

IBM Z

Customer Experiences:

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

Tracy Dean
IBM Offering Manager
tld1@us.ibm.com

June 2020



Agenda

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



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

Administration and Provisioning

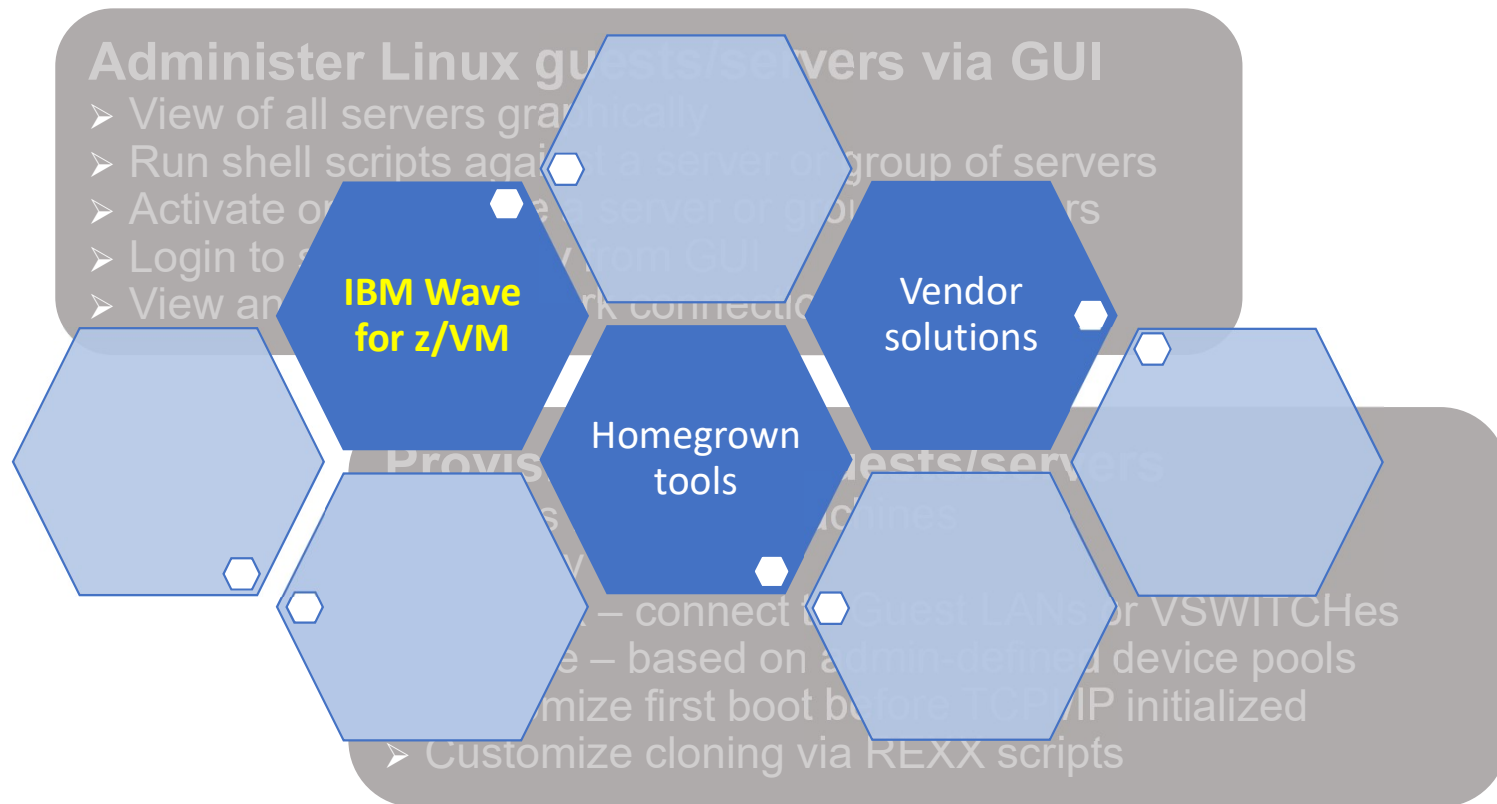
Administer Linux guests/servers via GUI

- 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

Administration and Provisioning



Performance Monitoring and Automation

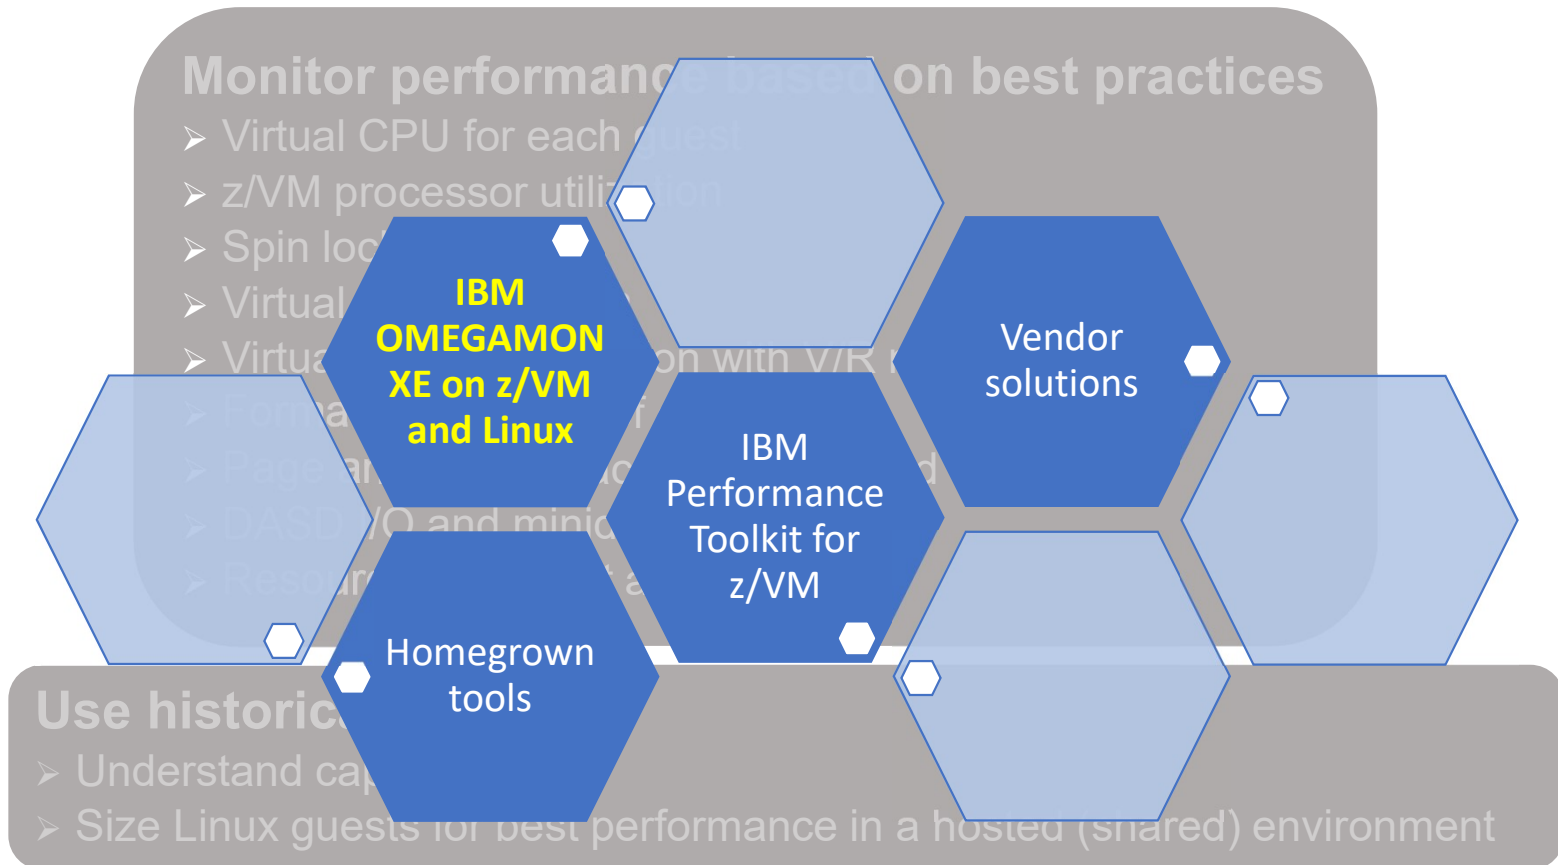
Monitor performance based on best practices

- 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



Operational Monitoring and Automation

Console monitoring and viewing, user and system events

- Operations staff monitoring a central console of alerts
- System programmers debugging a problem
- Log data available for audits or future reference

Gather Data

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

React

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

Prevent

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 audit and future reference

**IBM
Operations
Manager for
z/VM**

IBM Wakeup,
PROP,
*VMEVENT,
*MSG, etc

Vendor
solutions

**Homegrown
tools**

Schedule automated jobs

- Spool cleanup based on rules
- Minidisk cleanup (from I/O error)
- Orderly startup and shutdown
 - Relocation of critical guest to LPAR or CSI member
- Backups of z/VM system

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.

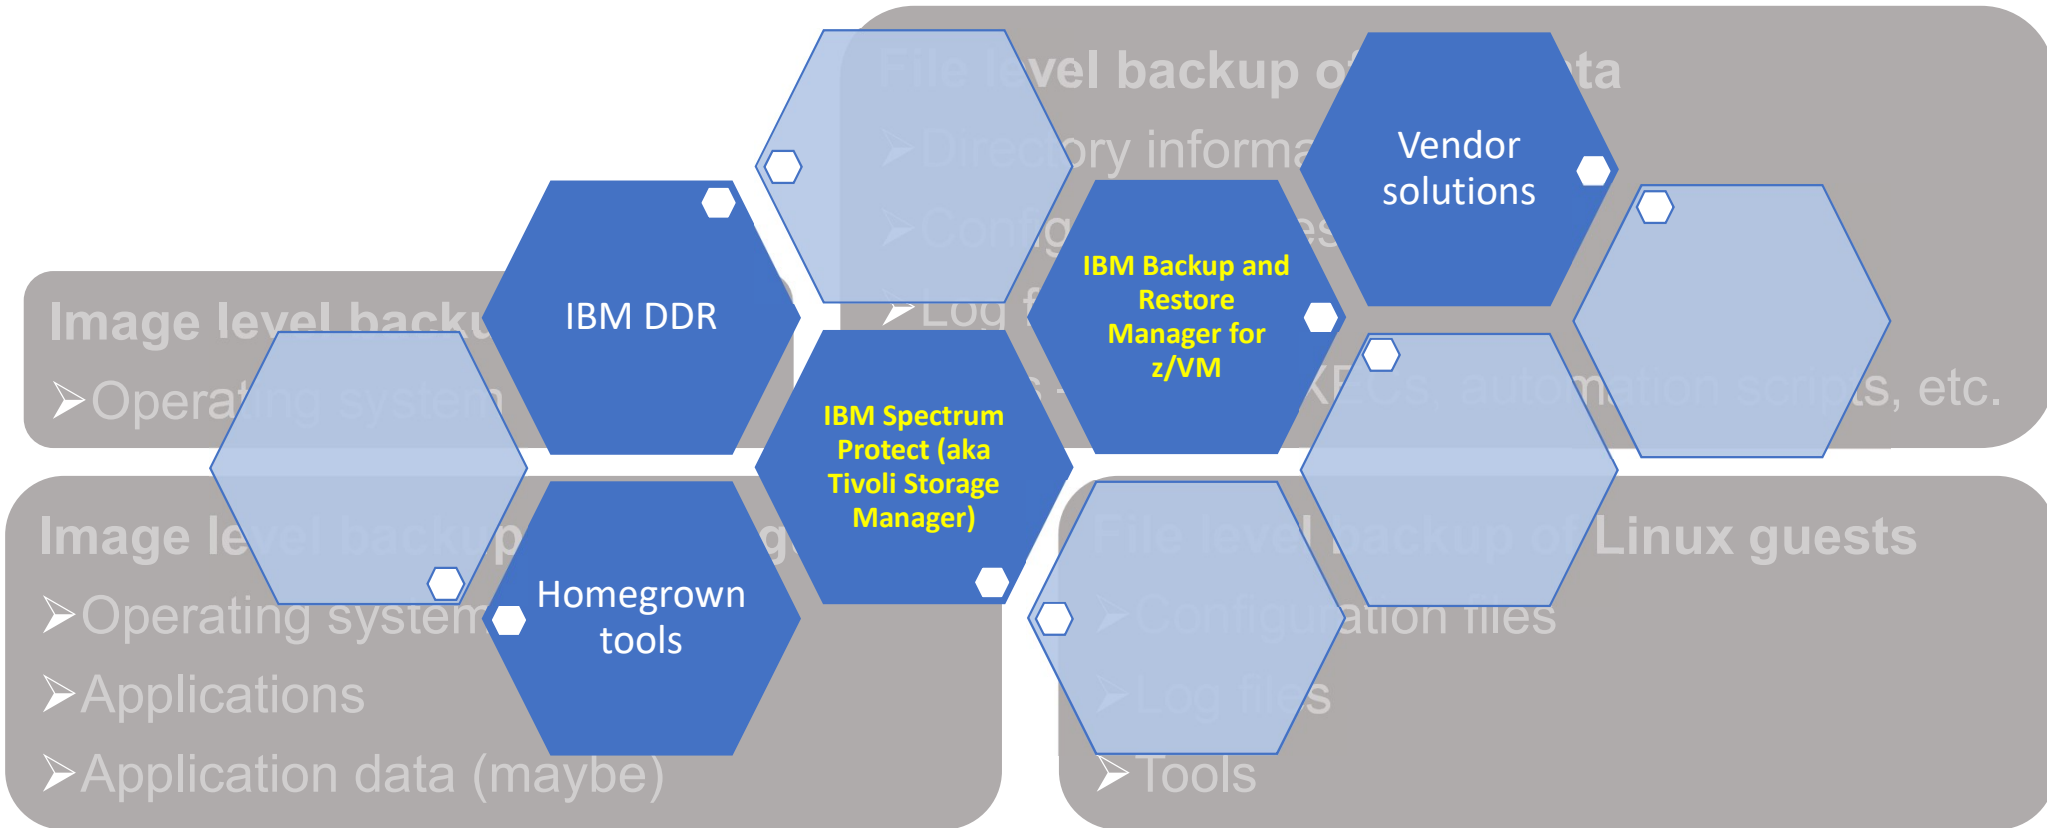
Image level backup of Linux guests

- 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
Administration and Provisioning
Backup and Recovery

Alert If z/VM System Not Responding



The Situation:

- z/VM system just works
- No staff monitoring z/VM consoles
- Central alert system used extensively across enterprise
- Want alerting if z/VM system is not responding

Operations
Manager

Initial solution

None

- Wait for application to fail
- Wait for the phone to ring

Final solution

Alerting via heartbeat checking

- Send a "heartbeat" from z/VM to central alert system every n minutes or seconds
 - Using SNMP
- Central alert system reacts if no heartbeat received as expected

Alert If z/VM System Not Responding

The Situation:

Customers with multiple z/VM systems

- Use this approach for the z/VM systems to monitor each other
- May not get alert if problem with the Z system overall

- Wait for application to fail
- Wait for the phone to ring

- Send a heartbeat from z/VM to central alert system every n minutes or seconds
 - Using SNMP
- Central alert system reacts if no heartbeat received as expected

Recovering From Linux Configuration Problems

The Situation:

- Network configuration change (or other error) prevents Linux server from connecting to the network
- Normal network-based access to the Linux guest (using ssh, for example) is not available



Wave

Initial solution

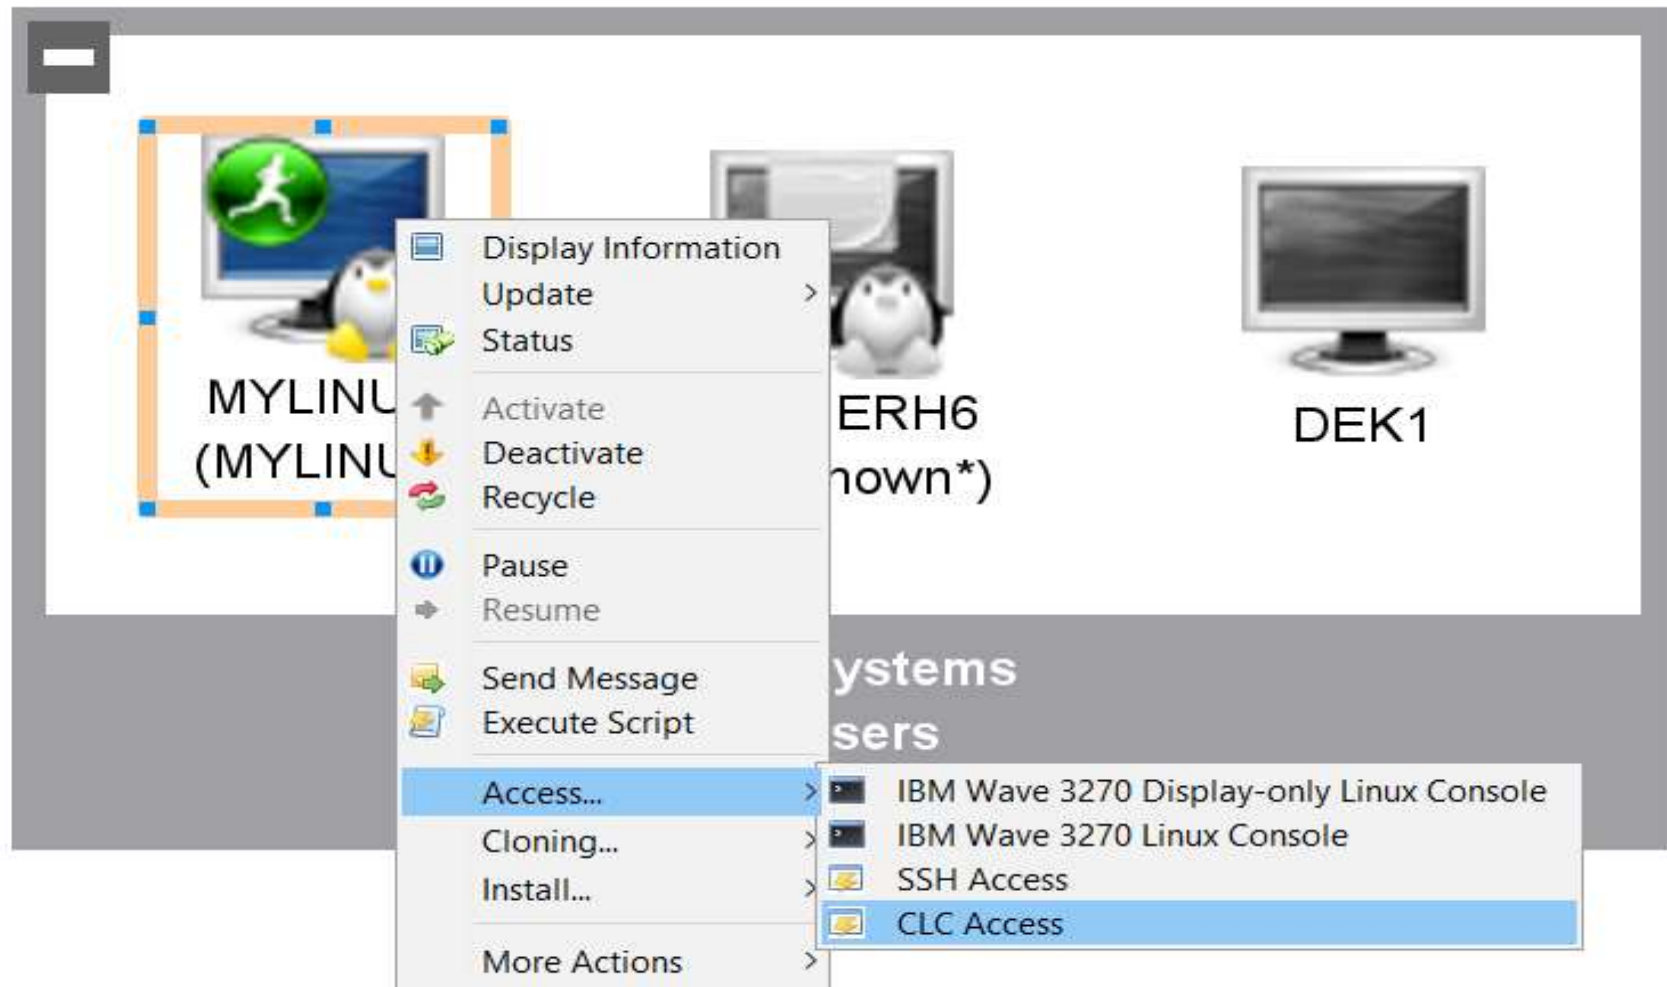
Logon to Linux guest using 3270

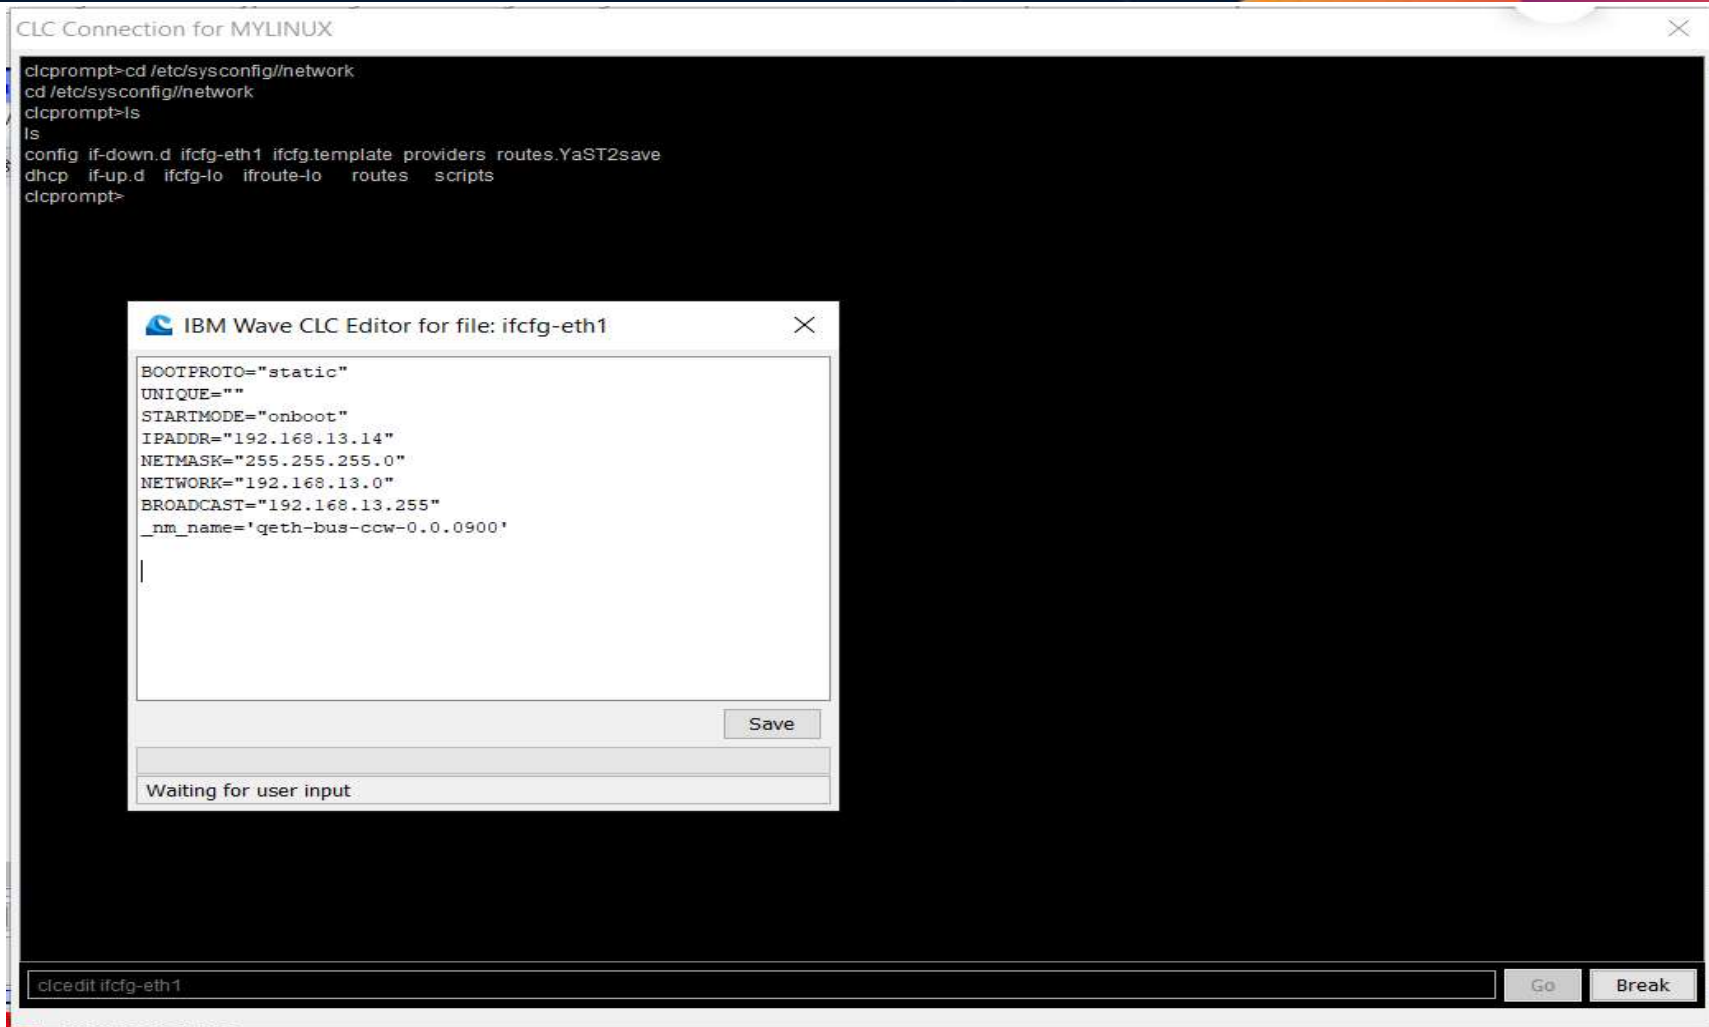
- 3215 line-mode interface
- Edit configuration files using sed
- Painful and error-prone

Final solution

Exploit Communication-Less Connection

- Familiar, scrollable environment
- Simple, easy to use WYSIWYG editor
 - CLCEDIT





The screenshot displays the IBM Wave CLC Editor interface. The main window, titled "CLC Connection for MYLINUX", shows a terminal session with the following commands and output:

```
clcprompt>cd /etc/sysconfig/network
cd /etc/sysconfig/network
clcprompt>ls
ls
config if-down.d ifcfg-eth1 ifcfg.template providers routes.YaST2save
dhcp if-up.d ifcfg-lo ifroute-lo routes scripts
clcprompt>
```

An inset window titled "IBM Wave CLC Editor for file: ifcfg-eth1" shows the configuration for the `ifcfg-eth1` file:

```
BOOTPROTO="static"
UNIQUE=""
STARTMODE="onboot"
IPADDR="192.168.13.14"
NETMASK="255.255.255.0"
NETWORK="192.168.13.0"
BROADCAST="192.168.13.255"
_nm_name='qeth-bus-ccw-0.0.0900'
```

The inset window includes a "Save" button and a status bar that reads "Waiting for user input". The main window's status bar at the bottom shows the command `clcedit ifcfg-eth1` and buttons for "Go" and "Break".

Coordinate Application Shutdown with z/OS

The Situation:

- Database on z/OS
- Application server on Linux on z
- Shutdown of database necessitates shutdown of application server

Operations
Manager

Initial solution

- Manual coordination of shutdown
- Inconvenient for system programmers/operations during non-business hours

Final solution

Console monitoring tools

- System Automation on z/OS sends message to z/VM
- Automation on z/VM triggers application server shutdown
- Automation on z/VM sends message to z/OS when ready

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

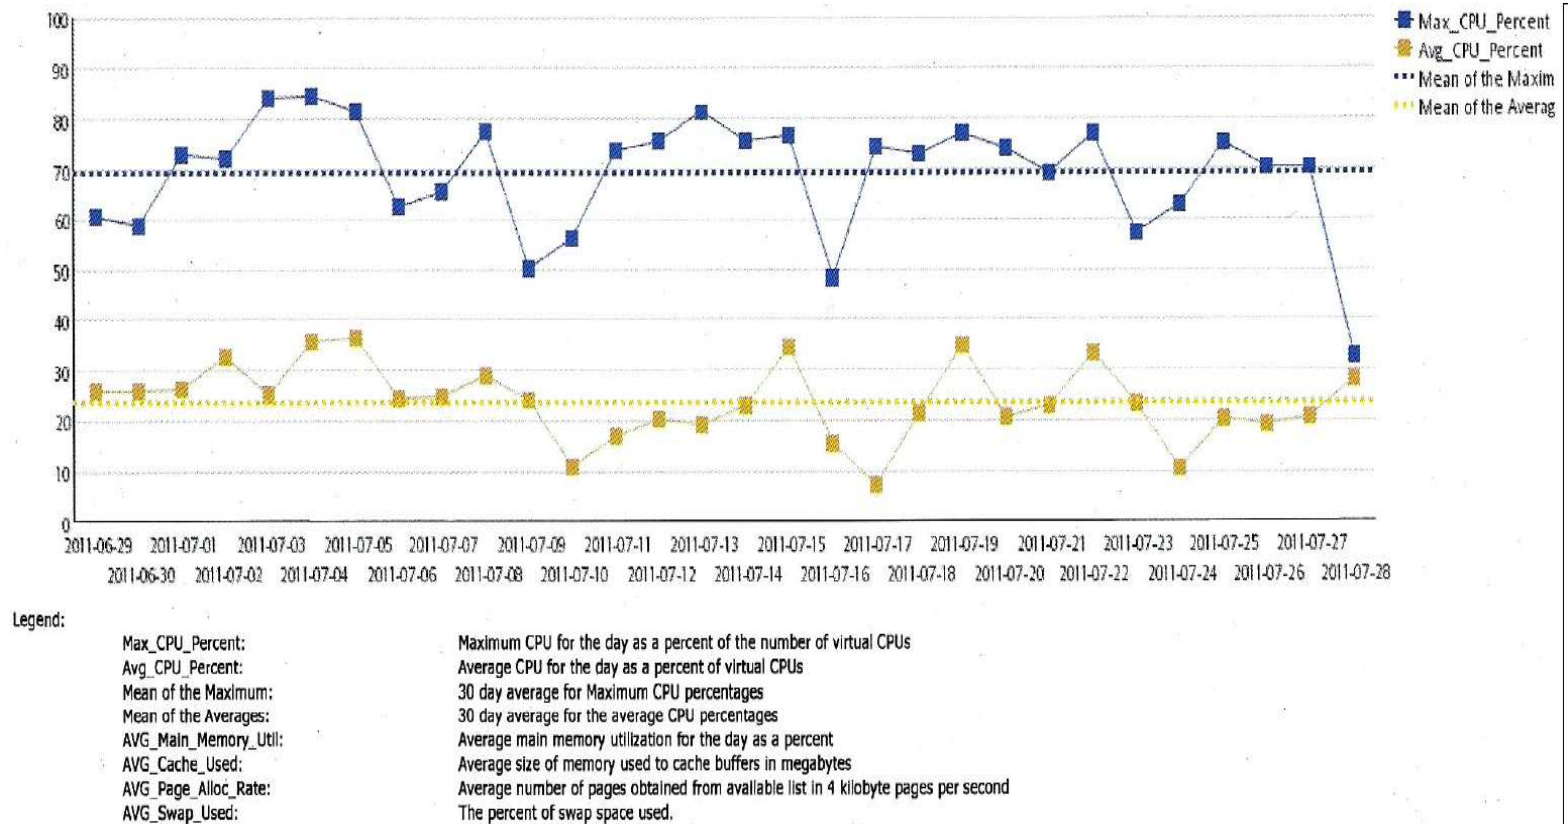
- Tools like TOP and others don't reflect the virtualized environment.
- Users get mixed information, wrong conclusions.
- Misunderstanding between application owners, Linux admins, and system providers

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



Average Linux Memory Breakdown Example



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)
- No obvious impact on applications



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

Send z/VM and Linux Alerts to z/OS

The Situation:

- Extensive **automation** for **alerts** already running on **z/OS**
 - Automation and operations teams trained there
- Want all **mainframe** alerts to be handled this way
- Need **z/VM** and **Linux** on Z alerts **included**



Operations
Manager

Initial solution

None

- z/VM and Linux alerts sent via email or to central console only
- No alerts sent to enterprise alert system
- Mainframe operations team not manage and automate z/VM and Linux alerts

Final solution

Monitoring/automation tool

- Trigger alerts for z/VM & Linux events, messages, etc.
- Send via syslog writer to z/OS USS syslog
- Configure USS syslog to send all alerts from z/VM to z/OS syslog
- Enable existing z/OS automation

Long Term Recovery of z/VM Files

The Situation:

- Backups of z/VM volumes done from z/OS
- Retention is only a few **weeks**
- New release of z/VM installed by less experienced staff
- Some local customizations/automation **not preserved**
- Not available via z/OS backups due to short retention



Backup and
Restore
Manager

Initial solution

Options


- Keep volume backups on z/OS for months instead of weeks
 - Additional unnecessary data retained longer on tape
 - Tedious file level recovery
- Re-do all customizations from memory (or with help from IBM)

Final solution

File level backup and recovery

- Weekly full backups and daily incrementals of all z/VM log files and customizations
- Retain months or years without large amounts of tape or DASD

Stopping and Restarting TCPIP



The Situation:

- Want to “bounce” TCPIP server on z/VM on dev/test system
- No access to HMC or system console
- If issue shutdown or FORCE for TCPIP then lose TN3270 access to system

**Operations
Manager**

Initial solution

Find and coordinate with on-site operations staff who have access to system console or HMC

Final solution


Monitoring & automation tool

- Monitor for CP event indicating TCPIP has logged off
- Automatically XAUTOLOG it (after 3-5 seconds)
- Easily bounce TCPIP as needed without relying on operations staff

Shared Monitoring and Automation Across LPARs

The Situation:

- Multiple z/VM LPARs not in same SSI cluster
- Similar monitoring and automation configuration on all LPARs
- Want to share monitoring and automation configuration across LPARs



Operations
Manager

Initial solution

Manual processing

- Common configuration information maintained on one system
- Shared within SSI cluster
- Manually copied and reloaded on other LPARs

Final solution

Automated real-time sharing

- Shared read/only disk across non-SSI members
- Update configuration from single LPAR
- Automatically trigger reload on all SSI and non-SSI systems via IP communications

Sending Security Messages to Analytics



The Situation:

- Enterprise policy of sending security-related messages to analytics platform
- z/VM logon/logoff and RACF login errors only logged in console log of OPERATOR
- Want z/VM security reporting to be “just like other platforms”

**Operations
Manager**

Initial solution

None

- No analytics and alerting of z/VM RACF-related activity

Final solution

Automation tool

- Automatically capture RACF logon/logoff messages on OPERATOR console
- Send to analytics platform in key/value pair format

Capturing Linux Log Data

The Situation:

- z/VM console data being captured
- No Linux console data
- Linux log data stored locally on each guest
- Linux server crashes and corrupts file system
- No log data to debug/analyze the problem



Operations Manager

Initial Solution

None

- No log data
- Concerned about too much data being captured on z/VM for Linux guests

Final solution

Capture Linux console & log data

- Use z/VM monitoring tool to capture console data and forward to Splunk
- Syslog data sent directly to Splunk

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

- Harden consolidated console data to disk in real-time
- One log file per day
- If have an issue, easily review console data written in the log file

Including Performance Data with z/OS Processes

The Situation:

- Collecting **performance data** on z/VM (Performance Toolkit)
- All mainframe performance data processed on **z/OS**
- Want to **include z/VM and Linux** data
- z/VM can not do FTP PUT to z/OS

Operations
Manager

Initial Solution

Manual processing each morning

Login and run commands to

- Summarize PerfKit data
- FTP file to z/OS
- Erase file from z/VM

Final solution

Automated processing overnight

- Schedule commands to summarize data on z/VM
- When complete, send message to z/OS
- z/OS: FTP GET file from z/VM
- z/OS: FTP message to z/VM indicating successful file retrieval
- z/VM: erase the file

Perform Weekly System Healthcheck



The Situation:

- Need to monitor system for various thresholds
 - **Spool space** filling up
 - **Paging space** filling up
 - **Disk full** for several z/VM service machines or guest

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

Linux Admin Access to Guests



The Situation:

- Linux system administrators want ability to add/remove disk space for their guests
- Linux admins don't have skills for z/VM commands

Wave

Initial solution

Requests via email/phone

- Linux admins call z/VM system programmers for each request
- Response time not ideal

Final solution


Tool for Linux Admins

- Graphical interface
- More intuitive for Linux admins
- Limited access to z/VM, DirMaint functions

Spool and Page Space Full

The Situation:

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



Operations Manager
or
OMEGAMON

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 for this one function

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
- Included in general monitoring/automation

Cloning Linux Servers



The Situation:

- Linux server creation taking too long (1.5. hours)
 - Faster than distributed but still manual
- Adding disk space to existing Linux servers taking too long
 - Also manual
- z/VM system programmers interrupted to stop or start Linux servers

Wave

Initial solution

Manual processes

z/VM system programmers using procedures involving many z/VM commands

Final solution

Graphical tool for z/VM system programmers and Linux administrators

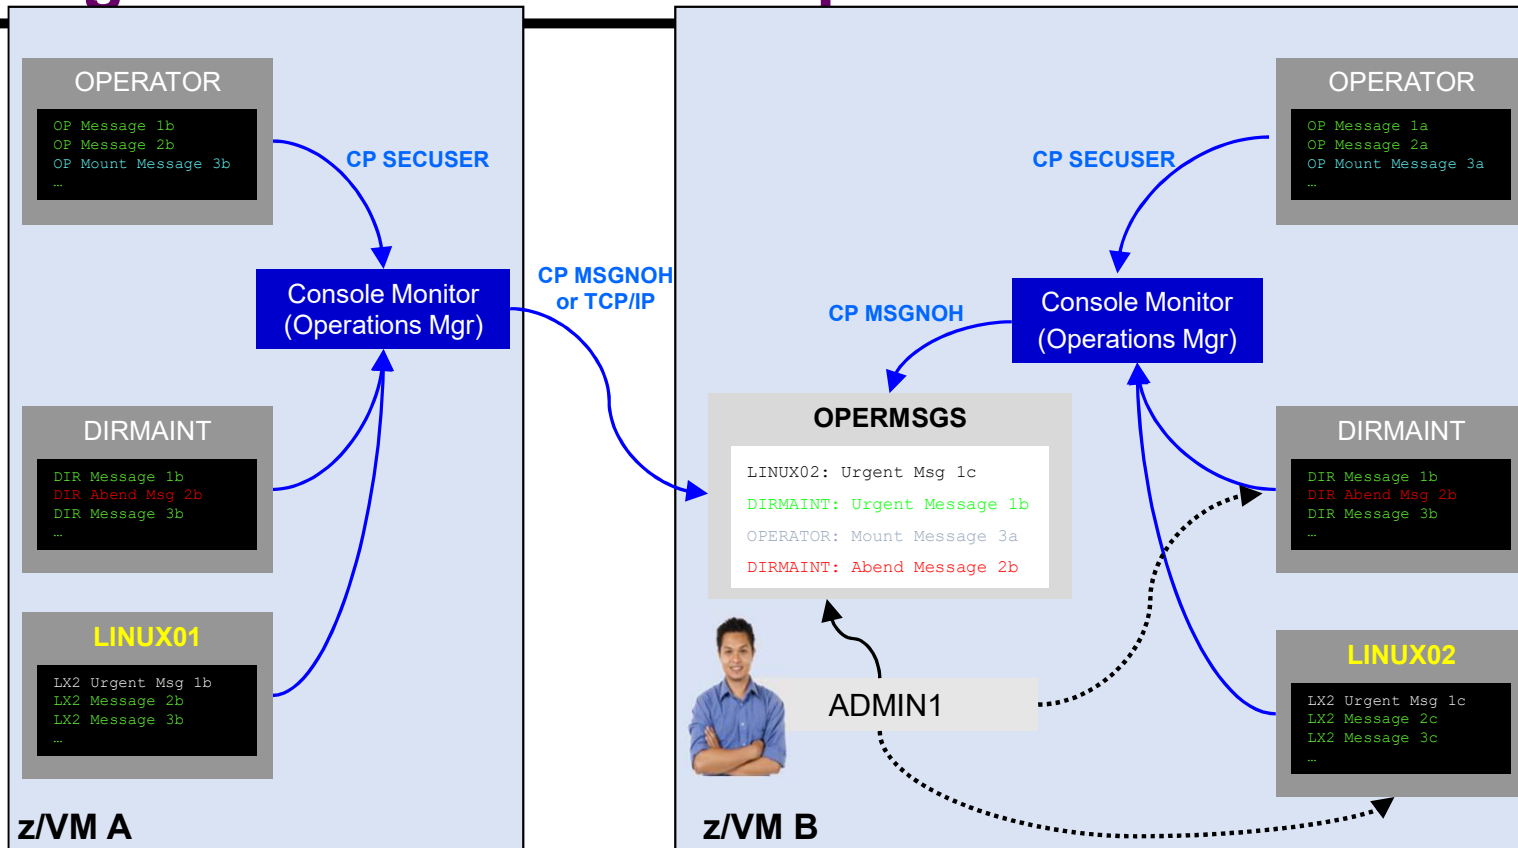
- More easily clone from golden image
- Reduced server creation time by 83%
- Reduced time to add disk space by 93%
- Linux admins given access to stop/start servers

Central Operations Console



- Already have z/OS console in operations center
 - Alerts, important messages
 - Operations staff watching consoles and taking actions
- 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



Single System Image (SSI) supported but not required

Unidentified Change in Performance



The Situation:

- System performed at “normal” level for period of time
 - CPU utilization
- Over several days, steady increase until “new normal”
- No new applications or virtual servers
- Unknown cause

OMEGAMON

Initial solution

Accept the new normal

Or

Dig through performance log data
(MONITOR records or Performance
ToolKit)

Final solution

Change control, historical data collection

- Lock in on the period of time of the increase
- Find specific servers contributing to increase
- Review change control records

On-Demand: Persistent Historical Views



Graceful Shutdown of z/VM from GDPS



The Situation:

- Shutdown of z/VM LPAR included in **GDPS** processing
- **Shutdown** of **Linux** guests handled by **GDPS**
- Need **graceful shutdown** of **z/VM** without triggering monitoring and automation

Operations
Manager

Initial solution

None

- GDPS handled shutdown of guests
- Shutdown of z/VM interfered with monitoring and automation

Final solution

Automated graceful shutdown

- GDPS signal triggers automation
- “Runtime” monitors and automation suspended/deactivated
- “Shutdown” monitors and automation resumed/activated

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 and
Restore
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
incremental backups 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 z/VM and Linux guest


- Performance data pointed to specific Linux guest
- 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 programmer why **application** was running so slow
- Application owner wants to research the problem

OMEGAMON

Drill down to each layer within a specified time window

Initial solution

Look at performance data for z/OS and Linux guest

- Performance data pointed to specific Linux guest
- No application data

Final solution

Performance monitoring solution for all layers

z/OS
Linux operating system
Application

Hypervisor (CP) Using 25% of CPU

The Situation:

- Most monitoring focuses on CPU utilization overall
- Missing focus on **hypervisor's % of CPU** as a separate metric
 - How much is the hypervisor using?
 - What's the "overhead" of the hypervisor
- Best Practice is to investigate if hypervisor using > 10% of CPU
- One morning found hypervisor using 25% of CPU



OMEGMAON

Initial solution

None

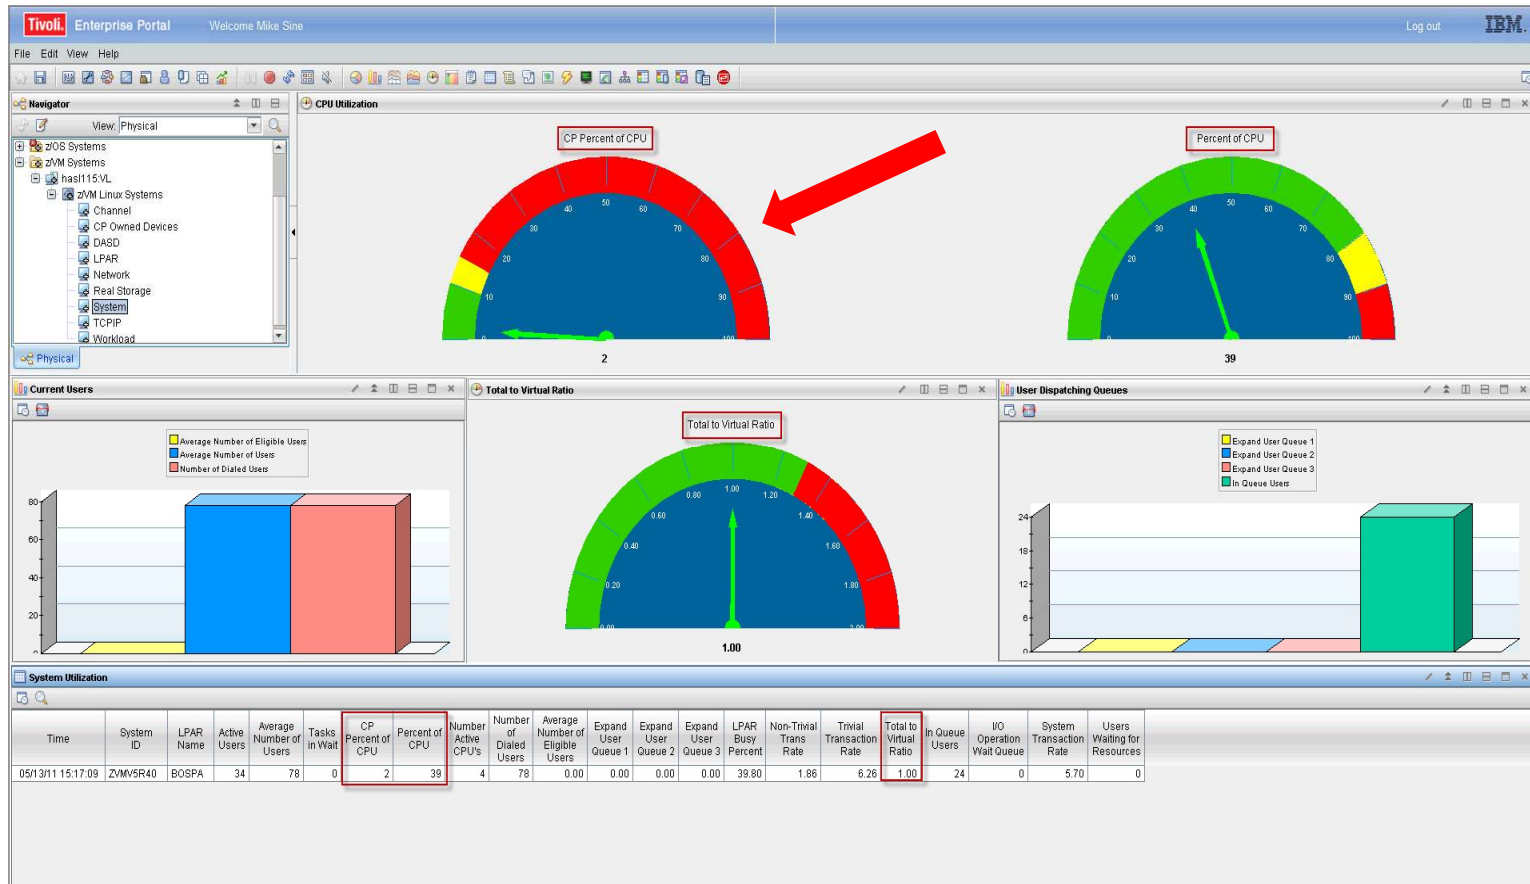
- System CPU measured, while hypervisor (CP) specific numbers omitted
- Reactive steps only taken when performance issue arose

Final solution

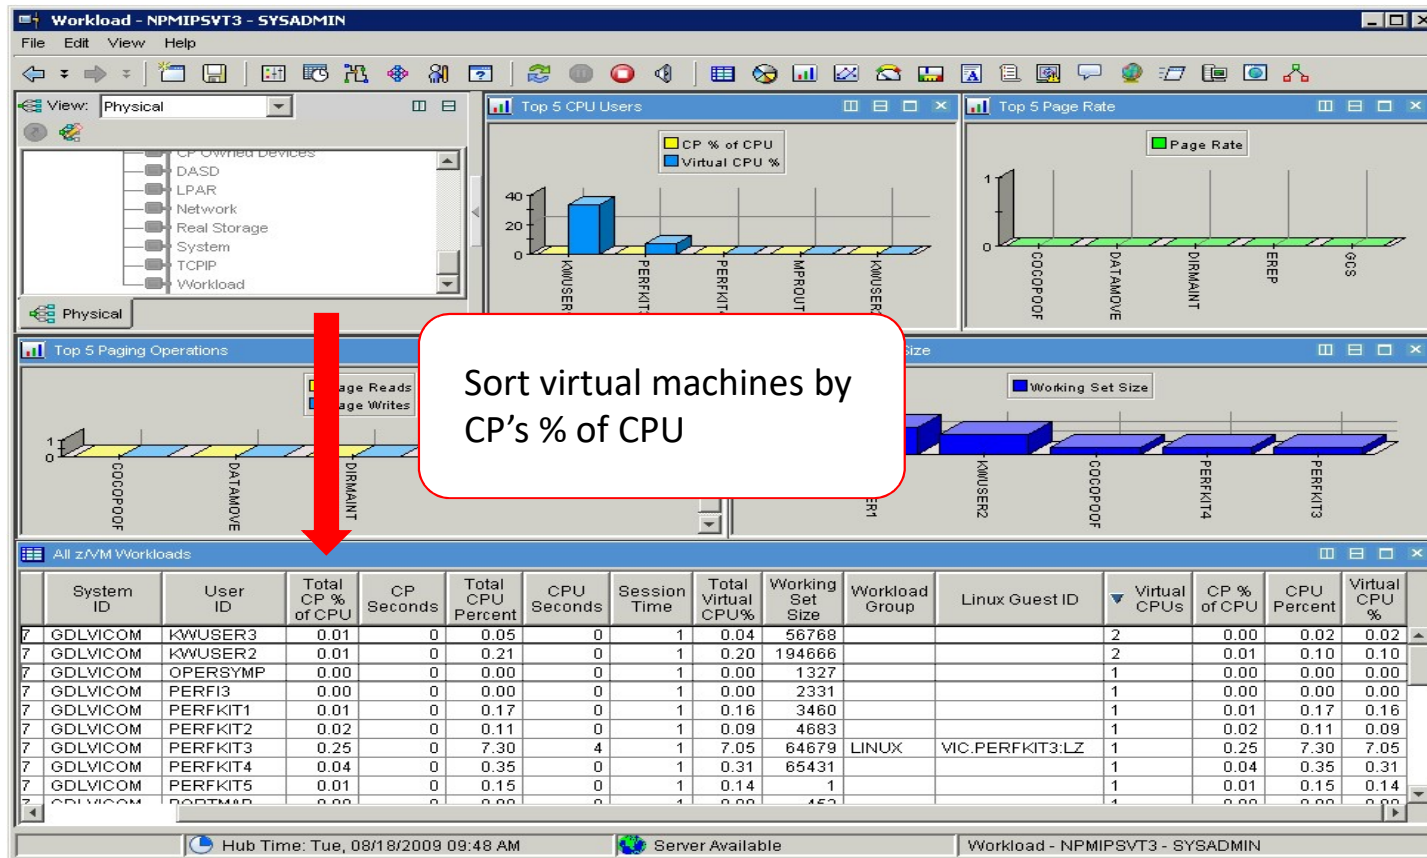
Monitoring tool

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

System Processor Utilization Workspace



z/VM Workload Workspace



Summary and Reference Information

- Production (and dev/test) systems need
 - Monitoring – operational and performance
 - Automation
 - Administration and provisioning
 - Backup and recovery
- Real situations need to be addressed
 - Learn from others
- Solutions exist
- Demos available
- Contact
 - Tracy Dean, tld1@us.ibm.com

IBM Infrastructure Suite for z/VM and Linux

- Bundle/suite of IBM products
- Tools needed to manage the z/VM and Linux on z Systems infrastructure
- Website:
 - <https://www.ibm.com/us-en/marketplace/infrastructure-suite-zvm-and-linux>
- **Videos on Youtube:**
 - <https://www.youtube.com/watch?v=tSWMdbyL49A&list=PLezLS0Tuqb-7gPpC56G1HzJ2893qhAcau>
- **IBMVM Mailing list:**
 - <http://listserv.uark.edu/archives/ibmvm.html>

Summary and Reference Information

- IBM Z IT Service Management newsletter
 - Arrives in your inbox once every two months
 - Announcements
 - Coming events
 - Resources
 - Includes: Operational Analytics, Automation, OMEGAMON, IMS Tools, CICS Tools, Storage Tools, z/VM Tools
 - Subscribe:
ibm.biz/zITSMNewsletterSubscribe

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. **Monitor service machines for logoff – and autolog them**
13. Send an email if page space approaches full
14. Monitor SSI connectivity between 2 cluster members
15. **Suppress passwords on Linux consoles**
16. Autolog a Linux guest and send message if doesn't start successfully
17. Monitor Linux file system and send email when approaching full
18. Send alerts to other tools via syslog
19. Non-SSI high availability environment: monitor LPAR CPU utilization – if too high, stop a guest and restart on another LPAR

Backup and Recovery – Demos Available

- Performing an incremental backup
- Restoring files from backup
- Back up and restore single and multiconfiguration users in an SSI environment
- Scheduling image backups of Linux guests
- Suspend and resume a guest as part of backup
- Pause and resume a long-running backup for system IPL
- Reviewing a disaster recovery backup
- Reviewing data in the backup catalog for recovery

धन्यवाद

Hindi

多謝

Traditional

감사합니다

Korean

Спасибо

Russian

Ndzi khense ngopfu

Tsonga

Gracias

Spanish

شكراً

Arabic

Grazie

Italian

Thank You

English

Obrigado

Brazilian Portuguese

Danke

German

多谢

Simplified Chinese

Merci

French

Ke a leboha

Tswana

நன்றி

Tamil

ありがとうございました

Japanese

ขอบพระคุณ

Thai