

# London OpenSolaris User Group


A wide-angle photograph of the London skyline at sunset. The sun is low on the horizon, casting a golden glow over the city. In the foreground, a bridge spans across the River Thames, with several boats visible underneath. The background features a mix of historic and modern architecture, including the dome of St. Paul's Cathedral and the distinctive curved shape of the Gherkin building. Construction cranes are visible against the sky, indicating ongoing development in the city.

## DTrace For Absolute Beginners

James MacFarlane  
Solaris RPE


# Why DTrace

- It's Dynamic
  - No need to re-compile, or restart the code being traced
- It's Flexible
  - Can trace (almost) anything
  - Can trace userland or kernel, or both in one go
- It's Discrete
  - Has virtually no impact on the observed code




# Alternative's ?

- truss(1m)
  - Userland function / system call tracing
  - Can attach to a running process
  - Slows the traced process execution down
  - Cannot view kernel calls
- tnf
  - Kernel tracing through shared buffers
  - Need to re-compile
  - A complete pain to use :-)




# Using TNF

- Define your probes
  - Set TNF\_PROBE\_[0-5] points in the code
  - Re-compile the driver
  - Install the module and reboot to activate it
- Start the tracing
  - Create kernel buffer
  - Enable probes
- Gather the results
  - Dump the kernel buffer




# Using DTrace

- Define your probes
  - On command line or from script file
- Read the output
- Modify your probes
- Read more output
  - And repeat ...




# DTrace Terms

- Provider
  - An application, library or kernel module that supplies probes for tracing
- Probe
  - A trigger point that can be used for taking readings
- Action
  - What to do when a probe is triggered
- Predicate
  - A method of restricting actions for a probe



# Providers & Probes

- Distributed across the kernel and userland
  - List all probes with 'dtrace -l'
- Probes specified as :  
    provider:module:function:name  
    eg syscall::ioctl:entry
- Can use wildcards to match multiple probes  
    eg syscall::open\*:entry



# Actions

- Actions are used to capture or display specific data from the probe
- They take effect every time the associated probe is triggered
- Examples :

stack


tracemem

printf

mutex\_owner

printa

dirname






# Predicates

- Predicates control the actions from a probe
- They can be based on PID, UID, GID, time etc
- Eg :

`syscall::entry`      Fires on every syscall made


`/ uid == 1001 /`      Restricts to syscalls for user 1001

`/ pid == 2579 /`      Restricts to process 2579 only



# Aggregations

- Allow us to produce more readable output
- Can be used to show :
  - Number of times called `count()`
  - Total value `sum()`
  - Average value `avg()`
  - Min / Max values `min()`, `max()`
  - Frequency distribution `quantize()`, `lquantize()`



# Writing D Scripts

- Can invoke DTrace directly or via a script
- Many “1-liners” available for common uses

```
dtrace -n 'syscall:::entry {}'
```


- Better to start with scripts

```
#!/usr/sbin/dtrace -s
```

```
syscall:::entry
```

```
{
```

```
}
```




# Script Layout

```
#!/usr/sbin/dtrace -s
```


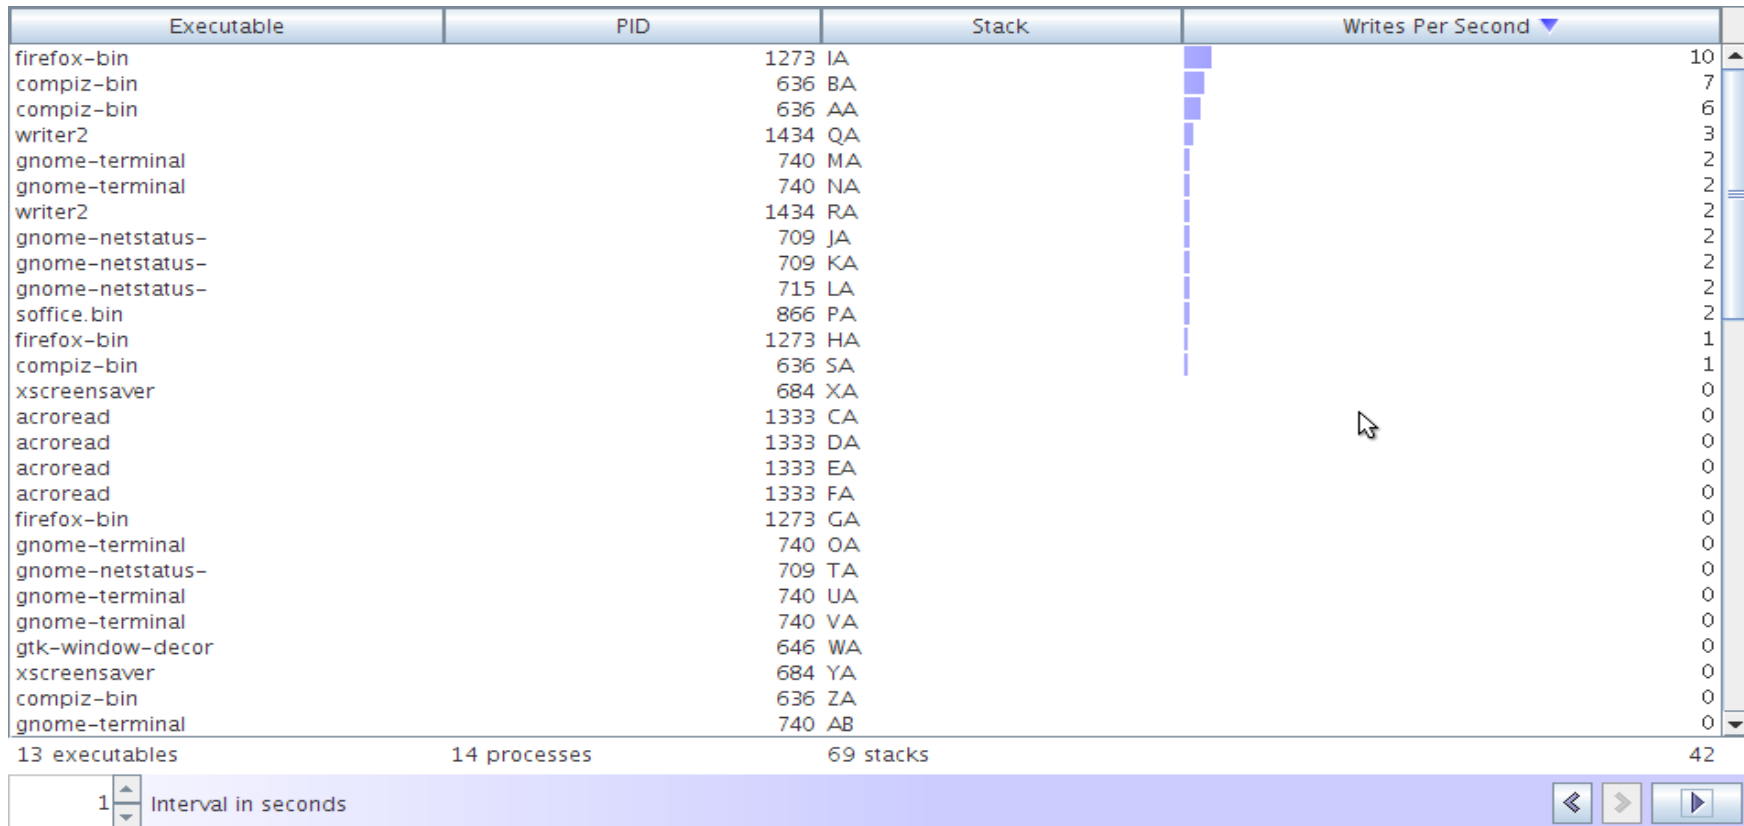
```
/*  
 * Comment Block  
*/
```

```
probe description  
/ predicate (optional) /  
{  
    action statements  
}
```





# Chime

- OpenSolaris DTrace GUI project.




# References

- Solaris Performance and Tools
  - Mauro, McDougall and Gregg
  - ISBN 978-0131568198



# References

- Dtrace Documentation
  - <http://wikis.sun.com/display/Dtrace/Documentation>
- Chime GUI
  - <http://hub.opensolaris.org/bin/view/Project+dtrace-chime/>
- Blogs
  - <http://blogs.sun.com/bmc>
  - <http://blogs.sun.com/ahl>
  - <http://blogs.sun.com/brendan>



# London OpenSolaris User Group

A panoramic view of the London skyline at sunset. The sun is low on the horizon, casting a golden glow over the city. In the foreground, a bridge spans across the River Thames. The background features a mix of historic and modern architecture, including St. Paul's Cathedral with its prominent dome, several skyscrapers, and construction cranes. The water of the river is calm, reflecting the warm light of the setting sun.

## DTrace For Absolute Beginners

James MacFarlane



# Example Scripts

```
#!/usr/sbin/dtrace -s
```

```
/*  
 * Basic truss functionality in DTrace  
 */
```

```
/*  
 * Trigger on entry to every syscall  
 * Add the name of the syscall to the appropriate array element  
 */
```

```
syscall:::entry  
/ pid == $1 /  
{  
    @Syscalls[probefunc] = count();  
}
```

```
/*  
 * Trigger every second  
 * Display and reset the contents of the array  
 */
```

```
profile:::tick-1s  
{  
    printa(@Syscalls);  
    trunc(@Syscalls);  
}
```

# Example Scripts

```
#!/usr/sbin/dtrace -s
```

```
/*  
 * Look at open / close calls on a specific process  
 */
```

```
syscall::open:entry  
/ pid == $1 /  
{  
    @Syscalls[probefunc] = count();  
    @Files[copyinstr(arg0)] = count();  
}
```

```
syscall::close:entry  
/ pid == $1 /  
{  
    @Syscalls[probefunc] = count();  
}
```

```
profile:::tick-5s  
{  
    printa(@Syscalls);  
    printa(@Files);  
    trunc(@Syscalls);  
    trunc(@Files);  
}
```