IBM Z

Operational Monitoring and Automation of z/VM and Linux on IBM Z

Tracy Dean

IBM Product Manager

January 2025

IBM

Agenda

- Introduction to recommended practices and examples
- IBM Operations Manager for z/VM
  - Overview
  - Customer experiences
- Considerations for z/VM Single System Image
- Summary
  - Reference information

IBM Solutions

- Security
  - RACF and zSecure Manager for z/VM
- Performance monitoring
  - OMEGAMON XE on z/VM and Linux
  - Performance Toolkit for z/VM
- Backup and recovery
  - Backup and Restore Manager for z/VM
  - Tape Manager for z/VM
  - Storage Protect (aka Spectrum Protect or Tivoli Storage Manager)
- Automation and operational monitoring
  - Operations Manager for z/VM
    - Including integration with existing monitoring and alert systems

Complete Solution for Management of the z/VM and Linux IBM Z or LinuxONE Environment

IBM Infrastructure Suite for z/VM and Linux V2

OMEGAMON XE on z/VM and Linux

Performance monitoring of z/VM hypervisor and Linux guests

Storage Protect

File level backup and recovery for Linux virtual machines

Operations Manager for z/VM

- Facilitate operational monitoring and automated operations
- Take action based on events

Backup and Restore Manager for z/VM

- Image and file level backup/restore of z/VM environment
- Image level backup/restore of Linux

Cloud Infrastructure Center

(optional separately priced feature)

IaaS offering that provides industry-standard user experience for both traditional and cloud infrastructure

Tape Manager for z/VM

(optional separately priced feature)

Support Backup and Restore Manager performing backups to and recovery from real or virtual tape systems

Single PID: 5698-K01 (S&S 5698-K02)

Recommended Practices – Operational Monitoring and Automation

Console monitoring and viewing – current state and historical

- 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

Gather Data

Keep monitoring close to the operating system

React

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

Monitor as you grow

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

Product Overview

IBM Operations Manager for z/VM

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

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

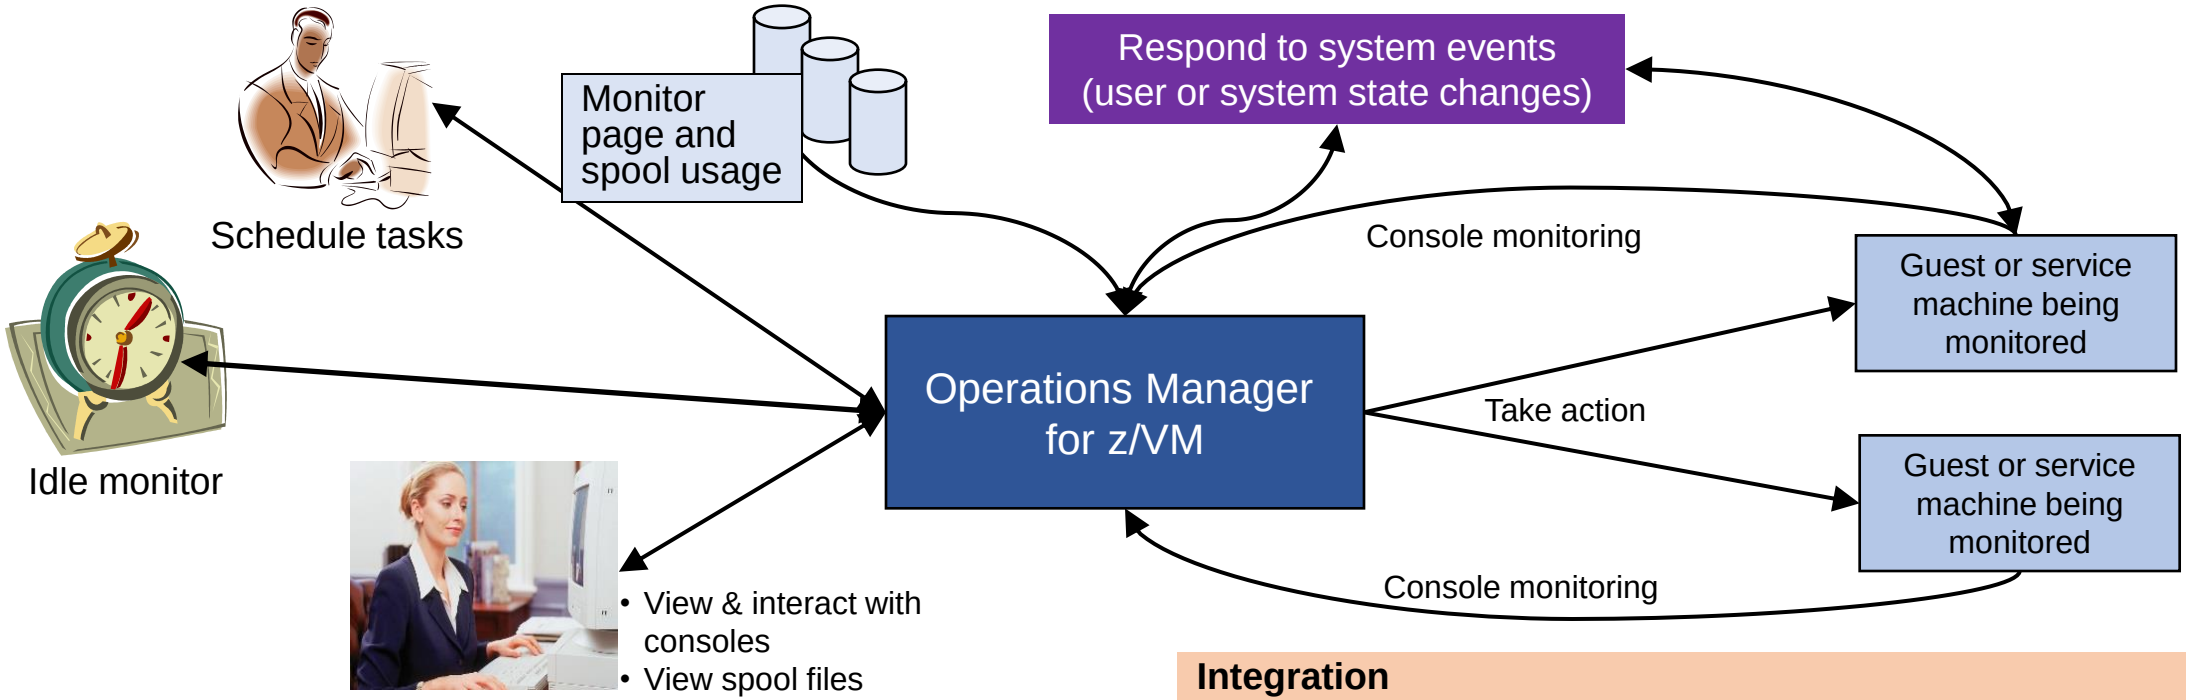

Automation
- Routine activities done more effectively with minimal operations staff
- Schedule tasks to occur on a regular basis

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

Executing Actions

- Specify action to take in response to
  - Console rule definition
  - Schedule
  - Spool monitor
  - Etc.
- Types of actions
  - Change color, highlight, hold, or suppress a console message
  - **CP or CMS commands**
  - **Rexx** EXECs, for example:
    - Send email
    - Send SNMP trap
    - Clean up a disk
  - **Write** data to a **TCP/IP** address/hostname and port
    - Send data to a syslog daemon/server
    - Send to any log analytics processor

Executing Actions

- **Dynamically include data** about the triggering event
  - Available to the action via substitution variables
- **Limit** the number of times an **action** is taken in a specified period of time
  - Avoid executing action repeatedly
  - Take a different action when the limit is reached
- Take multiple actions based on one message, event, schedule, etc.
  - Chain actions together
- Execute the action on another LPAR running Operations Manager
  - Communication is IP-based
  - **Does not require SSI**

Dynamic Configuration

- **Initial configuration** file loaded at startup
  - May imbed other configuration files
  - Filename can be a substitution variable for the system name
- Most **configuration options** can be **updated** while **Operations Manager is running**
  - Add, delete, or change:
    - Rules, actions, monitors, schedules, holidays, groups, user authorization
  - **Suspend or resume** rules, monitors, schedules
- Multiple methods
  - CMS command interface
  - (Re)load a new or updated configuration file
  - Commands in action routines
- **Sample configuration** files provided
  - Includes some of the demos in this presentation
    - Operations Manager configuration statements
    - **Sample Rexx** code

Features and Functions

- Monitor service machine consoles
- Monitor page space and spool usage
- Monitor system events
- Schedule events/actions
- Take actions automatically based on monitoring results
  - Includes taking actions on other z/VM systems with Operations Manager
- View and interact with monitored consoles from authorized user IDs
- Find and view spool files
- Dynamic configuration
- Separation of access control

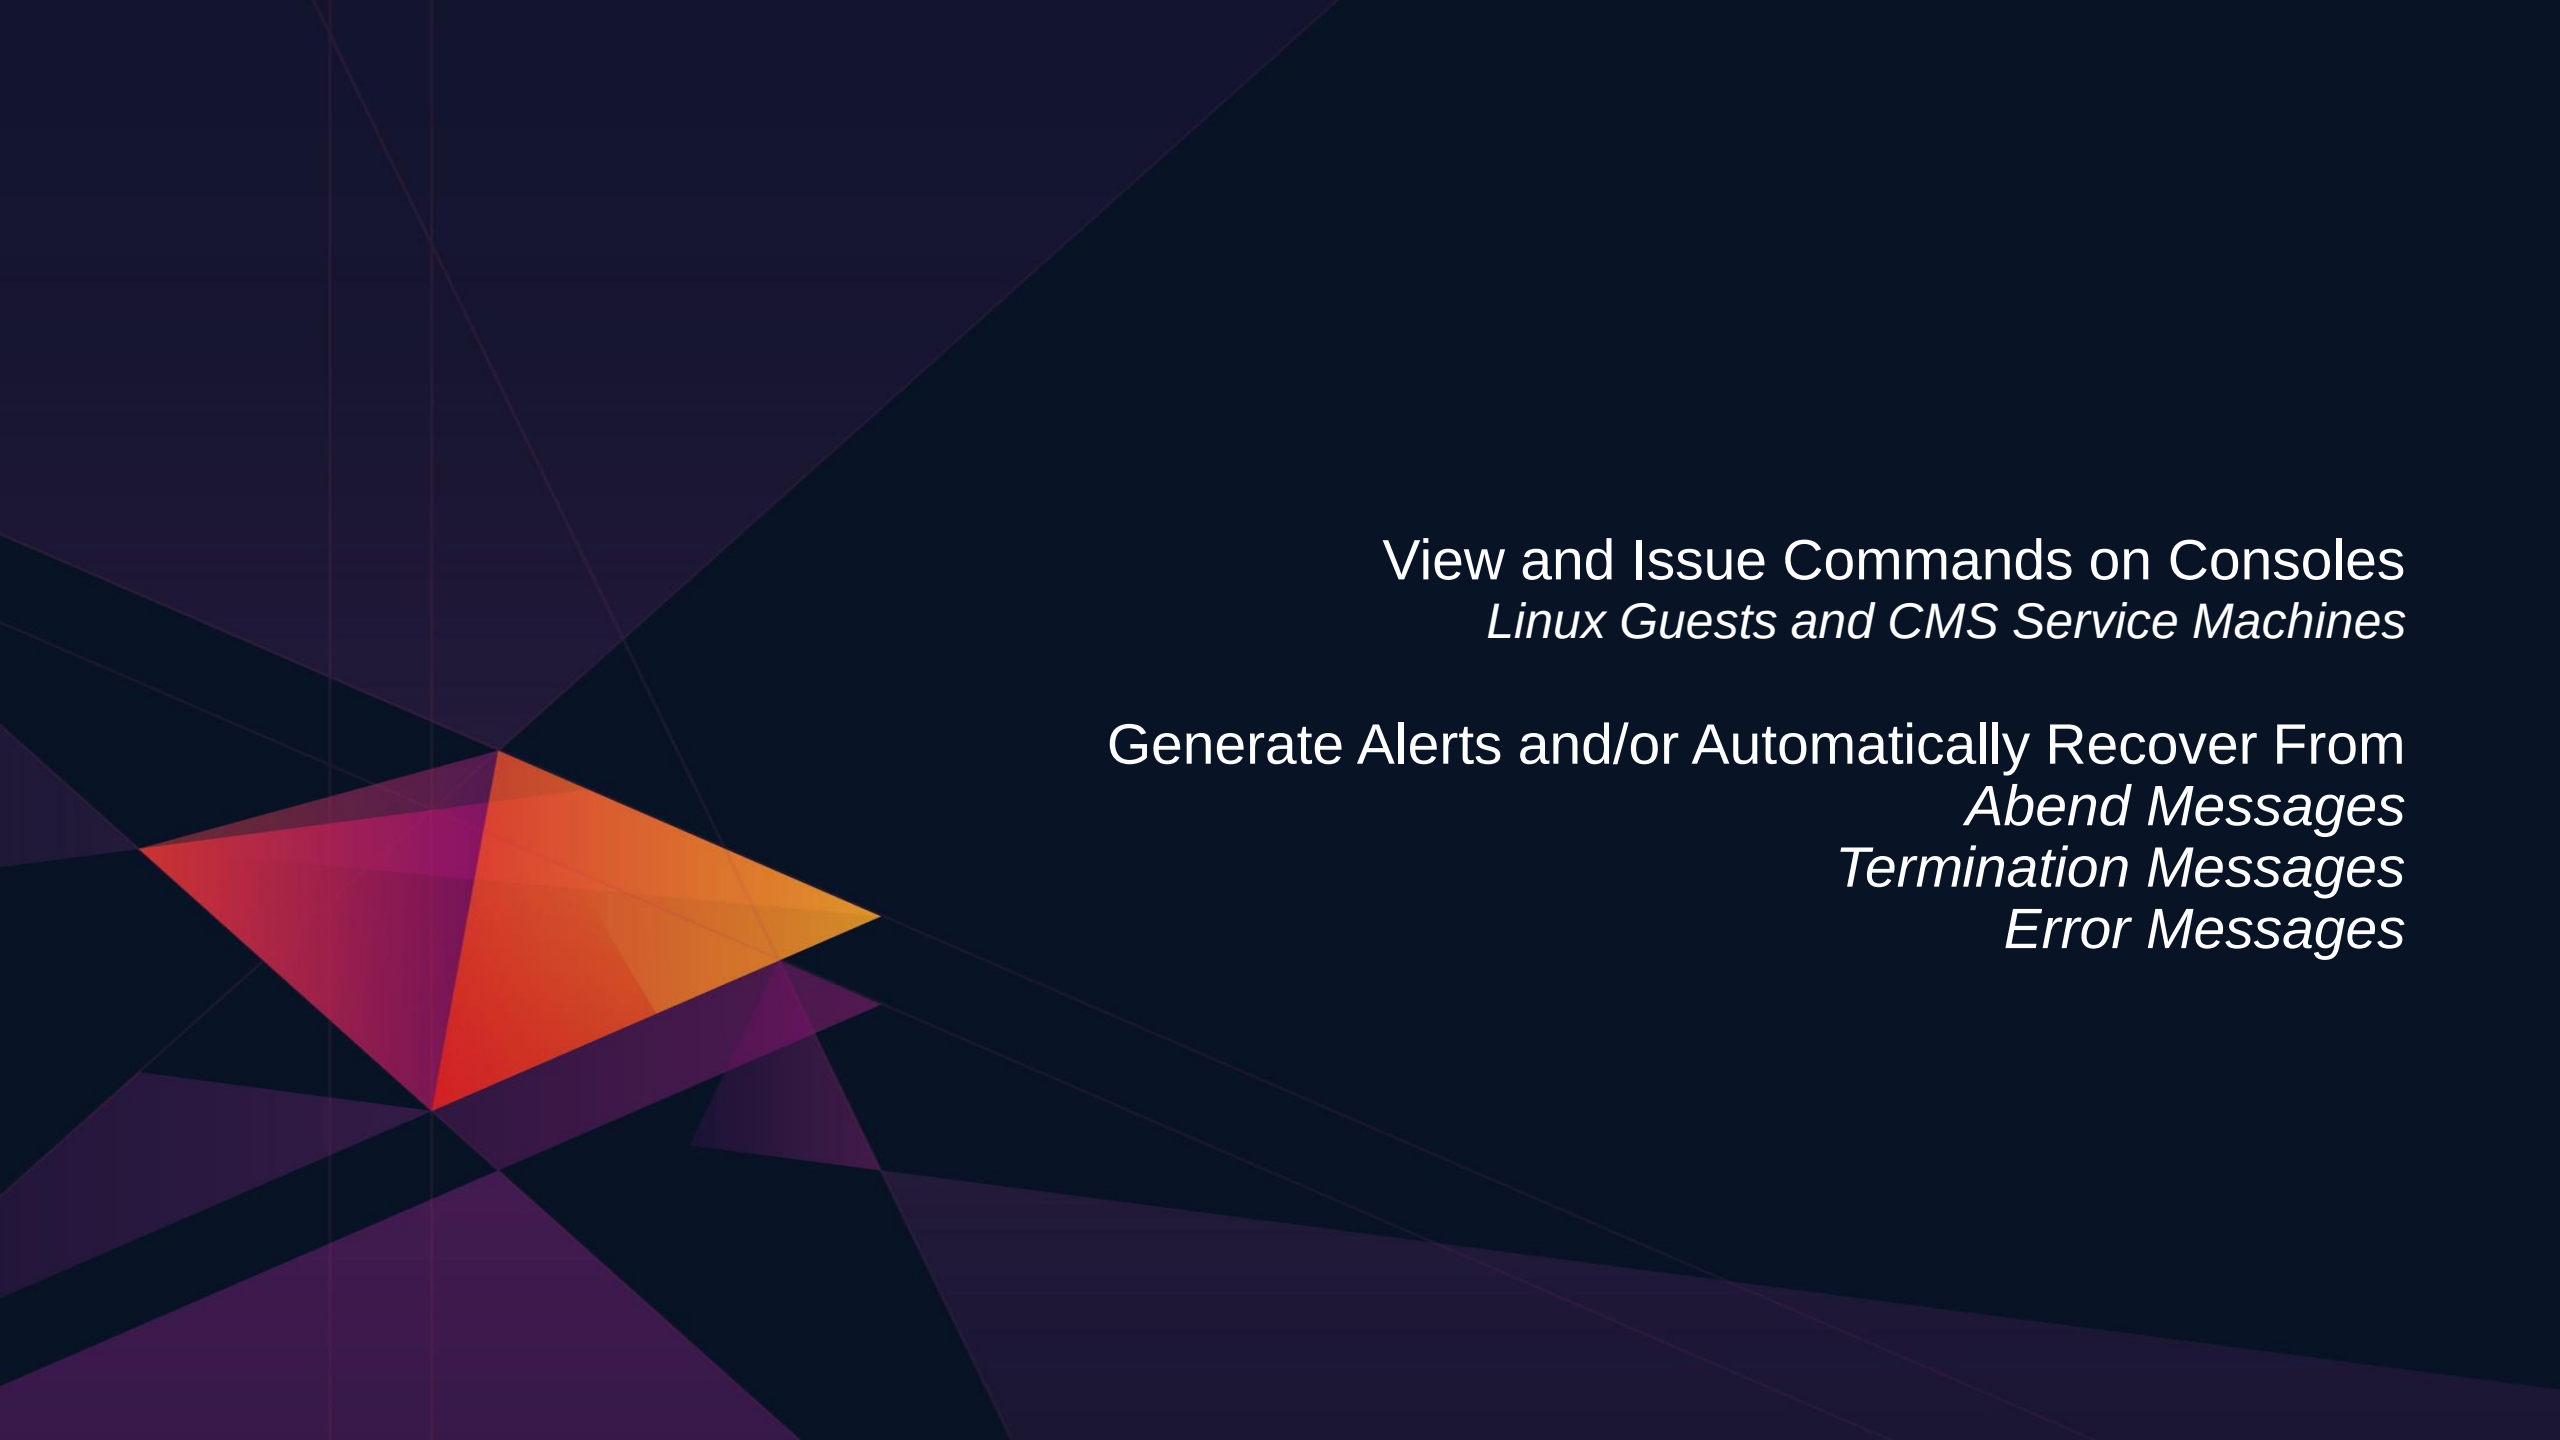
View and Issue Commands on Consoles
Linux Guests and CMS Service Machines
Generate Alerts and/or Automatically Recover From
Abend Messages
Termination Messages
Error Messages

Monitor Guest and Service Machine Consoles

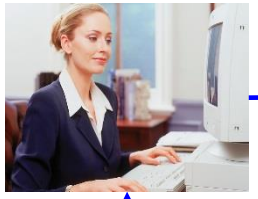

Operations Manager

slog Message 1
slog Message 2
slog Message 3
…

syslog data

OPER Message 1
OPER Message 2
OPER Message 3
…

OPERATOR

LNX Message 1
LNX Message 2
LNX Message 3
…

LINUX

TCP Message 1
TCP Message 2
TCP Message 3
…

TCPIP

TEST Message 1
TEST Message 2
TEST Message 3
…

Test data

Data space 1
TEST Message 1
TEST Message 2
…

Data space 2
LNX Message 1
LNX Message 3
…

Data space 3
OPER Message 1
OPER Message 2
…

Data space 4
TCP Message 1
TCP Message 2
…

Data space 5
slog Message 1
slog Message 2
slog Message 3
…

Filtered

Data space 6
OPER Message 1
LNX Message 1
LNX Message 2
TCP Message 1
slog Message 1
OPM Message 1
TEST Message 1
OPER Message 2
…

Daily log
OPER Message 1
LNX Message 1
LNX Message 2
TCP Message 1
slog Message 1
OPM Message 1

Unfiltered (with masking)

View and Interact with Consoles

- Authorized users can **view live consoles** of monitored service machines & guests
  - **Multiple users** can view the same console simultaneously
  - No need to logon to the user ID to see its console
    - No interruption of the user ID
  - No need to create and close console files of disjointed data
  - Test data and Linux syslog data treated as a “console”
  - Views can be defined to look at a group of consoles in one view
  - Can specify a date and time range for your view within currently available data
  - Can request a copy of the current console data for a user or set of users (disk or reader file)
  - Format of date in the view is based on requestor’s CP DATEFORMAT setting

- Full screen mode
  - **Scroll** up and down to view and search historical data
  - Auto scroll (on or off) as new output is displayed on the console
  - From command line, **issue commands** back to the monitored console
- Amount of data that is visible depends on specified or default data space size
  - Or date/time range specified

- Rules/actions may modify the view
  - **Suppress messages** from the console
  - **Hold or highlight messages** with color, blinking, etc.
- Authorized users can view the log file
  - Can also request a copy of the log file from today or a previous day

System Abend with No Console Data

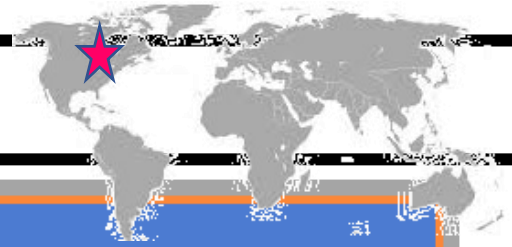

The Situation:

- **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

Capturing Linux Log Data

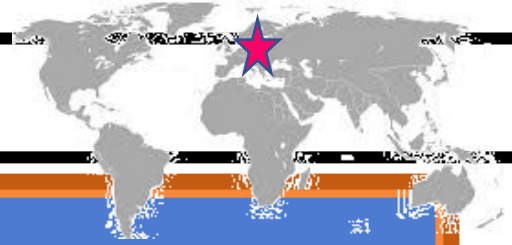

The Situation:

- 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

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

- Console data captured on z/VM and forwarded to Splunk
- Syslog data sent directly to Splunk

Monitor Service Machines

- Define rules to
  - Scan **console messages** for **text matching**
    - Includes column, wildcard, and exclusion support
    - Optionally restrict to specific user ID(s)
  - **Take actions** based on matches
- Multiple rules can apply to one message
  - Rules processed in order of definition in the configuration file
  - FINAL option available to indicate no additional rules should be evaluated

Error Messages on Linux IPL

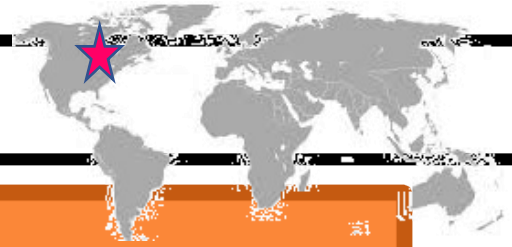

The Situation:

- 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

Coordinate Application Shutdown with z/OS

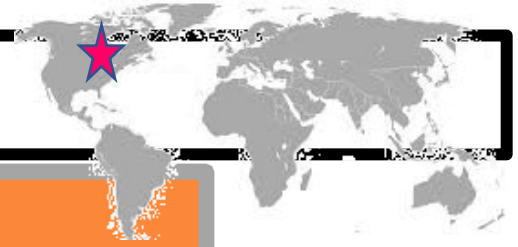

The Situation:

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

Initial solution

- 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

Sending Security Messages to Analytics

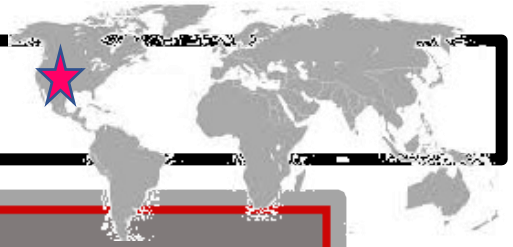

The Situation:

- 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

Final solution

Automation tool

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

Generate Alerts and/or Automatically Recover From

Critical User IDs or Guests Logging Off

Critical User IDs or Guests Enter Error State

Respond to System Events (Guest State Changes)

- Create monitors for z/VM system events (*VMEVENT)
  - Class 0, related to **user IDs**
    - 0 - Logon
    - 1 - **Logoff**
    - 2 - Failure condition (including **CP READ and Disabled Wait**)
    - 3 - Logoff timeout started
    - 4 - Forced sleep started
    - 5 - Runnable state entered (VM READ)
    - 6 - Free storage limit exceeded
    - 9 - Outbound relocation started
    - 10 - Inbound relocation started
    - 11 - **Outbound relocation complete**
    - 12 - Inbound relocation complete
    - 13 - **Outbound relocation terminated**
    - 14 - Inbound relocation terminated
    - 15 – Timebomb exploded
  - Optionally restrict to specific user ID(s)

Respond to System Events (System State Changes)

- Class 2 and 3, **related to SSI**
  - 7 – SSI Mode (Stable, Influx, Safe)
  - 8 – SSI Member State (Down, Joining, Joined, Leaving, Isolated, Suspended, Unknown)
- Class 4, related to **networking**
  - 16 – Device activated
  - 17 – Additional device activated
  - 18 – Device deactivated, connection to hardware still operational
  - 19 – Device deactivated, connection to hardware not operational
- Specify the **action** associated with the event
  - Actions specified are the same as those for schedules, console rules, and other monitors

Stopping and Restarting TCPIP

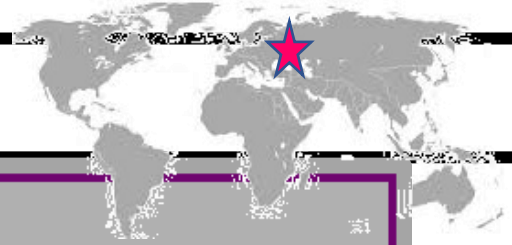

The Situation:

- 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

Initial solution

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

Final solution

Monitoring & automation tool

- Monitor for CP event indicating TCPIP has logged off
- Automatically XAUTOLOG it
- Easily bounce TCPIP as needed without relying on operations staff

Generate Alerts and/or Automatically Recover From

Spool Space Approaching Full

Page Space Approaching Full

Monitor Page and Spool Usage, View Spool Files

- Create page and spool space monitors to trigger actions when
  - Percent of spool usage falls within a **specified range**
  - Percent of spool usage increases at a specified rate
  - Percent of page space usage falls within a specified range
  - Percent of page space **usage increases** at a specified rate
- Actions triggered can be the same actions used by console monitoring
- For spool files, authorized users can
  - Use **full screen interface to list of spool files** based on one or more attributes
    - Owner
    - Size
    - Date created
  - From the list, the user can
    - **Sort** the list on any of the available columns
    - **View the contents** of an individual spool file
    - **Purge,** transfer, or change a spool file
  - Includes information on spool volume name(s) where each spool file is located
    - Easily find all spool files on a specific spool volume

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 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

Schedule Automated System Maintenance Procedures

Monitor for Rules, Monitors and Schedules Not Triggered

Spool Cleanup Based on Policies
Backups
Disk Cleanup
Orderly Startup and Shutdown

Schedule Events and Actions

- Define schedules
  - Hourly, daily, weekly, monthly, or yearly, nth weekday of the month
  - Once on specified month, day, year, and time
  - Based on ISO week definitions (week number; even, odd, first, last week)
  - At regular intervals
    - Every x hours and y minutes
  - Within a specified window of time
    - Specify start time
    - Specify conflicting schedules
    - Specify maximum time to defer this schedule
  - Within limits
    - Restrict to specific days of the week: Monday through Sunday plus holidays
    - Restrict to certain hours of the day
- Specify the action associated with the schedule
  - Actions specified are the same as those for console rules and all other monitors

Idle Monitors

- Define idle monitors
  - Watch for idle rules, schedules, and monitors
    - Rule, schedule, or monitor **not triggered *n*** number of times within specified period of time
- Specify the action associated with the idle monitor
  - Actions specified are the same as those for schedules, console rules, other monitors

SSI vs non-SSI Considerations

Operations Manager - non-SSI Environment

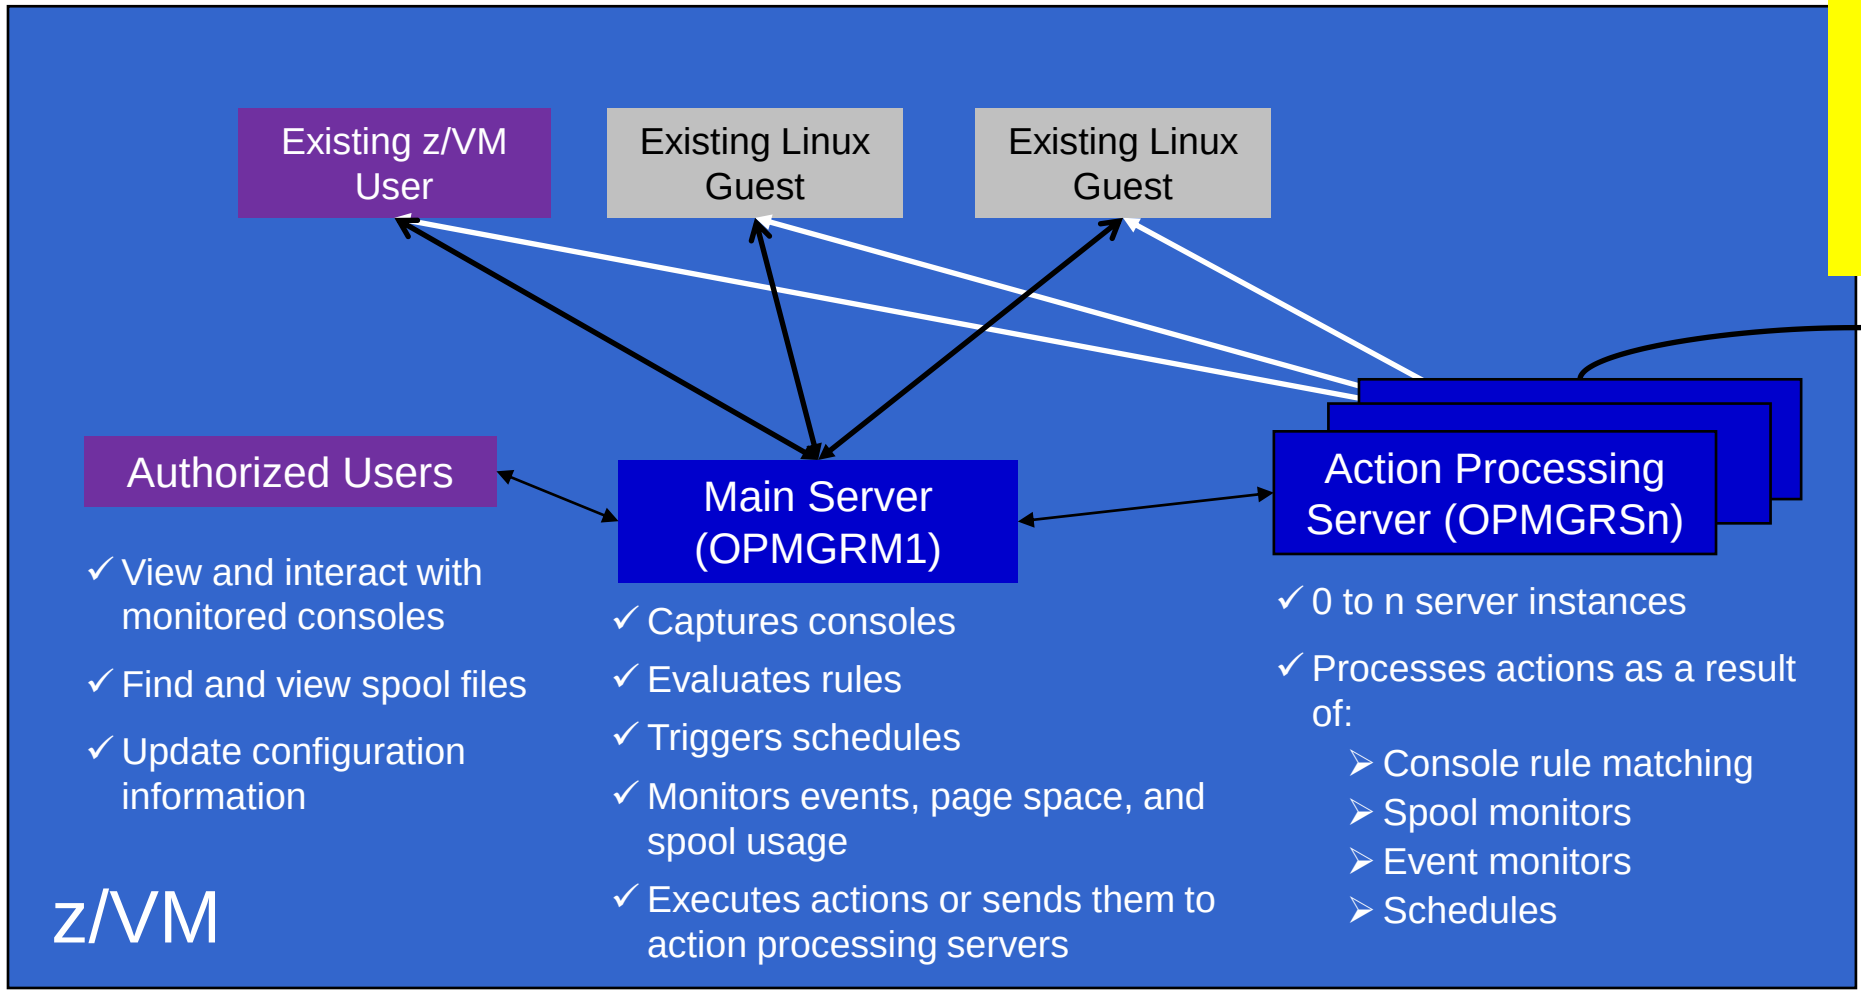

SSI Considerations

Console Monitoring

SSI Considerations for Console Monitoring

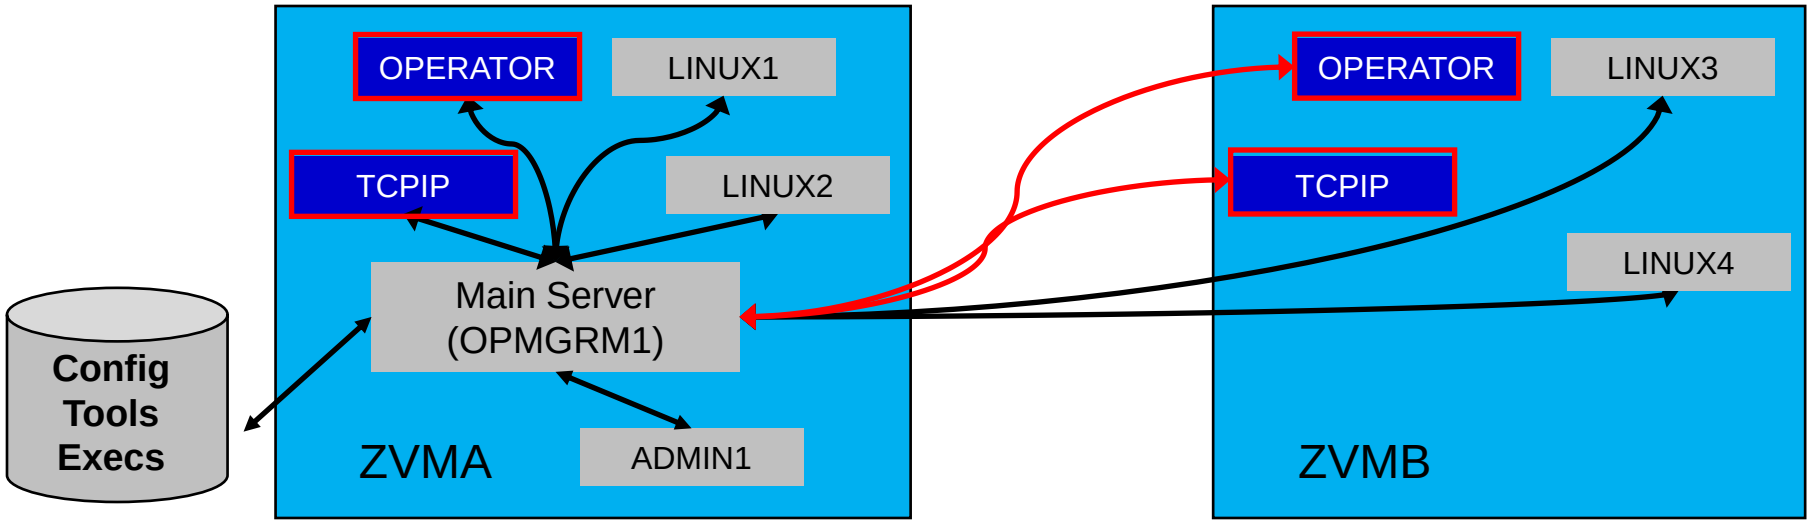

Option 1

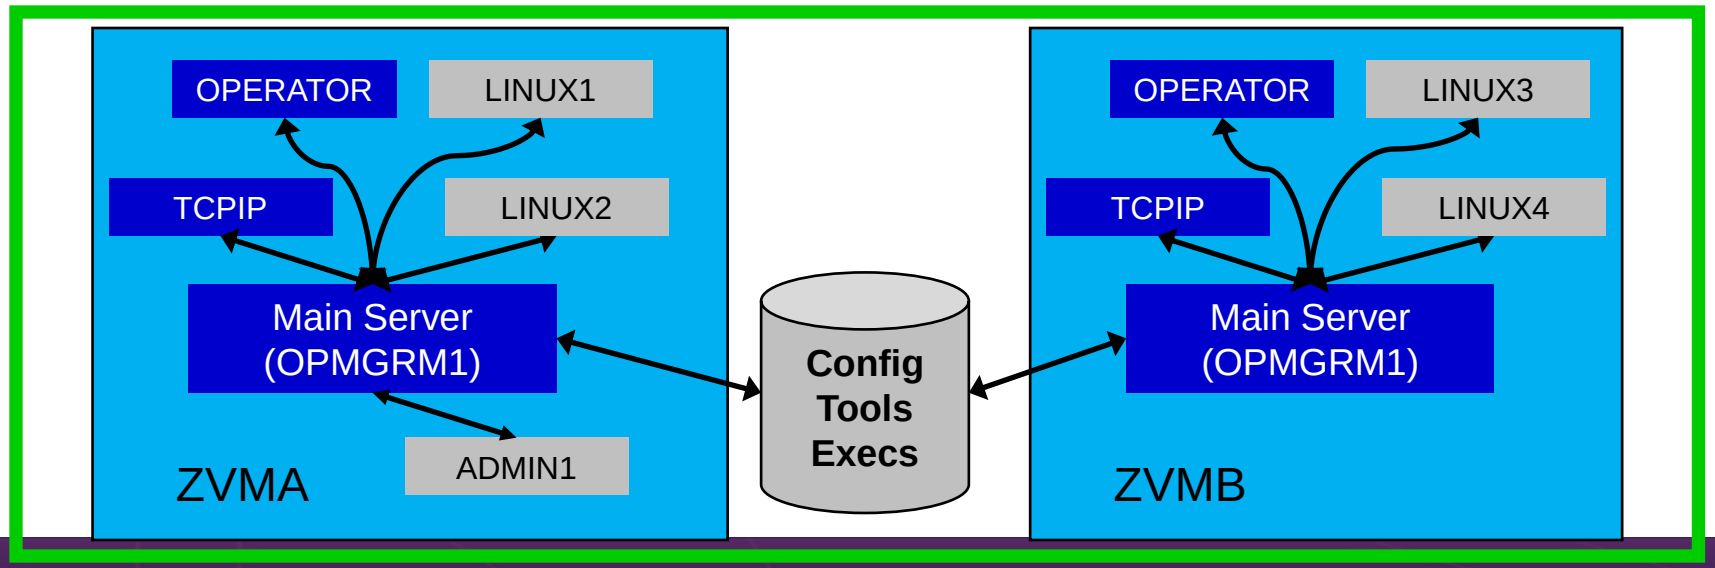

Option 2

Recommended

Single Config User

Multiconfig User

Operations Manager in SSI Cluster - Example

- Multiconfiguration users: OPMGRM1, OPMGRSn, OPERATOR, MAINT
- Single configuration users: LINUX1, LINUX2, OPERCONS, DEMOADMN
  - May relocate OPERCONS and DEMOADMN manually (supported) or via VMRELOCATE (unsupported, but you can make it work)

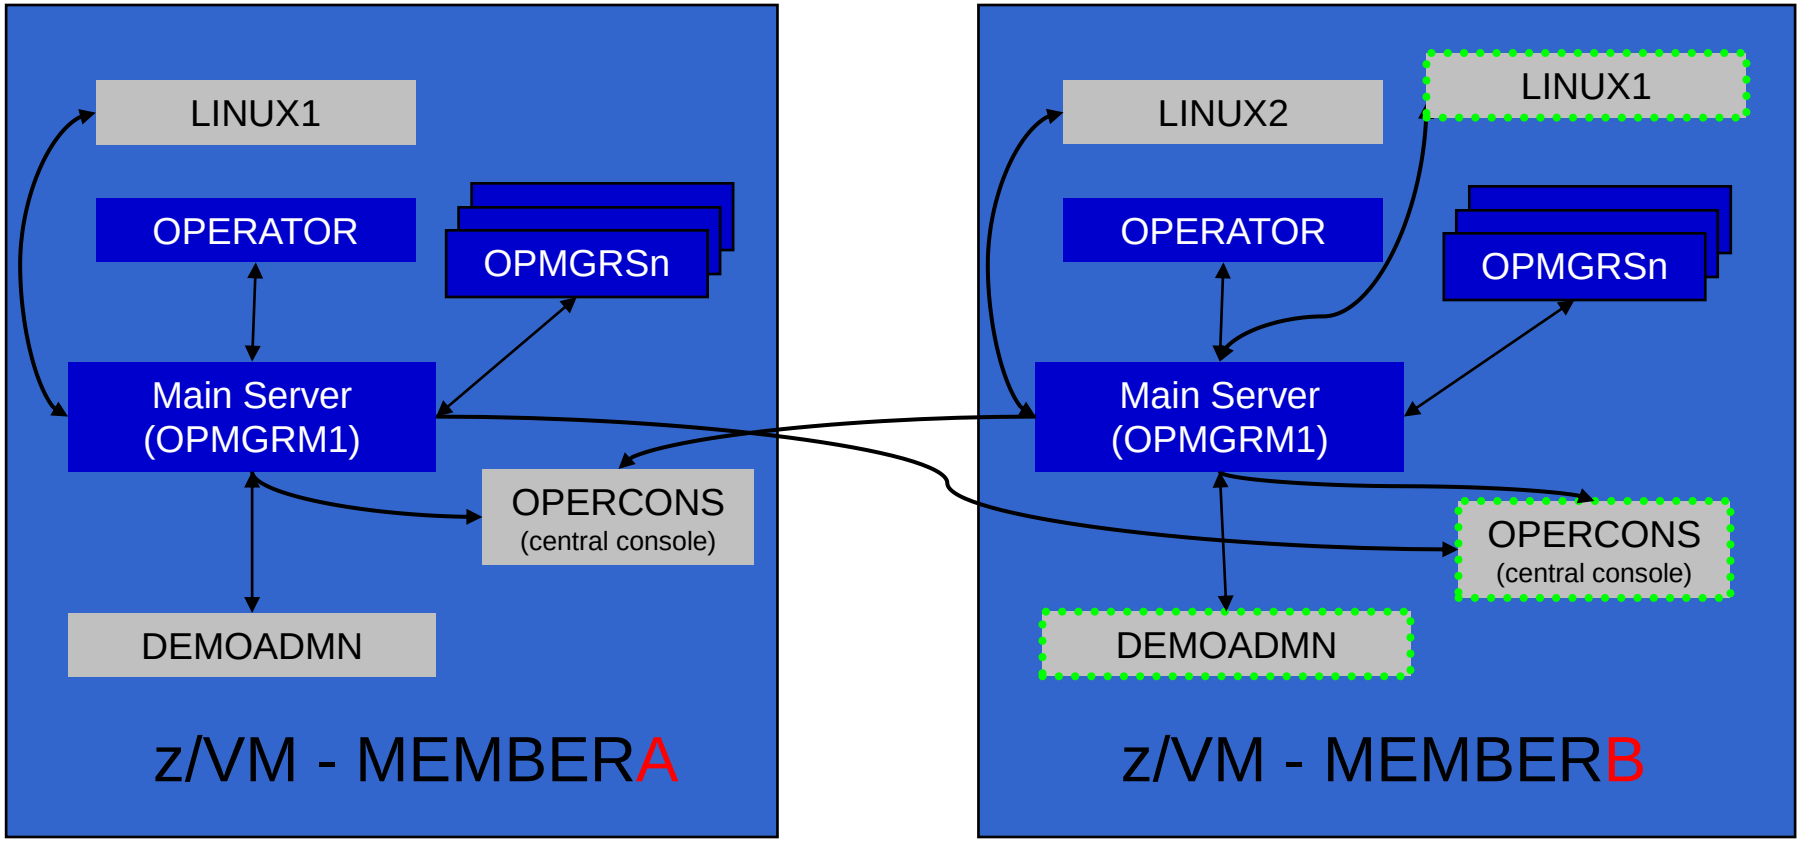

Relocating OPERCONS and DEMOADMN (CMS Users) …

- VMRELOCATE for CMS user IDs not officially supported
- Can be done for some CMS users
  - Create single configuration user ID for z/VM system disks
  - Copy MAINT 190, 19D, 19E to minidisks owned by this new user ID
    - Use DDR for 190 since it needs to be IPLable
  - Relocateable CMS user must IPL from identical NSS (CMS) or minidisk (190)
    - Use SPXTAPE to copy CMS NSS
      - VMRELOCATE uses checksum of NSS to determine if identical
      - CMS NSS includes date/time it was loaded
    - Or, have relocateable CMS users IPL 190 instead of IPL CMS

OPERCONS DIRECT

```
USER OPERCONS …
…
OPTION CHPIDVIRTUALIZATION ONE
…
IPL 190
…
LINK CMAINT 0190 0190 RR
LINK CMAINT 019D 019D RR
LINK CMAINT 019E 019E RR
…
```

PROFILE EXEC

```
/* PROFILE EXEC for OPERCONS */
…
'SET RELPAGE OFF'
…
```

… Relocating OPERSSI and DEMOADMN (CMS Users)

- Beware
  - It's worth repeating … VMRELOCATE for CMS user IDs not officially supported
  - All members of the cluster must be kept at same z/VM (or at least CMS) code level
  - If IPL 190, will use more memory as each user ID will have private copy of CMS
  - SET RELPAGE OFF may have a negative impact on overall system performance
  - Only works for "basic" CMS users
    - All relocation rules still apply
    - E.g. user IDs connecting to VMCF or IUCV can't relocate

4910004864

Monitor Service Machines - Considerations

- Consoles received by Operations Manager via SECUSER or OBSERVER
  - Prefer SECUSER
    - OBSERVER won't detect CP and VM READ messages
    - Output of actions on OBSERVEd console may not be viewable in console
  - OBSERVER allows Operations Manager to receive console output even if user is logged on
- SSI allows SECUSER and OBSERVER across members of cluster in some situations
  - Content does not contain member name information
  - Rules, actions, and users wouldn't be able to distinguish between IDENTITY users on multiple members
  - Creates single point of failure on one member
- Recommendation for z/VM Single System Image environments
  - Have all consoles monitored by an Operations Manager server on the same member as the monitored guest (i.e. all Operations Manager servers are IDENTITY users)
    - Requires action processing servers (OPMGRSn) to be on same member as main server
  - Share configuration data on 198 minidisk owned by OPMGRM1 but in IDENTITY section (not SUBCONFIG section)
    - OPMGRM1 links the disk read only, files updated from system programmer user IDs
    - Main configuration file unique to each member
    - Imbed common file(s) used by all members
  - Request a copy of the current console of a remote user
    - `SMSG OPMGRM1 at membername VIEWCON USER userid MODE RDR`

SSI Considerations

Page Space Monitoring
Spool Space Monitoring
Viewing and Managing Spool Files

SSI Considerations for Page and Spool Space Monitoring

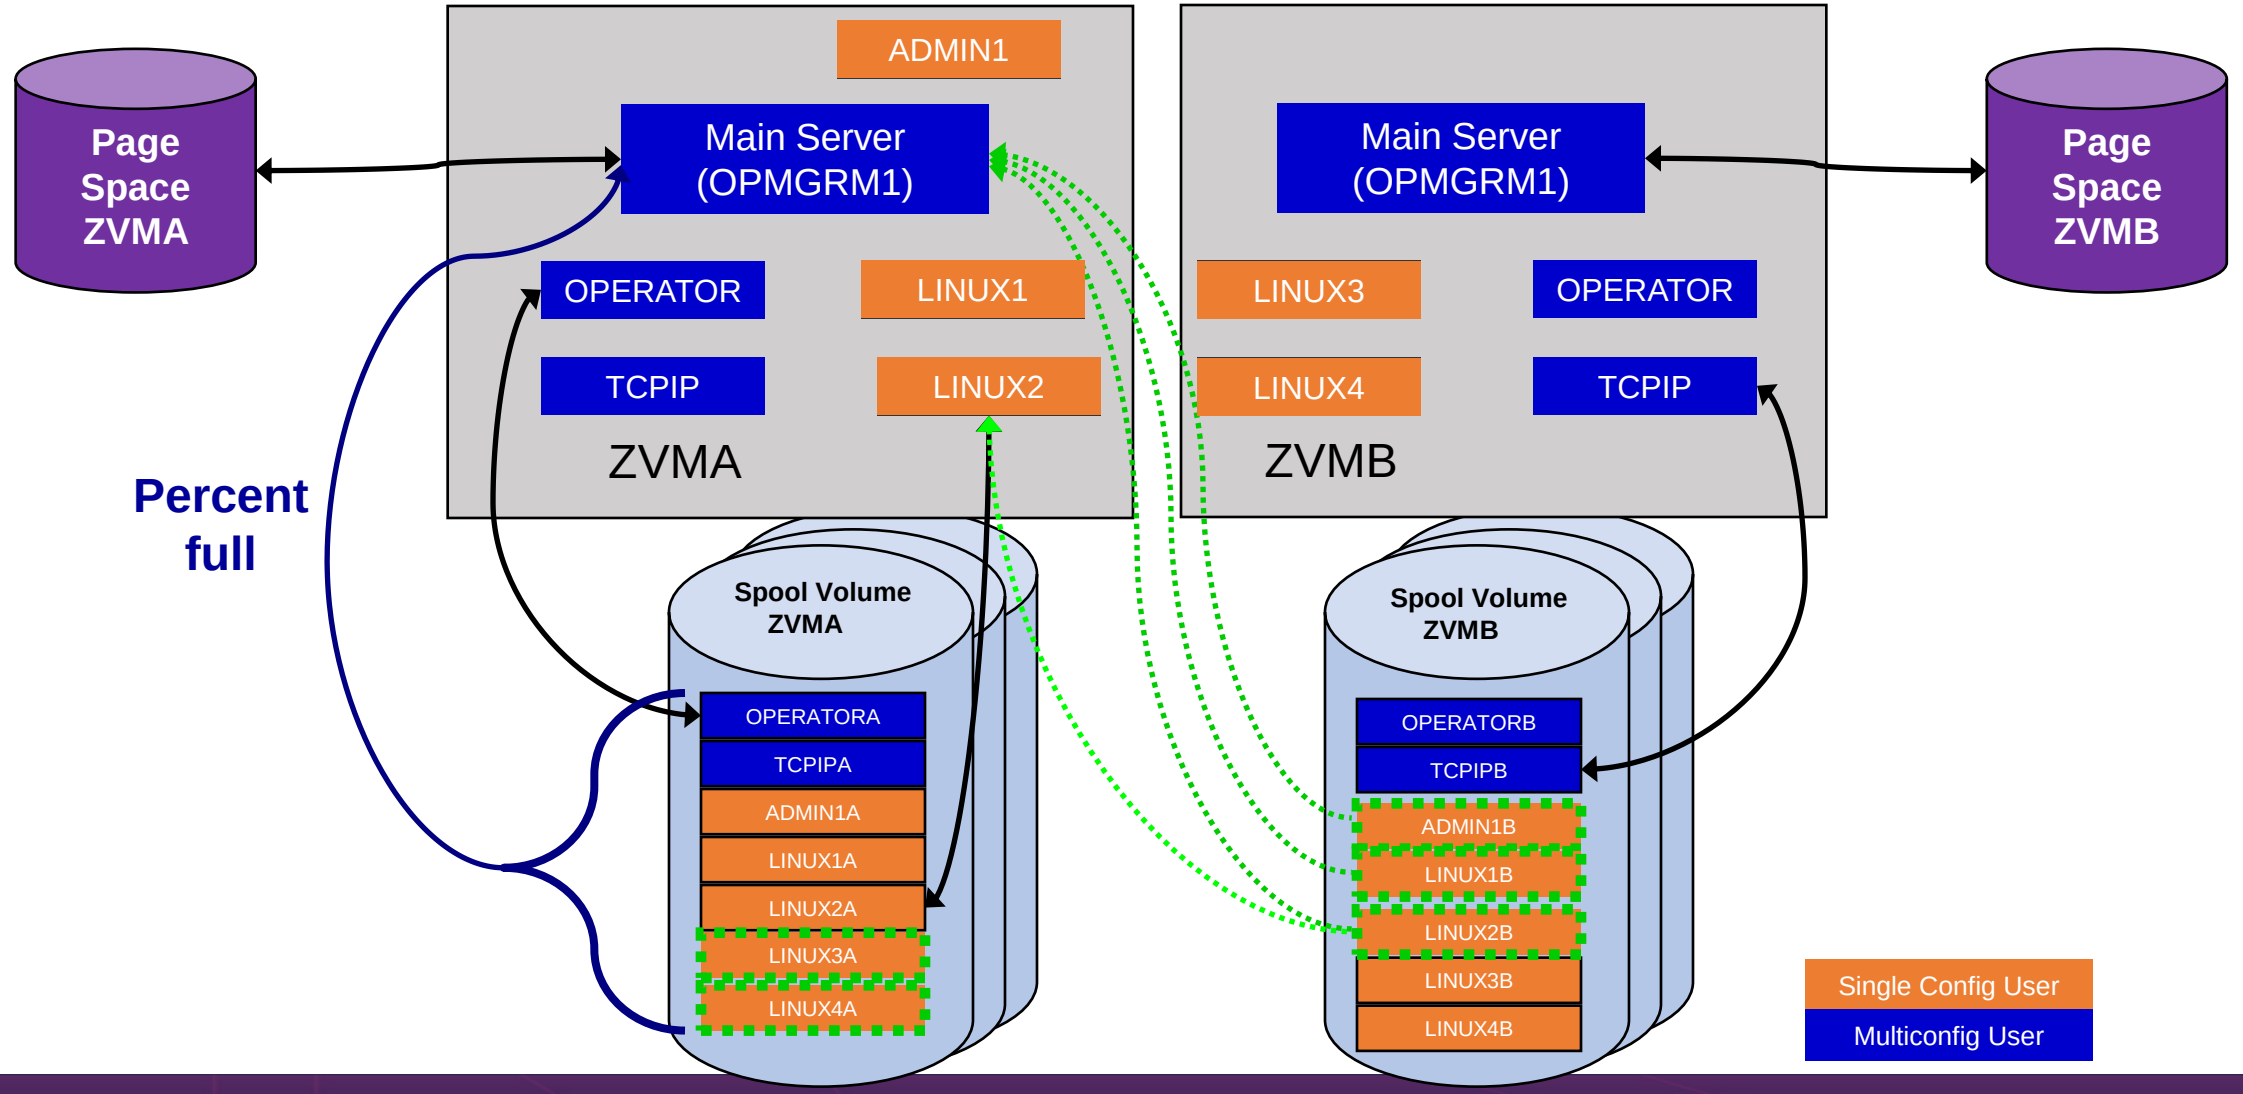

Spool and Page Space Monitoring - Considerations

- Page space is local
  - Separate space for each member and only visible to the local member
  - **No impact from SSI**
- Spool data
  - Spool files are placed on spool volumes owned by the member where the spool file was created
  - Users see their own spool data no matter where they are logged on and where the data was created

Spool and Page Space Monitoring - Considerations

Users and applications (like Operations Manager) who can see all spool files need to be aware:

- Spool data for **multiconfiguration** users
  - Only spool files owned by the local instance of that user are visible on the local member
  - No visibility to spool files owned by other instances of that user on other members
- Spool data for single configuration users:

Single configuration user status	All spool files created on **this** member	PRT/PUN files created on **other** members	RDR files created on **other** members
User logged off	Visible	Visible	Not visible
User logged onto **this** member	Visible	Visible	Visible (but not on local spool volumes)
User logged onto **another** member	Visible	Visible	Not visible

Spool and Page Space Monitoring - Considerations

- Recommendation
  - Have an Operations Manager server on each member to monitor spool and page space
  - Be aware of spool files visible in Operations Manager but not resident on this member's spool volumes
    - Indicated with "+" in VIEWSPL

SSI Considerations

Managing Configuration Files

Managing Configuration Files

- Put all configuration files on a shared disk
  - Default is 198 disk for OPMGRM1 – in **IDENTITY** section
  - Alternatively SFS
- Create a main configuration file with authorizations and system settings – shared by all members
  - All Operations Manager servers on all members load this file
- Create a common configuration file used by all members
- Imbed a unique configuration file based on the system name of this member

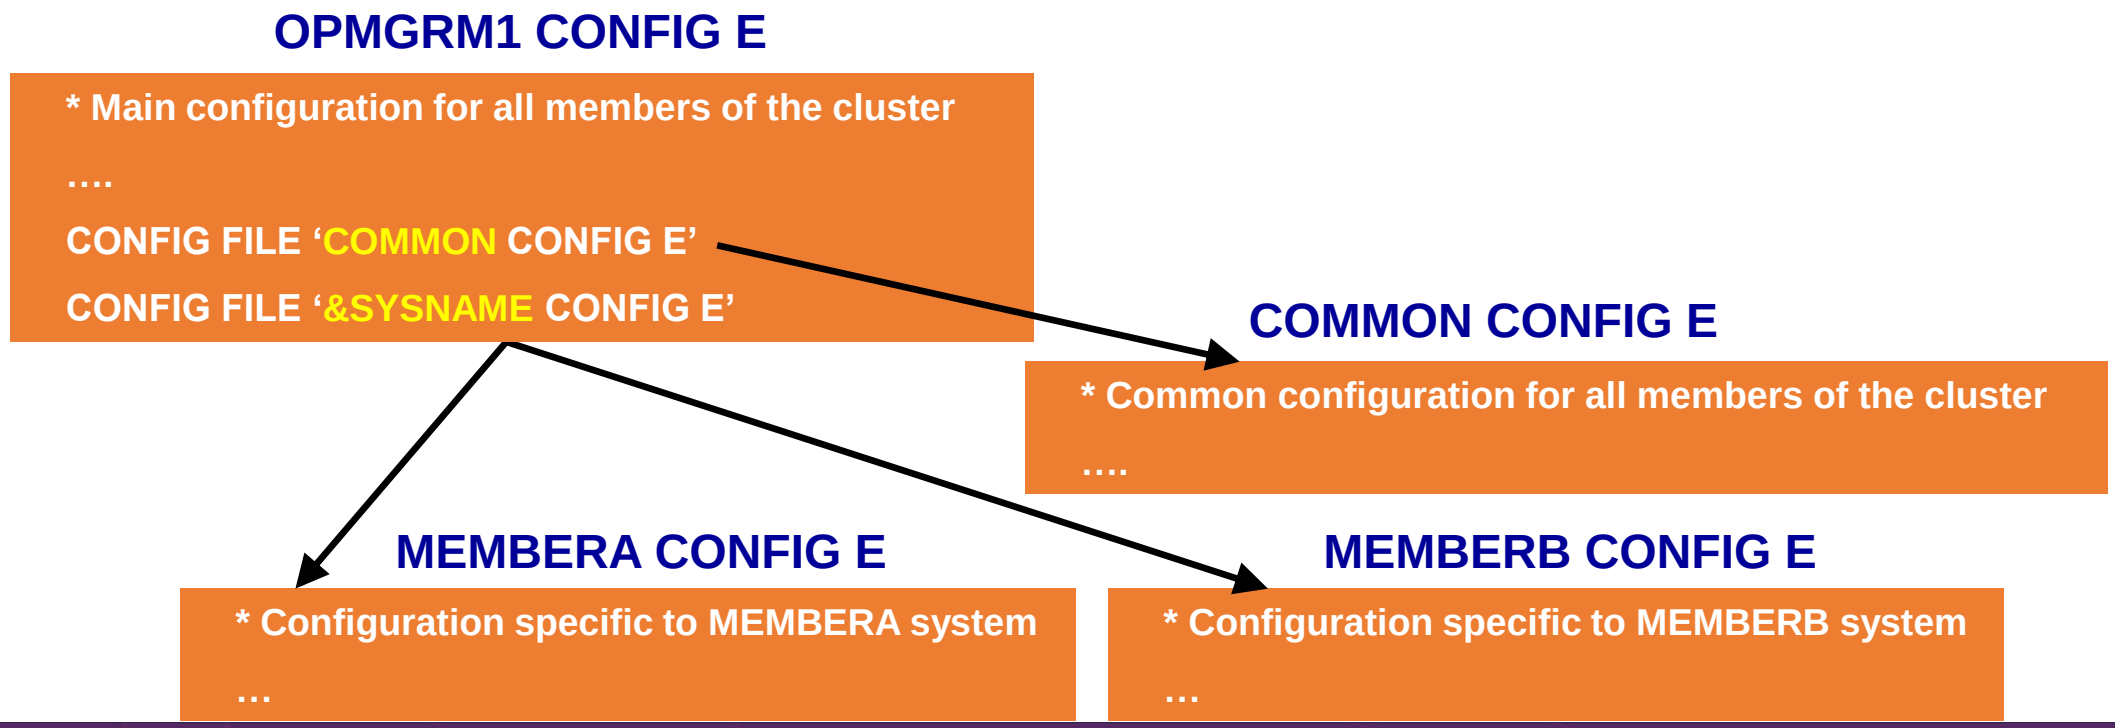

Other Scenarios

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

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

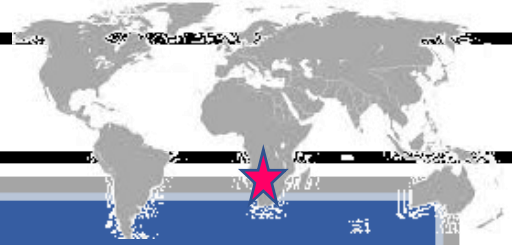

The Situation:

- 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 IBM Z alerts **included**

Initial solution

None

- z/VM and Linux alerts sent via email or to central console only
- Mainframe operations team not able to participate in enterprise solution

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

Shared Monitoring and Automation Across LPARs

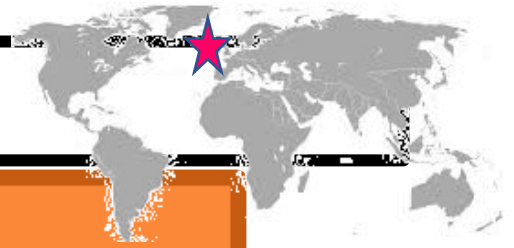

The Situation:

- 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

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 reload on all SSI and non-SSI systems**

Including Performance Data with z/OS Processes

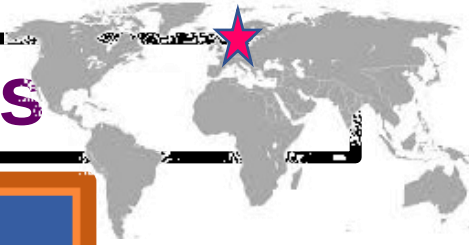

The Situation:

- Collecting **performance data** on z/VM (Performance Toolkit)
- All mainframe performance data processed on **z/OS**
- Want to **include z/VM and Linux** data

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**
- **When complete, send message to z/OS**
- **z/OS: FTP file from z/VM**
- **z/OS: FTP message to z/VM indicating successful file retrieval**
- **z/VM: erase the file**

Central Operations Console

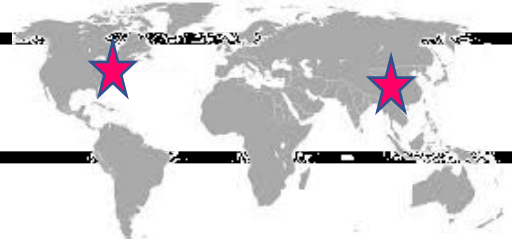

- Already have z/OS console in operations center
  - Alerts, important messages
  - Operations staff watching consoles and taking actions
- Want **one** console for all **z/VM** LPARs and **Linux** guests
  - Operations staff sees **only important messages** on central console
  - **When needed** can 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

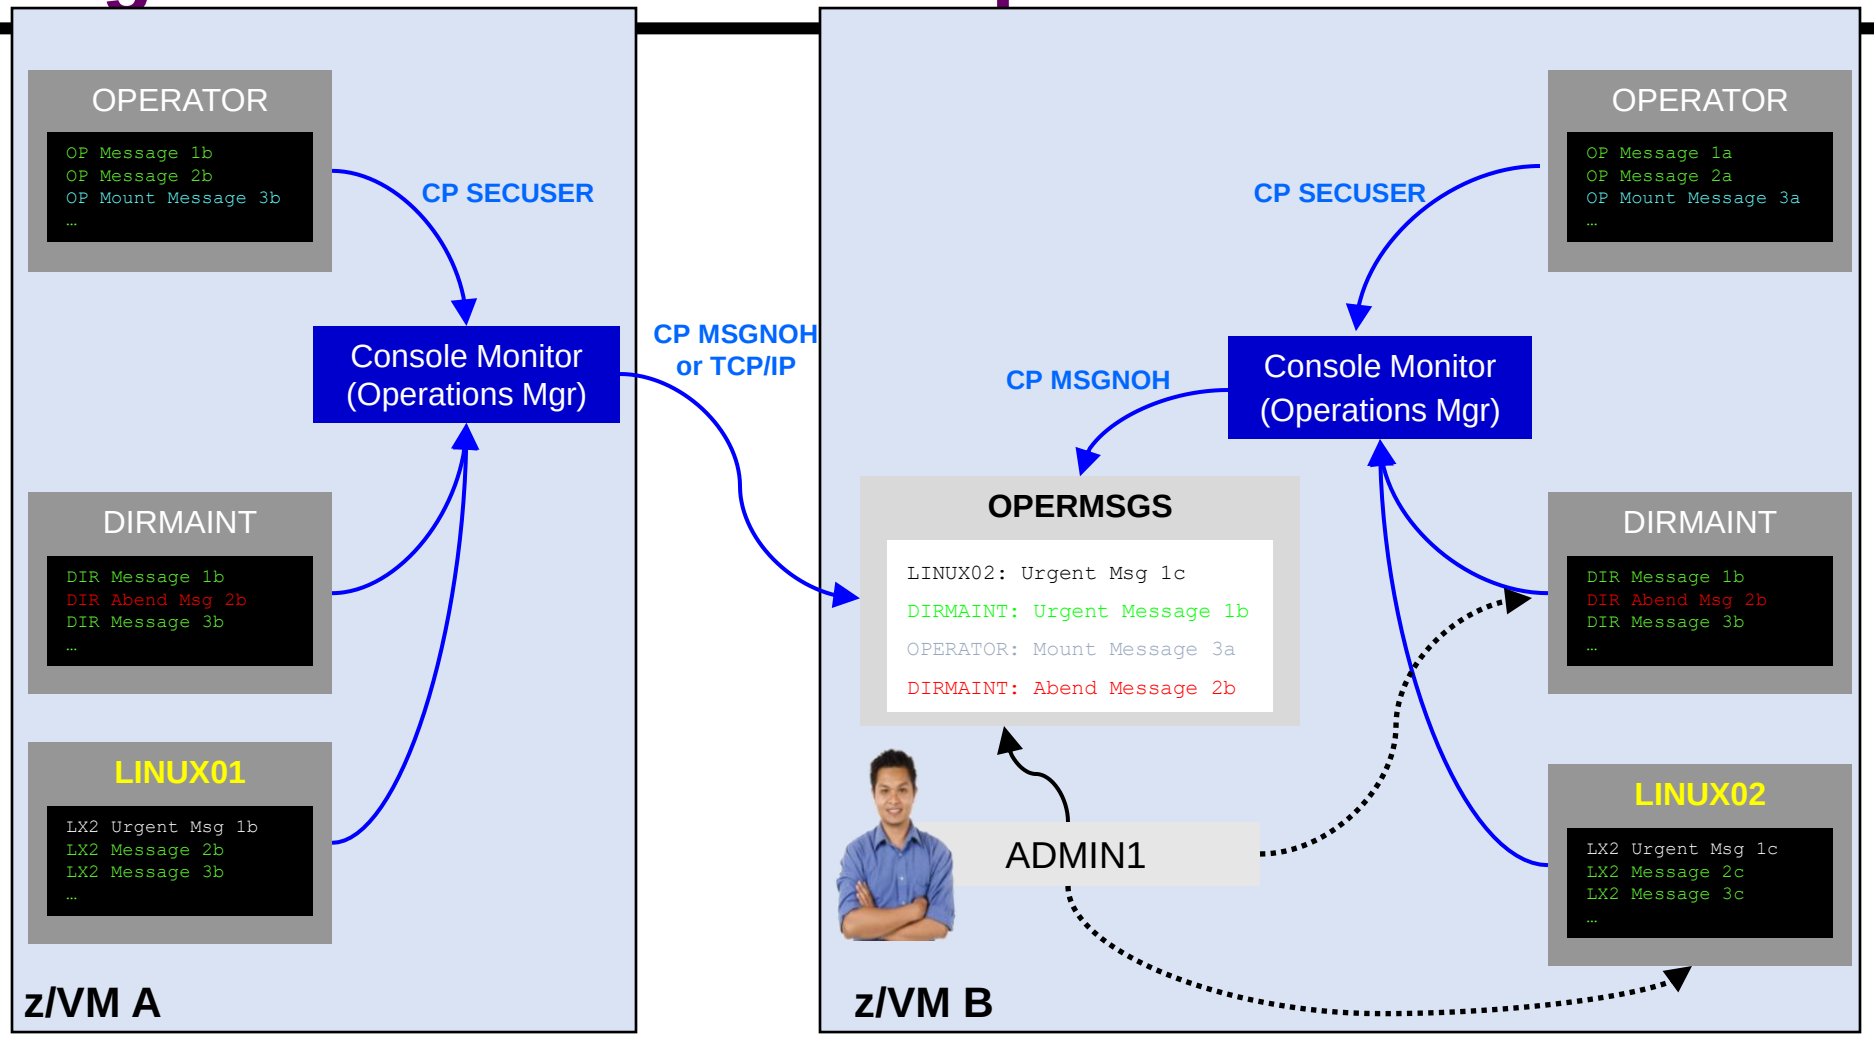

Single System Image (SSI) supported but not required

Graceful Shutdown of z/VM from GDPS

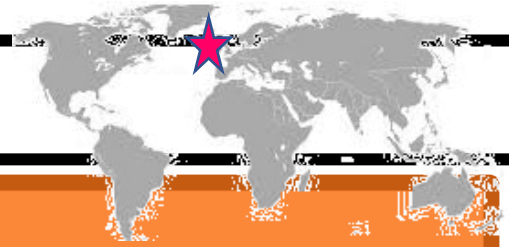

The Situation:

- 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

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**

Perform Weekly System Healthcheck

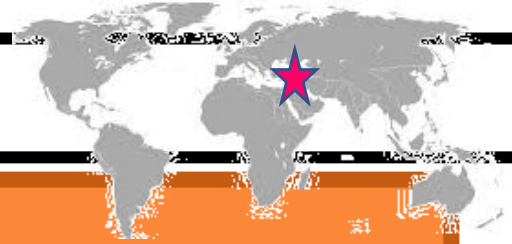

The Situation:

- 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

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

Summary and References

Recommended Practices – 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

VIEWCON
VIEWLOG
Log file

Rules
Event monitors
Spool/page monitors

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

Schedules
SFPURGER
Rules
Backup Manager

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

Summary

- Use Operations Manager to
  - **Automate** daily operations
  - **Integrate** your z/VM and Linux on IBM Z environment with existing enterprise monitoring and **alerting**
  - Prevent problems rather than react to them
  - Automate reactions to problems when they can't be prevented
  - **Improve problem determination** procedures
  - Increase programmer and operator productivity
  - Continue to monitor locally with improved management of clusters
- Sometimes several alternatives for monitoring for the same event
  - Console message (rules)
  - Scheduled healthchecks (schedules)
  - User ID status changes (event monitor)
- Actions allow integration with other platforms and products

Reference Information

- Web sites
  - Product page: https://www.ibm.com/products/operations-manager-for-zvm
    - Publications, presentation, white papers
    - Pre-requisites
    - Support
- White papers on Operations Manager website (Resources tab)
  - Routing Linux syslog data
  - Sending alerts from Operations Manager to Netcool/OMNIbus
  - Using Shared File System to store Operations Manager configuration files and automation EXECs
  - Automatically logging on a user at Linux system boot time for easier console management and action execution
- **IBMVM** Mailing list
  - http://listserv.uark.edu/archives/ibmvm.html

धन्यवाद

Hindi

多謝

Traditional

감사합니다

Korean

Спасибо

Russian

Ndzi khense ngopfu

Tsonga

Gracias

Spanish

Thank You

English

شكراً

Arabic

Grazie

Italian

Obrigado

Brazilian Portuguese

Danke

German

Ke a leboha

Tswana

多谢

Simplified Chinese

Merci

French

நன்றி

Tamil

ありがとうございました

Japanese

ขอบคุณ

Thai

Demonstration Scenarios

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