



HP-UX 11i v3 Knowledge-on-Demand

HP technical Webcast series: deployment optimization



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



GlancePlus

HP technical Webcast series: deployment optimization



Technology for better business outcomes

Introducing today's speaker

Doug Grumann is an expert on system performance analysis and the HP system performance tools. He worked on the original team that developed Glance on the HP-UX and other unix platforms over 15 years ago, and he has worked in this product domain since then, participating in various aspects of development, support, and marketing.



Contents

- Glance introduction
- Overview of system performance analysis
- Tour of character mode and motif mode user interfaces
- Customization
- Bottleneck analysis example
- Performance Tips
- Data flow and references

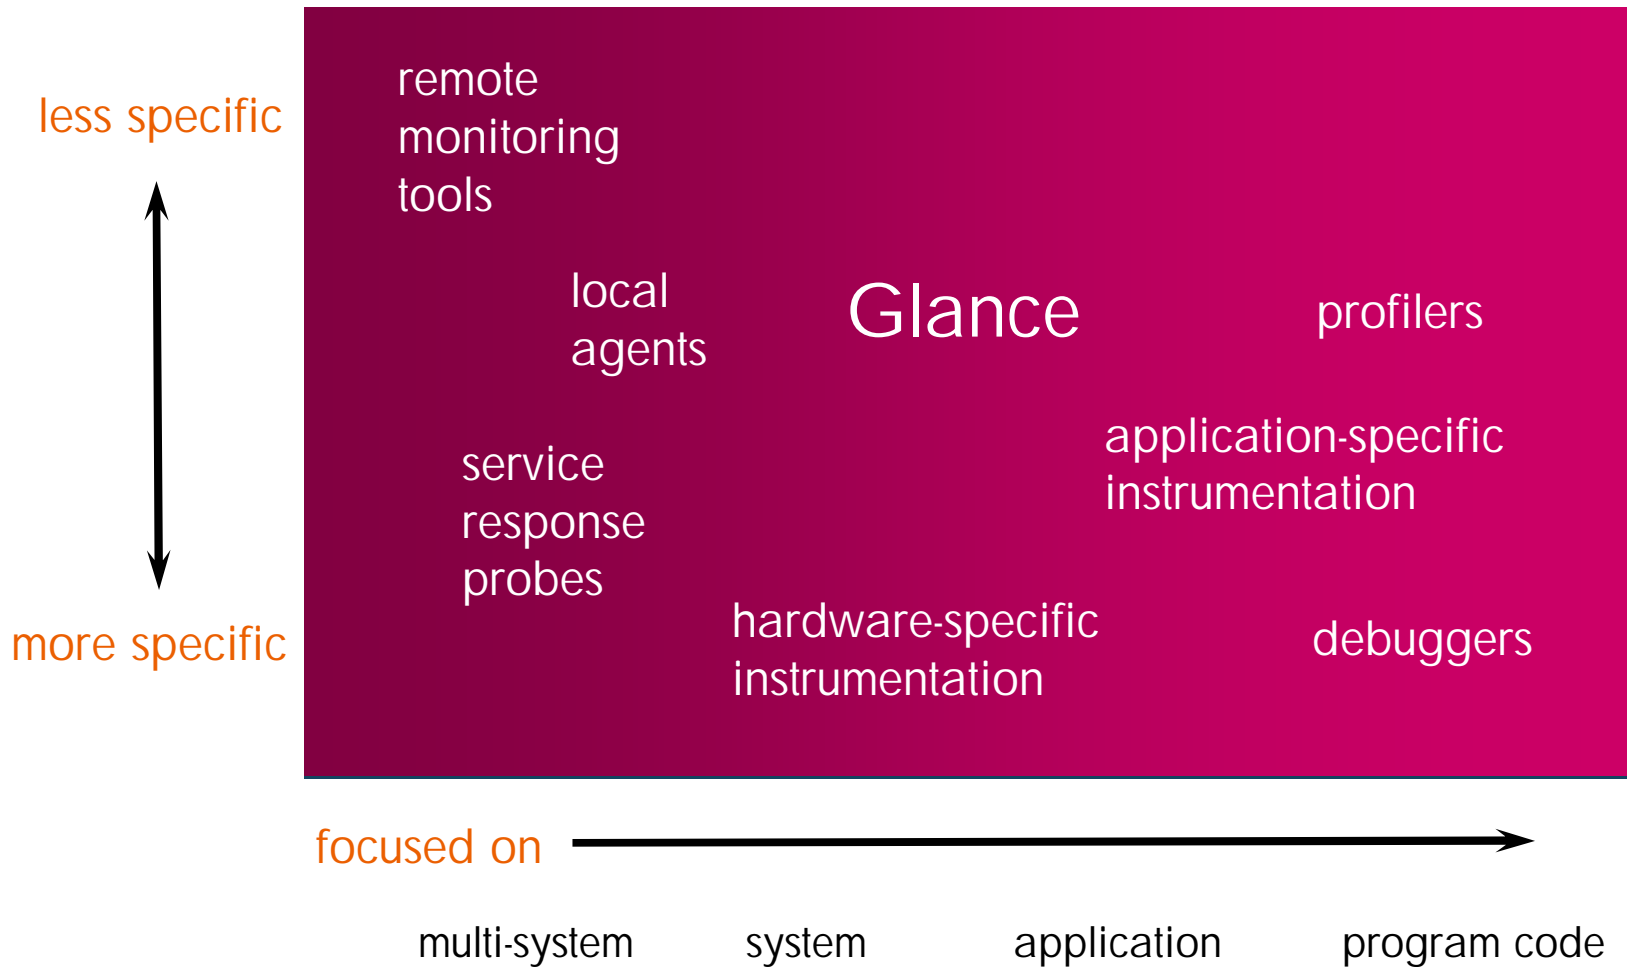
Glance at a glance

- Glance answers better than any other tool: “what's going on right now on my system?” It is the best-in-industry server performance diagnostic!
- Glance supports a drill-down approach to troubleshooting from bottleneck analysis to detailed visualization of over two thousand performance metrics
- Choose from two user interfaces: character mode glance, or the Motif mode (xglance / gpm) GUI

Where does it come from?

- Glance was one of the very first software products ever produced by HP. Since the 1990, Glance has been available on every release of HP-UX. Multivendor unix versions are also available
- Glance is included with the HP-UX Application Release media: trial bundles are available from that media even if you do not have the product license
- Glance is included with HP-UX Enterprise and Mission Critical Operating Environment bundles
- Glance is also available as separate product from HP Software
- The Glance Pak product bundle includes Glance along with the HP Software Performance Agent (MeasureWare)

Performance tool spectrum



The art of system performance

“why is our server slow???”


“is the I/O load
balanced across disks?”

“what apps are running?”

server systems of all types

“should I buy
more memory?”

“what’s this
program doing?”





“which resources
are bottlenecked?”

heterogeneous OSes and applications

System performance

– no answers without data!



global cpu utilization

per-process user mode cpu

service time

processor queues

syscall rate

by-disk I/O rate

thread list

application active procs

virtual memory size

Tips for effective system performance

A resource that is fully utilized and has a queue of processes waiting for it is bottlenecked. High utilization by itself is not necessarily a problem.

DO:

- Figure out what's "normal" on your systems
- Be willing to do the work to know what you're doing
- Change one thing at a time
- Keep in mind: **It Depends**

DO NOT:

- Ignore systems that are currently running well
- Forget to keep historical performance data
- Fix things that ain't broke
- Assume one solution fits all

Glance tour



Getting started

- Login to your unix server and check to see if `/opt/perf/bin/glance` is there.
- Version info available via:
`/opt/perf/bin/perfstat -v` (consider updating to the latest release – 4.6)
- Glance has a man-page, online help, Release Notes under `/opt/perf/ReleaseNotes/`, and metrics list under `/opt/perf/paperdocs/`
- Product info on the Web:
<http://www.openview.hp.com/products/gplus/index.html>

Initial glance screen

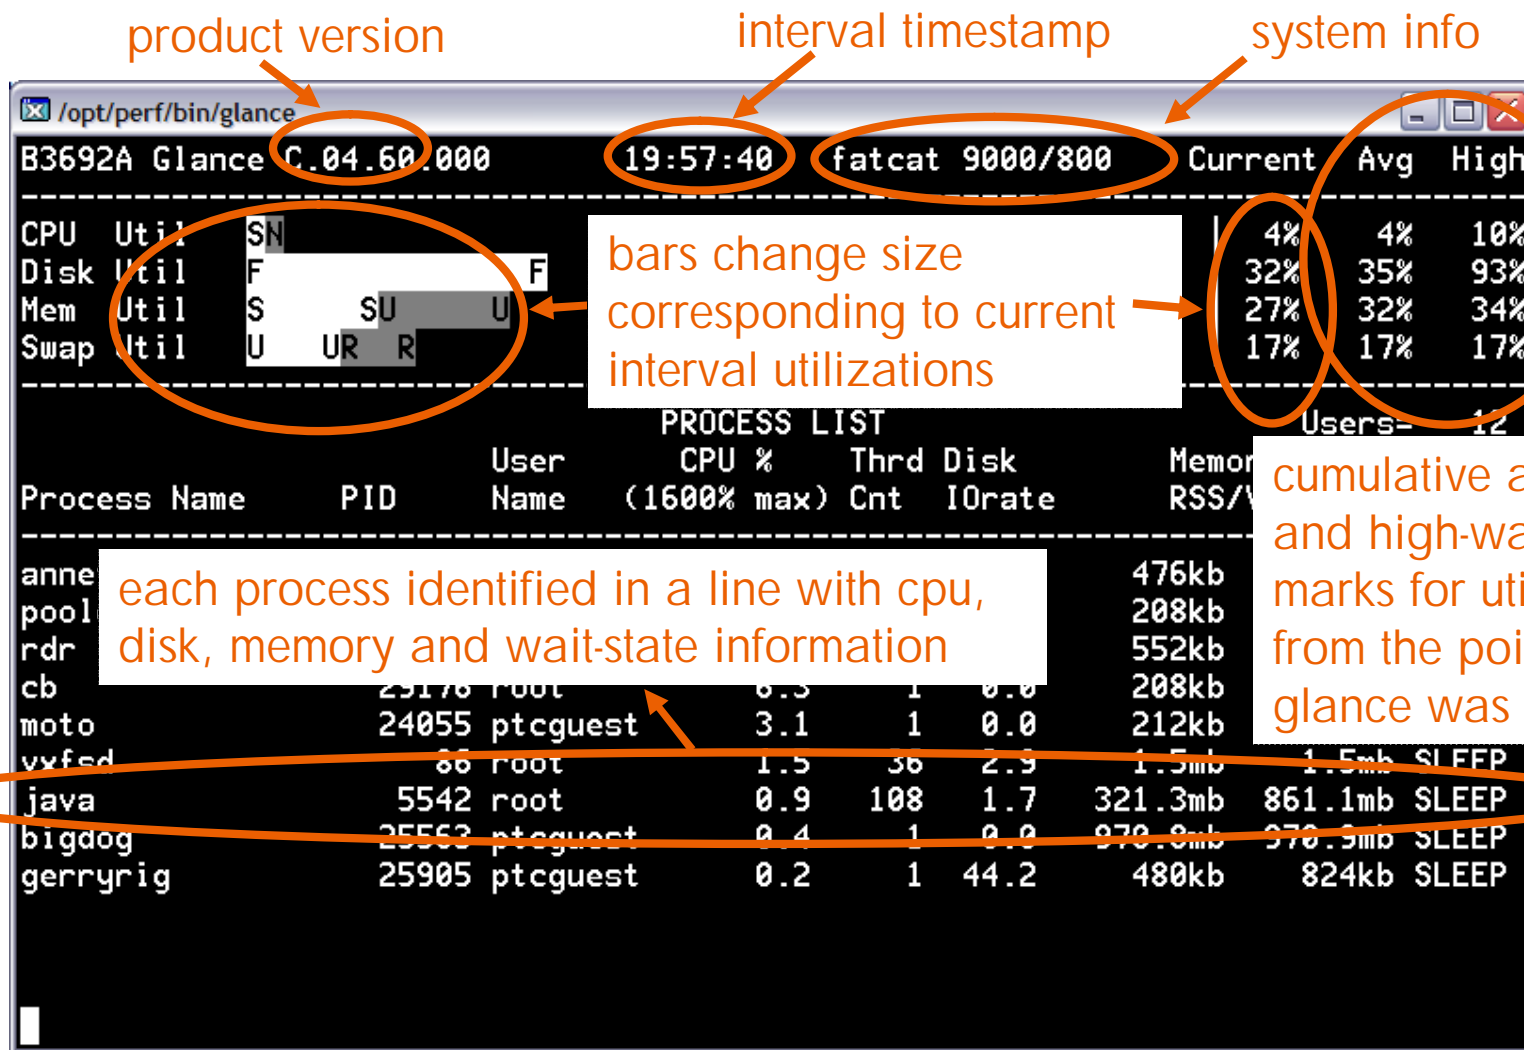
top part of glance screen shows "global" data reflecting overall system status

```
B3692A Glance C.04.60.000 19:57:40 fatcat 9000/800 Current Avg High
-----
CPU Util SN 4% 10%
Disk Util F 5% 93%
Mem Util S SU U 2% 34%
Swap Util U UR R 7% 17%
```

bottom part of glance screen shows most active individual processes

```
PROCESS LIST Users= 12
-----
Process Name PID User CPU % Thrd Disk Memory Block
(1600% max) Cnt IOrate RSS/VSS On
-----
annetest 19662 ptcgust 24.0 1 0.0 476kb 692kb SLEEP
poole 6212 root 11.4 1 0.0 208kb 236kb SLEEP
rdr 633 root 10.8 1 0.0 552kb 2.7mb PRI
cb 29176 root 6.3 1 0.0 208kb 236kb SLEEP
moto 24055 ptcgust 3.1 1 0.0 212kb 236kb SLEEP
vxfsd 86 root 1.5 36 2.9 1.5mb 1.5mb SLEEP
java 5542 root 0.9 108 1.7 321.3mb 861.1mb SLEEP
bigdog 25563 ptcgust 0.4 1 0.0 970.8mb 970.9mb SLEEP
gerryrig 25905 ptcgust 0.2 1 44.2 480kb 824kb SLEEP
```

Initial glance screen



Glance help screen

```

/opt/perf/bin/glance
B3692A Glance C.04.60.000      22:37:27  fatcat 9000/800  Current Avg High
-----
CPU  Util  SN  | 4%  4%  10%
Disk Util  F  | 32% 34% 100%
Mem  Util  S  SU  U  | 27% 29% 34%
Swap Util  U  UR  R  | 17% 17% 17%
-----
                          Glance Control Key Menu

? - 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

Enter command or function key:                               Page 2 of 2

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

```

Glance online help

The screenshot shows the Glance online help interface. On the left, a black box contains help text for the 'CPU Util Bar' metric. On the right, a table displays performance data for 'fatcat 9000/800'. Below the table is a menu of navigation options.

fatcat 9000/800

	Current	Avg	High
-----	4%	4%	10%
	35%	34%	100%
	27%	29%	34%
-----	17%	17%	17%

enu

- A - Application List
- P - PRM Group List
- stem Y - Global System Calls
- F - Process Open Files
- d M - Process Memory Regions
- Volume R - Process Resources
- tivity W - Process Wait States
- L - Process System Calls
- terface y - Renice Process
- s - Select Process
- J - Thread Wait

FS/Appl/Trans/Thread

Page 1

F7 or 'e' for prev menu

Enter command or function key:

One moment please...

Page 1 of 2

	Page Forward	Page Back	hpterm	Print Toggle	Exit Menu	Exit Help
--	--------------	-----------	--------	--------------	-----------	-----------

Reports on all system resources

The screenshot shows the output of the 'glance' utility. At the top, it displays system information: 'B3692A Glance C.04.60.000'. Below this is a summary table of resource utilization:

Resource	Current	Avg	High
CPU Util	4%	4%	10%
Disk Util	1%	34%	100%
Mem Util	7%	29%	34%
Swap Util	7%	17%	17%

A callout box with an orange border contains the text: 'example: memory report screen showing paging activity as well as summary data about overall memory utilization on the bottom'. Below the summary table is the 'MEMORY REPORT' section, which includes a table of memory-related events and their rates, followed by a summary of memory usage.

Event	Current	Cumulative	Current Rate	Cum Rate	High Rate
Page Faults	6542	1481658	109.2	129.7	1729.0
Page In	0	55367	0.0	4.8	635.0
Page Out	0	0	0.0	0.0	0.0
KB Paged In	0kb	216.3mb	0.0	19.3	13734.4
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	5491	0.0	0.4	45.8
VM Writes	0	5841	0.0	0.5	5.1

Summary of memory usage:

Total VM :	2.7gb	Sys Mem :	2.5gb	User Mem:	2.9gb	Phys Mem :	19.9gb
Active VM:	2.4gb	Buf Cache:	38mb	Free Mem:	14.5gb	FileCache:	876mb

Page 1 of 1

Drill down on per-process details

```

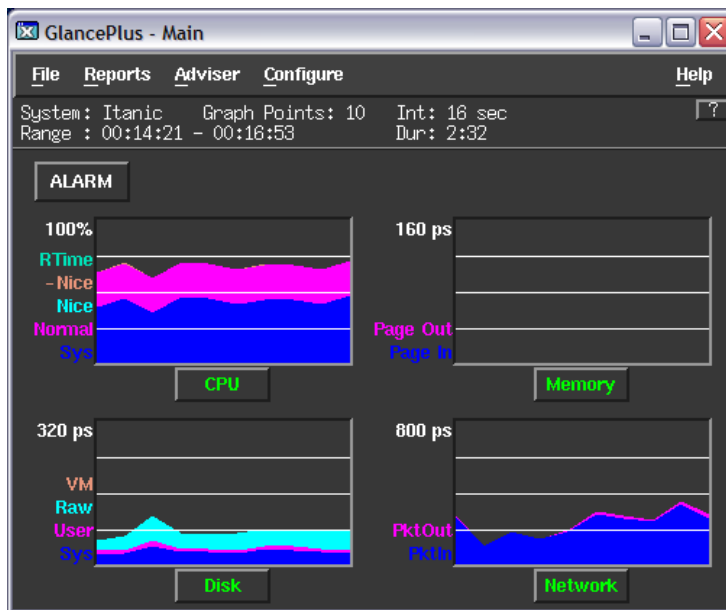
/opt/perf/bin/glance
B3692A Glance C.04.60.0
-----
CPU Util      SN  4% 10%
Disk Util     F  34% 100%
Mem Util      S  29% 34%
Swap Util     U  17% 17%
-----
Resources PID:      6212, poole      PPID:      6385  euid:      0  User:root
-----
CPU Usage (util):      11.5  Log Reads :      0  Wait Reason :      SLEEP
User/Nice/RT CPU:      4.5  Log Writes:     0  Total RSS/VSS : 208kb/ 236kb
System CPU :           7.0  Phy Reads :     0  Traps / Vfaults: 63/ 0
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 :             168  VM Reads :     0  Signals Recd :   60
Nice Value :           39  VM Writes :     0  Mesg Sent/Recd : 0/ 0
Dispatches :          198  Sys Reads :     0  Other Log Rd/Wt: 0/ 0
Forced CSwitch :       8  Sys Writes:     0  Other Phy Rd/Wt: 0/ 0
VoluntaryCSwitch:     60  Raw Reads :     0  Proc Start Time
Running CPU :          13  Raw Writes:     0      Wed Aug 8 19:10:16 2007
CPU Switches :         0  Bytes Xfer:    0kb
Argv1: -n
C - cum/interval toggle  % - pct/absolute toggle
-----
Process Wait Memory Open hpterm Next Process
Resource States Regions Files Keys Syscalls

```

metrics specific to each process include cpu, I/O, and memory activity: additional drilldowns are available to focus in even more!

xglance / gpm interface

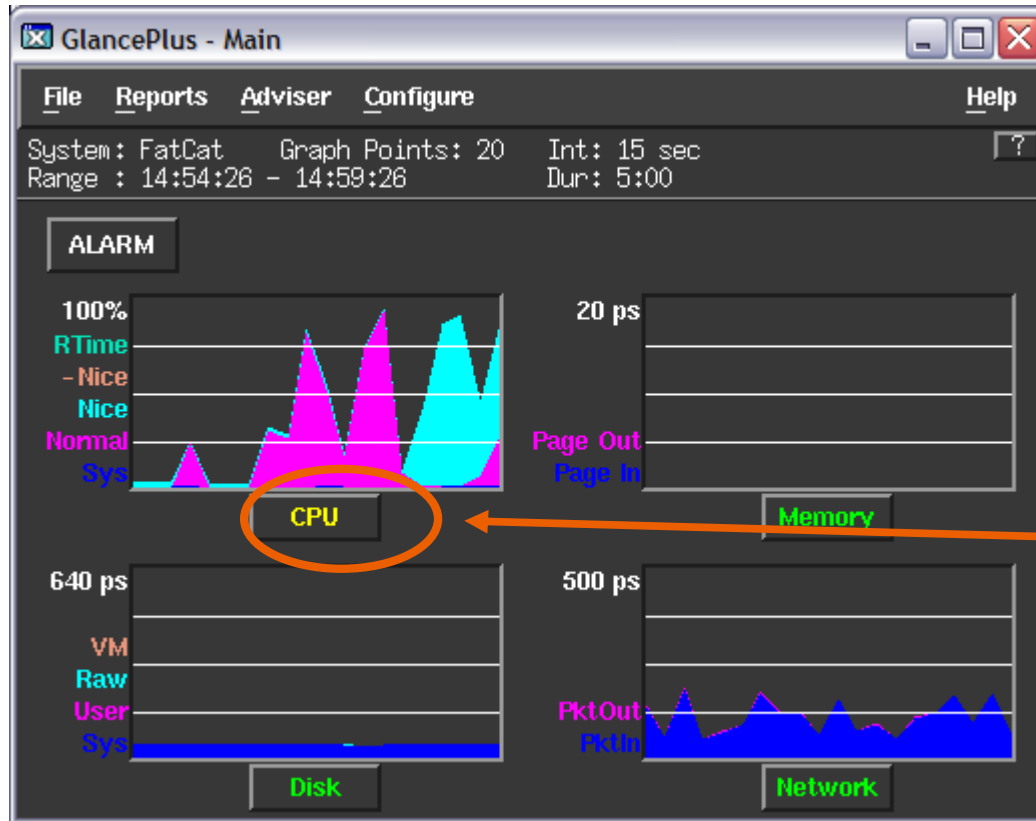
- When you have a X/Motif environment available, Glance's graphical user interface is useful.
- The xglance interface shows you the same basic reports as in character mode but with more options and more detail.



The screenshot shows the 'GlancePlus - Process List' window. It features a menu bar with 'File', 'Reports', 'Admin', 'Configure', 'Search', and 'Help'. Below the menu, system information is displayed: 'System: Itanic', 'Last Update: 00:17:39', 'Int: 15 sec'. It also shows 'Processes: 21 of 246 Selected' and 'Users: 15'. The main area contains a table of running processes.

Process Name	PID	CPU %	Phys IO Rt	Stop Reason	User Name
hpvmapp	1993	157.5	0.0	OTHER	root
annetest	28992	97.5	0.0	PRI	dougg
annetest	28978	97.3	0.0	PRI	dougg
hpvmapp	1531	91.1	0.0	OTHER	root
hpvmapp	167	82.8	0.0	OTHER	root
hpvmapp	1929	15.2	0.0	SLEEP	root
perfd	28320	1.7	1.6	SLEEP	root
glance	2751	0.9	1.3	TERM	ptcguest
vxfsd	57	0.9	12.2	SLEEP	root

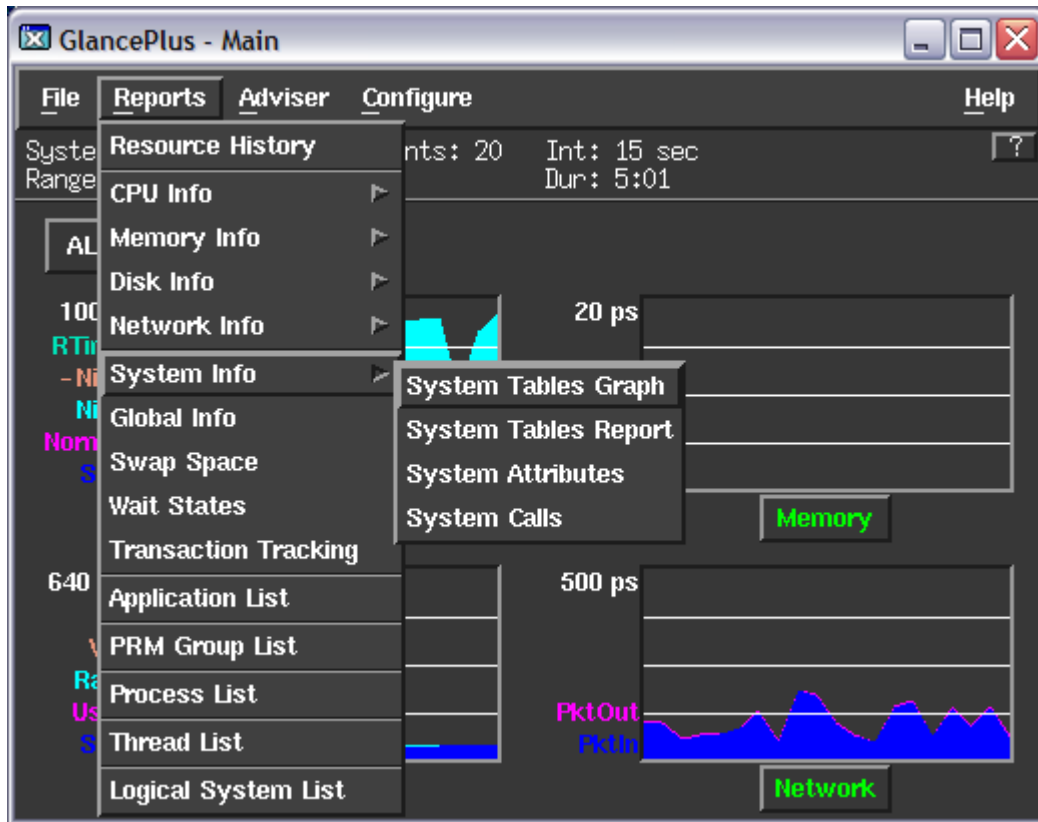
xglance main window



system resource graphs show cpu, memory paging, disk I/O and networking activity over time

the Glance "adviser" functionality runs continuously, detecting bottleneck conditions and alerting you to the area of concern by changing colors

Reports available



you can adjust fonts, measurement intervals, graph duration, limits and other display options

from the main window you can drill down into any system resources or bottleneck focus areas including applications, processes, and thread detail in various reports...

Reports!

The screenshot shows the GlancePlus System Calls report window. The window title is "GlancePlus - System Calls". The menu bar includes "File", "Reports", "Configure", and "Help". The main content area displays the following information:

System: FatCat Last Update: 15:58:24 Int: 10 sec
System Calls: All 168 Selected
Interval for collection: 10 sec
Time for collection : 32:41

SysCall Name	System Call ID	SysCall Call Rt	CPU Time
poll	269	788669.4	8.94193
sched_yield	341	103092.9	0.40910
getitimer	86	11237.4	0.18455
times	43	11184.1	0.16638
sigenable	492	109.6	0.00052
siginhibit	491	109.6	0.00059
close	6	97.9	0.01395
getmsg			
time			
putmsg			
read			
open			
write			
sigprocmask			
ioctl			
gettimeofday			
stat			
SYS_vxd			
SYS_schedcpu			
nanosleep			

A text box overlay on the right side of the screenshot contains the following text:

system calls report lists the kernel requests processes are making – you can change the sort field to focus on the highest syscall rate or total cpu time spent in that call...

More!

The image shows a stack of GlancePlus windows. The top window, 'GlancePlus - Process Memory Regions (bigdog)', displays memory usage for the 'bigdog' process (PID: 19056). A callout box points to the 'DATA' region with 969.0mb RSS, stating: 'memory region detail list for any selected process: data (heap) growing may be a symptom of an application memory leak!'. Below this, a table lists memory regions with columns for Type, RSS KB, and VSS KB. The 'DATA' row is circled in orange. The 'Process Wait States' window below it shows a table of wait states for 'gerrynig' (PID: 12070), with 'IO' circled in orange. A callout box says: 'see what resources processes are waiting for'. The 'Process Wait States' table includes columns for Process Name, PID, PPID, and various wait state categories like IPC, Job Control, Message, etc.

Type	RSS KB	VSS KB
DATA	969.0mb	969.1
MEMMAP	1.3mb	1.4
MEMMAP	184kb	204
MEMMAP	36kb	36
MEMMAP	36kb	44
UAREA	36kb	36
STACK	32kb	32kb
MEMMAP	20kb	20kb
MEMMAP	16kb	16kb
MEMMAP	16kb	16kb
TEXT	12kb	12kb
MEMMAP	12kb	12kb
MEMMAP	8kb	8kb
NULLDR	4kb	4kb
MEMMAP	4kb	4kb
UNUSED	4kb	4kb
MEMMAP	4kb	4kb

Process Name	PID	PPID
gerrynig	12070	12054

Wait State	Percentage	Category
IPC	0.00%	Cache
Job Control	0.00%	CDROM I/O
Message	0.00%	Disk I/O
Pipe	0.00%	Graphics
RPC	0.00%	Inode
Semaphore	0.00%	IO
Sleep	70.72%	LAN
Socket	0.00%	NFS
Stream	0.00%	Priority
Terminal	0.00%	System
		Virtual
		Other

Customizing Glance



xglance main graph limits

The screenshot shows the xglance utility window. At the top, it displays system information: Glance C.04,60,000, 22:01:02, ovruxta3, ia64. Below this is a table of resource usage:

	Current	Avg	High
CPU Util	34%	34%	34%
Disk Util	4%	4%	4%
Mem Util	89%	89%	89%
Swap Util	32%	32%	32%

Below the table is a section titled "INTERESTING PROCESS THRESHOLD OPTIONS" with a list of filters and their current thresholds:

Filter	Current	Thresholds
CPU Utilization	> 5	(1.0 %)
Disk I/O Rate	> 0.5	(1.0 IOs/sec)
Resident Set Size	> 100	(20 Mbytes)
Virtual Set Size	> 900	(500 Mbytes)
User name	=	(all)
Program name	=	(all)
TTY path name	=	(all)
Use match logic (and/or)	:	(or)
Sort key (name/cpu/disk/rss):	:	(cpu)

At the bottom of the window, there are buttons for "Read Sys", "All Defaults", "Perform Entries", "Task", "hpterm", "Set Sys Defaults", "Help", and "Cancel".

Below the main window, there is a control panel with buttons for "Sys", "Disk", and "Network". At the very bottom, there are "OK", "Reset", and "Help" buttons.

graph limits rise to track high water level seen by glance – and may become too high over time

xglance choose metrics

Select the metrics you wish to display in this window.

GlancePlus - Process List

Process Name	PID
tld	1
cld	
dano	
ram	
vxfsd	
dld	

GlancePlus Online Help

File Edit Bookmark Options Help

Contents Index Back Print Find

PROC_DISK_PHYS_IO_RATE
THREAD_DISK_PHYS_IO_RATE

The average number of [physical disk I/Os](#) per second made by the process or kernel thread during the interval.

For processes which run for less than the measurement interval, this metric is normalized over the measurement interval. For example, a process ran for 1 second and did 50 I/Os during its life. If the measurement interval is 5 seconds, it is reported as having done 10 I/Os per second. If the measurement interval is 60 seconds, it is reported as having done 50/60 or 0.83 I/Os per second.

"Disk" in this instance refers to any locally attached physical disk drives (that is, "spindles") that may hold file systems and/or swap. NFS mounted disks are not included in this list.

On HP-UX, since this value is reported by the drivers, multiple physical requests that have been collapsed to a single physical operation (due to driver IO merging) are only counted once.

On a threaded operating system, such as HP-UX 11.0 and beyond, process usage of a resource is calculated by summing the usage of that resource by its kernel threads. If this metric is reported for a kernel thread, the value is the resource usage by that single kernel thread. If this metric is reported for a process, the value is the sum of the resource usage by all of its kernel threads. Alive kernel threads and kernel threads that have died

"?" on-item help button will bring up help when cursor is clicked on a metric or window

Application grouping in the parm file

```
rerun glance after editing ~/.parm applications
Glance C.04.60.000          23:34:10  itania  ia64  Current Avg High
-----
CPU Util  SN [71%] N
Disk Util FFV [4%]
Mem Util  S [90%] SU U
Swap Util U [33%] UR R
-----
APPLICATION LIST
Users= 7
Idx Application      Num Active CPU AvgCPU Logl Phys Res Virt
Procs Procs Util Util IO IO Mem Mem
-----
1 other              0 0 0.0 0.0 0.0 0.0 0kb 0kb
2 testapp            15 7 70.5 66.4 103.8 10.1 388.4mb 448.9mb
3 weblogic           2 0
4 Oracle              5 0
5 other_user_root    148 9
-----
S - Select an Application
Page 1 of 1
```

to verify, you can then select an application to see what processes are included in each

Adviser customizations

```

/var/opt/perf/adviser.syntax
# The following symptoms are used by the default Alarm Window
# Bottleneck alarms. They are re-evaluated every interval and
# the probabilities are summed. These summed probabilities are
# checked by the bottleneck alarms. The buttons on the gpm
# main window will turn yellow when a probability exceeds 50%
# for an interval, and red when a probability exceeds 90% for
# an interval. You may edit these rules to suit your environment:

symptom CPU_Bottleneck type=CPU
rule GBL_CPU_TOTAL_UTIL > 75 prob 25
rule GBL_CPU_TOTAL_UTIL > 85 prob 25
rule GBL_CPU_TOTAL_UTIL > 90 prob 25
rule GBL_PRI_QUEUE > 3 prob 25

symptom Disk_Bottleneck type=DISK
rule GBL_DISK_UTIL_PEAK > 50 prob GBL_DISK_UTIL_PEAK
rule GBL_DISK_SUBSYSTEM_QUEUE > 3 prob 25

symptom Memory_Bottleneck type=MEMORY
rule GBL_MEM_QUEUE > 2 prob 20
rule GBL_MEM_PAGEOUT_RATE > 5 prob 20
rule GBL_MEM_PAGEOUT_RATE > 50 prob 20
rule GBL_DISK_VM_WRITE_RATE > 5 prob 20
"/var/opt/perf/adviser.syntax" [Read only] 203 lines, 7281 characters

```

character-mode glance reads its adviser rules for bottleneck definitions from the `adviser.syntax` file, `xglance` or `gpm` keeps its adviser syntax in a binary format and you need to edit that inside the tool itself

- in either case, you can change the adviser rules to suit your unique needs

“Programming” the adviser

```
processcpu adviser execution
> glance -adviser_only -syntax ./processcpu -j 30

Welcome to Glance

12:33:59 alive processes =171.0, active processes = 31.4, run queue = 1.1
pid      5795 cld          cpu util = 48.5
number of processes currently blocked waiting for cpu = 1.00

12:34:29 alive processes =176.0, active processes =
pid      5796 dano        cpu util = 18.4
number of processes currently blocked waiting for cpu

12:34:59 alive processes =182.0, active processes =
2.9
pid      5795 cld          cpu util = 0.5
pid      5796 dano        cpu util = 0.2
pid      18186 tld        cpu util = 47.0
pid      18190 tld        cpu util = 30.8
pid      18199 tld        cpu util = 28.0
pid      18203 tld        cpu util = 27.7
number of processes currently blocked waiting for cpu = 6.00
```

character-mode glance has runtime options to just run the adviser against the syntax file you created, updating at an interval of your choice.. “batch processing” for custom performance monitoring!

Bottleneck Analysis Example and Tips



Performance scenario

You manage a HP Virtual Machine server running four guest OSES. You have Glance and the Performance Agent (PA) installed on the server and the guests as well.

```
20:05 ALARM [5] START WARNING: HPVM CPU Bottleneck probability= 75%
20:10 ALARM [9] START CRITICAL: cld application response time >5s
```

Performance alarms may be generated by PA or Smart Plugins into Operations Manager, or they may arise via remote service probes. User complaints about application performance may come in via your IT service desk saying response has gone bad as well. Regardless of the origin of an alarm, the process of analyzing the root cause starts with an examination of key performance metrics. In this example, the application resides on your HPVM guests so your first step is to see what may be unusual about their behavior.

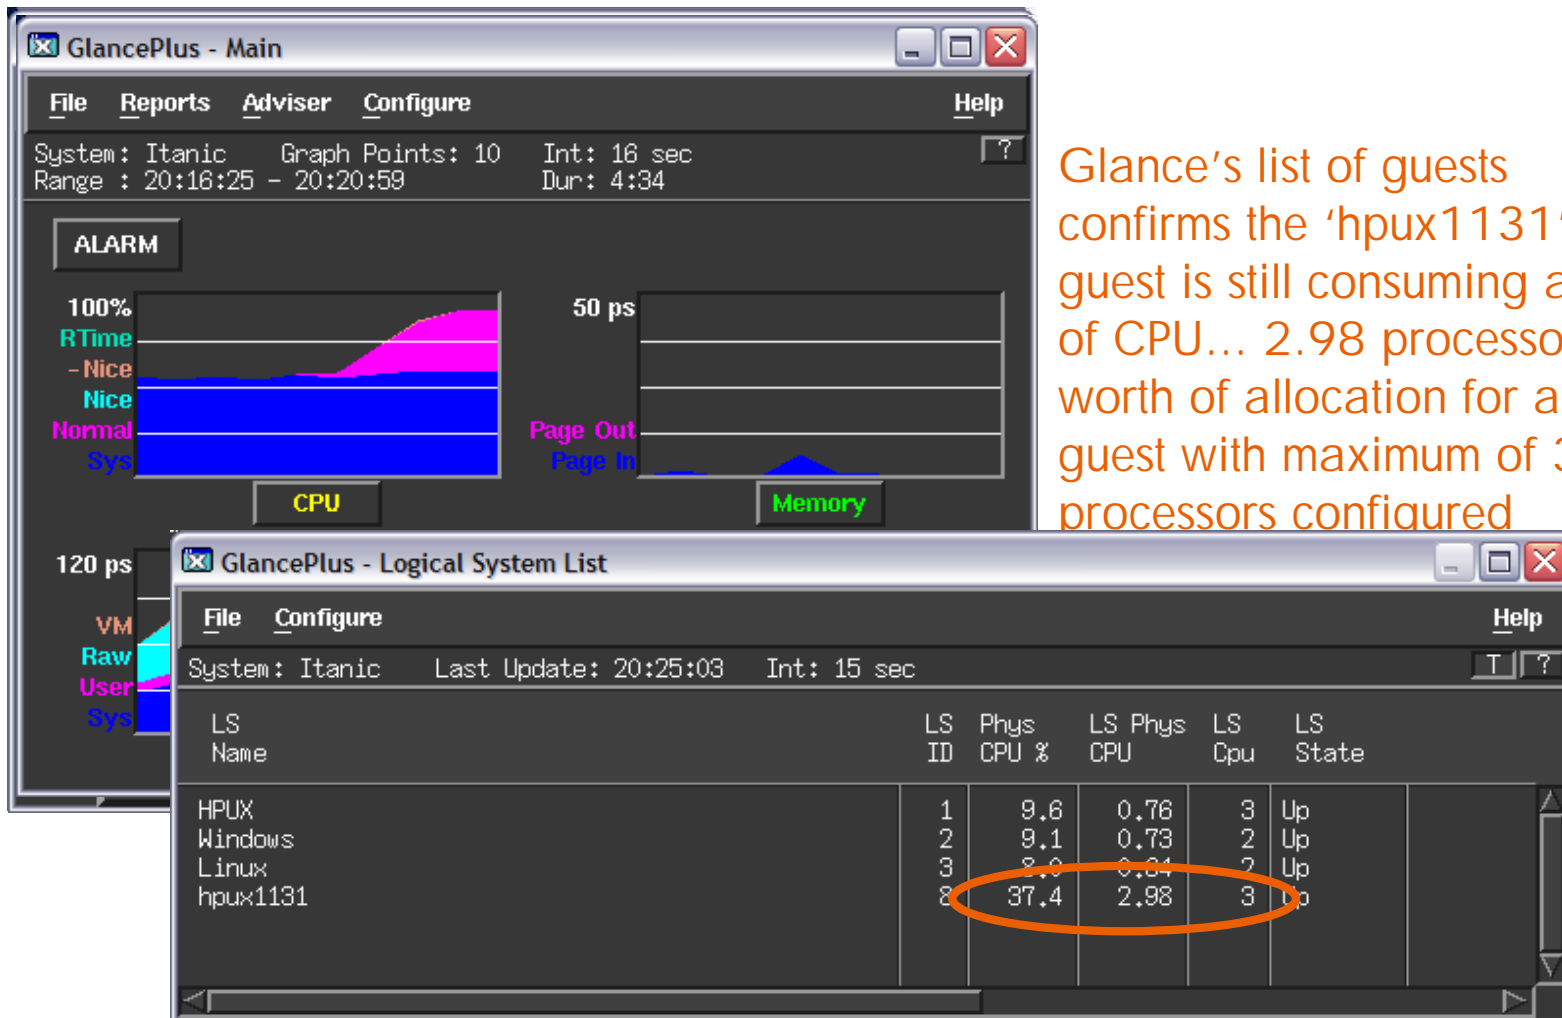
Performance Manager

The screenshot displays the Performance Manager Java Interface in a Microsoft Internet Explorer browser. The main window shows a 'Display' tab with a tree view of systems (itania, itanic, lia, linf) and a 'Graphs' section. A graph titled 'Utilization of Physical CPU Entitled' is visible, showing CPU utilization over time. A secondary window titled 'LS drilldown from itanic' is open, displaying a table of LS utilization data.

TIME	LS Name	LS ID	Physical CPU Util	LS State
8/15/07 8:07:00 PM	Windows	2	9.4	Up
8/15/07 8:07:00 PM	Linux	3	0.65	Up
8/15/07 8:07:00 PM	hpux1131	8	37.45	Up
8/15/07 8:08:00 PM	HPUX	1	8.35	Up
8/15/07 8:08:00 PM	Windows	2	9.94	Up
8/15/07 8:08:00 PM	Linux	3	0.66	Up
8/15/07 8:08:00 PM	hpux1131	8	37.32	Up
8/15/07 8:09:00 PM	HPUX	1	7.28	Up
8/15/07 8:09:00 PM	Windows	2	9.76	Up
8/15/07 8:09:00 PM	Linux	3	0.65	Up
8/15/07 8:09:00 PM	hpux1131	8	37.42	Up
8/15/07 8:10:00 PM	HPUX	1	7.74	Up
8/15/07 8:10:00 PM	Windows	2	9.38	Up
8/15/07 8:10:00 PM	Linux	3	0.64	Up
8/15/07 8:10:00 PM	hpux1131	8	37.37	Up

PM can zoom graphically or in table form.

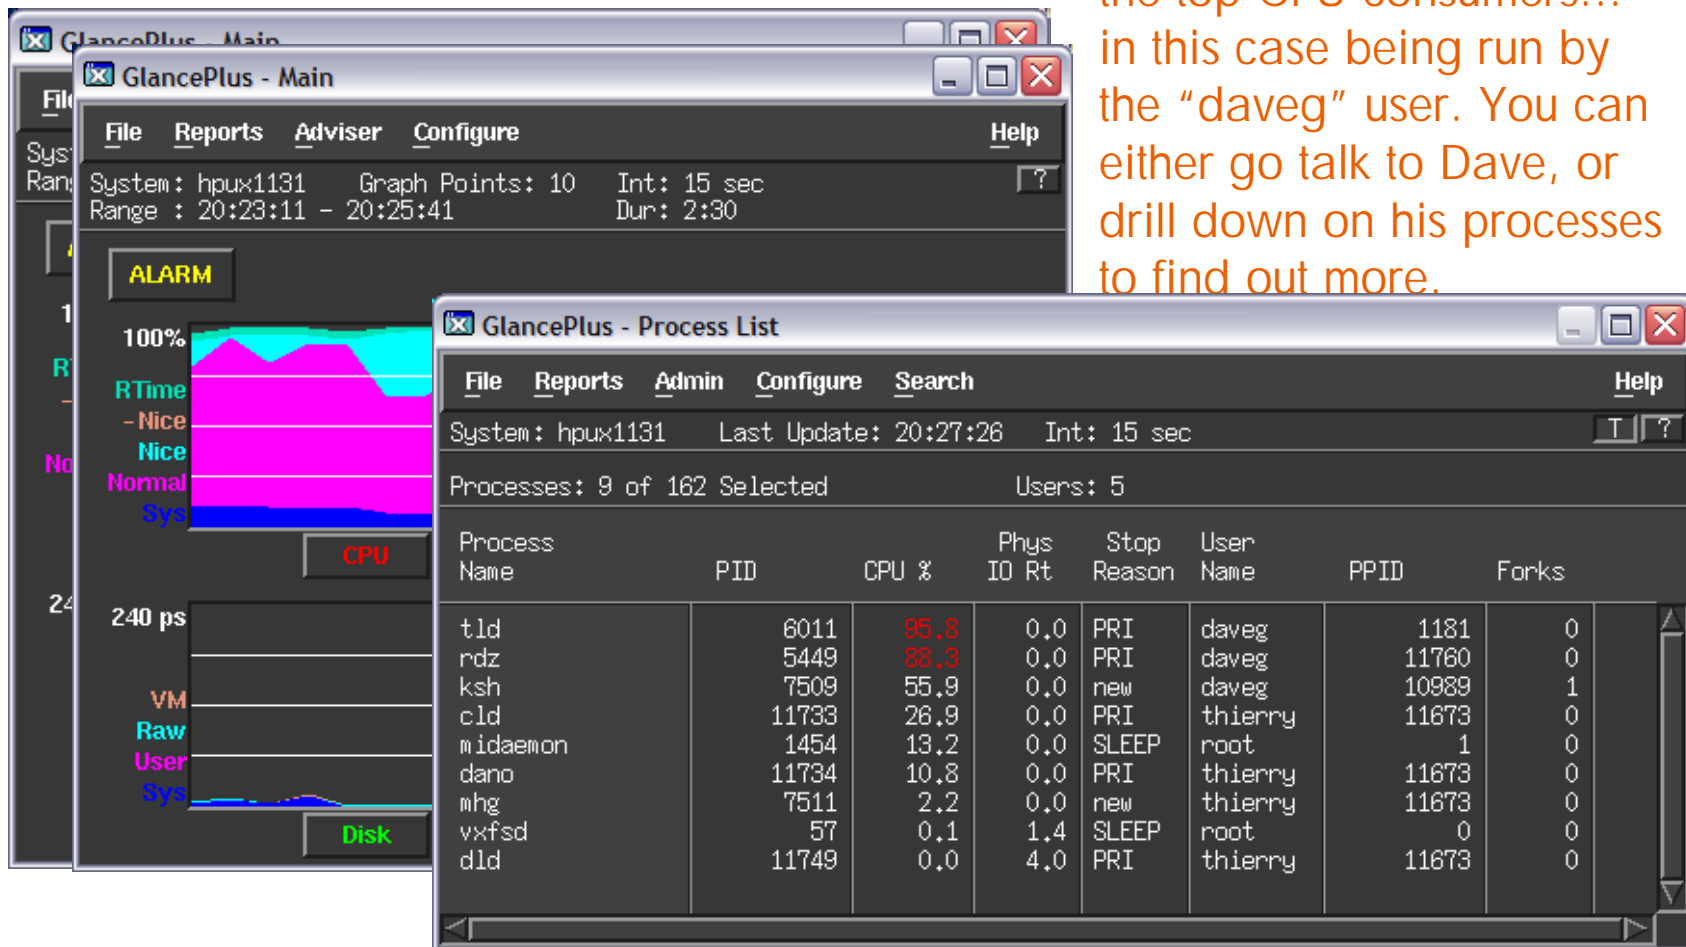
Glance on the server



Glance's list of guests confirms the 'hpux1131' guest is still consuming a lot of CPU... 2.98 processor's worth of allocation for a guest with maximum of 3 processors configured

Glance on the guest

The process list highlights the top CPU consumers... in this case being run by the "daveg" user. You can either go talk to Dave, or drill down on his processes to find out more.



Glance drilldown to process detail

The screenshot shows the GlancePlus interface with the 'Process System Calls' window open. The window displays the following information:

- System: hpux1131, Last Update: 20:38:38, Int: 15 sec
- Process Name: rdz, PID: 5449, PPID: 11760, User: daveg, State: active
- System Calls: All 1 Selected
- Interval for collection: 15 sec, Time for collection: 1:46
- Summary table:

SysCall Name	SysCall ID	SysCall Count	SysCall Rate	Elapsed Time
read	3	3766009	251067.0	8.89083

The 'SysCall Rate' value of 251067.0 is circled in orange. Below the table, the command line shows: `Cmd : ./rdz -v`

System call details on a specific process can give insight into the internal program behavior – in this case, Dave’s program is doing nothing but making over 200,000 read calls every second. Probably a mistake!

Tips for CPU bottlenecks

- When CPU is running higher than normal, and there is a queue (run queue or processes blocked on Priority), look for activity different from normal.
- Look at what applications and processes are doing (user / system mode CPU), and what they are waiting for (stop reasons).
- If you eliminate CPU hogs and runaways, look to adding horsepower to the system, or load balancing (psets, partitioning, PRM, WLM).
- Memory and I/O bottlenecks can masquerade as System CPU bottlenecks.

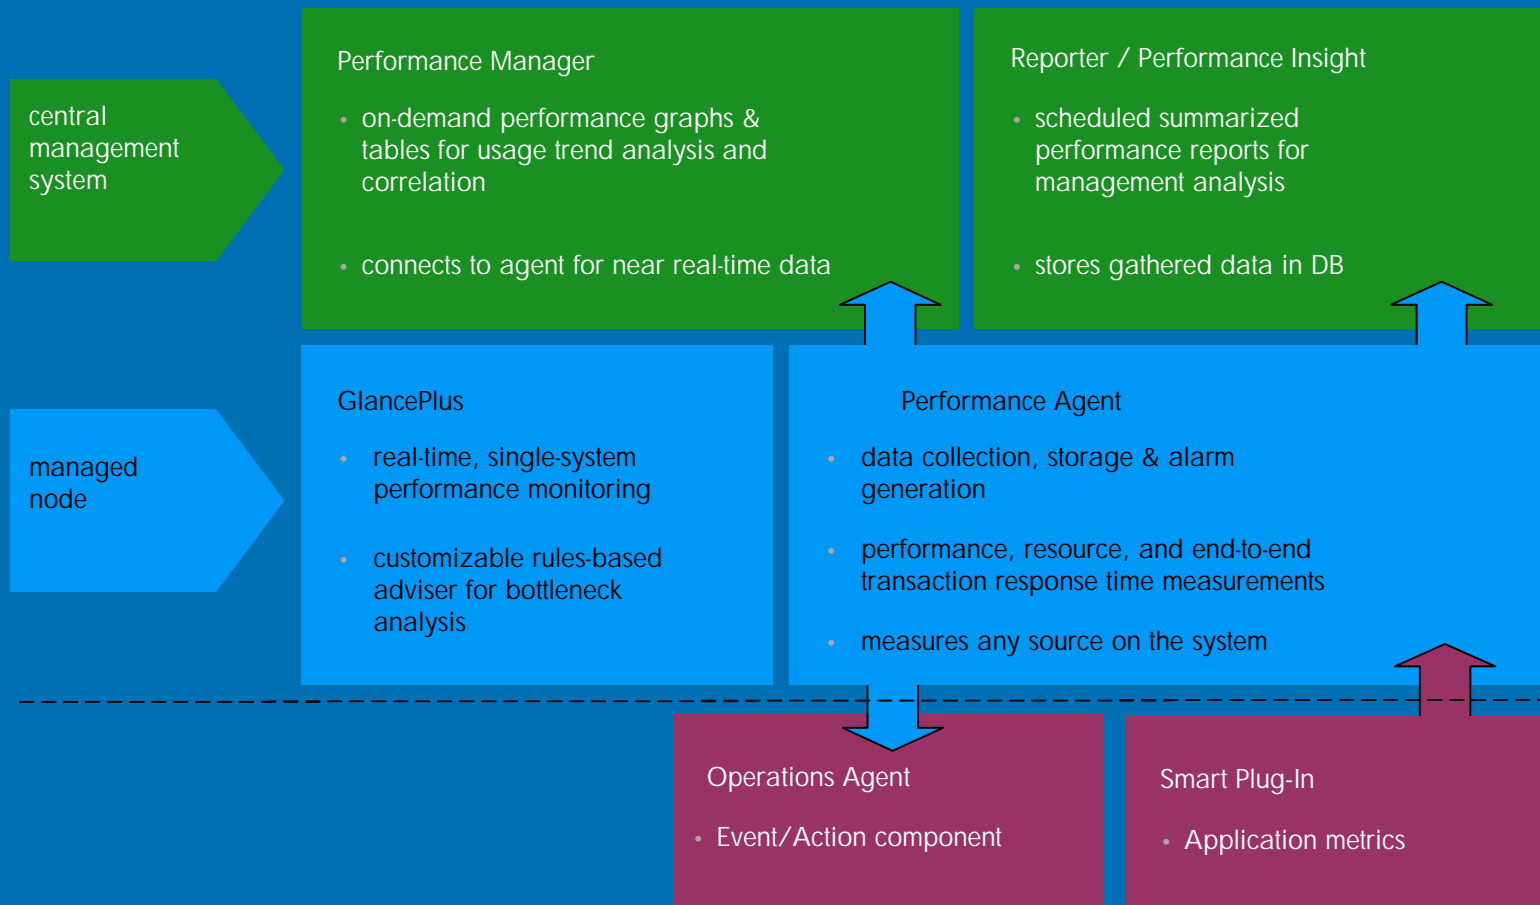
Tips for Memory bottlenecks

- High memory utilization **along with** pageout activity are signs of a memory bottleneck.
- Sort the glance process list by Virtual Set Size to catch “hogs”: watch for growth of DATA memory regions (heap growth).
- Prior to HP-UX 11.31, do not let Buffer Cache use too much memory. On HP-UX 11.31, monitor the File Cache (it is included in User Memory). Neither of these should ever be allowed to cause paging.
- More memory is often a wise investment for performance!

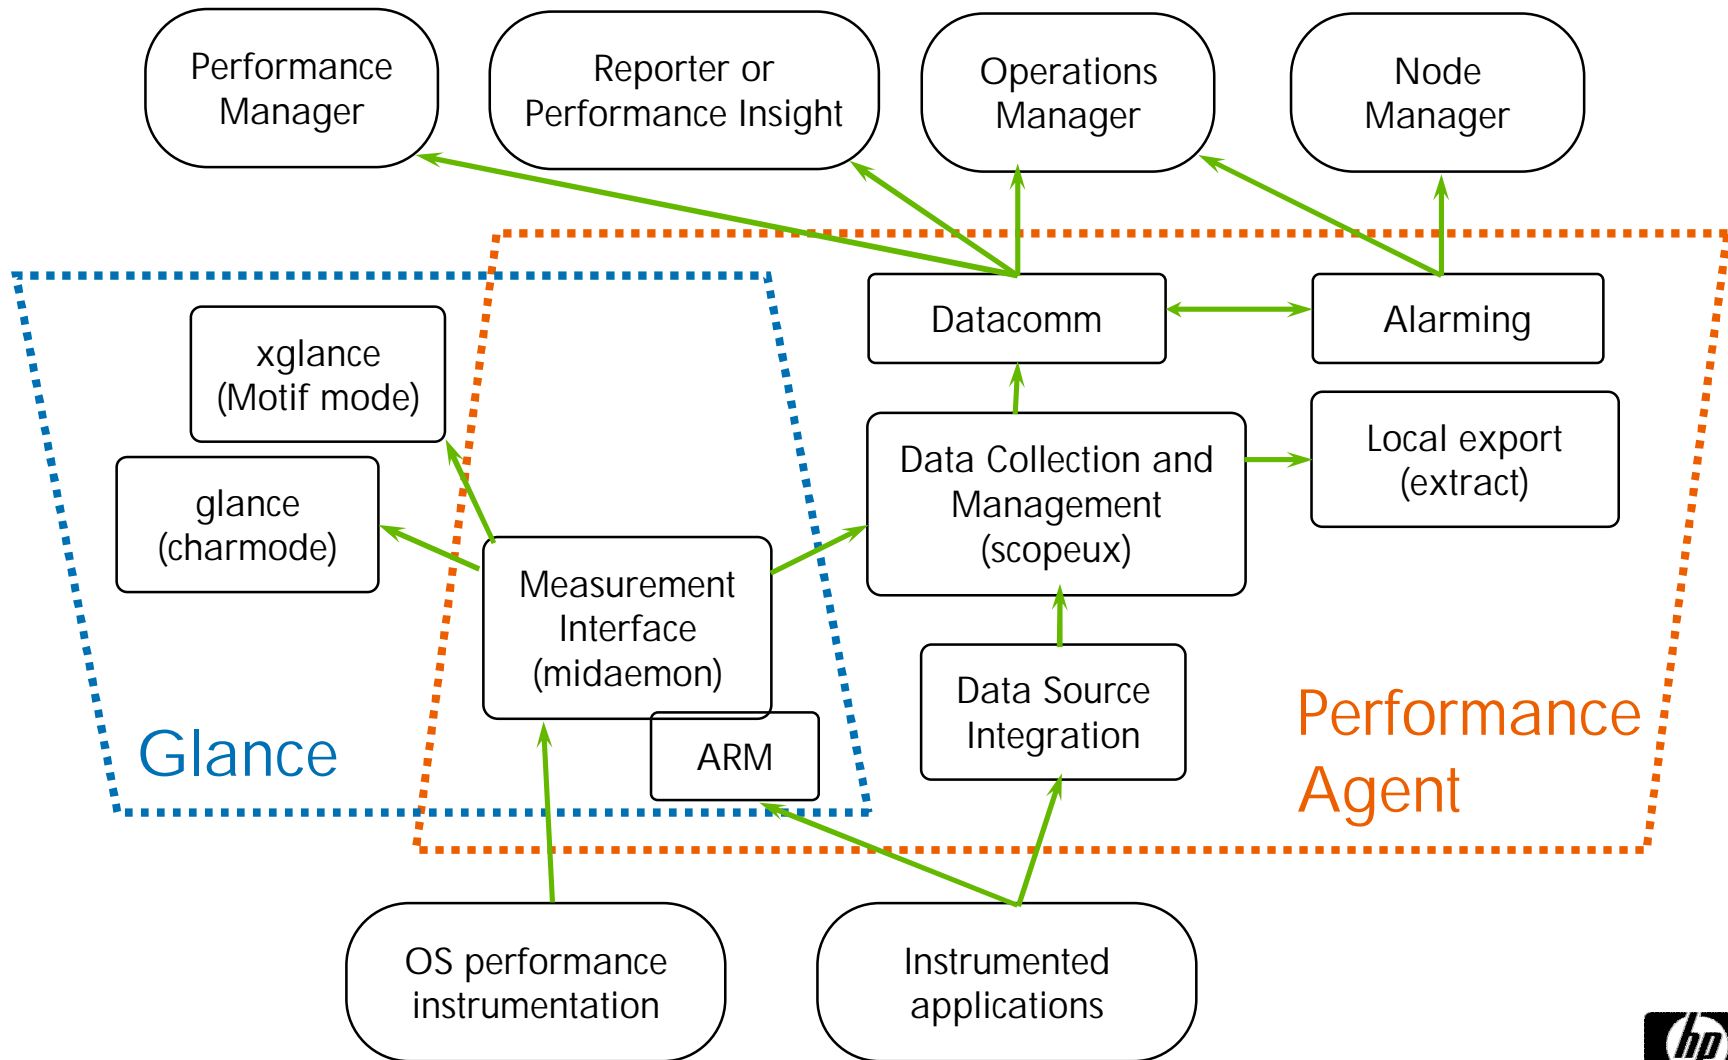
Tips for I/O bottlenecks

- In the Disk List, look for the busiest disks (over 50% busy utilization). In xglance, choose metrics to add the newer Service Time metrics. Consider changing configuration to offload “hot” devices.
- New for HP-UX 11.31 Glance: Investigate I/O by “HBA” (I/O Channel).
- Sort processes in Glance by Physical I/O rate to see top consumers. Drill down into their Open Files and see what System Calls are most active. Consider application tuning methods.

Dataflow and References



Glance and PA dataflow



Common measurement technology

- Both Glance and Performance Agent share common measurement software component (“nums” and midaemon)
 - shared metric set
 - cross-platform consistency
 - common bottleneck definitions
 - common parm file for application definitions
 - similar alarm syntax and default alarms
- On HP-UX we have unique performance measurement technology superior to all other system vendors in depth and efficiency

Glance and PA strengths

- Installed base larger than any performance tool competitor (over 1 million licenses sold, and >80% penetration rate on HP-UX servers).
- Available on major server platforms including HP-UX, Linux, Windows (PA-only), Solaris, AIX.
- Superior metrics both in depth (detail), usability (relevance), reliability (accuracy), and meaning (definitions).
- “Software Durability”; providing customer value for nearly two decades with >60 update releases on 6 platforms.
- Integration with HP management software ecosystem.

External web references

- Glance:
<http://managementsoftware.hp.com/products/gplus/>
- Performance Manager and Agent:
<http://managementsoftware.hp.com/products/ovperf/>
- Consolidated event and Performance Management:
http://managementsoftware.hp.com/solutions/ev_prf/
- Product docs: http://ovweb.external.hp.com/lp/doc_serv