


Technology for better business outcomes



HP-UX 11i v3 Knowledge-on-Demand

- Objective: Support developers, deployment personnel and customers in achieving better business outcomes with HP-UX 11i
- What HP is providing: A series of technical on-demand training Webcasts
 - Focused on helping people who deploy HP-UX Integrity solutions increase performance through optimizing their installations for HP-UX 11i v3 on HP Integrity servers
 - Available at www.hp.com/go/kod

HP-UX 11i v3 Knowledge-on-Demand webcasts – planned topics

- HP-UX 11i v3 operating system optimization
 - Dynamic nPartitions
 - HP-UX 11i v3 tunables
 - I/O optimization
 - System Management Homepage
 - Software Assistant
 - GlancePlus
 - Performance trouble-shooting on v3
- Optimizing high availability configurations
 - Serviceguard configuration and manageability
 - Configuring Serviceguard with Oracle RAC
 - Serviceguard delta training
- Optimizing virtualized configurations
 - Capacity advisor
 - Workload monitoring and management
 - System sizing with HP VM

Additional Webcasts to be published going forward!

Topics subject to change without notice.



Related HP-UX 11i v3 resources

- All deployment resources
 - HP-UX 11i developers content
www.hp.com/go/hpuxdev
 - HP-UX 11i v3 news, functionality, product download, and services resources
www.hp.com/go/hpux11i
 - HP Integrity server product information
www.hp.com/go/integrity

We hope you enjoy this
Knowledge-on-Demand topic!

Thank you for taking time to learn about HP-UX 11i
v3 and related technologies.

Please provide feedback on today's topic and/or
future topics by using our online HP-UX 11i
Knowledge-on-Demand Feedback form:

www.hp.com/go/kodfeedback



HP-UX 11i performance trouble-shooting

Tools and tips



Technology for better business outcomes

Introducing the speaker

- Ken Johnson has worked at HP for 20 years and is a senior engineer on the Worldwide Technology Expert Center (WTEC) Performance Team.
- He specializes in fixing customer performance problems in their HP-UX server environments.



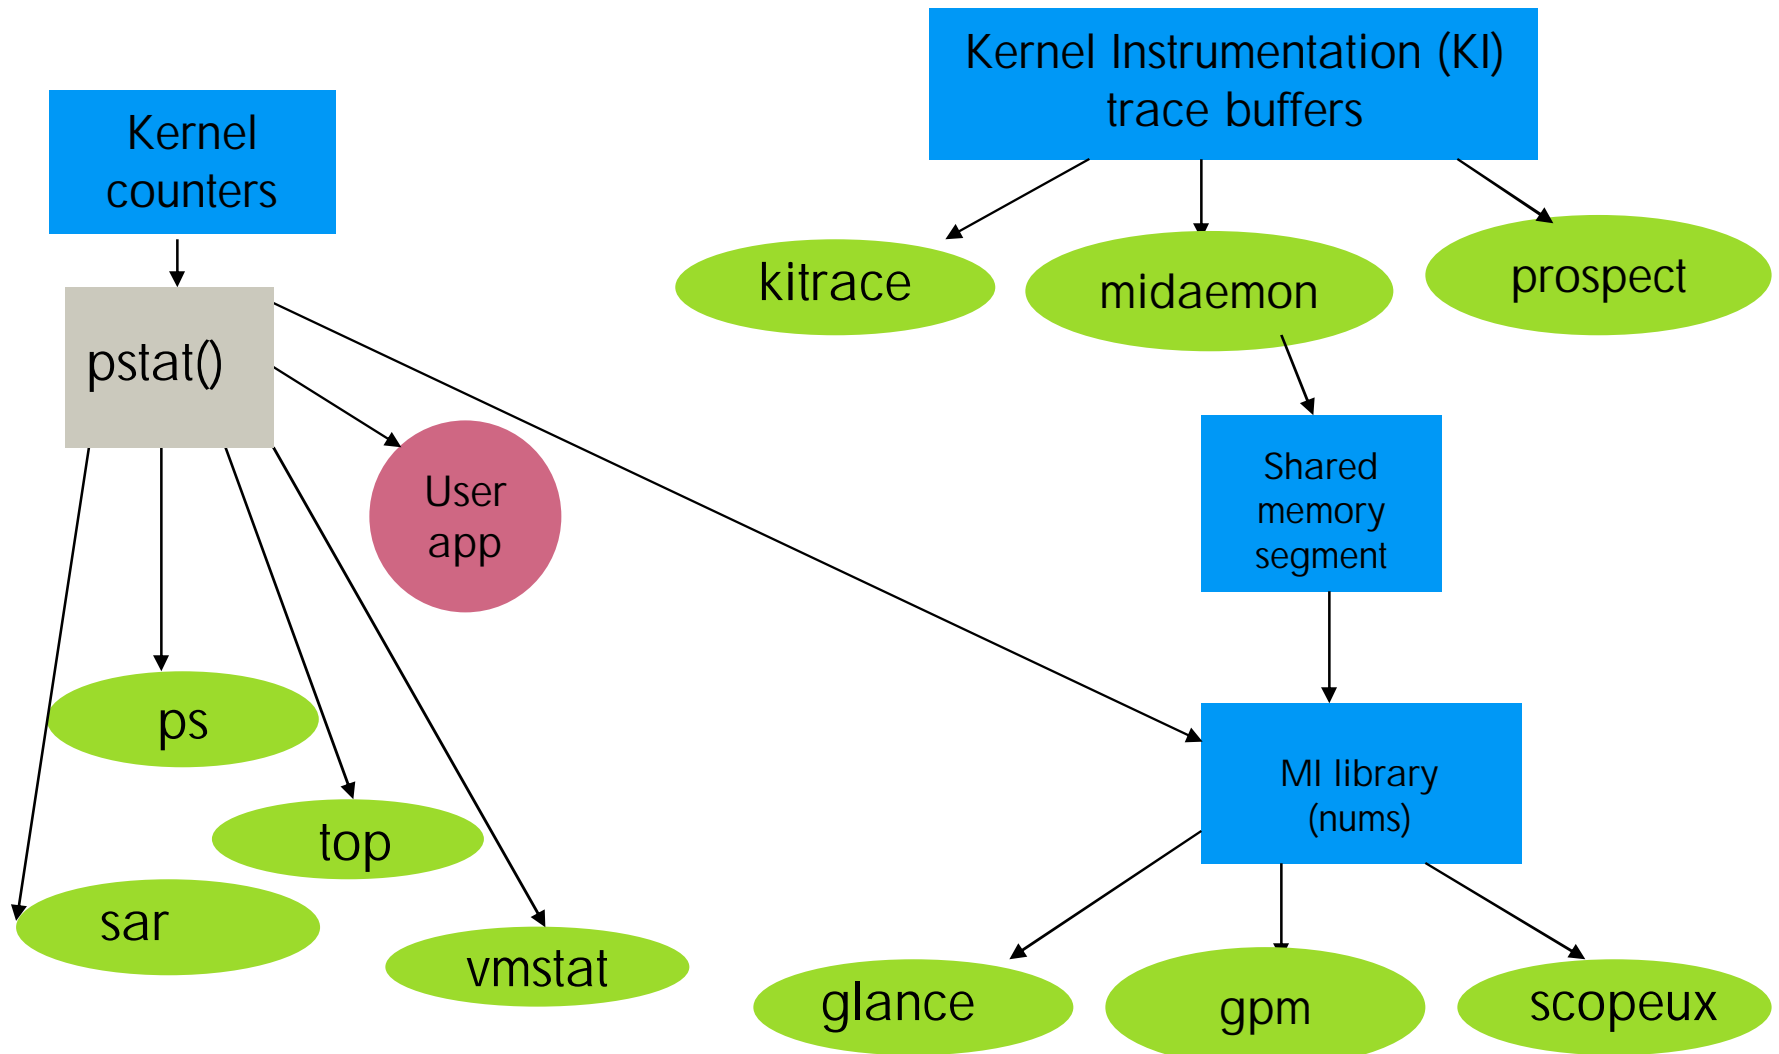
Agenda

- Explain how HP-UX performance metrics are collected
- Discuss tools in the context of CPU, memory and I/O performance
 - sar/vmstat/iostat/ps/uptime
 - Glance/GPM/OVPM
 - Caliper/Prospect
 - pstack
 - tusc/lsof

Tool and metric orientation

- Global performance tools
 - Needed for historical trending and capacity planning
 - Will show the obvious problems
- Lower level tools
 - Needed for troubleshooting non-obvious problems
 - They aren't always "performance" tools

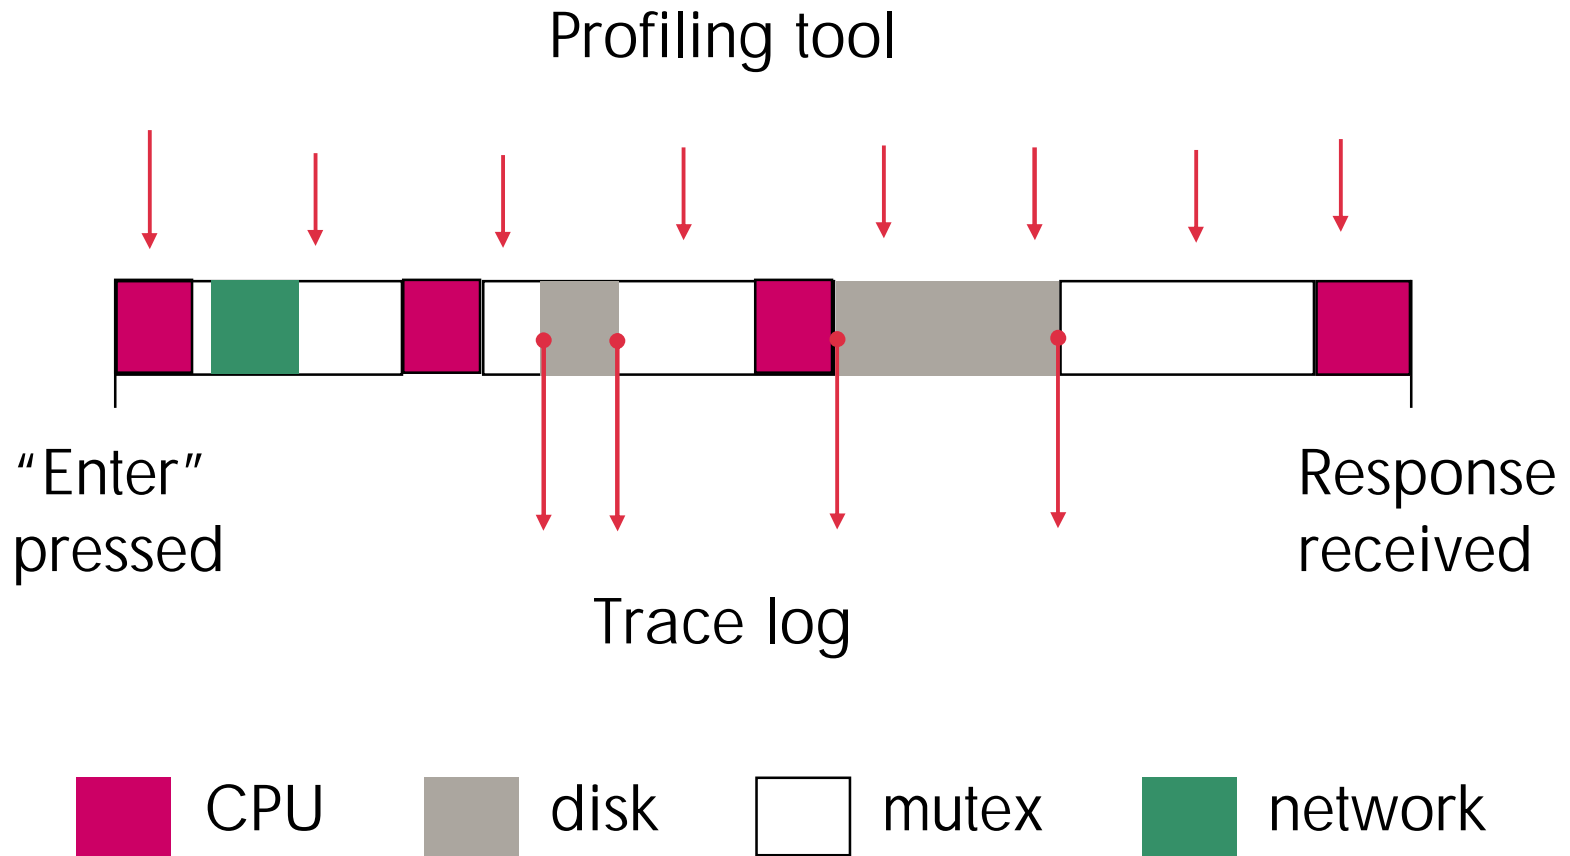
Performance tools data sources



Types of performance metrics

- Profile
 - Build up a picture of behavior over time
 - Cheap but not always accurate
 - Sample must be proportional to the problem
- Traces
 - Instrument interesting events
 - e.g. count number of occurrences, or time duration
 - Accurate
 - May introduce significant overhead

Profiling vs tracing



Numbers to notice

- CPU utilization > 80%
- System CPU > 20%
- Processor load unbalanced
- Load Average > 5
- Disk %busy > 60%
- Disk queueing > 5
- Disk wait/service times > 20ms
- Disks unbalanced
- Pageouts/Deactivations > 0

Behaviors to notice

- High fork()/exit() rate
- Busiest process
- Processes with big memory footprints
- High system memory utilization
- Long running system calls

Glance/GPM

- Included with HP-UX 11i Enterprise & Mission-Critical Operating Environments (OE)
- 60 day trial licenses available for older releases
- GPM is the graphical interface
- 'advisor mode' allows scripted, non-terminal interface for data collection
- Glance/GPM output can be saved to file or printed
- Fully supported by HP
- Safe and non-intrusive in a production environment
- Information extracted from the kernel via the measurement interface (MI) using midaemon and pstat() call

Reflection for HP

HP veloce.r1w - Reflection for HP

File Edit Connection Setup Macro Script Window Help

GlancePlus C.04.55.00 15:45:36 veloce ia64 Current Avg High

CPU Util	S	SU	U		21%	21%	31%
Disk Util	F				100%	98%	100%
Mem Util	S		SU		70%	69%	70%
Swap Util	U	UR			11%	11%	11%

PROCESS LIST Users= 3

Process Name	PID	User Name	CPU % (1600% max)	Thrd Cnt	Disk IOrate	Memory RSS/VSS	Block On
dbreaderb	4122	kenj	79.2	1	0.0	332kb 468kb	SLEEP
dbreaderb	4121	kenj	77.4	1	0.0	332kb 560kb	SLEEP
dbreadera	4117	kenj	76.8	1	0.0	328kb 468kb	OTHER
dbreadera	4118	kenj	76.8	1	0.0	328kb 468kb	SLEEP
midaemon	1614	root	7.7	18	0.0	135.5mb 139.6mb	SLEEP
vhand	2	root	6.7	1	0.4	64kb 72kb	SLEEP
logwriterb	4114	kenj	0.6	1	71.2	380kb 736kb	CACHE
logwritera	4113	kenj	0.6	1	72.0	380kb 740kb	CACHE
dbwriterb	4120	kenj	0.2	1	22.2	332kb 468kb	CACHE
dbwritera	4116	kenj	0.2	1	28.9	332kb 468kb	CACHE
dbwritera	4115	kenj	0.2	1	29.1	332kb 468kb	IO
dbwriterb	4119	kenj	0.0	1	22.2	332kb 468kb	CACHE
cimprovagt	1489	root	0.0	14	0.0	20.2mb 51.2mb	SLEEP
vxsvc	1875	root	0.0	13	0.0	26.2mb 67.5mb	SLEEP

C - cum/interval toggle Page 1 of 2

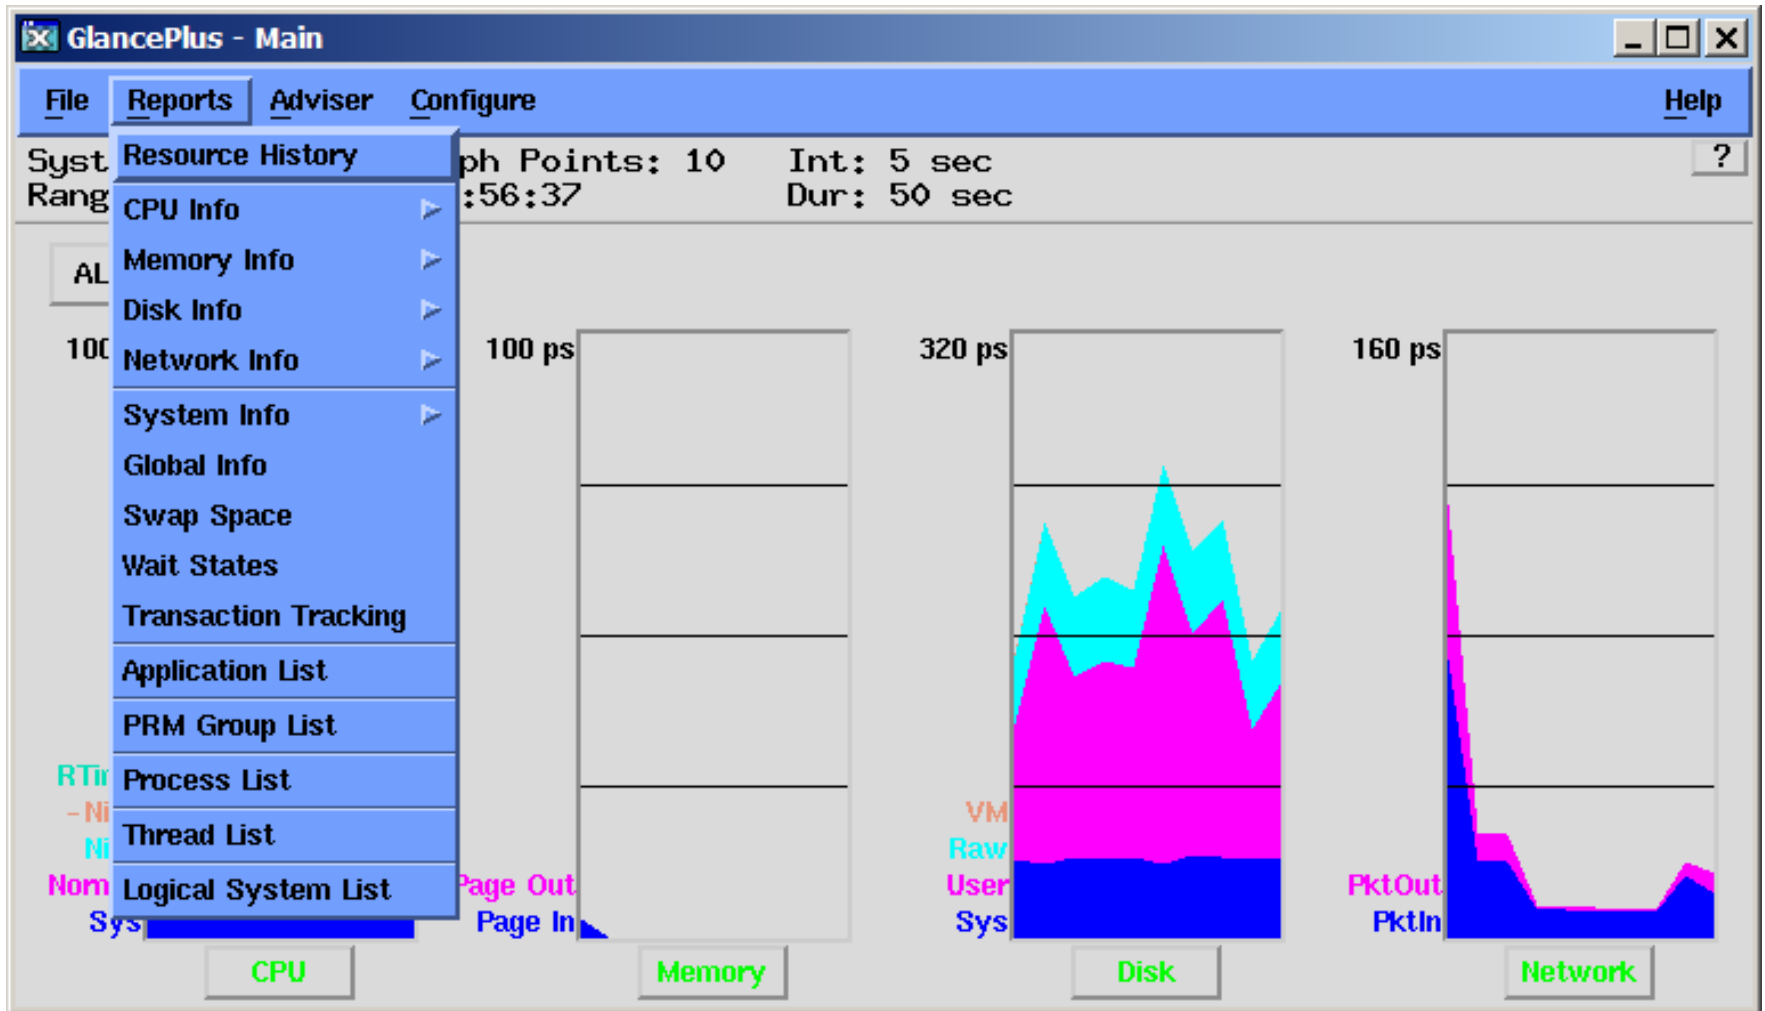
Process List CPU Report Memory Report Disk Report Next Keys Select Process Help Exit Glance

875, 1 HP70094 -- veloce.cup.hp.com via TELNET 02:54:11 Num

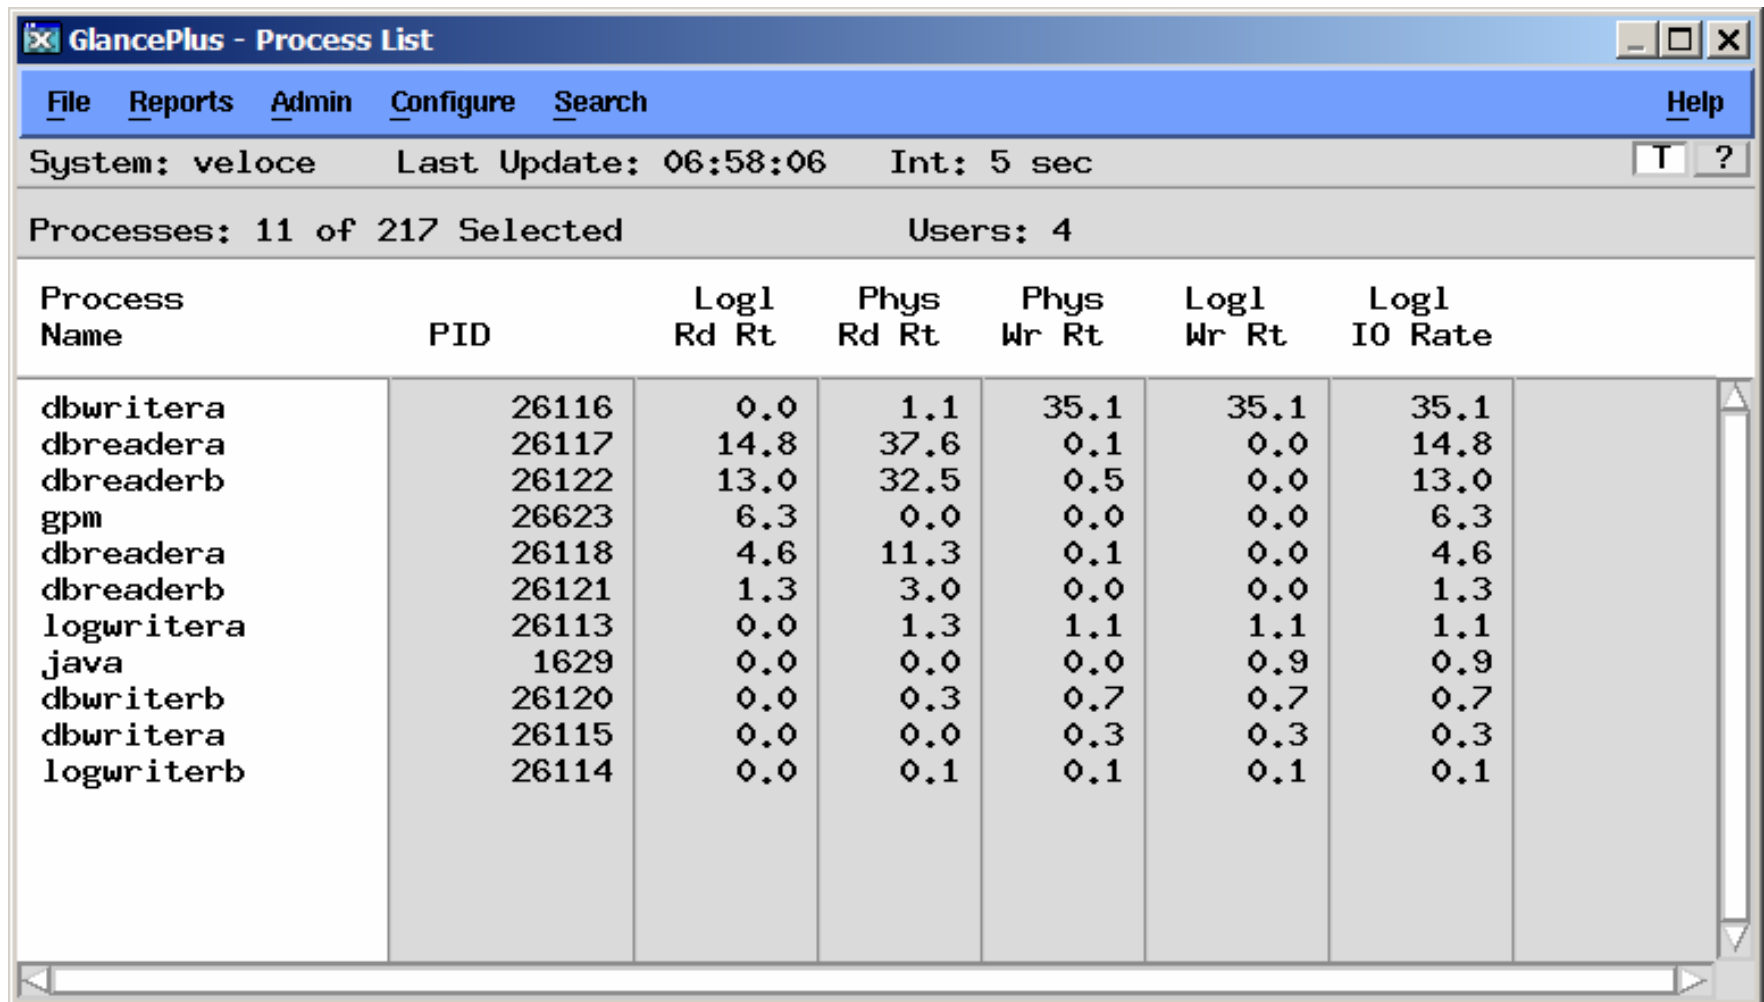
Glance commands menu

h - Online Help	q - exit (or e)	A - Application List
g - Process List	d - Disk Report	P - PRM Group List
a - CPU By Processor	i - IO By File System	Y - Global System Calls
c - CPU Report	u - IO By Disk	F - Process Open Files
m - Memory Report	v - IO By Logical Volume	M - Process Memory Regions
t - System Tables	N - NFS Global Activity	R - Process Resources
w - Swap Space	n - NFS By System	W - Process Wait States
B - Global Waits	l - Network By Interface	L - Process System Calls
Z - Global Threads	T - Trans Tracker	y - Renice Process
G - Process Threads	H - Alarm History	s - Select Process
I - Thread Resource	J - Thread Wait	
S - Select Disk/NFS/Appl/Trans/Thread		
? - Commands Menu	b - Page Backward (or -)	< - Display Previous Screen
! - Invoke Shell	f - Page Forward (or +,space)	> - Display Next Logical Scr
h - Online Help	q - exit (or e)	z - Reset Statistics to Zero
p - Print Toggle	r - Refresh Screen (or ^L)	<cr>- Update Current Screen
j - Adjust Interval	o - Threshold Screen Options	

GPM



GPM Process List



The screenshot shows a window titled "GlancePlus - Process List". The menu bar includes "File", "Reports", "Admin", "Configure", "Search", and "Help". The status bar shows "System: veloce", "Last Update: 06:58:06", and "Int: 5 sec". Below this, it indicates "Processes: 11 of 217 Selected" and "Users: 4". The main content is a table with the following columns: Process Name, PID, LogI Rd Rt, Phys Rd Rt, Phys Wr Rt, LogI Wr Rt, and LogI IO Rate.

Process Name	PID	LogI Rd Rt	Phys Rd Rt	Phys Wr Rt	LogI Wr Rt	LogI IO Rate
dbwritera	26116	0.0	1.1	35.1	35.1	35.1
dbreadera	26117	14.8	37.6	0.1	0.0	14.8
dbreaderb	26122	13.0	32.5	0.5	0.0	13.0
gpm	26623	6.3	0.0	0.0	0.0	6.3
dbreadera	26118	4.6	11.3	0.1	0.0	4.6
dbreaderb	26121	1.3	3.0	0.0	0.0	1.3
logwritera	26113	0.0	1.3	1.1	1.1	1.1
java	1629	0.0	0.0	0.0	0.9	0.9
dbwriterb	26120	0.0	0.3	0.7	0.7	0.7
dbwritera	26115	0.0	0.0	0.3	0.3	0.3
logwriterb	26114	0.0	0.1	0.1	0.1	0.1

Glance adviser scripts

- Run from the command line
- Output to stdout
- Provide easy access to ~2000 metrics
- Metrics listed and documented here:
`/opt/perf/paperdocs/gp/C/metrics.txt`
- Sample scripts here:
`/opt/perf/examples/adviser`
- Run like this:

```
# glance -j 5 -adviser_only -syntax myscript
```

Sample adviser script

```
headersprinted = 0

disk loop {

  if BYDSK_PHYS_IO_RATE > 0 then {
    if headersprinted == 0 then {
      print "      device      %util   queue   r+w/s  blks/s  read/s rblks/s write/s wblks/s avserv"
      headersprinted = 1
    }

    # sar shows average service time in milliseconds:
    avserv = ( BYDSK_UTIL / 100 ) / BYDSK_PHYS_IO_RATE * 1000

    # sar blks/s is 512-byte blocks per second (KB rate times 2):
    blks = BYDSK_PHYS_BYTE_RATE * 2
    readblks = BYDSK_PHYS_READ_BYTE_RATE * 2
    writeblks = BYDSK_PHYS_WRITE_BYTE_RATE * 2

    print GBL_STATTIME, "      ",BYDSK_DEVNAME|15, BYDSK_UTIL|7|2,
          BYDSK_REQUEST_QUEUE|8|2,
          BYDSK_PHYS_IO_RATE|8|0, blks|8|0,
          BYDSK_PHYS_READ_RATE|8|0, readblks|8|0,
          BYDSK_PHYS_WRITE_RATE|8|0, writeblks|8|0,
          avserv|7|2
  }
}

if headersprinted == 0 then
  print GBL_STATTIME, "      (no disk activity this interval)"
```

Sample adviser script output

```
# glance -j5 -adviser_only -syntax coldsk
```

```
Welcome to GlancePlus
```

	device	%util	queue	r+w/s	blks/s	read/s	rblks/s	write/s	wblks/s	avserv
13:41:04	0/0/0/3/0.6.0	0.76	0.00	1	19	0	0	1	19	5.85
13:41:09	0/0/0/3/0.6.0	1.11	0.00	1	22	0	0	1	22	8.54
13:41:14	0/0/0/3/0.6.0	0.95	0.00	3	37	0	0	3	37	3.80
13:41:19	0/0/0/3/0.6.0	0.19	0.00	0	3	0	0	0	3	19.00
13:41:24	0/0/0/3/0.6.0	0.44	0.00	0	7	0	0	0	7	11.00
13:41:29	0/0/0/3/0.6.0	1.91	0.00	3	46	0	0	3	46	6.37
13:41:29	1/0/0/3/0.6.0	0.19	0.00	0	0	0	0	0	0	19.00
13:41:34	(no disk activity this interval)									

CPU performance



CPU bottlenecks

- We have a CPU bottleneck if our application spends most of its time...
 - Waiting to execute on the CPU, or ..
 - Executing on the CPU
- If executing we need to know
 - Where CPU cycles are consumed
 - User space code or system/kernel code
 - How CPU cycles are consumed
 - Useful productive work or spinning for a resource

Waiting to execute ... uptime & sar

```
# uptime
```

```
2:30pm up 14days, 2:39, 33 users, load average: 1.71, 1.88, 1.80
```

```
# sar -q 3 3
```

```
HP-UX veloce B.11.31 U ia64 07/16/07
```

13:01:49	runq-sz	%runocc	swpq-sz	%swpocc
13:01:52	0.0	0	0.0	0
13:01:55	0.0	0	0.0	0
13:01:58	0.0	0	0.0	0
Average	0.0	0	0.0	0

Waiting to execute ...

Glance CPU report by processor

HP veloce.r1w - Reflection for HP

File Edit Connection Setup Macro Script Window Help

GlancePlus C.04.55.00 13:08:59 veloce ia64 Current Avg High

CPU Util		0%	0%	1%
Disk Util		0%	1%	7%
Mem Util	S SU U	25%	25%	25%
Swap Util	U UR	10%	10%	10%

CPU BY PROCESSOR Users= 1

CPU	State	Util	LoadAvg(1/5/15 min)			CSwitch	Last Pid
16	Disable	0.0	0.0/	0.0/	0.0	0	7126
8	Enable	0.0	0.0/	0.0/	1.3	10	89
24	Disable	0.0	0.0/	0.0/	0.0	0	91
0	Enable	0.4	0.0/	0.0/	0.0	31	89
7	Enable	0.2	0.0/	0.0/	0.3	20	91
1	Enable	0.0	0.0/	0.0/	0.0	21	87
12	Disable	0.0	0.0/	0.0/	0.0	0	89
19	Enable	0.2	0.0/	0.0/	0.0	69	2384
5	Enable	0.0	0.0/	0.0/	0.3	10	89
3	Enable	0.0	0.0/	0.0/	0.0	45	89
18	Enable	0.0	0.0/	0.0/	0.0	13	87
11	Disable	0.0	0.0/	0.0/	0.0	0	87
15	Disable	0.0	0.0/	0.0/	0.0	0	1614
14	Disable	0.0	0.0/	0.0/	0.0	0	89

Page 2 of 6

IO By File Sys IO By Disk IO By Logl Vol Swap Space Next Keys CPU By Processr Alarm History Select

771, 1 HP70094 -- veloce.cup.hp.com via TELNET 00:17:35 Num

Waiting to execute ...

Thread wait states

The screenshot shows the 'glanceplus' utility interface. At the top, it displays system information: 'GlancePlus C.03.85.00', time '14:16:54', host 'pcj18e', architecture 'ia64', and current user 'root'. Below this is a table of resource utilization:

Resource	Current	Avg	High
Cpu Util	100%	100%	100%
Disk Util	1%	1%	2%
Mem Util	12%	12%	12%
Swap Util	10%	10%	10%

Below the utilization table, it shows 'Wait States' for TID: 10663, newtm.32, PID: 6990, euid: 0, User: root. A table of wait states is displayed:

Event	%	Blocked On	%
IPC	0.0	Cache	0.0
Job Control	0.0	CDROM IO	0.0
Message	0.0	Disk IO	0.0
Pipe	0.0	Graphics	0.0
RPC	0.0	Inode	0.0
Semaphore	0.0	IO	0.0
Sleep	5.0	LAN	0.0
Socket	0.0	NFS	0.0
Stream	0.0	Priority	83.1
Terminal	0.0	System	0.0
Other	0.0	Virtual Mem:	0.0

The 'Priority' row is circled in red. At the bottom of the window, there are control buttons: 'Process Threads', 'Thread Resource', 'Thread Wait', 'Next Keys', 'Help', and 'Select'. The status bar at the bottom shows '63, 1', 'HP70092 -- pcj18e.atl.hp.com via TELNET', '00:04:51', and 'Num'.

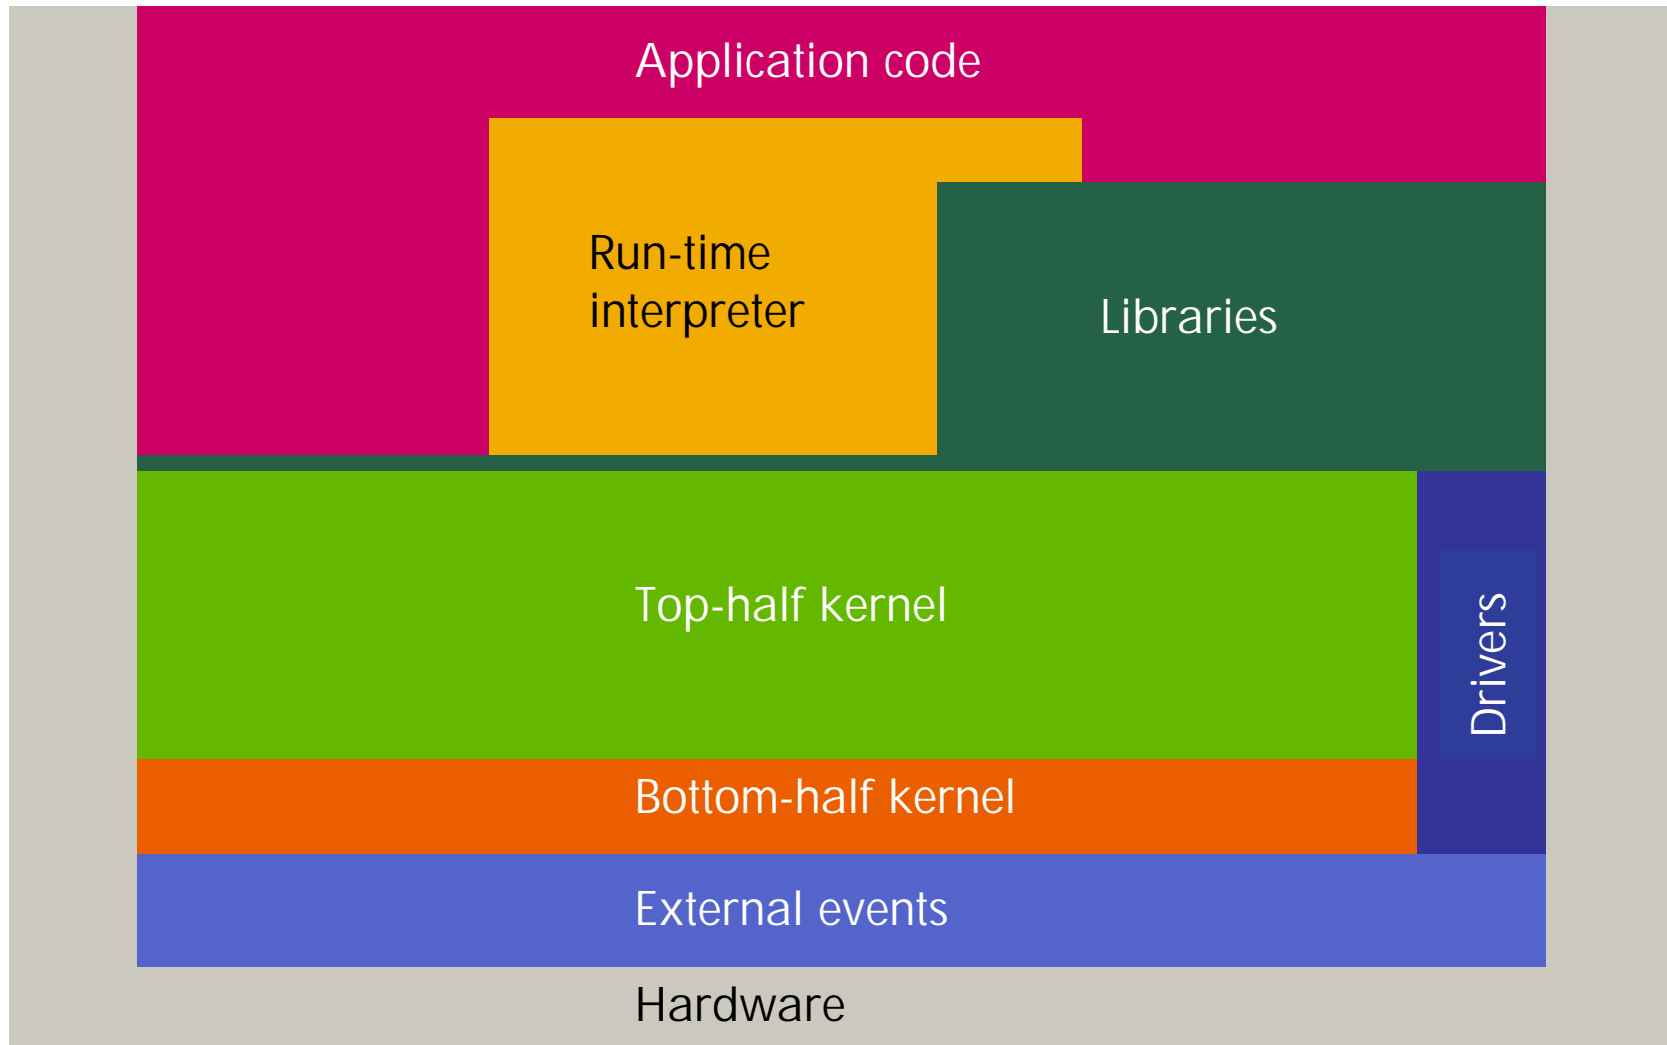
CPU execution ...

ps

```
# ps -ef
```

UID	PID	PPID	C	STIME	TTY	TIME	COMMAND
root	0	0	0	Jul 15	?	0:03	swapper
root	1	0	0	Jul 15	?	0:01	init
root	8	0	0	Jul 15	?	0:01	ksyncer_daemon
root	7	0	0	Jul 15	?	0:00	ioconfigd
root	4	0	0	Jul 15	?	0:00	unhashdaemon
root	73	0	0	Jul 15	?	0:00	awd
root	93	0	0	Jul 15	?	0:00	sblksched
root	92	0	0	Jul 15	?	0:00	sblksched
root	91	0	0	Jul 15	?	1:13	pagezerod
root	90	0	0	Jul 15	?	2:21	cmcd
root	87	0	0	Jul 15	?	12:50	vxfsd
root	80	0	0	Jul 15	?	0:00	ipmid

When it's running, what's it doing?



CPU execution ... Glance CPU report

HP veloce.r1w - Reflection for HP

File Edit Connection Setup Macro Script Window Help

GlancePlus C.04.55.00 13:06:07 veloce ia64 Current Avg High

CPU Util				0%	0%	0%
Disk Util	F			1%	2%	7%
Mem Util	S	SU	U	25%	25%	25%
Swap Util	U	UR		10%	10%	10%

CPU REPORT						Users= 1
State	Current	Average	High	Time	Cum Time	
User	0.0	0.0	0.0	0.00	0.00	
Nice	0.0	0.0	0.0	0.00	0.00	
Negative Nice	0.0	0.0	0.0	0.00	0.00	
RealTime	0.0	0.0	0.0	0.00	0.00	
System	0.0	0.1	0.2	0.00	0.05	
Interrupt	0.0	0.0	0.0	0.00	0.02	
ContextSwitch	0.0	0.0	0.0	0.00	0.00	
Traps	0.0	0.0	0.0	0.00	0.00	
Vfaults	0.0	0.0	0.0	0.00	0.00	
Idle	100.0	99.8	100.0	4.55	41.91	

Top CPU user: PID 87, vxfsd 0.2% cpu util
Active CPUs: 16

Page 1 of 2

Process List CPU Report Memory Report Disk Report Next Keys Select Process Help Exit Glance

773, 1 HP70094 -- veloce.cup.hp.com via TELNET 00:14:45 Num

CPU execution ... Glance process resource

```
File Edit Connection Setup Macro Window Help
[Icons]
GlancePlus C.04.50.00          09:15:56  snappy  ia64  Current Avg High
-----
CPU Util  S  SN  | 13% 13% 13%
Disk Util FF  | 3% 3% 3%
Mem Util  S  SUUB B  | 34% 34% 34%
Swap Util U  UR  | 11% 11% 11%
-----
Resources PID: 21510, reader      PPID: 5087  euid: 0  User: root
-----
CPU Usage (util): 66.7 Log Reads : 0 Wait Reason : MESSG
User/Nice/RT CPU: 4.0 Log Writes: 0 Total RSS/VSS : 1.9mb/ 1.9mb
System CPU : 62.0 Phy Reads : 0 Traps / Vfaults: 4/ 65160
Interrupt CPU : 0.0 Phy Writes: 0 Faults Mem/Disk: 0/ 0
Cont Switch CPU : 0.0 FS Reads : 0 Deactivations : 0
Scheduler : HPUX FS Writes : 0 Forks & Vforks : 0
Priority : 155 VM Reads : 0 Signals Recd : 0
Nice Value : 24 VM Writes : 0 Mesg Sent/Recd : 0/ 0
Dispatches : 137 Sys Reads : 0 Other Log Rd/Wt: 0/ 0
Forced CSwitch : 9 Sys Writes: 0 Other Phy Rd/Wt: 0/ 0
VoluntaryCSwitch: 64 Raw Reads : 0 Proc Start Time
Running CPU : 3 Raw Writes: 0 Wed May 2 09:15:47 2007
CPU Switches : 0 Bytes Xfer: 0kb
Argv1: 9998
_ C - cum/interval toggle % - pct/absolute toggle Page 1 of 1
Process Wait Memory Open Next Process
Resource States Regions Files Keys Syscalls
71, 1 HP70092 -- snappy.cup.hp.com via TELNET 20:52:12 Num
```

CPU execution ...

Glance process syscalls

GlancePlus C.04.50.00 00:46:33 snappy ia64 Current Avg High

CPU Util	S	SN							
Disk Util									
Mem Util	S	SUUB	B						
Swap Util	U	UR							

System Calls PID: 3339, reader PPID: 5087 euid: 0 User: root

System Call Name	ID	Count	Rate	Elapsed Time	Cum Ct	CumRate	Elapsed CumTime
msgsnd	161	185	35.5	0.00231	4055	33.5	0.05036
msgrcv	162	184	35.3	1.88601	4054	33.5	41.52027
shmget	163	184	35.3	0.00749	4054	33.5	0.16460
shmat	165	184	35.3	0.00095	4054	33.5	0.02100
shmdt	166	185	35.5	0.00116	4055	33.5	0.02585

Cumulative Interval: 121 secs

Page 1 of 1

Process Resource Wait States Memory Regions Open Files Next Keys Process Syscalls

1120, 1 HP70092 -- snappy.cup.hp.com via TELNET 36:22:46 Num

CPU execution ...

Prospect – PA-RISC processor-based systems

- Prospect provides:
 - System-wide summary
 - Per-process and per-thread data:
 - Execution/blocked
 - Blocked where and for how long
 - Userspace (application and library) and kernel profile
 - Even instruction-level profiling!
 - No special build required
 - Very low overhead – uses KI tracepoints
- <http://www.hp.com/go/prospect>

CPU execution ...

Prospect

Summary of Operation

Basic Invocation

```
# prospect -P
```

```
# prospect -V 2 -e -f prospect.out <command>
```

Where `-V` is the level of tracing

- V 2 the command

- V 3 the command and descendents

- V 4 all activity

- V 6 process profile

- V kernel

CPU execution ...

Using prospect

- Start the prospect daemon:

```
# prospect -P
```

- Collect data on everything:

```
# prospect -V4 -f pro.out sleep 60
```

- Collect data on program abc and its descendents:

```
# prospect -V3 -f pro.out abc <arguments>
```

- Collect a profile of PIDs 12345 and 67890:

```
# prospect -V6 12345,67890 -f pro.out sleep 60
```

- Collect stack samples of PID every 100ms:

```
# prospect --hprof --sampling-interval=100\  
-sampling-timer=profile -o pro.out -a PID
```


CPU execution ...

Prospect output – execution/blocked time

```

=====
bvsmgr          (25664)NM Birth (pre-exist) ps:= (bvsmgr -I 83 -a abcdef)
                Process did not END - use last SLEEP
-----
Section 1.0: Sorted Execution+Blocked=Elapsed Clocks for this Process, EPB
-----
Clock           Type           Seconds % Elapsed uSec/Count  Counts

ELAPSED         TOTAL           300.957616 100.000%   Total
BLOCKED         TOTAL           285.537890  94.876%   80888.9   3530
Blocked         misc            284.914237  94.669%   165551.6  1721
EXECUTION       TOTAL           15.419727   5.124%   Total
Execution       User            14.464175   4.806%   85.9     168334
Execution       Sys             0.784501    0.261%   4.9     158705
Blocked         PreempT        0.623653    0.207%   344.7    1809
Execution       IntUser        0.136142    0.045%   26.7     5099
Execution       CsW            0.019822    0.007%   5.7     3508
Execution       TraP           0.006761    0.002%   7.1     958
Execution       IntSys         0.005371    0.002%   23.4     230
Execution       Vfault        0.002956    0.001%   46.2     64

```

CPU execution ...

Prospect userspace profile

Section 2.1.5: Thread Ktid(5711) Intersections with 100Hz System Clock, UK

USER portion of profile:

581 hits = 5.81 seconds. KTC time was 6.423 seconds

pcnt	accum%	Hits	Secs	Address	Routine name	Filename
8%	8%	49	0.49	setWriteLockOwner		/distrib/hpux/lib/libsysssrv.sl
4%	12%	21	0.21	mutexUnlock		/distrib/hpux/lib/libsysssrv.sl
3%	15%	17	0.17	pthread_mutex_lock		/usr/lib/libpthread.1
3%	18%	15	0.15	statsEnabled		/distrib/hpux/lib/libsysssrv.sl
2%	20%	13	0.13	memLogUpdate		/distrib/hpux/lib/libsysssrv.sl
2%	22%	11	0.11	dtSERootFv		/test/beef
2%	24%	11	0.11	getTD		/distrib/hpux/lib/libobjsrv.sl
2%	25%	10	0.10	clearWriteLockOwner		/distrib/hpux/lib/libsysssrv.sl
2%	27%	10	0.10	ctProxyObjectFv		/distrib/hpux/lib/libobjsrv.sl
2%	29%	9	0.09	pthread_mutex_unlock		/usr/lib/libpthread.1
2%	30%	9	0.09	ctSERootFv		/distrib/hpux/lib/libbeef_impl.sl
2%	32%	9	0.09	collectLockStats		/distrib/hpux/lib/libsysssrv.sl
1%	33%	8	0.08	asSWSMOffset		/test/beef
1%	34%	8	0.08	pthread_mutex_trylock		/usr/lib/libpthread.1
1%	36%	8	0.08	__spin_lock		/usr/lib/libpthread.1

CPU execution...

Prospect kernel profile

KERNEL portion of profile:

121 hits = 1.21 seconds. KTC time was 0.448 seconds

pcnt	accum%	Hits	Secs	Routine name	Filename
26%	26%	32	0.32	_lwp_mutex_lw_unlock	/stand/vmunix
12%	38%	14	0.14	spinunlock	/stand/vmunix
11%	49%	13	0.13	syscall	/stand/vmunix
7%	55%	8	0.08	No-symbols, 4 unique Instructions	/stand/vmunix
5%	60%	6	0.06	lock_done	/stand/vmunix
5%	65%	6	0.06	lock_write	/stand/vmunix
4%	69%	5	0.05	\$syscallrtn	/stand/vmunix
3%	73%	4	0.04	syscallinit	/stand/vmunix
3%	76%	4	0.04	resume_cleanup	/stand/vmunix
3%	79%	4	0.04	copyin	/stand/vmunix
3%	83%	4	0.04	ki_accum_push_TOS	/stand/vmunix
2%	84%	2	0.02	\$syscallexit	/stand/vmunix
2%	86%	2	0.02	spluser	/stand/vmunix
2%	88%	2	0.02	thread_mutex_init	/stand/vmunix
1%	88%	1	0.01	setjmp	/stand/vmunix
1%	89%	1	0.01	ulbcopy_pcxu_method	/stand/vmunix

CPU execution ...

Caliper – Itanium-based systems only

- Performance [execution] measurement & analysis tool
- Itanium native code
- HP-UX and Linux
- Per-process (user + kernel), system-wide (kernel)
- Uses the Performance Monitor Unit (PMU) hardware and/or dynamic instrumentation [HP-UX only] as needed.
- Measures any application without special build
- <http://www.hp.com/go/caliper>

CPU execution ...

Using caliper

- Get help:

```
# caliper -H
```

- Collect a flat profile for specified command, show individual threads:

```
# caliper fprof -o fprof.out -t ./command arg1 arg2
```

- Collect a flat profile, assembly-level detail:

```
# caliper fprof -o fprof.out -r statement:instruction ./command arg1 arg2
```

- Attach to process 1234 to collect a sampled call graph of its behavior in user and kernel code:

```
# caliper scgprof -o fprof.out --event-defaults PLM=all -a 1234 -e 30
```

- Generate a 20-second kernel profile:

```
# caliper fprof -o kern.out --scope kernel -e 20
```

- Monitor CPU stats for mybm command; include help:

```
# caliper cpu --info -o cpu.out ./mybm
```

- Monitor memory usage for mybm command:

```
# caliper fprof --memory-usage=exit -o mem.out ./mybm
```

- Run caliper advisor on previously collected data:

```
# caliper advise
```

CPU execution ...

Caliper flat profile – load module

```
# caliper fprof -o fprof.out --event-defaults PLM=all --threads=all -e 50 22008
```

Load Module Summary

```
-----
```

% Total	Cumulat	IP	IP	Load Module
Samples	% of	Samples		
	Total			
50.00	50.00	74304		cvpipe
50.00	100.00	74294		libc.so.1
0.00	100.00	4		vmunix
100.00	100.00	148602		Total

```
-----
```

CPU execution ...

Caliper flat profile – function summary

Function Summary

```
-----
```

% Total	Cumulat	IP	Function	File
IP	% of	IP		
Samples	Total	Samples		
50.00	50.00	74304	cvpipe::main	cvpipe.c
50.00	100.00	74294	libc.so.1::pthread_cond_wait	pth_stubs1.c
0.00	100.00	2	vmunix::mpsched_hdlr	mpsched.c
0.00	100.00	1	vmunix::post_hdlr	hl_ivt.c
0.00	100.00	1	vmunix::resume_cleanup	pdk_swch.c

```
-----
```

CPU execution ...

Caliper flat profile – function details

Function Details

```
-----  
% Total          Line|  
  IP            IP   slot| >Statement|  
Samples         Samples Col,Offset Instruction  
-----  
50.00    [cvpipe::main, 0x4000ce0, cvpipe.c]  
          74304          53  Function Totals  
-----  
[ /home/patk/tech_forum/cvpipe.c ]  
          95  >  for (i=0; i<NITERS; i++) {  
          96  >  
          97  >  /* Wait for space in dwnstream que */  
          98  >          pthread_mutex_lock(mtx);  
          99  >          while (qs01 >= MAXQUEUE)  
          74304 100  >          pthread_cond_wait(cv, mtx);  
          101  >  
          102  >          /* Place item in queue */  
          103  >          qs01++;  
          104  >          pthread_cond_signal(cv);  
          105  >          pthread_mutex_unlock(mtx);
```

CPU execution ...

Caliper flat profile – instructions

% Total	Line			
IP	IP	Slot	>Statement	
Samples	Samples	Col,Offset	Instruction	
51.35	[test1::dequeue(queue_t*), 0x4002b80, lffifo.c]			
	4855	~324	Function Totals	

	[/home/col/lockfree/SIAC/lffifo.c]			
	(362)	~324	>{	
	362	~1,0x0000:0	M	adds sp=-64,sp
		:1	M	addp4 r11=0,r32
		:2	I	addp4 r10=8,r32
	(851)	~332	*> mytail = q->tail;	
	374	~3,0x0010:0	M	addp4 r18=16,r32
		:1	I_	nop.i 0 ;;
		:2	I_	adds r9=40,sp ;;
	(2887)	~333	> next = q->hp-	
	>base[myhead.fields.ptr].next;			
	349	~3,0x0020:0	M	ld4 r8=[r11]
		:1	M	ld8 r20=[r10]
		:2	I_	nop.i 0 ;;
	(851)	~332	*> mytail = q->tail;	
	477	~3,0x0030:0	M	ld8 r19=[r18]
		:1	M	st8 [r9]=r20
		:2	I	extr.u r17=r20,48,16

CPU execution

tusc – trace UNIX system calls

- Traces a process and optionally descendents
- Shows syscalls, arguments, return values
- See section 2 of the online manuals
- Safe, but will degrade performance of traced processes
- Useful options:
 - c Summary of syscall counts, errors and CPU time
 - ccc CPU time for every syscall
 - T %T Print a timestamp before every trace
 - rall Display read buffer for all reads
 - wall Display write buffer for all writes
 - f Follow fork()s

CPU execution ...

tusc

```
# tusc -vall cp /etc/group /dev/null
{420815} execve("/usr/bin/cp", 0x680f083c, 0x680f084c) ..... [entry]
  argv[0] @ 680f0014: "cp"
  argv[1] @ 680f0017: "/etc/group"
  argv[2] @ 680f0022: "/dev/null"
{420815} execve("/usr/bin/cp", 0x680f083c, 0x680f084c)..... = 0
{420815} stat64("/etc/group", 0x40012128) ..... = 0
{420815} stat64("/dev/null", 0x400121a0) ..... = 0
{420815} open("/etc/group", 0x800, 0322) ..... = 5
{420815} open("/dev/null", O_WRONLY|O_CREAT|O_TRUNC|0x800, 0)..... = 6
{420815} read(5, "r o o t : : 0 : r o o t \n o t h ".., 65536)..... = 210
{420815} write(6, "r o o t : : 0 : r o o t \n o t h ".., 210) ..... = 210
{420815} read(5, 0x40001f38, 65536) ..... = 0
{420815} close(5) ..... = 0
{420815} close(6) ..... = 0
{420815} exit(0) ..... WIFEXITED(0)
```

CPU execution ...

Pstack – Itanium-based systems only

- The stack trace into a hot function is often nice to know
- Std on 11.31 provided with linker patch on 11.23 (PHSS_34858+)
- Prints a stack trace for each thread in a process

```
----- lwpid : 64109 -----  
  
0: 60000000c02acc50 : __kakeup() + 0x30 (/usr/lib/hpux32/libc.so.1)  
1: 60000000c038fae0 : __mxn_wakeup() + 0x740 (/usr/lib/hpux32/libpthread.so.1)  
2: 60000000c03291e0 : pthread_cond_signal() + 0x480 (/usr/lib/hpux32/libpthread.so.1)  
3: 60000000c0328dc0 : pthread_cond_signal() + 0x60 (/usr/lib/hpux32/libpthread.so.1)  
4: 0000000004001790 : run_mid_thread() + 0x1e0 (./cvpipe)  
5: 60000000c0338d20 : __pthread_bound_body() + 0x170 (/usr/lib/hpux32/libpthread.so.1)
```

How to approach high %sys

- Identify the costly system calls – glance/tusc
- Identify the hot kernel functions – caliper/prospect
- Use syscall, function & lock names as keywords to search for solutions in the support database
- Try to compare a 'good' and 'bad' performance period

Summary

- When it's running, what's it doing?
 - Caliper fprof, Prospect, Glance, pstack, tusc, application logging
- When it's waiting, what's it waiting for?
 - Glance, Prospect, pstack, application logging (even better with timestamps)

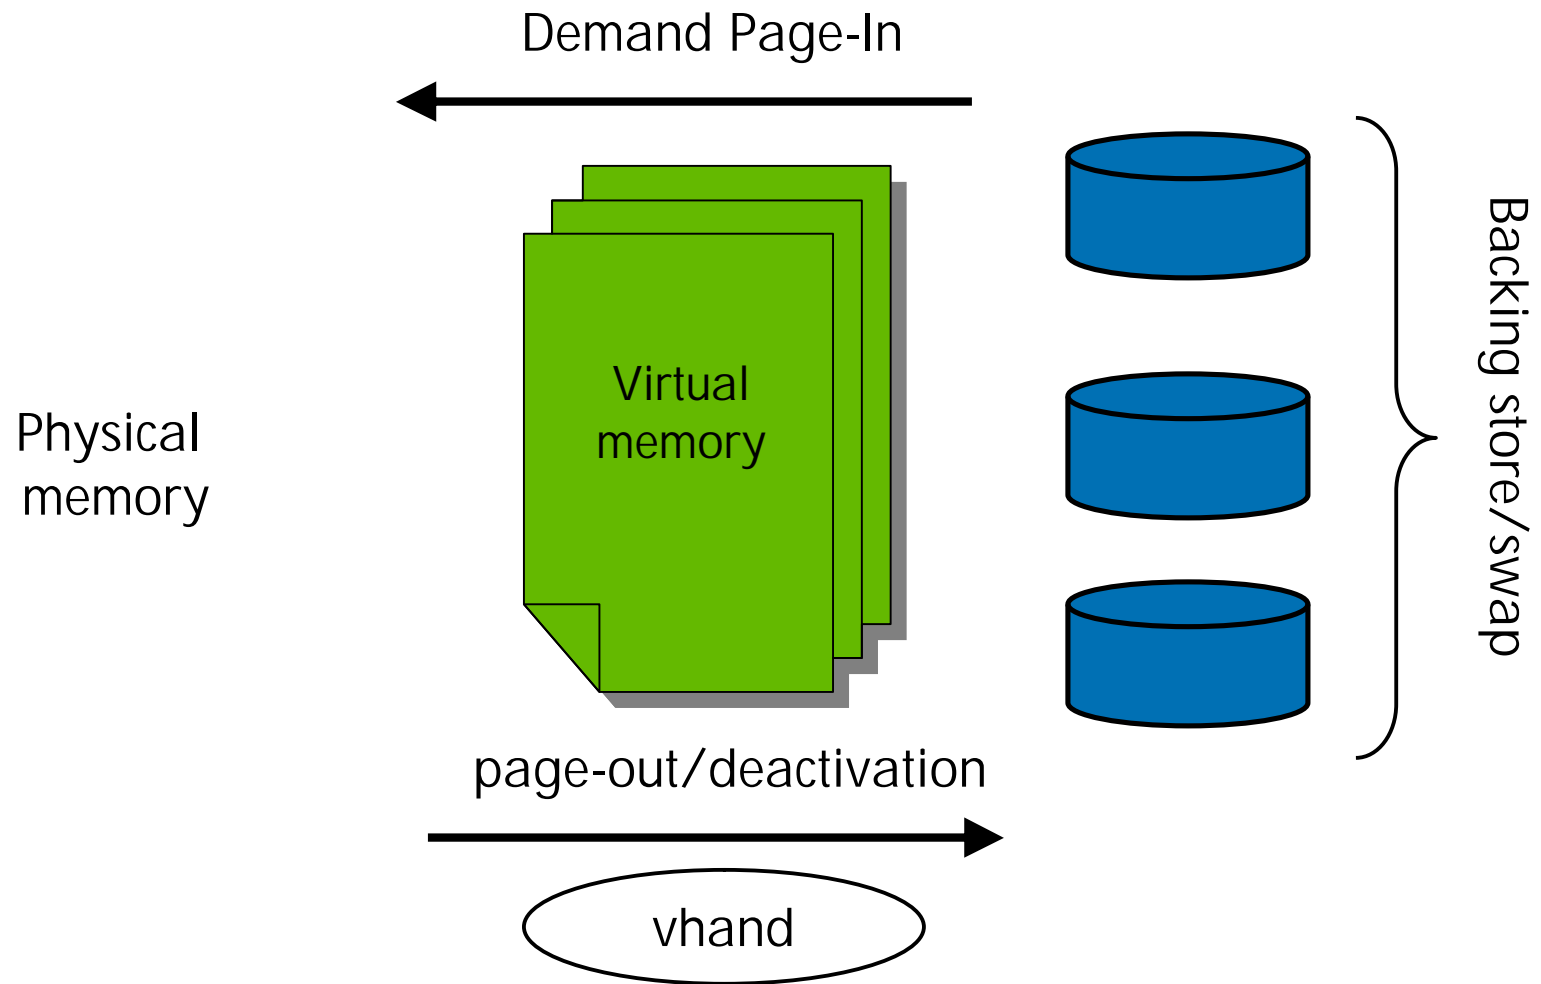
Memory performance



It generally comes down to...

- Am I waiting for memory?
 - Waiting for pages to be allocated
 - Waiting for pages to be brought in from binary
 - Waiting for pages to be brought in from swap
 - Waiting for slow memory
- What is occupying most memory?
 - Kernel or user-space?
 - Which kernel component?
 - Is there a memory leak?
 - Can we tune?

Virtual memory



Is there memory pressure?

- Glance
 - Memory Report (m)
 - Page In/Page Out/Reactivations/Deactivations
 - Memory totals at bottom of the screen
 - Select Process (s) – vhand
 - CPU Usage
 - VM Writes
 - IO by Logical Volume (v) – check the swap LVOL
- vmstat (-s)
- ps -elf – check flags
- sar -w

How much free memory is there?

```
# vmstat
procs
r   b   w   avm   free   re   at   pi   po   fr   de   sr   in   sy   cs   us sy id
1  13   0  42096 695372  3   1   0   0   0   0   0  1202 14517 223  4 27 69
```

Glance memory report

HP **veloce.r1w - Reflection for HP** [-] [] [X]

File Edit Connection Setup Macro Script Window Help

GlancePlus C.04.55.00 13:44:23 veloce ia64 Current Avg High

CPU Util					0%	0%	1%
Disk Util	F				1%	1%	19%
Mem Util	S	SU	U		25%	25%	25%
Swap Util	U	UR			10%	10%	10%

MEMORY REPORT Users= 1

Event	Current	Cumulative	Current Rate	Cum Rate	High Rate
Page Faults	2	58448	0.3	24.9	811.7
Page In	0	19640	0.0	8.3	339.6
Page Out	0	0	0.0	0.0	0.0
KB Paged In	0kb	1.2mb	0.0	0.5	170.6
KB Paged Out	0kb	0kb	0.0	0.0	0.0
Reactivations	0	0	0.0	0.0	0.0
Deactivations	0	0	0.0	0.0	0.0
KB Deactivated	0kb	0kb	0.0	0.0	0.0
VM Reads	0	177	0.0	0.0	12.1
VM Writes	0	0	0.0	0.0	0.0
Total VM :	648mb	Sys Mem : 3.1gb	User Mem: 830mb	Phys Mem : 15.9gb	
Active VM:	489mb	Buf Cache: 0mb	Free Mem: 11.9gb	FileCache: 76mb	

Page 1 of 1

IO By
File Sys

IO By
Disk

IO By
Logl Vol

Swap
Space

Next
Keys

CPU By
Processr

Alarm
History

Select

771, 1 HP70094 -- veloce.cup.hp.com via TELNET 00:52:59 Num

Is my thread waiting for memory?

Glance thread wait states screen

```

B3692A GlancePlus C.03.70.00    15:12:52  bokmaai 9000/820    Current Avg High
-----
CPU Util      ██████████ S | 32% 30% 100%
Disk Util     ██████████ F | 88% 25% 100%
Mem Util      █████ SU UB B | 75% 69% 100%
Swap Util     █████ UR R | 28% 31% 99%
    
```

```

Wait States TID:1497356, memtest    PID: 7237  eid: 0 User: root
    
```

Event	%	Blocked On	%
IPC	0.0	Cache	0.0
Job Control	0.0	CDROM IO	0.0
Message	0.0	Disk IO	0.0
Pipe	0.0	Graphics	0.0
RPC	0.0	Inode	0.0
Semaphore	0.0	IO	0.0
Sleep	27.5	LAN	0.0
Socket	0.0	NFS	0.0
Stream	0.0	Priority	0.0
Terminal	0.0	System	14.2
Other	0.0	Virtual Mem	45.0

█ C - cum/interval toggle % - pct/absolute toggle Page 1 of 1

Process	Thread	Thread		hpterm	Next		Help	Select
Threads	Resource	Wait			Keys			

How much time does my thread wait for memory?

Prospect

Section 1.0: Sorted Execution+Blocked=Elapsed Clocks for this Process, EPB

Clock	Type	Seconds	% Elapsed	uSec/Count	Counts
ELAPSED	TOTAL	55.170400	100.000%	Total	
EXECUTION	TOTAL	54.448278	98.691%	Total	
Execution	User	50.650799	91.808%	5413.1	9357
Execution	Vfault	2.431336	4.407%	39215.1	62
Execution	Sys	1.093497	1.982%	8963.1	122
BLOCKED	TOTAL	0.722122	1.309%	40117.9	18
Blocked	Preempt	0.481112	0.872%	53456.8	9
Blocked	Vmblock	0.447344	0.811%	1331.4	336
Execution	IntUser	0.207806	0.377%	32.4	6406
Execution	IntSys	0.034900	0.063%	68.3	511

How much cpu is vhand consuming?

```

B3692A GlancePlus C.03,70,00      13:12:01  bokmaai 9000/820      Current  Avg  High
-----
CPU Util      S          S          | 23%   11%  100%
Disk Util     V          V          |100%   13%  100%
Mem Util      S          SU         UB         B | 99%   64%  100%
Swap Util     U          UR         R          | 24%   13%   99%
-----

```

```

Resources PID:      2, vhand      PPID:      0 euid:      0 User: root
-----

```

```

CPU Usage (util):      18.2 Log Reads :      0 Wait Reason :      SLEEP
User/Nice/RT CPU:     0.0 Log Writes:     0 Total RSS/VSS :      32kb/   32kb
System CPU           :      17.1 Phy Reads :      0 Traps / Vfaults:     0/      0
Interrupt CPU        :      1.1 Phy Writes:    1343 Faults Mem/Disk:     0/      0
Cont Switch CPU      :      0.0 FS Reads :      0 Deactivations :      0
Scheduler            :      HPUX FS Writes :      0 Forks & Vforks :      0
Priority              :      128 VM Reads :      0 Signals Recd :      0
Nice Value           :      20 VM Writes :    1343 Mesg Sent/Recd :     0/      0
Dispatches           :      38 Sys Reads :      0 Other Log Rd/Wt:     0/      0
Forced CSwitch       :      0 Sys Writes:     0 Other Phy Rd/Wt:     0/      0
VoluntaryCSwitch    :      38 Raw Reads :      0 Proc Start Time
Running CPU          :      0 Raw Writes:     0   Tue Nov  9 15:16:29 2004
CPU Switches        :      3 Bytes Xfer:     0kb

```

```

| C - cum/interval toggle  % - pct/absolute toggle      Page 1 of 1

```

Process	Wait	Memory	Open	hpterm	Next	Process		
Resource	States	Regions	Files		Keys	Syscalls		

How actively is vhand stealing pages?

```
#vmstat -s
```

```
.
```

```
.
```

```
25066 page outs
```

```
102205 pages paged out
```

```
.
```

```
48 revolutions of the clock hand
```

```
121342 pages scanned for page out
```

```
7156 pages freed by the clock daemon
```

```
.
```

```
.
```

How frequently are processes being deactivated?

```
#vmstat -s  
2576 swap ins  
2604 swap outs  
335 pages swapped in  
1067 pages swapped out  
.  
.
```

Which processes are deactivated

ps command:

flags Flags (octal and additive) associated with the
 process:
 0 Swapped
 1 In core

The default heading for this column is F.

F	S	UID	PID	PPID	C	PRI	NI	ADDR	SZ	WCHAN	STIME	TTY	TIME	COMD
1	S	root	2157	1	0	138	0	426fb880	68	408c14c0	11:13:40	?	0:00	/sbin/krsd -i
0	S	root	2158	1	0	154	20	426fbb00	41	9d8f34	11:13:40	?	0:00	/sbin/sfd
1	S	root	2205	1756	0	154	10	4260eac0	136	919280	11:13:46	?	0:01	psmctd

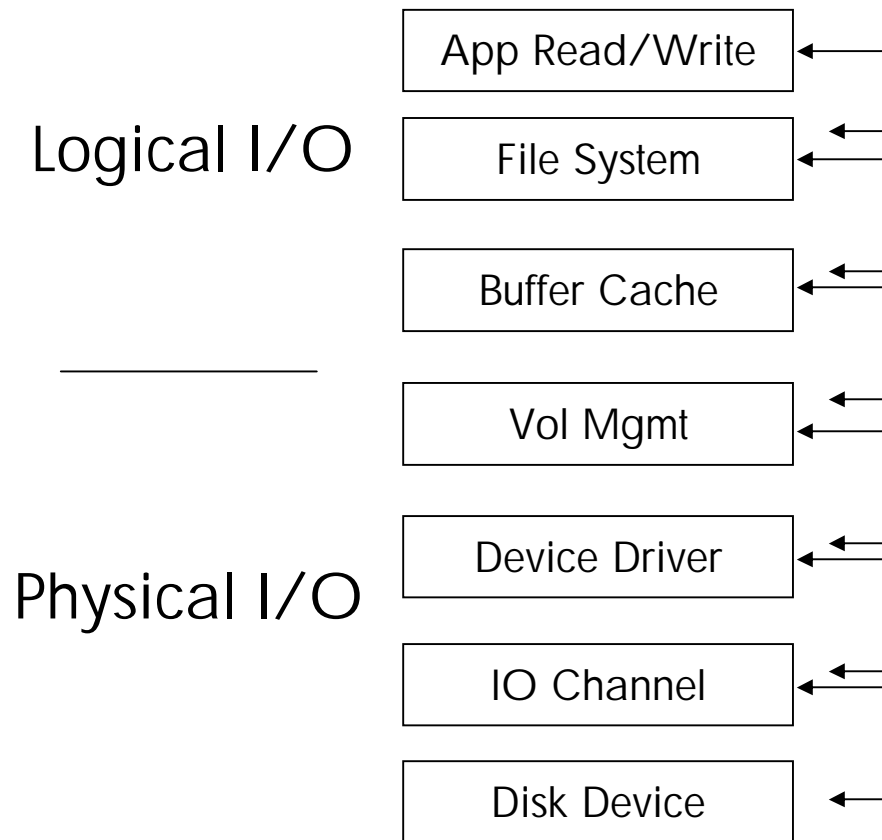
Where did the memory go?

- Glance
 - Memory Report (m) – global info
 - Memory Totals at bottom
 - Process List (g) – per process
 - RSS – Resident Set Size
 - Memory Resources (R) – per process
 - Memory Regions (M) – per process
- ps -elf (SZ field)
- kcusage (dbc_max_pct)
- ipcs -ma
- vxfsstat -b -i/(roughly 2Kb per Inode)

I/O performance



Elements of an I/O request



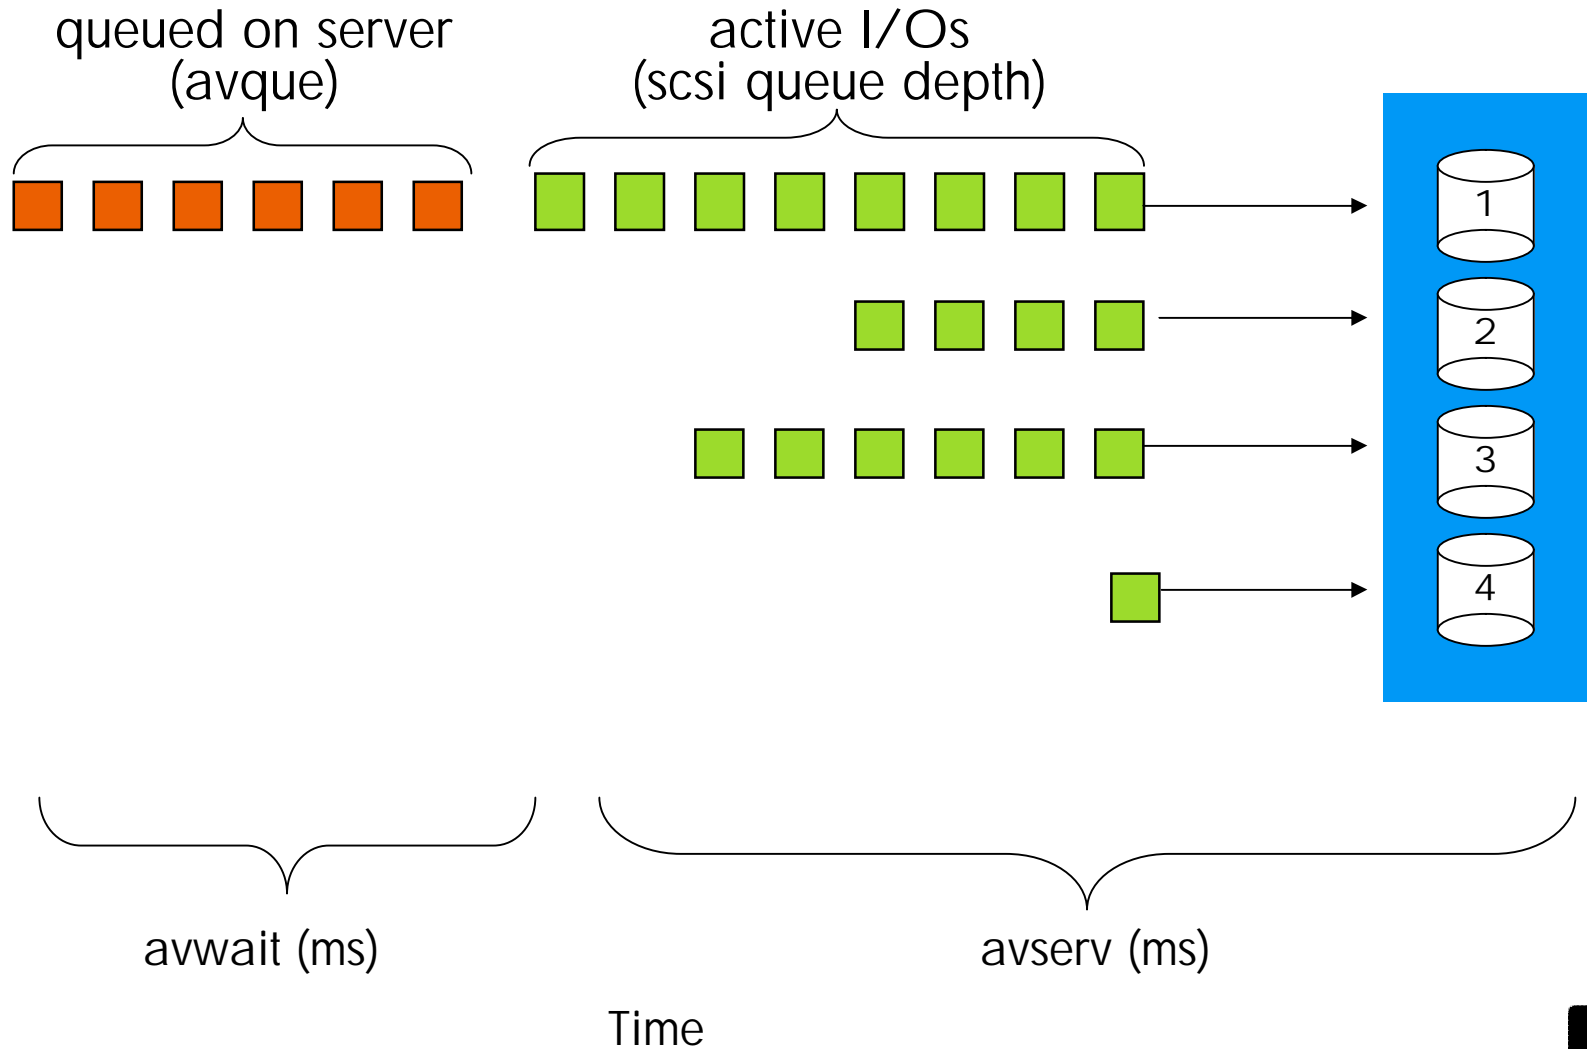
iostat

`iostat` iteratively reports I/O statistics for each active disk on the system. Disk data is arranged in a four-column format:

Column Heading	Interpretation
<code>device</code>	Device name
<code>bps</code>	Kilobytes transferred per second
<code>sps</code>	Number of seeks per second
<code>mtps</code>	Milliseconds per average seek

<code>device</code>	<code>bps</code>	<code>sps</code>	<code>mtps</code>
<code>c8t0d1</code>	0	0.0	1.0
<code>c10t0d1</code>	0	0.0	1.0
<code>c7t6d0</code>	4716	289.6	1.0

Disk I/O metrics



sar 11.31 – read and write

```
sar -d -R 3 2
```

```
HP-UX veloce B.11.31 U ia64 07/16/07
```

14:58:24	device	%busy	avque	r/s	w/s	blks/s	avwait	avserv
14:58:27	disk5	100.00	0.50	0	240	11383	0.04	21.33
	disk12	32.12	2.72	0	919	115751	1.17	2.29
14:58:30	disk5	100.00	0.59	0	248	11755	0.57	21.45
	disk12	32.33	2.72	0	924	116395	1.17	2.29
Average	disk5	100.00	0.53	0	246	11527	0.21	21.16
Average	disk12	25.28	2.71	0	724	91232	1.16	2.29

sar 11.31 by HBA

```
# sar -H 3 3
```

```
HP-UX veloce B.11.31 U ia64 07/16/07
```

	ctlr	util	t-put	IO/s	r/s	w/s	read	write	avque	await	avserv
		%age	MB/s	num	num	num	MB/s	MB/s	num	msec	msec
Average	fcd0	2	1.65	37	0	37	0.00	1.65	1	0	1
Average	fcd2	2	1.76	37	0	37	0.00	1.76	1	0	1

sar on HP-UX 11i v3 (11.31) by lunpath

```
# sar -L 3 3
```

	lunpath	%busy	avque	r/s	w/s	blks/s	await	avserv
		%age	num	num	num	num	msec	msec
Average	disk5_lunpath0	100.00	0.51	1	239	11812	0.09	21.84
Average	disk12_lunpath11	5.22	0.77	0	53	5409	0.06	1.51
Average	disk12_lunpath7	5.33	0.81	0	53	5275	0.07	1.55
Average	disk12_lunpath8	5.88	2.00	0	53	5260	0.37	2.02
Average	disk12_lunpath12	5.88	2.03	0	53	5251	0.38	2.10
Average	disk12_lunpath17	5.11	0.87	0	53	5281	0.08	1.47
Average	disk12_lunpath18	5.22	0.80	0	53	5146	0.07	1.44
Average	disk12_lunpath20	5.88	1.95	0	53	5278	0.37	2.06
Average	disk12_lunpath19	5.99	2.00	0	53	5220	0.38	2.15

Logical vs. physical I/O

- Logical I/O – what the process or thread requested through system calls such as read/write or pread/pwrite
 - reported by DB utilities (statspak, awr)
glance, tusc
- Physical I/O what the host requested from the storage to satisfy the logical request from the application
 - reported by glance
sar, iostat

Glance – process syscalls screen

GlancePlus C.04.50.00 13:10:19 snappy ia64 Current Avg High

Metric	Current	Avg	High
CPU Util	4%	5%	5%
Disk Util	100%	100%	100%
Mem Util	89%	89%	89%
Swap Util	22%	22%	22%

System Calls PID: 7647, dbworker PPID: 1 euid: 0 User: root

System Call Name	ID	Count	Rate	Elapsed Time	Cum Ct	CumRate	Elapsed Time
read	3	792	152.3	5.07162	1512	100.8	10.16259
nanosleep	353	7	1.3	0.15581	14	0.9	0.33272
lseek64	365	720	138.4	0.00103	1374	91.6	0.00201

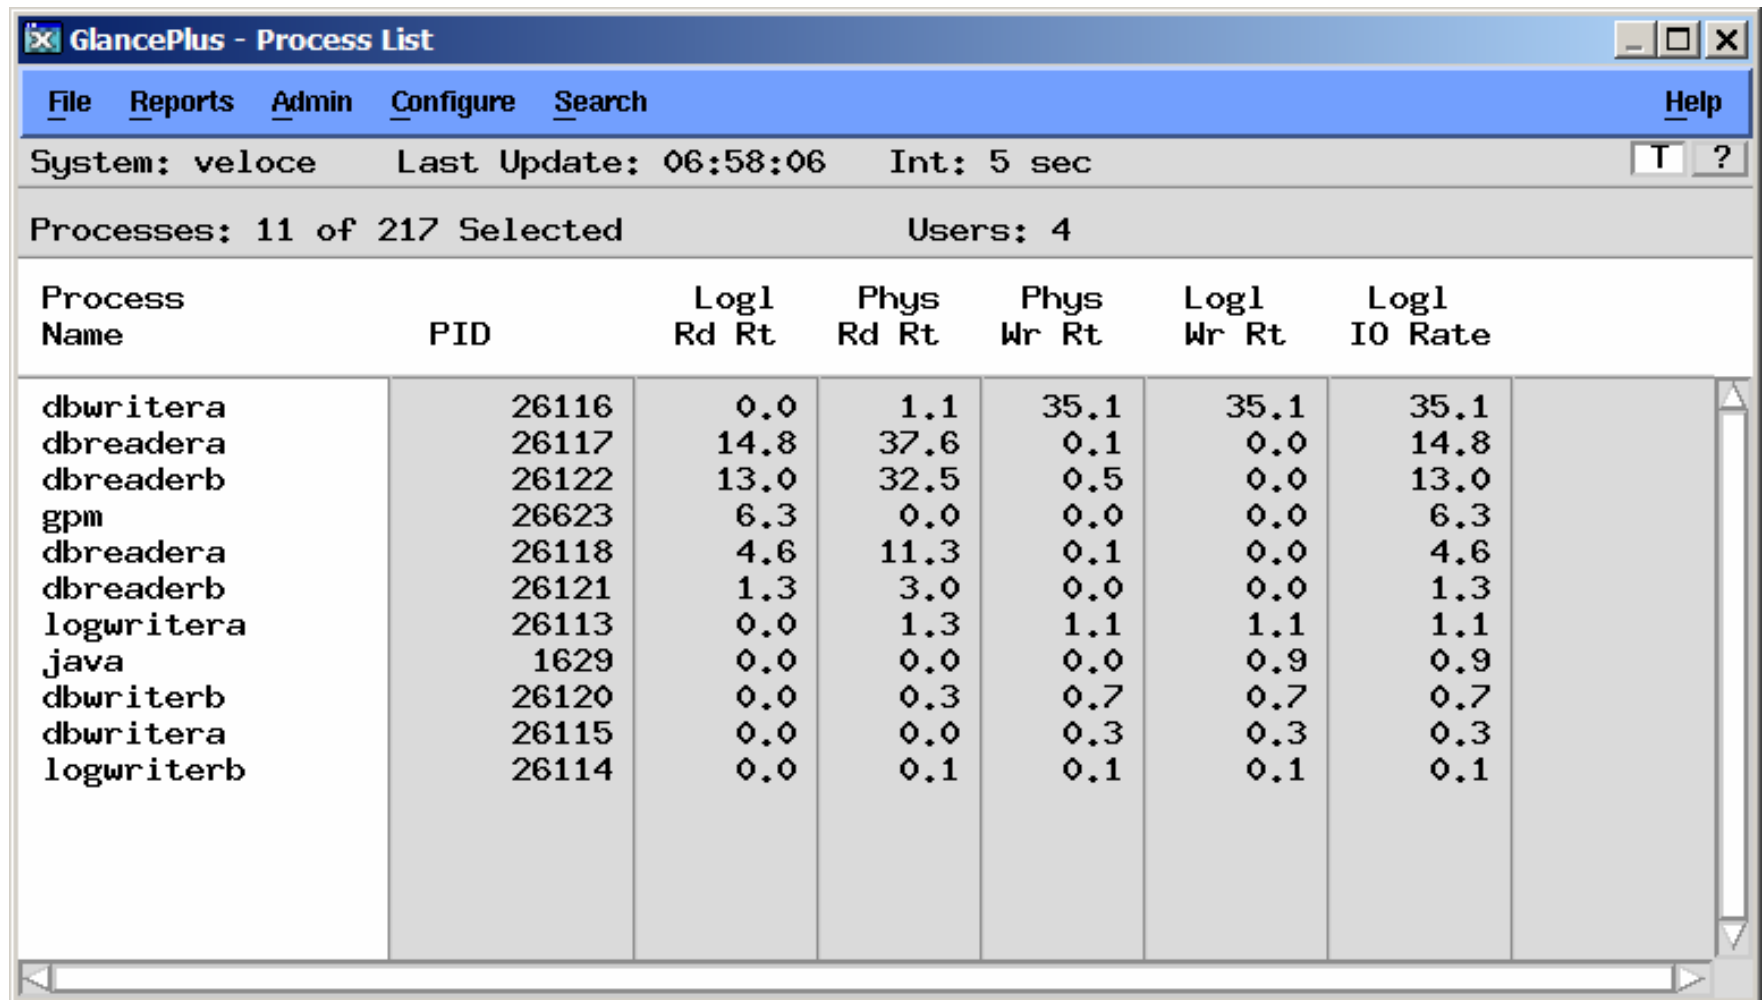
Cumulative Interval: 15 secs Page 1 of 1

Process Resource Wait States Memory Regions Open Files Next Keys Process Syscalls

Process resource screen

```
fleet.cup.hp.com.r1w - Reflection for HP
File Edit Connection Setup Macro Window Help
GlancePlus C.04.50.00      13:18:27  snappy  ia64  Current Avg High
-----
CPU Util      S S      | 5%  5%  10%
Disk Util    F          F | 100% 100% 100%
Mem Util     S          SUUB B | 89%  89%  89%
Swap Util    U          URR      | 22%  22%  22%
-----
Resources PID: 7648, dbworker      PPID: 1  euid: 0  User: root
-----
CPU Usage (util):      1.0  Log Reads : 1293  Wait Reason : IO
User/Nice/RT CPU:     0.0  Log Writes: 0      Total RSS/VSS : 296kb/ 516kb
System CPU :          0.8  Phy Reads : 954  Traps / Vfaults: 0/ 0
Interrupt CPU :       0.0  Phy Writes: 0    Faults Mem/Disk: 0/ 0
Cont Switch CPU :    0.0  FS Reads : 954  Deactivations : 0
Scheduler :          HPUX  FS Writes : 0    Forks & Vforks : 0
Priority :            148  VM Reads : 0   Signals Recd : 0
Nice Value :         20  VM Writes : 0  Mesg Sent/Recd : 0/ 0
Dispatches :        604  Sys Reads : 0  Other Log Rd/Wt: 0/ 0
Forced CSwitch :     1   Sys Writes: 0  Other Phy Rd/Wt: 0/ 0
VoluntaryCSwitch:   602  Raw Reads : 0  Proc Start Time
Running CPU :        2   Raw Writes: 0  Thu Apr 26 13:03:37 2007
CPU Switches :       0   Bytes Xfer: 12.2mb :
Cmd : ./dbworker
_  C - cum/interval toggle  % - pct/absolute toggle      Page 1 of 1
Process Wait Memory Open Next Process
Resource States Regions Files Keys Syscalls
1025, 1      HP70092 -- fleet.cup.hp.com via TELNET      01:24:32      Num
start 3 Jabber Mes... Internet Ex... Inbox - Microso... fleet.cup.hp.co... fleet.cup.hp.co... 2 Microsoft Of... 4:20 PM
```

GPM Process List



The screenshot shows a window titled "GlancePlus - Process List" with a menu bar (File, Reports, Admin, Configure, Search, Help) and a status bar (System: veloce, Last Update: 06:58:06, Int: 5 sec). Below the status bar, it indicates "Processes: 11 of 217 Selected" and "Users: 4". The main content is a table with the following columns: Process Name, PID, LogI Rd Rt, Phys Rd Rt, Phys Wr Rt, LogI Wr Rt, and LogI IO Rate. The table lists 11 processes, including dbwritera, dbreadera, dbreaderb, gpm, java, and logwritera, along with their respective PIDs and resource usage statistics.

Process Name	PID	LogI Rd Rt	Phys Rd Rt	Phys Wr Rt	LogI Wr Rt	LogI IO Rate
dbwritera	26116	0.0	1.1	35.1	35.1	35.1
dbreadera	26117	14.8	37.6	0.1	0.0	14.8
dbreaderb	26122	13.0	32.5	0.5	0.0	13.0
gpm	26623	6.3	0.0	0.0	0.0	6.3
dbreadera	26118	4.6	11.3	0.1	0.0	4.6
dbreaderb	26121	1.3	3.0	0.0	0.0	1.3
logwritera	26113	0.0	1.3	1.1	1.1	1.1
java	1629	0.0	0.0	0.0	0.9	0.9
dbwriterb	26120	0.0	0.3	0.7	0.7	0.7
dbwritera	26115	0.0	0.0	0.3	0.3	0.3
logwriterb	26114	0.0	0.1	0.1	0.1	0.1

Viewing the application requests

```
# tusc -fvpn -o tusc.out 7648
```

```
[./dbworker ][7648] lseek64(8, 713400320, SEEK_SET) = 713400320
[./dbworker ][7648] read(8, "\001T - \0\0\0\8Z 049 T - \0"..., 8192) = 8192
[./dbworker ][7648] lseek64(3, 37543936, SEEK_SET) = 37543936
[./dbworker ][7648] read(3, "\001c41f\0\0\0\> D 9 T 1f@ "..., 8192) = 8192
[./dbworker ][7648] lseek64(4, 30269440, SEEK_SET) = 30269440
[./dbworker ][7648] read(4, "\0\0? ] \0\0\0\b\0~ b] \0"..., 8192) = 8192
[./dbworker ][7648] lseek64(5, 618446848, SEEK_SET) = 618446848
[./dbworker ][7648] read(5, "\001& e6\0\0\0\b02M c6e6\0"..., 8192) = 8192
[./dbworker ][7648] lseek64(6, 1793851392, SEEK_SET) = 1793851392
[./dbworker ][7648] read(6, "\003W ` \0\0\0\b06aec ` \0"..., 8192) = 8192
[./dbworker ][7648] read(6, "\003W a \0\0\0\b06aecg a \0"..., 8192) = 8192
[./dbworker ][7648] lseek64(7, 1492762624, SEEK_SET) = 1492762624
[./dbworker ][7648] read(7, "\002c7ce\0\0\0\b058f9ce\0"..., 8192) = 8192
```

Possible I/O-related wait states

- VM – Virtual memory, direct I/O waiting for pages to be locked
- IO – Waiting for physical I/O
- Cache – Waiting for a buffer cache block, could be waiting for I/O through the cache or for a “hot” block in the cache
- System – Possible inode lock contention
- Job Control – will show up if the process has been tusc'd

Glance wait states

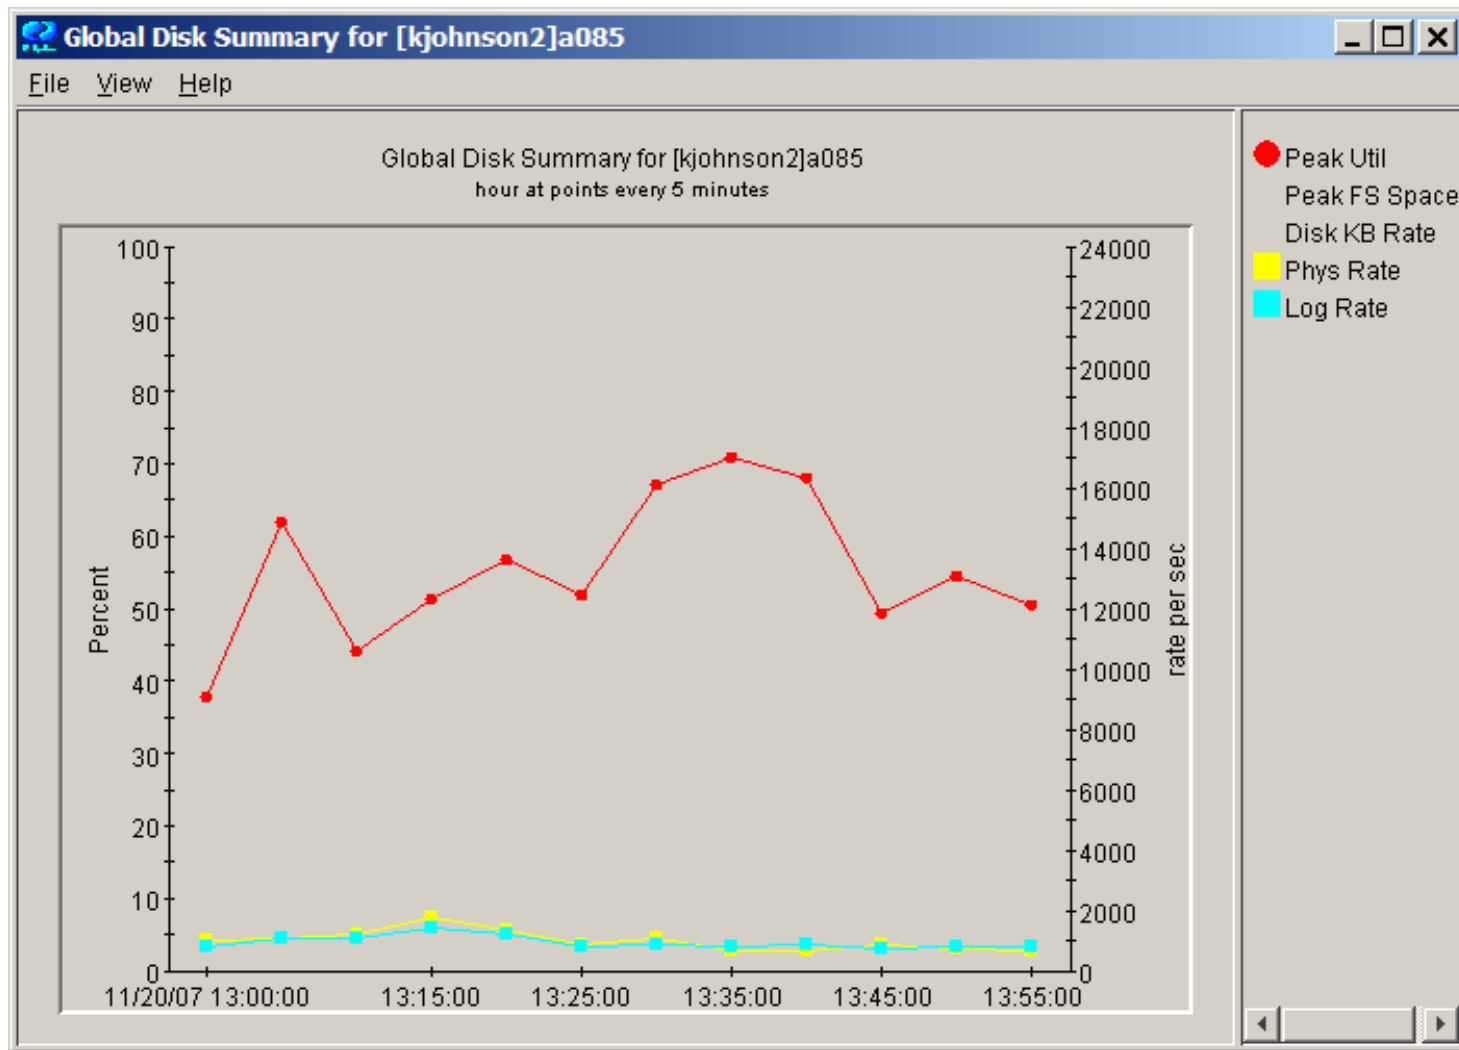
```

fleet.cup.hp.com.r1w - Reflection for HP
File Edit Connection Setup Macro Window Help
GlancePlus C.04.50.00      11:03:57  snappy  ia64  Current Avg High
-----
CPU Util   S | 1%  2%  16%
Disk Util  F | 100% 100% 100%
Mem Util   S | 89%  89%  90%
Swap Util  U | 22%  22%  22%
-----
Wait States PID: 18695, dbworker      PPID: 1 euid: 0 User: root

Event      %      Blocked On      %
-----
IPC        : 0.0      Cache           : 1.9      CPU Util   : 0.4
Job Control: 0.0      CDROM IO        : 0.0      Wait Reason: SYSTM
Message    : 0.0      Disk IO         : 0.0
Pipe       : 0.0      Graphics        : 0.0
RPC        : 0.0      Inode           : 0.0
Semaphore  : 0.0      IO              : 32.4
Sleep      : 0.0      LAN             : 0.0
Socket     : 0.0      NFS             : 0.0
Stream     : 0.0      Priority         : 0.0
Terminal   : 0.0      System          : 64.8
Other      : 0.6      Virtual Mem:    0.0

C - cum/interval toggle  % - pct/absolute toggle      Page 1 of 1
Process Wait Memory Open Next Process
Resource States Regions Files Keys Syscalls
1022, 1 | HP70092 -- fleet.cup.hp.com via TELNET | 05:41:03 | Num
start | Microsoft ... | Internet Ex... | Jabber Messe... | alicia.briggi | Reflection ... | Microsoft ... | 2:06 PM
  
```

OVPM



OVP

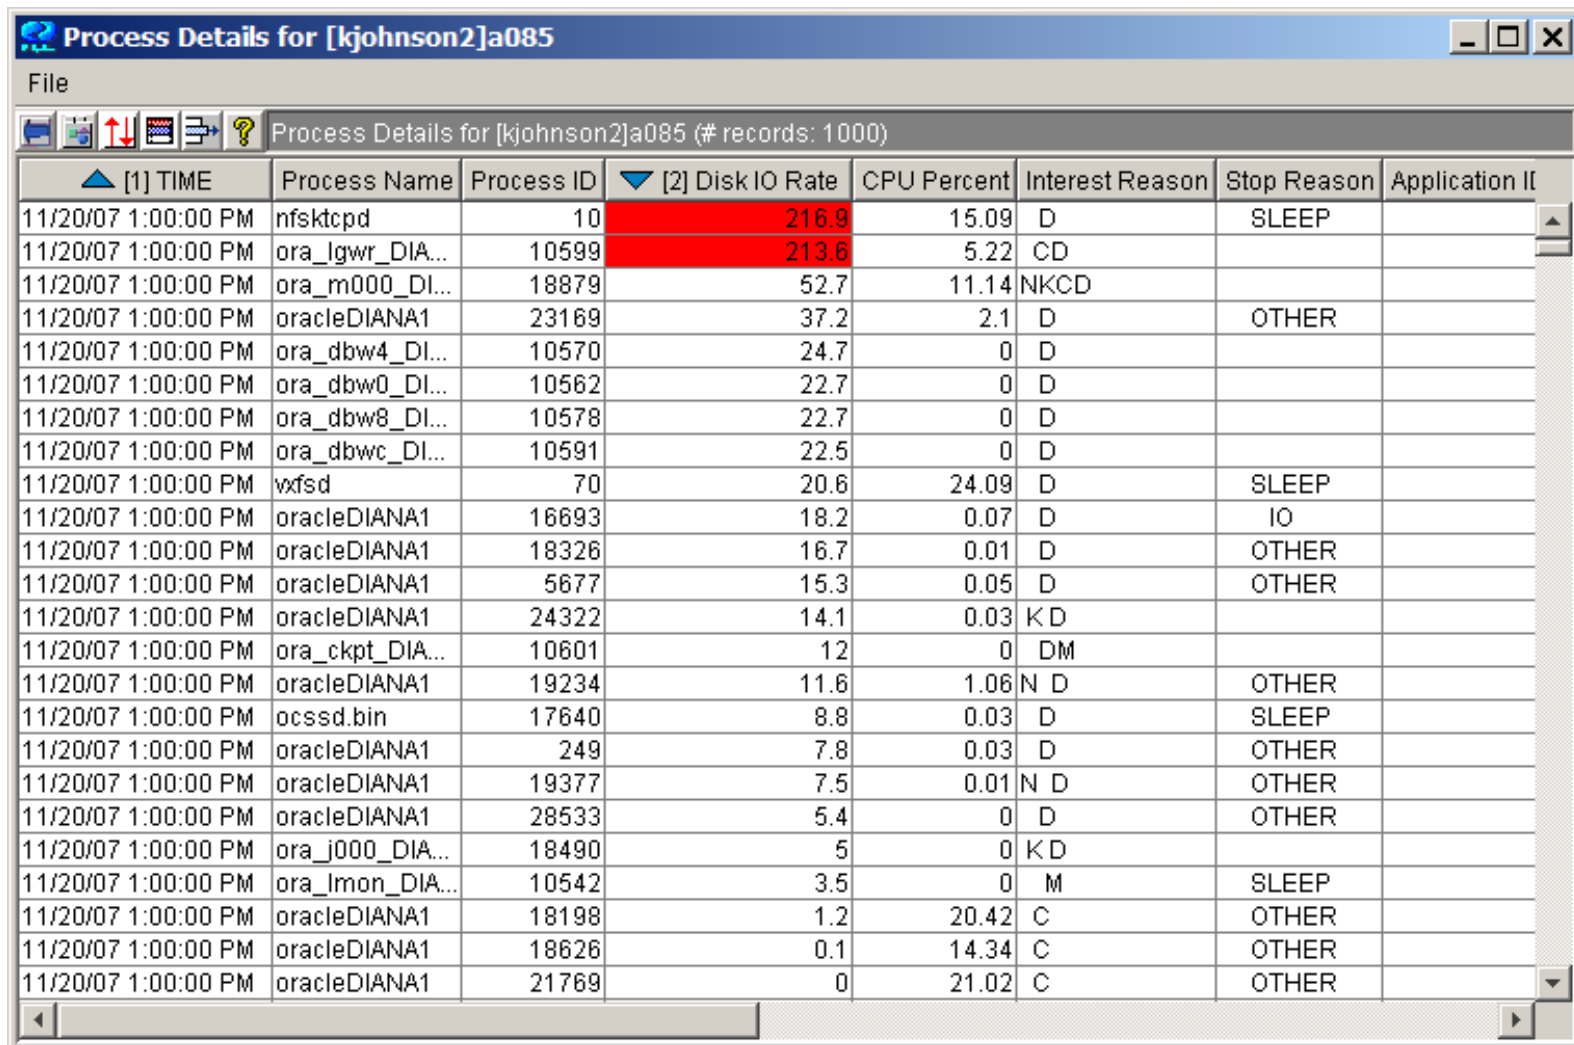
Disk Device Details for [kjohnson2]a085

File

Disk Device Details for [kjohnson2]a085 (# records: 781)

[1] TIME	Disk Device Name	[2] Read Rate	Write Rate	Disk Busy Utilization	Request Queue	Avg Service Time
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.7.5	68.5	0	21.82	0	4.89
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.7.6	49.3	0	13.65	0	4.08
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.2.7	29.6	2.4	20.58	0	6.52
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.2.3	29.5	2.5	21.89	0	6.92
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.2.5	29.1	2.5	20.14	0	6.45
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.2.6	28.8	2.6	19.3	0	6.2
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.2.0	28.6	2.3	19.82	0	6.49
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.2.4	28.1	2.3	21.44	0	7.13
11/20/07 1:00:00 PM	3/0/6/1/0/4/1.6.60.203.0.7.5	27.8	0	15.97	0	8.59
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.2.1	26.6	2.5	37.83	0	15.35
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.2.2	26.4	2.5	34.99	0	12.95
11/20/07 1:00:00 PM	3/0/6/1/0/4/1.6.60.203.0.7.6	18.4	0	9.37	0	7.45
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.1.3	12.4	7.3	8.77	0	4.46
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.1.7	12.1	7.4	8.12	0	4.18
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.1.6	12	7.7	7.98	0	4.04
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.1.2	11.7	6.9	12.84	0	6.96
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.1.0	9.5	6.6	7.64	0	4.76
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.1.5	9.5	7.1	8.01	0	4.83
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.1.1	9.2	7.2	14.98	0	9.34
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.1.4	9.2	6.7	8.34	0	5.24
11/20/07 1:00:00 PM	3/0/6/1/0/4/1.6.60.203.0.2.2	8	0.9	13.6	0	15.78
11/20/07 1:00:00 PM	3/0/6/1/0/4/1.6.60.203.0.2.1	7.7	0.9	18.51	0	25.94
11/20/07 1:00:00 PM	3/0/6/1/0/4/1.6.60.203.0.2.3	5.9	0.7	4.81	0	7.26
11/20/07 1:00:00 PM	2/0/6/1/0/4/1.6.60.228.0.3.3	5.6	3.3	0.94	0	1.04

OVP



Process Details for [kjohnson2]a085

File

Process Details for [kjohnson2]a085 (# records: 1000)

[1] TIME	Process Name	Process ID	[2] Disk IO Rate	CPU Percent	Interest Reason	Stop Reason	Application ID
11/20/07 1:00:00 PM	nfsktcpd	10	216.9	15.09	D	SLEEP	
11/20/07 1:00:00 PM	ora_lgwr_DIA...	10599	213.6	5.22	CD		
11/20/07 1:00:00 PM	ora_m000_DI...	18879	52.7	11.14	NKCD		
11/20/07 1:00:00 PM	oracleDIANA1	23169	37.2	2.1	D	OTHER	
11/20/07 1:00:00 PM	ora_dbw4_DI...	10570	24.7	0	D		
11/20/07 1:00:00 PM	ora_dbw0_DI...	10562	22.7	0	D		
11/20/07 1:00:00 PM	ora_dbw8_DI...	10578	22.7	0	D		
11/20/07 1:00:00 PM	ora_dbwc_DI...	10591	22.5	0	D		
11/20/07 1:00:00 PM	wfsd	70	20.6	24.09	D	SLEEP	
11/20/07 1:00:00 PM	oracleDIANA1	16693	18.2	0.07	D	IO	
11/20/07 1:00:00 PM	oracleDIANA1	18326	16.7	0.01	D	OTHER	
11/20/07 1:00:00 PM	oracleDIANA1	5677	15.3	0.05	D	OTHER	
11/20/07 1:00:00 PM	oracleDIANA1	24322	14.1	0.03	KD		
11/20/07 1:00:00 PM	ora_ckpt_DIA...	10601	12	0	DM		
11/20/07 1:00:00 PM	oracleDIANA1	19234	11.6	1.06	N D	OTHER	
11/20/07 1:00:00 PM	ocssd.bin	17640	8.8	0.03	D	SLEEP	
11/20/07 1:00:00 PM	oracleDIANA1	249	7.8	0.03	D	OTHER	
11/20/07 1:00:00 PM	oracleDIANA1	19377	7.5	0.01	N D	OTHER	
11/20/07 1:00:00 PM	oracleDIANA1	28533	5.4	0	D	OTHER	
11/20/07 1:00:00 PM	ora_j000_DIA...	18490	5	0	KD		
11/20/07 1:00:00 PM	ora_lmon_DIA...	10542	3.5	0	M	SLEEP	
11/20/07 1:00:00 PM	oracleDIANA1	18198	1.2	20.42	C	OTHER	
11/20/07 1:00:00 PM	oracleDIANA1	18626	0.1	14.34	C	OTHER	
11/20/07 1:00:00 PM	oracleDIANA1	21769	0	21.02	C	OTHER	

Performance trouble-shooting – Summary

- We've reviewed
 - how HP-UX performance metrics are collected
 - tools in the context of CPU, memory, and I/O performance
- In conclusion, ...

Performance trouble-shooting – Summary

When it's running, what's it doing?

When it's waiting, what's it waiting for?



i n v e n t