

System Analysis Toolkit (SAT) User's Guide



©2001–2014, QNX Software Systems Limited, a subsidiary of BlackBerry. All rights reserved.

QNX Software Systems Limited
1001 Farrar Road
Ottawa, Ontario
K2K 0B3
Canada

Voice: +1 613 591-0931
Fax: +1 613 591-3579
Email: info@qnx.com
Web: <http://www.qnx.com/>

QNX, QNX CAR, Neutrino, Momentics, Aviage, and Foundry27 are trademarks of BlackBerry Limited that are registered and/or used in certain jurisdictions, and used under license by QNX Software Systems Limited. All other trademarks belong to their respective owners.

Electronic edition published: Thursday, February 20, 2014

Table of Contents

About This Guide	7
Typographical conventions	8
Technical support	10
 Chapter 1: Introduction	 11
What can the SAT do for you?	12
Components of the SAT	14
Instrumented kernel	14
Kernel buffer management	15
Data-capture program (tracelogger)	15
Data interpretation (e.g. traceprinter)	16
Integrated Development Environment	16
 Chapter 2: Events and the Kernel	 19
Generating events: a typical scenario	20
Multithreaded example	20
Thread context-switch time	21
Restarting threads	21
Simple and combine events	22
Fast and wide modes	23
Classes and events	24
Communication class: _NTO_TRACE_COMM	24
Control class: _NTO_TRACE_CONTROL	25
Interrupt classes: _NTO_TRACE_INTENTER, _NTO_TRACE_INTEXIT, _NTO_TRACE_INT_HANDLER_ENTER, and _NTO_TRACE_INT_HANDLER_EXIT	26
Kernel-call classes: _NTO_TRACE_KERCALLEENTER, _NTO_TRACE_KERCALLEEXIT, and _NTO_TRACE_KERCALLINT	26
Process class: _NTO_TRACE_PROCESS	30
System class: _NTO_TRACE_SYSTEM	30
Thread class: _NTO_TRACE_THREAD	31
User class: _NTO_TRACE_USER	33
Virtual thread class: _NTO_TRACE_VTHREAD	34
 Chapter 3: Kernel Buffer Management	 37
Linked list size	38
Full buffers and the high-water mark	39
Buffer overruns	40

Chapter 4: Capturing Trace Data	41
Using tracelogger to control tracing	43
Managing trace buffers	43
tracelogger's modes of operation	43
Choosing between wide and fast modes	44
Filtering events	45
Specifying where to send the output	45
Using TraceEvent() to control tracing	46
Managing trace buffers	46
Modes of operation	47
Filtering events	48
Choosing between wide and fast modes	48
Inserting trace events	49
 Chapter 5: Filtering	 51
The static rules filter	53
The dynamic rules filter	56
Setting up a dynamic rules filter	56
Event handler	57
Removing event handlers	59
The post-processing facility	60
 Chapter 6: Interpreting Trace Data	 61
Using traceprinter and interpreting the output	63
Building your own parser	66
The traceparser library	66
Simple and combine events	67
The traceevent_t structure	67
Event interlacing	67
Timestamps	68
 Chapter 7: Tutorials	 69
The instrex.h header file	70
Gathering all events from all classes	71
Gathering all events from one class	74
Gathering five events from four classes	76
Gathering kernel calls	79
Event handling - simple	83
Inserting a user simple event	87
 Appendix A: Current Trace Events and Data	 89
Interpreting the table	90

Table of events	93
-----------------------	----

About This Guide

The QNX Neutrino System Analysis Toolkit *User's Guide* describes how to use the instrumented microkernel to obtain a detailed analysis of what's happening in an entire QNX Neutrino system. This guide contains the following sections and chapters:

To find out about:	Go to:
What the SAT is, what it can do for you, and how it works	Introduction (p. 11)
What generates events	Events and the Kernel (p. 19)
How the kernel buffers data	Kernel Buffer Management (p. 37)
How to save data	Capturing Trace Data (p. 41)
Different ways to reduce the amount of data	Filtering (p. 51)
What the data tells you	Interpreting Trace Data (p. 61)
Examples of filtering the trace data	Tutorials (p. 69)
What specific events return	Current Trace Events and Data

Typographical conventions

Throughout this manual, we use certain typographical conventions to distinguish technical terms. In general, the conventions we use conform to those found in IEEE POSIX publications.

The following table summarizes our conventions:

Reference	Example
Code examples	<code>if(stream == NULL)</code>
Command options	<code>-lR</code>
Commands	<code>make</code>
Environment variables	<i>PATH</i>
File and pathnames	<code>/dev/null</code>
Function names	<code>exit()</code>
Keyboard chords	Ctrl –Alt –Delete
Keyboard input	Username
Keyboard keys	Enter
Program output	login:
Variable names	<i>stdin</i>
Parameters	<i>parm1</i>
User-interface components	Navigator
Window title	Options

We use an arrow in directions for accessing menu items, like this:

You'll find the Other... menu item under **Perspective → Show View** .

We use notes, cautions, and warnings to highlight important messages:



Notes point out something important or useful.



Cautions tell you about commands or procedures that may have unwanted or undesirable side effects.



Warnings tell you about commands or procedures that could be dangerous to your files, your hardware, or even yourself.

Note to Windows users

In our documentation, we use a forward slash (/) as a delimiter in all pathnames, including those pointing to Windows files. We also generally follow POSIX/UNIX filesystem conventions.

Technical support

Technical assistance is available for all supported products.

To obtain technical support for any QNX product, visit the Support area on our website (www.qnx.com). You'll find a wide range of support options, including community forums.

Chapter 1

Introduction

In many computing environments, developers need to monitor a dynamic execution of realtime systems with emphasis on their key architectural components. Such monitoring can reveal hidden hardware faults and design or implementation errors, as well as help improve overall system performance.

In order to accommodate those needs, we provide sophisticated tracing and profiling mechanisms, allowing execution monitoring in real time or offline. Because it works at the operating system level, the SAT, unlike debuggers, can monitor applications without having to modify them in any way.

The main goals for the SAT are:

- ease of use
- insight into system activity
- high performance and efficiency with low overhead

What can the SAT do for you?

In a running system, many things occur behind the scenes:

- Kernel calls are being made.
- Messages are being passed.
- Interrupts are being handled.
- Threads are changing states— they're being created, blocking, running, restarting, and dying.

The results of this activity are changes to the system state that are normally hidden from developers. The SAT is capable of intercepting these changes and logging them. Each event is logged with a timestamp and the ID of the CPU that handled it.



For a full understanding of how the kernel works, see the QNX Neutrino Microkernel chapter in the *System Architecture* guide.

The SAT offers valuable information at all stages of a product's life cycle, from prototyping to optimization to in-service monitoring and field diagnostics.




Figure 1: The SAT view and the debugger view.

In complicated systems, the information provided by standard debugging programs may not be detailed enough to solve the problem. Or, the problem may not be a bug as much as a process that's not behaving as expected. Unlike the SAT, debuggers lack the execution history essential to solving the many complex problems involved in “application tuning.” In a large system, often consisting of many interconnected components or processes, traditional debugging, which lets you look at only a single module, can't easily assist if the problem lies in how the modules interact with each other. Where a debugger can view a single process, the SAT can view *all* processes at the same time. Also, unlike debugging, the SAT doesn't need code augmentation and can be used to track the impact of external, precompiled code.

Because it offers a system-level view of the internal workings of the kernel, the SAT can be used for performance analysis and optimization of large interconnected systems as well as single processes.

It allows realtime debugging to help pinpoint deadlock and race conditions by showing what circumstances led up to the problem. Rather than just a “snapshot”, the SAT offers a “movie” of what's happening in your system.

Because the instrumented version of the kernel runs with negligible performance penalties, you can optionally leave it in the final embedded system. Should any problems arise in the field, you can use the SAT for low-level diagnostics.

The SAT offers a nonintrusive method of instrumenting the code—programs can literally monitor themselves. In addition to passive/non-intrusive event tracing, you can proactively trace events by injecting your own “flag” events.

Components of the SAT

The QNX Neutrino System Analysis Toolkit (SAT) consists of the following main components:

- [Instrumented kernel](#) (p. 14)
- [Kernel buffer management](#) (p. 15)
- [Data-capture program \(tracelogger\)](#) (p. 15)
- [Data interpretation \(e.g. traceprinter\)](#) (p. 16)

You can also trace and analyze events under control of the [Integrated Development Environment](#) (p. 16).




Figure 2: Overall view of the SAT.

Instrumented kernel

The instrumented kernel is actually the regular QNX Neutrino microkernel with a small, highly efficient event-gathering module included. Except for the instrumentation, its operation is virtually indistinguishable—the instrumented kernel runs at 98% of the speed of our regular microkernel.

As threads run, the instrumented kernel continuously intercepts information about what the kernel is doing, generating time-stamped and CPU-stamped events that are stored in a circular linked list of buffers. Because the tracing occurs at the kernel level, the SAT can track the performance of *all* processes, including the data-capturing program.

To check to see if your system is running the instrumented kernel, type:

```
ls /proc/boot
```

and then look for a file whose name includes `procnto`. If the file name is `procnto-*instr`, you're running the instrumented kernel; if the file name doesn't include `instr`, you're running the noninstrumented kernel.

To switch to the instrumented kernel, do the following:

1. In your buildfile, replace the entry for `procnto` with the appropriate `procnto-*instr`. For more information, see the entry for `procnto` in the *Utilities Reference*.
2. Add `tracelogger` and `traceprinter` to your buildfile or target.
3. Run the `mkifs` utility to rebuild the image. For more information, see the entry for `mkifs` in the *Utilities Reference*.
4. Replace your current boot image with the new one.

Kernel buffer management

The kernel buffer is composed of many small buffers. Although the number of buffers is limited only by the amount of system memory, it's important to understand that this space must be managed carefully. If *all* of the events are being traced on an active system, the number of events can be quite large.

To allow the instrumented kernel to write to one part of the kernel buffer and store another part of it simultaneously, the kernel buffer is organized as a circular linked list. As the buffer data reaches a high-water mark (about 70% full), the instrumented kernel module sends a signal to the data-capture program with the address of the buffer. The data-capture program can then retrieve the buffer and save it in a storage location for offline processing or pass it to a data interpreter for realtime manipulation. In either case, once the buffer has been emptied, it's once again available for use by the kernel.

Data-capture program (`tracelogger`)

The QNX Neutrino RTOS includes a `tracelogger` that you can use to capture data. This daemon receives events from the instrumented kernel and saves them in a file or sends them to a device for later analysis.



The data-capture utilities require `root` privileges to allocate buffer memory or to use functions such as `InterruptHookTrace()` (see the QNX Neutrino *C Library Reference*). Data-capture utilities won't work properly without these privileges.

Because the `tracelogger` may write data at rates well in excess of 20 MB/minute, running it for prolonged periods or running it repeatedly can use up a large amount of space. If disk space is low, erase old log files regularly. (In its default mode, `tracelogger` overwrites its previous default file.)

You can also control tracing from your application (e.g. to turn tracing on just for a problematic area) with the `TraceEvent()` kernel call. This function has over 30 different commands that let you:

- create internal trace buffers

- set up filters
- control the tracing process
- insert user defined events

For more information, see the [Capturing Trace Data](#) (p. 41) chapter in this guide, the entry for `tracelogger` in the *Utilities Reference*, and the entry for `TraceEvent()` in the QNX Neutrino *C Library Reference*.

Data interpretation (e.g. `traceprinter`)

To aid in processing the binary trace event data, we provide the `libtraceparser` library. The API functions let you set up a series of functions that are called when complete buffer slots of event data have been received/read from the raw binary event stream.

We also provide a linear trace event printer (`traceprinter`) that outputs all of the trace events ordered linearly by their timestamp as they're emitted by the kernel. This utility uses the `libtraceparser` library. You can also modify the `traceprinter` source as a basis for your own custom parser or use the API to create an interface to do the following offline or in real time:

- perform analysis
- display results
- debug applications
- create a self-monitoring system
- show events ordered by process or by thread
- show thread states and transitions
- show currently running threads

The `traceparser` library provides an API for parsing and interpreting the trace events that are stored in the event file. The library simplifies the parsing and interpretation process by letting you easily:

- set up callback functions and associations for each event
- retrieve header and system information from the trace event file
- debug and control the parsing process

For more information, see the [Interpreting Trace Data](#) (p. 61) chapter in this guide, as well as the entry for `traceprinter` in the *Utilities Reference*.

Integrated Development Environment

The QNX Momentics Tool Suite's IDE provides a graphical interface that you can use to capture and examine tracing events. The IDE lets you filter events, zoom in on ranges of them, examine the associated data, save subsets of events, and more.

For more information, see the Analyzing Your System with Kernel Tracing chapter of the IDE *User's Guide*.

Chapter 2

Events and the Kernel

The QNX Neutrino microkernel generates events for more than just system calls. The following are some of the activities that generate events:

- kernel calls
- scheduling activity
- interrupt handling
- thread/process creation, destruction, and state changes

In addition, the instrumented kernel also inserts “artificial” events for:

- time events
- user events that may be used as “marker flags”

Also, single kernel calls or system activities may actually generate more than one event.

Generating events: a typical scenario

Processes that are running on QNX Neutrino can run multiple threads. Having more than one thread increases the level of complexity—the OS must handle threads of differing priorities competing with each other.

Multithreaded example

In our example we'll use two threads:

Thread	Priority
A	High
B	Low

Now we'll watch them run, assuming both start at the same time:



When logging starts, the instrumented kernel sends information about each thread. Existing processes will appear to be created during this procedure.

Time	Thread	Action	Explanation
t1	A	Create	Thread is created.
t2	A	Block	The thread is waiting for, say, I/O; it can't continue without it.
t3	B	Create	Rather than sit idle, the kernel runs next highest priority thread.
t4	B	Kernel Call	Thread B is working.
t4.5	N/A	N/A	I/O completed; Thread A is ready to run.
t5	B	Block	Thread A is now ready to run—it preempts thread B.
t6	A	Run	Thread A resumes.

Time	Thread	Action	Explanation
t7	A	Dies	Its task complete, the thread terminates.
t8	B	Runs	Thread B continues from where it left off.
t9	...	...	...

Thread context-switch time

Threads don't switch instantaneously—after one thread blocks or yields to another, the kernel must save the settings before running another thread. The time to save this state and restore another is known as *thread context-switch time*. This context-switch time between threads is small, but important.




Figure 3: Thread context switching.

In some cases, two or more threads may switch back and forth without actually accomplishing much. This is akin to two overly polite people each offering to let the other pass through a narrow door first—neither of them gets to where they're going on time (two aggressive people encounter a similar problem). This type of problem is exactly what the SAT can quickly and easily highlight. By showing the context-switch operations in conjunction with thread state transitions, you can quickly see why otherwise fast systems seem to “crawl.”

Restarting threads

In order to achieve maximum responsiveness, much of the QNX Neutrino microkernel is fully preemptible. In some cases, this means that when a thread is interrupted in a kernel call, it won't be able to restart exactly where it began. Instead, the kernel call will be restarted—it “rewinds” itself. The SAT tries to hide the spurious calls but may not succeed in suppressing them all. As a result, it's possible to see several events generated from a specific thread that has been preempted. If this occurs, the last event is the actual one.

Simple and combine events

Most events can be described in a single event buffer slot; we call these *simple events*. When there's too much information to describe the event in a single buffer slot, the event is described in multiple event buffer slots; we call this a *combine event*. The event buffer slots all look the same, so there's no need for the data-capture program to distinguish between them.

For more information about simple events and combine events, see the [Interpreting Trace Data](#) (p. 61) chapter.

Fast and wide modes

You can gather data for events in the following modes:

Wide mode

The instrumented kernel uses as many buffer slots as are necessary to fully log the event. The amount of space is theoretically unlimited and can span several kilobytes for a single event. Most of the time, it doesn't exceed *four* 16-byte spaces.

Fast mode

The instrumented kernel uses *only one* buffer slot per event.

In general, wide mode generates several times more data than fast mode.



Fast mode doesn't simply clip the tail end of the event data that you'd get in wide mode; fast mode summarizes the most important aspects of the event in a single buffer slot. Thus, the first element of an event in wide mode might not be the same as the same event in fast mode.

You can set fast and wide mode for all classes, specific classes, and even specific events in a class; some can be fast while others are wide. We'll describe how to set this in the [Capturing Trace Data](#) (p. 41) chapter.

For the specific output differences between fast and wide mode, see the [Current Trace Events and Data](#) appendix.

Classes and events

There can be a lot of events in even a small trace, so they're organized into *classes* to make them easier for you to manage:

- *Communication class:* `_NTO_TRACE_COMM` (p. 24)
- *Control class:* `_NTO_TRACE_CONTROL` (p. 25)
- *Interrupt classes:* `_NTO_TRACE_INTENTER`, `_NTO_TRACE_INTEXTIT`, `_NTO_TRACE_INT_HANDLER_ENTER`, and `_NTO_TRACE_INT_HANDLER_EXIT` (p. 26)
- *Kernel-call classes:* `_NTO_TRACE_KERCALLEENTER`, `_NTO_TRACE_KERCALLEEXIT`, and `_NTO_TRACE_KERCALLINT` (p. 26)
- *Process class:* `_NTO_TRACE_PROCESS` (p. 30)
- QNX Unified Intrumentation Platform class: `_NTO_TRACE_QUIP`
- *System class:* `_NTO_TRACE_SYSTEM` (p. 30)
- *Thread class:* `_NTO_TRACE_THREAD` (p. 31)
- *User class:* `_NTO_TRACE_USER` (p. 33)
- *Virtual thread class:* `_NTO_TRACE_VTHREAD` (p. 34)

(The `<sys/trace.h>` header file also defines an `_NTO_TRACE_EMPTY` class, but it's a placeholder and isn't currently used.)

The sections that follow list the events for each class, along with a description of when the events are emitted, as well as the labels that `traceprinter` and the IDE use to identify the events.

For information about the data for each event, see the [Current Trace Events and Data](#) appendix.

Communication class: `_NTO_TRACE_COMM`

The `_NTO_TRACE_COMM` class includes events related to communication via messages and pulses.

Event	traceprinter label	IDE label	Emitted when:
<code>_NTO_TRACE_COMM_ERROR</code>	<code>MSG_ERROR</code>	Error	A client is unblocked because of a call to <code>MsgError()</code>
<code>_NTO_TRACE_COMM_REPLY</code>	<code>REPLY_MESSAGE</code>	Reply	A reply is sent
<code>_NTO_TRACE_COMM_RMSG</code>	<code>REC_MESSAGE</code>	Receive Message	A message is received
<code>_NTO_TRACE_COMM_RPULSE</code>	<code>REC_PULSE</code>	Receive Pulse	A pulse is received

Event	traceprinter label	IDE label	Emitted when:
_NTO_TRACE_COMM_SIGNAL	SIGNAL	Signal	A signal is received
_NTO_TRACE_COMM_SMSG	SND_MESSAGE	Send Message	A message is sent
_NTO_TRACE_COMM_SPULSE	SND_PULSE	Send Pulse	A pulse is sent
_NTO_TRACE_COMM_SPULSE_DEA	SND_PULSE_DEA	Death Pulse	A _PULSE_CODE_COIDDEATH pulse is sent
_NTO_TRACE_COMM_SPULSE_DIS	SND_PULSE_DIS	Disconnect Pulse	A _PULSE_CODE_DISCONNECT pulse is sent
_NTO_TRACE_COMM_SPULSE_EXE	SND_PULSE_EXE	Sigevent Pulse	A SIGEV_PULSE is sent
_NTO_TRACE_COMM_SPULSE_QUN	SND_PULSE_QUN	QNet Unblock Pulse	A _PULSE_CODE_NET_UNBLOCK pulse is sent
_NTO_TRACE_COMM_SPULSE_UN	SND_PULSE_UN	Unblock Pulse	A _PULSE_CODE_UNBLOCK pulse is sent

Control class: _NTO_TRACE_CONTROL

The _NTO_TRACE_CONTROL class includes events related to the control of tracing itself.

Event	traceprinter label	IDE label	Emitted when:
_NTO_TRACE_CONTROLBUFFER	BUFFER	Buffer	The instrumented kernel starts filling a new buffer
_NTO_TRACE_CONTROLTIME	TIME	Time	The 32 Least Significant Bits (LSB) part of the 64-bit clock rolls over, or the kernel emits an _NTO_TRACE_CONTROLBUFFER event

The purpose of emitting _NTO_TRACE_CONTROLBUFFER events is to help tracelogger and the IDE track the buffers and determine if any buffers have been dropped. The instrumented kernel emits an _NTO_TRACE_CONTROLTIME event at the same time to keep the IDE in sync (in case a dropped buffer contained an _NTO_TRACE_CONTROLTIME event for a rollover of the clock).

Interrupt classes: `_NTO_TRACE_INTENTER`, `_NTO_TRACE_INTEXTIT`, `_NTO_TRACE_INT_HANDLER_ENTER`, and `_NTO_TRACE_INT_HANDLER_EXIT`

These classes track interrupts.

Class	traceprinter label	IDE label	Emitted when:
<code>_NTO_TRACE_INTENTER</code>	<code>INT_ENTR</code>	Entry	Overall processing of an interrupt begins
<code>_NTO_TRACE_INTEXTIT</code>	<code>INT_EXIT</code>	Exit	Overall processing of an interrupt ends
<code>_NTO_TRACE_INT_HANDLER_ENTER</code>	<code>INT_HANDLER_ENTR</code>	Handler Entry	Entering an interrupt handler
<code>_NTO_TRACE_INT_HANDLER_EXIT</code>	<code>INT_HANDLER_EXIT</code>	Handler Exit	Exiting an interrupt handler

The expected sequence is:

```
INTR_ENTER
  INTR_HANDLER_ENTER
  INTR_HANDLER_EXIT
  INTR_HANDLER_ENTER
  INTR_HANDLER_EXIT
INT_EXIT
```

`_NTO_TRACE_INT` is a pseudo-class that comprises all of the interrupt classes.

The “event” is an interrupt vector number, in the range from `_NTO_TRACE_INTFIRST` through `_NTO_TRACE_INTLAST`.

Kernel-call classes: `_NTO_TRACE_KERCALLEENTER`, `_NTO_TRACE_KERCALLEEXIT`, and `_NTO_TRACE_KERCALLINT`

These classes track kernel calls.

- `_NTO_TRACE_KERCALLEENTER` and `_NTO_TRACE_KERCALLEEXIT` track the entrances to and exits from kernel calls.
- `_NTO_TRACE_KERCALLINT` tracks interrupted kernel calls. When we exit the kernel, we check to see if the kernel call arguments are valid. If so, then we log an `_NTO_TRACE_KERCALLEEXIT` event with the parameters. If not, then we log an `_NTO_TRACE_KERCALLINT` event with no parameters. If you get an `EINTR` return code from your kernel call, you'll also see an `_NTO_TRACE_KERCALLINT` event in the trace log.

`_NTO_TRACE_KERCALL` is a pseudo-class that comprises all these classes.

The traceprinter labels for these classes are KER_CALL, KER_EXIT, and INT_CALL, followed by an uppercase version of the kernel call; the IDE labels consist of the kernel call, followed by Enter, Exit, or INT.

Most of the events in these classes correspond in a fairly obvious way to the kernel calls; some correspond to internal functions:

Event	Kernel call
__KER_BAD	N/A
__KER_CHANCON_ATTR	<i>ChannelConnectAttr()</i>
__KER_CHANNEL_CREATE	<i>ChannelCreate()</i>
__KER_CHANNEL_DESTROY	<i>ChannelDestroy()</i>
__KER_CLOCK_ADJUST	<i>ClockAdjust()</i>
__KER_CLOCK_ID	<i>ClockId()</i>
__KER_CLOCK_PERIOD	<i>ClockPeriod()</i>
__KER_CLOCK_TIME	<i>ClockTime()</i>
__KER_CONNECT_ATTACH	<i>ConnectAttach()</i>
__KER_CONNECT_CLIENT_INFO	<i>ConnectClientInfo()</i>
__KER_CONNECT_DETACH	<i>ConnectDetach()</i>
__KER_CONNECT_FLAGS	<i>ConnectFlags()</i>
__KER_CONNECT_SERVER_INFO	<i>ConnectServerInfo()</i>
__KER_INTERRUPT_ATTACH	<i>InterruptAttach()</i>
__KER_INTERRUPT_DETACH	<i>InterruptDetach()</i>
__KER_INTERRUPT_DETACH_FUNC	N/A
__KER_INTERRUPT_MASK	<i>InterruptMask()</i>
__KER_INTERRUPT_UNMASK	<i>InterruptUnmask()</i>
__KER_INTERRUPT_WAIT	<i>InterruptWait()</i>
__KER_MSG_CURRENT	<i>MsgCurrent()</i>
__KER_MSG_DELIVER_EVENT	<i>MsgDeliverEvent()</i>
__KER_MSG_ERROR	<i>MsgError()</i>
__KER_MSG_INFO	<i>MsgInfo()</i>
__KER_MSG_KEYDATA	<i>MsgKeyData()</i>

Event	Kernel call
__KER_MSG_READIOV	<i>MsgReadIov()</i>
__KER_MSG_READV	<i>MsgRead()</i> , <i>MsgReadv()</i>
__KER_MSG_READWRITEV	N/A
__KER_MSG_RECEIVEPULSEV	<i>MsgReceivePulse()</i> , <i>MsgReceivePulsev()</i>
__KER_MSG_RECEIVEV	<i>MsgReceive()</i> , <i>MsgReceivev()</i>
__KER_MSG_REPLYV	<i>MsgReply()</i> , <i>MsgReplyv()</i>
__KER_MSG_SENDEV	<i>MsgSend()</i> , <i>MsgSendv()</i> , and <i>MsgSendvs()</i>
__KER_MSG_SENDEVNC	<i>MsgSendnc()</i> , <i>MsgSendvnc()</i> , and <i>MsgSendvsnc()</i>
__KER_MSG_SEND_PULSE	<i>MsgSendPulse()</i>
__KER_MSG_VERIFY_EVENT	<i>MsgVerifyEvent()</i>
__KER_MSG_WRITEV	<i>MsgWrite()</i> , <i>MsgWritev()</i>
__KER_NET_CRED	<i>NetCred()</i>
__KER_NET_INFOSCOID	<i>NetInfoScoid()</i>
__KER_NET_SIGNAL_KILL	<i>NetSignalKill()</i>
__KER_NET_UNBLOCK	<i>NetUnblock()</i>
__KER_NET_VTID	<i>NetVtid()</i>
__KER_NOP	N/A
__KER_RING0 (not generated in QNX Neutrino 6.3.0 or later)	<i>__Ring0()</i>
__KER_SCHED_GET	<i>SchedGet()</i>
__KER_SCHED_INFO	<i>SchedInfo()</i>
__KER_SCHED_SET	<i>SchedSet()</i>
__KER_SCHED_YIELD	<i>SchedYield()</i>
__KER_SIGNAL_ACTION	<i>SignalAction()</i>
__KER_SIGNAL_FAULT	N/A
__KER_SIGNAL_KILL	<i>SignalKill()</i>
__KER_SIGNAL_PROCMASK	<i>SignalProcmask()</i>
__KER_SIGNAL_RETURN	<i>SignalReturn()</i>
__KER_SIGNAL_SUSPEND	<i>SignalSuspend()</i>

Event	Kernel call
__KER_SIGNAL_WAITINFO	<i>SignalWaitInfo()</i>
__KER_SYNC_CONDVAR_SIGNAL	<i>SyncCondvarSignal()</i>
__KER_SYNC_CONDVAR_WAIT	<i>SyncCondvarWait()</i>
__KER_SYNC_CREATE	<i>SyncCreate()</i> , <i>SyncTypeCreate()</i>
__KER_SYNC_CTL	<i>SyncCtl()</i>
__KER_SYNC_DESTROY	<i>SyncDestroy()</i>
__KER_SYNC_MUTEX_LOCK	<i>SyncMutexLock()</i>
__KER_SYNC_MUTEX_REVIVE	<i>SyncMutexRevive()</i>
__KER_SYNC_MUTEX_UNLOCK	<i>SyncMutexUnlock()</i>
__KER_SYNC_SEM_POST	<i>SyncSemPost()</i>
__KER_SYNC_SEM_WAIT	<i>SyncSemWait()</i>
__KER_SYS_CPUPAGE_GET	N/A
__KER_THREAD_CANCEL	<i>ThreadCancel()</i>
__KER_THREAD_CREATE	<i>ThreadCreate()</i>
__KER_THREAD_CTL	<i>ThreadCtl()</i>
__KER_THREAD_DESTROY	<i>ThreadDestroy()</i>
__KER_THREAD_DESTROYALL	N/A
__KER_THREAD_DETACH	<i>ThreadDetach()</i>
__KER_THREAD_JOIN	<i>ThreadJoin()</i>
__KER_TIMER_ALARM	<i>TimerAlarm()</i>
__KER_TIMER_CREATE	<i>TimerCreate()</i>
__KER_TIMER_DESTROY	<i>TimerDestroy()</i>
__KER_TIMER_INFO	<i>TimerInfo()</i>
__KER_TIMER_SETTIME	<i>TimerSettime()</i>
__KER_TIMER_TIMEOUT	<i>TimerTimeout()</i>
__KER_TRACE_EVENT	<i>TraceEvent()</i>

Process class: _NTO_TRACE_PROCESS

The _NTO_TRACE_PROCESS class includes events related to the creation and destruction of processes.

Event	traceprinter label	IDE label	Emitted when:
_NTO_TRACE_PROCCREATE	PROCCREATE	Create Process	A process is created
_NTO_TRACE_PROCCREATE_NAME	PROCCREATE_NAME	Create Process Name	A newly created process is given a name.
_NTO_TRACE_PROCDESTROY	PROCDESTROY	Destroy Process	A process is destroyed
_NTO_TRACE_PROCDESTROY_NAME	—	—	(Not currently used)
_NTO_TRACE_PROCTHREAD_NAME	PROCTHREAD_NAME	Thread Name	A name is assigned to a thread

System class: _NTO_TRACE_SYSTEM

The _NTO_TRACE_SYSTEM class includes events related to the system as a whole.

Event	traceprinter label	IDE label	Emitted when:
_NTO_TRACE_SYS_ADDRESS	ADDRESS	Address	A breakpoint is hit
_NTO_TRACE_SYS_APS_BNKR	APS_BANKRUPTCY	APS Bankruptcy	An adaptive partition exceeded its critical budget
_NTO_TRACE_SYS_APS_BUDGETS	APS_NEW_BUDGET	APS Budgets	<i>SchedCtl()</i> is called with a command of SCHED_APS_CREATE_PARTITION or SCHED_APS_MODIFY_PARTITION. Also emitted automatically when the adaptive partitioning scheduler clears a critical budget as part of handling a bankruptcy.
_NTO_TRACE_SYS_APS_NAME	APS_NAME	APS Name	<i>SchedCtl()</i> is called with a command of SCHED_APS_CREATE_PARTITION
_NTO_TRACE_SYS_COMPACTION	COMPACTION	Compaction	The memory defragmentation compaction_minimal algorithm is triggered (see “Defragmenting physical memory” in the Process Manager chapter of the <i>System Architecture</i> guide)
_NTO_TRACE_SYS_FUNC_ENTER	FUNC_ENTER	Function Enter	A function that's instrumented for profiling is entered

Event	traceprinter label	IDE label	Emitted when:
_NTO_TRACE_SYS_FUNC_EXIT	FUNC_EXIT	Function Exit	A function that's instrumented for profiling is exited
_NTO_TRACE_SYS_MAPNAME	MAPNAME	MMap Name	<i>dlopen()</i> is called
_NTO_TRACE_SYS_MMAP	MMAP	MMap	<i>mmap()</i> or <i>mmap64()</i> is called
_NTO_TRACE_SYS_MUNMAP	MUNMAP	MMUnmap	<i>munmap()</i> is called
_NTO_TRACE_SYS_PATHMGR	PATHMGR_OPEN	Path Manager	An operation involving a path name — such as <i>open()</i> — that's routed via the <i>libc</i> connect function occurs. The connect function sends a message to <i>procnto</i> to resolve the path and find the set of resource managers that could potentially match the path. It's upon receiving this message that <i>procnto</i> emits this event.
_NTO_TRACE_SYS_SLOG	SLOG	System Log	A message is written to the system log

You can use the following convenience functions to insert certain System events into the trace data:

trace_func_enter()

Insert an _NTO_TRACE_SYS_FUNC_ENTER event for a function

trace_func_exit()

Insert an _NTO_TRACE_SYS_FUNC_EXIT event for a function

trace_here()

Insert an _NTO_TRACE_SYS_ADDRESS event for the current address

Thread class: _NTO_TRACE_THREAD

The _NTO_TRACE_THREAD class includes events related to state changes for threads.

Event	traceprinter label	IDE label	Emitted when a thread:
_NTO_TRACE_THCONDVAR	THCONDVAR	Condvar	Enters the CONDVAR state
_NTO_TRACE_THCREATE	THCREATE	Create Thread	Is created
_NTO_TRACE_THDEAD	THDEAD	Dead	Enters the DEAD state

Event	traceprinter label	IDE label	Emitted when a thread:
_NTO_TRACE_THDESTROY	THDESTROY	Destroy Thread	Is destroyed
_NTO_TRACE_THINTR	THINTR	Interrupt	Enters the INTERRUPT state
_NTO_TRACE_THJOIN	THJOIN	Join	Enters the JOIN state
_NTO_TRACE_THMUTEX	THMUTEX	Mutex	Enters the MUTEX state
_NTO_TRACE_THNANOSLEEP	THNANOSLEEP	NanoSleep	Enters the NANOSLEEP state
_NTO_TRACE_THNET_REPLY	THNET_REPLY	NetReply	Enters the NET_REPLY state
_NTO_TRACE_THNET_SEND	THNET_SEND	NetSend	Enters the NET_SEND state
_NTO_TRACE_THREADY	THREADY	Ready	Enters the READY state
_NTO_TRACE_THRECEIVE	THRECEIVE	Receive	Enters the RECEIVE state
_NTO_TRACE_THREPLY	THREPLY	Reply	Enters the REPLY state
_NTO_TRACE_THRUNNING	THRUNNING	Running	Enters the RUNNING state
_NTO_TRACE_THSEM	THSEM	Semaphore	Enters the SEM state
_NTO_TRACE_THSEND	THSEND	Send	Enters the SEND state
_NTO_TRACE_THSIGSUSPEND	THSIGSUSPEND	SigSuspend	Enters the SIGSUSPEND state
_NTO_TRACE_THSIGWAITINFO	THSIGWAITINFO	SigWaitInfo	Enters the SIGWAITINFO state
_NTO_TRACE_THSTACK	THSTACK	Stack	Enters the STACK state
_NTO_TRACE_THSTOPPED	THSTOPPED	Stopped	Enters the STOPPED state
_NTO_TRACE_THWAITCTX	THWAITCTX	WaitCtx	Enters the WAITCTX state
_NTO_TRACE_THWAITPAGE	THWAITPAGE	WaitPage	Enters the WAITPAGE state
_NTO_TRACE_THWAITTHREAD	THWAITTHREAD	WaitThread	Enters the WAITTHREAD state

If your system includes the adaptive partitioning scheduler module, the data for these events includes the partition ID and scheduling flags (e.g., `AP_SCHED_BILL_AS_CRIT`). For more information, see the *Adaptive Partitioning User's Guide*.

For more information about thread states, see “Thread life cycle” in the QNX Neutrino Microkernel chapter of the *System Architecture* guide.

User class: `_NTO_TRACE_USER`

The `_NTO_TRACE_USER` class includes custom events that your program creates.

You can create these events by calling one of the following convenience functions:

trace_logb()

Insert a user combine trace event

trace_logf()

Insert a user string trace event

trace_logi()

Insert a user simple trace event

trace_nlogf()

Insert a user string trace event, specifying a maximum string length

trace_vnlogf()

Insert a user string trace event, using a variable argument list

or by calling `TraceEvent()` directly, with one of the following commands:

- `_NTO_TRACE_INSERTSUSEREVENT` to create a simple event containing a small amount of data
- `_NTO_TRACE_INSERTCUSEREVENT` to create a combine event containing an arbitrary amount of data



The *len* argument for the `_NTO_TRACE_INSERTCUSEREVENT` command is the number of *integers* (not bytes) in the passed buffer.

- `_NTO_TRACE_INSERTUSRSTREVENT` to create an event containing a null-terminated string

The event must be in the range from `_NTO_TRACE_USERFIRST` through `_NTO_TRACE_USERLAST`, but you can decide what each event means.

The `traceprinter` label for these events is `USREVENT`; the IDE label is `User Event`. In both cases, this label is followed by the event type, expressed as an integer.

Virtual thread class: `_NTO_TRACE_VTHREAD`

The `_NTO_TRACE_VTHREAD` class includes events related to state changes for *virtual threads*, special objects related to Transparent Distributed Processing (TDP) over Qnet.

The kernel often keeps pointers from different data structures to relevant threads. When those threads are off-node via Qnet, there isn't a local thread object to represent them, so the kernel creates a virtual thread object.

The events for virtual threads are similar to those for normal threads, but virtual threads don't go through the same set of state transitions that normal threads do:

Event	traceprinter label	IDE label	Emitted when a virtual thread:
<code>_NTO_TRACE_VTHCONDVAR</code>	<code>VTHCONDVAR</code>	VCondvar	Enters the CONDVAR state
<code>_NTO_TRACE_VTHCREATE</code>	<code>VTHCREATE</code>	Create VThread	Is created
<code>_NTO_TRACE_VTHDEAD</code>	<code>VTHDEAD</code>	VDead	Enters the DEAD state
<code>_NTO_TRACE_VTHDESTROY</code>	<code>VTHDESTROY</code>	Destroy VThread	Is destroyed
<code>_NTO_TRACE_VTHINTR</code>	<code>VTHINTR</code>	VInterrupt	Enters the INTERRUPT state
<code>_NTO_TRACE_VTHJOIN</code>	<code>VTHJOIN</code>	VJoin	Enters the JOIN state
<code>_NTO_TRACE_VTHMUTEX</code>	<code>VTHMUTEX</code>	VMutex	Enters the MUTEX state
<code>_NTO_TRACE_VTHNANOSLEEP</code>	<code>VTHNANOSLEEP</code>	VNanosleep	Enters the NANOSLEEP state
<code>_NTO_TRACE_VTHNET_REPLY</code>	<code>VTHNET_REPLY</code>	VNetReply	Enters the NET_REPLY state
<code>_NTO_TRACE_VTHNET_SEND</code>	<code>VTHNET_SEND</code>	VNetSend	Enters the NET_SEND state
<code>_NTO_TRACE_VTHREADY</code>	<code>VTHREADY</code>	VReady	Enters the READY state
<code>_NTO_TRACE_VTHRECEIVE</code>	<code>VTHRECEIVE</code>	VReceive	Enters the RECEIVE state
<code>_NTO_TRACE_VTHREPLY</code>	<code>VTHREPLY</code>	VReply	Enters the REPLY state
<code>_NTO_TRACE_VTHRUNNING</code>	<code>VTHRUNNING</code>	VRunning	Enters the RUNNING state
<code>_NTO_TRACE_VTHSEM</code>	<code>VTHSEM</code>	VSemaphore	Enters the SEM state
<code>_NTO_TRACE_VTHSEND</code>	<code>VTHSEND</code>	VSend	Enters the SEND state
<code>_NTO_TRACE_VTHSIGSUSPEND</code>	<code>VTHSIGSUSPEND</code>	VSigSuspend	Enters the SIGSUSPEND state
<code>_NTO_TRACE_VTHSIGWAITINFO</code>	<code>VTHSIGWAITINFO</code>	VSigWaitInfo	Enters the SIGWAITINFO state
<code>_NTO_TRACE_VTHSTACK</code>	<code>VTHSTACK</code>	VStack	Enters the STACK state
<code>_NTO_TRACE_VTHSTOPPED</code>	<code>VTHSTOPPED</code>	VStopped	Enters the STOPPED state

Event	traceprinter label	IDE label	Emitted when a virtual thread:
_NTO_TRACE_VTHWAITCTX	VTHWAITCTX	VWaitCtx	Enters the WAITCTX state
_NTO_TRACE_VTHWAITPAGE	VTHWAITPAGE	VWaitPage	Enters the WAITPAGE state
_NTO_TRACE_VTHWAITTHREAD	VTHWAITTHREAD	VWaitThread	Enters the WAITTHREAD state

Chapter 3

Kernel Buffer Management

As the instrumented kernel intercepts events, it stores them in a circular linked list of buffers.




Figure 4: The kernel buffers.

As each buffer fills, the instrumented kernel sends a signal to the data-capturing program that the buffer is ready to be read.

Each buffer is of a fixed size and is divided into a fixed number of slots:

- Event buffer slots per buffer: 1024
- Event buffer slot size: 16 bytes
- Buffer size: 16 KB

Some events are *single buffer slot events* (“simple events”) while others are *multiple buffer slot events* (“combine events”). In either case there is only *one* event, but the number of *event buffer slots* required to describe it may vary.

For details, see the [Interpreting Trace Data](#) (p. 61) chapter.

Linked list size

Although the size of the buffers is fixed, the maximum number of buffers used by a system is limited only by the amount of memory.

(The `tracelogger` utility uses a default setting of 32 buffers, or about 500 KB of memory.)

The buffers share kernel memory with the application(s), and the kernel automatically allocates memory at the request of the data-capture utility. The kernel allocates the buffers in contiguous physical memory space. If the data-capture program requests a larger block than is available contiguously, the instrumented kernel returns an error message.

For all intents and purposes, the number of events the instrumented kernel generates is infinite. Except for severe filtering or logging for only a few seconds, the instrumented kernel will probably exhaust the circular linked list of buffers, no matter how large it is. To allow the instrumented kernel to continue logging indefinitely, the data-capture program must continuously pipe (empty) the buffers.

Full buffers and the high-water mark

As each buffer becomes full, the instrumented kernel sends a signal to the data-capturing program to save the buffer. Because the buffer size is fixed, the kernel sends only the buffer address; the length is constant.

The instrumented kernel can't flush a buffer or change buffers within an interrupt. If the interrupt wasn't handled before the buffer became 100% full, some of the events may be lost. To ensure this never happens, the instrumented kernel requests a buffer flush at the high-water mark.

The high-water mark is set at an efficient, yet conservative, level of about 70%. Most interrupt routines require fewer than 300 event buffer slots (approximately 30% of 1024 event buffer slots), so there's virtually no chance that any events will be lost. (The few routines that use extremely long interrupts should include a manual buffer-flush request in their code.)

Therefore, in a normal system, the kernel logs about 715 events of the fixed maximum of 1024 events before notifying the capture program.

Buffer overruns

The instrumented kernel is both the very core of the system and the controller of the event buffers.

When the instrumented kernel is busy, it logs more events. The buffers fill more quickly, and the instrumented kernel requests that the buffers be flushed more often. The data-capture program handles each flush request; the instrumented kernel switches to the next buffer and continues logging events. In an extremely busy system, the data-capture program may not be able to flush the buffers as quickly as the instrumented kernel fills them.

In a three-buffer scenario, the instrumented kernel fills buffer 1 and signals the data-capture program that the buffer is full. The data-capture program takes “ownership” of buffer 1 and the instrumented kernel marks the buffer as “busy/in use.” If, say, the file is being saved to a hard drive that happens to be busy, then the instrumented kernel may fill buffer 2 and buffer 3 before the data-capture program can release buffer 1. In this case, the instrumented kernel skips buffer 1 and writes to buffer 2. The previous contents of buffer 2 are overwritten and the timestamps on the event buffer slots will show a discontinuity.

For more on buffer overruns, see the [Tutorials](#) (p. 69) chapter.

Chapter 4

Capturing Trace Data

The program that captures data is the “messenger” between the instrumented kernel and the filesystem.




Figure 5: Possible data capture configurations.

The main function of the data-capture program is to send the buffers given to it by the instrumented kernel to an output device (which may be a file or something else). In order to accomplish this function, the program must also:

- interface with the instrumented kernel
- specify data-filtering requirements the instrumented kernel will use

You must configure the instrumented kernel before logging. The instrumented kernel configuration settings include:

- buffer allocations (size)
- which events and classes of events to log (filtering)
- whether to log the events in wide mode or fast mode



The instrumented kernel retains the settings, and multiple programs access a single instrumented kernel configuration. Changing the settings in one process supersedes the settings made in another.

We've provided `tracelogger` as the default data-capture utility. Although you can write your own utility, there's little need to; if you do, the quickest approach is to tailor the `tracelogger` code to suit your own needs.

You can control the capture of data via `qconn` (under the control of the IDE), `tracelogger` (from the command line), or directly from your application. All three approaches use the `TraceEvent()` function to control the instrumented kernel:




Figure 6: Controlling the capture of trace data.

For information about controlling the trace from the IDE, see the Analyzing Your System with Kernel Tracing chapter of the IDE *User's Guide*.

Let's look first at using `tracelogger`, and then we'll describe how you can use `TraceEvent()` to control tracing from your application.



- Don't run more than one instance of `tracelogger` at a time. Similarly, don't run `tracelogger` and trace events under control of the IDE at the same time.
- On a multicore system, if the clocks on all processors aren't synchronized, then `tracelogger` will produce data with inconsistent timestamps, and the IDE won't be able to load the trace file. The IDE attempts to properly order the events in the trace file, and this can go awry if the timestamp data is incorrect.

The `traceprinter` utility doesn't have any issues with such traces because it doesn't attempt to reorder the data and interpret it; it simply dumps the contents of each event.

Using tracelogger to control tracing

The options that you use when you start `tracelogger` affect the way that the instrumented kernel logs events *and* how `tracelogger` captures them.

Managing trace buffers

You can use `tracelogger`'s command-line options to manage the instrumented kernel's buffers.

You can specify:

- the number of buffers
- whether or not to preserve the buffers in shared memory, to reuse later

You can also specify the number of buffers that `tracelogger` itself uses.

For more information, see the entry for `tracelogger` in the *Utilities Reference*.

`tracelogger`'s modes of operation

You can run `tracelogger` in several modes — depending on how and what you want to trace — by specifying the following command-line options:

`-n iterations`

In this mode, the kernel logs events, and `tracelogger` captures *iterations* buffers worth of data, and then terminates. This is the default mode, and default number of iterations is 32.

`-r`

Ring mode: the kernel stores all events in its circularly linked list of buffers without flushing them. The maximum time for which you can capture events (without overwriting earlier ones) is determined by the number of allocated buffers, as well as by the number of generated trace events.

In ring mode, `tracelogger` doesn't capture the events until it gets a `SIGINT` signal (e.g. you press **Ctrl-C**), or an application calls `TraceEvent()` with a command of `_NTO_TRACE_STOP`.

If you don't specify the `-r` option, `tracelogger` runs in linear mode; every filled-up buffer is captured and flushed immediately.

`-d1`

Daemon mode: the kernel doesn't log events, and `tracelogger` doesn't capture them until an application calls `TraceEvent()` with a command of

`_NTO_TRACE_START`. Logging continues until an application calls `TraceEvent()` with a command of `_NTO_TRACE_STOP`, or you terminate `tracelogger`.



In daemon mode, `tracelogger` ignores any other options, unless you also specify the `-E` option to use *extended daemon mode*.

-s seconds

The kernel logs events; `tracelogger` captures them over the specified time.

-c

The kernel logs events; `tracelogger` captures them, and continues to do so until you terminate it.

All of the above, except for daemon mode, constitute *normal mode*. In normal mode, you configure, start, and stop the tracing from the command line; in daemon (`-d1`) mode, your application must do everything from code. However, if you also use the `-E` option, you get the best of both modes: the command-line configuration of normal mode, and the full control of daemon mode.

Here's an outline of the strengths, weaknesses, and features of these modes:

Feature	Normal mode	Daemon mode	Extended daemon mode
tracelogger support	Full	Limited	Full
Controllability	Limited	Full	Full
Events recorded by default	All	None	All
Configuration difficulty	Easy	Harder	Easy
Configuration method	Command line only	User program, using calls to <code>TraceEvent()</code>	Command line and user program
Logging starts	Instantaneously	Through a user program; also calls to <code>TraceEvent()</code>	Through a user program; also calls to <code>TraceEvent()</code>

For a full description of the `tracelogger` utility and its options, see its entry in the *Utilities Reference*.

Choosing between wide and fast modes

By default, the instrumented kernel and `tracelogger` collect data in fast mode; to switch to wide mode, specify the `-w` option when you start `tracelogger`.

Filtering events

The `tracelogger` utility gives you some basic control over filtering by way of its `-F` option. This filtering is limited to excluding entire classes of events at a time; if you need a finer granularity, you'll need to use `TraceEvent()`, as described in the [Filtering](#) (p. 51) chapter in this guide.

By default, `tracelogger` captures all events from all classes, but you can disable the tracing of events from the classes as follows:

To disable this class:	Specify:
Kernel calls	-F1
Interrupt	-F2
Process	-F3
Thread	-F4
Virtual thread	-F5
Communication	-F6
System	-F7

You can specify more than one filter by using multiple `-F` options. Note that you can't disable the Control or User classes with this option. For more information about classes, see the [Events and the Kernel](#) (p. 19) chapter of this guide.

Specifying where to send the output

Because the circular linked list of buffers can't hope to store a complete log of event activity for any significant amount of time, the `tracebuffer` must be handed off to a data-capture program. Normally the data-capture program pipes the information to either an output device or a file.

By default, the `tracelogger` utility saves the output in the binary file `/dev/shmem/tracebuffer.kev`, but you can use the `-f` option to specify a different path. The `.kev` extension is short for “kernel events”; you can use a different extension, but the IDE recognizes `.kev` and automatically uses the System Profiler to open such files.

You can also map the file in shared memory (`-M`), but you must then also specify the maximum size for the file (`-S`).

Using *TraceEvent()* to control tracing

You don't have to use `tracelogger` to control all aspects of tracing; you can call *TraceEvent()* directly—which (after all) is what `tracelogger` does. Using *TraceEvent()* to control tracing means a bit more work for you, but you have much more control over specific details.

You could decide not to use `tracelogger` at all, and use *TraceEvent()* exclusively, but you'd then have to manage the buffers, collect the trace data, and save it in the appropriate form—a significant amount of work, although you can take advantage of the source code for `tracelogger` to help.

In practical terms you'll likely use `tracelogger` and *TraceEvent()* together. For example, you might run `tracelogger` in daemon mode, to take advantage of its management of the trace data, but call *TraceEvent()* to control exactly which events to trace.

The *TraceEvent()* kernel call takes a variable number of arguments. The first is always a command and determines what (if any) additional arguments are required.

For reference information about *TraceEvent()*, see the QNX Neutrino *C Library Reference*. The source code for `tracelogger` might also help you.

Managing trace buffers

As mentioned above, you can use *TraceEvent()* to manage the instrumented kernel's buffers, but it's probably easier to run `tracelogger` in daemon mode and let it look after the buffers. Nevertheless, here's a summary of how to do it with *TraceEvent()*:



In order to allocate or free the trace buffers, your application must have the `PROCmgr_AID_TRACE` ability enabled. For more information, see the entry *procmgr_ability()* in the QNX Neutrino *C Library Reference*.

- To allocate the buffers, use the `_NTO_TRACE_ALLOCBUFFER` command, specifying the number of buffers and a pointer to a location where *TraceEvent()* can store the physical address of the beginning of the circular linked list of allocated trace buffers:

```
TraceEvent(_NTO_TRACE_ALLOCBUFFER, uint bufnum, void** linkliststart);
```

Allocated trace buffers can store 1024 simple trace events.

- To free the buffers, use the `_NTO_TRACE_DEALLOCBUFFER` command. It doesn't take any additional arguments:

```
TraceEvent(_NTO_TRACE_DEALLOCBUFFER);
```

All events stored in the trace buffers are lost.

- To flush the buffer, regardless of the number of trace events it contains, use the `_NTO_TRACE_FLUSHBUFFER` command:

```
TraceEvent(_NTO_TRACE_FLUSHBUFFER);
```

- To get the number of simple trace events that are currently stored in the trace buffer, use the `_NTO_TRACE_QUERYEVENTS` command:

```
num_events = TraceEvent(_NTO_TRACE_QUERYEVENTS);
```

Modes of operation

TraceEvent() doesn't support the different modes of operation that `tracelogger` does; your application has to indicate when to start tracing, how long to trace for, and so on:

- To start tracing, use the `_NTO_TRACE_START` or `_NTO_TRACE_STARTNOSTATE` command:

```
TraceEvent(_NTO_TRACE_START);
TraceEvent(_NTO_TRACE_STARTNOSTATE);
```

These commands are similar, except that `_NTO_TRACE_STARTNOSTATE` suppresses the initial system state information (which includes thread IDs and the names of processes).

- To stop tracing, use the `_NTO_TRACE_STOP` command:

```
TraceEvent(_NTO_TRACE_STOP);
```

You can decide whether to trace until you've gathered a certain quantity of data, trace for a certain length of time, or trace only during an operation that's of particular interest to you. After stopping the trace, you should flush the buffer by calling:

```
TraceEvent(_NTO_TRACE_FLUSHBUFFER);
```

- To use ring mode, use the `_NTO_TRACE_SETRINGMODE` command:

```
TraceEvent(_NTO_TRACE_SETRINGMODE);
```

As described earlier in this chapter, in ring mode the kernel stores all events in a circular fashion inside the linked list without flushing them.

- To use linear mode (the default), use the `_NTO_TRACE_SETLINEARMODE` command:

```
TraceEvent(_NTO_TRACE_SETLINEARMODE);
```

When you use this mode, every filled-up buffer is captured and flushed immediately.

Filtering events

You can select events in an additive or subtractive manner; you can start with no events, and then add specific classes or events, or you can start with all events, and then exclude specific ones. We'll discuss using *TraceEvent()* to filter events in the [Filtering](#) (p. 51) chapter.

Choosing between wide and fast modes

TraceEvent() gives you much finer control over wide and fast mode than you can get with *tracelogger*, which can simply set the mode for *all* events in *all* traced classes. Using *TraceEvent()*, you can set fast and wide mode for all classes, a specific class, or a specific event in a class:

- To set the mode for all classes, use the `_NTO_TRACE_SETALLCLASSESWIDE` or `_NTO_TRACE_SETALLCLASSESFAST` command. These commands don't require any additional arguments:

```
TraceEvent( _NTO_TRACE_SETALLCLASSESWIDE );
TraceEvent( _NTO_TRACE_SETALLCLASSESFAST );
```

- To set the mode for all events in a class, use the `_NTO_TRACE_SETCLASSFAST` or `_NTO_TRACE_SETCLASSWIDE` command. These commands require a class as an additional argument:

```
TraceEvent(_NTO_TRACE_SETCLASSFAST, int class);
TraceEvent(_NTO_TRACE_SETCLASSWIDE, int class);
```

For example:

```
TraceEvent(_NTO_TRACE_SETCLASSWIDE, _NTO_TRACE_KERCALLENTER);
```

- To set the mode for a specific event in a class, use the `_NTO_TRACE_SETEVENTFAST` or `_NTO_TRACE_SETEVENTWIDE` command, specifying the class, followed by the event:

```
TraceEvent(_NTO_TRACE_SETEVENTFAST, int class, int event)
TraceEvent(_NTO_TRACE_SETEVENTWIDE, int class, int event)
```

For example:

```
TraceEvent(_NTO_TRACE_SETEVENTFAST, _NTO_TRACE_KERCALLENTER,
           __KER_INTERRUPT_ATTACH);
```

Inserting trace events

You can even use *TraceEvent()* to insert your own events into the trace data. You can call *TraceEvent()* directly (see below), but it's much easier to use the following convenience functions:

trace_func_enter()

Insert a trace event for the entry to a function

trace_func_exit()

Insert a trace event for the exit from a function

trace_here()

Insert a trace event for the current address

trace_logb()

Insert a user combine trace event

trace_logbc()

Insert a trace event of an arbitrary class and type with arbitrary data

trace_logf()

Insert a user string trace event

trace_logi()

Insert a user simple trace event

trace_nlogf()

Insert a user string trace event, specifying a maximum string length

trace_vnlogf()

Insert a user string trace event, using a variable argument list

If you want to call *TraceEvent()* directly, use one of the following commands:

- `_NTO_TRACE_INSERTCUSEREVENT`
- `_NTO_TRACE_INSERTEVENT`
- `_NTO_TRACE_INSERTSUSEREVENT`
- `_NTO_TRACE_INSERTUSRSTREVENT`

For more information, see the entry for *TraceEvent()* in the QNX Neutrino *C Library Reference*.

Chapter 5

Filtering

Gathering many events generates a lot of data, which requires *memory* and *processor time*. It also makes the task of interpreting the data more difficult.

Because the amount of data that the instrumented kernel generates can be overwhelming, the SAT supports several types of *filters* that you can use to reduce the amount of data to be processed:

Static rules filter

A simple filter that chooses events based on their type, class, or other simple criteria.

Dynamic rules filter

A more complex filter that lets you register a callback function that can decide — based on the state of your application or system, or on whatever criteria you choose — whether or not to log a given event.

Post-processing filter

A filter that you run after capturing event data. Like the dynamic rules filter, this can be as complex and sophisticated as you wish.

The static and dynamic rules filters affect the amount of data being logged into the kernel buffers; filtered data is discarded — you save processing time and memory, but there's a chance that some of the filtered data could have been useful.

In contrast, the post-processing facility doesn't discard data; it simply doesn't use it — if you've saved the data, you can use it later.




Figure 7: Overall view of the SAT and its filters.

Most of the events don't indicate what caused the event to occur. For example, an event for entering *MsgSendv()* doesn't indicate which thread in which process called it; you have to infer it during interpretation from a previous thread-running event. You have carefully choose what you filter to avoid losing this context.

The static rules filter

You can use the static rules filter to track or filter events for all classes, certain events in a class, or even events related to specific process and thread IDs. You can select events in an additive or subtractive manner; you can start with no events, and then add specific classes or events, or you can start with all events, and then exclude specific ones.

The static rules filter is the best, most efficient method of data reduction. It generally frees up the processor while significantly reducing the data rate. This filter is also useful for gathering large amounts of data periodically, or after many hours of logging without generating gigabytes of data in the interim.

You set up this filter using the following *TraceEvent()* commands:

`_NTO_TRACE_ADDALLCLASSES, _NTO_TRACE_DELALLCLASSES`

Emit or suppress tracing for *all* classes and events:

```
TraceEvent(_NTO_TRACE_ADDALLCLASSES);
TraceEvent(_NTO_TRACE_DELALLCLASSES);
```



The `_NTO_TRACE_DELALLCLASSES` command doesn't suppress the process- and thread-specific tracing that the `_NTO_TRACE_SETCLASSPID`, `_NTO_TRACE_SETCLASSTID`, `_NTO_TRACE_SETEVENTPID`, and `_NTO_TRACE_SETEVENTTID` commands set up. You need to clear their tracing separately, as shown below.

`_NTO_TRACE_ADDCLASS, _NTO_TRACE_DELCLASS`

Emit or suppress all trace events from a specific class:

```
TraceEvent(_NTO_TRACE_ADDCLASS, class);
TraceEvent(_NTO_TRACE_DELCLASS, class);
```

For information about the different classes, see “[Classes and events](#) (p. 24)” in the Events and the Kernel chapter of this guide.

`_NTO_TRACE_ADDEVENT, _NTO_TRACE_DELEVENT`

Emit or suppress a specific event in a specific class:

```
TraceEvent(_NTO_TRACE_ADDEVENT, class, event);
TraceEvent(_NTO_TRACE_DELEVENT, class, event);
```

`_NTO_TRACE_SETCLASSPID, _NTO_TRACE_CLRCLASSPID`

Emit or suppress *all* events from a specified process ID:

```
TraceEvent(_NTO_TRACE_SETCLASSPID, int class, pid_t pid);
TraceEvent(_NTO_TRACE_CLRCLASSPID, int class);
```

`_NTO_TRACE_SETCLASSTID, _NTO_TRACE_CLRCLASSTID`

Emit or suppress *all* events from the specified process and thread IDs:

```
TraceEvent(_NTO_TRACE_SETCLASSTID, int class, pid_t pid, tid_t tid);
TraceEvent(_NTO_TRACE_CLRCLASSTID, int class);
```

`_NTO_TRACE_SETEVENTPID, _NTO_TRACE_CLREVENTPID`

Emit or suppress a specific event for a specified process ID:

```
TraceEvent(_NTO_TRACE_SETEVENTPID, int class, int event, pid_t pid);
TraceEvent(_NTO_TRACE_CLREVENTPID, int class, int event);
```

`_NTO_TRACE_SETEVENTTID, _NTO_TRACE_CLREVENTTID`

Emit or suppress a specific event for the specified process and thread IDs:

```
TraceEvent(_NTO_TRACE_SETEVENTTID, int class, int event,
           pid_t pid, tid_t tid);
TraceEvent(_NTO_TRACE_CLREVENTTID, int class, int event);
```

The `_NTO_TRACE_SETCLASSPID`, `_NTO_TRACE_SETCLASSTID`, `_NTO_TRACE_SETEVENTPID`, and `_NTO_TRACE_SETEVENTTID` commands apply only to these classes:



- `_NTO_TRACE_COMM`
- `_NTO_TRACE_KERCALL`, `_NTO_TRACE_KERCALLEENTER`,
`_NTO_TRACE_KERCALLEEXIT`
- `_NTO_TRACE_SYSTEM`
- `_NTO_TRACE_THREAD`
- `_NTO_TRACE_VTHREAD`

The instrumented kernel retains these settings, so you should be careful not to make any assumptions about the settings that are in effect when you set up your filters. For example, you might want to start by turning off all filtering:

```
TraceEvent(_NTO_TRACE_DELALLCLASSES);
TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_KERCALL);
TraceEvent(_NTO_TRACE_CLRCLASSTID, _NTO_TRACE_KERCALL);
TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_THREAD);
TraceEvent(_NTO_TRACE_CLRCLASSTID, _NTO_TRACE_THREAD);
TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_VTHREAD);
TraceEvent(_NTO_TRACE_CLRCLASSTID, _NTO_TRACE_VTHREAD);
TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_SYSTEM);
TraceEvent(_NTO_TRACE_CLRCLASSTID, _NTO_TRACE_SYSTEM);
```

```
TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_COMM);  
TraceEvent(_NTO_TRACE_CLRCLASSTID, _NTO_TRACE_COMM);
```

You can select events in an additive or subtractive manner; you can start with no events, and then add specific classes or events, or you can start with all events, and then exclude specific ones.

For an example using the static filter, see the [five_events.c](#) (p. 76) example in the [Tutorials](#) (p. 69) chapter.

The dynamic rules filter

The dynamic rules filter can do all the filtering that the static filter does—and more—but it isn't as quick. This filter lets you register functions (event handlers) that decide whether or not to log a given event.



If you want to use a dynamic rules filter, be sure that you've also set up a static rules filter that logs the events you want to examine. For example, if you want to dynamically examine events in the `_NTO_TRACE_THREAD` class, also call:

```
TraceEvent(_NTO_TRACE_ADDCLASS, _NTO_TRACE_THREAD);
```

For an example of using the dynamic rules filter, see the [eh_simple.c](#) (p. 83) example in the [Tutorials](#) (p. 69) chapter.

Setting up a dynamic rules filter

Before you set up dynamic filtering, you must:

- have the `PROCMGR_AID_TRACE` and `PROCMGR_AID_IO` abilities enabled. For more information, see the entry for *procmgr_ability()* in the QNX Neutrino *C Library Reference*.
- request I/O privileges by calling *ThreadCtl()* with the `_NTO_TCTL_IO` flag:

```
if (ThreadCtl(_NTO_TCTL_IO, 0) != EOK) {  
    fprintf(stderr, "argv[0]: Failed to obtain I/O privileges\n");  
    return (-1);  
}
```

Then call *TraceEvent()* with one of these commands:

`_NTO_TRACE_ADDCLASSEVHANDLER`

Register a function to call whenever an event for the given class is emitted:

```
TraceEvent(_NTO_TRACE_ADDCLASSEVHANDLER, class,  
           int (*event_hdlr)(event_data_t*),  
           event_data_t* data_struct);
```

`_NTO_TRACE_ADDEVENTHANDLER`

Register a function to call whenever an event for the given class and event type is emitted:

```
TraceEvent(_NTO_TRACE_ADDEVENTHANDLER, class, event,  
           int (*event_hdlr)(event_data_t*),  
           event_data_t* data_struct);
```

The additional arguments are:

event_hdlr

A pointer to the function that you want to register. The prototype for the function is:

```
int event_hdlr (event_data_t *event_data);
```

data_struct

A pointer to a locally defined data structure, of type `event_data_t`, where the kernel can store event data to pass to the event handler (see below).

Event handler

The dynamic filter is an event handler that works like an interrupt handler. When this filter is used, a section of your custom code is executed. The code can test for a set of conditions before determining whether the event should be stored.



The only library functions that you can call in your event handler are those that are safe to call from an interrupt handler. For a list of these functions, see the Full Safety Information appendix in the *QNX Neutrino C Library Reference*. If you call an unsafe function — such as `printf()` — in your event handler, you'll crash your entire system.

If you want to log the current event, return a non-zero value; to discard the event, return 0. Here's a very simple event handler that says to log all of the given events:

```
int event_handler(event_data_t* dummy_pt)
{
    return(1);
}
```



If you use both types of dynamic filters (event handler and class event handler), and they both apply to a particular event, the event is logged if *both* event handlers return a non-zero value.

In addition to deciding whether or not the event should be logged, you can use the dynamic rules filter to output events to external hardware or to perform other tasks—it's up to you because it's your code. Naturally, you should write the code as efficiently as possible in order to minimize the overhead.

You can access the information about the intercepted event within the event handler by examining the `event_data_t` structure passed as an argument to the event

handler. The layout of the `event_data_t` structure (declared in `<sys/trace.h>`) is as follows:

```
/* event data filled by an event handler */
typedef struct
{
    __traceentry header; /* same as traceevent header */
    uint32_t* data_array; /* initialized by the user */
    uint32_t ei_num; /* number of elements returned */
    void* area; /* user data */
    uint32_t feature_mask; /* bits indicate valid features */
    uint32_t feature[_NTO_TRACE_FI_NUM]; /* feature array
                                         - additional data */
} event_data_t;
```

The `event_data_t` structure includes a pointer to an array for the data arguments of the event. You *must* provide an array, and it must be large enough to hold the data for the event or events that you're handling (see the [Current Trace Events and Data](#) appendix). For example:



```
event_data_t e_d_1;
uint32_t data_array_1[20]; /* 20 elements for potential args. */
e_d_1.data_array = data_array_1;
```

If you don't provide the data array, or it isn't big enough, your data segment could become corrupted.

You can use the following macros, defined in `<sys/trace.h>`, to work with the header of an event:

`_NTO_TRACE_GETEVENT_C(c)`

Get the class.

`_NTO_TRACE_GETEVENT(c)`

Get the type of event.

`_NTO_TRACE_GETCPU(h)`

Get the number of the CPU that the event occurred on.

`_NTO_TRACE_SETEVENT_C(c,cl)`

Set the class in the header `c` to be `cl`.

`_NTO_TRACE_SETEVENT(c, e)`

Set the event type in the header `c` to be `e`.

The bits of the `feature_mask` member are related to any additional features (arguments) that you can access inside the event handler. All standard data arguments — the ones that correspond to the data arguments of the trace event — are delivered without changes within the `data_array`.

There are two constants associated with each additional feature:

- `_NTO_TRACE_FM***` — feature parameter masks
- `_NTO_TRACE_FI***` — feature index parameters

The currently defined features are:

Feature	Parameter mask	Index
Process ID	<code>_NTO_TRACE_FMPID</code>	<code>_NTO_TRACE_FIPID</code>
Thread ID	<code>_NTO_TRACE_FMTID</code>	<code>_NTO_TRACE_FITID</code>

If any particular bit of the *feature_mask* is set to 1, then you can access the feature corresponding to this bit within the *feature* array. Otherwise, you must not access the feature. For example, if the expression:

```
feature_mask & _NTO_TRACE_FMPID
```

is TRUE, then you can access the additional feature corresponding to identifier `_NTO_TRACE_FMPID` as:

```
my_pid = feature[_NTO_TRACE_FIPID];
```

Removing event handlers

To remove event handlers, call *TraceEvent()* with these commands:

`_NTO_TRACE_DELCLASSEVHANDLER`

Remove the function for the given class and event type:

```
TraceEvent(_NTO_TRACE_DELCLASSEVHANDLER, class);
```

`_NTO_TRACE_DELEVENTHANDLER`

Remove the function for the given class and event type:

```
TraceEvent(_NTO_TRACE_DELEVENTHANDLER, class, event);
```

The post-processing facility

The post-processing facility is different from the other filters in that it reacts to the events without permanently discarding them (or having to choose not to). Because the processing is done on the captured data, often saved as a file, you could make multiple passes on the same data without changing it—one pass could count the number of thread state changes, another pass could display all the kernel events.

The post-processing facility is really a collection of callback functions that decide what to do for each event. One example of post-processing is the `traceprinter` utility itself. It prints all the events instead of filtering them, but the principles are the same.

We'll look at `traceprinter` in more detail in the [Interpreting Trace Data](#) (p. 61) chapter.

Chapter 6

Interpreting Trace Data

Once the data has been captured, you may process it, either in real time or offline.




Figure 8: Possible data interpretation configurations.

The best tool (by far) for interpreting the copious amounts of trace data is the Integrated Development Environment. It provides a sophisticated and versatile user interface that lets you filter and examine the data.




Figure 9: Examining trace data in the IDE's System Profiler perspective.

For more information, see the Analyzing Your System with Kernel Tracing chapter of the IDE *User's Guide*.

We also provide a `traceprinter` utility that simply prints a plain-text version of the trace data, sending its output to `stdout` or to a file.

You can also build your own, custom interpreter, using the `traceparser` library.

Using traceprinter and interpreting the output

The simplest way to turn the tracing data into a form that you can analyze is to pass the .kev file through traceprinter. For details, see its entry in the *Utilities Reference*.

Let's take a look at an example of the output from traceprinter. This is the output from “[Gathering all events from all classes](#) (p. 71)” in the Tutorials chapter.

The output starts with some information about how you ran the trace:

```
TRACEPRINTER version 1.02
TRACEPARSER LIBRARY version 1.02
-- HEADER FILE INFORMATION --
  TRACE_FILE_NAME:: all_classes.kev
  TRACE_DATE:: Wed Jun 24 10:52:58 2009
  TRACE_VER_MAJOR:: 1
  TRACE_VER_MINOR:: 01
  TRACE_LITTLE_ENDIAN:: TRUE
  TRACE_ENCODING:: 16 byte events
  TRACE_BOOT_DATE:: Tue Jun 23 11:47:46 2009
  TRACE_CYCLES_PER_SEC:: 736629000
  TRACE_CPU_NUM:: 1
  TRACE_SYSNAME:: QNX
  TRACE_NODENAME:: localhost
  TRACE_SYS_RELEASE:: 6.4.1
  TRACE_SYS_VERSION:: 2009/05/20-17:35:56EDT
  TRACE_MACHINE:: x86pc
  TRACE_SYSPAGE_LEN:: 2264
TRACE_TRACELOGGER_ARGS:: tracelogger -d1 -n 3 -f all_classes.kev
```

The next section includes information about all the processes in existence when the trace started:

```
-- KERNEL EVENTS --
t:0x4f81e320 CPU:00 CONTROL: BUFFER sequence = 33, num_events = 714
t:0x4f81e320 CPU:00 CONTROL :TIME msb:0x000037b0 lsb(offset):0x4f81e014
t:0x4f82017a CPU:00 PROCESS :PROCCREATE_NAME
  ppid:0
  pid:1
  name:proc/boot/procnto-smp-instr
t:0x4f820f9a CPU:00 THREAD :THCREATE pid:1 tid:1
t:0x4f821358 CPU:00 THREAD :THREADY pid:1 tid:1
t:0x4f821698 CPU:00 THREAD :THCREATE pid:1 tid:2
t:0x4f821787 CPU:00 THREAD :THRECEIVE pid:1 tid:2
t:0x4f8219ca CPU:00 THREAD :THCREATE pid:1 tid:3
t:0x4f821ac6 CPU:00 THREAD :THRECEIVE pid:1 tid:3
t:0x4f821c94 CPU:00 THREAD :THCREATE pid:1 tid:4
t:0x4f821d90 CPU:00 THREAD :THRECEIVE pid:1 tid:4
t:0x4f821f6c CPU:00 THREAD :THCREATE pid:1 tid:5
t:0x4f82205b CPU:00 THREAD :THRECEIVE pid:1 tid:5
t:0x4f8222aa CPU:00 THREAD :THCREATE pid:1 tid:7
t:0x4f822399 CPU:00 THREAD :THRECEIVE pid:1 tid:7
t:0x4f8225bd CPU:00 THREAD :THCREATE pid:1 tid:8
t:0x4f8226ac CPU:00 THREAD :THRECEIVE pid:1 tid:8
t:0x4f8228ca CPU:00 THREAD :THCREATE pid:1 tid:10
t:0x4f8229b9 CPU:00 THREAD :THRECEIVE pid:1 tid:10
t:0x4f822b7d CPU:00 THREAD :THCREATE pid:1 tid:11
t:0x4f822c6c CPU:00 THREAD :THRECEIVE pid:1 tid:11
t:0x4f822dd7 CPU:00 THREAD :THCREATE pid:1 tid:12
t:0x4f822ec6 CPU:00 THREAD :THRECEIVE pid:1 tid:12
t:0x4f8230ac CPU:00 THREAD :THCREATE pid:1 tid:15
t:0x4f82319b CPU:00 THREAD :THRECEIVE pid:1 tid:15
t:0x4f8233ca CPU:00 THREAD :THCREATE pid:1 tid:20
t:0x4f8234b9 CPU:00 THREAD :THRECEIVE pid:1 tid:20
t:0x4f823ad0 CPU:00 PROCESS :PROCCREATE_NAME
  ppid:1
  pid:2
  name:sbin/tinit
t:0x4f823f38 CPU:00 THREAD :THCREATE pid:2 tid:1
t:0x4f82402e CPU:00 THREAD :THREPLY pid:2 tid:1
t:0x4f82447d CPU:00 PROCESS :PROCCREATE_NAME
  ppid:2
  pid:4099
  name:proc/boot/pci-bios
t:0x4f824957 CPU:00 THREAD :THCREATE pid:4099 tid:1
t:0x4f824a4d CPU:00 THREAD :THRECEIVE pid:4099 tid:1
t:0x4f824ff8 CPU:00 PROCESS :PROCCREATE_NAME
  ppid:2
  pid:4100
  name:proc/boot/slogger
```

You can suppress this initial information by passing the `_NTO_TRACE_STARTNOSTATE` command to `TraceEvent()`, but you'll likely need the information (including process IDs and thread IDs) to make sense out of the actual trace data.

The sample above shows the creation and naming of the instrumented kernel `procnto-smp-instr` (process ID 1) and its threads (thread ID 1 is the idle thread), followed by `tinit` (process ID 2), `pci-bios`, and `slogger`. Some of these are the processes that were launched when you booted your system.

This continues for a while, culminating in the creation of the `tracelogger` process and our own program, `all_classes` (process ID 1511472):

```
t:0x4f852aa8 CPU:00 PROCESS :PROCCREATE_NAME
                        ppid:426015
                        pid:1507375
                        name:usr/sbin/tracelogger
t:0x4f853360 CPU:00 THREAD  :THCREATE      pid:1507375 tid:1
t:0x4f853579 CPU:00 THREAD  :THRECEIVE     pid:1507375 tid:1
t:0x4f85392a CPU:00 THREAD  :THCREATE      pid:1507375 tid:2
t:0x4f853a19 CPU:00 THREAD  :THSIGWAITINFO pid:1507375 tid:2
t:0x4f853d96 CPU:00 PROCESS :PROCCREATE_NAME
                        ppid:426022
                        pid:1511472
                        name:./all_classes
t:0x4f854048 CPU:00 THREAD  :THCREATE      pid:1511472 tid:1
t:0x4f854140 CPU:00 THREAD  :THRUNNING     pid:1511472 tid:1
```

Next is the exit from our program's call to `TraceEvent()`:

```
t:0x4f854910 CPU:00 KER_EXIT:TRACE_EVENT/01 ret_val:0x00000000 empty:0x00000000
```

Why doesn't the trace doesn't include the *entry* to `TraceEvent()`? Well, `tracelogger` didn't log anything until our program *told* it to — by calling `TraceEvent()`!

So far, so good, but now things get more complicated:

```
t:0x4f856aac CPU:00 KER_CALL:THREAD_DESTROY/47 tid:-1 status_p:0
t:0x4f857dca CPU:00 KER_EXIT:THREAD_DESTROY/47 ret_val:0x00000030 empty:0x00000000
t:0x4f8588d3 CPU:00 KER_CALL:THREAD_DESTROYALL/48 empty:0x00000000 empty:0x00000000
t:0x4f858ed7 CPU:00 THREAD  :THDESTROY     pid:1511472 tid:1
t:0x4f8598b9 CPU:00 THREAD  :THDEAD        pid:1511472 tid:1
t:0x4f859c4c CPU:00 THREAD  :THRUNNING     pid:1 tid:1
```

You can see that a thread is being destroyed, but which one? The *tid* of -1 refers to the current thread, but which process does it belong to? As mentioned earlier, most of the events don't indicate what caused the event to occur; you have to infer from a previous thread-running event. In this case, it's our own program (process ID 1511472) that's ending; it starts the tracing, and then exits. Thread 1 of `procnto-smp-instr` (the idle thread) runs.

The trace continues like this:

```
t:0x4f85c6e3 CPU:00 COMM      :SND_PULSE_EXE scoid:0x40000002 pid:1
t:0x4f85cecd CPU:00 THREAD  :THRUNNING     pid:1 tid:12
t:0x4f85d5ad CPU:00 THREAD  :THREADY       pid:1 tid:1
t:0x4f85e5b3 CPU:00 COMM      :REC_PULSE    scoid:0x40000002 pid:1
t:0x4f860ee2 CPU:00 KER_CALL:THREAD_CREATE/46 func_p:f0023170 arg_p:eff6e000
t:0x4f8624c7 CPU:00 THREAD  :THCREATE      pid:1511472 tid:1
t:0x4f8625ff CPU:00 THREAD  :THWAITTHREAD pid:1 tid:12
t:0x4f8627b4 CPU:00 THREAD  :THRUNNING     pid:1511472 tid:1
t:0x4f8636fd CPU:00 THREAD  :THREADY       pid:1 tid:12
t:0x4f865c34 CPU:00 KER_CALL:CONNECT_SERVER_INFO/41 pid:0 coid:0x00000000
t:0x4f866836 CPU:00 KER_EXIT:CONNECT_SERVER_INFO/41 coid:0x00000000 info->nd:0
t:0x4f86735e CPU:00 KER_CALL:TIMER_TIMEOUT/75 timeout_flags:0x00000050 ntime(sec):30
t:0x4f868445 CPU:00 KER_EXIT:TIMER_TIMEOUT/75 prev_timeout_flags:0x00000000 otime(sec):0
t:0x4f8697d3 CPU:00 INT_ENTR:0x00000000 (0) IP:0xf008433e
t:0x4f86a276 CPU:00 INT_HANDLER_ENTR:0x00000000 (0) PID:126997 IP:0x080b7334 AREA:0x0812a060
t:0x4f86afa7 CPU:00 INT_HANDLER_EXIT:0x00000000 (0) SIGEVENT:NONE
t:0x4f86b304 CPU:00 INT_HANDLER_ENTR:0x00000000 (0) PID:1 IP:0xf0056570 AREA:0x00000000
t:0x4f86ca12 CPU:00 INT_HANDLER_EXIT:0x00000000 (0) SIGEVENT:NONE
t:0x4f86cff6 CPU:00 INT_EXIT:0x00000000 (0) inkernel:0x00000f01
t:0x4f86e276 CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000000 msg:"" (0x0040116)
t:0x4f86e756 CPU:00 COMM      :SND_MESSAGE rcvid:0x0000004f pid:159762
t:0x4f86f84a CPU:00 THREAD  :THREPLY       pid:1511472 tid:1
t:0x4f8705dd CPU:00 THREAD  :THREADY       pid:159762 tid:1
t:0x4f8707d4 CPU:00 THREAD  :THRUNNING     pid:159762 tid:1
```

```

t:0x4f870bff CPU:00 COMM      :REC_MESSAGE   rcvid:0x0000004f pid:159762
t:0x4f878b6c CPU:00 KER_CALL:MSG_REPLYV/15 rcvid:0x0000004f status:0x00000000
t:0x4f878f4b CPU:00 COMM      :REPLY_MESSAGE tid:1 pid:1511472
t:0x4f8798d2 CPU:00 THREAD    :THREADY       pid:1511472 tid:1

```

The `SND_PULSE_EXE` event indicates that a `SIGEV_PULSE` was sent to the server connection ID `0x40000002` of `procnto-smp-instr`, but what is it, and who sent it? Thread 12 of the kernel receives it, and then surprisingly creates a new thread 1 in our process (ID 1511472), and starts chatting with it. What we're seeing here is the teardown of our process. It delivers a death pulse to the kernel, and then one of the kernel's threads receives the pulse and creates a thread in the process to clean up.

In the midst of this teardown, an interrupt occurs, its handler runs, and a message is sent to the process with ID 159762. By looking at the initial system information, we can determine that process ID 159762 is `devc-pty`.

Farther down in the trace is the actual death of our `all_classes` process:

```

t:0x4f8faa68 CPU:00 THREAD    :THRUNNING      pid:1 tid:20
t:0x4f8fb09f CPU:00 COMM      :REC_PULSE      scoid:0x40000002 pid:1
t:0x4f8ff1a5 CPU:00 PROCESS   :PROCDESTROY    ppid:426022 pid:1511472

```

As you can tell from a very short look at this trace, wading through a trace can be time-consuming, but can give you a great understanding of what exactly is happening in your system.

You can simplify your task by terminating any processes that you don't want to include in the trace, or by filtering the trace data.

Building your own parser

If you want to create your own parser, consider the structure of `traceprinter` as a starting point. This utility consists of a long list of callback definitions, followed by a fairly simple parsing procedure. Each of the callback definitions is for printing.

The following sections give a brief introduction to the building blocks to the parser, and some of the issues you'll need to handle.

The `traceparser` library

The `traceparser` library provides a front end to facilitate the handling and parsing of events received from the instrumented kernel and the data-capture utility.

The library serves as a thin middle layer to:

- assemble multiple buffer slots into a single event
- perform data parsing to execute user-defined callbacks triggered by certain events

You typically use the `traceparser` functions as follows:

1. Initialize the `traceparser` library by calling `traceparser_init()`. You can also use this function to get the state of your parser.
2. Set the `traceparser` debug mode and specify a `FILE` stream for the debugging output by calling `traceparser_debug()`.
3. Set up callbacks for processing the trace events that you're interested in:

traceparser_cs()

Attach a callback to an event

traceparser_cs_range()

Attach a callback to a range of events

When you set up a callback with either of these functions, you can provide a pointer to arbitrary user data to be passed to the callback.

4. Start parsing your trace data by calling `traceparser()`
5. Destroy your parser by calling `traceparser_destroy()`

You can get information about your parser at any time by calling `traceparser_get_info()`.

For more information about these functions, see their entries in the QNX Neutrino *C Library Reference*.

Simple and combine events

A simple event is an event that can be described in a *single event buffer slot*; a combine event is an event that is larger and can be fully described only in *multiple event buffer slots*. Both simple and combine events consist of only *one* kernel event.

Each event buffer slot is an opaque `traceevent_t` structure.

The `traceevent_t` structure



The `traceevent_t` structure is *opaque*—although some details are provided, the structure shouldn't be accessed without the `libtraceparser` API.

The `traceevent_t` structure is only 16 bytes long, and only half of that describes the event. This small size reduces instrumentation overhead and improves granularity. Where the information required to represent the event won't fit into a single `traceevent_t` structure, it spans as many `traceevent_t` structures as required, resulting in a *combine event*. A combine event isn't actually several events combined, but rather a single, long event requiring a combination of `traceevent_t` elements to represent it.

In order to distinguish regular events from combine events, the `traceevent_t` structure includes a 2-bit flag that indicates whether the event is a single event or whether it's the first, middle, or last `traceevent_t` structure of the event. The flag is also used as a rudimentary integrity check. The timestamp element of the combine event is identical in each buffer slot; no other event will have the same timestamp.

Adding this “thin” protocol doesn't burden the instrumented kernel and keeps the `traceevent_t` structure small. The trade-off is that it may take many `traceevent_t` structures to represent a single kernel event.

Event interlacing

Although events are timestamped immediately, they may not be written to the buffer in one single operation (atomically). When *multiple buffer slot events* (“combine events”) are being written to the buffer, the process is frequently interrupted in order to allow other threads and processes to run. Events triggered by higher-priority threads are often written to the buffer first. Thus, events may be *interlaced*. Although events may not be contiguous, they are *not* scrambled (unless there's a buffer overrun.) The sequential order of the combine event is always correct, even if it's interrupted with a different event.

In order to maintain speed during runtime, the instrumented kernel writes events unsorted as quickly as possible; reassembling the combine events must be done in post-processing. The `libtraceparser` API transparently reassembles the events.

Timestamps

The timestamp is the 32 Least Significant Bits (LSB) part of the 64-bit clock. Whenever the 32-bit portion of the clock rolls over, a `_NTO_TRACE_CONTROLTIME` control event is issued. Although adjacent events will never have the same exact timestamp, there may be some timestamp duplication due to the clock's rolling over.

The rollover control event includes the 32 Most Significant Bits (MSB), so you can reassemble the full clock time, if required. The timestamp includes only the LSB in order to reduce the amount of data being generated. (A 1-GHz clock rolls over every 4.29 seconds—an eternity compared to the number of events generated.)



Although the majority of events are stored chronologically, you shouldn't write code that depends on events being retrieved chronologically. Some *multiple buffer slot events* (combine events) may be interlaced with others with *leading* timestamps. In the case of buffer overruns, the timestamps will definitely be scrambled, with entire blocks of events out of chronological order. Spurious gaps, while theoretically possible, are very unlikely.

Chapter 7

Tutorials

This chapter leads you through some tutorials to help you learn how to use `TraceEvent()` to control event tracing.

These tutorials all follow the same general procedure:

1. Compile the specified C program into a file of the same name, without the `.c` extension.
2. Run the specified `tracelogger` command. Because we're running `tracelogger` in daemon mode, it doesn't start logging events until our program tells it to. This means that you don't have to rush to start your C program; the tracing waits for you.

The default number of buffers is 32, which produces a rather large file, so we'll use the `-n` option to limit the number of buffers to a reasonable number. Feel free to use the default, but expect a large file.

3. In a *separate* terminal window, run the compiled C program. Some examples use options.
4. Watch the first terminal window; a few seconds after you start your C program, `tracelogger` will finish running.
5. If you run the program, it generates its own sample result file. The “tracebuffer” files are binary files that can be interpreted only with the `libtraceparser` library, which the `traceprinter` utility uses.

If you don't want to run the program, take a look at our `traceprinter` output. (Note that different versions and systems will create slightly different results.)



You may include these samples in your code as long as you comply with the license agreement.

The instrex.h header file

To reduce repetition and keep the programs simple, we've put some functionality into a header file, `instrex.h`:

```
/*
 * $QNXLicenseC:
 * Copyright 2007, QNX Software Systems. All Rights Reserved.
 *
 * You must obtain a written license from and pay applicable license fees to QNX
 * Software Systems before you may reproduce, modify or distribute this software,
 * or any work that includes all or part of this software. Free development
 * licenses are available for evaluation and non-commercial purposes. For more
 * information visit http://licensing.qnx.com or email licensing@qnx.com.
 *
 * This file may contain contributions from others. Please review this entire
 * file for other proprietary rights or license notices, as well as the QNX
 * Development Suite License Guide at http://licensing.qnx.com/license-guide/
 * for other information.
 * $
 */

/*
 * instrex.h instrumentation examples - public definitions
 */

#ifdef __INSTREX_H_INCLUDED

#include <errno.h>
#include <stdio.h>
#include <string.h>

/*
 * Supporting macro that intercepts and prints a possible
 * error state during calling TraceEvent(...)
 *
 * Call TRACE_EVENT(TraceEvent(...)) <=> TraceEvent(...)
 */
#define TRACE_EVENT(prog_name, trace_event) \
if((int)((trace_event))=(-1)) \
{ \
    (void) fprintf \
    ( \
        stderr, \
        "%s: line:%d function call TraceEvent() failed, errno(%d): %s\n", \
        prog_name, \
        __LINE__, \
        errno, \
        strerror(errno) \
    ); \
    \
    return (-1); \
}

/*
 * Prints error message
 */
#define TRACE_ERROR_MSG(prog_name, msg) \
(void) fprintf(stderr, "%s: %s\n", prog_name, msg)

#define __INSTREX_H_INCLUDED
#endif
```

You'll have to save `instrex.h` in the same directory as the C code in order to compile the programs.

Gathering all events from all classes

In our first example, we'll set up daemon mode to gather *all* events from *all* classes. Here's the source, `all_classes.c`:

```
/*
 * $QNXLicenseC:
 * Copyright 2007, QNX Software Systems. All Rights Reserved.
 *
 * You must obtain a written license from and pay applicable license fees to QNX
 * Software Systems before you may reproduce, modify or distribute this software,
 * or any work that includes all or part of this software. Free development
 * licenses are available for evaluation and non-commercial purposes. For more
 * information visit http://licensing.qnx.com or email licensing@qnx.com.
 *
 * This file may contain contributions from others. Please review this entire
 * file for other proprietary rights or license notices, as well as the QNX
 * Development Suite License Guide at http://licensing.qnx.com/license-guide/
 * for other information.
 * $
 */

#ifdef __USAGE
%C - instrumentation example

%C - example that illustrates the very basic use of
the TraceEvent() kernel call and the instrumentation
module with tracelogger in a daemon mode.

All classes and their events are included and monitored.

In order to use this example, start the tracelogger
in the daemon mode as:

tracelogger -n iter_number -d1

with iter_number = your choice of 1 through 10

After you start the example, the tracelogger (daemon)
will log the specified number of iterations and then
terminate. There are no messages printed upon successful
completion of the example. You can view the intercepted
events with the traceprinter utility.

See accompanied documentation and comments within
the sample source code for more explanations.
#endif

#include <sys/trace.h>
#include "instrex.h"

int main(int argc, char **argv)
{
    /*
     * Just in case, turn off all filters, since we
     * don't know their present state - go to the
     * known state of the filters.
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_DELALLCLASSES));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_KERCALL));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_KERCALL));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_THREAD));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_THREAD));

    /*
     * Set fast emitting mode for all classes and
     * their events.
     * Wide emitting mode could have been
     * set instead, using:
     *
     * TraceEvent(_NTO_TRACE_SETALLCLASSESWIDE)
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_SETALLCLASSESFAST));

    /*
     * Intercept all event classes
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_ADDALLCLASSES));

    /*
     * Start tracing process
     *
     * During the tracing process, the tracelogger (which
     * is being executed in a daemon mode) will log all events.
     * You can specify the number of iterations (i.e. the
     * number of kernel buffers logged) when you start tracelogger.
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_START));

    /*
     * The main() of this execution flow returns.
     */
}
```

```

    * However, the main() function of the tracelogger
    * will return after registering the specified number
    * of events.
    */
    return (0);
}

```

Compile it, and then run tracelogger in one window:

```
tracelogger -d1 -n 3 -f all_classes.kev
```

and run the compiled program in another:

```
./all_classes
```

Despite how quickly the program ran, the amount of data it generated is rather overwhelming.

The trace data is in `all_classes.kev`; to examine it, type:

```
traceprinter -f all_classes.kev | less
```

The output from traceprinter will look something like this:

```

TRACEPRINTER version 1.02
TRACEPARSER LIBRARY version 1.02
-- HEADER FILE INFORMATION --
    TRACE_FILE_NAME:: all_classes.kev
    TRACE_DATE:: Wed Jun 24 10:52:58 2009
    TRACE_VER_MAJOR:: 1
    TRACE_VER_MINOR:: 01
    TRACE_LITTLE_ENDIAN:: TRUE
    TRACE_ENCODING:: 16 byte events
    TRACE_BOOT_DATE:: Tue Jun 23 11:47:46 2009
    TRACE_CYCLES_PER_SEC:: 736629000
    TRACE_CPU_NUM:: 1
    TRACE_SYSNAME:: QNX
    TRACE_NODENAME:: localhost
    TRACE_SYS_RELEASE:: 6.4.1
    TRACE_SYS_VERSION:: 2009/05/20-17:35:56EDT
    TRACE_MACHINE:: x86pc
    TRACE_SYSPAGE_LEN:: 2264
TRACE_TRACELOGGER_ARGS:: tracelogger -d1 -n 3 -f all_classes.kev
-- KERNEL EVENTS --
t:0x4f81e320 CPU:00 CONTROL: BUFFER sequence = 33, num_events = 714
t:0x4f81e320 CPU:00 CONTROL: TIME msb:0x000037b0 lsb(offset):0x4f81e014
t:0x4f82017a CPU:00 PROCESS: PROC_CREATE_NAME
    ppid:0
    pid:1
    name:proc/boot/procnto-smp-instr
...
t:0x4f854048 CPU:00 THREAD: THCREATE pid:1511472 tid:1
t:0x4f854140 CPU:00 THREAD: THRUNNING pid:1511472 tid:1
t:0x4f854910 CPU:00 KER_EXIT: TRACE_EVENT/01 ret_val:0x00000000 empty:0x00000000
t:0x4f856aac CPU:00 KER_CALL: THREAD_DESTROY/47 tid:-1 status_p:0
t:0x4f857dca CPU:00 KER_EXIT: THREAD_DESTROY/47 ret_val:0x00000030 empty:0x00000000
t:0x4f8588d3 CPU:00 KER_CALL: THREAD_DESTROYALL/48 empty:0x00000000 empty:0x00000000
t:0x4f858ed7 CPU:00 THREAD: THDESTROY pid:1511472 tid:1
t:0x4f8598b9 CPU:00 THREAD: THDEAD pid:1511472 tid:1
t:0x4f859c4c CPU:00 THREAD: THRUNNING pid:1 tid:1
t:0x4f85c6e3 CPU:00 COMM: SND_PULSE_EXE scoid:0x40000002 pid:1
t:0x4f85cecd CPU:00 THREAD: THRUNNING pid:1 tid:12
t:0x4f85d5ad CPU:00 THREAD: THREADY pid:1 tid:1
t:0x4f85e5b3 CPU:00 COMM: REC_PULSE scoid:0x40000002 pid:1
t:0x4f860ee2 CPU:00 KER_CALL: THREAD_CREATE/46 func_p:f0023170 arg_p:eff6e000
t:0x4f8624c7 CPU:00 THREAD: THCREATE pid:1511472 tid:1
t:0x4f8625ff CPU:00 THREAD: THWAITTHREAD pid:1 tid:12
t:0x4f8627b4 CPU:00 THREAD: THRUNNING pid:1511472 tid:1
t:0x4f8636fd CPU:00 THREAD: THREADY pid:1 tid:12
t:0x4f866c34 CPU:00 KER_CALL: CONNECT_SERVER_INFO/41 pid:0 coid:0x00000000
t:0x4f866836 CPU:00 KER_EXIT: CONNECT_SERVER_INFO/41 coid:0x00000000 info->nd:0
t:0x4f86735e CPU:00 KER_CALL: TIMER_TIMEOUT/75 timeout_flags:0x00000050 ntime(sec):30
t:0x4f868445 CPU:00 KER_EXIT: TIMER_TIMEOUT/75 prev_timeout_flags:0x00000000 otime(sec):0
t:0x4f8697d3 CPU:00 INT_ENTR: 0x00000000 (0) IP:0xf008433e
t:0x4f86a276 CPU:00 INT_HANDLER_ENTR: 0x00000000 (0) PID:126997 IP:0x080b7334 AREA:0x0812a060
t:0x4f86afa7 CPU:00 INT_HANDLER_EXIT: 0x00000000 (0) SIGEVENT:NONE
t:0x4f86b304 CPU:00 INT_HANDLER_ENTR: 0x00000000 (0) PID:1 IP:0xf0056570 AREA:0x00000000
t:0x4f86ca12 CPU:00 INT_HANDLER_EXIT: 0x00000000 (0) SIGEVENT:NONE
t:0x4f86cff6 CPU:00 INT_EXIT: 0x00000000 (0) inkernel:0x00000f01
t:0x4f86e276 CPU:00 KER_CALL: MSG_SENDRV/11 coid:0x00000000 msg:"" (0x00040116)
t:0x4f86e756 CPU:00 COMM: SND_MESSAGE rcvid:0x0000004f pid:159762
t:0x4f86f84a CPU:00 THREAD: THREPLY pid:1511472 tid:1
t:0x4f8705dd CPU:00 THREAD: THREADY pid:159762 tid:1
t:0x4f8707d4 CPU:00 THREAD: THRUNNING pid:159762 tid:1
t:0x4f870bfb CPU:00 COMM: REC_MESSAGE rcvid:0x0000004f pid:159762
t:0x4f878b6c CPU:00 KER_CALL: MSG_REPLYV/15 rcvid:0x0000004f status:0x00000000
t:0x4f878f4b CPU:00 COMM: REPLY_MESSAGE tid:1 pid:1511472
t:0x4f8798d2 CPU:00 THREAD: THREADY pid:1511472 tid:1
t:0x4f879db8 CPU:00 KER_EXIT: MSG_REPLYV/15 ret_val:0 empty:0x00000000
t:0x4f87a84f CPU:00 KER_CALL: MSG_RECEIVEV/14 chid:0x00000001 rparts:1
...

```

This example demonstrates the capability of the trace module to capture *huge* amounts of data about the events. The first part of the trace data is the initial state information about all the running processes; to suppress it, start the tracing with `_NTO_TRACE_STARTNOSTATE` instead of `_NTO_TRACE_START`.

While it's good to know how to gather everything, we'll clearly need to be able to refine our search.

Gathering all events from one class

Now we'll gather *all* events from only *one* class, `_NTO_TRACE_THREAD`. This class is arbitrarily chosen to demonstrate filtering by classes; there's nothing particularly special about this class versus any other. For a full list of the possible classes, see "[Classes and events](#) (p. 24)" in the Events and the Kernel chapter in this guide.

Here's the source, `one_class.c`:

```
/*
 * $QNXLicenseC:
 * Copyright 2007, QNX Software Systems. All Rights Reserved.
 *
 * You must obtain a written license from and pay applicable license fees to QNX
 * Software Systems before you may reproduce, modify or distribute this software,
 * or any work that includes all or part of this software. Free development
 * licenses are available for evaluation and non-commercial purposes. For more
 * information visit http://licensing.qnx.com or email licensing@qnx.com.
 *
 * This file may contain contributions from others. Please review this entire
 * file for other proprietary rights or license notices, as well as the QNX
 * Development Suite License Guide at http://licensing.qnx.com/license-guide/
 * for other information.
 * $
 */

#ifdef __USAGE
%C - instrumentation example

%C - example that illustrates the very basic use of
the TraceEvent() kernel call and the instrumentation
module with tracelogger in a daemon mode.

Only events from the thread class (_NTO_TRACE_THREAD)
are monitored (intercepted).

In order to use this example, start the tracelogger
in the daemon mode as:

tracelogger -n iter_number -d1

with iter_number = your choice of 1 through 10

After you start the example, the tracelogger (daemon)
will log the specified number of iterations and then
terminate. There are no messages printed upon successful
completion of the example. You can view the intercepted
events with the traceprinter utility.

See accompanied documentation and comments within
the sample source code for more explanations.
#endif

#include <sys/trace.h>
#include "instrex.h"

int main(int argc, char **argv)
{
    /*
     * Just in case, turn off all filters, since we
     * don't know their present state - go to the
     * known state of the filters.
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_DEALLCLASSES));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_KERCALL));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_KERCALL));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_THREAD));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_THREAD));

    /*
     * Intercept only thread events
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_ADDCLASS, _NTO_TRACE_THREAD));

    /*
     * Start tracing process
     *
     * During the tracing process, the tracelogger (which
     * is being executed in daemon mode) will log all events.
     * You can specify the number of iterations (i.e. the
     * number of kernel buffers logged) when you start tracelogger.
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_START));

    /*
     * The main() of this execution flow returns.
     */
}
```

```

    * However, the main() function of the tracelogger
    * will return after registering the specified number
    * of events.
    */
    return (0);
}

```

Compile it, and then run tracelogger in one window:

```
tracelogger -d1 -n 3 -f one_class.kev
```

and run the compiled program in another:

```
./one_class
```

The trace data is in one_class.kev; to examine it, type:

```
traceprinter -f one_class.kev | less
```

The output from traceprinter will look something like this:

```

TRACEPRINTER version 1.02
TRACEPARSER LIBRARY version 1.02
-- HEADER FILE INFORMATION --
    TRACE_FILE_NAME:: one_class.kev
    TRACE_DATE:: Wed Jun 24 10:55:05 2009
    TRACE_VER_MAJOR:: 1
    TRACE_VER_MINOR:: 01
    TRACE_LITTLE_ENDIAN:: TRUE
    TRACE_ENCODING:: 16 byte events
    TRACE_BOOT_DATE:: Tue Jun 23 11:47:46 2009
    TRACE_CYCLES_PER_SEC:: 736629000
    TRACE_CPU_NUM:: 1
    TRACE_SYSNAME:: QNX
    TRACE_NODENAME:: localhost
    TRACE_SYS_RELEASE:: 6.4.1
    TRACE_SYS_VERSION:: 2009/05/20-17:35:56EDT
    TRACE_MACHINE:: x86pc
    TRACE_SYSPAGE_LEN:: 2264
TRACE_TRACELOGGER_ARGS:: tracelogger -d1 -n 3 -f one_class.kev
-- KERNEL EVENTS --
t:0x002c4d55 CPU:00 CONTROL: BUFFER sequence = 37, num_events = 714
t:0x002c4d55 CPU:00 THREAD :THCREATE      pid:1 tid:1
t:0x002c5531 CPU:00 THREAD :THREADY       pid:1 tid:1 priority:0 policy:1
t:0x002c5bbe CPU:00 THREAD :THCREATE      pid:1 tid:2
t:0x002c5cd2 CPU:00 THREAD :THRECEIVE     pid:1 tid:2 priority:255 policy:2
t:0x002c6185 CPU:00 THREAD :THCREATE      pid:1 tid:3
t:0x002c6272 CPU:00 THREAD :THRECEIVE     pid:1 tid:3 priority:255 policy:2
t:0x002c64eb CPU:00 THREAD :THCREATE      pid:1 tid:4
t:0x002c65d8 CPU:00 THREAD :THRECEIVE     pid:1 tid:4 priority:10 policy:2
t:0x002c67fc CPU:00 THREAD :THCREATE      pid:1 tid:5
t:0x002c68ea CPU:00 THREAD :THRECEIVE     pid:1 tid:5 priority:255 policy:2
t:0x002c6bae CPU:00 THREAD :THCREATE      pid:1 tid:7
t:0x002c6c9b CPU:00 THREAD :THRECEIVE     pid:1 tid:7 priority:10 policy:2
t:0x002c6f03 CPU:00 THREAD :THCREATE      pid:1 tid:8
t:0x002c6ff0 CPU:00 THREAD :THRECEIVE     pid:1 tid:8 priority:10 policy:2
t:0x002c72ec CPU:00 THREAD :THCREATE      pid:1 tid:10
t:0x002c73d9 CPU:00 THREAD :THRECEIVE     pid:1 tid:10 priority:10 policy:2
t:0x002c7665 CPU:00 THREAD :THCREATE      pid:1 tid:11
t:0x002c7752 CPU:00 THREAD :THRECEIVE     pid:1 tid:11 priority:10 policy:2
t:0x002c7a9d CPU:00 THREAD :THCREATE      pid:1 tid:12
t:0x002c7b8a CPU:00 THREAD :THRECEIVE     pid:1 tid:12 priority:10 policy:2
t:0x002c7e44 CPU:00 THREAD :THCREATE      pid:1 tid:15
t:0x002c7f31 CPU:00 THREAD :THRECEIVE     pid:1 tid:15 priority:10 policy:2
t:0x002c81a2 CPU:00 THREAD :THCREATE      pid:1 tid:20
t:0x002c828f CPU:00 THREAD :THRECEIVE     pid:1 tid:20 priority:10 policy:2
t:0x002c88e3 CPU:00 THREAD :THCREATE      pid:2 tid:1
t:0x002c89d3 CPU:00 THREAD :THREPLY       pid:2 tid:1 priority:10 policy:3
t:0x002c8fad CPU:00 THREAD :THCREATE      pid:4099 tid:1
t:0x002c909a CPU:00 THREAD :THRECEIVE     pid:4099 tid:1 priority:10 policy:3
t:0x002c95b7 CPU:00 THREAD :THCREATE      pid:4100 tid:1
t:0x002c96a4 CPU:00 THREAD :THRECEIVE     pid:4100 tid:1 priority:10 policy:3
t:0x002c9b6e CPU:00 THREAD :THCREATE      pid:4101 tid:1
t:0x002c9ccd CPU:00 THREAD :THSIGWAITINFO pid:4101 tid:1 priority:10 policy:3
...

```

Notice that we've significantly reduced the amount of output.

Gathering five events from four classes

Now that we can gather specific classes of events, we'll refine our search even further and gather only five specific types of events from four classes.

Here's the source, `five_events.c`:

```
/*
 * $QNXLicenseC:
 * Copyright 2007, QNX Software Systems. All Rights Reserved.
 *
 * You must obtain a written license from and pay applicable license fees to QNX
 * Software Systems before you may reproduce, modify or distribute this software,
 * or any work that includes all or part of this software. Free development
 * licenses are available for evaluation and non-commercial purposes. For more
 * information visit http://licensing.qnx.com or email licensing@qnx.com.
 *
 * This file may contain contributions from others. Please review this entire
 * file for other proprietary rights or license notices, as well as the QNX
 * Development Suite License Guide at http://licensing.qnx.com/license-guide/
 * for other information.
 * $
 */

#ifdef __USAGE
%C - instrumentation example

%C - example that illustrates the very basic use of
the TraceEvent() kernel call and the instrumentation
module with tracelogger in a daemon mode.

Only five events from four classes are included and
monitored. Class _NTO_TRACE_KERCALL is intercepted
in a wide emitting mode.

In order to use this example, start the tracelogger
in the daemon mode as:

tracelogger -n iter_number -d1

with iter_number = your choice of 1 through 10

After you start the example, the tracelogger (daemon)
will log the specified number of iterations and then
terminate. There are no messages printed upon successful
completion of the example. You can view the intercepted
events with the traceprinter utility.

See accompanied documentation and comments within
the example source code for more explanations.
#endif

#include <sys/trace.h>
#include <sys/kercalls.h>

#include "instrex.h"

int main(int argc, char **argv)
{
    /*
     * Just in case, turn off all filters, since we
     * don't know their present state - go to the
     * known state of the filters.
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_DEALLCLASSES));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_KERCALL));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSTID, _NTO_TRACE_KERCALL));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_THREAD));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSTID, _NTO_TRACE_THREAD));

    /*
     * Set wide emitting mode
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_SETALLCLASSESWIDE));

    /*
     * Intercept two events from class _NTO_TRACE_THREAD
     */
    TRACE_EVENT
    (
        argv[0],
        TraceEvent(_NTO_TRACE_ADDEVENT, _NTO_TRACE_THREAD, _NTO_TRACE_THRUNNING)
    );
    TRACE_EVENT
    (
        argv[0],
        TraceEvent(_NTO_TRACE_ADDEVENT, _NTO_TRACE_THREAD, _NTO_TRACE_THCREATE)
    );
}
```

```

    /* Intercept one event from class _NTO_TRACE_PROCESS
    */
    TRACE_EVENT
    (
        argv[0],
        TraceEvent(_NTO_TRACE_ADDEVENT, _NTO_TRACE_PROCESS, _NTO_TRACE_PROCCREATE_NAME)
    );

    /*
    * Intercept one event from class _NTO_TRACE_INTENTER
    */
    TRACE_EVENT
    (
        argv[0],
        TraceEvent(_NTO_TRACE_ADDEVENT, _NTO_TRACE_INTENTER, _NTO_TRACE_INTFIRST)
    );

    /*
    * Intercept one event from class _NTO_TRACE_KERCALLEXIT,
    * event __KER_MSG_READV.
    */
    TRACE_EVENT
    (
        argv[0],
        TraceEvent(_NTO_TRACE_ADDEVENT, _NTO_TRACE_KERCALLEXIT, __KER_MSG_READV)
    );

    /*
    * Start tracing process
    *
    * During the tracing process, the tracelogger (which
    * is being executed in a daemon mode) will log all events.
    * You can specify the number of iterations (i.e. the
    * number of kernel buffers logged) when you start tracelogger.
    */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_START));

    /*
    * The main() of this execution flow returns.
    * However, the main() function of the tracelogger
    * will return after registering the specified number
    * of events.
    */
    return (0);
}

```

Compile it, and then run `tracelogger` in one window:

```
tracelogger -d1 -n 3 -f five_events.kev
```

and run the compiled program in another:

```
./five_events
```

The trace data is in `five_events.kev`; to examine it, type:

```
traceprinter -f five_events.kev | less
```

The output from `traceprinter` will look something like this:

```

TRACEPRINTER version 1.02
TRACEPARSER LIBRARY version 1.02
-- HEADER FILE INFORMATION --
    TRACE_FILE_NAME:: five_events.kev
    TRACE_DATE:: Wed Jun 24 10:56:04 2009
    TRACE_VER_MAJOR:: 1
    TRACE_VER_MINOR:: 01
    TRACE_LITTLE_ENDIAN:: TRUE
    TRACE_ENCODING:: 16 byte events
    TRACE_BOOT_DATE:: Tue Jun 23 11:47:46 2009
    TRACE_CYCLES_PER_SEC:: 736629000
    TRACE_CPU_NUM:: 1
    TRACE_SYSNAME:: QNX
    TRACE_NODENAME:: localhost
    TRACE_SYS_RELEASE:: 6.4.1
    TRACE_SYS_VERSION:: 2009/05/20-17:35:56EDT
    TRACE_MACHINE:: x86pc
    TRACE_SYSPAGE_LEN:: 2264
TRACE_TRACELOGGER_ARGS:: tracelogger -d1 -n 3 -f five_events.kev
-- KERNEL EVENTS --
t:0x1a14da34 CPU:00 CONTROL: BUFFER sequence = 41, num_events = 714
t:0x1a14da34 CPU:00 PROCESS :PROCCREATE_NAME
    pid:0
    pid:1
    name:proc/boot/procnto-smp-instr
t:0x1a14ea7d CPU:00 THREAD :THCREATE pid:1 tid:1
t:0x1a14ed04 CPU:00 THREAD :THCREATE pid:1 tid:2
...
t:0x1a1cc345 CPU:00 THREAD :THRUNNING pid:1 tid:15 priority:10 policy:2
t:0x1a1d01b9 CPU:00 THREAD :THRUNNING pid:8200 tid:5 priority:10 policy:3
t:0x1a1d6424 CPU:00 INT_ENTR:0x00000000 (0) IP:0xb8229890
t:0x1a1ed261 CPU:00 THREAD :THRUNNING pid:1 tid:4 priority:10 policy:2
t:0x1a1f0016 CPU:00 THREAD :THRUNNING pid:426022 tid:1 priority:10 policy:2

```

```
...  
t:0x2e77ebc5 CPU:00 THREAD :THRUNNING      pid:1613871 tid:1 priority:10 policy:2  
t:0x2e78598d CPU:00 THREAD :THRUNNING      pid:8200 tid:5 priority:10 policy:3  
t:0x2e7ac4fc CPU:00 INT_ENTR:0x00000000 (0)  IP:0xb8229f61  
t:0x2e7cec3b CPU:00 KER_EXIT:MSG_READV/16  
      rbytes:22540  
      rmsg:"" (0x1a15080f 0x6e696273 0x6e69742f)  
t:0x2e7da478 CPU:00 THREAD :THRUNNING      pid:1003562 tid:1 priority:10 policy:2  
t:0x2e7dc288 CPU:00 THREAD :THRUNNING      pid:1 tid:15 priority:10 policy:2  
...
```

We've now begun to selectively pick and choose events—the massive amount of data is now much more manageable.

Gathering kernel calls

The kernel calls are arguably the most important class of calls. This example shows not only filtering, but also the arguments intercepted by the instrumented kernel. In its base form, this program is similar to the `one_class.c` example that gathered only one class.

Here's the source, `ker_calls.c`:

```
/*
 * $QNXLicenseC:
 * Copyright 2007, QNX Software Systems. All Rights Reserved.
 *
 * You must obtain a written license from and pay applicable license fees to QNX
 * Software Systems before you may reproduce, modify or distribute this software,
 * or any work that includes all or part of this software. Free development
 * licenses are available for evaluation and non-commercial purposes. For more
 * information visit http://licensing.qnx.com or email licensing@qnx.com.
 *
 * This file may contain contributions from others. Please review this entire
 * file for other proprietary rights or license notices, as well as the QNX
 * Development Suite License Guide at http://licensing.qnx.com/license-guide/
 * for other information.
 * $
 */

#ifdef __USAGE
%C - instrumentation example

%C - [-n num]

%C - example that illustrates the very basic use of
the TraceEvent() kernel call and the instrumentation
module with tracelogger in a daemon mode.

All thread states and all/one (specified) kernel
call number are intercepted. The kernel call(s)
is(are) intercepted in wide emitting mode.

Options:
-n <num> kernel call number to be intercepted
(default is all)

In order to use this example, start the tracelogger
in the daemon mode as:

tracelogger -n iter_number -d1

with iter_number = your choice of 1 through 10

After you start the example, the tracelogger (daemon)
will log the specified number of iterations and then
terminate. There are no messages printed upon successful
completion of the example. You can view the intercepted
events with the traceprinter utility.

See accompanied documentation and comments within
the sample source code for more explanations.
#endif

#include <sys/trace.h>
#include <unistd.h>
#include <stdlib.h>

#include "instrex.h"

int main(int argc, char **argv)
{
    int arg_var; /* input arguments parsing support */
    int call_num=(-1); /* kernel call number to be intercepted */

    /* Parse command line arguments
     * - get optional kernel call number
     */
    while((arg_var=getopt(argc, argv, "n:"))!=(-1)) {
        switch(arg_var)
        {
            case 'n': /* get kernel call number */
                call_num = strtoul(optarg, NULL, 10);
                break;
            default: /* unknown */
                TRACE_ERROR_MSG
                (
                    argv[0],
                    "error parsing command-line arguments - exiting\n"
                );
        }
    }
}
```

```

        return (-1);
    }
}

/*
 * Just in case, turn off all filters, since we
 * don't know their present state - go to the
 * known state of the filters.
 */
TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_DEALLCLASSES));
TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_KERCALL));
TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSTID, _NTO_TRACE_KERCALL));
TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_THREAD));
TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSTID, _NTO_TRACE_THREAD));

/*
 * Set wide emitting mode for all classes and
 * their events.
 */
TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_SETALLCLASSES));

/*
 * Intercept _NTO_TRACE_THREAD class
 * We need it to know the state of the active thread.
 */
TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_ADDCLASS, _NTO_TRACE_THREAD));

/*
 * Add all/one kernel call
 */
if(call_num != (-1)) {
    TRACE_EVENT
    (
        argv[0],
        TraceEvent(_NTO_TRACE_ADDEVENT, _NTO_TRACE_KERCALL, call_num)
    );
} else {
    TRACE_EVENT
    (
        argv[0],
        TraceEvent(_NTO_TRACE_ADDCLASS, _NTO_TRACE_KERCALL)
    );
}

/*
 * Start tracing process
 *
 * During the tracing process, the tracelogger (which
 * is being executed in a daemon mode) will log all events.
 * You can specify the number of iterations (i.e. the
 * number of kernel buffers logged) when you start tracelogger.
 */
TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_START));

/*
 * The main() of this execution flow returns.
 * However, the main() function of the tracelogger
 * will return after registering the specified number
 * of events.
 */
return (0);
}

```

Compile it, and then run tracelogger in one window:

```
tracelogger -d1 -n 3 -f ker_calls.all.kev
```

and run the compiled program in another:

```
./ker_calls
```

The trace data is in ker_calls.all.kev; to examine it, type:

```
traceprinter -f ker_calls.all.kev | less
```

The output from traceprinter will look something like this:

```

TRACEPRINTER version 1.02
TRACEPARSER LIBRARY version 1.02
-- HEADER FILE INFORMATION --
  TRACE_FILE_NAME:: ker_calls.all.kev
  TRACE_DATE:: Wed Jun 24 10:57:01 2009
  TRACE_VER_MAJOR:: 1
  TRACE_VER_MINOR:: 01
  TRACE_LITTLE_ENDIAN:: TRUE
  TRACE_ENCODING:: 16 byte events
  TRACE_BOOT_DATE:: Tue Jun 23 11:47:46 2009
  TRACE_CYCLES_PER_SEC:: 736629000
  TRACE_CPU_NUM:: 1
  TRACE_SYSNAME:: QNX
  TRACE_NODENAME:: localhost
  TRACE_SYS_RELEASE:: 6.4.1
  TRACE_SYS_VERSION:: 2009/05/20-17:35:56EDT
  TRACE_MACHINE:: x86pc
  TRACE_SYSPAGE_LEN:: 2264

```

```

TRACE_TRACELOGGER_ARGS:: tracelogger -d1 -n 3 -f ker_calls.all.kev
-- KERNEL EVENTS --
t:0x463ad541 CPU:00 CONTROL: BUFFER sequence = 45, num_events = 714
t:0x463ad541 CPU:00 THREAD :THCREATE pid:1 tid:1
t:0x463adbe1 CPU:00 THREAD :THREADY pid:1 tid:1 priority:0 policy:1
t:0x463adfa8 CPU:00 THREAD :THCREATE pid:1 tid:2
t:0x463ae098 CPU:00 THREAD :THRECEIVE pid:1 tid:2 priority:255 policy:2
t:0x463ae375 CPU:00 THREAD :THCREATE pid:1 tid:3
...
t:0x463d59b6 CPU:00 THREAD :THSIGWAITINFO pid:1658927 tid:2 priority:10 policy:2
t:0x463d5cb2 CPU:00 THREAD :THCREATE pid:1663024 tid:1
t:0x463d5da7 CPU:00 THREAD :THRUNNING pid:1663024 tid:1 priority:10 policy:2
t:0x463d666e CPU:00 KER_EXIT:TRACE_EVENT/01 ret_val:0x00000000 empty:0x00000000
t:0x463d8e65 CPU:00 KER_CALL:THREAD_DESTROY/47
tid:-1
priority:-1
status_p:0
t:0x463da615 CPU:00 KER_EXIT:THREAD_DESTROY/47 ret_val:0x00000030 empty:0x00000000
t:0x463daf0a CPU:00 KER_CALL:THREAD_DESTROYALL/48 empty:0x00000000 empty:0x00000000
t:0x463db531 CPU:00 THREAD :THDESTROY pid:1663024 tid:1
t:0x463dc114 CPU:00 THREAD :THDEAD pid:1663024 tid:1 priority:10 policy:2
t:0x463dc546 CPU:00 THREAD :THRUNNING pid:1 tid:1 priority:0 policy:1
t:0x463df45d CPU:00 THREAD :THRUNNING pid:1 tid:4 priority:10 policy:2
t:0x463dfa7f CPU:00 THREAD :THREADY pid:1 tid:1 priority:0 policy:1
t:0x463e36b4 CPU:00 KER_CALL:THREAD_CREATE/46
pid:1663024
func_p:f0023170
arg_p:eff4e000
flags:0x00000000
stacksize:10116
stackaddr_p:eff4e264
exitfunc_p:0
policy:0
sched_priority:0
sched_curpriority:0
param.__ss_low_priority:0
param.__ss_max_repl:0
param.__ss_repl_period.tv_sec:0
param.__ss_repl_period.tv_nsec:0
param.__ss_init_budget.tv_sec:0
param.__ss_init_budget.tv_nsec:0
param.empty:0
param.empty:0
guardsize:0
empty:0
empty:0
empty:0
t:0x463e50b0 CPU:00 THREAD :THCREATE pid:1663024 tid:1
t:0x463e51d0 CPU:00 THREAD :THWAITTHREAD pid:1 tid:4 priority:10 policy:2
t:0x463e5488 CPU:00 THREAD :THRUNNING pid:1663024 tid:1 priority:10 policy:2
t:0x463e6408 CPU:00 THREAD :THREADY pid:1 tid:4 priority:10 policy:2
...

```

The `ker_calls.c` program takes a `-n` option that lets us view only one type of kernel call. Let's run this program again, specifying the number 14, which signifies `__KER_MSG_RECEIVE`. For a full list of the values associated with the `-n` option, see `/usr/include/sys/kercalls.h`.

Run `tracelogger` in one window:

```
tracelogger -d1 -n 3 -f ker_calls.14.kev
```

and run the program in another:

```
./ker_calls -n 14
```

The trace data is in `ker_calls.14.kev`; to examine it, type:

```
traceprinter -f ker_calls.14.kev | less
```

The output from `traceprinter` will look something like this:

```

TRACEPRINTER version 1.02
TRACEPARSER LIBRARY version 1.02
-- HEADER FILE INFORMATION --
TRACE_FILE_NAME:: ker_calls.14.kev
TRACE_DATE:: Wed Jun 24 13:39:20 2009
TRACE_VER_MAJOR:: 1
TRACE_VER_MINOR:: 01
TRACE_LITTLE_ENDIAN:: TRUE
TRACE_ENCODING:: 16 byte events
TRACE_BOOT_DATE:: Tue Jun 23 11:47:46 2009
TRACE_CYCLES_PER_SEC:: 736629000
TRACE_CPU_NUM:: 1
TRACE_SYSNAME:: QNX
TRACE_NODENAME:: localhost
TRACE_SYS_RELEASE:: 6.4.1
TRACE_SYS_VERSION:: 2009/05/20-17:35:56EDT

```

```

TRACE_MACHINE:: x86pc
TRACE_SYSPAGE_LEN:: 2264
TRACE_TRACELOGGER_ARGS:: tracelogger -d1 -n 3 -f ker_calls.14.kev
-- KERNEL_EVENTS --
t:0x73bf28d0 CPU:00 CONTROL: BUFFER sequence = 62, num_events = 714
t:0x73bf28d0 CPU:00 THREAD :THCREATE      pid:1 tid:1
t:0x73bf2e16 CPU:00 THREAD :THREADY       pid:1 tid:1 priority:0 policy:1
t:0x73bf3203 CPU:00 THREAD :THCREATE      pid:1 tid:2
...

t:0x73c21746 CPU:00 THREAD :THRUNNING     pid:1 tid:1 priority:0 policy:1
t:0x73c24352 CPU:00 THREAD :THRUNNING     pid:1 tid:15 priority:10 policy:2
t:0x73c247e0 CPU:00 THREAD :THREADY       pid:1 tid:1 priority:0 policy:1
t:0x73c2547b CPU:00 KER_EXIT:MSG_RECEIVEV/14
      rcvid:0x00000000
      rmsg:"" (0x00000000 0x00000081 0x001dd030)
      info->nd:0
      info->srcnd:0
      info->pid:1953840
      info->tid:1
      info->chid:1
      info->scoid:1073741874
      info->coid:1073741824
      info->msglen:0
      info->srcmsglen:56
      info->dstmsglen:24
      info->priority:10
      info->flags:0
      info->reserved:0
t:0x73c29270 CPU:00 THREAD :THCREATE      pid:1953840 tid:1
t:0x73c293ca CPU:00 THREAD :THWAITTHREAD pid:1 tid:15 priority:10 policy:2
t:0x73c2964a CPU:00 THREAD :THRUNNING     pid:1953840 tid:1 priority:10 policy:2
t:0x73c2a36c CPU:00 THREAD :THREADY       pid:1 tid:15 priority:10 policy:2
t:0x73c2fccc CPU:00 THREAD :THREPLY      pid:1953840 tid:1 priority:10 policy:2
t:0x73c30f6b CPU:00 THREAD :THREADY       pid:159762 tid:1 priority:10 policy:3
t:0x73c311b0 CPU:00 THREAD :THRUNNING     pid:159762 tid:1 priority:10 policy:3
t:0x73c31835 CPU:00 KER_EXIT:MSG_RECEIVEV/14
      rcvid:0x0000004f
      rmsg:"" (0x00040116 0x00000000 0x00000004)
      info->nd:0
      info->srcnd:0
      info->pid:1953840
      info->tid:1
      info->chid:1
      info->scoid:1073741903
      info->coid:0
      info->msglen:4
      info->srcmsglen:4
      info->dstmsglen:0
      info->priority:10
      info->flags:0
      info->reserved:0
t:0x73c3a359 CPU:00 THREAD :THREADY       pid:1953840 tid:1 priority:10 policy:2
t:0x73c3af50 CPU:00 KER_CALL:MSG_RECEIVEV/14 chid:0x00000001 rparts:1
t:0x73c3b25e CPU:00 THREAD :THRECEIVE     pid:159762 tid:1 priority:10 policy:3
...

```

Event handling - simple

In this example, we intercept two events from two different classes. Each event has an event handler attached to it; the event handlers are closing and opening the stream. Here's the source, `eh_simple.c`:

```

/*
 * $QNXLicenseC:
 * Copyright 2007, QNX Software Systems. All Rights Reserved.
 *
 * You must obtain a written license from and pay applicable license fees to QNX
 * Software Systems before you may reproduce, modify or distribute this software,
 * or any work that includes all or part of this software. Free development
 * licenses are available for evaluation and non-commercial purposes. For more
 * information visit http://licensing.qnx.com or email licensing@qnx.com.
 *
 * This file may contain contributions from others. Please review this entire
 * file for other proprietary rights or license notices, as well as the QNX
 * Development Suite License Guide at http://licensing.qnx.com/license-guide/
 * for other information.
 * $
 */

#ifdef __USAGE
%C - instrumentation example

%C - example that illustrates the very basic use of
the TraceEvent() kernel call and the instrumentation
module with tracelogger in a daemon mode.

Two events from two classes are included and monitored
interchangeably. The flow control of monitoring the
specified events is controlled with attached event
handlers.

In order to use this example, start the tracelogger
in the daemon mode as:

tracelogger -n 1 -d1

After you start the example, the tracelogger (daemon)
will log the specified number of iterations and then
terminate. There are no messages printed upon successful
completion of the example. You can view the intercepted
events with the traceprinter utility.

See accompanied documentation and comments within
the sample source code for more explanations.
#endif

#include <unistd.h>
#include <sys/trace.h>
#include <sys/kercalls.h>
#include "instrex.h"

/*
 * Prepare event structure where the event data will be
 * stored and passed to an event handler.
 */
event_data_t e_d_1;
uint32_t data_array_1[20]; /* 20 elements for potential args. */

event_data_t e_d_2;
uint32_t data_array_2[20]; /* 20 elements for potential args. */

/*
 * Global state variable that controls the
 * event flow between two events
 */
int g_state;

/*
 * Event handler attached to the event __KER_MSG_SENDV
 * from the _NTO_TRACE_KERCALL class.
 */
int call_msg_send_eh(event_data_t* e_d)
{
    if(g_state) {
        g_state = !g_state;
        return (1);
    }

    return (0);
}

/*
 * Event handler attached to the event _NTO_TRACE_THRUNNING
 * from the _NTO_TRACE_THREAD (thread) class.
 */

```

```

int thread_run_eh(event_data_t* e_d)
{
    if(!g_state) {
        g_state = !g_state;
        return (1);
    }

    return (0);
}

int main(int argc, char **argv)
{
    /*
     * First fill arrays inside event data structures
     */
    e_d_1.data_array = data_array_1;
    e_d_2.data_array = data_array_2;

    /*
     * Just in case, turn off all filters, since we
     * don't know their present state - go to the
     * known state of the filters.
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_DEALLCLASSES));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_KERCALL));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSTID, _NTO_TRACE_KERCALL));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_THREAD));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSTID, _NTO_TRACE_THREAD));

    /*
     * Set fast emitting mode
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_SETALLCLASSESFAST));

    /*
     * Obtain I/O privileges before adding event handlers
     */
    if (ThreadCtl(_NTO_TCTL_IO, 0) != EOK) { /* obtain I/O privileges */
        (void) fprintf(stderr, "argv[0]: Failed to obtain I/O privileges\n");
        return (-1);
    }

    /*
     * Intercept one event from class _NTO_TRACE_KERCALL,
     * event __KER_MSG_SENDDV.
     */
    TRACE_EVENT
    (
        argv[0],
        TraceEvent(_NTO_TRACE_ADDEVENT, _NTO_TRACE_KERCALLENTER, __KER_MSG_SENDDV)
    );

    /*
     * Add event handler to the event __KER_MSG_SENDDV
     * from _NTO_TRACE_KERCALL class.
     */
    TRACE_EVENT
    (
        argv[0],
        TraceEvent(_NTO_TRACE_ADDEVENTHANDLER, _NTO_TRACE_KERCALLENTER,
                    __KER_MSG_SENDDV, call_msg_send_eh, &e_d_1)
    );

    /*
     * Intercept one event from class _NTO_TRACE_THREAD
     */
    TRACE_EVENT
    (
        argv[0],
        TraceEvent(_NTO_TRACE_ADDEVENT, _NTO_TRACE_THREAD, _NTO_TRACE_THRUNNING)
    );

    /*
     * Add event handler to the _NTO_TRACE_THRUNNING event
     * from the _NTO_TRACE_THREAD (thread) class.
     */
    TRACE_EVENT
    (
        argv[0],
        TraceEvent(_NTO_TRACE_ADDEVENTHANDLER, _NTO_TRACE_THREAD,
                    _NTO_TRACE_THRUNNING, thread_run_eh, &e_d_2)
    );

    /*
     * Start tracing process
     *
     * During the tracing process, the tracelogger (which
     * is being executed in a daemon mode) will log all events.
     * The number of iterations has been specified as 1.
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_START));

    /*
     * During one second collect all events
     */
    (void) sleep(1);

    /*
     * Stop tracing process by closing the event stream.
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_STOP));

    /*

```

```

    * Flush the internal buffer since the number
    * of stored events could be less than
    * "high water mark" of one buffer (715 events).
    *
    * The tracelogger will probably terminate at
    * this point, since it has been executed with
    * one iteration (-n 1 option).
    */
TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_FLUSHBUFFER));

/*
 * Delete event handlers before exiting to avoid execution
 * in the missing address space.
 */
TRACE_EVENT
(
    argv[0],
    TraceEvent(_NTO_TRACE_DELEVENTHANDLER, _NTO_TRACE_KERCALLENTER, __KER_MSG_SENDV)
);
TRACE_EVENT
(
    argv[0],
    TraceEvent(_NTO_TRACE_DELEVENTHANDLER, _NTO_TRACE_THREAD, _NTO_TRACE_THRUNNING)
);

/*
 * Wait one second before terminating to hold the address space
 * of the event handlers.
 */
(void) sleep(1);

return (0);
}

```

Compile it, and then run tracelogger in one window:

```
tracelogger -d1 -n 1 -f eh_simple.kev
```



For this example, we've specified the number of iterations to be 1.

Run the compiled program in another window:

```
./eh_simple
```

The trace data is in eh_simple.kev; to examine it, type:

```
traceprinter -f eh_simple.kev | less
```

The output from traceprinter will look something like this:

```

TRACEPRINTER version 1.02
TRACEPARSER LIBRARY version 1.02
-- HEADER FILE INFORMATION --
  TRACE_FILE_NAME:: eh_simple.kev
  TRACE_DATE:: Wed Jun 24 10:58:41 2009
  TRACE_VER_MAJOR:: 1
  TRACE_VER_MINOR:: 01
  TRACE_LITTLE_ENDIAN:: TRUE
  TRACE_ENCODING:: 16 byte events
  TRACE_BOOT_DATE:: Tue Jun 23 11:47:46 2009
  TRACE_CYCLES_PER_SEC:: 736629000
  TRACE_CPU_NUM:: 1
  TRACE_SYSNAME:: QNX
  TRACE_NODENAME:: localhost
  TRACE_SYS_RELEASE:: 6.4.1
  TRACE_SYS_VERSION:: 2009/05/20-17:35:56EDT
  TRACE_MACHINE:: x86pc
  TRACE_SYSPAGE_LEN:: 2264
TRACE_TRACELOGGER_ARGS:: tracelogger -d1 -n 1 -f eh_simple.kev
-- KERNEL EVENTS --
t:0x33139a74 CPU:00 CONTROL: BUFFER sequence = 53, num_events = 482
t:0x33139a74 CPU:00 THREAD :THRUNNING pid:1749040 tid:1
t:0x362d0710 CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000003 msg:"" (0x00100102)
t:0x362d1d04 CPU:00 THREAD :THRUNNING pid:217117 tid:1
t:0x362e8e3e CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000000 msg:"" (0x00200113)
t:0x362ea264 CPU:00 THREAD :THRUNNING pid:4102 tid:8
t:0x362f1248 CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000003 msg:"" (0x00000106)
t:0x362f1c67 CPU:00 THREAD :THRUNNING pid:217117 tid:1
t:0x362fd08b CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000003 msg:"" (0x00100102)
t:0x362fd949 CPU:00 THREAD :THRUNNING pid:217117 tid:1
t:0x36305424 CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000003 msg:"" (0x00000106)
t:0x36305e35 CPU:00 THREAD :THRUNNING pid:217117 tid:1
t:0x3630a572 CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000003 msg:"" (0x00000106)
t:0x3630aeb7 CPU:00 THREAD :THRUNNING pid:217117 tid:1
t:0x3631bd5b CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000003 msg:"" (0x00100102)
t:0x3631c6aa CPU:00 THREAD :THRUNNING pid:217117 tid:1
t:0x363253bb CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000003 msg:"" (0x00000106)
t:0x36325d95 CPU:00 THREAD :THRUNNING pid:217117 tid:1
t:0x369b2349 CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000003 msg:"" (0x00000106)
t:0x369b2bbe CPU:00 THREAD :THRUNNING pid:217117 tid:1

```

```
t:0x369b88d8 CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000007 msg:"" (0x00100106)
t:0x369b974a CPU:00 THREAD :THRUNNING pid:1 tid:15
t:0x369c48ab CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000008 msg:"" (0x00100106)
t:0x369c53db CPU:00 THREAD :THRUNNING pid:126997 tid:2
t:0x369cee17 CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000008 msg:"" (0x00100106)
t:0x369cf533 CPU:00 THREAD :THRUNNING pid:126997 tid:2
t:0x369d82b6 CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000009 msg:"" (0x00100106)
t:0x369d9178 CPU:00 THREAD :THRUNNING pid:8200 tid:10
t:0x369eae84 CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000006 msg:"" (0x00020100)
t:0x369eb687 CPU:00 THREAD :THRUNNING pid:1 tid:15
t:0x369f56b1 CPU:00 KER_CALL:MSG_SENDV/11 coid:0x00000006 msg:"" (0x00020100)

...
```

This is an important example because it demonstrates the use of the dynamic rules filter to perform tasks beyond basic filtering.

Inserting a user simple event

This example demonstrates the insertion of a user event into the event stream. Here's the source, `usr_event_simple.c`:

```

/*
 * $QNXLicenseC:
 * Copyright 2007, QNX Software Systems. All Rights Reserved.
 *
 * You must obtain a written license from and pay applicable license fees to QNX
 * Software Systems before you may reproduce, modify or distribute this software,
 * or any work that includes all or part of this software. Free development
 * licenses are available for evaluation and non-commercial purposes. For more
 * information visit http://licensing.qnx.com or email licensing@qnx.com.
 *
 * This file may contain contributions from others. Please review this entire
 * file for other proprietary rights or license notices, as well as the QNX
 * Development Suite License Guide at http://licensing.qnx.com/license-guide/
 * for other information.
 * $
 */

#ifdef __USAGE
%C - instrumentation example

%C - example that illustrates the very basic use of
the TraceEvent() kernel call and the instrumentation
module with tracelogger in a daemon mode.

All classes and their events are included and monitored.
Additionally, four user-generated simple events and
one string event are intercepted.

In order to use this example, start the tracelogger
in the daemon mode as:

tracelogger -n iter_number -d1

with iter_number = your choice of 1 through 10

After you start the example, the tracelogger (daemon)
will log the specified number of iterations and then
terminate. There are no messages printed upon successful
completion of the example. You can view the intercepted
events with the traceprinter utility. The intercepted
user events (class USREVENT) have event IDs
(EVENT) equal to: 111, 222, 333, 444 and 555.

See accompanied documentation and comments within
the sample source code for more explanations.
#endif

#include <sys/trace.h>
#include <unistd.h>

#include "instrex.h"

int main(int argc, char **argv)
{
    /*
     * Just in case, turn off all filters, since we
     * don't know their present state - go to the
     * known state of the filters.
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_DEALLCLASSES));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_KERCALL));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSTID, _NTO_TRACE_KERCALL));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSPID, _NTO_TRACE_THREAD));
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_CLRCLASSTID, _NTO_TRACE_THREAD));

    /*
     * Set fast emitting mode for all classes and
     * their events.
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_SETALLCLASSESFAST));

    /*
     * Intercept all event classes
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_ADDALLCLASSES));

    /*
     * Start tracing process
     *
     * During the tracing process, the tracelogger (which
     * is being executed in a daemon mode) will log all events.
     * You can specify the number of iterations (i.e. the
     * number of kernel buffers logged) when you start tracelogger.
     */
    TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_START));

    /*

```

```

    * Insert four user-defined simple events and one string
    * event into the event stream. The user events have
    * arbitrary event IDs: 111, 222, 333, 444, and 555
    * (possible values are in the range 0...1023).
    * The user events with ID=(111, ..., 444) are simple events
    * that have two numbers attached: ({1,11}, ..., {4,44}).
    * The user string event (ID 555) includes the string,
    * "Hello world".
    */
TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_INSERTSUSEREVENT, 111, 1, 11));
TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_INSERTSUSEREVENT, 222, 2, 22));
TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_INSERTSUSEREVENT, 333, 3, 33));
TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_INSERTSUSEREVENT, 444, 4, 44));
TRACE_EVENT(argv[0], TraceEvent(_NTO_TRACE_INSERTUSRSTREVENT, 555, "Hello world" ));

/*
 * The main() of this execution flow returns.
 * However, the main() function of the tracelogger
 * will return after registering the specified number
 * of events.
 */
return (0);
}

```

Compile it, and then run `tracelogger` in one window:

```
tracelogger -d1 -n 3 -f usr_event_simple.kev
```

and run the compiled program in another:

```
./usr_event_simple
```

The trace data is in `usr_event_simple.kev`; to examine it, type:

```
traceprinter -f usr_event_simple.kev | less
```

The output from `traceprinter` will look something like this:

```

TRACEPRINTER version 1.02
TRACEPARSER LIBRARY version 1.02
-- HEADER FILE INFORMATION --
  TRACE_FILE_NAME:: usr_event_simple.kev
    TRACE_DATE:: Wed Jun 24 10:59:34 2009
    TRACE_VER_MAJOR:: 1
    TRACE_VER_MINOR:: 01
  TRACE_LITTLE_ENDIAN:: TRUE
  TRACE_ENCODING:: 16 byte events
  TRACE_BOOT_DATE:: Tue Jun 23 11:47:46 2009
  TRACE_CYCLES_PER_SEC:: 736629000
    TRACE_CPU_NUM:: 1
    TRACE_SYSNAME:: QNX
    TRACE_NODENAME:: localhost
  TRACE_SYS_RELEASE:: 6.4.1
  TRACE_SYS_VERSION:: 2009/05/20-17:35:56EDT
    TRACE_MACHINE:: x86pc
    TRACE_SYSPAGE_LEN:: 2264
TRACE_TRACELOGGER_ARGS:: tracelogger -d1 -n 3 -f usr_event_simple.kev
-- KERNEL EVENTS --
t:0x254620e4 CPU:00 CONTROL: BUFFER sequence = 54, num_events = 714
...

t:0x25496c81 CPU:00 PROCESS :PROCCREATE_NAME
                        ppid:426022
                        pid:1810480
                        name:./usr_event_simple
t:0x2549701a CPU:00 THREAD :THCREATE      pid:1810480 tid:1
t:0x25497112 CPU:00 THREAD :THRUNNING    pid:1810480 tid:1
t:0x2549793a CPU:00 KER_EXIT:TRACE_EVENT/01 ret_val:0x00000000 empty:0x00000000
t:0x25497f48 CPU:00 USREVENT:EVENT:111, d0:0x00000001 d1:0x0000000b
t:0x254982c5 CPU:00 USREVENT:EVENT:222, d0:0x00000002 d1:0x00000016
t:0x25498638 CPU:00 USREVENT:EVENT:333, d0:0x00000003 d1:0x00000021
t:0x25498996 CPU:00 USREVENT:EVENT:444, d0:0x00000004 d1:0x0000002c
t:0x25499451 CPU:00 USREVENT:EVENT:555 STR:"Hello world"
t:0x2549bde5 CPU:00 KER_CALL:THREAD_DESTROY/47 tid:-1 status_p:0
t:0x2549d0d6 CPU:00 KER_EXIT:THREAD_DESTROY/47 ret_val:0x00000030 empty:0x00000000
t:0x2549d8ae CPU:00 KER_CALL:THREAD_DESTROYALL/48 empty:0x00000000 empty:0x00000000
...

```

Inserting your own events lets you flag events or bracket groups of events to isolate them for study. It's also useful for inserting internal, customized information into the event stream.

Appendix A

Current Trace Events and Data

This appendix provides a table that lists all the trace events and summarizes the data included for each in both wide and fast modes.

Interpreting the table

As you examine the table, note the following:

- Some of the functions listed below (e.g. *InterruptDetachFunc()*, *SignalFault()*) are internal ones that you won't find documented in the QNX Neutrino *C Library Reference*.
- If a function has a restartable version (with a `_r` in its name), the events for both versions are as listed for the function without the `_r`.
- If a kernel call fails, the exit trace event includes the return code and the error code (e.g. an *errno* value), instead of the data listed below.

As an example, let's look at the events for *MsgSend()*, *MsgSendv()*, and *MsgSendvs()*. As mentioned above, the information is the same for the restartable versions of these functions too.

Here's what the table gives for the entry (`_NTO_TRACE_KERCALLENTER`) to these functions:

```
Class: _NTO_TRACE_KERCALLENTER
Event: __KER_MSG_SENDEV
Fast: coid, msg
Wide: coid, sparts, rparts, msg[0], msg[1], msg[2]
Call: MsgSend,MsgSendv,MsgSendvs
#Args: MSG_SENDEV, fHcoid, Dsparts, Drparts, fSmsg, s, s
```

This part describes the `__KER_MSG_SENDEV` trace event that's emitted on entry to the function. In fast mode, the event includes the following data:

Fast mode data	Number of bytes for the event
Connection ID	4 bytes
Message data	4 bytes (the first 4 bytes usually comprise the header)
	Total emitted: 8 bytes

In wide mode, the event includes the following data:

Wide mode data	Number of bytes for the event
Connection ID	4 bytes
# of parts to send	4 bytes
# of parts to receive	4 bytes
Message data	4 bytes (the first 4 bytes usually comprise the header)
Message data	4 bytes

Wide mode data	Number of bytes for the event
Message data	4 bytes
	Total emitted: 24 bytes

The second (`_NTO_TRACE_KERCALLEEXIT`) part describes the `__KER_MSG_SENDEV` event that's emitted on exit from the function:

```
Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_MSG_SENDEV
Fast: status, rmsg[0]
Wide: status, rmsg[0], rmsg[1], rmsg[2]
Call: MsgSend,MsgSendv,MsgSendvs
#Args: MSG_SENDEV, fDstatus, fSrmsg, s, s
```

In fast mode, the event includes the following data if the kernel call was successful:

Fast mode data	Number of bytes for the event
Exit status	4 bytes
Message data	4 bytes (the first 4 bytes usually comprise the header)
	Total emitted: 8 bytes

In wide mode, the event includes the following data if the kernel call was successful:

Wide mode data	Number of bytes for the event
Exit status	4 bytes
Message data	4 bytes (the first 4 bytes usually comprise the header)
Message data	4 bytes
Message data	4 bytes
	Total emitted: 16 bytes

In both fast and wide mode, the event includes the following data if the kernel call failed:

Fast and wide mode data	Number of bytes for the event
Exit status	4 bytes
Error code	4 bytes
	Total emitted: 8 bytes

For many of the events, you'll see a comment like this:

```
#Args: MSG_SENDEV, fHcoid, Dsparts, Drparts, fSmsg, s, s
```

This line indicates how `traceprinter` displays the data associated with the event. The format codes are as follows:

Code	Format
H	Hexadecimal (32 bit)
D	Decimal (32 bit)
X	Hexadecimal (64 bit)
E	Decimal (64 bit)
S	Begin a character string
s	Continue with a character string
P	Pointer
N	Named string
f	Fast mode prefix

For example, `fHcoid` indicates that the connection ID (`coid`) is displayed as a 32-bit hexadecimal number, and it's included in fast mode (and wide mode).

Table of events

```
# The format of this file is as follows
# #'s are comments
# Event blocks are delimited by a blank line
# Class -> The external class description
# Event -> The external event description
# Fast -> Comma delimited list of wide argument subset
# Wide -> Full arguments emitted
# Call -> If the event is associated with a kernel call, put them here
```

Note - all _NTO_TRACE_KERCALLEXIT calls now have errno as the second parameter if the kercall failed (i.e. the retval is -1)

```
Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_BAD
Fast: empty, empty
Wide: empty, empty
Call: N/A
#Args: BAD, fHempty, fHempty
```

```
Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_BAD
Fast: ret_val, empty
Wide: ret_val, empty
Call: N/A
#Args: BAD, fHret_val, fHempty
```

```
Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CACHE_FLUSH
Fast: addr, nlines
Wide: addr, nlines, flags, index
Call: CacheFlush
#Args: CACHE_FLUSH, fHaddr, fHnlines, fHflags, fDindex
```

```
Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CACHE_FLUSH
Fast: ret_val, empty
Wide: ret_val, empty
Call: CacheFlush
#Args: CACHE_FLUSH, fHret_val, fHempty
```

```
Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CHANNEL_CREATE
Fast: flags, empty
Wide: flags, empty
Call: ChannelCreate
#Args: CHANNEL_CREATE, fHflags, fHempty
```

```
Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CHANNEL_CREATE
Fast: chid, empty
Wide: chid, empty
Call: ChannelCreate
#Args: CHANNEL_CREATE, fHchid, fHempty
```

```
Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CHANNEL_DESTROY
Fast: chid, empty
Wide: chid, empty
Call: ChannelDestroy
#Args: CHANNEL_DESTROY, fHchid, fHempty
```

```
Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CHANNEL_DESTROY
Fast: ret_val, empty
Wide: ret_val, empty
Call: ChannelDestroy
#Args: CHANNEL_DESTROY, fHret_val, fHempty
```

```
Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CHANCON_ATTR
Fast: chid, flags
Wide: chid, flags, new_attr
```

```
Call: ChannelConnectAttr

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CHANCON_ATTR
Fast: chid, flags
Wide: chid, flags, old_attr
Call: ChannelConnectAttr

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_CLOCK_ADJUST
Fast: id, new->tick_count
Wide: id, new->tick_count, new->tick_nsec_inc
Call: ClockAdjust
#Args: CLOCK_ADJUST, fDid, fDnew->tick_count, Dnew->tick_nsec_inc

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CLOCK_ADJUST
Fast: ret_val, old->tick_count
Wide: ret_val, old->tick_count, old->tick_nsec_inc
Call: ClockAdjust
#Args: CLOCK_ADJUST, fDret_val, fDold->tick_count, Dold->tick_nsec_inc

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_CLOCK_ID
Fast: pid, tid
Wide: pid, tid
Call: ClockId
#Args: CLOCK_ID, fDpid, fDt看id

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CLOCK_ID
Fast: ret_val, empty
Wide: ret_val, empty
Call: ClockId
#Args: CLOCK_ID, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_CLOCK_PERIOD
Fast: id, new->nsec
Wide: id, new->nsec, new->fract
Call: ClockPeriod
#Args: CLOCK_PERIOD, fDid, fDnew->nsec, Dnew->fract

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CLOCK_PERIOD
Fast: ret_val, old->nsec
Wide: ret_val, old->nsec, old->fract
Call: ClockPeriod
#Args: CLOCK_PERIOD, fDret_val, fDold->nsec, Dold->fract

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_CLOCK_TIME
Fast: id, new(sec)
Wide: id, new(sec), new(nsec)
Call: ClockTime
#Args: CLOCK_TIME, fDid, fDnew(sec), Dnew(nsec)

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CLOCK_TIME
Fast: ret_val, old(sec)
Wide: ret_val, old(sec), old(nsec)
Call: ClockTime
#Args: CLOCK_TIME, fDret_val, fDold(sec), Dold(nsec)

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_CONNECT_ATTACH
Fast: nd, pid
Wide: nd, pid, chid, index, flags
Call: ConnectAttach
#Args: CONNECT_ATTACH, fHnd, fDpid, Hchid, Dindex, Hflags

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CONNECT_ATTACH
Fast: coid, empty
Wide: coid, empty
Call: ConnectAttach
#Args: CONNECT_ATTACH, fHcoid, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_CONNECT_CLIENT_INFO
```

```

Fast: scoid, ngroups
Wide: scoid, ngroups
Call: ConnectClientInfo
#Args: CONNECT_CLIENT_INFO, fHscoid, fDngroups

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CONNECT_CLIENT_INFO
Fast: ret_val, info->nd
Wide: ret_val, info->nd, info->pid, info->sid, flags, info->ruid, info->euid,
      info->suid, info->rgid, info->egid, info->sgid, info->ngroups,
      info->grouplist[0], info->grouplist[1], info->grouplist[2],
      info->grouplist[3], info->grouplist[4], info->grouplist[5],
      info->grouplist[6], info->grouplist[7]
Call: ConnectClientInfo
#Args: CONNECT_CLIENT_INFO, fHscoid, fDngroups

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CONNECT_DETACH
Fast: coid, empty
Wide: coid, empty
Call: ConnectDetach
#Args: CONNECT_DETACH, fHcoid, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CONNECT_DETACH
Fast: ret_val, empty
Wide: ret_val, empty
Call: ConnectDetach
#Args: CONNECT_DETACH, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CONNECT_FLAGS
Fast: coid, bits
Wide: pid, coid, masks, bits
Call: ConnectFlags
#Args: CONNECT_FLAGS, fDpid, fHcoid, fHmasks, fHbits

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CONNECT_FLAGS
Fast: old_flags, empty
Wide: old_flags, empty
Call: ConnectFlags
#Args: CONNECT_FLAGS, fHold_flags, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CONNECT_SERVER_INFO
Fast: pid, coid
Wide: pid, coid
Call: ConnectServerInfo
#Args: CONNECT_SERVER_INFO, fDpid, fHcoid

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_CONNECT_SERVER_INFO
Fast: coid, info->nd
Wide: coid, info->nd, info->srcnd, info->pid, info->tid, info->chid,
      info->scoid, info->coid, info->msglen, info->srcmsglen, info->dstmsglen,
      info->priority, info->flags, info->reserved
Call: ConnectServerInfo
#Args: CONNECT_SERVER_INFO, fDpid, fHcoid

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_INTERRUPT_ATTACH
Fast: intr, flags
Wide: intr, handler_p, area_p, areaseize, flags
Call: InterruptAttach
#Args: INTERRUPT_ATTACH, fDintr, fPhandler_p, fParea_p, fDareaseize, fHflags

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_INTERRUPT_ATTACH
Fast: int_fun_id, empty
Wide: int_fun_id, empty
Call: InterruptAttach
#Args: INTERRUPT_ATTACH, fHint_fun_id, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_INTERRUPT_DETACH
Fast: id, empty
Wide: id, empty
Call: InterruptDetach
#Args: INTERRUPT_DETACH, fDid, fHempty

```

```
Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_INTERRUPT_DETACH
Fast: ret_val, empty
Wide: ret_val, empty
Call: InterruptDetach
#Args: INTERRUPT_DETACH, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_INTERRUPT_DETACH_FUNC
Fast: intr, handler_p
Wide: intr, handler_p
Call: N/A
#Args: INTERRUPT_DETACH_FUNC, fDintr, fPhandler_p

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_INTERRUPT_DETACH_FUNC
Fast: ret_val, empty
Wide: ret_val, empty
Call: N/A
#Args: INTERRUPT_DETACH_FUNC, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_INTERRUPT_MASK
Fast: intr, id
Wide: intr, id
Call: InterruptMask
#Args: INTERRUPT_MASK, fDintr, fDid

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_INTERRUPT_MASK
Fast: mask_level, empty
Wide: mask_level, empty
Call: InterruptMask
#Args: INTERRUPT_MASK, fHmask_level, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_INTERRUPT_UNMASK
Fast: intr, id
Wide: intr, id
Call: InterruptUnmask
#Args: INTERRUPT_UNMASK, fDintr, fDid

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_INTERRUPT_UNMASK
Fast: mask_level, empty
Wide: mask_level, empty
Call: InterruptUnmask
#Args: INTERRUPT_UNMASK, fHmask_level, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_INTERRUPT_WAIT
Fast: flags, timeout_tv_sec
Wide: flags, timeout_tv_sec, timeout_tv_nsec
Call: InterruptWait
#Args: INTERRUPT_WAIT, fHflags, fDtimeout_tv_sec, Dtimeout_tv_nsec

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_INTERRUPT_WAIT
Fast: ret_val, empty
Wide: ret_val, empty
Call: InterruptWait
#Args: INTERRUPT_WAIT, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_INTERRUPT_CHARACTERISTIC
Fast: id, type, new
Wide: id, type, new
Call: InterruptCharacteristic
#Args: INTERRUPT_CHARACTERISTIC, fDid, fHtype, fDnew

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_INTERRUPT_CHARACTERISTIC
Fast: id, old
Wide: id, old
Call: InterruptCharacteristic
#Args: INTERRUPT_CHARACTERISTIC, fDid, fDold

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_MSG_CURRENT
```

```

Fast: rcvid, empty
Wide: rcvid, empty
Call: MsgCurrent
#Args: MSG_CURRENT, fHrcvid, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_CURRENT
Fast: empty, empty
Wide: empty, empty
Call: MsgCurrent
#Args: MSG_CURRENT, fHempty, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_PAUSE
Fast: rcvid, cookie
Wide: rcvid, cookie
Call: MsgPause
#Args: MSG_CURRENT, fHrcvid, fHcookie

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_PAUSE
Fast: empty, empty
Wide: empty, empty
Call: MsgPause
#Args: MSG_CURRENT, fHempty, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_DELIVER_EVENT
Fast: rcvid, event->sigeval_notify
Wide: rcvid, event->sigeval_notify, event->sigeval_notify_function_p,
      event->sigeval_value, event->sigeval_notify_attributes_p
Call: MsgDeliverEvent

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_DELIVER_EVENT
Fast: ret_val, empty
Wide: ret_val, empty
Call: MsgDeliverEvent
#Args: MSG_DELIVER_EVENT, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_ERROR
Fast: rcvid, err
Wide: rcvid, err
Call: MsgError
#Args: MSG_ERROR, fHrcvid, fDerr

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_ERROR
Fast: ret_val, empty
Wide: ret_val, empty
Call: MsgError
#Args: MSG_ERROR, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_INFO
Fast: rcvid, info_p
Wide: rcvid, info_p
Call: MsgInfo
#Args: MSG_INFO, fHrcvid, fPinfo_p

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_INFO
Fast: ret_val, info->nd
Wide: ret_val, info->nd, info->srcnd, info->pid, info->tid, info->chid,
      info->scooid, info->coid, info->msglen, info->srcmsglen, info->dstmsglen,
      info->priority, info->flags, info->reserved
Call: MsgInfo

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_KEYDATA
Fast: rcvid, op
Wide: rcvid, op
Call: MsgKeyData
#Args: MSG_KEYDATA, fHrcvid, fHop

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_KEYDATA
Fast: ret_val, newkey
Wide: ret_val, newkey

```

```
Call: MsgKeyData
#Args: MSG_KEYDATA, fHret_val, fDnewkey

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_MSG_READIOV
Fast: rcvid, offset
Wide: rcvid, parts, offset, flags
Call: MsgReadiov
#Args: MSG_READIOV, fHrcvid, Dparts, fHoffset, Hflags

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_READIOV
Fast: rbytes, rmsg[0]
Wide: rbytes, rmsg[0], rmsg[1], rmsg[2]
Call: MsgReadiov
#Args: MSG_READIOV, fDrbytes, fSrmsg, s, s

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_MSG_READV
Fast: rcvid, offset
Wide: rcvid, rmsg_p, rparts, offset
Call: MsgRead,MsgReadv
#Args: MSG_READV, fHrcvid, Prmsg_p, Drparts, fHoffset

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_READV
Fast: rbytes, rmsg[0]
Wide: rbytes, rmsg[0], rmsg[1], rmsg[2]
Call: MsgRead,MsgReadv
#Args: MSG_READV, fDrbytes, fSrmsg, s, s

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_MSG_READWRITEV
Fast: src_rcvid, dst_rcvid
Wide: src_rcvid, dst_rcvid
Call: N/A
#Args: MSG_READWRITEV, fHsrc_rcvid, fHdst_rcvid

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_READWRITEV
Fast: msglen, empty
Wide: msglen, empty
Call: N/A
#Args: MSG_READWRITEV, fDmsglen, fHempty

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_MSG_RECEIVEPULSEV
Fast: chid, rparts
Wide: chid, rparts
Call: MsgReceivePulse,MsgReceivePulsev
#Args: MSG_RECEIVEPULSEV, fHchid, fDrparts

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_RECEIVEPULSEV
Fast: rcvid, rmsg[0]
Wide: rcvid, rmsg[0], rmsg[1], rmsg[2], info->nd, info->srcnd, info->pid,
info->tid, info->chid, info->scoid, info->coid, info->msglen,
info->srcmsglen, info->dstmsglen, info->priority, info->flags,
info->reserved
Call: MsgReceivePulse,MsgReceivePulsev

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_MSG_RECEIVEV
Fast: chid, rparts
Wide: chid, rparts
Call: MsgReceive,MsgReceivev
#Args: MSG_RECEIVEV, fHchid, fDrparts

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_RECEIVEV
Fast: rcvid, rmsg[0]
Wide: rcvid, rmsg[0], rmsg[1], rmsg[2], info->nd, info->srcnd, info->pid,
info->tid, info->chid, info->scoid, info->coid, info->msglen,
info->srcmsglen, info->dstmsglen, info->priority, info->flags,
info->reserved
Call: MsgReceive,MsgReceivev

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_MSG_REPLYV
Fast: rcvid, status
```

```

Wide: rcvid, sparts, status, smsg[0], smsg[1], smsg[2]
Call: MsgReply,MsgReplyv
#Args: MSG_REPLYV, fHrcvid, Dsparts, fHstatus, Ssmg, s, s

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_REPLYV
Fast: ret_val, empty
Wide: ret_val, empty
Call: MsgReply,MsgReplyv
#Args: MSG_REPLYV, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_SEND_PULSE
Fast: coid, code
Wide: coid, priority, code, value
Call: MsgSendPulse
#Args: MSG_SEND_PULSE, fHcoid, Dpriority, fHcode, Hvalue

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_SEND_PULSE
Fast: status, empty
Wide: status, empty
Call: MsgSendPulse
#Args: MSG_SEND_PULSE, fDstatus, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_SENDV
Fast: coid, msg
Wide: coid, sparts, rparts, msg[0], msg[1], msg[2]
Call: MsgSend,MsgSendv,MsgSendvs
#Args: MSG_SENDV, fHcoid, Dsparts, Drparts, fSmsg, s, s

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_SENDV
Fast: status, rmsg[0]
Wide: status, rmsg[0], rmsg[1], rmsg[2]
Call: MsgSend,MsgSendv,MsgSendvs
#Args: MSG_SENDV, fDstatus, fSrmsg, s, s

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_SENDVNC
Fast: coid, msg
Wide: coid, sparts, rparts, msg[0], msg[1], msg[2]
Call: MsgSendnc,MsgSendvnc,MsgSendvsn
#Args: MSG_SENDVNC, fHcoid, Dsparts, Drparts, fSmsg, s, s

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_SENDVNC
Fast: ret_val, rmsg[0]
Wide: ret_val, rmsg[0], rmsg[1], rmsg[2]
Call: MsgSendnc,MsgSendvnc,MsgSendvsn
#Args: MSG_SENDVNC, fHret_val, fSrmsg, s, s

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_VERIFY_EVENT
Fast: rcvid, event->sigeval_notify
Wide: rcvid, event->sigeval_notify, event->sigeval_notify_function_p,
event->sigeval_value, event->sigeval_notify_attribute_p
Call: MsgVerifyEvent

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_VERIFY_EVENT
Fast: status, empty
Wide: status, empty
Call: MsgVerifyEvent
#Args: MSG_VERIFY_EVENT, fDstatus, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_WRITEV
Fast: rcvid, offset
Wide: rcvid, sparts, offset, msg[0], msg[1], msg[2]
Call: MsgWrite,MsgWritev
#Args: MSG_WRITEV, fHrcvid, Dsparts, fHoffset, Smsg, s, s

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_MSG_WRITEV
Fast: wbytes, empty
Wide: wbytes, empty
Call: MsgWrite,MsgWritev
#Args: MSG_WRITEV, fDwbytes, fHempty

```

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_NET_CRED
Fast: coid, info_p
Wide: coid, info_p
Call: NetCred
#Args: NET_CRED, fHcoid, fPinfo_p

Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_NET_CRED
Fast: ret_val, info->nd
Wide: ret_val, info->nd, info->pid, info->sid, info->flags, info->ruid,
info->euid, info->suid, info->rgid, info->egid, info->sgid, info->ngroups,
info->grouplist[0], info->grouplist[1], info->grouplist[2],
info->grouplist[3], info->grouplist[4], info->grouplist[5],
info->grouplist[6], info->grouplist[7]
Call: NetCred

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_NET_INFOSCOID
Fast: scoid, infoscoid
Wide: scoid, infoscoid
Call: NetInfoScoid
#Args: NET_INFOSCOID, fHscoid, fHinfoscoid

Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_NET_INFOSCOID
Fast: ret_val, empty
Wide: ret_val, empty
Call: NetInfoScoid
#Args: NET_INFOSCOID, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_NET_SIGNAL_KILL
Fast: pid, signo
Wide: cred->ruid, cred->euid, nd, pid, tid, signo, code, value
Call: NetSignalKill
#Args: NET_SIGNAL_KILL, fDstatus, fHempty

Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_NET_SIGNAL_KILL
Fast: status, empty
Wide: status, empty
Call: NetSignalKill
#Args: NET_SIGNAL_KILL, fDstatus, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_NET_UNBLOCK
Fast: vtid, empty
Wide: vtid, empty
Call: NetUnblock
#Args: NET_UNBLOCK, fHvtid, fHempty

Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_NET_UNBLOCK
Fast: ret_val, empty
Wide: ret_val, empty
Call: NetUnblock
#Args: NET_UNBLOCK, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_NET_VTID
Fast: vtid, info_p
Wide: vtid, info_p, tid, coid, priority, srcmsglen, keydata, srcnd, dstmsglen
Call: NetVtid

Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_NET_VTID
Fast: ret_val, empty
Wide: ret_val, empty
Call: NetVtid
#Args: NET_VTID, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_NOP
Fast: dummy, empty
Wide: dummy, empty
Call: N/A
#Args: NOP, fHdummy, fHempty

```

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_NOP
Fast: empty, empty
Wide: empty, empty
Call: N/A
#Args: NOP, fHempty, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_RING0
Fast: func_p, arg_p
Wide: func_p, arg_p
Call: __Ring0
#Args: RING0, fPfunc_p, fParg_p

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_RING0
Fast: ret_val, empty
Wide: ret_val, empty
Call: __Ring0
#Args: RING0, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SCHED_GET
Fast: pid, tid
Wide: pid, tid
Call: SchedGet
#Args: SCHED_GET, fDpid, fDt看id

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SCHED_GET
Fast: ret_val, sched_priority
Wide: ret_val, sched_priority, param.ss_low_priority, param.ss_max_repl,
      param.ss_repl_period.tv_sec, param.ss_repl_period.tv_nsec,
      param.ss_init_budget.tv_sec, param.ss_init_budget.tv_nsec
Call: SchedGet
#Args: SCHED_GET, fDpid, fDt看id

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SCHED_INFO
Fast: pid, policy
Wide: pid, policy
Call: SchedInfo
#Args: SCHED_INFO, fDpid, fDpolicy

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SCHED_INFO
Fast: ret_val, priority_max
Wide: ret_val, priority_min, priority_max, interval_sec, interval_nsec, priority_priv
Call: SchedInfo
#Args: SCHED_INFO, fDpid, fDpolicy

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SCHED_SET
Fast: pid, sched_priority
Wide: pid, tid, policy, sched_priority, sched_curpriority,
      param.ss_low_priority, param.ss_max_repl, param.ss_repl_period.tv_sec,
      param.ss_repl_period.tv_nsec, param.ss_init_budget.tv_sec,
      param.ss_init_budget.tv_nsec
Call: SchedSet
#Args: SCHED_SET, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SCHED_SET
Fast: ret_val, empty
Wide: ret_val, empty
Call: SchedSet
#Args: SCHED_SET, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SCHED_YIELD
Fast: empty, empty
Wide: empty, empty
Call: SchedYield
#Args: SCHED_YIELD, fHempty, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SCHED_YIELD
Fast: ret_val, empty
Wide: ret_val, empty
Call: SchedYield

```

```
#Args: SCHED_YIELD, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_SIGNAL_FAULT
Fast: sigcode, addr
Wide: sigcode, addr
Call: N/A
#Args: SIGNAL_FAULT, fDsigcode, fPaddr

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SIGNAL_FAULT
Fast: ret_val, reg_1
Wide: ret_val, reg_1, reg_2, reg_3, reg_4, reg_5
Call: N/A
#Args: SIGNAL_FAULT, fHret_val, fHreg_1, Hreg_2, Hreg_3, Hreg_4, Hreg_5

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_SIGNAL_KILL
Fast: pid, signo
Wide: nd, pid, tid, signo, code, value
Call: SignalKill
#Args: SIGNAL_KILL, Hnd, fDpid, Dtid, fDsigno, Hcode, Hvalue

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SIGNAL_KILL
Fast: ret_val, empty
Wide: ret_val, empty
Call: SignalKill
#Args: SIGNAL_KILL, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_SIGNAL_RETURN
Fast: s_p, empty
Wide: s_p, empty
Call: SignalReturn
#Args: SIGNAL_RETURN, fPs_p, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SIGNAL_RETURN
Fast: ret_val, empty
Wide: ret_val, empty
Call: SignalReturn
#Args: SIGNAL_RETURN, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_SIGNAL_SUSPEND
Fast: sig_blocked->bits[0], sig_blocked->bits[1]
Wide: sig_blocked->bits[0], sig_blocked->bits[1]
Call: SignalSuspend
#Args: SIGNAL_SUSPEND, fHsig_blocked->bits[0], fHsig_blocked->bits[1]

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SIGNAL_SUSPEND
Fast: ret_val, sig_blocked_p
Wide: ret_val, sig_blocked_p
Call: SignalSuspend
#Args: SIGNAL_SUSPEND, fHret_val, fPsig_blocked_p

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_SIGNAL_WAITINFO
Fast: sig_wait->bits[0], sig_wait->bits[1]
Wide: sig_wait->bits[0], sig_wait->bits[1]
Call: SignalWaitinfo
#Args: SIGNAL_WAITINFO, fHsig_wait->bits[0], fHsig_wait->bits[1]

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SIGNAL_WAITINFO
Fast: sig_num, si_code
Wide: sig_num, si_signo, si_code, si_errno, p[0], p[1], p[2], p[3], p[4], p[5], p[6]
Call: SignalWaitinfo
#Args: SIGNAL_WAITINFO, fHsig_wait->bits[0], fHsig_wait->bits[1]

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_SYNC_CONDVAR_SIGNAL
Fast: sync_p, all
Wide: sync_p, all, sync->count, sync->owner
Call: SyncCondvarSignal
#Args: SYNC_CONDVAR_SIGNAL, fPsync_p, fDall, Dsync->count, Dsync->owner

Class: _NTO_TRACE_KERCALLEXIT
```

```

Event: __KER_SYNC_CONDVAR_SIGNAL
Fast: ret_val, empty
Wide: ret_val, empty
Call: SyncCondvarSignal
#Args: SYNC_CONDVAR_SIG, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_SYNC_CONDVAR_WAIT
Fast: sync_p, mutex_p
Wide: sync_p, mutex_p, sync->count, sync->owner, mutex->count, mutex->owner
Call: SyncCondvarWait
#Args: SYNC_CONDVAR_WAIT, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_SYNC_CONDVAR_WAIT
Fast: ret_val, empty
Wide: ret_val, empty
Call: SyncCondvarWait
#Args: SYNC_CONDVAR_WAIT, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_SYNC_CREATE
Fast: type, sync_p
Wide: type, sync_p, count, owner, protocol, flags, prioceiling, clockid
Call: SyncCreate
#Args: SYNC_CREATE, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_SYNC_CREATE
Fast: ret_val, empty
Wide: ret_val, empty
Call: SyncCreate
#Args: SYNC_CREATE, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_SYNC_CTL
Fast: cmd, sync_p
Wide: cmd, sync_p, data_p, count, owner
Call: SyncCtl
#Args: SYNC_CTL, fDcmd, fPsync_p, Pdata_p, Dcount, Downer

Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_SYNC_CTL
Fast: ret_val, empty
Wide: ret_val, empty
Call: SyncCtl
#Args: SYNC_CTL, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_SYNC_DESTROY
Fast: sync_p, owner
Wide: sync_p, count, owner
Call: SyncDestroy
#Args: SYNC_DESTROY, fPsync_p, Dcount, fDowner

Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_SYNC_DESTROY
Fast: ret_val, empty
Wide: ret_val, empty
Call: SyncDestroy
#Args: SYNC_DESTROY, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_SYNC_MUTEX_LOCK
Fast: sync_p, owner
Wide: sync_p, count, owner
Call: SyncMutexLock
#Args: SYNC_MUTEX_LOCK, fPsync_p, Dcount, fDowner

Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_SYNC_MUTEX_LOCK
Fast: ret_val, empty
Wide: ret_val, empty
Call: SyncMutexLock
#Args: SYNC_MUTEX_LOCK, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_SYNC_MUTEX_REVIVE
Fast: sync_p, owner
Wide: sync_p, count, owner

```

```
Call: SyncMutexRevive
#Args: SYNC_MUTEX_REVIVE, fPsync_p, Dcount, fDowner

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SYNC_MUTEX_REVIVE
Fast: ret_val, empty
Wide: ret_val, empty
Call: SyncMutexRevive
#Args: SYNC_MUTEX_REVIVE, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_SYNC_MUTEX_UNLOCK
Fast: sync_p, owner
Wide: sync_p, count, owner
Call: SyncMutexUnlock
#Args: SYNC_MUTEX_UNLOCK, fPsync_p, Dcount, fDowner

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SYNC_MUTEX_UNLOCK
Fast: ret_val, empty
Wide: ret_val, empty
Call: SyncMutexUnlock
#Args: SYNC_MUTEX_UNLOCK, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_SYNC_SEM_POST
Fast: sync_p, count
Wide: sync_p, count, owner
Call: SyncSemPost
#Args: SYNC_SEM_POST, fPsync_p, fDcount, Downer

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SYNC_SEM_POST
Fast: ret_val, empty
Wide: ret_val, empty
Call: SyncSemPost
#Args: SYNC_SEM_POST, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_SYNC_SEM_WAIT
Fast: sync_p, count
Wide: sync_p, try, count, owner
Call: SyncSemWait
#Args: SYNC_SEM_WAIT, fPsync_p, Dtry, fDcount, Downer

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SYNC_SEM_WAIT
Fast: ret_val, empty
Wide: ret_val, empty
Call: SyncSemWait
#Args: SYNC_SEM_WAIT, fDret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_SYS_CPUPAGE_GET
Fast: index, empty
Wide: index, empty
Call: N/A
#Args: SYS_CPUPAGE_GET, fDindex, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SYS_CPUPAGE_GET
Fast: ret_val, empty
Wide: ret_val, empty
Call: N/A
#Args: SYS_CPUPAGE_GET, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLEENTER
Event: __KER_THREAD_CANCEL
Fast: tid, canstub_p
Wide: tid, canstub_p
Call: ThreadCancel
#Args: THREAD_CANCEL, fDtid, fPcanstub_p

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_THREAD_CANCEL
Fast: ret_val, empty
Wide: ret_val, empty
Call: ThreadCancel
#Args: THREAD_CANCEL, fHret_val, fHempty
```

```

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_THREAD_CREATE
Fast: func_p, arg_p
Wide: pid, func_p, arg_p, flags, stacksize, stackaddr_p, exitfunc_p, policy,
      sched_priority, sched_curpriority, param.ss_low_priority, param.ss_max_repl,
      param.ss_repl_period.tv_sec, param.ss_repl_period.tv_nsec,
      param.ss_init_budget.tv_sec, param.ss_init_budget.tv_nsec
Call: ThreadCreate
#Args: THREAD_CREATE, fHthread_id, fHowner

Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_THREAD_CREATE
Fast: thread_id, owner
Wide: thread_id, owner
Call: ThreadCreate
#Args: THREAD_CREATE, fHthread_id, fHowner

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_THREAD_CTL
Fast: cmd, data_p
Wide: cmd, data_p
Call: ThreadCtl
#Args: THREAD_CTL, fHcmd, fPdata_p

Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_THREAD_CTL
Fast: ret_val, empty
Wide: ret_val, empty
Call: ThreadCtl
#Args: THREAD_CTL, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_THREAD_DESTROY
Fast: tid, status_p
Wide: tid, priority, status_p
Call: ThreadDestroy
#Args: THREAD_DESTROY, fDtid, Dpriority, fPstatus_p

Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_THREAD_DESTROY
Fast: ret_val, empty
Wide: ret_val, empty
Call: ThreadDestroy
#Args: THREAD_DESTROY, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_THREAD_DESTROYALL
Fast: empty, empty
Wide: empty, empty
Call: N/A
#Args: THREAD_DESTROYALL, fHempty, fHempty

Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_THREAD_DESTROYALL
Fast: ret_val, empty
Wide: ret_val, empty
Call: N/A
#Args: THREAD_DESTROYALL, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_THREAD_DETACH
Fast: tid, empty
Wide: tid, empty
Call: ThreadDetach
#Args: THREAD_DETACH, fDtid, fHempty

Class: _NTO_TRACE_KERCALLEEXIT
Event: __KER_THREAD_DETACH
Fast: ret_val, empty
Wide: ret_val, empty
Call: ThreadDetach
#Args: THREAD_DETACH, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_THREAD_JOIN
Fast: tid, status_p
Wide: tid, status_p
Call: ThreadJoin
#Args: THREAD_JOIN, fDtid, fPstatus_p

```

```
Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_THREAD_JOIN
Fast: ret_val, status_p
Wide: ret_val, status_p
Call: ThreadJoin
#Args: THREAD_JOIN, fHret_val, fPstatus_p

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_TIMER_CREATE
Fast: timer_id, event->sigev_notify
Wide: timer_id, event->sigev_notify, event->sigev_notify_function_p,
     event->sigev_value, event->sigev_notify_attributes_p
Call: TimerCreate
#Args: TIMER_CREATE, fHtimer_id, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_TIMER_CREATE
Fast: timer_id, empty
Wide: timer_id, empty
Call: TimerCreate
#Args: TIMER_CREATE, fHtimer_id, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_TIMER_DESTROY
Fast: id, empty
Wide: id, empty
Call: TimerDestroy
#Args: TIMER_DESTROY, fHid, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_TIMER_DESTROY
Fast: ret_val, empty
Wide: ret_val, empty
Call: TimerDestroy
#Args: TIMER_DESTROY, fHret_val, fHempty

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_TIMER_INFO
Fast: pid, id
Wide: pid, id, flags, info_p
Call: TimerInfo
#Args: TIMER_INFO, fDpid, fHid, Hflags, Pinfo_p

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_TIMER_INFO
Fast: prev_id, info->itime.nsec
Wide: prev_id, info->itime.nsec, info->itime.interval_nsec, info->otime.nsec,
     info->otime.interval_nsec, info->flags, info->tid, info->notify,
     info->clockid, info->overruns, info->event.sigev_notify,
     info->event.sigev_notify_function_p, info->event.sigev_value,
     info->event.sigev_notify_attributes_p
Call: TimerInfo
#Args: TIMER_INFO, fDpid, fHid, Hflags, Pinfo_p

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_TRACE_EVENT
Fast: mode, class
Wide: mode, class, event, data_1, data_2
Call: TraceEvent
#Args: TRACE_EVENT, fHmode, fHclass[header], Hevent[time_off], Hdata_1, Hdata_2

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_TRACE_EVENT
Fast: ret_val, empty
Wide: ret_val, empty
Call: TraceEvent
#Args: TRACE_EVENT, fHret_val, fHempty

Class: _NTO_TRACE_INTENTER
Event: _NTO_TRACE_INTFIRST - _NTO_TRACE_INTLAST
Fast: IP, kernel_flag
Wide: interrupt_number, kernel_flag
Call: N/A

Class: _NTO_TRACE_INTEXIT
Event: _NTO_TRACE_INTFIRST - _NTO_TRACE_INTLAST
Fast: interrupt_number, kernel_flag
Wide: interrupt_number, kernel_flag
Call: N/A
```

```

Class: _NTO_TRACE_INT_HANDLER_ENTER
Event: _NTO_TRACE_INTFIRST - _NTO_TRACE_INTLAST
Fast: pid, interrupt_number, ip, area
Wide: pid, interrupt_number, ip, area
Call: N/A

Class: _NTO_TRACE_INT_HANDLER_EXIT
Event: _NTO_TRACE_INTFIRST - _NTO_TRACE_INTLAST
Fast: interrupt_number, sigevent
Wide: interrupt_number, sigevent
Call: N/A

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_SIGNAL_ACTION
Fast: signo, act->sa_handler_p
Wide: pid, sigstubs_p, signo, act->sa_handler_p, act->sa_flags,
      act->sa_mask.bits[0], act->sa_mask.bits[1]
Call: SignalAction

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SIGNAL_ACTION
Fast: ret_val, act->sa_handler_p
Wide: ret_val, act->sa_handler_p, act->sa_flags, act->sa_mask.bits[0],
      act->sa_mask.bits[1]
Call: SignalAction

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_SIGNAL_PROCMASK
Fast: pid, tid
Wide: pid, tid, how, sig_blocked->bits[0], sig_blocked->bits[1]
Call: SignalProcmask

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_SIGNAL_PROCMASK
Fast: ret_val, sig_blocked->bits[0]
Wide: ret_val, sig_blocked->bits[0], sig_blocked->bits[1]
Call: SignalProcmask

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_TIMER_SETTIME
Fast: clock_id, itime->nsec
Wide: clock_id, flags, itime->nsec, itime->interval_nsec
Call: TimerSettime

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_TIMER_SETTIME
Fast: ret_val, itime->nsec
Wide: ret_val, itime->nsec, itime->interval_nsec
Call: TimerSettime

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_TIMER_ALARM
Fast: clock_id, itime->nsec
Wide: clock_id, itime->nsec, itime->interval_nsec
Call: TimerAlarm

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_TIMER_ALARM
Fast: ret_val, itime->nsec
Wide: ret_val, itime->nsec, itime->interval_nsec
Call: TimerAlarm

Class: _NTO_TRACE_KERCALLENTER
Event: __KER_TIMER_TIMEOUT
Fast: clock_id, timeout_flags, ntime
Wide: clock_id, timeout_flags, ntime, event->sigev_notify,
      event->sigev_notify_function_p, event->sigev_value,
      event->sigev_notify_attributes_p
Call: TimerTimeout

Class: _NTO_TRACE_KERCALLEXIT
Event: __KER_TIMER_TIMEOUT
Fast: prev_timeout_flags, otime
Wide: prev_timeout_flags, otime
Call: TimerTimeout

# Control Events

Class: _NTO_TRACE_CONTROL
Event: _NTO_TRACE_CONTROLTIME

```

Fast: msbtime, lsbtime
Wide: msbtime, lsbtime
Call: N/A

Class: _NTO_TRACE_CONTROL
Event: _NTO_TRACE_CONTROLBUFFER
Fast: buffer sequence number, num events
Wide: buffer sequence number, num events
Call: N/A

Process Events

Class: _NTO_TRACE_PROCESS
Event: _NTO_TRACE_PROCCREATE
Fast: ppid, pid
Wide: ppid, pid
Call: N/A

Class: _NTO_TRACE_PROCESS
Event: _NTO_TRACE_PROCCREATE_NAME
Fast: ppid, pid, name
Wide: ppid, pid, name
Call: N/A

Class: _NTO_TRACE_PROCESS
Event: _NTO_TRACE_PROCDESTROY
Fast: ppid, pid
Wide: ppid, pid
Call: N/A

Class: _NTO_TRACE_PROCESS
Event: _NTO_TRACE_PROCDESTROY_NAME
Fast: ppid, pid, name
Wide: ppid, pid, name
Call: N/A

Class: _NTO_TRACE_PROCESS
Event: _NTO_TRACE_PROCTHREAD_NAME
Fast: pid, tid, name
Wide: pid, tid, name
Call: N/A

#Thread state changes

Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THDEAD
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
(incl APS_SCHED_* critical bit def'd in kermacros.h)
** note: partition id and sched_flags only present if APS scheduler module
is loaded.
Call: N/A

Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THRUNNING
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
(incl APS_SCHED_* critical bit def'd in kermacros.h)
** note: partition id and sched_flags only present if APS scheduler module
is loaded.
Call: N/A

Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THREADY
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
(incl APS_SCHED_* critical bit def'd in kermacros.h)
** note: partition id and sched_flags only present if APS scheduler module
is loaded.
Call: N/A

Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THSTOPPED
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
(incl APS_SCHED_* critical bit def'd in kermacros.h)
note: partition id and sched_flags only present if APS scheduler module
is loaded.
** note: partition id and sched_flags only present if APS scheduler module
is loaded.

Call: N/A

Class: _NTO_TRACE_THREAD

Event: _NTO_TRACE_THSEND

Fast: pid, tid

Wide: pid, tid, priority, policy, partition id, sched_flags

(incl APS_SCHED_* critical bit def'd in kermacros.h)

note: partition id and sched_flags only present if APS scheduler module is loaded.

** note: partition id and sched_flags only present if APS scheduler module is loaded.

Call: N/A

Class: _NTO_TRACE_THREAD

Event: _NTO_TRACE_THRECEIVE

Fast: pid, tid

Wide: pid, tid, priority, policy, partition id, sched_flags

(incl APS_SCHED_* critical bit def'd in kermacros.h)

note: partition id and sched_flags only present if APS scheduler module is loaded.

** note: partition id and sched_flags only present if APS scheduler module is loaded.

Call: N/A

Class: _NTO_TRACE_THREAD

Event: _NTO_TRACE_THREPLY

Fast: pid, tid

Wide: pid, tid, priority, policy, partition id, sched_flags

(incl APS_SCHED_* critical bit def'd in kermacros.h)

note: partition id and sched_flags only present if APS scheduler module is loaded.

** note: partition id and sched_flags only present if APS scheduler module is loaded.

Call: N/A

Class: _NTO_TRACE_THREAD

Event: _NTO_TRACE_THSTACK

Fast: pid, tid

Wide: pid, tid, priority, policy, partition id, sched_flags

(incl APS_SCHED_* critical bit def'd in kermacros.h)

note: partition id and sched_flags only present if APS scheduler module is loaded.

** note: partition id and sched_flags only present if APS scheduler module is loaded.

Call: N/A

Class: _NTO_TRACE_THREAD

Event: _NTO_TRACE_THWAITTHREAD

Fast: pid, tid

Wide: pid, tid, priority, policy, partition id, sched_flags

(incl APS_SCHED_* critical bit def'd in kermacros.h)

note: partition id and sched_flags only present if APS scheduler module is loaded.

** note: partition id and sched_flags only present if APS scheduler module is loaded.

Call: N/A

Class: _NTO_TRACE_THREAD

Event: _NTO_TRACE_THWAITPAGE

Fast: pid, tid

Wide: pid, tid, priority, policy, partition id, sched_flags

(incl APS_SCHED_* critical bit def'd in kermacros.h)

note: partition id and sched_flags only present if APS scheduler module is loaded.

** note: partition id and sched_flags only present if APS scheduler module is loaded.

Call: N/A

Class: _NTO_TRACE_THREAD

Event: _NTO_TRACE_THSIGSPEND

Fast: pid, tid

Wide: pid, tid, priority, policy, partition id, sched_flags

(incl APS_SCHED_* critical bit def'd in kermacros.h)

note: partition id and sched_flags only present if APS scheduler module is loaded.

** note: partition id and sched_flags only present if APS scheduler module is loaded.

Call: N/A

Class: _NTO_TRACE_THREAD

Event: _NTO_TRACE_THSIGWAITINFO
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
 (incl APS_SCHED_* critical bit def'd in kermacros.h)
 note: partition id and sched_flags only present if APS scheduler module
 is loaded.
 ** note: partition id and sched_flags only present if APS scheduler module
 is loaded.
Call: N/A

Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THNANOSLEEP
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
 (incl APS_SCHED_* critical bit def'd in kermacros.h)
 note: partition id and sched_flags only present if APS scheduler module
 is loaded.
 ** note: partition id and sched_flags only present if APS scheduler module
 is loaded.
Call: N/A

Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THMUTEX
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
 (incl APS_SCHED_* critical bit def'd in kermacros.h)
 note: partition id and sched_flags only present if APS scheduler module
 is loaded.
 ** note: partition id and sched_flags only present if APS scheduler module
 is loaded.
Call: N/A

Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THCONDVAR
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
 (incl APS_SCHED_* critical bit def'd in kermacros.h)
 note: partition id and sched_flags only present if APS scheduler module
 is loaded.
 ** note: partition id and sched_flags only present if APS scheduler module
 is loaded.
Call: N/A

Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THJOIN
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
 (incl APS_SCHED_* critical bit def'd in kermacros.h)
 note: partition id and sched_flags only present if APS scheduler module
 is loaded.
 ** note: partition id and sched_flags only present if APS scheduler module
 is loaded.
Call: N/A

Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THINTR
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
 (incl APS_SCHED_* critical bit def'd in kermacros.h)
 note: partition id and sched_flags only present if APS scheduler module
 is loaded.
 ** note: partition id and sched_flags only present if APS scheduler module
 is loaded.
Call: N/A

Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THSEM
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
 (incl APS_SCHED_* critical bit def'd in kermacros.h)
 note: partition id and sched_flags only present if APS scheduler module
 is loaded.
 ** note: partition id and sched_flags only present if APS scheduler module
 is loaded.
Call: N/A

Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THWAITCTX
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags

```
(incl APS_SCHED_* critical bit def'd in kermacros.h)
note: partition id and sched_flags only present if APS scheduler module
is loaded.
** note: partition id and sched_flags only present if APS scheduler module
is loaded.
Call: N/A
```

```
Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THNET_SEND
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
      (incl APS_SCHED_* critical bit def'd in kermacros.h)
      note: partition id and sched_flags only present if APS scheduler module
      is loaded.
      ** note: partition id and sched_flags only present if APS scheduler module
      is loaded.
Call: N/A
```

```
Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THNET_REPLY
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
      (incl APS_SCHED_* critical bit def'd in kermacros.h)
      note: partition id and sched_flags only present if APS scheduler module
      is loaded.
      ** note: partition id and sched_flags only present if APS scheduler module
      is loaded.
Call: N/A
```

```
Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THCREATE
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
      (incl APS_SCHED_* critical bit def'd in kermacros.h)
      note: partition id and sched_flags only present if APS scheduler module
      is loaded.
      ** note: partition id and sched_flags only present if APS scheduler module
      is loaded.
Call: N/A
```

```
Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THDESTROY
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
      (incl APS_SCHED_* critical bit def'd in kermacros.h)
      note: partition id and sched_flags only present if APS scheduler module
      is loaded.
      ** note: partition id and sched_flags only present if APS scheduler module
      is loaded.
Call: N/A
```

```
Class: _NTO_TRACE_THREAD
Event: _NTO_TRACE_THNET_REPLY
Fast: pid, tid
Wide: pid, tid, priority, policy, partition id, sched_flags
      (incl APS_SCHED_* critical bit def'd in kermacros.h)
      note: partition id and sched_flags only present if APS scheduler module
      is loaded.
      ** note: partition id and sched_flags only present if APS scheduler module
      is loaded.
Call: N/A
```

#VThread state changes

```
Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHDEAD
Fast: pid, tid
Wide: pid, tid
Call: N/A
```

```
Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHRUNNING
Fast: pid, tid
Wide: pid, tid
Call: N/A
```

```
Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHREADY
Fast: pid, tid
Wide: pid, tid
```

Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHSTOPPED
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHSEND
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHRECEIVE
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHREPLY
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHSTACK
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHWAITTHREAD
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHWAITPAGE
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHSIGSUSPEND
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHSIGWAITINFO
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHNANOSLEEP
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHMUTEX
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHCONDVAR
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHJOIN
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHINTR
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHSEM
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHWAITCTX
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHNET_SEND
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHNET_REPLY
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHCREATE
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHDESTROY
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_VTHREAD
Event: _NTO_TRACE_VTHNET_REPLY
Fast: pid, tid
Wide: pid, tid
Call: N/A

Class: _NTO_TRACE_COMM
Event: _NTO_TRACE_COMM_SMSG
Fast: rcvid, pid
Wide: rcvid, pid
Call: N/A

Class: _NTO_TRACE_COMM
Event: _NTO_TRACE_COMM_RMSG
Fast: rcvid, pid
Wide: rcvid, pid
Call: N/A

Class: _NTO_TRACE_COMM
Event: _NTO_TRACE_COMM_REPLY
Fast: tid, pid
Wide: tid, pid
Call: N/A

Class: _NTO_TRACE_COMM
Event: _NTO_TRACE_COMM_ERROR
Fast: tid, pid
Wide: tid, pid
Call: N/A

Class: _NTO_TRACE_COMM
Event: _NTO_TRACE_COMM_SPULSE
Fast: scoid, pid
Wide: scoid, pid
Call: N/A

Class: _NTO_TRACE_COMM
Event: _NTO_TRACE_COMM_RPULSE

Fast: scoid, pid
Wide: scoid, pid
Call: N/A

SIGEV_PULSE delivered
Class: _NTO_TRACE_COMM
Event: _NTO_TRACE_COMM_SPULSE_EXE
Fast: scoid, pid
Wide: scoid, pid
Call: N/A

_PULSE_CODE_DISCONNECT pulse delivered
Class: _NTO_TRACE_COMM
Event: _NTO_TRACE_COMM_SPULSE_DIS
Fast: scoid, pid
Wide: scoid, pid
Call: N/A

_PULSE_CODE_COIDDEATH pulse delivered
Class: _NTO_TRACE_COMM
Event: _NTO_TRACE_COMM_SPULSE_DEA
Fast: scoid, pid
Wide: scoid, pid
Call: N/A

_PULSE_CODE_UNBLOCK delivered
Class: _NTO_TRACE_COMM
Event: _NTO_TRACE_COMM_SPULSE_UN
Fast: scoid, pid
Wide: scoid, pid
Call: N/A

_PULSE_CODE_NET_UNBLOCK delivered
Class: _NTO_TRACE_COMM
Event: _NTO_TRACE_COMM_SPULSE_QUN
Fast: scoid, pid
Wide: scoid, pid
Call: N/A

Class: _NTO_TRACE_COMM
Event: _NTO_TRACE_COMM_SIGNAL
Fast: si_signo, si_code
Wide: si_signo, si_code, si_errno, __data.__pad[0-6]
Call: N/A

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_PATHMGR
Fast: pid, tid, pathname
Wide: pid, tid, pathname
Call: Any pathname operation (routed via libc connect)

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_APS_NAME
Fast: partition id, partition name
Wide: partition id, partition name
Call: SchedCtl with sched_aps.h:: SCHED_APS_CREATE_PARTITION

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_APS_BUDGETS
Fast: partition id, new percentage cpu budget, new critical budget ms
Wide: partition id, new percentage cpu budget, new critical budget ms
Call: SchedCtl with sched_aps.h SCHED_APS_CREATE_PARTITION or
SCHED_APS_MODIFY_PARTITION. Also emitted automatically when APS scheduler
clears a critical budget as part of handling a bankruptcy.

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_APS_BNKR
Fast: suspect pid, suspect tid, partition id
Wide: suspect pid, suspect tid, partition id
Call: automatically when a partition exceeds its critical budget.

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_MMAP
Fast: pid, addr (64), len (64), flags
Wide: pid, addr (64), len (64), flags, prot, fd, align (64), offset (64), name
Call: mmap/mmap64

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_MUNMAP
Fast: pid, addr (64), len (64)

Wide: pid, addr (64), len (64)
Call: munmap

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_MAPNAME
Fast: pid, addr (32), len (32), name
Wide: pid, addr (32), len (32), name
Call: dlopen

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_ADDRESS
Fast: addr(32), <null>
Wide: addr(32), <null>
Call: whenever a breakpoint is hit

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_FUNC_ENTER
Fast: thisfn(32), call_site(32)
Wide: thisfn(32), call_site(32)
Call: whenever a function is entered (and it is instrumented)

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_FUNC_EXIT
Fast: thisfn(32), call_site(32)
Wide: thisfn(32), call_site(32)
Call: whenever a function is exited (and it is instrumented)

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_SLOG
Fast: opcode(32), severity(32), message
Wide: opcode(32), severity(32), message
Call: when the kernel wants to note an unusual occurrence

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_DEFRAG_START
Fast: block_size(32)
Wide: block_size(32)
Call: when the memory defragmentation compaction_minimal algorithm is triggered

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_RUNSTATE
Fast: bitset(32) 0x1 - CPU is on/offline, 0x2 - CPU manually requested
on/offline, 0x4 CPU is dynamically offline-able or not, 0x8 system is
in runstate burst mode
Wide: same as above
Call: when the runstate for a CPU changes

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_POWER
Fast: bitset(32), mode(32)
Wide: same as above
Call: idle mode entry/exit, CPU frequency change

The bitset field holds the following:

```
/* Power event flags */
#define _NTO_TRACE_POWER_CPUMASK 0x0000ffffu
#define _NTO_TRACE_POWER_IDLE 0x00010000u
#define _NTO_TRACE_POWER_START 0x00020000u
#define _NTO_TRACE_POWER_IDLE_REACHED 0x00040000u /* for _NTO_TRACE_POWER_IDLE */
#define _NTO_TRACE_POWER_VFS_OVERDRIVE 0x00040000u /* for !_NTO_TRACE_POWER_IDLE */
#define _NTO_TRACE_POWER_VFS_DYNAMIC 0x00080000u /* for !_NTO_TRACE_POWER_IDLE */
#define _NTO_TRACE_POWER_VFS_STEP_UP 0x00100000u /* for !_NTO_TRACE_POWER_IDLE */
```

The bottom 16 bits is the CPU that the mode change applies to. For idle events, this will always be the same as the CPU in the event header. For frequency changes, they may be different (e.g. CPU 0 changes CPU 1's frequency).

If the POWER_IDLE bit is on, this an idle event, if off the event is a frequency change.

If the POWER_START bit is on, it means that we're starting a power event: idle is being entered, we're kicking off a frequency change request. If the bit is off: we're coming out of idle, the frequency change has been completed.

On the idle exit event, the IDLE_REACHED bit indicates that the CPU achieved the requested sleep mode.

For frequency entry events, the VFS_OVERDRIVE bit indicates that the change was being requested by the reception of an overdrive sigevent. The VFS_DYNAMIC bit indicates that the DVFS algorithm is requesting a change due to CPU loading. If neither is on, it's a change due to powerman's list of allowed modes no longer including the frequency that we were previously running at.

The second word of the event is the mode of the power event. For idle events, this is the number given by the "sleep=?" characteristic in the powerman configuration file. For frequency events, this is the value given by the "throughput=?" characteristic (usually the CPU frequency).

Note that for frequency events, the second word for the entry event and exit event may be different. E.g. powerman might request CPU 0 to be run at 300MHz, but CPU 0 & CPU 1 frequencies are tied together and CPU 1 wants to run at 800MHz. In that case the CPU specific code may decide to run CPU 0 at 800MHz instead of the requested 300 and will report the fact in the exit event. Treat the frequency entry as the requested mode and the exit as the actual mode.

Due to interrupt preemptions, you can not be guaranteed that for each entry event there will be a matching exit event and vis versa. E.g. there might be multiple idle entries before an idle exit or vis versa.

Relatively shortly after the start of tracing, powerman will dump a series of frequency exit events giving the current frequencies of each of the CPU's. You should make the assumption that CPU was running in that mode at the start of the trace.

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_IPI
Fast: ipicmd(32), interrupted ip(32)
Wide: same as above
Call: when an inter-processor-interrupt is received

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_PAGEWAIT
Fast: pid(32), tid(32), ip(32), vaddr(32)
Wide: pid(32), tid(32), ip(32), vaddr(32), fault type(32), mmap_flags(32), object_offset(64), object_name(string)
Call: during page fault handling

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_SYS_TIMER
Fast: pid(32), tid(32), timer_id(32), flags(32)
Wide: same as above
Call: on timer expiry
Note - the timer_id will be -1 for timer_timeout expiry

Class: _NTO_TRACE_SYSTEM
Event: _NTO_TRACE_DEFRAG_END
Fast: rc(32), freemem(32), maxblock(32)
Wide: rc(32), freemem(32), maxblock(32)
Call: on completion of the defragmentation process (whether successful or not)

Index

- `__KER_*` events 26
- `__KER_MSG_SENDRV` 90, 91
- `__Ring0()` 27
- `_NTO_TCTL_IO` 56
- `_NTO_TRACE_ALLOCBUFFER` 46
- `_NTO_TRACE_COMM*` 24
- `_NTO_TRACE_COMM*` events 24
- `_NTO_TRACE_CONTROL` 25
- `_NTO_TRACE_CONTROL*` events 25
- `_NTO_TRACE_CONTROLTIME` 68
- `_NTO_TRACE_DEALLOCBUFFER` 46
- `_NTO_TRACE_EMPTY` (not currently used) 24
- `_NTO_TRACE_FIPID` 59
- `_NTO_TRACE_FITID` 59
- `_NTO_TRACE_FLUSHBUFFER` 47
- `_NTO_TRACE_FMPID` 59
- `_NTO_TRACE_FMTID` 59
- `_NTO_TRACE_GETCPU()` 58
- `_NTO_TRACE_GETEVENT_C()` 58
- `_NTO_TRACE_GETEVENT()` 58
- `_NTO_TRACE_INSERTCUSEREVENT` 33, 49
- `_NTO_TRACE_INSERTEVENT` 49
- `_NTO_TRACE_INSERTSUSEREVENT` 33, 49
- `_NTO_TRACE_INSERTUSRSTREVENT` 33, 49
- `_NTO_TRACE_INT` 26
- `_NTO_TRACE_INT_HANDLER_ENTER` 26
- `_NTO_TRACE_INT_HANDLER_EXIT` 26
- `_NTO_TRACE_INT*` events 26
- `_NTO_TRACE_INTENTER` 26
- `_NTO_TRACE_INTEXIT` 26
- `_NTO_TRACE_INTFIRST` 26
- `_NTO_TRACE_INTLAST` 26
- `_NTO_TRACE_KERCALL` 26
- `_NTO_TRACE_KERCALLEENTER` 26, 90
- `_NTO_TRACE_KERCALLEEXIT` 26, 91
- `_NTO_TRACE_KERCALLINT` 26
- `_NTO_TRACE_PROC*` events 30
- `_NTO_TRACE_PROCESS` 30
- `_NTO_TRACE_QUERYEVENTS` 47
- `_NTO_TRACE_SETALLCLASSESFAST` 48
- `_NTO_TRACE_SETALLCLASSESWIDE` 48
- `_NTO_TRACE_SETCLASSFAST` 48
- `_NTO_TRACE_SETCLASSWIDE` 48
- `_NTO_TRACE_SETEVENT_C()` 58
- `_NTO_TRACE_SETEVENT()` 58
- `_NTO_TRACE_SETEVENTFAST` 48
- `_NTO_TRACE_SETEVENTWIDE` 48
- `_NTO_TRACE_SETLINEARMODE` 47
- `_NTO_TRACE_SETRINGMODE` 47
- `_NTO_TRACE_START` 44, 47, 73
- `_NTO_TRACE_STARTNOSTATE` 47, 73
- `_NTO_TRACE_STOP` 43, 44, 47
- `_NTO_TRACE_SYS*` events 30
- `_NTO_TRACE_SYSTEM` 30
- `_NTO_TRACE_TH*` events 31
- `_NTO_TRACE_THREAD` 31

- `_NTO_TRACE_USER` 33
- `_NTO_TRACE_USERFIRST` 33
- `_NTO_TRACE_USERLAST` 33
- `_NTO_TRACE_VTH*` events 34
- `_NTO_TRACE_VTHREAD` 34
- `_PULSE_CODE_COIDDEATH` 24
- `_PULSE_CODE_DISCONNECT` 24
- `_PULSE_CODE_NET_UNBLOCK` 24
- `_PULSE_CODE_UNBLOCK` 24
- `.kev` extension 45

A

- adaptive partitioning 30, 32
 - event data for 32
- Address (IDE event label) 30
- ADDRESS (traceprinter event label) 30
- APS Bankruptcy (IDE event label) 30
- APS Budgets (IDE event label) 30
- APS Name (IDE event label) 30
- APS_BANKRUPTCY (traceprinter event label) 30
- APS_NAME (traceprinter event label) 30
- APS_NEW_BUDGET (traceprinter event label) 30

B

- bankruptcy (adaptive partitions) 30
- Buffer (IDE event label) 25
- BUFFER (traceprinter event label) 25
- buffers, kernel 14, 25, 37, 43, 46, 47, 67
 - circular linked list 37
 - events concerning 25
 - flushing 47
 - managing 43, 46
 - specifications 37

C

- ChannelConnectAttr() 27
- ChannelCreate() 27
- ChannelDestroy() 27
- circular linked list 37
- classes 24, 25, 26, 30, 31, 33, 34, 48, 58
 - `_NTO_TRACE_EMPTY` (not currently used) 24
 - Communication 24
 - Control 25
 - Interrupt 26
 - Kernel-call 26
 - Process 30
 - pseudo 26
 - `_NTO_TRACE_INT` 26
 - `_NTO_TRACE_KERCALL` 26
 - setting fast or wide mode for 48
 - System 30
 - Thread 31

- classes (*continued*)
 - type, extracting from the header 58
 - User 33
 - Virtual thread 34
- ClockAdjust() 27
- ClockId() 27
- ClockPeriod() 27
- clocks, importance of synchronizing on multicore systems 42
- ClockTime() 27
- combine events 22, 33, 67
 - user-defined 33
- communication, events concerning 24
- Compaction (IDE event label) 30
- COMPACTION (traceprinter event label) 30
- Condvar (IDE event label) 31
- configuring 41
 - data capture 41
 - instrumented kernel 41
- ConnectAttach() 27
- ConnectClientInfo() 27
- ConnectDetach() 27
- ConnectFlags() 27
- ConnectServerInfo() 27
- control of tracing, events concerning 25
- CPU index, extracting from an event 58
- Create Process (IDE event label) 30
- Create Process Name (IDE event label) 30
- Create Thread (IDE event label) 31
- Create VThread (IDE event label) 34
- critical budgets, exceeding (adaptive partitions) 30
- custom events 33

D

- daemon mode 44
- data capture 14, 41
- data interpretation 14, 63, 66, 67, 68
- data reduction 51
- Dead (IDE event label) 31
- Death Pulse (IDE event label) 24
- defragmentation of memory 30
- Destroy Process (IDE event label) 30
- Destroy Thread (IDE event label) 31
- Destroy VThread (IDE event label) 34
- Disconnect Pulse (IDE event label) 24
- dlopen() 30
- dynamic rules filter 51, 56

E

- Enter (IDE event label) 27
- Entry (IDE event label) 26
- Error (IDE event label) 24
- error codes, included in trace event data for kernel calls 90
- event_data_t 57
- events 19, 22, 24, 25, 26, 30, 31, 33, 34, 47, 48, 49, 56, 57, 58, 59, 67, 89
 - classes 24, 25, 26, 30, 31, 33, 34
 - Communication 24
 - Control 25
 - Interrupt 26
 - Kernel calls 26

- events (*continued*)
 - classes (*continued*)
 - Process 30
 - System 30
 - Thread 31
 - User 33
 - Virtual thread 34
 - combine 22
 - data for 89
 - getting the number of in a trace buffer 47
 - handlers 56, 57, 59
 - adding 56
 - functions safe to use within 57
 - removing 59
 - inserting 49
 - interpreting 67
 - setting fast or wide mode for 48
 - simple 22
 - type, extracting from the header 58
- examples of tracing 69
- Exit (IDE event label) 26, 27
- extended daemon mode 44

F

- fast mode 23, 44, 48
 - setting with TraceEvent() 48
 - setting with tracelogger 44
- filters 45, 51
- FUNC_ENTER (traceprinter event label) 30
- FUNC_EXIT (traceprinter event label) 30
- Function Enter (IDE event label) 30
- Function Exit (IDE event label) 30
- functions 30, 57
 - instrumented for profiling 30
 - safe to use in an event handler 57

H

- Handler Entry (IDE event label) 26
- Handler Exit (IDE event label) 26

I

- I/O 56
 - privileges, requesting 56
- I/O privileges 56
- initial state information, suppressing 73
- instrumented (for profiling) functions 30
- instrumented kernel 14, 37, 41, 68
 - configuring 41
- Int (IDE event label) 27
- INT_CALL (traceprinter event label) 27
- INT_ENTR (traceprinter event label) 26
- INT_EXIT (traceprinter event label) 26
- INT_HANDLER_ENTR (traceprinter event label) 26
- INT_HANDLER_EXIT (traceprinter event label) 26
- Integrated Development Environment (IDE) 16, 24, 25, 26, 27, 30, 31, 33, 34, 41, 45, 61
 - event labels 24, 25, 26, 27, 30, 31, 33, 34
 - Address 30

Integrated Development Environment (IDE) (*continued*)event labels (*continued*)

APS Bankruptcy 30
 APS Budgets 30
 APS Name 30
 Buffer 25
 Compaction 30
 Condvar 31
 Create Process 30
 Create Process Name 30
 Create Thread 31
 Create VThread 34
 Dead 31
 Death Pulse 24
 Destroy Process 30
 Destroy Thread 31
 Destroy VThread 34
 Disconnect Pulse 24
 Enter 27
 Entry 26
 Error 24
 Exit 26, 27
 Function Enter 30
 Function Exit 30
 Handler Entry 26
 Handler Exit 26
 Int 27
 Interrupt 31
 Join 31
 MMap 30
 MMap Name 30
 MMUnmap 30
 Mutex 31
 NanoSleep 31
 NetReply 31
 NetSend 31
 Path Manager 30
 QNet Unblock Pulse 24
 Ready 31
 Receive 31
 Receive Message 24
 Receive Pulse 24
 Reply 24, 31
 Running 31
 Semaphore 31
 Send 31
 Send Message 24
 Send Pulse 24
 Sigevent Pulse 24
 Signal 24
 SigSuspend 31
 SigWaitInfo 31
 Stack 31
 Stopped 31
 System Log 30
 Thread Name 30
 Time 25
 Unblock Pulse 24
 User Event 33
 VCondvar 34
 VDead 34
 VInterrupt 34

Integrated Development Environment (IDE) (*continued*)event labels (*continued*)

VJoin 34
 VMutex 34
 VNanoSleep 34
 VNetReply 34
 VNetSend 34
 VReady 34
 VReceive 34
 VReply 34
 VRunning 34
 VSemaphore 34
 VSend 34
 VSigSuspend 34
 VSigWaitInfo 34
 VStack 34
 VStopped 34
 VWaitCtx 34
 VWaitPage 34
 VWaitThread 34
 WaitCtx 31
 WaitPage 31
 WaitThread 31

recognizes the .kev extension 45

interlacing 67

Interrupt (IDE event label) 31

InterruptAttach() 27

InterruptDetach() 27

InterruptMask() 27

interrupts, events concerning 26

InterruptUnmask() 27

InterruptWait() 27

J

Join (IDE event label) 31

K

KER_CALL (traceprinter event label) 27

KER_EXIT (traceprinter event label) 27

kernel 56

calls 56

ThreadCtl(), ThreadCtl_r() 56

kernel buffers 14, 25, 37, 67

circular linked list 37

events concerning 25

specifications 37

kernel calls 26, 90

events concerning 26

trace event data on failure 90

L

library 66

linear mode 43

log 45

M

MAPNAME (traceprinter event label) 30

- memory defragmentation 30
- memory, tracelogger output in shared 45
- messages 24
- MMap (IDE event label) 30
- MMap (traceprinter event label) 30
- MMap Name (IDE event label) 30
- mmap(), mmap64() 30
- MMUnmap (IDE event label) 30
- MSG_ERROR (traceprinter event label) 24
- MsgCurrent() 27
- MsgDeliverEvent() 27
- MsgError() 24, 27
- MsgInfo() 27
- MsgKeyData() 27
- MsgRead() 27
- MsgReadIov() 27
- MsgReadv() 27
- MsgReceive() 27
- MsgReceivePulse() 27
- MsgReceivePulsev() 27
- MsgReceivev() 27
- MsgReply() 27
- MsgReplyv() 27
- MsgSend() 27, 90
- MsgSendnc() 27
- MsgSendPulse() 27
- MsgSendv() 27, 90
- MsgSendvnc() 27
- MsgSendvs() 27, 90
- MsgSendvsnc() 27
- MsgVerifyEvent() 27
- MsgWrite() 27
- MsgWritev() 27
- multicore systems 42, 58
 - extracting the CPU index from an event 58
 - importance of synchronizing clocks 42
- MUNMAP (traceprinter event label) 30
- munmap() 30
- Mutex (IDE event label) 31

N

- NanoSleep (IDE event label) 31
- NetCred() 27
- NetInfoScoid() 27
- NetReply (IDE event label) 31
- NetSend (IDE event label) 31
- NetSignalKill() 27
- NetUnblock() 27
- NetVtid() 27
- normal mode 44

O

- open() 30

P

- partitions, adaptive 30
- path manager 30
- Path Manager (IDE event label) 30

- PATHMGR_OPEN (traceprinter event label) 30
- post-processing filter 51, 60
- PROCCREATE (traceprinter event label) 30
- PROCCREATE_NAME (traceprinter event label) 30
- PROCDESTROY (traceprinter event label) 30
- processes 56
 - I/O privileges, requesting 56
- processes, events concerning 30
- PROCMGR_AID_IO 56
- PROCMGR_AID_TRACE 46, 56
- PROCTHREAD_NAME (traceprinter event label) 30
- profiling, functions instrumented for 30
- pseudo-classes 26
 - _NTO_TRACE_INT 26
 - _NTO_TRACE_KERCALL 26
- pulses 24

Q

- qconn 41
- Qnet 34
- QNet Unblock Pulse (IDE event label) 24

R

- Ready (IDE event label) 31
- REC_MESSAGE (traceprinter event label) 24
- REC_PULSE (traceprinter event label) 24
- Receive (IDE event label) 31
- Receive Message (IDE event label) 24
- Receive Pulse (IDE event label) 24
- Reply (IDE event label) 24, 31
- REPLY_MESSAGE (traceprinter event label) 24
- ring mode 43
- Running (IDE event label) 31

S

- SAT 11, 12, 14
- SCHED_APS_CREATE_PARTITION 30
- SCHED_APS_MODIFY_PARTITION 30
- SchedGet() 27
- SchedInfo() 27
- SchedSet() 27
- SchedYield() 27
- Semaphore (IDE event label) 31
- Send (IDE event label) 31
- Send Message (IDE event label) 24
- Send Pulse (IDE event label) 24
- shared memory, tracelogger output in 45
- SIGEV_PULSE 24
- Sigevent Pulse (IDE event label) 24
- SIGINT 43
- Signal (IDE event label) 24
- SIGNAL (traceprinter event label) 24
- SignalAction() 27
- SignalKill() 27
- SignalProcmask() 27
- SignalReturn() 27
- SignalSuspend() 27
- SignalWaitInfo() 27

- SigSuspend (IDE event label) 31
- SigWaitInfo (IDE event label) 31
- simple events 22, 33, 67
 - user-defined 33
- SLOG (traceprinter event label) 30
- SND_MESSAGE (traceprinter event label) 24
- SND_PULSE (traceprinter event label) 24
- SND_PULSE_DEA (traceprinter event label) 24
- SND_PULSE_DIS (traceprinter event label) 24
- SND_PULSE_EXE (traceprinter event label) 24
- SND_PULSE_QUN (traceprinter event label) 24
- SND_PULSE_UN (traceprinter event label) 24
- Stack (IDE event label) 31
- state information, suppressing initial 73
- states 31, 34
 - threads 31
 - virtual threads 34
- static rules filter 51, 53
- Stopped (IDE event label) 31
- string events, user-defined 33
- structures 67
- SyncCondvarSignal() 27
- SyncCondvarWait() 27
- SyncCtl() 27
- SyncDestroy() 27
- SyncMutexLock() 27
- SyncMutexRevive() 27
- SyncMutexUnlock() 27
- SyncSemPost() 27
- SyncSemWait() 27
- SyncTypeCreate() 27
- system information, suppressing initial 73
- system log 30
- System Log (IDE event label) 30
- system, events concerning 30

T

- TDP (Transparent Distributed Processing) 34
- Technical support 10
- THCONDVAR (traceprinter event label) 31
- THCREATE (traceprinter event label) 31
- THDEAD (traceprinter event label) 31
- THDESTROY (traceprinter event label) 31
- THINTR (traceprinter event label) 31
- THJOIN (traceprinter event label) 31
- THMUTEX (traceprinter event label) 31
- THNANOSLEEP (traceprinter event label) 31
- THNET_REPLY (traceprinter event label) 31
- THNET_SEND (traceprinter event label) 31
- Thread Name (IDE event label) 30
- ThreadCancel() 27
- ThreadCreate() 27
- ThreadCtl() 27, 56
- ThreadCtl(), ThreadCtl_r() 56
- ThreadDestroy() 27
- ThreadDetach() 27
- ThreadJoin() 27
- threads, events concerning 20, 31, 34
- THREADY (traceprinter event label) 31
- THRECEIVE (traceprinter event label) 31
- THREPLY (traceprinter event label) 31
- THRUNNING (traceprinter event label) 31
- THSEM (traceprinter event label) 31
- THSEND (traceprinter event label) 31
- THSIGSUSPEND (traceprinter event label) 31
- THSIGWAITINFO (traceprinter event label) 31
- THSTACK (traceprinter event label) 31
- THSTOPPED (traceprinter event label) 31
- THWAITCTX (traceprinter event label) 31
- THWAITPAGE (traceprinter event label) 31
- THWAITTHREAD (traceprinter event label) 31
- Time (IDE event label) 25
- TIME (traceprinter event label) 25
- TimerAlarm() 27
- TimerCreate() 27
- TimerDestroy() 27
- TimerInfo() 27
- TimerSettime() 27
- TimerTimeout() 27
- timestamps 25, 42, 68
 - importance of synchronizing on multicore systems 42
- trace_func_enter() 31, 49
- trace_func_exit() 31, 49
- trace_here() 31, 49
- trace_logb() 33, 49
- trace_logbc() 49
- trace_logf() 33, 49
- trace_logi() 33, 49
- trace_nlogf() 33, 49
- trace_vnlogf() 33, 49
- traceevent_t 67
- TraceEvent() 27, 33, 41, 43, 44, 46, 47, 48, 49, 69
 - controlling tracing with 41, 44, 46
 - creating user events 33
 - examples of use 69
 - inserting events with 49
 - managing trace buffers 46
 - modes of operation 47
 - ring mode 43
 - wide and fast modes 48
- tracelog 45
- tracelogger 41, 43, 44, 45, 69
 - .kev extension 45
 - controlling tracing with 41
 - directing the output from 45
 - examples of use 69
 - filtering 45
 - managing trace buffers 43
 - modes 43
 - running 43
 - wide and fast modes 44
- traceparser_cs_range() 66
- traceparser_cs() 66
- traceparser_debug() 66
- traceparser_destroy() 66
- traceparser_get_info() 66
- traceparser_init() 66
- traceparser() 66
- traceprinter 16, 24, 25, 26, 27, 30, 31, 33, 34, 60, 63, 66
 - as the basis for your own parser 66
 - event labels 24, 25, 26, 27, 30, 31, 33, 34
 - ADDRESS 30
 - APS_BANKRUPTCY 30

traceprinter (*continued*)event labels (*continued*)

- APS_NAME 30
- APS_NEW_BUDGET 30
- BUFFER 25
- COMPACTION 30
- FUNC_ENTER 30
- FUNC_EXIT 30
- INT_CALL 27
- INT_ENTR 26
- INT_EXIT 26
- INT_HANDLER_ENTR 26
- INT_HANDLER_EXIT 26
- KER_CALL 27
- KER_EXIT 27
- MAPNAME 30
- MMAP 30
- MSG_ERROR 24
- MUNMAP 30
- PATHMGR_OPEN 30
- PROCCREATE 30
- PROCCREATE_NAME 30
- PROCDESTROY 30
- PROCTHREAD_NAME 30
- REC_MESSAGE 24
- REC_PULSE 24
- REPLY_MESSAGE 24
- SIGNAL 24
- SLOG 30
- SND_MESSAGE 24
- SND_PULSE 24
- SND_PULSE_DEA 24
- SND_PULSE_DIS 24
- SND_PULSE_EXE 24
- SND_PULSE_QUN 24
- SND_PULSE_UN 24
- THCONDVAR 31
- THCREATE 31
- THDEAD 31
- THDESTROY 31
- THINTR 31
- THJOIN 31
- THMUTEX 31
- THNANOSLEEP 31
- THNET_REPLY 31
- THNET_SEND 31
- THREADY 31
- THRECEIVE 31
- THREPLY 31
- THRUNNING 31
- THSEM 31
- THSEND 31
- THSIGSUSPEND 31
- THSIGWAITINFO 31
- THSTACK 31
- THSTOPPED 31
- THWAITCTX 31
- THWAITPAGE 31
- THWAITTHREAD 31
- TIME 25
- USREVENT 33
- VTHCONDVAR 34

traceprinter (*continued*)event labels (*continued*)

- VTHCREATE 34
- VTHDEAD 34
- VTHDESTROY 34
- VTHINTR 34
- VTHJOIN 34
- VTHMUTEX 34
- VTHNANOSLEEP 34
- VTHNET_REPLY 34
- VTHNET_SEND 34
- VTHREADY 34
- VTHRECEIVE 34
- VTHREPLY 34
- VTHRUNNING 34
- VTHSEM 34
- VTHSEND 34
- VTHSIGSUSPEND 34
- VTHSIGWAITINFO 34
- VTHSTACK 34
- VTHSTOPPED 34
- VTHWAITCTX 34
- VTHWAITPAGE 34
- VTHWAITTHREAD 34
- interpreting the output 63
- post-processing filter 60
- tracing 25, 41
 - control of, events concerning 25
 - controlling 41
- Transparent Distributed Processing (TDP) 34
- tutorials 69
- Typographical conventions 8

U

- Unblock Pulse (IDE event label) 24
- User Event (IDE event label) 33
- user-defined events 33
- USREVENT (traceprinter event label) 33

V

- VCondvar (IDE event label) 34
- VDead (IDE event label) 34
- VInterrupt (IDE event label) 34
- virtual threads, events concerning 34
- VJoin (IDE event label) 34
- VMutex (IDE event label) 34
- VNanosleep (IDE event label) 34
- VNetReply (IDE event label) 34
- VNetSend (IDE event label) 34
- VReady (IDE event label) 34
- VReceive (IDE event label) 34
- VReply (IDE event label) 34
- VRunning (IDE event label) 34
- VSemaphore (IDE event label) 34
- VSend (IDE event label) 34
- VSigSuspend (IDE event label) 34
- VSigWaitInfo (IDE event label) 34
- VStack (IDE event label) 34
- VStopped (IDE event label) 34
- VTHCONDVAR (traceprinter event label) 34

VTHCREATE (traceprinter event label) 34
VTHDEAD (traceprinter event label) 34
VTHDESTROY (traceprinter event label) 34
VTHINTR (traceprinter event label) 34
VTHJOIN (traceprinter event label) 34
VTHMUTEX (traceprinter event label) 34
VTHNANOSLEEP (traceprinter event label) 34
VTHNET_REPLY (traceprinter event label) 34
VTHNET_SEND (traceprinter event label) 34
VTHREADY (traceprinter event label) 34
VTHRECEIVE (traceprinter event label) 34
VTHREPLY (traceprinter event label) 34
VTHRUNNING (traceprinter event label) 34
VTHSEM (traceprinter event label) 34
VTHSEND (traceprinter event label) 34
VTHSIGSUSPEND (traceprinter event label) 34
VTHSIGWAITINFO (traceprinter event label) 34

VTHSTACK (traceprinter event label) 34
VTHSTOPPED (traceprinter event label) 34
VTHWAITCTX (traceprinter event label) 34
VTHWAITPAGE (traceprinter event label) 34
VTHWAITTHREAD (traceprinter event label) 34
VWaitCtx (IDE event label) 34
VWaitPage (IDE event label) 34
VWaitThread (IDE event label) 34

W

WaitCtx (IDE event label) 31
WaitPage (IDE event label) 31
WaitThread (IDE event label) 31
wide mode 23, 44, 48
 setting with TraceEvent() 48
 setting with tracelogger 44

