

ORACLE[®]



Systemtap Usage at Oracle

Elena Zannoni
Senior Development Manager, Linux Engineering

Some Oracle DB notes

- Linux, Solaris, HPUX, AIX, Windows, Mac OS, and many more platforms
- Database compiled with full optimization, uses PBO (profile based optimization),
- Linked binary has minimal symbol information
- Written in C and C++
- Multi-process (on Linux)
- Very few shared libraries
- On disk binary size: 145Mb

Systemtap with Oracle DB

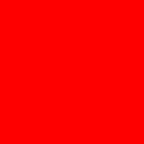
- Visibility into block layer behavior and into scheduler.
- Monitoring of i/o device drivers for cciss, networking layers in the kernel and the drivers for e1000/tg3
- Track killed Oracle process (killed process, process that signalled the kill)
- Tracking of user mode asynchronous I/O requests and synchronous I/O requests
- Tracking resources tuned via aio_nr and aio_max_nr
- Timer probes
- Tracked how long it takes for a killed process to be reaped by the OS
- OCFS2 tapset

Wish list for Systemtap

- User space probing
- Java tracing
- Kernel Markers, port to older kernels
- Debugging information related problems: size, and granularity. Microsoft Visual Studio compiler allows to select the amount of debugging information emitted by the compiler
- A way to use systemtap without requiring kernel-devel to be installed (other than cross-compiling)
- List of available markers and list of non-probeable functions.
- Script portability across kernel versions issue needs to be addressed

Kernel Stress Testing using .ko's

- Systemtap can be used to generate the .ko's
- Memory Fragmentation
 - <http://oss.oracle.com/projects/codefragments/src/trunk/fragment-slab>
- Memory Pressure
 - <http://oss.oracle.com/projects/codefragments/src/trunk/bufferheads>
- Code is available at those URLs.
- See article: Philip Copeland, “Using Kernel Modules to Help Test Applications”, Enterprise Open Source Magazine, August 2006
 - <http://extentech.sys-con.com/read/284266.htm>



The preceding is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.