


# **z/VSE Tape-less operation in a HA and D/R mindset**



Wilhelm Mild  
IT Architect  
IBM Germany  
[mildw@de.ibm.com](mailto:mildw@de.ibm.com)



# Trademarks

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

AIX*	IBM*	PR/SM	WebSphere*	z/OS*
BladeCenter*	IBM (logo)*	System Storage*	XIV*	z/VM*
DataPower*	NetWeaver*	System x*	z9*	z/VSE
DB2*	Parallel Sysplex*	System z*	z10 EC	
FICON*	POWER*	System z9*	zEnterprise	
FlashCopy*	POWER7*	System z10*		
GDPS*	Power Systems			

\* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

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.

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

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

InfiniBand is a trademark and service mark of the InfiniBand Trade Association.

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.

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

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

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

\* All other products may be trademarks or registered trademarks of their respective companies.

## Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

## Resiliency – often called business continuity

### Data Center News:

#### The mainframe in business resiliency

nochan, Contributor  
12 Apr 2006 | SearchDataCenter.com

### Business resiliency

- the ability of the enterprise to continue to function:
  - as effectively as possible
  - man-made problems and technical errors
  - in the face of natural disasters

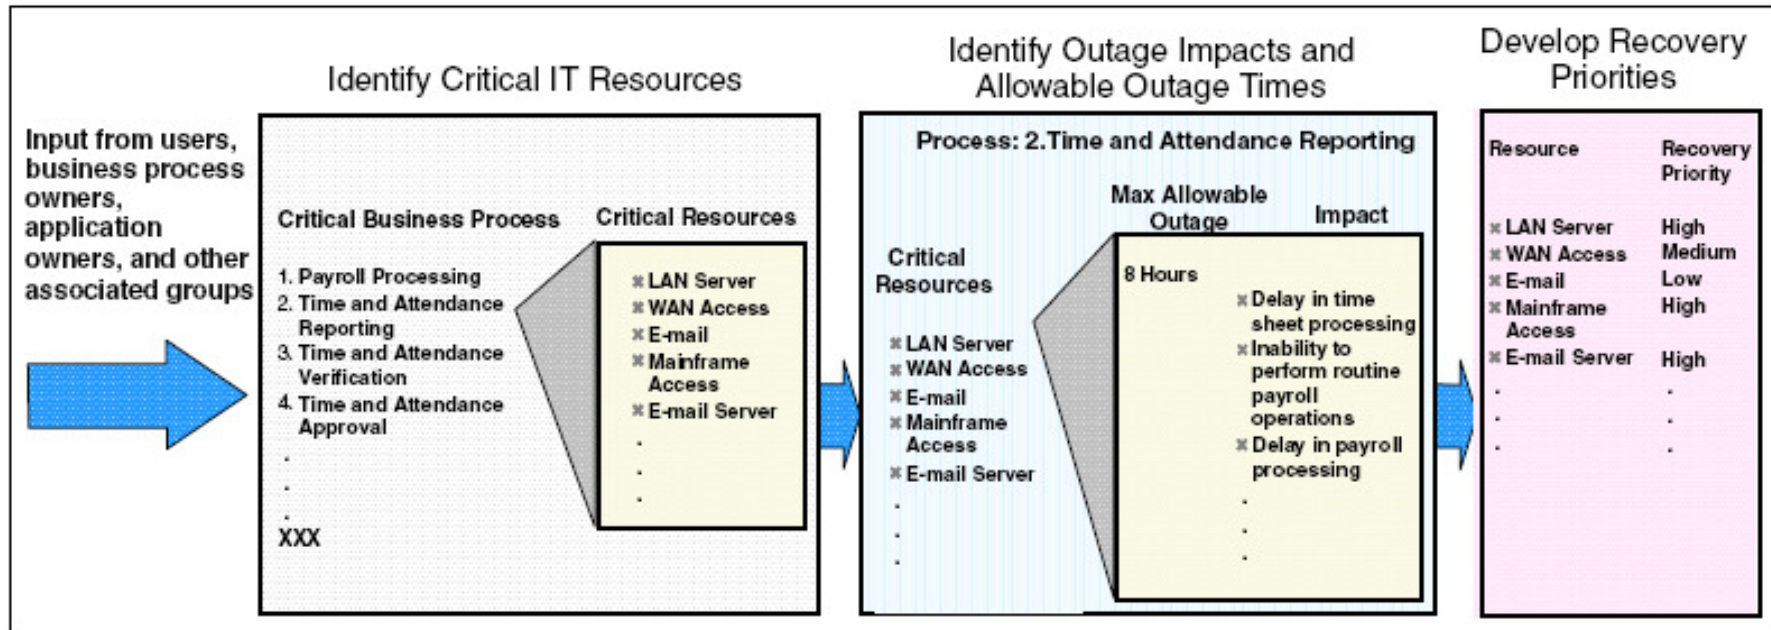
### -High Importance:

- recent natural and man-made disasters,
- plus new requirements for business compliance,
- have increased the importance of business resiliency to the point where even SMBs (Small to Medium sized Businesses) must plan and implement resiliency strategies, with input from the highest levels of the organization.

[http://searchdatacenter.techtarget.com/news/article/0,289142,sid80\\_gci1179879,00.html](http://searchdatacenter.techtarget.com/news/article/0,289142,sid80_gci1179879,00.html)

## The Business impact analysis (BIA)

- IT Resource relation and priorities for Recovery
- Consider all environments
- Prioritize based on business importance




*Example of the Business Impact Analysis process*

# Identify RTO, RPO und NRO




Business Resiliency Plan



## Recovery Time Objective (RTO)

What time difference can be between Failure and a total productional run level ?



## Recovery Point Objective (RPO)

What is the toleration for data loss?

RPO = "0" means, NULL data loss acceptable

RPO = "5" means, data loss in last 5 min acceptable

TREND: RPO = 0

## Network Recovery Objective (NRO)

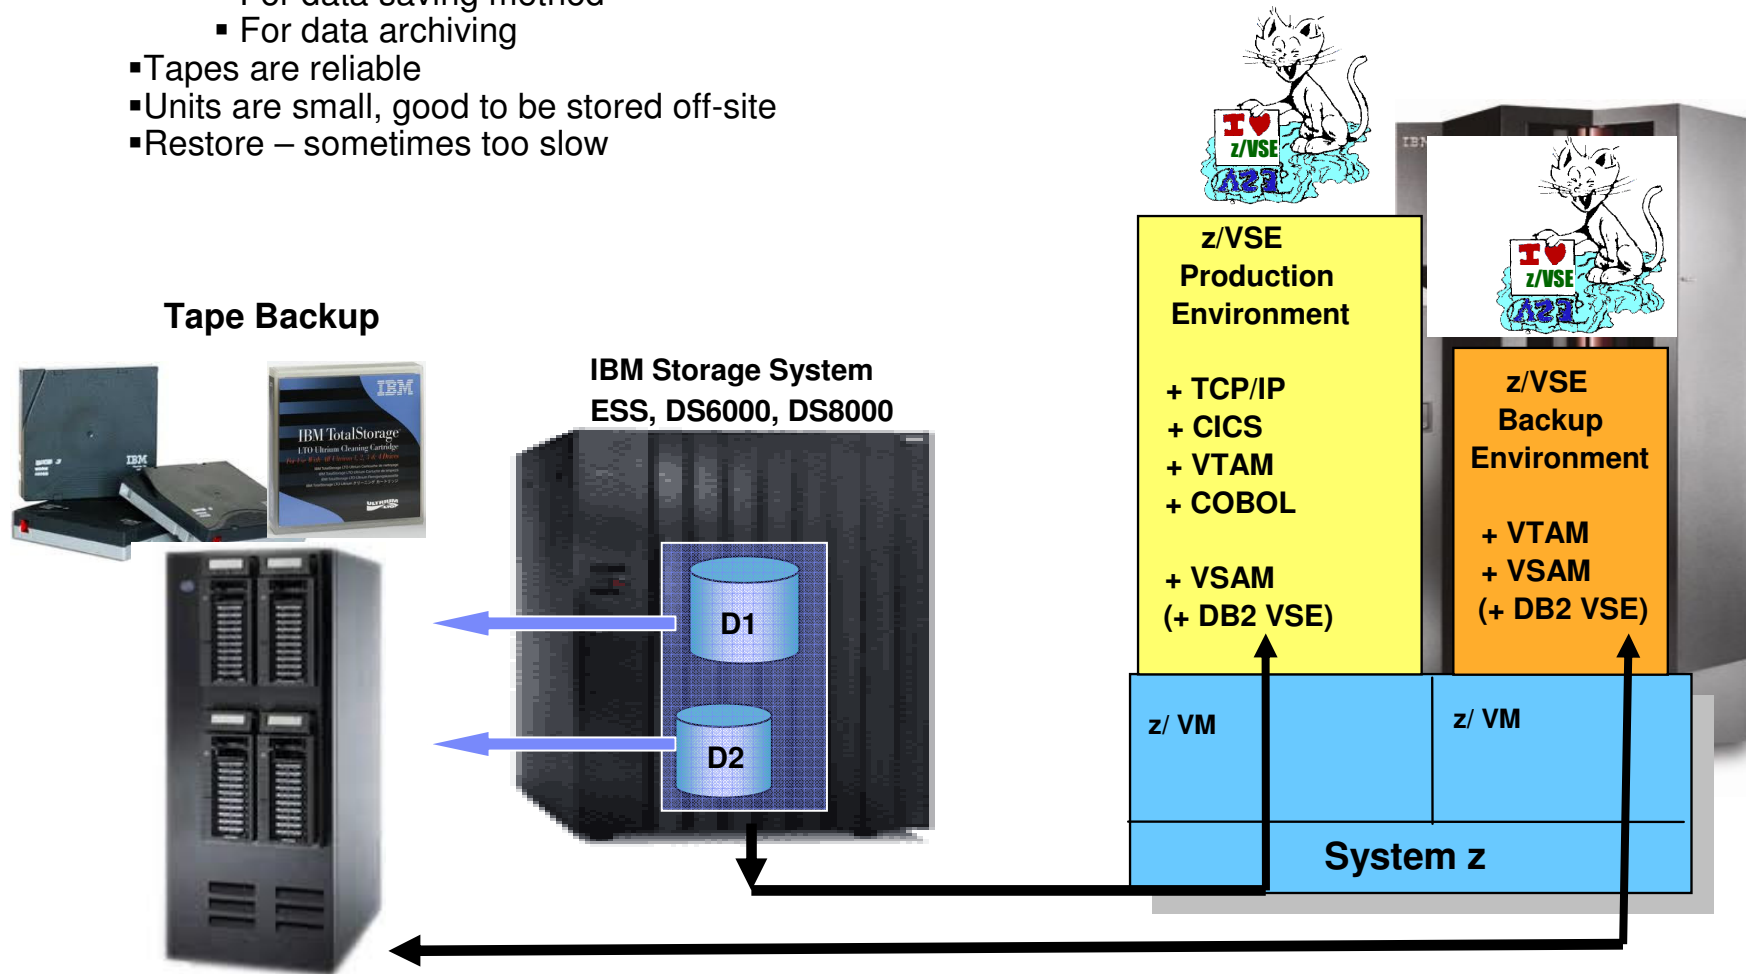
Time requirements for network availability.



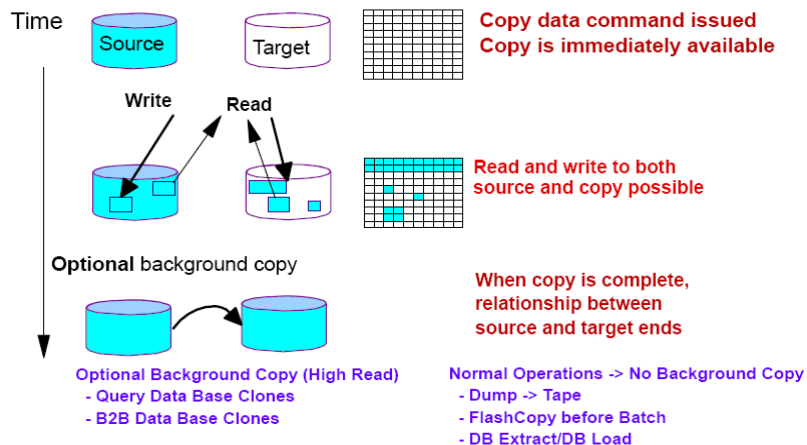
## Backup Procedures with Tapes

### Traditional Backup to tape:

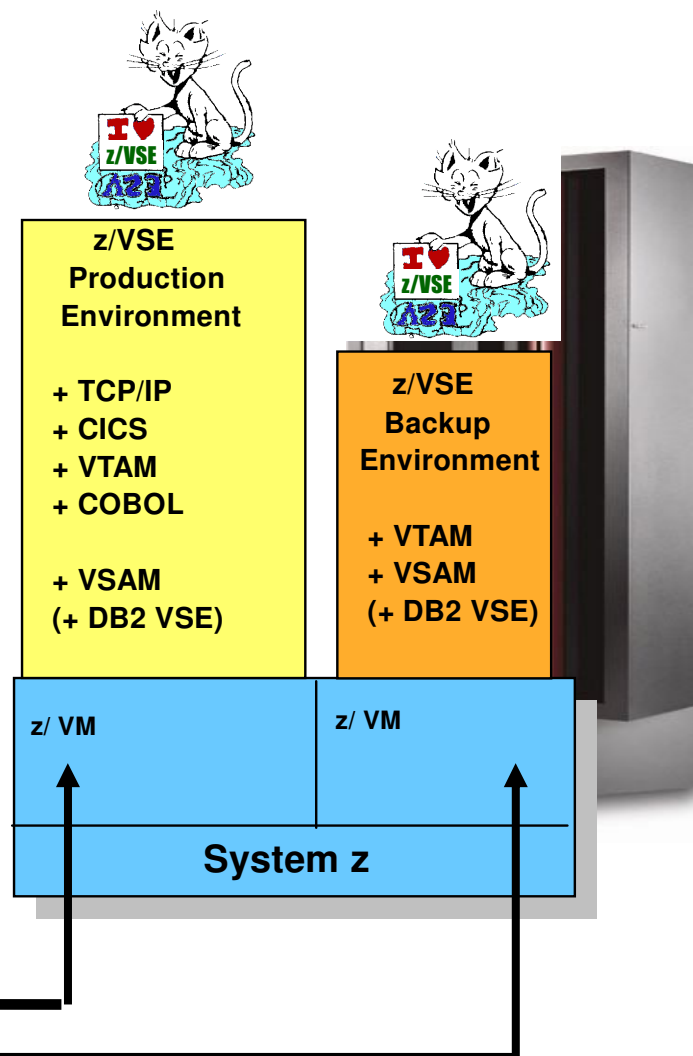
- Before batch run
- After batch run
- Before system updates
- For data saving method
- For data archiving
- Tapes are reliable
- Units are small, good to be stored off-site
- Restore – sometimes too slow



### DS8000 / DS6000 / ESS Flash Copy

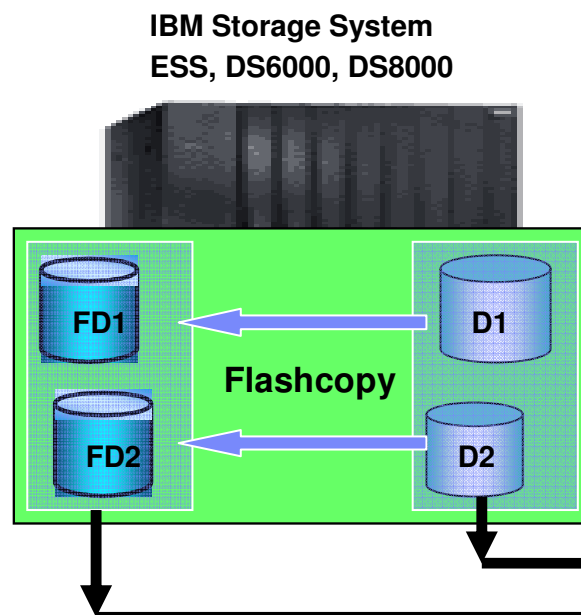


## IT Environment for almost 24x7 Availability



### Flashcopy:

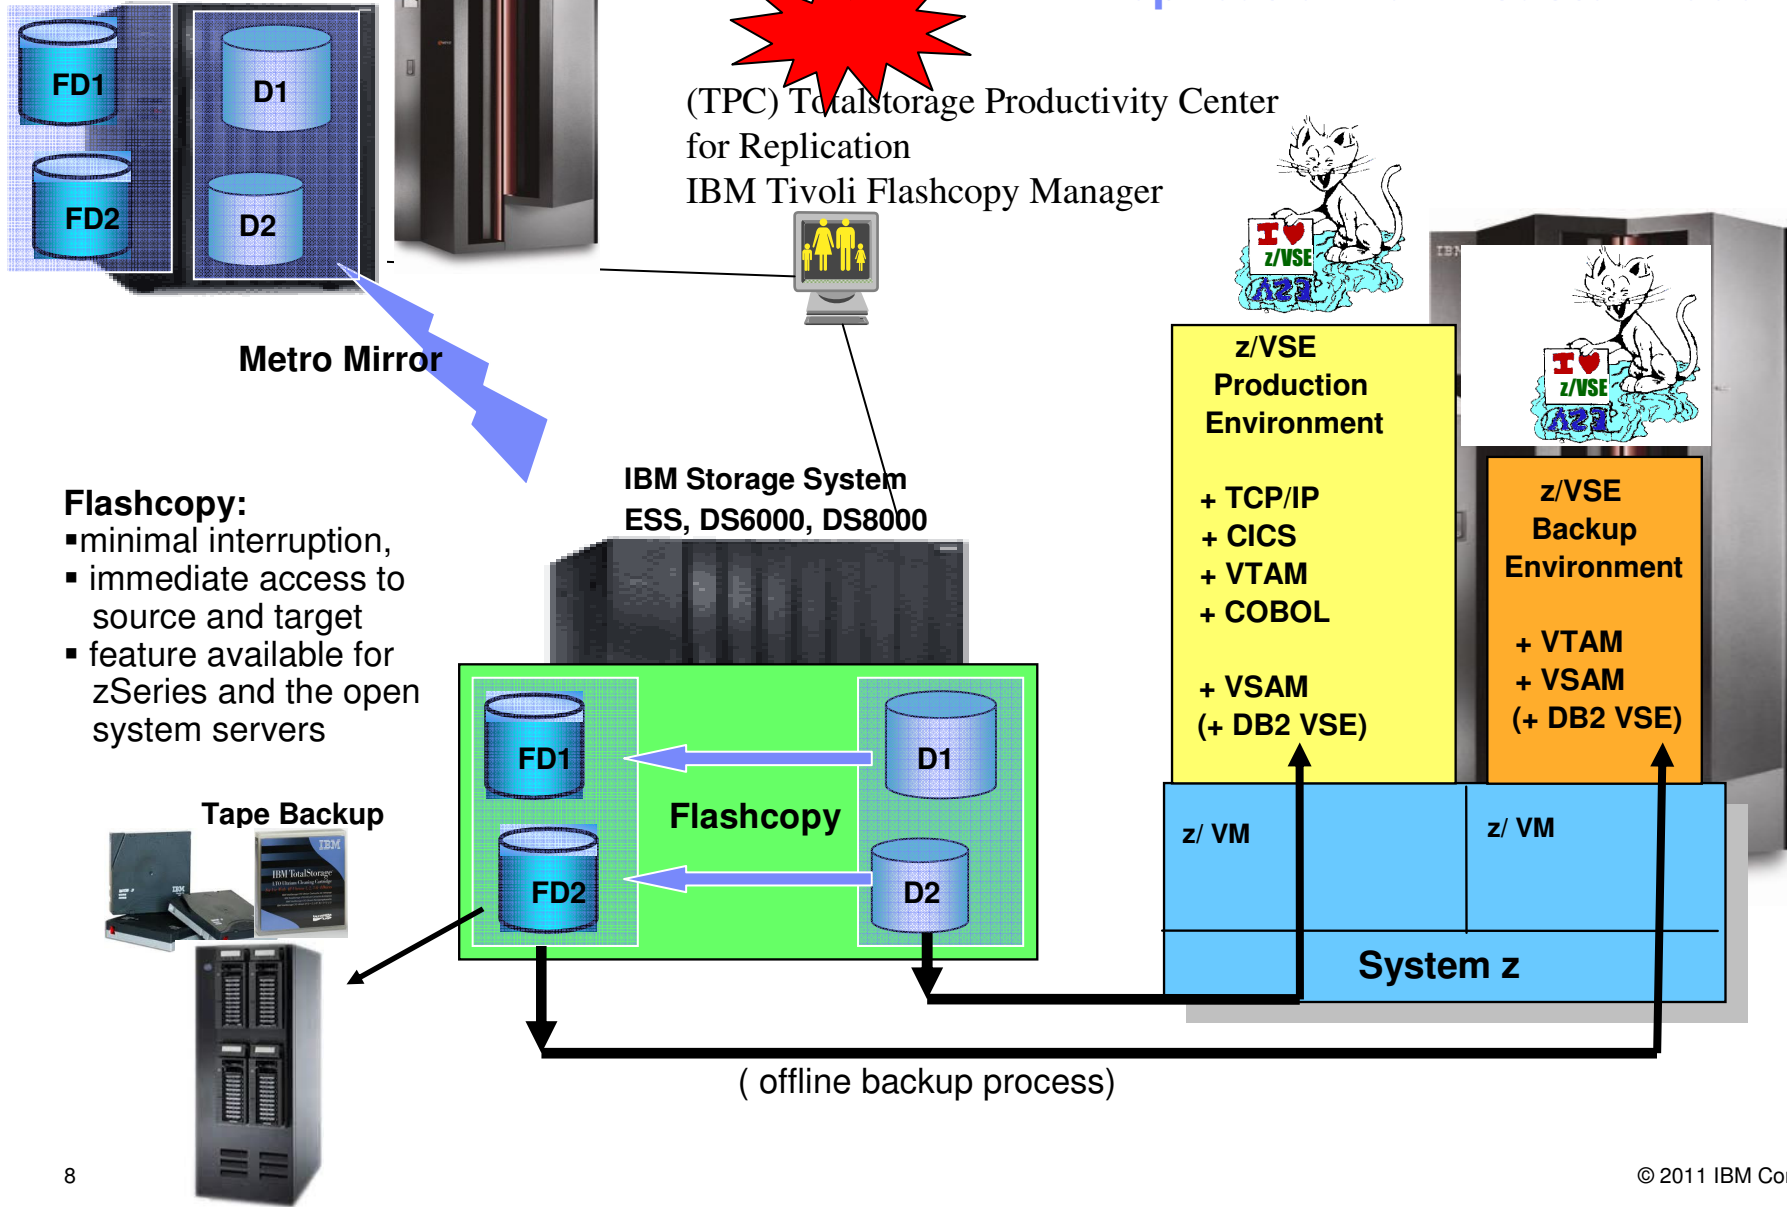
- minimal interruption,
- immediate access to source and target
- feature available for System z and the open system servers



( offline backup process )

### Disaster Recovery Site - IBM Storage System

### Backups with 2-site Real-Time Replication for Disaster Recovery



- Flashcopy:**
- minimal interruption,
  - immediate access to source and target
  - feature available for zSeries and the open system servers

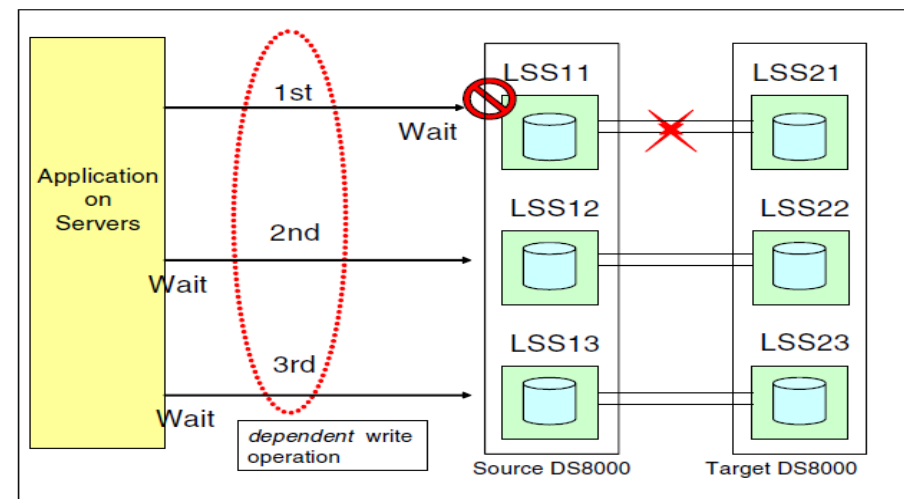
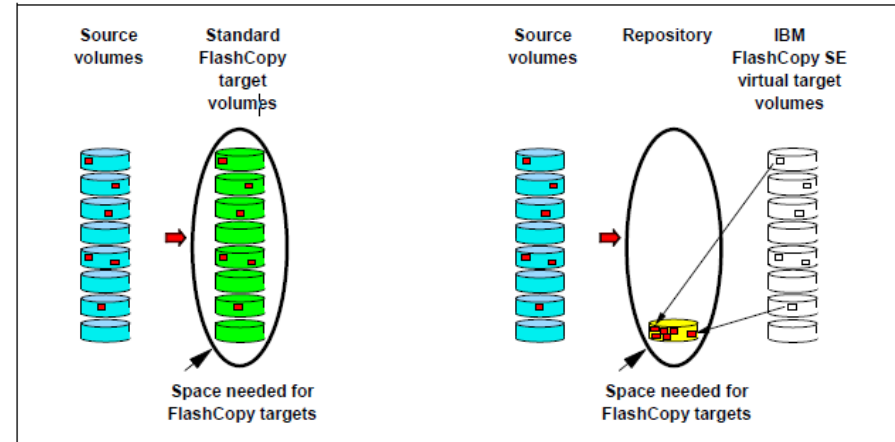


# z/VSE – System Storage Support

## DASD

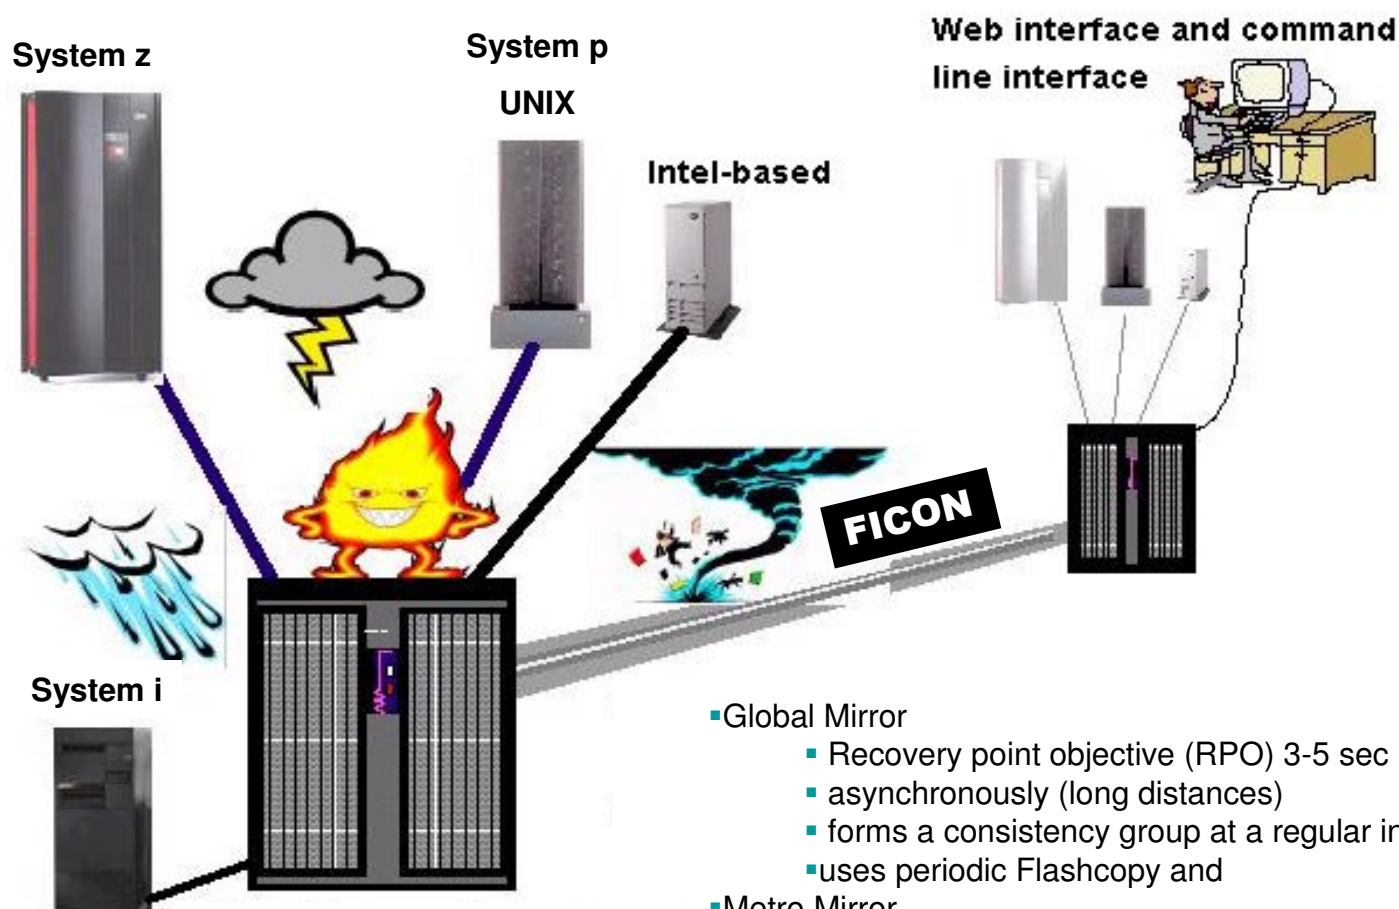
### FlashCopy

- **FlashCopy Space Efficient (SE)** - z/VSE 4.2 +
  - volumes do not occupy physical capacity when created but are seen as a virtual volumes
  - space gets allocated from a repository when data is actually written to the volume – **NOCOPY only**
  
- **consistency group** (in z/VSE 4.3)
  - applications have spread their data over multiple volumes and must be kept at a consistent level (like DB2, VSAM).
  - manage the consistency of dependent writes by FREEZE the source volume for up to 2 min



# Enterprise Storage solutions

– Disaster Recovery (DR) and the ‘Peer to Peer Remote Copy’ (PPRC) methods



**TotalStorage System families:**

- DS8000, DS6000, ESS Series
- SVC, XIV, Storwise 7000

- Global Mirror
  - Recovery point objective (RPO) 3-5 sec
  - asynchronously (long distances)
  - forms a consistency group at a regular interval
  - uses periodic Flashcopy and
- Metro Mirror
  - synchronous (< 300 km)
  - instant FlashCopy

# z/VSE – System Storage Support


## Storage Virtualization

### SCSI : (since z/VSE 3.1)

- FCP attached SCSI disks
- SCSI disks are seen in z/VSE as FBA
- DS8000 disk controller


### SVC (SAN Volume Controller) : (since z/VSE 4.2)

- SAN Volume Controller (SVC) creates a single pool of SCSI disk capacity
- Disk storage options include IBM DS8000, DS4000, etc. XIV, Storvise 7000 plus qualified systems from various non-IBM vendors




## High Availability with Disk swapping: Implementation in Linux on z

- Linux device mapper (dm) is used for logical volume mapping-LVM (target is multipath)
- Multipath tools are used to set up the device pairs during IPL automatically
- Explicit control commands to the DASD devices and the device mapper are used to do the 'HyperSwap'



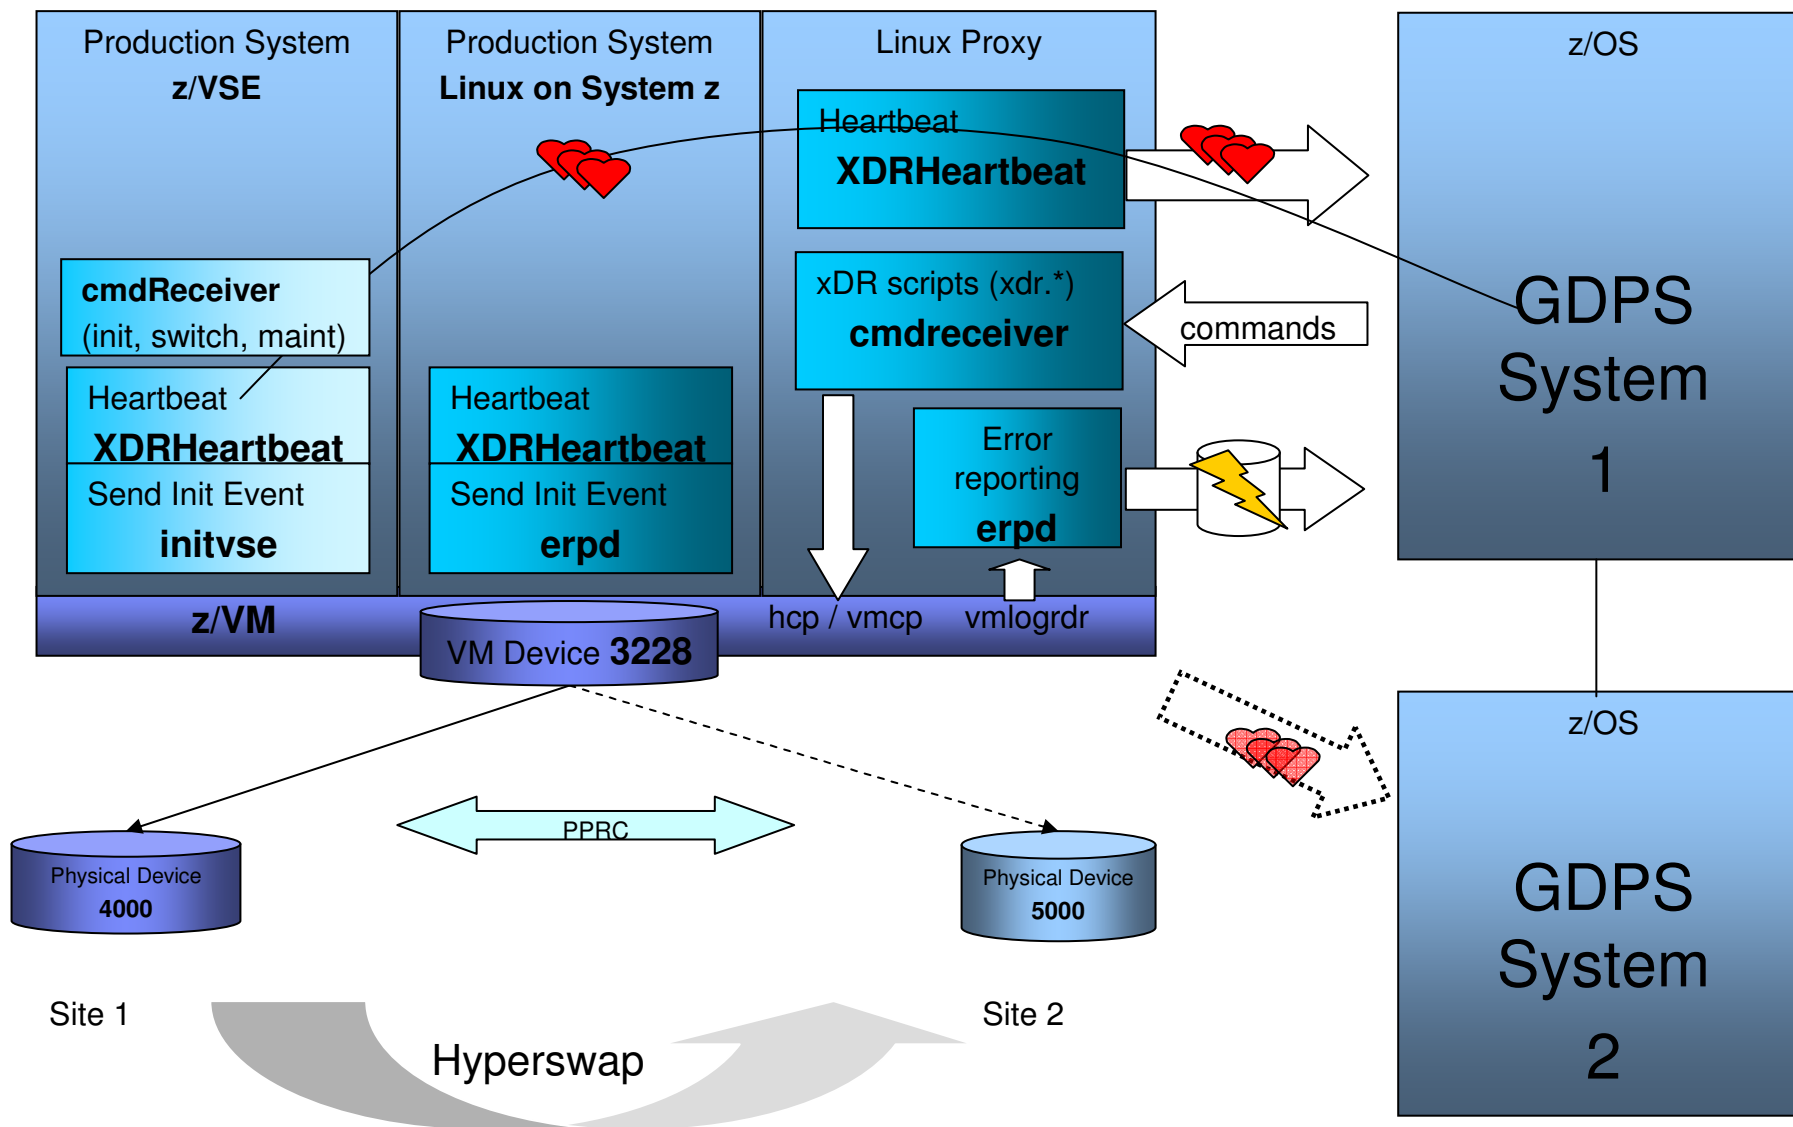
## HA with HyperSwap: Implementation in z/VM

- Standard VM mechanisms are used for logical mapping
- VM provides a CP interface that allows to configure PPRCed devices and CP commands to do the HyperSwap





# GDPS and xDR Support for z/VSE as active guest under z/VM



---

## Purpose of Tapes – why tapes are still important today

### Tapes can be REAL TAPES or Virtual Tapes

- Per usage definition tapes are for another media to keep data
- Backup of data
  - For ensured alternative long life media
  - For archiving data over a long period
  - For storing data in a extreme secured place (i.e.bunker)
- Historically shown as very reliable
- Many procedures that ensure Backup of:
  - vital data,
  - subsystems,
  - Entire DASD images
- Tape replacements have to be compatible with existing procedures
  - Virtual tapes
  - Virtual Tape Libraries (VTL, VTS)

## Tape options for z/VSE

### Standalone tape drives :

- (3480)
- 3490, 3490E
- 3590 Model A, E and H
- 3592 J1A , E05 (TS1120) and E06 (TS1130)  
(where E05 and E06 are encryption capable drives)



### Tape Library : physical

- TS3500



### Tape Library : logical

- TS7680 Protect Tier Deduplication Gateway
- TS7700 Virtualization Engine



## TS7700 Virtualization Engine - Family

- **TS7720 (disk only )** - supported with z/VSE 4.2
- **TS7740 (attaches to a TS3500)** supported with z/VSE 3.1
- Single Cluster GRID support only up to z/VSE 4. 3
- Multi Cluster GRID support with z/VSE 5. 1
- Copy Export support with z/VSE 5.1



**!New!** Disaster Recovery (DR) support with the TS7700 Virtualization Engine

## TS7700 Virtualization Engine

### TS7720 Virtualization Engine

- **disk-only** configuration
- Tape Volume Cache capacity up to 70 TB of uncompressed data
- NO physical tape library for back-end processing
- The number of logical volumes is limited by the size of the cache





## TS7700 Virtualization Engine


### **TS7740** Virtualization Engine (attaches to a TS3500)

- Maximum of 256 virtual drives (3490E)  
and 1,000,000 virtual volumes
- Web-based management tools
- up to 6 TB native tape volume cache
- Supports TS1120 / TS1130 tape drive-based encryption
- Supports logical WORM (Write Once Read Many) ,  
in z/VSE 4.3 and newer

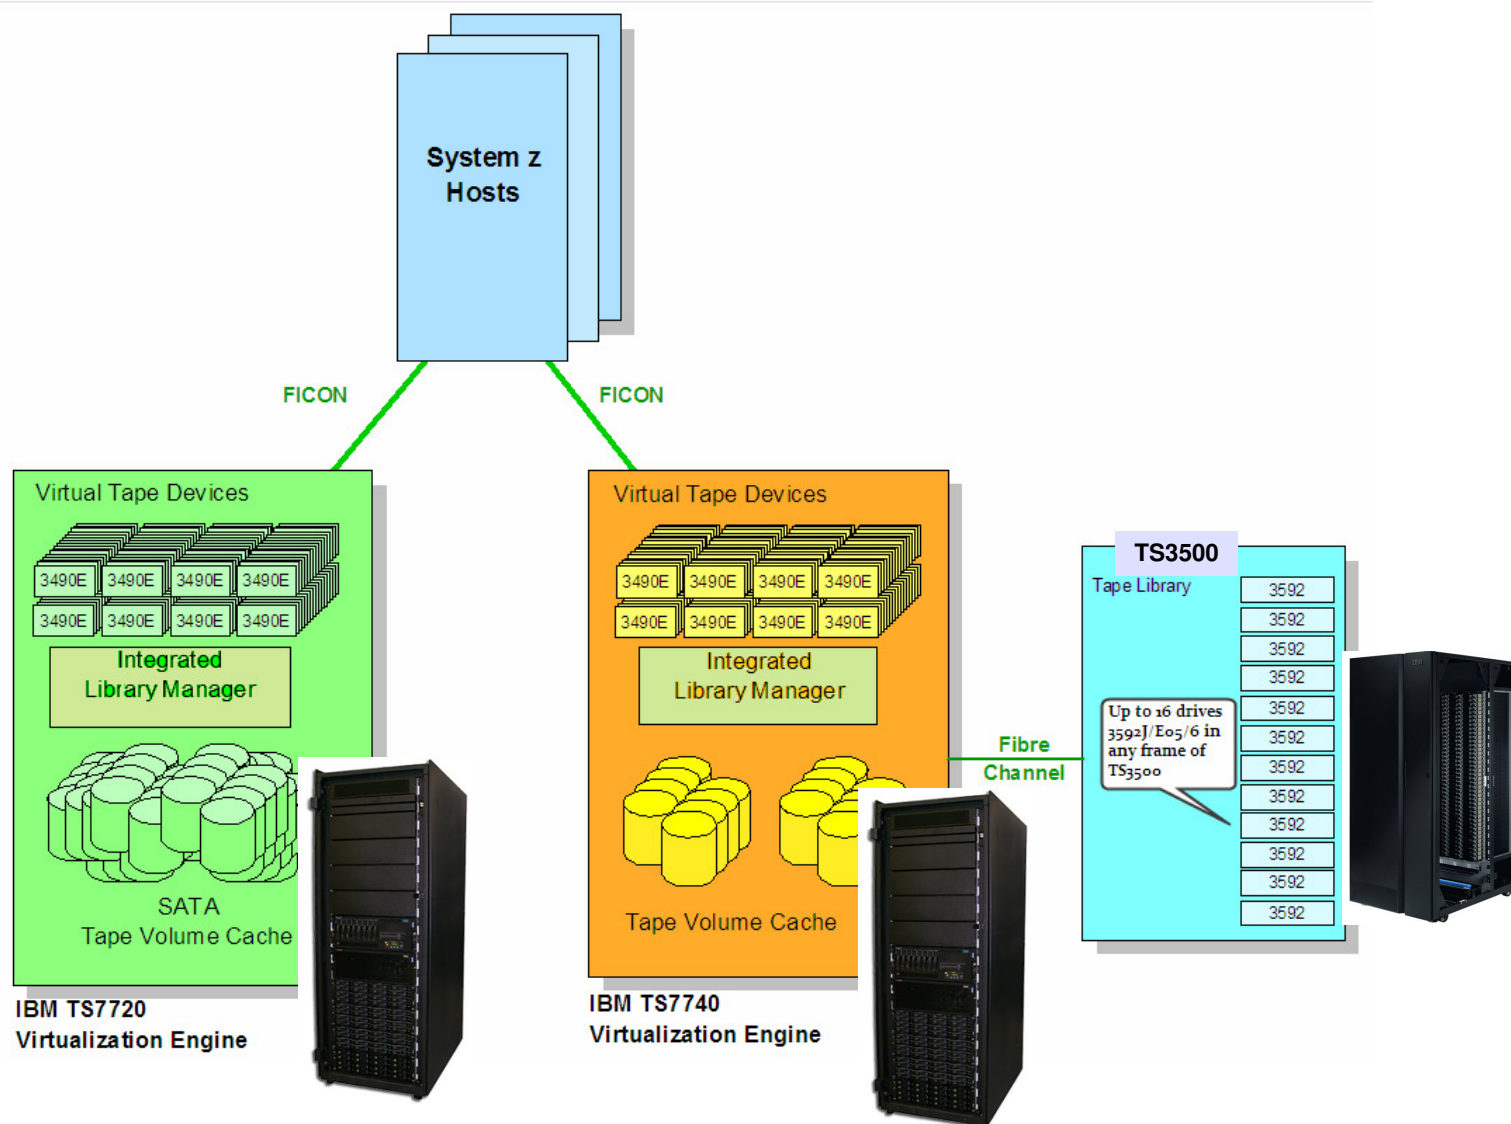


## TS3500 Tape Library – attachment option for TS7740

- The physical tape drives of the **TS3500** are managed by the **TS7740** Virtualization Engine internal management software.
- They cannot be accessed from any other attached host
- These drives are used exclusively by the TS7740 Virtualization Engine
  - for the mounts required for copying virtual volumes to stacked volumes,
  - recalling virtual volumes into the cache,
  - reclaiming stacked volume space

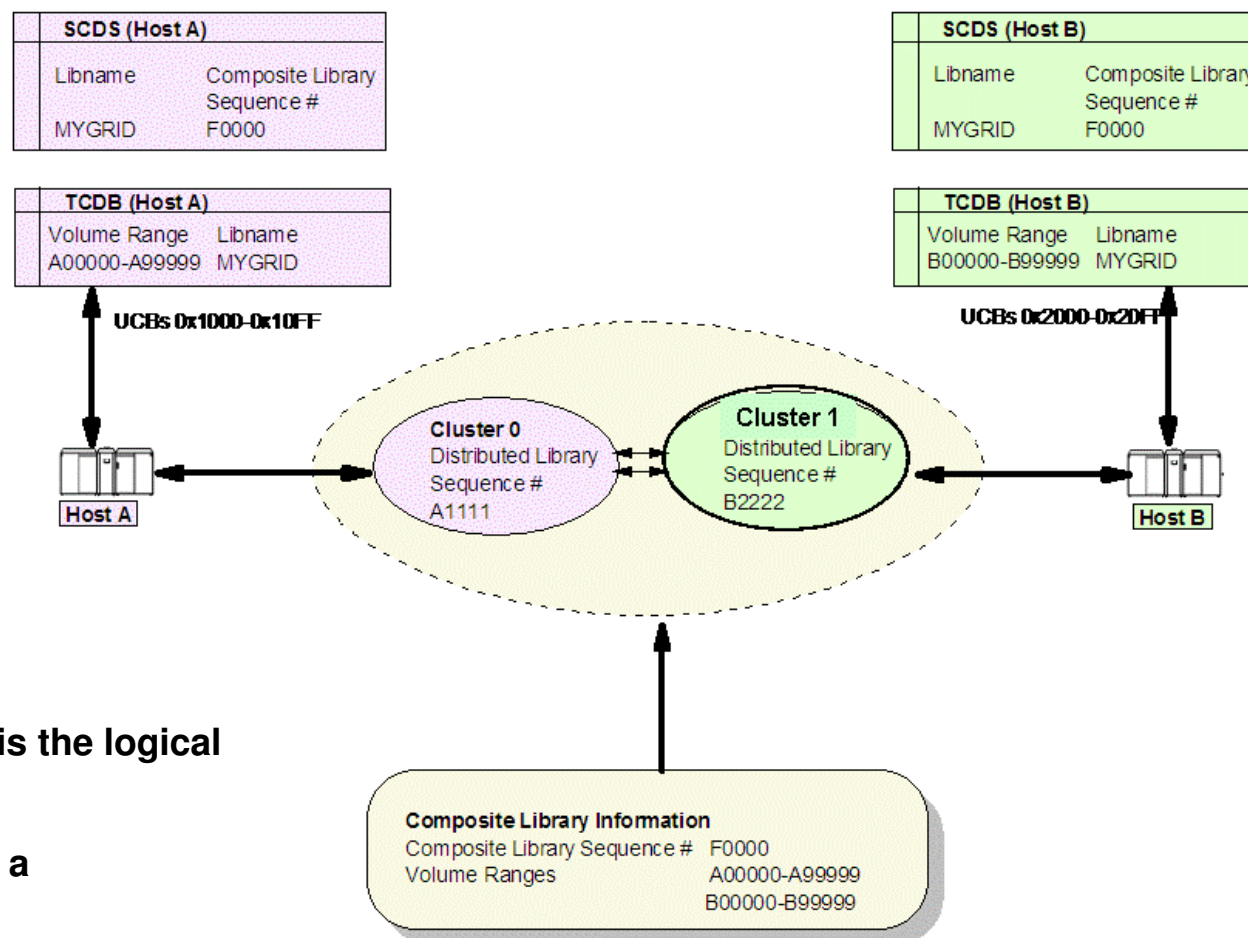


# TS7700 Virtualization Engine - Summary



# TS7700 Virtualization Engine - Grid

## TS7700 : two - Cluster GRID

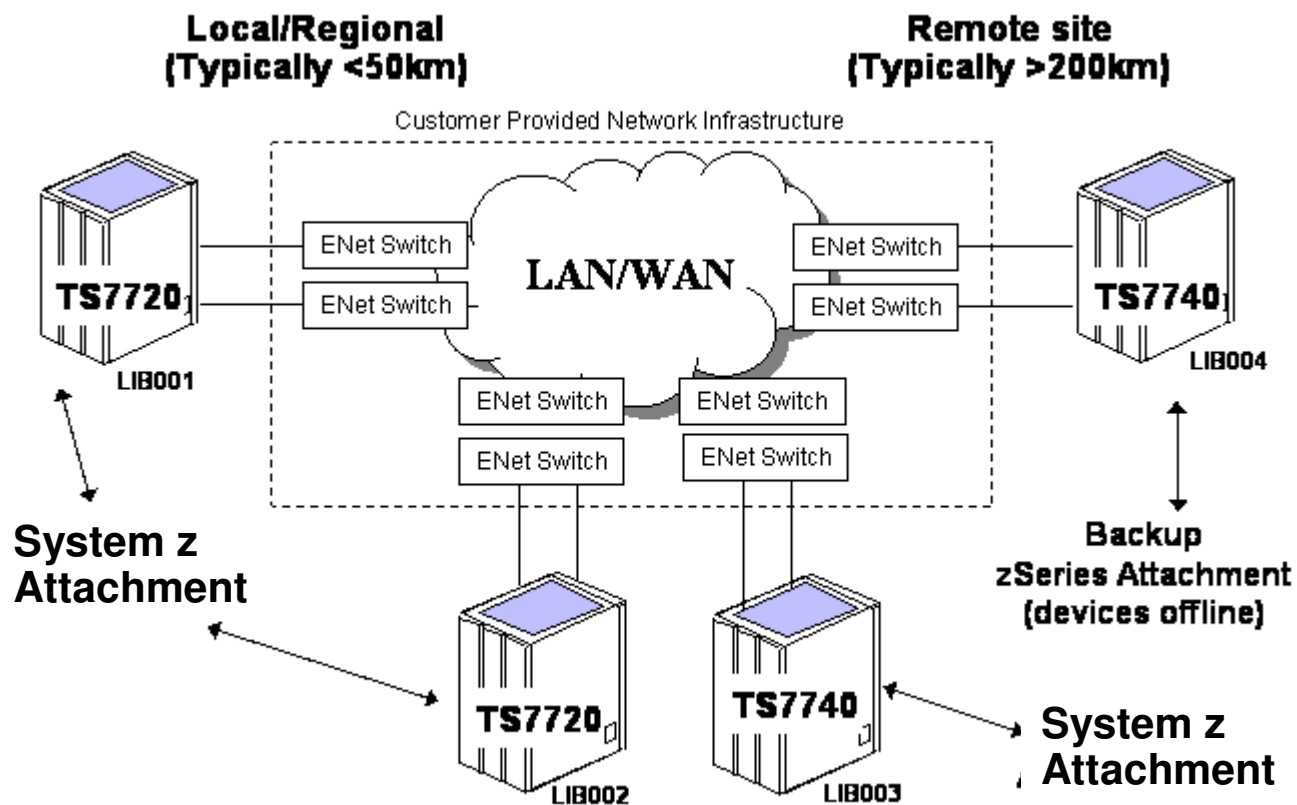


The Composite Library is the logical image of the grid

Each cluster in a grid is a Distributed Library

## TS7700 Virtualization Engine –Hybrid Grid

**Hybrid Grid** A Hybrid Grid describes a Multi Cluster Grid with an intermix of TS7720 Virtualization Engine and TS7740 Virtualization Engine clusters





---

## TS7700 Virtualization Engine – z/VSE Grid support

### **TS7700 Multi Cluster GRID support introduced in z/VSE 5.1**

- Two, three or four TS7700 Virtualization Engines can be interconnected through Ethernet links to form a *Multi Cluster Grid configuration*.
- Any data replicated between the clusters is accessible through any other cluster in a Grid configuration.
- Through remote volume access, you can reach any virtual volume through any virtual device.
- By setting up policies on the TS7700 Virtualization Engines Management Interface (MI), you define where and when you want to have multiple copies of your data.
- A Grid configuration looks like a single storage subsystem to the hosts.
- Whether a single- or multi-cluster configuration, the entire subsystem appears as a single tape library to the attached hosts. This can be described as a composite library with underlying distributed libraries.

The distributed libraries are not seen from the host.

---

## TS7700 Virtualization Engine

### **TS7700 setup in z/VSE :**

#### **Running VSE in LPAR mode :**

- Tailor and submit TLSDEF.PROC skeleton in ICCF LIB 59 (define devices and library name)

#### **Running VSE as guest under VM :**

- Run DFSMS , RMS under VM with VGS (VSE Guest server) where devices and library name are defined

#### **Both :**

- Add SYS ATL= TLS | VM
- ADD devices as 3490E in IPL proc
- Submit JCL LIBSERV MOUNT commands to get volumes mounted
- Recommendation: Make use of a tape management system , like CA DYNAM/T , BIM-EPIC or BVS to handle Tape Library commands under the cover.

---

## TS7700 Virtualization Engine

With z/VSE 5.1 TS7700 Multi Cluster **GRID support** :

- support of the Asynchronous Operator Messages (AOM) on the VSE console :

(Important AOM messages are passed from now on also with DEBUG = OFF )

AOMAP00I LIBRARY INFORMATION CUU=0C9F, **LIB=BARR88 ,CLUSTER=01 | COMPOSITE**

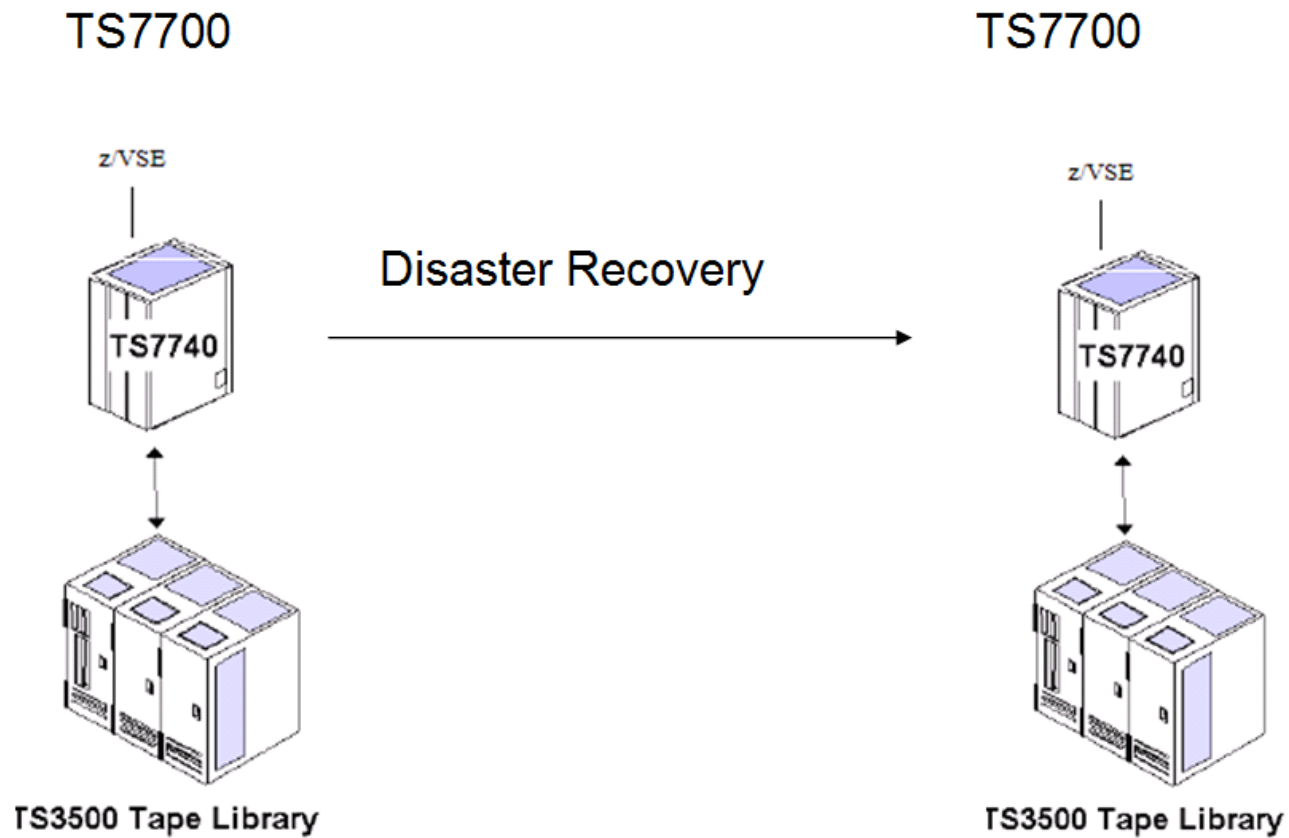
new messages for the 'extended operational states' in AOMAP13I:

### **AOMAP13I OPERATIONAL STATE CHANGE:**

- COPY OPERATIONS DISABLED
- VTS OPERATIONS DEGRADED
- IMMED COPY COMPL DEFERRD
- SERVICE PREPARATION
- FORCED PAUSE
- GRID LINKS DEGRADED
- HOST DISABLED COPY OPER
- LIMITED CACHE FREE SPACE
- OUT OF CACHE RESOURCES

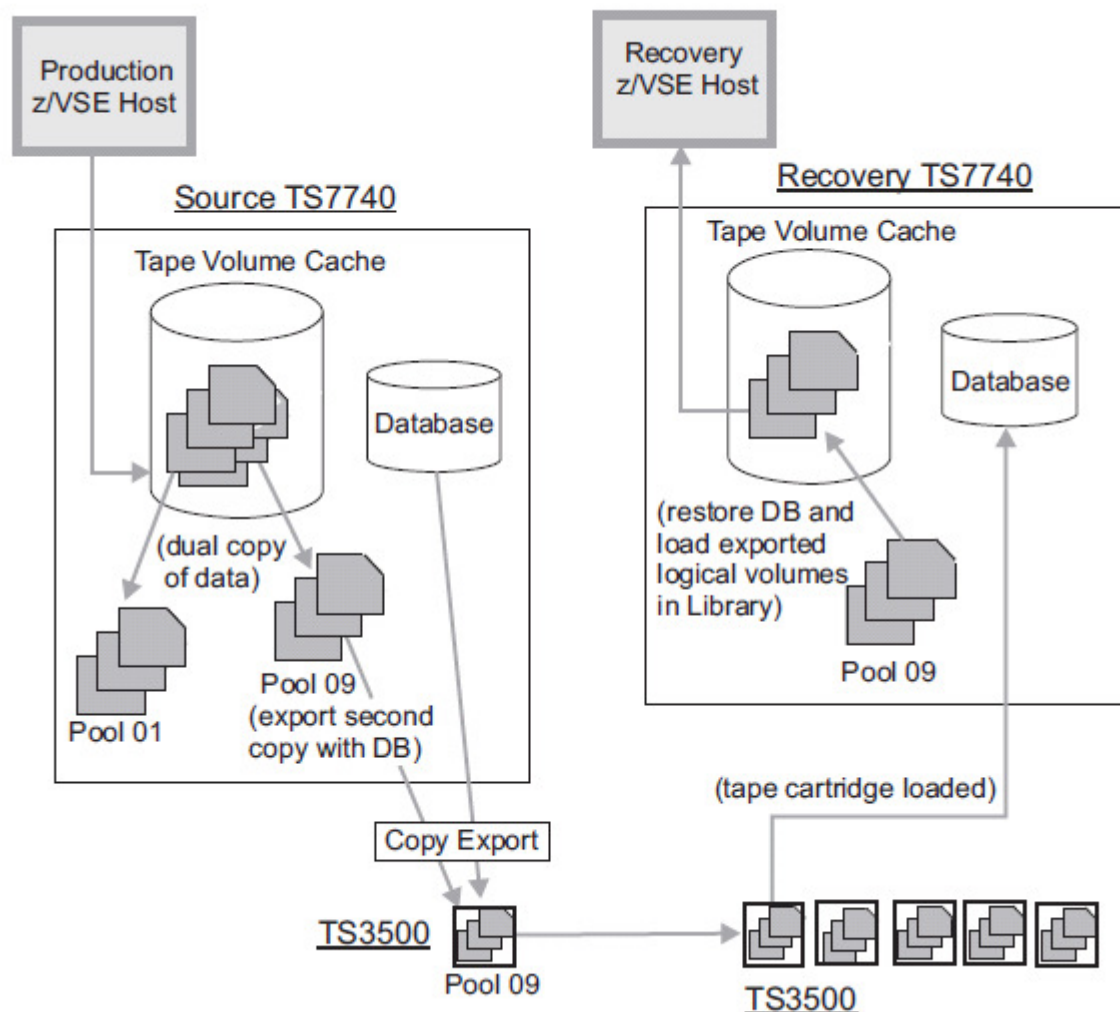
# TS7700 Virtualization Engine – D/R scenario

## TS7700 Disaster Recovery setup



# TS7700 Virtualization Engine – Copy Export

## TS7700 Copy Export support



---

## TS7700 Virtualization Engine

### **TS7700 Copy Export support** with z/VSE 5.1

- Copy export provides a new function that allows a copy of selected logical volumes written to the TS7700 to be removed and taken offsite for disaster recovery (DR) purposes.
- Since the data being exported is a copy of the logical volume, the logical volume data remains accessible by the production host systems.
- During the Copy Export operation, a copy of the current TS7700's database is written to the exported physical volumes.
- To restore access to the data on the physical volumes removed, all exported physical volumes for a source TS7700 are placed into a library that is attached to an empty TS7700.
- A disaster recovery procedure is then performed that restores access using the latest copy of the database.

---

## TS7700 Virtualization Engine

### TS7700 Copy Export support

A Copy Export operation is performed using one of these methods:

- By tailoring and then submitting job SKCOPYEX from ICCF LIB 59
- By executing the JCL LIBSERV COPYEX command using a pre-initialized logical volume

Copy Export procedure :

1. assign a range of logical volumes to a management class (pool) at the hardware Management Interface (MI)
2. change and run job skeleton SKCOPYEX ( ICCF lib 59) to create an export list file volume on a logical volume and execute JCL COPYEX command to start the Copy Export function
3. job ends with : AOMAP17I COPYEX OPERATION COMPLETE for VOLID=..... RC=
4. change and run job skeleton SKCPEXRD ( ICCF lib 59) to read the logical volume file 3 contents ( contains list of exported logical volumes)
5. insert physical volumes into the DR site TS7700 , run DR procedure at the MI



For more information, please see the z/VSE web site:  
<http://www.ibm.com/zvse/>

The screenshot shows the IBM z/VSE website interface. At the top, there's a navigation bar with the IBM logo, a search box, and a 'United States [change]' dropdown. Below this is a secondary navigation menu with links for Home, Solutions, Services, Products, Support & downloads, and My IBM. A personalized welcome message for 'Dr. Klaus Göbel' is displayed. The main content area features a breadcrumb trail: IBM Systems > Mainframe servers > Operating systems >. The central focus is the 'z/VSE' section, which includes a 'z/VSE V5.1 Preview' banner with a globe graphic and a 'Learn more' section with links to 'About z/VSE', 'News', and 'History of z/VSE'. A 'Preview: z/VSE V5.1 offers 64-bit virtual addressing for future workloads' section follows, detailing the new version's features and planned availability in the fourth quarter of 2011. A 'Note' section provides information on how to stay updated on product announcements. The left sidebar contains a vertical menu of links for various aspects of z/VSE, from 'About z/VSE' to 'Contact z/VSE'. The right sidebar offers user assistance options like 'E-mail us', 'Get the latest news about z/VSE through Twitter', and 'Mark your calendar' for the WAVV 2011 event in Colorado Springs, CO.

---

## More Information

- **TS7700 Virtualization Engine**

- Redbook TS7700 R1.7

<http://www.redbooks.ibm.com/redbooks/pdfs/sg247712.pdf>

- z/VSE Administration
- z /VSE Planning
- z/VSE System Control Statements

- **VTAPE**

- **z/VSE VTAPE Usage**

<http://www-03.ibm.com/systems/z/os/zvse/documentation/teconf.html>

# Questions?



**Wilhelm Mild**  
IBM IT Architect



IBM Deutschland Research  
& Development GmbH  
Schönaicher Strasse 220  
71032 Böblingen, Germany

Office: +49 (0)7031-16-3796  
mildw@de.ibm.com