

# Solutions for Real Customer Challenges Managing z/VM and Linux

Tracy Dean

Product Manager, z/VM Tools and IMS Tools

June 2026



# Agenda

01

Introduction

02

What does  
“managing” include?

What tools or products can  
you use ?

03

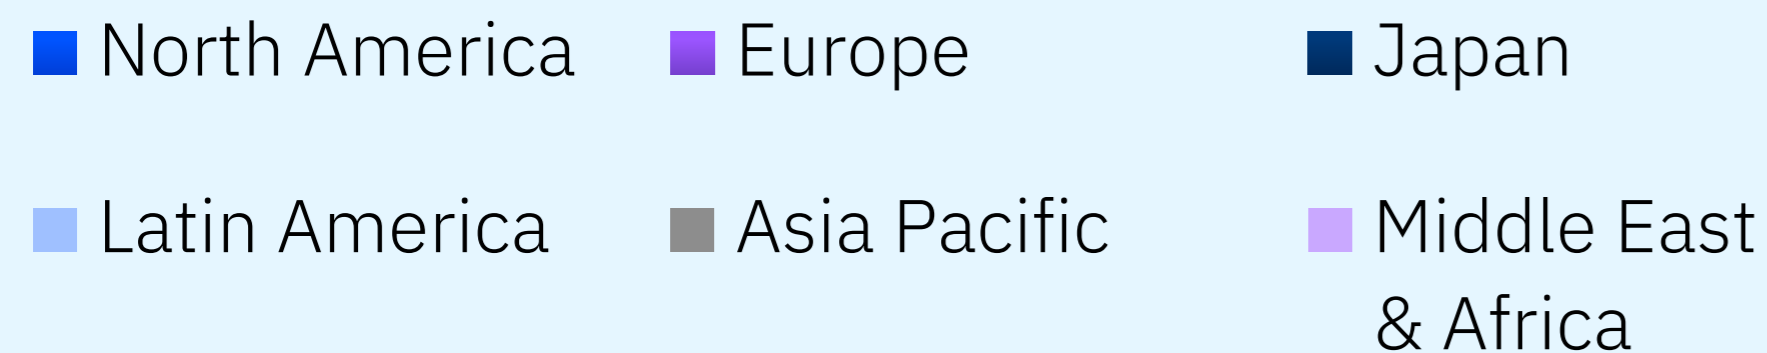
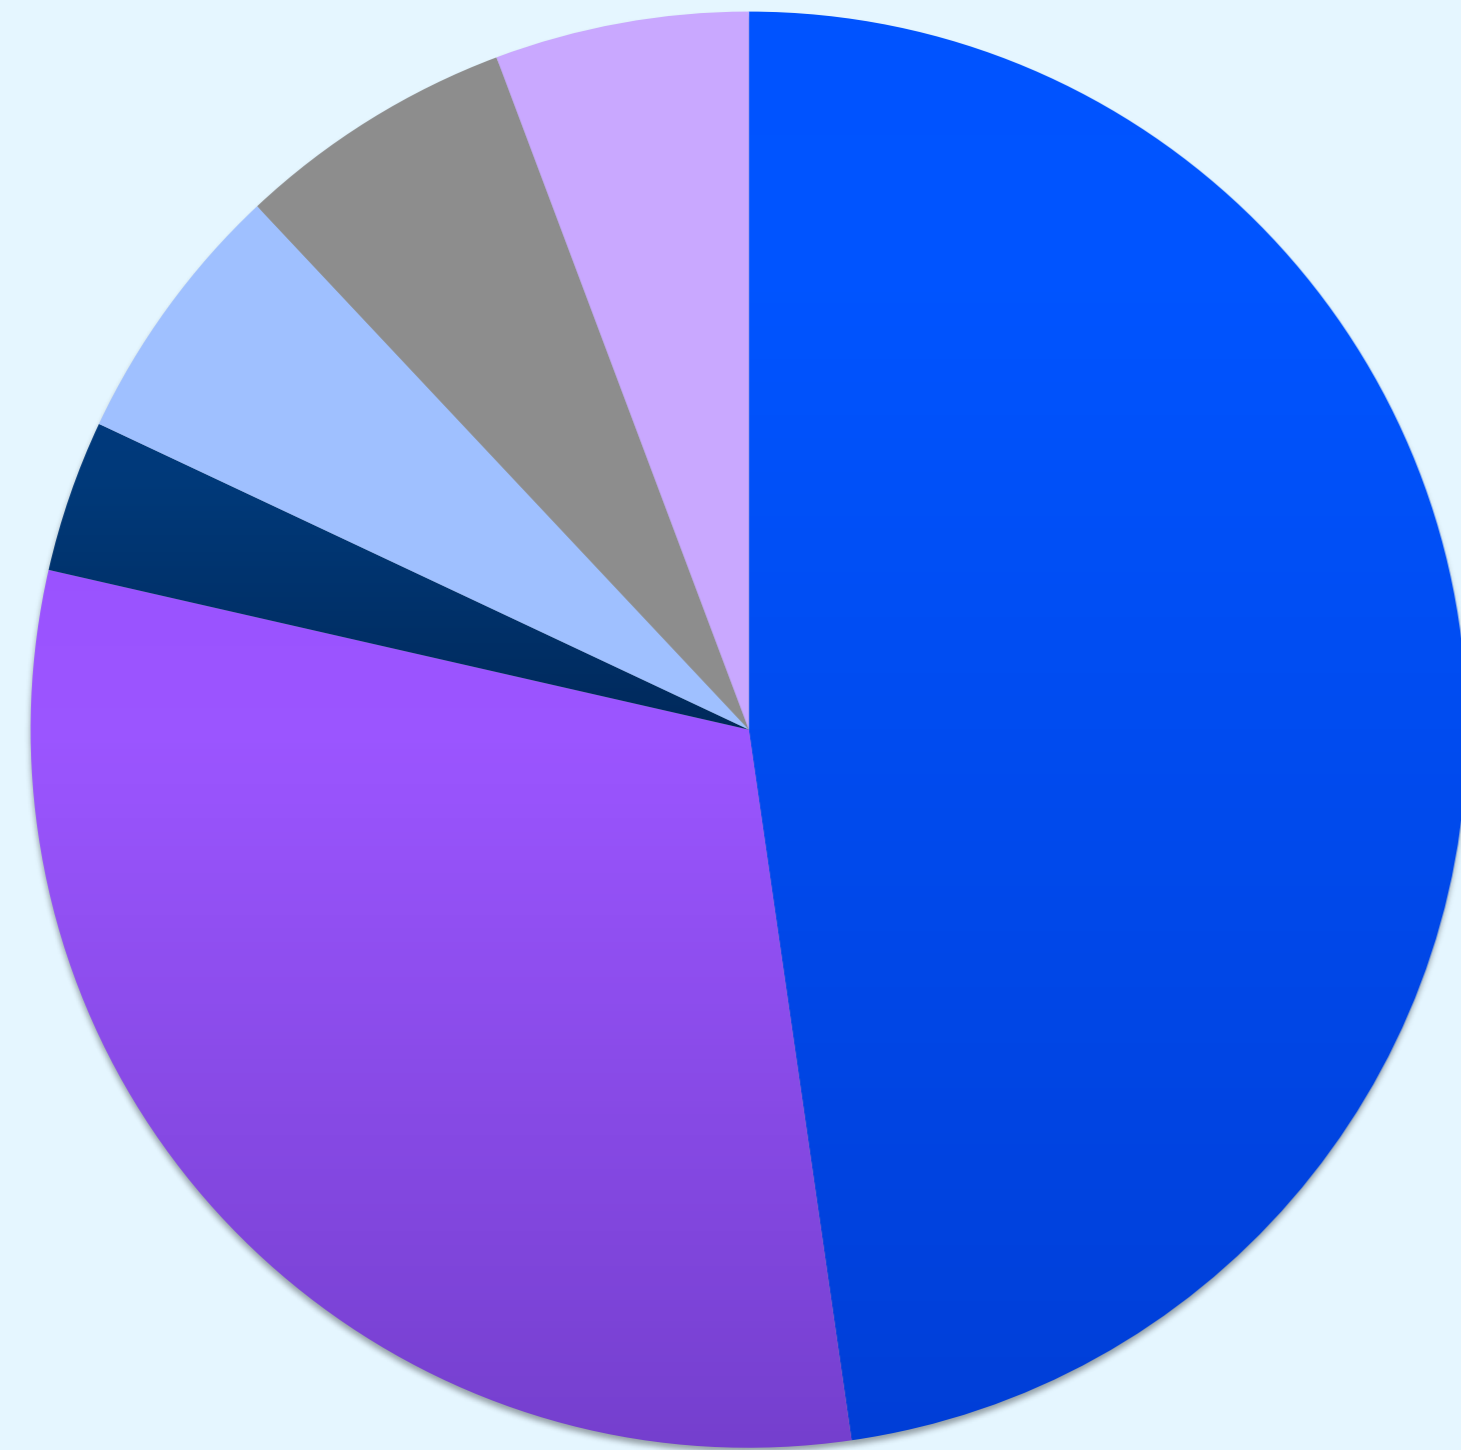
Customer scenarios

- Operational monitoring and automation
- Performance monitoring
- Backup and recovery

04

Summary and  
reference  
information

## z/VM Customers



## Customers worldwide using z/VM

- Most to host Linux workloads
  - VSE<sup>n</sup>, z/TPF, etc.
- Historical z/VM workload and skills centered in North America and Europe
- New personnel joining the platform worldwide
- Customers need to show value quickly
- Time, resources and skills not available to write management and monitoring tools in-house

## Unique value hosting Linux on IBM Z 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
- Capture and use historical data
- Solutions
  - Multiple Vendors
  - IBM
    - *IBM Performance Toolkit for z/VM*
    - *IBM OMEGAMON XE on z/VM and Linux*

## Operational monitoring and automation

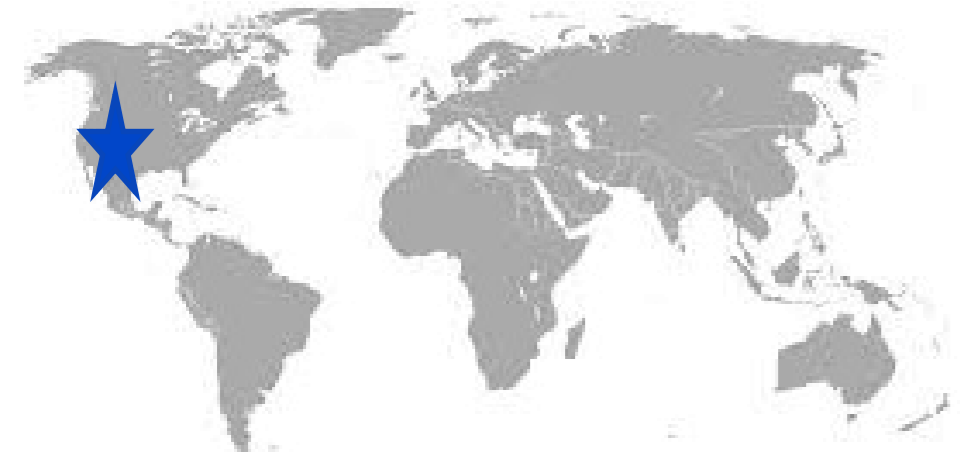
- Collect and store operational data: consoles, user and system events
- Monitor minidisk space, SFS space, spool, etc.
- React to issues by generating alerts or executing commands or scripts
- Prevent issues by creating regularly scheduled actions
- Solutions
  - *Multiple vendors*
  - *IBM Operations Manager for z/VM*

## Backup and Recovery

- Image level backup of z/VM and Linux
- File level backup of z/VM and Linux
- Solution
  - Multiple vendors
  - *IBM Backup and Restore Manager for z/VM*
  - *IBM Storage Protect*

Customer scenarios

# 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” (and less capable) 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
  - By next day, updated to send in key/value pair format

Operations Manager for z/VM

# 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

# 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

## Initial solution: Considered two options

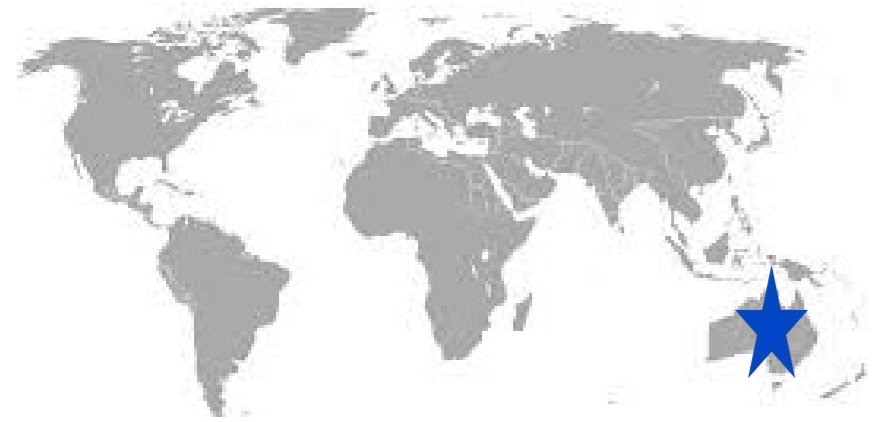
- Keep volume backups on z/OS for months instead of weeks
  - Additional unnecessary data retained longer on tape
  - Tedious file level recovery
- Re-do all customizations from memory (or with help from IBM)

## 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

**Backup and Restore Manager for z/VM**  
**Tape Manager for z/VM (optional)**

# Sending console and event data to Grafana Loki



## The situation

- Starter system for z/VM and Linux on IBM Z
- During debugging, want to view console, user and system event data off-platform
- Do not want to create z/VM user IDs for every user authorized to view the data

## Initial solution: None

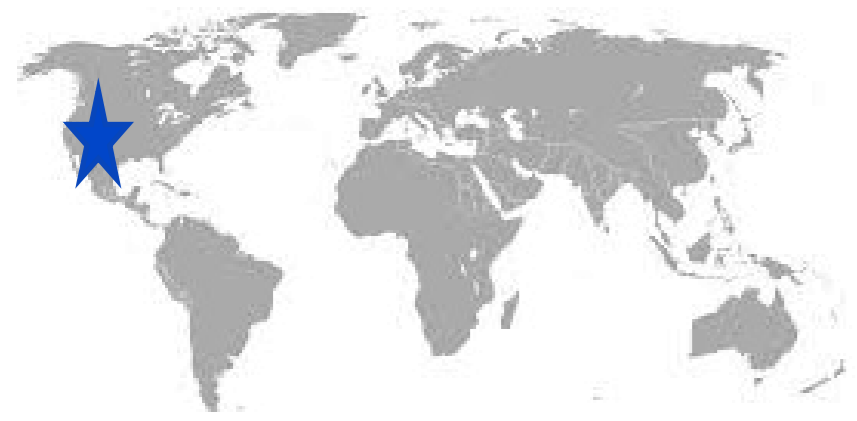
- To view consoles and event data
  - Users require 3270 access
  - Users require z/VM user IDs

## Final solution:

Send console data off platform on regular schedule

- Define schedule on desired interval
- Use TCP writer to send daily log file to target host and port
  - Log file at current point in time
  - Contains all console and event data (and more)

# Capturing multi-line messages



## The situation

- Using Operations Manager to capture and react to messages
- z/VM issues some messages as multi-line messages
  - RACF  
ICH408I USER(MNT2 ) GROUP(\$SERVADM)  
NAME(#####)
  - OPMGRM1.198 CL(VMMDISK)  
INSUFFICIENT ACCESS AUTHORITY  
ACCESS INTENT(CONTROL) ACCESS ALLOWED(NONE )
- Unable to send full message text as part of response/action

## Initial solution: None

- Only able to capture the first line of the message to include in
  - A notification
  - Streaming to analytics

## Final solution: New IBM Idea

- Request new capability from development
- When message rule is triggered, capture the next  $n$  lines of messages
  - Make them available to the associated action
  - Specify a time limit for  $n$  messages to appear
- New function APAR/PTF: PH68258/U005223

# Report of backup tapes used



## The situation

- Running backups of z/VM and Linux environment
- Using virtual tape library
- Want to capture the tapes used for specific backup jobs
  - Make list available for Disaster Recovery without DASD mirroring

## Initial solution: Manual commands

- Wait for jobs to complete
- Login and view Backup Manager console output to find the tapes as they are being mounted
- Tedious and need to remember to do this after each backup job completes

## Final solution: Automation

- Define a rule to trigger on the backup complete message, capturing the job name and instance
- When rule is triggered
  - Call an action to issue command to Backup Manager to return list of tapes associated with specific job name and instance
  - Create and send email with list of tapes
- New Operations Manager sample: APAR/PTF: PH68258/U005223

**Backup and Restore Manager for z/VM**

**Tape Manager for z/VM**

**Operations Manager for z/VM**

# Coordinate application shutdown with z/OS



## The situation

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

## Initial solution: None

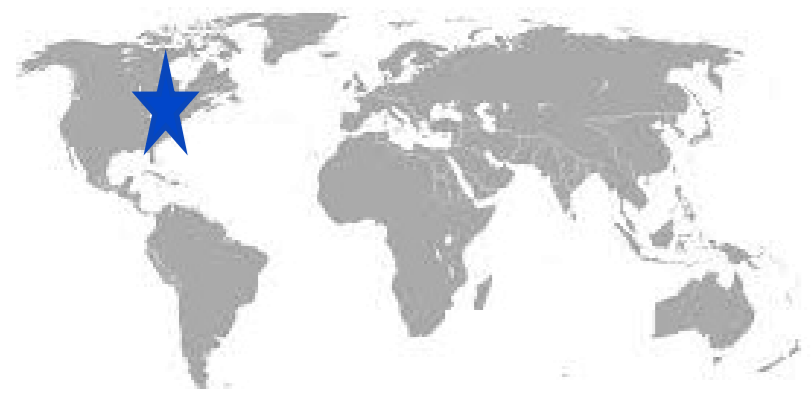
- Manual coordination of shutdown
- Inconvenient for z/VM system programmers and operations during non-business hours

## 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 servers are shut down
- Reverse process for re-start

Operations Manager for z/VM

# Stopping and restarting TCP/IP



## The situation

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

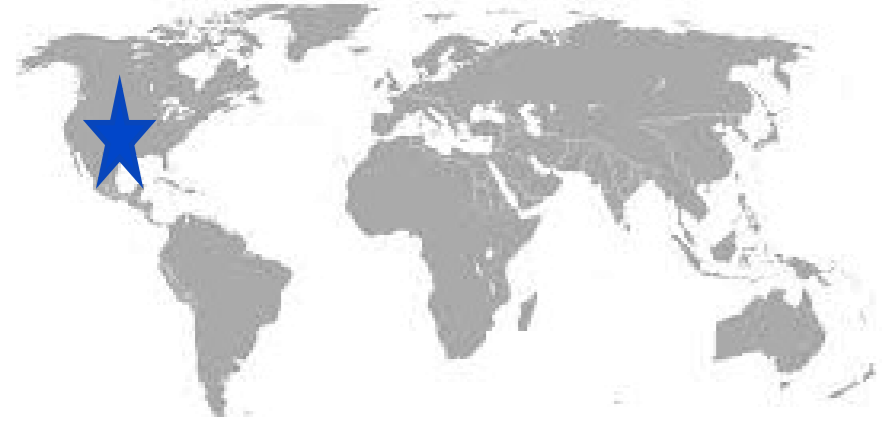
## Initial solution: Manual coordination

- 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

# 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
- VMRELOCATE not available for SFS servers
- Server must be restarted on another member so services are still available to remaining members
  - Need minimal downtime

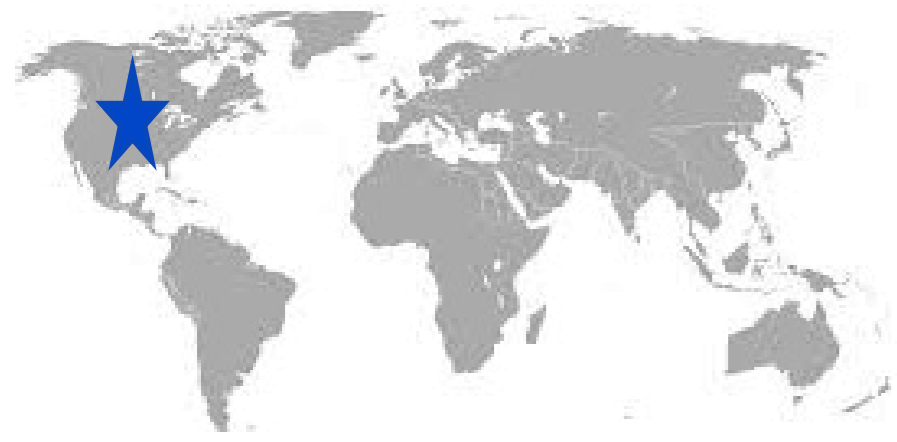
## Initial solution: Manual procedures

- Shut down SFS servers on the LPAR being taken down (LPAR1)
- Restart SFS servers on another LPAR in the cluster (LPAR2)
- Shut down LPAR1
- SFS server unavailable during manual process

## 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

# Separate backups and tapes for Disaster Recovery 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
  - Allow reading production tapes
  - Prevent writing and expiring production tapes
  - Use separate range of DR test volsers for write processing

## 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 system
  - Update Backup Manager to use DR test tape pool
- Real DR not an issue
  - It is production running at a different site

[Backup and Restore Manager for z/VM](#)

[Tape Manager for z/VM](#)

[TS7700](#)

[White paper](#)

# Send z/VM and Linux alerts to z/OS



## The situation

- Minimal z/VM staff
- 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
- z/VM team managing and automating z/VM and Linux alerts
  - Prefer mainframe operations team does this

## Final solution: Monitoring and automation

- Trigger alerts for z/VM and 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

Operations Manager for z/VM

# Error messages on Linux IPL



## The situation

- During boot process, Linux file system is sometimes **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
- Long term: use logs to debug and find root cause

## 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

# 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 collection of historical performance data

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

# 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

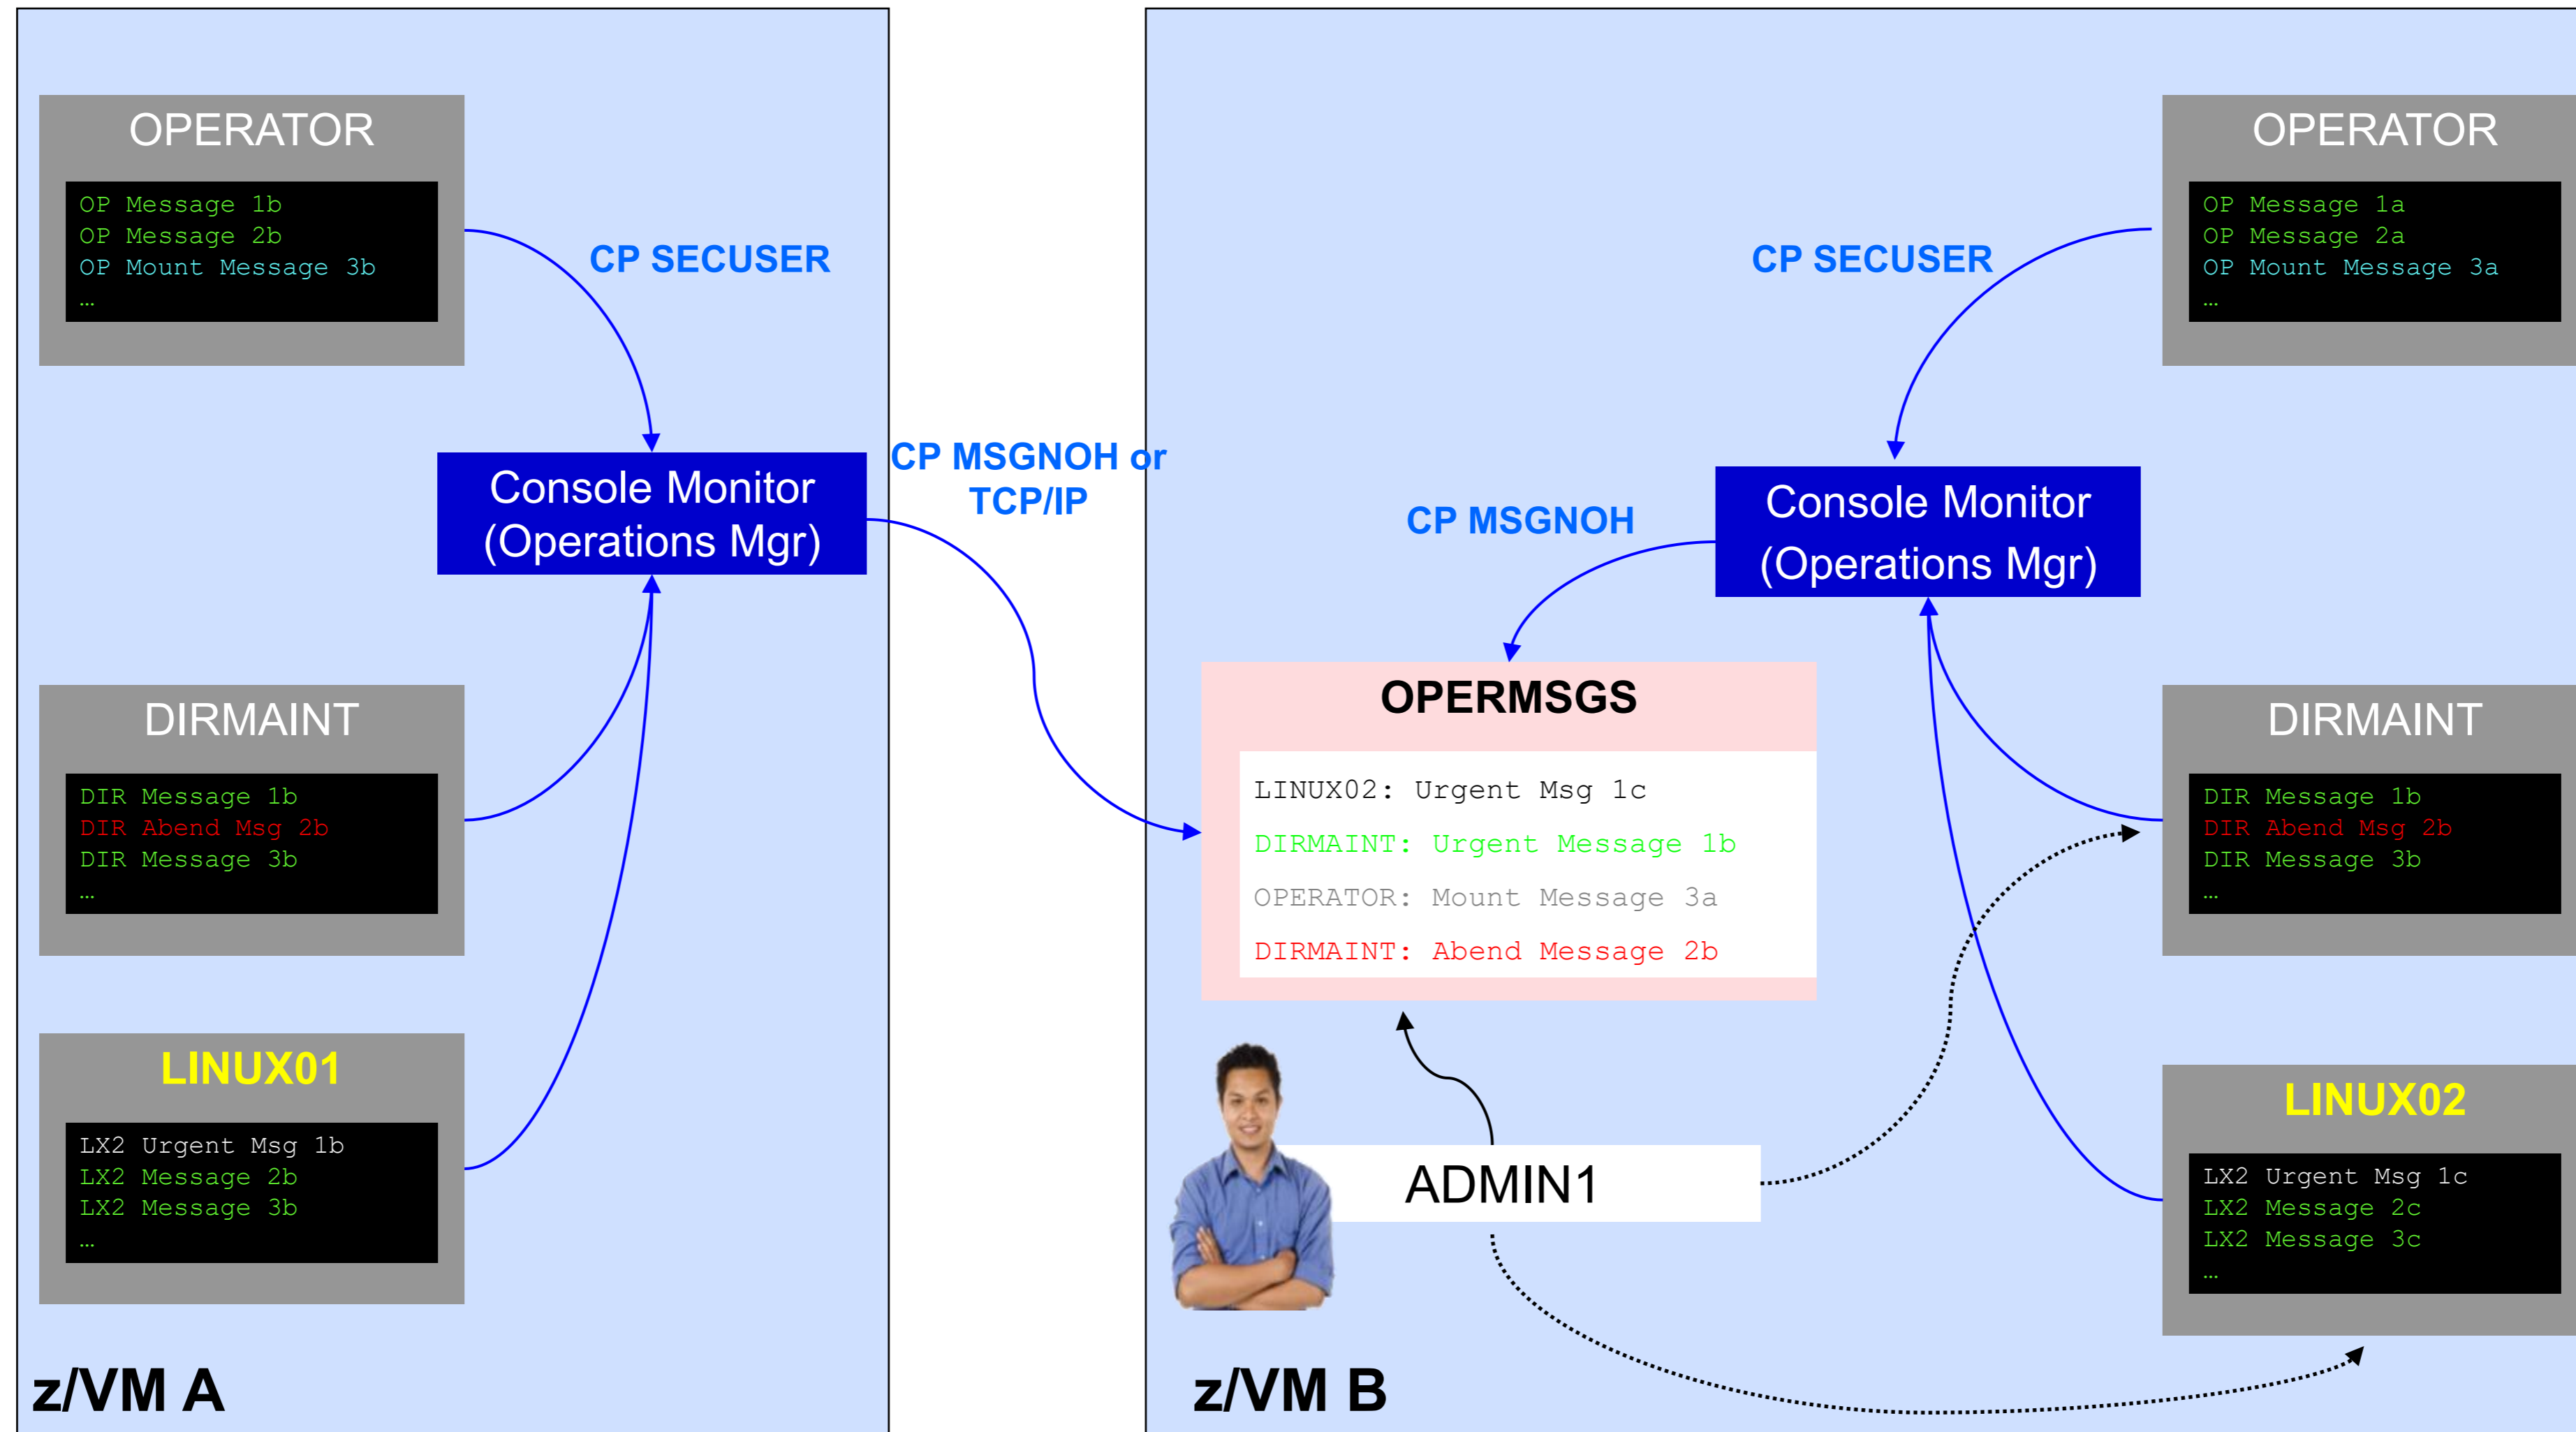
- Multiple consoles to monitor
  - Inhibits z/VM and Linux growth
- Does not include all alerts or important messages
- Includes unnecessary informational messages

## 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 message to single 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

Operations Manager for z/VM

# Creating a Central Console Operations Console



**Single System Image (SSI) supported but not required**

# 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 or 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 analytics platform
- Configure Linux to send syslog data directly to analytics platform

# 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 GET 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 PUT message to z/VM indicating successful file retrieval
- z/VM: erase the file

Operations Manager for z/VM

# 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

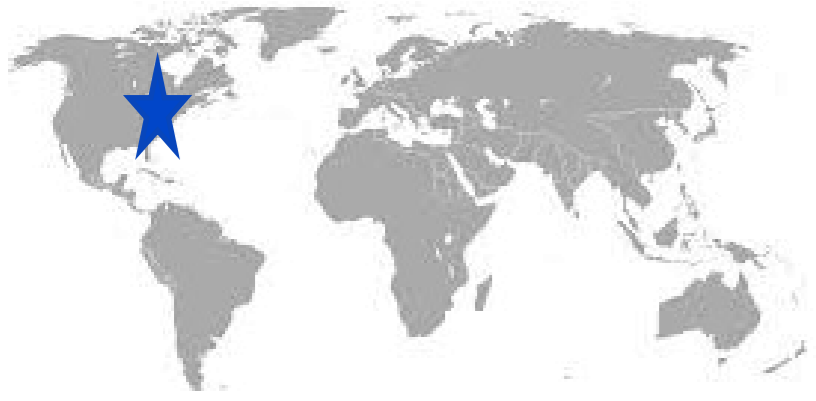
- Common configuration information maintained on one system
- Shared within SSI cluster
- Manually copied and reloaded on LPARs or SSI clusters

## 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

Operations Manager for z/VM

# 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

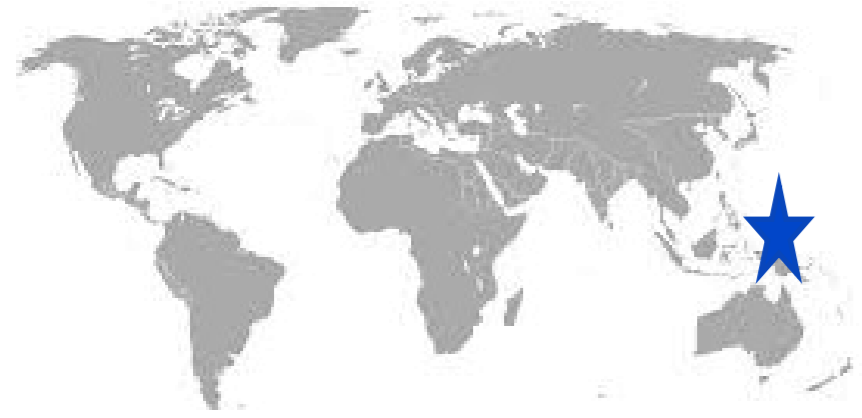
- Train people to make backup copies before updating a file

## 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

Backup and Restore Manager for z/VM

# 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 mobile phone
- Show IBM and ask for help
- **EREP & Accounting disks full**
- No knowledge of impact

## Final solution:

Turn off accounting – not used

## Automate disk monitoring

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

**Operations Manager for z/VM**

# 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

Operations Manager for z/VM

# 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**

## 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

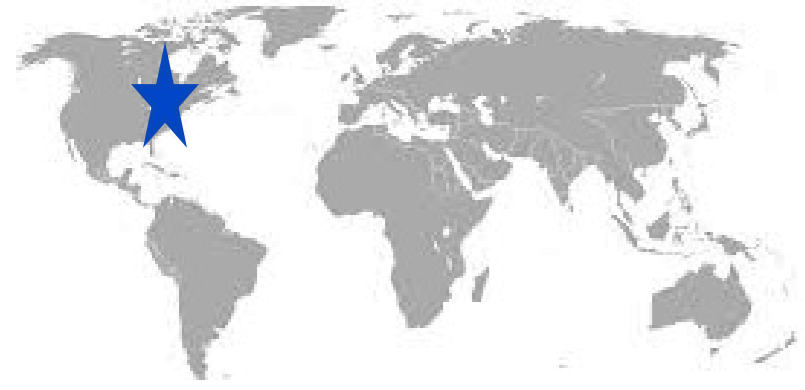
- Simple monitor setup
- Watch for percent full to be within (multiple) specified threshold ranges
- Watch for sudden growth
- Take action
- Easily add or change threshold or frequency
- Included in general monitoring/automation

Operations Manager for z/VM

or

OMEGAMON XE on z/VM and Linux

# Why was an application running slow?



## The situation

- Application owner asks z/VM system programmer why **application** was running **slow yesterday** afternoon
- Application owner does not 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

OMEGAMON XE on z/VM and Linux

# 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

Operations Manager for z/VM

# 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
- Learn from others
- Solutions exist
- Demos are available

## 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)

Karishma Deole

[Karishma.Deole@ibm.com](mailto:Karishma.Deole@ibm.com)

## Website:

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