Tutorial: Profiling in Xen

J. Renato Santos

HP Labs

Xen Summit September 7-8, 2006

Preparing the environment

- **Xen version requirement**
  - latest xen-unstable or xen 3.0.3 and above
- **Download oprofile 0.9.1 source (oprofile-0.9.1.tar.gz) from http://oprofile.sourceforge.net**
- **Download oprofile patch for Xen (oprofile-0.9.1-xen-rc2.patch) from http://xenoprof.sourceforge.net**
- **Compile and install oprofile modified for xen (as root)**

```
> tar –zxvf oprofile-0.9.1.tar.gz
> cd oprofile-0.9.1
> patch –p1 < ../oprofile-0.9.1-xen-rc2.patch
> ./configure –with-kernel-support
> make
> make install
```

Preparing the environment (cont)

- **Verify if we are using the right OProfile**
  - opcontrol –help
  - If successful, above command should list --xen as a valid option
- **Edit .config file for dom0 and domU to enable OProfile module (do not use kernel builtin option)**
  - CONFIG_PROFILING=y
    CONFIG_OPROFILE=m
- **Re-compile and re-install Xen in dom0 file system**
- **Copy Xen and kernel images (in /boot) to domU file system**
- **Copy domU modules (in /lib/modules) to domU file system**
- **We are now ready to start profiling**

Active domain profiling (1)

Example: Profiling domain 0, domain 1, and domain 3

- **Remove any samples from previous runs (in dom0 and domU's)**

  dom0> opcontrol --reset

  dom1> opcontrol –reset

  dom3> opcontrol --reset

- **Define profiling session parameters in dom0**
  - dom0> opcontrol --start-daemon --event=GLOBAL_POWER_EVENTS:1000000 --xen=/boot/xen-syms-3.0-unstable –vmlinux=/boot/vmlinux-syms-2.6.16.13-xen0 --active-domains=1,3

    --start-daemon : start oprofile daemon with given parameters but do not start profiling yet.

    --event : event used for profiling (non-halted clock cycles, cache misses, etc). can be specified multiple times

    --xen : Xen image (uncompressed version)

    --vmlinux : kernel image for dom0 (uncompressed version)

    --active-domains : list of domain ids to be profiled in active mode

Active domain profiling (2)

- **Start OProfile in all active domU's**

  dom1> opcontrol --start --event=GLOBAL_POWER_EVENT:1000000:1:1:1
  --xen=/boot/xen-syms-3.0-unstable –vmlinux=/boot/vmlinux-syms-2.6.16.13-xenU

  dom3> opcontrol --start --event=GLOBAL_POWER_EVENT:1000000:1:1:1
  --xen=/boot/xen-syms-3.0-unstable –vmlinux=/boot/vmlinux-syms-2.6.16.13-xenU

  --event : need to specify same event as specified in dom0.

- **Start profiling (in dom0)**

  dom0> opcontrol --start

- **Run experiment to be profiled**

- **Stop profiling**

  dom0> opcontrol --stop

- **Shutdown OProfile daemon**

  dom0> opcontrol --shutdown
  dom1> opcontrol --shutdown
  dom3> opcontrol --shutdown

Passive domain profiling

- **Remove any samples from previous runs (dom0)**

  dom0> opcontrol --reset

- **Define profiling session parameters (dom0)**
  - dom0> opcontrol --start-daemon --event=GLOBAL_POWER_EVENTS:1000000 --xen=/boot/xen-syms-3.0-unstable --vmlinux=/boot/vmlinux-syms-2.6.16.13-xen0 --passive-domains=1,3 --passive-images=/boot/vmlinux-syms-2.6.16.13-xenU,/boot/vmlinux-syms-2.6.16.13-xenU

    --passive-images : kernel images for passive domains (same order as --passive-domains)

- **Start profiling (in dom0)**

  dom0> opcontrol --start

- **Run experiment to be profiled**
- **Stop profiling**

  dom0> opcontrol --stop

- **Shutdown OProfile daemon**

  dom0> opcontrol --shutdown

Obtaining profile results (1)

- **Command to get per binary image profile**

  > opreport

```
CPU: P4 / Xeon with 2 hyper-threads, speed 2794.57 MHz (estimated)
Counted GLOBAL_POWER_EVENTS events (time during which processor is not stopped) with a unit
 mask of 0x01 (mandatory) count 1000000
GLOBAL_POWER_E...|
  samples|      %|
------------------
    82499 67.9306   vmlinux-syms-2.6.16.13-xenU
    38297 31.5342   xen-syms-3.0-unstable
      435  0.3582    rcv
       93  0.0766    oprofiled
       41  0.0338    libcrypto.so.0.9.7f
       39  0.0321    libc-2.3.6.so
       23  0.0189    ld-2.3.6.so
        7  0.0058    bash
        4  0.0033    nifd
        3  0.0025    sshd
```

Obtaining profile results (2)

- **Command to get per symbol (function) profile**

  > opreport -l

```
CPU: P4 / Xeon with 2 hyper-threads, speed 2794.57 MHz (estimated)
Counted GLOBAL_POWER_EVENTS events (time during which processor is not stopped) with a unit mask of 0x01 (mandatory) count 10(
samples  %        app name                 symbol name
6913     5.4236   xen-syms-3.0-unstable    do_grant_table_op
5922     4.6461   vmlinux-syms-2.6.16-xen0 net_tx_action
5419     4.2515   xen-syms-3.0-unstable    find_domain_by_id
5066     3.9745   xen-syms-3.0-unstable    gnttab_transfer
4870     3.8208   vmlinux-syms-2.6.16-xen0 net_rx_action
4311     3.3822   vmlinux-syms-2.6.16-xen0 nf_iterate
4032     3.1633   vmlinux-syms-2.6.16-xen0 hypercall_page
3555     2.7891   vmlinux-syms-2.6.16-xen0 eth_type_trans
3397     2.6651   xen-syms-3.0-unstable    hypercall
2820     2.2124   vmlinux-syms-2.6.16-xen0 e1000_intr
2607     2.0453   xen-syms-3.0-unstable    alloc_domheap_pages
2549     1.9998   vmlinux-syms-2.6.16-xen0 nf_hook_slow
2264     1.7762   xen-syms-3.0-unstable    evtchn_set_pending
2249     1.7645   vmlinux-syms-2.6.16-xen0 e1000_clean_rx_irq
2216     1.7386   xen-syms-3.0-unstable    __copy_from_user_ll
2178     1.7088   vmlinux-syms-2.6.16-xen0 dev_queue_xmit
2160     1.6946   xen-syms-3.0-unstable    get_page_type
                                           ⋮
44       0.0345   libc-2.3.6.so            malloc
44       0.0345   vmlinux-syms-2.6.16-xen0 e1000_update_stats
42       0.0330   vmlinux-syms-2.6.16-xen0 dma_map_page
41       0.0322   vmlinux-syms-2.6.16-xen0 raise_softirq_irqoff
38       0.0298   xen-syms-3.0-unstable    do_arch_sched_op
37       0.0290   vmlinux-syms-2.6.16-xen0 ktime_get
37       0.0290   vmlinux-syms-2.6.16-xen0 rcu_pending
```

Obtaining profile results (3)

- **Active domains**
  - Multiple reports: one per active domain
- **Passive domains**
  - Passive domain samples reported in domain 0 report

CPU: P4 / Xeon with 2 hyper-threads, speed 2794.57 MHz (estimated)
Counted GLOBAL_POWER_EVENTS events (time during which processor is not stopped) with a unit mask of 0x01 (mandatory) count

samples	%	image name	app name	symbol name
30353	12.0434	domain1-kernel	domain1-kernel	copy_to_user_ll
7174	2.8465	xen-syms-3.0-unstable	xen-syms-3.0-unstable	do_grant_table_op
6040	2.3965	vmlinux-syms-2.6.16.13-xen0	vmlinux-syms-2.6.16.13-xen0	net_tx_action
5508	2.1854	xen-syms-3.0-unstable	xen-syms-3.0-unstable	find_domain_by_id
4944	1.9617	xen-syms-3.0-unstable	xen-syms-3.0-unstable	gnttab_transfer
4848	1.9236	domain1-xen	domain1-xen	evtchn_set_pending
4631	1.8375	vmlinux-syms-2.6.16.13-xen0	vmlinux-syms-2.6.16.13-xen0	net_rx_action
4322	1.7149	domain1-kernel	domain1-kernel	tcp_v4_rcv
4145	1.6446	xen-syms-3.0-unstable	xen-syms-3.0-unstable	hypercall
4005	1.5891	domain1-xen	domain1-xen	guest_remove_page
3644	1.4459	vmlinux-syms-2.6.16.13-xen0	vmlinux-syms-2.6.16.13-xen0	hypercall_page
3589	1.4240	vmlinux-syms-2.6.16.13-xen0	vmlinux-syms-2.6.16.13-xen0	eth_type_trans
3425	1.3590	domain1-xen	domain1-xen	get_page_from_l1e
2846	1.1292	domain1-kernel	domain1-kernel	eth_type_trans
2770	1.0991	vmlinux-syms-2.6.16.13-xen0	vmlinux-syms-2.6.16.13-xen0	e1000_intr
75	0.0298	domain1-apps	domain1-apps	(no symbols)
69	0.0274	domain1-kernel	domain1-kernel	ns_to_timespec
69	0.0274	vmlinux-syms-2.6.16.13-xen0	vmlinux-syms-2.6.16.13-xen0	delay_tsc
68	0.0270	domain1-kernel	domain1-kernel	do_IRQ
66	0.0262	oprofiled	oprofiled	odb_insert

Other useful OProfile commands

- **List hardware performance events for current CPU model**
  - opcontrol -l
- **Unload oprofile module**
  - Opcontrol –deinit

- **More information**
  - OProfile user manual (http://oprofile.sourceforge.net)
  - OProfile man pages
  - XenOprofile user guide (http://xenoprof.sourceforge.net)


hp
invent