

DTrace Quick Reference

DTrace Providers

| Provider | Description |
|------------------------|--|
| <code>dtrace</code> | Provides several probes related to DTrace itself. Use these probes to initialize state before tracing begins, process state after tracing has completed, and handle unexpected execution errors in other probes. |
| <code>lockstat</code> | Provides probes that can be used to discern lock contention statistics, or to understand virtually any aspect of locking behavior. |
| <code>profile</code> | Provides probes associated with a time-based interrupt firing every fixed, specified time interval. |
| <code>fbt</code> | Provides probes associated with the entry to and return from most functions in the Solaris kernel. |
| <code>syscall</code> | Provides a probe at the entry to and return from every system call in the system. |
| <code>sdt</code> | Provides probes at sites that a programmer has formally designated. This provider allows programmers to choose locations of interest to DTrace users and convey some knowledge about each location through the probe name. |
| <code>sysinfo</code> | Provides probes that correspond to kernel statistics classified by the name <code>sys</code> . |
| <code>vminfo</code> | Provides probes that correspond to the <code>vm</code> kernel statistics. |
| <code>proc</code> | Provides probes pertaining to the following activities: process creation and termination, LWP creation and termination, execution of new program images, and sending and handling signals. |
| <code>sched</code> | Provides probes related to CPU scheduling. Because a CPU is the one resource that all threads must consume, the <code>sched</code> provider is very useful for understanding systemic behavior. |
| <code>io</code> | Provides probes related to disk input and output. |
| <code>mib</code> | Provides probes that correspond to counters in the Solaris management information bases (MIBs). |
| <code>fpuinfo</code> | Provides probes that correspond to the simulation of floating-point instructions on SPARC microprocessors. |
| <code>pid</code> | Allows for tracing of the entry and return of any function in a user process as well as any instruction as specified by an absolute address or function offset. |
| <code>plockstat</code> | Provides probes that can be used to observe the behavior of user-level synchronization primitives, including lock contention and hold times. |
| <code>fasttrap</code> | Allows for tracing at specific, preprogrammed user process locations. |

DTrace Functions

| Name | Prototype | Description |
|-----------------------|--|---|
| <code>trace</code> | <code>void trace (expression)</code> | Takes a D expression as argument and traces the result to the directed buffer. |
| <code>tracemem</code> | <code>void tracemem (address, size_t nbytes)</code> | Takes the memory address specified by <code>address</code> into the directed buffer for the length specified by <code>nbytes</code> . <code>Address</code> is a D expression. |
| <code>printf</code> | <code>void printf (string format, ...)</code> | The arguments are a format string followed by a variable number of arguments. The arguments are formatted for output according to the specified format string. |
| <code>printa</code> | <code>void printa (aggregation) void printa (string</code> | Enables displaying and formatting of aggregations. If a format is not specified, the default format is used. |

| | | |
|------------|---|--|
| | format, aggregation) | |
| stack | void stack (int nframes), void stack (void) | Records a kernel stack trace, nframes in depth. The number specified by the stackframes option is used if nframes is not specified. May also be used as a key to an aggregation. |
| ustack | void ustack (int nframes, int strsize) void ustack (int nframes) void ustack (void) | Records a user stack trace, nframes in depth. The number specified by the ustackframes option is used if nframes is unspecified. If strsize is specified and non-zero, ustack() will allocate the specified amount of string space and use it to perform address-to-symbol translation directly from the kernel. |
| jstack | void jstack (int nframes, int strsize) void jstack (int nframes) void jstack (void) | Alias for ustack() that uses the jstackframes option for the stack frame value and jstackstrsize for the string space size. |
| stop | void stop (void) | Forces the process that fires the enabled probe to stop when it next leaves the kernel. |
| raise | void raise (int signal) | Sends the specified signal to the currently running process. |
| copyout | void copyout (void *buf, uintptr_t addr, size_t nbytes) | Copies nbytes from the buffer buf to the address addr in the address space of the process associated with the current thread. |
| copyoutstr | void copyoutstr (string str, uintptr_t addr, size_t maxlen) | Copies the string str to the address addr in the address space of the process associated with the current thread. The string length is limited to the value set by the strsize option. |
| system | void system (string program, ...) | Causes program to be executed as if it were given to the shell as input. Program may contain any of the printf/printa formats. Other arguments must match the specified format in program. |
| breakpoint | void breakpoint (void) | Induces a kernel breakpoint, causing the system to stop and transfer control to the kernel debugger. |
| panic | void panic (void) | Causes a kernel panic. Should be used to force a system crash dump at a time of interest. |
| chill | void chill (int nanoseconds) | Causes DTrace to spin for the given nanoseconds. For system safety, DTrace will refuse to execute the chill action for more than 500 milliseconds in each 1-second on any CPU. |
| exit | void exit (int status) | Immediately stops tracing, notifies DTrace consumer to cease tracing, performs any final processing, and calls exit with the specified status. |
| alloca | void *alloca (size_t size) | Allocates size bytes out of scratch space and returns a pointer to the allocated memory. |
| basename | string basename (char *str) | Creates a string that consists of a copy of the specified string, but without any prefix that ends in /. |
| bcopy | void bcopy (void *src, void *dest, | Copies size bytes from the memory pointed to by src, to the memory pointed to by dest. All source memory must |

| | | |
|---------------------|--|---|
| | size_t size) | lie outside of scratch memory, and all destination memory must lie within it. |
| cleanpath | string cleanpath (char *str) | Creates a string that consists of a copy of the path indicated by <i>str</i> , but with redundant elements eliminated. This might result in shorter invalid paths being returned. |
| copyin | void *copyin (uintptr_t addr, size_t size) | Copies the specified size in bytes from the specified user address into a DTrace scratch buffer and returns the address of this buffer. The resulting buffer pointer is 8-byte aligned. |
| copyinstr | string copyinstr (uintptr_t addr) | Copies a null-terminated C string from the specified user address into a DTrace scratch buffer, and returns the address of this buffer. The <i>strsize</i> option limits the string length. |
| copyinto | void copyinto (uintptr_t addr, size_t size, void *dest) | Copies the specified size in bytes from the specified user address into the DTrace scratch buffer specified by <i>dest</i> . |
| dirname | string dirname (char *str) | Creates a string that consists of all but the last level of the path name specified by <i>str</i> . |
| msgdsz | size_t msgdsz (mblk_t *mp) | Returns the number of bytes in the data message pointed to by <i>mp</i> . |
| msgsz | size_t msgsz (mblk_t *mp) | Returns the number of bytes in the message pointed to by <i>mp</i> . |
| mutex_owned | int mutex_owned (kmutex_t *mutex) | Returns non-zero if the calling thread currently holds the specified kernel mutex, or zero if the specified adaptive mutex is currently unowned. |
| mutex_owner | kthread_t *mutex_owner (kmutex_t *mutex) | Returns the thread pointer of the current owner of the specified adaptive kernel mutex. Returns NULL if the specified adaptive mutex is currently unowned, or if the specified mutex is a spin mutex. |
| mutex_type_adaptive | int mutex_type_adaptive (kmutex_t *mutex) | Returns non-zero if the specified kernel mutex is of type MUTEX_ADAPTIVE, or zero if it is not. |
| progenyof | int progenyof (pid_t pid) | Returns non-zero if the calling process is among the progeny of the specified process ID. |
| rand | int rand (void) | Returns a pseudo-random integer. |
| rw_iswriter | int rw_iswriter(krwlock_ t *rwlock) | Returns non-zero if the specified reader-writer lock is either held or desired by a writer. Returns zero if the lock is held only by readers, no writer is blocked, or the lock is not held at all. |
| rw_write_held | int rw_write_held (krwlock_t *rwlock) | Returns non-zero if the specified reader-writer lock is currently held by a writer. Returns zero if the lock is held only by readers or not held at all. |
| speculation | int speculation (void) | Reserves a speculative trace buffer for use with <i>speculate()</i> and returns an identifier for this buffer. |
| strjoin | string strjoin(char *str1 char *str2) | Creates a string that consists of <i>str1</i> concatenated with <i>str2</i> . |
| strlen | size_t strlen(string str) | Returns the length of the specified string in bytes, excluding the terminating null byte. |

DTrace Aggregating Functions

| Name | Arguments | Result |
|-----------|---|---|
| count | none | Number of times called. |
| sum | scalar expression | Total value of the specified expressions. |
| avg | scalar expression | Arithmetic average of the specified expressions. |
| min | scalar expression | Smallest value among the specified expressions. |
| max | scalar expression | Largest value among the specified expressions. |
| lquantize | scalar expression, lower bound, upper bound, step value | A linear frequency distribution, sized by the specified range, of the values of the specified expressions. Increments the value in the highest bucket that is less than the specified expression. |
| quantize | scalar expression | A power-of-two frequency distribution of the values of the specified expressions. Increments the value in the highest power-of-two bucket that is less than the specified expression. |

DTrace Variables

| Variable | Description | Usage |
|-------------------|--|--|
| Scalar | Used to represent fixed-size data objects. They could be individual such as pointers and integers or composite such as arrays. | <code>x=123</code> (scalar variable <code>x</code> of type integer) |
| Associative array | Used to represent collections of data elements that can be retrieved by specifying a key. There is no predefined limit of the number of elements. Elements can be indexed by any tuple, and elements are not stored in preallocated consecutive storage locations. | <code>a[123, "hello"] = 56</code> (associative array <code>a</code> with key, <code>[int, string]</code>) |
| Thread-Local | Used to declare variable storage that is local to each operating system thread. | <code>self->x= 45</code> (thread-local variable <code>x</code> of type integer) |
| Clause-Local | This variable is active for the lifetime of a given probe clause and its storage is reused for each D program clause. | <code>this->c='D'</code> (character clause-local variable <code>c</code>) |
| Built-in | All these variables are scalar global variables. | -- |
| External | The backquote character (<code>`</code>) is a scoping operator for accessing variables that are defined in the OS and not in your D program. | <code>`kmem_flags</code> (accessing a C variable in the kernel source code) |

DTrace Built-in Variables

| Type and Name | Description |
|--------------------------------------|---|
| <code>int64_t arg0, ..., arg9</code> | The first 10 input arguments to a probe represented as raw 64-bit integers. If fewer than 10 arguments are passed to the current probe, the remaining variables return zero. |
| <code>args[]</code> | The typed arguments to the current probe, if any. <code>args[]</code> array is accessed using an integer index, but each element is defined to be the type corresponding to the given probe argument. |
| <code>uintptr_t caller</code> | The program counter location of the current thread just before entering the current probe. |
| <code>chipid_t chip</code> | The CPU chip identifier for the current physical chip. |
| <code>processorid_t cpu</code> | The CPU identifier for the current CPU. |
| <code>cpuinfo_t *curcpu</code> | The CPU information for the current CPU. |
| <code>lwpsinfo_t *curlwpsinfo</code> | The lightweight process (LWP) state of the LWP associated with the current thread. |
| <code>psinfo_t *curpsinfo</code> | The process state of the process associated with the current thread. |
| <code>kthread_t</code> | The address of the operating system kernel's internal data structure for the current |

| | |
|-------------------------------------|---|
| <code>*curthread</code> | thread, the <code>kthread_t</code> . <code>kthread_t</code> is defined in <code><sys/thread.h></code> . |
| <code>string cwd</code> | The name of the current working directory of the process associated with the current thread. |
| <code>uint_t epid</code> | The enabled probe ID (EPID) for the current probe. This integer uniquely identifies a particular probe that is enabled with a specific predicate and set of actions. |
| <code>int errno</code> | The error value returned by the last system call executed by this thread. |
| <code>string execname</code> | The name that was passed to <code>exec(2)</code> to execute the current process. |
| <code>gid_t gid</code> | The real group ID of the current process. |
| <code>uint_t id</code> | The probe ID for the current probe. This ID is the system-wide unique identifier for the probe as published by DTrace. |
| <code>uint_t ipl</code> | The interrupt priority level (IPL) on the current CPU at probe firing time. |
| <code>lgrp_id_t lgrp</code> | The latency group ID for the latency group of which the current CPU is a member. |
| <code>pid_t pid</code> | The process ID of the current process. |
| <code>pid_t ppid</code> | The parent process ID of the current process. |
| <code>string probefunc</code> | The function name portion of the current probe's description. |
| <code>string probemod</code> | The module name portion of the current probe's description. |
| <code>string probename</code> | The name portion of the current probe's description. |
| <code>string probeprovider</code> | The provider name portion of the current probe's description. |
| <code>psetid_t pset</code> | The processor set ID for the processor set containing the current CPU. |
| <code>string root</code> | The name of the root directory of the process associated with the current thread. |
| <code>uint_t stackdepth</code> | The current thread's stack frame depth at probe firing time. |
| <code>id_t tid</code> | The thread ID of the current thread. For threads associated with user processes, this value is equal to the result of a call to <code>pthread_self</code> . |
| <code>uint64_t timestamp</code> | The current value of a nanosecond timestamp counter. This counter increments from an arbitrary point in the past and should only be used for relative computations. |
| <code>uid_t uid</code> | The real user ID of the current process. |
| <code>uint64_t uregs[]</code> | The current thread's saved user-mode register values at probe firing time. |
| <code>uint64_t vtimestamp</code> | The current value of a nanosecond timestamp counter that is the amount of time the current thread has been running on a CPU, minus the time spent in DTrace predicates and actions. |
| <code>uint64_t walltimestamp</code> | The current number of nanoseconds since 00:00 Universal Coordinated Time, January 1, 1970. |