Performance Troubleshooting

Ken Johnson
Escalation Engineer

Hewlett-Packard
100 Mayfield Ave MS 37UM
Mountain View CA 94043

fax: (650) 691-3187
knjohnson@hp.com

Purpose of session

To share the strategies and tactics used by the HP Crisis Management Team to resolve performance escalations - using real world examples and case studies.

We will not deal with system tuning, capacity planning or benchmarking.

The CMT Perspective: Emergency Room

- We do triage to stop the blood

- Fix the system quickly - otherwise ship the patient out

- We have a system perspective

  Understand interactions between HW/OS/Network/DB/Application

Define the performance problem

- Is this a user/business impacting problem or a metric-only issue ?

- Isolate the problem

system wide	-or-	particular application
all the time	-or-	specific time of day
network access	-or-	local access
nfs mounts	-or-	local disks
consistent	-or-	erratic

- Quantify whenever possible

- Measure the objective effect of changes

Know what your thermometers are measuring

- Different tools may use the same terminology

- Identify what a metric is really measuring

  wait time / service time
  page out / swap out
  run queue / load average

- Always have more than one data point

Is the work necessary ?

- Is the IO demand efficient ?
  example: missing index after reorg

- Are the cpu cycles necessary ?
  example: spinning while waiting for shared resource
  unnecessary context switching

- Is the application efficient ?
  example: excessive semop calls in a db

- Is memory utilization necessary ?
  example: maxuser set high
  buffer cache set to default 50%

Look for anomalies

- System call rates / cpu utilization
  example: shell script in a loop - vfork high

- IO patterns by device and time of day
  example: database tracing - IO rates to trace files

- Wait states - both global and per process
  example: semop waits for database process

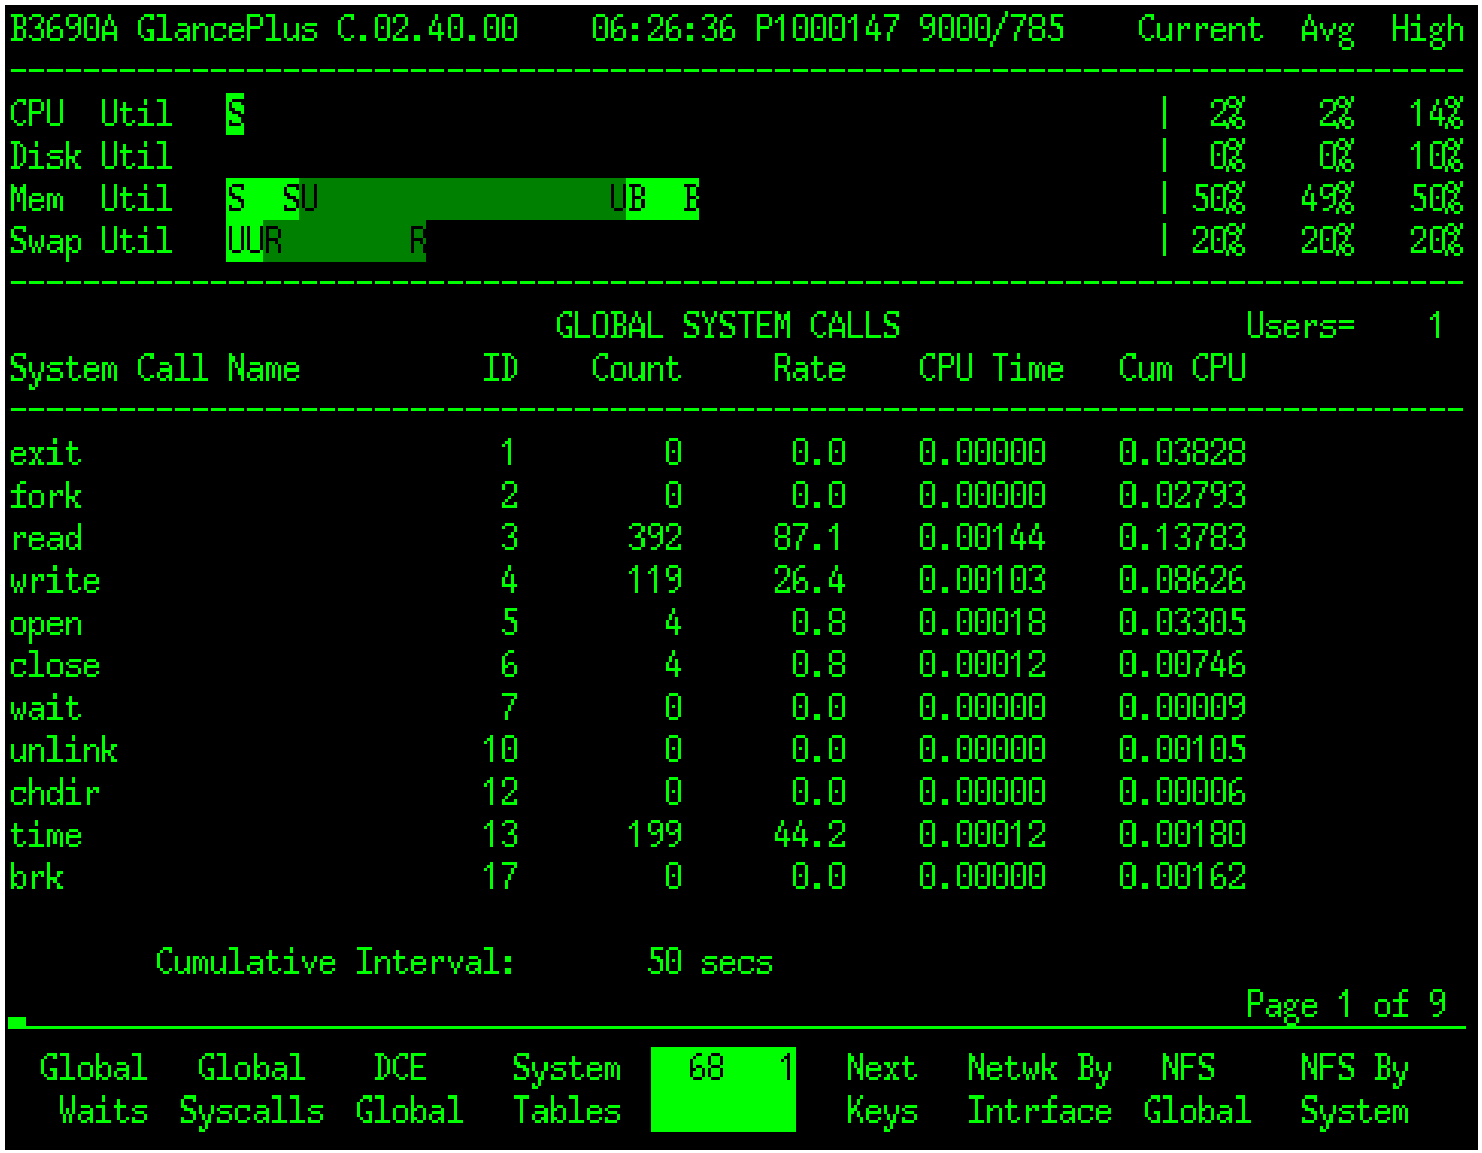
B3690A GlancePlus C.02.40.00 06:26:36 P1000147 9000/785 Current Avg High
CPU Util | 2% 2% 14%
Disk Util | 0% 0% 10%
Mem Util | 50% 49% 50%
Swap Util | 20% 20% 20%
GLOBAL SYSTEM CALLS Users= 1
System Call Name ID Count Rate CPU Time Cum CPU
exit 1 0 0.0 0.00000 0.03828
fork 2 0 0.0 0.00000 0.02793
read 3 392 87.1 0.00144 0.13783
write 4 119 26.4 0.00103 0.08626
open 5 4 0.8 0.00018 0.03305
close 6 4 0.8 0.00012 0.00746
wait 7 0 0.0 0.00000 0.00009
unlink 10 0 0.0 0.00000 0.00105
chdir 12 0 0.0 0.00000 0.00006
time 13 199 44.2 0.00012 0.00180
brk 17 0 0.0 0.00000 0.00162
Cumulative Interval: 50 secs
Page 1 of 9
Global Waits
Global Syscalls
DCE Global
System Tables
68 1
Next Keys
Netwk By Intrface
NFS Global
NFS By System

Look for anomalies

- System call rates / cpu utilization
  example: shell script in a loop - vfork high

- IO patterns by device and time of day
  example: database tracing - IO rates to trace files

- Wait states - both global and per process
  example: semop waits for database process

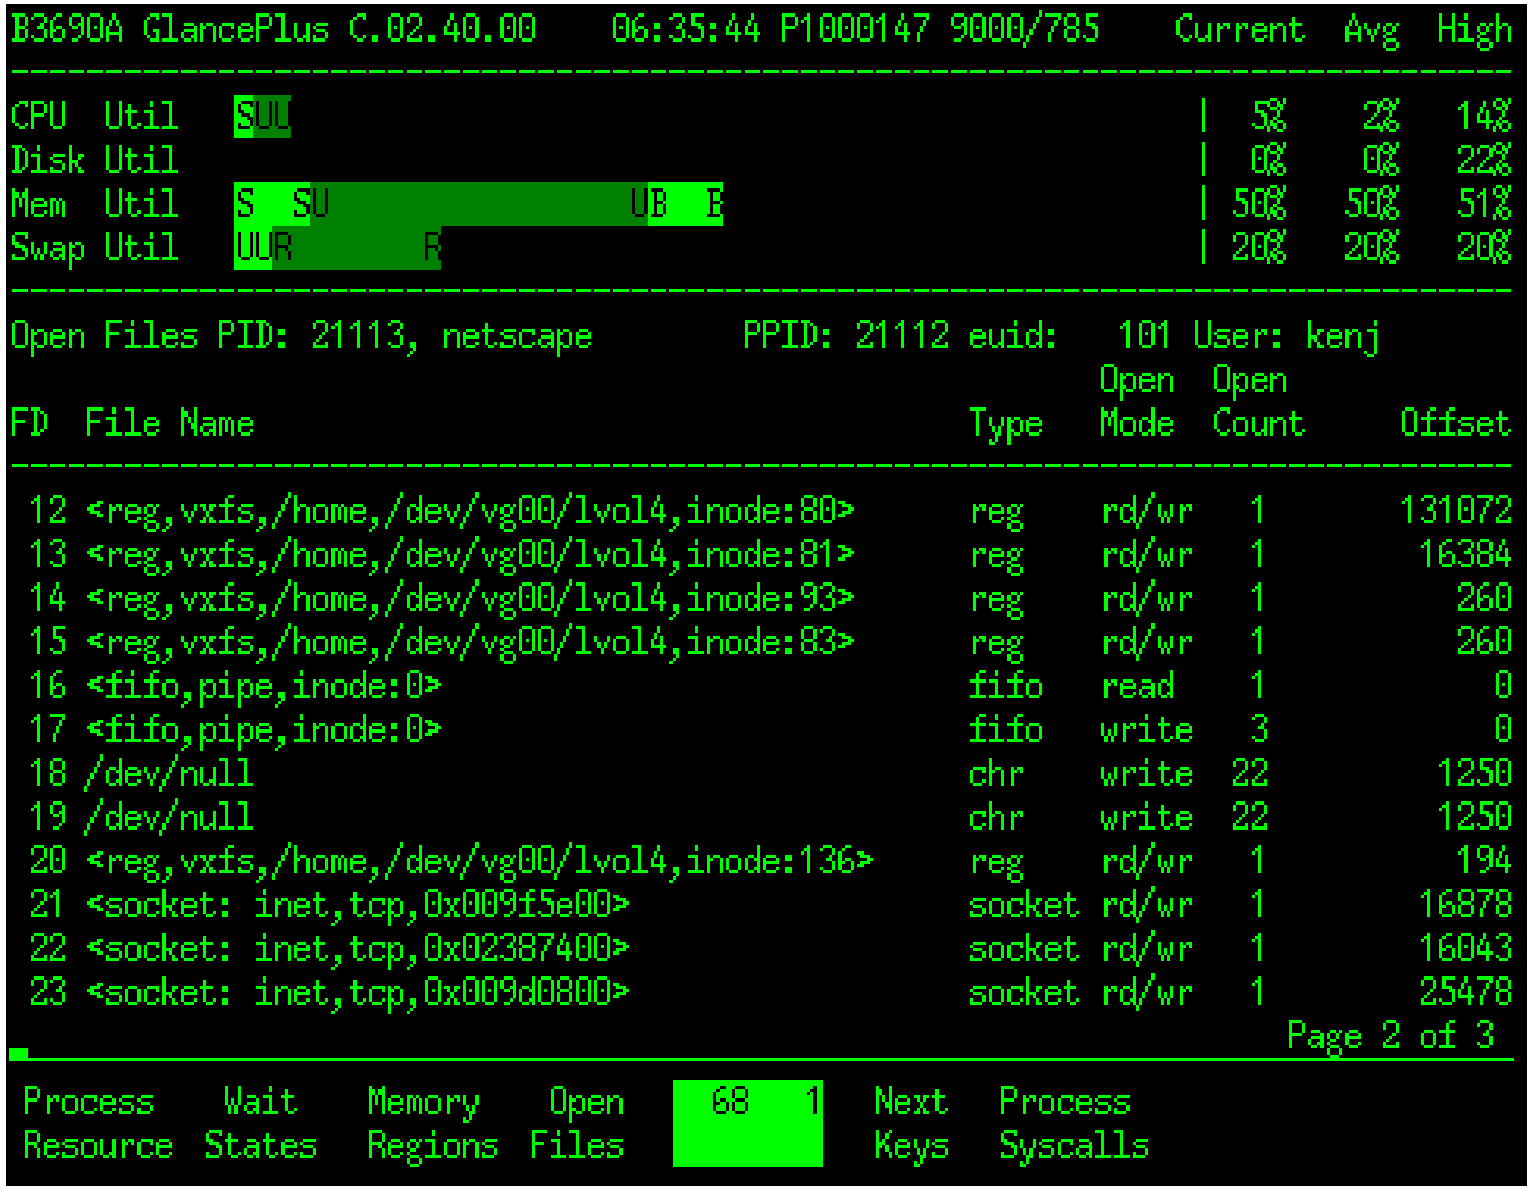
B3690A GlancePlus C.02.40.00 06:35:44 P1000147 9000/785 Current Avg High
CPU Util | 5% 2% 14%
Disk Util | 0% 0% 22%
Mem Util | 50% 50% 51%
Swap Util | 20% 20% 20%
Open Files PID: 21113, netscape PPID: 21112 euid: 101 User: kenj
FD File Name Type Open Mode Open Count Offset
12 <reg,vxfs,/home,/dev/vg00/lvol4,inode:80> reg rd/wr 1 131072
13 <reg,vxfs,/home,/dev/vg00/lvol4,inode:81> reg rd/wr 1 16384
14 <reg,vxfs,/home,/dev/vg00/lvol4,inode:93> reg rd/wr 1 260
15 <reg,vxfs,/home,/dev/vg00/lvol4,inode:83> reg rd/wr 1 260
16 <fifo,pipe,inode:0> fifo read 1 0
17 <fifo,pipe,inode:0> fifo write 3 0
18 /dev/null chr write 22 1250
19 /dev/null chr write 22 1250
20 <reg,vxfs,/home,/dev/vg00/lvol4,inode:136> reg rd/wr 1 194
21 <socket: inet,tcp,0x009f5e00> socket rd/wr 1 16878
22 <socket: inet,tcp,0x02387400> socket rd/wr 1 16043
23 <socket: inet,tcp,0x009d0800> socket rd/wr 1 25478
Page 2 of 3
Process Resource
Wait States
Memory Regions
Open Files
68 1
Next Keys
Process Syscalls

Look for anomalies

- System call rates / cpu utilization
  example: shell script in a loop - vfork high

- IO patterns by device and time of day
  example: database tracing - IO rates to trace files

- Wait states - both global and per process
  example: semop waits for database process

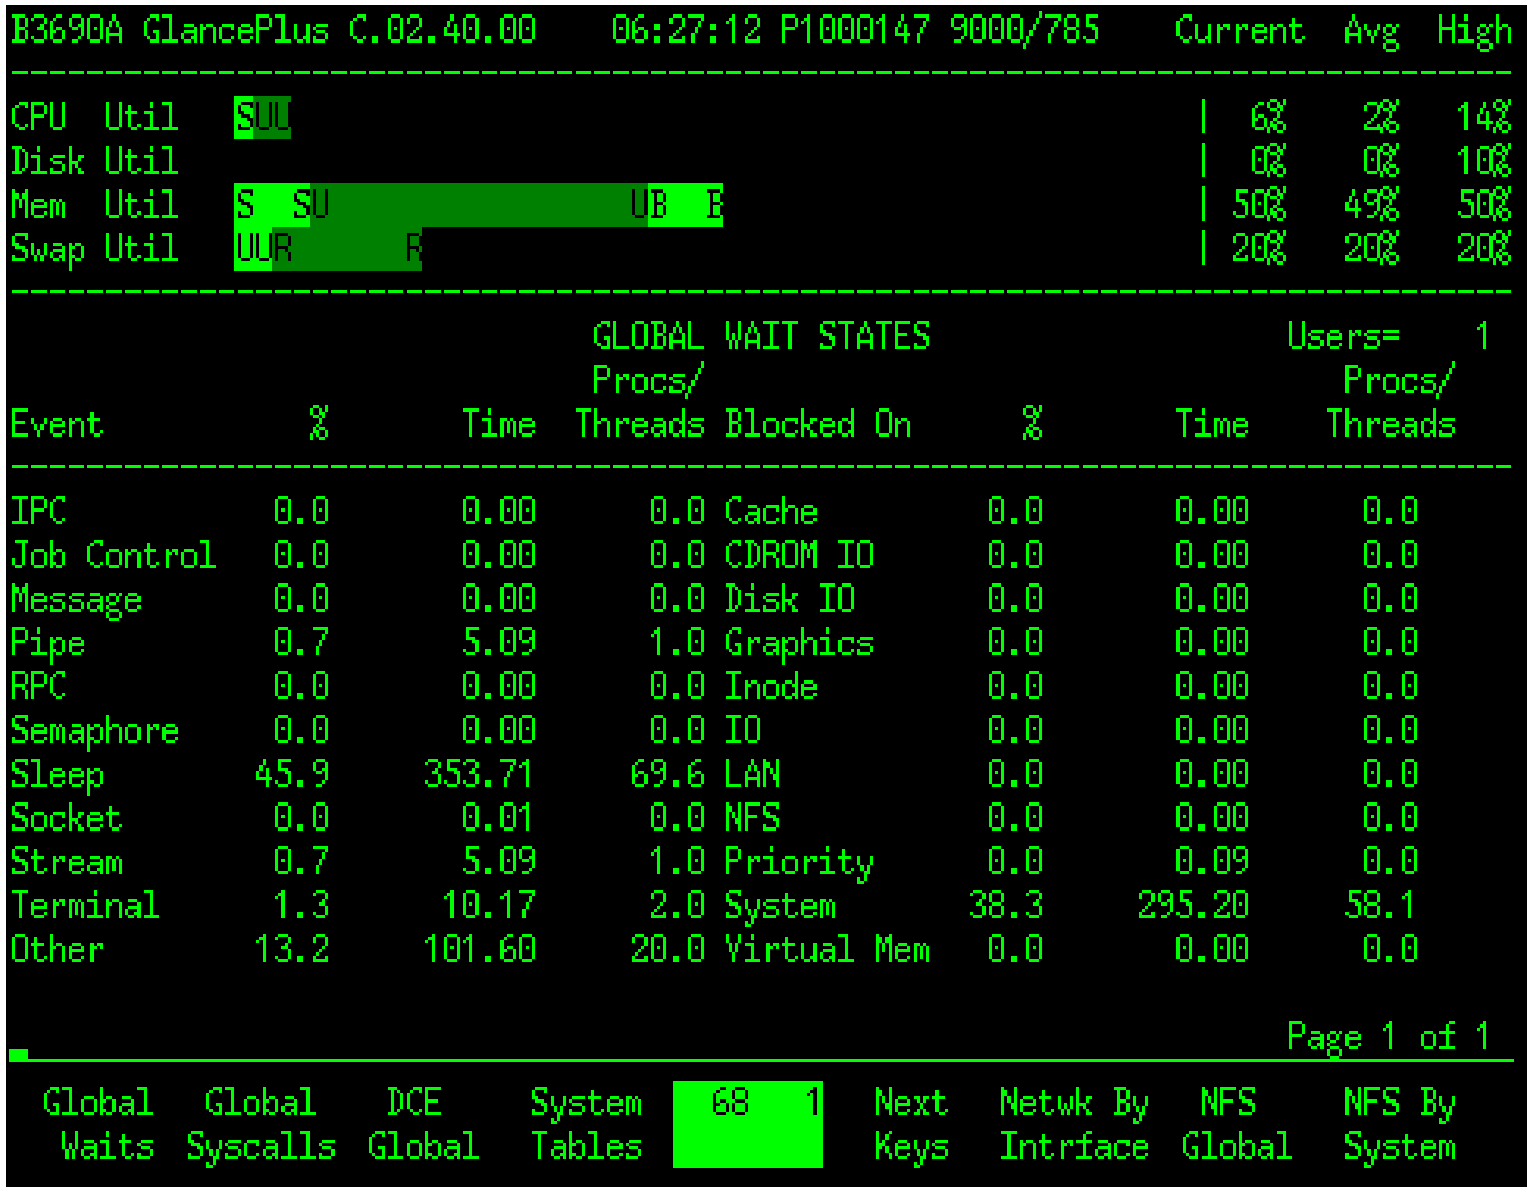
B3690A GlancePlus C.02.40.00 06:27:12 P1000147 9000/785 Current Avg High
CPU Util | 6% 2% 14%
Disk Util | 0% 0% 10%
Mem Util | 50% 49% 50%
Swap Util | 20% 20% 20%
GLOBAL WAIT STATES
Users= 1
Event % Time Procs/ Threads Blocked On % Time Procs/ Threads
IPC 0.0 0.00 0.0 Cache 0.0 0.00 0.0
Job Control 0.0 0.00 0.0 CDROM IO 0.0 0.00 0.0
Message 0.0 0.00 0.0 Disk IO 0.0 0.00 0.0
Pipe 0.7 5.09 1.0 Graphics 0.0 0.00 0.0
RPC 0.0 0.00 0.0 Inode 0.0 0.00 0.0
Semaphore 0.0 0.00 0.0 IO 0.0 0.00 0.0
Sleep 45.9 353.71 69.6 LAN 0.0 0.00 0.0
Socket 0.0 0.01 0.0 NFS 0.0 0.00 0.0
Stream 0.7 5.09 1.0 Priority 0.0 0.09 0.0
Terminal 1.3 10.17 2.0 System 38.3 295.20 58.1
Other 13.2 101.60 20.0 Virtual Mem 0.0 0.00 0.0
Page 1 of 1
Global Waits
Global Syscalls
DCE Global
System Tables
68 1
Next Keys
Netwk By Intrface
NFS Global
NFS By System

System/User cpu ratio

- What is system cpu ? Why is it important ?

- High system cpu can point to:

  - o high number of system calls
  - o memory / IO problems
  - o thrashing / spinning in kernel

- Identify any changes in individual system calls
  example: full directory - stat() system call

- CMT has utilities do kernel profiling on production systems

Isolate the components

- Isolate Buffer Cache effects from device IO rates
  example: measure IO rates to raw devices

- Take the network out of the picture
  example: make local queries rather than remote client queries

- Omniback debugging techniques
  example: isolate disk IO, network, tape IO, data compressibility

Know Your System

- Internals knowledge of UX

- Internals knowledge of the application/database

Know Your System

- Transaction Reporting
  example: SAP instrumentation
  ARM instrumentation

- Maintain a history (sar/vmstat/scope/application measures)

- Develop an intuition for your systems

- Watch it closely when its healthy

- Know the performance pattern over the day/week/month

Identify benchmark processes on your system

- Make sure you are measuring the essential business processes

- Create a benchmark that can be run independently of the application
  example: IO starvation

CPU *Rules of Thumb*

- System cpu <= 30%

- Total cpu < 80%

- Small load average

MEMORY *Rules of Thumb*

- Never page out

- Never deactivate processes

- Buffer Cache < 500mb

IO *Rules of Thumb*

- Utilization < 50% on any drive

- Minimal queuing < 4

- Response time ~10 milliseconds

NETWORK *Rules of Thumb*

- Rarely see network bottlenecks

- DNS issues

  - example: 15 second DNS delay prevented SG startup
  - example: DNS failure causing slow db/application startup