

Monitoring VM for Performance and Control



Laura J. Knapp
IBM Technical Evangelist
1-919-224-2205
laura@lauraknapp.com
www.lauraknapp.com

Baseline Your Network

Gather device/protocol/application software versions

Gather device/protocol/application configuration

Gather device/protocol/application statistics

Have a logical view of your network

Understand the protocols running

Be informed of recent changes in devices/protocols/applications

Be aware of future plans that are currently in test



Causes of Connectivity Problems

Faulty hardware or media

Bugs

Backhoes cutting fiber

Power outages

Not enough resources

Hardware inoperable

Protocol mismatches

Handshake/bind problems

Blocked resources

Down resources



Causes of Performance Problems

Network congestion

Less desirable route

Underpowered network devices

Network faults such as spanning tree loops

Network noise

Network errors

Insufficient resources

Resource configuration

Over-subscription



Increasing Importance of Performance

Performance Management

The practice of managing network service response time, consistency and quality for individual services and services overall

Performance Related Risks

- Network degradation and failure
- Application timeouts and failure
- Application degradation

Loss of Customers



The Performance Problem

Over-provisioning

- Lots of provisions (rare)
- More resources than can be consumed
 - Food on a cruise
 - Congressional parking spaces
 - AOL CD-ROMs

Over-subscribing - lots of subscribers

- Lots of subscribers (common)
- Many users consume all the resources
 - Batteries, chain saws, interstate lanes during a hurricane
 - Phone calls on Mothers' Day
 - Many to few: whenever there's a bottleneck or funnel
 - Fast to slow: things will back up



Agenda

Introduction and background

Performance Methodologies

Areas of Concern

Greed -- Am I miss-using buffer space?

Sloth - What are response times?

Envy - What's going on in the IP Stack?

Lust - What is system availability?

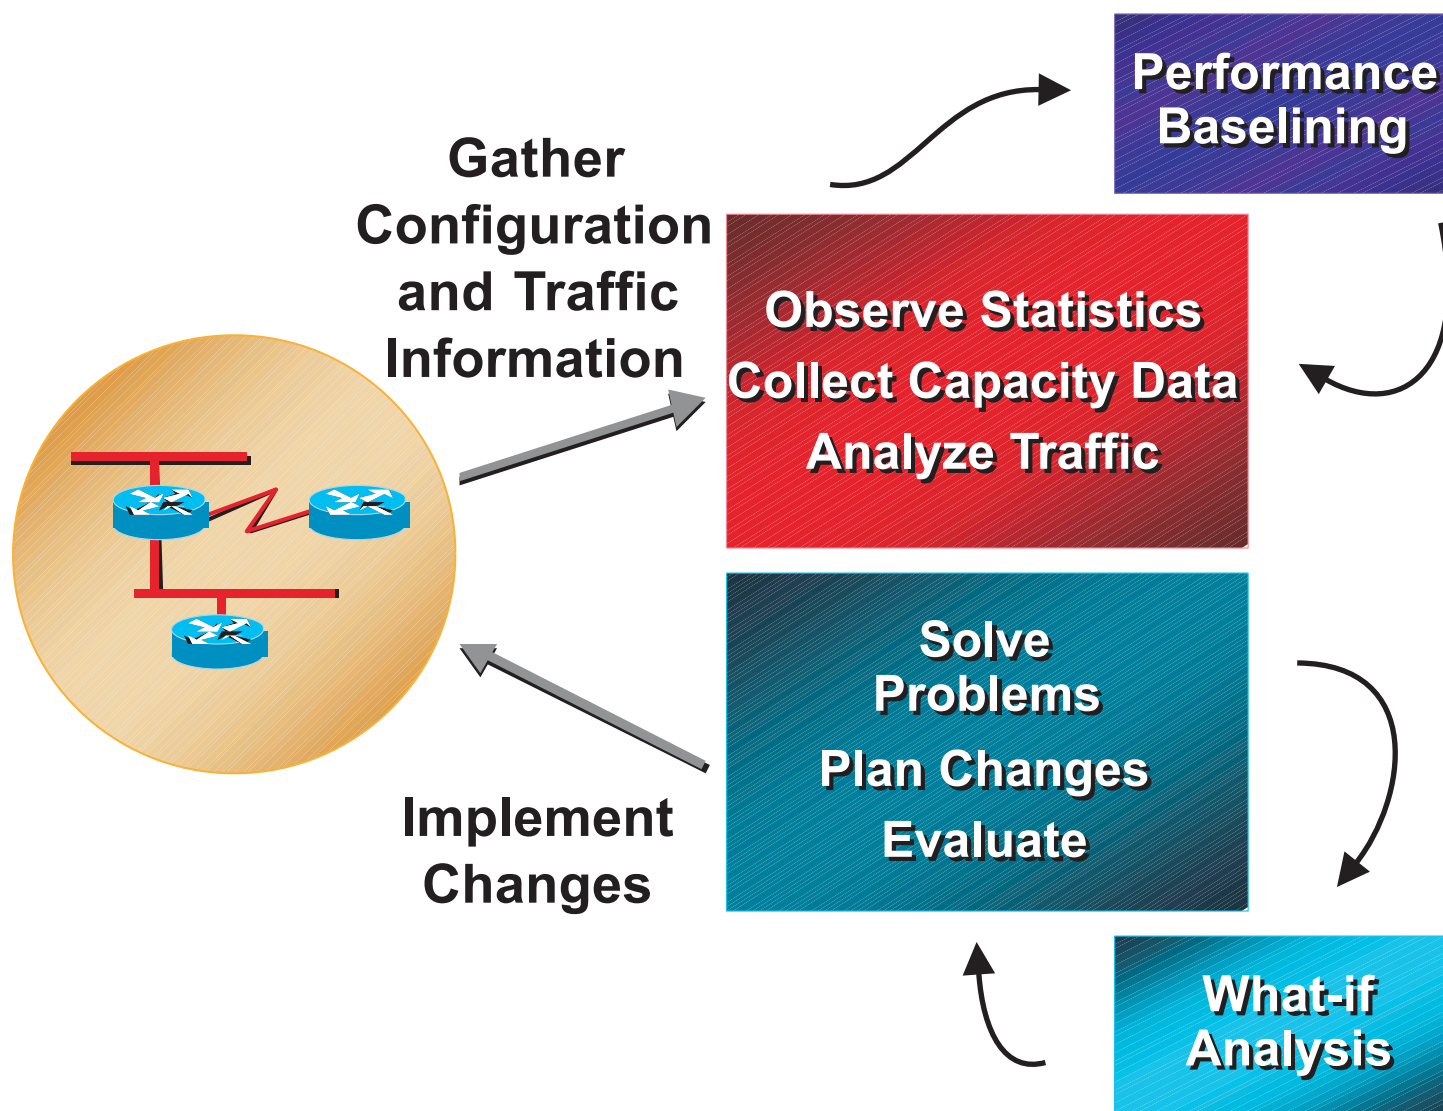
Anger - Who is using my resources?

Gluttony - Who is hogging resources?

Pride - What connections are there available?



Effective Performance Management



IP Resource Bottlenecks

CPU

Memory

Buffering, queuing, and latency

Interface and pipe sizes

Network Capacity

Speed and distance


Application characteristics

Results in

Network capacity problems

Utilization overload

Application failure



Information to Collect

Link/segment utilization

CPU utilization

Memory utilization

Response time

Queue/buffer drops

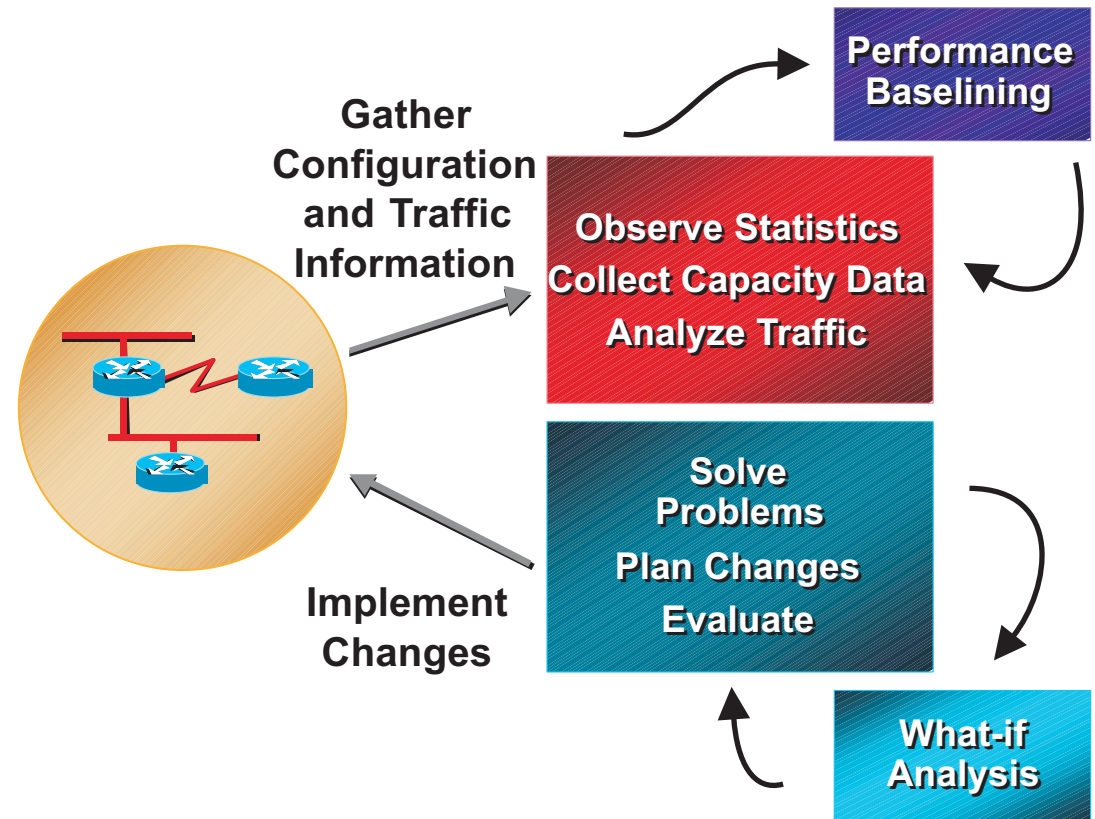
Broadcast volumes

Traffic shaping parameters

RMON statistics

Packet/frame drops/loss

Environment specific



Performance Plan

Develop information collection plan

Define parameters to be monitored/measured and the threshold

Acquire proper authority to change threshold

Determine frequency of monitoring and reporting

Determine frequency of alerting mechanism

Define parameters that trigger alert mechanism

Define performance areas of interest

Report and interpret results

Determine tools for collecting information



Agenda

Introduction and background

Performance Methodologies

Areas of Concern

Am I miss-using buffer space?

What are response times?

What's going on in the IP Stack?

What is system availability?

Who is using my resources?

Who is hogging resources?

What connections are there available?



The Problem

Buffer Utilization

Buffers are critical component of any operating system
Buffers are critical components of any application

Running low or out of buffers on any system can cause

- immediate application failure**
- system slowdown impacting all applications**
- need to restart system**

Running low or out of buffers on any application can cause

- immediate application failure**
- domino effect on related resources and applications**
- intermittent application oddities**

Tuning buffer utilization is important

How do you know hat you are using?

How do you know you are about ready to exceed limits?



Elements

Do you have your buffers pools properly set ?

What are you currently using?


What buffer areas are in expansion?

What is the expansion increments?

If I reallocate buffers can I alert if a buffer reaches a certain utilization point?

Can alerts be forwarded to my operations console, whichever one I choose?

What are the totals for my system and are leaks occurring?



Understanding Buffer Storage

```
A - 3270 (wlaa) - wlaa3279.WS
Z01 CMS 12/09/03 13:58:37 B
> Q. Help PF1
=====
SYSTEM OVERVIEW
=====
SV z/VM V4 R3.0 SLU 0201
CPU: 2064 #09147E-19147E
=====
TRND
+ CPU SUPR PROB STOR PAGE INQ LOG
+ BUSY STATE STATE UTIL RATE USER USER
+current 1% 1% % 67% 8 31
+average 3% 1% 2% 67% 9 31
+peak 8% 1% 7% 68% 9 31
=====
For detail on a user, move cursor to user ID and press PF11.
=====
ZRUS
+ User CPU Virt Pg/s Urec I/O's WSS Res Share Status
+ SYSTEM .0 .1 .1 .0 .0 .0 0 0 0 Suspended
+ VMLNX2 .2 .1 .1 .0 .0 .0 128K 128K 20% 512M Tst Idle
+ VMLNX1 .1 .1 .1 .0 .0 .0 128K 128K 20% 512M Tst Idle
+ VMLNX3 .1 .0 .1 .0 .0 .0 32K 32K 20% 128M Tst Idle
+ GNORR .1 .0 .1 .0 .0 .0 731 1130 100 64M Run'g Q0
+ VTAM .0 .0 .0 .0 .0 .0 865 1329 7500 48M Tst Idle
+ VMUTIL .0 .0 .0 .0 .0 .0 148 254 100 32M Idle
+ VMSEVRU .0 .0 .0 .0 .0 .0 1191 1192 1500 32M Idle
+ VMSEVR3 .0 .0 .0 .0 .0 .0 1182 1183 1500 32M Idle
+ VMSEVR4 .0 .0 .0 .0 .0 .0 1180 1181 1500 32M Idle
=====
For detail on an exception, move cursor to exception name and press PF11.
=====
LEXSY OMEGAMON/VM Exception Analysis
+ SPFU Warning: MAINT using 33 Spool Files
+ SPFU Warning: OMLOGON using 246 Spool Files
+ SPFU Warning: *NSS using 30 Spool Files
+ XMVM Warning: VMLNX4 not running in the system
+ VMID CVGEN D Has been IDLE for 13:58 Hours
+ VMID EREP D Has been IDLE for 16:48 Hours
+ VMID GCS D Has been IDLE for 556:14 Hours
+ VMID OMGEN D Has been IDLE for 13:58 Hours
+ VMID OPERSYMP D Has been IDLE for 556:15 Hours
MA 3 01/001
```

The Problem Response Time

Web users expect 2 to 5 second response time

SNA users expect sub-second response time

No one is ever happy with what they get

External customers may go elsewhere

Where is the problem?

Network? DASD?

Router have long ques?

Is the Lan to slow?

Is the route long?

Operating system?

Too long to queue for transmit?

Application?

Protocol?

Window size improperly set?

MTU size improperly set?



Elements

What are overall response times in my network?

What are response times for different size frames?


Can I look at a specific address and determine its response time?

Are both real time and historical views available?

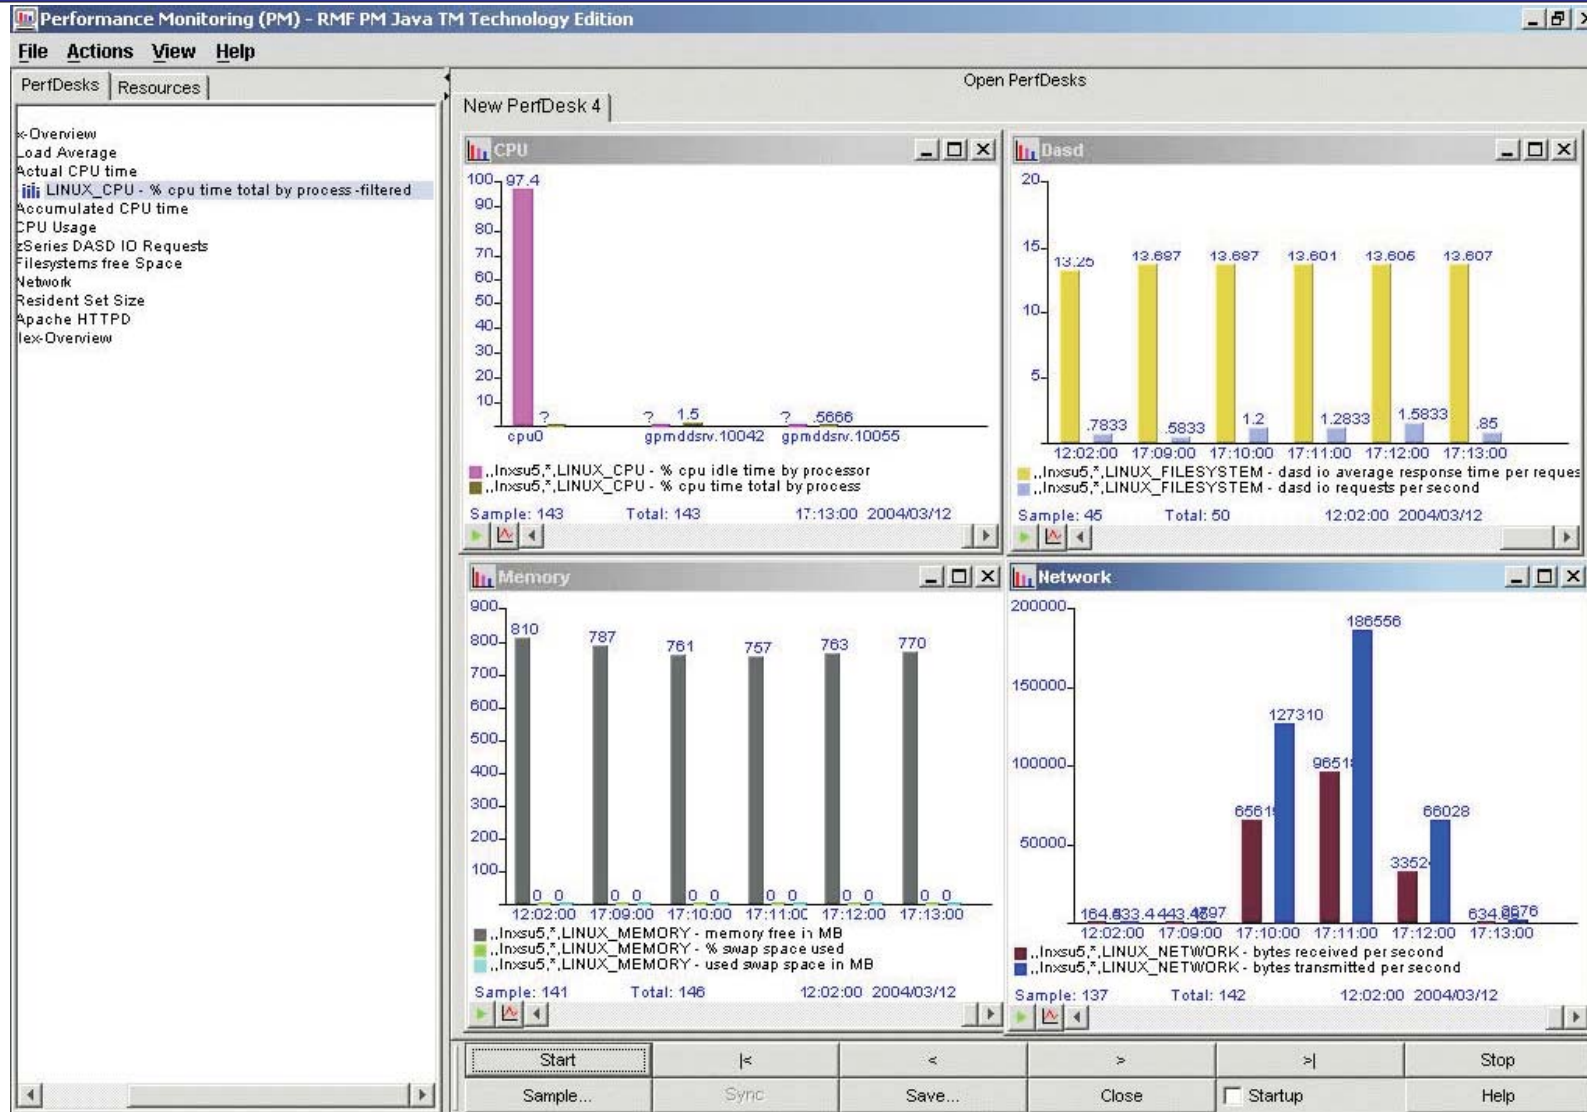
Are both graphical and tabular views available?

Can I set thresholds?

Can I send alerts?



Response Time



The Problem


System Utilization

Since you cannot over-provision your system (add as much memory as you want, as much DASD, etc) you need to optimize

Determining what is currently being used on the system will assist in determining how much you can grow the system

An application behaving poorly may be due to improper design, improper setting of system resources to use, or application configuration

Sluggishness of a system may be due to not enough CPU, I/O overloads, or queue latencies



Elements


What system resources are major address spaces using?

Can I select the address space to view?

How does the utilization change within a given time?

Can I freeze a screen?

Do I have the raw data available to me?



System Utilization

```
A - 3270 (wlaa) - wlaa3279.WS
Z01 CMS OM/VM V610.99 CMS 12/09/03 13:58:37 B
> Q. Help PF1 Back PF3 Up PF7 Down PF8 Zoom PF11
=====
> SYSTEM OVERVIEW
>
> SYS >> z/VM V4 R3.0 SLU 0201
> + CPU: 2064 #09147E-19147E
>
> =====
> TRND
> + CPU SUPR PROB STOR PAGE INQ LOG
> + BUSY STATE STATE UTIL RATE USER USER
> +current 1% 1% % 67% 8 31
> +average 3% 1% 2% 67% 9 31
> +peak 8% 1% 7% 68% 9 31
> =====
> For detail on a user, move cursor to user ID and press PF11.
>
> ZRUSR
> + Userid CPU CP Virt Pg/s Urec I/O's USS Res Share Size Status
> + ---
> + SYSTEM .3 .3 .0 .0 .0 .0 0 3016 0 0 Suspnded
> + VMLNX2 .2 .1 .1 .0 .0 3.4 128K 128K 20xs 512M Tst Idle
> + VMLNX1 .1 .1 .1 .0 .0 .0 128K 128K 20xs 512M Tst Idle
> + VMLNX3 .1 .0 .1 .0 .0 .0 32K 32K 20xs 128M Tst Idle
> + GNORR .1 .0 .1 .0 .0 4.0 731 1130 100 64M Run'g Q0
> + VTAM .0 .0 .0 .0 .0 2.8 865 1329 7500 48M Tst Idle
> + VMUTIL .0 .0 .0 .0 .0 .0 148 254 100 32M Idle
> + VMSERVU .0 .0 .0 .0 .0 .0 1180 1181 1500 32M Idle
> + VMSERVS .0 .0 .0 .0 .0 .0 1180 1181 1500 32M Idle
> + VMSERVV .0 .0 .0 .0 .0 .0 1180 1181 1500 32M Idle
> =====
> For detail on an exception, move cursor to exception name and press PF11.
>
> OMEGAMON/VM Exception Analysis
> + SPFU Warning: MAINT using 33 Spool Files
> + SPFU Warning: OMLOGON using 246 Spool Files
> + SPFU Warning: *NSS using 30 Spool Files
> + XMVM Warning: VMLNX4 not running in the system
> + VMID CVGEN D Has been IDLE for 13:58 Hours
> + VMID EREP D Has been IDLE for 16:48 Hours
> + VMID GCS D Has been IDLE for 556:14 Hours
> + VMID OMGEN D Has been IDLE for 13:58 Hours
> + VMID OPERSYMP D Has been IDLE for 556:15 Hours
```


The Problem Availability

Resources, applications, network components that are not available impact many aspects of your system

IP is especially prone to this due to the 'non-configurable' operations

Critical resources can come and go with no 'network-wide' configuration, but this may impact other systems

Five steps may occur in a process before you realize that the six step requires a resource that is no longer available



The Elements

Can you get a quickview of overall availability?

Can you define critical resources?


Have alerts been sent?

Is the system not available because the system is down or because a resource like a router is having problems (traceroute)?

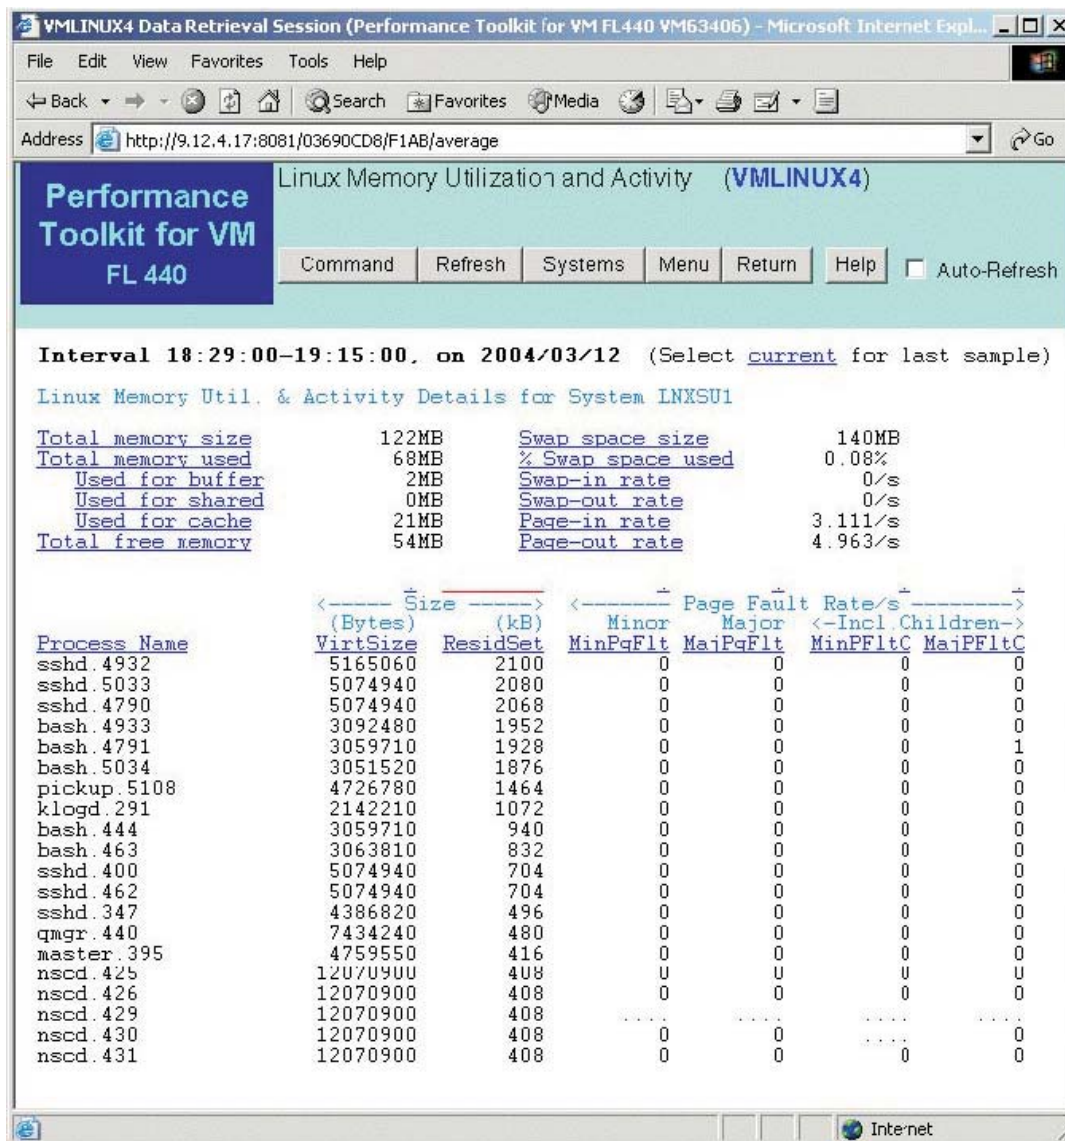
Can I tell if the route is not the normal route taken?

Has the situation cleared itself up?

How can I get more details from an offending intermediary system?



Real Time Availability



The Problem

Resource Utilization

Application usage by end users is very unpredictable in IP. What was valid last week may not be valid today

An application installed on a system and active not being utilized by end users is taking system resources that could be used by other applications

Sometimes it is appropriate to block users after a given number have logged onto an application in order to conserve existing resources

Knowing who is using what on a given system can help determine long term capacity planning needs for the system



The Elements

For a given system, can you determine the applications being used?


Can you tell for each application the session or user counts?

Can you tell for each application the number of bytes transferred?

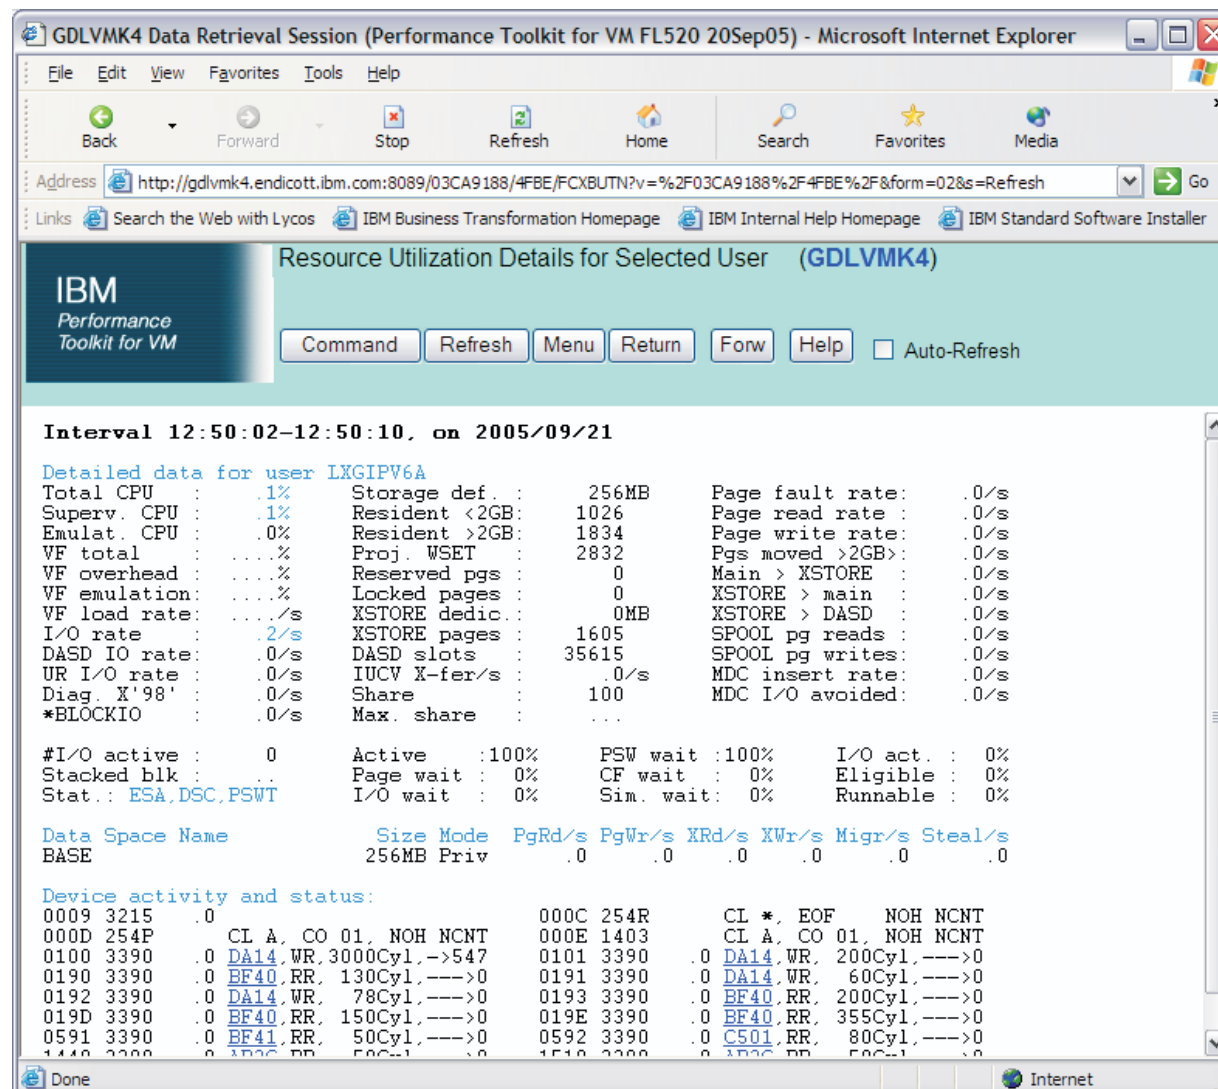
Can you get more details on a historical basis? Determine patterns like a 3% monthly growth in usage?

Can you alert on miss-use of an application?

Can you shut a user out of an application in real time?



General User Resource Utilization




The Problem

Who is Hogging Resources

Excessive backups to a production server by end users can impact production applications

A continuous backup can reduce system resources available for other functions

Your expensive DASD may be used via FTP to hold trivial end user data (like games)



The Elements

Can you determine sessions by applications or bytes by applications?

Can you determine top 10 clients bytes transferred?

For selected clients can you determine bytes transferred?

Can you determine past history?

Can you determine if alerts have been sent?

Can you view not only the IP address but the DNS name?

Can you set refresh rate?



Application

```
A - 3270 (wlaa) - wlaa3279.WS
Z01 CMS OM/VM V610.99 CMS 12/09/03 13:58:37 B
> Q. Help PF1 Back PF3 Up PF7 Down PF8 Zoom PF11
=====
> SYSTEM OVERVIEW
=====
SYS >> z/VM V4 R3.0 SLU 0201
+ CPU: 2064 #09147E-19147E
=====
TRND
+ CPU SUPR PROB STOR PAGE INQ LOG
+ BUSY STATE STATE UTIL RATE USER USER
+current 1% 1% % 67% 8 31
+average 3% 1% 2% 67% 9 31
+peak 8% 1% 7% 68% 0 31
=====
> For details, move cursor to user ID and PF1.
=====
2RUSER
+ Userid CPU CP Vint Pg/s Urec I/O's WSS Res Share Size Status
+ SYSTEM .3 .3 .0 .0 .0 .0 0 3016 0 0 Suspended
+ VMLNX2 .2 .1 .1 .0 .0 .0 3.4 128K 128K 20xs 512M Tst Idle
+ VMLNX1 .1 .1 .1 .0 .0 .0 128K 128K 20xs 512M Tst Idle
+ VMLNX3 .1 .0 .1 .0 .0 .0 32K 32K 20xs 128M Tst Idle
+ GNORR .1 .0 .1 .0 .0 .0 731 1130 100 64M Run'g Q0
+ VTAM .0 .0 .0 .0 .0 .0 2.8 865 1329 7500 48M Tst Idle
+ VMUTIL .0 .0 .0 .0 .0 .0 148 254 100 32M Idle
+ VMSERVU .0 .0 .0 .0 .0 .0 1191 1192 1500 32M Idle
+ VMSERVS .0 .0 .0 .0 .0 .0 1182 1183 1500 32M Idle
+ VMSRVR .0 .0 .0 .0 .0 .0 1180 1181 1500 32M Idle
=====
> For details on an exception, move cursor to exception name and PF1.
=====
LEXSY OMESY Exception Analysis
+ SPFU Warning: MAIN using 248 Spool Files
+ SPFU Warning: OMLOGON using 30 Spool Files
+ SPFU Warning: *NSS using 30 Spool Files
+ XMVM Warning: VMLNX4 not running in the system
+ VMID CVGEN D Has been IDLE for 13:58 Hours
+ VMID EREP D Has been IDLE for 16:48 Hours
+ VMID GCS D Has been IDLE for 556:14 Hours
+ VMID OMGEN D Has been IDLE for 13:58 Hours
+ VMID OPERSYMP D Has been IDLE for 556:15 Hours
MA a 01/001
```


The Problem

A System Perspective

You are dealing with a system, not just a standalone computer. Other tools in the network may give you the views you want, but without access, the information is not readily available to you

Standalone CPU based tools are not expandable to view the outside components

Others will need to be involved as you delve into problems, but the tools at your disposal need to give you basic information in order to proceed



The Elements

Can you determine information on not only the CPU involved, but also other network components that may be impacting the problem?

Can you determine availability and response times for the users of your CPU?

**Real time and historical data is needed
.. One to solve immediate problems
... One to allow capacity planning**

Are commands access provided as well as alerting to operational consoles

Can a new employee quickly learn the system?



Performance Management Pride

Performance Toolkit for VM FL 440

Initial Performance Data Selection Menu (VMLINUX4)
Select performance screen

Command Refresh Systems Forw Help ☐ Auto-Refresh

General System Data	I/O Data	History Data (by Time)
1. CPU load and trans.	11. Channel load	31. Graphics selection
2. Storage utilization	12. Control units	32. History data files*
3. Storage subpools	13. I/O device load*	33. Benchmark displays*
4. Priv. operations	14. CP owned disks*	34. Correlation coeff.
5. System counters	15. Cache extend. func.*	35. System summary*
6. CP IUCV services	16. DASD I/O assist	36. Auxiliary storage
7. SPOOL file display*	17. DASD seek distance*	37. CP communications*
8. LPAR data	18. I/O prior. queueing*	38. DASD load
9. Shared segments	19. I/O configuration	39. Minidisk cache*
A. Shared data spaces	1A. I/O config. changes	3A. Paging activity
B. Virt. disks in stor.		3B. Prcc. load & config*
C. Transact. statistics	User Data	3C. Logical part. load
D. Monitor data	21. User resource usage*	3D. Response time (all)*
E. Monitor settings	22. User paging load*	3E. RSK data menu*
F. System settings	23. User wait states*	3F. Scheduler queues
G. System configuration	24. User response time*	3G. Scheduler data
H. VM Resource Manager	25. Resources/transact.*	3H. SFS/BFS logs menu*
I. Exceptions	26. User communication*	3I. System log
K. User defined data*	27. Multitasking users*	3K. TCP/IP data menu*
	28. User configuration*	3L. User communication
	29. Linux systems*	3M. User wait states

Service Level Management

Define performance requirements
Define upgrade criteria by performance
Measure performance
Review thresholds and baseline

Threshold	WAN	LAN
CPU	75-90%	75-90%
Link	80-90%	40-90%
Memory	50%	50%
Output Queue	200	25
Buffer Misses	Any	Any
Broadcast Vol	10/Sec	300/Sec
FECN/BECN	10/Sec	N/A

Performance Summary

You never solve performance problems.....You just keep moving them

The basic performance issues remain the same.....But QoS adds a new view

Emerging applications need higher levels of performance

Performance data readily available.....but the interpretation and action plans are lax

Complexity

Expect change and new ideas to emerge

Policy systems required to ease administration complexity

