

# Solutions for Real Customer Challenges

## Managing z/VM and Linux

---

Tracy Dean, Product Manager, z/VM Tools & IMS Tools

June 2022

# Agenda

Introduction

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

## z/VM Customers



- North America
- Europe
- Greater China
- Japan
- Latin America
- Asia Pacific
- Middle East & Africa

## Customers worldwide using z/VM

- Most to host Linux workloads
- Historical z/VM workload and skills centered in North America and Europe
- Continued growth requires worldwide adoption
- Customers need to show value quickly
- Time, resources and skills not available to write management and monitoring tools in-house

## Unique value hosting Linux workloads

- Requires integration with existing enterprise solutions
- “Just another platform”
- Must be monitored, automated, and recoverable

**What is managing and  
what tools can I use?**

# Performance monitoring

## 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

## Solutions

- Multiple vendors
- IBM
  - Performance Toolkit for z/VM
  - OMEGAMON XE on z/VM and Linux

# Operational monitoring & automation

## Collect and store data

- Consoles
- User and system events
- View live and store on disk in real time

## React

- Generate alerts or automatically recover
- 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

## Prevent

- 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

## Solutions

- Multiple vendors
- IBM
  - Operations Manager for z/VM

# Backup and recovery

## 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

## Solutions

- Multiple vendors
- IBM
  - Backup and Restore Manager for z/VM
  - Spectrum Protect

# Customer scenarios



# 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 not responding

## Initial solution: None

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

## Final solution: Alert via heartbeat check

- 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
- Can be used for multiple z/VM systems to monitor each other

# Coordinate application shutdown with z/OS



## The situation

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

## Initial solution: None

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

## Final solution: Console monitoring

- 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
- Reverse process for re-start

# Separate backups and tapes for DR test



## The situation

- Using Backup and Restore Manager for z/VM and Tape Manager for z/VM in production to perform backups to tape
- Production DASD mirrored to DR site
- TS7700 tape grid – separate cluster at DR site
- During DR **test**, need to
  - Prevent writing and expiring production tapes
  - Use separate range of DR test volsers for write processing
  - Allow reading production tapes

Backup and Restore Manager for z/VM

Tape Manager for z/VM

TS7700

## Initial solution: Treating DR test like real DR

- Using all production tape volsers in DR test as if in production
- Have two tape catalogs that are not synchronized talking to the same tape grid
- Can overwrite production tapes at DR test

## Final solution: Separate library categories and tape pools for DR test

- Define once in production (mirrored to DR test)
  - Separate scratch category in TS7700 for DR test
  - Separate range of volsers for DR test
  - Separate Tape Manager tape pool for DR test
- At DR test
  - Put all volsers in production tape pool(s) on hold
    - Prevents writing or expiring from DR test
  - Update Backup Manager to use DR test tape pool
- Real DR not an issue
  - It is production running at a different site

# 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: Create homegrown tool

- Scan logs on a daily basis
- Search for error messages

## Final solution: Console monitoring

- Write a rule looking for error message that occurs during boot process
- 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

## 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: Disk monitoring

- Simple monitor setup
- Automatically monitor percent full
- Email someone
- Follow documented procedures to save data
- Follow-up action to automate procedures

# 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** alerts **included**

## 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 managing and automating z/VM and Linux alerts

## Final solution: Monitoring and automation

- 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
- Enables 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 for z/VM  
Tape Manager for z/VM (optional)

## Initial solution: Considered two 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 incremental backups of all z/VM log files and customizations
- Retain these backups for months or years without large amounts of tape or DASD
- Can optionally do image level backups from z/VM instead of from z/OS

# Consolidating all console data



## The situation

- Using Operations Manager to capture console data and store it on disk
- **Missing data** from users autologged before Operations Manager during z/VM IPL
  - OPERATOR
  - RACF

## Initial solution: Separate data in spool and on disk

- Spooled consoles for users logged on before Operations Manager
- Content of monitored consoles on disk
- For same user, data in spool file separate from Operations Manager data on disk

## Final solution: Combined data on disk

- At Operations Manager startup
  - Receive spool data for users already running
  - Add spool data to console of running user
- Automatically added to Operations Manager log and visible in VIEWCON
- New sample code in PTF UI77841



# 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”

## Initial solution: None

- No analytics and alerting of z/VM RACF-related activity
- z/VM looks “different” to management

## Final solution: Automation

- Automatically capture RACF logon/logoff messages on OPERATOR console
- Use TCP writer to send messages to analytics platform in syslog format
  - Next day follow-up to send in key/value pair format

# 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


## Initial solution: Two options

- Accept the new normal
- Dig through performance log data
  - MONITOR records
  - Performance Toolkit for z/VM

## Final solution: Change control and 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



# 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
  - z/OS must retrieve or GET

## 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 each night

- 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

# Capturing Linux log data



## The situation

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

## Initial solution: None

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

## Final solution: Capture and save Linux console and log data

- Use z/VM monitoring tool to capture console data, store on disk, and forward to Splunk
- Configure Linux to send syslog data directly to Splunk

# 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

- **Train people** to make backup copies before updating a file

## Final solution: File level backup and recovery

- Weekly full backups of all z/VM customizable or customer created files
- Daily incremental backups of the same z/VM files

# Graceful z/VM shutdown 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
  - Avoid false alerts
  - Keep monitoring and automation running

## 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

# Spool and page space full



## The situation

- z/VM “just works”
- Little monitoring in place
- System grows
- **Spool and page space fill up**
- System abends
- **Unplanned outage**

Operations Manager for z/VM  
or  
OMEGAMON XE on z/VM and Linux

## 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 specific function

## Final solution: Consolidated monitoring

- Simple monitor setup
- Watch for percent full to be within (multiple) specified threshold ranges
- Watch for sudden growth
- Take action
- Easily add or change threshold or frequency
- Included in general monitoring/automation



# 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 needed to research the problem

Initial solution: Look at performance data for z/VM and Linux guest

- Performance data pointed to specific Linux guest
- No performance data at the application level

Final solution: One performance monitoring tool for all layers

- Hypervisor
- Linux operating system
- Application
- Drill down to each layer within locked-in specified period of time

# Restarting shared servers on another member of SSI cluster



## The situation

- Servers such as Shared File System can be used across an SSI cluster
- One member of the cluster goes down for maintenance
- Server must be restarted on another member so services are still available to remaining members
  - Need minimal downtime

## Initial solution: Manual procedures

- Shut down SFS servers on the LPAR being taken down (LPAR1)
- Restart SFS servers on another LPAR in the cluster (LPAR2)
- Shut down LPAR1

## Final solution: Automation

- Detect shut down of SFS server on LPAR1
- Trigger an action on LPAR1 to tell LPAR2 to start the SFS server
- Action on LPAR2 starts the server
- Minimal downtime and no typos or human delays
- Some servers can span beyond SSI cluster
  - Methodology still works
  - SSI not required

# Central operations console



## The situation

- Already have one or more **z/OS consoles** in operations center
  - Alerts, important messages
  - Operations staff watching consoles and taking actions
- Want to include console(s) for all **z/VM LPARs** and **Linux** guests

## Initial solution: OPERATOR console on each LPAR

- Multiple consoles to monitor
- Does not include all alerts or important messages

## Final solution: Single operations console for multiple z/VM LPARs

- Create one user ID as operations console
- Monitor consoles of all service machines and Linux guests on each LPAR
- Create rules for important messages
- Take action to send alert or message to operations console
  - Within the LPAR: via MSG or MSGNOH
  - Across LPARs: via TCP/IP, **SSI not required**
- Operations team sees only important messages
  - When needed, can view full console of any service machine or guest
  - Easily expand to include more LPARs

# Stopping and restarting TCP/IP



## 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

## Initial solution: Manual coordination

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

## Final solution: Automated process

- 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
  - Avoid maintaining duplicate definitions

## Initial solution: Manual processing

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

## Final solution: Automated real-time sharing of configuration information

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

# Perform system health checks regularly



## 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 guests

## Initial solution: Login weekly and perform checks manually

- Simple checklist
  - Check disk space
  - Check page space
  - Check spool space

## Final solution: Automate regular monitoring and alerts

- Perform checks more often
- Notify team if any are out of bounds

# System abend with no console data



## The situation

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

## Initial solution: Two options

- IPL cold start and hope for the best
- IPL cold start and dig through dump data (if you have it)

## Final solution: Console monitoring

- 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

# Summary

Real situations  
need to be  
addressed

Solutions exist

Demos available

Learn from others

## Production (and dev/test) systems need

- Monitoring: operational and performance
- Automation
- Backup and recovery

## Contact:

Tracy Dean

[tld1@us.ibm.com](mailto:tld1@us.ibm.com)

## Website:

[www.ibm.com/products/infrastructure-suite-zvm-and-linux](http://www.ibm.com/products/infrastructure-suite-zvm-and-linux)



THANK YOU!

QUESTIONS?

# Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

IBM\*

ibm.com\*

IBM logo\*

## \* Registered trademarks of IBM Corporation

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

BMC is a registered trademark of BMC, Inc in the United States, other countries, or both.

Broadcom, CA are registered trademarks of Broadcom, Inc in the United States, other countries, or both.

IntelliMagic is a registered trademark of Intellimagic, Inc in the United States, other countries, or both.

IT Infrastructure Library is a Registered Trade Mark of AXELOS Limited.

ITIL is a Registered Trade Mark of AXELOS Limited.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Zowe™, the Zowe™ logo and the Open Mainframe Project™ are trademarks of The Linux Foundation.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

MXG is a registered trademark of Merril Systems Group, Inc in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Rocket is a registered trademark of Rocket Software, Inc in the United States, other countries, or both.

SAS is a registered trademark of SAS, Inc in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Zetaly is a registered trademark of Zetaly, Inc in the United States, other countries, or both.

Other product and service names might be trademarks of IBM or other companies.

© IBM Corporation 2021. All Rights Reserved.

The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this publication or any other materials. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

IBM