

SystemTap Tapset Reference Manual

William Cohen <wcohen@redhat.com>

SystemTap Tapset Reference Manual

by William Cohen

Copyright © 2008 Red Hat, Inc.

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. Memory Tapset	1
probe vm.pagefault	2
probe vm.pagefault.return	3
addr_to_node	4
probe vm.write_shared	5
probe vm.write_shared_copy	6
probe vm.mmap	7
probe vm.munmap	8
probe vm.brk	9
probe vm.oom_kill	10

Chapter 1. Memory Tapset

Name

probe vm.pagefault — Records that a page fault occurred.

probe vm.pagefault

Values

<i>write_access</i>	Indicates whether this was a write.
<i>address</i>	The address of the faulting memory access.

Context

The process which triggered the fault

Name

probe vm.pagefault.return — Records type of fault that occurred.

```
probe vm.pagefault.return
```

Values

fault_type 0 (VM_FAULT_OOM), 1 (VM_FAULT_SIGBUS), 2 (VM_FAULT_MINOR),
and 3 (VM_FAULT_MAJOR)

Name

`addr_to_node` — Returns which NUMA node has the given address.

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

```
addr:long;
```

Arguments

addr The address of the faulting memory access.

Name

probe vm.write_shared — Write to shared page.

probe vm.write_shared

Values

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 vm.write_shared_copy.

Name

probe vm.write_shared_copy — Page copy for shared page write.

```
probe vm.write_shared_copy
```

Values

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

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 vm.shared_write.

Name

probe vm.mmap — Fires when an mmap is requested.

probe vm.mmap

Values

length the length of the memory segment

address the requested address

Context

The process calling mmap.

Name

probe vm.munmap — Fires when an munmap is requested.

```
probe vm.munmap
```

Values

length the length of the memory segment

address the requested address

Context

The process calling munmap.

Name

probe vm.brk — Fires when a brk is requested (resizing a heap).

probe vm.brk

Values

None

Description

address - the requested address *length* - the length of the memory segment

Context

The process calling brk.

Name

probe vm.oom_kill — Fires when a thread is targetted by the OOM killer.

```
probe vm.oom_kill
```

Values

task the task being killed

Context

The process that tried to consume more memory, and thus triggered the OOM. (correct?)