

## System z Expo

October 13 – 17, 2008 – Las Vegas, Nevada



# Mainframe Resiliency with z/VSE

**zEG10**

Wilhelm Mild

Authorized

**IBM** | **Training**

23-Sep-08

© 2008 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](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.

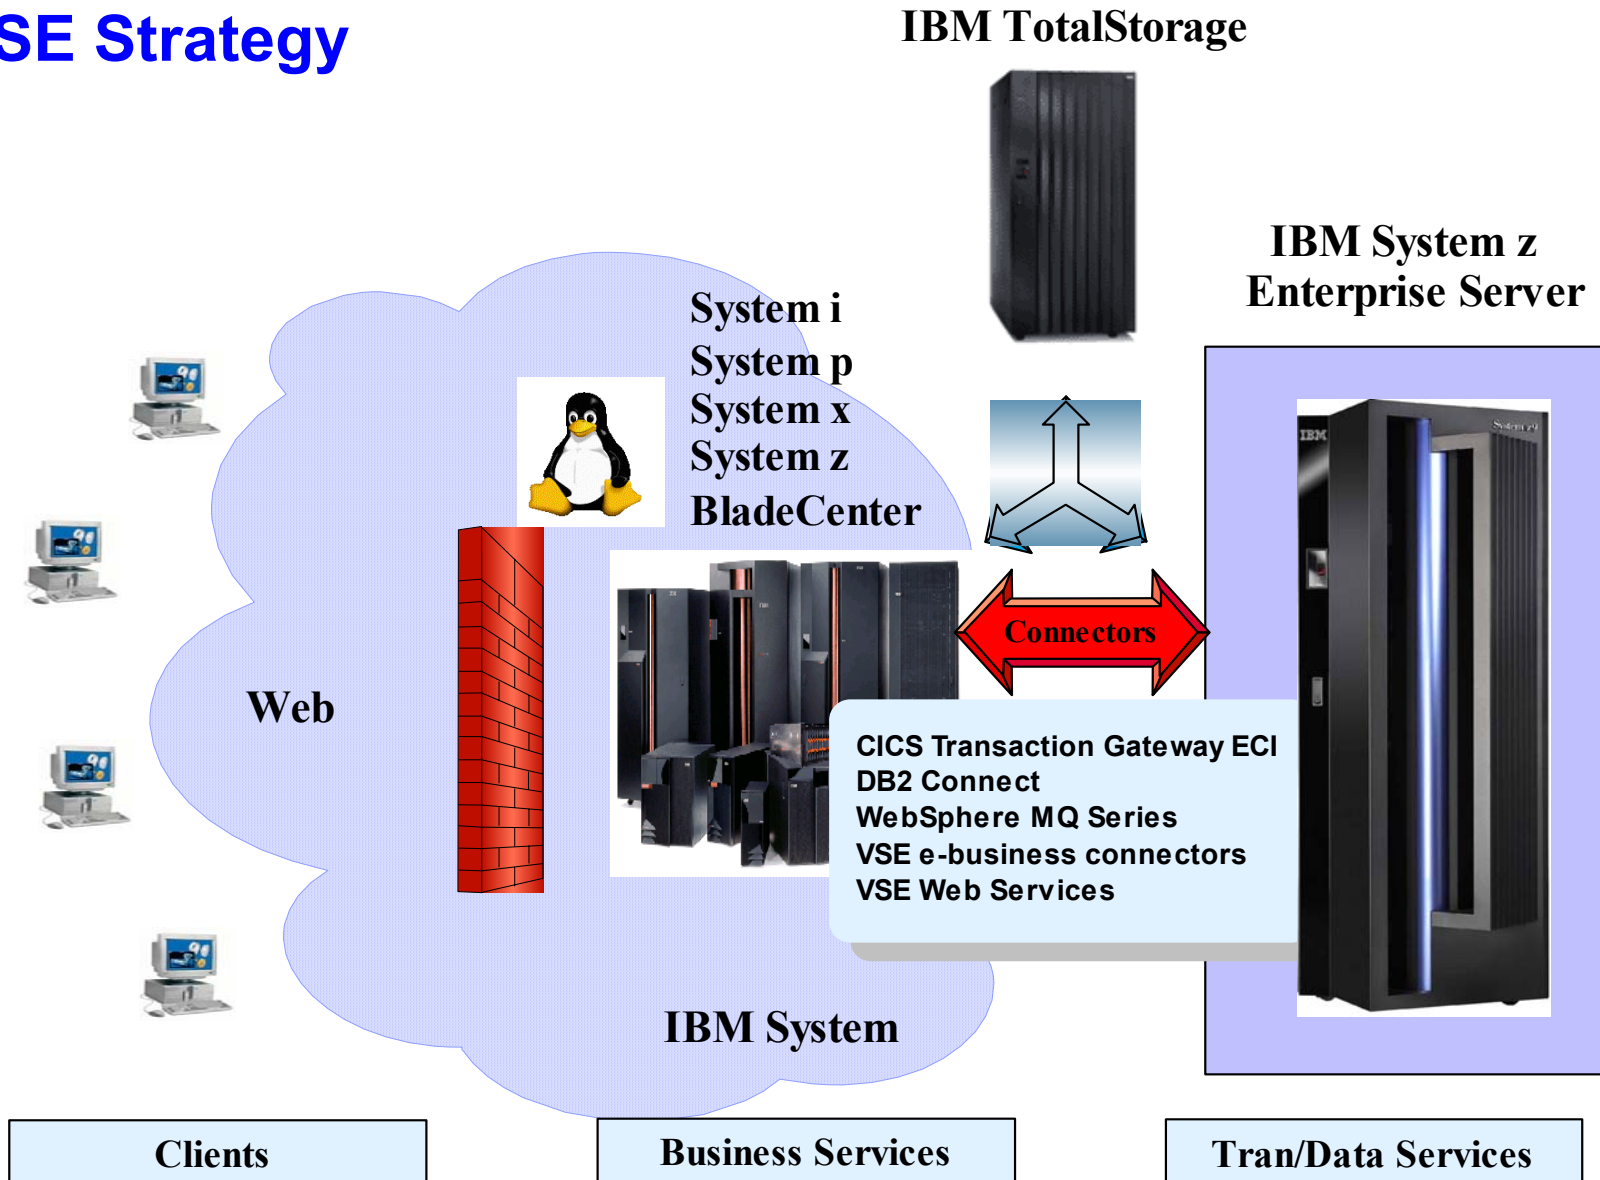
**Together, these two tasks can make the mainframe again the focus of users' business resiliency strategies.**

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

# VSE Strategy





## System z Storage Options for z/VSE



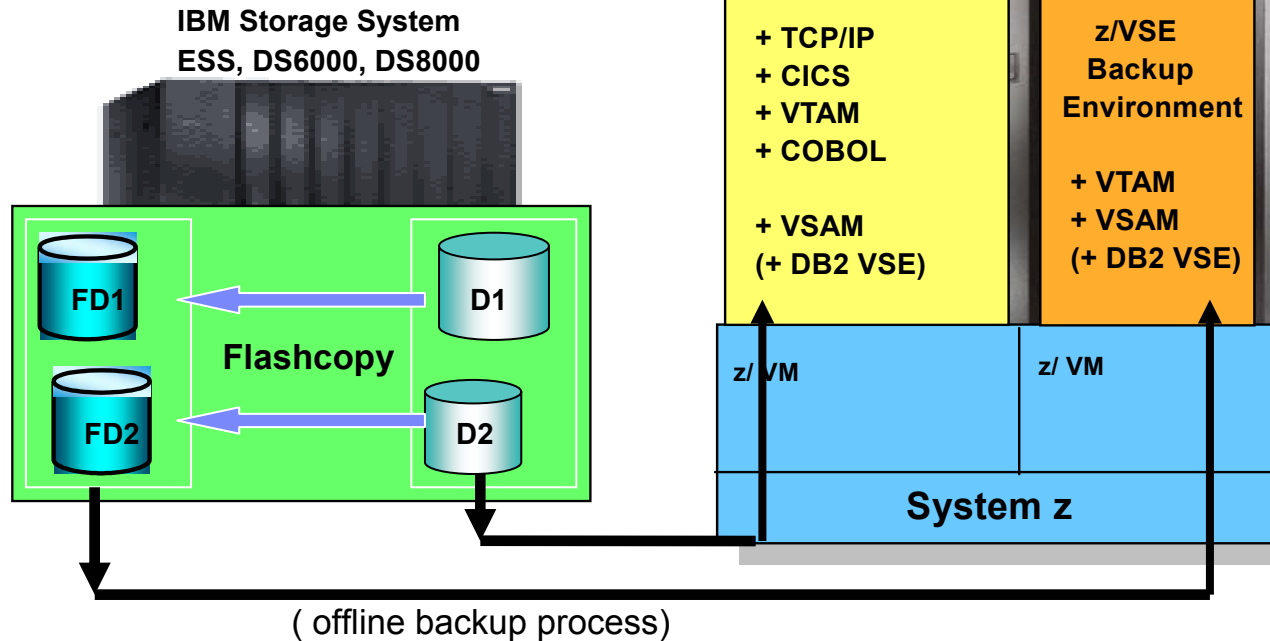
IBM TotalStorage	DS6000	ESS 750, 800, 800Turbo	DS8000
ESCON	Not Avail	Yes	Yes
FICON	Yes	Yes	Yes
FCP/SCSI	Yes	Yes	Yes

# IT Environment for 24x7 Availability



### Flashcopy:

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






## IT Environment Needs for 24x7 Availability

**x** inhibitors of **online** processing time

☞ backup-window

☞ batch-window

Typical processing time-line:

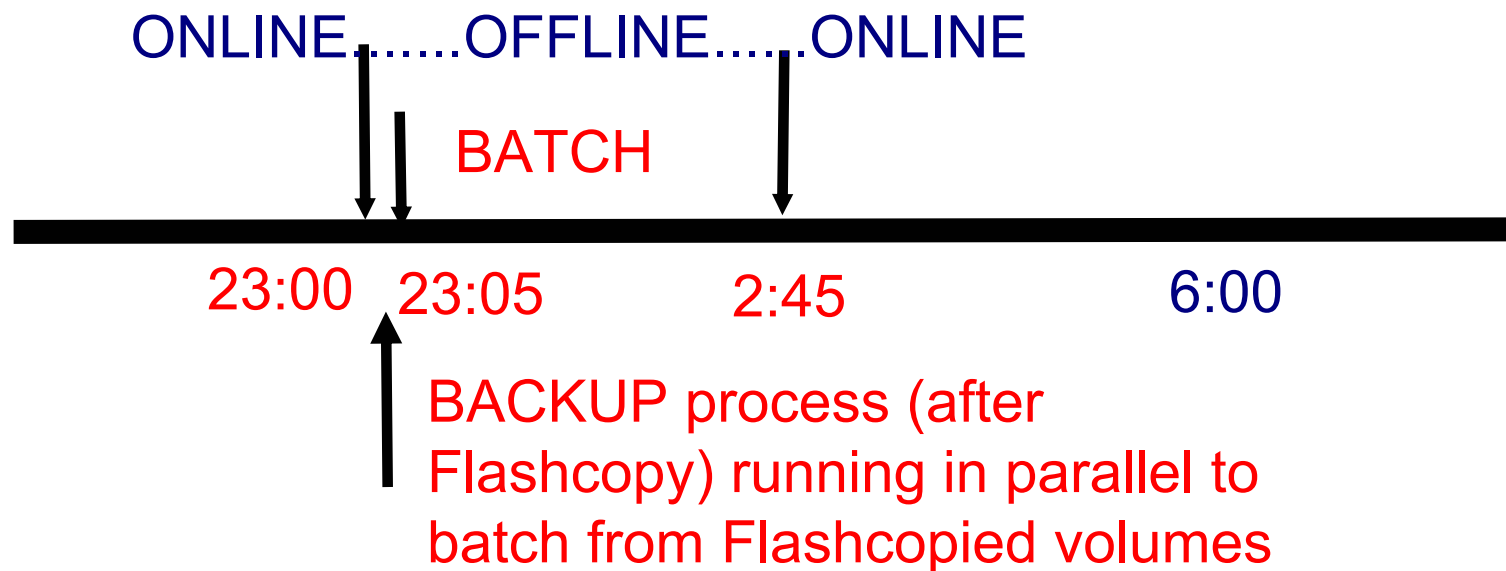


**BACKUP** process (after Flashcopy) running in parallel to batch from Flashcopied volumes

## IT Environment Needs for 24x7 Availability

- modern Storage solutions can reduce OFFLINE time:
  - eliminate backup window – using FLASHCOPY

Typical processing time-line:

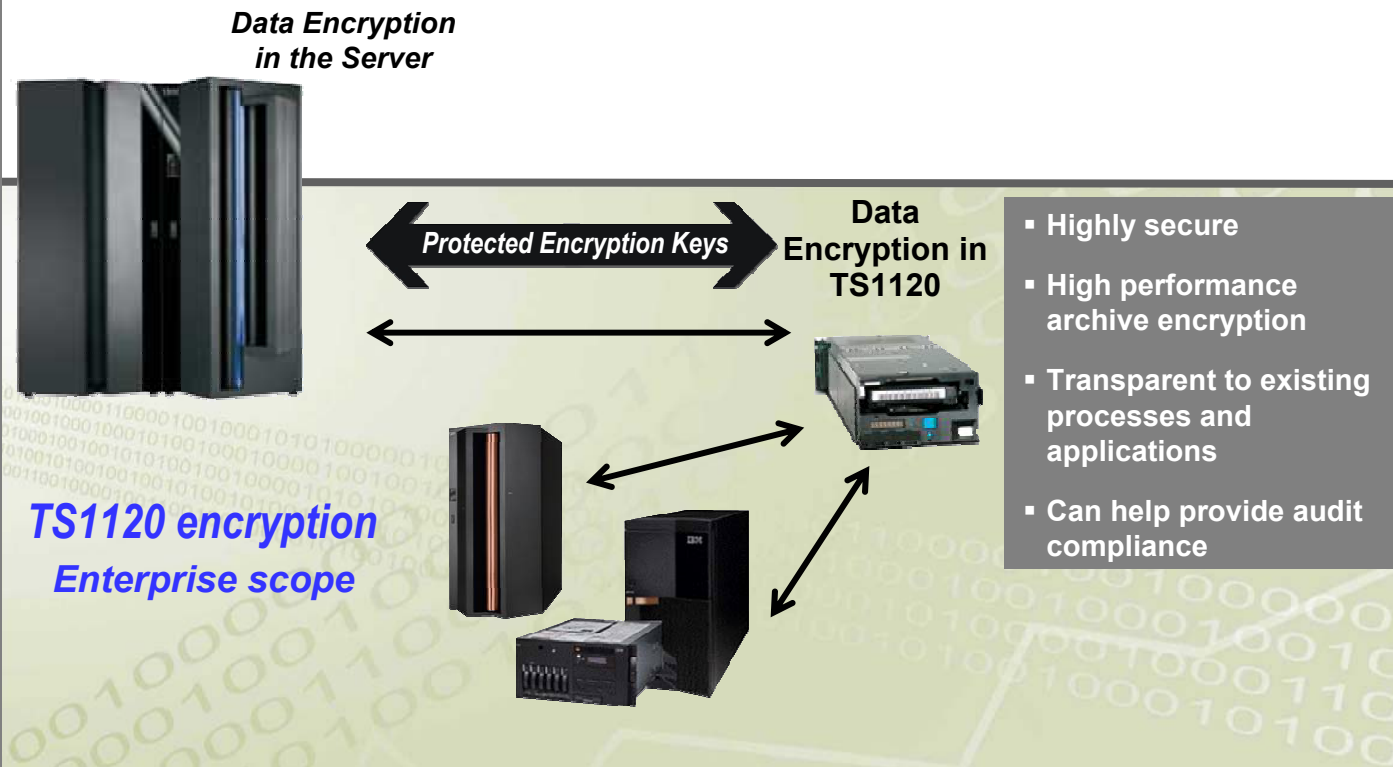


# IBM TS1120 Tape Drive Encryption for z/VSE

## Centralized key management

- Help protect and manage encryption keys
  - Highly secure and available key data store
  - Long term key management
  - Disaster recovery capabilities
- Single point of control
  - Non-VSE, Java-based platform
  - TCP/IP connection to tape control unit

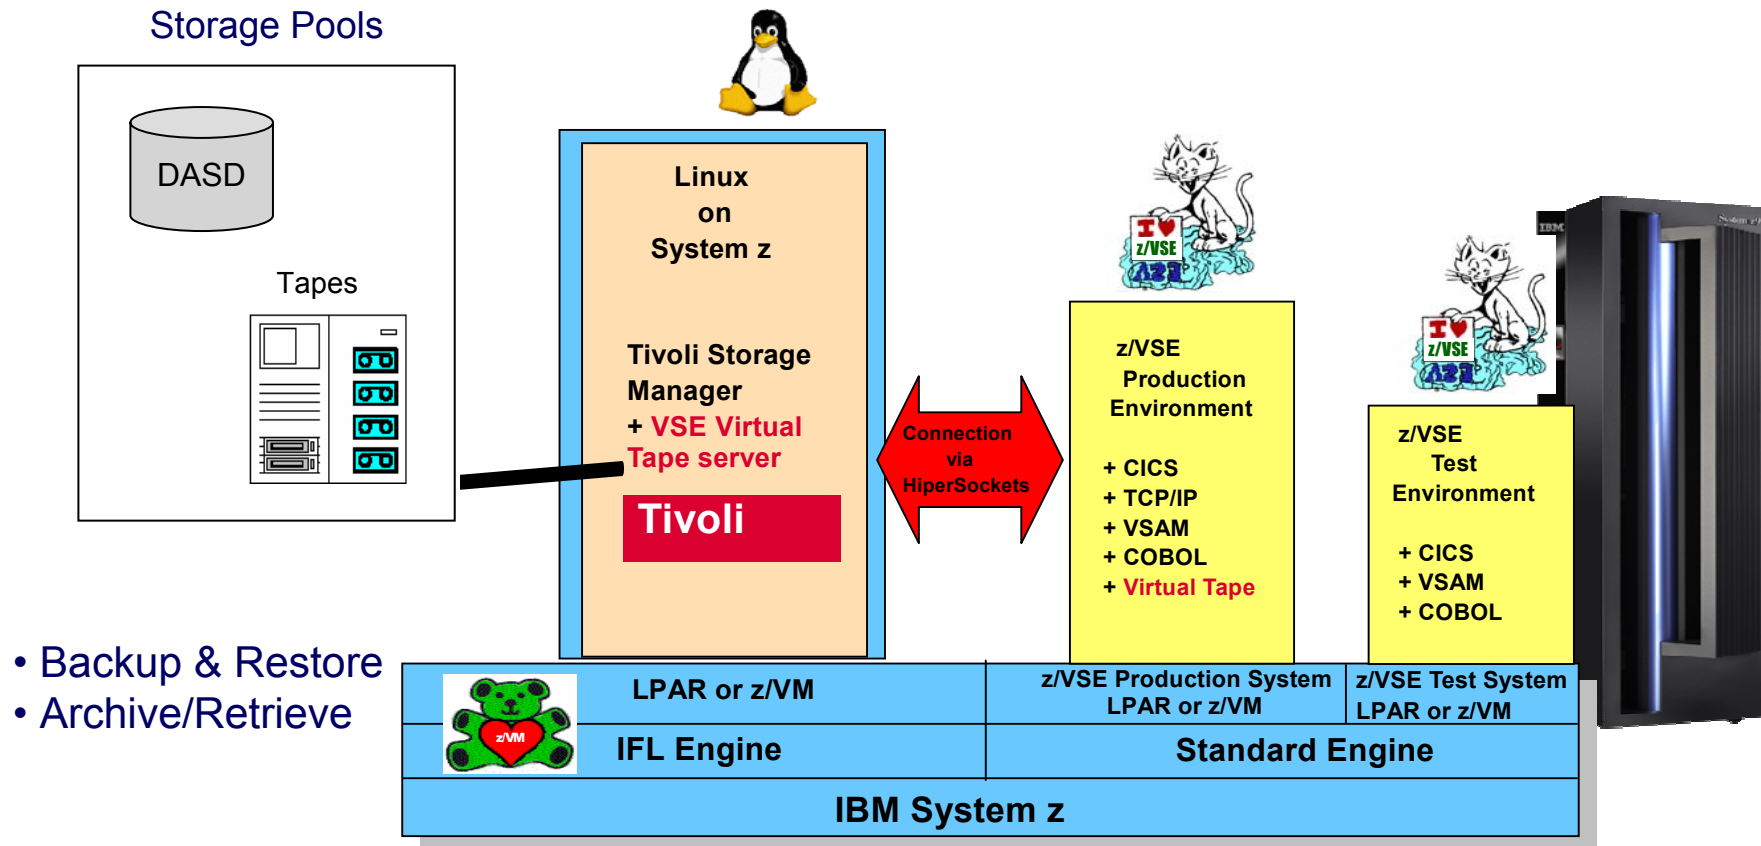
*z/VSE V3.1 and Z/VSE 4.1 support of the TS1120 Tape Drive with encryption. z/VSE support requires the Encryption Key Manager component running on another operating system other than z/VSE using an out-of-band connection.”*



- Highly secure
- High performance archive encryption
- Transparent to existing processes and applications
- Can help provide audit compliance

# z/VSE 4.1 – Backup integration

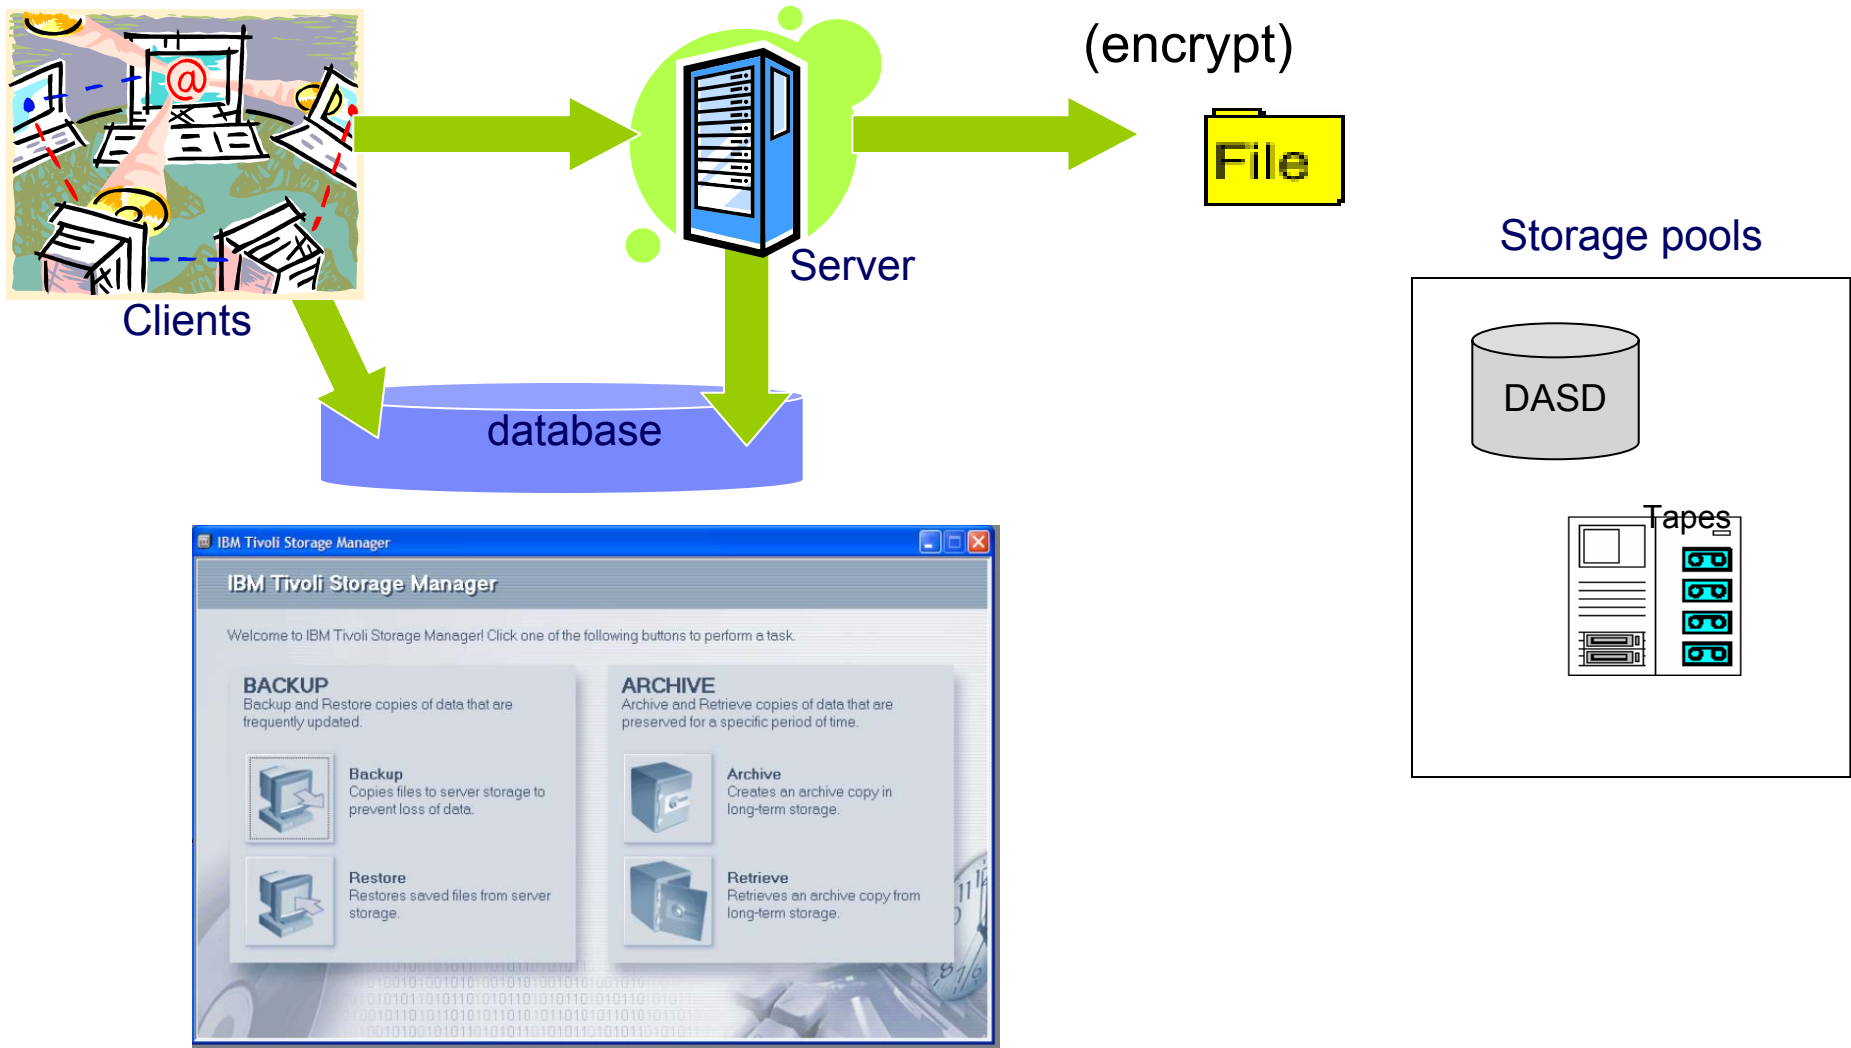
## Integration of z/VSE with Tivoli Storage Manager (TSM)



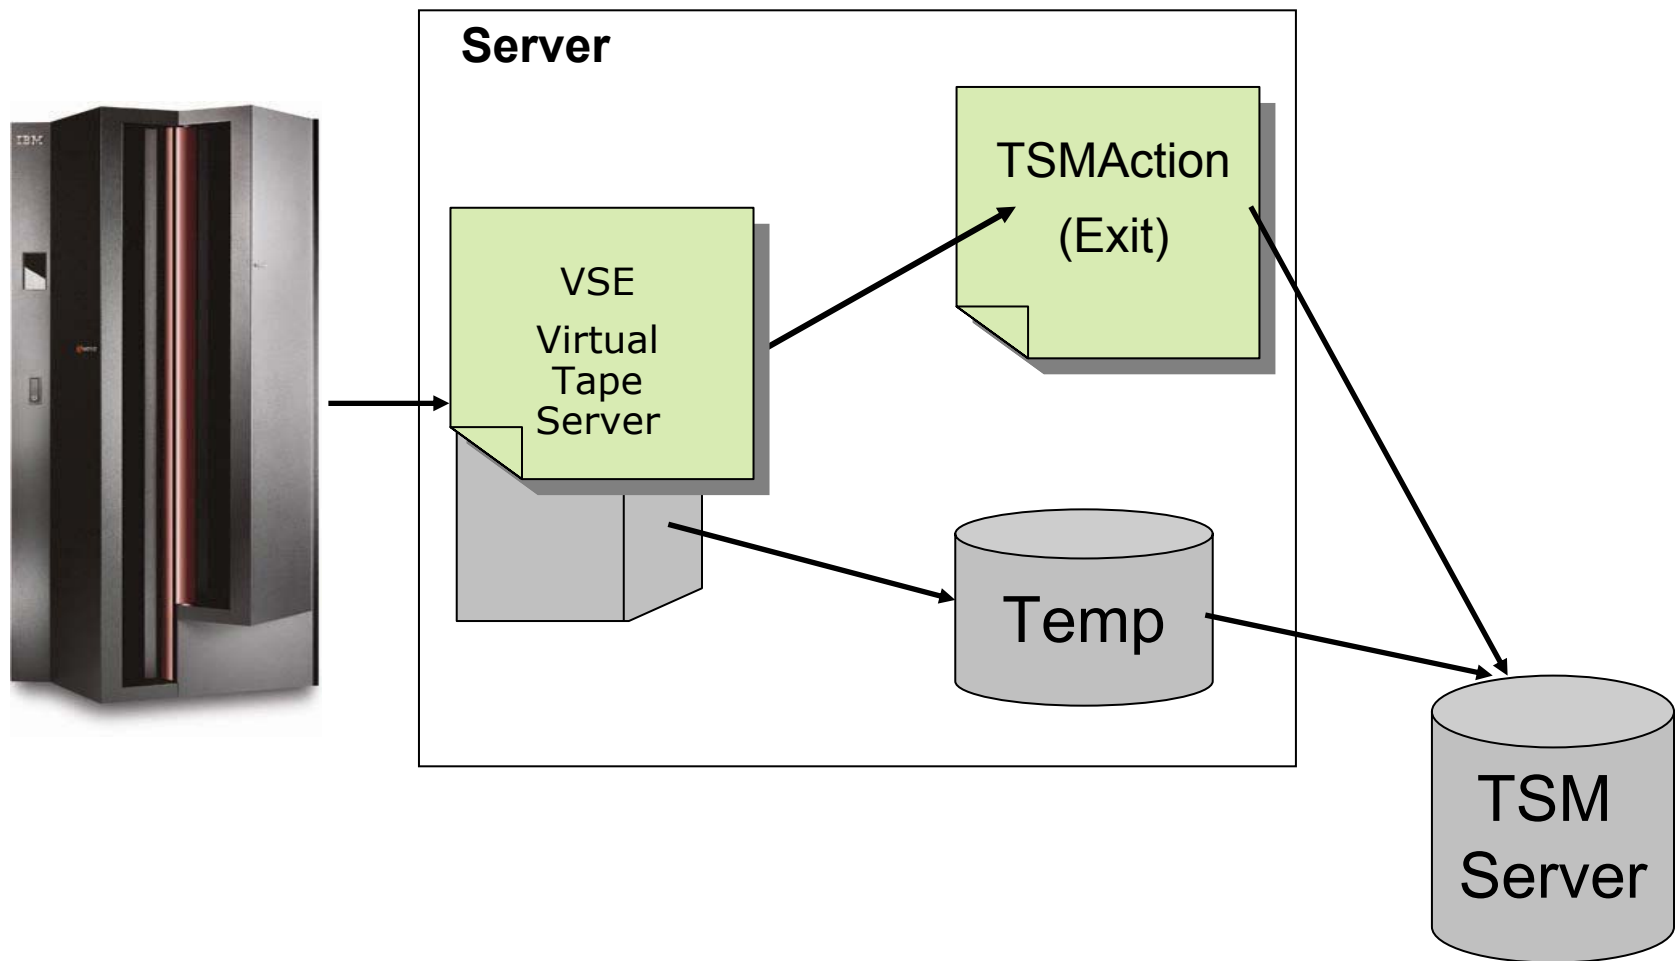
## z/VSE Backups with Tivoli Storage Manager

- New with z/VSE 4.1
- 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 Manager - Architecture



## 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

```

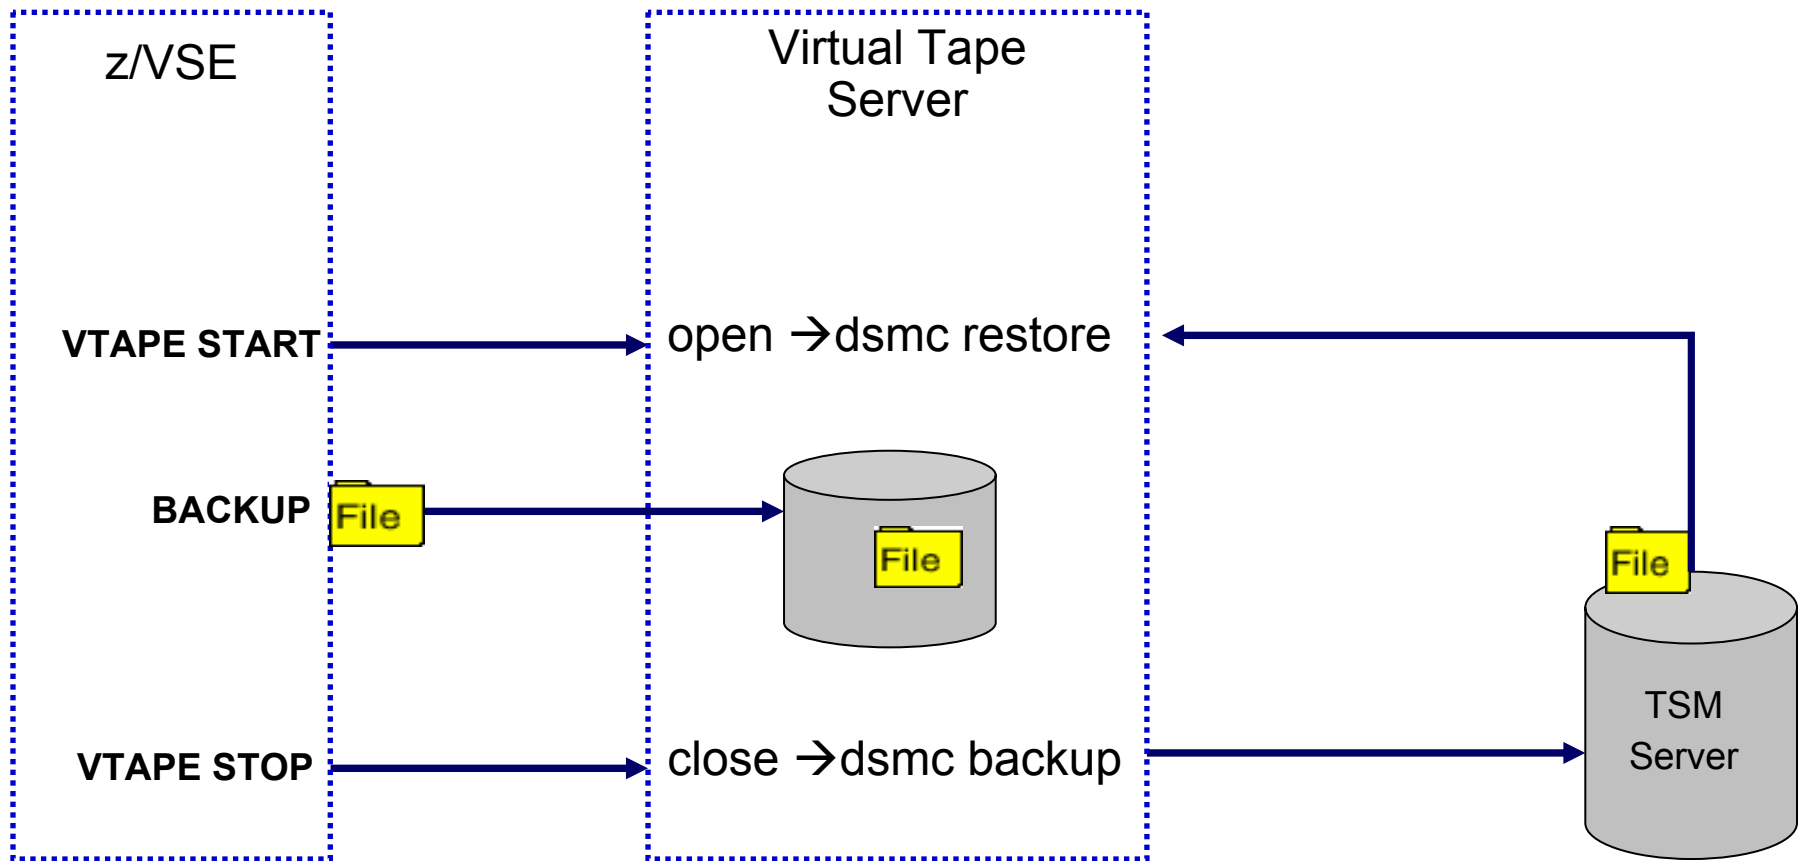
#### Syntax:

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



## Data Resiliency - Tapes

- Thus, for example, many businesses continue to believe that backup/restore to tape is their best strategy for data resiliency, whereas backup/restore may not meet the availability (it takes a long time to back up and restore) and recoverability (some transactions may be lost) requirements of the business for some applications.

# Enterprise Storage solutions – disaster recovery (Peer to Peer remote Copy - PPRC)



## 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

## Concepts of Disaster Recovery with VSE

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

## 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

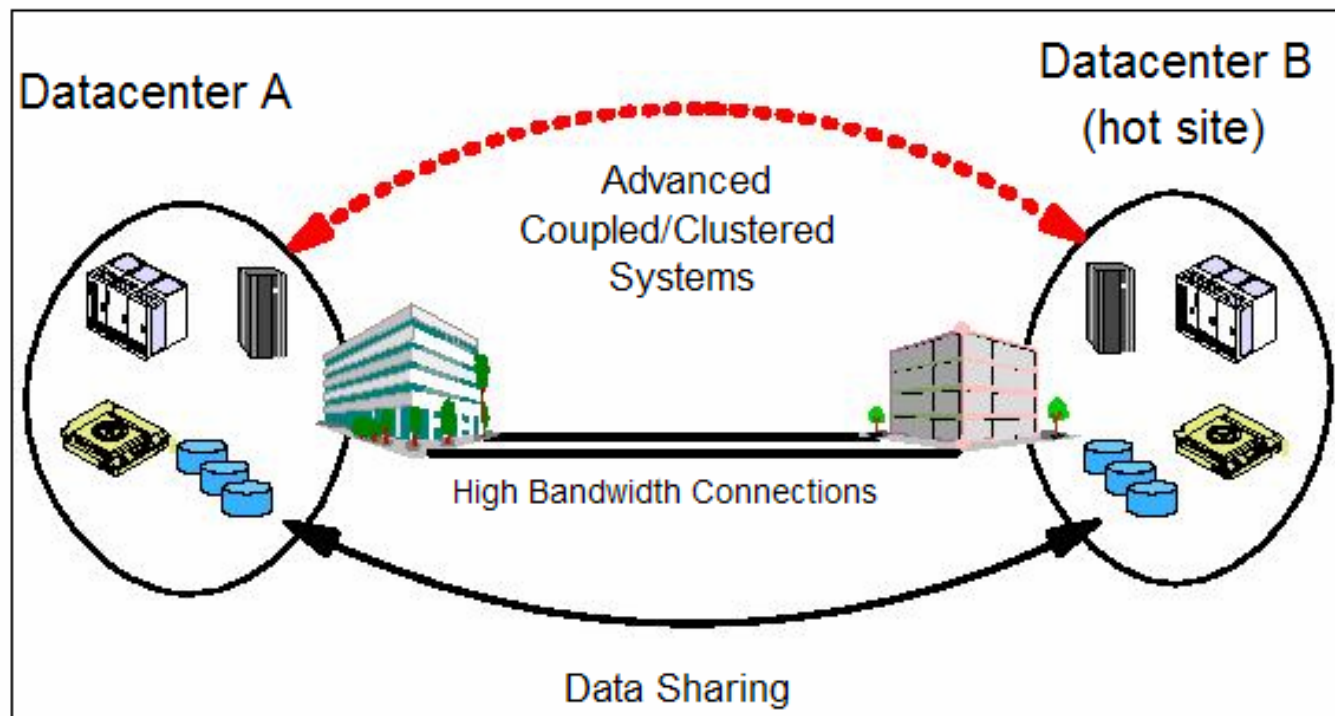
## Objectives for Disaster Recovery with VSE

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




*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 normally a few minutes.






*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 usually less than 12 hours.



*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 a day.

## 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.

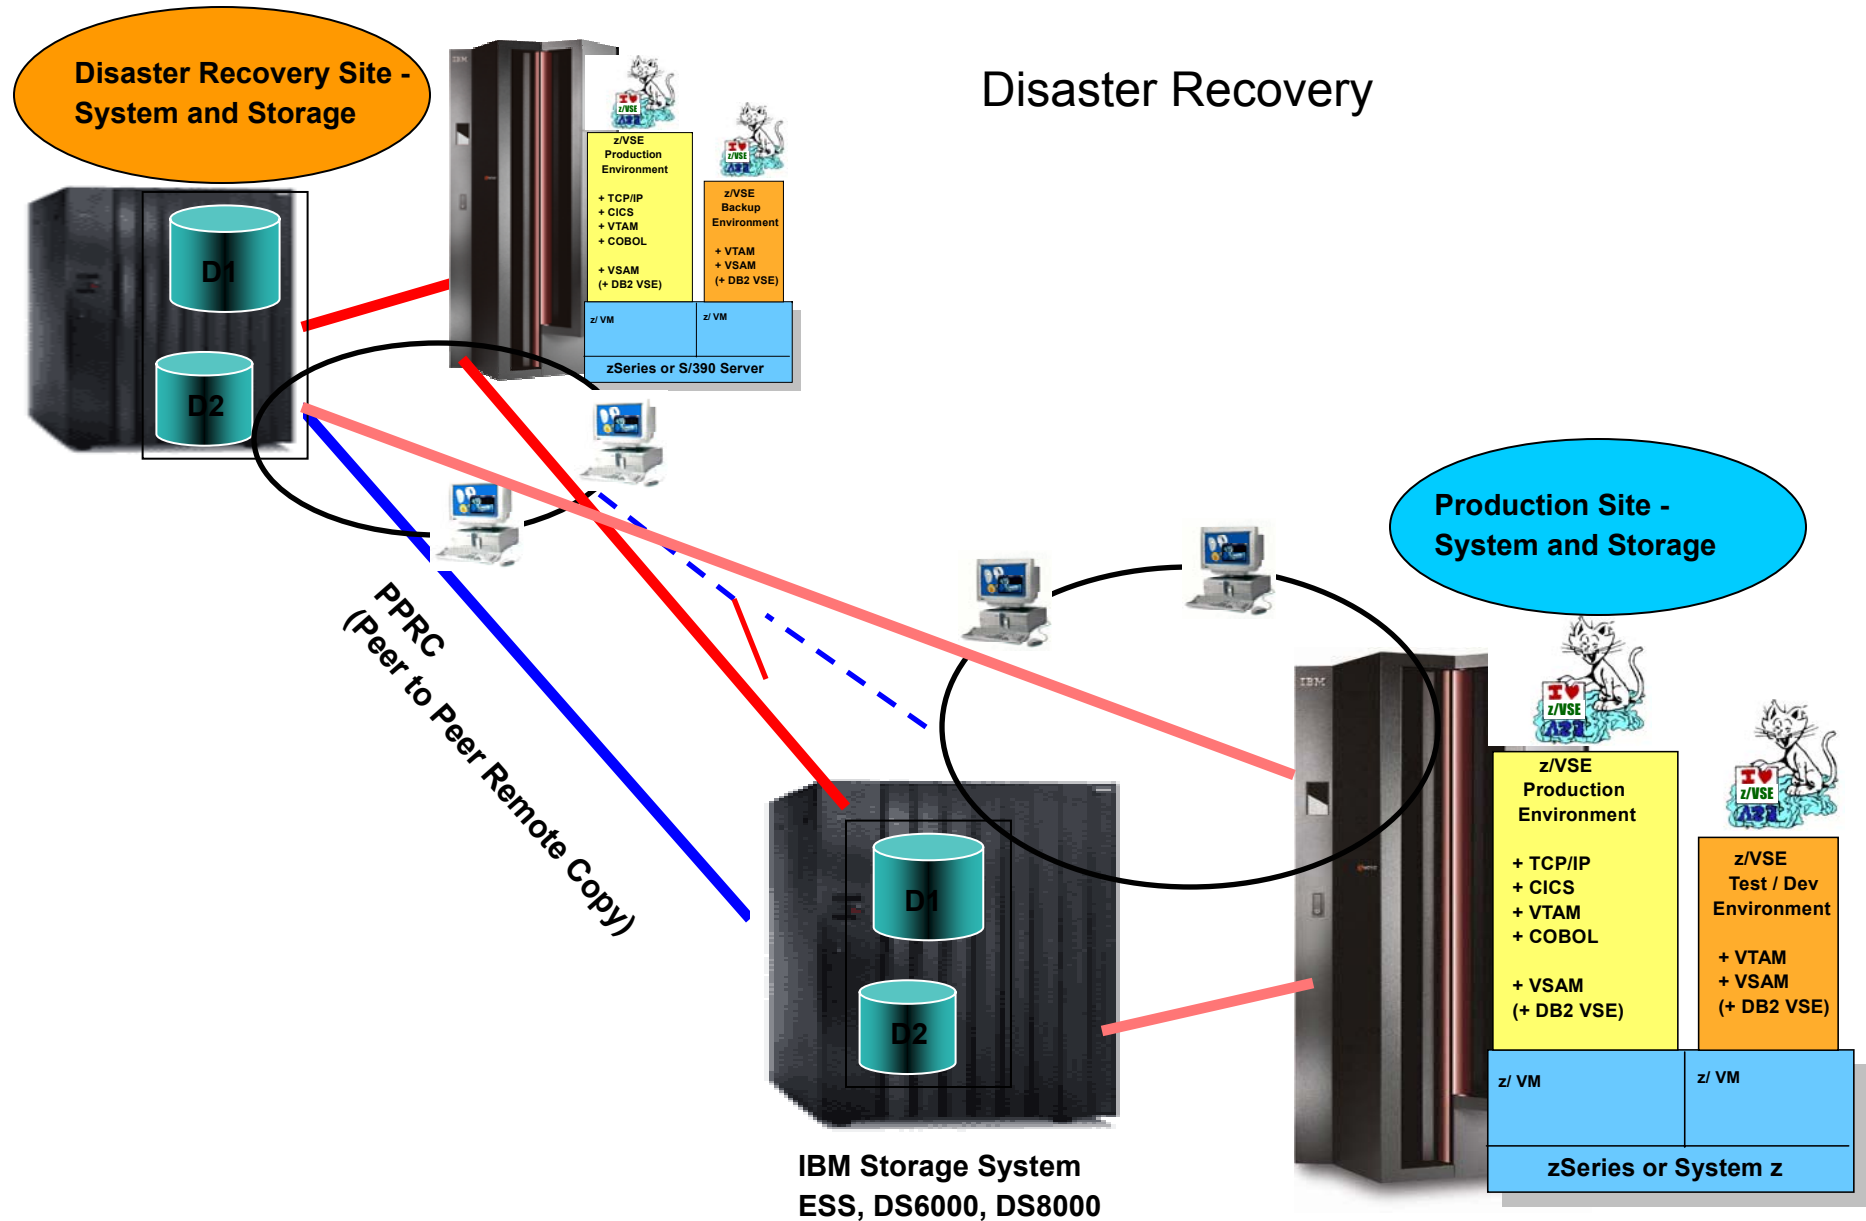
## System environment Agreements for DR – continued (2)

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.




# Disaster Recovery

**Disaster Recovery Site - IBM Storage System**

Disaster Recovery rules definition and infrastructure automation



(TPC) Totalstorage Productivity Center for Replication with Two Site Business Continuity



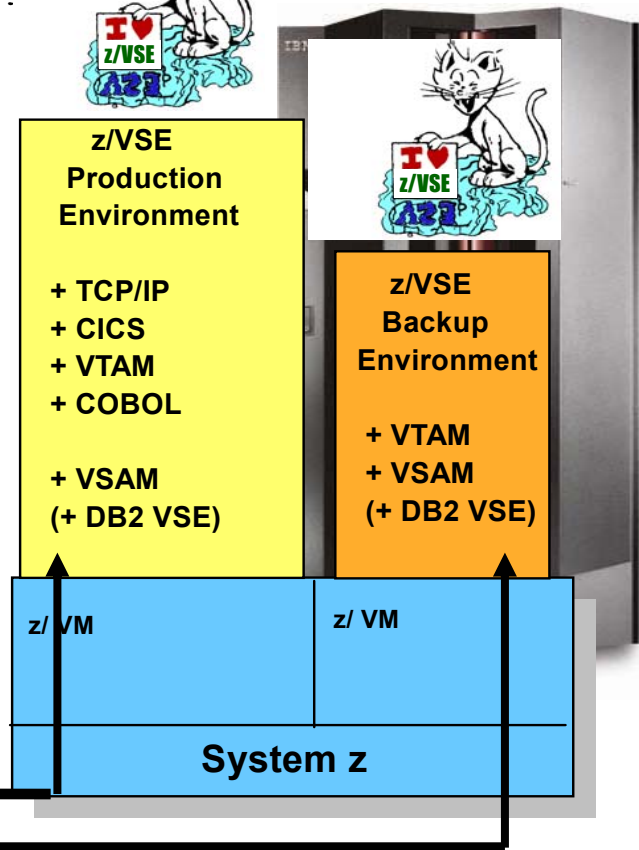

**PPRC (Peer to Peer Remote Copy)**



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 only and one for DR

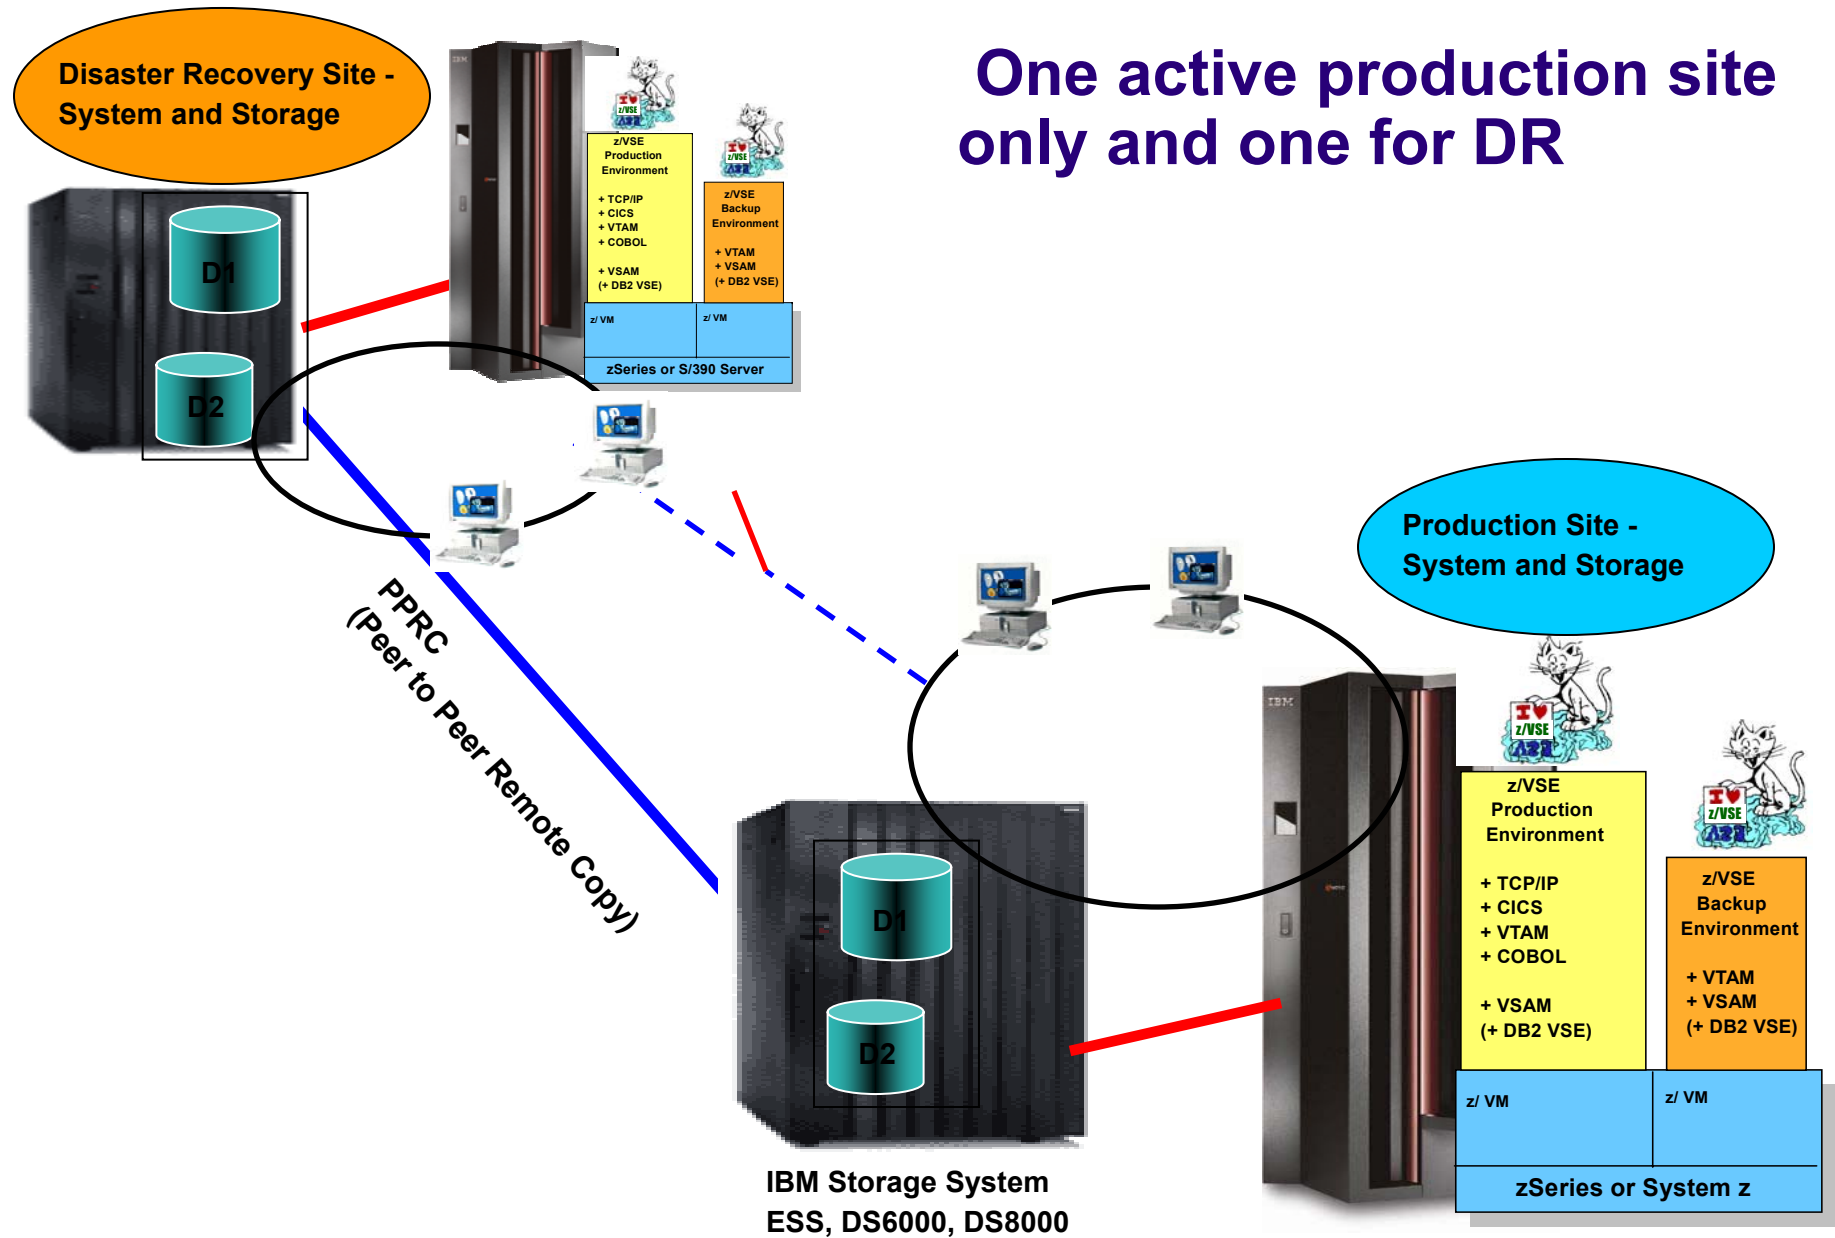
(3) Two active sites with production and test

(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 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**

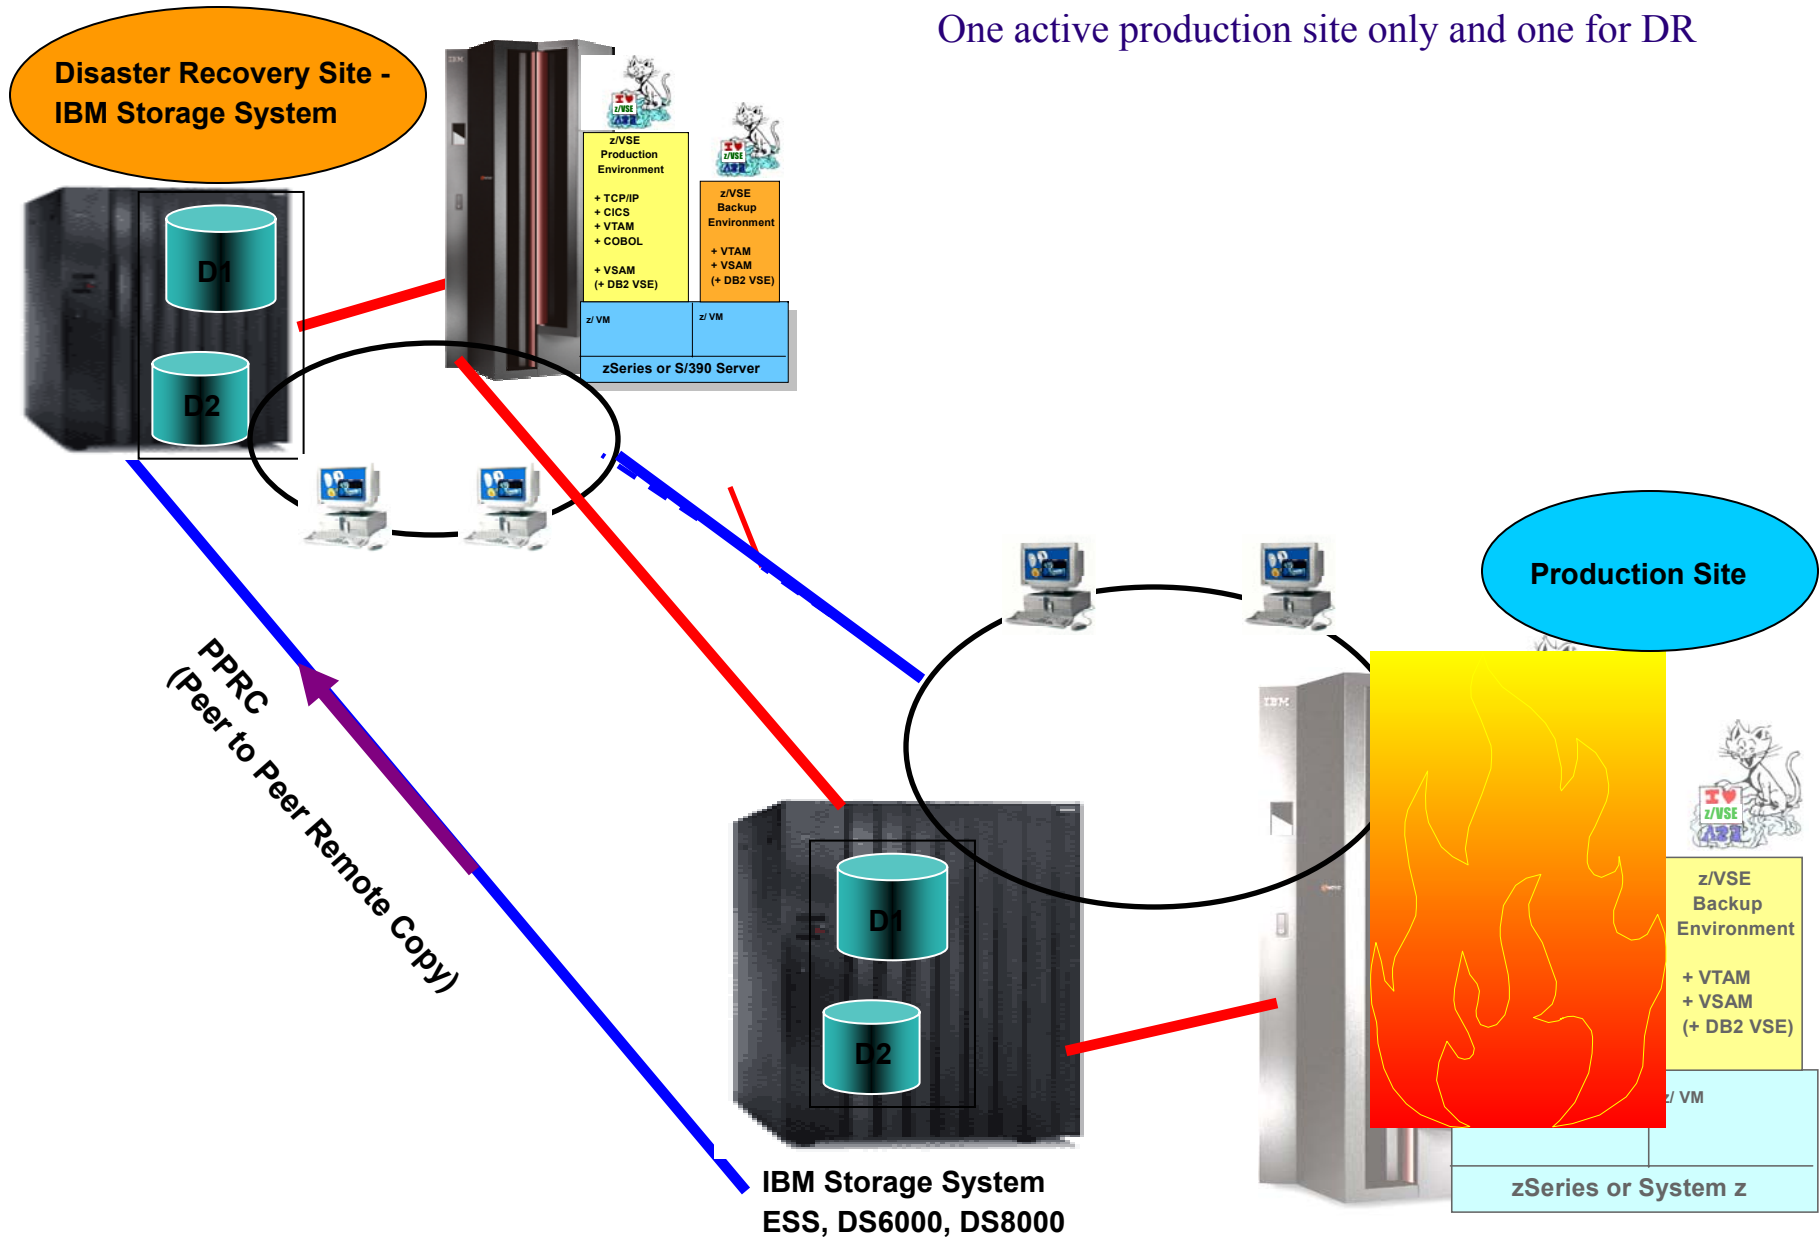


# One active production site only and one for DR

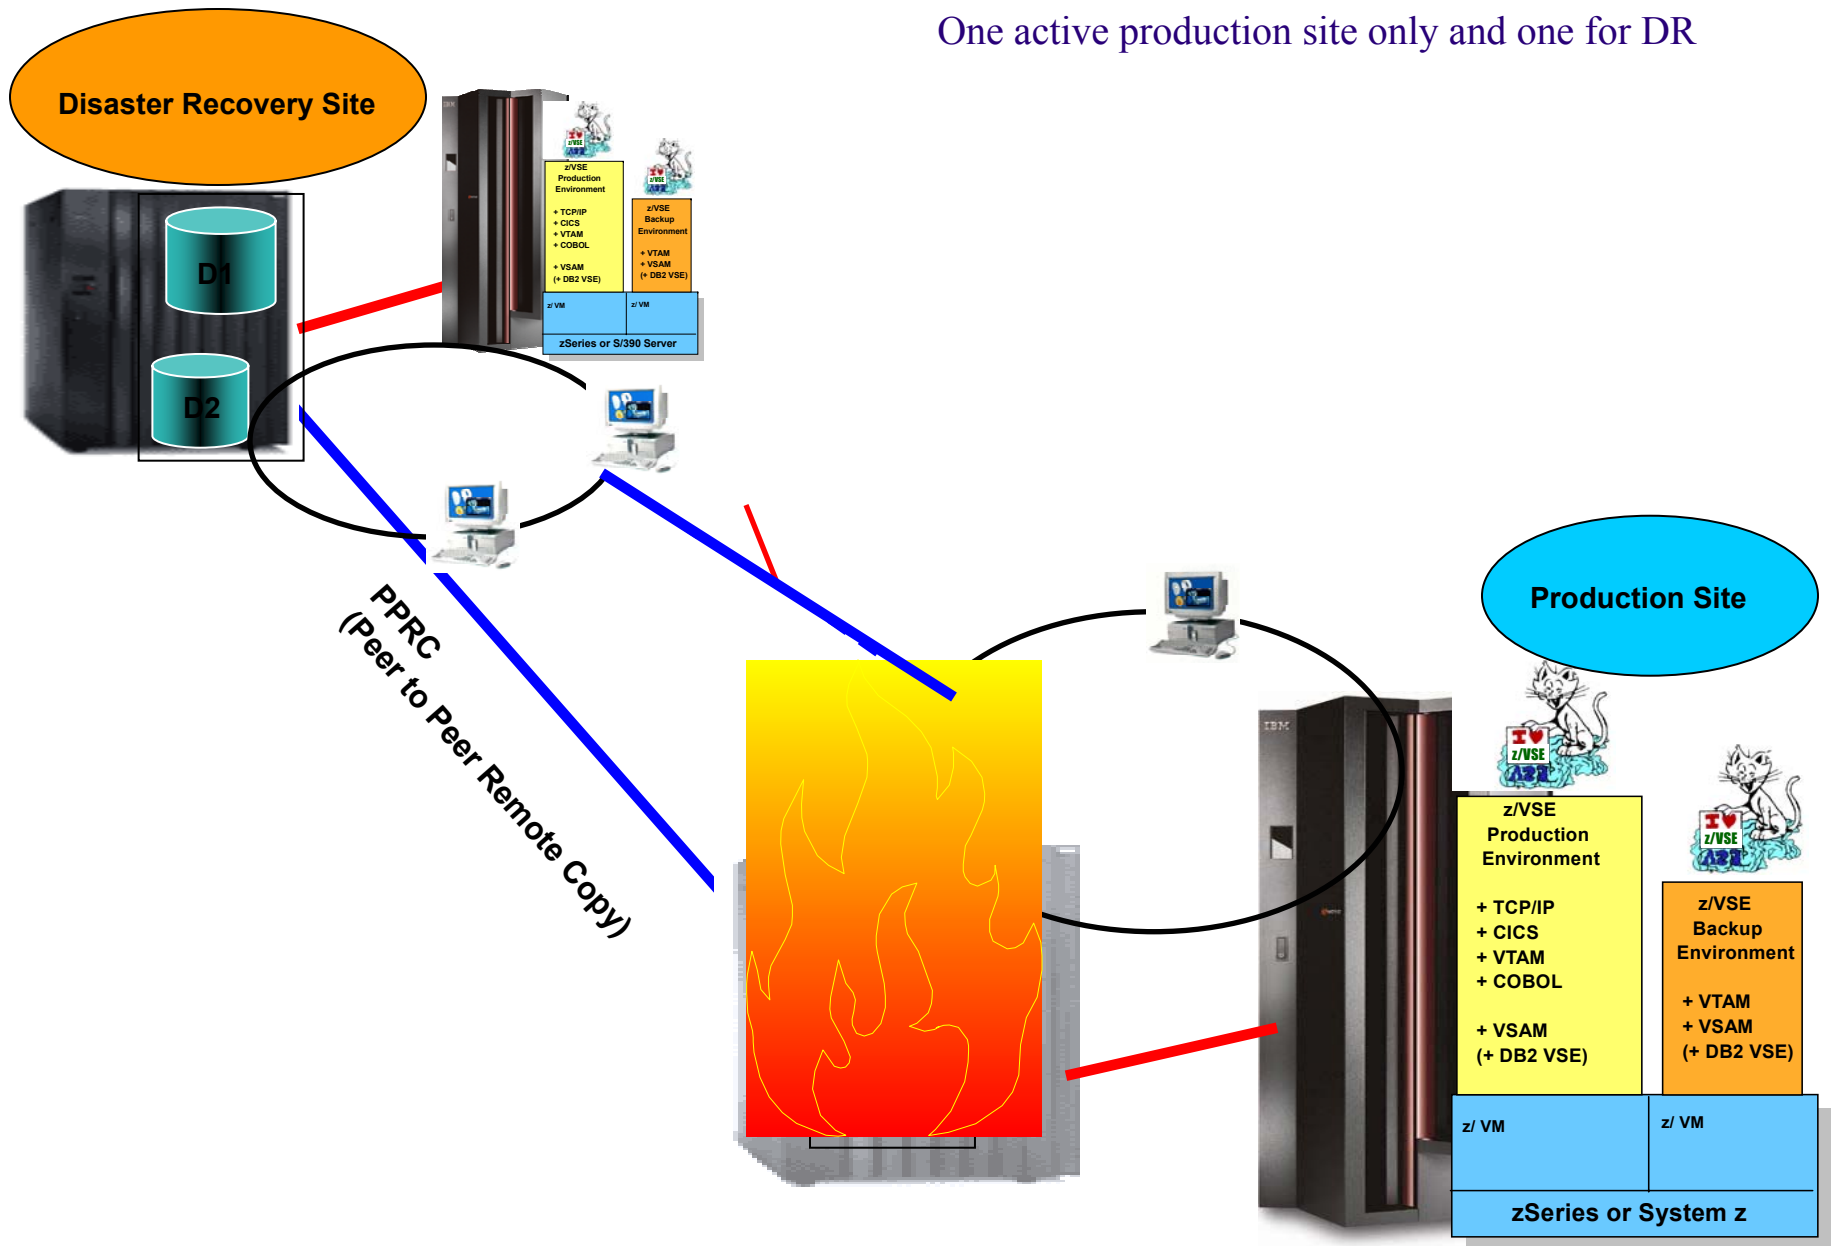
# One active production site only and one for DR




One active production site only and one for DR



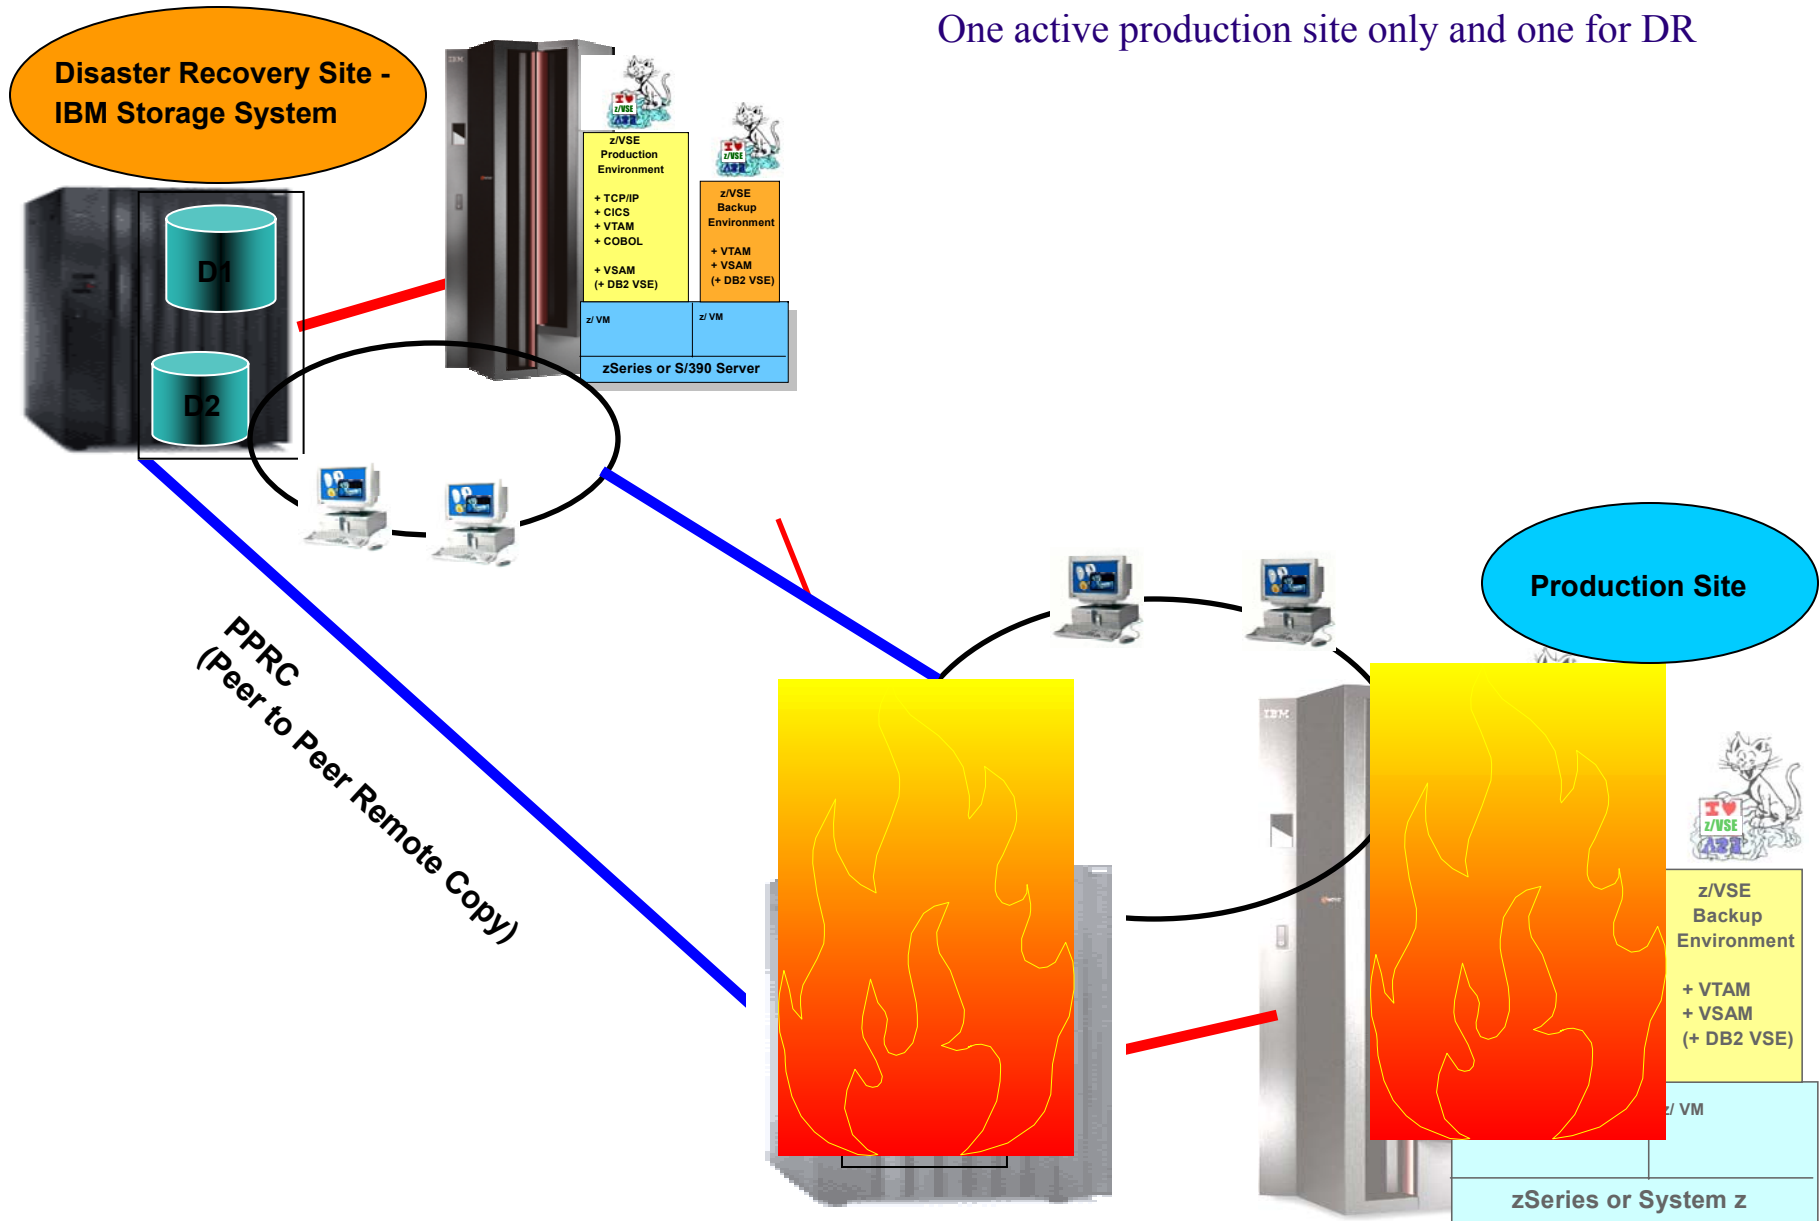
One active production site only and one for DR



One active production site only and one for DR



One active production site only and one for DR





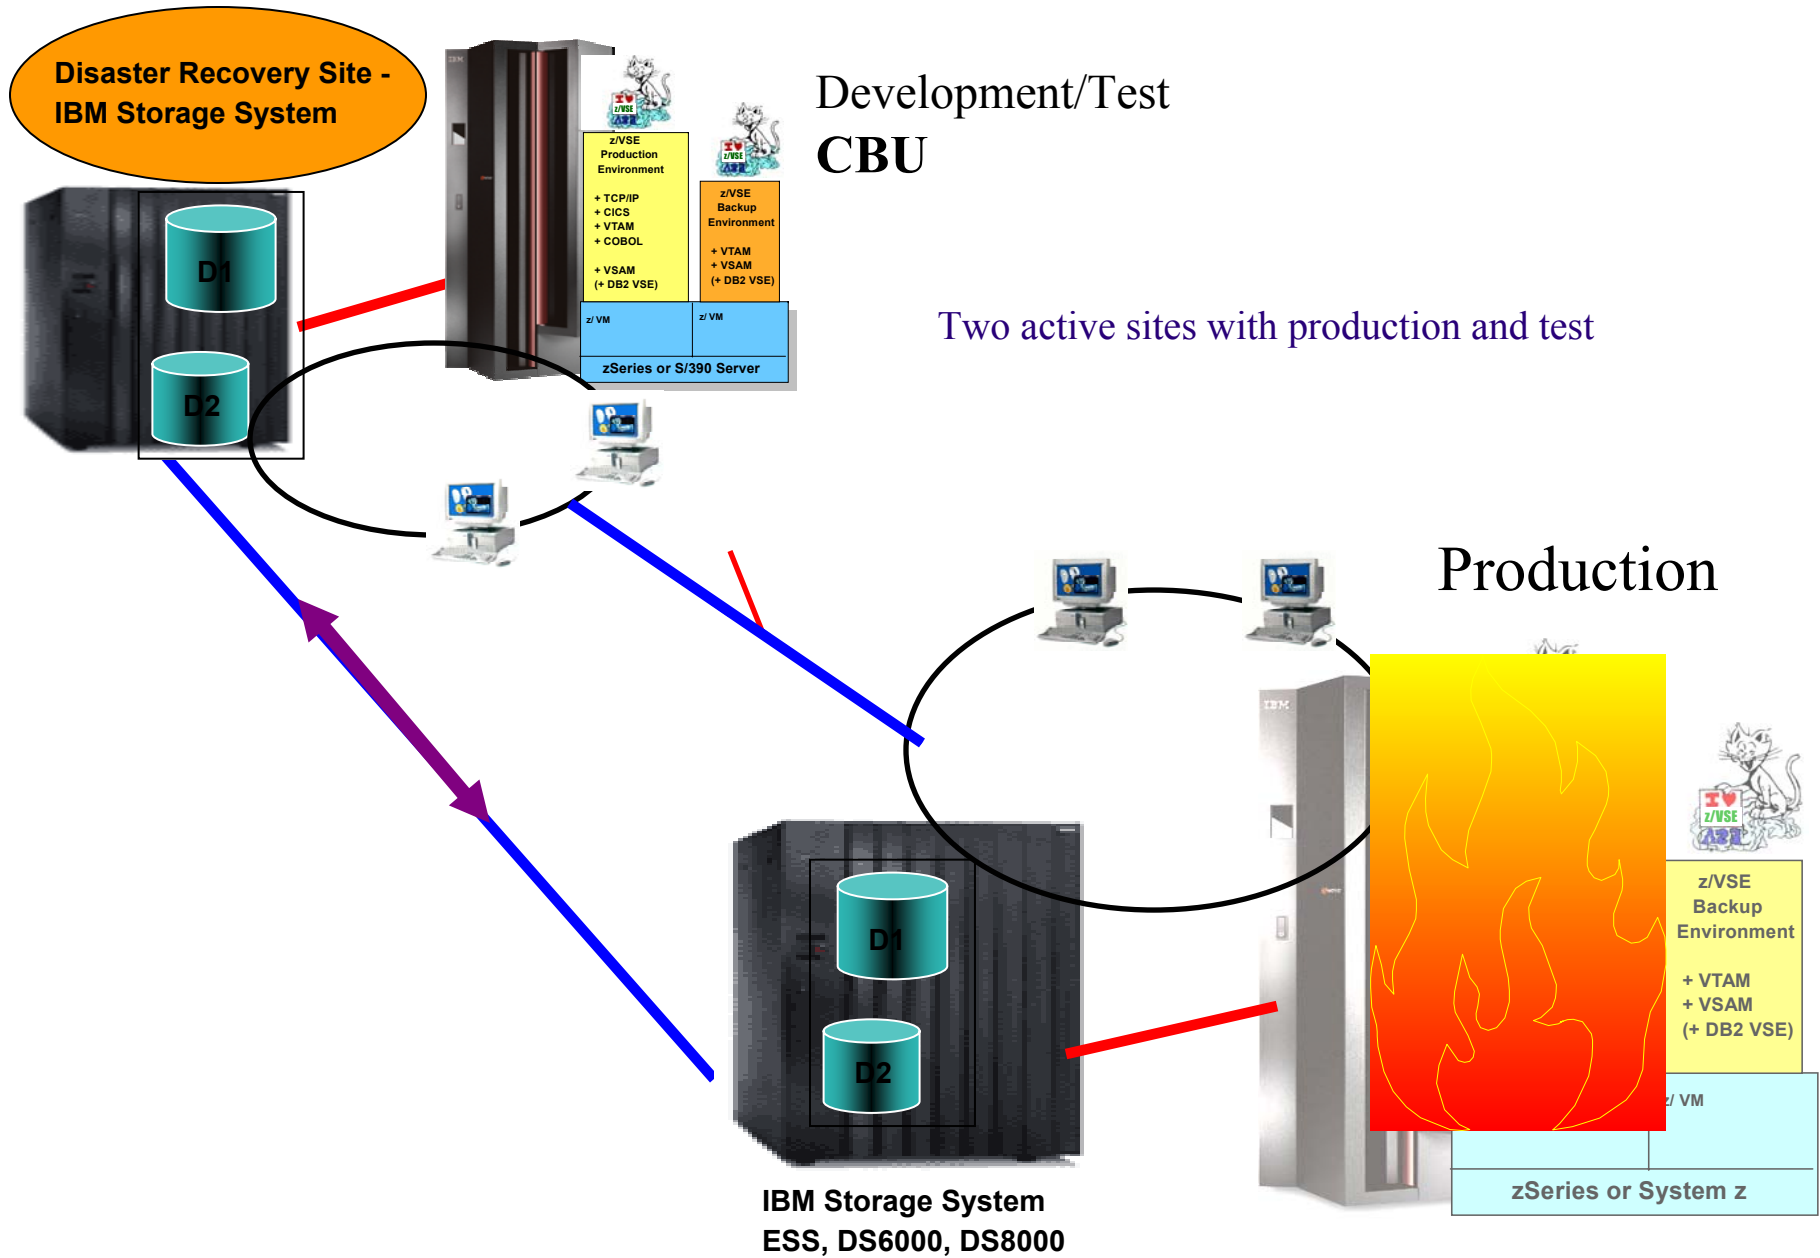
## Steps in case of a disaster Recovery

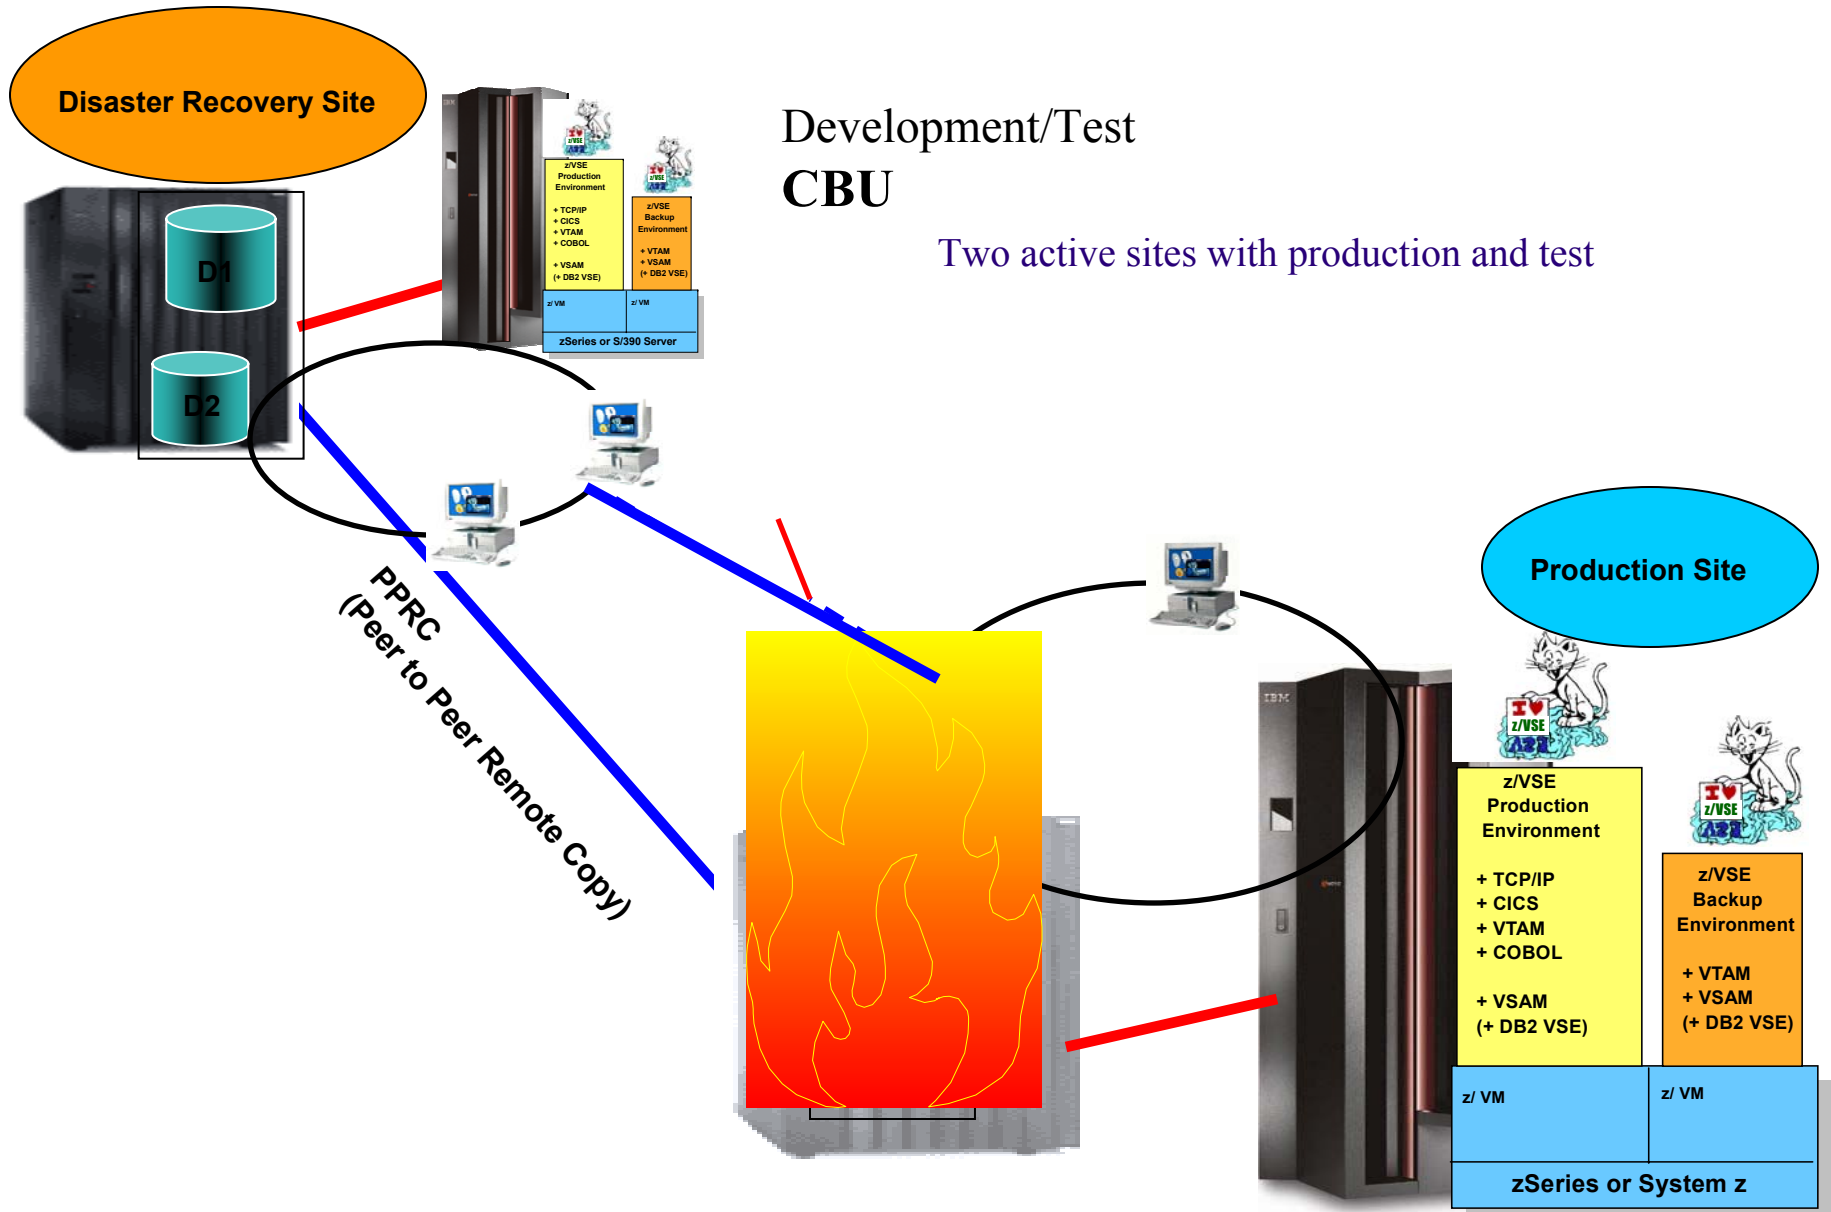
1. **Emergency phones and messaging methods have to be enabled**
2. **Start z/VM on the Recovery Site (on a COLD environment)**
  1. **Start the CBU (Capacity Backup Upgrade) if defined – to accelerate start of VSE systems**
3. **Switch the OSA Adapter Network Connectivity**
4. **Start Online VSE machines (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.**

## 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 2<sup>nd</sup> 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

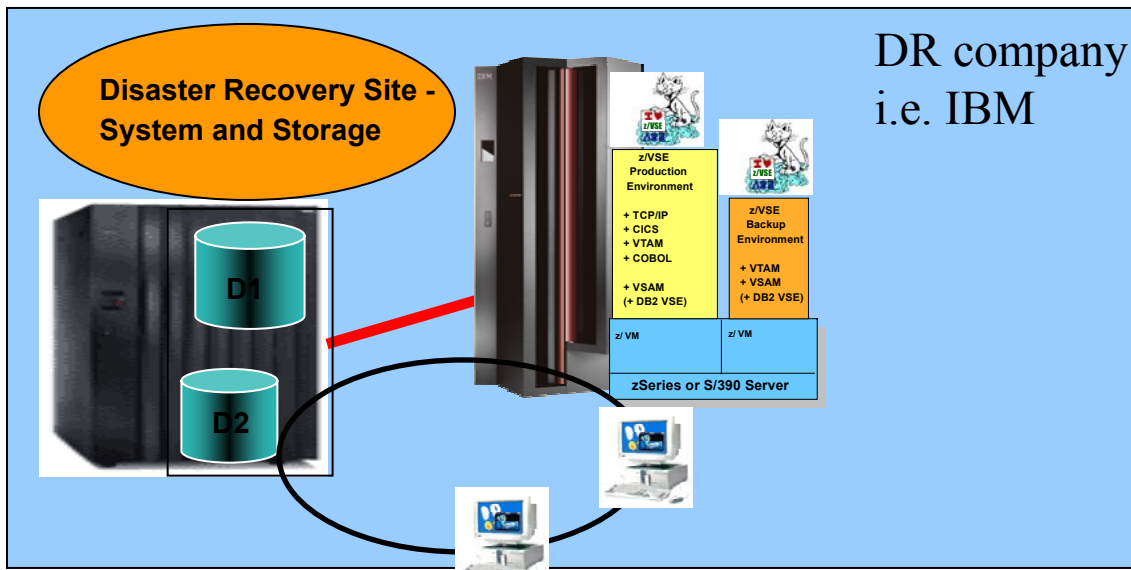


## Off-site Disaster Recovery

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

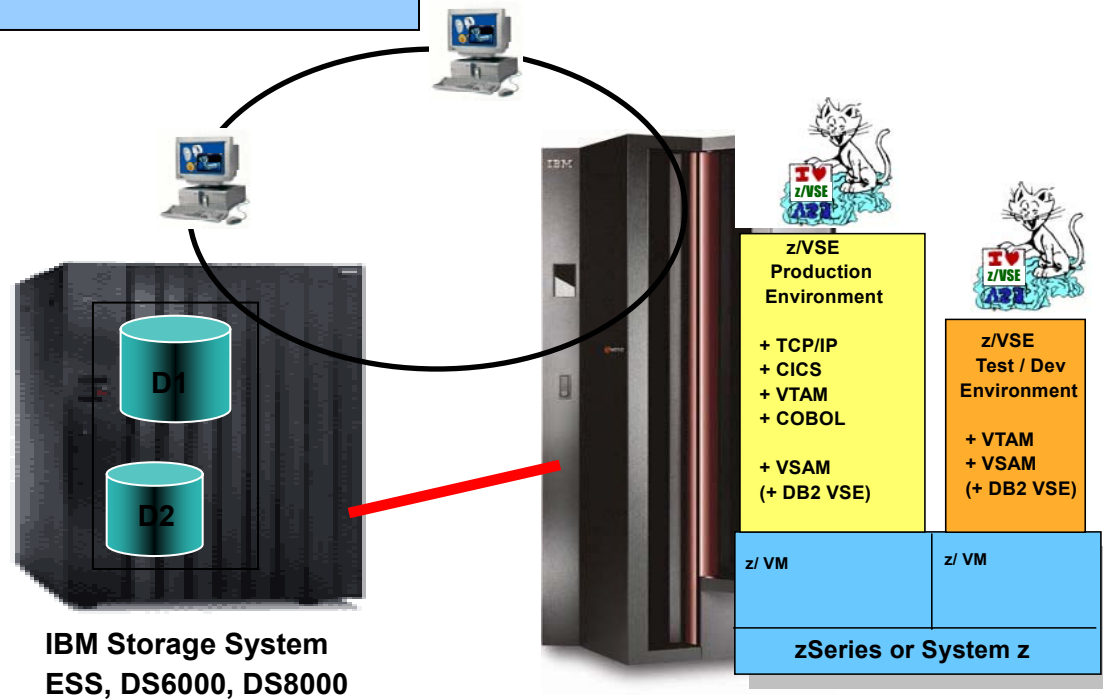
Necessary Agreements are required:

- 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.



## z/VSE on the web

z/VSE Solutions



The screenshot shows the IBM z/VSE website interface. The address bar displays `http://www.ibm.com/servers/eserver/zseries/zvse/`. The navigation menu on the left includes links for Home, Products, Services & solutions, Support & downloads, and My account. The 'Solutions' link is circled in red. The main content area features a large heading 'z/VSE' and a paragraph describing the product's capabilities. Below this is a banner for 'Announcing z/VSE V3.1' with a '40 YEARS' logo. Further down, there is a section titled 'Redesigned z/VSE homepage' and another for 'z/VSE Version3 Release 1' which lists supported hardware and software configurations.

New Web presence: [ibm.com/servers/eserver/zseries/zvse/solutions](http://ibm.com/servers/eserver/zseries/zvse/solutions)

## ***Additional Informations***

- z/VSE Home Page  
<http://www.ibm.com/servers/eserver/zseries/zvse/>
- z/VSE Solutions and Utilities  
<http://www-1.ibm.com/servers/eserver/zseries/zvse/solutions/>



- e-business Solutions for VSE/ESA SG24-5662
- e-business Connectivity for VSE/ESA SG24-5950
- CICS Transaction Server for VSE/ESA  
CICS Web Support SG24-5997-00
- ***WebSphere Handbook (Connectors to z/OS and VSE) SG24-7042***

z/VSE Contact: [zvse@de.ibm.com](mailto:zvse@de.ibm.com)