IBM **z Systems**

Customer Experiences:

Monitoring and Managing z/VM, Linux on z Systems and LinuxONE

Mike Sine
IBM
sine@us.ibm.com

October 2016

Agenda

- What does “managing” include?
  - What tools or products can you use?
- Customer scenarios
  - Operational monitoring and automation
  - Performance monitoring
  - Backup and recovery

What is “Managing” and What Tools Can I Use?

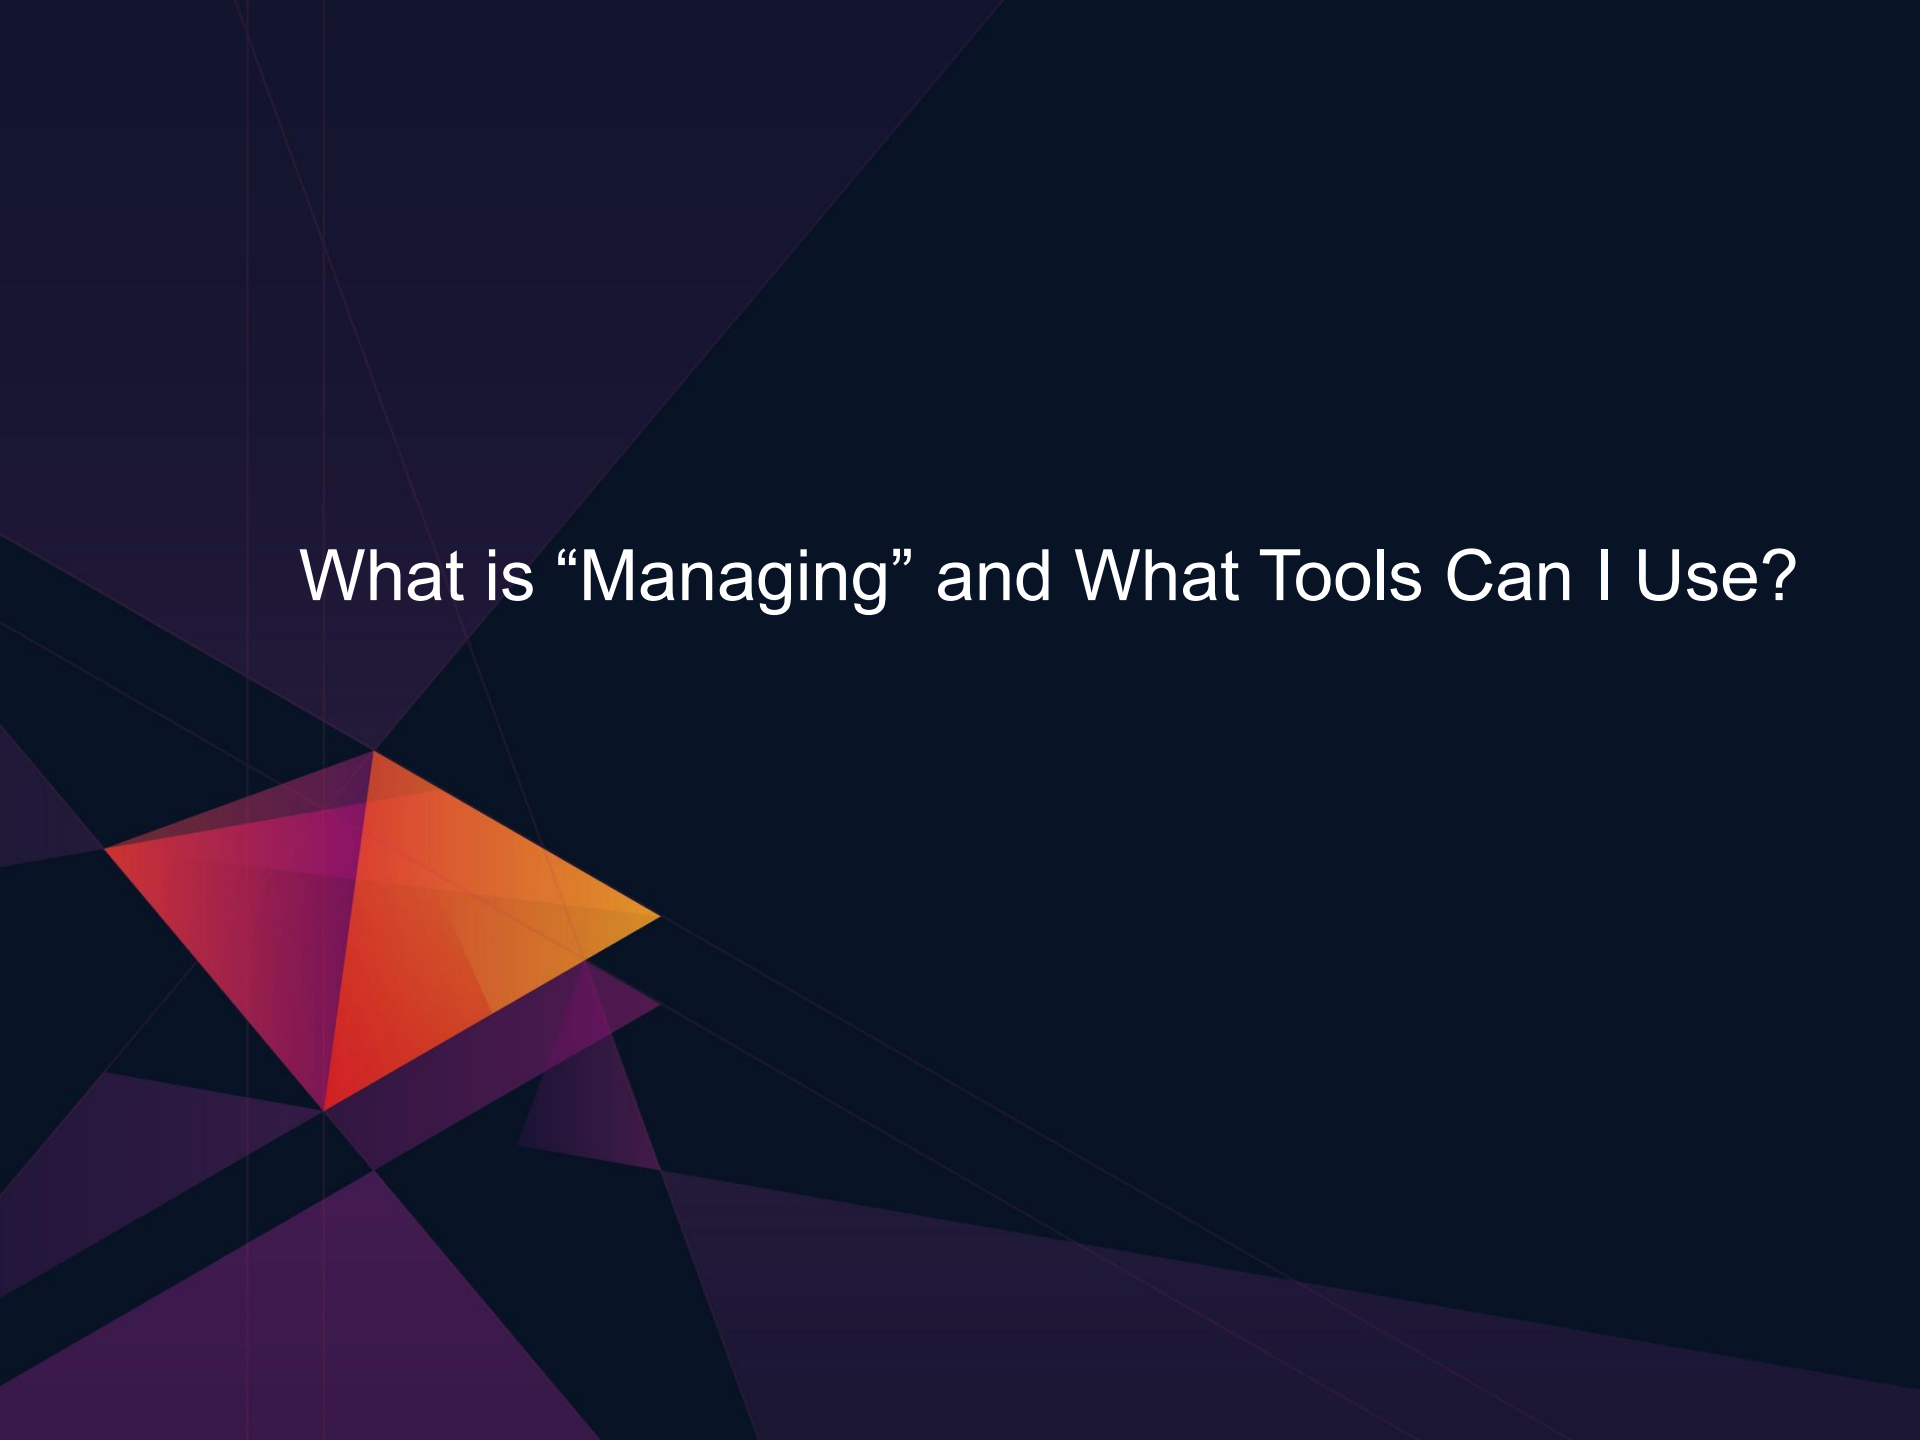

Administration and Provisioning

Administer Linux guests/servers via GUI
- View of all servers graphically
- Run shell scripts against a server or group of servers
- Activate or deactivate a server or group of servers
- Login to server directly from GUI
- View and modify network connections

Provision Linux guests/servers
- Across LPARs or machines
- Memory and CPU
- Network – connect to Guest LANs or VSWITCHes
- Storage – based on admin-defined device pools
- Customize first boot before TCPI/IP initialized
- Customize cloning via REXX scripts

Real time monitoring
- High level view of system status via dashboard gauges
- View storage utilization

Administration and Provisioning

Manage and administer Linux guests/servers via GUI
- View of all servers graphically
- Run shell scripts against a server or group of servers
- Activate or deactivate server or group of servers
- Login to server directly from GUI
- View and modify connections

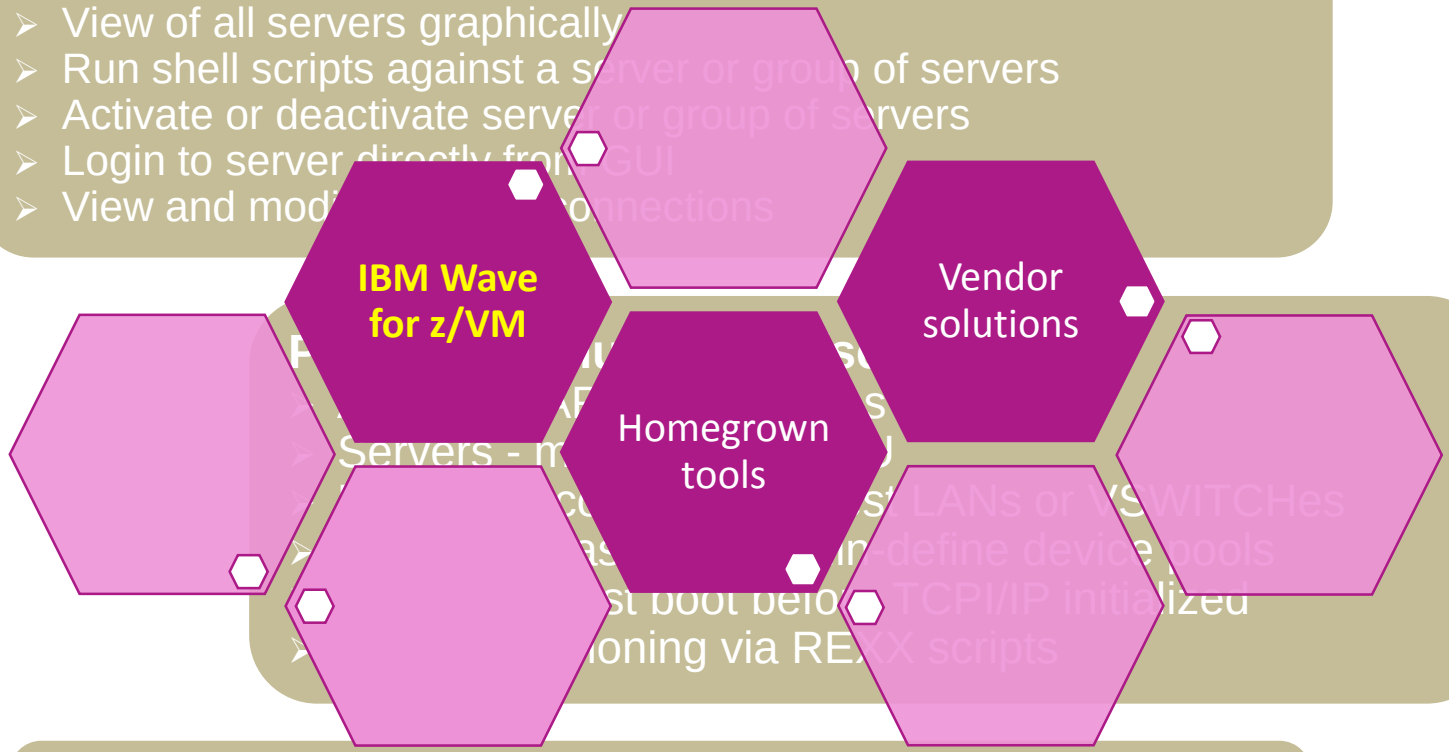

- Servers - m
- st boot befo
- ioning via RE

Real time monitoring
- High level view of system status via dashboard gauges
- View storage utilization

Performance Monitoring and Automation

Monitor performance based on best practices

- Virtual CPU for each guest
- z/VM processor utilization
- Spin lock wait
- Virtual disk utilization
- Virtual storage utilization with V/R memory ratio
- Formation and size of eligible list
- Page and spool space utilization and I/O rates
- DASD I/O and minidisk cache usage
- Resource constraint analysis

Use historical data to

- Understand capacity
- Size Linux guests for best performance in a hosted (shared) environment

Performance Monitoring and Automation

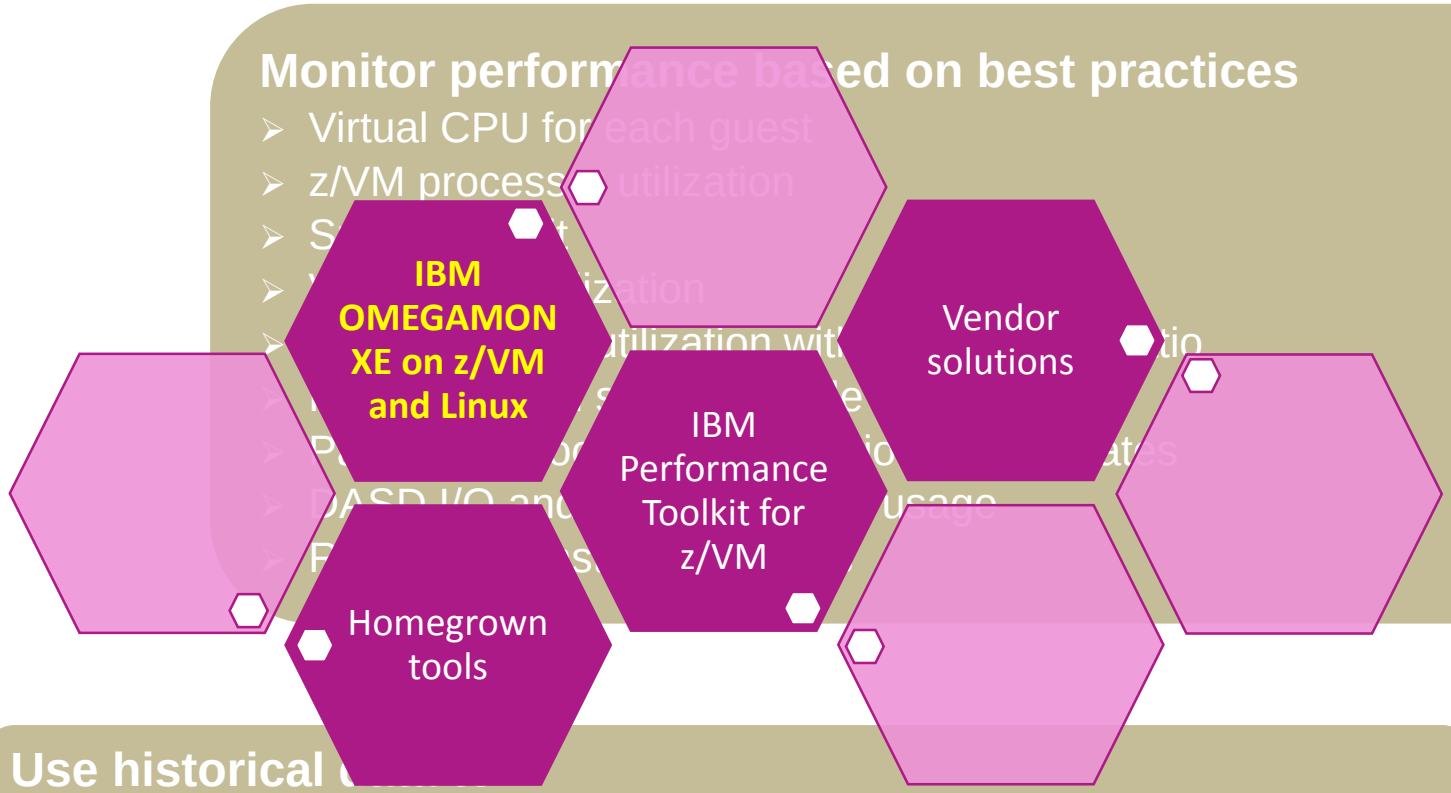

Monitor performance based on best practices

- Virtual CPU for each guest
- z/VM process utilization
- S
- V ization
- utilization wit tio
- P e
- P o ates
- DASD I/O and usage
- R s

Use historical

- Understand capacity
- Size Linux guests for best performance in a hosted environment

Operational Monitoring and Automation

Console monitoring and viewing

- Operations staff monitoring a central console of alerts
- System programmers debugging a problem on a guest or service machine
- Console log data available for audits or future reference

Generate alerts and/or automatically recover from

- Abend, termination, or error messages
- Service machine disks approaching full
- Critical user IDs or guests being logged off or entering error state
- Spool and/or page space approaching full

Schedule automated system maintenance procedures

- Spool cleanup based on policies
- Minidisk cleanup (from logs), including archiving
- Orderly startup and shutdown
  - Relocation of critical guests to another SSI member
- Backups of z/VM system

Operational Monitoring and Automation

View & issue commands on consoles of Linux guests and CMS service machines

- Operations staff monitoring multiple consoles or a central console of alerts

System programmers debugging a problem on a guest or service machine

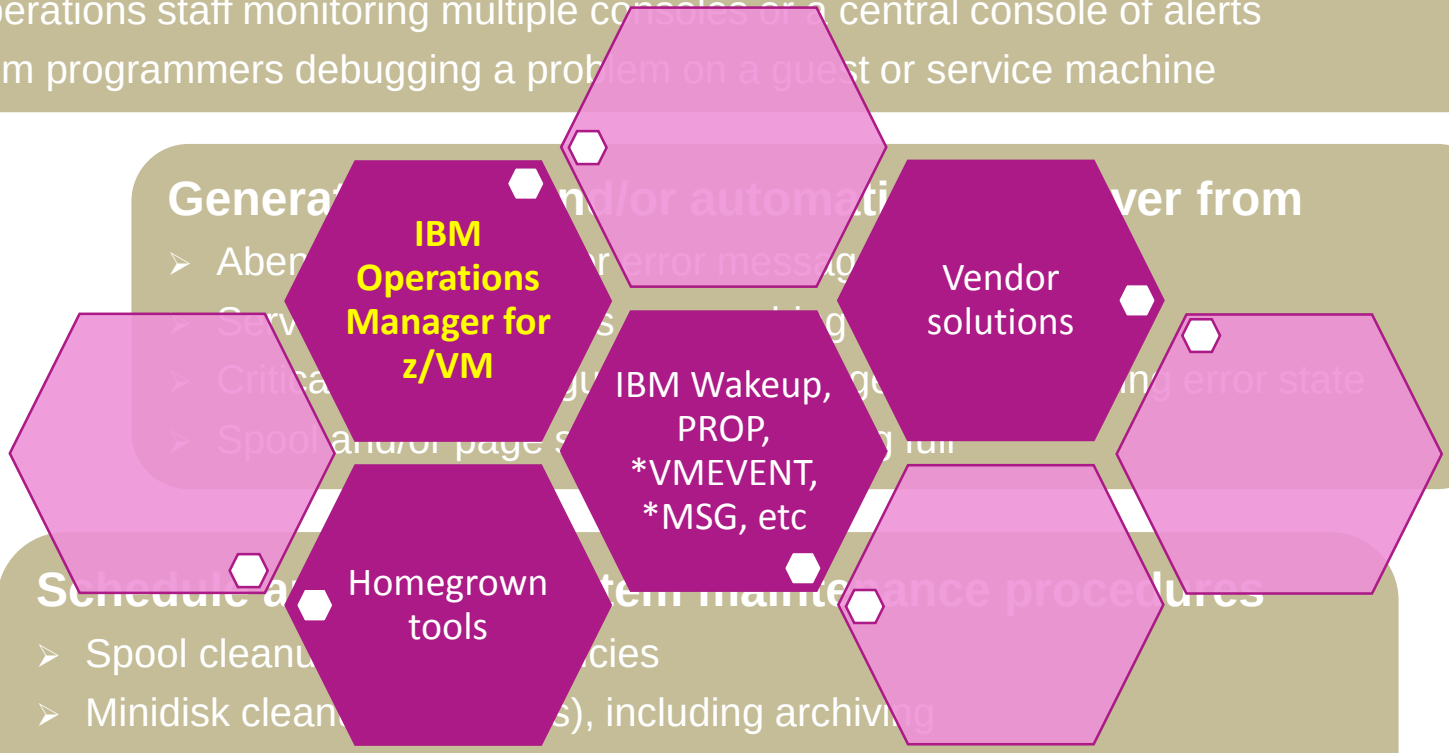

Schedule and automate system maintenance procedures

- Spool cleanup policies
- Minidisk cleanup (DIRMAINT), including archiving
- Orderly startup and shutdown
  - Relocation of critical guests to another SSI member
- Backups of z/VM system

Backup and Recovery of z/VM and Linux

File level backup of z/VM data
- Directory information
- Configuration files
- Log files
- Tools – REXX EXECs, automation scripts, etc.

Image level backup of z/VM
- Operating system

Image level backup of Linux guests
- Operating system
- Applications
- Application data (maybe)

File level backup of Linux guests
- Configuration files
- Log files
- Tools

Backup and Recovery of z/VM and Linux

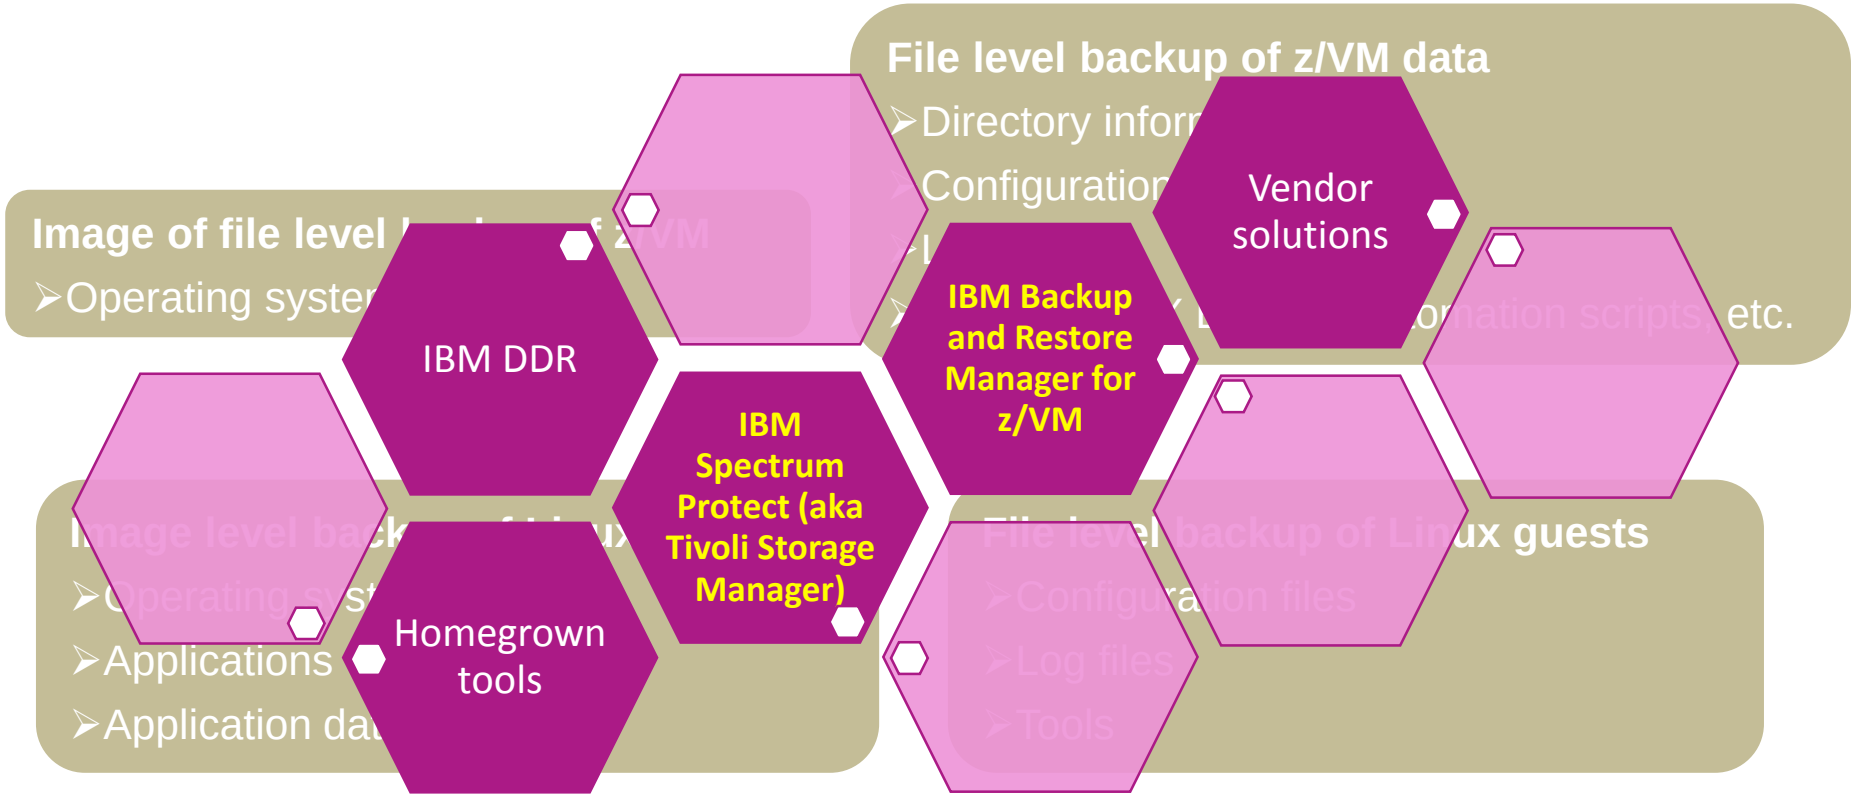

Customer Scenarios

Operational Monitoring and Automation

Performance Monitoring and Troubleshooting

Backup and Recovery

Error Messages on Linux IPL

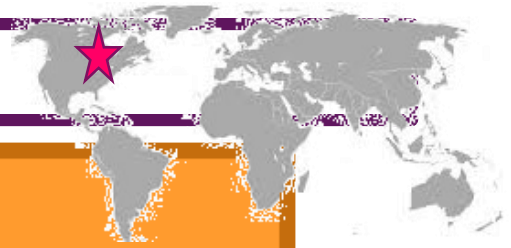

The Situation:

- During boot process, Linux file system is **read-only**
- Application needs read/write
  - But sometimes not until hours or days after boot
- Error discovered **hours or days later** when application fails

Operations Manager

Initial solution

Write homegrown tool

Scan logs on a daily basis looking for error messages

Final solution

Console monitoring tool

Write a rule looking for error message during boot process and take action immediately

Error Message on z/VM IPL

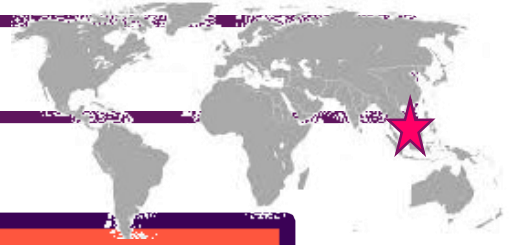

The Situation:

- Error messages on z/VM IPL
- **Reason unknown** to customer (new to z/VM)

Operations Manager

Initial solution

None

- Took photo of HMC with smartphone
- Show IBM and ask for help
  - **EREP & Accounting disks full**
- No knowledge of impact

Final solution

Monitoring tool

- Simple monitor setup
- Automatically monitor percent full
- Email someone who can follow documented procedures to save/archive data

Hipervisor Using 25% of CPU

The Situation:

- Most monitoring focuses on CPU utilization overall
- Missing focus on **CP's % of CPU** as a separate metric
  - How much is the hipervisor using?
- Best Practice is to investigate if hipervisor using > 10% of CPU
- One morning found CP% at 25%, simple drill down revealed cause

OMEGAMON

Initial solution

None

- System CPU measured, while CP specific numbers omitted
- Only reactive steps taken when performance issue arose

Final solution

Monitoring tool

- Automatically monitor CP % for threshold of 10%
- Once threshold is alerted, simple proactive drill down in tool reveals impact often before downstream performance impact is noticed

System Processor Utilization Workspace

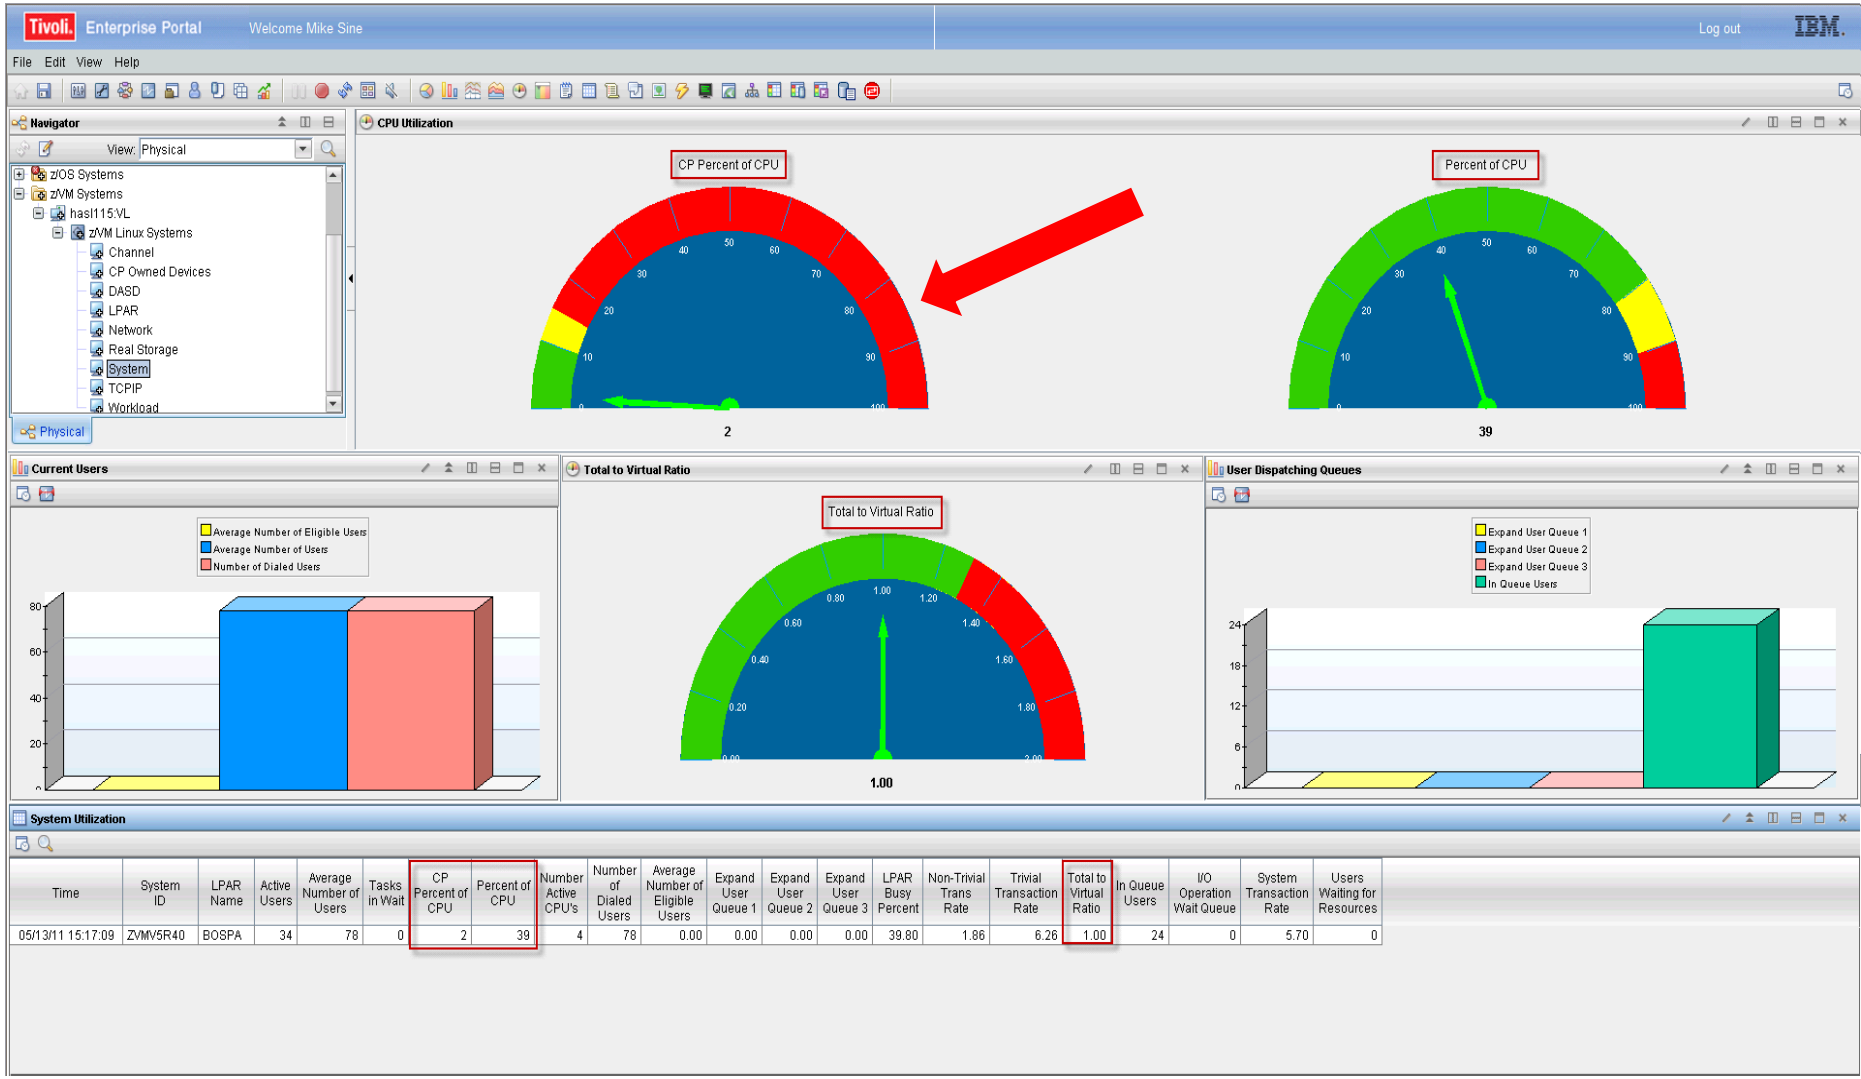

Time	System ID	LPAR Name	Active Users	Average Number of Users	Tasks in Wait	CP Percent of CPU	Percent of CPU	Number Active CPU's	Number of Dialed Users	Average Number of Eligible Users	Expand User Queue 1	Expand User Queue 2	Expand User Queue 3	LPAR Busy Percent	Non-Trivial Trans Rate	Trivial Transaction Rate	Total to Virtual Ratio	In Queue Users	I/O Operation Wait Queue	System Transaction Rate	Users Waiting for Resources
05/13/11 15:17:09	ZVMV5R40	BOSPA	34	78	0	2	39	4	78	0.00	0.00	0.00	0.00	39.80	1.86	6.26	1.00	24	0	5.70	0

z/VM Workload Workspace

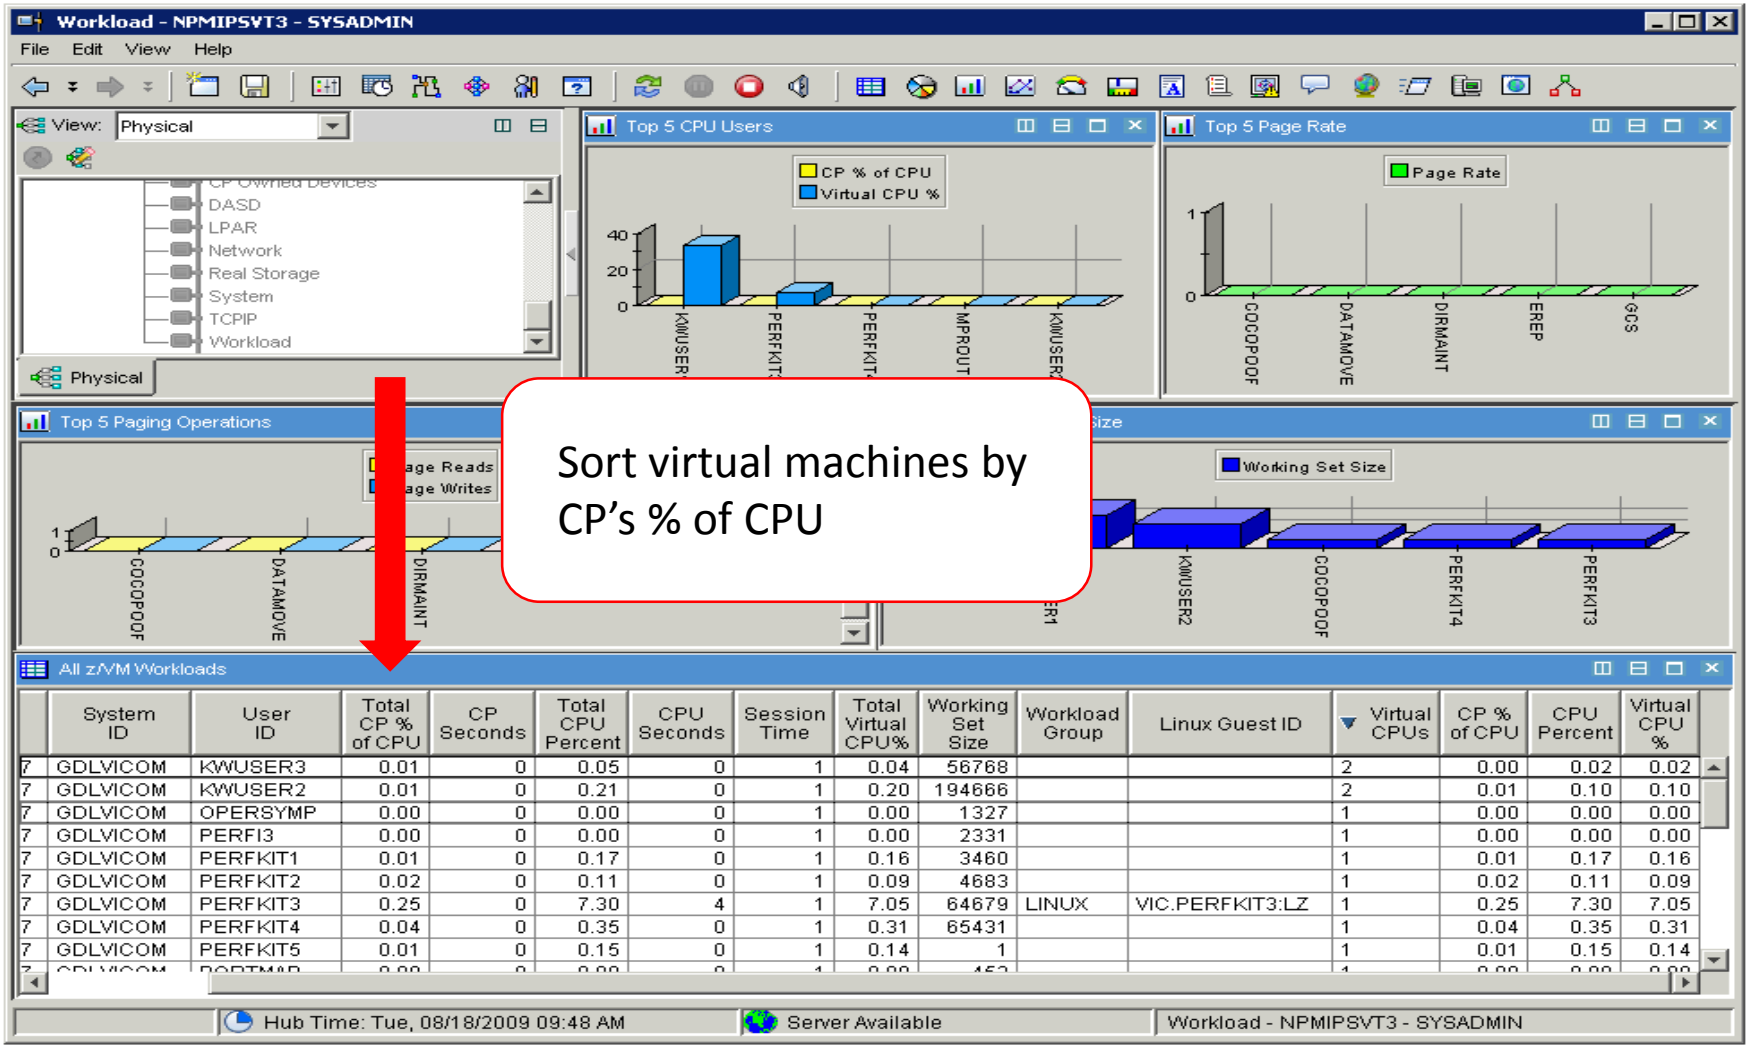

	System ID	User ID	Total CP % of CPU	CP Seconds	Total CPU Percent	CPU Seconds	Session Time	Total Virtual CPU%	Working Set Size	Workload Group	Linux Guest ID	Virtual CPUs	CP % of CPU	CPU Percent	Virtual CPU %
7	GDLVICOM	KWUSER3	0.01	0	0.05	0	1	0.04	56768			2	0.00	0.02	0.02
7	GDLVICOM	KWUSER2	0.01	0	0.21	0	1	0.20	194666			2	0.01	0.10	0.10
7	GDLVICOM	OPERSYMP	0.00	0	0.00	0	1	0.00	1327			1	0.00	0.00	0.00
7	GDLVICOM	PERFI3	0.00	0	0.00	0	1	0.00	2331			1	0.00	0.00	0.00
7	GDLVICOM	PERFKIT1	0.01	0	0.17	0	1	0.16	3460			1	0.01	0.17	0.16
7	GDLVICOM	PERFKIT2	0.02	0	0.11	0	1	0.09	4683			1	0.02	0.11	0.09
7	GDLVICOM	PERFKIT3	0.25	0	7.30	4	1	7.05	64679	LINUX	VIC.PERFKIT3:LZ	1	0.25	7.30	7.05
7	GDLVICOM	PERFKIT4	0.04	0	0.35	0	1	0.31	65431			1	0.04	0.35	0.31
7	GDLVICOM	PERFKIT5	0.01	0	0.15	0	1	0.14	1			1	0.01	0.15	0.14

System ABEND with No Console Data

The Situation:

- **Legacy** best practice of **spooling consoles**
- System abends
- IPL with warm start unsuccessful or not possible
- **No console data** to review what happened leading up to abend
- Dump data only

Operations Manager

Initial solution

IPL cold start and hope for the best

Or

IPL cold start and dig through dump data

Final solution

Console monitoring tool

IPL cold start and review console data written in one log file on disk

Spool and Page Space Full

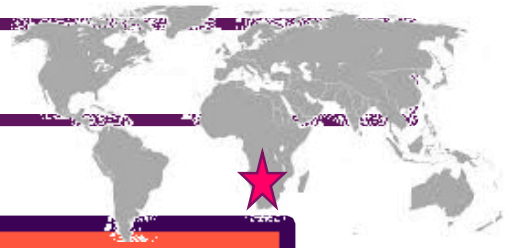

The Situation:

- **Spool** and **page** space **fill up**
- System abends
- **Unplanned outage**

Operations Manager

Initial solution

Homegrown tool

- Create a service machine running WAKEUP
- Check spool and page space percent full on regular intervals
- Maintain service machine and code

Final solution

Monitoring tool

- Simple monitor setup
- Watch for percent full to be within threshold range
- Watch for sudden growth
- Take action
- Easily add or change threshold or frequency

Resource Utilization Reports

OMEGAMON

The Situation:

- Linux admins misinterpret utilization of their virtual servers
- Overwhelm support with (unnecessary) demands for additional resources
- Sysadmin tools don't show correct breakdown in a virtual server

Initial solution

SysAdmin Tools

- Tools like TOP and others don't reflect the virtualized environment.
- Users get mixed information and make wrong conclusions.
- Misunderstanding between application owners, Linux admins, and system providers

Final solution

Monitoring tool

- Develop reports
  - CPU utilization max and average
  - Monthly memory utilization breakdown
- Linux admins and application owners satisfied they are getting necessary resources

Maximum and Average CPU example

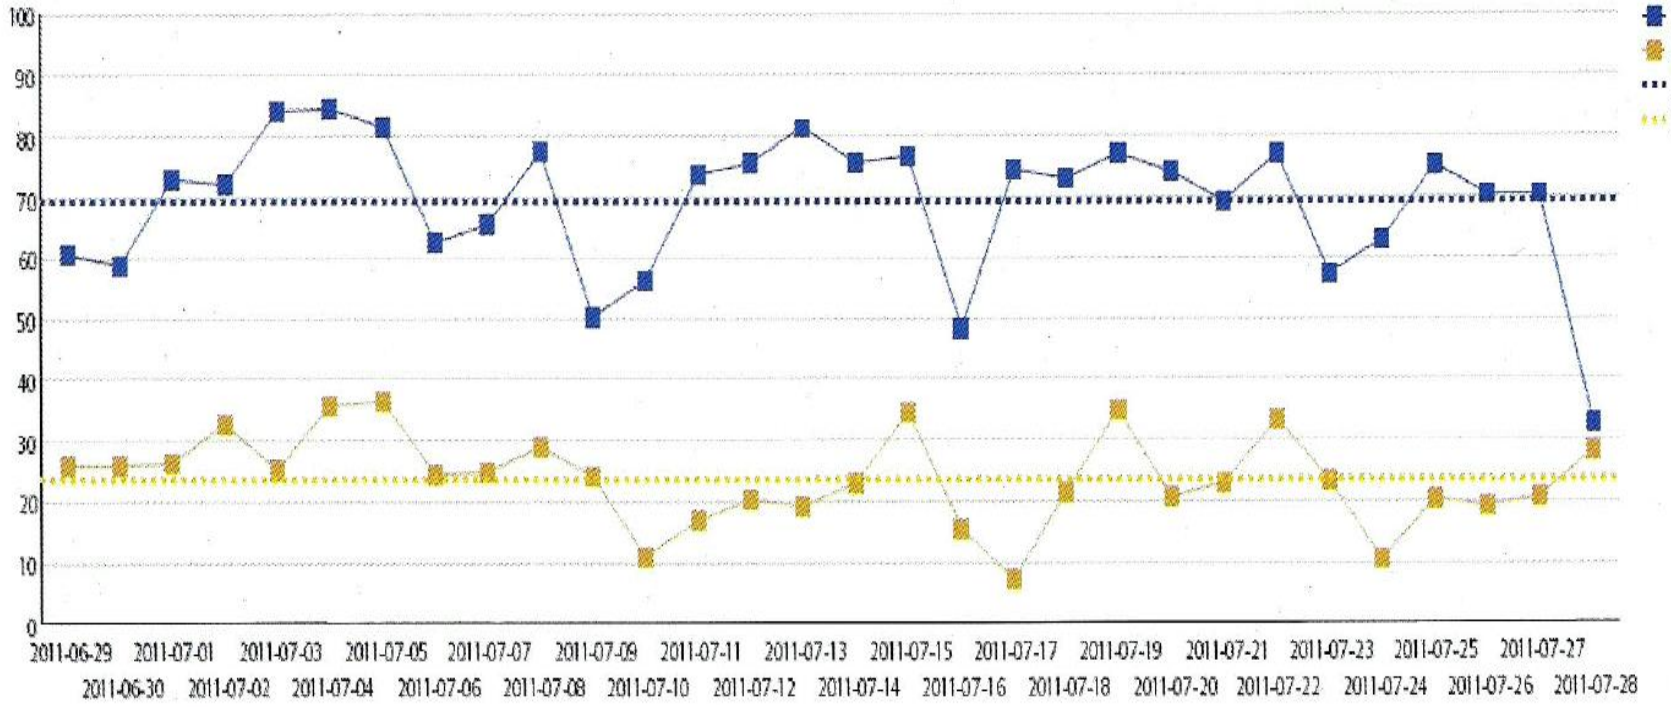

Legend:

Max_CPU_Percent:	Maximum CPU for the day as a percent of the number of virtual CPUs
Avg_CPU_Percent:	Average CPU for the day as a percent of virtual CPUs
Mean of the Maximum:	30 day average for Maximum CPU percentages
Mean of the Averages:	30 day average for the average CPU percentages
AVG_Main_Memory_Util:	Average main memory utilization for the day as a percent
AVG_Cache_Used:	Average size of memory used to cache buffers in megabytes
AVG_Page_Alloc_Rate:	Average number of pages obtained from available list in 4 kilobyte pages per second
AVG_Swap_Used:	The percent of swap space used.

Average Linux Memory Breakdown Example

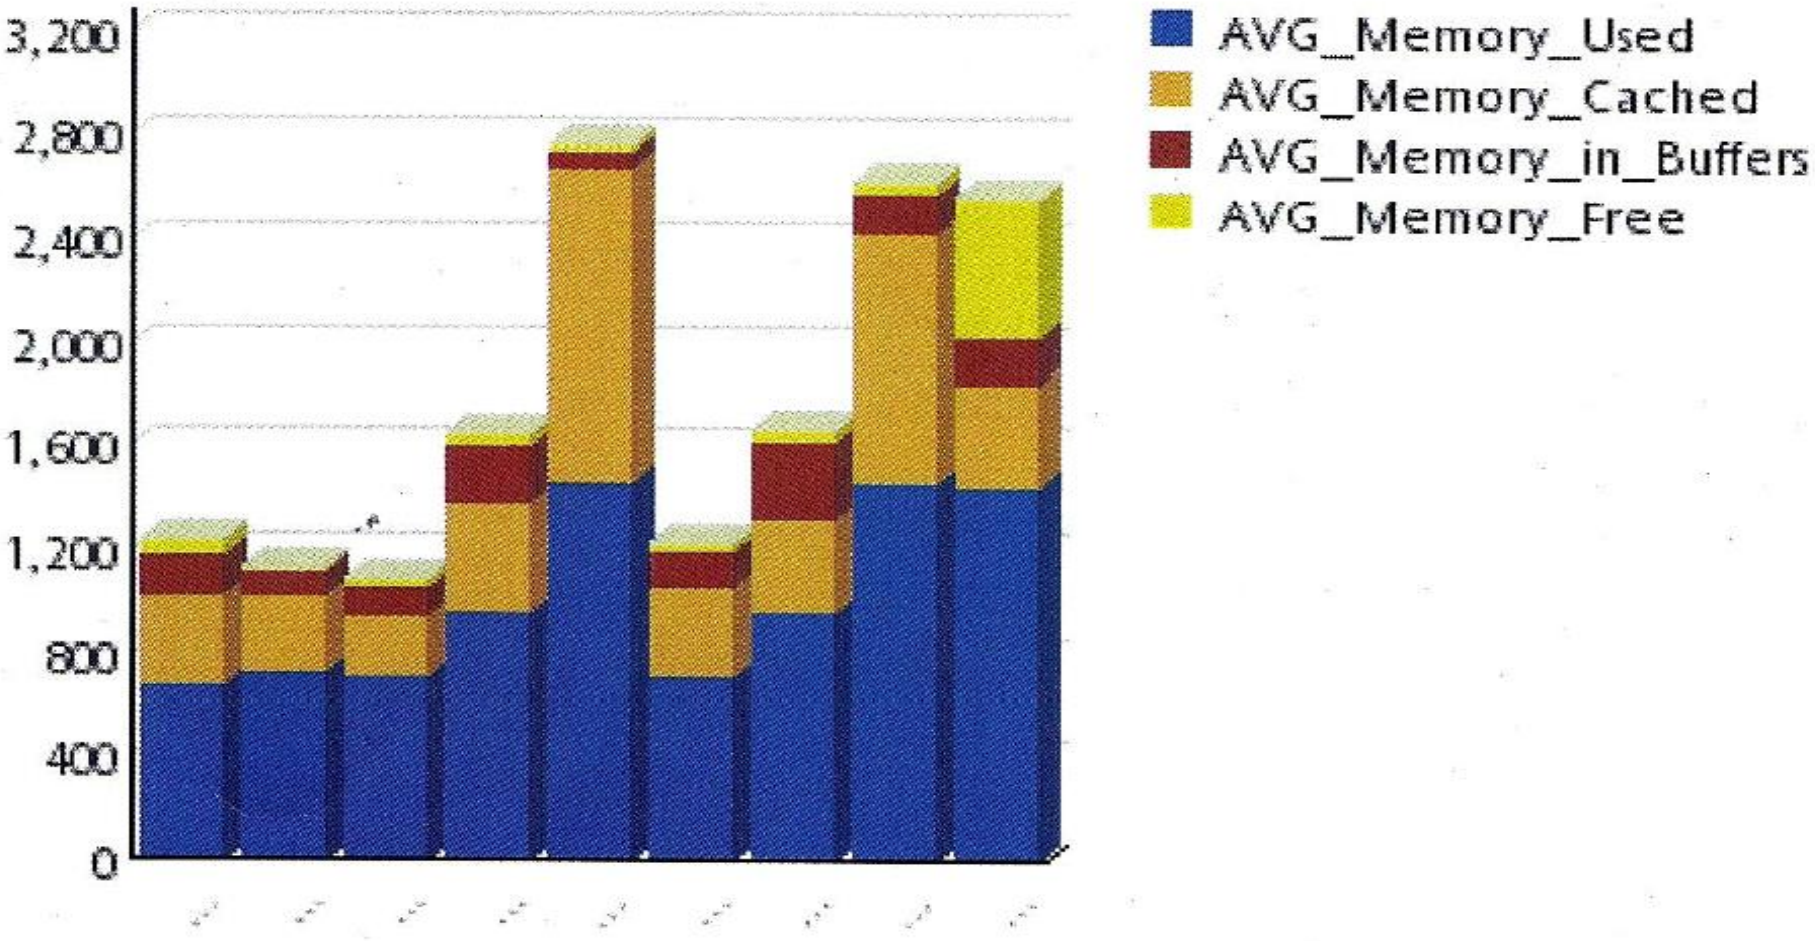

Painful Recovery of Critical z/VM Files

The Situation:

- Backups of z/VM volumes done from **z/OS**
- Operational issue (aka user error) **corrupts** a configuration file
- Recovery is **tedious** and error-prone process
  - Restoring whole volume
  - Mapping a new minidisk to the right location on the volume
- Recovery **very** difficult if corrupted file is **USER DIRECT**

Backup Manager

Initial solution

Train people to make backup copies before updating a file

Final solution

File level backup and recovery

Weekly full backups and daily incrementals of all z/VM files

Why Was an Application Running Slow

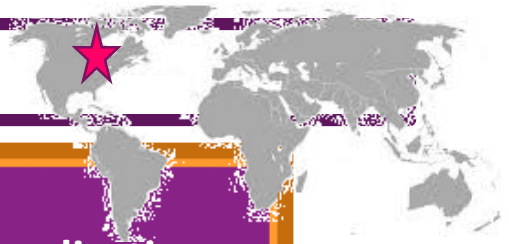

The Situation:

- Application owner asks z/VM system programmer why **application** was running **slow yesterday** afternoon
- Application owner doesn't have the data he needs to research the problem

OMEGAMON

Initial Solution

Look at performance data for the Linux guest

- **Performance data in logs for the Linux operating system**
- **No application data**

Final solution

One performance monitoring solution for all layers

- **Hipervisor**
- **Linux operating system**
- **Application**

Why Was an Application Running Slow

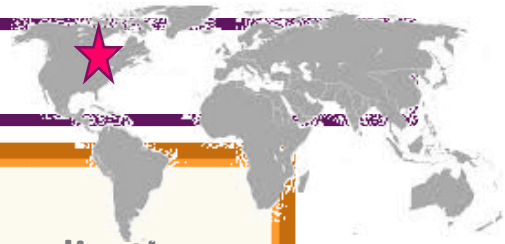

The Situation:

- Application owner asks z/ rammer why **application** was running **slow y**
- Application owne ls to research the problem

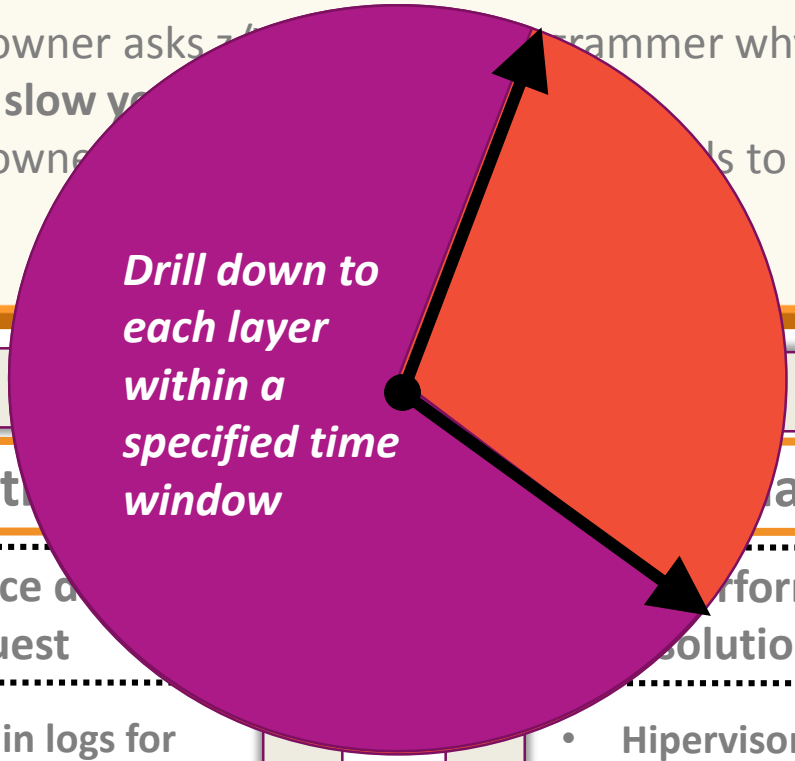

Initial Solut

Look at performance d the Linux guest

- Performance data in logs for the Linux operating system
- No application data

al solution

rformance monitoring olution for all layers

- Hipervisor
- Linux operating system
- Application

Why Was an Application Running Slow

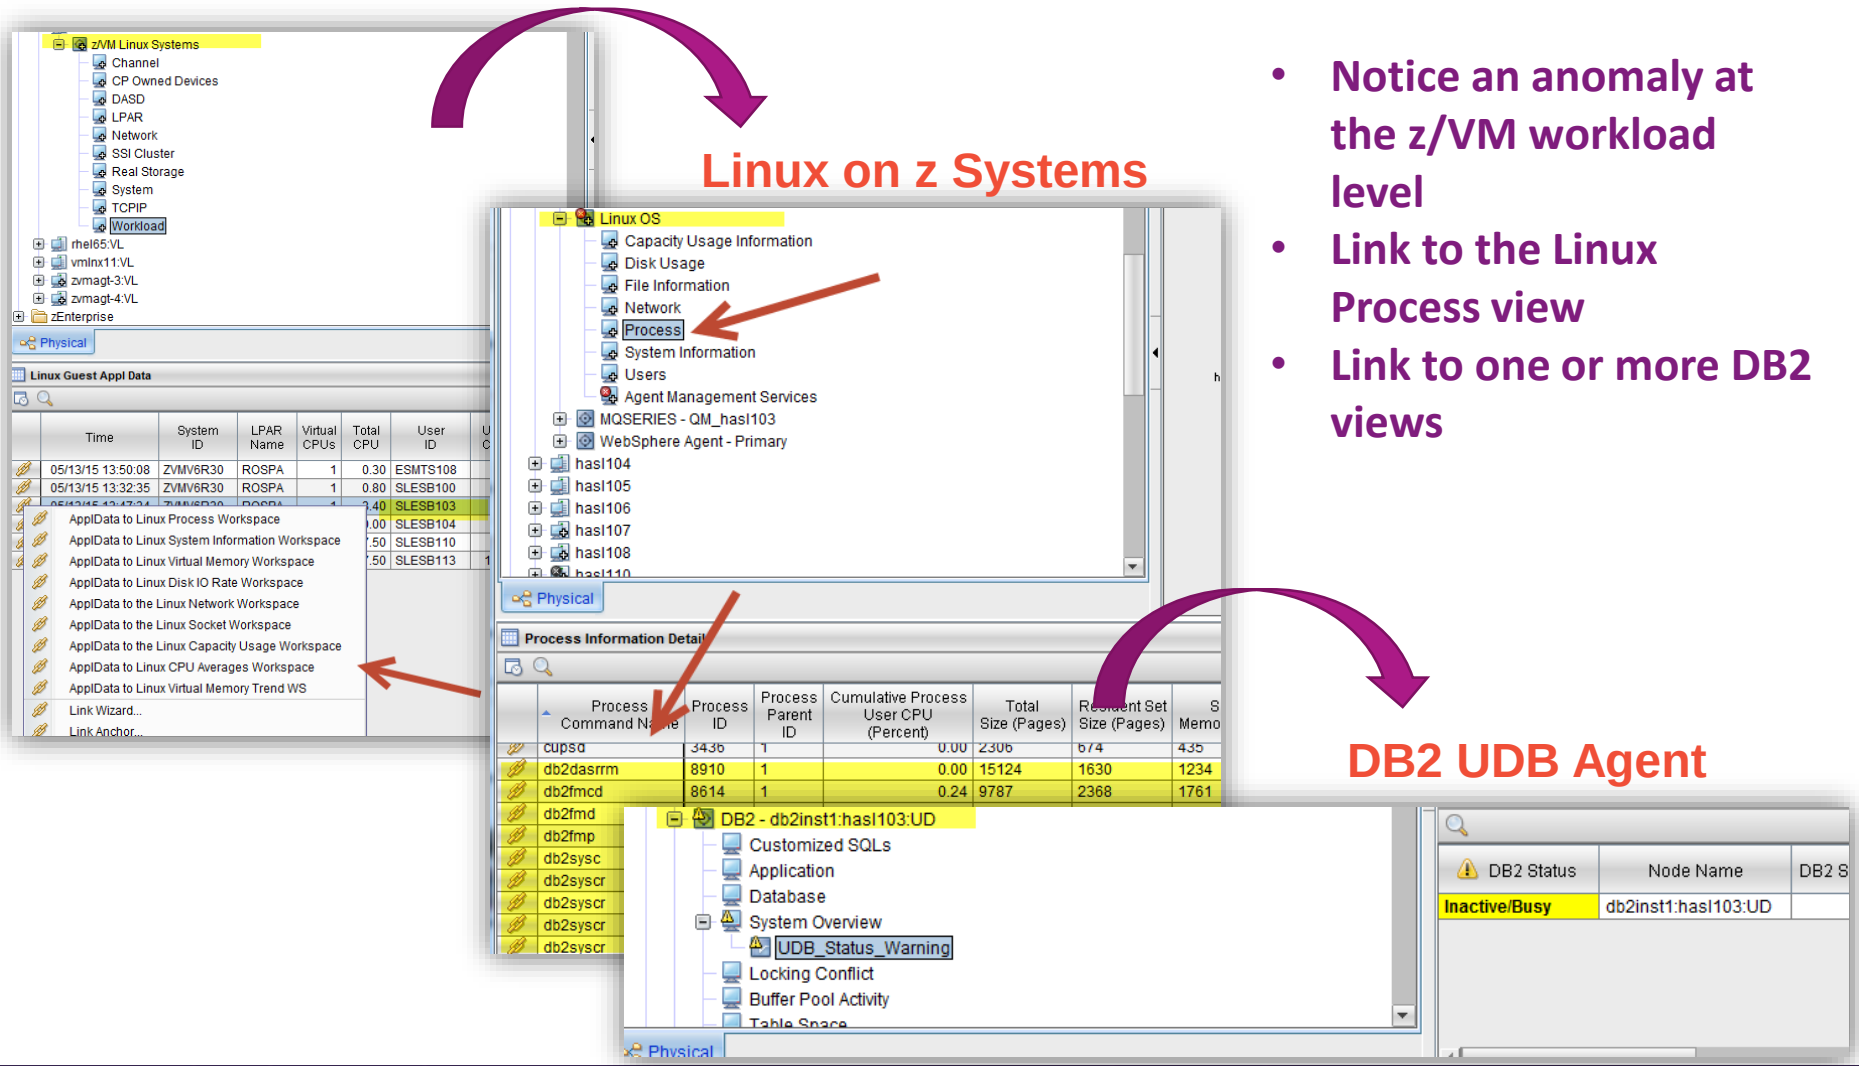

- Notice an anomaly at the z/VM workload level
- Link to the Linux Process view
- Link to one or more DB2 views

Perform Weekly System Healthcheck

The Situation:

Need to monitor system to verify not approaching a threshold

- **Spool space** filling up
- **Paging space** filling up
- **Disk full** for several z/VM service machines or guests

EREP
SMTP
DIRMAINT
…

Operations Manager

Initial solution

Logon weekly and go through checklist manually

Check disk space
Check page space
Check spool space

Final solution

Automate regular monitoring and alerts

Email team if anything approaches threshold

Perform Weekly System Healthcheck

The Situation:

- Need to monitor system to verify not approaching a threshold
  - **Disk full** for several z/VM service machines or guests

- Add additional automation to automatically clean up the disk
  - Back up or archive data
  - Erase files

Initial solution

Logon weekly and go through checklist manually

Check disk space
Check page space
Check spool space

Final solution

Automate regular monitoring and alerts

Email team if anything approaches threshold

Central Operations Console

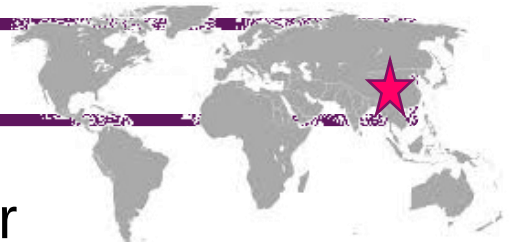

- Already have z/OS console in operations center
  - Alerts, important messages, etc. for operations staff
- Want **one** console for all **z/VM** LPARs and **Linux** guests
  - Operations staff sees **only important messages** on central console
  - **When needed** can also look at **full console** of any specific user ID or guest
  - Can expand to include more LPARs as environment grows
    - Still a **single** console

Creating a Central Console Operations Console

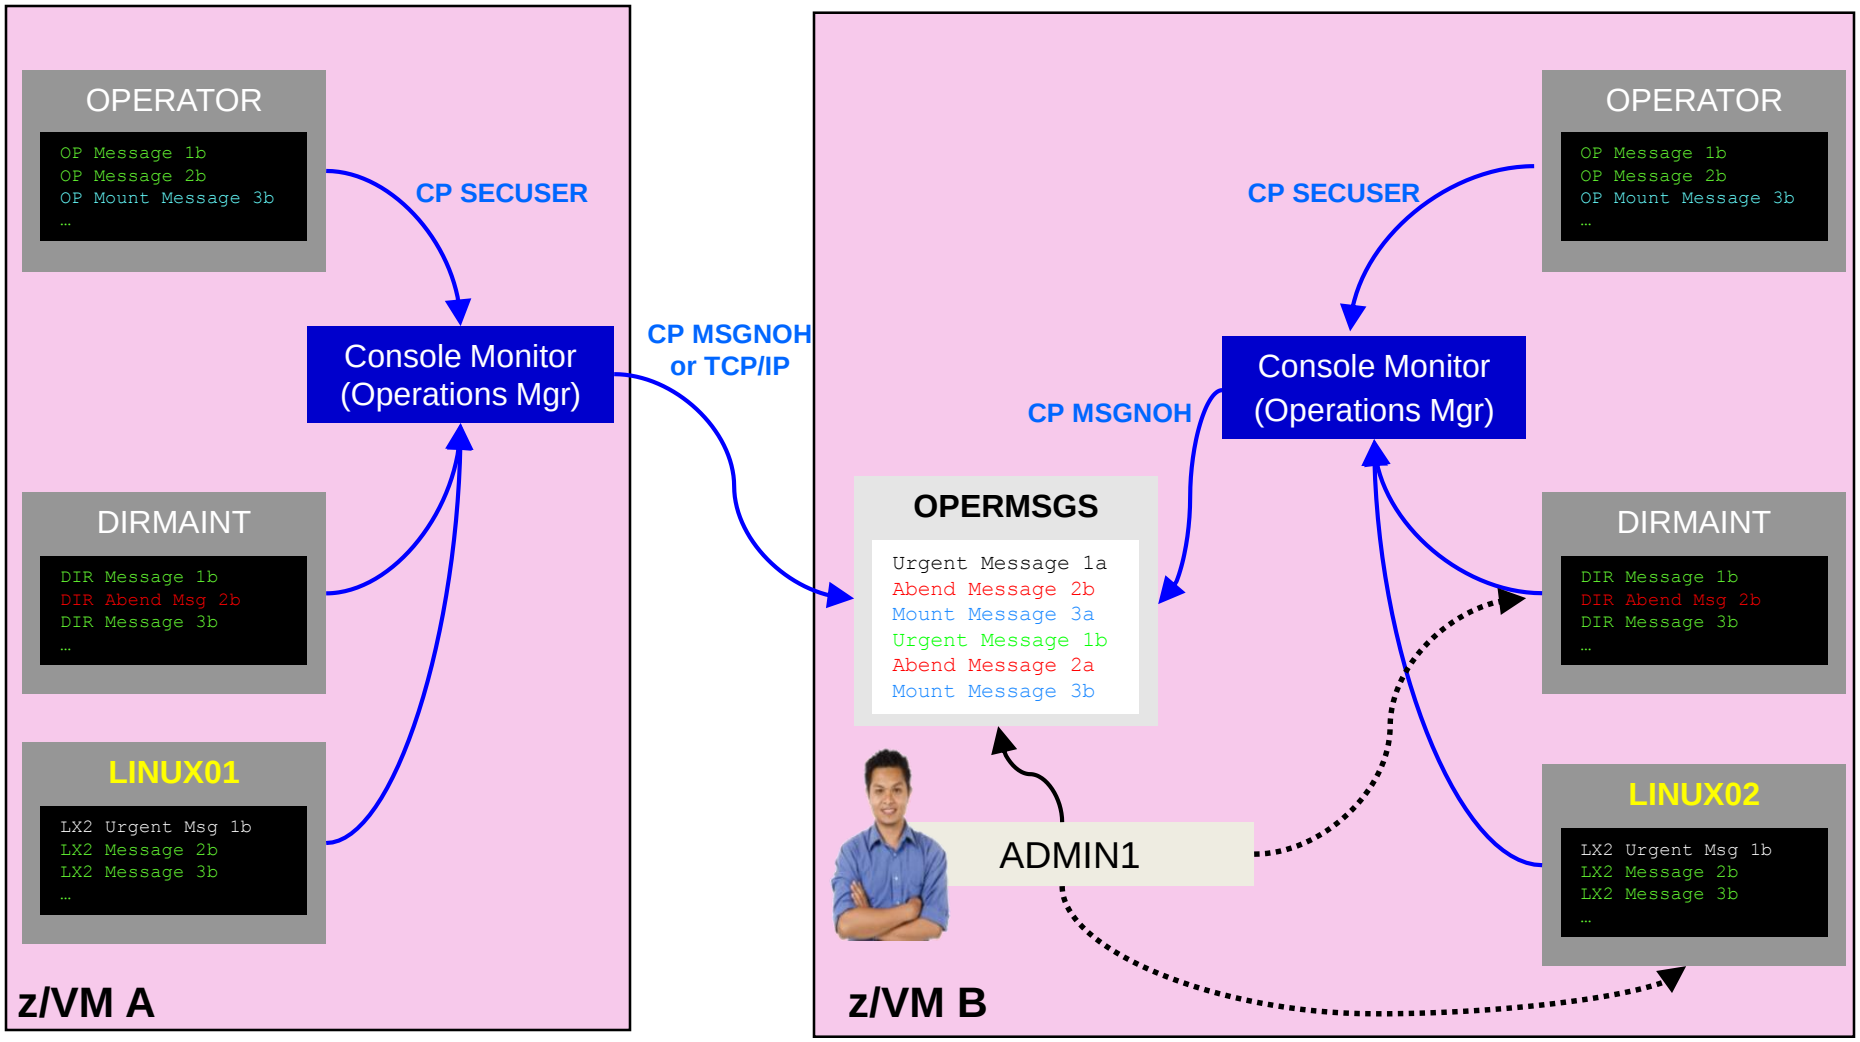

History On-Demand with Persistent Historical Views

This makes it easier to see anomalies, or match spikes. Capturing performance data as a base line is a must:

- General history data – business as usual.
- Detailed raw monitor data prior to and following any major changes.
- Ability to review attributes of a past incident through the enterprise view!
- On-Demand through the Portal or Batch

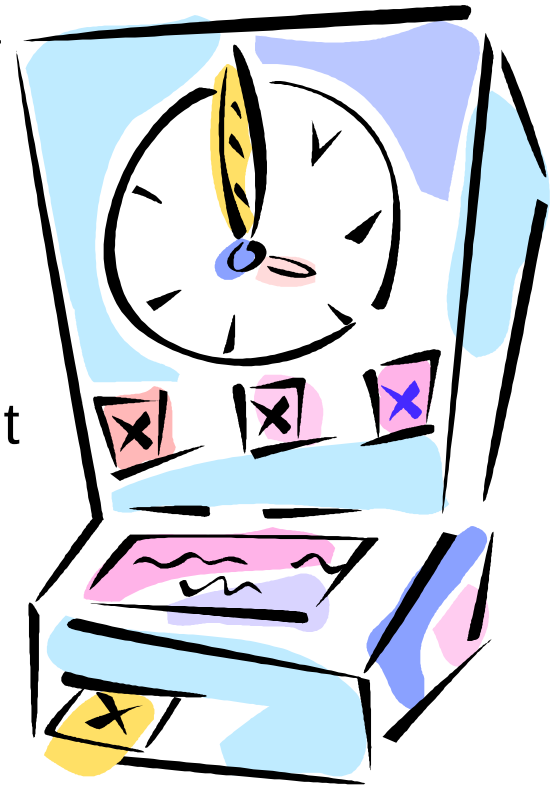

OMEGAMON

On-Demand: Persistent Historical Views

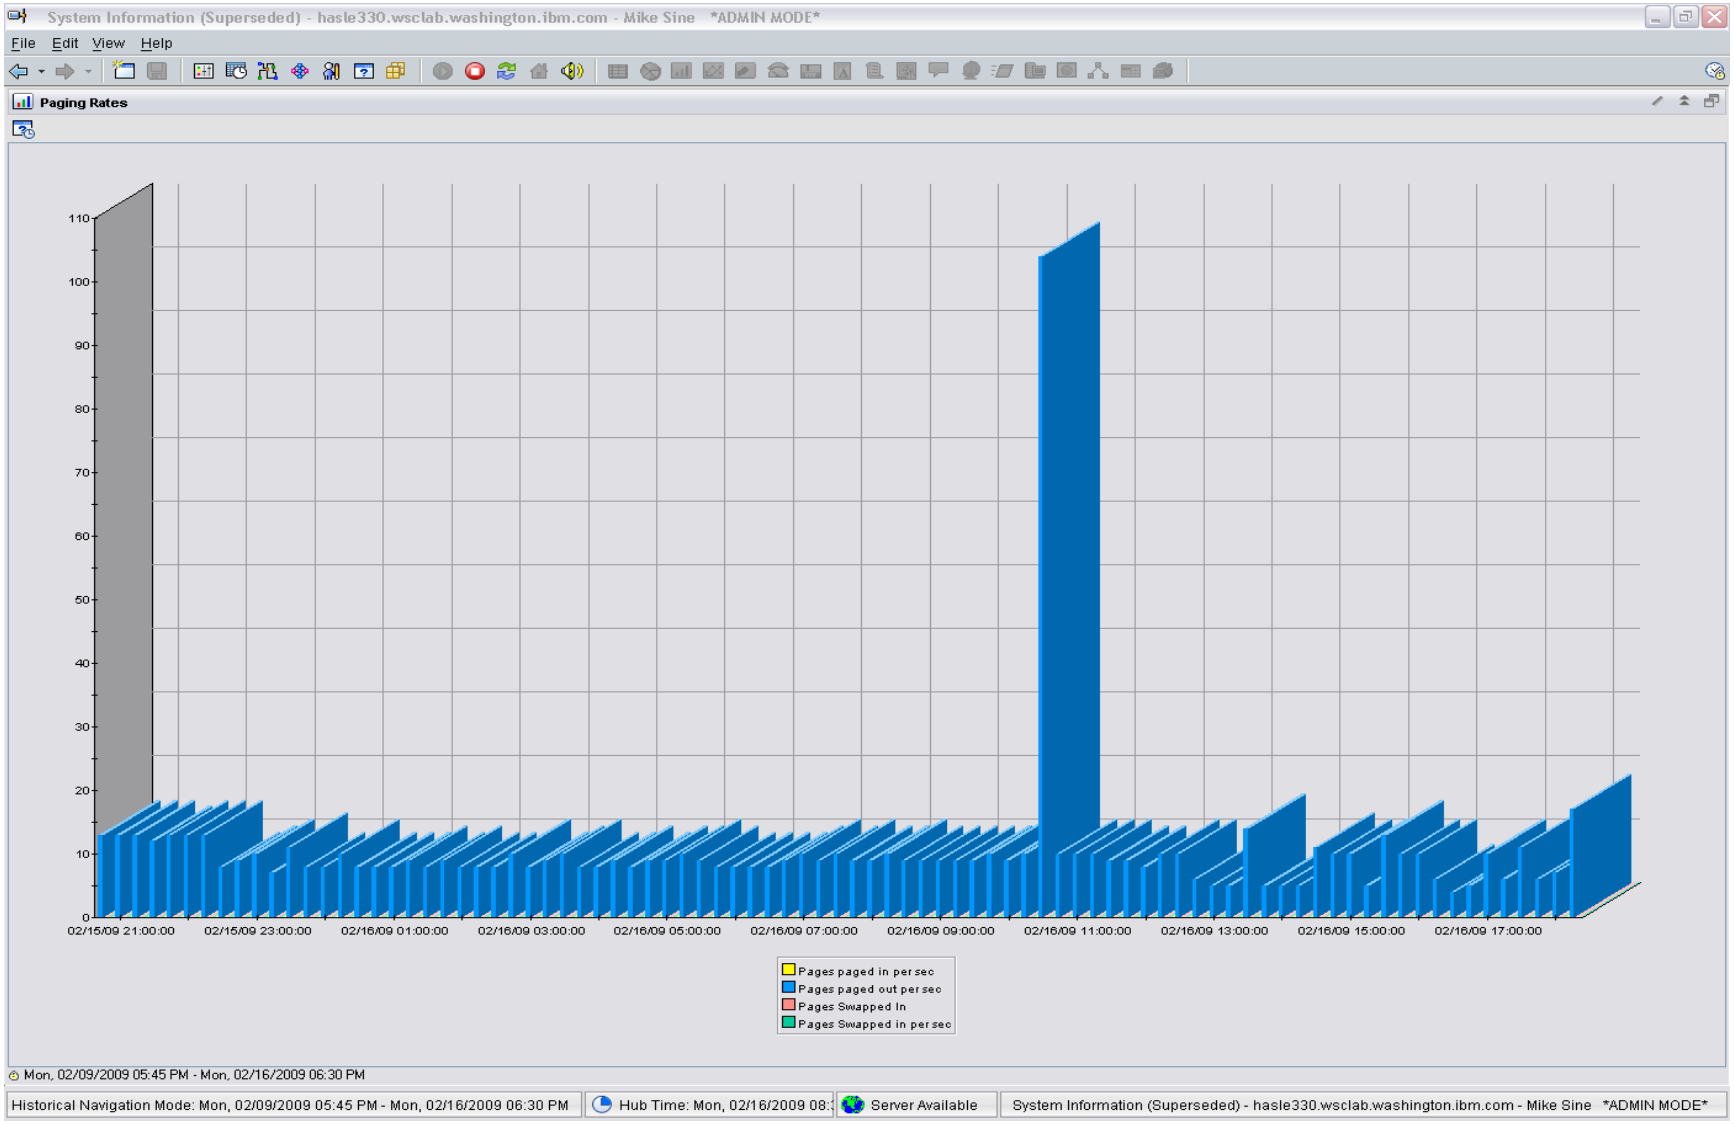

On-Demand: Persistent Historical Views

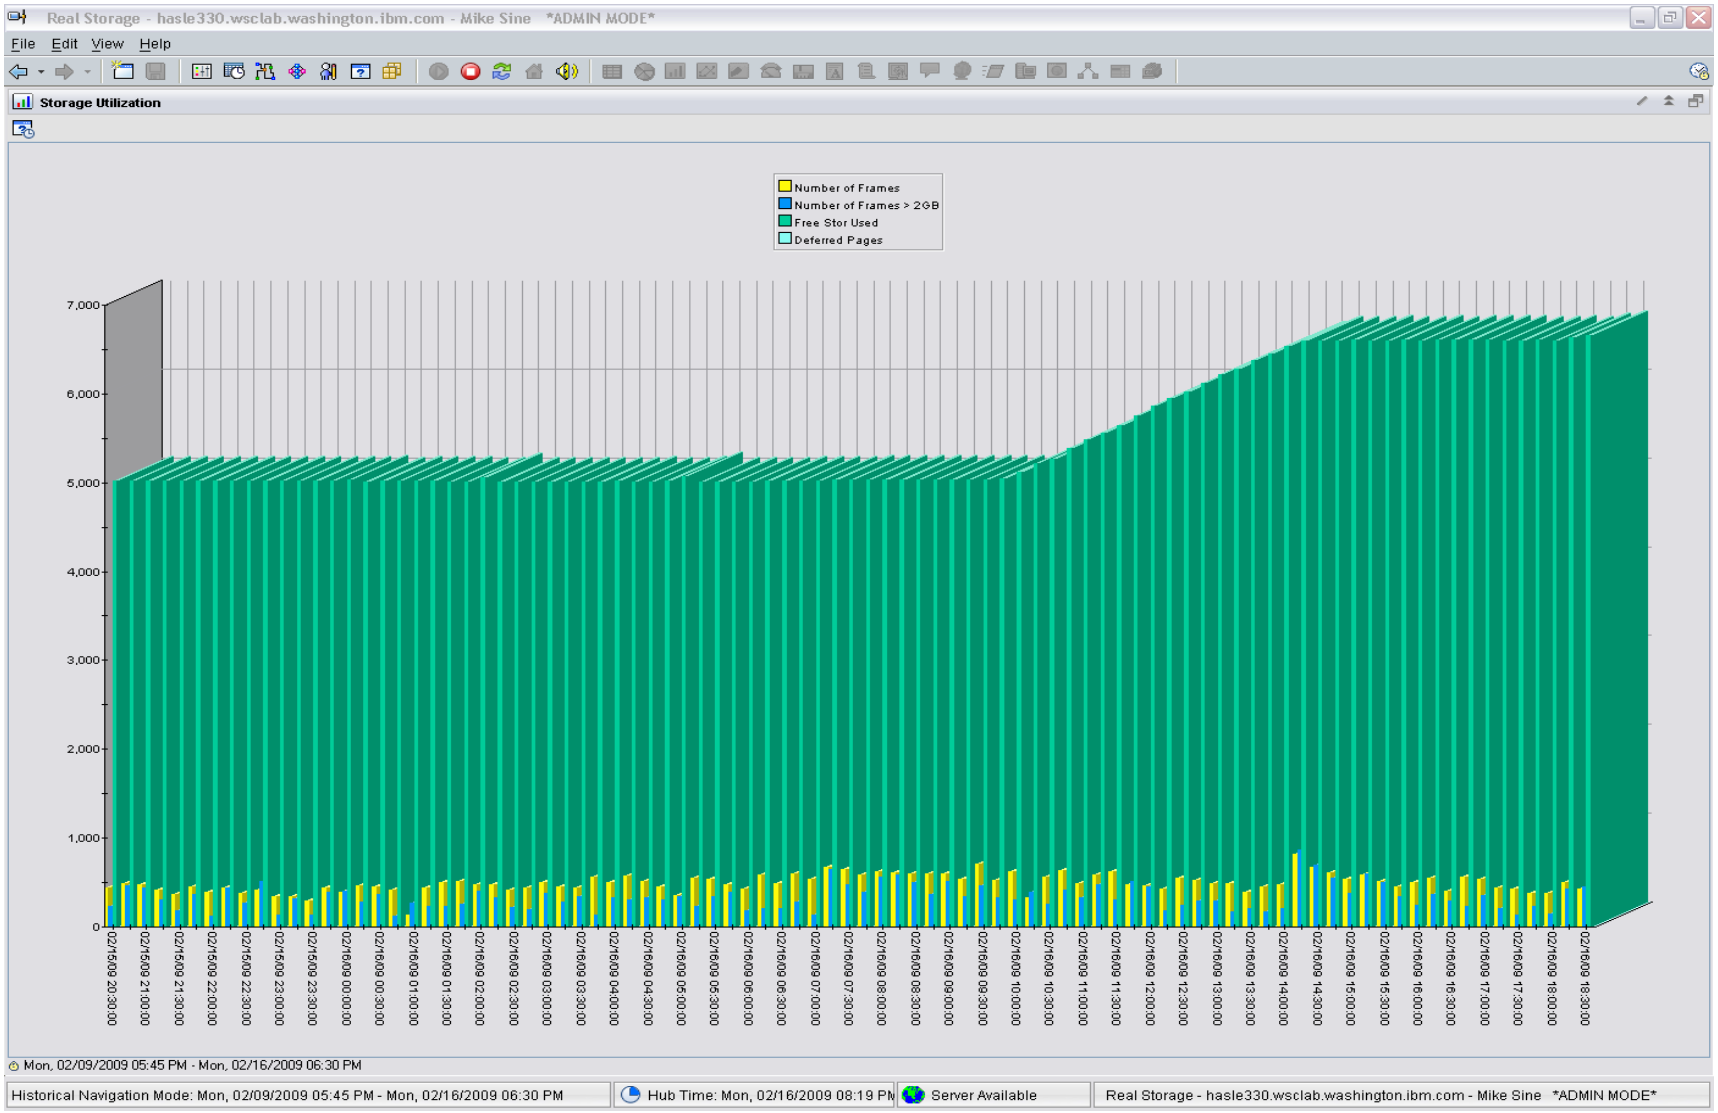

IBM Infrastructure Suite for z/VM and Linux

- Bundle/suite of IBM products
- Announced and available September 2014
- Tools needed to manage the z/VM and Linux on z Systems infrastructure
  - Wave for z/VM
  - OMEGAMON XE on z/VM and Linux
  - Operations Manager for z/VM
  - Backup and Restore Manager for z/VM
    - Order Tape Manager for z/VM separately if plan to back up to tape
  - Spectrum Protect Extended Edition (previously Tivoli Storage Manager)
- Discounted price as a bundle
- Website:
  - http://www.ibm.com/software/products/en/ibm-infrastructure-suite-for-zvm-and-linux
- DeveloperWorks Wiki **– videos of product use/demos**
  - http://ibm.biz/Bd4up3

Summary and Reference Information

- Production systems need
  - Monitoring – operational and performance
  - Automation
  - Backup and recovery
- Real situations need to be addressed
  - Learn from others
- Solutions exist
- Demos available
- Contacts
  - Tracy Dean, tld1@us.ibm.com
  - Mike Sine, sine@us.ibm.com

धन्यवाद

Hindi

多謝

Traditional Chinese

감사합니다

Korean

Спасибо

Russian

Gracias

Spanish

Thank You

English

Obrigado

Brazilian Portuguese

شكرا

Arabic

Danke

German

Grazie

Italian

多谢

Simplified Chinese

Merci

French

நன்றி

Tamil

ありがとうございました

Japanese

ขอบคุณ

Thai