



IBM System z Technical University



October 4–8, 2010 — Boston, MA

Disaster Scenarios and Backup procedures for z/VSE, z/VM and Linux on System z

Session ID: zDL03

Wilhelm Mild, IBM

Authorized

IBM | **Training**

28-Sep-10

© 2010 IBM Corporation

Trademarks

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

CICS*	IBM*	Virtual Image Facility
DB2*	IBM logo*	VM/ESA*
DB2 Connect	IMS	VSE/ESA
DB2 Universal Database	Intelligent Miner	z/VSE
e-business logo*	Multiprise*	VisualAge*
Enterprise Storage Server	MQSeries*	VTAM*
HiperSockets	OS/390*	WebSphere*
	S/390*	xSeries
	SNAP/SHOT*	z/Architecture
		z/VM
		z/VSE
		zSeries
		System z

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

LINUX is a registered trademark of Linus Torvalds

Tivoli is a trademark of Tivoli Systems Inc.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries

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

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

Intel is a registered trademark of Intel Corporation.

Resiliency – often called business continuity

Data Center News:

The mainframe in business resiliency

By Wayne Kernochan, Contributor
SearchDataCenter.com

Business resiliency

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

Resiliency – often called business continuity

Data Center News:

The mainframe in business resiliency

By Wayne Kernochan, Contributor
SearchDataCenter.com

“Today the mainframe is a necessary but not sufficient condition for good business resiliency.”

The business strategist usually considers the **mainframe** as first among equals:

A platform whose resiliency can be counted on,

but which must be integrated with other systems via resiliency software and hardware in order to achieve the quality-of-service, availability, and recoverability goals that the enterprise now needs.

The mainframe can do more for business resiliency strategies than simply be the best at what it does.

- In the **first place**, the mainframe can act as the testing ground for new resiliency technologies and strategies that will then flow downwards to the rest of an enterprise's integrated resiliency infrastructure.
- In the **second place**, the mainframe can act as a data resiliency hub, supervising key data resiliency tasks.

Together, these two tasks can make the mainframe again the focus of users' business resiliency strategies.

http://searchdatacenter.techtarget.com/news/article/0,289142,sid80_gci1179879,00.html

Example of a business resiliency matrix

	Operational resiliency	Disaster resiliency
Application resiliency	How quickly can I fail over an application from one system to another -- seconds, minutes, or hours?	Which applications are important to keep running when disaster strikes, and which can wait for an hour, a day, or a week?
Data resiliency	How fast and how close to the time of failure can I recover data from backups if a business-critical system fails?	What is my tradeoff between costs of mirroring to a site 150 or more miles away and risks if I use less storage or a closer site?
Network resiliency	Does my intranet have enough alternate pathways in case an electrical outage occurs?	What do I do if power lines between the local site and the disaster recovery site are knocked out?
People resiliency	Who can substitute for my systems administrator if he/she is sick?	Who can act temporarily for the (CEO, COO, CFO) if he/she has a heart attack?

Business Continuity Planning (BCP)

- is an enterprise wide planning process
- creates detailed procedures in the case of a disaster.
- BCP Plans take into account
 - processes
 - people
 - facilities
 - systems
 - external elements which could affect an organization to function
- BCP objective is to maintain critical business processes in the event of an unplanned outage
- the scope of BCP involves the IT infrastructure
- Disaster Recovery Planning is an IT-centric logical subset of the BCP process.

Disaster Recovery Planning (DRP)

- functions as a logical subset to the Business Continuity Planning (BCP) process.
- DRP process ensures continuity of operations in the event of a wide variety of disaster scenarios.
- IT operations handles DRP and BCP functions as a closely coupled process.
- The published DRP is typically an IT focused plan
 - designed to provide continuity of:
 - operations for applications,
 - databases,
 - system,
 - networks,
 - telephony,
 - staff,
 - supporting infrastructure (power, cooling, space).

Considerations for Disaster Recovery

(1) Concepts of Disaster Recovery (DR)

(2) One active production site and one for DR

(3) Two active sites with production and test

(4) Borrowed Resources for Disaster Recovery

Concepts of Disaster Recovery

A Disaster Recovery is needed if the main systems are unable to work.

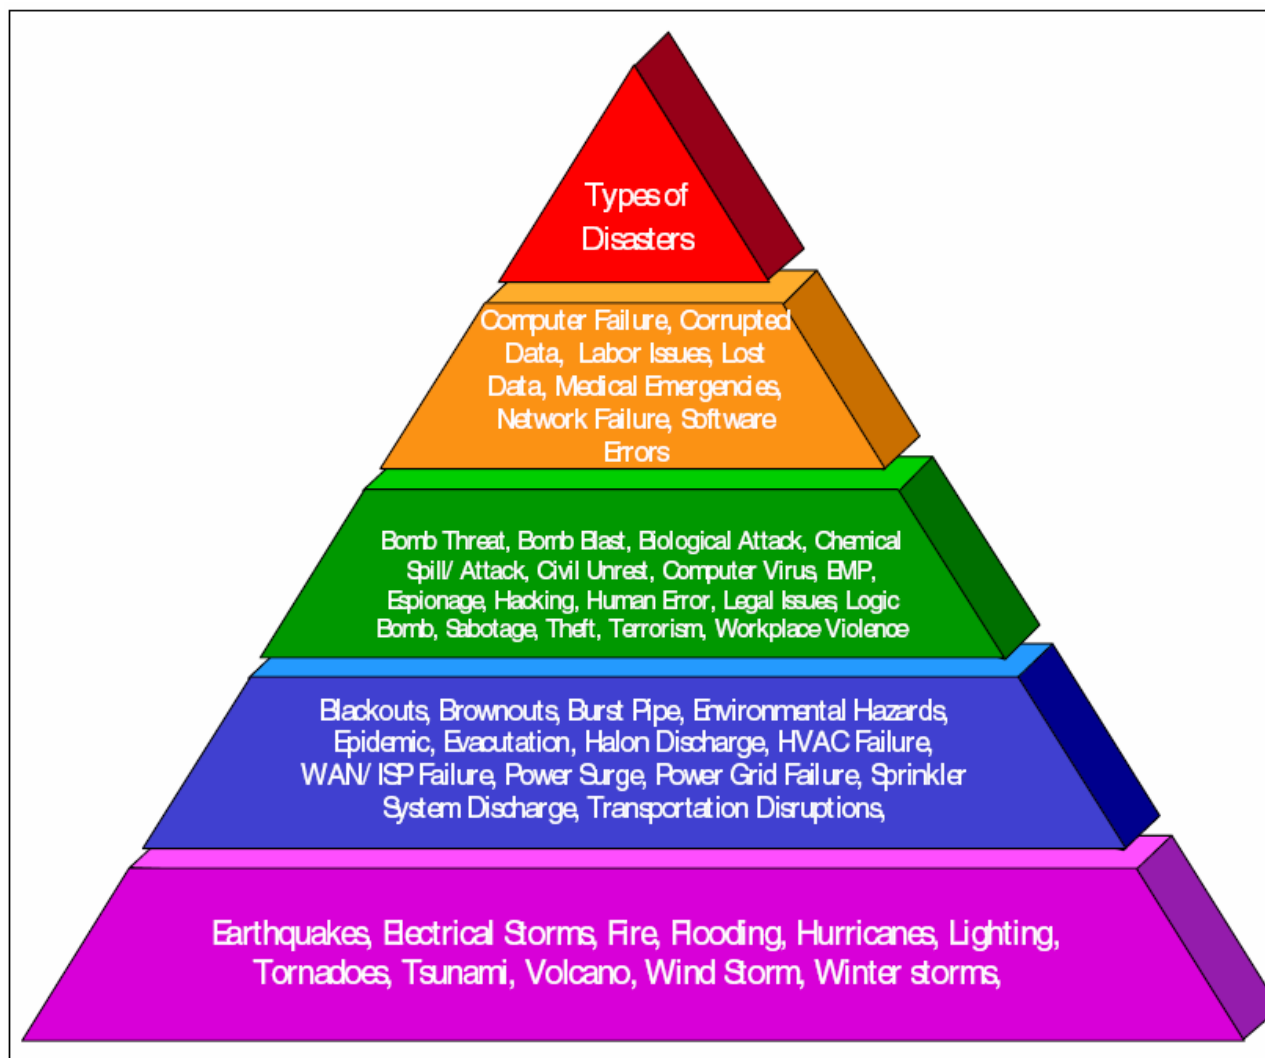
- Main machines
- Storage subsystems
- Communication of people with Data Center

Reasons for failures:

- Outage of power
- Natural catastrophe (Water, Wind, earthquake,...)
- Technical failures
 - Human error
 - Hardware errors and outages
- Political (terror)

Impact: Inability to be productive – loss of money

Types of disasters

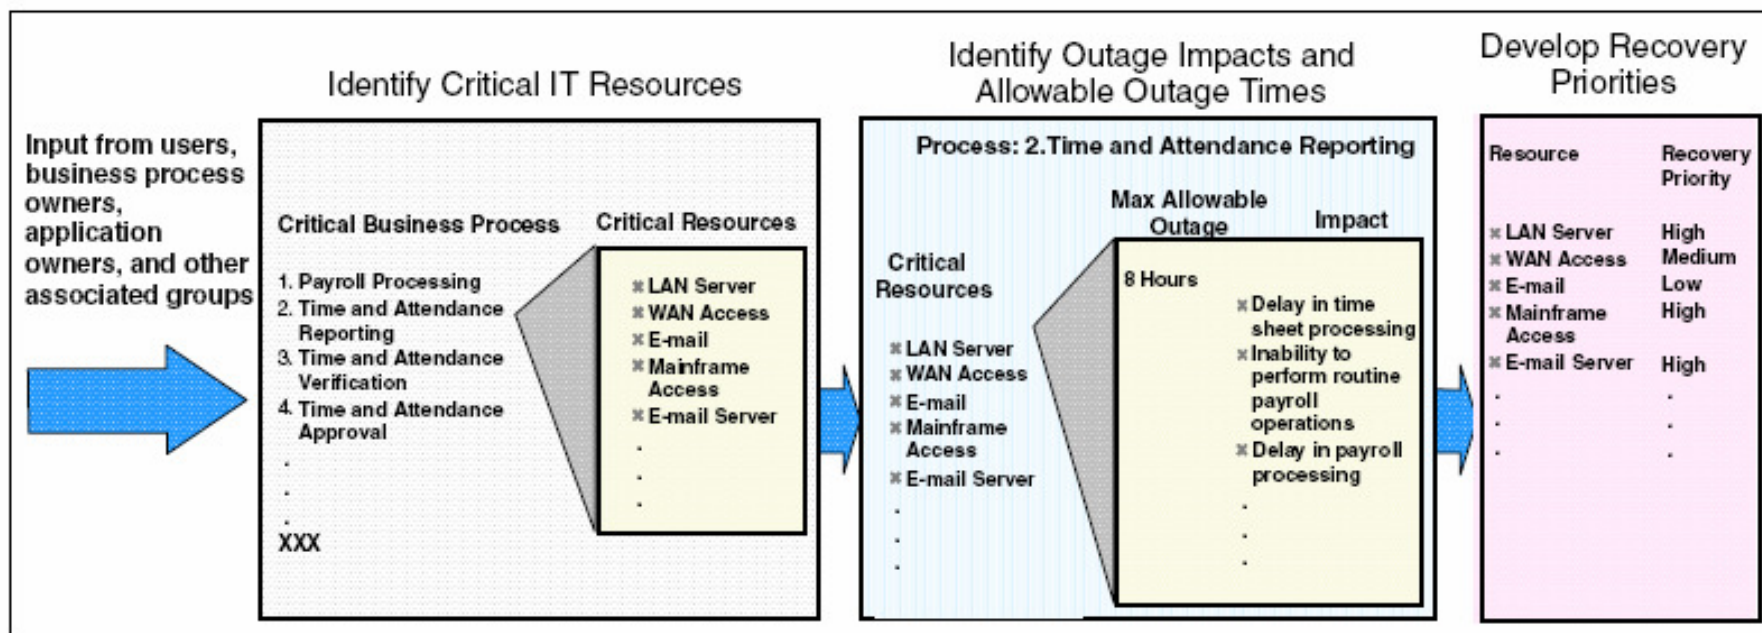


Major discussion areas

- Possible Systems affected
 - Type of systems, relation, how many systems participate in the DR scenario
- System positions – Geographically
 - Distance between them for data mirroring
- Connectivity and attachments
 - Ability to replace each other w/o application/user adjustments
- Separation of Data Stores
 - Logical connected data should reside on same side
- Network topology
 - Types of networks to be interconnected
- Operating Systems and application Landscape
 - Application execution based on operating systems

The Business impact analysis (BIA)

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



Example of the Business Impact Analysis process

Objectives for Disaster Recovery

Following Objectives are the same for Systems and Storage

- Minimize time of outage
- Minimize affected systems in case of a disaster
- Minimize effort for a restart


Required knowledge in case of a DR:

- Special Communication hardware for the DR case – to avoid busy lines from users
- Documentation of DR 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

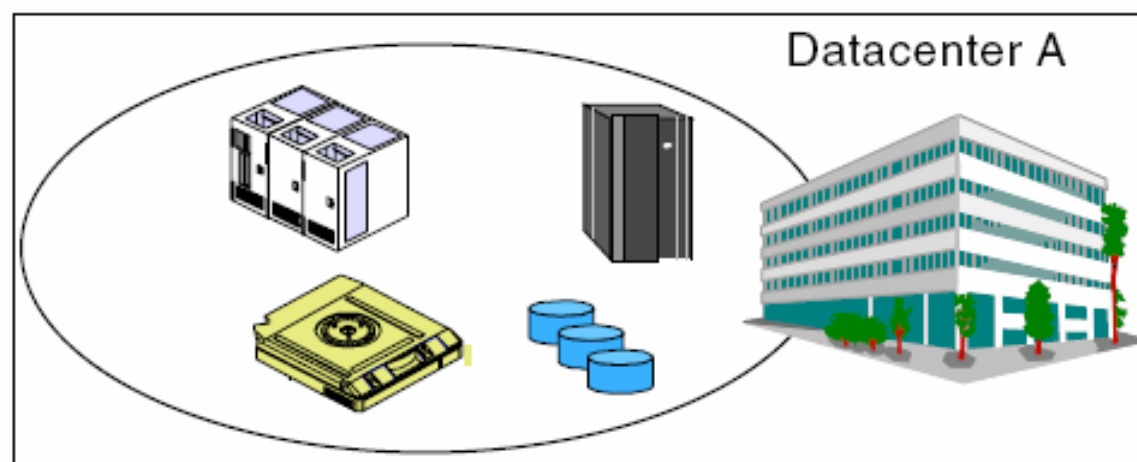
TREND: RPO = 0

Network Recovery Objective (NRO)

Time requirements for network availability.

Tier 0 – NO DR scenario in place

- Standalone datacenter
- NO DR site at all
- A failure would mean to use existing Backups and rebuild the environment

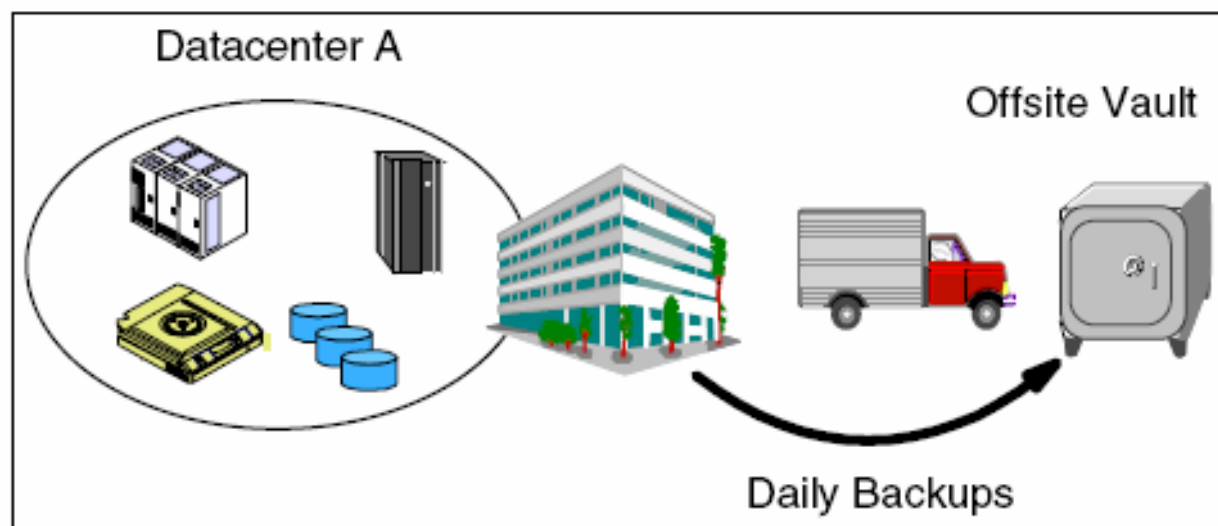


Tier 0 - Do nothing, no offsite data

Note: The typical length of recovery time in this instance is unpredictable. In many cases complete recovery of applications, systems, and data is never possible.

Tier 1: Offsite store of Backups procedure

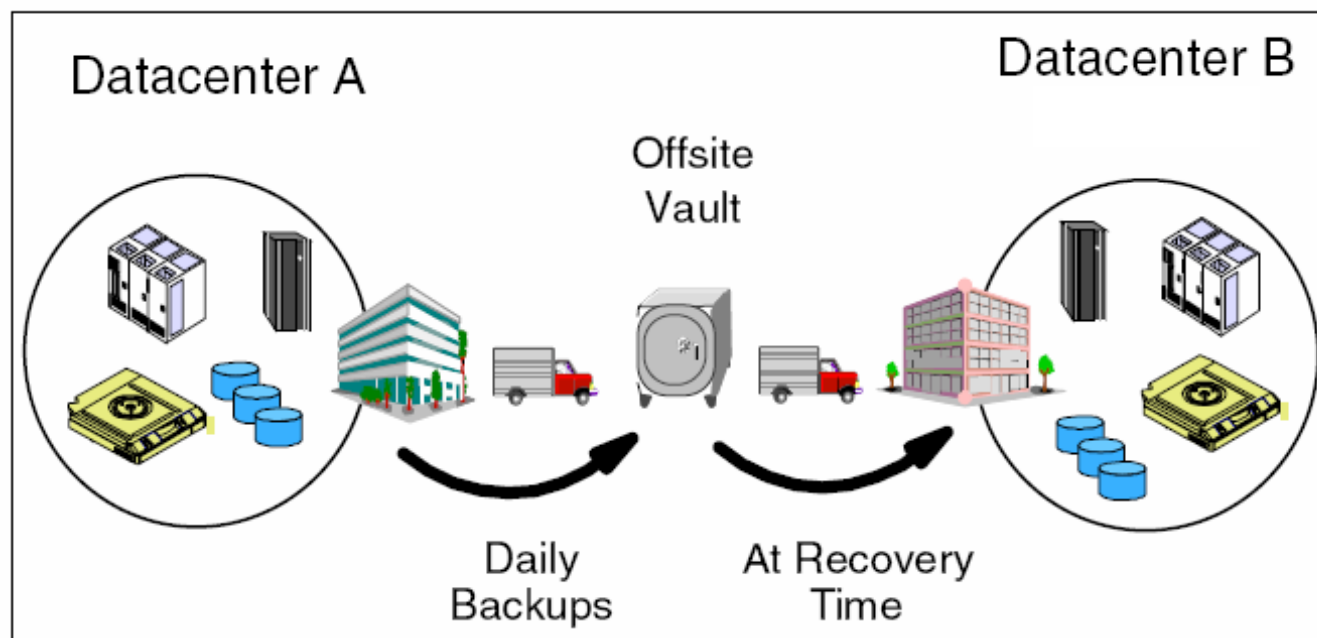
- Pickup Truck Access Method (PTAM) because vaulting and retrieval of data is typically handled by couriers
- this is a relatively inexpensive option
- used by many sites



Tier 1 - Offsite vaulting (PTAM)

Note: The typical length of time for recovery depends on the agreements.

Tier 2: Offsite data center for DR situations

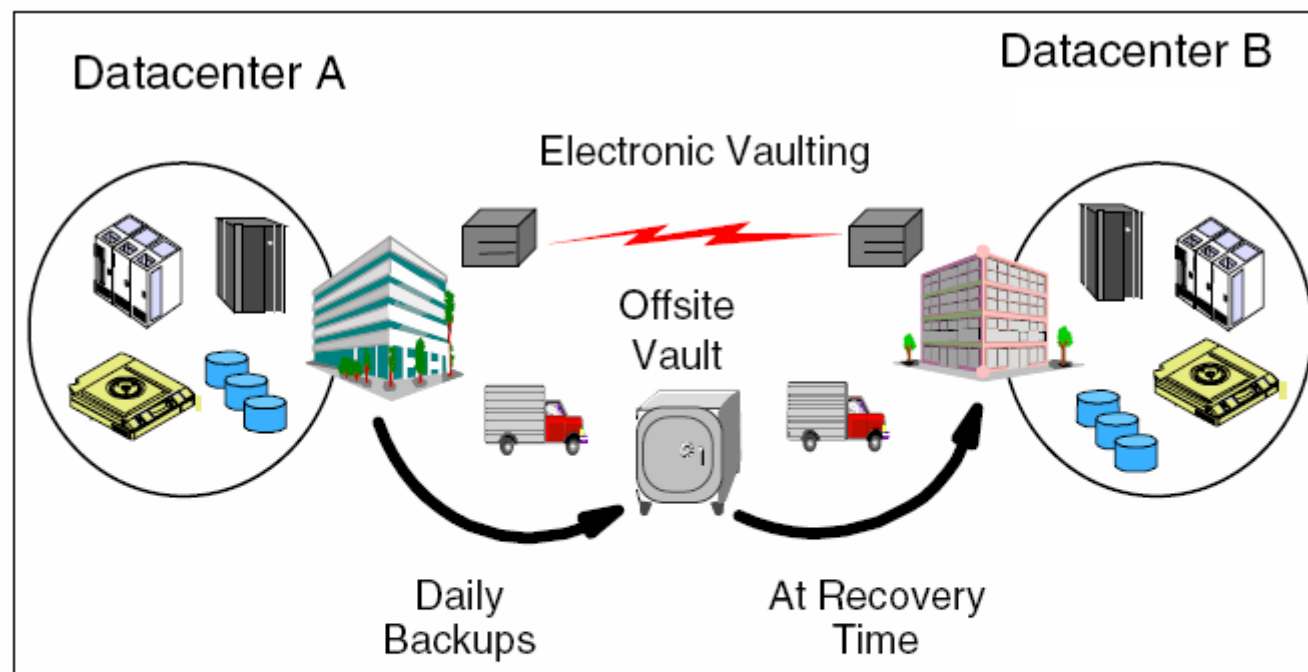


Tier 2 - Offsite vaulting with a hot site (PTAM + hot site)

Tier 2 installations rely on a courier (PTAM) to get data to an offsite storage facility. In the event of a disaster, the data at the offsite storage facility is moved to the hot site and restored onto the backup hardware provided. Moving to a hot site increases the cost but reduces the recovery time significantly. The key to the hot site is that appropriate hardware to recover the data (for example, a compatible tape device) is present and operational.

Note: The typical length of time for recovery is normally more than one day.

Tier 3: Electronic synchronization and regular Backups

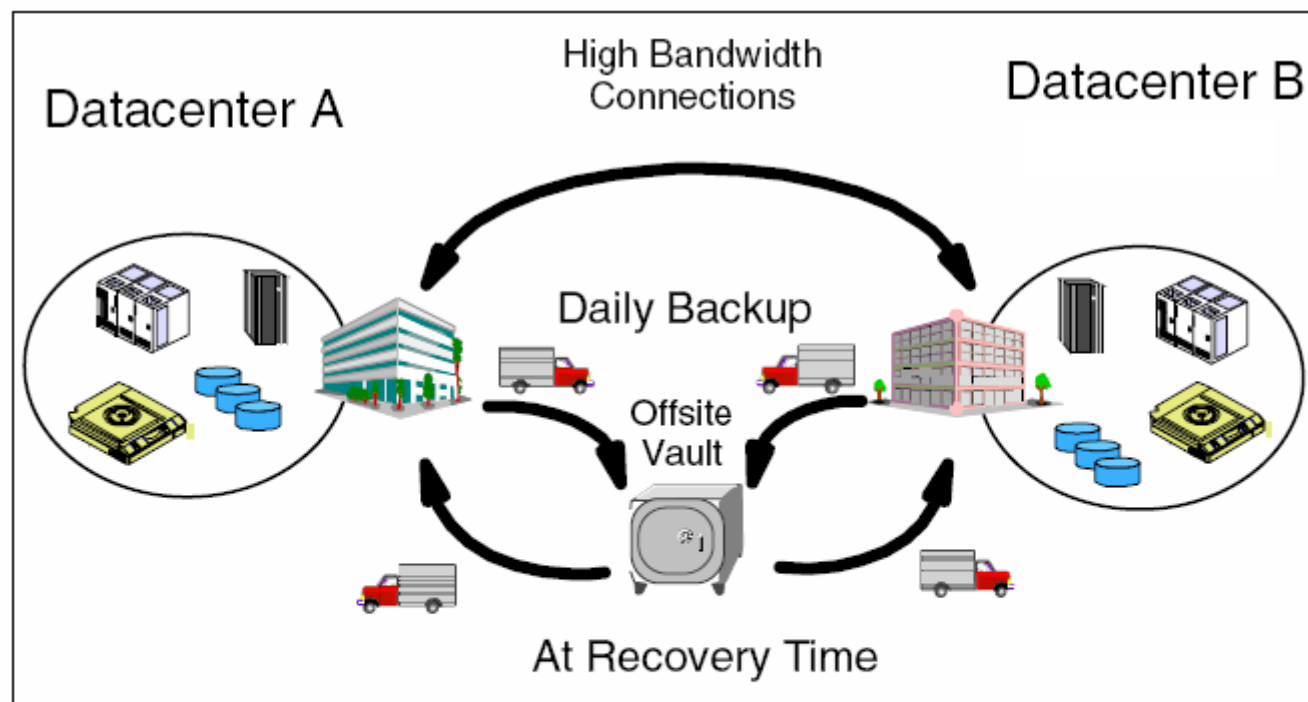


Tier 3 - Offsite electronic vaulting

The hot site is kept running permanently, thereby increasing the cost. As the critical data is already being stored at the hot site, the recovery time is once again significantly reduced. Often, the hot site is a second data center operated by the same firm or a Storage Service Provider.

Note: The typical length of time for recovery is normally about one day

Tier 4: Mixed electronic synchronization and offsite backup




Tier 4 - Electronic vaulting with hot site (active secondary site)

In this scenario, the workload may be shared between the two sites. There is a continuous transmission of data between the two sites with copies of critical data available at both sites. Any other non-critical data still needs to be recovered from the offsite vault via courier in the event of a disaster.

Note: The typical length of time for recovery is normally less than one day

Tier 5: Modern Global/Metro Mirror synchronization




Tier 5 - Two-site, two-phase commit

Tier 5 also requires partially or fully dedicated hardware on the secondary platform with the ability to automatically transfer the workload over to the secondary platform. We now have a scenario where the data between the two sites is synchronized by remote two-phase commit. The critical data and applications are therefore present at both sites and only the in-flight data is lost during a disaster. With a minimum amount of data to recover and reconnection of the network to implement, recovery time is reduced significantly.

Note: The typical length of time for recovery is minutes to max a couple of hours

Tier 6: Par Excellence: Parallel Sysplex scenario



Tier 6 - Zero data loss (advanced coupled systems)

This is the most expensive Disaster Recovery solution as it requires coupling or clustering applications, additional hardware to support data replication, and high bandwidth connections over extended distances. However, it also offers the speediest recovery by far.

Note: The typical length of time for recovery is zero – automatic failover takes place

System environment Agreements for DR

IBM special Agreements for Recovery:

- IBM Customer Agreement (ICA),
 - IBM Agreement for Programs (IAP),
 - International Program License Agreement (IPLA)
-
- The level of use acquired is documented in a Proof of Entitlement (PoE)
 - “one install”, (w/o other restrictions), allows a copy of the program on more than one machine under the customer’s control, but only one program is authorized to be in use at any given time. Or customer may use the program **temporarily** on another machine, if the Designated Machine is inoperable.

It applies to all programs licensed under these agreements for:

- Backup use,
- Disaster Recovery (DR),
- BRS when a backup and recovery service is involved

System environment Agreements for DR

IBM defines 3 types of situations for programs running or resident on backup machines: "cold"; "warm"; and "hot".

Accepted actions concerning the copy of the program used for backup purposes:

- ❖ cold - a copy of the program may be stored for backup purposes on a machine as long as the program has not been started.
 - ❖ There is no charge for this copy.
- ❖ warm - a copy of the program may reside for backup purposes on a machine and is started, but is "idling", and is **not doing any work of any kind**.
 - ❖ There is no charge for this copy.
- ❖ hot - a copy of the program may reside for backup purposes on a machine, is started and is doing work. However, this program must be ordered.
 - ❖ There is a charge for this copy.

System environment Agreements for DR - continued

For the 'warm' situation - "Doing Work", includes:

- production,
 - development,
 - program maintenance,
 - testing
 - mirroring of transactions,
 - updating of files,
 - synchronization of programs, data or other resources (e.g., active linking with another machine, program, data base or other resource, etc.)
 - any activity or configurability that would allow an active hot-switch or other synchronized switch-over between programs, data bases, or other resources to occur.
- A scheduled hardware outage, such as preventive maintenance or installation of upgrades, is NOT considered a backup situation.

System environment Agreements for DR – continued (2)

Preparation for emergency backup situations requires periodic tests – based on the requirements of system availability.

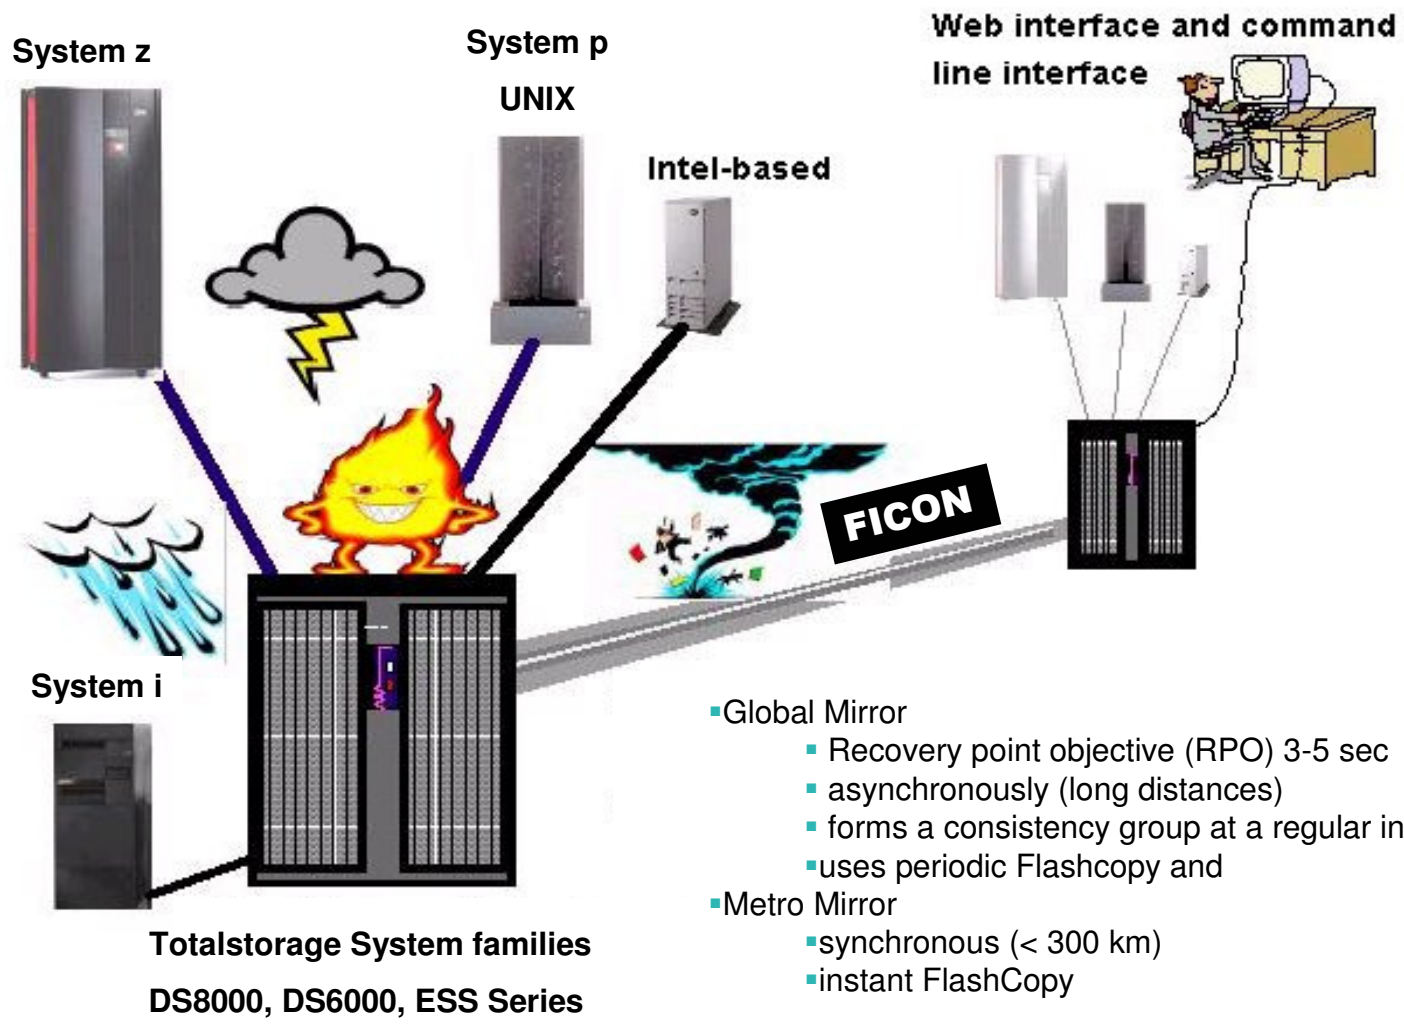
No extra program charges apply for these tests if:

- The number is appropriate (e.g., 1-3 tests per year)
- The duration is adequate, (e.g. 2 to 3 days per test).
- For more frequent tests required (e.g. for on-line systems running 24x7 critical customer business operation)
 - a shorter duration without exceeding the total hours of above guidelines.

There can be no productive output or work done from the tests and no development, program maintenance or testing as part of the tests. IBM has the right to review the customer's rationale for not licensing the IBM Program copy for the backup environment.

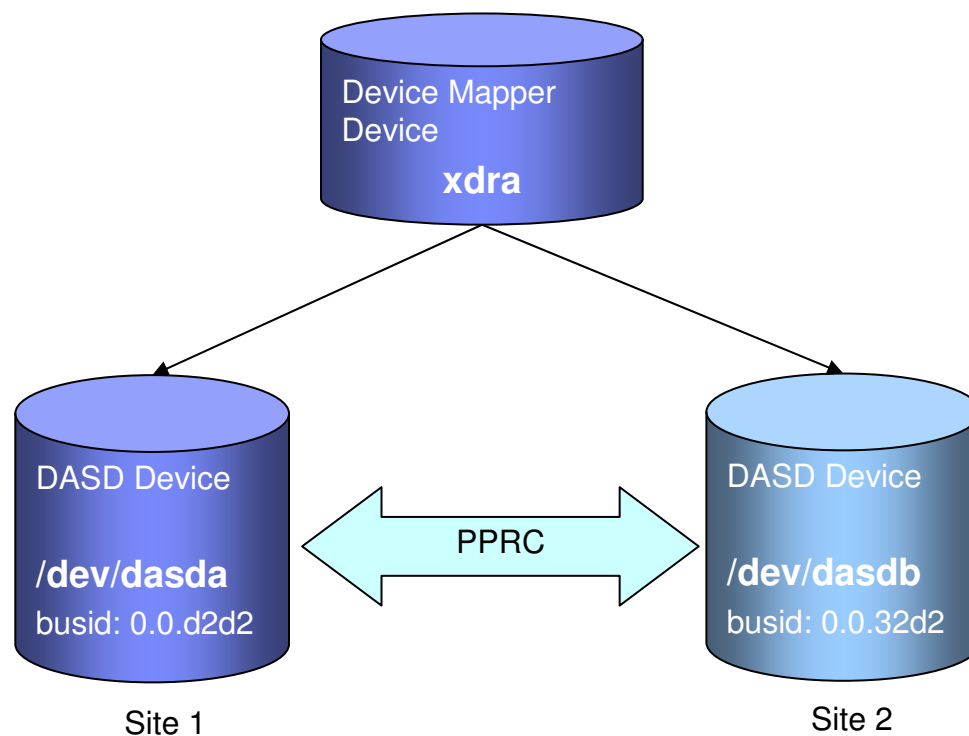
Enterprise Storage solutions

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



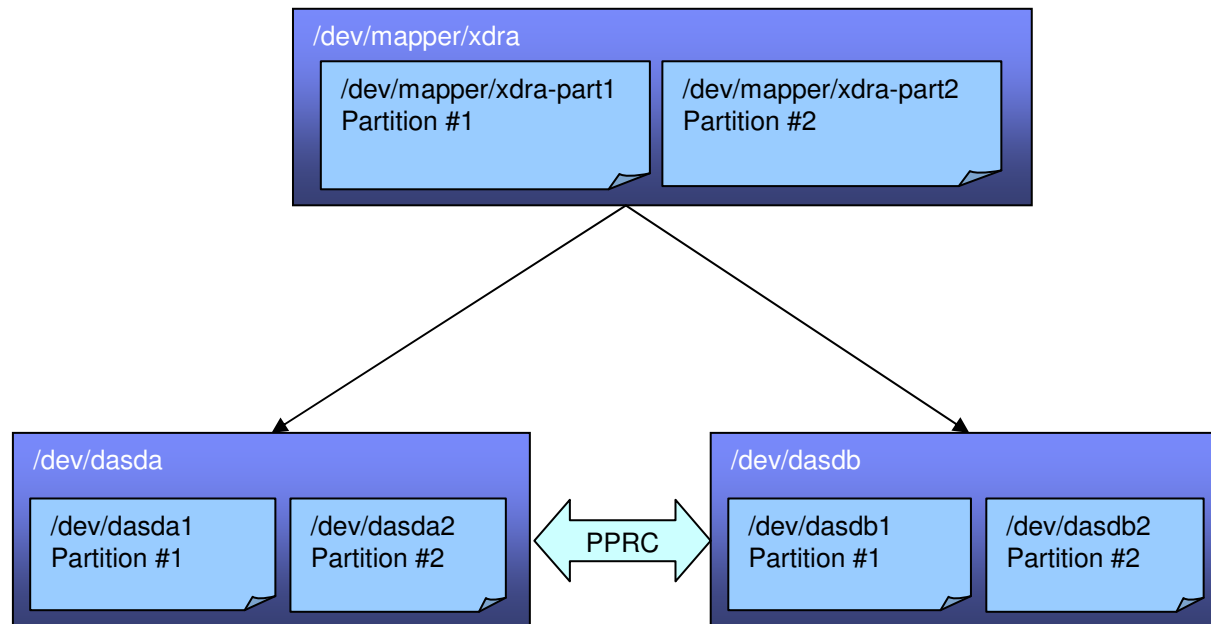
HyperSwap: Implementation in Native Mode for Linux

- Linux device mapper is used for logical mapping (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



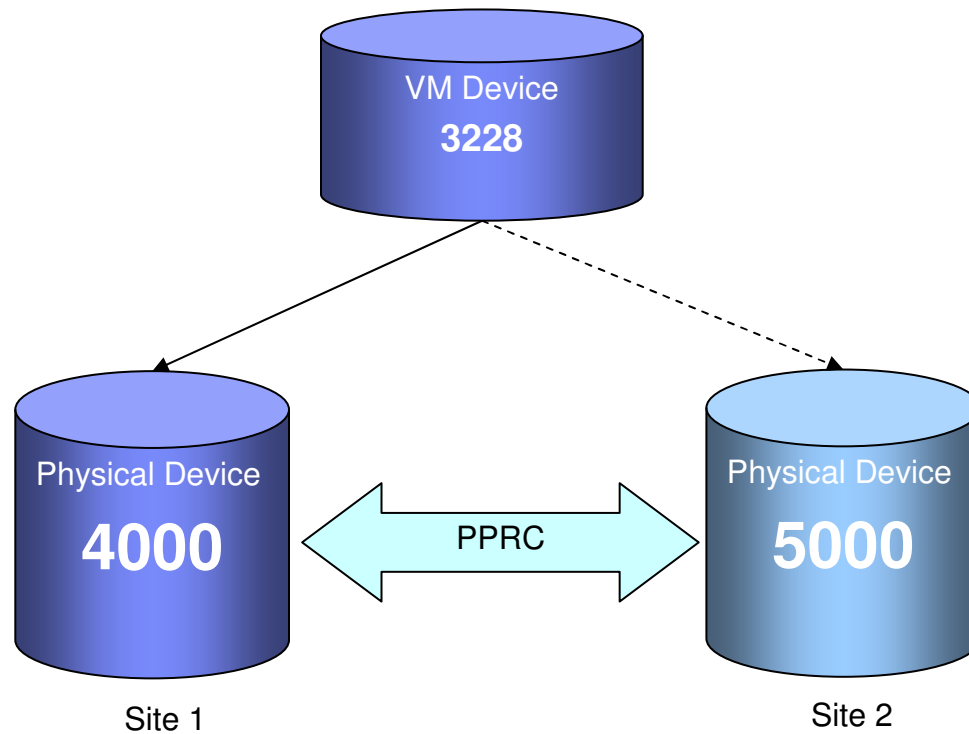
Multipath: Mapping of Partitions

- The multipath target maps the complete device
- Partitions are mapped using linear targets which point to the multipath device
- Naming convention is to append `-part<partition number>` to the device name

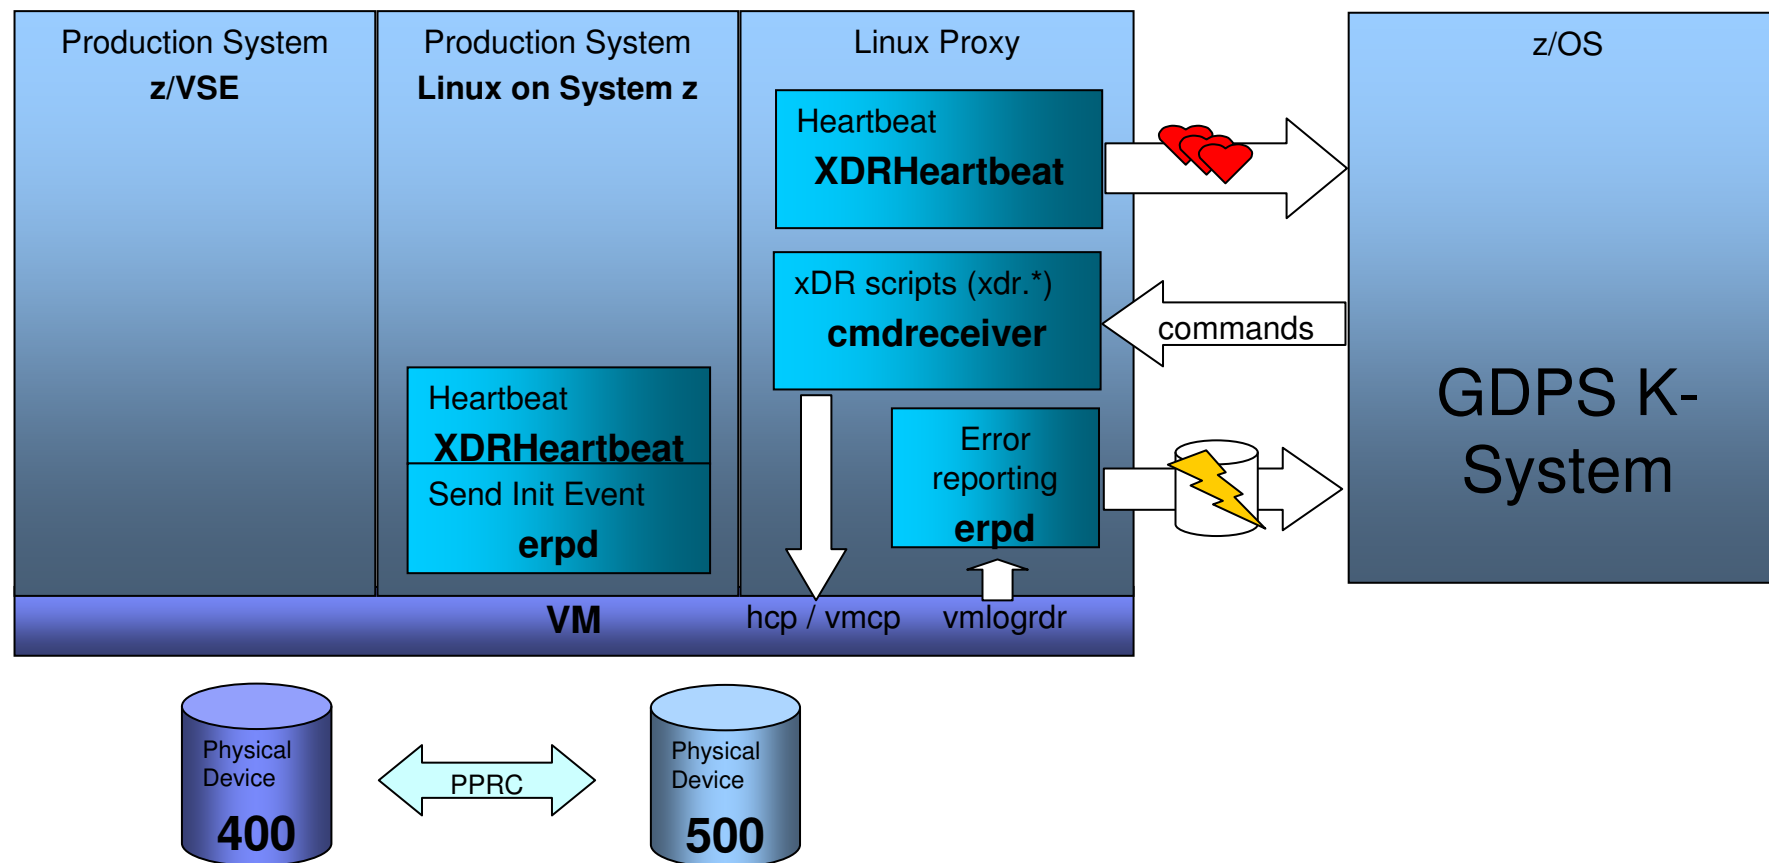


HyperSwap: Implementation in z/VM

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



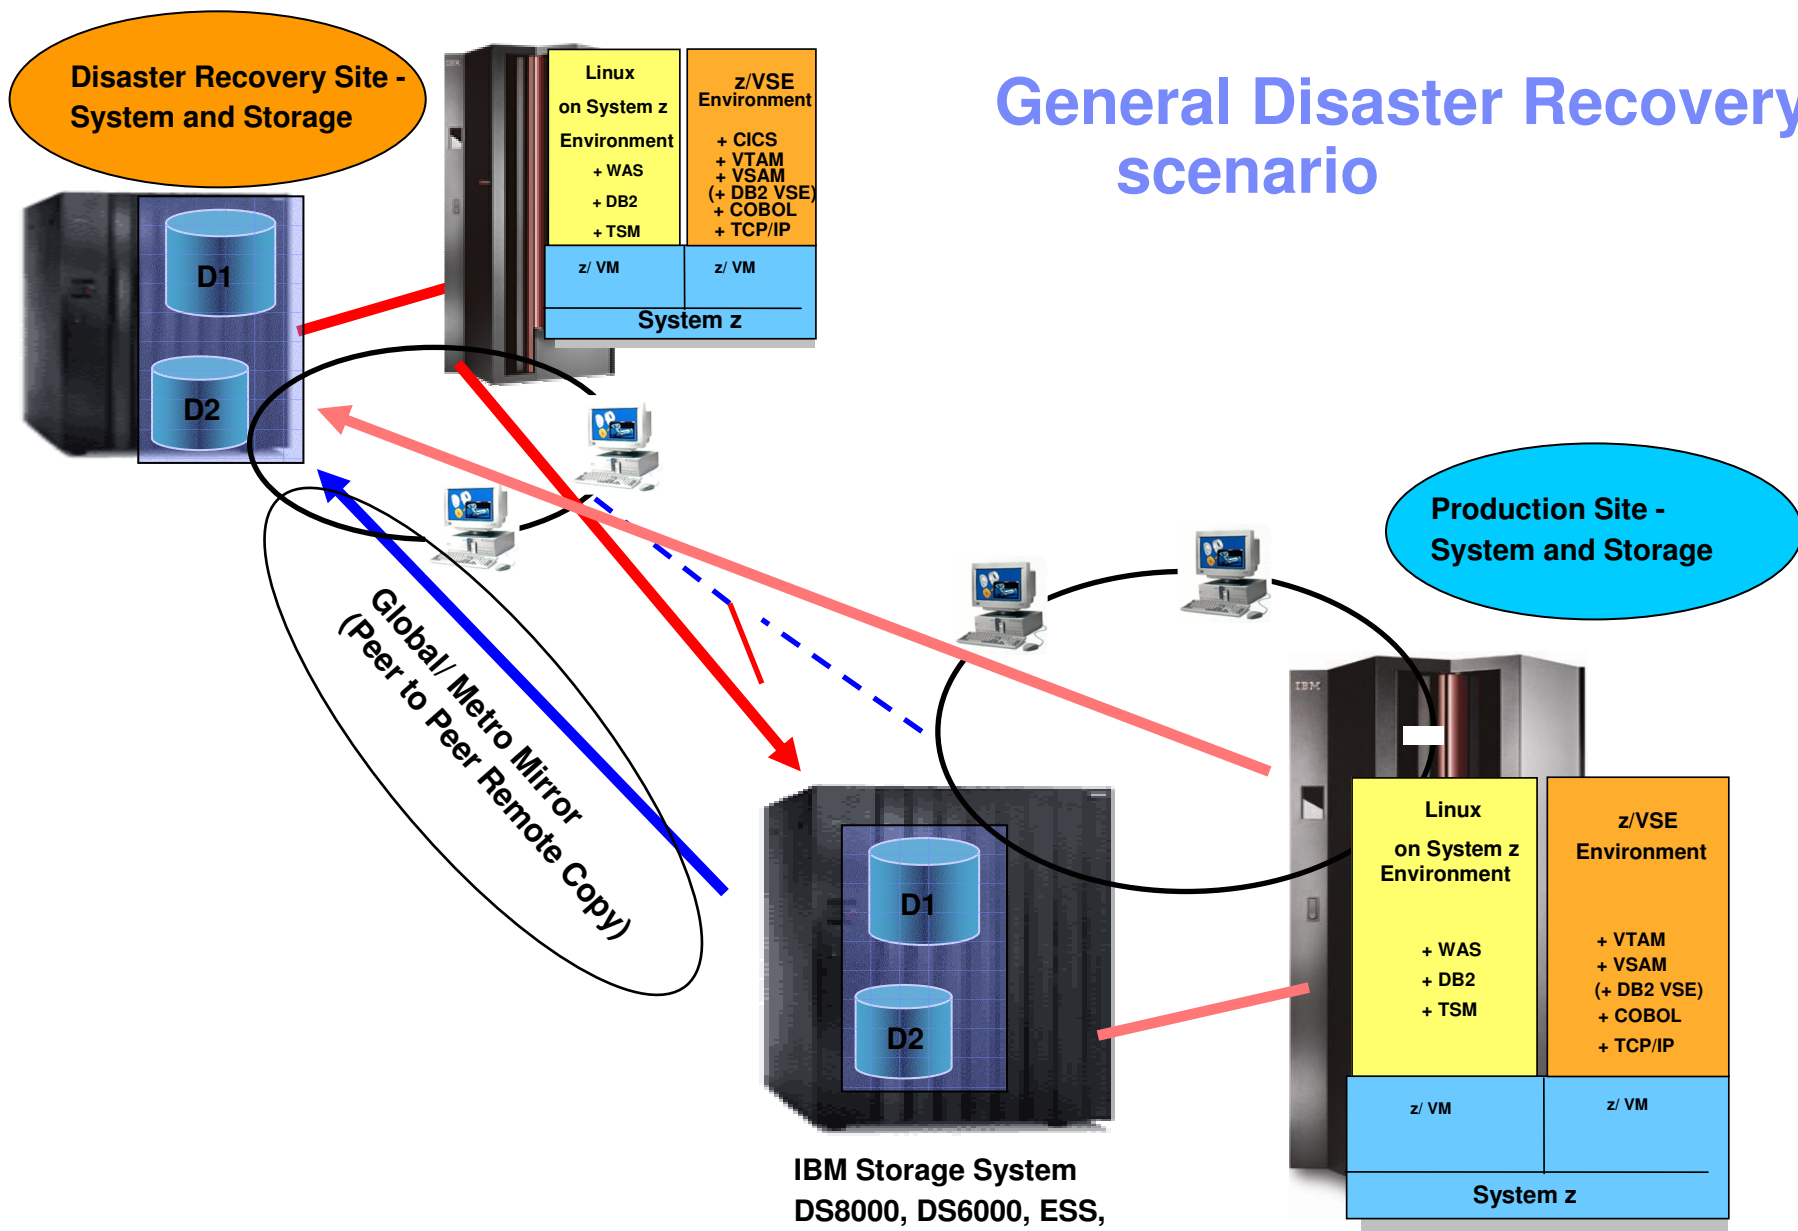
1. xDR Support for z/VSE as passive guest in z/VM



1. z/VSE as passive guest under z/VM and Linux Proxy


- z/VSE is passive for GDPS
- transparent HyperSwap of disks via z/VM triggered by GDPS (vmlogdr)
- communication with GDPS via Linux Proxy only

General Disaster Recovery scenario



Disaster Recovery automation


Disaster Recovery Site -
IBM Storage System



(TPC) Totalstorage Productivity Center
for Replication


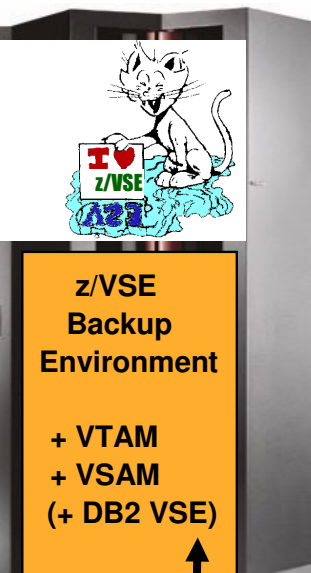
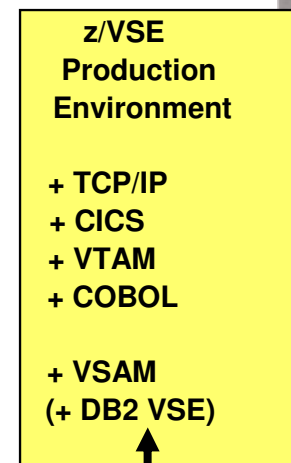
Metro Mirror

IBM Storage System
ESS, DS6000, DS8000



Flashcopy:

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



(offline backup process)

Scenarios for Disaster Recovery with VSE

(1) Concepts of Disaster Recovery (DR)

(2) One active production site and one for DR

(3) Two active sites with production and test

(4) Borrowed Resources for Disaster Recovery

(1) One active production site only and one for DR

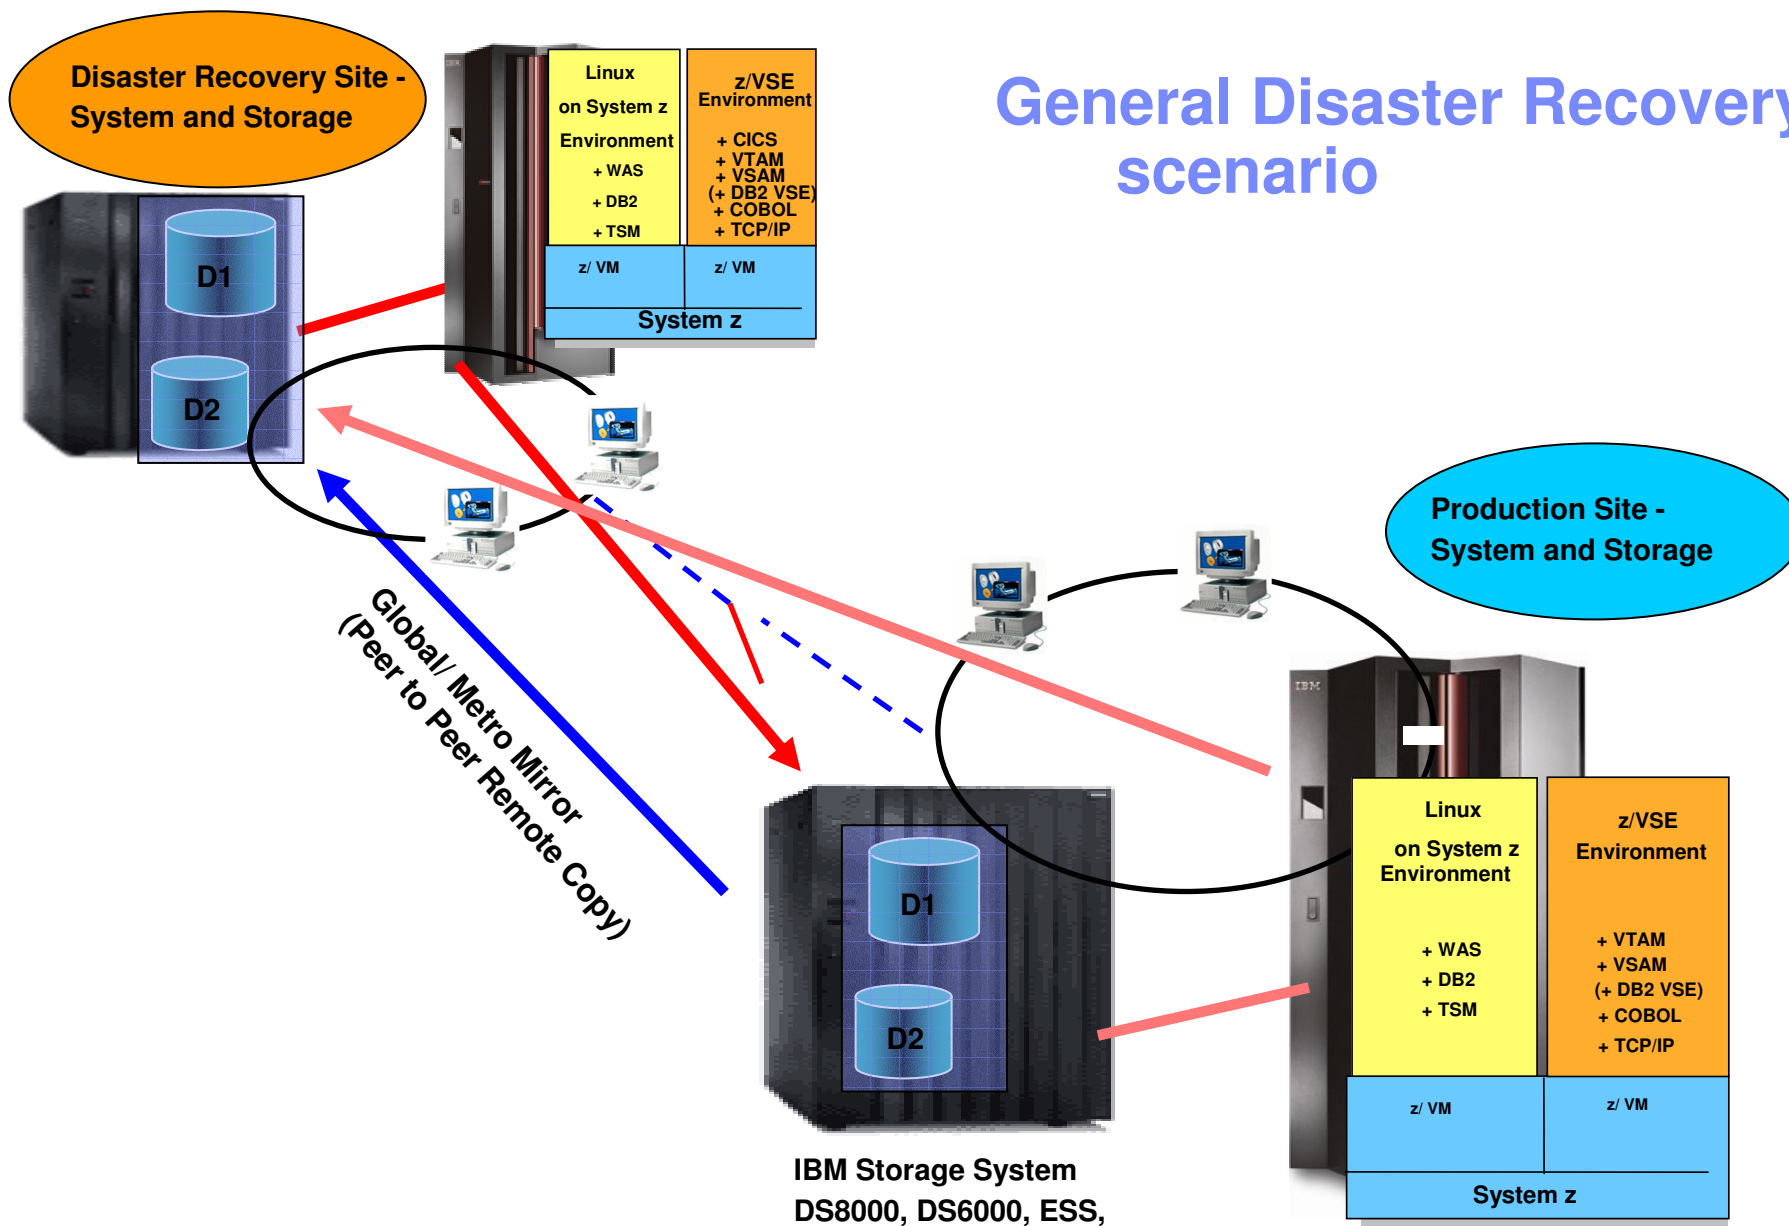
Environment setup for disaster Recovery

- ❖ **DR System**
 - ❖ **An IBM agreement is done to start this machine with the same power as the production site in Case of Recovery**
 - ❖ **An additional agreement can be made for increased capacity, to shorten the startup time of the VSE systems**
 - ❖ **A COLD environment setup - the System is switched off**
 - ❖ **A WARM environment setup - the System is idling**
 - ❖ **Both Systems are able to connect to both Storage subsystems**
 - ❖ **(on the production and DR site)**

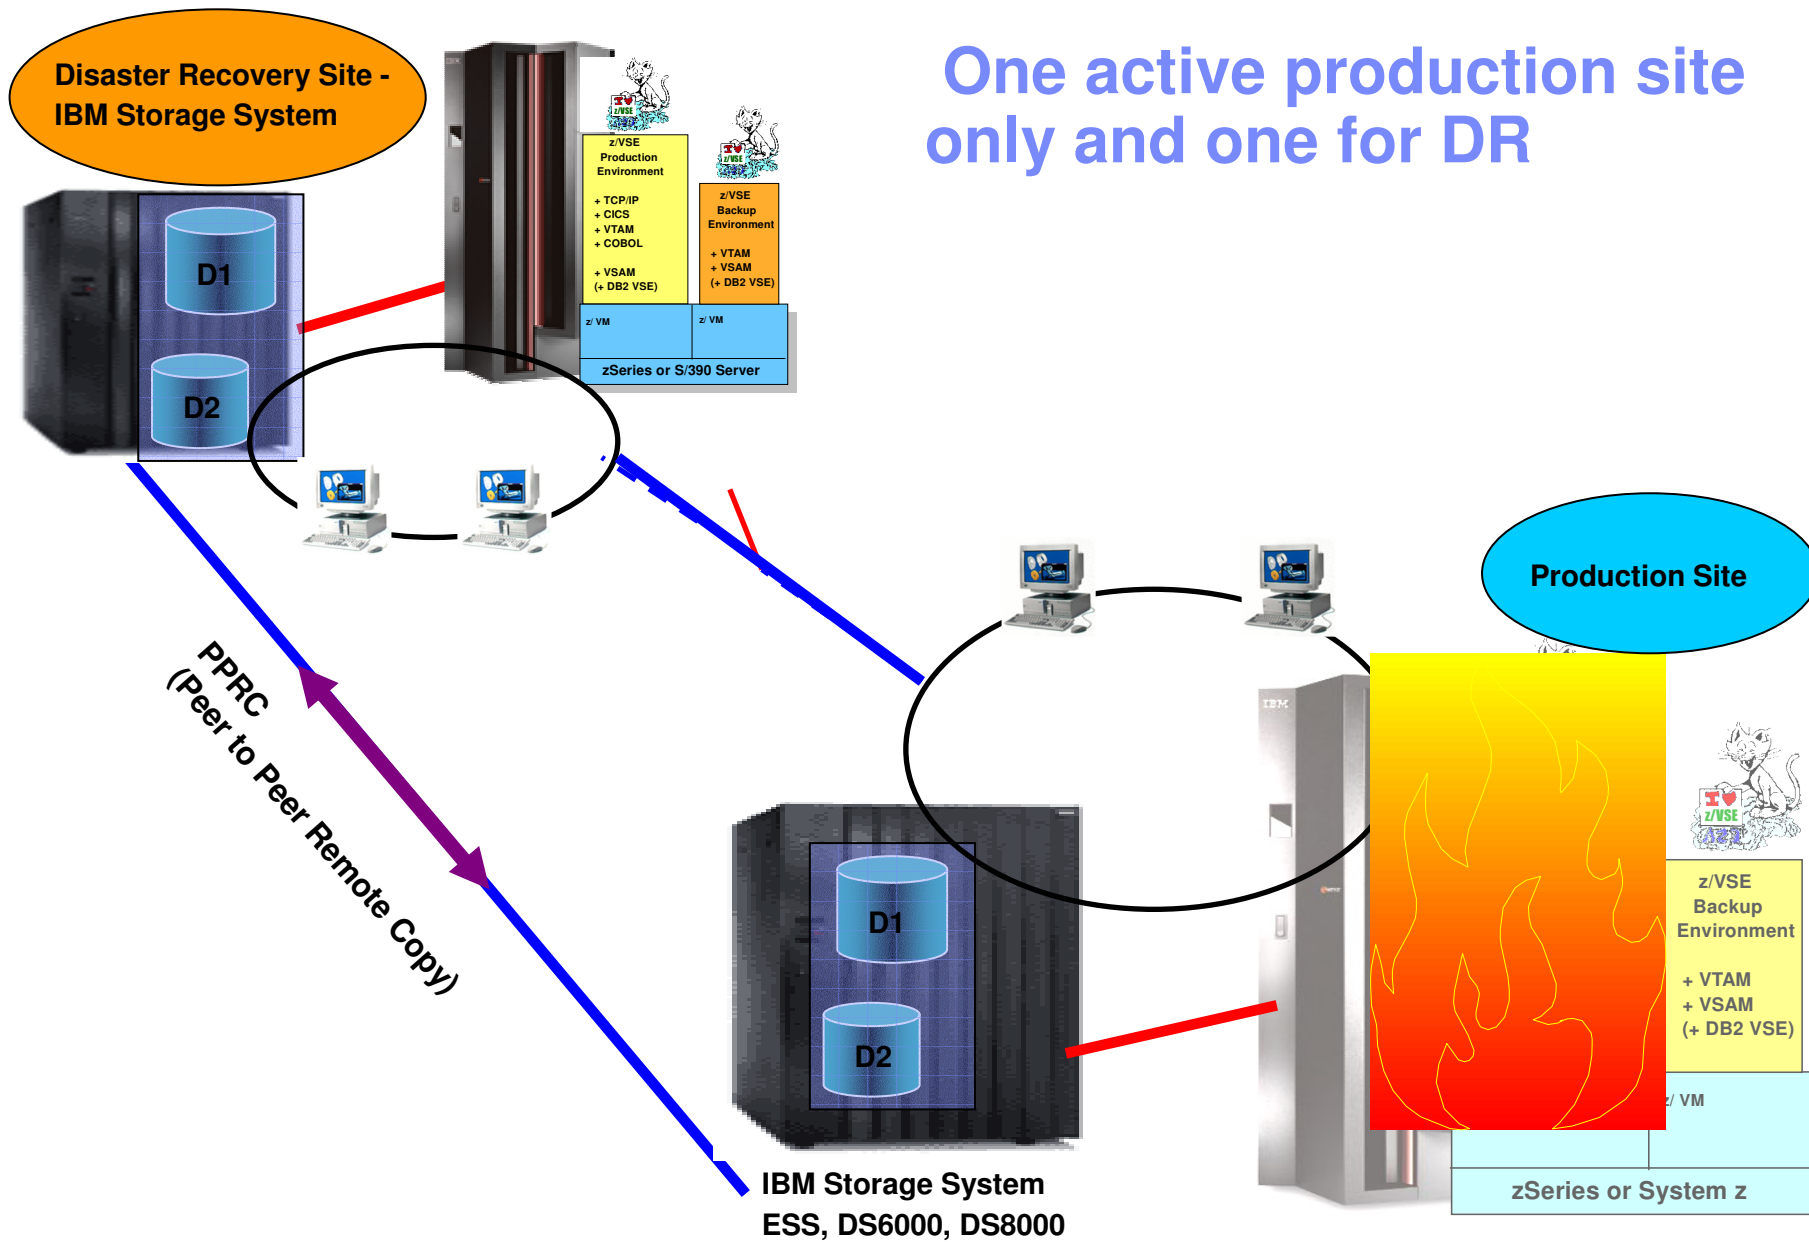
- ❖ **Storage Systems**
 - ❖ **The Production Storage system is cross connected to the one for DR**
 - ❖ **Data is mirrored via PPRC (real time or asynchronous)**
 - ❖ **Enablement to switch the PPRC direction**

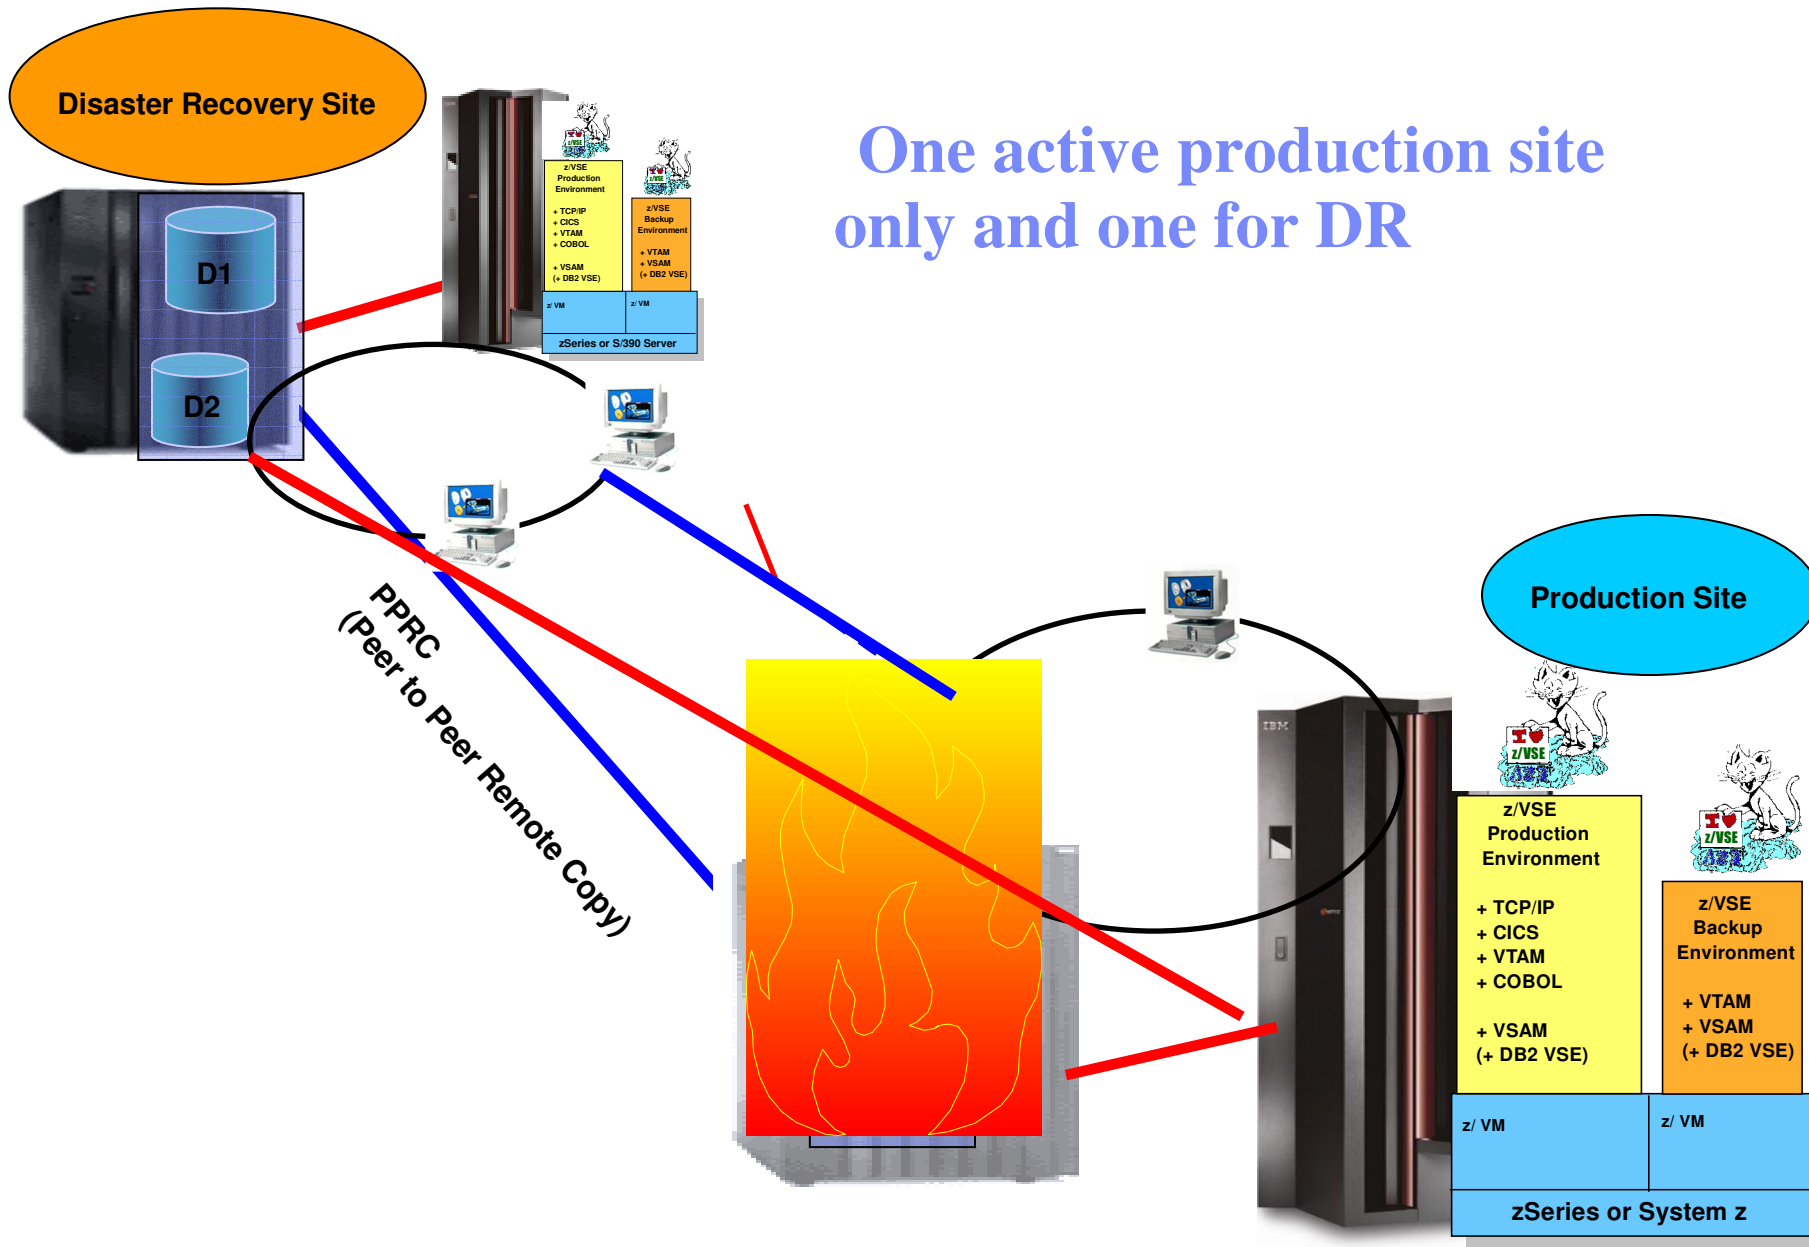
- ❖ **Network**
 - ❖ **Possibility to switch between the productional and DR network**

General Disaster Recovery scenario





One active production site only and one for DR






Disaster Recovery Site -
IBM Storage System



One active production site
only and one for DR

PPRC
(Peer to Peer Remote Copy)



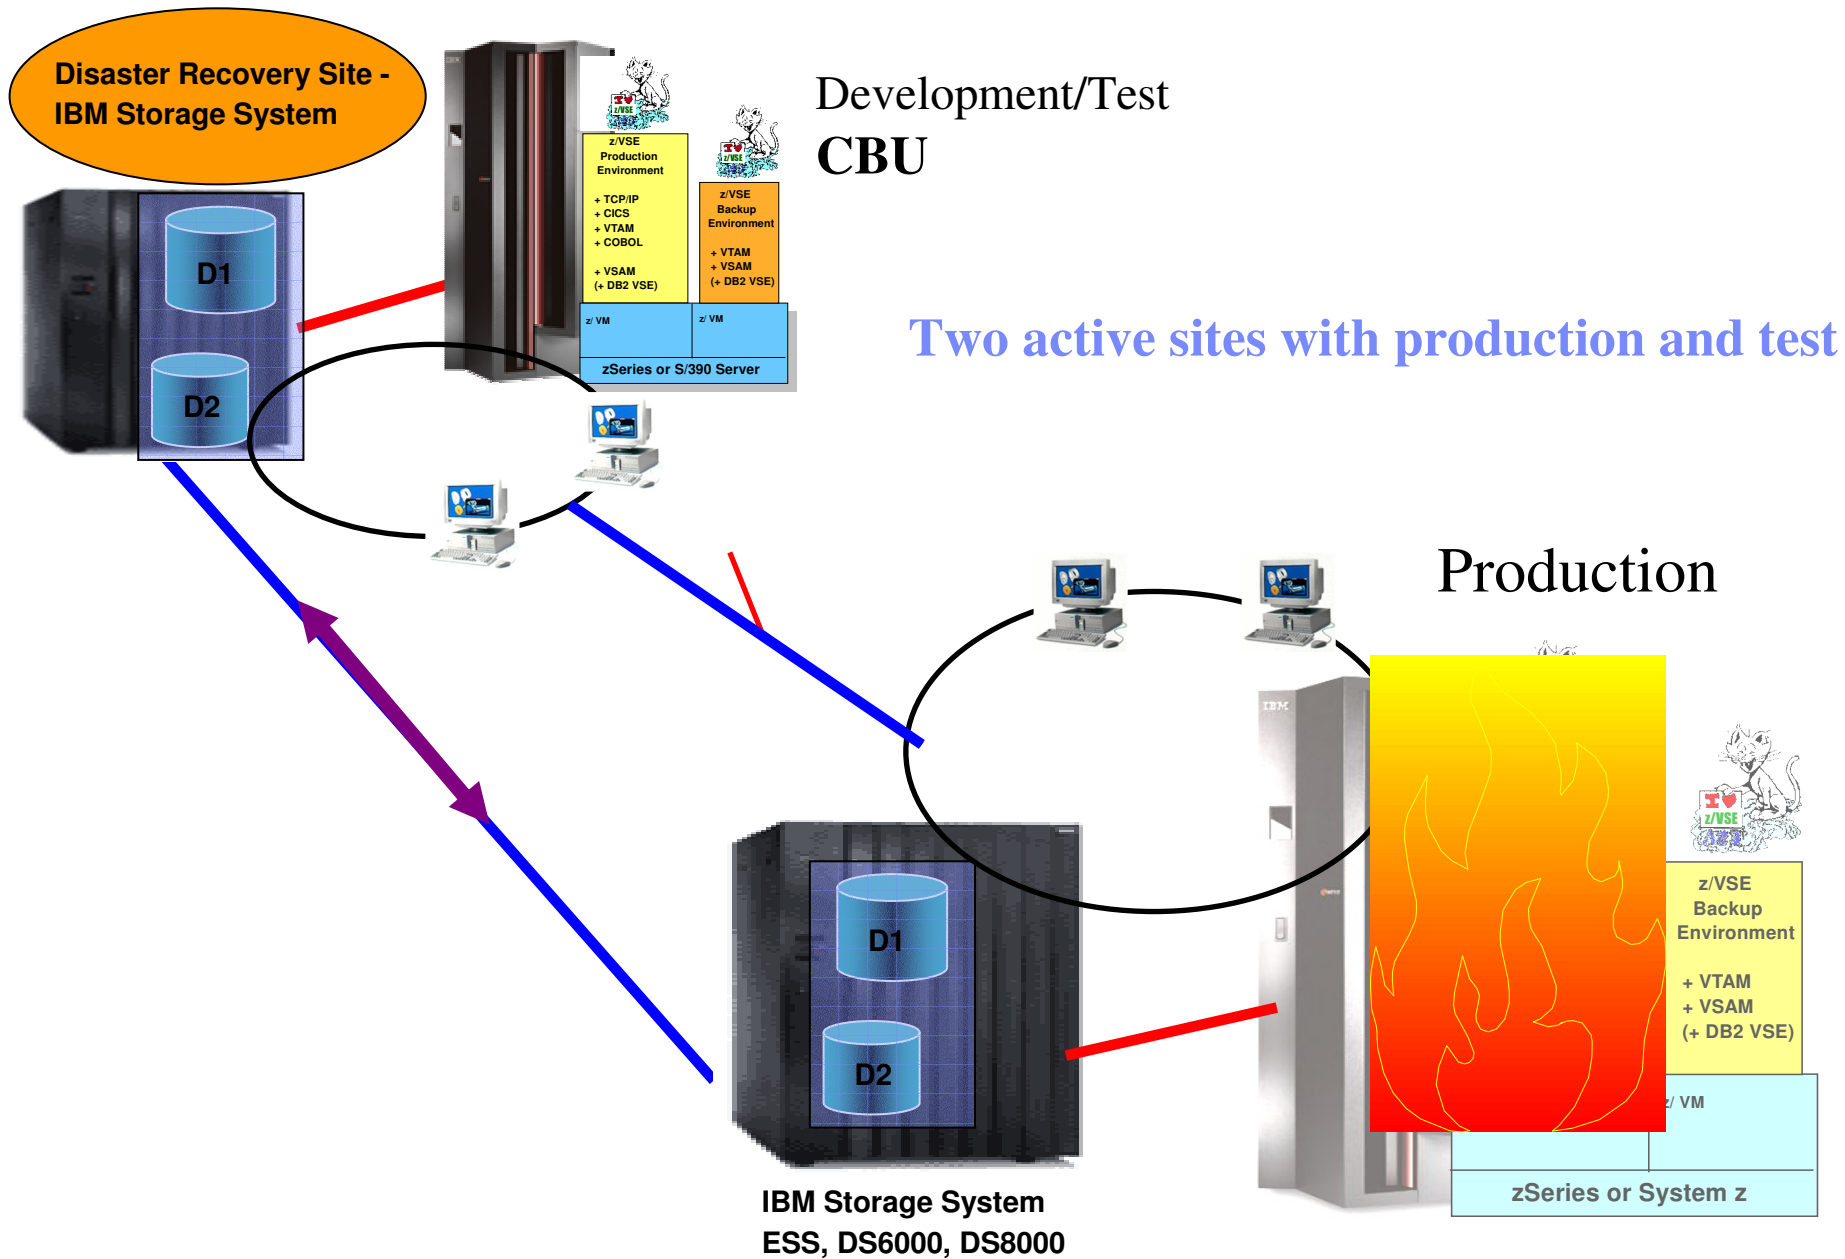
Steps in case of a disaster Recovery

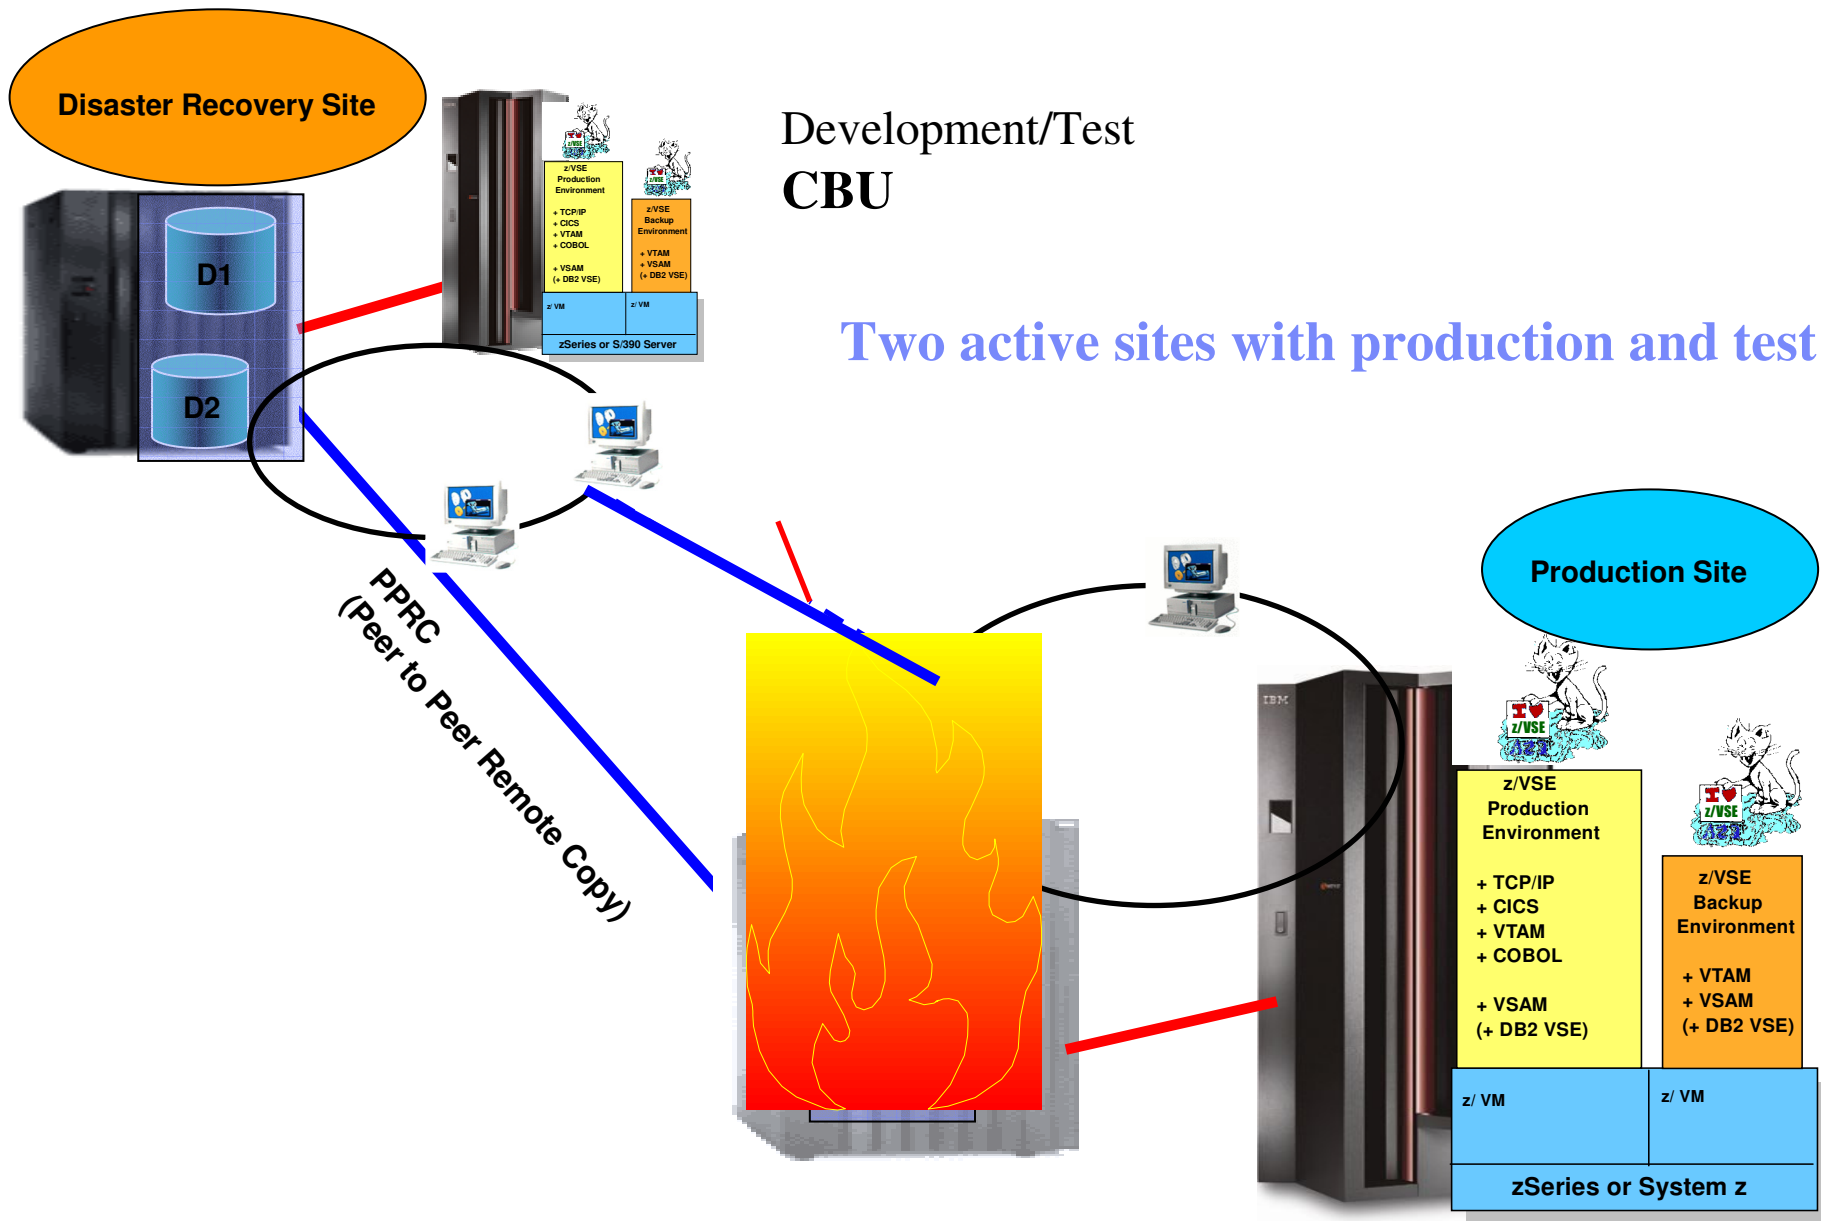
1. **Emergency phones and messaging methods have to be enabled**
2. **Start z/VM on the Recovery Site (on a COLD environment)**
 1. **Start the CBU (Capacity Backup Upgrade) if defined – to accelerate start of VSE systems**
3. **Switch the OSA Adapter Network Connectivity**
4. **Start Online VSE machines (all CICS partitions should start automatically)**
5. **After all productional machines are running – the capacity can be reduced to the normal productional capacity**

Note: These Steps must be tested and trained periodically to have a well functioning process in case of a disaster Recovery failure.

Scenarios for Disaster Recovery with VSE

- (1) Concepts of Disaster Recovery (DR)
- (2) One active production site and one for DR
- (3) Two active sites with production and test
- (4) Borrowed Resources for Disaster Recovery





(1) Two active sites with production and test Environment setup for disaster Recovery

- ❖ **DR System**
 - ❖ **An IBM agreement is done to increase the machine for DR capacity with the power of the production site, using CBU (Capacity Backup Upgrade)**
 - ❖ **In a WARM environment setup - the System is idling**
 - ❖ **In a HOT Environment setup – the system is very fast ready to take over the production workload**
 - ❖ **Both Systems are able to connect to both Storage subsystems**
 - ❖ **(on the production and DR site)**
- ❖ **Storage Systems**
 - ❖ **The Production Storage system is connected to the one for DR**
 - ❖ **The DR Storage system is connected to the production Storage**
 - ❖ **Data is mirrored via PPRC (real time or asynchronous)**
 - ❖ **Enablement to switch the PPRC direction**
- ❖ **Network**
 - ❖ **Possibility to switch between the productional and DR network**

Steps in case of a disaster Recovery

1. **Emergency phones and messaging methods have to be enabled**
2. **Start the CBU (Capacity Backup Upgrade)**
3. **Switch the OSA Adapter Network Connectivity**
4. **Start the Online VSE machines if not already started (all CICS partitions should start automatically)**
5. **After all productional machines are running – the capacity can be reduced to the normal productional capacity**

Note: These Steps must be tested and trained periodically to have a well functioning process in case of a disaster Recovery failure.

z/VSE Customer example – distributor Germany

- mainframe on D/R site is Power-On-Reset
- Activities in case of disaster:
 - Switch attachments for channel attached printer
 - IPL VM
 - CBU on HMC 5-7 min (conform IBM doc about 30 min)
 - activate 2nd CPU in VM
 - VSE IPL – start all business applications

For starting all CICS environments and applications they need 25 min

Scenarios for Disaster Recovery with VSE

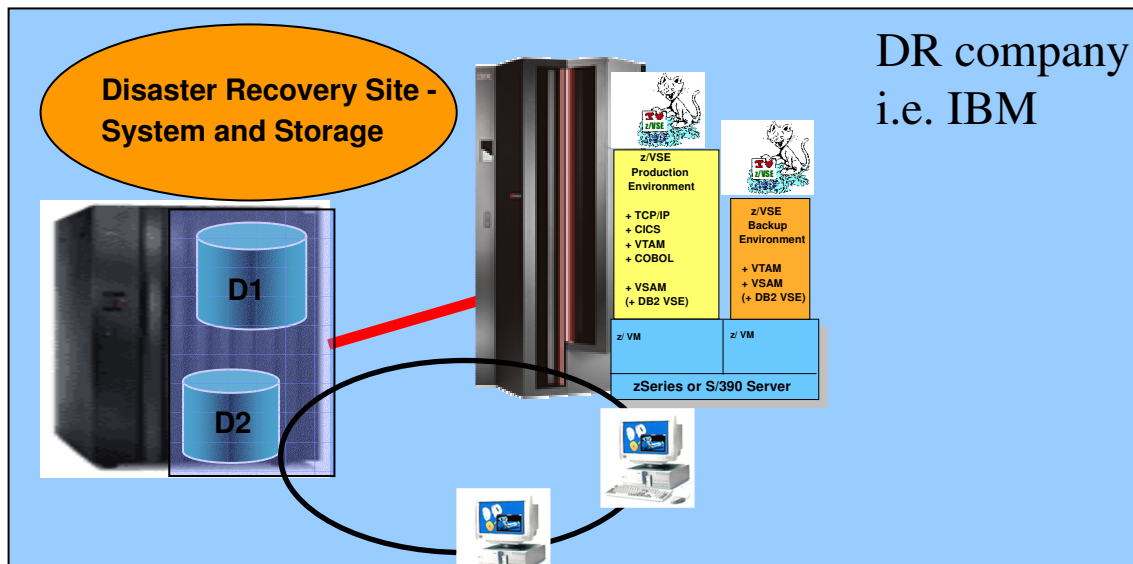
- (1) Concepts of Disaster Recovery (DR)
- (2) One active production site and one for DR
- (3) Two active sites with production and test
- (4) Borrowed Resources for Disaster Recovery

Off-site Disaster Recovery

A Disaster Recovery Site can be made offsite on other customers with IBM equipment.

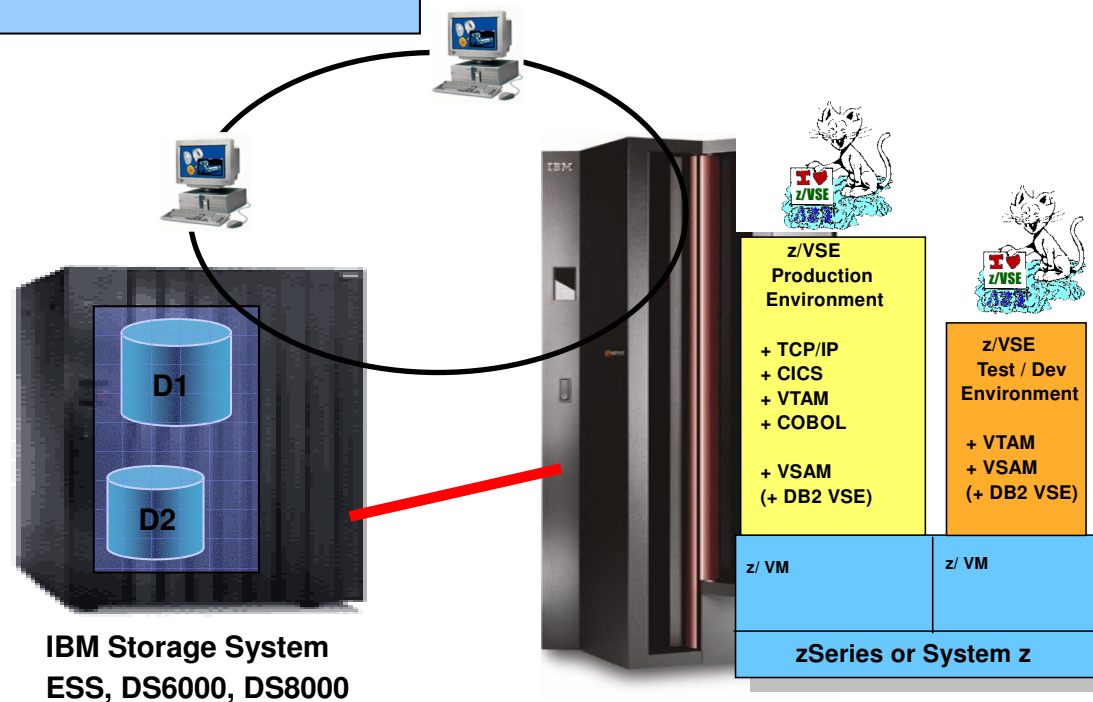
Necessary Agreements are required:

- An contract with HW details
- The DR procedure must be well defined and described
- Data for the DR case are provided periodically to the DR Center
- Training is done periodically and the DR procedure is verified



OFF-Site Disaster Recovery

Depending on failure, not all actual running processes can be redone entirely.



Scenarios for Disaster Recovery with VSE

(1) Disaster Recovery (DR)

(2) Backup Concepts

Business Continuity, Disaster Recovery, and TSM

- Levels of interdependence
 - Business requirements for availability
 - continuity of IT operations
 - Business Continuity Planning (BCP)
 - Disaster Recovery Planning (DRP)
- Storage management processes and infrastructure impact
 - recovery times for operations
 - overall continuity of operations.
- Backup procedures important factors for:
 - availability of data
 - restore procedures and the documented and planned execution
- Bare Metal Backup

Disaster Recovery and TSM

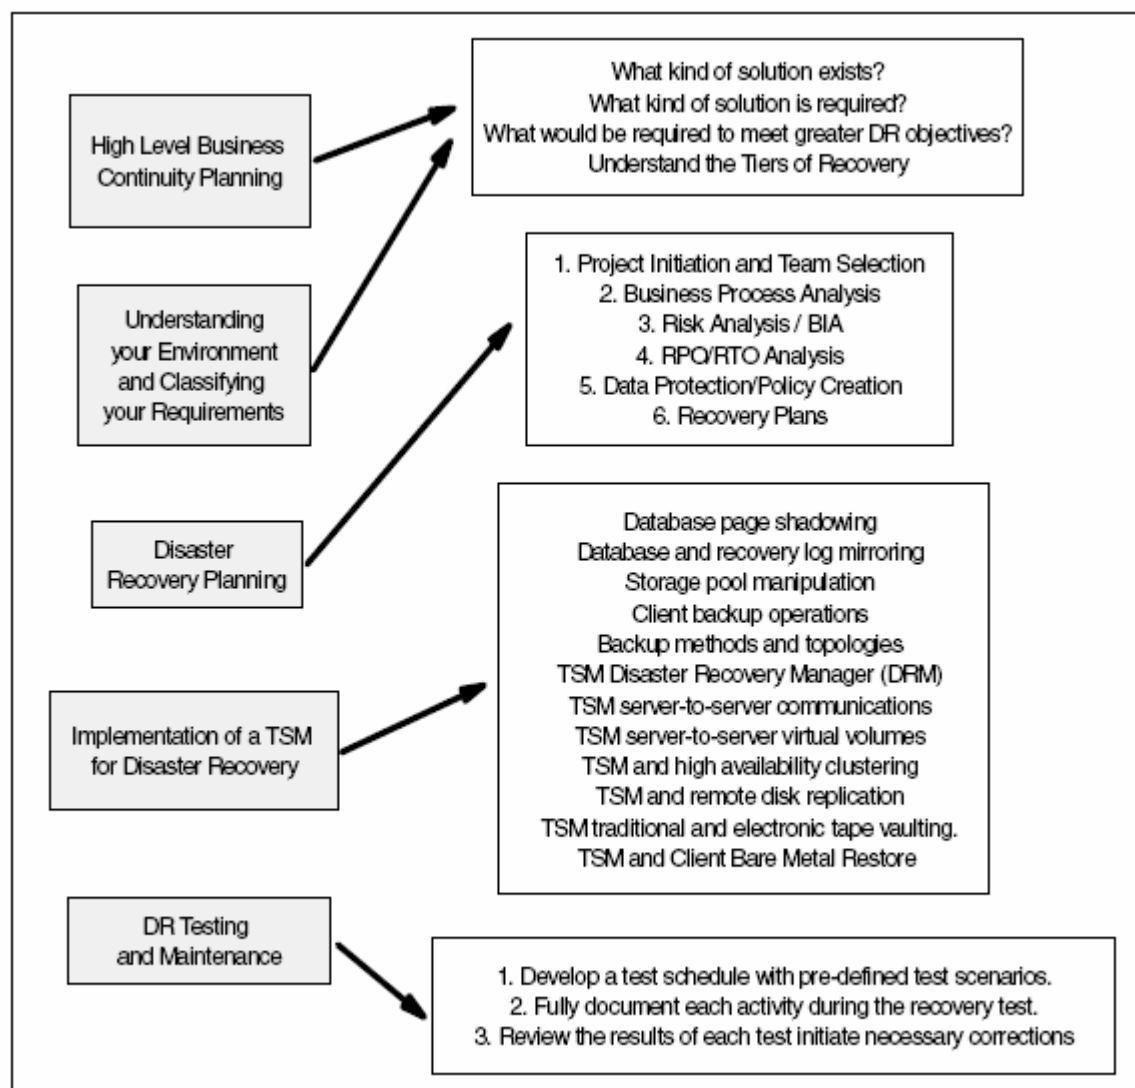
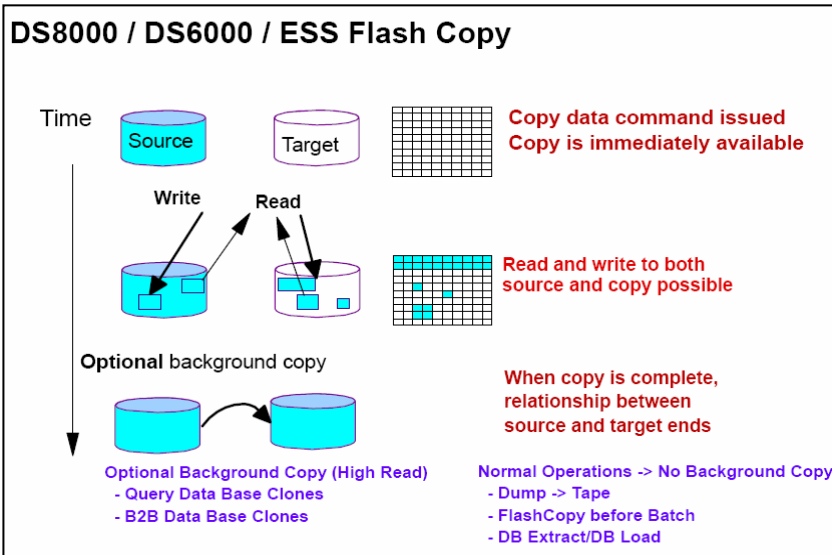
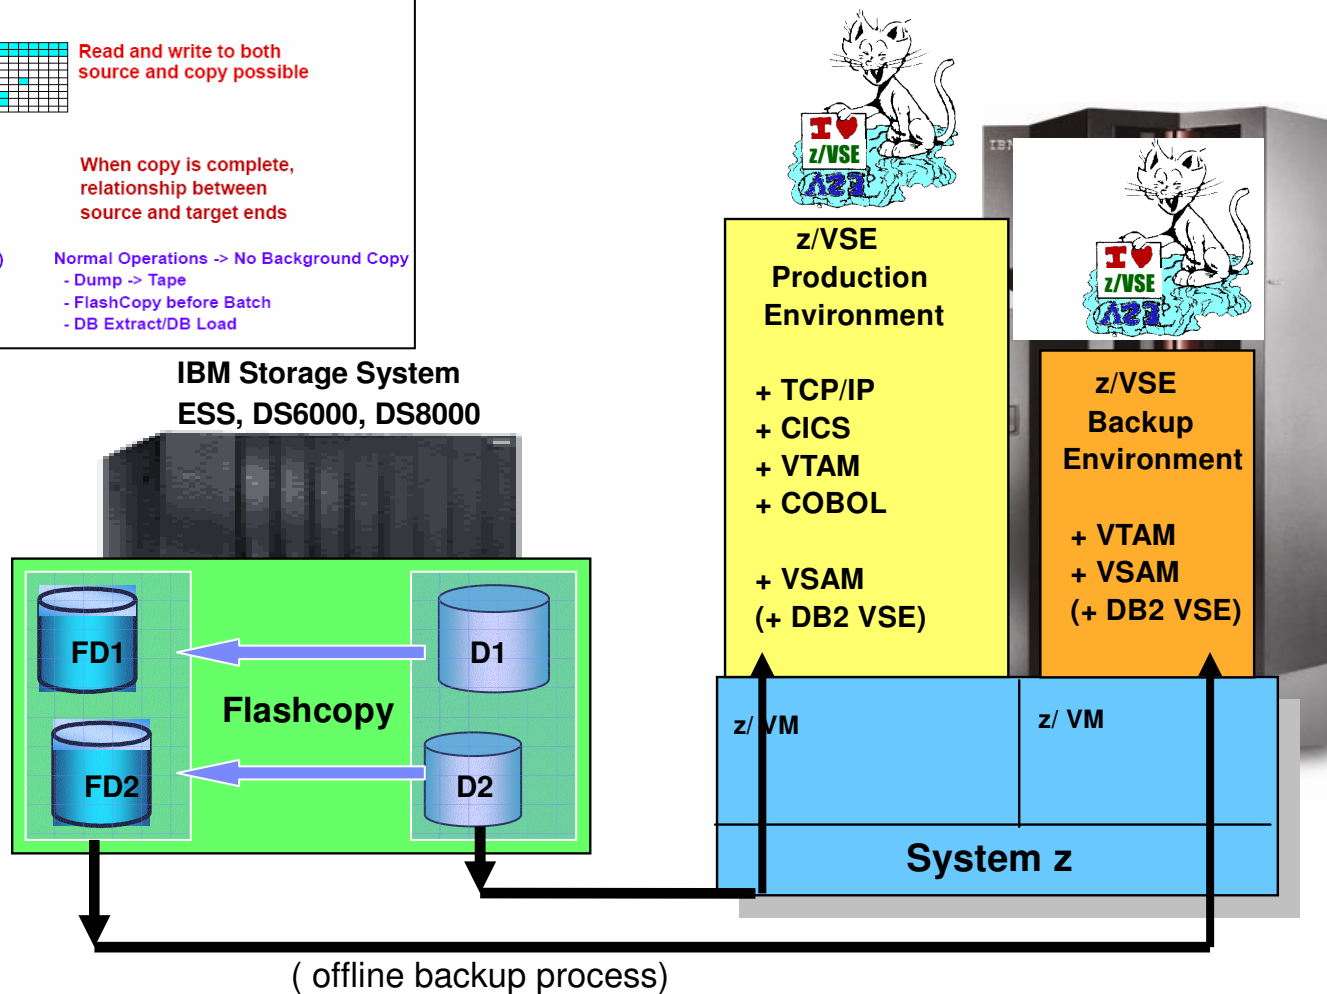


Figure 14-12 Summary of TSM and Disaster Recovery considerations



IT Environment for almost 24x7 Availability




IT Environment Needs for 24x7 Availability

✘ inhibitors of **online** processing time

☞ backup-window



☞ batch-window

Typical processing time-line:



BACKUP process (after Flashcopy) running in parallel to batch from Flashcopied volumes

Tivoli Storage Manager - Architecture



DR Plan Generation for TSM Server

- Generate the recovery plan using the PREPARE command in TSM
- The plan is stored in a time-stamped file in the local directory as defined in **SET DRMPLANPREFIX**

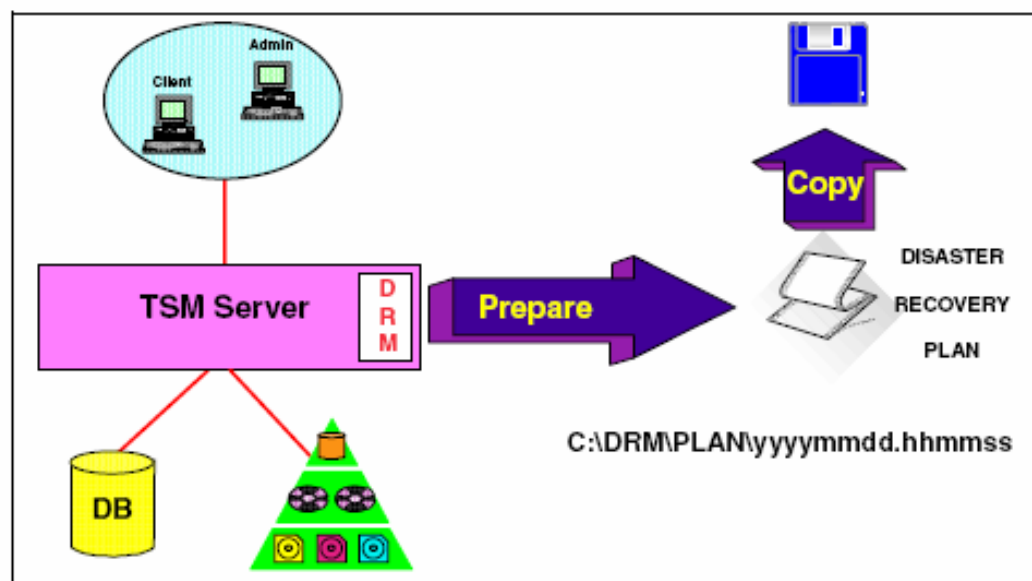


Figure 8-12 Disaster Recovery Plan generation

TSM with warm DR site

- DRM output sent to the remote site using a secure network using FTP
- The TSM database backups can be manually vaulted along with the copy storage pool data to the warm site environment for disaster

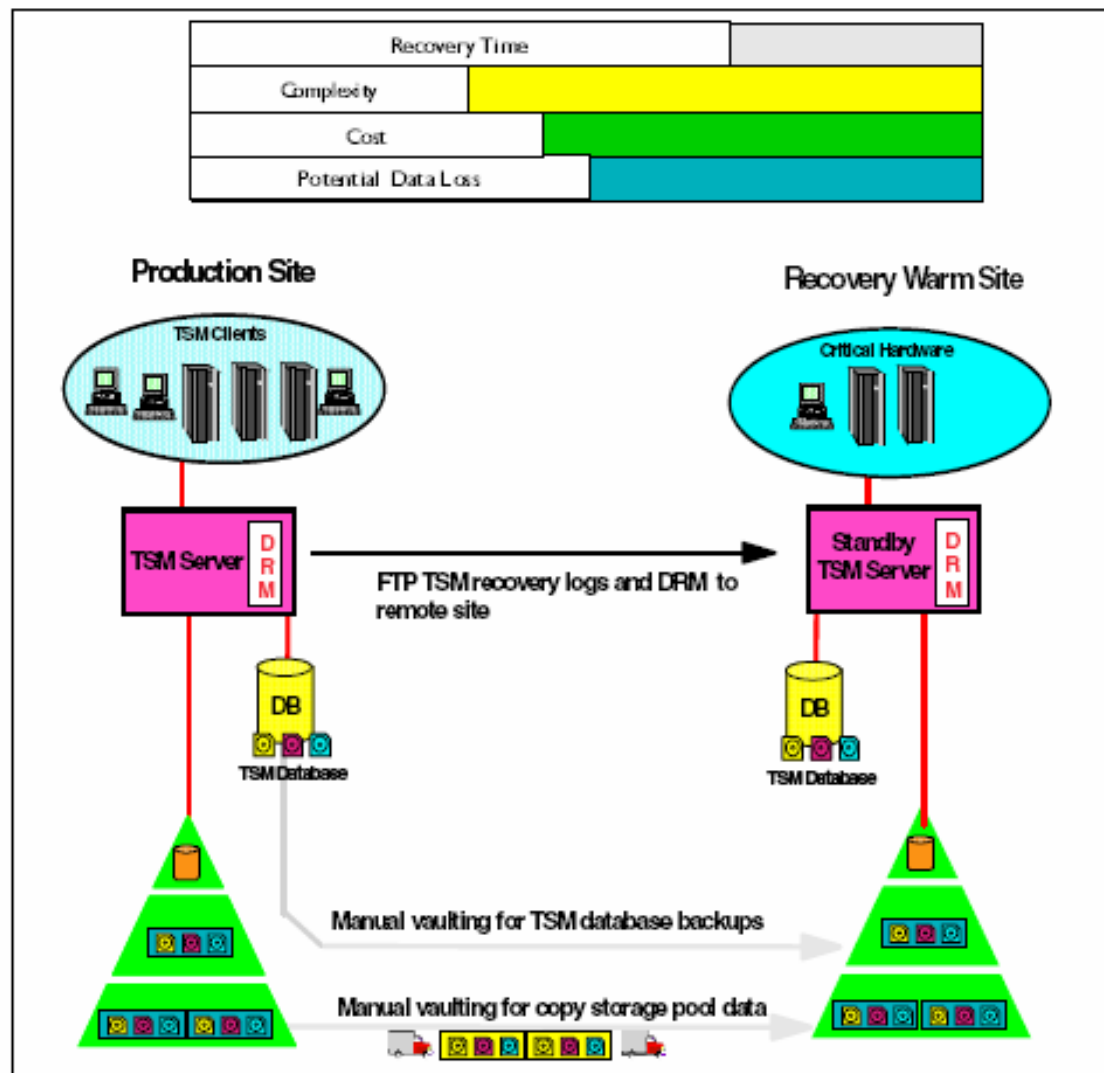
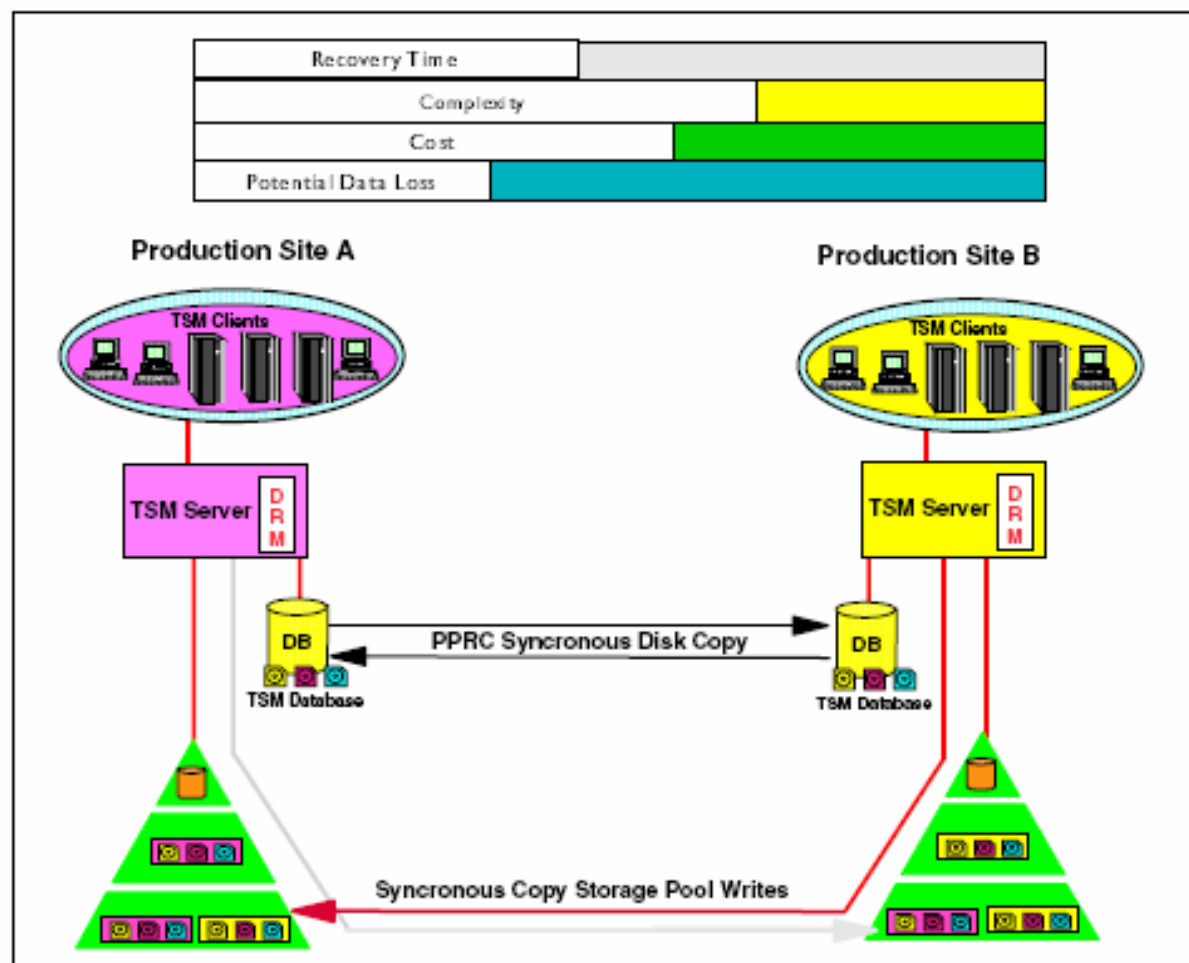


Figure 14-6 Electronic vaulting of TSM DRM, manual DB, and copy pool vaulting

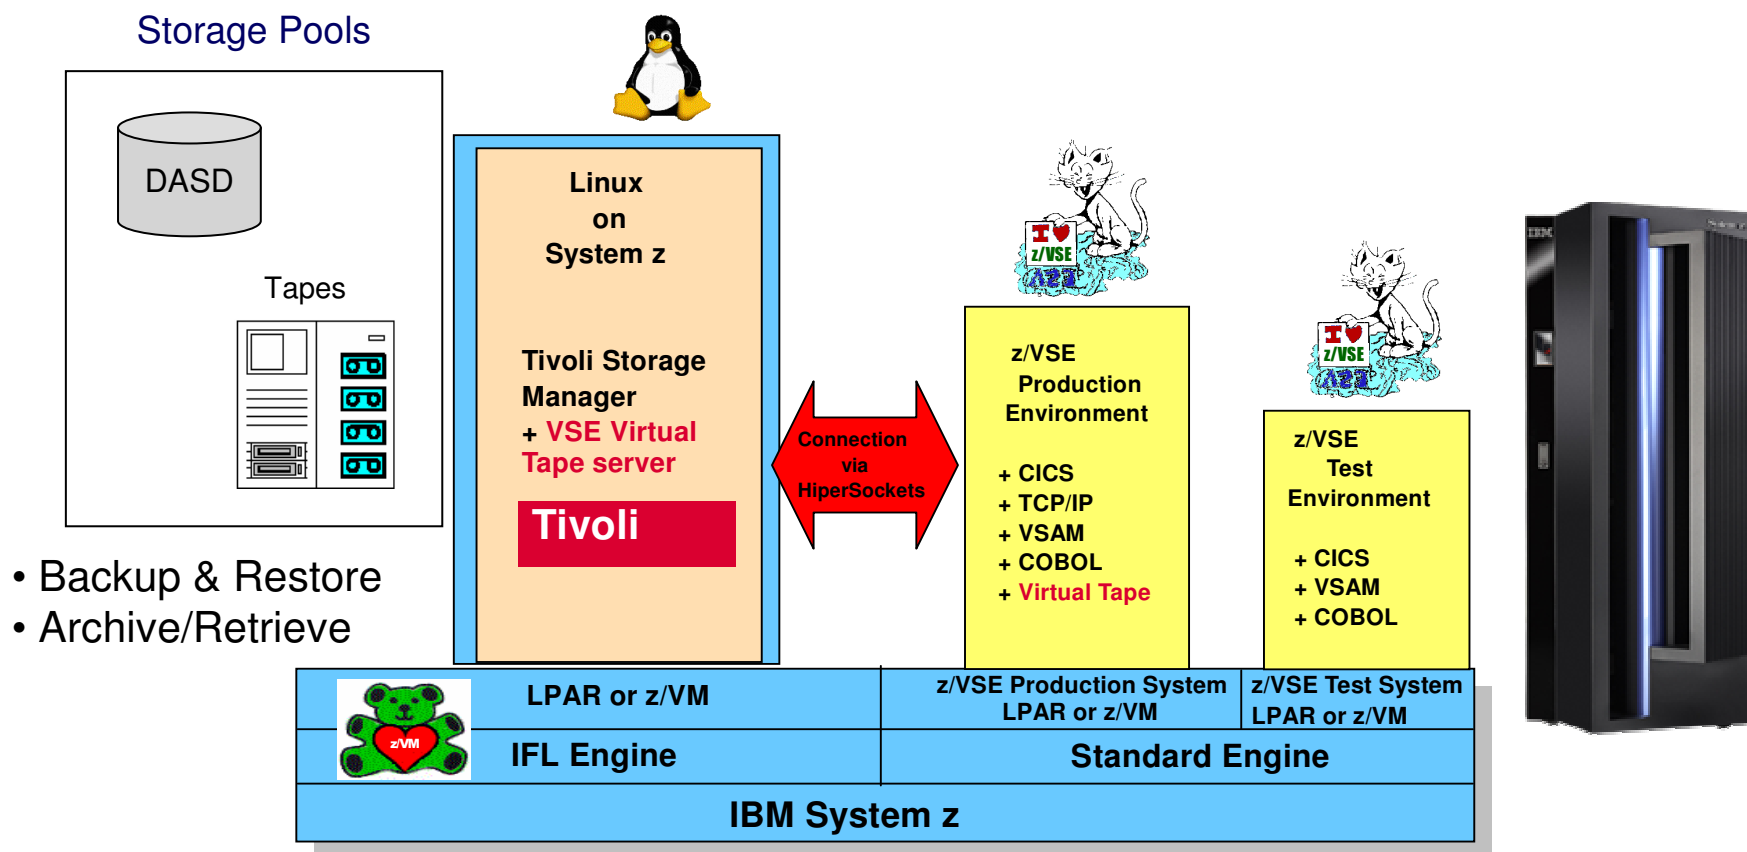
DR mirrored site with TSM Backup setup

- each TSM environment functions independently,
- it vaults its data to the alternate TSM site.
- In the case of a disaster, the existing TSM environment is used to recover the lost TSM environment and associated



Dual production sites, electronic TSM DB and copy pool vaulting


z/VSE – Backup integration with Tivoli Storage Manager (TSM)



z/VSE Backups with Tivoli Storage Manager

- New with z/VSE 4.1 and newer
- Uses the TSM Command-Line interface (DSMC)
- Based on the VSE VTAPE Functionality
 - entire tape images will be stored via TSM
 - VTAPE OPEN/CLOSE Exit (are Actions)
 - On OPEN the tape image will be restored per TSM to the TSM server and can be accessed by VSE
 - On CLOSE the tape image will be saved per TSM to the corresponding storage pool (dasd or tape)

Tivoli Storage Managers – Connection to z/VSE



Tivoli Storage Managers – used with z/VSE Backup

Backup of VSAM Clusters with TSM

```

* $$ JOB JNM=VSAMBKUP,DISP=L,CLASS=0
// JOB VSAMBKUP
// LIBDEF PHASE,SEARCH=IJSYSRS.SYSLIB
* THIS JOB BACKS UP VSAM DATASETS
// DLBL IJSYSUC, 'VSESP.USER.CATALOG',,VSAM
*
* THIS FUNCTION USES A VTAPE FOR OUTPUT
VTAPE START,UNIT=181,LOC=9.152.216.105,FILE='TSM:VSAM.AWS (BACKUP)',SCRATCH
// ASSGN SYS005,181
// EXEC IDCAMS,SIZE=AUTO
        BACKUP ( -
                VSAM.CONN.SAMPLE.DATA -
                ... -
                ) -
        REW -
        NOCOMPACT -
        BUFFERS(3)
/*
// ASSGN SYS005,UA
VTAPE STOP,UNIT=181
/&
* $$ EOJ


```

Syntax:

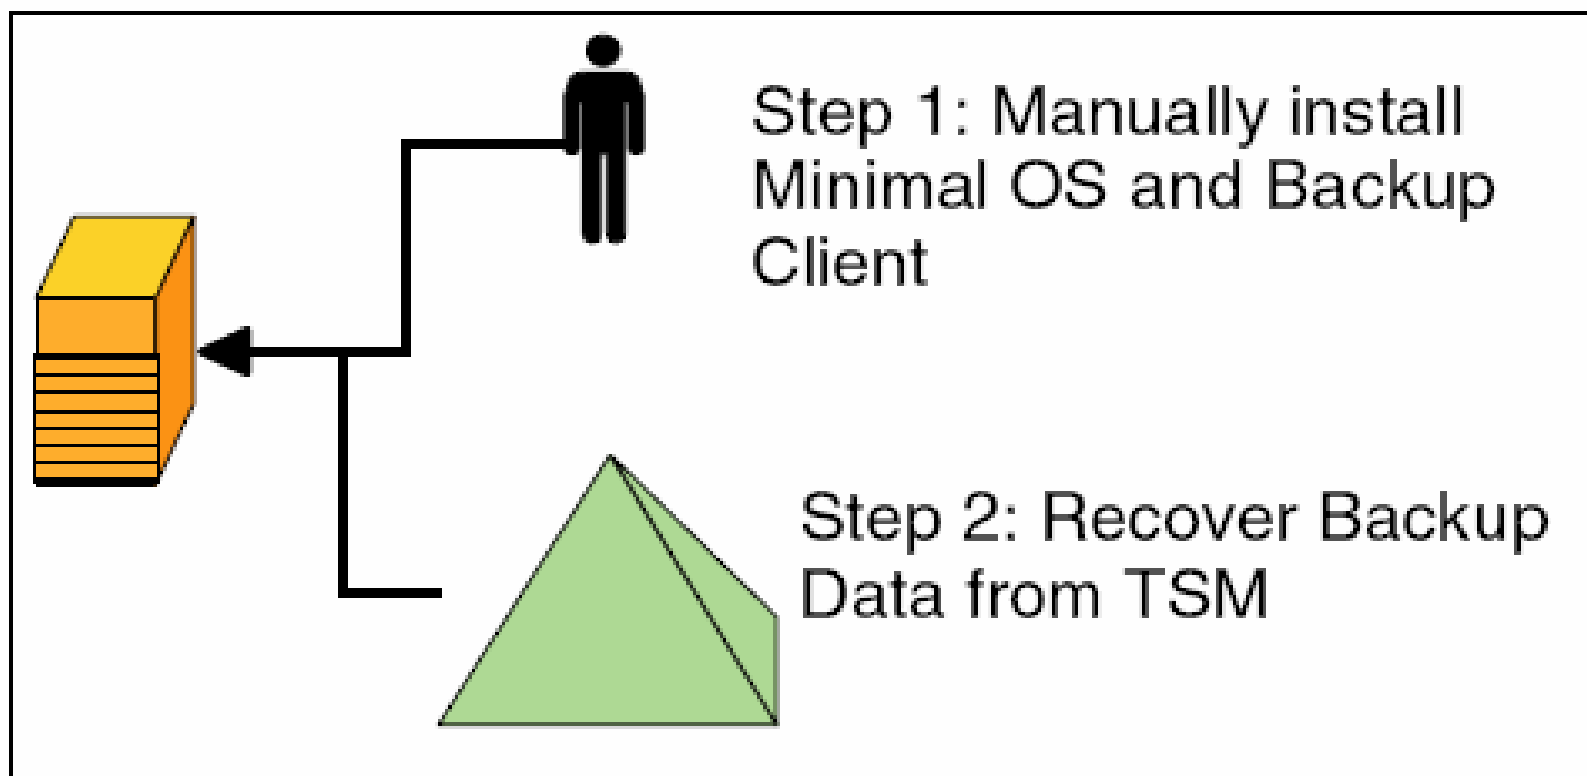
TSM:<name>(<mode>,<optionset>,
 <fromdate>,<fromtime>)

mode	- BACKUP or ARCHIVE
optionset	- Name of configuration
fromdate	- date (for Restore)
fromtime	- time (for Restore)

Tivoli Storage Managers – Connection to z/VSE

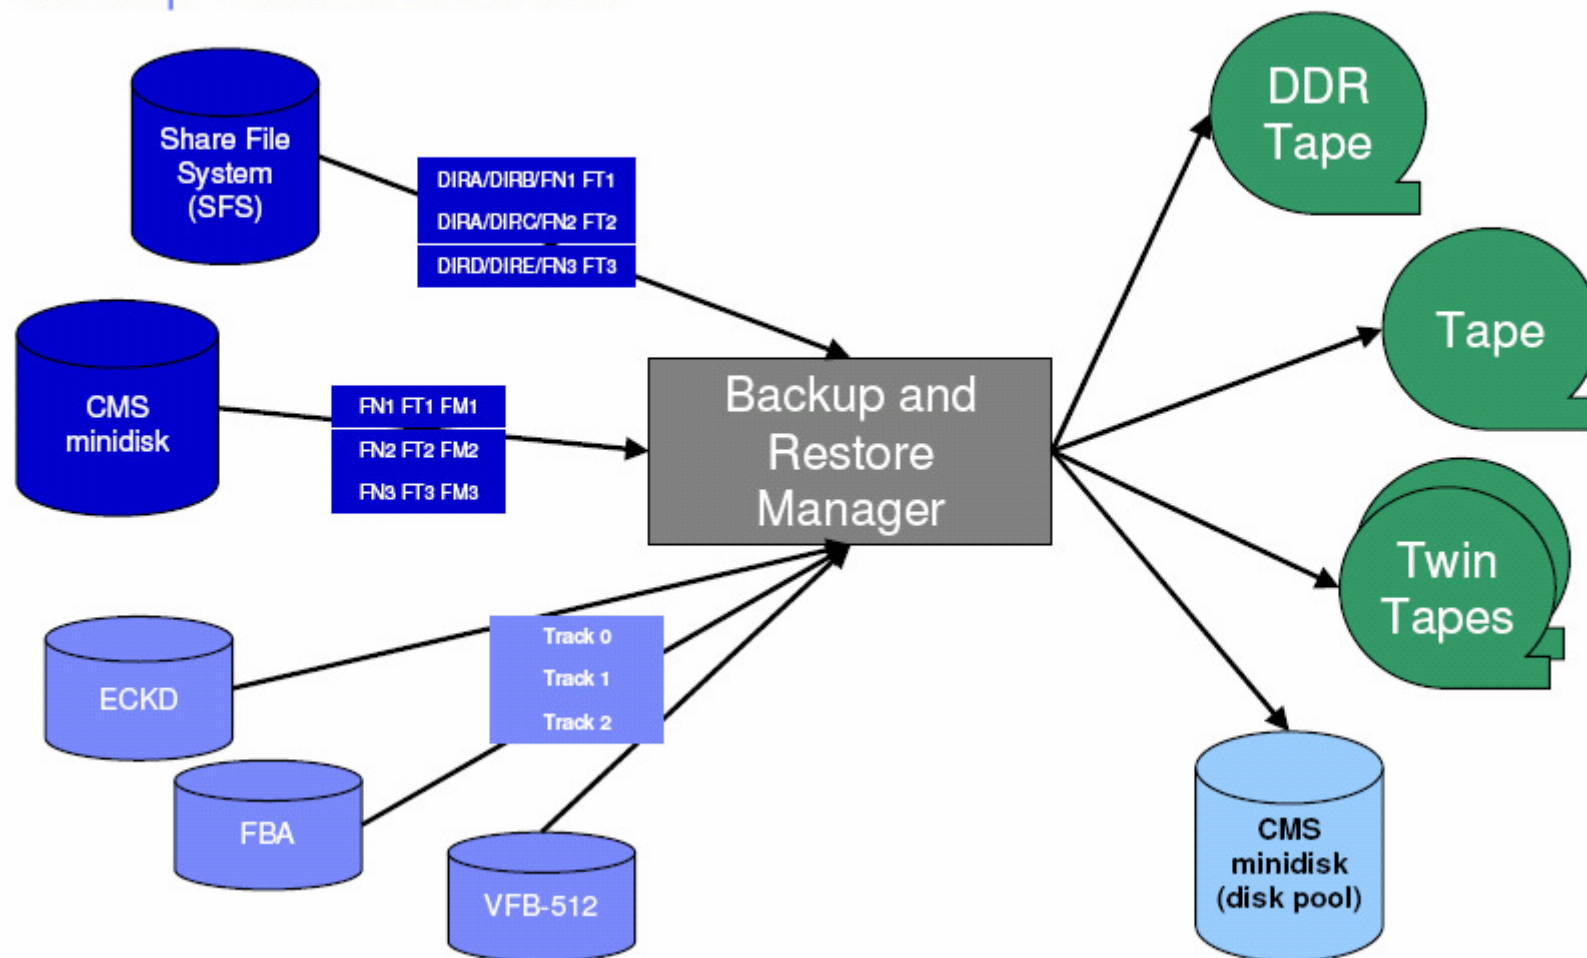


Bare Metal Recovery (BMR)



z/VM Backup / Restore Manager - BACKUP

Backup Data and Media

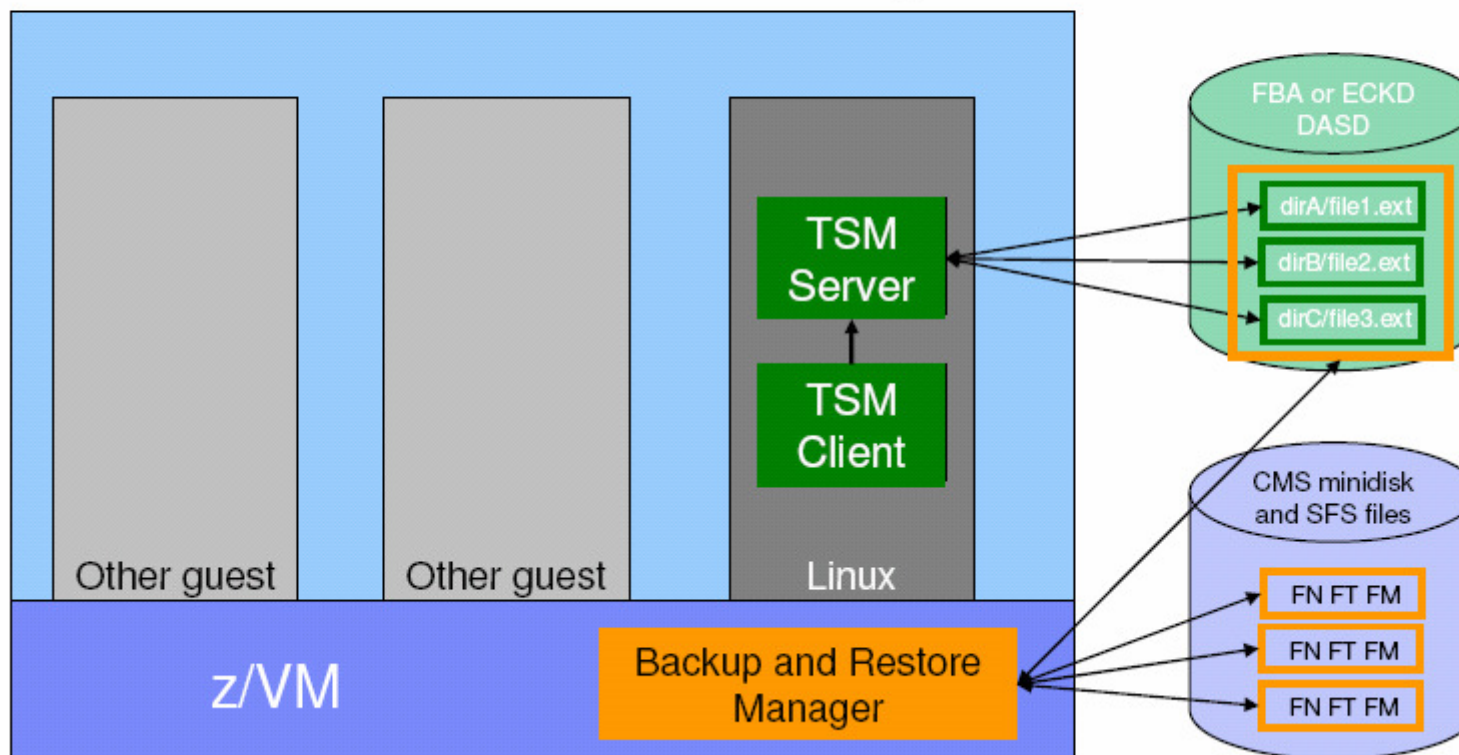


z/VM Backup / Restore Manager

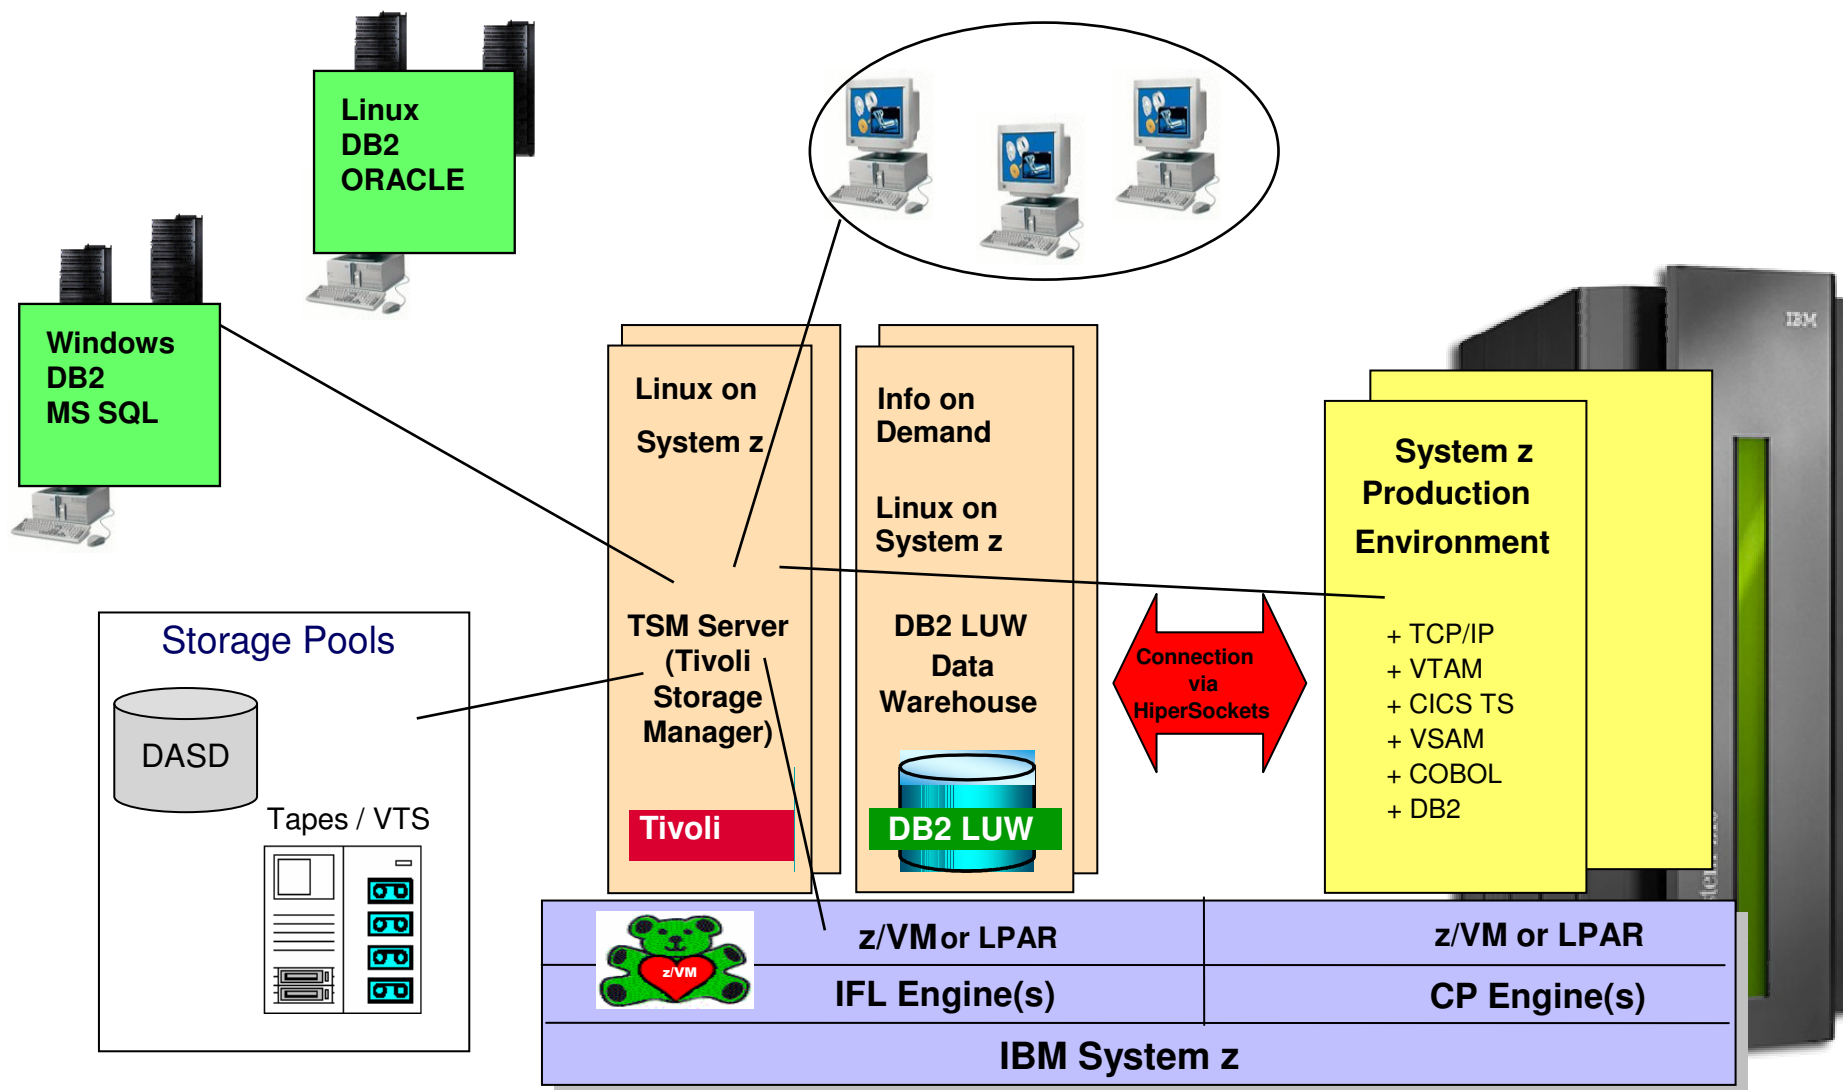
Backup and Restore Manager and Linux Guests

Using Backup and Restore Manager with Tivoli Storage Manager

Choose the solution that meets your needs



Implement TSM on Linux on System z as central Backup Hub



Strategies to achieve tiers of DR using TSM

- IBM Tivoli Storage Manager (TSM) can be used to help provide all the tiers of DR
- Use of Disaster Recovery Manager (DRM) which can automate the TSM server recovery process and manage offsite volumes.
- Vaulting of TSM database, recovery log, volume history information, device configuration information, DRP file (if using DRM) and copy pools for storage at an offsite location.
- Use of TSM server-to-server communications to enable enterprise configuration (multiple TSM servers), enterprise event logging and monitoring, and command routing.
- TSM servers installed at multiple locations, optionally setup as peer to peer servers (that is, each server able to recover at the alternate site).
- Use of TSM virtual volumes over TCP/IP connection to allow storage of TSM entities (TSM database backups, recovery log backups, and primary and copy storage pools, DRM plan files) on remote target servers.
- Use of high bandwidth connections and data replication technology (such as IBM PPRC, EMC SRDF) to support asynchronous/synchronous data replication of TSM databases backups, recovery log backups, TSM database and recovery log mirrors, and storage pools.
- Use of remote electronic tape vaulting of TSM database and recovery log backups, primary or copy storage pools. Extended distances can be achieved by using distance technologies, for example, extended SAN, DWDM, IP/WAN channel extenders.

IBM Tivoli Storage Manager at a glance

- **Tivoli Storage Manager** — automates data backup and restore (B/R) functions, on a broad range of platforms and storage devices
- **Tivoli Storage Manager Extended Edition** — expands B/R with data de-duplication and disaster recovery functionality
- **IBM Tivoli Storage Manager for Mail** — helps secure IBM Lotus® Domino® and Microsoft Exchange data
- **IBM Tivoli Storage Manager for Databases** — helps secure IBM Informix®, Oracle and Microsoft SQL data
- **IBM Tivoli Storage Manager HSM for Windows** — provides Hierarchical Storage Management with a policy-based management system
- **IBM Tivoli Storage Manager for Advanced Copy Services** — protects your mission-critical data that requires 24x7 availability with snapshot backup
- **IBM Tivoli Storage Manager for Copy Services** — high-efficiency B/R of data and applications, eliminating backup-related performance impacts.
- **IBM Tivoli Storage Manager for Enterprise Resource Planning** — helps protect vital SAP R/3 system data efficiently, consistently and reliably.
- **IBM Tivoli Storage Manager for Space Management** — automatically moves inactive data to free online disk space for important active data.
- **IBM Tivoli Storage Manager for Storage Area Networks** — for SAN-connected Tivoli Storage Manager servers and client computers
- **IBM Tivoli Storage Manager for System Backup and Recovery** — offers a comprehensive system B/R and reinstallation tool with bare-metal restore
- **IBM Tivoli Storage Manager FastBack™** — provides a continuous data protection and recovery management platform for Microsoft Windows servers.
- **IBM Tivoli Storage Manager FastBack for Microsoft Exchange** — provides the ability to quickly and easily recover granular Microsoft Exchange data
- **IBM Tivoli Storage Manager FastBack for Bare Machine Recovery** — restores entire systems, whether to comparable hardware, dissimilar hardware, or a virtual machine.
- **IBM Tivoli Storage Manager FastBack Center** — combines the features of the IBM Tivoli Storage Manager FastBack family of products into one solution.
- **IBM Tivoli Continuous Data Protection for Files** — provides continuous, automated backup of desktop and laptop workstations.