

IBM z Systems

Customer Experiences: Managing the z/VM and Linux on z Systems Infrastructure

Tracy Dean
IBM
tld1@us.ibm.com

August 2015




Agenda

- What does “managing” include?
 - What tools or products can you use?
- Customer scenarios
 - Operational monitoring and automation
 - Performance monitoring
 - Backup and recovery
- Demos
- Summary and reference information
- Hands-on Lab
 - Managing a z/VM and Linux on z Systems Environment Using IBM Solutions
 - Tuesday: 04:30 PM - 05:30 PM, Asia 5


The background features a dark blue gradient with abstract geometric shapes in shades of purple and orange. A prominent shape in the lower-left quadrant is a multi-faceted polygon with a color gradient from purple to orange. The overall aesthetic is modern and professional.

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

Three Dimensions of Systems Management



Three Dimensions of Systems Management



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 TCP/IP initialized
- Customize cloning via REXX scripts

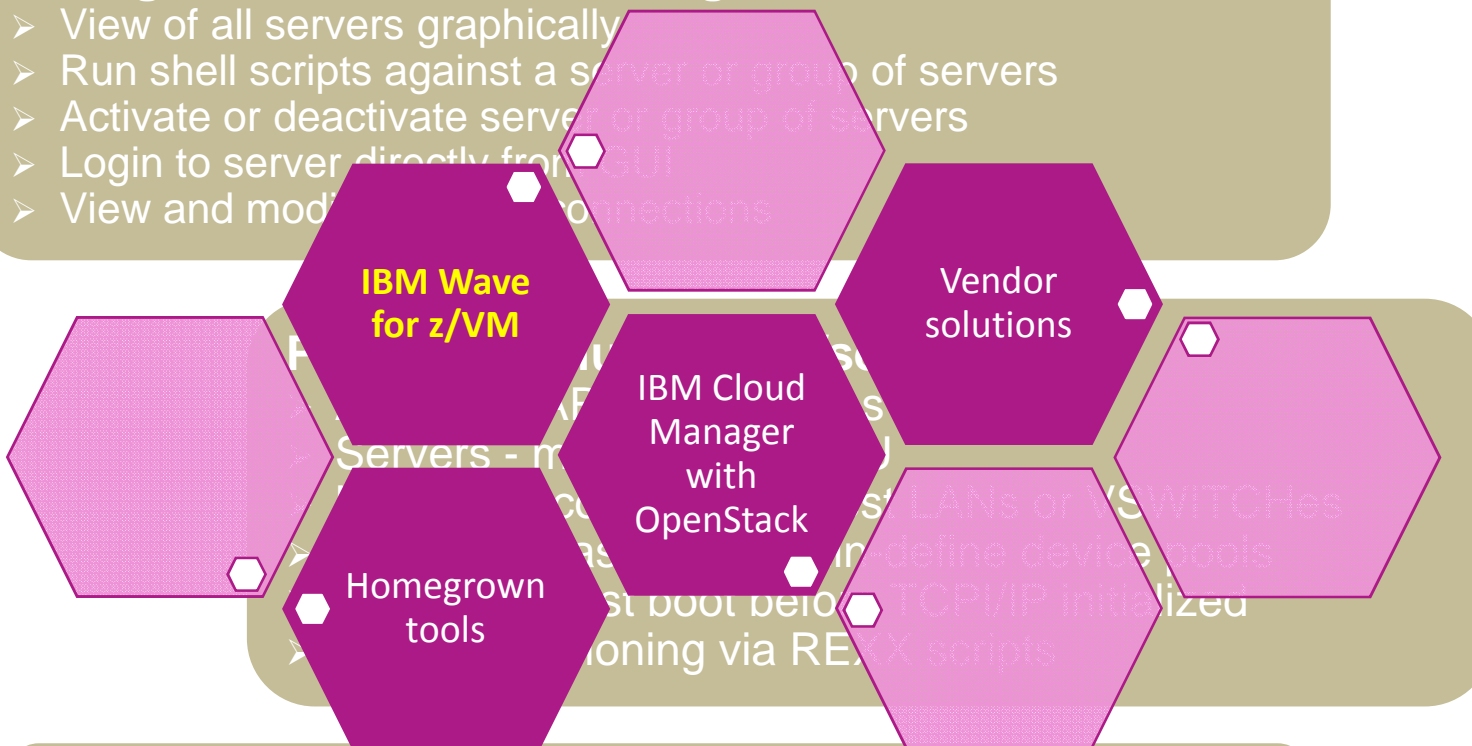
Real time monitoring

- 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



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... utilization
- IBM OMEGAMON XE on z/VM and Linux
- Vendor solutions
- IBM Performance Toolkit for z/VM
- Homegrown tools

Use historical c...

- 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

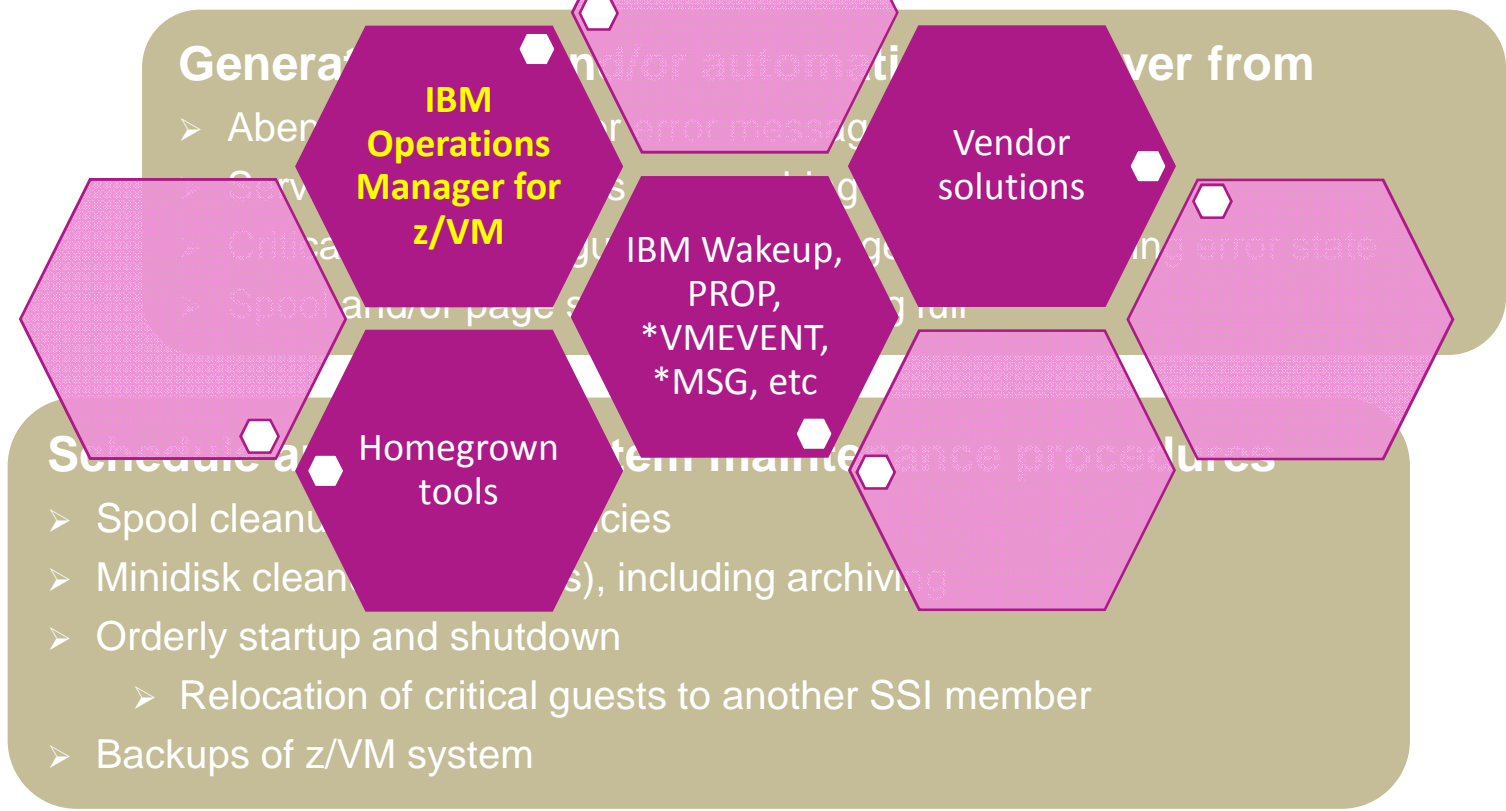
Schedule automated system maintenance procedures

- 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 on a central console of alerts
- System programmers debugging a problem on a guest or service machine



Backup and Recovery of z/VM and Linux

Image level backup of z/VM

- Operating system

File level backup of z/VM data

- Directory information
- Configuration files
- Log files
- Tools – REXX EXECs, automation scripts, etc.

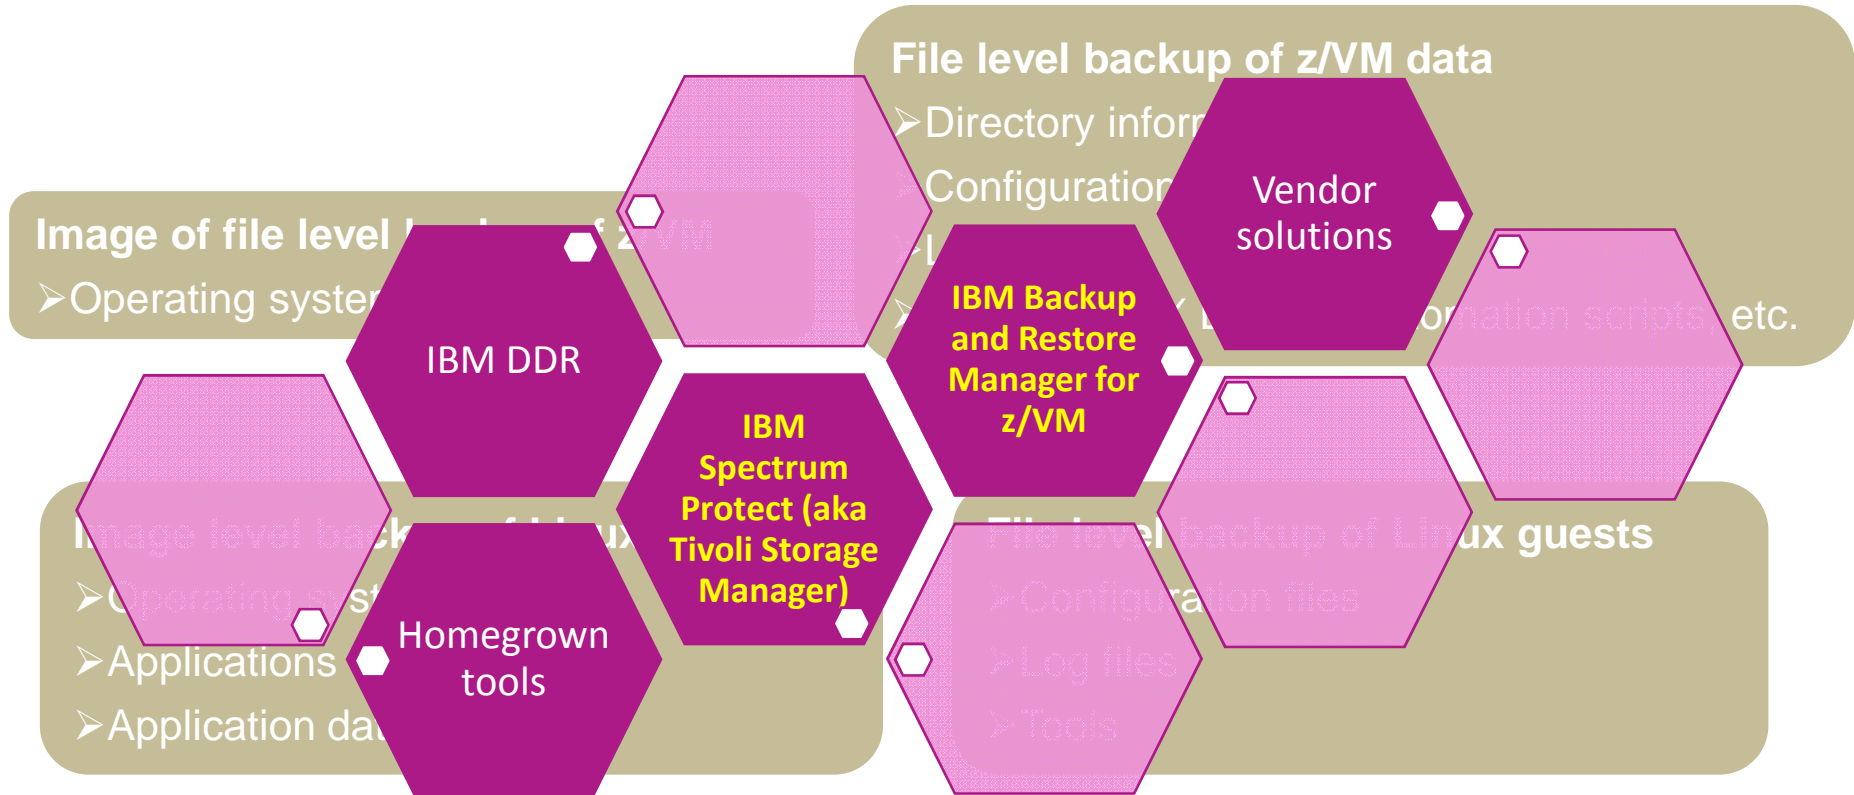
Image level backup of Linux guests

- 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

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
- **EREP disk full**
- **Accounting disk full**

Initial solution

None


- Took photo of HMC with smartphone
- Show IBM and ask for help
- No knowledge of impact of the message

Final solution

Monitoring tool

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

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

Initial solution

IPL cold start and hope for the best

Or

IPL cold start and dig through dump data

Final solution

Console monitoring tool

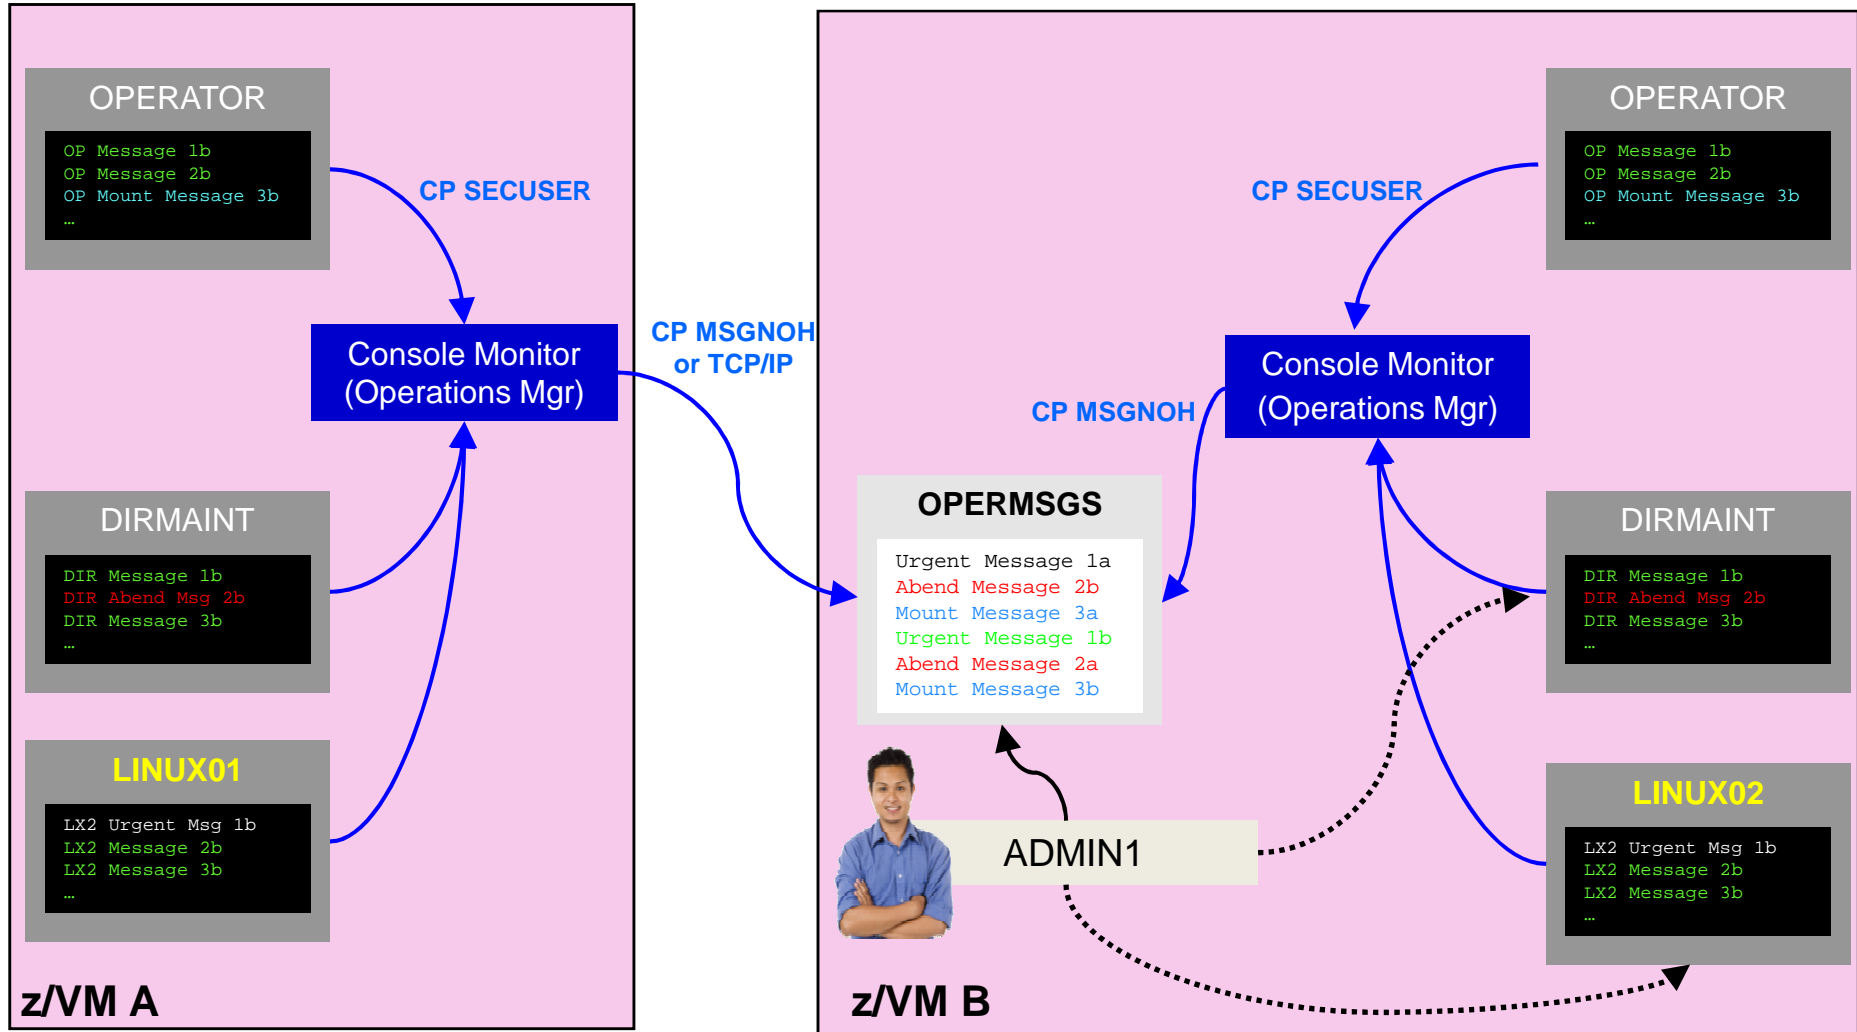
IPL cold start and review console data written in one log file on disk

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



Spool and Page Space Full



The Situation:

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

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 Constraint Analysis



The Situation:

- Performance monitor says **CPU utilization** for system is **high**
- Is that a **problem**?
- What's the **impact** on the **applications**?



Initial solution

Guess

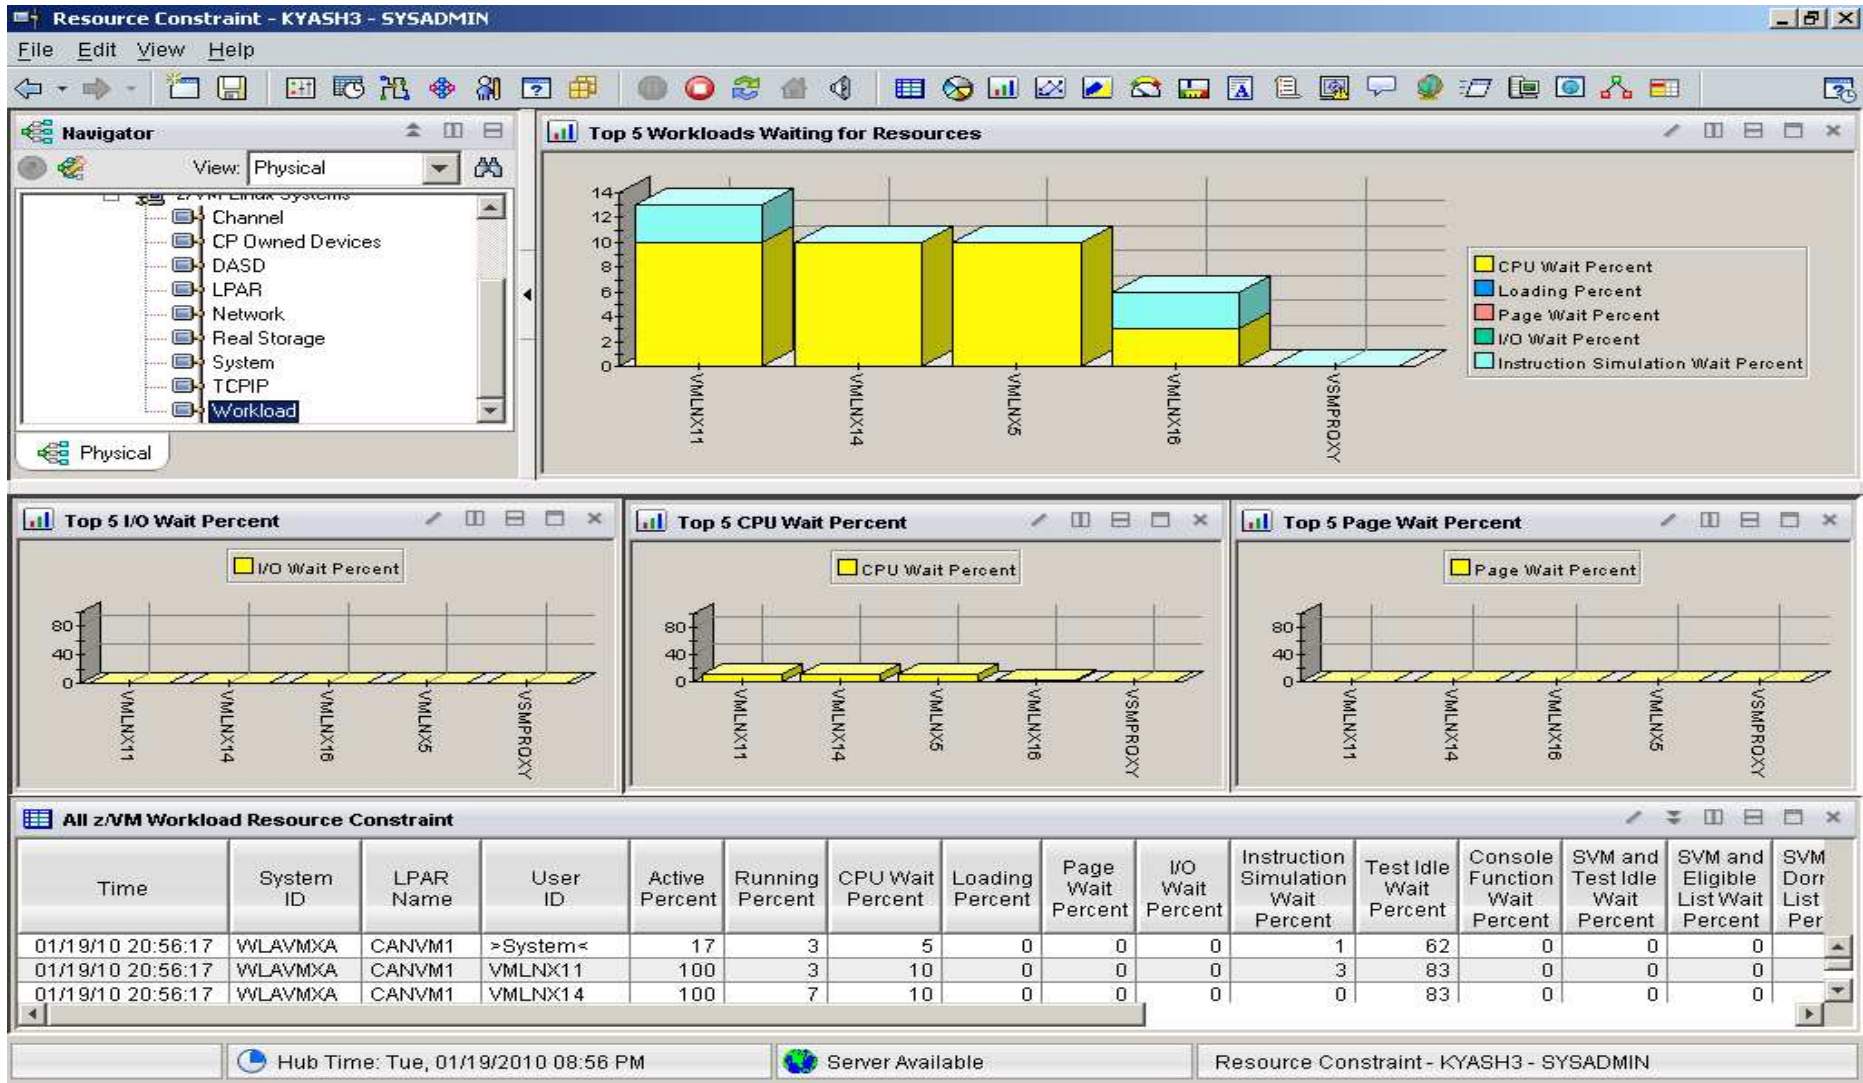
- Wait and see if anyone complains
- Dig around and look at several other metrics
- Move or stop guests
- Add more hardware if consistently high

Final solution


Monitoring tool

- Look at resource constraint analysis
- Which guests are waiting on which resources
- Is a critical guest impacted?
- Efficient use of resources?

Resource Constraint Analysis



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**

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

Provide Service Offering to Agencies



Hmmm..

Challenges – *staff-up* or *tool up*?

Very limited in-house z/VM expertise – must leverage existing z Systems skills



Yep it's true

No budget to staff a new z/VM team, or train potential new customers on z/VM



Yep again

Robust disaster recovery solution already in place; this new service must be 100% compatible

z Systems automation principles in place – Automate, Automate, Automate as much as possible



Uh, deal with it!

Monitoring and Management

*First take care of the simple,
yet important stuff*

Monitor and manage the infrastructure
How?



Scheduler

Integrated, powerful, rule based – very important!



True system event monitoring

No polling, no heart-beats, no agents



Take actions

Commands, REXX scripts, chain other take actions (automation de-dupe!)



Notification

Today e-mail, next up: SNMP alerts to inform Security Information and Event Management (SIEM) system

Backup and Recovery

Backing up this new environment

Backup and Restore Requirements



Must integrate into existing z Systems environment

VTS grid across two datacenters



Native full and incremental backup & restore

z/VM file level backup & restore a must!



Flexible selection criteria

Define once and let naming standards take care of growth (new Linux servers, new DASD, new minidisks)



Integrate with Automation


Schedule, monitor, notification



Easy!

So easy to use that this guy could do it


Provide Service Offering to Agencies



Challenges – *staff-up* or *tool up*?

Tool up!

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

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

- Hypervisor
- Linux operating system
- Application

Why Was an Application Running Slow



The Situation:

- Application owner asks z/OS programmer why application was running **slow** yesterday
- Application owner asks z/OS programmer to research the problem

Drill down to each layer within a specified time window

Initial Solution

Look at performance of the Linux guest

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

Final solution

Performance monitoring solution for all layers

- Hypervisor
- Linux operating system
- Application

Why Was an Application Running Slow

z/VM

z/VM Linux Systems

- Channel
- CP Owned Devices
- DASD
- LPAR
- Network
- SSI Cluster
- Real Storage
- System
- TCP/IP
- Workload

Time	System ID	LPAR Name	Virtual CPUs	Total CPU	User ID	U
05/13/15 13:50:08	ZVMV6R30	ROSPA	1	0.30	ESMTS108	
05/13/15 13:32:35	ZVMV6R30	ROSPA	1	0.80	SLESB100	
05/13/15 13:32:34	ZVMV6R30	ROSPA	4	1.40	SLESB103	
					SLESB104	
					SLESB110	
					SLESB113	1

Linux on z Systems

Linux OS

- Capacity Usage Information
- Disk Usage
- File Information
- Network
- Process
- System Information
- Users
- Agent Management Services
- MQSERIES - QM_has1103
- WebSphere Agent - Primary
- has1104
- has1105
- has1106
- has1107
- has1108
- has1110

Process Command Name	Process ID	Process Parent ID	Cumulative Process User CPU (Percent)	Total Size (Pages)	Resident Set Size (Pages)	8
cupsd	3436	1	0.00	2306	674	435
db2dasrrm	8910	1	0.00	15124	1630	1234
db2fmc	8614	1	0.24	9787	2368	1761

Process Information Detail

- DB2 - db2inst1.has1103:UD
 - Customized SQLs
 - Application
 - Database
 - System Overview
 - UDB_Status_Warning
 - Locking Conflict
 - Buffer Pool Activity
 - Table Space

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

DB2 UDB Agent

DB2 Status	Node Name	DB2 S
Inactive/Busy	db2inst1.has1103:UD	

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
...

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

The background features a complex geometric pattern of overlapping triangles and polygons. The color palette is dominated by dark blues and purples, with a prominent triangular shape in the lower-left quadrant transitioning from red to orange to yellow. The text 'The Solutions' is positioned in the upper-right area of the image.

The Solutions

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
 - Tivoli Storage Manager Extended Edition (now Spectrum Protect)
- Discounted price as a bundle
- Website:
 - <http://www.ibm.com/software/products/en/ibm-infrastructure-suite-for-zvm-and-linux>
- DeveloperWorks Wiki
 - https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/W9b511b099ded_4e32_abfb_ed8ce4da5b17


Operations Manager for z/VM

Increase productivity

- Authorized users to view and interact with monitored virtual machines without logging onto them
- Multiple users view/interact with a virtual machine simultaneously

Improve system availability

- Monitor virtual machines and processes
- Take automated actions based on console messages
- Reduce problems due to operator error



Automation

- Routine activities done more effectively with minimal operations staff
- Schedule tasks to occur on a regular basis

Integration

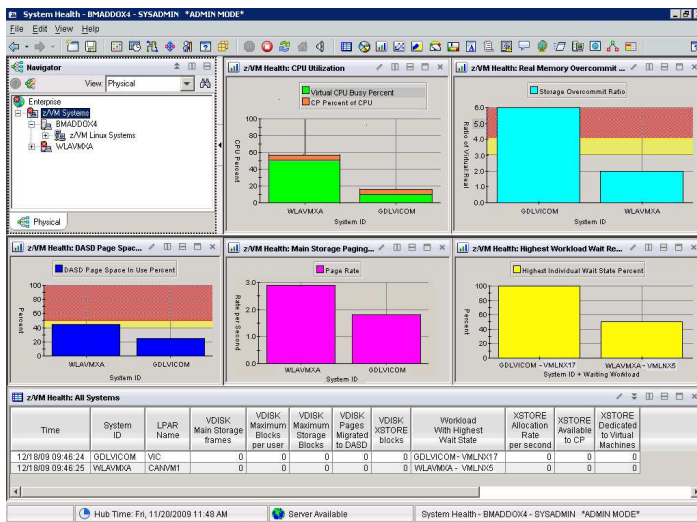
- Fulfill take action requests from performance monitoring products (e.g. OMEGAMON XE on z/VM and Linux)
- Send alerts to email, central event management systems (e.g. Netcool/OMNIBus), etc.

OMEGAMON XE on z/VM and Linux

Bringing z/VM and Linux monitoring into the Enterprise View

Increased Performance & Availability

Enterprise-ready
cloud monitoring

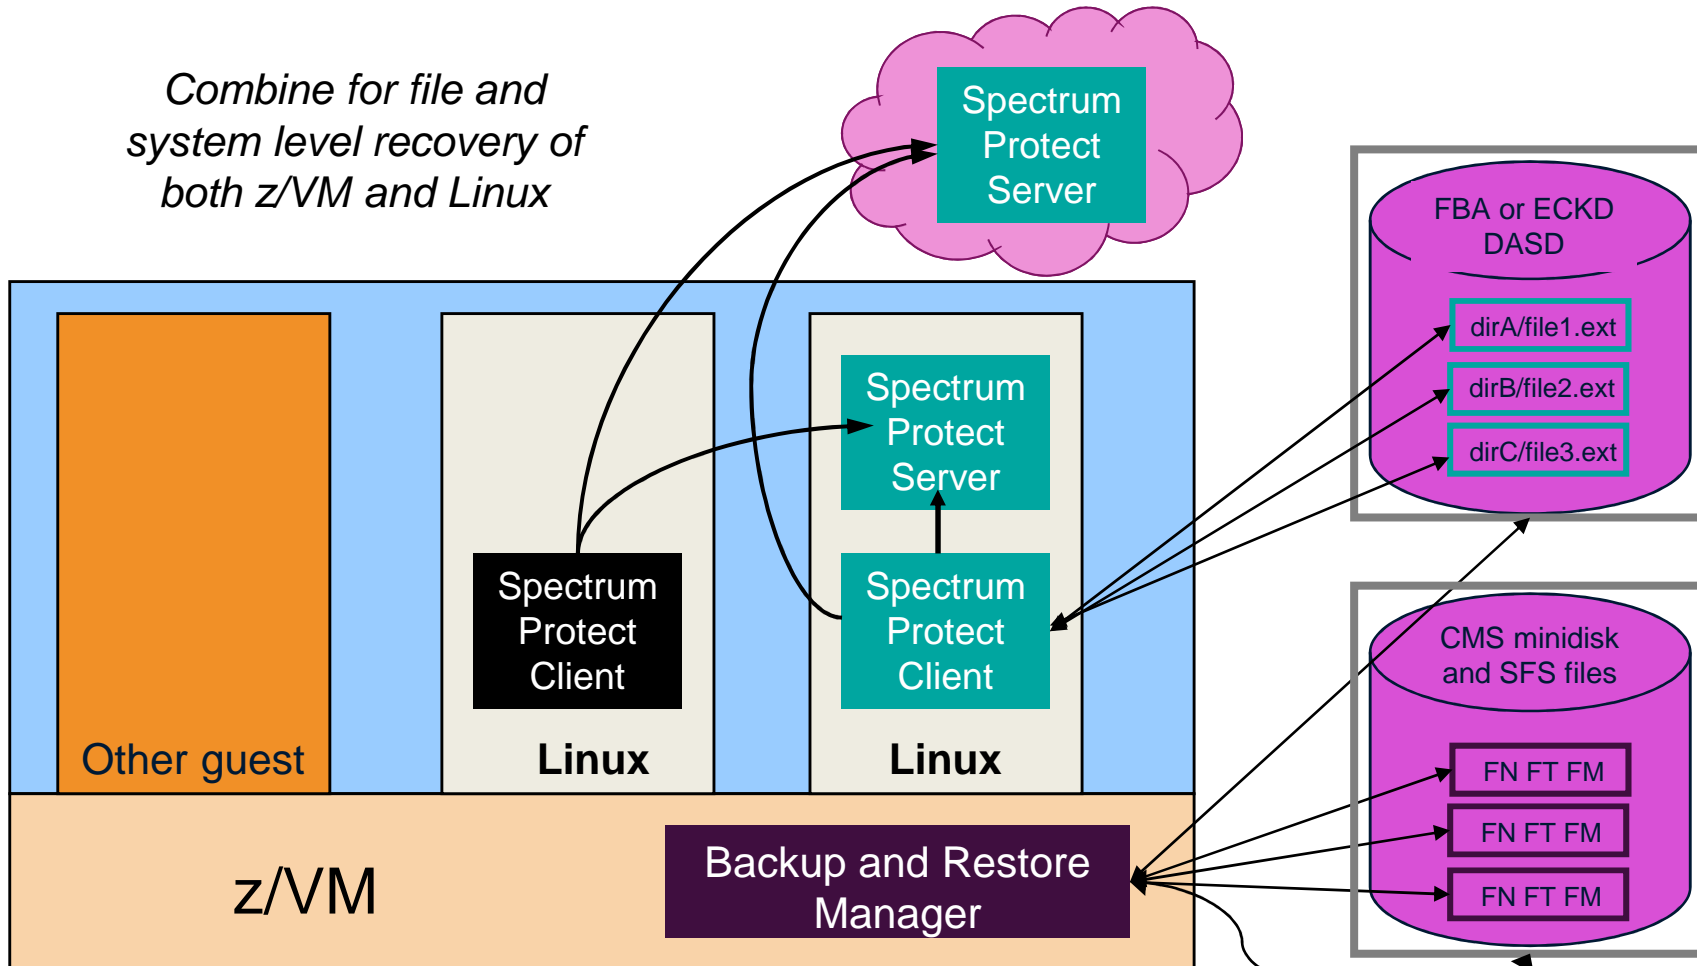


- Provides **insight** into the health and performance of z/VM and Linux
- Rich collections of attributes monitor thresholds for z/VM and Linux best practices
- **Reflex automation** provides timely resolution and/or notification
- Lightweight **visibility** to the z/VM hypervisor, Linux OS, and Linux Log data in one tool
- **Deep integration** with IBM Tivoli Monitoring and OMEGAMON family, bringing z/VM and Linux data to the Enterprise view (cross platform workflow management)
- Persistent **historical views** allows management of real and virtual resources across peak periods and downtimes for **clear view** of resource usage and constraints

Backup and Restore Manager for z/VM and Spectrum Protect

*Using Backup and Restore Manager with Spectrum Protect
(formerly Tivoli Storage Manager)*

Combine for file and system level recovery of both z/VM and Linux



Tape Manager for z/VM

- Manage tapes
 - Define tapes in a catalog
 - Free or used
 - Retention/expiration information
 - ATL/VTs or manual mount
 - Data Security Erase
 - Group tapes together into pools
 - Ownership and access control
 - Media type
- Manage devices
 - ATL/VTs
 - Manual mount
 - Tape grid/cluster
- Share devices with other systems
- Support for multiple vendors
 - IBM
 - Oracle STK
 - EMC

- **Manage mount requests**
- **Volume specific and scratch requests**
 - Standard Label
 - Non-Label
 - Bypass Label Processing

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
 - https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/W9b511b099ded_4e32_abfb_ed8ce4da5b17
- Contact
 - Tracy Dean, tld1@us.ibm.com

The background features a complex geometric pattern of overlapping triangles and polygons. The color palette is primarily dark blue and purple, with a prominent orange and red triangle in the lower-left quadrant. The overall effect is a modern, abstract digital aesthetic.

Live Demos

Automation Demos Available

1. View consoles of Linux guests, Linux syslog data, and CMS user IDs or service machines
2. Send an e-mail based on a console message
3. **Send an alert to Netcool/OMNIBus based on a console message, hold and unhold messages**
 - a. Using POSTZMSG interface to Netcool/OMNIBus
 - b. **Using SNMP interface to Netcool/OMNIBus**
4. **Send a message or email if spool approaches full**
 - a. Send a message if spool usage is too high on any member of an SSI Cluster – see how spool files appear in SSI
 - b. **Send an email if spool usage is too high on a single system**
5. View and clean up spool files
6. Automated spool cleanup
7. Archiving DIRMAINT's log files when disk gets full
8. Process a file of test messages as a console
9. Process Linux syslog data as a console
10. Create a central operations console on one z/VM system
11. Create a central operations console across multiple z/VM systems
 - a. When the systems are in an SSI cluster
 - b. When the systems are not in an SSI cluster
12. Integration with OMEGAMON XE on z/VM and Linux - take action based on CPU usage of Linux guest
13. Monitor service machines for logoff – and autolog them
14. Send an email if page space approaches full
15. Monitor SSI connectivity between 2 cluster members
16. Suppress passwords on Linux consoles
17. Autolog a Linux Guest and Send Message if Doesn't Start Successfully

Scenario 3b: Send an Alert to OMNIbus – Using SNMP

- Watch all monitored consoles for an error message that includes the word “read-only”
- If this word appears on a console
 - Change the message to red and hold it
 - Send an alert to OMNIbus, using SNMPTRAP command on z/VM
 - Automatically unhold the message after 4 minutes
- Dynamically include in the alert
 - IP address of the z/VM system where the error occurred
 - User ID that received the error message
 - Text of the error message

Scenario 4b: Send an Email if Spool Usage is Too High

- Operations Manager monitors the spool usage (percent full)
 - Demo monitor requires spool to only be 5% full or higher
 - Usage exceeds the specified limit
 - Automatically send an e-mail to someone who can evaluate and take action
 - For demo purposes
 - Spool monitor is currently suspended
 - Dynamically resume (re-activate) the spool monitor
 - Suspend (de-activate) the spool monitor when complete

धन्यवाद

Hindi

多謝

Traditional Chinese

감사합니다

Korean

Спасибо

Russian

Gracias

Spanish

شكراً

Arabic

Thank
You

English

Obrigado

Brazilian Portuguese

Grazie

Italian

多谢

Simplified Chinese

Danke
German

Merci

French



நன்றி

Tamil

ありがとうございました

Japanese

ขอบคุณ

Thai