



IBM IT Education Services

E52

Wilhelm Mild

IBM Boeblingen Laboratory

VSEESA@de.ibm.com


**Hardware Options with
VSE and ESS**

VSE Technical Conference

November 10 - 12, 2003 | Hilton, Las Vegas, NV

© 2003 IBM Corporation

Enterprise Storage Server



IBM @server. For the next generation of e-business.

Enterprise Storage Server - Overview


Powerful Seascape Architecture



RAID 5 and
RAID 10




Storage sharing among
different platforms:
IBM e-business servers
and non-IBM servers



SSA loops with
faster and greater
capacity disk drives




zSeries
parallel I/O



More powerful
SMP processors
with Turbo option

Point in time
copy solutions



Fault tolerant
system

Disaster Recovery solutions:
synchronous and non-synchronous



High capacity scalable storage
- up to 55.9 TB physical capacity



IBM @server. For the next generation of e-business.

ESS Connectivity Options

Host adapter bays

- 4 bays
- 4 host adapters per bay

64-bit ESCON host adapters


- Up to 32 ESCON links
- 2 ESCON links per host adapter

2 Gb Fibre Channel / FICON host adapters

- Up to 16 Fibre Channel / FICON links
- 1 FICON link per host adapter
- Long wave or short wave
- Auto speed detection - 1 Gb or 2 Gb

SCSI host adapters

- Up to 32 SCSI bus connections
- 2 SCSI ports per host adapter



Adapters can be intermixed

- Any combination of host adapter cards up to a maximum of 16

ESS Cache access

NVS

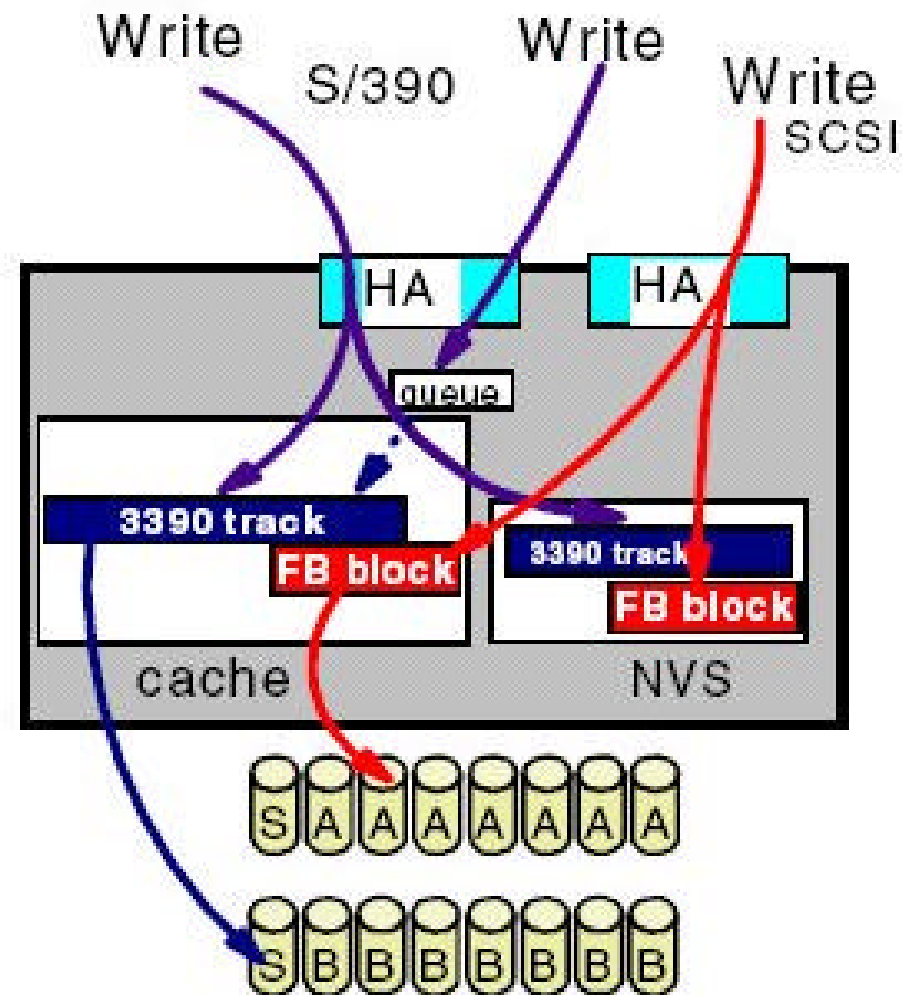
- 192MB/ Cluster
- Battery backed up for 7 days
- 4K segments

100% Fast Write hit

- Data written to NVS first
- I/O complete when data in NVS
- Common to CKD and FB

Destaging

- Managed by LRU
- Idle destage

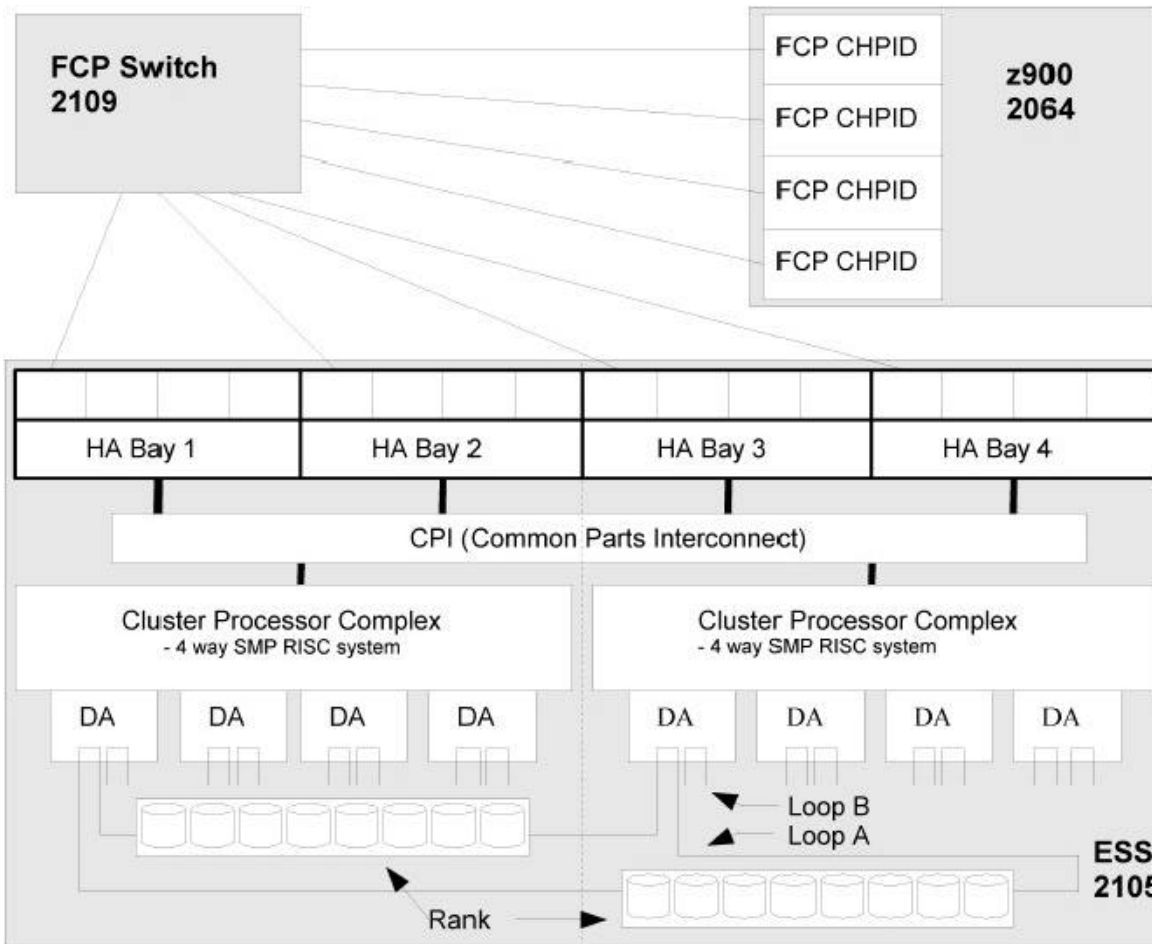


ESS – Enterprise Storage System

- Don't treat ESS as a black box, understand its structure
- The default is close to worst case:
- You ask for 16 disks and your SysAdmin gives you
- addresses 5100-510F
- What's wrong with that?

ESS – Architecture

Let's have a deeper look to the elements of the scenario:



➤ **CHPIDs**

➤ **Host Adapter (HA) supporting FCP (FCP port)**

- 16 Host Adapters, organized in 4 bays

➤ **Device Adapter Pairs (DA)**

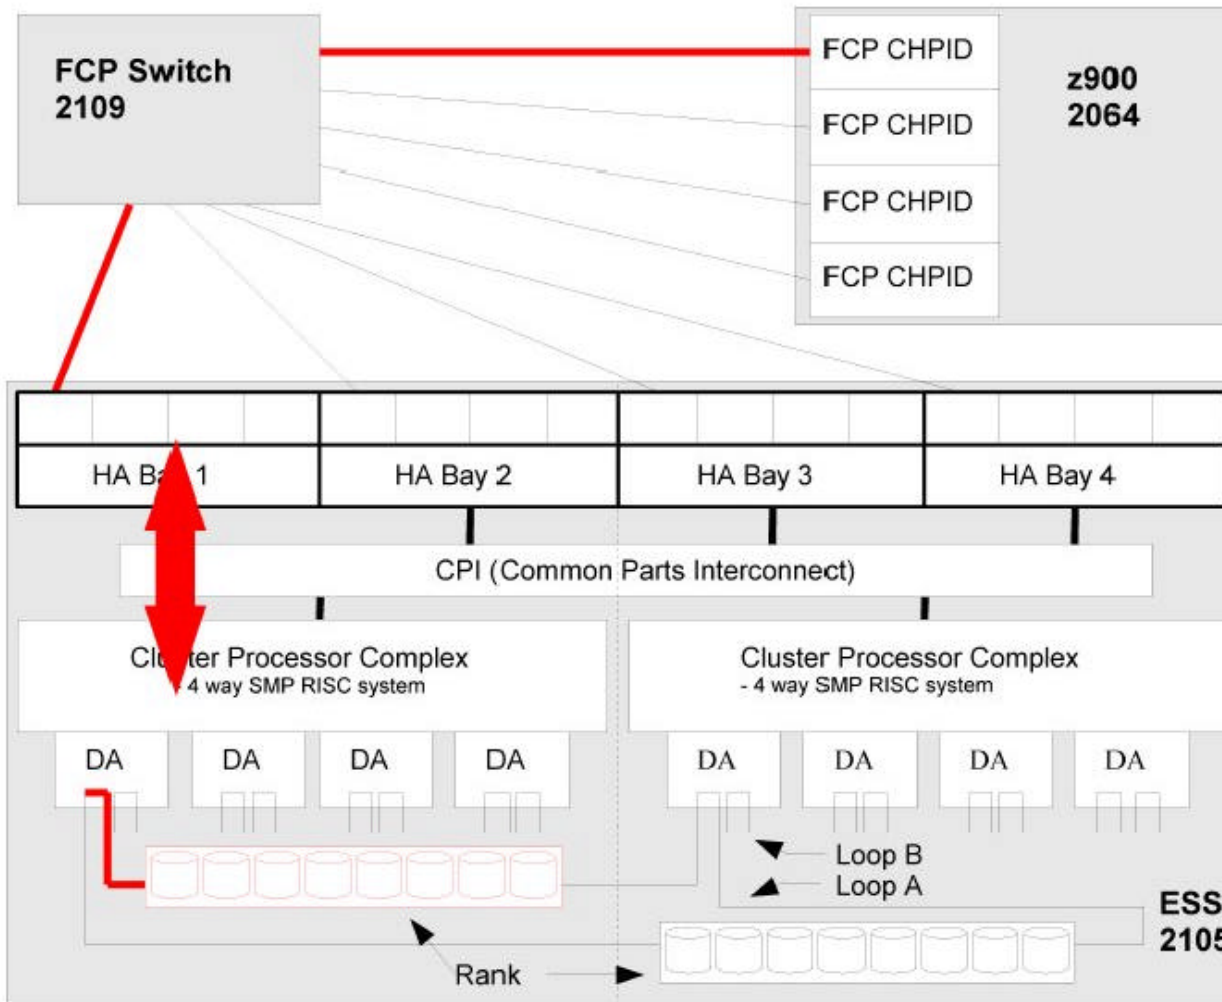
- each one supports two loops

➤ **Disks are organized in ranks**

- each rank (8 physical disks) implement one RAID 5 array (with logical disks)

ESS – Architecture

Scenarios: single disk, single rank



➤ **CHPIDs**

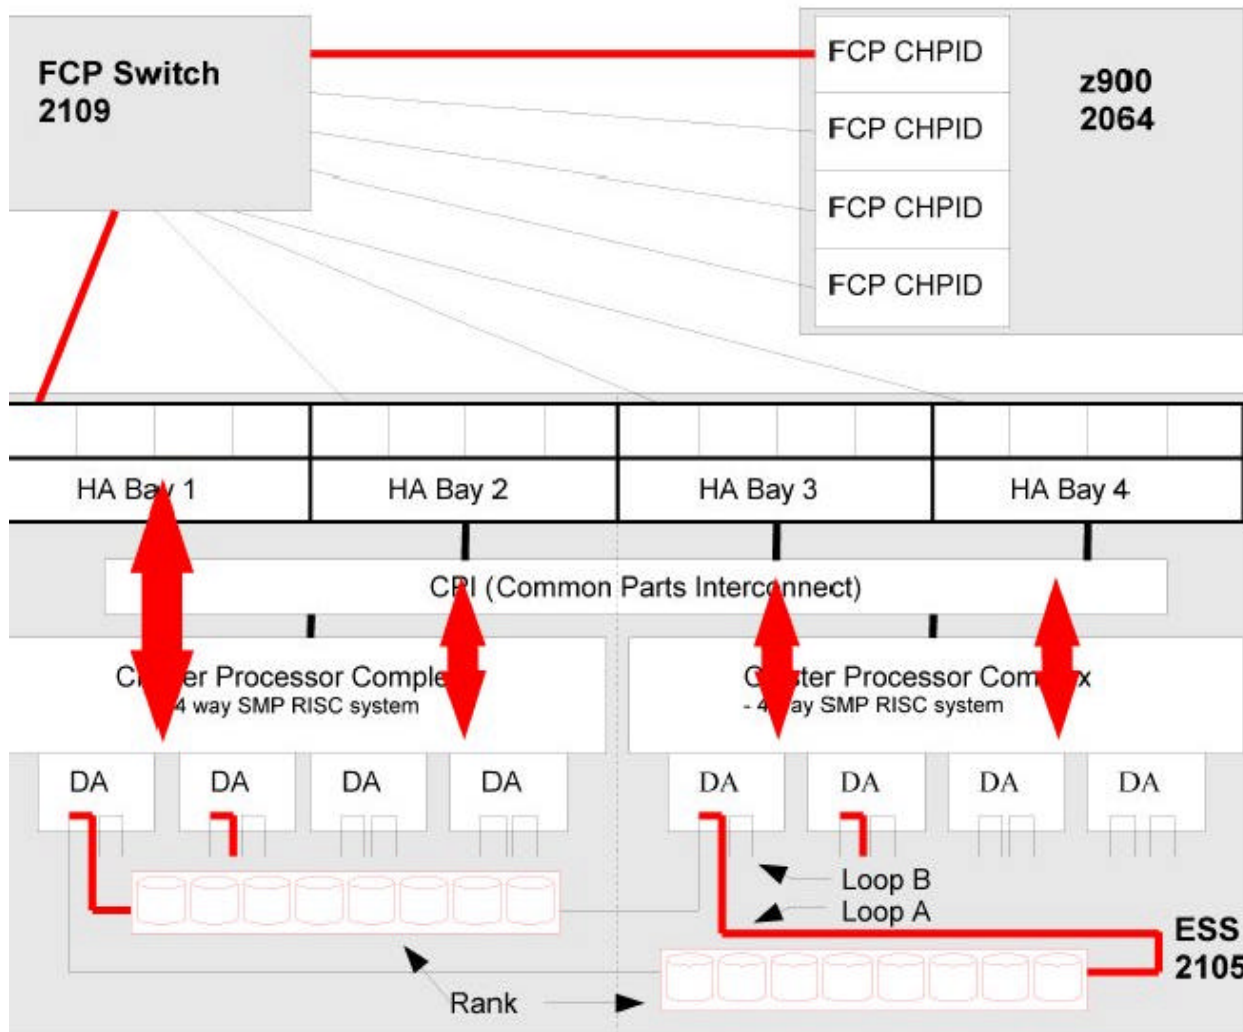
➤ **Host Adapter (HA) supporting FCP (FCP port)**
- 16 Host Adapters, organized in 4 bays, 4 ports each

➤ **Device Adapter Pairs (DA)**
- each one supports two loops

➤ **Disks are organized in ranks**
- each rank (8 physical disks) implement one RAID 5 array (with logical disks)

ESS – Architecture

Scenario: single host adapter



➤ **CHPIDs**

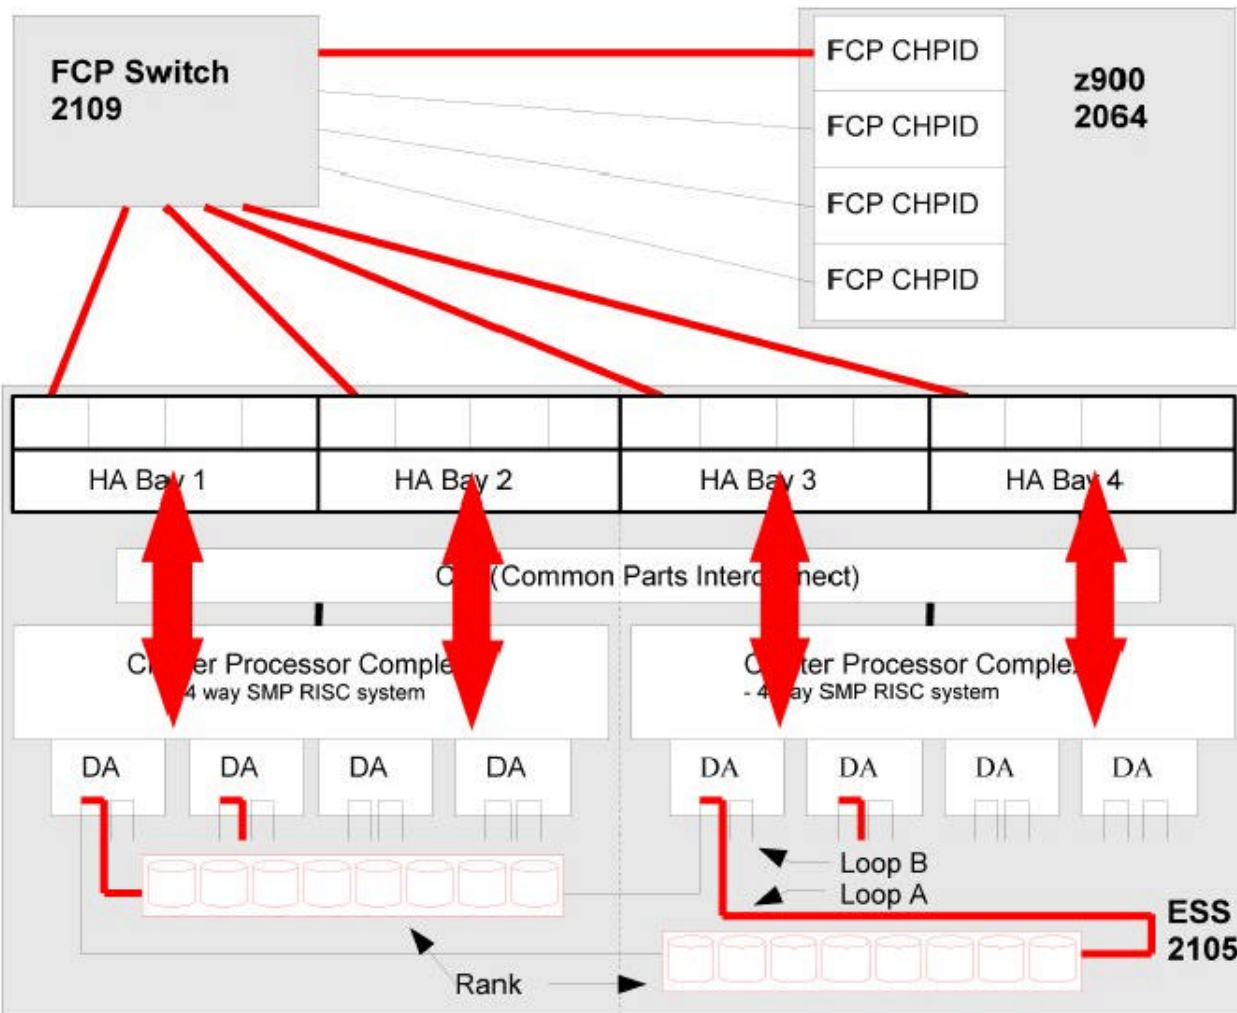
➤ **Host Adapter (HA) supporting FCP (FCP port)**
- 16 Host Adapters, organized in 4 bays, 4 ports each

➤ **Device Adapter Pairs (DA)**
- each one supports two loops

➤ **Disks are organized in ranks**
- each rank (8 physical disks) implements one RAID 5 array (with logical disks)

ESS – Architecture

Scenario: single CHPID



➤ **CHPIDs**

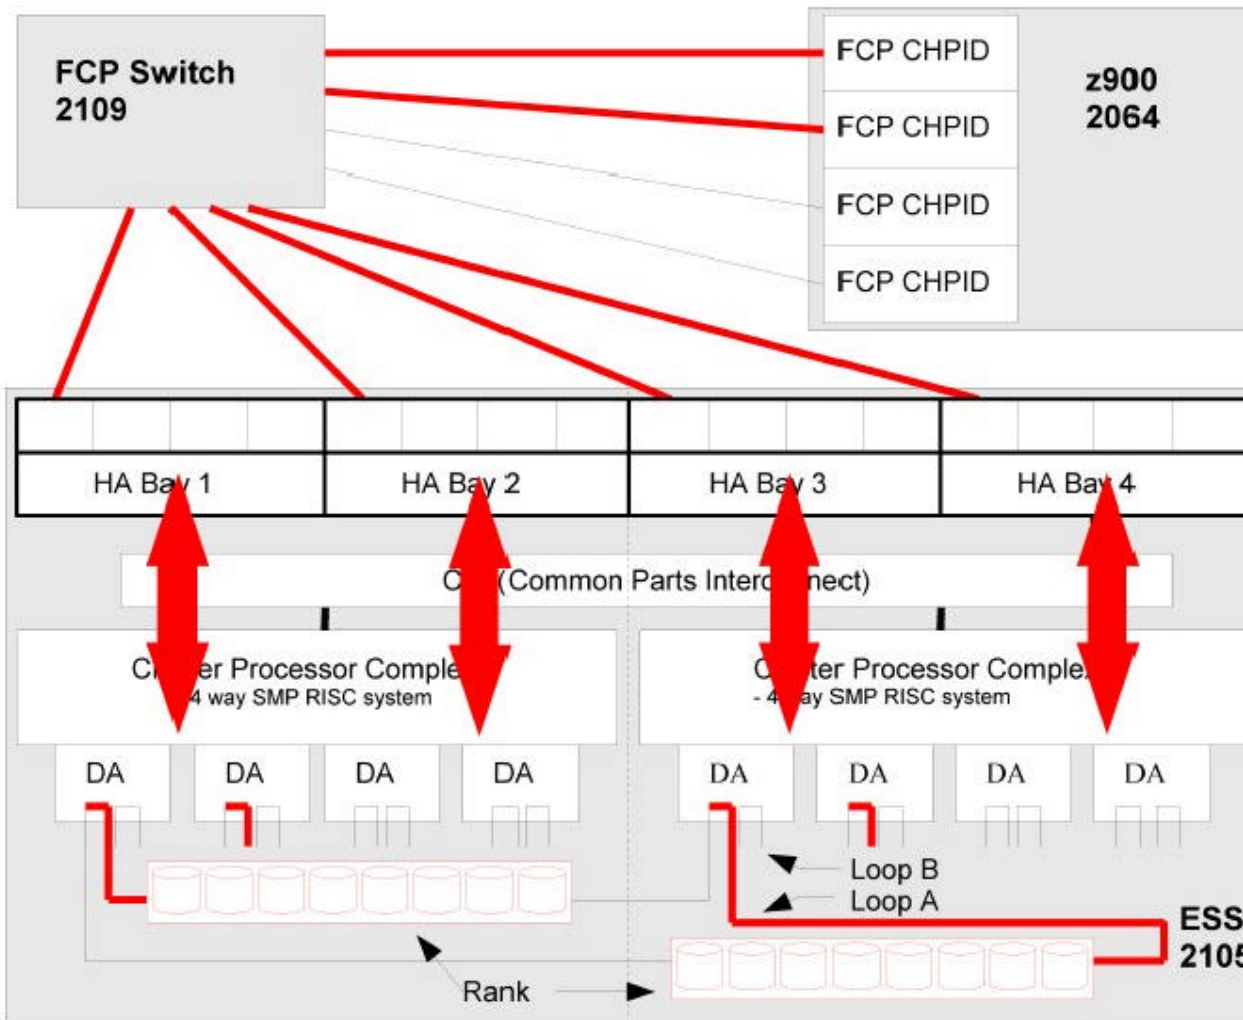
➤ **Host Adapter (HA) supporting FCP (FCP port)**
- 16 Host Adapters, organized in 4 bays, 4 ports each

➤ **Device Adapter Pairs (DA)**
- each one supports two loops

➤ **Disks are organized in ranks**
- each rank (8 physical disks) implement one RAID 5 array (with logical disks)

ESS – Architecture

Scenario: two CHPIDs



➤ **CHPIDs**

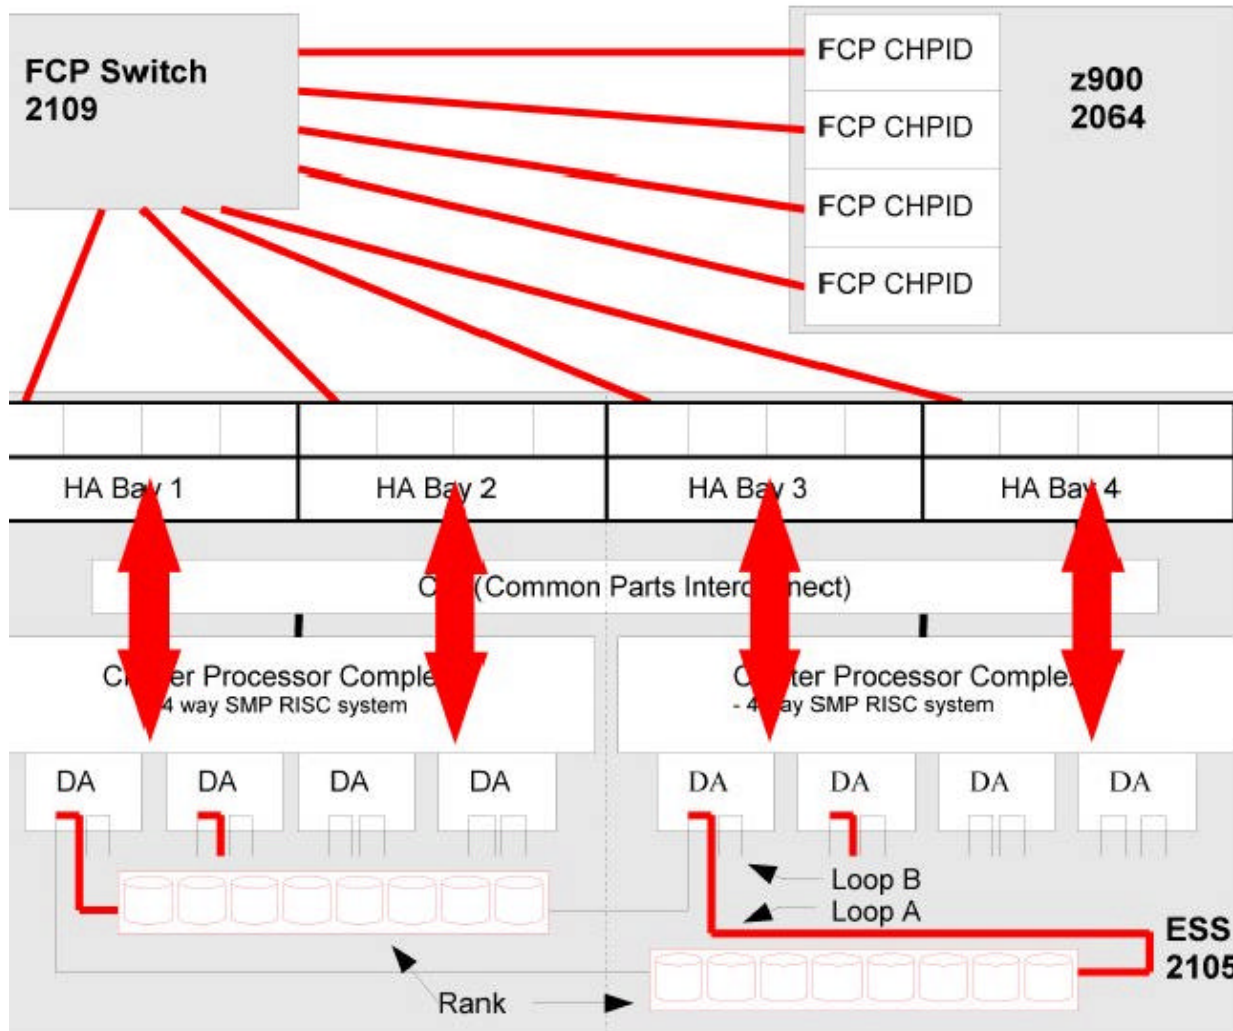
➤ **Host Adapter (HA) supporting FCP (FCP port)**
- 16 Host Adapters, organized in 4 bays, 4 ports each

➤ **Device Adapter Pairs (DA)**
- each one supports two loops

➤ **Disks are organized in ranks**
- each rank (8 physical disks) implements one RAID 5 array (with logical disks)

ESS – Architecture

Scenario: four CHPIDs (4C4H4R ESS 2105)



➤ **CHPIDs**

➤ **Host Adapter (HA) supporting FCP (FCP port)**
 -16 Host Adapters, organized in 4 bays, 4 ports each

➤ **Device Adapter Pairs (DA)**
 - each one supports two loops

➤ **Disks are organized in ranks**
 - each rank (8 physical disks) implements one RAID 5 array (with logical disks)

Hardware Setup for ESS tests

- 2064-216, 917 Mhz, 256MB LPAR
 - 4 FICON Express channels used for FCP (IOCDS: type FCP)
 - 6 FICON Express Channels used for FICON (IOCDS: type FC)
- 2109-F16 FCP switch
- ESS 2105-F20:
 - 16GB cache, 4 FCP host adapters, 6 FICON host adapters
 - 4 device adapter pairs
 - only A-loops contain disks (36.4 GB, 10,000 RPM):
 - - 4 ranks for FB (fixed block) disks used for FCP
 - - 8 ranks ECKD disks used for FICON measurements

FCP Measurement

- Summary of the Scenarios:


| Scenario | used resources | | | | limiting resource |
|--|----------------|----|-------|-------|-------------------|
| | CHPIDs | HA | Ranks | Disks | |
| single Disk | 1 | 1 | 1 | 1 | 1 host adapter |
| single Rank | 1 | 1 | 1 | 8 | 1 host adapter |
| single Host Adapter | 1 | 1 | 4 | 8 | 1 host adapter |
| single CHPID | 1 | 4 | 4 | 16 | 1 CHPID |
| two CHPIDs | 2 | 4 | 4 | 16 | 2 CHPIDs |
| maximum available = 4C4H4R ESS 2105 | 4 | 4 | 4 | 16 | 4 host adapters |

- Benchmark used for measuring: **iozone** (<http://www.iozone.org>)
 - multi process sequential file system I/O
 - each process writes and reads a 350 MB file on a separate disk
 - System: LPAR, 4 CPUs, 128 MB main memory, Linux 2.4.17 with hz timer of
- scaling was: 1, 2, 4, 8, 16 processes
the maximum throughput values were taken as result

IBM @server. For the next generation of e-business.

ESS – Architecture

Results – Maximum Throughput




- 1 HA limits to 40MB/s write and 65 MB/s read, regardless of the number of ranks
- 4 HA are limiting to 125 MB/s write and 240 MB/s read, but 4 CHPIDs are required to make use of it
- 31 bit and 64 bit difference is small
- it is expected that the values further increase using more ranks, HA, CHPIDs

IBM @server. For the next generation of e-business.

ESS – Architecture

- this makes it **slow**:
 - ❖ when all disks are from one rank and accessed via the same path
- this makes it **fast**:
 - ✓ use many host adapters
 - ✓ spread the host adapters used across all host adapter bays
 - ✓ use as much CHPIDs as possible and
 - ✓ access each disk through all CHPIDs, if possible
(FICON, LVM1-mp)
 - ✓ don't use more than two HAs per CHPID
 - ✓ spread the disks used over all ranks equally
 - ✓ avoid using the same resources (CHPID, host adapter, rank)
as long as possible
- These apply to FCP and FICON

ESS Copy Functions



IBM @server. For the next generation of e-business.

ESS Copy Functions

With the ESS we have copy functions for copying data either:

- ❖ remotely, (that is copying data to a separate remote-site ESS)
- ❖ locally, (that is within the same on-site ESS)

The ESS Copy Services can also be classified as:

- dynamic

Dynamic copy functions constantly update the secondary copy as applications make changes to the primary data source.

- point-in-time

Point-in-time copy functions provide an instantaneous copy or view of what the original data looked like at a specific point in time.

ESS Copy Functions

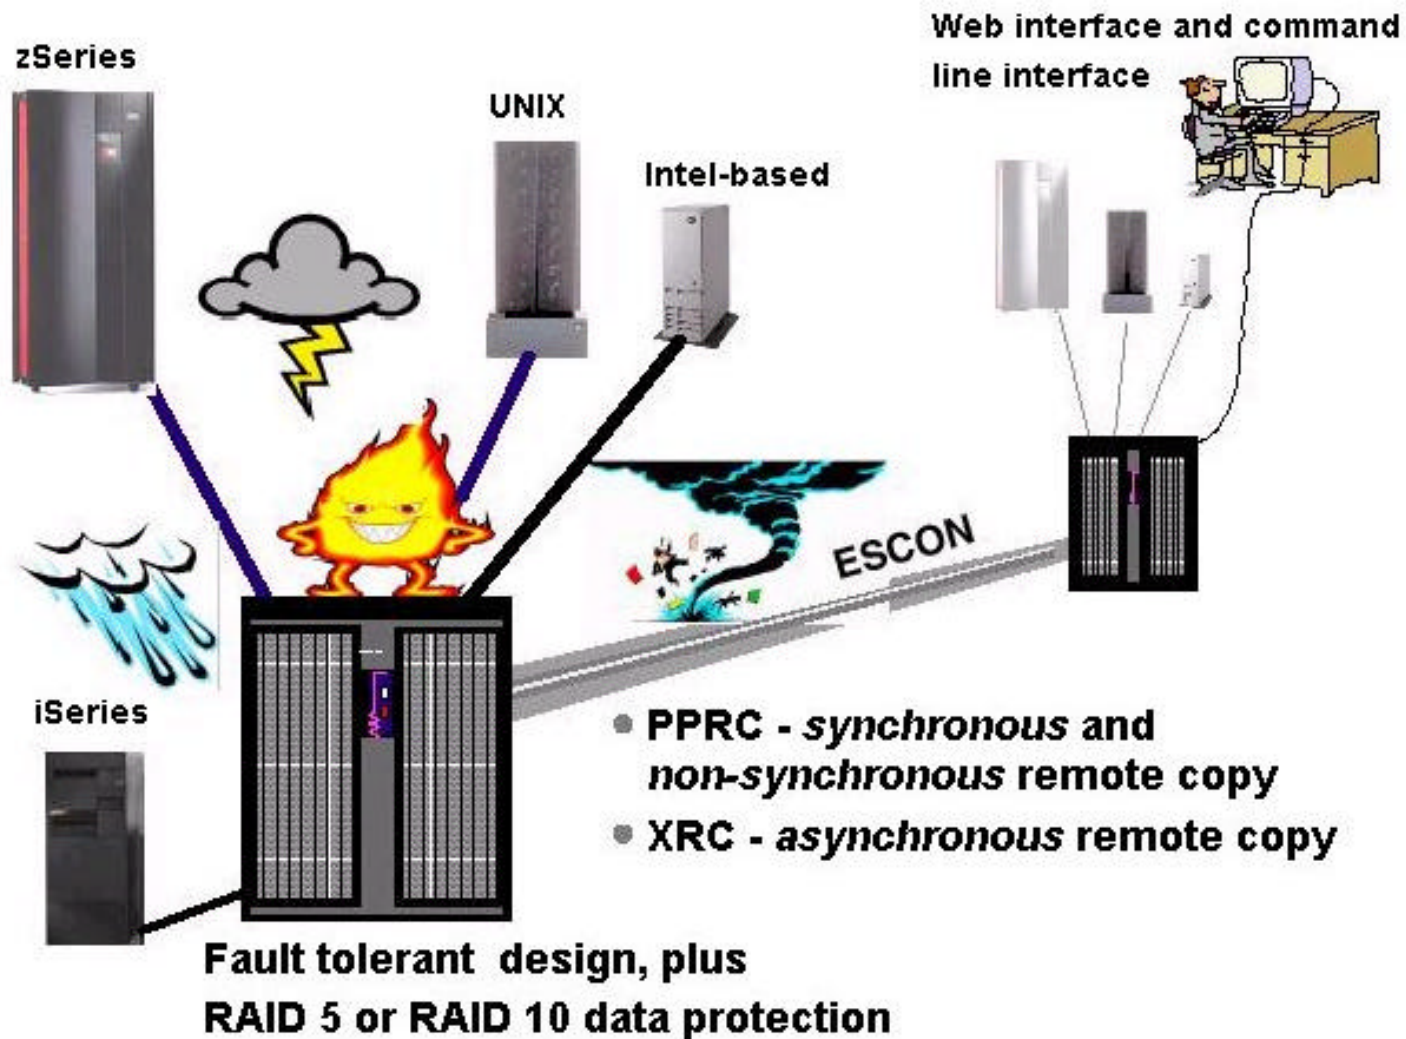
Remote copy category:

- Extended remote copy (XRC)
 - based on the powerful SDM software that runs on the zSeries servers with z/OS
 - asynchronous (not affected by any distance implications)

- Peer-to-peer remote copy (PPRC)
 - designed to constantly maintain a current copy of the local application data at a remote site, in a synchronous manner.
 - it is an ESS hardware-based mirroring solution

- Peer-to-peer remote copy Extended Distance (PPRC-XD)
 - non-synchronous implementation of PPRC.
 - it is an ESS hardware-based solution for remote backup copy and data migration
 - can be used over very long distances.

ESS Copy Functions



IBM @server. For the next generation of e-business.

ESS Copy Functions (2)

Local copy category:

– Concurrent Copy

- offers another method for creating a point-in-time copy in zSeries and S/390 environments with z/OS SDM function (system data mover),
- with the source data fully available for access and update after initiation of the copy operation.
- Concurrent Copy (CC) is the combined function of the ESS either a real instant copy or a Flashcopy will be performed (if possible).

– FlashCopy

- provides a point-in-time copy for the ESS volumes
- creates a physical point-in-time copy of the data, with minimal interruption to applications,
- makes it possible to access both the source and target copies immediately
- optional feature on the ESS for the open system servers, zSeries and iSeries servers.

FlashCopy – the difference

FlashCopy Version 1

- original implementation of FlashCopy in the ESS
- is working at volume level only
- source and target volumes needed to be in the same ESS logical subsystem (LSS).
- a source and a target volume could only be involved in one FlashCopy relationship at a time.
- FlashCopy Version 1 is available as an optional feature of the ESS.

FlashCopy Version 2

- Version 1, plus the following enhancements:
- In addition to volume copy, extent FlashCopy is also possible
- Multiple FlashCopy relationships can be present on a volume and extent level.
- When doing extent FlashCopy, target tracks can be on the same volume as the source tracks, or even in another location.
- Source and target volumes can be on the same or different LSSs.
- Incremental copies of established FlashCopy relationships can be done.
- Consistency groups can be enabled when establishing FlashCopy relationships.
- Current FlashCopy Version 1 users can upgrade to Version 2, provided they install ESS Licensed Internal Code (LIC) 2.2.0 or later

Storage requirements for today's IT

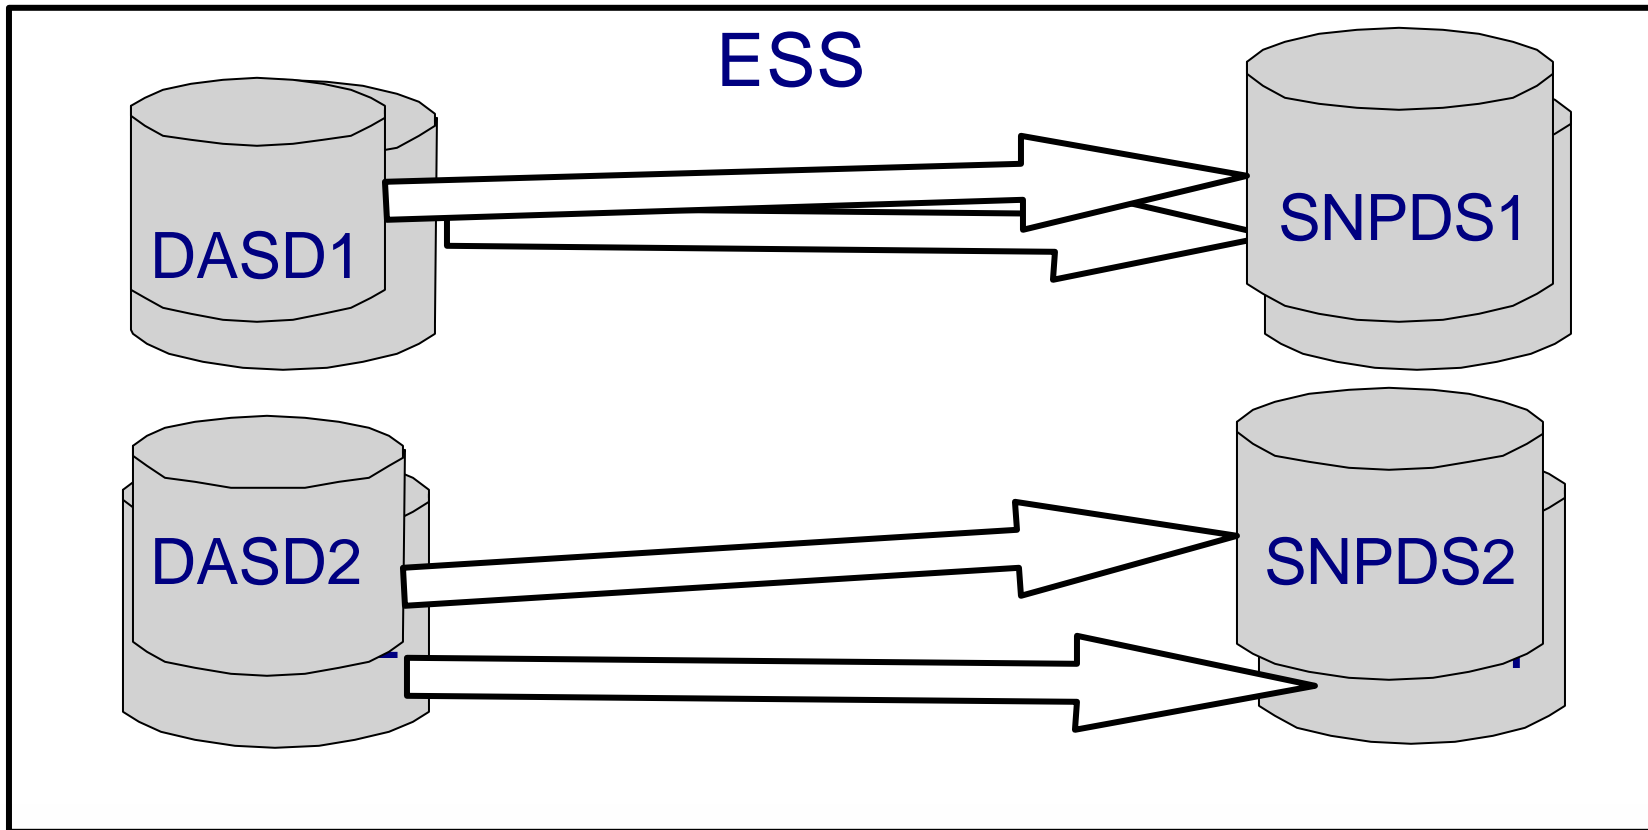
Today, data processing centers increasingly require applications to be available on-line

- ✓ 24 hours per day,
- ✓ seven days per week
- ✓ (24x7)
- ✓ without compromising recoverability in the event of a failure.

FlashCopy allows you to move towards such solutions.

Flashcopy - logical copy of DASD's

FlashCopy is an ESS Copy Services function that can quickly copy data from source volumes to target volumes that are within the same ESS logical subsystem (LSS).



Flashcopy functions

The ESS Copy Services Web user interface— the ESS establishes a FlashCopy relationship between the target volumes and the source volumes.

Establish phase

- establish of FlashCopy relationships (ESS microcode builds the metadata)
- allows the ESS to have target volumes that reflect the source volumes (T0 copy),
- Very short time (establish relation) - after that FlashCopy is said to be complete.
 - logically the target volume has been created,
- once FlashCopy (relationship) is completed, the source and target volumes are immediately available for applications use.
- A source and a target volume can be involved in only one FlashCopy relationship at a time.
- establishing a FlashCopy relationship,
- the execution of the background copy is optional

Flashcopy phases

Copy Phase

- physical copy phase begins (after establishing the relationship between source and target volumes)
- the FlashCopy relationship will end
 - automatically
 - when the physical background copy task completes
 - when all the tracks on the source or the target volume have been updated by a host
 - 'NO background copy' was specified
 - when the FlashCopy relationship is explicitly removed
 - by command
 - using the ESS Copy Services Web User Interface

Note: When background copy is requested with the Persistent FlashCopy option, then the relationship is not automatically removed when the physical background copy completes —it must be explicitly removed

IBM @server. For the next generation of e-business.

FlashCopy Phases

Copy Phase (2)

- background physical copy of the volumes is optional
 - background copy specified (default)
 - - the background copy task copies data from the source to the target volume
 - - minimal impact on normal application I/O processing.
 - - application reads some data
 - - from the target that has not yet been copied to the target volume,
 - - the data is read from the source
 - - otherwise, the read can be satisfied from the target volume.
 - - updating a not-yet-copied track on the source volume,
 - an on-demand copy of the track to the target volume is done
 - no delay at all to the host I/O operation (asynchronous)
 - NOCOPY specified
 - only updated tracks from source volume are copied to the target.
 - no delay to the host I/O operation (asynchronous)
 - the T0 copy of the source volume will be available for use as long as the source-target FlashCopy relationship exists.

FlashCopy Phases

Copy Phase (2)

- background physical copy of the volumes is optional
 - background copy specified (default)
 - - the background copy task copies data from the source to the target volume
 - - minimal impact on normal application I/O processing.
 - - application reads some data
 - - from the target that has not yet been copied to the target volume,
 - - the data is read from the source
 - - otherwise, the read can be satisfied from the target volume.
 - - updating a not-yet-copied track on the source volume,
 - an on-demand copy of the track to the target volume is done
 - no delay at all to the host I/O operation (asynchronous)
 - NOCOPY specified
 - only updated tracks from source volume are copied to the target.
 - no delay to the host I/O operation (asynchronous)
 - the T0 copy of the source volume will be available for use as long as the source-target FlashCopy relationship exists.

FlashCopy Phases

Copy Phase (2)

➤DDSR

Delete Data Space Request is a function that ensures that the 'logical' deletion of a dataset by the VSE system results in the 'physical' deletion of the data on the subsystem. DDSR, if specified without any additional operand allows the operator to delete the VTOC entries and all physical space for all non-secured files residing on VSE managed RVA devices whose associated expiration date has been reached and which have been created by VSE.

ESS Copy Services Web User Interface



IBM @server. For the next generation of e-business.

Flashcopy support in zSeries OS's

❖ FlashCopy Invocation in z/OS and OS/390

❖ FlashCopy invocation in z/OS and OS/390 can be done in four different ways:

- ❖ DFSMSDss utility
- ❖ TSO commands
- ❖ ANTRQST API interface
- ❖ ESS Copy Services Web user interface

❖ VSE/ESA support

❖ The FlashCopy support for ESS is part of VSE/ESA 2.5.

❖ FlashCopy support for VSE is provided by three products:

- ❖ IXFP SNAP command.
- ❖ VSAM SYNONYM command. This has not been retro-fitted.
- ❖ VSE/Fast Copy.
- ❖ The ESS Copy Services Web user interface

Flashcopy support in zSeries OS's

VM support

CP support

- FlashCopy provides the ability to make copies between any mix of
- Log Structured Array (LSA) and non-LSA (Home Address) source and target DASD.
- FlashCopy support includes the CP FLASHCOPY command
(copy a virtual volume to another virtual volume in the same LSS).
- The format is:

FLASHCopy vdev1 scyl1 ecyl1 TO vdev2 scyl2 ecyl2

For example for a whole volume as defined in VM:

FLASHCOPY F100 0 END to F109 0 END

Note:

The virtual volumes must reside on the same ESS

The target volume must be in write mode

If the whole volume is to be copied then use END for the ending cylinder number

If cylinder 0 is copied, the target device may become a duplicate of the source device

Whole cylinders must be copied —track extents are not allowed

The default 'background copy' option is used — 'no background copy' cannot be used

Flashcopy support in zSeries OS's

VM Guest System support

VM/ESA and z/VM provide FlashCopy support for guest systems for

- attached volumes
- for full-pack minidisks.

VSE

- VSE FlashCopy can be run either using IXFP Commands or the ESS Specialist.

MVS and z/OS

- The target should be offline at the completion of the copy and
- then it can be varied online by the operator.
- The source and target volumes for the copy must be in the guest's virtual configuration.
- the guest must have write authority to the target
- if the target is reserved, the guest must hold the reserve.

TPF

- Flashcopy can be run from the ESS Copy Services WUI or, if a
- from the ESS Copy Services Command Line interface (CLI) – for supported server .

Flashcopy Requirements

- ▲ Source and the target volumes must be in the same ESS logical subsystem (LSS).
- ▲ Source and the target volumes must have the same track format.
- ▲ The target volume must be at least as large as the source volume.
- ▲ A volume can be in only one FlashCopy relationship at a time.
- ▲ There is a limit of 4096 volumes —or 2048 relationships— per ESS Copy Services server when using the Web user interface. (for all FlashCopy, PPRC and PPRC-XD pairs)
 - ▲ This limit does not apply to host command invocations (i.e zSeries Operating system commands)
- ▲ The FlashCopy feature is an optional ESS feature that must be enabled.

FlashCopy reduces Online downtime

x inhibitors of online processing time

☞ backup-window

☞ batch-window



✓ Eliminate the Backup-window

VSAM backup using
FlashCopy (ESS)

What is "FlashCopy" for VSE

- ▶ The DASD architecture of Enterprise Storage Server (ESS) allow copy of DASD's with the utility "*FlashCopy*" .
- ▶ The COPY process takes few seconds instead of hours !
- ▶ From OP system view the copy is a real copy of data.
- ▶ From the DASD controller view it is a virtual copy of data.

VSAM-Restrictions in a VSE system

- ▶ Duplicate VOLIDs (DASD names) not allowed on a VSE System !
- ▶ Duplicate VSAM Catalog names not allowed on a VSE System !

FlashCopy for VSAM Datasets
would mean:

- ▶ duplicate VOLID'S and/or duplicate VSAM Catalog names on the system




unpredictable Results!



Difficulties in using FlashCopy for VSAM Datasets

- ▶ or many changes required in the copied VSAM catalog regarding new VOLID'S (i.e. volume list for each dataset) and Catalog names.



performance


increased error risk

possible lose of data



Support for FlashCopy for VSAM Datasets with VSE/ESA 2.5

1. IDCAMS SNAP Utility

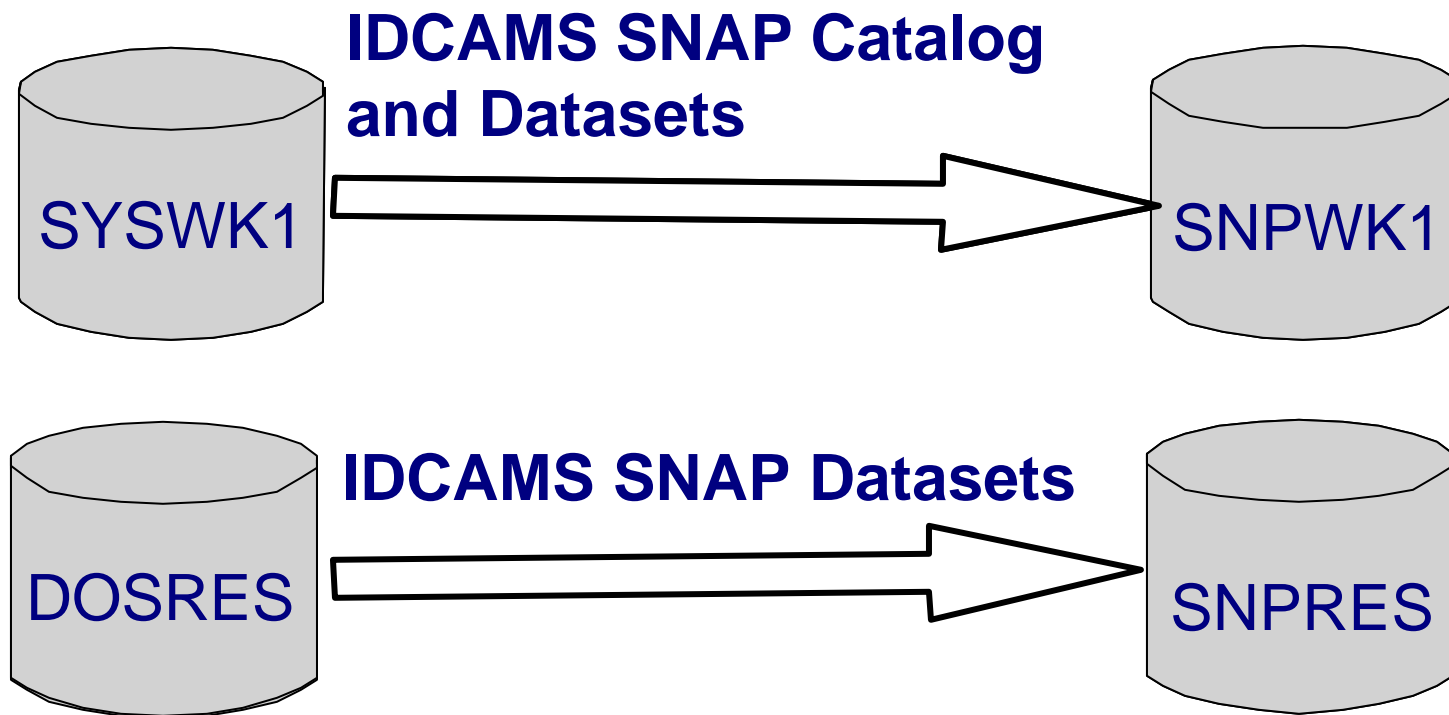


2. IDCAMS "Synonym" BACKUP

1. IDCAMS SNAP Utility Program

- ▶ IDCAMS SNAP makes copies of entire DASD volumes (VSAM catalog- and VSAM data- volumes).
- ▶ IDCAMS SNAP changes the names of the copied volumes (VOLID).
- ▶ After IDCAMS SNAP all copied volumes are **ONLINE** available for backup!

Step 1: IDCAMS SNAP - copy all DASD's and give new Volid's



After Step 1, the DASD's and catalogs copied are identical, but cannot be used.

Few seconds later:

- IDCAMS SNAP finished.
- Online applications can be restarted (CICS).
- The catalog and the datasets on the snapped (copied) volumes are identical with the original volumes (only VOLID's are different),




but the copied datasets can not be used, because

- ▶ Duplicate VOLIDs not allowed on a VSE System !
 - SNAP changed the VOLID'S
- ▶ Duplicate VSAM Catalog names not allowed on a VSE System !

Step 2: IMPORT CONNECT a new catalog name

The catalog on the snaped volume needs a new name.




we simulate a new catalog name with IDCAMS IMPORT CONNECT ,



a synonym catalog name.


Step 3: Backup VSAM datasets from snapped volumes (the VSE system is online)

only "Synonym Backup" can read VSAM data from SNPWK1 and SNPRES !



After Step3: a "normal Backup medium" was created

What is "Synonym Backup"? (1)



a synonym list is used, to redirect VSAM to the snaped volumes (with the synonym catalog name) and execute the BACKUP from the copied Datasets.

- VSAM controls the "synonym connection" to the snaped (renamed) catalog and datasets.
- Only "Synonym Backup" can read the VSAM datasets from the copied volumes.

What is "Synonym Backup"? (2)

- ▶ With the exception of using the new synonym list, the backup process is unchanged.
- ▶ That means, all functions of IDCAMS BACKUP can be used
- ▶ IDCAMS BACKUP produces a normal Backup-Medium for IDCAMS RESTORE.

Backup with synonym list

EXEC IDCAMS

BACKUP (.....) -

SYNONYMLIST -

(SOURCEVOLUMES (SYSWK1 , DOSRES) -

TARGETVOLUMES (SNPWK1 , SNPRES) -

-

CATALOG (VSESP . USER . CATALOG) -

SYNCAT (VSESP . SNAP . CATALOG))

Sample job: SNAP AND VSAM BACKUP

```
// JOB SNAP AND BACKUP FROM SNAPPED VOLUMES
// ASSGN SYS005,180
// DLBL IJSYSUC,'VSESP.SNAP.CATALOG',,VSAM
// EXEC IDCAMS,SIZE=AUTO
/* STEP 1: DO THE SNAPSHOT */ -
  SNAP -
    SOURCEVOLUMES(SYSWK1,DOSRES) -
      TARGETVOLUMES(SNPWK1,SNPRES)
/* AFTER STEP 1 THE ONLINE SYSTEM MAY BE STARTED */
/* STEP 2: SYNONYM NAME FOR THE SNAPPED CATALOG */-
  IMPORT CONNECT OBJECTS((VSESP.SNAP.CATALOG -
    VOLUMES(SNPWK1) DEVT(3390))) -
    CATALOG(VSAM.MASTER.CATALOG)
/* STEP 3: BACKUP FROM SNAPPED VOLUMES */ -
  BACKUP (*) -
    SYNONYMLIST( -
      SOURCEVOLUMES(SYSWK1,DOSRES) -
      TARGETVOLUMES(SNPWK1,SNPRES) -
      CATALOG(VSESP.USER.CATALOG) -
      SYNONYMCATALOG(VSESP.SNAP.CATALOG) )
/*
```


/& **IBM** @server. For the next generation of e-business.

Conclusion FlashCopy

Steps for online VSAM Backup using FlashCopy

- ▶ **Close online applications (shutdown CICS)**
- ▶ **FlashCopy the DASD's (datasets/databases, catalogs)**
 - ▶ eventually run batch job streams
- ▶ **restart CICS and the online applications**
- ▶ **Backup your VSAM data during Production**

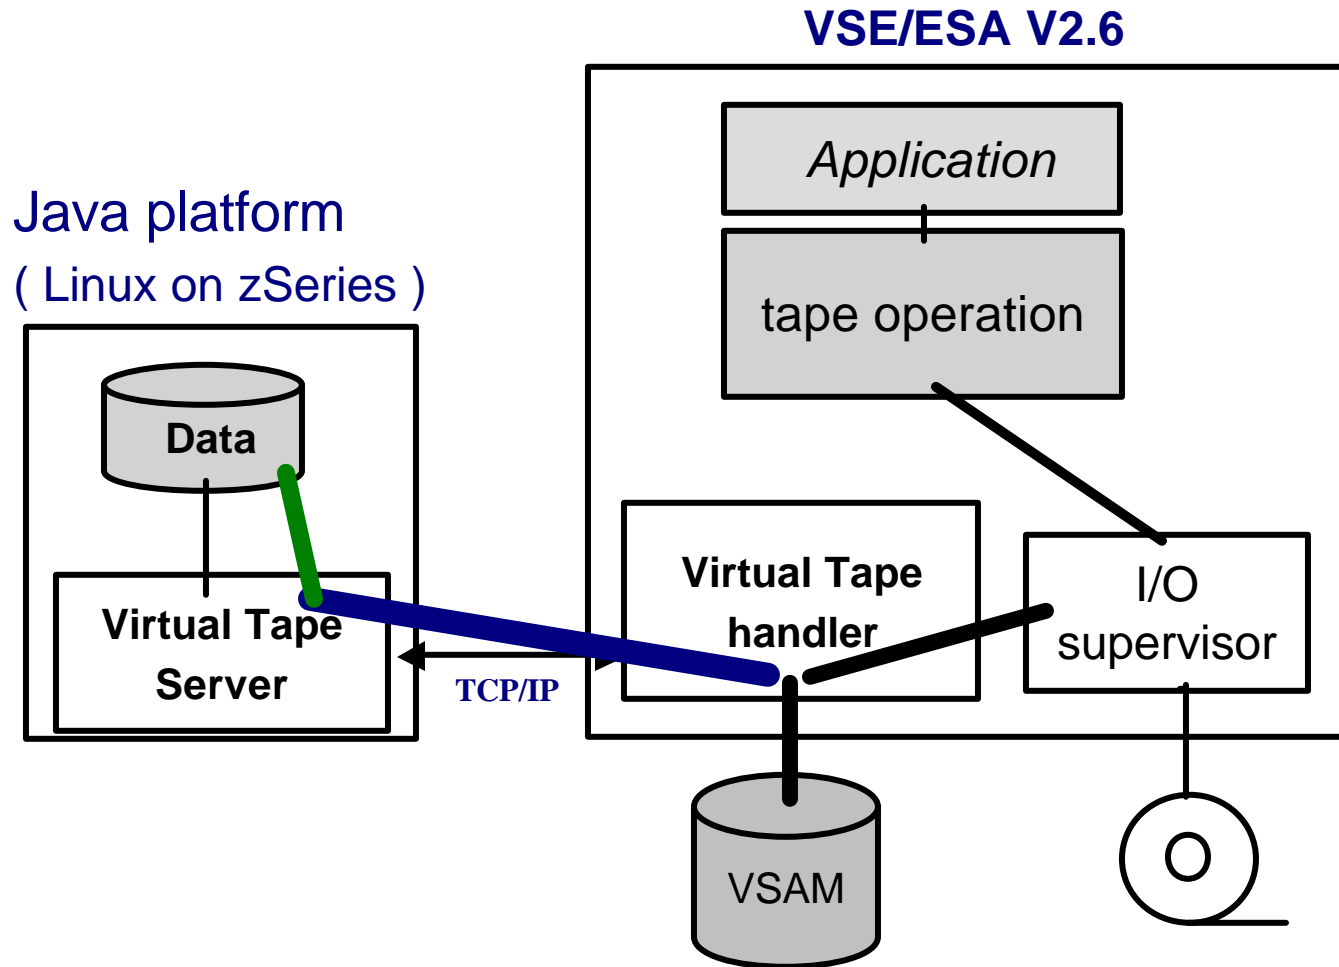
VSAM Backup and Online system in parallel with VSE/ESA 2.5 and 2.6



✓ Integrate VSE Backups in standard processes

Use of VSE/ESA 2.6 Virtual tape support to integrate VSE Backup media into general, automatic Backup processes

Virtual Tape support

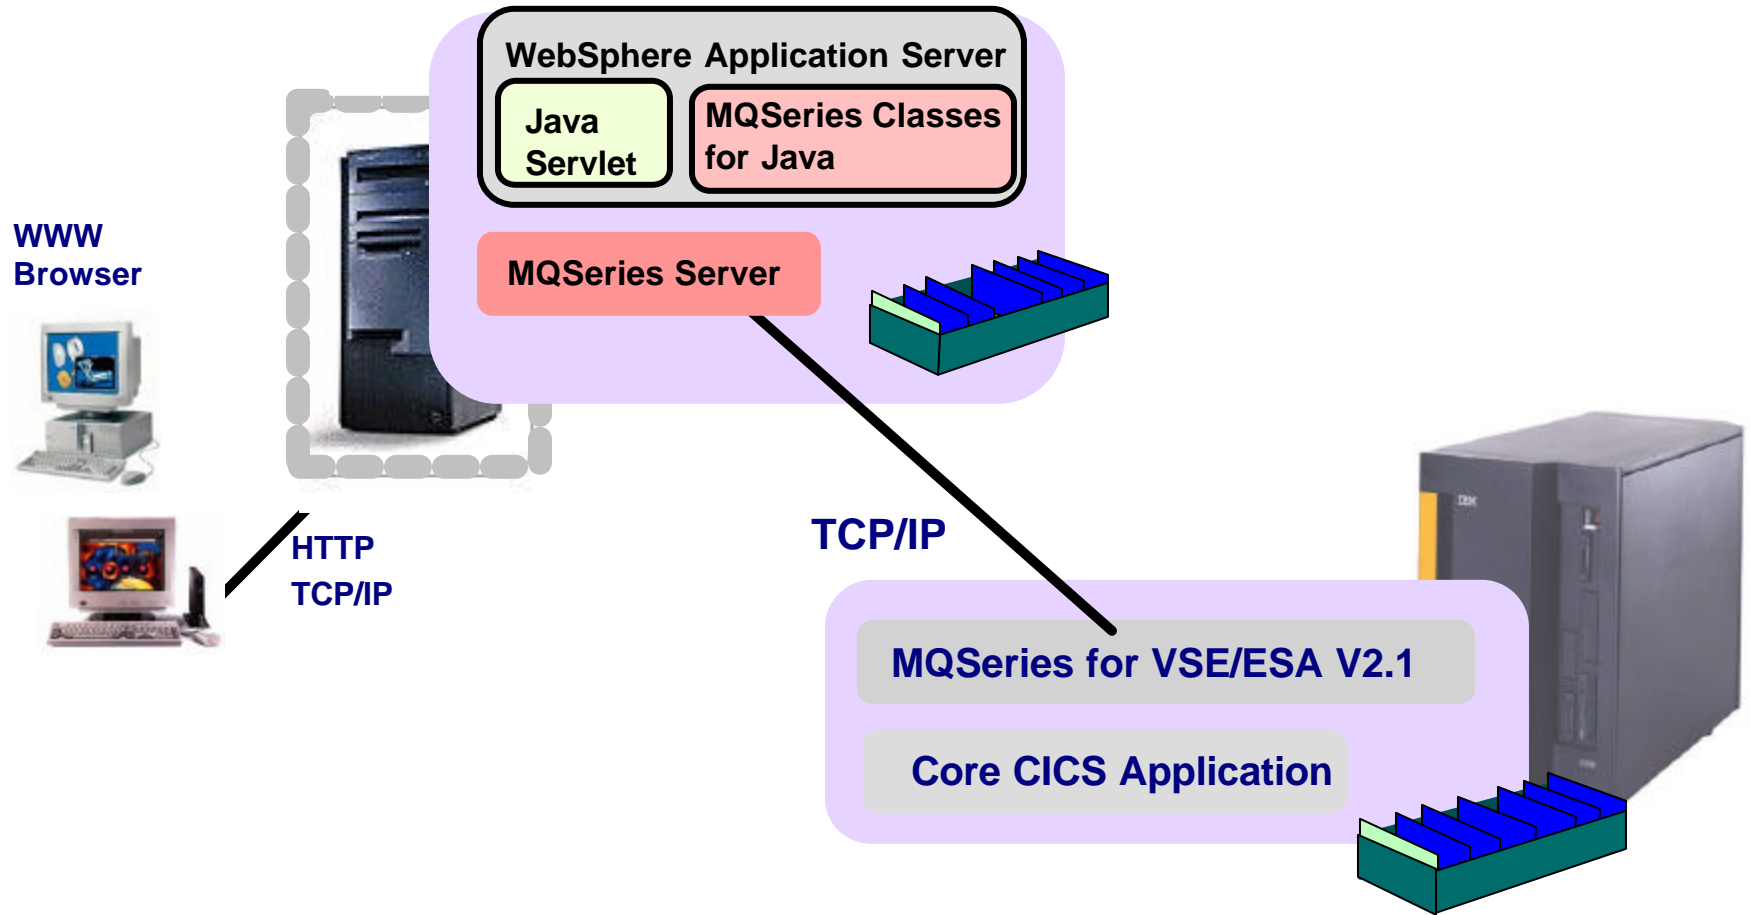


- ▶ simulates a real tape (tape operation supported)
- ▶ transparent for applications

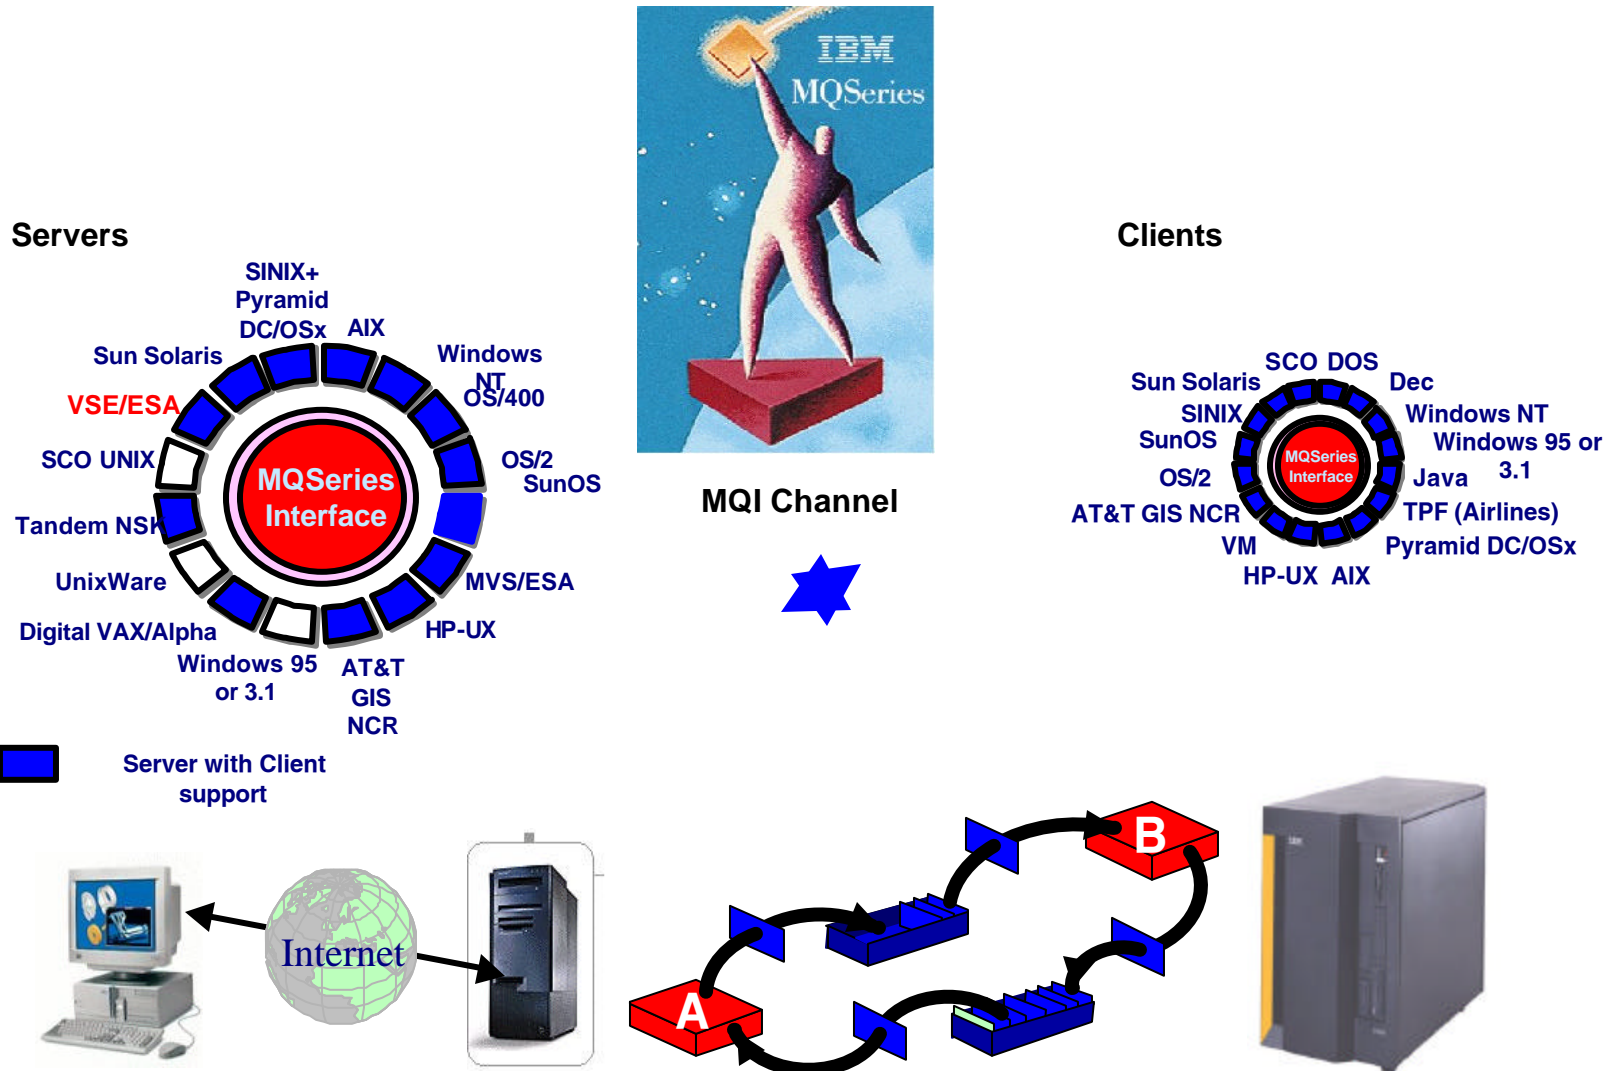
✓ Batch-window solutions

Use of MQ Series and
the new e-business connectors
to avoid Production downtime

Asynchronous work with MQSeries

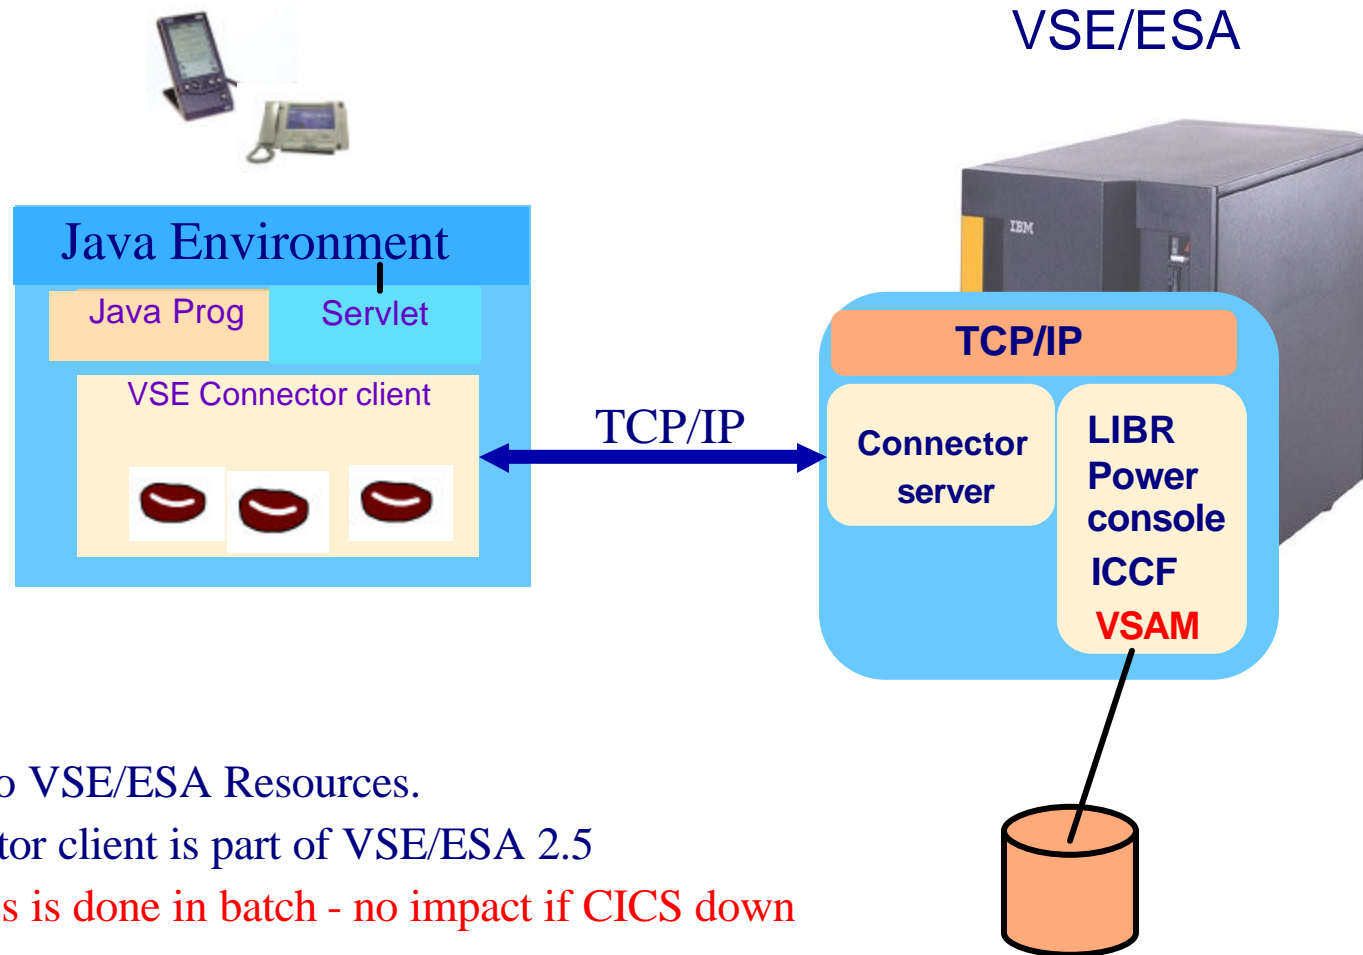


MQSeries and MQSeries Clients




IBM @server. For the next generation of e-business.

Java-based Connector



- ▶ Java access to VSE/ESA Resources.
- ▶ VSE Connector client is part of VSE/ESA 2.5
- ▶ **VSAM** access is done in batch - no impact if CICS down



IBM @server. For the next generation of e-business.

More Information

IBM Redbooks

- IBM TotalStorage Solutions for Disaster Recovery, SG24-6547
- IBM TotalStorage Enterprise Storage Server Model 800, SG24-6424-01
- Implementing ESS Copy Services on UNIX and Windows NT/2000, SG24-5757

Other resources

- IBM TotalStorage Enterprise Storage Server Web Interface User's Guide, SC26-7448

Web sites

- ESS support information
<http://www.storage.ibm.com/hardsoft/products/ess/supserver.htm>
- ESS Web site
<http://www.storage.ibm.com/hardsoft/products/ess/ess.htm>

VSEESA@de.ibm.com

IBM @server. For the next generation of e-business.