Soluzioni di Business Continuity per l'Information Technology

The New Enterprise Data Center
IBM Information Infrastructure

Stefano Ricci
IT Architect
IBM Italia
Systems & Technology Group
21 Novembre 2008

Agenda

- **Principi e caratteristiche funzionali di una soluzione di Business Continuity**
- **Architettura di una soluzione di BC :**
  - **Tiers della Business Continuity**
  - **Componenti e tematica multi-site**
- **Infrastruttura tecnologica a supporto delle soluzioni BC**
  - **Soluzioni di replica “Storage Based“**
  - **Clustering**
  - **Automazione e scenari di recovery**

IT Business Continuity is more than Recovery

Business Continuity Disciplines ……

High Availability
Fault-tolerant, failure-resistant infrastructure supporting continuous application processing

Continuous Operations
Non-disruptive backups and system maintenance coupled with continuous availability of applications

Disaster Recovery
Protection against unplanned outages such as disasters through reliable, predictable recovery

Protection of critical Business data
Recovery is predictable and reliable

Operations continue after a disaster
Costs are predictable and manageable

IT Business Continuity is more than Recovery

Business Continuity Disciplines ……

*Contrary to popular belief, **hardware** failures account for a minority of system outages*

- Several studies place the proportion between 18% and 45%
- Human error, software error and planned maintenance cause the majority of service outages

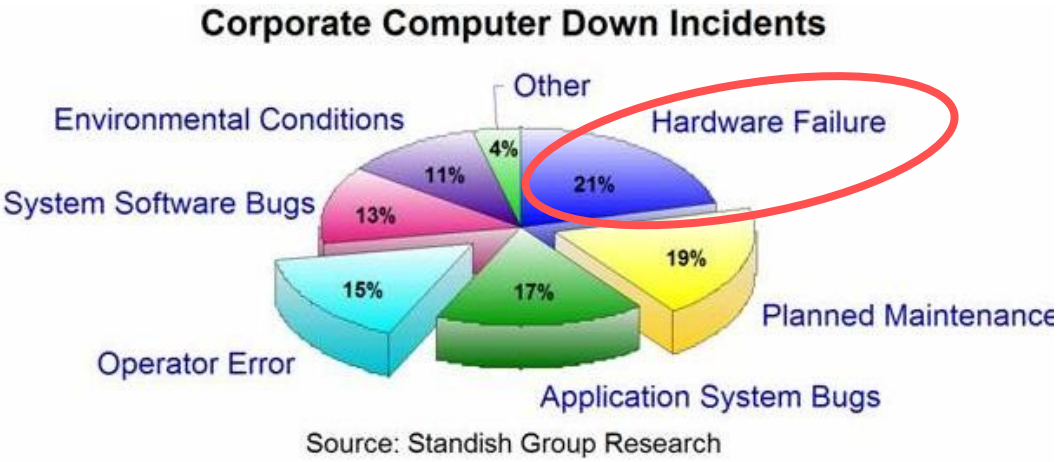

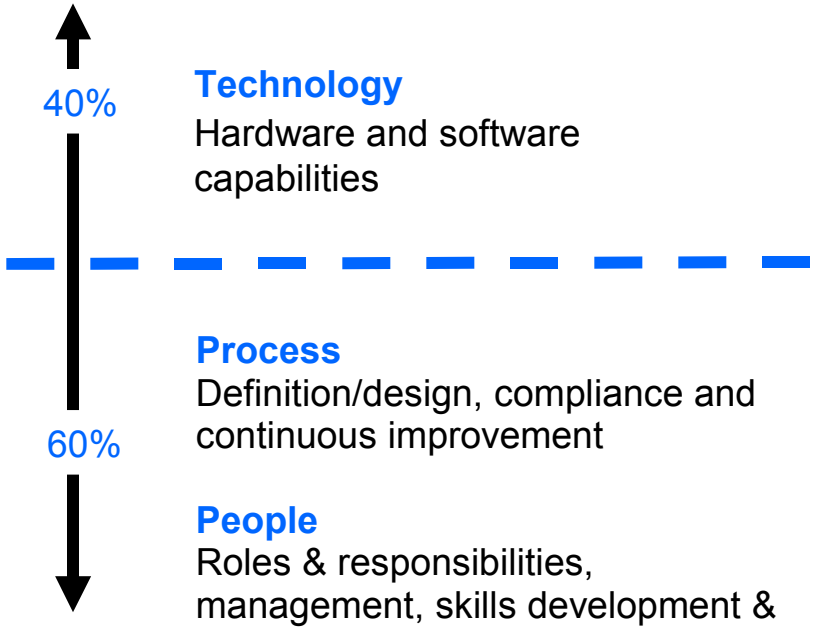

Key Elements of Business Continuance Plan

Strategic planning

Linee guida alla continuità operativa nella Pubblica Amministrazione

STUDIO/ANALISI DEL CONTESTO

Elementi di valutazione di una soluzione di BC:

Data protection

Data availability

Data recoverability

Manageability and Operability

Total Cost of Ownership (TCO)

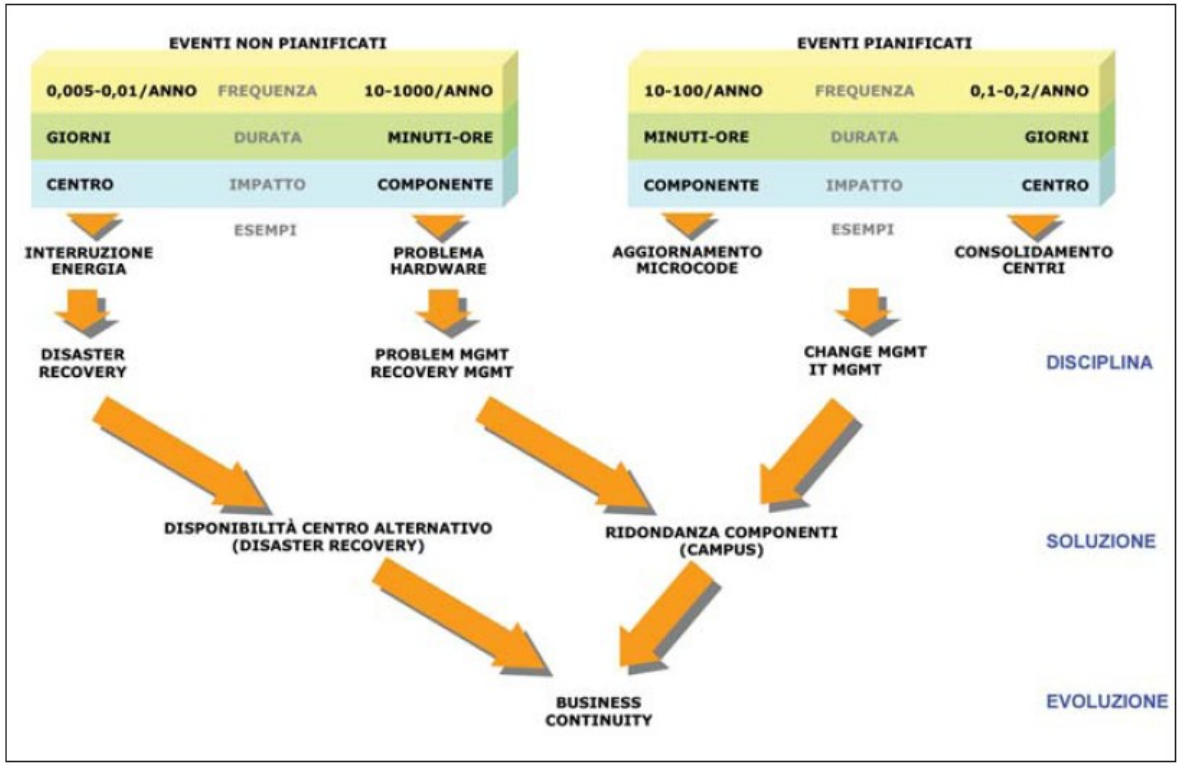

Key Elements of Business Continuance Plan
Metrics

Principal BC/DR metrics :

- **Recovery Time Objective (RTO)** - measured in hours (or minutes).
  - How long can the business tolerate the IT systems being unavailable?
- **Recovery Point Objective (RPO)** - measured in seconds (or minutes).
  - How much data can the business afford to recover (or lose)?
- **Network Recovery Time (NRO)** - measured in minutes.
  - How long will it take to reestablish and / or restart the network?

RTO and RPO considerations :

Determine RTO and RPO as a trade-off between Cost and LoB requirements

- Business Impact Analysis (BIA)
- What Host platforms do you use – zSeries, Open Systems, mixed data
- How many Data Centers - Three site topology required ? Evaluate distance and bandwidht
  - Telecom bandwidth often the largest TCO component (often 50% - 70%) on a three year basis

Timeline of an IT Recovery

Mirroring provides a *point in time, data consistent copy of data* at the remote site.

Recovery Point Objective (RPO). How much data must be recreated?

RPO | Assess | Execute hdw., o/s, and data integrity recovery

Production Site

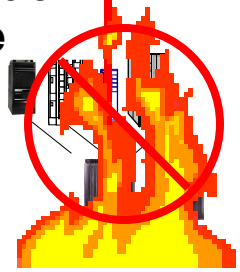

Outage!

✓ Management Control

Recovery Site

✓ Data

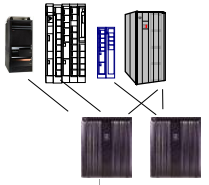

✓ Operating System

✓ Telecom Network

✓ Physical Facilities

✓ Operations Staff

✓ Applications Staff

✓ Network Staff

Recovery Time Objective (RTO) of hardware data integrity.

RPO | Software transaction integrity recovery

Δ Data

✓ Applications

Recovery Time Objective (RTO) of transaction integrity.

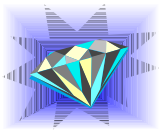

Now we're done!

Categories of IT Components that must be recovered

A comprehensive architecture with the blending of five IT components, resulting in *best* solution:

- **Servers**
- **Storage**
- **Software and Automation**
- **Infrastructure and Networking**
- **Skills and Services**

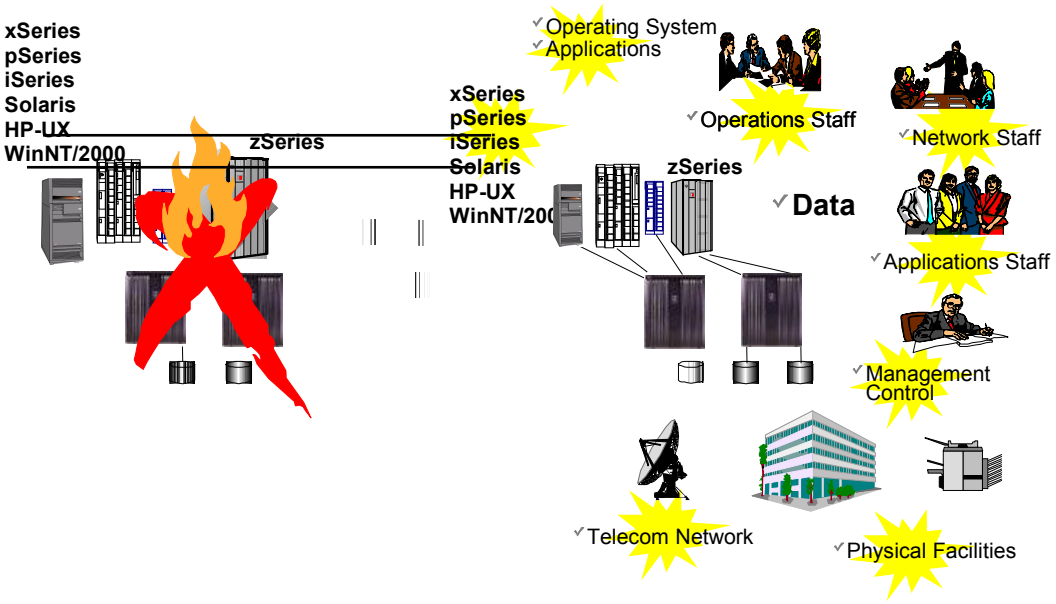

Business Continuity Tiers

Balancing recovery time objective with cost / value

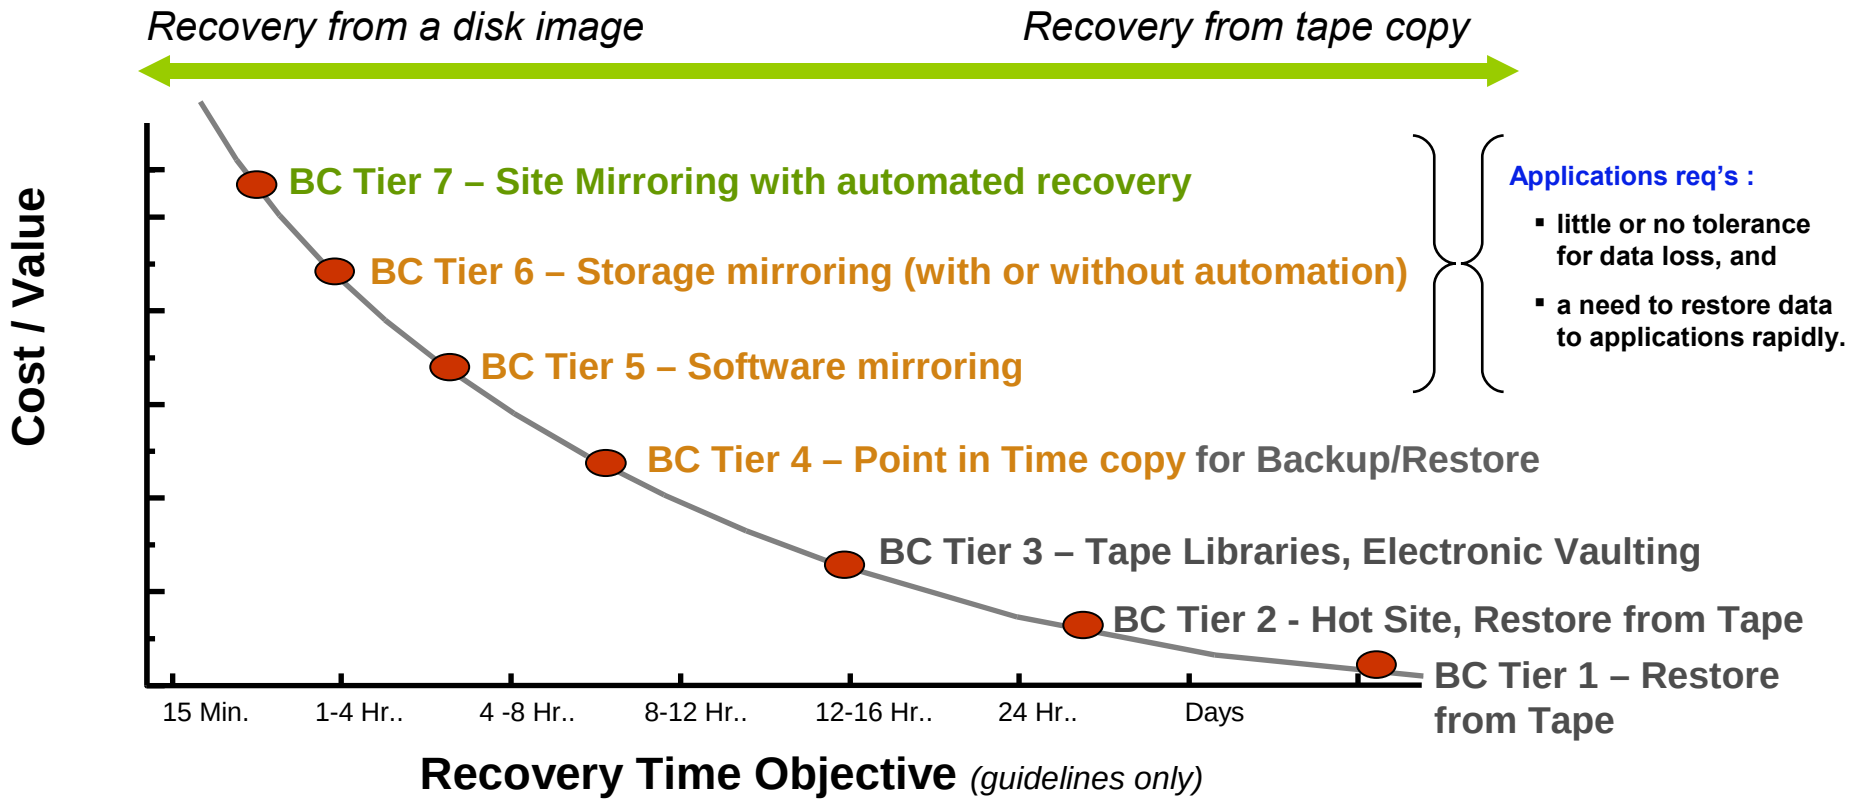

Business Continuity Tiers - business process segments

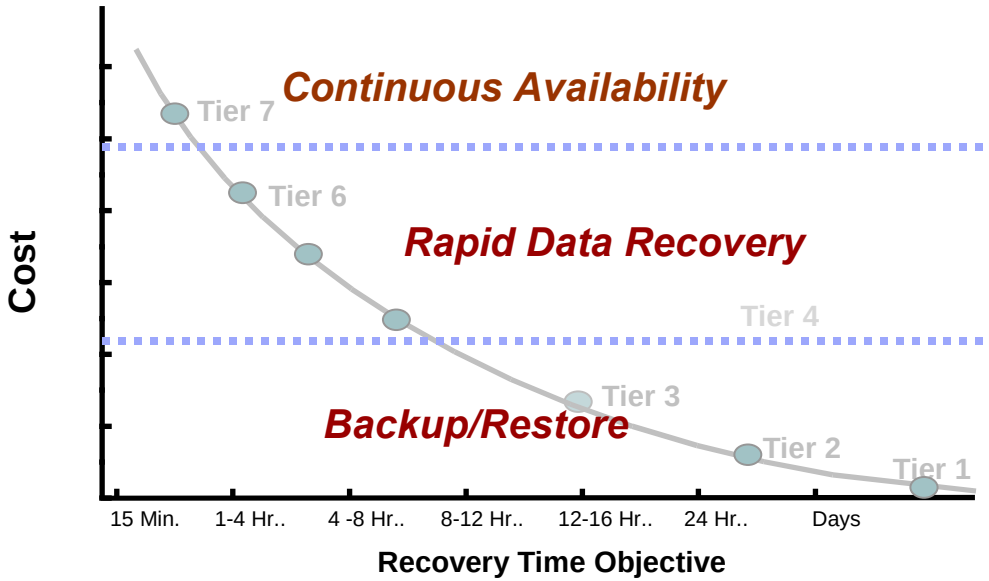

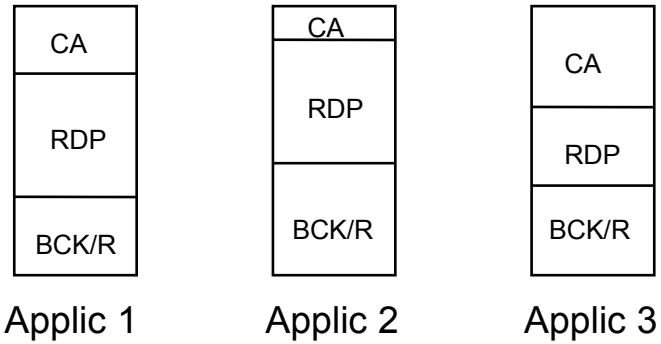

Some examples …….

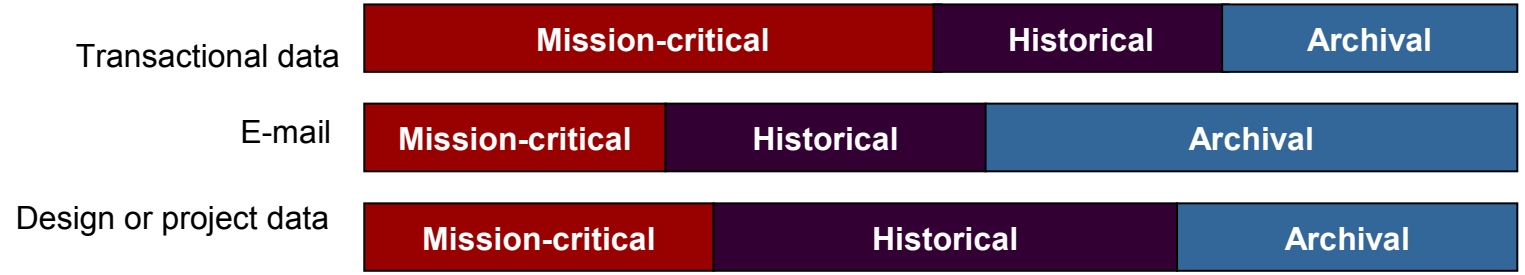

Server - High Availability / Disaster Recovery: Clustering Solution

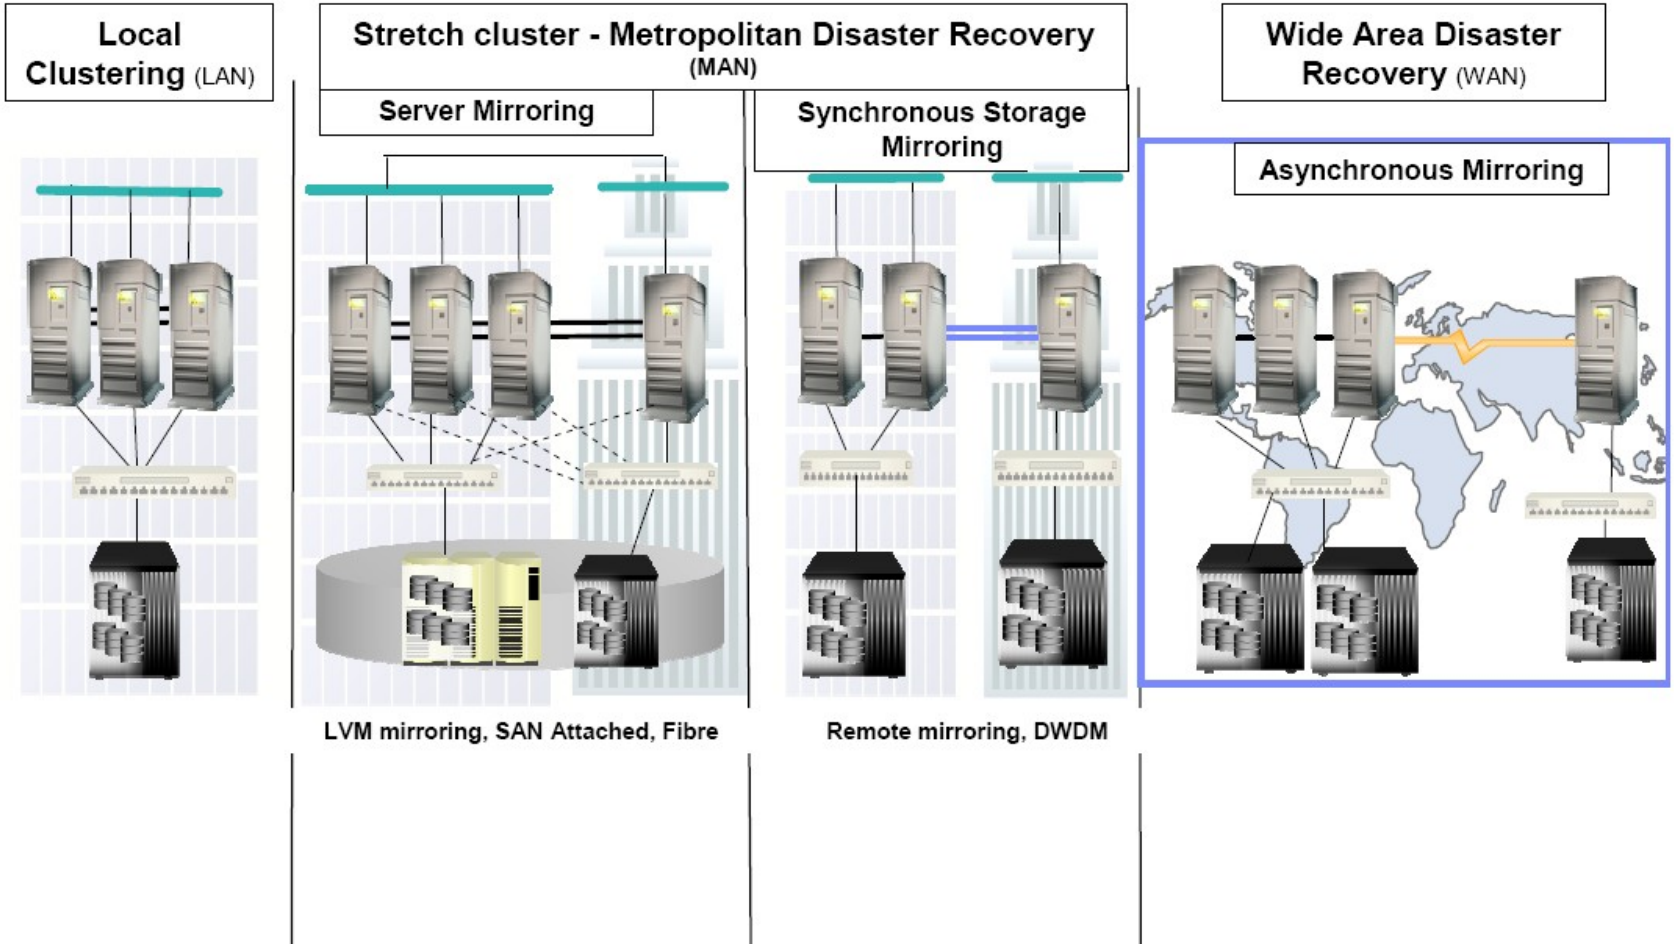

Server : zServer Continuous Availability

Single System

- Built In Redundancy
- Capacity Upgrade on Demand
- Capacity Backup
- Hot Pluggable I/O

Parallel Sysplex

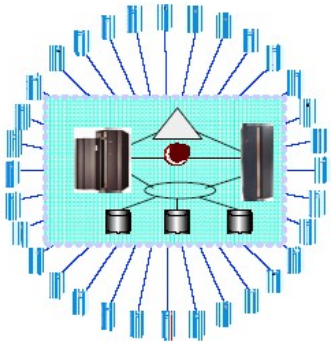

1 to 32 Systems

- Addresses Planned/Unplanned HW/SW Outages
- Flexible, Nondisruptive Growth
  - Capacity beyond largest CEC
  - Scales better than SMPs
- Dynamic Workload/Resource Management

GEO Plex

GDPS

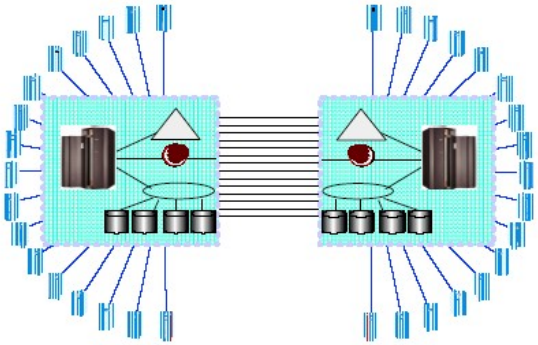

Site 1 Site 2

- Addresses Site Failure/Maintenance
- Sync/Async Data Mirroring
  - Eliminates Tape/Disk SPOF
  - No/Some Data Loss
- Application Independent

Storage - Where to place data replication?

Host-based (Server) replication AIX Logical Volume Manager (LVM), AIX Geographic Logical Volume Manager (GLVM), General Parallel File System (GPFS) – or **Application replication**

Hardware (Storage Array) based replication: IBM Metro Mirror (Disk Storage synchronous mirroring), IBM Global Mirror (Disk Storage asynchronous mirroring), IBM Virtual Tape Remote Replication

Network (SAN) based replication : Brocade and CISCO Switch Routing appliance

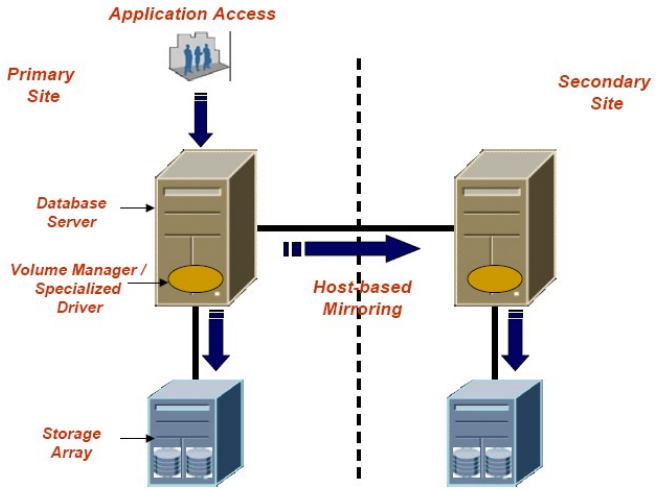

Fig. 1: Host-based Remote Mirroring

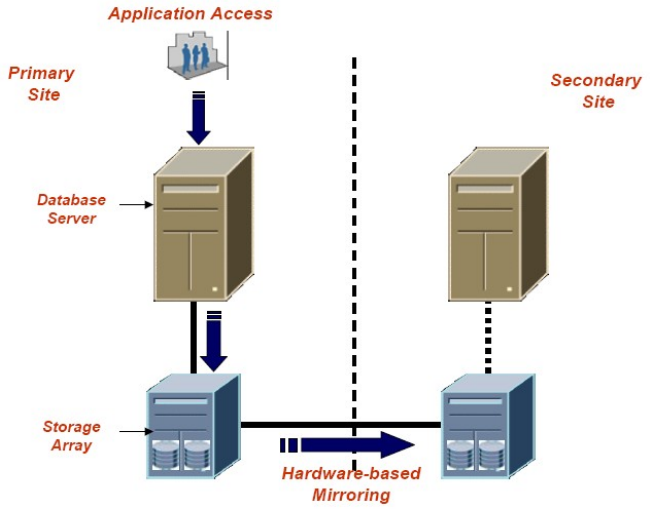

Fig. 2: Hardware-based Remote Mirroring

Storage - Where to place data replication?

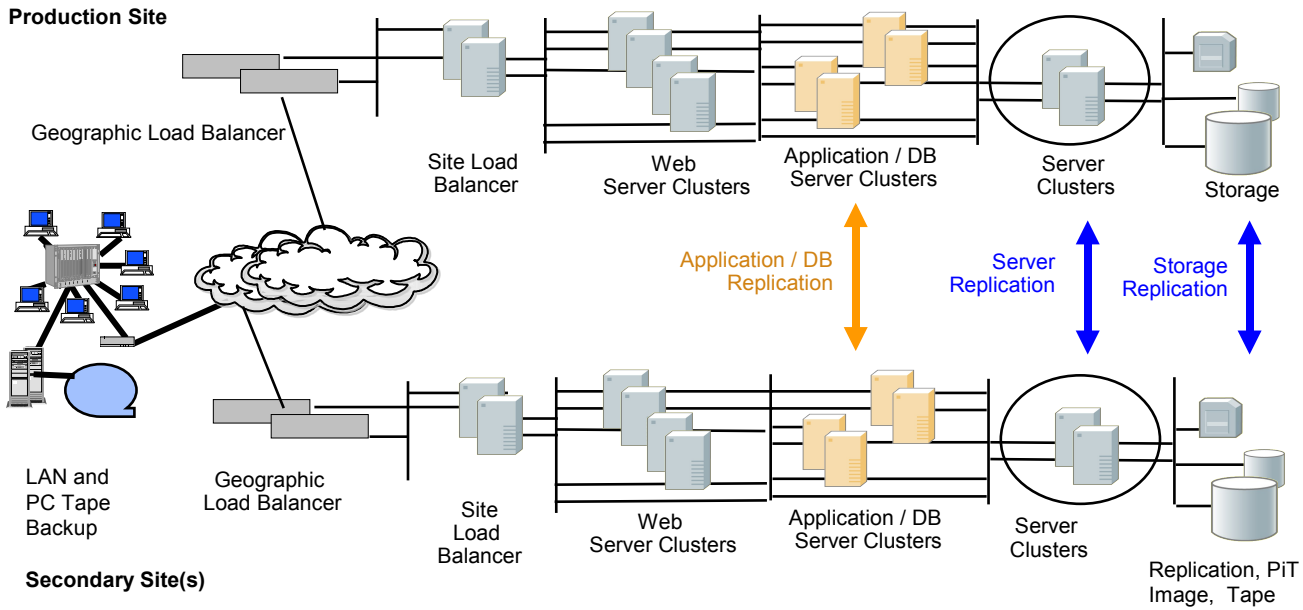

More complexity

- HOST Replication - Application / Database replication based
  - Requires less bandwidth, usually asynchronous mirroring only
  - Span of consistency is the application or database only
  - More complex implementation, must implement for each application
- HOST Replication - Server based
  - Relatively simple implementation, application independent
  - Uses server cycles, span limited to that OS
- HW Replication – Storage Array based
  - Simpler implementation, mirror logical disks, multiple heterogeneous systems, application independent
  - Requires more bandwidth
  - Consistent Repeatable RPO (RTO depends on Server/Data/Workload/Network)

More expensive

Application Based Replication: DB2 HADR

DB2 for z/OS implements DB2 data sharing on a parallel sysplex

DB2 for LUW (DB2 for Linux, Unix, and Windows) systems employs a "shared nothing" architecture. DB2 for LUW can restart quickly with use of **High Availability Disaster Recovery (HADR)** where the primary DB2 instance is continuously sending log records to the hot standby DB2 instance keeping its database copy synchronized.

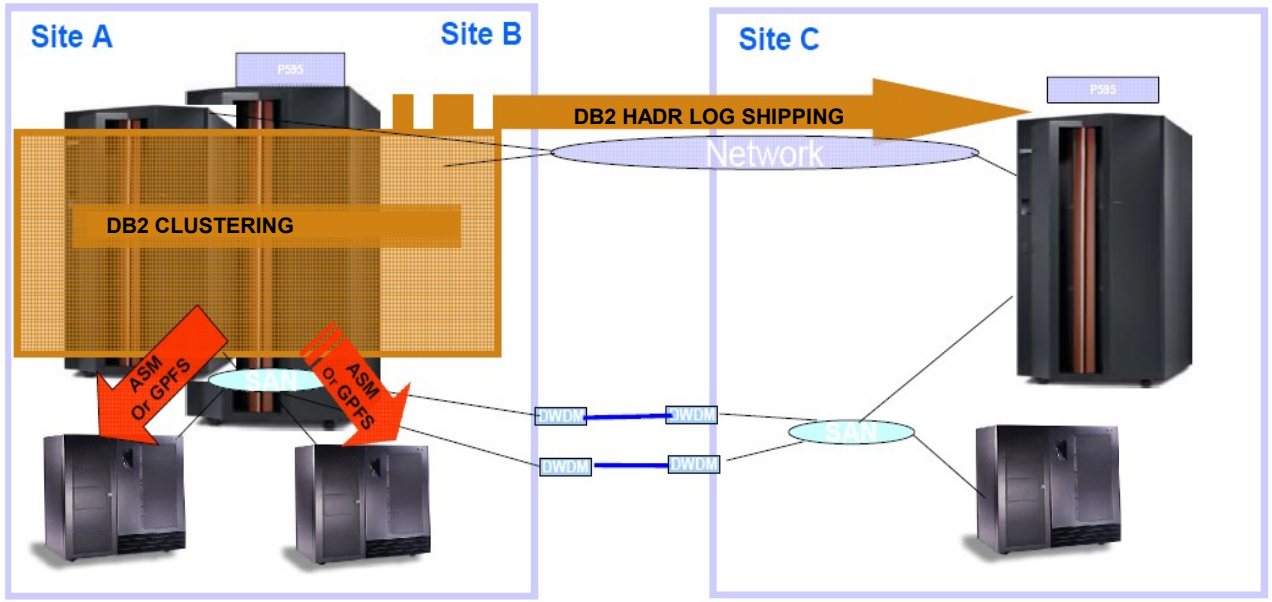

Application Based Replication: Oracle Dataguard

Oracle Data Guard is a software solution that maintains the consistency between the primary and standby databases by synchronously or asynchronously transmitting transactional redo data from the primary to the standby databases.

Stretch Cluster (or Extended Clusters, or Distance Clusters, or Geo-Clusters) allows these nodes to be located at different sites while still maintaining a unified cluster configuration over a single database.

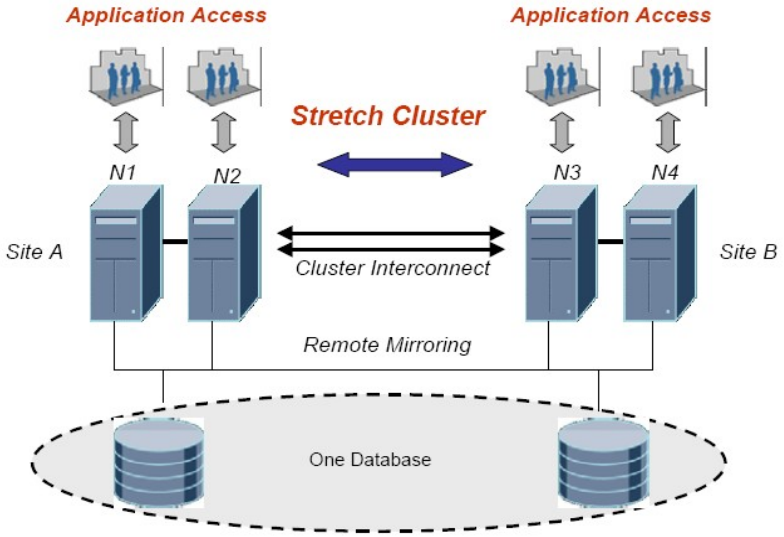

Fig. 3: A 4-node RAC Stretch Cluster Configuration

Disk System Copy functions

Copia Locale

E' una copia dei dati istantanea (copia dei puntatori ai blocchi di dati) all'interno dello stesso sistema di Storage

Copia Remota tra 2 siti

Sincrona : copia dei dati (senza perdita transazioni) tra sistemi di Storage su distanza metropolitana
Asincrona : copia dei dati (con perdita transazioni) tra sistemi di Storage su distanza geografica

Copia Remota tra 3 siti

Copia dei dati sincrona verso il sito vicino (per garantire la Continuita' operativa) ed asincrona verso il sito piu' lontano (per finalita' di Disaster Recovery)

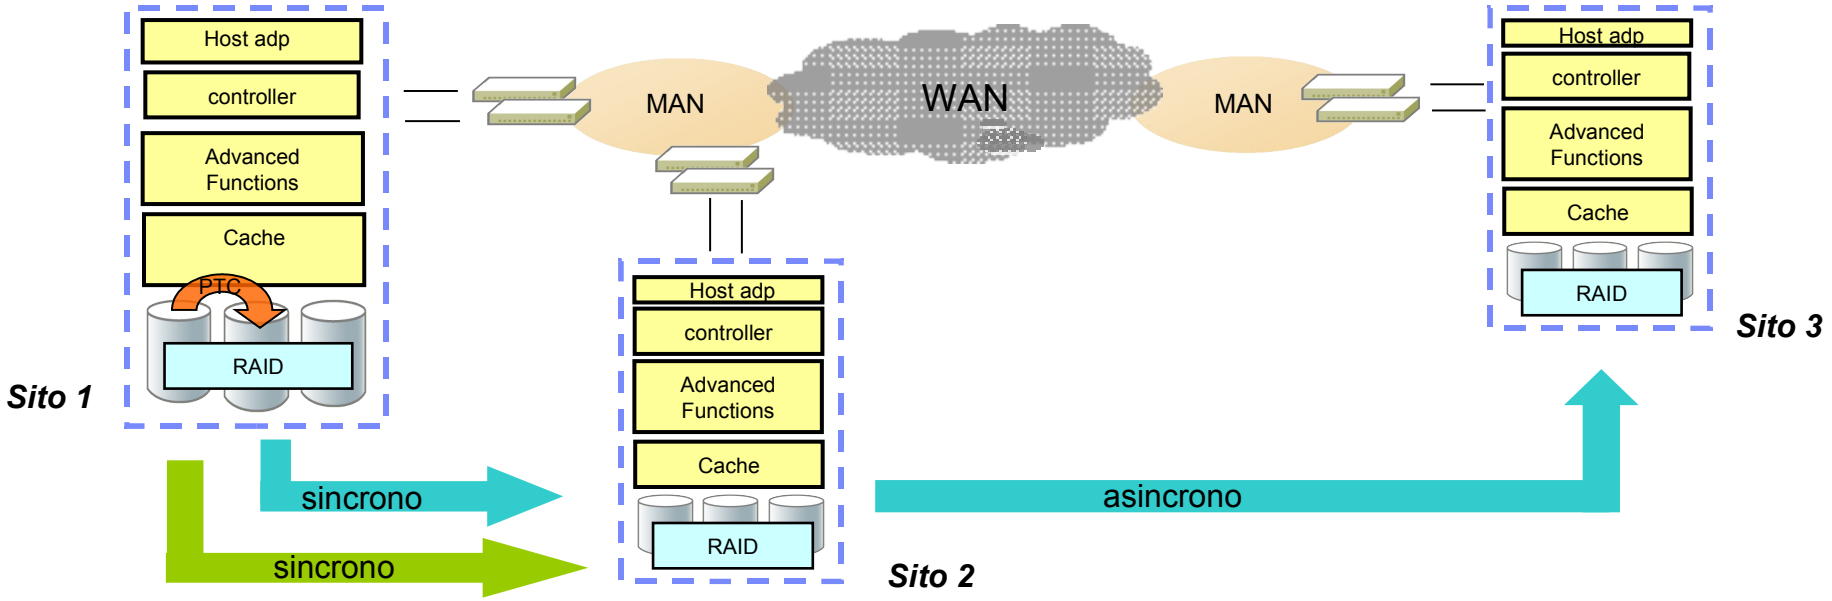

Disk System Copy functions

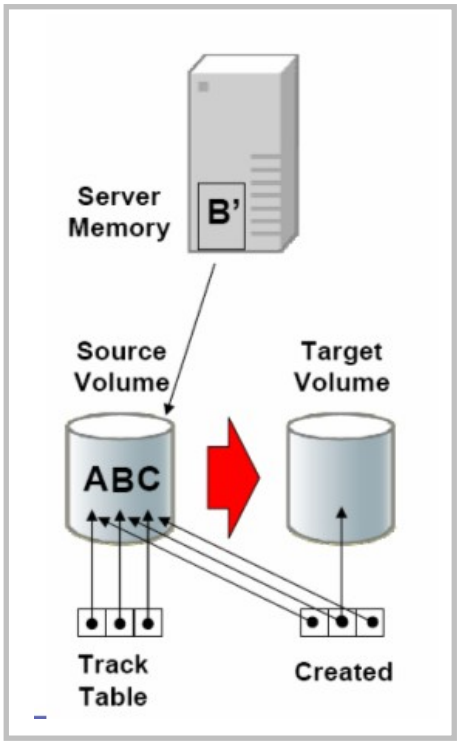

Local Copy
Point in time Copy

Synchronous
Remote Copy

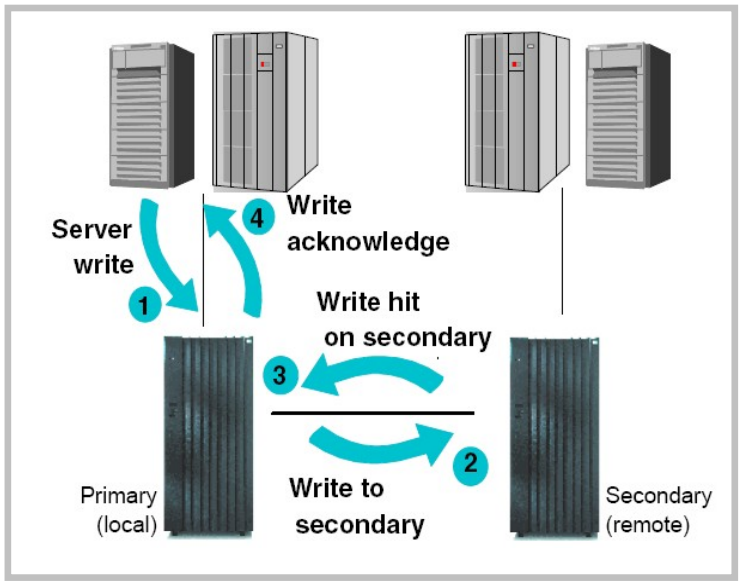

Asynchronous
Remote Copy

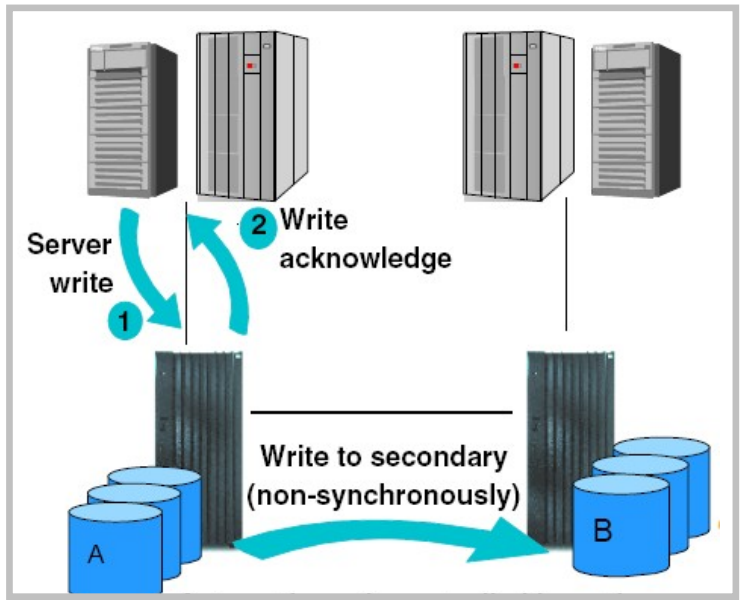

General Disk Replication Guidance

What Customers Want to Do

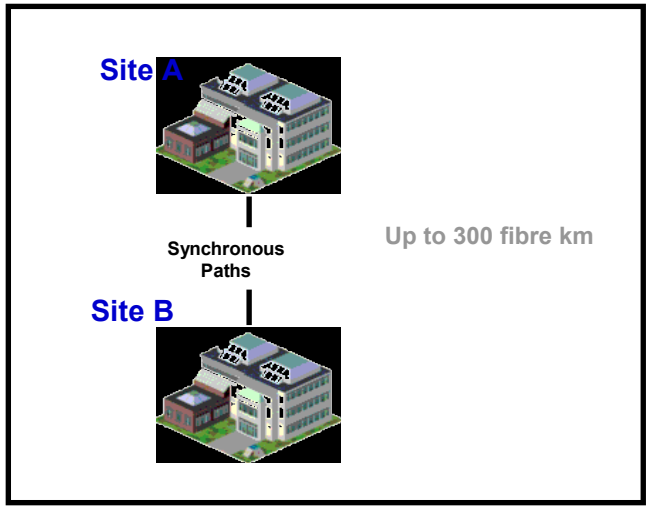

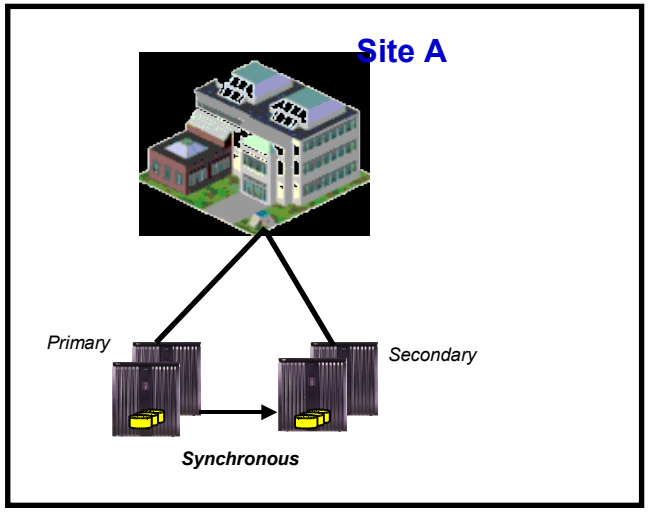

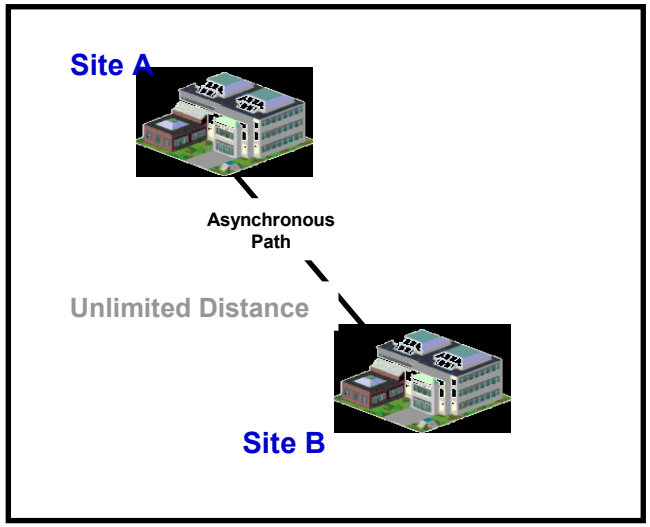

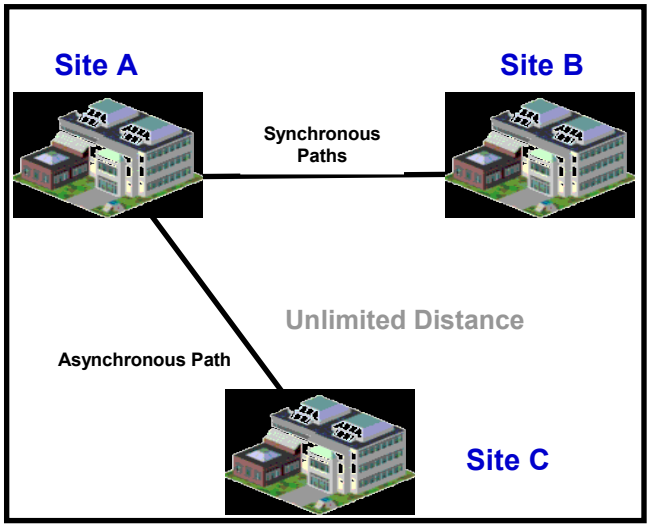

Remote Data Replication Technologies

- **Synchronous**
  - *IBM Metro Mirror*
  - *EMC SRDF/S*
  - *HDS TrueCopy*
- **Asynchronous**
  - *IBM Global Mirror*
  - *EMC SRDF/A*
  - *HDS Universal Replicator*
- **z/OS Asynchronous**
  - IBM *zOS Global Mirror (XRC)*

Two sites

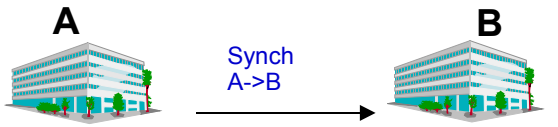

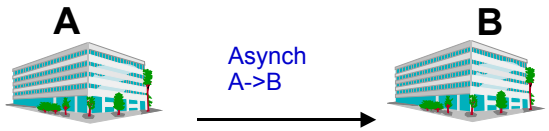

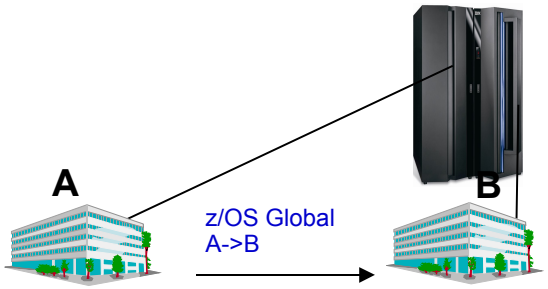

Three sites

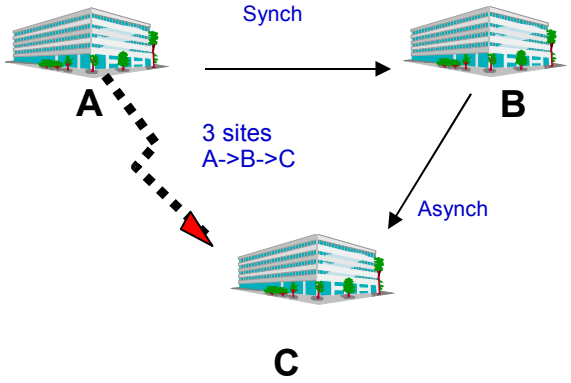

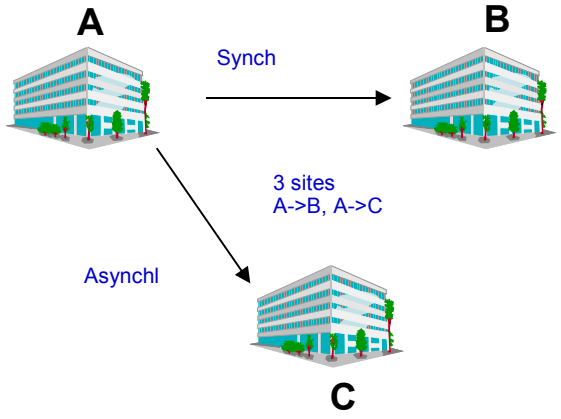

IBM *z/OS Global Mirror* (XRC) concepts - Journaling

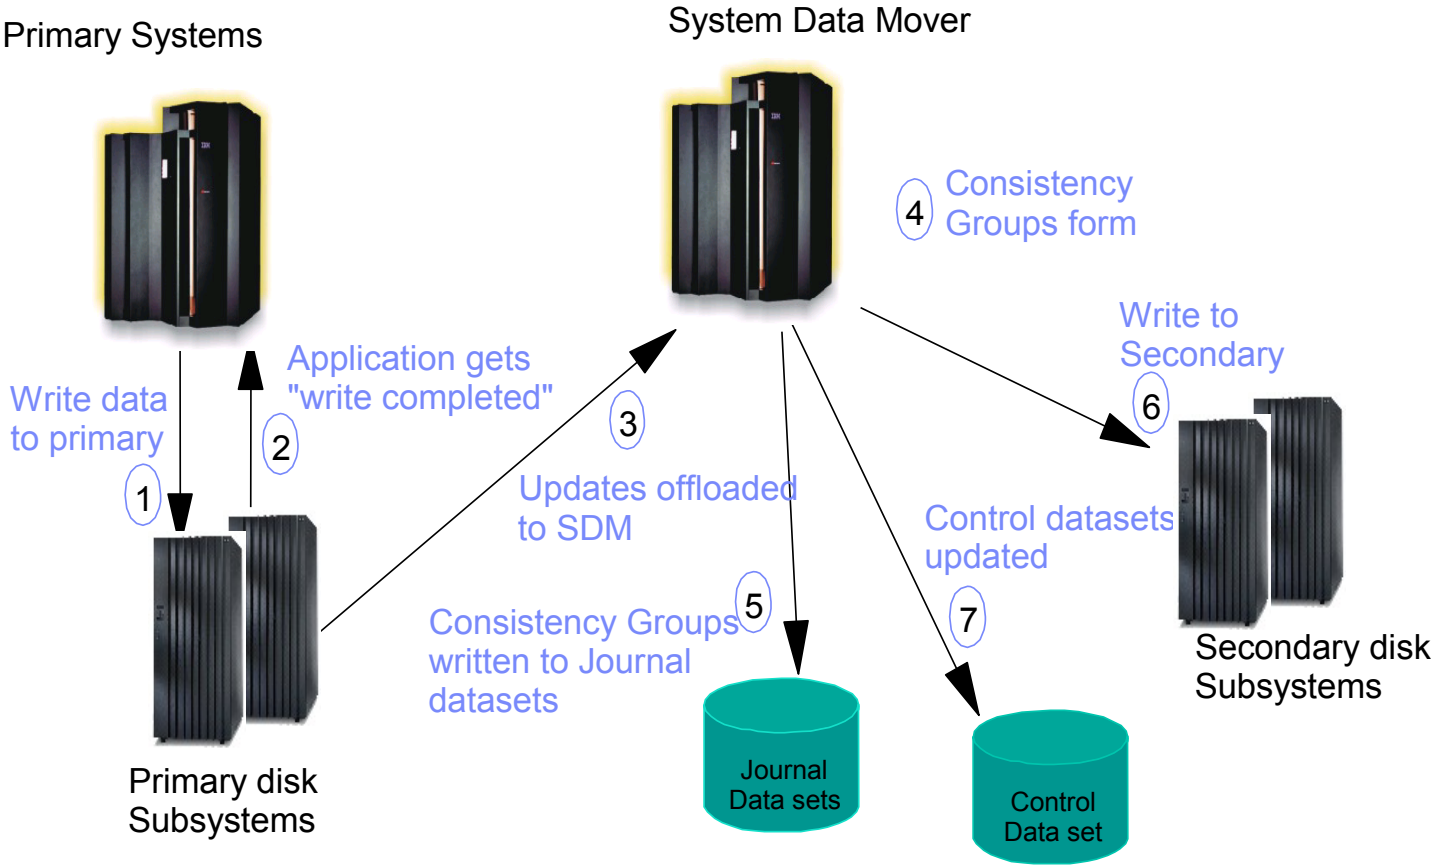

3 sites Disk replication

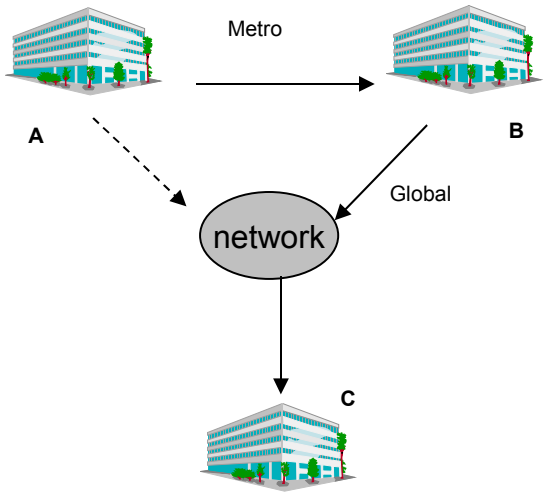

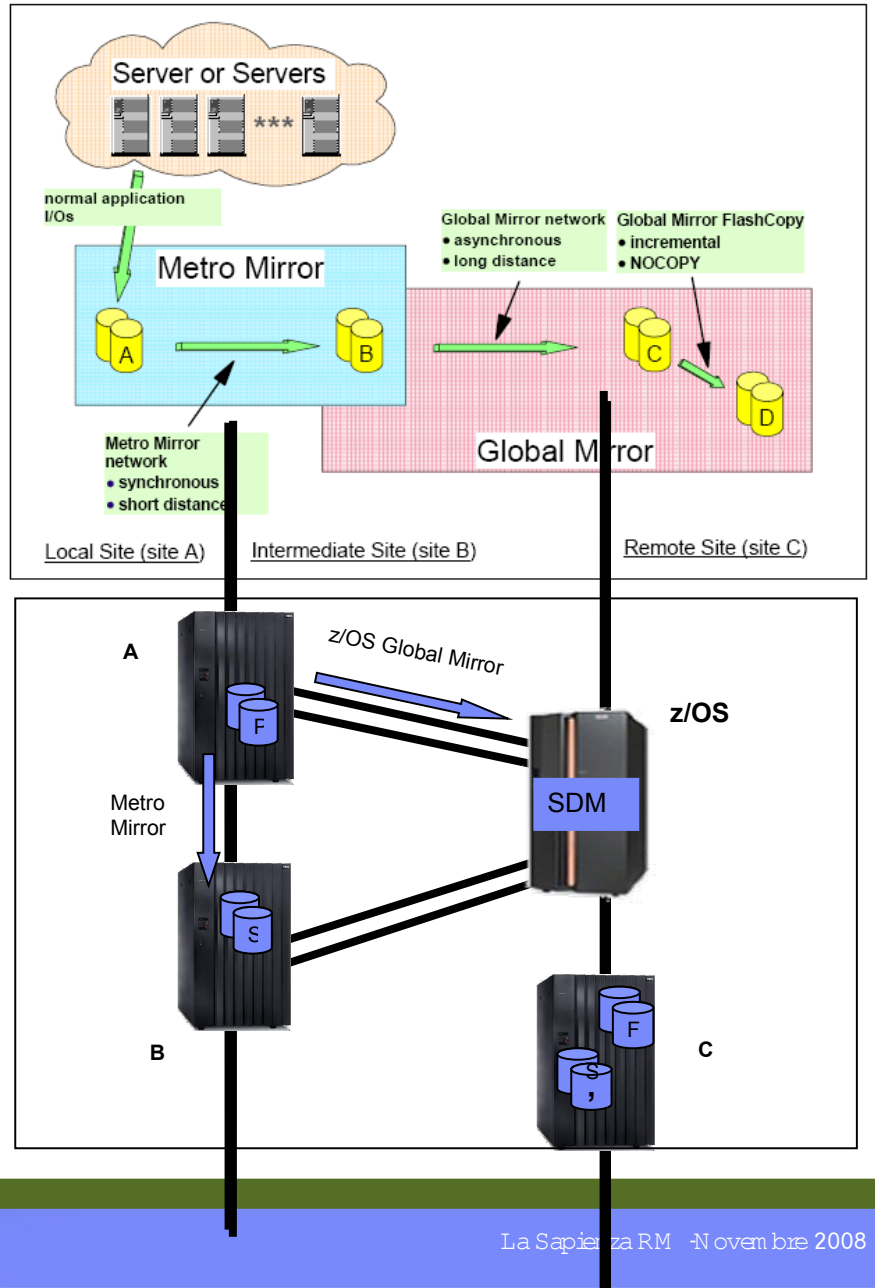

The three site replication requirements are:

- Fast Failover / Failback to any site
- Fast re-establishment of 3 site recovery, without production outages
- Quickly resynchronize any site with incremental changes only (Links and bandwidth assumed between all sites)
- Cascading vs Star configuration

Le Soluzioni a Nastro

Il nastro rimane il dispositivo di memorizzazione col miglior rapporto costo/prestazioni

Operazioni su tape automatizzate attraverso librerie automatiche (ATL) munite di accessori robotici

In ambienti di tipo Large Enterprise le soluzioni Tape vengono virtualizzate con un caching a disco

Il concetto di deduplicazione dei dati consente la riduzione estrema dei dati memorizzati

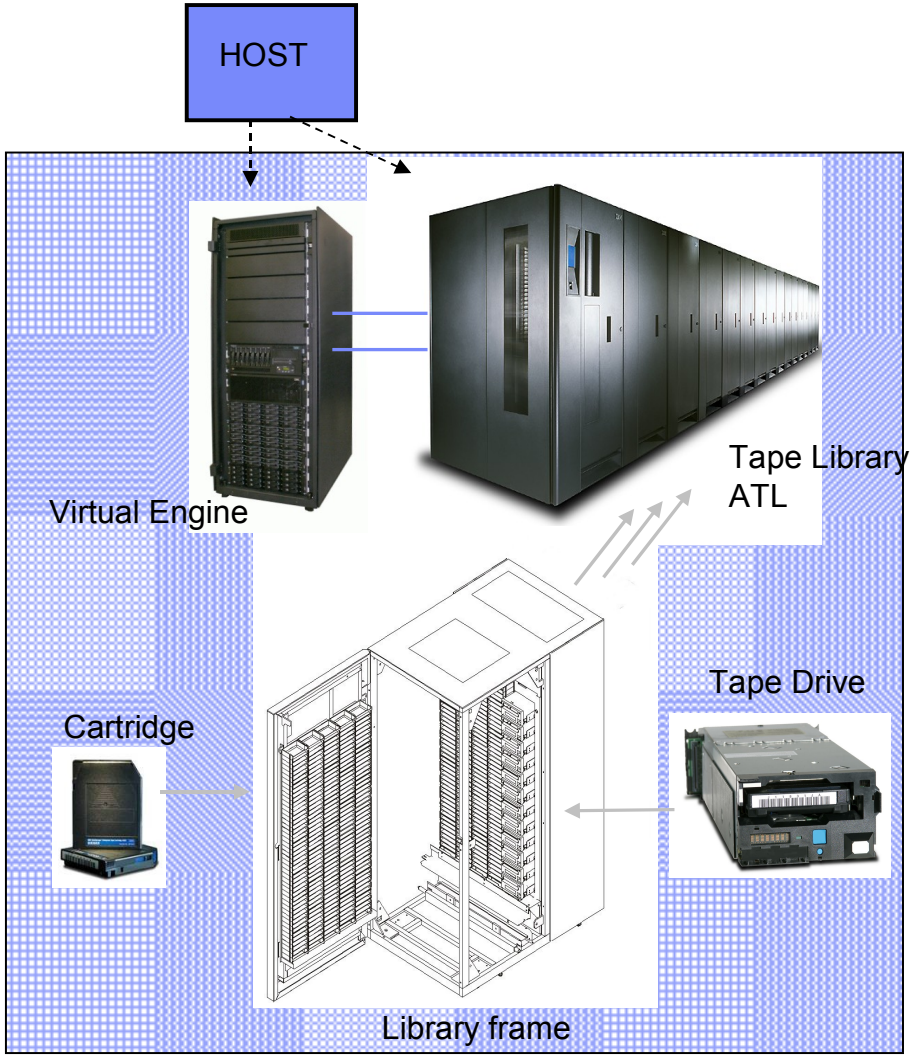

Il concetto di Virtual Tape

I principali benefici apportati da un apparato di virtualizzazione di nastri:

- Simula allo host un numero variabile di tape drive virtuali
- Elimina i ritardi indotti dall'accesso al tape fisico poiche' i dati vengono scritti in cache (memoria dei server interni). I dati vengono scaricati in asincrono dalla cache ai nastri
- Disegnato per utilizzare a pieno la capacita' delle cartucce e delle librerie di nastri. Fa lo stack di piu' volumi logici sulle cartucce in backend permettendo di avere un numero ridotto di tape drives reali

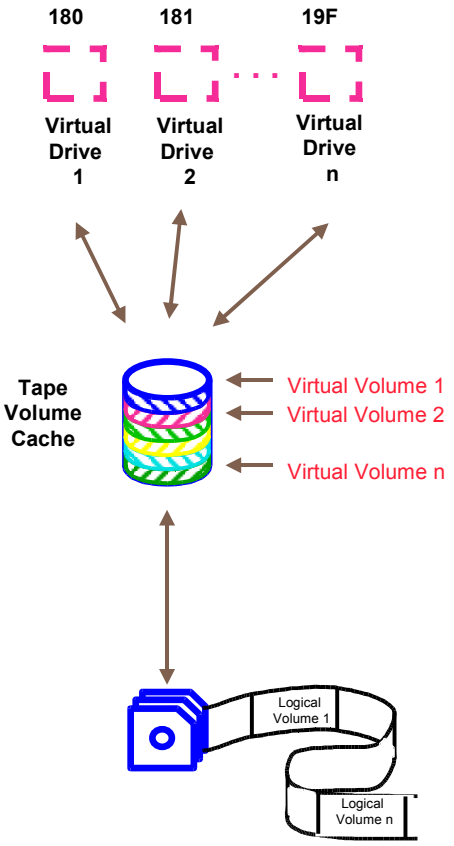

Enterprise Tape solutions: Three Site Grid Configuration

- Virtual Tape Libraries …..
  - Utilize tape assets more efficiently
  - Improves and speeds tape backup process
  - Improves RTO and RPO for tape-based data

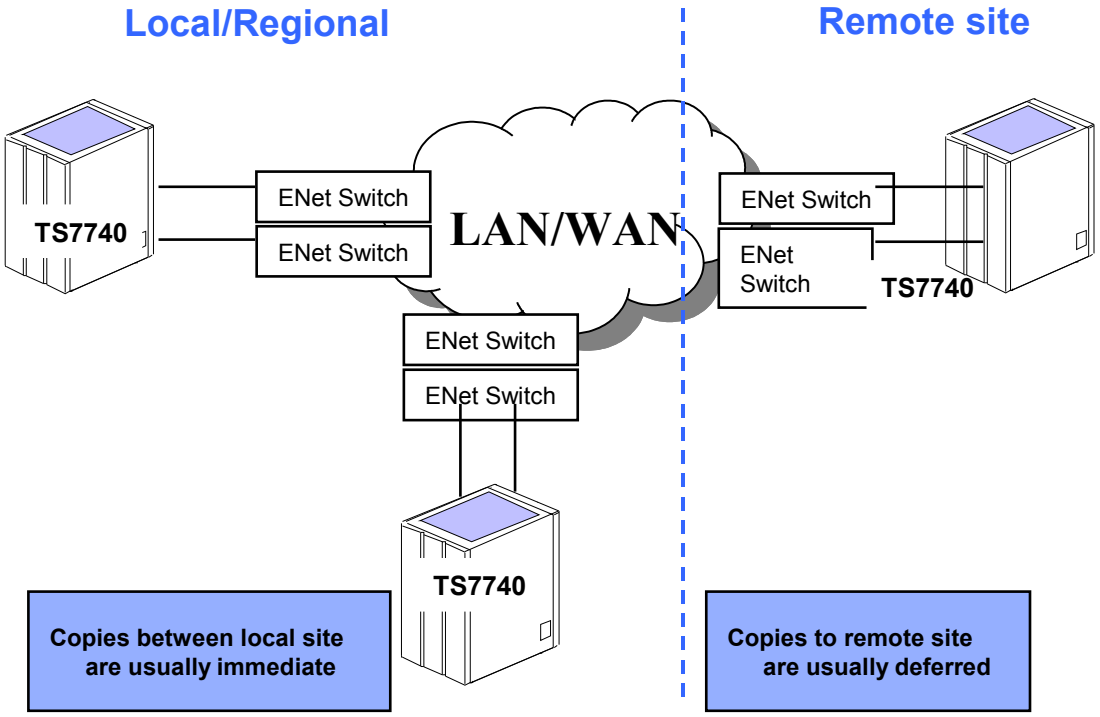

Virtual Tape Libraries and Deduplication

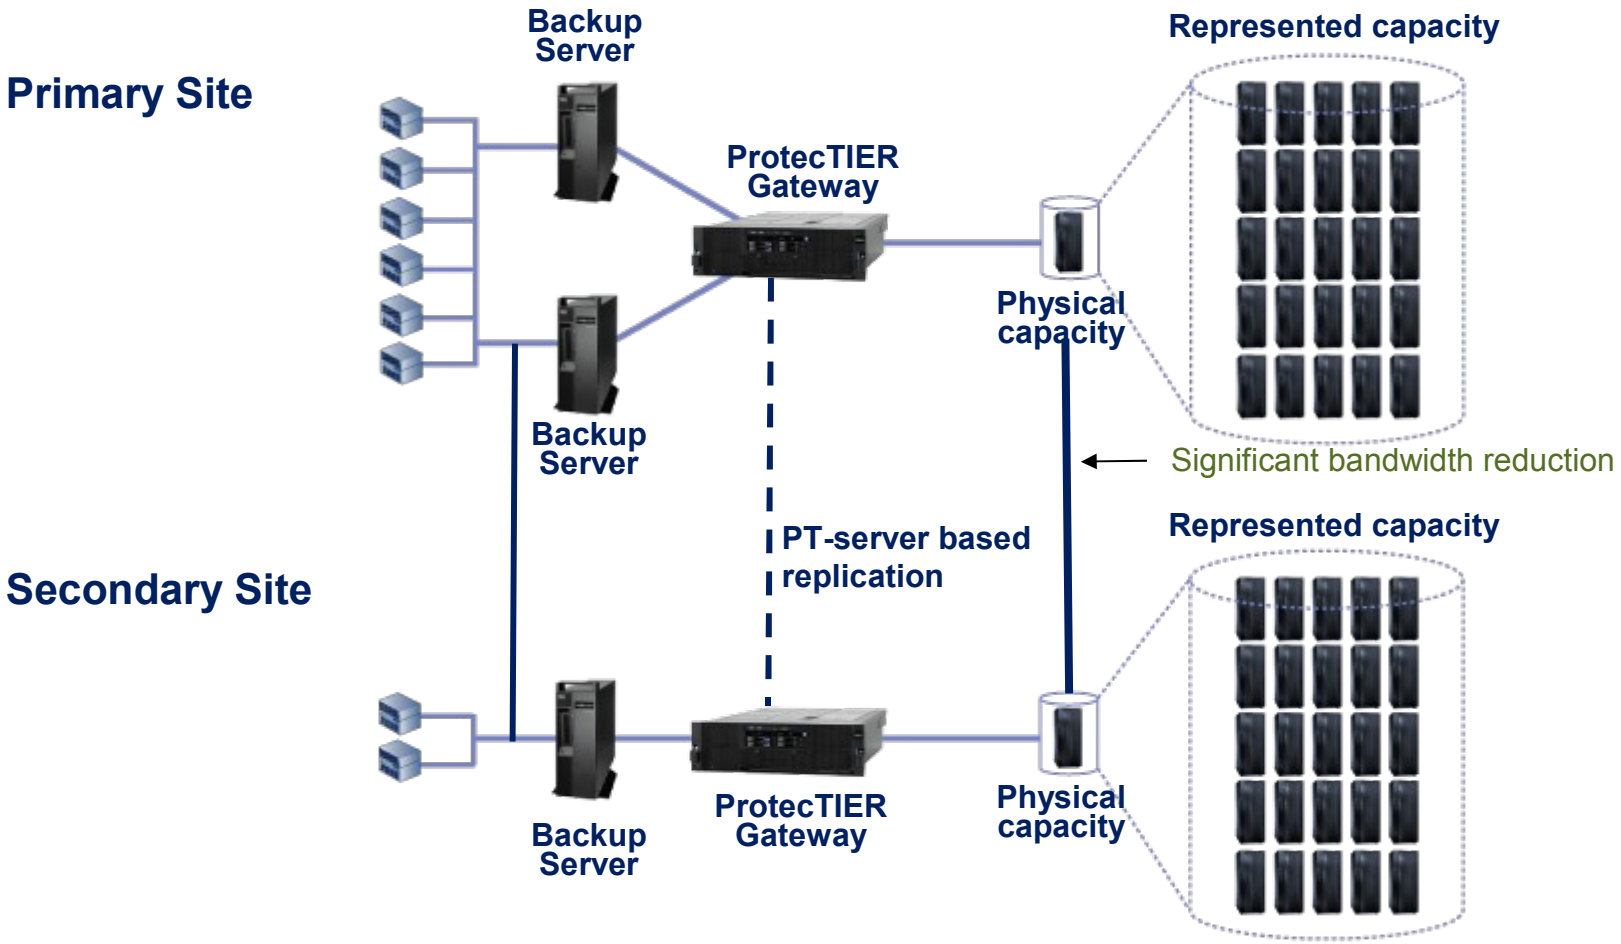

Connettivita' Fiber Channel

Fiber Channel Architecture

Set di regole (da FC-0 a FC-4) che definisce lo standard da implementare per il trasferimento di dati tra computer e device su fibra ottica

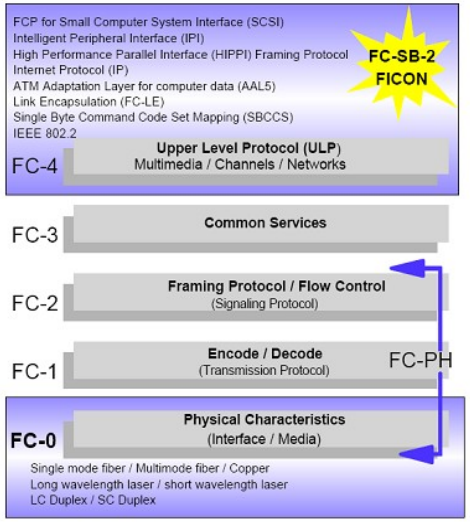

FICON

E' il protocollo che implementa in ambiente zSeries la Fiber Channel Architecture. E' uno standard industriale (FC-SB-2).

FCP

E' il Fiber Channel Protocol per SCSI e' permette di mappare i comandi del protocollo SCSI sulla Fiber Channel Architecture.

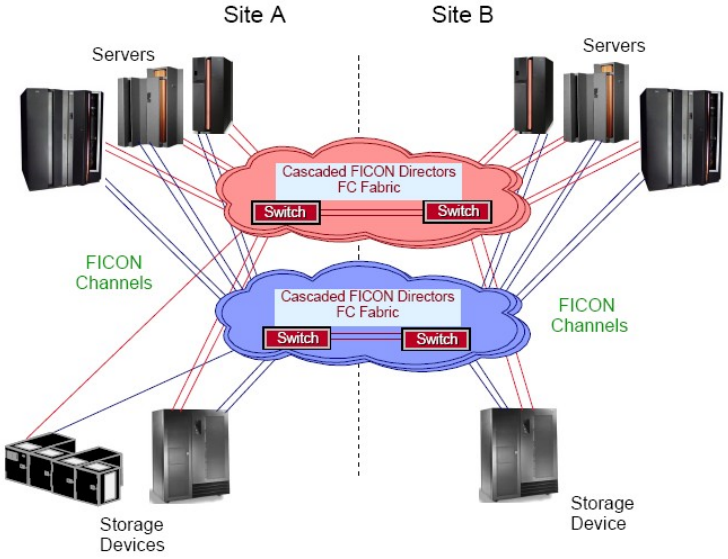

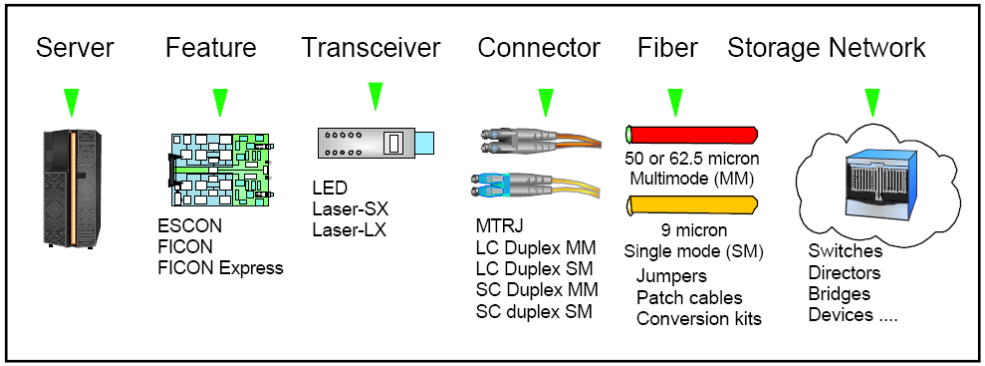

Tiers of Disaster Recovery: Positioning Automation

Automation Benefits

- Provides High Availability solution for large multi-platform zOS, Windows, UNIX and Linux environments.
- Support a wide variety of data replication methods: storage replication, server-based replication, database or software replication
- Reduce / avoid human errors during the recovery process
- Enables regular testing to validate that BC/DR procedures are:repeatable – reliable – scalable – auditable
- Helps reduce infrastructure management costs
- Helps maintain recovery readiness.

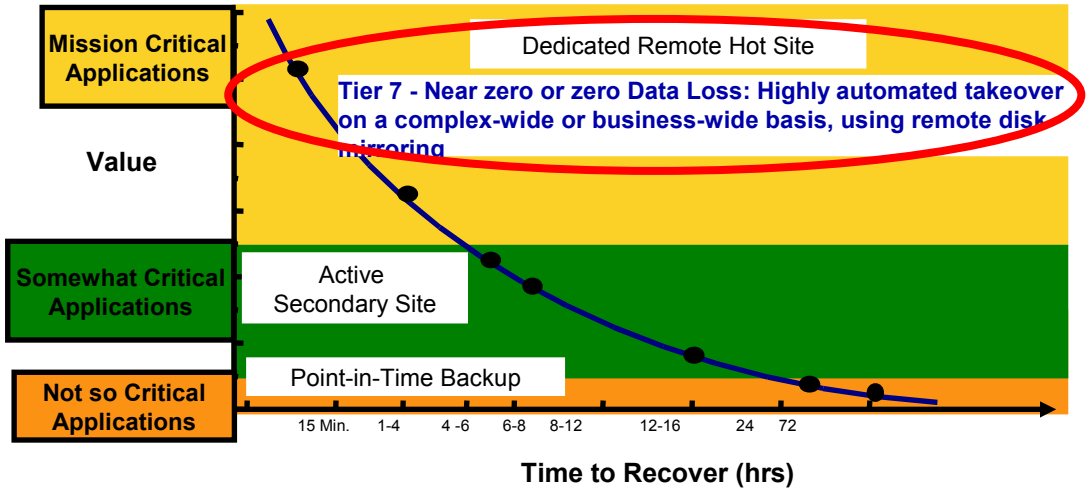

Automation is a principal enabler for successful BC/DR plans