 - Command-line interface i-iii, 1-1, 3-27, 4-1, 4-3, 5-1, 5-3, 7-2, 7-101, 7-106, 7-123, 7-155, 8-1, 9-1, 9-19, Glossary-2
 - Command-line user interface 7-127
 - Commands
 - ! 7-31
 - apply on dialogue** 7-30
 - apply on program** 7-43
 - attach** 6-1, 7-36
 - backtrace** 7-81
 - breakpoint** 7-95

cd 7-71
checkpoint 7-45
clear 7-104
commands 7-105
condition 7-106
continue 7-114
core-file 7-38
data-display 3-17, 7-88
debug 7-24
define 7-154
delay 7-129
delete 7-107
detach 7-37
directory 7-74
disable 7-108
display 7-88
down 7-126
echo 7-87
edit 7-75
enable 7-108
exec-file 7-40
family 7-45
finish 7-119
forward-search 7-76
frame 7-124
handle 7-122
heapcheck 3-32, 7-53, 7-148
heapdebug 3-12, 3-28, 3-29, 3-32, 7-50, 7-103, 7-149, 9-31, 9-32
heappoint 3-12, 3-32, 7-53, 7-102, 7-103, 7-136
help 7-127
ignore 7-109
info address 7-150
info args 7-149
info breakpoint 7-132
info convenience 7-139
info declaration 7-152
info dialogue 7-146
info directories 7-139
info display 7-139
info eventpoint 7-131
info family 7-146
info files 7-153
info frame 7-138
info functions 7-151
info heappoint 7-136
info history 7-140
info limits 7-140
info line 7-153
info locals 7-149
info log 7-131
info macros 7-158
info memory 7-142
info monitorpoint 7-135
info name 7-146
info on dialogue 7-147
info on program 7-147
info on restart 7-147
info patchpoint 7-134
info process 7-141
info registers 7-140
info representation 7-152
info signal 7-141
info sources 7-151
info tracepoint 7-133
info types 7-151
info variables 7-150
info watchpoint 7-137
info whatis 7-152
interest 7-63
jump 7-120
kill 7-37
list 7-72
load 7-91
login 7-22
logout 7-28
mcontrol 7-103
monitorpoint 7-101
mreserve 7-49
name 7-94
next 7-117
nexti 7-119
nodebug 7-25
notify 7-36
on dialogue 7-28
on program 7-41
on restart 7-44
output 7-87
patchpoint 7-96
print 7-81, 7-91
pwd 7-71
quit 7-21
redisplay 7-90
refresh 7-128
rerun 7-34
resume 7-115
reverse-search 7-76
run 7-34
select-context 7-126
set 7-83
set-auto-frame 7-65
set-children 7-47
set-debug-file-directory 7-25
set-disassembly 7-69
set-download 7-68
set-editor 7-67
set-exit 7-47
set-history 7-58

- set-language** 7-56
 - set-limits** 7-58
 - set-local** 7-61
 - set-log** 7-56
 - set-notify** 7-35
 - set-overload** 7-66
 - set-patch-area-size** 7-62
 - set-preallocate** 7-67
 - set-prompt** 7-59
 - set-qualifier** 7-57
 - set-restart** 7-61
 - set-safety** 7-60
 - set-search** 7-66
 - set-show** 7-32
 - set-terminator** 7-60
 - set-trace** 7-98
 - shell** 7-128
 - show** 7-33
 - signal** 7-121
 - source** 7-129
 - step** 7-116
 - stepi** 7-118
 - stop** 7-120
 - symbol-file** 7-38
 - tbreak** 7-110
 - tpatch** 7-111
 - tracepoint** 7-99
 - translate-object-file** 7-26
 - undisplay** 7-90
 - up** 7-125
 - wait** 7-48
 - watchpoint** 7-112
 - x** 7-84
 - commands** 3-11, 4-23, 7-92, 7-105
 - Commands attached to a dialogue 7-28
 - Commands attached to a program 7-42, 7-44
 - Commands on breakpoint 3-17, 5-22, 7-95, 7-96, 7-105, 9-30, Glossary-1
 - Commands on eventpoint 9-30
 - Commands on monitorpoint 3-17, 7-101, 7-105, 9-30
 - Commands on watchpoint 3-17
 - Comments 7-2, 7-154
 - Compilation
 - by debugger 3-10, 3-35
 - Compiling 1-2, 2-2, 3-12, 3-37, 4-2, 5-3
 - Condition
 - breakpoint 3-17, 5-20, 7-132, 7-133, 9-29, Glossary-2
 - eventpoint 3-10, 3-17, 7-106, 7-110, 7-132, 9-29
 - heappoint 3-17, 7-132, 7-137, 9-29
 - monitorpoint 3-17, 7-132, 7-136, 9-29
 - patchpoint 3-17, 7-132, 7-135, 9-29
 - tracepoint 3-17, 7-132, 7-134, 9-29
 - watchpoint 3-13, 3-17, 7-132, 7-138, 9-29
 - condition** 7-92, 7-98, 7-106
 - Condition removal 7-107, 9-29
 - conditional-expression* 3-13, 3-21, 7-95, 7-99, 7-106, 7-110, 7-112
 - config option 6-1
 - config** option 6-1
 - Configuration 6-1, 9-3, 9-4, 9-5
 - Context 3-19, 3-23, 3-37, 7-4, 7-56, 7-73, 7-119, 7-142, 7-151, 7-152, 7-153, Glossary-2, Glossary-8
 - Context Menu 9-3
 - context panel 5-7, 5-12, 5-13
 - Context-sensitive help 9-2, 9-17
 - continue** 1-5, 3-4, 3-15, 4-10, 4-14, 4-29, 7-114, 7-115, 7-116, 7-121
 - Continuing execution 1-5, 2-5, 2-6, 3-37, 4-10, 4-14, 4-20, 4-29, 5-10, 5-13, 5-19, 5-27, 7-114, 7-115, 7-119, 7-120, 7-121, 7-127, 9-55
 - Convenience variable 3-35, 7-3, 7-5, 7-83, Glossary-3
 - global 3-35, 7-62, 7-139
 - predefined 3-24, 7-3, 7-5, 7-132, 7-133, 7-134, 7-135, 7-136, 7-137, 7-140, 7-153
 - process local 3-35, 7-5, 7-62, 7-139
 - Core file 3-3, 3-4, 3-19, 6-1, 7-39, 7-153, 9-24, Glossary-3
 - core option 6-1
 - core-file** 7-38, 7-71
 - CPU bias 7-23
 - CPU hang 3-39
 - Cross reference
 - help 5-5
 - Crossing count Glossary-3
 - Current frame 3-24, 7-116, 7-118, 7-119, 9-20
 - Current Process 9-3
 - Current process Glossary-3
 - Current source file 7-73, 7-76, 7-77
 - Current stack frame 3-23, 3-24, 4-18, 4-19, 4-29, 5-17, 5-19, 5-27, 7-79, 7-95, 7-97, 7-99, 7-101, 7-102, 7-105, 7-111, 7-119, 7-124, 7-125, 7-126, 7-138, 7-141, Glossary-3
 - Current working directory 7-71
- ## D
- d (delete)** 7-107, D-7
 - d key 9-55
 - Data definitions
 - global 7-92
 - static 7-92
 - Data Display** button 5-23, 9-22
 - Data item 9-60, Glossary-3
 - Data** menu 9-14
 - Data Panel

- saving contents 9-15, 9-71
- Data panel 9-14
- Data Panel Add Heap Errors Dialog Box 9-73
- Data Panel Add Heap Information Dialog Box 9-73
- Data Panel Add Heap Leaks Dialog Box 9-73
- Data Panel Add Monitorpoints Dialog Box 9-74
- Data Panel Add Processes Dialog Box 9-74
- Data Panel Add Shells Dialog Box 9-74
- Data Panel Add Still Allocated Blocks Dialog Box 9-74
- Data Panel Condition Filter Expression Dialog Box 9-78
- Data Panel Describe Dialog Box 9-76
- data panel layout Glossary-3
- Data Panel Linked List Expression Dialog Box 9-78
- Data type
 - printing 7-151, 7-152
- Data Window 3-17, 5-23, 5-25, 5-28, 7-88, 9-22, 9-60
 - loading layout 9-16, 9-71
 - saving layout 9-16, 9-71
- Data Window Add Expression Dialog Box 9-72
- Data Window Add Local Variables Dialog Box 9-73
- Data Window Add Registers Dialog Box 9-74
- Data Window Add Stack Dialog Box 9-75
- Data Window Add Threads Dialog Box 9-75
- Data Window Call Stack Frames Dialog Box 9-75
- Data Window Edit Expression Dialog Box 9-75
- Data Window Expand Tree Dialog Box 9-75
- Data Window Item Dialog Box 9-72
- Data Window layout 9-16, 9-71, 9-76
- Data Window Load Layout Dialog Box 9-76
- Data Window Pointer Array Dimension Dialog Box 9-76
- Data Window Save Layout Dialog Box 9-76
- Data Window Save Snapshot Dialog Box 9-77
- Data Window Subscript Array Dialog Box 9-77
- data-display** 3-17, 7-88
- debug** 3-2, 3-5, 7-18, 7-24, 7-25, 7-146
- Debug command area 2-3, 2-6, 5-8, 5-17, 5-18, 5-19
- Debug Eventpoint Dialog Boxes 9-27
- Debug Eventpoint menu 5-20, 5-22, 5-25, 5-26, 9-27
- Debug Eventpoint Summarize/Change Dialog Box 5-20, 5-26
- Debug file directory 7-25
- Debug File Selection Dialog Box 9-27
- Debug group area 5-10, 5-11, 5-12, 5-16, 5-20, 5-24, 5-28
- Debug identification area 2-3, 5-12
- Debug information 4-3, 5-3, 7-38, 7-68
- Debug Interrupt button 9-21
- Debug menu bar 9-4
- Debug message area 2-3, 2-6, 5-8, 5-9, 5-10, 5-12, 5-13, 5-14, 5-18, 5-21, 5-22, 5-24, 5-26, 9-27
- Debug NightView menu 2-6, 5-29
- Debug session Glossary-4
- Debug source display 2-3, 2-5, 2-6, 5-9, 5-10, 5-11, 5-14, 5-15, 5-16, 5-18, 5-19, 5-21, 5-22, 5-26, 9-10, 9-11, 9-12, 9-22, 9-27, 9-28, 9-48, 9-51, 9-52, 9-53
- Debug source file name 2-3, 5-13, 5-18, 5-19, 5-28
- Debug **Source** menu 5-14, 9-10, 9-26, 9-49, 9-55
- Debug Source Selection Dialog Box 9-10, 9-26, 9-51
- Debug status area 2-3, 5-10, 5-11, 5-16, 5-20, 5-24, 5-28
- Debug status message 9-56
- Debug table 7-149, 7-151
- Debug **View** menu 5-18, 5-28
- Debug Window 2-3, 5-9, 5-10, 5-11, 5-14, 5-16, 5-20, 5-24, 5-25, 9-4, 9-23, 9-26, 9-27
- Debugger 3-1, Glossary-4
 - gdb 1-1, 1-4
 - NightView i-iii, 3-1
 - symbolic 3-1
- Debugging
 - multiple processes 3-2
 - real-time 3-6
 - single process 3-2
- Declaration
 - printing 7-152
- Default font Glossary-4
- define** 7-154, 7-156
- Defining a macro 7-155
- delay** 7-129
- delete** 4-21, 7-92, 7-105, 7-107
- Delete button 9-31
- Deleting breakpoints 4-21, 5-20, 7-107, 9-31
- Deleting eventpoints 7-107, 9-31
- Deleting heappoints 7-107, 9-31
- Deleting monitorpoints 7-107, 9-31
- Deleting patchpoints 7-107, 9-31
- Deleting tracepoints 7-107, 9-31
- Deleting watchpoints 7-107, 9-31
- detach** 7-25, 7-28, 7-37, 9-9
- Detaching 3-2, 3-18, 7-21, 7-28, 7-37, 7-100, 9-9, Glossary-4
- Dialog Box 2-3, 2-6, 9-2
 - Attach 3-3, 9-26
 - Breakpoint 5-21, 5-22, 9-12, 9-21, 9-27
 - Data Panel Add Heap Errors 9-73
 - Data Panel Add Heap Information 9-73
 - Data Panel Add Heap Leaks 9-73
 - Data Panel Add Monitorpoints 9-74
 - Data Panel Add Processes 9-15, 9-70, 9-74
 - Data Panel Add Shells 9-15, 9-70, 9-74
 - Data Panel Add Still Allocated Blocks 9-74
 - Data Panel Condition Filter Expression 9-78
 - Data Panel Describe 9-76
 - Data Panel Linked List Expression 9-78
 - Data Window Add Expression 9-14, 9-70, 9-72
 - Data Window Add Heap Errors 9-15, 9-71
 - Data Window Add Heap Information 9-15, 9-70

- Data Window Add Heap Leaks 9-15, 9-71
- Data Window Add Local Variables 9-14, 9-70, 9-73
- Data Window Add Registers 9-14, 9-70, 9-74
- Data Window Add Stack 9-14, 9-70, 9-75
- Data Window Add Still Allocated Blocks 9-15, 9-71
- Data Window Add Threads 9-15, 9-16, 9-70, 9-71, 9-75
- Data Window Call Stack Frames 9-75
- Data Window Edit Expression 9-75
- Data Window Expand Tree 9-75
- Data Window Item 9-72
- Data Window Load Layout 9-76
- Data Window Pointer Array Dimension 9-76
- Data Window Save Layout 9-76
- Data Window Save Snapshot 9-77
- Data Window Subscript Array 9-77
- Debug Eventpoint 9-27
- Debug Eventpoint Summarize/Change 5-20, 5-26
- Debug File Selection 9-27
- Debug Heap 9-32
- Debug Source Selection 9-10, 9-26, 9-51
- Eventpoint 9-27
- File Selection 9-11, 9-52
- Heappoint 9-13, 9-22, 9-27
- Monitorpoint 9-12, 9-21, 9-27
- Monitorpoint Update Interval 9-80
- Patchpoint 5-25, 9-13, 9-21, 9-27
- Remote Login 9-33
- Select a Function/Unit 5-15, 9-10, 9-26, 9-51
- Select a Source File 9-10, 9-26, 9-51, 9-52
- Tracepoint 9-13, 9-22, 9-27
- Warning 5-29, 9-2, 9-5, 9-8, 9-9, 9-21
- Watchpoint 9-13, 9-22, 9-27
- Dialogue 3-4, 3-5, 3-6, 4-4, 5-4, 5-11, 5-13, 7-15, 7-22, 7-23, 7-36, 7-147, 9-2, 9-8, 9-33, E-1, Glossary-4
 - commands on 7-28
 - local 3-5, 6-1, 7-1
 - local - with on dialogue 7-29, 7-30
 - printing 7-146
 - remote 3-6, 7-23, 9-8, 9-33
 - starting 7-22
 - terminating 7-28
- Dialogue command area 5-4
- Dialogue I/O area 2-2, 2-3, 3-5, 5-6, 5-11, 5-24, 5-28
- Dialogue input 7-30, 7-31, 7-34, 7-60
- Dialogue message area 5-4
- Dialogue name 7-24
- Dialogue NightView menu 2-6
- Dialogue output 3-5, 7-30, 7-32
- Dialogue prompt 6-2
- Dialogue shell 1-2, 2-2, 3-4, 3-5, 5-6, 9-55, E-1
 - login 7-22
 - logout 7-28
- Dialogue Window 2-3, 5-11
- Dimmed button 5-20, 9-27, 9-31
- Dimmed label 9-29
- Directory
 - current 7-71
 - debug file 7-25
- directory** 3-17, 7-71, 7-73, 7-74
- Directory searching 3-17, 7-73, 7-74, 7-139
- disable** 4-26, 7-92, 7-108
- Disabled button 5-20, 9-27, 9-31
- Disabling a breakpoint 4-22, 4-26, 5-21, 5-26, 7-108, 7-109, 9-29
- Disabling a heappoint 7-109, 9-29
- Disabling a monitorpoint 7-109, 9-29
- Disabling a patchpoint 7-109, 9-29
- Disabling a tracepoint 7-109, 9-29
- Disabling a watchpoint 7-109, 9-29
- Disabling an eventpoint 7-108, 7-109, 9-29
- Disassembly 7-69, 7-85, 7-117, 7-118, 9-48, Glossary-4
- Disassembly display mode 9-11, 9-48, 9-49, 9-53, 9-55
- Dismiss** button 9-2
- display** 3-17, 4-24, 7-4, 7-88, 7-90, 7-118, 7-139
- Display addresses limits 7-58, 7-140
- Display arguments 7-142, 7-149
- Display array 7-58
- Display breakpoint 4-27, 5-26, 7-131, 7-132
- Display checkpoint information 7-147
- Display convenience variables 7-139
- Display declaration 7-152
- Display dialogue information 7-146
- Display display variables 7-139
- DISPLAY environment variable 6-2
- Display eventpoint 4-27, 5-26, 7-131, 7-146
- Display expression 7-139, 7-152, 9-14, 9-70
- Display expression limits 7-140
- Display family information 7-146
- Display file names 7-153
- Display function names 7-151
- Display global variable 7-150
- Display heap errors 9-15, 9-71
- Display heap information 9-15, 9-70
- Display heap leaks 9-15, 9-71
- Display heappoint 7-131, 7-136
- Display item Glossary-4
- Display line number 7-153
- Display local variables 7-149, 9-14, 9-61, 9-70
- Display log file information 7-131
- Display macro 7-158
- Display mode
 - disassembly 9-11, 9-48, 9-49, 9-53, 9-55
 - source 9-48
- Display monitorpoint 7-131, 7-135
- Display monitorpoint values 9-15, 9-71

- Display on `program` commands 7-147
- Display on `restart` commands 7-147
- Display patchpoint 4-27, 5-26, 7-131, 7-134
- Display process information 7-141, 9-15, 9-70
- Display registers 9-62
- Display search path 7-139
- Display shell information 9-15, 9-70
- Display source file 4-8, 4-10, 7-72, 7-128
- Display source file names 7-151
- Display stack frame
 - all 4-17, 5-16, 7-81, 7-88, 9-14, 9-70
 - one 7-138
- Display still allocated blocks 9-15, 9-71
- Display string limits 7-58
- Display thread information 9-15, 9-70
- Display tracepoint 7-131, 7-133
- Display type 7-152
- Display type information 7-151
- Display value history 7-140
- Display variable 1-5, 2-6, 3-17, 4-15, 4-19, 4-24, 5-14, 5-18, 5-23, 7-81, 7-152
- Display variable address 7-150
- Display watchpoint 7-131, 7-137
- Documentation
 - online 1-1, 1-4, 2-1, 2-3, 4-4, 5-4, 7-127, 9-1, Glossary-6, Glossary-9
- Double clicking 5-14, 5-18, 9-27, 9-76, 9-77
- down** 4-19, 7-125, 7-126, 7-127, 9-20, 9-55
- Down button 5-19
- Down button 9-20
- DWARF 3-37, 7-56, Glossary-4
- dynamic library translations 7-40
- Dynamic linker 3-39
- Dynamically loaded library 3-4, 3-20, 3-39, 7-26

E

- e** key 9-55
- echo** 7-20, 7-87, 7-157
- edit** 7-75
- Edit source file 7-75, 9-11, 9-53, 9-55
- Editor
 - emacs 7-67, 8-2
 - gmacs 7-67, 8-2
 - vi 7-67, 8-2
- EDITOR environment variable 8-2
- ELF 3-37, 7-74, Glossary-4
- emacs editor 7-67, 8-2
- enable** 7-92, 7-94, 7-108, 7-109, 9-20, 9-55
- Enabling a breakpoint 7-108, 9-29
- Enabling a heappoint 7-108, 9-29
- Enabling a monitorpoint 7-108, 9-29

- Enabling a patchpoint 7-108, 9-29
- Enabling a tracepoint 7-108, 9-29
- Enabling a watchpoint 3-13, 7-108, 9-29
- Enabling an eventpoint 7-108, 9-29
- end define** 7-154, 7-155
- end on dialogue** 7-28
- end on program** 7-42
- end on restart** 7-44
- Enter key 9-55
- Entry point 7-92
- Environment variable
 - DISPLAY 6-2
 - EDITOR 8-2
 - NIGHTVIEW_ENV 3-5, 3-7
 - NSLM_SERVER A-2
 - PATH 3-9
 - SHELL 7-129
 - TERM 3-5, 8-1
 - VISUAL 8-2
- Error
 - abort 3-33
 - caution 3-33
 - warning 3-33
- Error message 7-138, 9-56
- Errors 1-4, 2-3, 3-33, 4-4, 5-4, 9-18
- Evaluation of expressions 3-19, 7-83, 7-88
- Event notification 7-35, 7-36
- Event-driven command streams 3-34
- Event-map file 7-98, 7-99, Glossary-5
- Eventpoint 3-9, 3-16, 3-17, 3-18, 7-16, 7-92, 7-98, 7-108, 7-133, 7-134, 7-135, 7-136, 9-12, 9-27, 9-57, Glossary-5
 - changing 9-28
 - clearing 7-104, 9-22
 - commands on 9-30
 - condition on 3-10, 3-17, 3-35, 7-106, 7-110, 7-132, 9-29
 - deleting 7-107, 9-31
 - disabling 7-108, 9-29
 - displaying 4-27, 5-26, 7-131
 - enabling 7-108, 9-29
 - hitting 7-132, 9-29, Glossary-7
 - ignoring 3-10, 3-17, 7-109, 7-132, 9-30
 - inserted 3-10, 3-11, 3-13, 3-14, 3-19, 3-35, 7-5, Glossary-7
 - named 3-17, 9-30
 - naming 7-16, 7-94
 - printing 7-146
 - removing 7-105
 - saving 3-17
 - setting 3-3, 9-28
 - state 7-131, 9-29
- Eventpoint crossing count Glossary-3
- Eventpoint Dialog Boxes 9-27

- Eventpoint ID 9-29
 - Eventpoint menu 9-12, 9-27
 - Debug 9-27
 - Eventpoint modifier 7-94, Glossary-5
 - /delete 7-94
 - /disabled 7-94, Glossary-5
 - Eventpoint number 3-9, 9-29
 - Eventpoint state 9-57
 - Eventpoint summary 5-20, 5-22, 5-26, 9-27, 9-57
 - Eventpoints
 - changing 5-22, 5-26, 9-27, 9-28, 9-57
 - Event-triggered commands 7-29, 7-42
 - Exception handling 7-122
 - exec 3-17, 4-7, 5-7, 7-37, 7-42, 7-44, 7-47, 7-48, 9-24
 - exec-file** 3-9, 3-40, 7-26, 7-39, 7-40, 7-42, 7-71
 - Executable
 - stripped 7-26, 7-38
 - Executable and linking format Glossary-4
 - Executable file 3-1, 7-38, 7-39, 7-40, 7-149, 7-153
 - Execution
 - continuing 1-5, 2-5, 2-6, 3-37, 4-10, 4-14, 4-20, 4-29, 5-10, 5-13, 5-19, 5-27, 7-114, 7-115, 7-119, 7-120, 7-121, 7-127, 9-55
 - restarting 3-16, 3-17, 7-42, 7-44, 7-45, 7-48, Glossary-1, Glossary-10
 - resuming 1-5, 2-5, 2-6, 4-10, 4-14, 4-20, 4-29, 5-10, 5-13, 5-19, 5-27, 7-113, 7-114, 9-19
 - starting 1-2, 2-2, 3-16
 - stopping 1-4, 2-5, 3-37, 4-9, 4-21, 4-22, 5-9, 5-20, 5-21, 7-95, 7-112
 - Exit messages 5-29
 - Exiting 1-6, 2-6, 3-17, 4-30, 5-29, 7-21, 7-48
 - Expression 3-23, 7-57
 - conditional 3-21, 4-21, 5-20, 7-96, 7-99, 7-100
 - displaying 7-88, 7-139, 9-14, 9-22, 9-70
 - evaluation 3-19, 3-20, 7-83, 7-88
 - floating-point 3-20
 - insertion 7-96, 7-111
 - language 7-89, 7-90, 7-91, 7-97
 - logging 7-99, 9-31
 - memory address 7-89
 - patchpoint 7-99, 9-30
 - printing 1-5, 4-15, 4-19, 4-24, 5-14, 5-18, 5-23, 7-81, 7-89, 7-135, 7-152, 9-22
 - regular 5-15, 7-16, 7-28, 7-66, 7-76, 7-140, 7-146, 7-149, 7-150, 7-151, 7-152, 7-158, 9-10, 9-26, 9-51
 - regular examples 7-18
 - syntax 7-3
 - Expression Evaluation 3-23
 - Expression limits 7-140
 - Expressions
 - monitoring 7-101
 - External data definitions 7-92
 - External variable 7-125
 - printing 7-150
- ## F
- f (frame)** 7-124, D-9
 - f key 9-55
 - F1 key 9-1, 9-2
 - fact program 1-1, 2-1
 - Family 3-2, 3-18, 4-12, 4-14, 4-20, 7-14, 7-22, 7-46, 7-58, 7-94, 7-120, 7-146, Glossary-5
 - printing 7-146
 - family** 4-12, 4-14, 7-39, 7-41, 7-45
 - File
 - .NightViewrc 3-36, 7-27, 7-29, 7-30
 - .profile 3-7
 - commands 7-129
 - core 3-3, 3-4, 3-19, 6-1, 7-39, 7-153, 9-24, Glossary-3
 - event-map 7-98, 7-99, Glossary-5
 - executable 3-1, 7-38, 7-39, 7-40, 7-149, 7-153
 - initialization 6-2, 6-3, 7-129, Glossary-7
 - library 3-1
 - log 7-56, 7-131
 - object 3-1, 3-12, 7-91
 - ReadyToDebug 1-3, 3-9, 4-4, 5-4
 - source 3-1, 4-8, 4-10, 5-1, 5-7, 5-10, 5-11, 5-13, 5-14, 5-15, 5-16, 5-18, 5-19, 5-21, 5-22, 5-26, 7-72, 7-73, 7-74, 7-75, 7-129, 7-139, 7-151, 7-153, 9-10, 9-48, 9-51
 - symbol 7-38, 7-41
 - trace event-map 7-98, 7-99, Glossary-5
 - File access 3-1, 3-7
 - File menu 9-4
 - File name
 - printing 7-153
 - File Selection Dialog Box 9-11, 9-52
 - Filter
 - PID 9-26
 - Program 9-26
 - User 9-26
 - Filter button 9-26
 - finish** 4-29, 7-119, 7-127, 9-20, 9-24, 9-55
 - Finish button 9-20
 - Fixed licenses A-1
 - Floating licenses A-1
 - Floating-point expressions 3-20
 - fo (forward-search)** 7-76, D-6
 - Focus
 - keyboard 9-1, Glossary-5, Glossary-8
 - Font
 - default Glossary-4

forbid safety level 6-2, 7-21, 7-28, 7-37, 7-46, 7-61,
7-107, 7-108, 7-109
Forking 3-2, 3-3, 4-7, 4-14, 5-7, 5-13, 7-37, 9-23,
Glossary-5
Formal argument
 macro 7-154, 7-156
Fortran i-iii, 3-10, 3-22, 3-23, 4-1, 4-3, 4-15, 4-16, 4-17,
4-18, 4-19, 4-20, 4-22, 4-23, 4-25, 4-28, 5-1,
5-3, 5-9, 5-12, 5-16, 5-17, 5-18, 5-19, 5-21,
5-22, 5-27, 7-56, 7-82, 7-97, 7-111, 7-155,
Glossary-2
forward-search 7-17, 7-73, 7-76
Frame
 displaying 7-138
 stack 3-23, 3-24, 4-18, 4-19, 4-29, 5-17, 5-19, 5-27,
7-4, 7-73, 7-79, 7-92, 7-95, 7-97, 7-99,
7-101, 7-102, 7-105, 7-111, 7-119, 7-124,
7-125, 7-126, 7-138, 7-141, 7-151,
Glossary-3, Glossary-5, Glossary-11
 stack - printing 4-17, 5-16, 7-81, 9-62
frame 3-24, 7-116, 7-119, 7-124, 7-127, 9-55
Frame address 7-6, 7-124, 7-138
Frame pointer 7-6, 7-139
Frame zero 7-6, 7-64, 7-79, 7-121, 7-124, 7-125, 7-126,
7-141
Frames
 hidden 7-6, 7-64, 7-79, 7-121
Full-screen interface i-iii, 1-1, 3-27, 6-2, 7-2, 7-101,
7-106, 7-128, 7-129, 7-155, 8-1, 8-2, 9-1,
Glossary-6
Full-screen user interface 7-127
Function 4-10, 4-16, 5-10, 5-15, 7-116, 7-117, 7-118,
7-119, 7-139
 static - location of 7-13
Function arguments
 printing 7-149
Function name
 list 7-151

G

-g option 1-2, 2-2, 4-2, 5-3
gdb 1-1, 1-4
GID 3-38
Global data definitions 7-92
Global variable 3-19, 7-125
 printing 7-150
gmacs editor 7-67, 8-2
Graphical user interface i-iii, 2-1, 3-27, 3-34, 6-1, 6-2,
7-1, 7-37, 7-123, 7-127, 9-1, Glossary-6
Group ID 3-38
GUI 2-1, 3-27, 3-34, 6-1, 6-2, 7-1, 7-37, 7-123, 9-1,

Glossary-6

Guide
 command summary C-1

H

H (source line decoration) 7-78
h key 9-55
handle 3-15, 3-17, 4-7, 5-8, 7-114, 7-121, 7-122,
7-141
hd (heapdebug) 7-50, D-4
Heap 7-142
heapcheck 3-32, 7-53, 7-148
heapdebug 3-12, 3-28, 3-29, 3-32, 7-50, 7-103, 7-149,
9-31, 9-32
Heappoint 3-9, 7-37, 7-78, 7-92, Glossary-5
 changing 9-28
 clearing 7-104
 condition on 3-17, 7-106, 7-132, 7-137, 9-29
 deleting 7-107, 9-31
 disabling 7-108, 9-29
 displaying 7-131, 7-136
 enabling 7-108, 9-29
 hitting 7-137, 9-29
 ignoring 3-17, 7-109, 7-132, 7-137, 9-30,
Glossary-7
 named 3-17, 7-94, 7-102, 9-30
 saving 3-17
 setting 7-92, 9-28
 state 7-136, 9-29
 temporary 9-29
heappoint 3-12, 3-32, 7-53, 7-102, 7-103, 7-136
Heappoint crossing count Glossary-3
Heappoint Dialog Box 9-13, 9-22, 9-27
Help
 context-sensitive 9-2, 9-17
 cross reference 5-5
help 1-1, 1-4, 3-33, 4-4, 7-127, 9-1, 9-2, 9-18,
Glossary-6, Glossary-9
Help button 9-2, 9-26, 9-27, 9-31, 9-72, 9-77
Help menu 2-1, 2-3, 2-4, 5-4, 5-5, 5-7, 9-1, 9-2, 9-17
-help option 6-1
Help system
 movement 9-1, 9-80
Help Window 2-1, 2-3, 2-4, 5-5, 5-8, 9-1, 9-2, 9-80,
Glossary-6
Help window
 exiting 2-4, 2-5
Hidden frames 7-6, 7-64, 7-79, 7-121
History
 command 3-36
 value 3-36, 4-16, 5-14, 7-3, 7-58, 7-83, 7-87, 7-140,

- Glossary-12
- Hit a breakpoint 9-29
- Hit a heappoint 9-29
- Hit a monitorpoint 9-29
- Hit a patchpoint 9-29
- Hit a tracepoint 9-29
- Hit a watchpoint 9-29
- Hit an eventpoint 9-29
- hold (mcontrol hold)** 7-103
- hold (mcontrol-hold)** D-7
- Hold button 9-60
- Hollerith data 7-82
- I**
- i b (info breakpoint)** 7-132, D-10
- I/O 3-5
- IA-32 registers 7-6
- Iconifying windows 9-1
- ID
- group 3-38
 - process 3-3, 3-5, 3-18, 4-6, 5-6, 5-13, 7-15, 9-26, Glossary-9
 - trace-event 7-98, 7-99, 9-31, Glossary-12
 - user 3-38
- ignore** 4-22, 7-92, 7-98, 7-109, Glossary-7
- Ignore count Glossary-7
- Ignoring breakpoints 3-17, 4-22, 5-21, 7-96, 7-109, 7-114, 7-132, 7-133, 9-30, Glossary-7
- Ignoring eventpoints 3-10, 3-17, 7-109, 7-132, 9-30
- Ignoring heappoints 3-17, 7-109, 7-132, 7-137, 9-30, Glossary-7
- Ignoring monitorpoints 3-17, 7-109, 7-136, 9-30, Glossary-7
- Ignoring patchpoints 3-17, 7-98, 7-109, 7-132, 7-135, 9-30, Glossary-7
- Ignoring tracepoints 3-17, 7-100, 7-109, 7-132, 7-134, 9-30, Glossary-7
- Ignoring watchpoints 3-17, 7-109, 7-132, 7-138, 9-30, Glossary-7
- info address** 7-150
- info args** 7-149
- info breakpoint** 7-131, 7-132
- info convenience** 7-139
- info declaration** 7-152
- info dialogue** 7-25, 7-63, 7-146
- info directories** 7-74, 7-139
- info display** 7-90, 7-91, 7-139
- info eventpoint** 4-27, 7-108, 7-131
- info family** 7-146
- info files** 7-153
- info frame** 7-4, 7-138
- info functions** 7-151
- info heappoint** 7-136
- info history** 7-140
- info limits** 7-59, 7-140
- info line** 7-74, 7-153
- info locals** 7-149
- info log** 7-131
- info macros** 7-158
- info memory** 7-63, 7-142
- info monitorpoint** 7-131, 7-135
- info name** 7-146
- info on dialogue** 7-29, 7-147
- info on program** 7-43, 7-147
- info on restart** 3-17, 7-44, 7-64, 7-147
- info patchpoint** 7-131, 7-134
- info process** 7-141
- info registers** 3-24, 7-9, 7-13, 7-139, 7-140
- info representation** 7-152
- info signal** 7-121, 7-141
- info sources** 7-151
- info tracepoint** 7-131, 7-133
- info types** 7-151
- info variables** 7-150
- info watchpoint** 7-137
- info whatis** 7-152
- Initialization file 6-2, 6-3, 7-129, Glossary-7
- Initialize tracing 7-98
- Inline interest level 7-64
- Inline subprograms 3-25, 7-117, 7-119
- Input
- dialogue 7-31, 7-34, 7-60
 - program 1-3, 2-3, 3-5, 4-11, 5-11, 7-31, 7-60
 - shell 9-55
- Input area 5-21, 5-22, 5-23, 5-25
- Input command 7-129
- Input terminator 7-31, 7-60
- Inserted eventpoints 3-10, 3-11, 3-13, 3-14, 3-19, 3-35, 7-5, Glossary-7
- Instruction
- branch 7-97, 7-111
- interest** 3-26, 7-63, 7-66, Glossary-7
- Interest level
- inline 7-64
 - justlines 3-17, 3-26, 7-64
 - nodebug 3-17, 3-26, 7-64
 - subprogram 3-17, 3-25, 7-63, Glossary-7
- Interest level threshold 3-17, 3-26, 7-64, Glossary-7
- Interesting subprograms 3-17, 3-24, 3-25, 4-13, 5-13, 7-6, 7-64, 7-117, 7-119, Glossary-7
- Interface
- command-line i-iii, 1-1, 3-27, 4-1, 4-3, 5-1, 5-3, 7-2, 7-101, 7-106, 7-123, 7-155, 8-1, 9-1, 9-19, Glossary-2
 - full-screen i-iii, 1-1, 3-27, 6-2, 7-2, 7-101, 7-106,

7-128, 7-129, 7-155, 8-1, 8-2, 9-1,

Glossary-6

graphical user i-iii, 2-1, 3-27, 3-34, 6-1, 6-2, 7-1,

7-37, 7-123, 9-1, Glossary-6

Interrupt button 9-21

Interrupt the debugger 9-21

Interrupting the debugger 3-27, 3-34, 7-104, 9-21

Interrupts 3-39

Invoking NightStar Tools 9-16

Invoking the debugger 1-2, 2-2, 3-37, 4-3, 5-3, 6-1

IPL register 3-39

J

Job control 3-18

jump 7-120

Justlines interest level 3-17, 3-26, 7-64

K

Key

= 9-55

> 9-55

b 9-55

d 9-55

e 9-55

Enter 9-55

f 9-55

F1 9-1, 9-2

h 9-55

N 9-55

n 9-55

p 9-55

r 9-55

Return 4-4, 4-8, 4-12, 4-25, 5-4, 5-6, 5-8, 5-11,

5-17, 5-18, 5-19, 5-21, 5-22, 5-23, 5-25,

7-19, 7-20, 7-73, 8-1, 9-26

S 9-55

s 9-55

Space 4-4, 4-12

u 9-55

Keyboard focus 9-1, Glossary-8

Kill 9-20

kill 3-14

kill 7-37, 9-9, 9-20, 9-21

Kill button 7-37, 9-20

Killing processes 3-17, 7-21, 7-37, 9-9, 9-21

L

l (list) 1-4, 7-72, D-5

Label

dimmed 9-29

Language 7-100, 7-120

machine 3-1, 3-37

Language expression 7-89, 7-90, 7-91, 7-97

Language support i-iii, 3-20, 3-23, 4-1, 5-1, 7-56, 7-142

Library

dynamically loaded 3-4, 3-20, 3-39, 7-26

shared 3-4, 3-20, 3-39, 7-26

Library file 3-1

licences 6-3

License A-1

firewall configurations A-3, A-5

fixed A-1

floating A-3

installation A-1

keys A-1

modes A-1

`ns1m_admin` A-1, A-3

report A-3

requests A-2

server A-2, A-3

support A-6

License manager 6-3

Limits

addresses 7-140

expression 7-140

Line decorations 2-5, 4-8, 4-10, 4-11, 4-14, 4-16, 5-7,

5-10, 5-11, 5-12, 5-13, 5-16, 5-18, 5-19, 5-21,

5-26, 7-74, 7-78, 7-117, 7-118, 9-11, 9-20,

9-21, 9-22, 9-49, 9-53

Line number

printing 7-153

Linking 1-2, 2-2, 4-2, 5-3, 7-100

dynamic Glossary-4

Linux 3-37

list 1-4, 4-8, 4-10, 7-19, 7-71, 7-72, 7-74, 7-76

List assembly code 9-11, 9-12, 9-53

List function names 7-151

List mode 9-48

List source file 4-8, 4-10, 7-72

List source file names 7-151

load 3-12, 7-26, 7-91

Loading Data Window layout 9-16, 9-71

local dialogue 1-3, 3-5, 4-4, 6-1, 7-1, 7-43

with on dialogue 7-29, 7-30

Local system 3-6

Local variable 3-19, 7-4, 7-88, 9-14, 9-61, 9-70

printing 7-149

Location

- in executable program 7-13
 - printing 7-58, 7-150
 - Location specifier 3-23, 7-13, 7-16, 9-12, 9-20, 9-22, 9-28, 9-30
 - Log
 - dialogue 7-32
 - Log file 7-56, 7-131
 - Logging
 - session 7-56
 - login** 3-7, 7-22
 - logout** 7-28, 9-8
- M**
- M (source line decoration) 7-78
 - Machine language 3-1, 3-37
 - Macro 3-18, 3-35, 7-19, 7-31, 7-153, Glossary-8
 - actual arguments 7-156
 - argument 7-154, 7-156
 - definition 7-155
 - example 7-19, 7-20
 - formal arguments 7-154, 7-156
 - printing 7-158
 - recursion 7-155
 - referencing 7-156
 - replacing 7-155
 - restart_begin_hook 3-18, 7-42
 - restart_end_hook 3-18
 - string 7-156
 - Macro body 3-35, 7-154, 7-159
 - Macro expansion 7-29, 7-42
 - Manual
 - online 1-1, 1-4, 2-1, 2-3, 4-4, 5-4, 7-127, 9-1, Glossary-6, Glossary-9
 - Manual section 4-4, 5-4, 7-127
 - mcontrol** 3-26, 7-103, 8-2
 - Memory 7-125
 - output 7-86, 7-89
 - shared 7-62, E-1
 - static 7-151
 - Memory address expression 7-89
 - Memory layout 7-142
 - Memory mapped I/O 3-39
 - Menu
 - context 9-3
 - Data** 9-14
 - Debug Eventpoint 5-20, 5-22, 5-25, 5-26, 9-27
 - Debug NightView 2-6, 5-29
 - Debug Source 5-14, 9-10, 9-26, 9-49, 9-55
 - Debug View 5-18, 5-28
 - Dialogue NightView 2-6
 - Eventpoint 9-12, 9-27
 - File 9-4
 - Help 2-1, 2-3, 2-4, 5-4, 5-5, 5-7, 9-1, 9-2, 9-17
 - Process 9-8
 - Shell 9-8
 - Source 9-10, 9-26, 9-49, 9-55
 - Tools 9-16
 - View 9-5
 - Menu bar 5-5
 - Debug 9-4
 - Message
 - error 7-138, 9-56
 - exit 5-29
 - process status 9-56
 - Message panel 9-56
 - message panel 5-6
 - mmap E-1
 - Mnemonic Glossary-8
 - A 9-5, 9-6, 9-9, 9-11, 9-15, 9-52, 9-66, 9-68, 9-69
 - B 9-7, 9-12, 9-50, 9-67
 - C 9-5, 9-15, 9-17, 9-69
 - D 9-4, 9-6, 9-7, 9-9, 9-12, 9-14, 9-53, 9-58, 9-68, 9-69
 - E 2-3, 9-6, 9-12, 9-13, 9-14, 9-18, 9-50, 9-58, 9-68
 - F 9-4, 9-5, 9-52, 9-56, 9-68, 9-69
 - G 9-6, 9-52, 9-56, 9-58, 9-70
 - H 2-1, 2-3, 2-4, 9-6, 9-8, 9-9, 9-13, 9-15, 9-17, 9-50, 9-51, 9-67, 9-69
 - I 9-7, 9-11, 9-53, 9-58, 9-72
 - K 9-9, 9-15, 9-18, 9-66
 - k 9-54
 - L 9-5, 9-8, 9-14, 9-18, 9-50, 9-58, 9-65
 - M 9-6, 9-11, 9-12, 9-18, 9-50, 9-53, 9-58
 - m 2-4
 - N 2-6, 9-7, 9-54, 9-58, 9-67, 9-68
 - n 2-1
 - O 9-4, 9-11, 9-53, 9-69
 - P 9-5, 9-8, 9-12, 9-15, 9-16, 9-50
 - Q 9-18
 - R 9-6, 9-8, 9-14, 9-52, 9-66
 - S 9-4, 9-6, 9-9, 9-10, 9-14, 9-51, 9-59, 9-67, 9-69
 - T 9-7, 9-8, 9-13, 9-15, 9-16, 9-17, 9-18, 9-51, 9-65, 9-66
 - U 9-10, 9-13, 9-17, 9-18, 9-51, 9-68, 9-69
 - V 9-5, 9-15, 9-19, 9-66
 - W 9-13, 9-51
 - X 2-6, 9-5, 9-65
 - Y 9-53
 - Z 9-7, 9-59
 - Mode
 - list 9-48
 - Monitor refresh rate 7-104
 - Monitor Window 3-26, 7-103, 8-2
 - simple full-screen 8-2
 - Monitoring expressions 7-101

Monitorpoint 3-9, 3-11, 3-12, 3-19, 3-21, 3-26, 3-39,
7-37, 7-62, 7-78, 7-92, 7-101, 7-103,
Glossary-5, Glossary-8
changing 9-28
clearing 7-104
commands on 3-17, 7-105, 7-132, 7-136, 9-30
condition on 3-17, 7-106, 7-132, 7-136, 9-29
deleting 7-107, 9-31
disabling 7-108, 9-29
displaying 7-131, 7-135
enabling 7-108, 9-29
hitting 7-136, 9-29
ignoring 3-17, 7-109, 7-132, 7-136, 9-30,
Glossary-7
named 3-17, 7-94, 7-101, 9-30
saving 3-17
setting 7-92, 7-101, 9-28
state 7-135, 9-29
temporary 9-29
monitorpoint 3-11, 3-26, 7-101, 7-135, Glossary-8
Monitorpoint crossing count Glossary-3
Monitorpoint Dialog Box 9-12, 9-21, 9-27
Monitorpoint Update Interval Dialog Box 9-80
Motif 9-1
Mouse button 1 2-1, 5-2, 5-14, 5-18, 9-1, 9-2, 9-60
Mouse button 3 9-60
mreserve 7-49
msg program 4-3, 4-7, 4-10, 4-20, 5-3, 5-7, 5-10, 5-19
Multiple processes i-iii, 3-2, 3-4, 7-124, 7-138, 7-139,
7-141, 7-151, 8-1

N

n (next) 7-117, D-8
N key 9-55
n key 9-55
name 7-92, 7-94, 7-97, 7-111
Named breakpoint 3-17, 7-94, 7-95, 7-110, 9-30
Named eventpoint 3-17, 7-16, 7-94, 9-30
Named heappoint 3-17, 7-102, 9-30
Named monitorpoint 3-17, 7-101, 9-30
Named patchpoint 3-17, 7-97, 7-111, 9-30
Named process 4-12, 4-14
Named tracepoint 3-17, 7-99, 9-30
Named watchpoint 3-17, 7-112
Newline 7-19, 7-73
next 4-10, 7-116, 7-117, 7-119, 7-127, 9-19, 9-55
Next button 5-10, 9-19
nexti 7-116, 7-118, 7-119, 7-127, 9-20, 9-55
Nexti button 9-20
NFS 3-2
ni (nexti) 7-119, D-9

nice value 7-23
NightProbe 9-16
NightStar Licence Manager 6-3
NightStar tool set 9-16
NightProbe 9-16
NightTrace 9-17
NightTune 9-17
NightTrace 3-6, 3-12, 7-99, 7-100, 9-17, Glossary-9
NightTune 9-17
NightView i-iii, 3-1, Glossary-9
NightView version 1-2, 4-4, 6-2, 9-19, G-1
NIGHTVIEW_ENV environment variable 3-5, 3-7
NLSM 6-3
nodebug 3-2, 3-5, 7-18, 7-21, 7-24, 7-25, 7-28, 7-29
Nodebug interest level 3-17, 3-26, 7-64
-nogui option 1-2, 4-3, 6-1
-nolocal option 6-1
Notification of events 7-32, 7-35, 7-36, 7-123, 8-1
notify 7-36
nslm_admin A-1, A-3
NSLM_SERVER A-2
ntrace 3-12, 7-98
ntraceud 3-38, 7-100
nview
invoking 6-1
nview
exiting 1-6, 2-6, 4-30, 5-29, 7-21
invoking 1-2, 2-2, 3-37, 4-3, 5-3
nview 1-2, 2-2, 8-1
nview option
-config 6-1
-core 6-1
-help 4-3, 5-3
-nogui 1-2, 4-3, 6-1
-nolocal 6-1
-simplescreen 8-1
nview options 6-1

O

Object activation 5-2, 5-20, 5-22, 5-26
Object file 3-1, 7-91
Object filename translations 7-26, 7-38, 7-40, 7-92,
7-146
Object selection 5-2, 5-5, 5-7, 5-14, 5-20, 5-21, 5-22,
5-25, 5-26, 5-29
OK button 2-6, 5-15, 5-21, 5-22, 5-23, 5-29, 9-27, 9-31,
9-72, 9-73, 9-74, 9-75, 9-76, 9-77, 9-78, 9-80
on dialogue
with local dialogue 7-29, 7-30
on dialogue 7-27, 7-28, 7-30, 7-147
on program 3-18, 7-18, 7-41, 7-43, 7-44, 7-147

- on restart** 3-16, 3-18, 7-42, 7-44, 7-147
 - Online documentation 1-1, 1-4, 2-1, 2-3, 4-4, 5-4, 7-127, 9-1, Glossary-6, Glossary-9
 - Online help system Glossary-9
 - Optimization 3-37, 7-117, 7-118
 - Option
 - attach 6-1
 - config 6-1
 - core 6-1
 - g 1-2, 2-2, 4-2, 5-3
 - help 4-3, 5-3
 - nogui 1-2, 4-3, 6-1
 - nolocal 6-1
 - simplescreen 8-1
 - x 6-3
 - xrm 6-2
 - Options
 - nview 6-1
 - Output 3-5
 - buffered 3-35
 - dialogue 3-5, 7-32
 - logging 3-36
 - memory 7-86, 7-89
 - program 9-55
 - session 7-56
 - shell 9-55
 - suppressed 8-1
 - text 7-87
 - output** 7-87
 - Output addresses limits 7-58
 - Output array 7-58
 - Output string limits 7-58
 - Output variable 1-5, 2-6, 4-15, 4-19, 4-24, 5-14, 5-18, 5-23, 7-81
 - Overloading 3-23, Glossary-9
- P**
- p (print)** 1-5, 7-81, D-6
 - P (source line decoration) 7-79
 - p key 9-55
 - Panel
 - data 9-14
 - message 9-56
 - panel
 - context 5-7, 5-12, 5-13
 - shell 5-4, 5-11
 - Parent process 4-14, 5-13, 9-26
 - Patch Glossary-9
 - patch area Glossary-9
 - Patchpoint 3-9, 3-21, 3-39, 7-37, 7-79, 7-80, 7-92, Glossary-5, Glossary-9
 - changing 9-28
 - clearing 7-104
 - condition on 3-17, 7-98, 7-106, 7-132, 7-135, 9-29
 - deleting 7-107, 9-31
 - disabling 7-108, 9-29
 - displaying 4-27, 5-26, 7-131, 7-134
 - enabling 7-108, 9-29
 - hitting 7-135, 9-29
 - ignoring 3-17, 7-98, 7-109, 7-132, 7-135, 9-30, Glossary-7
 - named 3-17, 7-94, 7-111, 9-30
 - saving 3-17
 - setting 4-25, 5-25, 7-92, 7-96, 7-111, 9-28
 - state 7-134, 9-29
 - temporary 7-111, 9-29
 - patchpoint** 3-12, 4-25, 7-96, 7-134
 - Patchpoint crossing count Glossary-3
 - Patchpoint Dialog Box 5-25, 9-13, 9-21, 9-27
 - Patchpoints named 7-97
 - PATH environment variable 3-9
 - Pattern
 - wildcard 7-18, 7-24, 7-42, 7-44, 7-151, 9-10, 9-26, 9-51
 - wildcard examples 7-19
 - Pattern matching 7-16, 7-24, 7-25, 7-66, 7-76, 7-146, Glossary-9
 - Pattern matching examples 7-18
 - PID 3-3, 3-5, 3-18, 4-6, 5-6, 5-13, 7-15, 9-26, Glossary-9
 - Pipelines 3-2, 3-4
 - Pointer
 - question mark 5-5, 9-2
 - Predefined convenience variable 3-24, 7-3, 7-5, 7-132, 7-133, 7-134, 7-135, 7-136, 7-137, 7-140, 7-153
 - Preferences 9-5
 - print
 - command attached to monitorpoint 7-102
 - print** 1-5, 3-36, 4-15, 4-19, 7-57, 7-81, 7-83, 7-87, 7-89, 7-91, 7-101, 7-125, 7-151, 7-156, 9-22, 9-55
 - Print addresses limits 7-58, 7-140
 - Print arguments 7-142, 7-149
 - Print array 7-58
 - Print breakpoint 4-27, 5-26, 7-131, 7-132
 - Print button 2-6, 5-14, 5-18, 9-22, 9-55
 - Print checkpoint information 7-147
 - Print convenience variables 7-139
 - Print current directory 7-71
 - Print declaration 7-152
 - Print dialogue information 7-146
 - Print display variables 7-139
 - Print eventpoint 4-27, 5-26, 7-131
 - Print eventpoint information 7-146

- Print expression 1-5, 4-15, 4-19, 4-24, 5-14, 5-18, 5-23, 7-81, 7-89, 7-152
- Print expression limits 7-140
- Print family information 7-146
- Print file names 7-153
- Print function names 7-151
- Print global variable 7-150
- Print heappoint 7-131, 7-136
- Print line number 7-153
- Print local variables 7-149
- Print log file information 7-131
- Print macro 7-158
- Print monitorpoint 7-135
- Print on `dialogue` commands 7-147
- Print on `program` commands 7-147
- Print on `restart` commands 7-147
- Print patchpoint 4-27, 5-26, 7-131, 7-134
- Print process information 7-141
- Print registers 7-140
- Print search path 7-139
- Print signal 7-141
- Print source file names 7-151
- Print stack frame
 - all 4-17, 5-16, 7-81
 - one 7-138
- Print string limits 7-58
- Print text 7-87
- Print tracepoint 7-131, 7-133
- Print type information 7-151, 7-152
- Print value history 7-140
- Print variable 7-152
- Print variable address 7-150
- Print watchpoint 7-131, 7-137
- Print Window 9-5
- printf** 7-155, 7-157
- Procedure 4-10, 4-16, 5-10, 5-15, 7-116, 7-117, 7-118, 7-119, 7-139, Glossary-9
- Procedure arguments
 - printing 7-149
- Procedure call 7-97
- Procedure name
 - list 7-151
- Process 3-2, 9-8, E-1, Glossary-10
 - abnormal termination 7-39
 - attaching to 6-1, 7-36, 9-9, 9-26
 - background 7-129
 - child 3-2, 3-3, 4-1, 4-7, 4-14, 5-1, 5-7, 5-13, 7-25, 7-37, 7-47, Glossary-1
 - current 9-3
 - exiting 3-17, 3-19, 7-48
 - killing 3-17, 7-21
 - multiple 3-2, 3-4, 8-1
 - naming 4-12, 4-14
 - parent 4-14, 5-13, 9-26
 - printing 7-141
 - pseudo 3-3, 3-19, 6-1, 7-39, 7-41
 - running 3-18, 3-20
 - single 3-2
 - stopped 3-18, 3-20, 3-23, 3-24
 - stopping 7-120
 - stopping debugging 7-37
 - terminated 3-19
 - terminating 3-17, 3-18
- Process families 3-2
- Process ID 3-3, 3-5, 3-18, 4-6, 5-6, 5-13, 7-15, 9-26, Glossary-9
- Process** menu 9-8
- Process state 3-18, 7-142, Glossary-10
- Processes
 - multiple i-iii, 7-124, 7-138, 7-139, 7-141, 7-151
- Program 3-2, Glossary-10
 - commands on 7-42
 - compiling 1-2, 2-2, 3-37, 4-2, 5-3
 - fact 1-1, 2-1
 - msg 4-3, 4-7, 4-10, 4-20, 5-3, 5-7, 5-10, 5-19
 - restarting 3-16, 3-17, 7-42, 7-44, 7-45, 7-48, Glossary-1, Glossary-10
 - running 1-2, 4-6, 5-6, 7-31, 7-34
 - setuid 3-3
 - starting 3-16
- Program arguments 6-1
- Program counter 3-20, 3-23, 3-24, 7-5, 7-121, 7-140, Glossary-10
- Program I/O E-1
- Program input 1-3, 2-3, 3-5, 4-11, 5-11, 7-31, 7-60
- Program location
 - specifying 7-13
- Program name 1-1, 2-1, 3-9, 4-3, 4-7, 5-3, 5-7, 6-2, 9-23
- Program output 3-5, 3-35, 3-36, 9-55
- Program, run 9-8
- Prologue 7-13
- Prompt 7-2, 7-59
 - \$ 1-3, 4-4, 5-4
 - (local) 1-3, 4-4
 - > 7-101, 7-106, 7-155
 - command 4-4, 4-11
 - dialogue 6-2, 7-2
 - shell 1-3, 4-4, 5-4
- ps 3-3
- Pseudo process 3-3, 3-19, 6-1, 7-39, 7-41
- ptrace** 3-3, E-1
- ptype (info declaration)** 7-152, D-11
- pwd** 7-71

Q

q (quit) 1-6, 7-2, 7-21, D-1
 Qualifier 3-4, 4-14, 4-20, 7-1, 7-14, 7-58, Glossary-10
 Quick command summary C-1
 quit 6-3
quit 1-6, 4-30, 7-2, 7-21, 9-5
 Quitting 1-6, 2-6, 4-30, 5-29, 7-21

R

r key 9-55
 Radio button 9-30, 9-73
 ReadyToDebug 1-3, 3-9, 4-4, 5-4
 Real-time debugging 3-6
 Recursion
 macro 7-155
redisplay 7-90
 Referencing macros 7-156
refresh 7-128
 Refreshing terminal 7-128
regexp 7-16, 7-76, 7-140, 7-146, 7-149, 7-150, 7-151,
 7-152, 7-158
 Register
 IPL 3-39
 Register variable 7-4
 Registers 3-1, 3-20, 3-23, 3-24, 7-3, 7-5, 7-88, 7-121,
 7-125, 7-139, 7-151, 9-14, 9-62, 9-70,
 Glossary-10
 AMD64 7-9
 display 9-62
 IA-32 7-6
 printing 7-140
 Regular expression 5-15, 7-16, 7-28, 7-66, 7-76, 7-140,
 7-146, 7-149, 7-150, 7-151, 7-152, 7-158, 9-10,
 9-26, 9-51
 Regular expression examples 7-18
release (mcontrol release) 7-103, D-7
 Release button 9-60
 Remote dialogue 3-6, 7-23, 9-8, 9-33, Glossary-10
 Remote file access 3-7
 Remote Login Dialog Box 9-33
 Remote system 3-6
 Repeating commands 4-14, 7-1, 7-19, 7-73
 Replacing commands 7-157
representation (info representation)
 7-152, D-11
rerun 7-34, 9-21
 Rerun button 9-21
 Rerunning a program 3-16, 3-17, 7-42, 7-44, 7-45, 7-48
 Restart

 commands on 7-44
 restart_begin_hook macro 3-18, 7-42
 restart_end_hook macro 3-18
 Restarting a program 3-17, 7-42, 7-48
 Restarting execution 3-16, 3-17, 7-42, 7-44, 7-45, 7-48,
 Glossary-1, Glossary-10
resume 3-15, 3-37, 4-20, 7-106, 7-113, 7-115, 9-19,
 9-20, 9-55
 Resume button 2-5, 2-6, 5-10, 5-11, 5-20, 5-24, 5-28,
 9-19
 Resuming execution 1-5, 2-5, 2-6, 4-10, 4-14, 4-20,
 4-29, 5-10, 5-13, 5-19, 5-27, 7-113, 7-114, 9-19
 Return key 4-4, 4-8, 4-12, 4-25, 5-4, 5-6, 5-8, 5-11,
 5-17, 5-18, 5-19, 5-21, 5-22, 5-23, 5-25, 7-19,
 7-20, 7-73, 8-1, 9-26
reverse-search 7-17, 7-73, 7-76
 rlogin 7-23
 Routine 4-10, 4-16, 5-10, 5-15, 7-116, 7-117, 7-118,
 7-119, 7-139, Glossary-11
 trace_open_thread 7-100
 trace_start 7-100
 Routine arguments
 printing 7-149
 Routine name
 list 7-151
 Routine replacement 7-92
 run 3-7, 9-33
run 1-2, 3-5, 4-6, 7-31, 7-34, 7-157
 Run a program 1-2, 3-16, 4-6, 5-6, 7-31, 7-34, 9-8, 9-21
 Run to Here button 9-20, 9-55

S

s (step) 7-116, D-8
 S key 9-55
 s key 9-55
 Safety level
 forbid 6-2, 7-21, 7-28, 7-37, 7-46, 7-61, 7-107,
 7-108, 7-109
 unsafe 3-34, 6-2, 7-20, 7-28, 7-37, 7-61, 7-107,
 7-108, 7-109
 verify 6-2, 7-20, 7-21, 7-28, 7-37, 7-46, 7-61,
 7-107, 7-108, 7-109
 Saving breakpoints 3-17
 Saving Data Panel contents 9-15, 9-71
 Saving Data Window layout 9-16, 9-71
 Saving eventpoints 3-17
 Saving heappoints 3-17
 Saving monitorpoints 3-17
 Saving patchpoints 3-17
 Saving tracepoints 3-17
 Saving watchpoints 3-17

- Scope 3-23, 3-35, 7-95, 7-100, 7-103, 7-150, Glossary-11
 - Script
 - debugger 7-129, 7-130
 - Scroll bar 2-4, 2-5, 5-5, 5-7, 9-27
 - Search button 9-26
 - Searching
 - function 9-10, 9-51
 - path 7-73, 7-74, 7-139
 - regular expression 5-15, 7-76, 9-26
 - wildcard pattern 9-10, 9-26, 9-51
 - Section
 - manual 4-4, 5-4, 7-127
 - Select a Function/Unit Dialog Box 5-15, 9-10, 9-26, 9-51
 - Select a Source File Dialog Box 9-10, 9-26, 9-51, 9-52
 - select-context** 3-37, 7-119, 7-126
 - Selection
 - object 5-2, 5-5, 5-7, 5-14, 5-20, 5-21, 5-22, 5-25, 5-26, 5-29
 - Selection policy
 - Browse 9-27
 - Semicolon 7-97, 7-111
 - Session
 - debug Glossary-4
 - Session logging 7-56
 - set** 7-83
 - set-auto-frame** 7-65
 - set-children** 3-2, 3-17, 4-7, 5-7, 7-47
 - set-debug-file-directory** 7-25
 - set-disassembly** 7-69
 - set-download** 7-68
 - set-editor** 7-67, 8-2
 - set-exit** 7-47
 - set-history** 7-58
 - set-language** 3-17, 7-56, 7-142
 - set-limits** 7-58, 7-83, 7-85, 7-131, 7-132, 7-133, 7-134, 7-135, 7-136
 - set-local** 3-35, 7-61
 - set-log** 7-56, 7-131
 - set-notify** 7-35
 - set-overload** 3-23, 7-66
 - set-patch-area-size** 7-62, 7-146, E-1, Glossary-9
 - set-preallocate** 7-67
 - set-prompt** 7-2, 7-59
 - set-qualifier** 4-20, 7-14, 7-57
 - set-restart** 3-17, 7-61
 - set-safety** 7-29, 7-42, 7-60
 - set-search** 7-66
 - set-show** 3-5, 7-32, 7-33, 7-56, 7-131
 - set-terminator** 7-60
- Setting a breakpoint 1-4, 2-5, 4-9, 4-21, 4-22, 5-9, 5-20, 5-21, 7-95, 7-110, 9-21, 9-28, 9-55
 - Setting a conditional breakpoint 3-10, 5-20, 7-132, 7-133, 9-29, Glossary-2
 - Setting a conditional eventpoint 7-106, 7-110, 7-132, 9-29
 - Setting a conditional heappoint 7-132, 7-137, 9-29
 - Setting a conditional monitorpoint 7-132, 7-136, 9-29
 - Setting a conditional patchpoint 7-132, 7-135, 9-29
 - Setting a conditional tracepoint 7-132, 7-134, 9-29
 - Setting a conditional watchpoint 7-132, 7-138, 9-29
 - Setting a heappoint 9-28
 - Setting a monitorpoint 7-101, 9-28
 - Setting a patchpoint 4-25, 5-25, 7-96, 7-111, 9-28
 - Setting a tracepoint 7-100, 9-28
 - Setting a watchpoint 7-112, 9-28
 - Setting an eventpoint 9-28
 - set-trace** 7-98, 7-99
 - Setuid programs 3-3
 - Shared library 3-4, 3-20, 3-39, 7-26
 - Shared memory 7-62, E-1
 - Shell Glossary-11
 - dialogue 3-4, 3-5, 5-6, 9-55, E-1
 - shell** 3-3, 3-38, 7-21, 7-71, 7-128, 7-157, 8-1, 9-19
 - SHELL environment variable 7-129
 - Shell I/O 9-55
 - Shell menu 9-8
 - shell panel 5-4, 5-11
 - Shell prompt 1-3, 4-4, 5-4
 - show** 3-5, 7-32, 7-33
 - si (steppi)** 7-118, D-8
 - SIGALRM 7-123
 - siginfo 3-15
 - SIGINT 7-123
 - signal** 3-15, 7-106, 7-114, 7-121
 - Signals 3-14, 3-17, 4-1, 4-7, 5-1, 5-8, 7-35, 7-43, 7-89, 7-114, 7-115, 7-116, 7-117, 7-118, 7-119, 7-120, 7-121, 7-122, Glossary-11
 - printing 7-141
 - SIGQUIT 7-123
 - SIGSTOP 7-114, 7-115, 7-122, 7-123
 - SIGTRAP 3-15, 7-120
 - SIGUSR1 4-1, 4-7, 4-25, 4-30, 5-1, 5-8
 - Simple full-screen interface i-iii, 1-1, 3-27, 6-2, 7-2, 7-101, 7-106, 7-128, 7-129, 7-155, 8-1, 8-2, 9-1, Glossary-6
 - editing commands 8-2
 - simplscreen option 8-1
 - Single process 3-2
 - Single stepping 3-15, 3-37, 3-39, 4-10, 4-16, 5-10, 5-15, 7-116, 7-117, 7-118, 7-119, 7-120, 9-20, 9-23, 9-55
 - source** 3-17, 3-34, 3-36, 7-19, 7-129, 7-148
 - Source display area 9-10, 9-27
 - Source display mode 9-48
 - Source file 3-1, 5-1, 5-7, 5-10, 5-11, 5-13, 5-14, 5-15,

- 5-16, 5-18, 5-19, 5-21, 5-22, 5-26, 7-129,
7-153, 9-10, 9-48, 9-51
- current 7-73, 7-76, 7-77
- displaying 4-8, 4-10, 7-72, 7-128
- editing 7-75, 9-11, 9-53, 9-55
- list 7-151
- search path for 7-73, 7-74, 7-139
- Source line decorations 2-5, 4-8, 4-10, 4-11, 4-14, 4-16,
5-7, 5-10, 5-11, 5-12, 5-13, 5-16, 5-18, 5-19,
5-21, 5-26, 7-74, 7-78, 7-117, 7-118, 9-11,
9-20, 9-21, 9-22, 9-49, 9-53
- Source listing 1-4, 2-3, 4-8, 4-10, 7-19, 7-72
- Source menu 9-10, 9-26, 9-49, 9-55
 - Debug 9-10, 9-26, 9-49, 9-55
- source panel 5-7, 5-13
- Space key 4-4, 4-12
- Stack Glossary-11
- Stack examination 1-5, 2-6, 4-17, 5-16, 7-81
- Stack frame 7-4, 7-73, 7-92, 7-151, Glossary-5,
Glossary-11
 - current 3-23, 3-24, 4-18, 4-19, 4-29, 5-17, 5-19,
5-27, 7-79, 7-95, 7-97, 7-99, 7-101, 7-102,
7-105, 7-111, 7-119, 7-124, 7-125, 7-126,
7-138, 7-141, Glossary-3
 - displaying 7-88, 7-138, 9-14, 9-70
 - printing 4-17, 5-16, 7-81
- Stack pointer 7-6, 7-140
- Stack variable 3-19
- Stale data indicator 3-27, 8-2, Glossary-12
- Starting execution 1-2, 2-2, 3-16
- Starting the debugger 1-2, 2-2, 3-37, 4-3, 5-3, 6-1
- Starting tracing 7-98
- State
 - breakpoint 7-133, 9-29
 - eventpoint 7-131, 9-29
 - heappoint 7-136, 9-29
 - monitorpoint 7-135, 9-29
 - patchpoint 7-134, 9-29
 - process 3-18, 7-142, Glossary-10
 - tracepoint 7-133, 9-29
 - watchpoint 7-137, 9-29
- Static data definitions 7-92
- Static function
 - specifying location of 7-13
- Static memory 7-151
- Static variable 3-19, 7-4
- Status bar 5-6
- status bar 5-7, 5-13
- Status information 7-130
- step** 3-25, 3-37, 4-16, 5-15, 7-113, 7-116, 7-117,
7-118, 7-127, 9-20, 9-55
- Step button 9-20
- stepi** 7-116, 7-118, 7-127, 9-20, 9-55
- Stepi button 9-20
- stop** 7-120, 9-19, 9-24
- Stop button 9-19
- Stopping a process 7-120, 9-19
- Stopping execution 1-4, 2-5, 3-37, 4-9, 4-21, 4-22, 5-9,
5-20, 5-21, 7-95, 7-112
- Stream
 - command Glossary-2
- String
 - C 3-38, 7-83
 - character 7-83, 7-140
 - macro 7-156
- String limits
 - printing 7-58
- strip 7-38
- Stripped executable 7-26, 7-38
- stty 8-1
- Subprogram 4-10, 4-16, 5-10, 5-15, 7-116, 7-117,
7-118, 7-119, 7-139
- Subprogram arguments
 - printing 7-149
- Subprogram interest level 3-17, 3-25, 7-63, Glossary-7
- Subprogram name
 - list 7-151
- Subprograms
 - inline 3-25
 - interesting 3-17, 3-24, 3-25, 4-13, 5-13, 7-6, 7-64,
7-117, 7-119, Glossary-7
 - uninteresting 3-17, 3-24, 3-25, 4-13, 5-13, 7-6,
7-64, 7-117, 7-119
- Subroutine 4-10, 4-16, 5-10, 5-15, 7-116, 7-117, 7-118,
7-119, 7-139
- Subroutine arguments
 - printing 7-149
- Subroutine call 3-22, 7-97
- Subroutine name
 - list 7-151
- Substitution
 - text 7-154
- Summary of commands 5-5, 9-18, C-1
- Summary of eventpoints 5-20, 5-22, 5-26, 9-27, 9-57
- Switch To button 5-12
- Symbol file 3-17, 7-38, 7-41, Glossary-12
- Symbol table 7-92, 7-130, 7-149, 7-151, 7-153
- symbol-file** 3-7, 3-17, 7-26, 7-38, 7-39, 7-71
- Symbolic debug information 7-38
- Symbolic debugger i-iii, 3-1
- Symbols
 - undefined 7-92
- Syntax
 - command 7-1
 - expression 7-3
 - qualifier 7-1
- System
 - local 3-6

remote 3-6
system 3-2, 7-47
System crash 3-39

T

T (source line decoration) 7-79
Tag
 trace-event 7-98, 7-99, Glossary-12
Task 3-37, 7-119, 7-126
tbreak 7-92, 7-110
Temporary breakpoint 7-110, 9-29
Temporary heappoint 9-29
Temporary monitorpoint 9-29
Temporary patchpoint 7-111, 9-29
Temporary tracepoint 9-29
Temporary watchpoint 9-29
TERM environment variable 3-5, 8-1
Terminal refresh 7-128
Terminating a process 3-17, 3-18, 7-37, 9-9, 9-21
Termination
 abnormal 7-39
Terminator
 input 7-31, 7-60
Text
 printing 7-87
Text cursor 9-20, 9-22
Text input area 5-21, 5-22, 5-23, 5-25, 9-28, 9-29, 9-30
Text substitution 7-154
Thread 3-37, 7-119, 7-126, 9-62, Glossary-12
Thread process 3-37, 7-15, 7-119, 7-126, 7-148, 9-62,
 Glossary-12
Threshold
 interest level 3-17, 3-26, 7-64
Toggle button 9-30
TOOLS menu 9-16
tpatch 7-92, 7-111
Trace Glossary-12
Trace initialization 7-98
trace_open_thread routine 7-100
trace_start routine 7-100
Trace-event ID 7-98, 7-99, 9-31, Glossary-12
Trace-event map file 7-98, 7-99, Glossary-5
Trace-event tag 7-98, 7-99, Glossary-12
Tracepoint 3-9, 3-12, 3-19, 3-21, 3-39, 7-37, 7-79, 7-80,
 7-92, 7-98, Glossary-5, Glossary-12
 changing 9-28
 clearing 7-104
 condition on 3-17, 7-106, 7-132, 7-134, 9-29
 deleting 7-107, 9-31
 disabling 7-108, 9-29
 displaying 7-131, 7-133

 enabling 7-108, 9-29
 hitting 7-134, 9-29
 ignoring 3-17, 7-100, 7-109, 7-132, 7-134, 9-30,
 Glossary-7
 named 3-17, 7-94, 7-99, 9-30
 saving 3-17
 setting 7-92, 7-100, 9-28
 state 7-133, 9-29
 temporary 9-29
tracepoint 3-12, 3-38, 7-99, 7-133
Tracepoint crossing count Glossary-3
Tracepoint Dialog Box 9-13, 9-22, 9-27
Tracing 3-6, 3-12, 3-38, 3-39
translate-object-file 3-7, 7-26, 7-38, 7-40,
 7-41
Translations
 dynamic library 7-40
 object filename 7-26, 7-38, 7-40, 7-92, 7-146
Tutorial
 command-line 4-1, 5-1
Type definition
 printing 7-151, 7-152

U

u key 9-55
UID 3-38
ulimit 3-5
Undefined symbols 7-92
undisplay 7-90
Uninteresting subprograms 3-17, 3-24, 3-25, 4-13, 5-13,
 7-6, 7-64, 7-117, 7-119
unsafe safety level 3-34, 6-2, 7-20, 7-28, 7-37, 7-61,
 7-107, 7-108, 7-109
up 3-24, 4-18, 7-4, 7-125, 7-127, 9-20, 9-55
Up button 5-17
Up button 9-20
User 7-146
User ID 3-38
User interface
 command-line i-iii, 1-1, 3-27, 4-1, 4-3, 5-1, 5-3,
 7-2, 7-101, 7-106, 7-123, 7-127, 7-155,
 8-1, 9-1, 9-19, Glossary-2
 full-screen i-iii, 1-1, 3-27, 6-2, 7-2, 7-101, 7-106,
 7-127, 7-128, 7-129, 7-155, 8-1, 8-2, 9-1,
 Glossary-6
 graphical i-iii, 2-1, 3-27, 3-34, 6-1, 6-2, 7-1, 7-37,
 7-123, 7-127, 9-1, Glossary-6

V

Value history 3-36, 4-16, 5-14, 7-3, 7-58, 7-83, 7-87, 7-140, Glossary-12

Variable

- assignment 3-21
- convenience 3-35, 7-3, 7-5, 7-62, 7-83, 7-139, Glossary-3
- declaration 3-21
- global 3-19, 7-125
- local 3-19, 3-23, 3-24, 7-4, 7-88, 9-14, 9-61, 9-70
- predefined convenience 3-24, 7-3, 7-5, 7-132, 7-133, 7-134, 7-135, 7-136, 7-137, 7-140, 7-153
- printing 7-149, 7-150, 7-152
- register 3-19, 7-4
- static 3-19, 7-4

verify safety level 6-2, 7-20, 7-21, 7-28, 7-37, 7-46, 7-61, 7-107, 7-108, 7-109

Version

NightView 1-2, 4-4, 6-2, 9-19, G-1

vi editor 7-67, 8-2

View menu 9-5

Virtual address space 7-142

Virtual memory 7-142

VISUAL environment variable 8-2

W

wait 7-48

Warning Dialog Box 5-29, 9-2, 9-5, 9-8, 9-9, 9-21

Warnings 3-33

Watchpoint

- changing 9-28
- commands on 3-17
- condition on 3-13, 3-17, 7-132, 7-138, 9-29
- deleting 7-107, 9-31
- disabling 9-29
- displaying 7-131, 7-137
- enabling 3-13, 7-108, 9-29
- hitting 7-138, 9-29
- ignoring 3-17, 7-109, 7-132, 7-138, 9-30, Glossary-7
- named 3-17, 7-112
- saving 3-17
- setting 7-112, 9-28
- state 7-137, 9-29
- temporary 9-29

watchpoint 7-112

Watchpoint Dialog Box 9-13, 9-22, 9-27

whatis(**info whatis**) 7-152, D-11

Wildcard pattern 7-18, 7-24, 7-42, 7-44, 7-151, 9-10, 9-26, 9-51

Wildcard pattern examples 7-19

wildcard_pattern 7-151

Window

Data 3-17, 5-23, 5-25, 5-28, 7-88, 9-22, 9-60

Debug 2-3, 5-9, 5-10, 5-11, 5-14, 5-16, 5-20, 5-24, 5-25, 9-4, 9-23, 9-26, 9-27

Dialogue 2-3, 5-11

Help 2-1, 2-3, 2-4, 5-5, 5-8, 9-1, 9-2, 9-80, Glossary-6

iconifying 9-1

Monitor 3-26, 7-103, 8-2

print 9-5

X

X 9-1

x 7-5, 7-19, 7-84, 7-89, 7-132, 7-133, 7-134, 7-135, 7-136, 7-137, 7-153

-x option 6-3

X Window System 3-33, 9-1

x1(**translate-object-file**) 7-26, D-1

-xrm option 6-2

