



**i n v e n t**

Presenter

# Crisis Management Team Performance Troubleshooting

Ken Johnson  
Escalation Engineer

Hewlett-Packard Company  
100 Mayfield Avenue, MS 37UM  
Mountain View, CA 94043

E-mail: [ken\\_johnson@hp.com](mailto:ken_johnson@hp.com)  
Fax: (650) 691-3187



Purpose

# Crisis Management Team Performance Troubleshooting

- To share the strategies and tactics used by the HP Crisis Management Team (CMT) to resolve performance escalations – using real world examples and case studies
- We will not deal with system tuning, capacity planning or benchmarking

# Crisis Management Team Performance Troubleshooting

1. The CMT Perspective: Emergency Room
2. Initial Steps in a Performance Escalation
3. Is the Work Necessary?
4. Looking for Anomalies
5. Isolating Components
6. Knowing Your System
7. Rules of Thumb

# The C M T Perspective : Emergency Room

- ER the TV show -we do triage to stop the bleeding
- Stabilize the system as fast as possible
- Quickly identify the first steps
  - Is this an HP defect / config issue / 3<sup>rd</sup> party issue ?
  - Often our value-add is to point in the right direction
- We have a system perspective
  - Understand interactions between HW , OS , Network , DB , Application
  - We train our engineers for a system perspective

# Initial Steps in a Performance Escalation

- Defining the Performance Problem
- Metrics – Know What the Thermometers are Measuring
- First Metric: System / User CPU Ratio
- Is There a Bottleneck?

## Defining the Performance Problem

Things we want to know in the first minutes

- When did the performance problem start?
- How do you know you have a problem ?
  - Is this a user/business impacting problem or a metric-only issue ?
  - Is the problem quantified ?
- Is the hardware and OS base stable and consistent?
- What changed ?

## Defining the Performance Problem

- Quantification
  - Allows you to measure the objective effect of changes
  - Define the current state and the goal
- Changing only one thing at a time
- Characterize 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?



## Metrics – Know What the Thermometers are Measuring

- Metrics are simply statistics produced by software
  - Some of our escalations are with performance tools
- Be sure what a metric is really measuring
  - Wait time/service time
  - Page out/swap out
  - Run queue/load average
  - Inode table utilization
- Always have more than one data point and always use more than one tool
- Your tools can affect the environment

## First Metric: System User CPU Ratio

- What is system CPU?
- Why is it important?
  - Points in initial directions to pursue root cause
  - HP owns this code
- High system CPU can point to:
  - High number of system calls
  - Memory/I/O problems
  - Thrashing/spinning in the kernel
- CMT has visibility into system CPU utilization
  - There are utilities we use to do kernel profiling on production systems

## Is There a Bottleneck?

- This is the supply side of performance
- Easiest to look at - easiest to fix
- IO
  - Is there queuing on any drives?
  - Are there long service wait times on any drives?
- CPU
  - Is there a significant load average?
  - Is system CPU high?
  - Are processes priority waited?
- Memory
  - Is there any paging or deactivations?
  - Is there significant swap utilization?

## Is the Work Necessary?

- Is the I/O Demand Efficient?
- Are the CPU Cycles Necessary?
- Is the Application Efficient?
- Is the Memory Utilization Necessary?

## Is the IO Demanded Efficient?

System :	K580 4-way 11.0 2 GB memory Database server	Manufacturing
Symptoms:	2 year installation Suddenly batch jobs taking much longer to execute No changes to programs or database settings No system bottlenecks Elevated IO rate but no queuing and fast service times	
Diagnosis:	Cost based plan had been used for key queries in DB Several of the queries started doing serial IO	

## Are the CPU Cycles Necessary?

System :	T600 8-way 10.20 2 GB memory Development system Compiling and source code management
Symptoms:	3 year installation Recently seeing slow overall performance Intermittent High system CPU and high context switch rates
Diagnosis:	Files used for compiling were located in one directory Large number of files and very volatile Contention around the directory file itself (25 MB) Spinning while waiting for shared resource caused unnecessary context switching

## Is the Application Efficient?

System :                    N 4000 4-way 11.0  
                                  4 GB memory  
                                  Web server

Symptoms:                New installation  
                                  Server throughput was never acceptable  
                                  High CPU utilization with mostly user CPU  
                                  Load average was reasonable and good  
                                  system response time

Diagnosis:                Identified large # of sem op calls in bolt-on application  
                                  Application was in the critical path for the server  
                                  Allowed vendor to identify configuration problem

## Is the Memory Utilization Necessary?

System : V2500 16-way 11.0  
4 GB memory  
Database server

Symptoms : Memory utilization at 100%  
High page out and deactivation rates

Diagnosis : Default 50% buffer cache had been used  
Max user had been set very high - affects many other kernel variables  
Final solution was to add memory and to tune kernel variables



## Looking for Anomalies

- System Call Rates/CPU Utilization
- IO Patterns
  - By Device
  - By Time of Day
  - By Process
- Wait States
  - Global and Per Process

## System Call Rates / CPU Utilization

System :                    K460 4-way 10.20  
                              2 GB memory  
                              Legacy shellscript-based application  
                              Files find in , processed , then put in a directory for  
                              pickup

Symptoms:                 Suddenly application throughput was down  
                              No changes to the application  
                              System CPU way up

Module 4

```

B3690A GlancePlus C.02.40.00 06:26:36 P1000147 9000/785 Current Avg High
-----
CPU Util 5 | 2% 2% 14%
Disk Util | 0% 0% 10%
Mem Util S SU UB E | 50% 49% 50%
Swap Util UUR R | 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 Global DCE System 68 1 Next Netwk By NFS NFS By
Waits Syscalls Global Tables Keys Interface Global System

```



## System Call Rates / CPU Utilization

- System : K460 4-way 10.20  
2 GB memory  
Legacy shellscript-based application  
Files found in, processed, then put in a directory for pickup
- Symptoms: Suddenly application throughput was down  
No changes to the application  
System CPU way up
- Diagnosis: vfork() was very large CPU consumer  
Identified shellscript that was in a loop

# I/O Patterns - By Device, By Time of Day, By Process

System :            N4000 4-way 11.0  
                      4 GB memory  
                      Database server for web front-end

Symptoms:        New installation  
                      System response was good  
                      Unacceptable database performance  
                      DB connections were short-lived  
                      Analysis showed that delay was in DB disconnect

Module 4

```

B3690A GlancePlus C.02.40.00    06:35:44 P1000147 9000/785    Current Avg High
-----
CPU Util    SUU                    | 5%    2%    14%
Disk Util   | 0%    0%    22%
Mem Util    S  SU                    UB  E    | 50%   50%   51%
Swap Util   UUR                    F    | 20%   20%   20%
-----
Open Files PID: 21113, netscape      PPID: 21112  euid: 101  User: kenj
                                     Open  Open
FD  File Name                          Type  Mode  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  Wait  Memory  Open  68  1  Next  Process
Resource States  Regions  Files  Keys  Syscalls

```



# IO Patterns - By Device, By Time of Day, By Process

- System : N4000 4-way 11.0  
4 GB memory  
Database server for web front-end
- Symptoms: New installation  
System response was good  
Unacceptable database performance  
DB connections were short-lived  
Analysis showed delay was in DB disconnect
- Diagnosis: Used Glance to observe when user disconnected  
Found high rates of IO during disconnect  
IO was to 2 database trace files

# Wait States – Both Global and Per Process

System : V2250 8-way 11.0  
8 GB memory  
Database server

Symptom : New installation  
Slow database throughput  
No system bottlenecks or high utilization



Module 4

```

B3690A GlancePlus C.02.40.00 06:27:12 P1000147 9000/785 Current Avg High
-----
CPU Util  SUU | 6% 2% 14%
Disk Util | 0% 0% 10%
Mem Util  S  SU          UB  E | 50% 49% 50%
Swap Util  UUR          R | 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 Global DCE System 68 1 Next Netwk By NFS NFS By
Waits Syscalls Global Tables Keys Intrface Global System

```



# Wait States – Both Global and Per Process

System : V2250 8-way 11.0  
8 GB memory  
Database server

Symptom : New installation  
Slow database throughput  
No system bottlenecks or high utilization

Diagnosis: Identified high semaphore waits  
Database tuning required

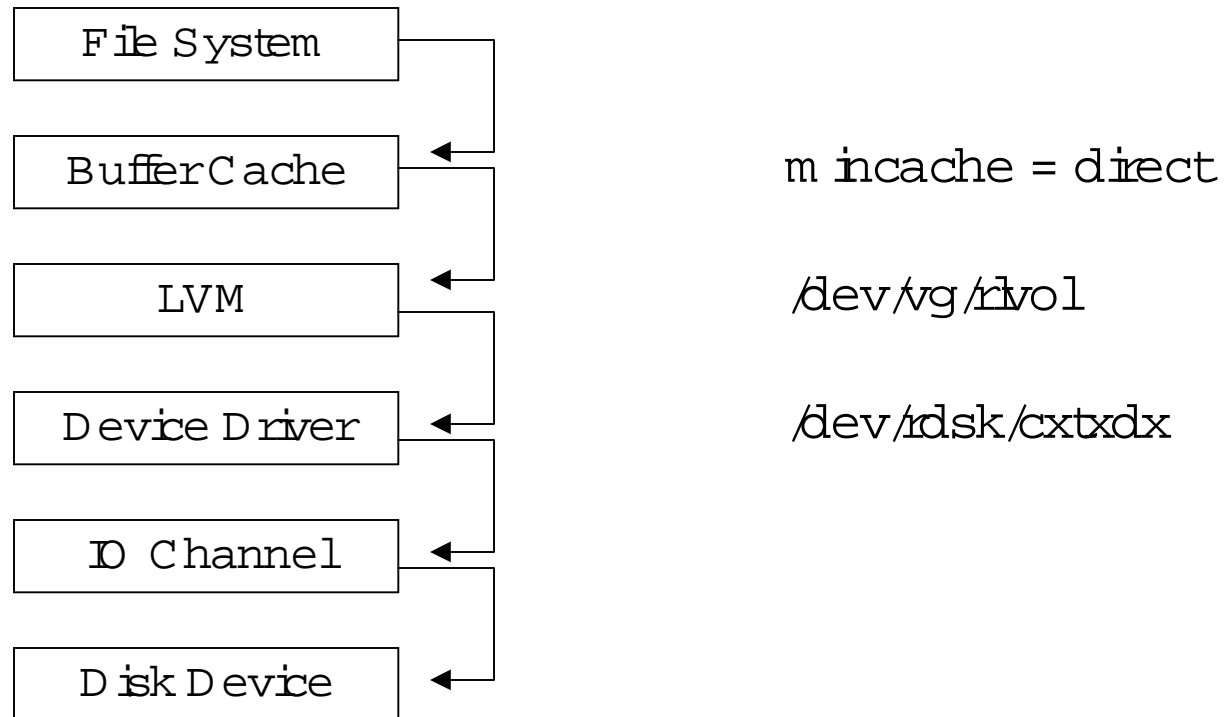
## Isolating Components

Make everything into a black box

Define and manipulate inputs and outputs

- The Discrete Elements of an IO Request
- Taking the Network Out of the Picture
- OmniBack Performance Debugging Techniques

# The Discrete Elements of an IO Request



# The Discrete Elements of an IO Request

- Bottlenecks can happen at any of the layers in either direction
- Isolate the IO test at one layer
  - `mountcachefilesystem = direct`
  - `/dev/vg/rvol`
  - `/dev/sdsk/cxtxdx`
- Only test reading or writing

# The Discrete Elements of an IO Request

Code Fragment for Timing IO Requests

```
#include <sys/time.h>

#define delta_tv(tv_0, tv_1) \
    (tv_1.tv_sec - tv_0.tv_sec + (tv_1.tv_usec - tv_0.tv_usec)/1000000.0)

struct timeval xtv0, xtv1;
struct tzone tz;
double rdt = 0.0;

main()
{
    gettimeofday(&xtv0, &tz);
    read(fd, buf, bufsize);
    gettimeofday(&xtv1, &tz);

    rdt = delta_tv(xtv0, xtv1);

    printf("milliseconds for read: % .3fms\n", 1000*rdt);
}
```

## Taking the Network Out of the Picture

- Multitiered applications (e.g. SAP) have large network components which can have a large impact on overall throughput
- Database access is often through sockets
- Techniques for isolation
  - Make local queries rather than client queries
  - With system issues execute problematic commands at the console
  - Use programs/benchmarks similar to those used for IO testing

## OmniBack Debugging Techniques

- Understand the capabilities of each component in the configuration
  - Examples:
    - T5xx ID backplane overwhelmed by more than one FW SCSI card
    - Modem DLTs need dedicated FW SCSI card/several disk-readers
- Isolate
  - Disk ID
  - Network
  - Tape ID
  - Updates to OmniBack database
  - Data compressibility



# OmniBack Debugging Techniques

## Measuring Data Compression

```
cat file | compress -v > /dev/null
```

```
$ cat dedbg | compress -v > /dev/null
```

```
Compression: 82.43%
```

```
$ cat dedbg | compress -v > dedbg.Z
```

```
Compression: 82.43%
```

## Knowing 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 it's healthy
- Know the performance pattern over the day/week/month
- Internals knowledge of the application/database
- Internals knowledge of the OS

## Rules of Thumb

- CPU
- MEMORY
- IO

## CPU Rules of Thumb

- System CPU  $\leq 30\%$
- Total CPU  $< 80\%$
- Small bad average

## MEMORY Rules of Thumb

- Never page out
- Never deactivate processes
- Buffer cache < 500 MB

## IO Rules of Thumb

- Utilization < 50% on any drive
- Minimal queuing < 4
- Response time ~10 milliseconds