

# Unix System Performance Analyzing and Understanding Unix System Resources

*Using today's data to plan for future,  
hassle-free processing.*

Lund Performance Solutions

240 Second Avenue SW  
Albany, OR 97321

Phone: (541) 812-7600

Fax: (541) 926-7723

[www.lund.com](http://www.lund.com)

Presented by:

D. Scott Pierson

Craig Lalley

Chris Reimer



# Analyzing System Resources

- Introduction
- Client-Server Environment
- Major Resources
  - CPU
  - Memory
  - Disk I/O
  - Network
- Questions & Answers

# Data Collection & Analysis Tools

- Standard Unix tools
  - sar, top, iostat, vmstat, netstat, uptime
- HP Tools
  - GlancePlus, MeasureWare, PerfView
- Lund Tools
  - SOS/9000, SOSLOGD, SOSLOGX, Performance Gallery Gold

# Workload Definition

- In MeasureWare, workloads are configured in `/var/opt/perf/parm`, and are called applications
- In SOS/9000, workloads are configured in `/etc/opt/lps/cfg/workdefs`

# Workload Definitions – Using MeasureWare

```
# Add user applications here

application = network
file = nfs*,biop,automount,inetd,snmp*,rpc*,llbd,netfmt,portmap
file = rbootd,telnet*,ftp*,*rlogin*,remsh*,rcp,nktl*,nvsisr,ttisr
file = losp,gosp,strmen,strweld,vtdaemon,mib*,trapdest*,*web*,xntpd,yp*
file = hp_unixagt,ntl*,pty*

application = perf tools
file=midaemon, glance, gpm, rx, scopeux, pv, extract, utility, mwa
file=alarmgen, rep_server, perflbd, agdbserver

application = memory management
file = swapper,vhand,syncer,pageout,fsflush,vxfsd

application = other user root
user = root
```

# Workload Definitions – Using SOS

```
# Define workload for performance tools
Perf_Tools
mix
PROG=lpsmid
PROG=sos
PROG=soslogd
PROG=soslogx
PROG=sar
PROG=iostat
PROG=vmstat

# Define INTERACT as all interactive processes.
INTERACT
INTERACT

# Define BATCH as non-interactive processes w/low nice value.
# This should be refined by the user.
BATCH
BATCH
NICE=21-39

# Define DAEMON as all remaining non-interactive processes.
# This should be refined by the user.
DAEMON
DAEMON
```

# Data Logging Interval

- MeasureWare has a fixed 5 minute interval for logging Global, Application, and Device data. Interesting Process data is logged every minute.
- SOS/900 has a user-defined logging interval, with a default value of 5 minutes. A snapshot of process data is taken when the other data is logged.

# Resources vs. Performance

- We usually measure *resource utilization* and say that we are measuring *performance*
- But there are only two ways to really measure Performance
  - Response Time
    - By Process
    - By Workload
    - System-wide
  - Throughput



# Automation

- Automatic reporting
- Analysis instead of reporting
- Export to web page

## Bobcat File Server Activity

Welcome to Bobcat's performance information page!

This page contains up-to-the-hour performance data about a live file server at LUND. It uses the automation features of one of our products, [Performance Gallery Gold](#), to automatically display performance data to the web.

See the [Q & A](#) section of this page for details on how this page was constructed.

### Quick Links

<a href="#">CPU Utilization</a>	<a href="#">File Server Statistics</a>	<a href="#">Session Information</a>	<a href="#">Disk Queue Length</a>
<a href="#">Network Traffic</a>	<a href="#">Paging Information</a>	<a href="#">Virtual Memory</a>	<a href="#">Automated Index Page</a>

### CPU Utilization

[Back to Top](#)

#### Bobcat File Server CPU Utilization (NT) 11/14/2001 03:01 - 11/14/2001 09:00

Time	CPU Utilization (%)
04:00	~5
05:00	~5
06:00	~5
07:00	~5
08:00	~5
09:00	~10


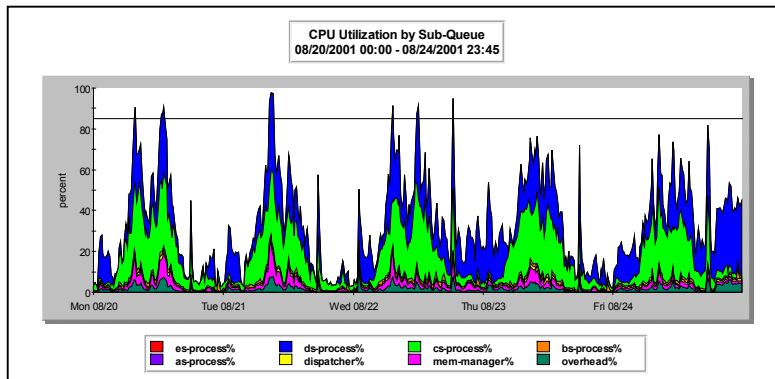
Bobcat's CPU utilization is generally quite low and rarely exceeds 20% during the course of an average work day.

Bobcat currently runs an Intel Pentium III 500 Mhz processor.


# Analysis Method — Implementing Change

- Identify issue
- Determine possible solutions
- Implement solution
- Test results

← Before After →



# Trending



# Client-Server Environment


- Definitions
- Client-Server Styles
- 3 Tier Architecture
- Problem Analysis
- Components of Response Time

# Client Server – Definitions

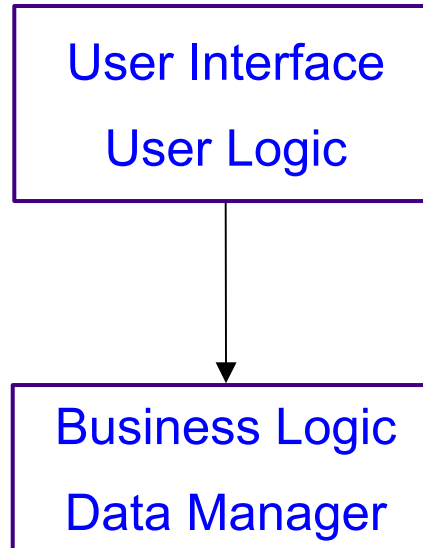
- Client-Server
  - When a transaction or function has components that execute on 2 or more computers
- Response Time
  - The elapsed time between hitting the ENTER or RETURN key and getting results back on the screen

# Client Server – Styles

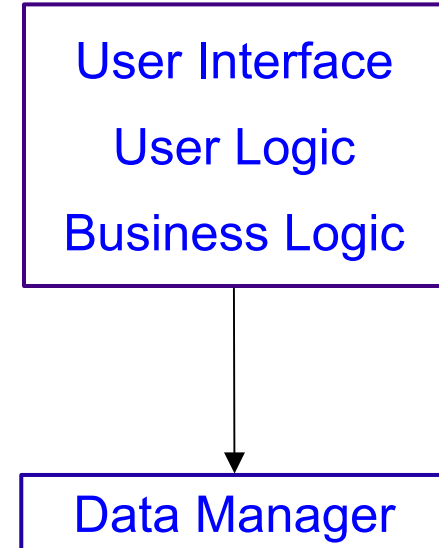
## Distributed Presentation



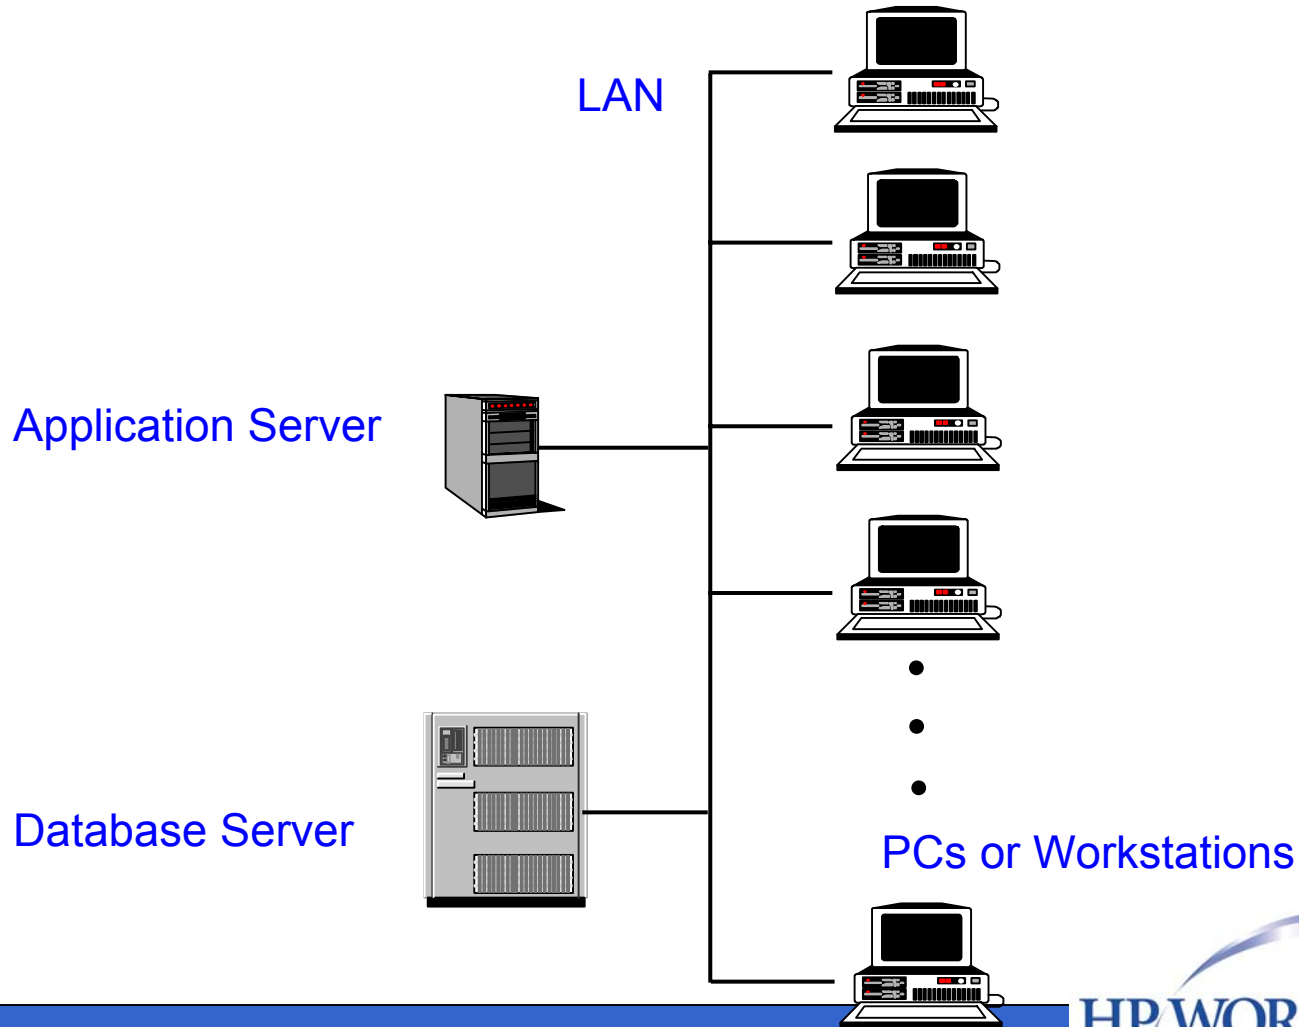
## Remote SQL



## Business Transaction



# Client Server – 3 Tier Architecture



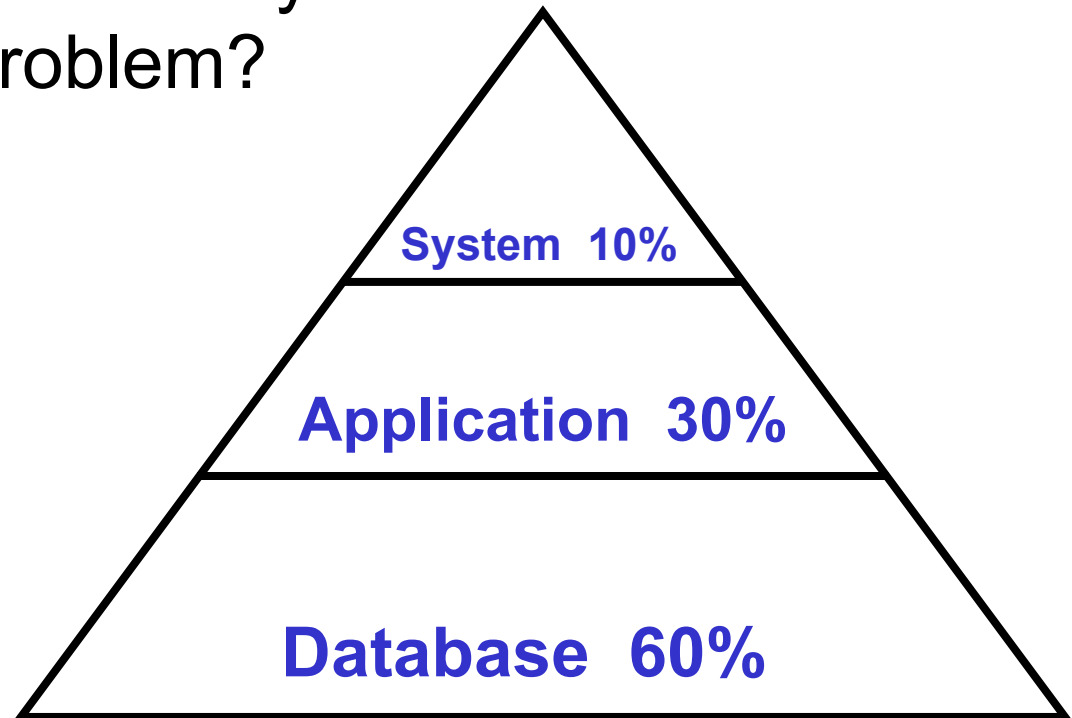
# Client Server – Problem Analysis

- Components to Analyze on Each System
  - CPU
  - Memory
  - Disk I/O
  - Network



# Client Server – Problem Analysis

Where are you most likely to find the source of a performance problem?




# Client Server – Problem Analysis

- Once you identify a performance problem, there are only two ways to resolve it:
  - Increase the resource
    - More Processors, Memory, Disk
    - Faster Processors, Disk
  - Decrease the demand
    - Tune the application
    - Tune the Operating System
    - Redistribute the work

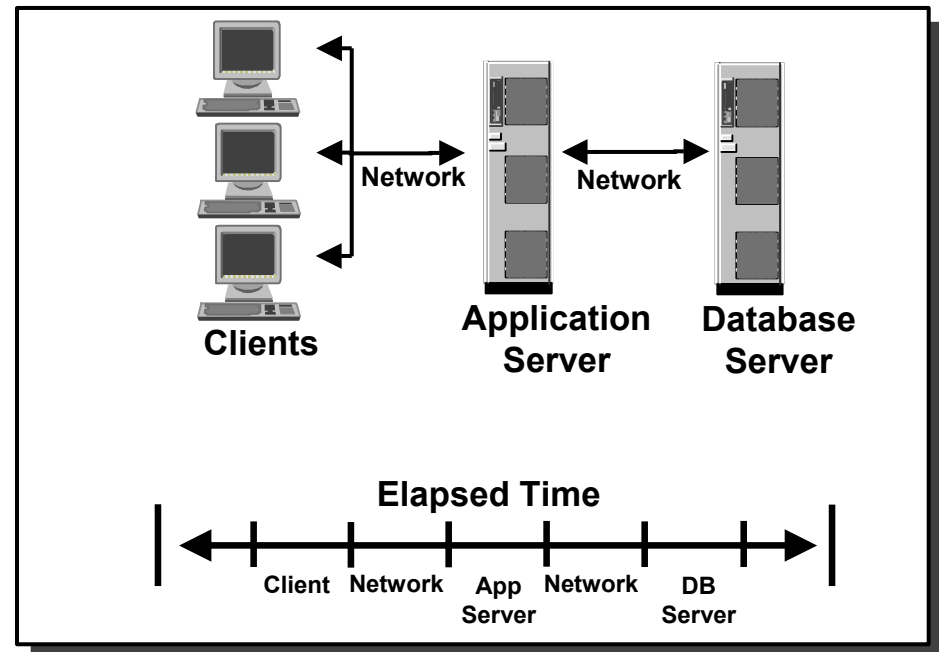
# Client Server – Response Time Component

## Centralized Environment



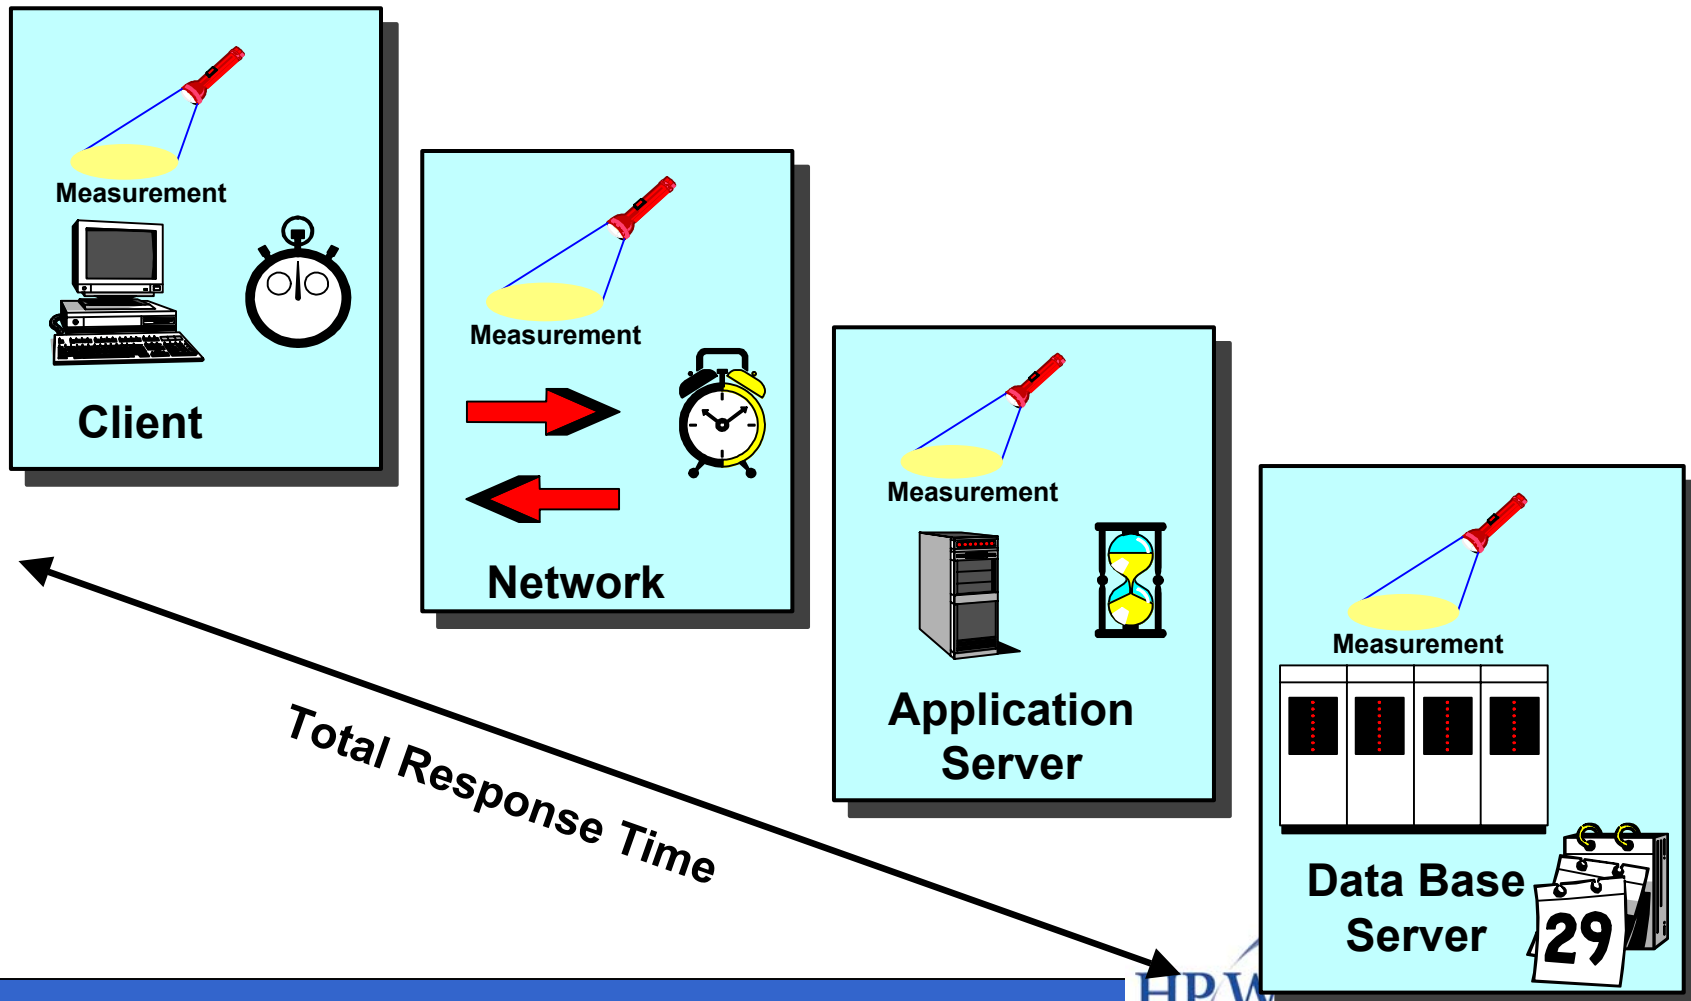
Data from one business transaction is typically fully contained

## Client Server Distributed Environment



Data from One Business Transaction is not fully contained

# Client Server – Response Time Component



# Global Summary

```
SOS B.01y eagle WED, 30 JAN 2002, 08:45 E: 00:22:00 I: 00:00
----- CPU UTILIZATION ----- CPU MISC -----
TOTAL BUSY: 61.3[ 1] HIGH PRI: 61.3[ 1] | Capture Ratio 0.8[ <]
User 0[ <] Sys 27.9[ <] Mem 0[ <] | RunQ Avg 0[ 0]
Real 0[ 0] Intr 6.3[ <] Idle 38.7[99] | 5/15 Min RunQ Avg 0/ 0
Nice 0[ 0] C SW 0[ <] | RunQ Busy % 0[ <]
NNice 27.0[ <] Trap 0[ <] |
----- MEM/VM -----
Read Hit % 100.0[100] Page Outs 0[ <]/s Mem Used % 67.1[ 67]
Write Hit % 100.0[ 55] Deact Byte 0[ 0]/s VM Used % 25.9[ 25]
----- MISC -----
#Sessions: 1 #Procs: 52 #Wait I/O: 0 Transactions: 33.0[ 9.2]/s
#Active: 0 #Active: 2 #Deact: 0 Avg Response Time: <
----- DISK -----
Disk IO/s IO% QLen | Disk IO/s IO% QLen | Disk IO/s IO% QLen
c0t5d0 0 100 0 | c0t2d0 0 0 0 | c0t6d0 0 0 0
----- PROCESS SUMMARY -----
PID Name User Name TTY CPU% Nice Pri RSS/Size #Rd #Wr Wait Resp
The CPU was used a total of 61.3 of its capacity during this interval <CI01>
This interval's 'hog' process is (PID 3644) with 68.1% of the CPU <PI01>
This interval's highest disk I/O user was (PID 3644) with 1 I/O's <PI02>
FLOCKS too high at 4 <GE03>

Enter command: (Showing lines 4 - 8 of 8)
UPDATE RESET SCREEN HELP MORE FREEZE PRINT EXIT
```

# System Tables

```

SOS B.01y eagle WED, 30 JAN 2002, 08:50 E: 00:26:56 I: 00:59
----- MISC TABLE/CACHE SUMMARY -----
                Entries      Used      Used %      High
File Table      920        197        21.4        202
File Lock Table 200          4          2.0         4
Pseudo TTY Table 60           1          1.7         1
Process Table   276         54         19.6        55
Inode Cache     476        322        67.6        476
----- IPC TABLE/CACHE SUMMARY -----
                Size      Entries      Used      Used %      High
Message Table   N/A         50          2          4.0         2
Message Buffer Cache 800k      N/A          0           0           0
Semaphore Table N/A         70          10         14.3        10
Shared Memory Table N/A        200         5           2.5         5
Shared Memory   200g      N/A        326         0          326
----- DNLC CACHE SUMMARY -----
Entries      Hit %      Lookups/s
1500         94.5      4.9
----- SYSV BUFFER CACHE SUMMARY -----
                Min      Max      DATA CACHE
                Size      Size      Size      Used      Used %      High      Hit %
Header Table   Entries Used  Used % |
27104 27104 100.0 | 26212k 256m 207m 207m 100.0 207m 72.7
-----
Enter command:
[UPDATE] [RESET] [SCREEN] [HELP] [MORE] [FREEZE] [PRINT] [EXIT]

```

# System Resources

- CPU
- Memory
- Disk I/O
- Network

# Central Processing Unit

- Metrics
- Process Scheduling
- Metric Thresholds
- Analysis Desktops
- Trending



# CPU – Metrics

General Metrics	User Mode	System Mode
CPU Busy %	CPU User %	CPU Sys %
Run Queue Length	CPU Nice %	CPU CSw %
CPU Busy % by Workload	CPU Neg Nice %	CPU Intr %
(CPU High-Pri Busy %)	CPU Real %	CPU Trap %
		CPU Vflt %

# CPU – Summary Screen

```

SOS B.01y eagle WED, 30 JAN 2002, 08:46 E: 00:23:22 I: 00:00
----- CPU SUMMARY -----
USER REAL NICE NNICE SYS INTR C SW TRAP MEM IDLE TOTAL HIGH PRI
          0 0 0 16.0 23.7 5.3 0.8 0 0 54.2 45.8 45.8
----- RUNQ STATISTICS -----
Interval Avg: 0[ <] 1/5/15 Min Avg: 0.1/ </ < Occ %: 0[ 1]
----- MISC STATISTICS -----
Forks: 0[ <]/s C Sws: <[ 11]/s
Intrs: <[ 333]/s Traps: <[ <]/s Sys C: <[ 121]/s
----- PER CPU UTILIZATION -----
TOTAL
CPU BUSY USER REAL NICE NNICE SYS INTR OHD IDLE 1/5/15 Min Last
          RunQ Avgs PID
-----
1 45.8 0 0 0 16.0 23.7 5.3 0.8 54.2 0.1/ </ < 3644
   [ 2] [ <] [ 0] [ 0] [ 1] [ 1] [ <] [ <] [98]

Enter command:
[ UPDATE ] [ RESET ] [ SCREEN ] [ HELP ] [ MORE ] [ FREEZE ] [ PRINT ] [ EXIT ]

```

# CPU – Process Scheduling

## Real Time Scheduling

System

| 0-----127 |

Highest

## Time Share Scheduling

System

| 128-----177 |

User

| 178-----255 |

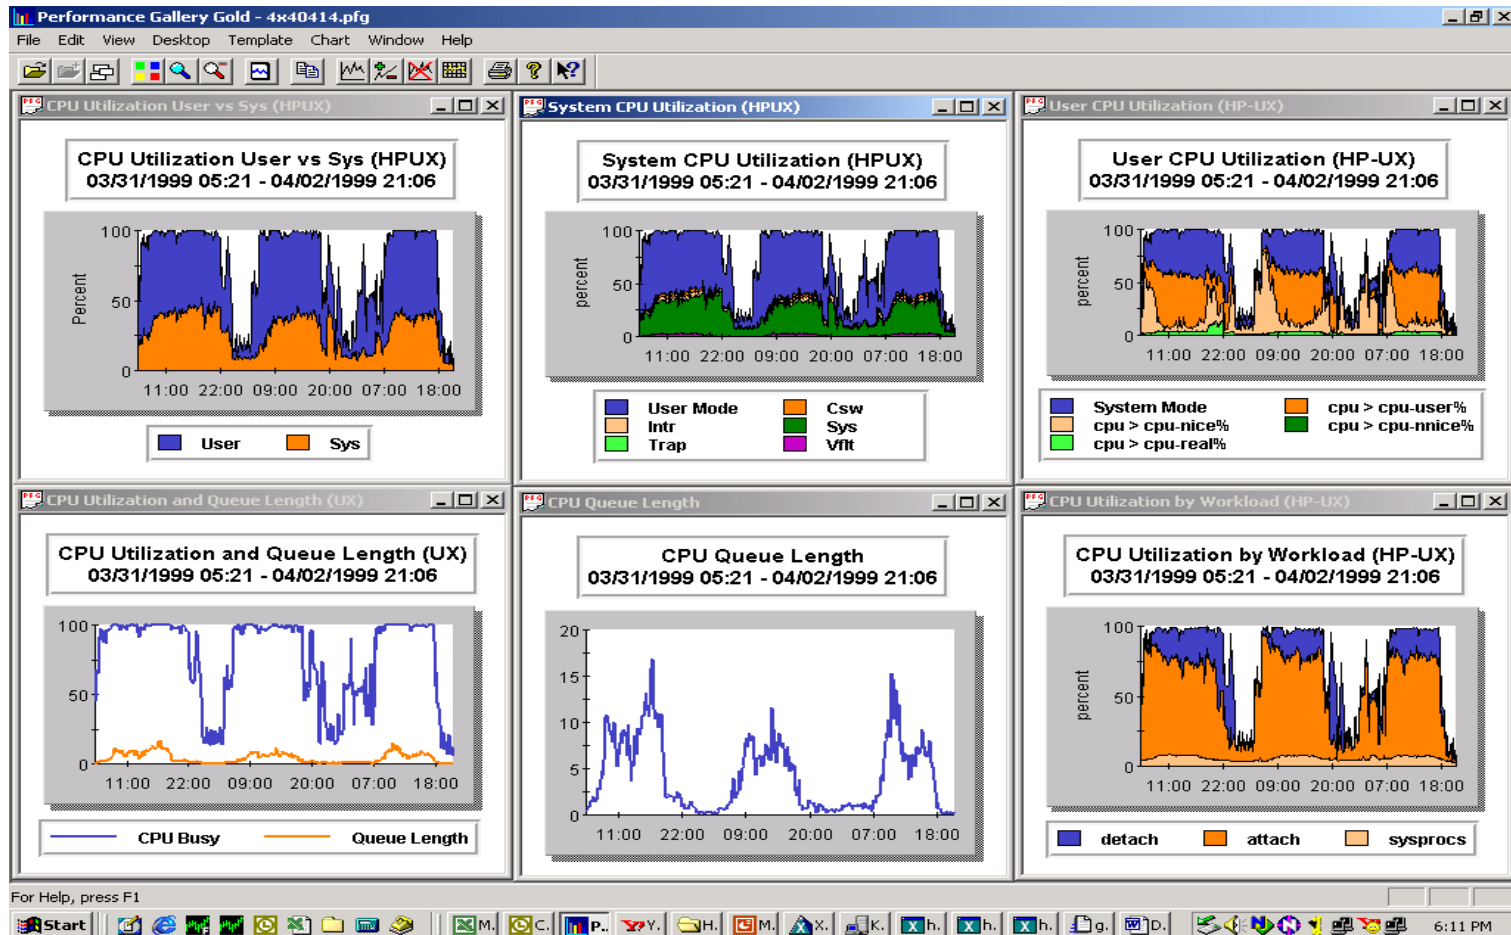
Lowest

- Nice lowers the priority, but increases the Unix priority number
- Negative nice raises the priority, but decreases the Unix priority number


# CPU – Thresholds

CPU Busy.....	85%
CPU Real Time.....	5%
System (Kernel) Mode.....	30%
Run Queue Size.....	10

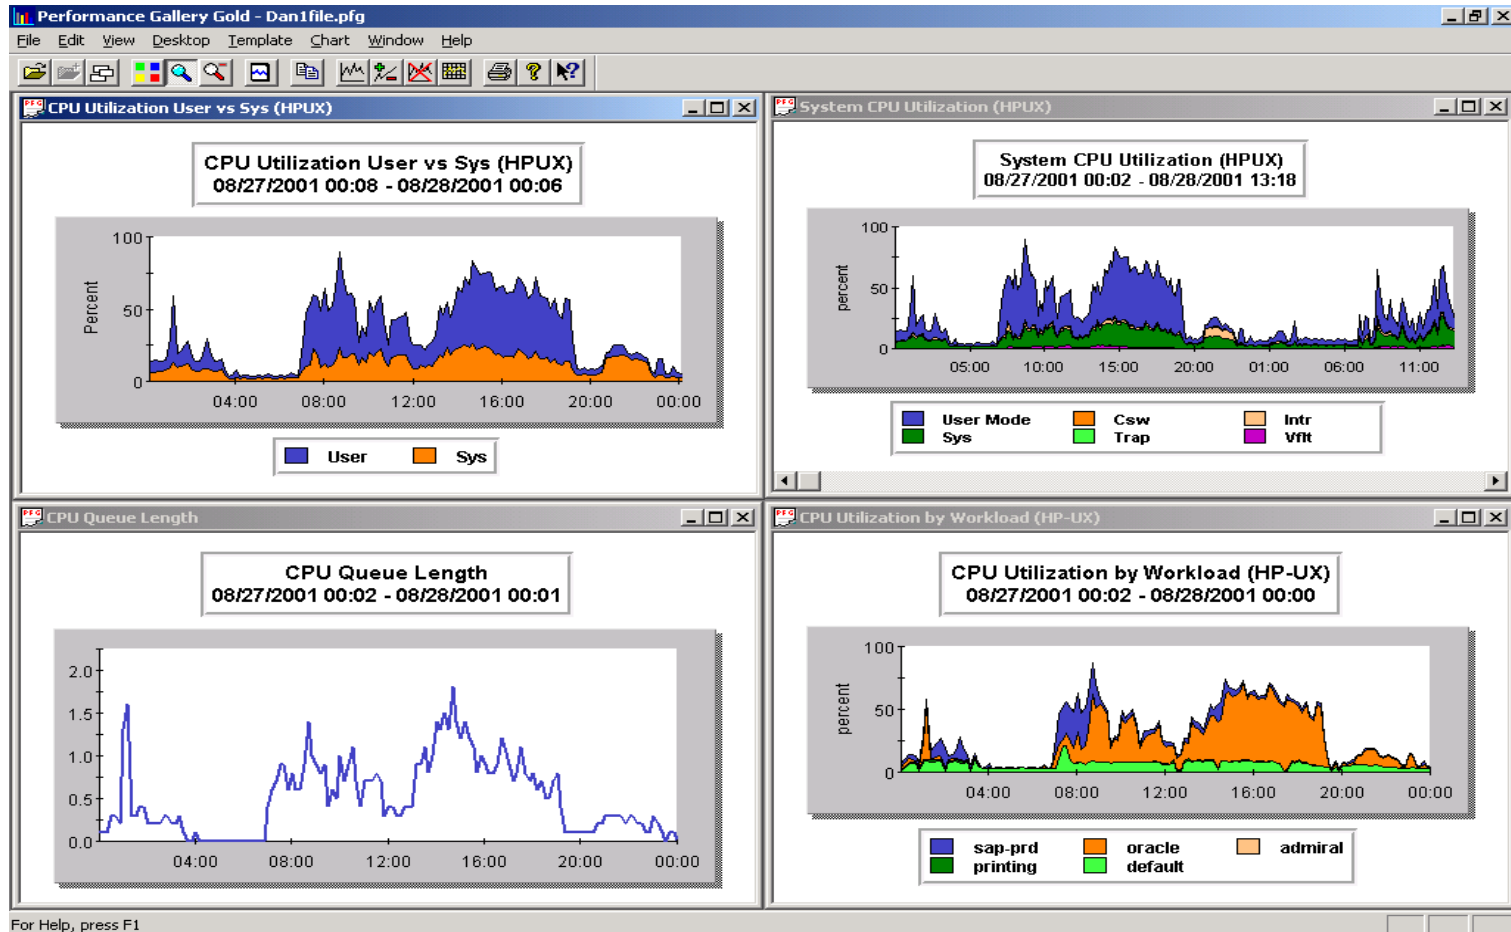
# CPU – Analysis Views



# CPU – Trending

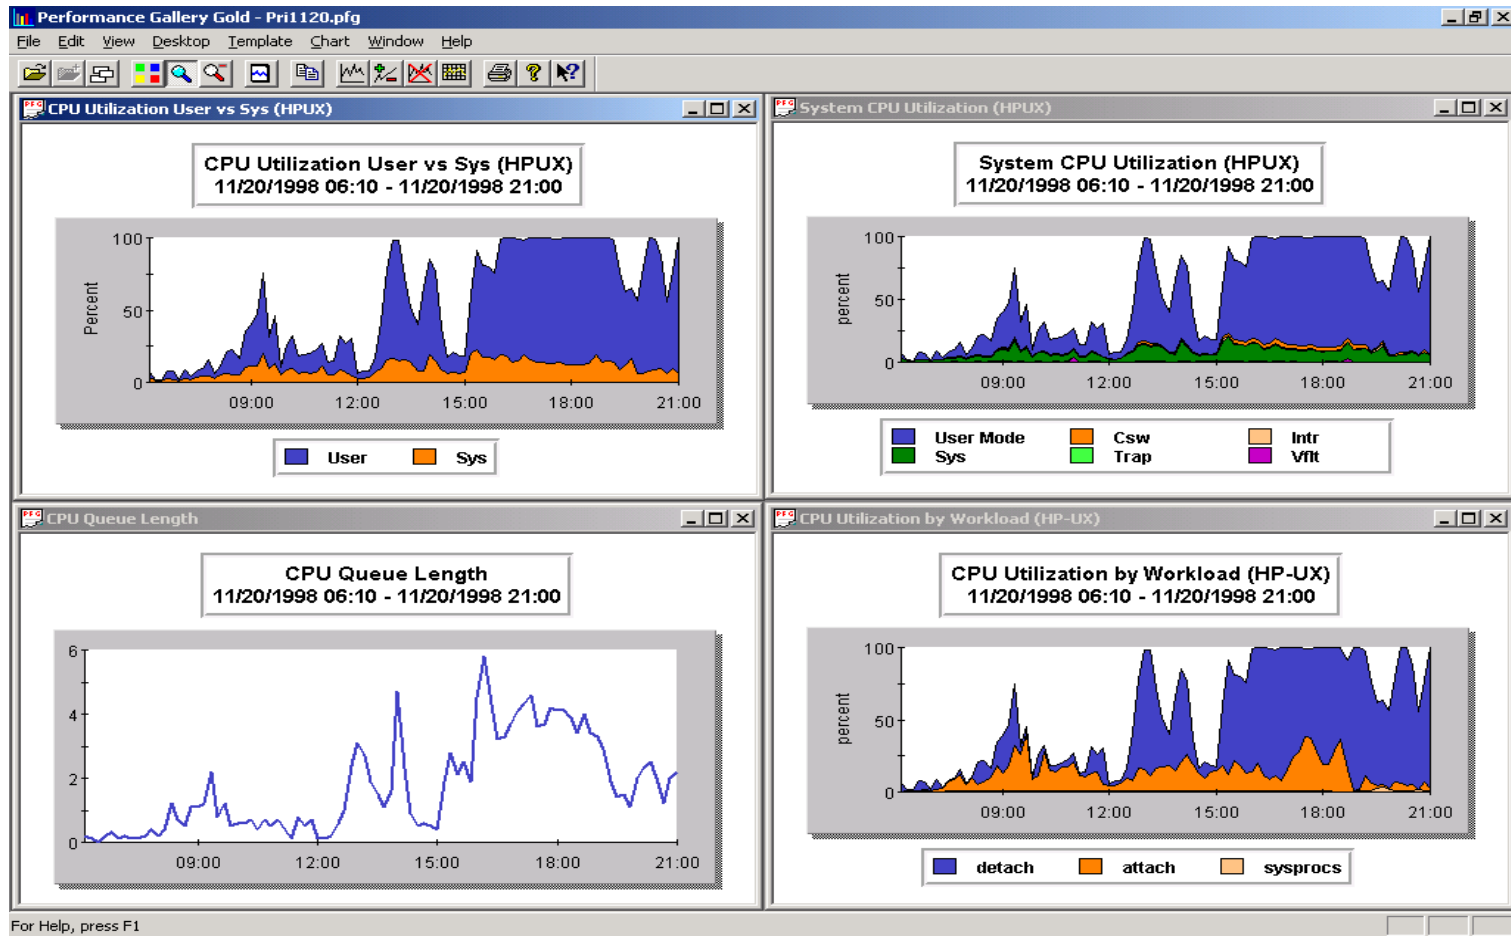


# CPU – Minimal Queueing



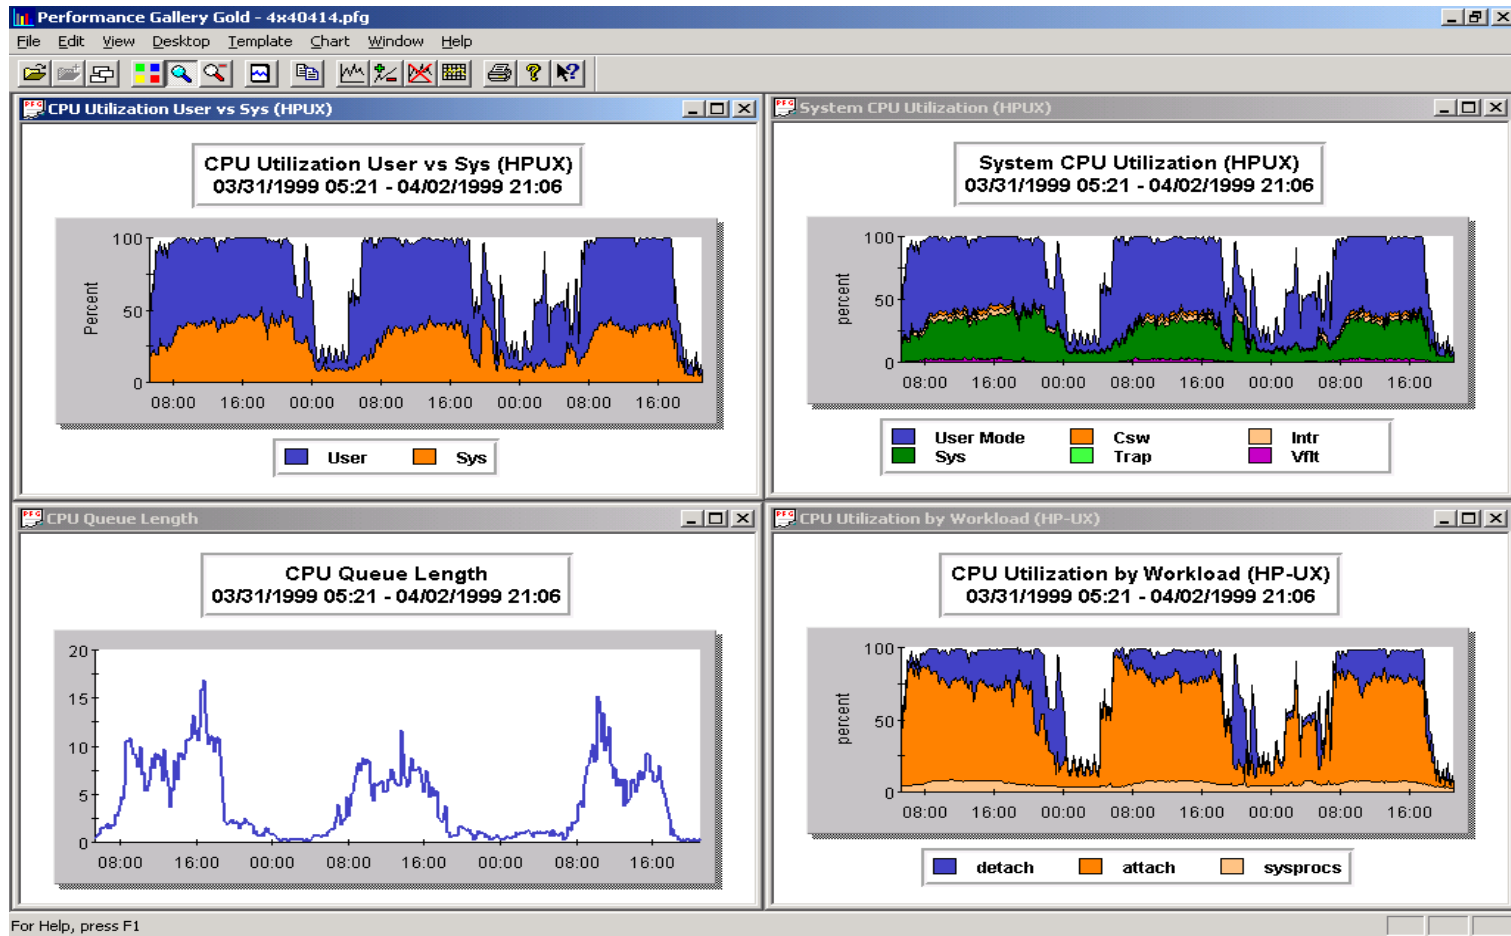
For Help, press F1

# CPU – Moderate Queueing





# CPU – Excessive Queueing



# Memory

- Metrics
- Memory Management
- Metric Thresholds
- Analysis Desktops
- Trending

# Memory – Metrics

Memory Utilization	Virtual Memory I/O
Used Memory %	Page In Rate
System Memory	Page Out Rate
Buffer Cache	Deactivation Rate
User Memory	
Free Memory	

# Memory – Summary Screen

```

SOS B.01y eagle WED, 30 JAN 2002, 08:48 E: 00:24:53 I: 00:01
----- MEM/VM ALLOCATION ----- PROC MEM STATUS -----
      Size      User      Sys      Buffer      Free      |      Run      Sleep      Total
Mem    512m     131m     35m     177m     167m     |      0         55         55
VM    1146m     301m     ---      ---      845m     |      Deact         0         0
----- PAGING -----
      In(/s)      Out(/s)      In(byte/s)      Out(byte/s)      #In      #Out
Pg Flts      0[ 2]
Pages        0[ <]      0[ <]      0[ 30]      0[ 16]      0[ 4]      0[ 6]
(De)act      0[ 0]      0[ 0]      0[ 0]      0[ 0]      0[ 0]      0[ 0]
VM I/O       0[ <]      0[ <]      0[ 30]      0[ 16]      0[ 4]      0[ 6]
Forks        0[ <]
              0[ 5272]
              0[ 34]
----- PAGE SCANNER -----
Page Recs      0[ 1]/s
Page Scans      0[ 0]/s
----- SYSV BUFFER CACHE -----
Read Cache Hit % 100.0[100]
headers 22834 size 177m DBC min size 26212k DBC max size 256m
Write Cache Hit % 98.2[ 99]
----- MEMORY MANAGEMENT CONFIG -----
lotsfree: 27m desfree: 4096k minfree: 1024k
umem: 27m page size: 4096
-----
Enter command:
[UPDATE] [RESET] [SCREEN] [HELP] [MORE] [FREEZE] [PRINT] [EXIT]

```


# Memory – Management

- There are 3 kernel parameters that affect paging and deactivations
  - LOTSFREE (Lots Free)
  - DESFREE (Desired Free)
  - MINFREE (Minimum Free)
- Memory Manager tries to keep free memory above DESFREE

# Memory – Management

- Paging begins when free memory drops between LOTSFREE and DESFREE
- Deactivations occur (paging stops) while free memory is below MINFREE
- Paging continues until free memory is above LOTSFREE

# Memory – Management

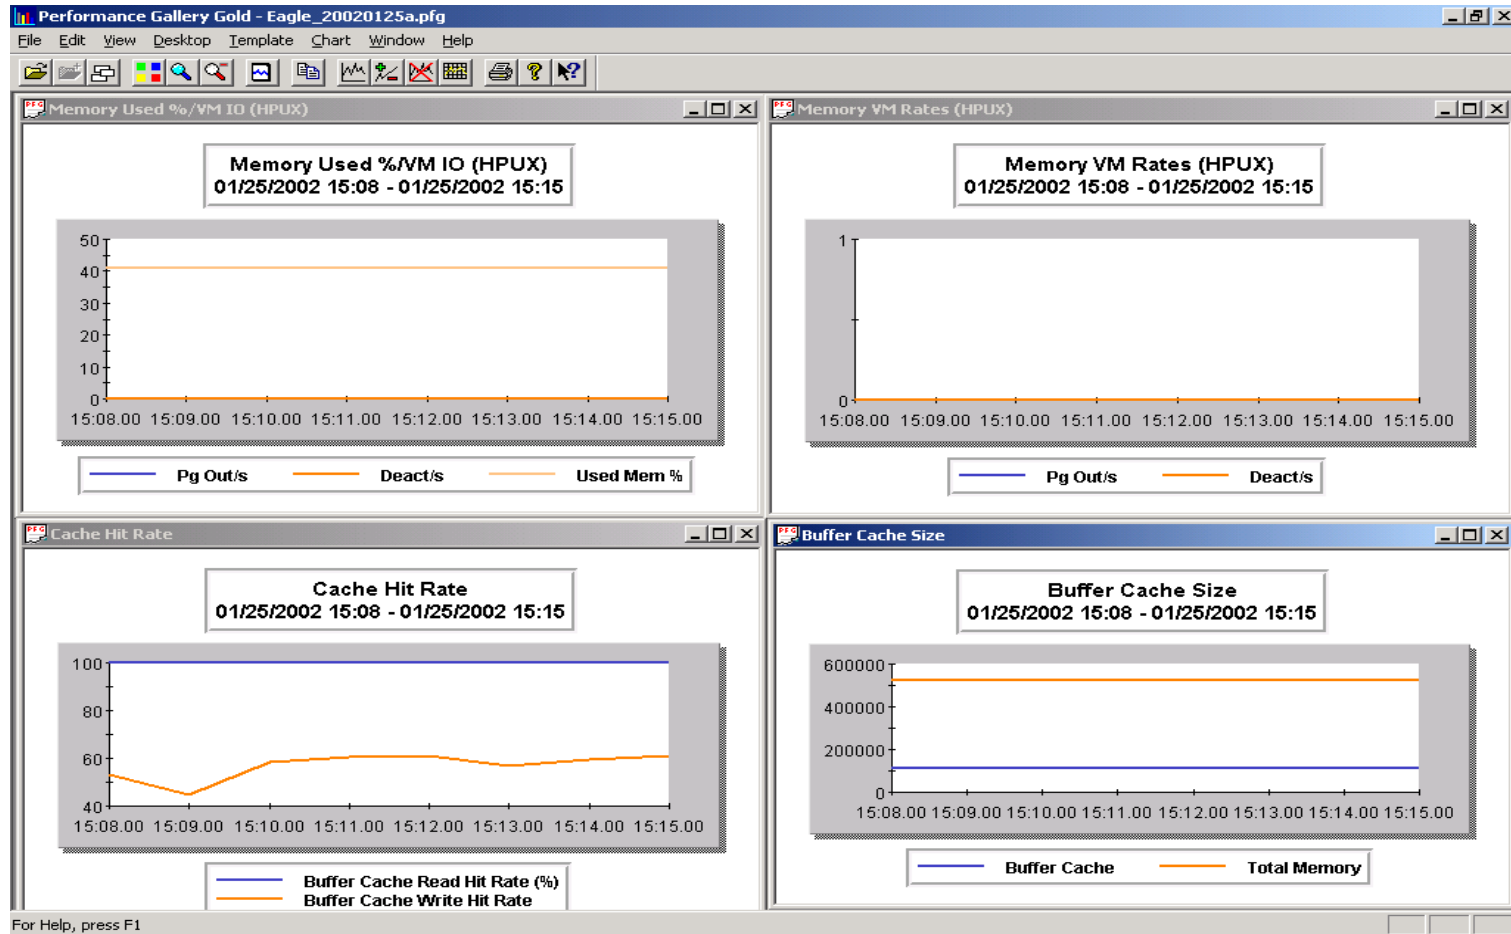


# Memory – Thresholds


Used Memory %.....	95%
Free Memory %.....	5%
Deactivation Rate.....	0



# Memory – Analysis Views



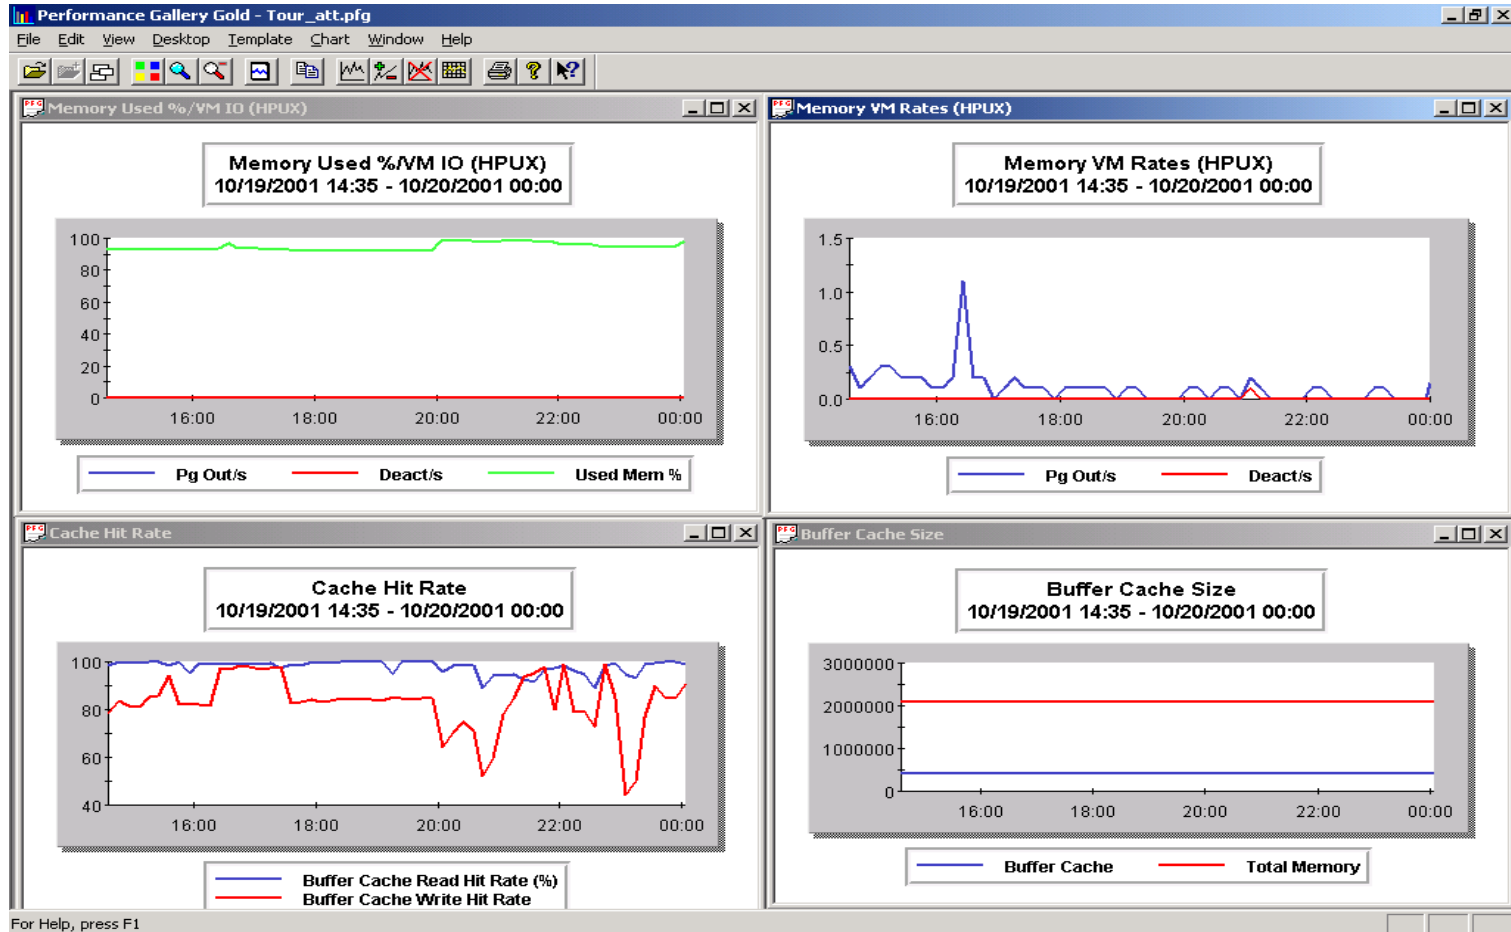
# Memory – Trending



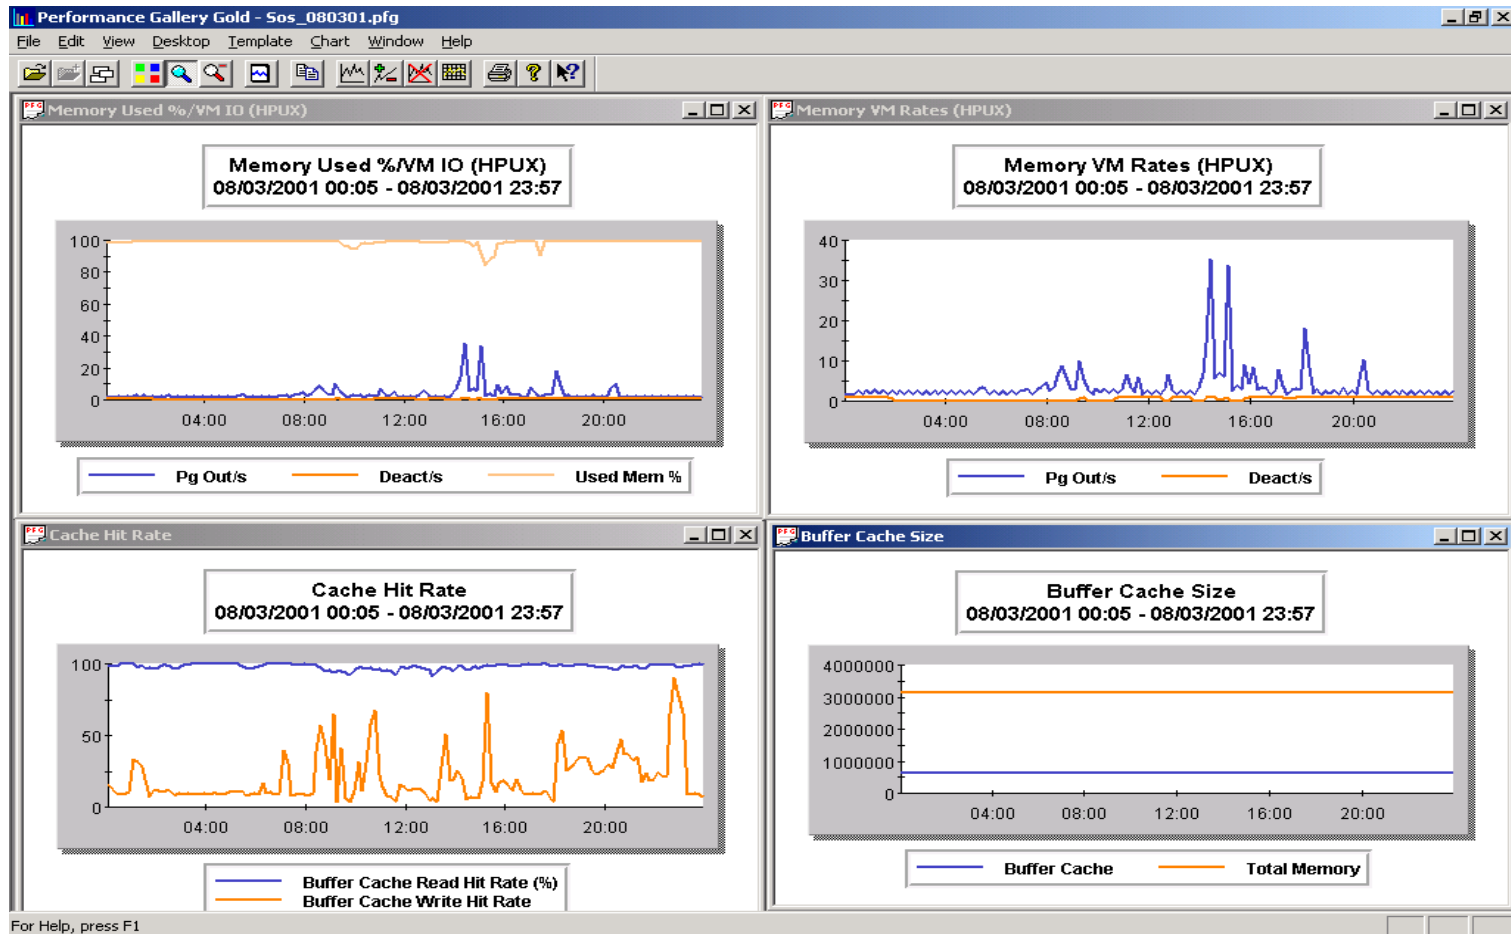
# Memory – No Pressure



# Memory – Minimal Pressure



# Memory – Excessive Pressure



# Disk

- Metrics
- Buffer Cache
- Metric Thresholds
- Analysis Desktops
- Threshold Filtering
- Trending

# Disk – Metrics

General	I/O Rates
Queue Length	Logical I/O (FS)
Wait Times	Physical I/O
Service Times	User I/O (FS)
Utilization %	Raw I/O
Buffer Cache Hit Rate %	VM I/O
	System I/O

# Disk –

## Metrics Global I/O Detail

```

SOS B.01y eagle WED, 30 JAN 2002, 08:49 E: 00:25:56 I: 00:02
----- GLOBAL I/O DETAIL -----

```

		Reads		Writes		Rate (/s)		Avg Size (kb)	
						Read	Write	Read	Write
Local	Logical	28	0	0	0	14.0	0	<k	0
		[10258]	[ 90]	[ 0]	[ 0]	[6.6]	[0.1]	[ <k]	[ <k]
	Physical	0	9	0	0	0	4.5	0	3
		[ 7]	[ 807]	[ 0]	[ 0]	[ 0]	[0.5]	[ 8]	[ 4]
	User FS	0	0	0	0	0	0	0	0
		[ 3]	[ 10]	[ 0]	[ 0]	[ 0]	[ 0]	[ 5]	[ 4]
	Sys FS	0	9	0	0	0	4.5	0	3
		[ 0]	[ 791]	[ 0]	[ 0]	[ 0]	[0.5]	[ 0]	[ 4]
	Virt Mem	0	0	0	0	0	0	0	0
		[ 4]	[ 6]	[ 0]	[ 0]	[ 0]	[ 0]	[ 11]	[ 4]
	Raw	0	0	0	0	0	0	0	0
		[ 0]	[ 0]	[ 0]	[ 0]	[ 0]	[ 0]	[ 0]	[ 0]
Remote	Logical	0	0	0	0	0	0	0	0
		[ 7251]	[ 144k]	[ 0]	[ 0]	[4.7]	[ 93]	[ 9]	[ <k]
	Physical	1	192	0	0	0.5	96.0	1	<
		[ 351]	[ 9067]	[ 0]	[ 0]	[0.2]	[5.8]	[ <k]	[ <]

Enter command:

UPDATE	RESET	SCREEN	HELP	MORE	FREEZE	PRINT	EXIT
--------	-------	--------	------	------	--------	-------	------



# Disk – Summary Screen

```
SOS B.01y eagle WED, 30 JAN 2002, 08:49 E: 00:25:56 I: 00:02
```

----- DISK I/O SUMMARY -----

Dev	I/O%	Qlen	Util%	Wait	Service	Rates (/s)		Avg Size(kb)	
				Time(ms)	Time(ms)	Read	Write	Read	Write
c0t6d0	0	0	0	0	0	0	0	0	0
c0t2d0	0	0	0	0	0	0	0	0	0
c0t5d0	100	0	3.9	0	8.6	0	4.5	0	3
TOTALS	100	0	1.3	0	8.6	0	4.5	0	3

Enter command:

UPDATE	RESET	SCREEN	HELP	MORE	FREEZE	PRINT	EXIT
--------	-------	--------	------	------	--------	-------	------

# Disk – Buffer Cache

- A pool of buffers in memory to maintain data in memory to avoid disk access
- Fixed vs. Dynamic (default)
- Default Kernel Parameters
  - `bufpages = 0`
  - `dbc_min_pct = 5`
  - `dbc_max_pct = 50`

# Disk –

## Buffer Cache Physical Disk I/O

Physical Disk Read =  
User (FS) Read +  
Raw Read +  
VM Read +  
System Read

Physical Disk Write =  
User (FS) Write +  
Raw Write +  
VM Write +  
System Write

**Disk –**

**Buffer Cache  
Hit Rate**

Buffer Cache Hit Rate %

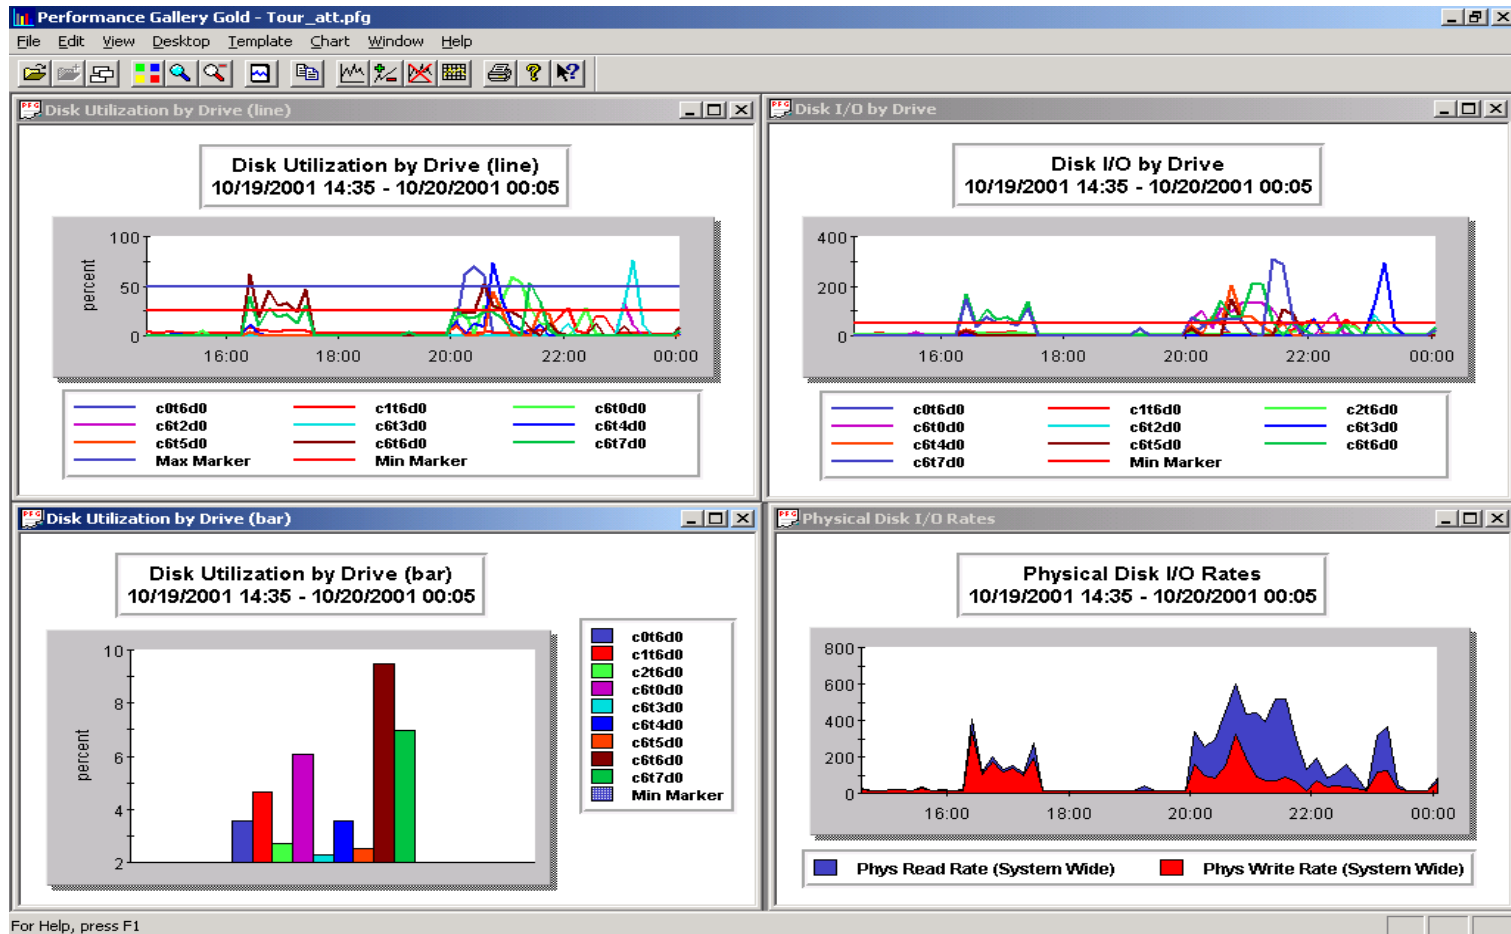
$$= \frac{\text{Logical Reads} - \text{Physical Reads}}{\text{Logical Reads}} * 100$$

$$= \frac{\text{Logical Rd Rate} - \text{User Rd Rate}}{\text{Logical Rd Rate}} * 100$$

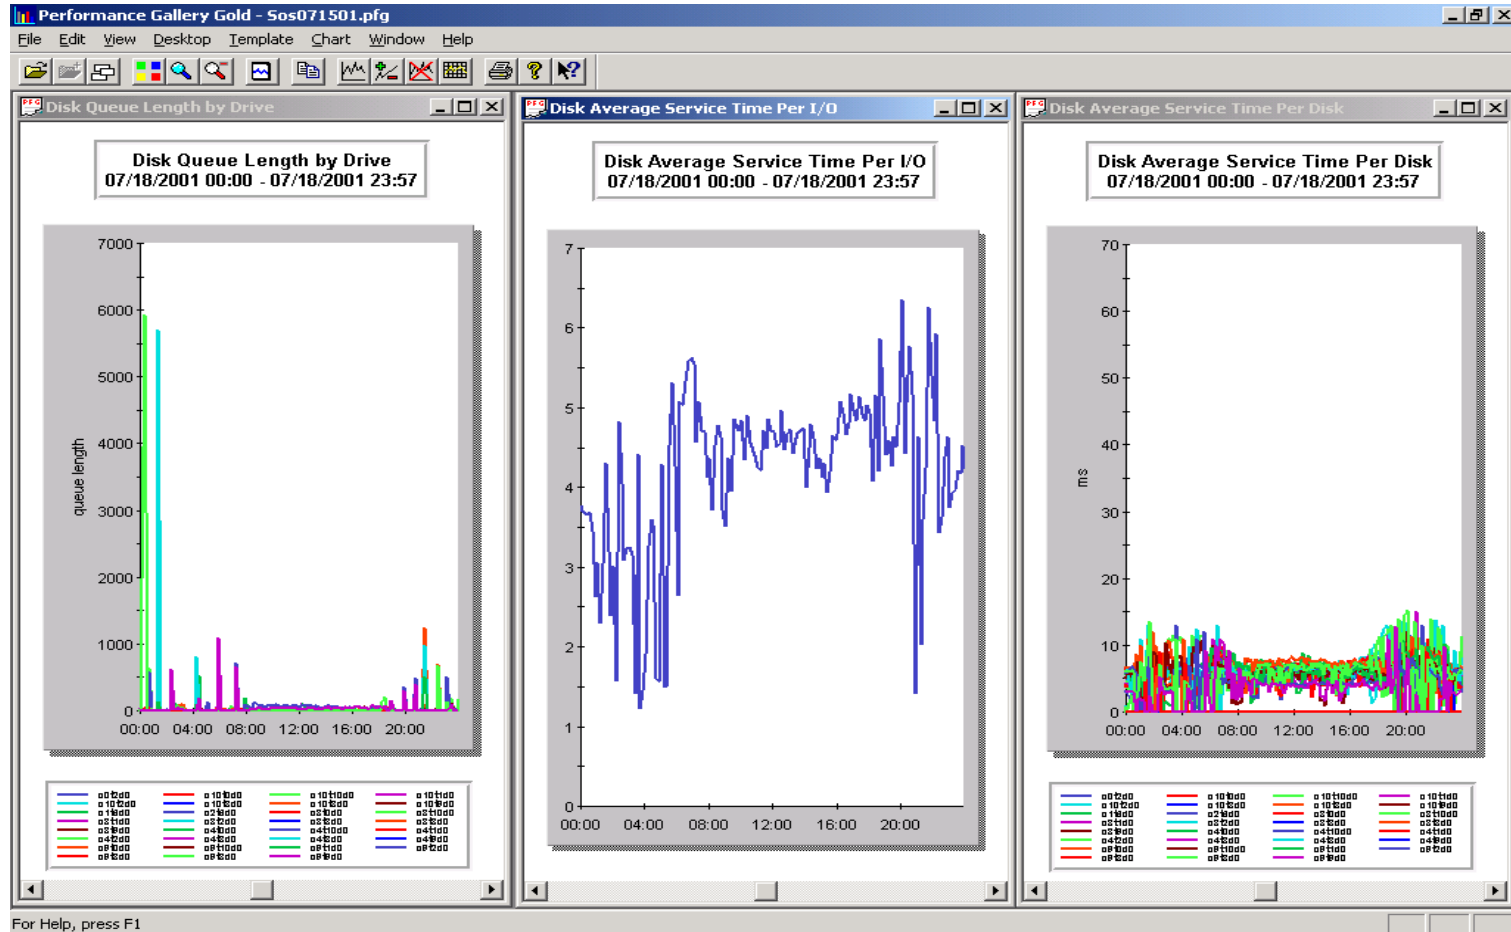
# Disk – Thresholds

Disk Utilization %.....	25-50%
Buffer Cache Hit Rate %.....	90-95%
Disk Queue Length.....	3

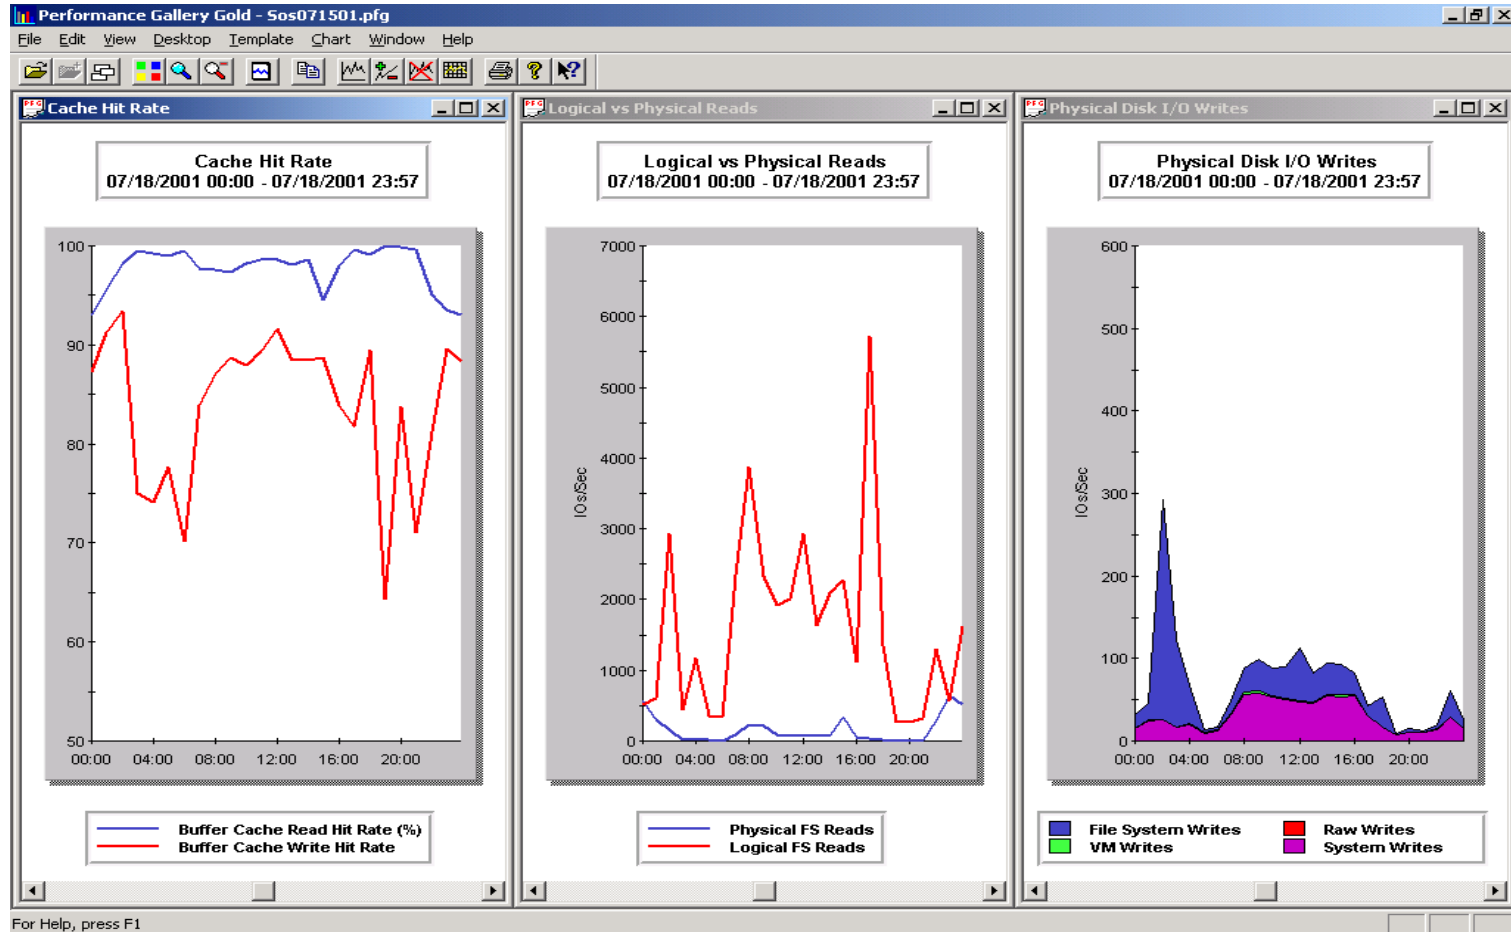
# Disk – Analysis Views



# Disk – Analysis Views




# Disk – Analysis Views






# Disk –

## Threshold Filtering Before Filter



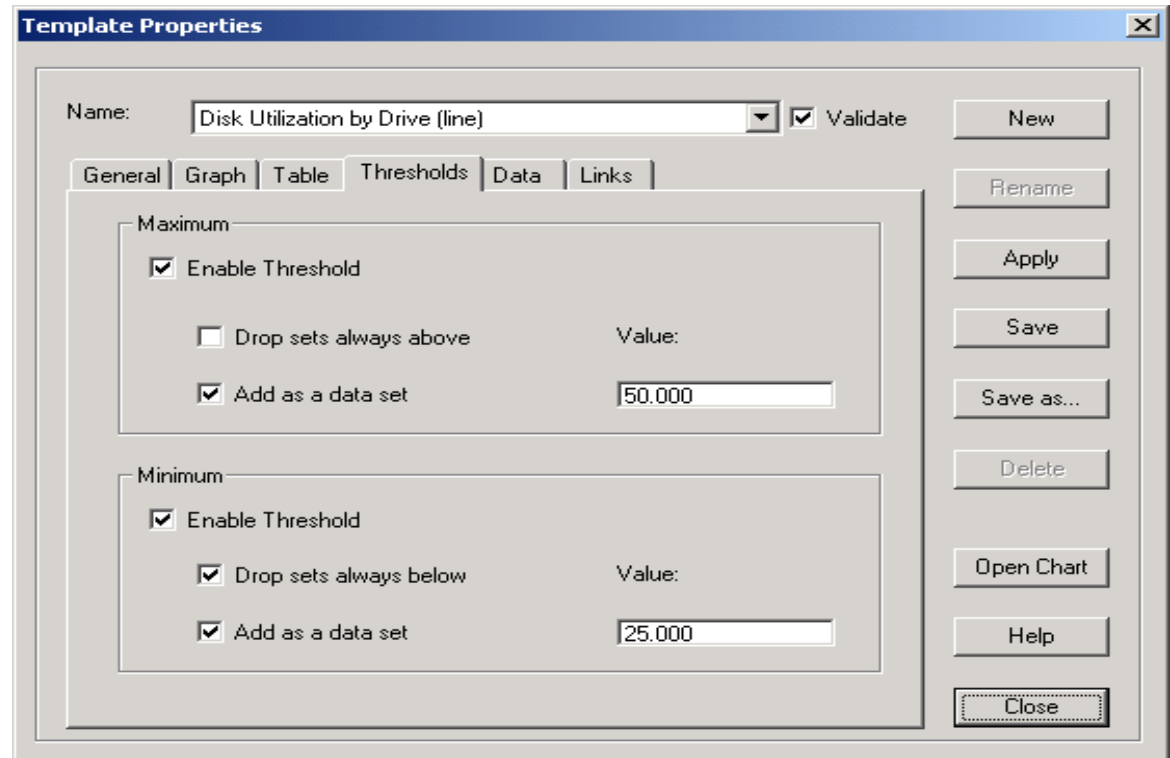
# Disk –

## Threshold Filtering After Filter




# Disk – Threshold Filtering


Location In Lund's Performance Gallery Gold to filter data by threshold settings



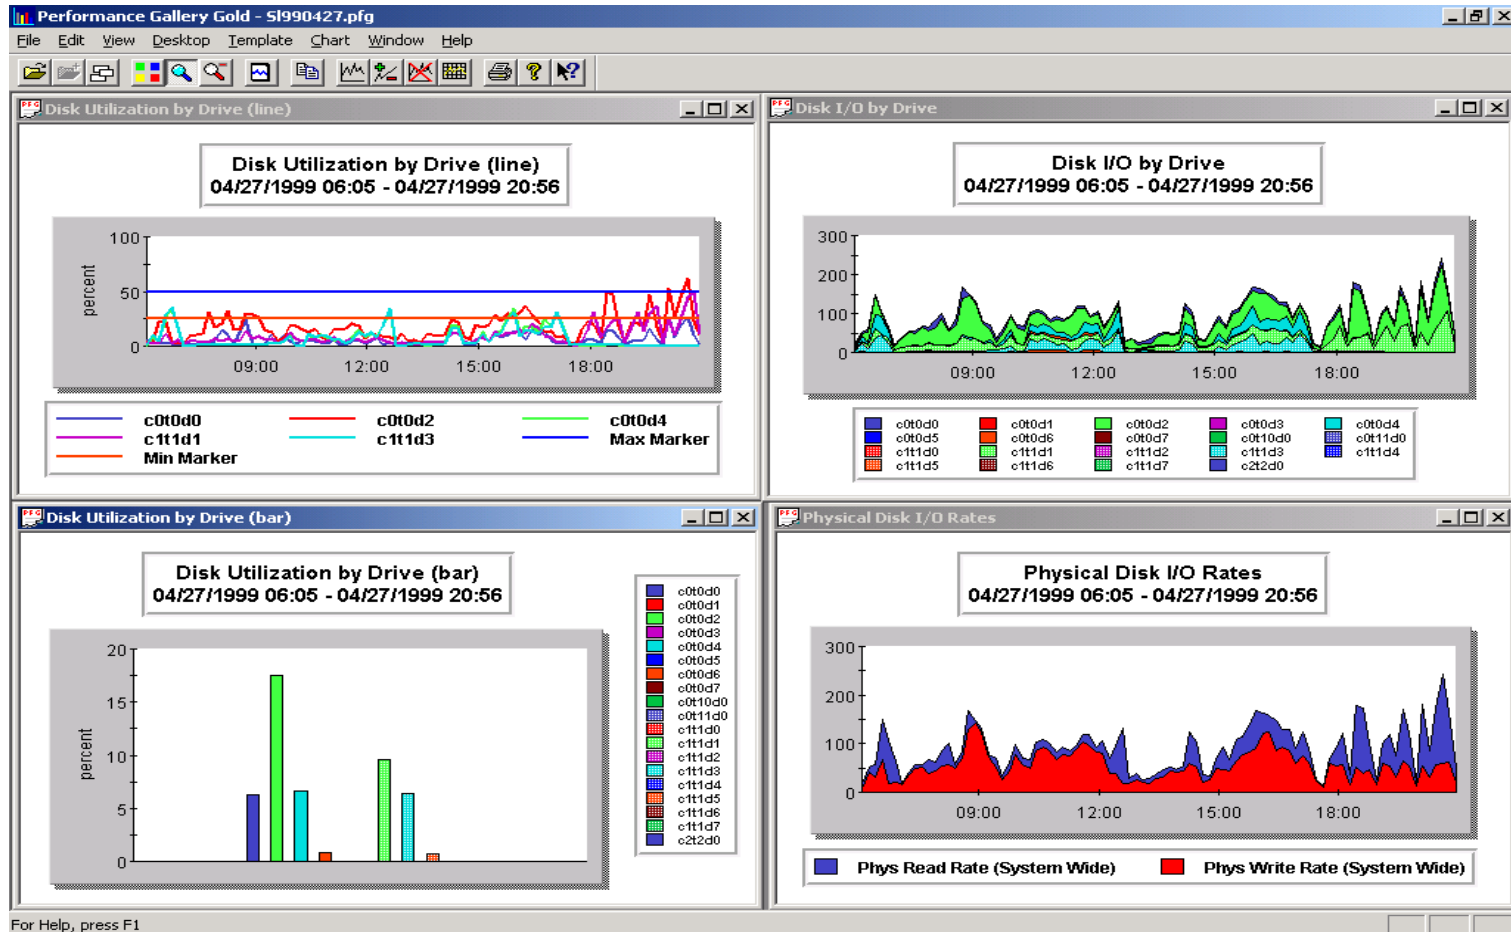
# Disk – Trending




# Disk – File System I/O



# Disk – Reasonable Utilization



# Disk – Excessive Utilization



# Network

- Metrics
- Metric Thresholds
- Analysis Desktops



# Network – Metrics

- TCP packets in & out
- UDP packets in & out
- Collisions
- Errors

# Network – Summary Screen

```
SOS B.01y eagle WED, 30 JAN 2002, 08:50 E: 00:26:56 I: 00:59
```

----- NETWORK SUMMARY -----

Protocol	Packets In/s	Packets Out/s	Errors In%	Errors Out%
IP	6.0[ 13]	15.5[ 32]	0[ 0]	0[ 0]
TCP/IP	5.5[ 7]	14.9[ 25]	0[ 0]	0.91[ 0.65]
ICMP	<[ <]	0[ 0]	0[ 0]	0[ 0]
UDP	0.5[ 6]	0.5[ 6]	0[ 0]	N/A[ N/A]

----- NETWORK INTERFACES -----

Interface	Packets In/s	Packets Out/s	Collision%	Error In%	Error Out%
lan0	6.0[ 13]	15.5[ 60]	0[ 0.68]	0[ 0]	0[ 0]
lo0	0[ 0]	0[ 0]	0[ 0]	0[ 0]	0[ 0]
TOTALS	6.0[ 13]	15.5[ 60]	0[ 0.68]	0[ 0]	0[ 0]

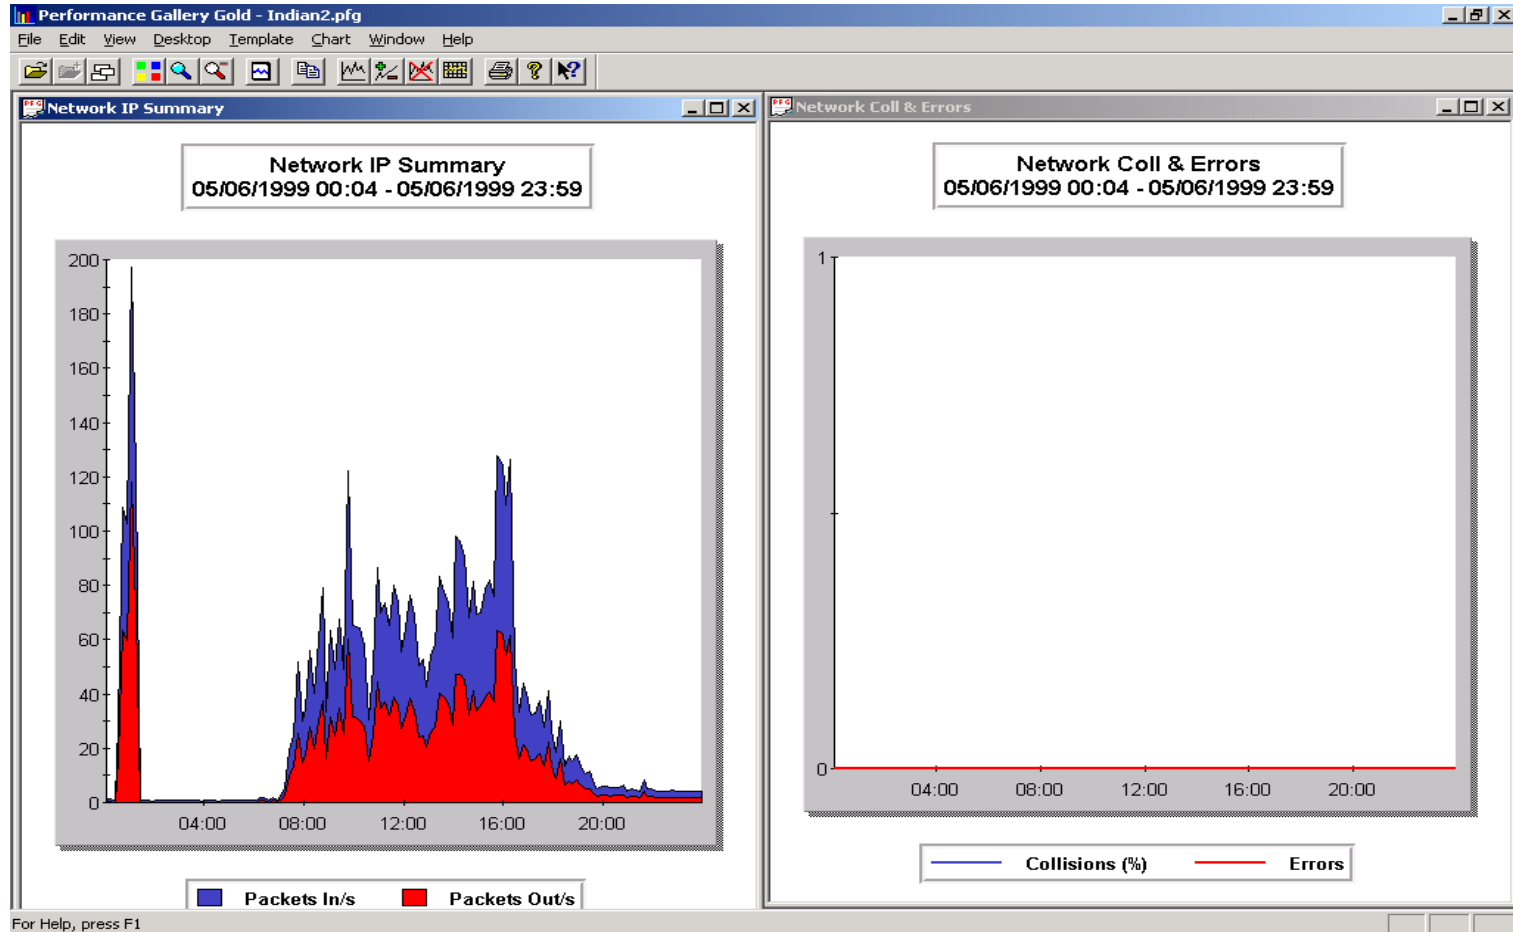
Enter command:

UPDATE	RESET	SCREEN	HELP	MORE	FREEZE	PRINT	EXIT
--------	-------	--------	------	------	--------	-------	------


# Network – Thresholds

Collisions %.....	10%
Errors .....	0

# Network – Analysis Views



# Network – Trending



# UNIX Performance (in review)

- Perform analysis on a regular schedule
- You must know what *normal* is, in order to recognize a problem
- Resolve each pressure symptom as it is identified to maintain good performance
- Develop a standard method of reporting

# Questions and Answers