

# **SystemTap Tapset Reference Manual**

**SystemTap**

---

# **SystemTap Tapset Reference Manual**

by SystemTap

Copyright © 2008-2009 Red Hat, Inc. and others

This documentation is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License version 2 as published by the Free Software Foundation.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

For more details see the file COPYING in the source distribution of Linux.

---

# Table of Contents

1. Introduction .....	1
Tapset Name Format .....	1
2. Context Functions .....	2
print_regs .....	3
execname .....	4
pid .....	5
tid .....	6
ppid .....	7
pexecname .....	8
gid .....	9
egid .....	10
uid .....	11
euid .....	12
cpu .....	13
pp .....	14
registers_valid .....	15
user_mode .....	16
is_return .....	17
target .....	18
stack_size .....	19
stack_used .....	20
stack_unused .....	21
uaddr .....	22
print_stack .....	23
probefunc .....	24
probemod .....	25
modname .....	26
symname .....	27
symdata .....	28
print_backtrace .....	29
backtrace .....	30
caller .....	31
caller_addr .....	32
3. Timestamp Functions .....	33
get_cycles .....	34
4. Memory Tapset .....	35
vm_fault_contains .....	36
vm.pagefault .....	37
vm.pagefault.return .....	38
addr_to_node .....	39
vm.write_shared .....	40
vm.write_shared_copy .....	41
vm.mmap .....	42
vm.munmap .....	43
vm.brk .....	44
vm.oom_kill .....	45
5. IO Scheduler Tapset .....	46
ioscheduler.elv_next_request .....	47
ioscheduler.elv_next_request.return .....	48
ioscheduler.elv_add_request .....	49
ioscheduler.elv_completed_request .....	50
6. SCSI Tapset .....	51
scsi.ioentry .....	52
scsi.iodispatching .....	53
scsi.iodone .....	54
scsi.iocompleted .....	55

SystemTap	
Tapset	
Ref-	
e-	
nce	
Man-	
ual	
7. Networking Tapset	56
netdev.receive	57
netdev.transmit	58
tcp.sendmsg	59
tcp.sendmsg.return	60
tcp.recvmsg	61
tcp.recvmsg.return	62
tcp.disconnect	63
tcp.disconnect.return	64
tcp.setsockopt	65
tcp.setsockopt.return	66
tcp.receive	67
udp.sendmsg	68
udp.sendmsg.return	69
udp.recvmsg	70
udp.recvmsg.return	71
udp.disconnect	72
udp.disconnect.return	73
ip_ntop	74
8. Socket Tapset	75
socket.send	76
socket.receive	77
socket.sendmsg	78
socket.sendmsg.return	79
socket.recvmsg	80
socket.recvmsg.return	81
socket.aio_write	82
socket.aio_write.return	83
socket.aio_read	84
socket.aio_read.return	85
socket.writev	86
socket.writev.return	87
socket.ready	88
socket.readv.return	89
socket.create	90
socket.create.return	91
socket.close	92
socket.close.return	93
sock_prot_num2str	94
sock_prot_str2num	95
sock_fam_num2str	96
sock_fam_str2num	97
sock_state_num2str	98
sock_state_str2num	99
9. Kernel Process Tapset	100
kprocess.create	101
kprocess.start	102
kprocess.exec	103
kprocess.exec_complete	104
kprocess.exit	105
kprocess.release	106
10. Signal Tapset	107
signal.send	108
signal.send.return	109
signal.checkperm	110
signal.checkperm.return	111
signal.wakeup	112
signal.check_ignored	113

signal.check_ignored.return .....	114
signal.force_segv .....	115
signal.force_segv.return .....	116
signal.syskill .....	117
signal.syskill.return .....	118
signal.sys_tkill .....	119
signal.systkill.return .....	120
signal.sys_tckill .....	121
signal.sys_tkill.return .....	122
signal.send_sig_queue .....	123
signal.send_sig_queue.return .....	124
signal.pending .....	125
signal.pending.return .....	126
signal.handle .....	127
signal.handle.return .....	128
signal.do_action .....	129
signal.do_action.return .....	130
signal.procmask .....	131
signal.flush .....	132

---

# Chapter 1. Introduction

SystemTap provides free software (GPL) infrastructure to simplify the gathering of information about the running Linux system. This assists diagnosis of a performance or functional problem. SystemTap eliminates the need for the developer to go through the tedious and disruptive instrument, recompile, install, and reboot sequence that may be otherwise required to collect data.

SystemTap provides a simple command line interface and scripting language for writing instrumentation for a live running kernel. The instrumentation makes extensive use of the probe points and functions provided in the *tapset* library. This document describes the various probe points and functions.

## Tapset Name Format

In this guide, tapset definitions appear in the following format:

```
name:return (parameters)
definition
```

The *return* field specifies what data type the tapset extracts and returns from the kernel during a probe (and thus, returns). Tapsets use 2 data types for *return*: *long* (tapset extracts and returns an integer) and *string* (tapset extracts and returns a string).

In some cases, tapsets do not have a *return* value. This simply means that the tapset does not extract anything from the kernel. This is common among asynchronous events such as timers, exit functions, and print functions.

---

# **Chapter 2. Context Functions**

The context functions provide additional information about where an event occurred. These functions can provide information such as a backtrace to where the event occurred and the current register values for the processor.

## Name

print\_regs -- Print a register dump.

## Synopsis

```
print_regs()
```

## Arguments

None

## Name

execname -- Returns the execname of a target process (or group of processes).

## Synopsis

```
execname:string()
```

## Arguments

None

## Name

pid -- Returns the ID of a target process.

## Synopsis

```
pid:long()
```

## Arguments

None

## Name

tid -- Returns the thread ID of a target process.

## Synopsis

```
tid:long()
```

## Arguments

None

## Name

ppid -- Returns the process ID of a target process's parent process.

## Synopsis

```
ppid:long()
```

## Arguments

None

## Name

pexecname -- Returns the execname of a target process's parent process.

## Synopsis

```
pexecname:string()
```

## Arguments

None

## Name

gid -- Returns the group ID of a target process.

## Synopsis

```
gid:long()
```

## Arguments

None

## Name

egid -- Returns the effective gid of a target process.

## Synopsis

```
egid:long()
```

## Arguments

None

## Name

uid -- Returns the user ID of a target process.

## Synopsis

```
uid:long()
```

## Arguments

None

## Name

euid -- Return the effective uid of a target process.

## Synopsis

```
euid:long()
```

## Arguments

None

## Name

cpu -- Returns the current cpu number.

## Synopsis

```
cpu:long()
```

## Arguments

None

## Name

pp -- Return the probe point associated with the currently running probe handler,

## Synopsis

```
pp:string()
```

## Arguments

None

## Description

including alias and wildcard expansion effects

## Context

The current probe point.

## Name

registers\_valid -- Determines validity of <command>register</command> and <command>u\_register</command> in current context.

## Synopsis

```
registers_valid:long()
```

## Arguments

None

## Description

Return 1 if register and u\_register can be used in the current context, or 0 otherwise. For example, <command>registers\_valid</command> returns 0 when called from a begin or end probe.

## Name

user\_mode -- Determines if probe point occurs in user-mode.

## Synopsis

```
user_mode:long()
```

## Arguments

None

## Description

Return 1 if the probe point occurred in user-mode.

## Name

is\_return -- Determines if probe point is a return probe.

## Synopsis

```
is_return:long()
```

## Arguments

None

## Description

Return 1 if the probe point is a return probe. <emphasis>Deprecated.</emphasis>

## Name

target -- Return the process ID of the target process.

## Synopsis

```
target:long()
```

## Arguments

None

## Name

stack\_size -- Return the size of the kernel stack.

## Synopsis

```
stack_size:long()
```

## Arguments

None

## Name

stack\_used -- Returns the amount of kernel stack used.

## Synopsis

```
stack_used:long()
```

## Arguments

None

## Description

Determines how many bytes are currently used in the kernel stack.

## Name

stack\_unused -- Returns the amount of kernel stack currently available.

## Synopsis

```
stack_unused:long()
```

## Arguments

None

## Description

Determines how many bytes are currently available in the kernel stack.

## Name

uaddr -- User space address of current running task. EXPERIMENTAL.

## Synopsis

```
uaddr:long()
```

## Arguments

None

## Description

Returns the address in userspace that the current task was at when the probe occurred. When the current running task isn't a user space thread, or the address cannot be found, zero is returned. Can be used to see where the current task is combined with usymname or symdata. Often the task will be in the VDSO where it entered the kernel. FIXME - need VDSO tracking support #10080.

## Name

print\_stack -- Print out stack from string.

## Synopsis

```
print_stack(stk:string)
```

## Arguments

*stk*    *stk*

String with list of hexidecimal addresses.

## Description

Perform a symbolic lookup of the addresses in the given string, which is assumed to be the result of a prior call to <command>backtrace</command>.

Print one line per address, including the address, the name of the function containing the address, and an estimate of its position within that function. Return nothing.

## Name

probefunc -- Return the probe point's function name, if known.

## Synopsis

```
probefunc:string()
```

## Arguments

None

## Name

probemod -- Return the probe point's module name, if known.

## Synopsis

```
probemod:string()
```

## Arguments

None

## Name

modname -- Return the kernel module name loaded at the address.

## Synopsis

```
modname:string (addr:long)
```

## Arguments

*addr*      *addr*

The address.

## Description

Returns the module name associated with the given address if known. If not known it will return the string “<unknown>”. If the address was not in a kernel module, but in the kernel itself, then the string “kernel” will be returned.

## Name

symname -- Return the symbol associated with the given address.

## Synopsis

```
symname:string (addr:long)
```

## Arguments

*addr*      *addr*  
The address to translate.

## Description

Returns the (function) symbol name associated with the given address if known. If not known it will return the hex string representation of addr.

## Name

symdata -- Return the symbol and module offset for the address.

## Synopsis

```
symdata:string(addr:long)
```

## Arguments

*addr*      *addr*  
The address to translate.

## Description

Returns the (function) symbol name associated with the given address if known, plus the module name (between brackets) and the offset inside the module, plus the size of the symbol function. If any element is not known it will be omitted and if the symbol name is unknown it will return the hex string for the given address.

## Name

print\_backtrace -- Print stack back trace

## Synopsis

```
print_backtrace()
```

## Arguments

None

## Description

Equivalent to <command>print\_stack(backtrace)</command>, except that deeper stack nesting may be supported. Return nothing.

## Name

backtrace -- Hex backtrace of current stack

## Synopsis

```
backtrace:string()
```

## Arguments

None

## Description

Return a string of hex addresses that are a backtrace of the stack. Output may be truncated as per maximum string length.

## Name

caller -- Return name and address of calling function

## Synopsis

```
caller:string()
```

## Arguments

None

## Description

Return the address and name of the calling function. <emphasis>Works only for return probes at this time.</emphasis>

## Name

caller\_addr -- Return caller address

## Synopsis

```
caller_addr:long()
```

## Arguments

None

## Description

Return the address of the calling function. <emphasis> Works only for return probes at this time.</emphasis>

---

# Chapter 3. Timestamp Functions

Each timestamp function returns a value to indicate when a function is executed. These returned values can then be used to indicate when an event occurred, provide an ordering for events, or compute the amount of time elapsed between two time stamps.

## Name

get\_cycles -- Processor cycle count.

## Synopsis

```
get_cycles:long()
```

## Arguments

None

## Description

Return the processor cycle counter value, or 0 if unavailable.

---

# **Chapter 4. Memory Tapset**

This family of probe points is used to probe memory-related events. It contains the following probe points:

## Name

vm\_fault\_contains -- Test return value for page fault reason

## Synopsis

```
vm_fault_contains:long(value:long,test:long)
```

## Arguments

*value*      *value*

The fault\_type returned by vm.page\_fault.return

*test*      *test*

The type of fault to test for (VM\_FAULT\_OOM or similar)

## Name

vm.pagefault -- Records that a page fault occurred.

## Synopsis

vm.pagefault

## Values

<i>write_access</i>	<i>write_access</i>	
		Indicates whether this was a write or read access; <command>1</command> indicates a write, while <command>0</command> indicates a read.
<i>address</i>	<i>address</i>	
		The address of the faulting memory access; i.e. the address that caused the page fault.

## Context

The process which triggered the fault

## Name

vm.pagefault.return -- Indicates what type of fault occurred.

## Synopsis

```
vm.pagefault.return
```

## Values

*fault\_type*      *fault\_type*

Returns either <command>0</command> (VM\_FAULT\_OOM) for out of memory faults, <command>2</command> (VM\_FAULT\_MINOR) for minor faults, <command>3</command> (VM\_FAULT\_MAJOR) for major faults, or <command>1</command> (VM\_FAULT\_SIGBUS) if the fault was neither OOM, minor fault, nor major fault.

## Name

addr\_to\_node -- Returns which node a given address belongs to within a NUMA system.

## Synopsis

```
addr_to_node:long(addr:long)
```

## Arguments

*addr*      *addr*  
The address of the faulting memory access.

## Name

vm.write\_shared -- Attempts at writing to a shared page.

## Synopsis

vm.write\_shared

## Values

*address*      *address*  
The address of the shared write.

## Context

The context is the process attempting the write.

## Description

Fires when a process attempts to write to a shared page. If a copy is necessary, this will be followed by a <command>vm.write\_shared\_copy</command>.

## Name

vm.write\_shared\_copy -- Page copy for shared page write.

## Synopsis

```
vm.write_shared_copy
```

## Values

*zero*      *zero*

Boolean indicating whether it is a zero page (can do a clear instead of a copy).

*address*      *address*

The address of the shared write.

## Context

The process attempting the write.

## Description

Fires when a write to a shared page requires a page copy. This is always preceded by a <command>vm.shared\_write</command>.

## Name

vm.mmap -- Fires when an <command>mmap</command> is requested.

## Synopsis

vm.mmap

## Values

*length*      *length*  
The length of the memory segment

*address*      *address*  
The requested address

## Context

The process calling <command>mmap</command>.

## Name

vm.munmap -- Fires when an <command>munmap</command> is requested.

## Synopsis

vm.munmap

## Values

*length*      *length*  
The length of the memory segment

*address*      *address*  
The requested address

## Context

The process calling <command>munmap</command>.

## Name

vm.brk -- Fires when a <command>brk</command> is requested (i.e. the heap will be resized).

## Synopsis

vm.brk

## Values

*length*      *length*  
The length of the memory segment

*address*      *address*  
The requested address

## Context

The process calling <command>brk</command>.

## Name

vm.oom\_kill -- Fires when a thread is selected for termination by the OOM killer.

## Synopsis

vm.oom\_kill

## Values

*task*    *task*  
The task being killed

## Context

The process that tried to consume excessive memory, and thus triggered the OOM. <remark>(is this correct?)</remark>

---

# **Chapter 5. IO Scheduler Tapset**

This family of probe points is used to probe IO scheduler activities. It contains the following probe points:

## Name

ioscheduler.elv\_next\_request -- Fires when a request is retrieved from the request queue

## Synopsis

```
ioscheduler.elv_next_request
```

## Values

<i>elevator_name</i>	<i>elevator_name</i>
----------------------	----------------------

The type of I/O elevator currently enabled

## Name

ioscheduler.elv\_next\_request.return -- Fires when a request retrieval issues a return signal

## Synopsis

```
ioscheduler.elv_next_request.return
```

## Values

*req\_flags*            *req\_flags*  
Request flags

*req*                *req*  
Address of the request

*disk\_major*        *disk\_major*  
Disk major number of the request

*disk\_minor*        *disk\_minor*  
Disk minor number of the request

## Name

ioscheduler.elv\_add\_request -- A request was added to the request queue

## Synopsis

ioscheduler.elv\_add\_request

## Values

<i>req_flags</i>	<i>req_flags</i>
Request flags	
<i>req</i>	<i>req</i>
Address of the request	
<i>disk_major</i>	<i>disk_major</i>
Disk major number of the request	
<i>elevator_name</i>	<i>elevator_name</i>
The type of I/O elevator currently enabled	
<i>disk_minor</i>	<i>disk_minor</i>
Disk minor number of the request	

## Name

ioscheduler.elv\_completed\_request -- Fires when a request is completed

## Synopsis

```
ioscheduler.elv_completed_request
```

## Values

<i>req_flags</i>	<i>req_flags</i>	Request flags
<i>req</i>	<i>req</i>	Address of the request
<i>disk_major</i>	<i>disk_major</i>	Disk major number of the request
<i>elevator_name</i>	<i>elevator_name</i>	The type of I/O elevator currently enabled
<i>disk_minor</i>	<i>disk_minor</i>	Disk minor number of the request

---

# **Chapter 6. SCSI Tapset**

This family of probe points is used to probe SCSI activities. It contains the following probe points:

## Name

scsi.ioentry -- Prepares a SCSI mid-layer request

## Synopsis

scsi.ioentry

## Values

*disk\_major*                    *disk\_major*  
The major number of the disk (-1 if no information)

*device\_state*                *device\_state*  
The current state of the device.

*disk\_minor*                *disk\_minor*  
The minor number of the disk (-1 if no information)

## Name

scsi.iodispatching -- SCSI mid-layer dispatched low-level SCSI command

## Synopsis

scsi.iodispatching

## Values

<i>lun</i>	<i>lun</i>	The lun number
<i>req_bufflen</i>	<i>req_bufflen</i>	The request buffer length
<i>host_no</i>	<i>host_no</i>	The host number
<i>device_state</i>	<i>device_state</i>	The current state of the device.
<i>dev_id</i>	<i>dev_id</i>	The scsi device id
<i>channel</i>	<i>channel</i>	The channel number
<i>data_direction</i>	<i>data_direction</i>	The data_direction specifies whether this command is from/to the device. 0 (DMA_BIDIRECTIONAL), 1 (DMA_TO_DEVICE), 2 (DMA_FROM_DEVICE), 3 (DMA_NONE)
<i>request_buffer</i>	<i>request_buffer</i>	The request buffer address

## Name

scsi.iodone -- SCSI command completed by low level driver and enqueued into the done queue.

## Synopsis

scsi.iodone

## Values

<i>lun</i>	<i>lun</i>	The lun number
<i>host_no</i>	<i>host_no</i>	The host number
<i>device_state</i>	<i>device_state</i>	The current state of the device
<i>dev_id</i>	<i>dev_id</i>	The scsi device id
<i>channel</i>	<i>channel</i>	The channel number
<i>data_direction</i>	<i>data_direction</i>	The data_direction specifies whether this command is from/to the device.

## Name

scsi.iocompleted -- SCSI mid-layer running the completion processing for block device I/O requests

## Synopsis

scsi.iocompleted

## Values

<i>lun</i>	<i>lun</i>	The lun number
<i>host_no</i>	<i>host_no</i>	The host number
<i>device_state</i>	<i>device_state</i>	The current state of the device
<i>dev_id</i>	<i>dev_id</i>	The scsi device id
<i>channel</i>	<i>channel</i>	The channel number
<i>data_direction</i>	<i>data_direction</i>	The data_direction specifies whether this command is from/to the device
<i>goodbytes</i>	<i>goodbytes</i>	The bytes completed.

---

# **Chapter 7. Networking Tapset**

This family of probe points is used to probe the activities of the network device and protocol layers.

## Name

netdev.receive -- Data received from network device.

## Synopsis

netdev.receive

## Values

*protocol*      *protocol*  
Protocol of received packet.

*dev\_name*      *dev\_name*  
The name of the device. e.g: eth0, ath1.

*length*      *length*  
The length of the receiving buffer.

## Name

netdev.transmit -- Network device transmitting buffer

## Synopsis

netdev.transmit

## Values

*protocol*      *protocol*  
The protocol of this packet.

*dev\_name*      *dev\_name*  
The name of the device. e.g: eth0, ath1.

*length*      *length*  
The length of the transmit buffer.

*truesize*      *truesize*  
The size of the the data to be transmitted.

## Name

tcp.sendmsg -- Sending a tcp message

## Synopsis

`tcp.sendmsg`

## Values

*name name*  
Name of this probe

*size size*  
Number of bytes to send

*sock sock*  
Network socket

## Context

The process which sends a tcp message

## Name

tcp.sendmsg.return -- Sending TCP message is done

## Synopsis

tcp.sendmsg.return

## Values

*name*    *name*  
Name of this probe

*size*    *size*  
Number of bytes sent or error code if an error occurred.

## Context

The process which sends a tcp message

## Name

tcp.recvmsg -- Receiving TCP message

## Synopsis

`tcp.recvmsg`

## Values

*saddr*      *saddr*

A string representing the source IP address

*daddr*      *daddr*

A string representing the destination IP address

*name*      *name*

Name of this probe

*sport*      *sport*

TCP source port

*dport*      *dport*

TCP destination port

*size*      *size*

Number of bytes to be received

*sock*      *sock*

Network socket

## Context

The process which receives a tcp message

## Name

tcp.recvmsg.return -- Receiving TCP message complete

## Synopsis

```
tcp.recvmsg.return
```

## Values

*saddr*      *saddr*

A string representing the source IP address

*daddr*      *daddr*

A string representing the destination IP address

*name*      *name*

Name of this probe

*sport*      *sport*

TCP source port

*dport*      *dport*

TCP destination port

*size*      *size*

Number of bytes received or error code if an error occurred.

## Context

The process which receives a tcp message

## Name

tcp.disconnect -- TCP socket disconnection

## Synopsis

tcp.disconnect

## Values

*saddr*      *saddr*

A string representing the source IP address

*daddr*      *daddr*

A string representing the destination IP address

*flags*      *flags*

TCP flags (e.g. FIN, etc)

*name*      *name*

Name of this probe

*sport*      *sport*

TCP source port

*dport*      *dport*

TCP destination port

*sock*      *sock*

Network socket

## Context

The process which disconnects tcp

## Name

tcp.disconnect.return -- TCP socket disconnection complete

## Synopsis

tcp.disconnect.return

## Values

*ret*      *ret*  
Error code (0: no error)

*name*    *name*  
Name of this probe

## Context

The process which disconnects tcp

## Name

tcp.setsockopt -- Call to setsockopt

## Synopsis

tcp.setsockopt

## Values

*optstr*      *optstr*  
Resolves optname to a human-readable format

*level*      *level*  
The level at which the socket options will be manipulated

*optlen*      *optlen*  
Used to access values for setsockopt

*name*      *name*  
Name of this probe

*optname*      *optname*  
TCP socket options (e.g. TCP\_NODELAY, TCP\_MAXSEG, etc)

*sock*      *sock*  
Network socket

## Context

The process which calls setsockopt

## Name

tcp.setsockopt.return -- Return from setsockopt

## Synopsis

tcp.setsockopt.return

## Values

*ret*      *ret*  
Error code (0: no error)

*name*    *name*  
Name of this probe

## Context

The process which calls setsockopt

## Name

tcp.receive -- Called when a TCP packet is received

## Synopsis

tcp.receive

## Values

*urg*        *urg*  
TCP URG flag

*psh*        *psh*  
TCP PSH flag

*rst*        *rst*  
TCP RST flag

*dport*      *dport*  
TCP destination port

*saddr*      *saddr*  
A string representing the source IP address

*daddr*      *daddr*  
A string representing the destination IP address

*ack*        *ack*  
TCP ACK flag

*syn*        *syn*  
TCP SYN flag

*fin*        *fin*  
TCP FIN flag

*sport*      *sport*  
TCP source port

## Name

udp.sendmsg -- Fires whenever a process sends a UDP message

## Synopsis

udp.sendmsg

## Values

*name*    *name*

The name of this probe

*size*    *size*

Number of bytes sent by the process

*sock*    *sock*

Network socket used by the process

## Context

The process which sent a UDP message

## Name

udp.sendmsg.return -- Fires whenever an attempt to send a UDP message is completed

## Synopsis

udp.sendmsg.return

## Values

*name name*

The name of this probe

*size size*

Number of bytes sent by the process

## Context

The process which sent a UDP message

## Name

udp.recvmsg -- Fires whenever a UDP message is received

## Synopsis

udp.recvmsg

## Values

*name name*

The name of this probe

*size size*

Number of bytes received by the process

*sock sock*

Network socket used by the process

## Context

The process which received a UDP message

## Name

udp.recvmsg.return -- Fires whenever an attempt to receive a UDP message received is completed

## Synopsis

udp.recvmsg.return

## Values

*name*    *name*

The name of this probe

*size*    *size*

Number of bytes received by the process

## Context

The process which received a UDP message

## Name

udp.disconnect -- Fires when a process requests for a UDP disconnection

## Synopsis

udp.disconnect

## Values

*flags*      *flags*  
Flags (e.g. FIN, etc)

*name*      *name*  
The name of this probe

*sock*      *sock*  
Network socket used by the process

## Context

The process which requests a UDP disconnection

## Name

udp.disconnect.return -- UDP has been disconnected successfully

## Synopsis

udp.disconnect.return

## Values

*ret*      *ret*  
Error code (0: no error)

*name*      *name*  
The name of this probe

## Context

The process which requested a UDP disconnection

## Name

ip\_ntop -- returns a string representation from an integer IP number

## Synopsis

```
ip_ntop:string(addr:long)
```

## Arguments

*addr*      *addr*  
the ip represented as an integer

---

# **Chapter 8. Socket Tapset**

This family of probe points is used to probe socket activities. It contains the following probe points:

## Name

socket.send -- Message sent on a socket.

## Synopsis

socket.send

## Values

*success*            *success*  
Was send successful? (1 = yes, 0 = no)

*protocol*        *protocol*  
Protocol value

*flags*            *flags*  
Socket flags value

*name*            *name*  
Name of this probe

*state*            *state*  
Socket state value

*size*            *size*  
Size of message sent (in bytes) or error code if success = 0

*type*            *type*  
Socket type value

*family*          *family*  
Protocol family value

## Context

The message sender

## Name

socket.receive -- Message received on a socket.

## Synopsis

socket.receive

## Values

*success*            *success*  
Was send successful? (1 = yes, 0 = no)

*protocol*        *protocol*  
Protocol value

*flags*            *flags*  
Socket flags value

*name*            *name*  
Name of this probe

*state*            *state*  
Socket state value

*size*            *size*  
Size of message received (in bytes) or error code if success = 0

*type*            *type*  
Socket type value

*family*          *family*  
Protocol family value

## Context

The message receiver

## Name

socket.sendmsg -- Message is currently being sent on a socket.

## Synopsis

socket.sendmsg

## Values

*protocol*      *protocol*  
Protocol value

*flags*      *flags*  
Socket flags value

*name*      *name*  
Name of this probe

*state*      *state*  
Socket state value

*size*      *size*  
Message size in bytes

*type*      *type*  
Socket type value

*family*      *family*  
Protocol family value

## Context

The message sender

## Description

Fires at the beginning of sending a message on a socket via the the `sock_sendmsg` function

## Name

socket.sendmsg.return -- Return from <command>socket.sendmsg</command>.

## Synopsis

```
socket.sendmsg.return
```

## Values

*success*            *success*  
Was send successful? (1 = yes, 0 = no)

*protocol*        *protocol*  
Protocol value

*flags*            *flags*  
Socket flags value

*name*            *name*  
Name of this probe

*state*            *state*  
Socket state value

*size*            *size*  
Size of message sent (in bytes) or error code if success = 0

*type*            *type*  
Socket type value

*family*          *family*  
Protocol family value

## Context

The message sender.

## Description

Fires at the conclusion of sending a message on a socket via the `sock_sendmsg` function

## Name

socket.recvmsg -- Message being received on socket

## Synopsis

socket.recvmsg

## Values

*protocol*      *protocol*  
Protocol value

*flags*      *flags*  
Socket flags value

*name*      *name*  
Name of this probe

*state*      *state*  
Socket state value

*size*      *size*  
Message size in bytes

*type*      *type*  
Socket type value

*family*      *family*  
Protocol family value

## Context

The message receiver.

## Description

Fires at the beginning of receiving a message on a socket via the `sock_recvmsg` function

## Name

socket.recvmsg.return -- Return from Message being received on socket

## Synopsis

```
socket.recvmsg.return
```

## Values

*success*            *success*  
Was receive successful? (1 = yes, 0 = no)

*protocol*        *protocol*  
Protocol value

*flags*            *flags*  
Socket flags value

*name*            *name*  
Name of this probe

*state*            *state*  
Socket state value

*size*            *size*  
Size of message received (in bytes) or error code if success = 0

*type*            *type*  
Socket type value

*family*          *family*  
Protocol family value

## Context

The message receiver.

## Description

Fires at the conclusion of receiving a message on a socket via the `sock_recvmsg` function.

## Name

socket.aio\_write -- Message send via `sock_aio_write`

## Synopsis

```
socket.aio_write
```

## Values

*protocol*        *protocol*  
Protocol value

*flags*        *flags*  
Socket flags value

*name*        *name*  
Name of this probe

*state*        *state*  
Socket state value

*size*        *size*  
Message size in bytes

*type*        *type*  
Socket type value

*family*        *family*  
Protocol family value

## Context

The message sender

## Description

Fires at the beginning of sending a message on a socket via the `sock_aio_write` function

## Name

socket.aio\_write.return -- Conclusion of message send via `sock_aio_write`

## Synopsis

```
socket.aio_write.return
```

## Values

*success*            *success*  
Was receive successful? (1 = yes, 0 = no)

*protocol*        *protocol*  
Protocol value

*flags*            *flags*  
Socket flags value

*name*            *name*  
Name of this probe

*state*            *state*  
Socket state value

*size*            *size*  
Size of message received (in bytes) or error code if success = 0

*type*            *type*  
Socket type value

*family*          *family*  
Protocol family value

## Context

The message receiver.

## Description

Fires at the conclusion of sending a message on a socket via the `sock_aio_write` function

## Name

socket.aio\_read -- Receiving message via `sock_aio_read`

## Synopsis

`socket.aio_read`

## Values

*protocol*      *protocol*  
Protocol value

*flags*      *flags*  
Socket flags value

*name*      *name*  
Name of this probe

*state*      *state*  
Socket state value

*size*      *size*  
Message size in bytes

*type*      *type*  
Socket type value

*family*      *family*  
Protocol family value

## Context

The message sender

## Description

Fires at the beginning of receiving a message on a socket via the `sock_aio_read` function

## Name

socket.aio\_read.return -- Conclusion of message received via `sock_aio_read`

## Synopsis

```
socket.aio_read.return
```

## Values

*success*            *success*  
Was receive successful? (1 = yes, 0 = no)

*protocol*        *protocol*  
Protocol value

*flags*            *flags*  
Socket flags value

*name*            *name*  
Name of this probe

*state*            *state*  
Socket state value

*size*            *size*  
Size of message received (in bytes) or error code if success = 0

*type*            *type*  
Socket type value

*family*          *family*  
Protocol family value

## Context

The message receiver.

## Description

Fires at the conclusion of receiving a message on a socket via the `sock_aio_read` function

## Name

socket.writev -- Message sent via socket\_writev

## Synopsis

socket.writev

## Values

*protocol*      *protocol*  
Protocol value

*flags*      *flags*  
Socket flags value

*name*      *name*  
Name of this probe

*state*      *state*  
Socket state value

*size*      *size*  
Message size in bytes

*type*      *type*  
Socket type value

*family*      *family*  
Protocol family value

## Context

The message sender

## Description

Fires at the beginning of sending a message on a socket via the `sock_writev` function

## Name

socket.writev.return -- Conclusion of message sent via socket\_writev

## Synopsis

```
socket.writev.return
```

## Values

*success*            *success*  
Was send successful? (1 = yes, 0 = no)

*protocol*        *protocol*  
Protocol value

*flags*            *flags*  
Socket flags value

*name*            *name*  
Name of this probe

*state*            *state*  
Socket state value

*size*            *size*  
Size of message sent (in bytes) or error code if success = 0

*type*            *type*  
Socket type value

*family*          *family*  
Protocol family value

## Context

The message receiver.

## Description

Fires at the conclusion of sending a message on a socket via the `sock_writev` function

## Name

socket.readv -- Receiving a message via `sock_readv`

## Synopsis

`socket.readv`

## Values

*protocol*      *protocol*  
Protocol value

*flags*      *flags*  
Socket flags value

*name*      *name*  
Name of this probe

*state*      *state*  
Socket state value

*size*      *size*  
Message size in bytes

*type*      *type*  
Socket type value

*family*      *family*  
Protocol family value

## Context

The message sender

## Description

Fires at the beginning of receiving a message on a socket via the `sock_readv` function

## Name

socket.readv.return -- Conclusion of receiving a message via `sock_readv`

## Synopsis

```
socket.readv.return
```

## Values

*success*            *success*  
Was receive successful? (1 = yes, 0 = no)

*protocol*        *protocol*  
Protocol value

*flags*            *flags*  
Socket flags value

*name*            *name*  
Name of this probe

*state*            *state*  
Socket state value

*size*            *size*  
Size of message received (in bytes) or error code if success = 0

*type*            *type*  
Socket type value

*family*          *family*  
Protocol family value

## Context

The message receiver.

## Description

Fires at the conclusion of receiving a message on a socket via the `sock_readv` function

## Name

socket.create -- Creation of a socket

## Synopsis

socket.create

## Values

*protocol*            *protocol*  
Protocol value

*name*                *name*  
Name of this probe

*requester*          *requester*  
Requested by user process or the kernel (1 = kernel, 0 = user)

*type*                *type*  
Socket type value

*family*              *family*  
Protocol family value

## Context

The requester (see requester variable)

## Description

Fires at the beginning of creating a socket.

## Name

socket.create.return -- Return from Creation of a socket

## Synopsis

```
socket.create.return
```

## Values

*success*                   *success*  
Was socket creation successful? (1 = yes, 0 = no)

*protocol*               *protocol*  
Protocol value

*err*                     *err*  
Error code if success == 0

*name*                   *name*  
Name of this probe

*requester*             *requester*  
Requested by user process or the kernel (1 = kernel, 0 = user)

*type*                   *type*  
Socket type value

*family*                *family*  
Protocol family value

## Context

The requester (user process or kernel)

## Description

Fires at the conclusion of creating a socket.

## Name

socket.close -- Close a socket

## Synopsis

socket.close

## Values

*protocol*      *protocol*  
Protocol value

*flags*      *flags*  
Socket flags value

*name*      *name*  
Name of this probe

*state*      *state*  
Socket state value

*type*      *type*  
Socket type value

*family*      *family*  
Protocol family value

## Context

The requester (user process or kernel)

## Description

Fires at the beginning of closing a socket.

## Name

socket.close.return -- Return from closing a socket

## Synopsis

```
socket.close.return
```

## Values

*name*    *name*  
Name of this probe

## Context

The requester (user process or kernel)

## Description

Fires at the conclusion of closing a socket.

## Name

sock\_prot\_num2str -- Given a protocol number, return a string representation.

## Synopsis

```
sock_prot_num2str:string(proto:long)
```

## Arguments

*proto*    *proto*  
The protocol number.

## Name

sock\_prot\_str2num -- Given a protocol name (string), return the corresponding protocol number.

## Synopsis

```
sock_prot_str2num:long(proto:string)
```

## Arguments

*proto*    *proto*  
The protocol name.

## Name

sock\_fam\_num2str -- Given a protocol family number, return a string representation.

## Synopsis

```
sock_fam_num2str:string(family:long)
```

## Arguments

*family*      *family*  
The family number.

## Name

sock\_fam\_str2num -- Given a protocol family name (string), return the corresponding

## Synopsis

```
sock_fam_str2num:long(family:string)
```

## Arguments

*family*      *family*

The family name.

## Description

protocol family number.

## Name

sock\_state\_num2str -- Given a socket state number, return a string representation.

## Synopsis

```
sock_state_num2str:string(state:long)
```

## Arguments

*state*      *state*

The state number.

## Name

sock\_state\_str2num -- Given a socket state string, return the corresponding state number.

## Synopsis

```
sock_state_str2num:long(state:string)
```

## Arguments

*state*      *state*

The state name.

---

# **Chapter 9. Kernel Process Tapset**

This family of probe points is used to probe process-related activities. It contains the following probe points:

## Name

kprocess.create -- Fires whenever a new process is successfully created

## Synopsis

```
kprocess.create
```

## Values

*new\_pid*      *new\_pid*  
The PID of the newly created process

## Context

Parent of the created process.

## Description

Fires whenever a new process is successfully created, either as a result of <command>fork</command> (or one of its syscall variants), or a new kernel thread.

## Name

kprocess.start -- Starting new process

## Synopsis

```
kprocess.start
```

## Values

None

## Context

Newly created process.

## Description

Fires immediately before a new process begins execution.

## Name

kprocess.exec -- Attempt to exec to a new program

## Synopsis

`kprocess.exec`

## Values

*filename*      *filename*  
The path to the new executable

## Context

The caller of exec.

## Description

Fires whenever a process attempts to exec to a new program.

## Name

kprocess.exec\_complete -- Return from exec to a new program

## Synopsis

`kprocess.exec_complete`

## Values

*success*      *success*

A boolean indicating whether the exec was successful

*errno*      *errno*

The error number resulting from the exec

## Context

On success, the context of the new executable. On failure, remains in the context of the caller.

## Description

Fires at the completion of an exec call.

## Name

kprocess.exit -- Exit from process

## Synopsis

```
kprocess.exit
```

## Values

*code*    *code*  
The exit code of the process

## Context

The process which is terminating.

## Description

Fires when a process terminates. This will always be followed by a kprocess.release, though the latter may be delayed if the process waits in a zombie state.

## Name

kprocess.release -- Process released

## Synopsis

```
kprocess.release
```

## Values

*pid*      *pid*  
PID of the process being released

*task*      *task*  
A task handle to the process being released

## Context

The context of the parent, if it wanted notification of this process' termination, else the context of the process itself.

## Description

Fires when a process is released from the kernel. This always follows a kprocess.exit, though it may be delayed somewhat if the process waits in a zombie state.

---

# **Chapter 10. Signal Tapset**

This family of probe points is used to probe signal activities. It contains the following probe points:

## Name

signal.send -- Signal being sent to a process

## Synopsis

```
signal.send
```

## Values

<i>send2queue</i>	<i>send2queue</i>
Indicates whether the signal is sent to an existing <command>sigqueue</command>	
<i>name</i>	<i>name</i>
The name of the function used to send out the signal	
<i>task</i>	<i>task</i>
A task handle to the signal recipient	
<i>sinfo</i>	<i>sinfo</i>
The address of <command>siginfo</command> struct	
<i>si_code</i>	<i>si_code</i>
Indicates the signal type	
<i>sig_name</i>	<i>sig_name</i>
A string representation of the signal	
<i>sig</i>	<i>sig</i>
The number of the signal	
<i>shared</i>	<i>shared</i>
Indicates whether the signal is shared by the thread group	
<i>sig_pid</i>	<i>sig_pid</i>
The PID of the process receiving the signal	
<i>pid_name</i>	<i>pid_name</i>
The name of the signal recipient	

## Context

The signal's sender.

## Name

signal.send.return -- Signal being sent to a process completed

## Synopsis

```
signal.send.return
```

## Values

<i>retstr</i>	<i>retstr</i>
The return value to either <command>__group_send_sig_info</command>, <command>specific_send_sig_info</command>, or <command>send_sigqueue</command>	
<i>send2queue</i>	<i>send2queue</i>
Indicates whether the sent signal was sent to an existing <command>sigqueue</command>	
<i>name</i>	<i>name</i>
The name of the function used to send out the signal	
<i>shared</i>	<i>shared</i>
Indicates whether the sent signal is shared by the thread group.	

## Context

The signal's sender. <remark>(correct?)</remark>

## Description

Possible <command>\_\_group\_send\_sig\_info</command> and <command>specific\_send\_sig\_info</command> return values are as follows;

<command>0</command> -- The signal is sucessfully sent to a process, which means that <1> the signal was ignored by the receiving process, <2> this is a non-RT signal and the system already has one queued, and <3> the signal was successfully added to the <command>sigqueue</command> of the receiving process.

<command>-EAGAIN</command> -- The <command>sigqueue</command> of the receiving process is overflowing, the signal was RT, and the signal was sent by a user using something other than <command>kill</command>.

Possible <command>send\_group\_sigqueue</command> and <command>send\_sigqueue</command> return values are as follows;

<command>0</command> -- The signal was either sucessfully added into the <command>sigqueue</command> of the receiving process, or a <command>SI\_TIMER</command> entry is already queued (in which case, the overrun count will be simply incremented).

<command>1</command> -- The signal was ignored by the receiving process.

<command>-1</command> -- (<command>send\_sigqueue</command> only) The task was marked <command>exiting</command>, allowing \* <command>posix\_timer\_event</command> to redirect it to the group leader.

## Name

signal.checkperm -- Check being performed on a sent signal

## Synopsis

signal.checkperm

## Values

<i>name</i>	<i>name</i>
Name of the probe point; default value is <command>signal.checkperm</command>	
<i>task</i>	<i>task</i>
A task handle to the signal recipient	
<i>sinfo</i>	<i>sinfo</i>
The address of the <command>siginfo</command> structure	
<i>si_code</i>	<i>si_code</i>
Indicates the signal type	
<i>sig_name</i>	<i>sig_name</i>
A string representation of the signal	
<i>sig</i>	<i>sig</i>
The number of the signal	
<i>pid_name</i>	<i>pid_name</i>
Name of the process receiving the signal	
<i>sig_pid</i>	<i>sig_pid</i>
The PID of the process receiving the signal	

## Name

signal.checkperm.return -- Check performed on a sent signal completed

## Synopsis

signal.checkperm.return

## Values

*retstr*      *retstr*  
Return value as a string

*name*      *name*  
Name of the probe point; default value is <command>signal.checkperm</command>

## Name

signal.wakeup -- Sleeping process being wakened for signal

## Synopsis

signal.wakeup

## Values

*resume*                  *resume*

Indicates whether to wake up a task in a <command>STOPPED</command> or <command>TRACED</command> state

*state\_mask*                  *state\_mask*

A string representation indicating the mask of task states to wake. Possible values are <command>TASK\_INTERRUPTIBLE</command>, <command>TASK\_STOPPED</command>, <command>TASK\_TRACED</command>, and <command>TASK\_INTERRUPTIBLE</command>.

*pid\_name*                  *pid\_name*  
Name of the process to wake

*sig\_pid*                  *sig\_pid*  
The PID of the process to wake

## Name

signal.check\_ignored -- Checking to see signal is ignored

## Synopsis

signal.check\_ignored

## Values

*sig\_name*      *sig\_name*  
A string representation of the signal

*sig*      *sig*  
The number of the signal

*pid\_name*      *pid\_name*  
Name of the process receiving the signal

*sig\_pid*      *sig\_pid*  
The PID of the process receiving the signal

## Name

signal.check\_ignored.return -- Check to see signal is ignored completed

## Synopsis

signal.check\_ignored.return

## Values

*retstr*      *retstr*  
Return value as a string

*name*      *name*  
Name of the probe point; default value is <command>signal.checkperm</command>

## Name

signal.force\_segv -- Forcing send of <command>SIGSEGV</command>

## Synopsis

signal.force\_segv

## Values

*sig\_name*      *sig\_name*  
A string representation of the signal

*sig*      *sig*  
The number of the signal

*pid\_name*      *pid\_name*  
Name of the process receiving the signal

*sig\_pid*      *sig\_pid*  
The PID of the process receiving the signal

## Name

signal.force\_segv.return -- Forcing send of <command>SIGSEGV</command> complete

## Synopsis

```
signal.force_segv.return
```

## Values

*retstr*      *retstr*  
Return value as a string

*name*      *name*  
Name of the probe point; default value is <command>force\_sigsegv</command>

## Name

signal.syskill -- Sending kill signal to a process

## Synopsis

signal.syskill

## Values

*sig sig*

The specific signal sent to the process

*pid pid*

The PID of the process receiving the signal

## Name

signal.syskill.return -- Sending kill signal completed

## Synopsis

```
signal.syskill.return
```

## Values

None

## Name

signal.sys\_tkill -- Sending a kill signal to a thread

## Synopsis

`signal.sys_tkill`

## Values

*sig\_name*      *sig\_name*  
The specific signal sent to the process

*sig*      *sig*  
The specific signal sent to the process

*pid*      *pid*  
The PID of the process receiving the kill signal

## Description

The `<command>tkill</command>` call is analogous to `<command>kill(2)</command>`, except that it also allows a process within a specific thread group to be targetted. Such processes are targetted through their unique thread IDs (TID).

## Name

signal.systkill.return -- Sending kill signal to a thread completed

## Synopsis

signal.systkill.return

## Values

None

## Name

signal.sys\_tgkill -- Sending kill signal to a thread group

## Synopsis

`signal.sys_tgkill`

## Values

*sig\_name*      *sig\_name*  
A string representation of the signal

*sig*                *sig*  
The specific kill signal sent to the process

*pid*                *pid*  
The PID of the thread receiving the kill signal

*tgid*                *tgid*  
The thread group ID of the thread receiving the kill signal

## Description

The <command>`tgkill`</command> call is similar to <command>`tkill`</command>, except that it also allows the caller to specify the thread group ID of the thread to be signalled. This protects against TID reuse.

## Name

signal.sys\_tgkill.return -- Sending kill signal to a thread group completed

## Synopsis

```
signal.sys_tgkill.return
```

## Values

None

## Name

signal.send\_sig\_queue -- Queuing a signal to a process

## Synopsis

signal.send\_sig\_queue

## Values

*sigqueue\_addr*                    *sigqueue\_addr*  
The address of the signal queue

*sig\_name*                        *sig\_name*  
A string representation of the signal

*sig*                              *sig*  
The queued signal

*pid\_name*                      *pid\_name*  
Name of the process to which the signal is queued

*sig\_pid*                        *sig\_pid*  
The PID of the process to which the signal is queued

## Name

signal.send\_sig\_queue.return -- Queuing a signal to a process completed

## Synopsis

```
signal.send_sig_queue.return
```

## Values

*retstr*      *retstr*  
Return value as a string

## Name

signal.pending -- Examining pending signal

## Synopsis

signal.pending

## Values

*sigset\_size*      *sigset\_size*  
The size of the user-space signal set

*sigset\_add*      *sigset\_add*  
The address of the user-space signal set (<command>sigset\_t</command>)

## Description

This probe is used to examine a set of signals pending for delivery to a specific thread. This normally occurs when the <command>do\_sigpending</command> kernel function is executed.

## Name

signal.pending.return -- Examination of pending signal completed

## Synopsis

signal.pending.return

## Values

*retstr*      *retstr*  
Return value as a string

## Name

signal.handle -- Signal handler being invoked

## Synopsis

signal.handle

## Values

<i>regs</i>	<i>regs</i>
The address of the kernel-mode stack area	
<i>sig_code</i>	<i>sig_code</i>
The <command>si_code</command> value of the <command>siginfo</command> signal	
<i>sig_mode</i>	<i>sig_mode</i>
Indicates whether the signal was a user-mode or kernel-mode signal	
<i>sinfo</i>	<i>sinfo</i>
The address of the <command>siginfo</command> table	
<i>oldset_addr</i>	<i>oldset_addr</i>
The address of the bitmask array of blocked signals	
<i>sig</i>	<i>sig</i>
The signal number that invoked the signal handler	
<i>ka_addr</i>	<i>ka_addr</i>
The address of the <command>k_sigaction</command> table associated with the signal	

## Name

signal.handle.return -- Signal handler invocation completed

## Synopsis

```
signal.handle.return
```

## Values

*retstr*      *retstr*  
Return value as a string

## Name

signal.do\_action -- Examining or changing a signal action

## Synopsis

signal.do\_action

## Values

*sa\_mask*

*sa\_mask*

The new mask of the signal

*oldsigact\_addr*

*oldsigact\_addr*

The address of the old <command>`sigaction`</command> struct associated with the signal

*sig*

*sig*

The signal to be examined/changed

*sa\_handler*

*sa\_handler*

The new handler of the signal

*sigact\_addr*

*sigact\_addr*

The address of the new <command>`sigaction`</command> struct associated with the signal

## Name

signal.do\_action.return -- Examining or changing a signal action completed

## Synopsis

```
signal.do_action.return
```

## Values

*retstr*      *retstr*  
Return value as a string

## Name

signal.procmask -- Examining or changing blocked signals

## Synopsis

signal.procmask

## Values

<i>how</i>	<i>how</i>
	Indicates how to change the blocked signals; possible values are <command>SIG_BLOCK=0</command> (for blocking signals), <command>SIG_UNBLOCK=1</command> (for unblocking signals), and <command>SIG_SETMASK=2</command> for setting the signal mask.
<i>oldsigset_addr</i>	<i>oldsigset_addr</i>
	The old address of the signal set (<command>sigset_t</command>)
<i>sigset</i>	<i>sigset</i>
	The actual value to be set for <command>sigset_t</command> <remark>(correct?)</remark>
<i>sigset_addr</i>	<i>sigset_addr</i>
	The address of the signal set (<command>sigset_t</command>) to be implemented

## Name

signal.flush -- Flusing all pending signals for a task

## Synopsis

signal.flush

## Values

*task*            *task*

The task handler of the process performing the flush

*pid\_name*        *pid\_name*

The name of the process associated with the task performing the flush

*sig\_pid*        *sig\_pid*

The PID of the process associated with the task performing the flush