

IBM z Systems

Backup Strategies for z/VM and Linux on z Systems

Tracy Dean
IBM
tld1@us.ibm.com

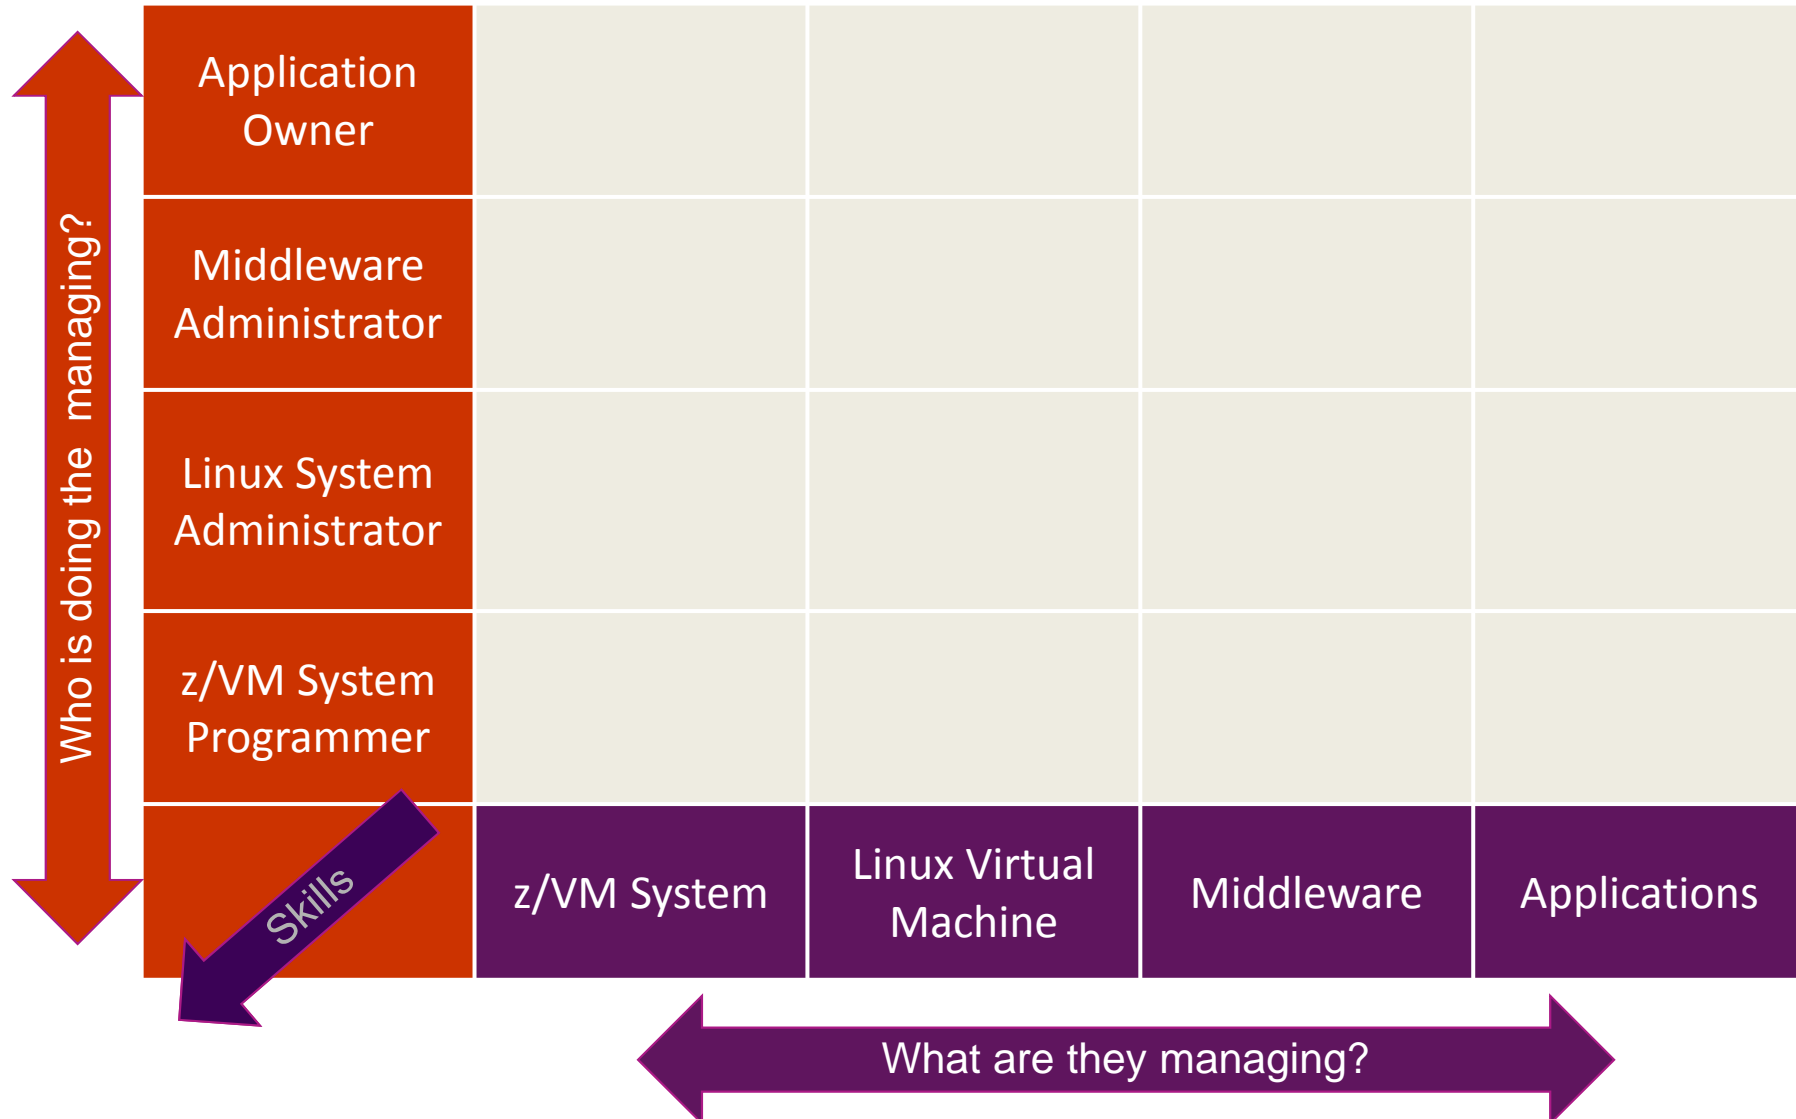
April 2015




Agenda

- Positioning
- Recommended practices and available options
 - Backing up and restoring z/VM
 - Backing up and restoring Linux on z Systems
- Backing up and restoring data in a z/VM SSI cluster
- Overview of IBM products
 - Backup and Restore Manager for z/VM
 - Tape Manager for z/VM
- Backup scenarios
 - Live demos
 - Configuration options and sample code
- Summary and reference information

Three Dimensions of Systems Management



Three Dimensions of Systems Management



IBM z/VM Management 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
 - New release (V1.3) announced February 24, 2015
 - Tape Manager for z/VM
 - Tivoli Storage Manager
- Automation and operational monitoring
 - Operations Manager for z/VM
 - Including integration with existing monitoring and alert systems
- Interactive provisioning and system resource management
 - IBM Wave for z/VM

IBM Infrastructure Suite for z/VM and Linux

- New IBM bundle/suite
- Announced and available September 2014
- Tools needed to manage the z/VM and Linux on z Systems infrastructure
 - Wave for z/VM
 - OMEGAMON XE on z/VM and Linux
 - Operations Manager for z/VM
 - Backup and Restore Manager for z/VM
 - Order Tape Manager for z/VM separately if plan to back up to tape
 - Tivoli Storage Manager Extended Edition
- Discounted price as a bundle
- Website:
 - <http://www.ibm.com/software/products/en/ibm-infrastructure-suite-for-zvm-and-linux>
- DeveloperWorks Wiki
 - https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/W9b511b099ded_4e32_abfb_ed8ce4da5b17

The background features a dark blue gradient with abstract geometric shapes in shades of purple and orange. A prominent orange and yellow triangle is located in the lower-left quadrant, pointing towards the center. The overall aesthetic is modern and professional.

Recommended Practices and Available Options

Recommended Practices – 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 (some?) Linux guests

- Operating system
- Applications
- Application data (maybe)

File level backup of Linux guests

- Configuration files
- Log files
- Tools

Recovery of z/VM system, including Linux guests


- Dependence on z/OS
versus
- Independent recovery

High Availability

Location A



Location B



High Availability and Backup/Recovery are **NOT** the Same


Location A



Location B



Does not address operational recovery needs



Recommended Practices – 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 (some?) Linux guests

- Operating system
- Applications
- Application data (maybe)


File level backup of Linux guests

- Configuration files
- Log files
- Tools

Recovery of z/VM system, including Linux guests

- Dependence on z/OS
versus
- Independent recovery


Image Level Backup/Recovery of z/VM and Linux Guests from z/OS



- **Image level backup and recovery of DASD volumes from z/OS**

- Existing z/OS procedures and tools in place
- Use existing tape devices
- Fast
- Doesn't include FCP-attached DASD
- Linux should be down
 - Flashcopy can minimize downtime
- Dependent on z/OS for recovery and DR
 - Is Linux workload critical – recovery required in parallel with z/OS in event of disaster?
- Using z/OS cycles (on general purpose processors) to back up z/VM and Linux

Image Level Backup/Recovery of z/VM and Linux Guests from z/VM



- **Image level backup and recovery of DASD volumes from z/VM**
 - **Low risk if z/VM is running – but not zero risk**
 - **Includes FCP-attached DASD (defined to z/VM as EDEVICES)**
 - **Volumes can not be DEDICATED to guest**
 - **Linux should be down**
 - **Flashcopy can minimize downtime**
 - **Recovery of z/VM and Linux independent from recovery of z/OS**
 - **Critical Linux workload recovered in parallel with z/OS in event of disaster**
 - **Faster recovery of z/VM and Linux overall**
 - **Backup software required on z/VM**
 - **Use z/VM cycles on IFL processors to back up z/VM and Linux**
 - **Requires mainframe attached tape devices**
 - **Share tape devices with z/OS – does not require both systems to be up**

What About DDR?

- DDR - DASD Dump Restore utility in z/VM
- Basic ability to copy data from one location to another
 - Command driven
 - Specify a source location
 - Specify a target location (disk or tape)
- Useful when copying/cloning minidisks or volumes
 - No ability to do file level backup/recovery
 - Be aware of “changing data” on active disks or volumes
- Very limited in terms of production level backup and recovery
- Advantages of Backup and Restore Manager for z/VM over DDR
 - File level backup and recovery
 - Incremental backups of z/VM (CMS and SFS) files
 - Cataloging of what has been backed up
 - Including full screen interfaces for finding backup data and restoring it
 - Automated expiration processing of catalog data and backup data on disk or tape
 - Flexibility to define a job once using wildcarding – future invocations of that job will back up any new data that meets the criteria
 - Invoke multiple service machines to share the backup task – completing the backup sooner
 - Integration with a tape management system – no need to manage tapes and tape mounts manually

Do I Need to Back Up **Every** Linux Guest ?

- It depends ...
- Is each guest image unique?
 - Are logs or other output stored within each guest?
 - Is configuration of each guest automated?
- Can a new guest be recreated from a golden image more easily than restoring it?

Is backing up just the “golden images” sufficient?

Recommended Practices – 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 (some?) Linux guests

- Operating system
- Applications
- Application data (maybe)

File level backup of Linux guests

- Configuration files
- Log files
- Tools

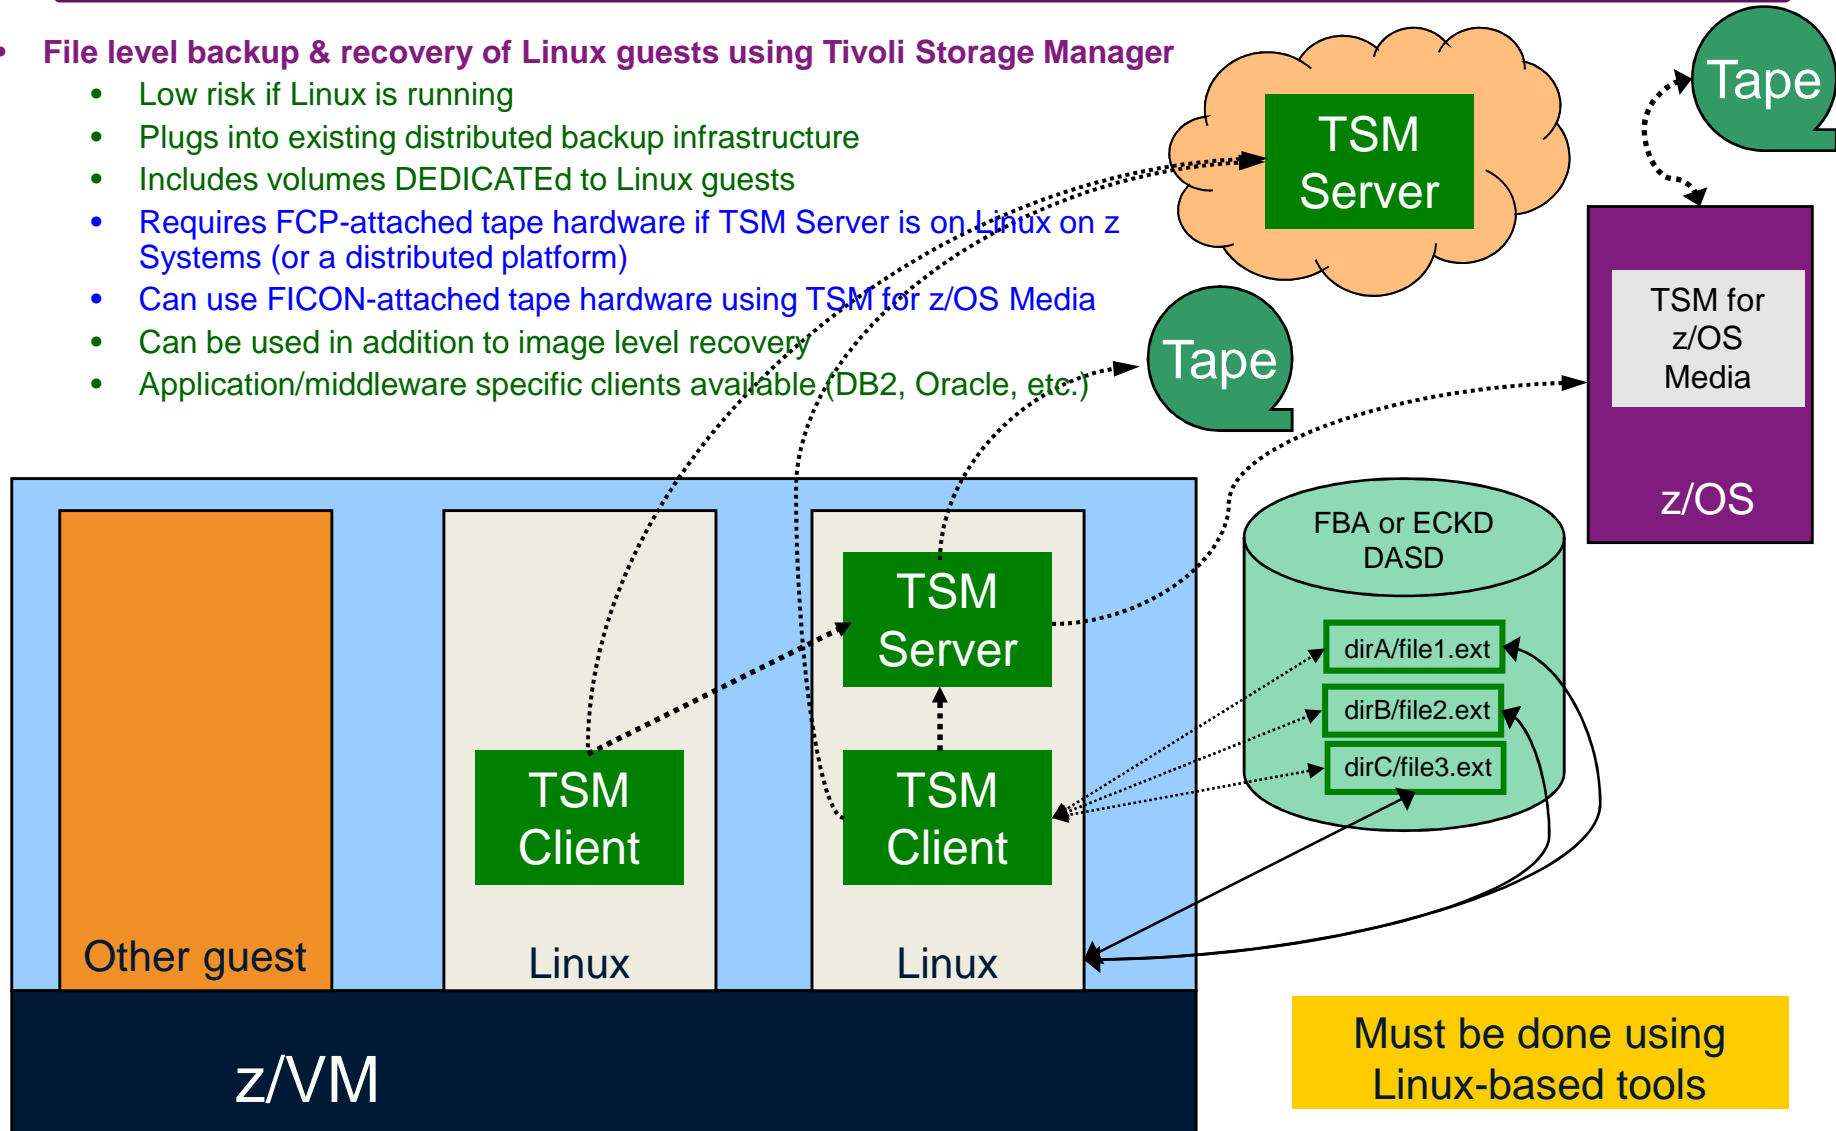
Recovery of z/VM system, including Linux guests

- Dependence on z/OS
- versus
- Independent recovery

File Level Backup and Recovery of Linux Guests

- File level backup & recovery of Linux guests using Tivoli Storage Manager

- Low risk if Linux is running
- Plugs into existing distributed backup infrastructure
- Includes volumes DEDICATED to Linux guests
- Requires FCP-attached tape hardware if TSM Server is on Linux on z Systems (or a distributed platform)
- Can use FICON-attached tape hardware using TSM for z/OS Media
- Can be used in addition to image level recovery
- Application/middleware specific clients available (DB2, Oracle, etc.)



Recommended Practices – 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 (some?) Linux guests

- Operating system
- Applications
- Application data (maybe)

File level backup of Linux guests

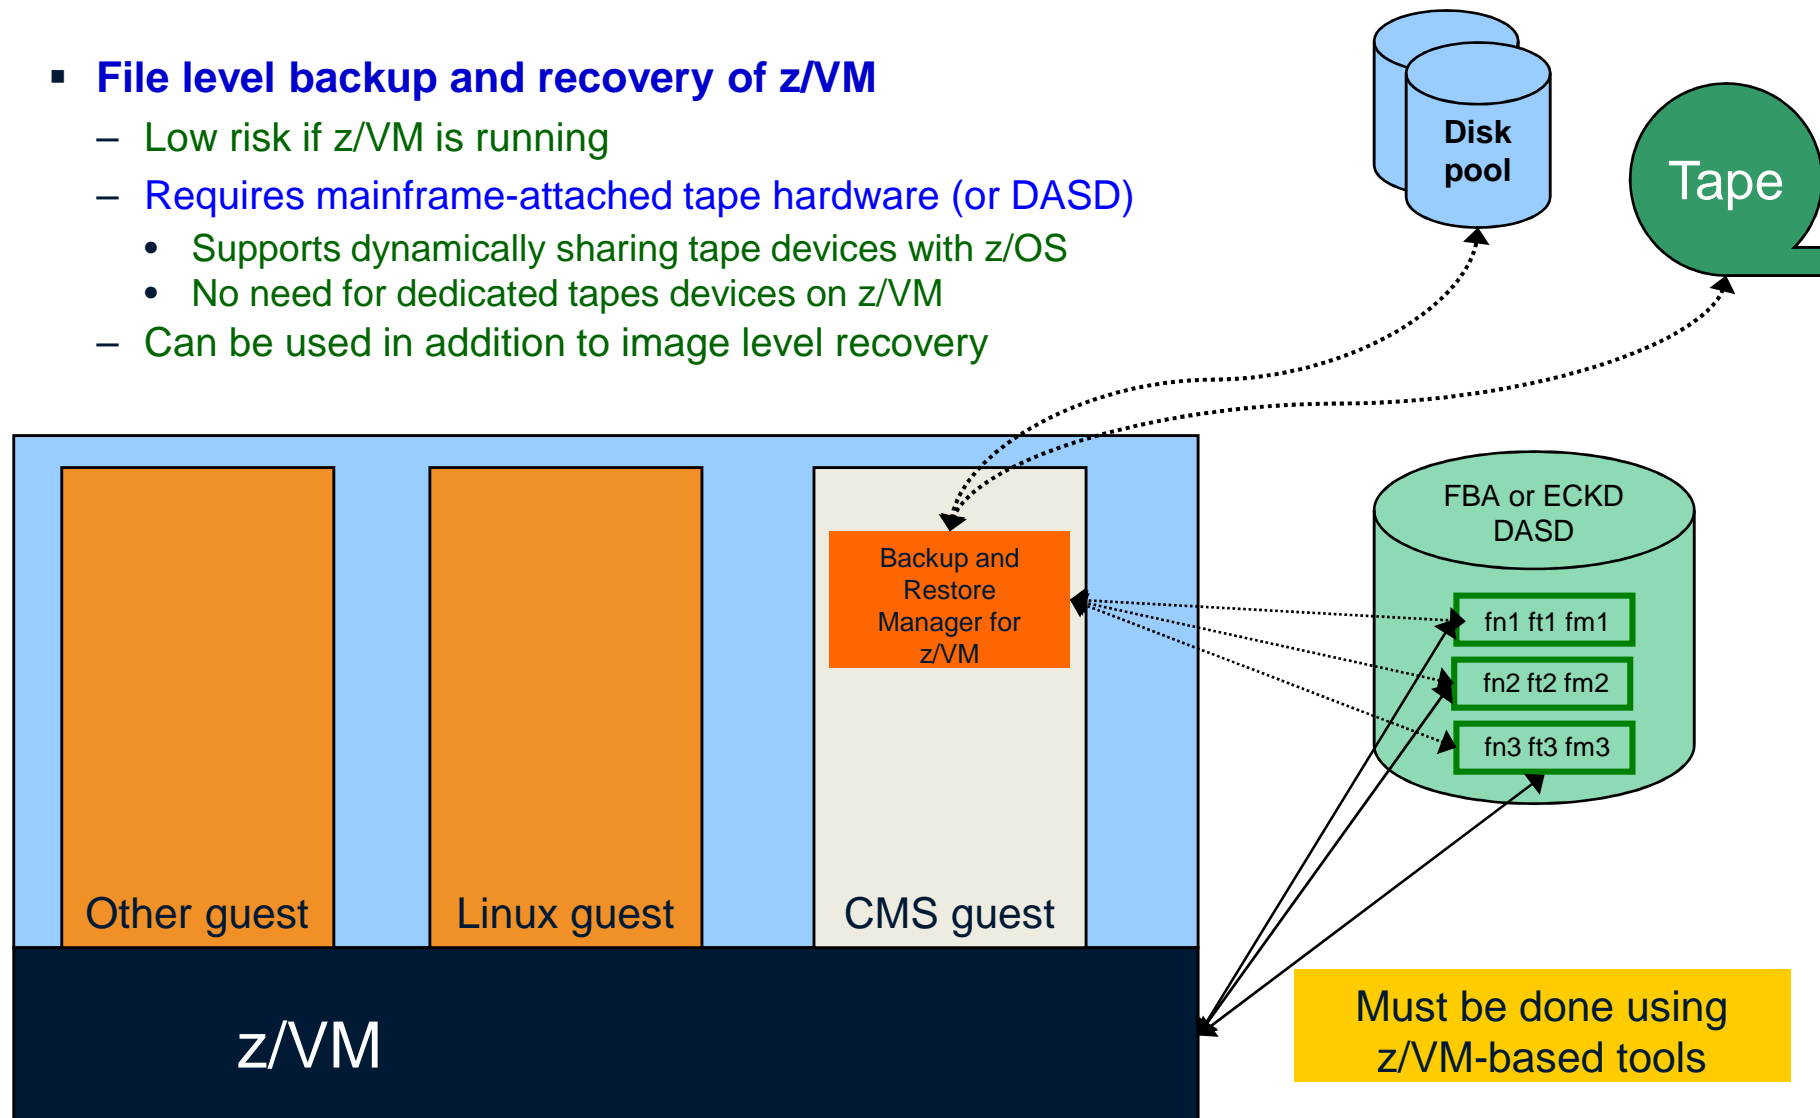
- Configuration files
- Log files
- Tools

Recovery of z/VM system, including Linux guests

- Dependence on z/OS
- versus
- Independent recovery

File Level Backup and Recovery of z/VM

- **File level backup and recovery of z/VM**
 - Low risk if z/VM is running
 - Requires mainframe-attached tape hardware (or DASD)
 - Supports dynamically sharing tape devices with z/OS
 - No need for dedicated tapes devices on z/VM
 - Can be used in addition to image level recovery



Where and How to Back Up z/VM and Linux Guests

- Using z/OS to back up and restore z/VM and Linux
 - Useful during Linux on z Systems POC or early stages of Linux roll-out
 - Easy and fast to implement for existing z/OS customers
 - Provides disaster/volume level recovery (not file level recovery)
 - **Concerns or issues long term as Linux workload grows or becomes critical**
 - Doesn't support FCP-attached DASD
 - File level recovery of z/VM or Linux data is time consuming and manual
 - Backups only contain volume images
 - In disaster situation, z/VM and Linux must wait for z/OS recovery before beginning their recovery
 - Increased use of z/OS CPU cycles to support z/VM and Linux
- Using native z/VM and Linux solutions for backup and recovery
 - **Supports operational errors and disaster situations**
 - File level backup and recovery of both z/VM and Linux
 - Image level backup and recovery of FCP and FICON-attached DASD (z/VM and Linux)
 - **Independent of z/OS**
 - Backups run on (less expensive) IFLs
 - Recovery in parallel with z/OS
 - Dynamically sharing of tape devices with z/OS is still possible
 - Does not require both systems to be up

Backing Up Linux – Should the Guest Be Up or Down?

- Linux keeps pending I/O's in memory when possible
 - Designed for distributed platforms where I/O is assumed to be slow

- Backup solutions that read Linux DASD volumes but run outside Linux don't have a view of these pending I/Os
 - Data on DASD may be in inconsistent state due to pending I/Os
 - Restoring data that has been backed up while Linux is running may not yield usable results
 - SYNC command exists to force all I/Os to be processed
 - Linux will immediately start caching new I/Os
 - Dependent on type of application running on Linux
 - Similar to pulling the plug on a distributed Linux server, then restarting it
 - But worse – backup occurs over a period of time
 - DASD A backed up, then while backing up DASD B, DASD A changes again

Backing Up Linux – Should the Guest Be **Up** or **Down**?

- Reduce risk by
 - “Right-sizing” Linux guests – don’t give more memory than needed
 - Recommended for performance reasons anyway
 - Using Flashcopy to flash the disks and back up the flashed copy
- For guaranteed recovery, shut down or suspend the guest before backing it up from z/VM or z/OS
 - Your experience may (will) vary
 - Evaluate the risk based on the application
 - Use Flashcopy to reduce the downtime

Using **Suspend** Before Backing Up Linux Guests

...

- SUSPEND/RESUME functions available in Linux on z Systems distributions
- Similar to hibernate function in Windows
 - Suspend
 - Completes all pending I/Os
 - Writes memory to disk
 - Resume
 - Detects suspend state
 - Reads memory from disk to restore previous state of the guest
- Requires setup and planning
 - Verify the effort is worth it for each type of guest
 - Otherwise, use shutdown instead of suspend

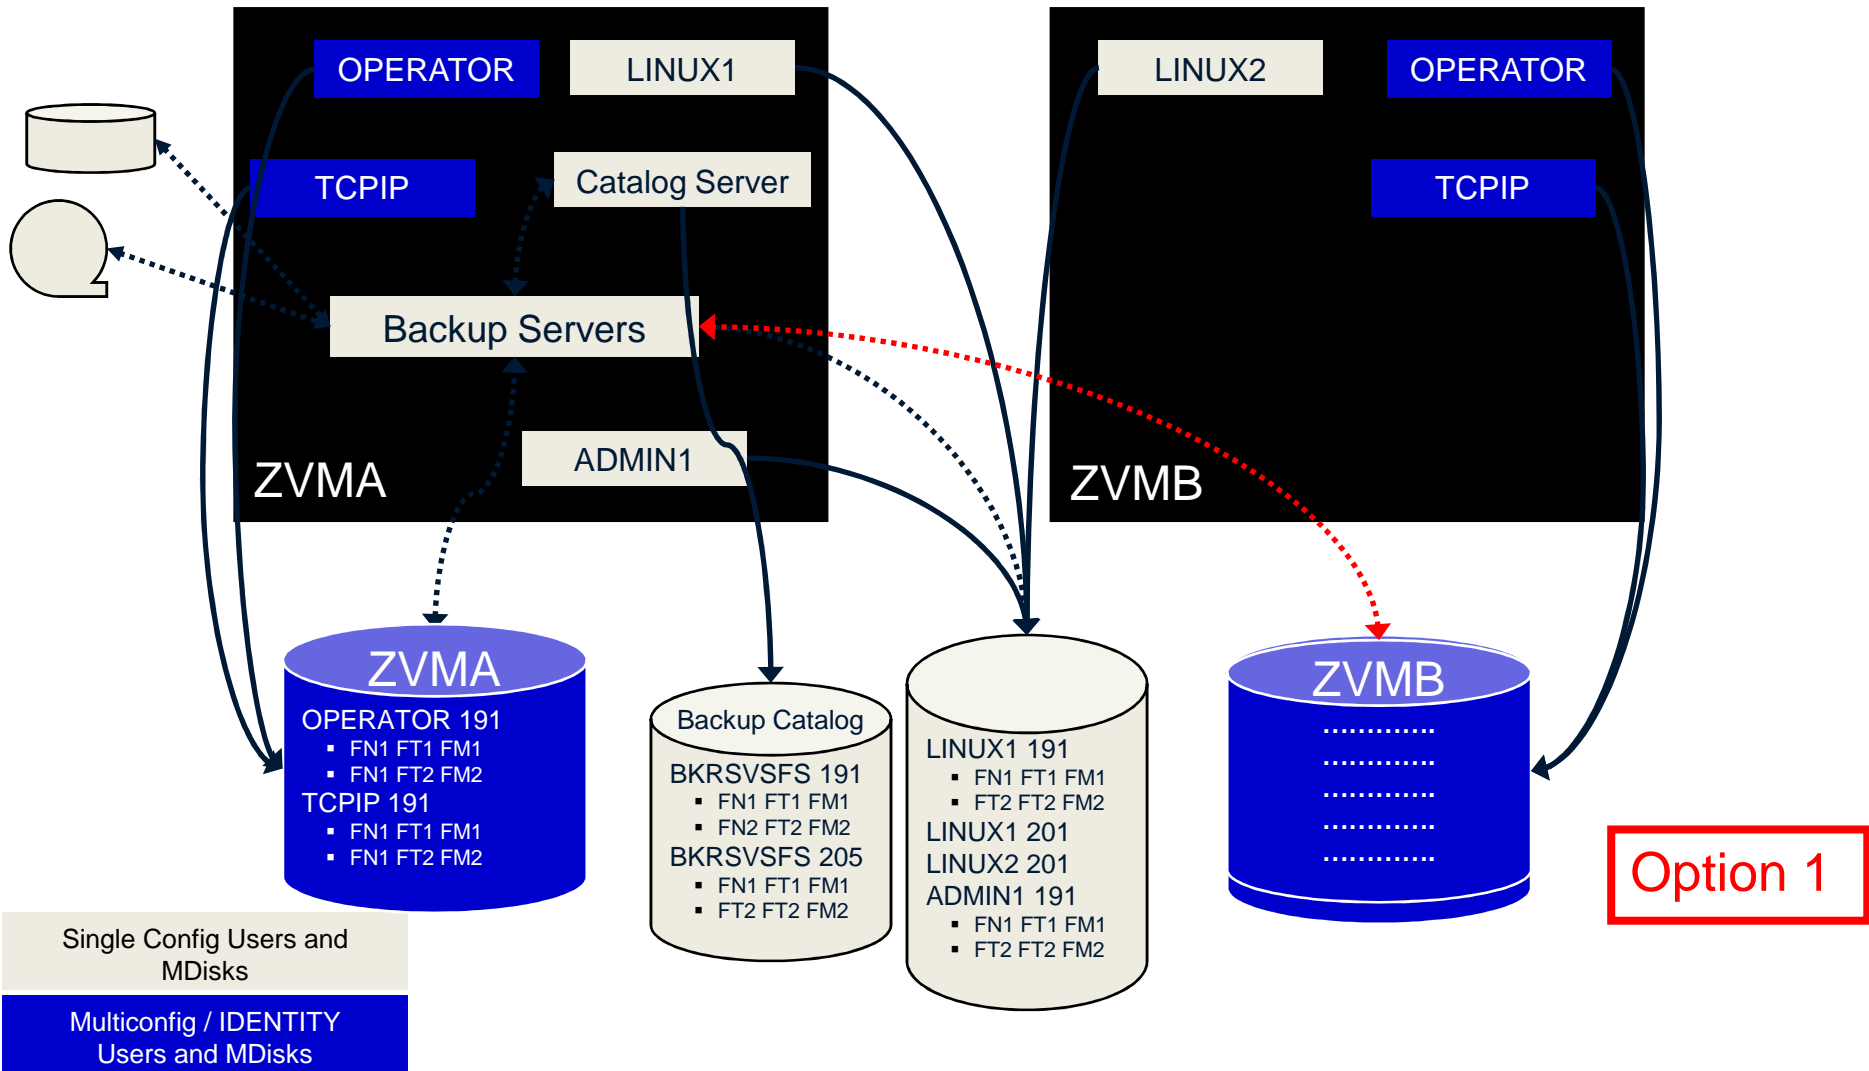
... Using **Suspend** Before Backing Up Linux Guests

- Setup
 - Specify swap disk in zipl.conf
 - Example: resume=/dev/disk/by-path/ccw-0.0.010f-part1
 - In list of swap disks
 - Specify this one with lowest priority
 - Use real disk (not VDISK)
 - Needs to have enough room for all memory of Linux guest + swap space
- Issue suspend via one of the following:
 - echo disk > /sys/power/state
 - CP SIGNAL SHUTDOWN
 - Must update config file on Linux to specify suspend rather than kill in response to signal shutdown
- Reference:
 - White paper – “Methods to pause a z/VM guest: Optimize the resource utilization of idling servers”
 - <http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101981>

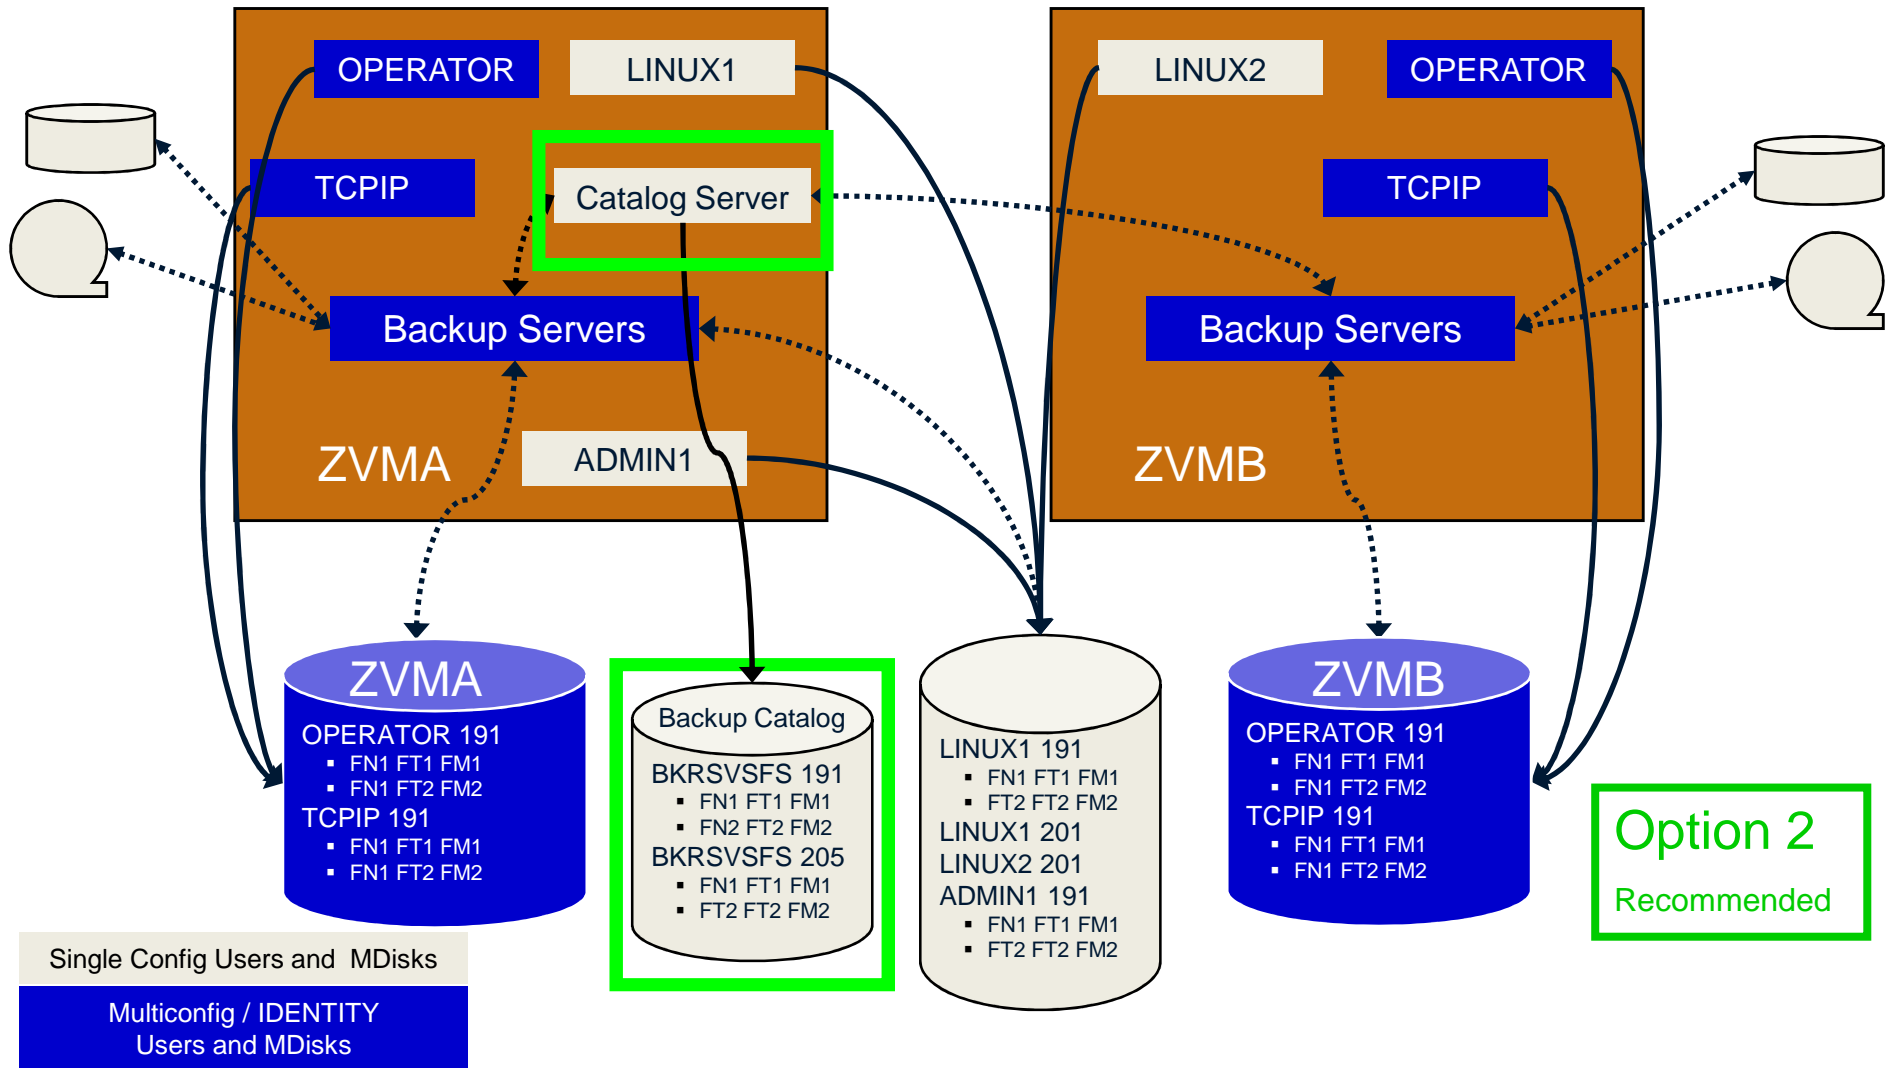


Backing up and Restoring Data in a z/VM SSI Cluster

SSI Considerations for Backup and Restore



SSI Considerations for Backup and Restore



SSI Considerations for Backup and Restore

- Backup service machines on any member can see all minidisks of **single configuration users**
- Backup service machines on any member can see all minidisks of **local multiconfiguration** (IDENTITY) users
 - Can **not** see minidisks of **IDENTITY** users on **other members**
 - Can **only** see DASD volumes (if shared/available) of IDENTITY users on other members
- Recommendation
 - Create Backup service machines as IDENTITY users on each member
 - For IBM Backup and Restore Manager: BKRBKUP, BKRCATLG, BKRWRKnn
 - If backup catalog is in SFS, create one single configuration user for SFS server/filepool
 - Configure as SSI (or REMOTE) in DMSPARMS file
 - Allows single configuration users to restore their own data when logged onto any member
 - Create multiple backup jobs
 - Separate job(s) for single configuration users – only run them from one member
 - For multiconfiguration (IDENTITY) users
 - One job per member
 - Use a unique job name on each member
 - Run the member specific job on that member's backup server



Backup and Recovery
IBM Backup and Restore Manager for z/VM

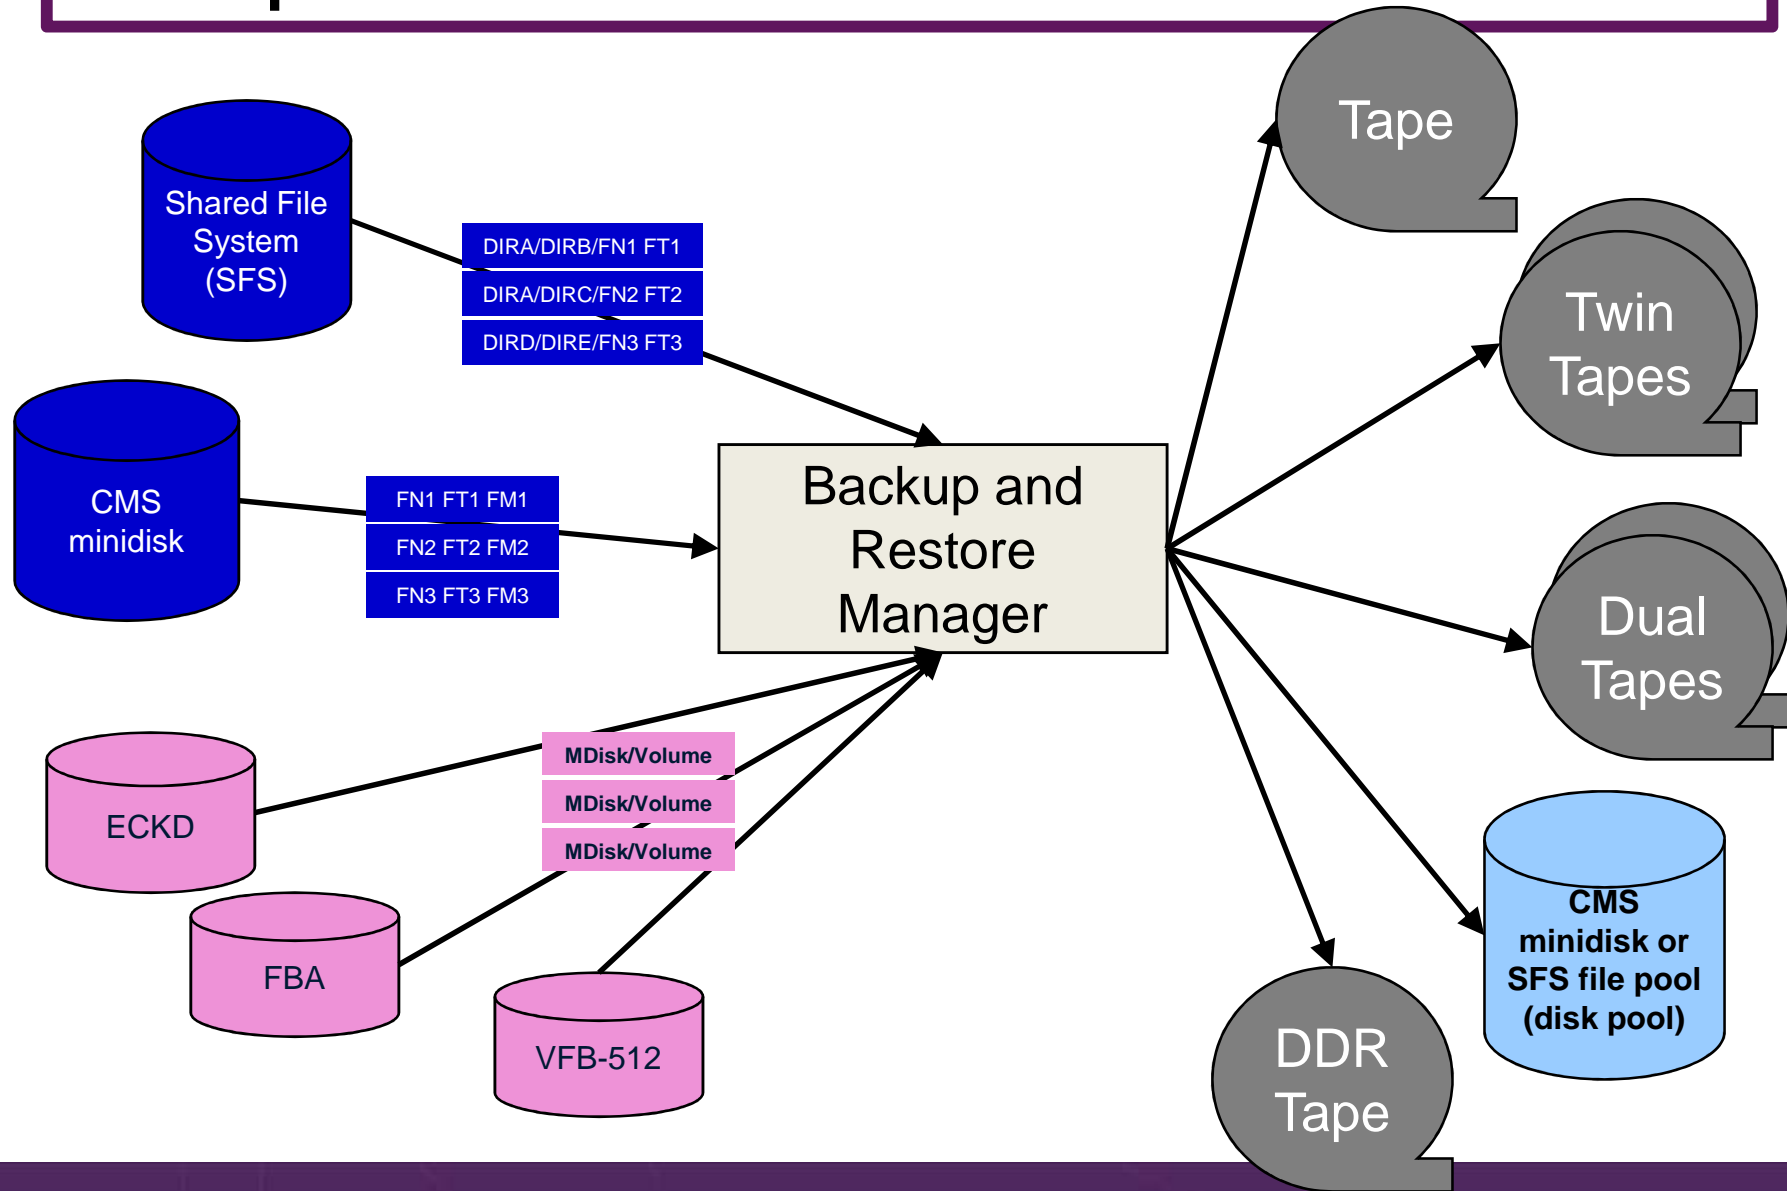
Product Overview

- Backup
 - Requested by administrators
 - Full or incremental
 - Flexible selection of disks and files to back up
 - Review job before submitting for backup
- Restore
 - Restore data via full screen interface or commands
 - Performed by users for their own data
 - Extended to other users available via exit
 - Performed by administrators for any data

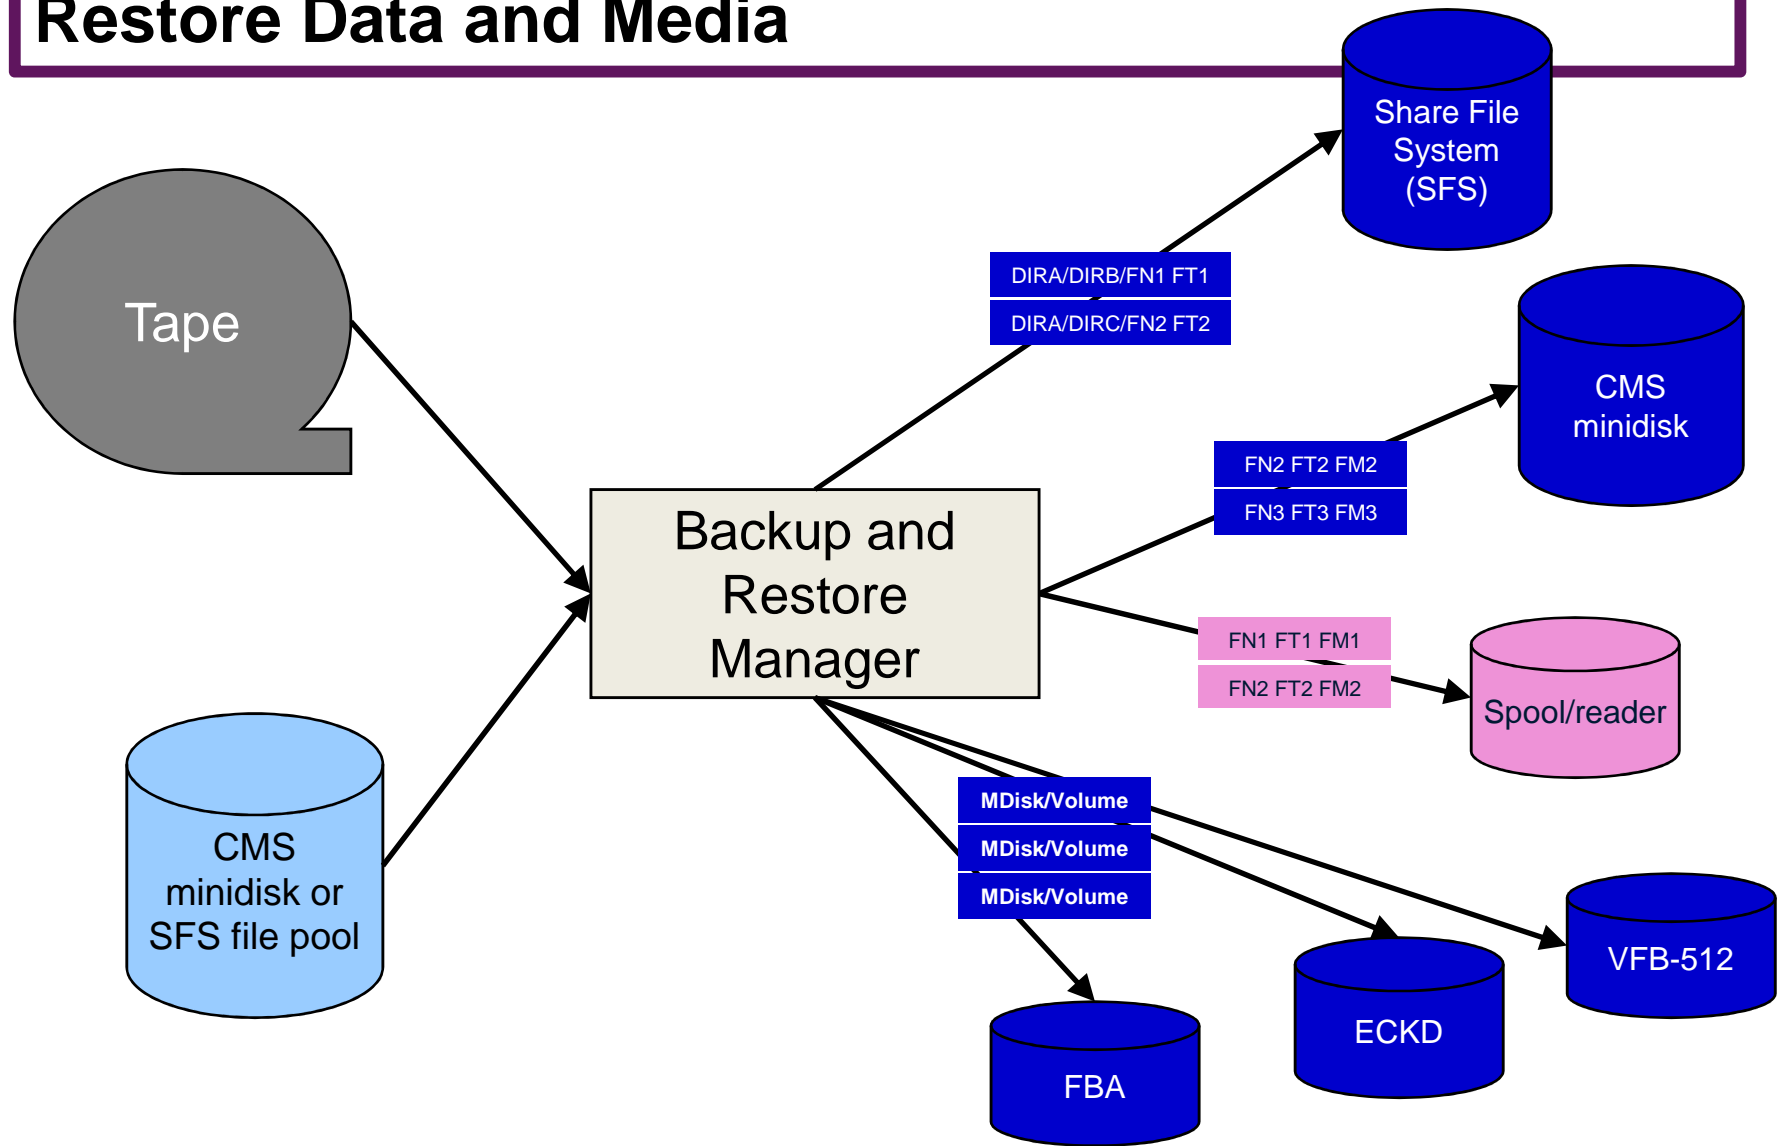
Catalog in Shared File System (SFS) – presentation on web site for installation and setup

- **Integration with Tape Manager for z/VM**
- **Optional compression of data during backup via exits**
 - Call your own compression algorithm
 - Use IBM provided routine
- **Encryption available via exits**
 - Call your own routine
 - Use vendor-written routine, such as V/Soft Software's Encrypt/Backup for z/VM
 - Use encryption capable tape devices

Backup Data and Media




Restore Data and Media



Backup and Restore Manager and Linux Guests

Using Backup and Restore Manager with Tivoli Storage Manager

Choose the solution that meets your needs – or combine for file recovery and DR



Key Benefits

- System backups available for Disaster Recovery
 - Option to restore using DDR or Backup and Restore Manager
 - Manage retention of DR backups
 - Retrieve a list of tapes associated with a specific backup
 - Pull list for movement to off-site storage
- Guest backups available for restoring to a previous state or level
- Backups of user data available for
 - Restoring to a previous state or level
 - Replacing files accidentally erased or corrupted
- Users restore their own data
 - No administrator interaction required

Key Benefits Cont...

- Flexible selection of data to back up
 - Include/exclude
 - Minidisks, SFS directories
 - Real device addresses or volsers
 - Extents
 - Mask by filename, filetype, or SFS path
 - Review a defined backup job before submission
- Management of backup data
 - Retention set as part of the backup job
 - Automatic aging and pruning of the backup catalog
 - Including associated tapes and disk pools (if backed up to disk)
 - View/query the list of expired backups
- Reduced backup window with concurrent processing
 - Multiple worker service machines sharing the job
 - Suggest one worker service machine for each available tape drive
 - Or minidisk in disk pool

Defining a Backup Job

```

/* Include/Exclude definitions */
/*****/
  FUNCTION  MEDIATYPE  OWNER      VDEV VOLUME DEVTYPE      START      END      SIZE
  -----|-----|-----|---|-----|-----|-----|---|-----|---|-----|
INCLUDE    MINIDISK    *          = *    *    *          = *          = *          = *
EXCLUDE    MINIDISK    *LNX*      = *    *    *          = *          = *          = *
EXCLUDE    MINIDISK    MAINT      = 0123 *    *    *          = *          = *          = *
EXCLUDE    MINIDISK    MAINT      = 0124 *    *    *          = *          = *          = *
EXCLUDE    MINIDISK    *          = *    *    *    *          = *          = END        = *
EXCLUDE    MINIDISK    *          = *    *    *    *          = *          = *          > 3300
INCLUDE    MINIDISK    MAINT      = 012* *    *    *          = *          = *          = *

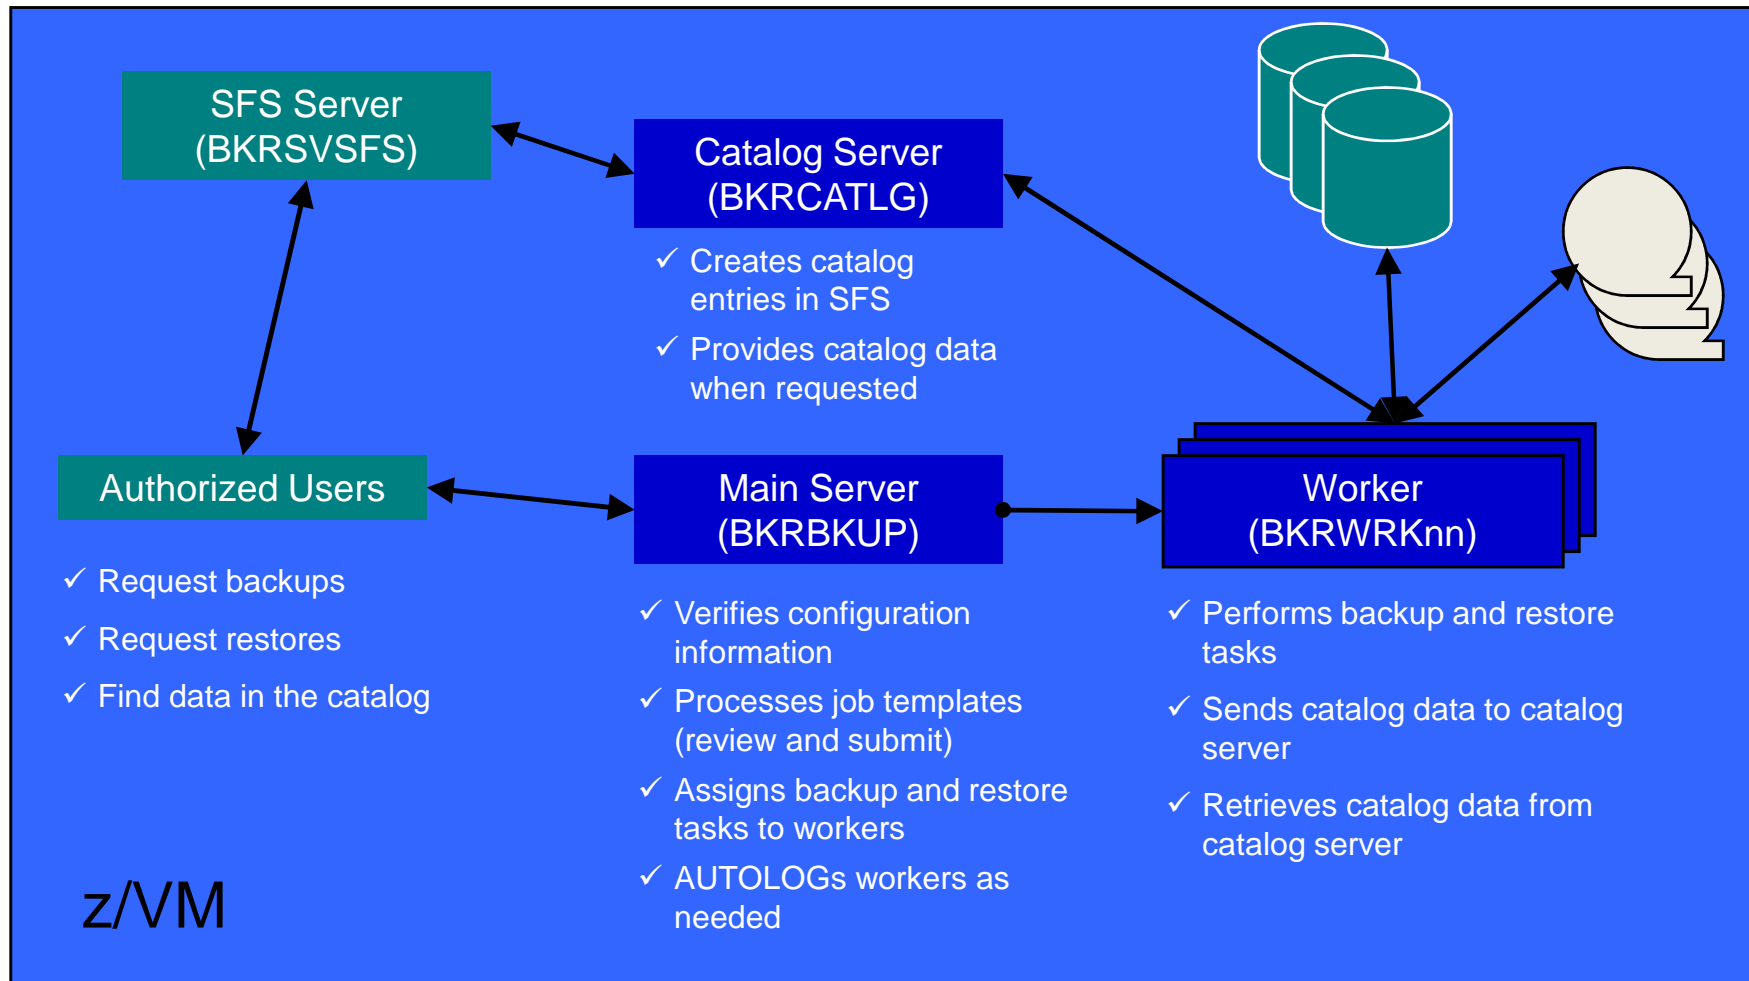
  FUNCTION  MEDIATYPE  ADDRESS
  -----|-----|-----|
INCLUDE    RDEVICE    900-90F
EXCLUDE    RDEVICE    *B

  FUNCTION  MEDIATYPE  VOLSER
  -----|-----|-----|
INCLUDE    RDEVVOL    630*

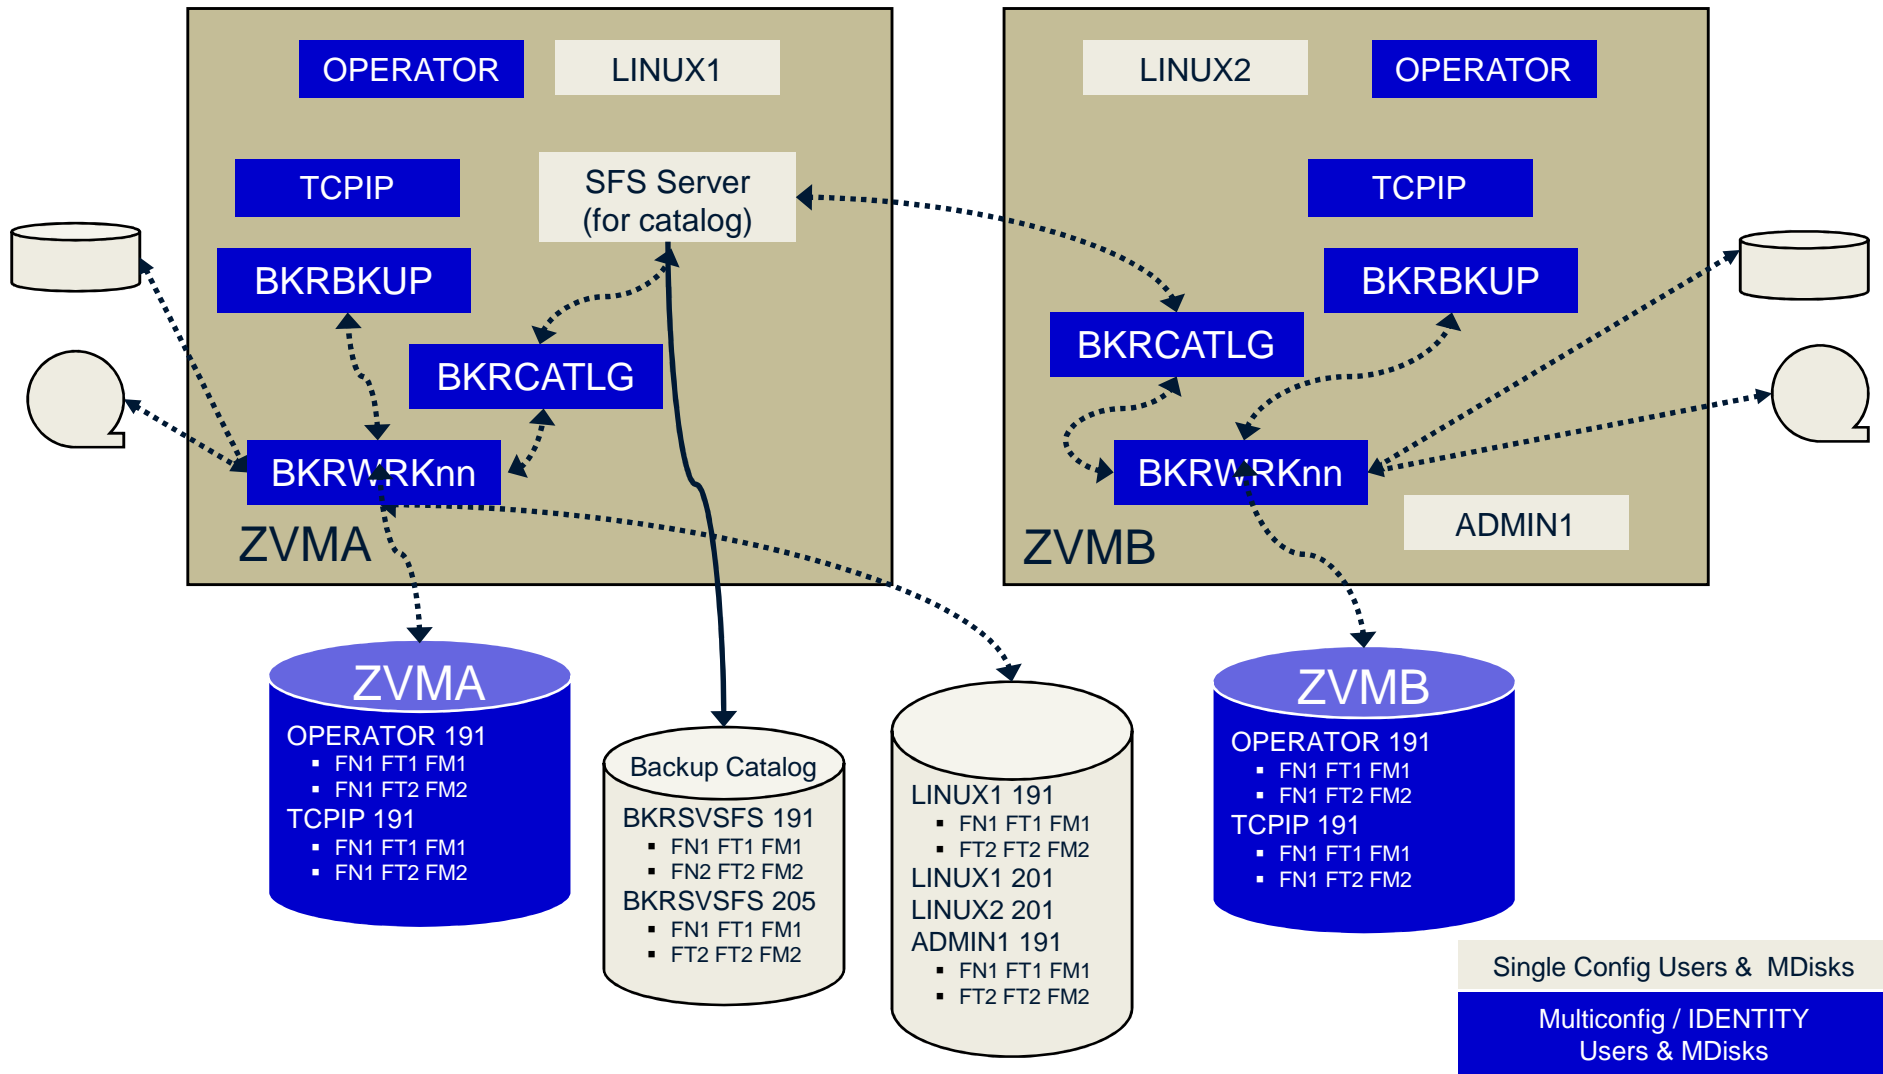
  FUNCTION  MEDIATYPE  POOLNAME  OWNER    FS
  -----|-----|-----|-----|---|
INCLUDE    SFS        VMSYSU:  *        SFS
EXCLUDE    SFS        VMSYSU:  VMSERVU SFS

```

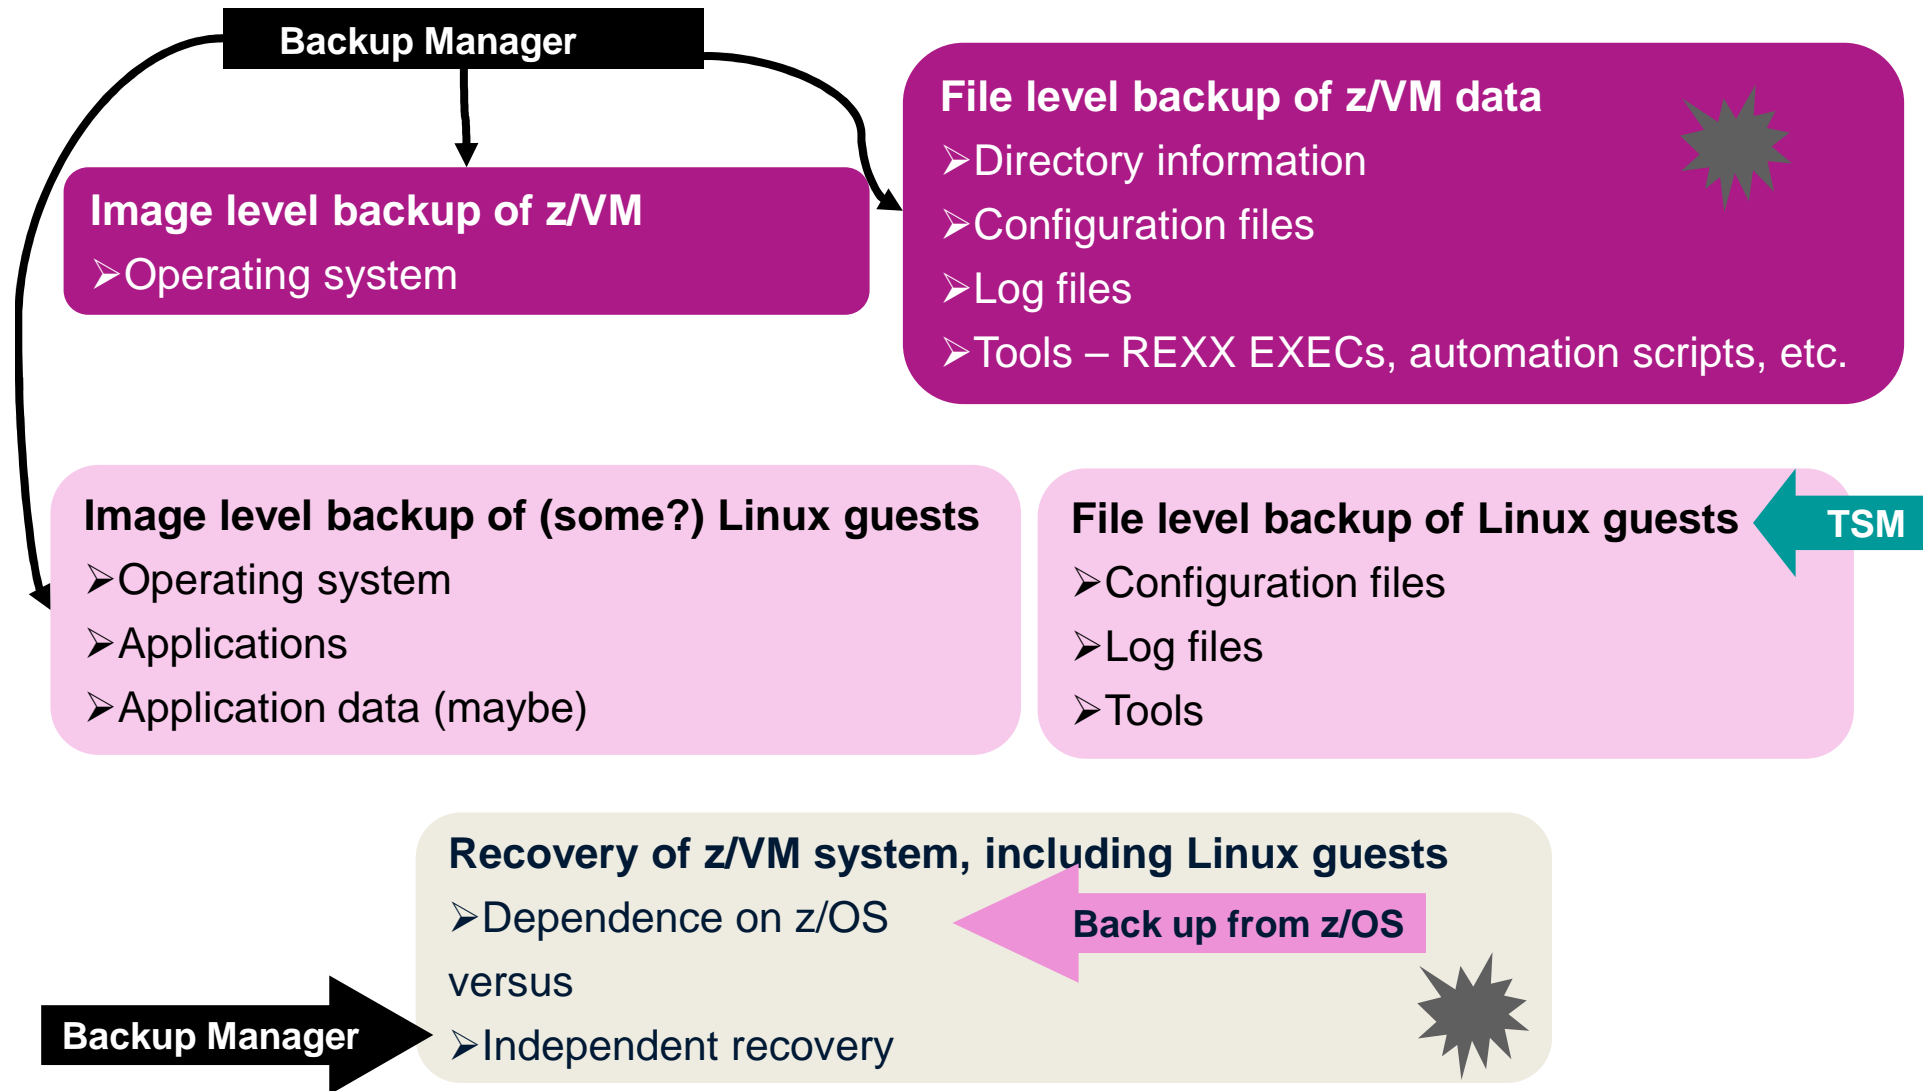
Backup and Restore Manager Architecture – non-SSI



Backup and Restore Manager Architecture – SSI



Recommended Practices – Backup and Recovery



Summary

- Use Backup and Restore Manager to
 - Perform file-level backups of z/VM data
 - Perform image level backups of non-z/VM guest data
 - Use Tivoli Storage Manager for file level backups of Linux
 - Perform disaster recovery backups of entire system
 - Easily find and restore data as needed
 - Automatically manage retention of backup data
 - Carefully plan for SSI configurations



Managing Tapes and Tape Devices
Tape Manager for z/VM

Product Overview

- Manage tapes
 - Define tapes in a catalog
 - Free or used
 - Retention/expiration information
 - ATL/VTS or manual mount
 - Data Security Erase
 - Group tapes together into pools
 - Ownership and access control
 - Media type
- Manage devices
 - Define available devices
 - Dedicated or assignable
 - Group devices together into device pools
 - ATL/VTS or manual mount
 - Any other grouping you choose (read only vs. write, location, etc.)
 - Share devices with other systems

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

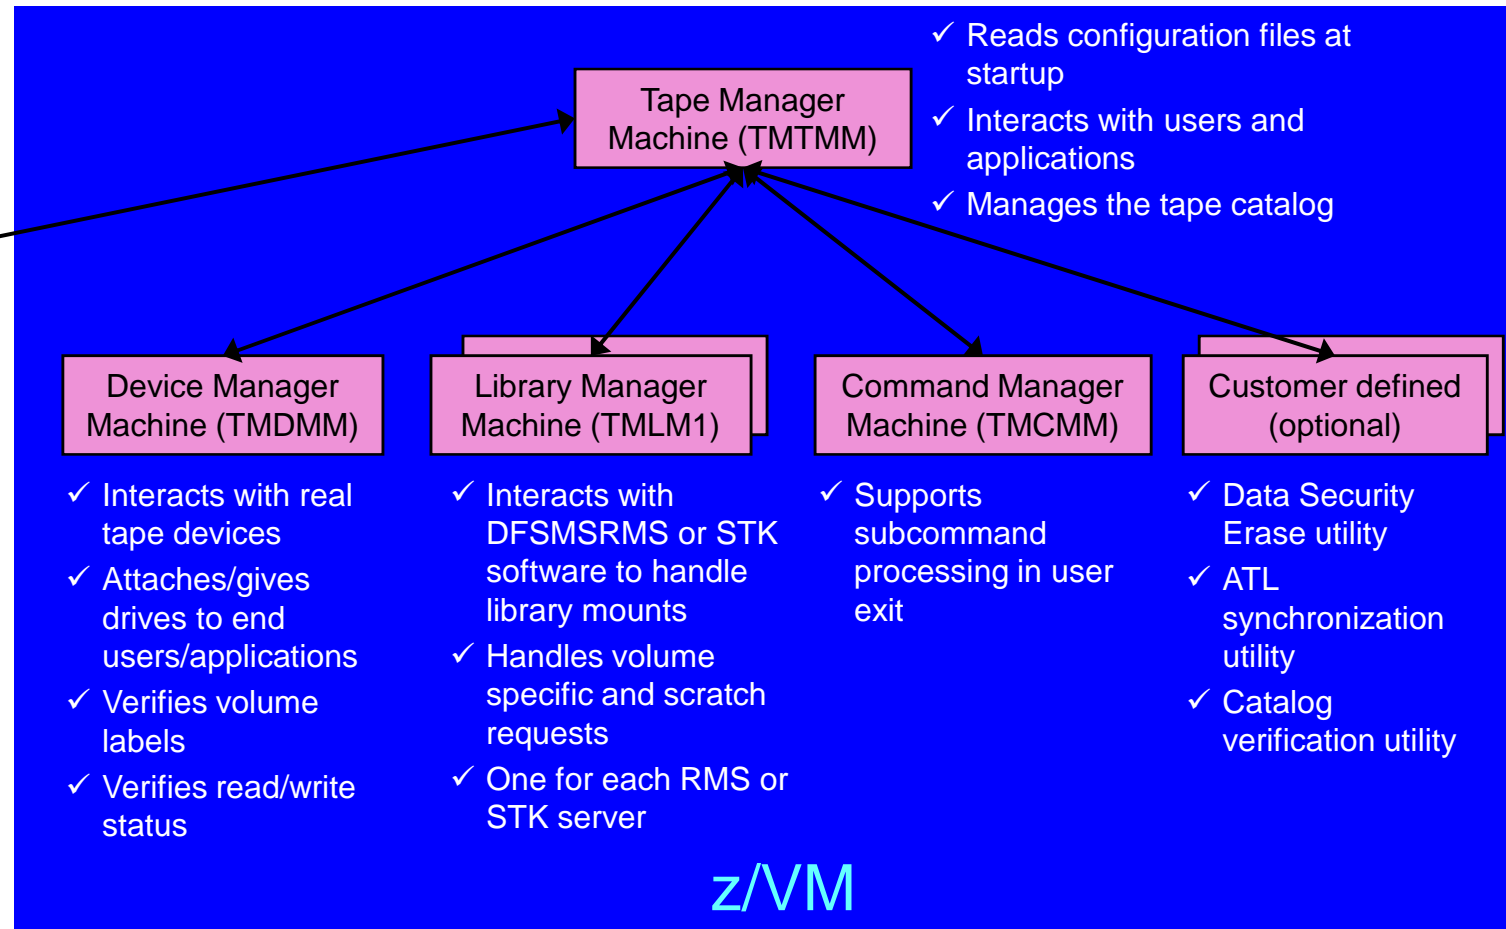
Key Benefits

- Effective management of tapes in ATL or VTS
 - Granular access control
 - Expiration processing
 - Notification for low threshold for tape resources
 - IBM libraries supported through DFSMSRMS on z/VM
 - STK libraries supported through STK Host Software Component for VM, or STK VM Client
 - EMC libraries supported through standard CCW interface
- Improved accuracy of manual tape processing
 - Granular access control
 - Automated interface to Operator for manual mounts
 - Internal label verification at attach/give and detach (SL only)
 - Read/Write verification at attach/give
- Integrated management of z/OS and z/VM tapes using DFSMSrmm on z/OS
 - Optionally use RMM on z/OS as the tape catalog for z/VM and z/OS tapes
 - Tapes, access control, and retention managed by the existing RMM catalog
 - Accessible via Tape Manager on z/VM
 - Tapes managed by RMM
 - Devices managed by Tape Manager – sharing devices with z/OS is discussed later
 - Not available for STK libraries


Data Security Erase (DSE)

- Erase (sensitive) data before tape is reused
- Option to enable DSE at tape pool or individual tape level
 - DSE-enabled flag included in each catalog entry
- DSE-enabled tapes marked as DSE-ready when freed
- Tape Manager DSE utility (TMDSE) executed on a separate user ID
 - Started manually or automatically with Operations Manager
 - Queries the catalog to find all tapes with DSE-ready flag on
 - Mounts each tape
 - Verifies volume label if possible
 - Configuration option to perform DSE on NL tapes or not
 - Erases tape
 - Turns off DSE-ready flag in catalog
 - Tape is now available for scratch unless its HOLD flag is on

Tape Manager in Standard Mode



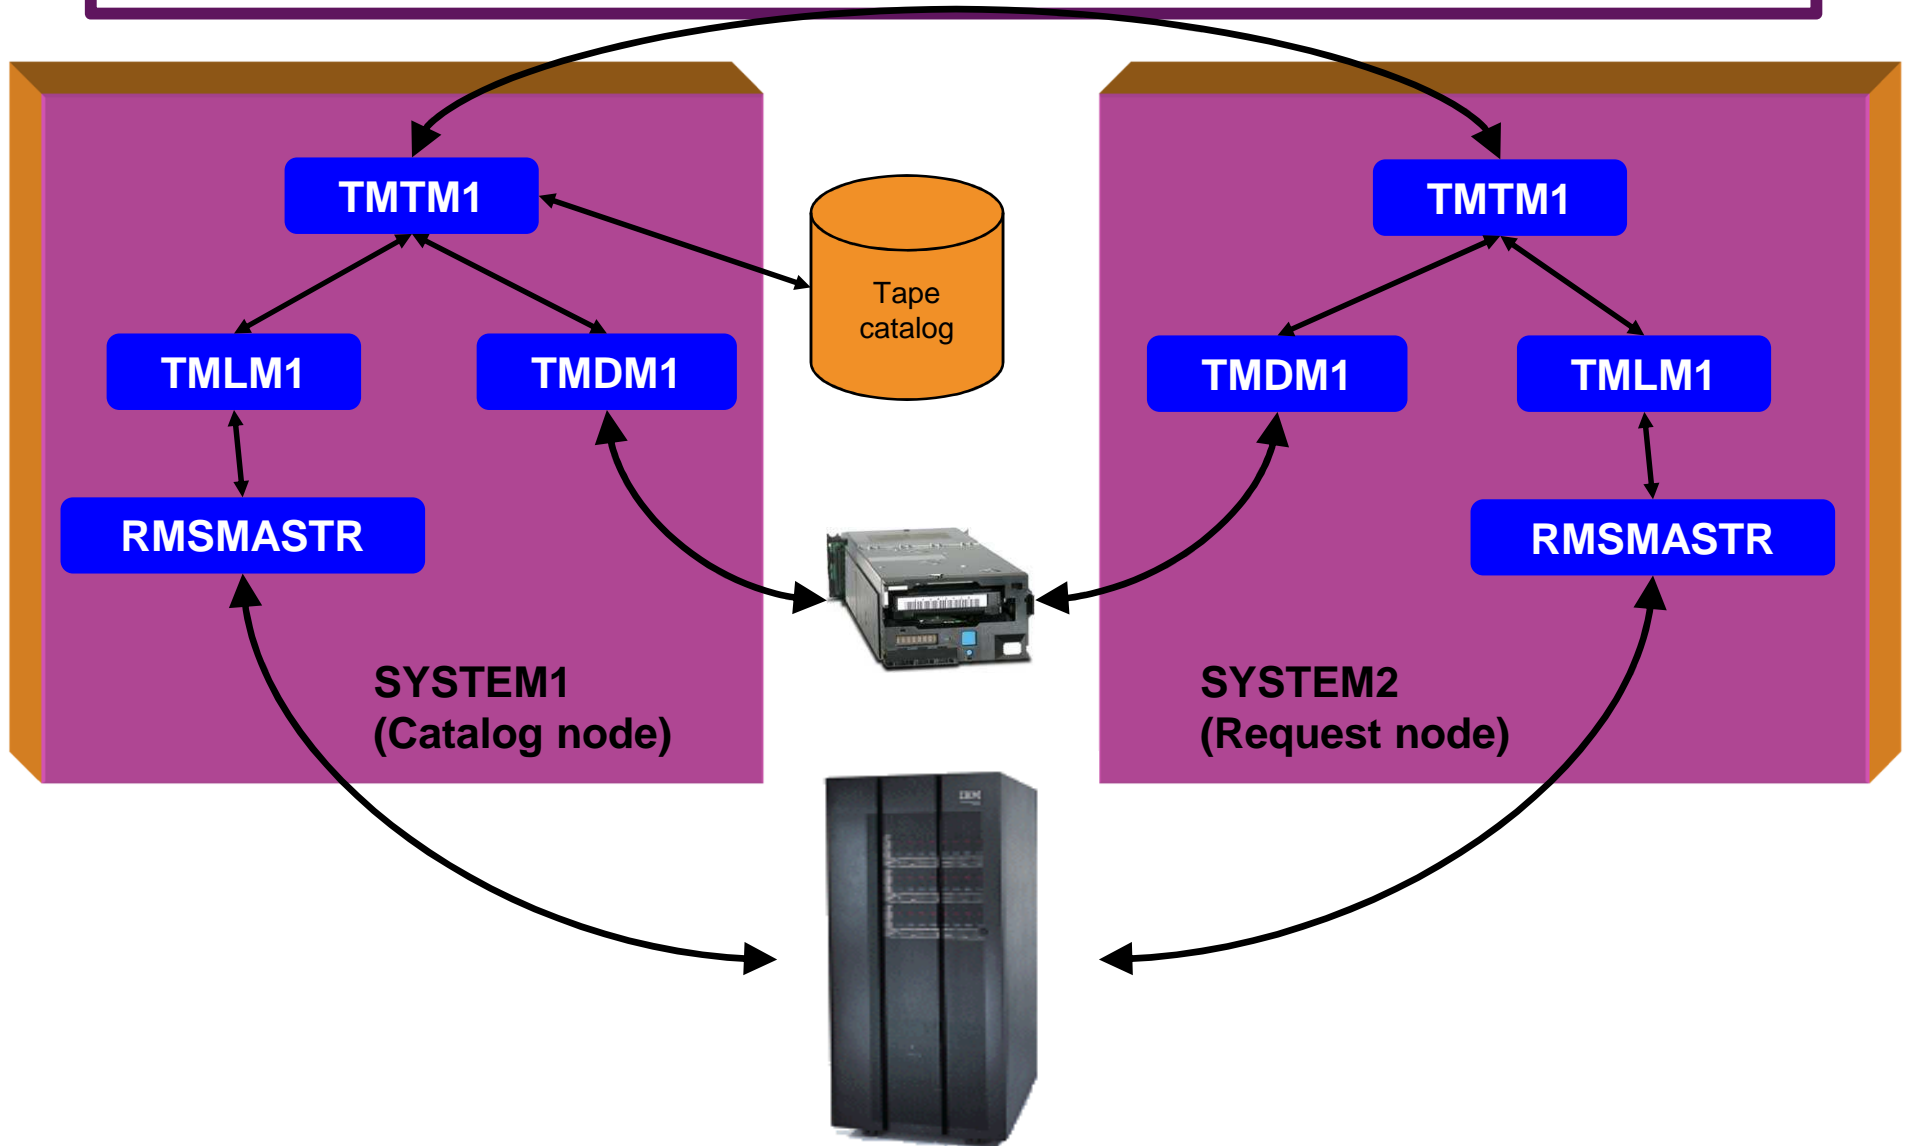
Tape Manager in RMM Mode




Support for **One Tape Catalog** Across **Multiple z/VM** Systems

- One “catalog node”
 - Responsible for the tape catalog contents
- Multiple “request nodes”
 - Manage requests on the local system
 - Communicate with catalog node to read or update catalog data
- One catalog used by multiple z/VM systems
 - No longer need to create a catalog on each z/VM system, each with its own range of volsers
 - All z/VM systems share one catalog
- IP used for communication between systems

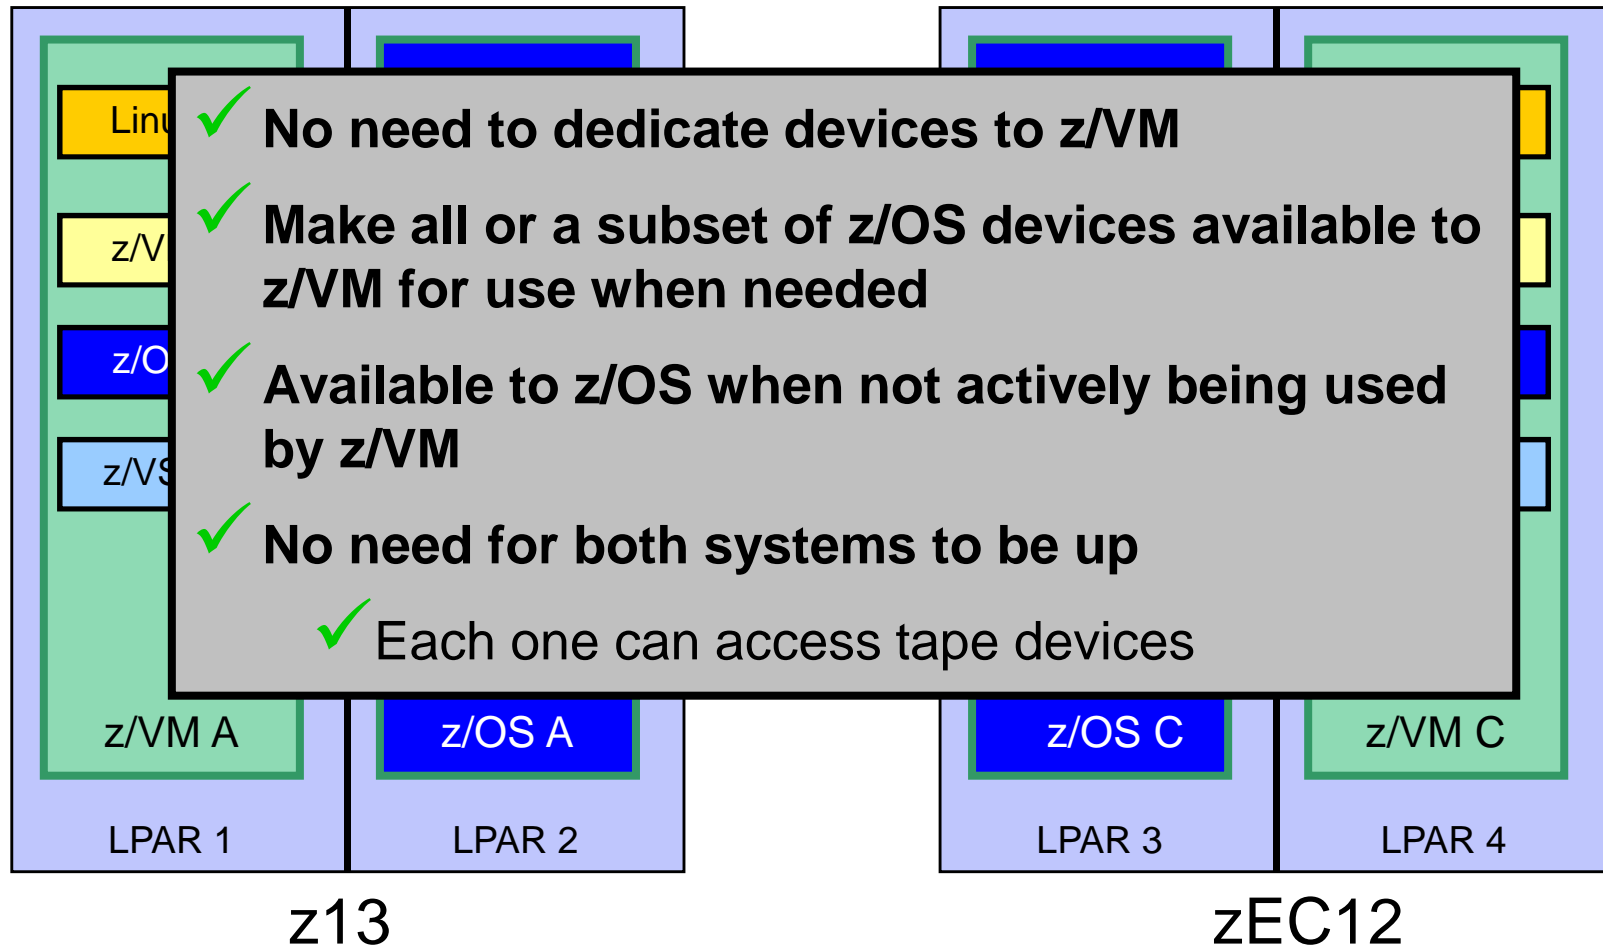
Communication Between Service Machines and Systems



Dynamically Share Real Tape Devices



Dynamically Share Tape Devices



Tape Manager for z/VM - Summary

- Use Tape Manager to
 - Manage and share devices
 - Manage tape volumes
 - Access control
 - Retention
 - Data Security
 - Improve accuracy of mount requests

Summary

- Management of z/VM systems with Linux guests requires monitoring and management tools
- IBM solutions exist
 - OMEGAMON XE on z/VM and Linux
 - zSecure Manager for z/VM
 - Operations Manager for z/VM
 - Wave for z/VM
 - Tape Manager for z/VM
 - Backup and Restore Manager for z/VM
 - Archive Manager for z/VM
- Demos are available

Reference Information

- Product Web site
 - Start at <http://www.ibm.com/software/products/en/backup-restore-manager-for-zvm>
 - Start at <http://www.ibm.com/software/products/en/tape-manager-for-zvm>
 - Product pages include
 - Publications
 - Pre-requisites
 - Presentations
 - White papers
 - Support
- e-mail
 - Mike Sine, sine@us.ibm.com, Technical Marketing
 - Tracy Dean, tld1@us.ibm.com, Product Manager
- White papers and presentations on Backup and Restore Manager and Tape Manager websites (Resources tab)
 - Getting Started with Installation, including SFS server creation and installation of Backup Mgr
 - z/VM V6.2 and later
 - z/VM V5.4
 - Backing up z/VM and Linux on System z – Tivoli Storage Manager vs Backup Manager
 - Pausing (including SUSPENDing) a Linux Guest
 - Enabling the FACILITY Class for Use by RACF for z/VM

The background of the slide is a dark blue gradient. On the left side, there is a complex geometric pattern of overlapping triangles and polygons. The colors in this pattern range from deep purple and indigo to bright orange and yellow. The shapes are semi-transparent, creating a layered effect. The overall aesthetic is modern and technical.

Demonstration Scenarios

Backup and Recovery – Demos Available

- A. Performing an incremental backup
- B. Restoring files from backup
- C. Back up and restore single and multiconfiguration users in an SSI environment
- D. Scheduling image backups of Linux guests
- E. Suspend and resume a guest as part of backup
- F. Reviewing a disaster recovery backup
- G. Reviewing data in the backup catalog for recovery

धन्यवाद

Hindi

多謝

Traditional Chinese

감사합니다

Korean

Спасибо

Russian

Gracias

Spanish

شكراً

Arabic

Thank
You

English

Obrigado

Brazilian Portuguese

Grazie

Italian

多谢

Simplified Chinese

Danke
German

Merci

French

நன்றி

Tamil

ありがとうございました

Japanese

ขอบคุณ

Thai