



## Solutions for BC/DR (as seen through Storage)

Klemen Bačak  
[klemen.bacak@si.ibm.com](mailto:klemen.bacak@si.ibm.com)

7.11.2007, Banja Luka

## Agenda....

- 09:30 - 10:00 Registracija
- 10:00 - 10:05 Otvaranje i dobrodošlica
- 10:05 - 10:35 Smanjivanje troškova pomoću konsolidovane infrastruktur
- 10:35 - 11:05 Lenovo novosti
- 11:05 - 11:30 Visokoškolske inicijative
- 11:30 - 11:45 Pauza
- 11:45 - 12:15 IBM BladeCenter / Pregled IBM SW proizvoda
- 12:15 - 12:45 Storage za mala i srednja preduzeća / IBM Tivoli SW proizvodi
- 12:45 - 13:30 Ručak
- 13:30 - 14:00 **Storage rješenja za Disaster Recovery** / IBM Lotus Notes / Domino 8
- 14:00 - 14:30 IBM rješenja za virtualizaciju / IBM Lotus Quickr
- 14:30 - 15:00 System p, POWER 6 i virtualizacija / Praktična demonstracija Notes / Quickr
- 15:00 - 15:15 Izvlačenje nagrada i završetak

# Definitions

Business Continuity

Recovery Time Objective (RTO)

Recovery Point Objective (RPO)



**Useful Disaster Recovery Definitions and Concepts**

***<http://www.redbooks.ibm.com/abstracts/tips0047.html>***

# Business Continuity

- High Availability
- Continuous Operations
- Disaster Recovery




# High Availability



*The ability to provide continuous operations*

# Continuous Operations

Systems designed to tolerate **local** outages.



# Disaster Recovery

Systems designed to tolerate **site** outages.



**Hurricane Katrina in the Gulf of Mexico, courtesy NASA**


# Recovery Time Objective (RTO)

How long until your systems are back up and running  
(available for production work)

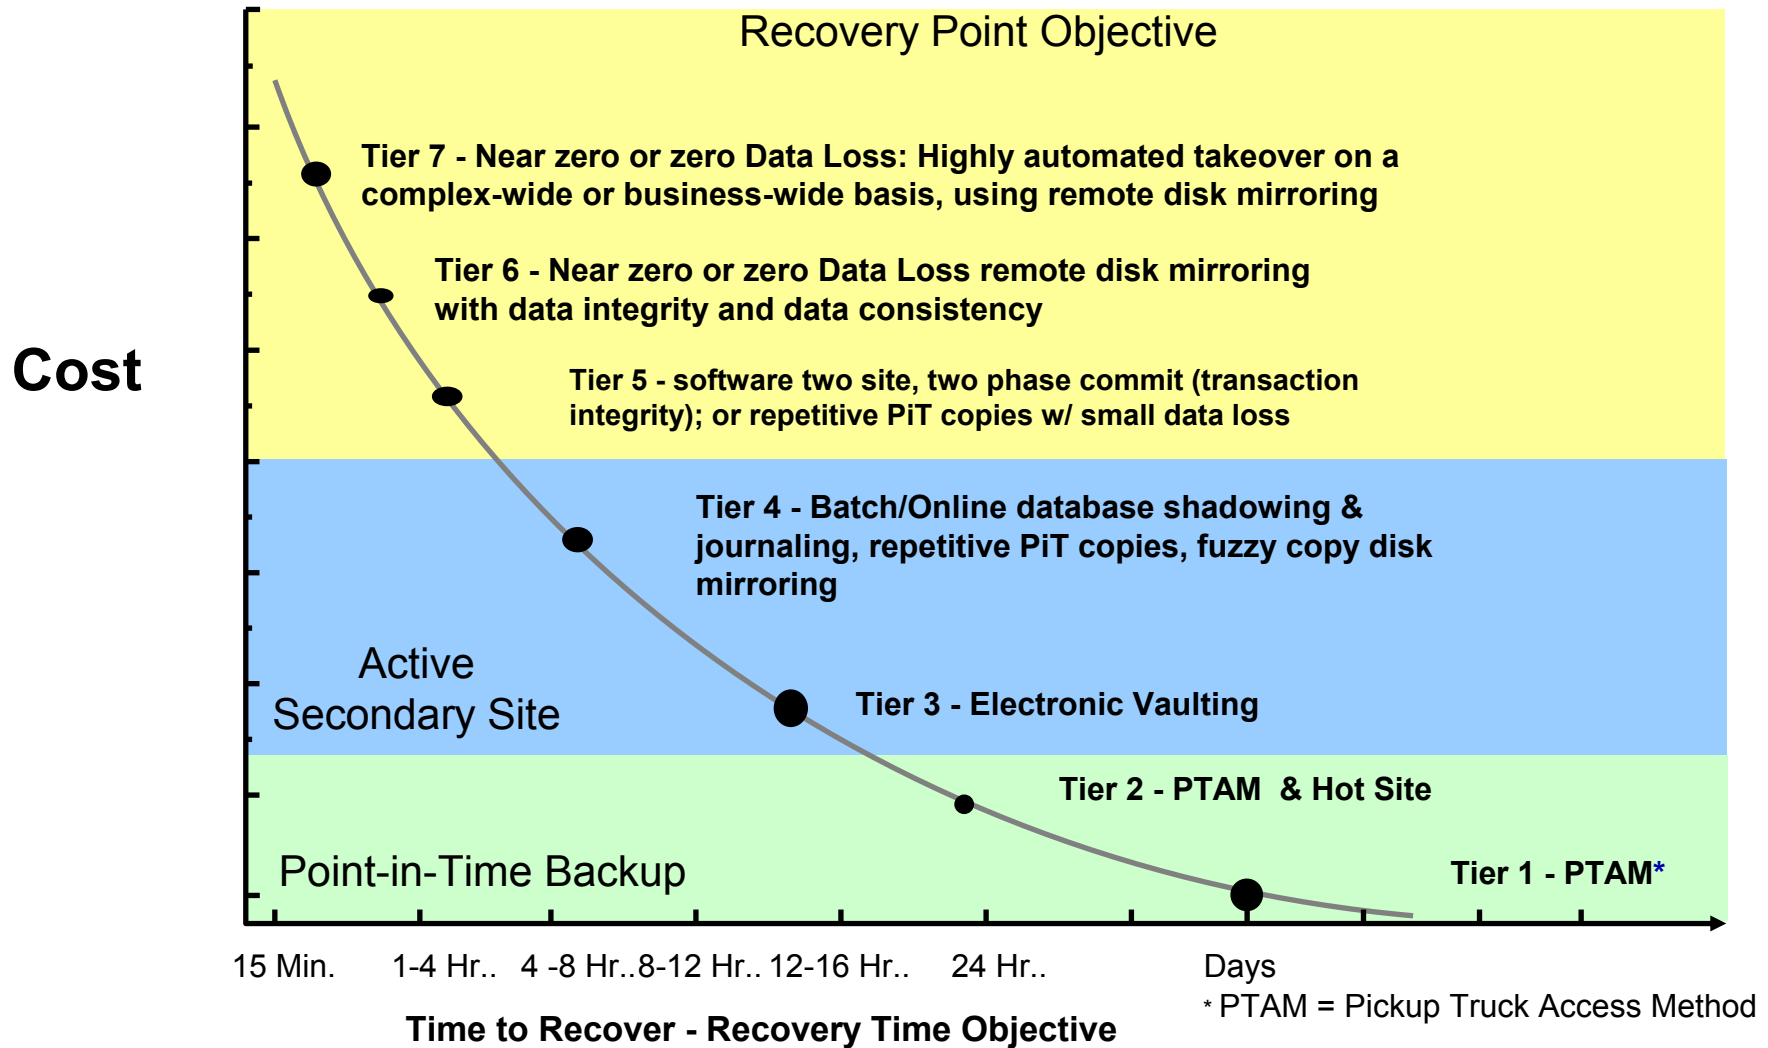


# Recovery Point Objective (RPO)

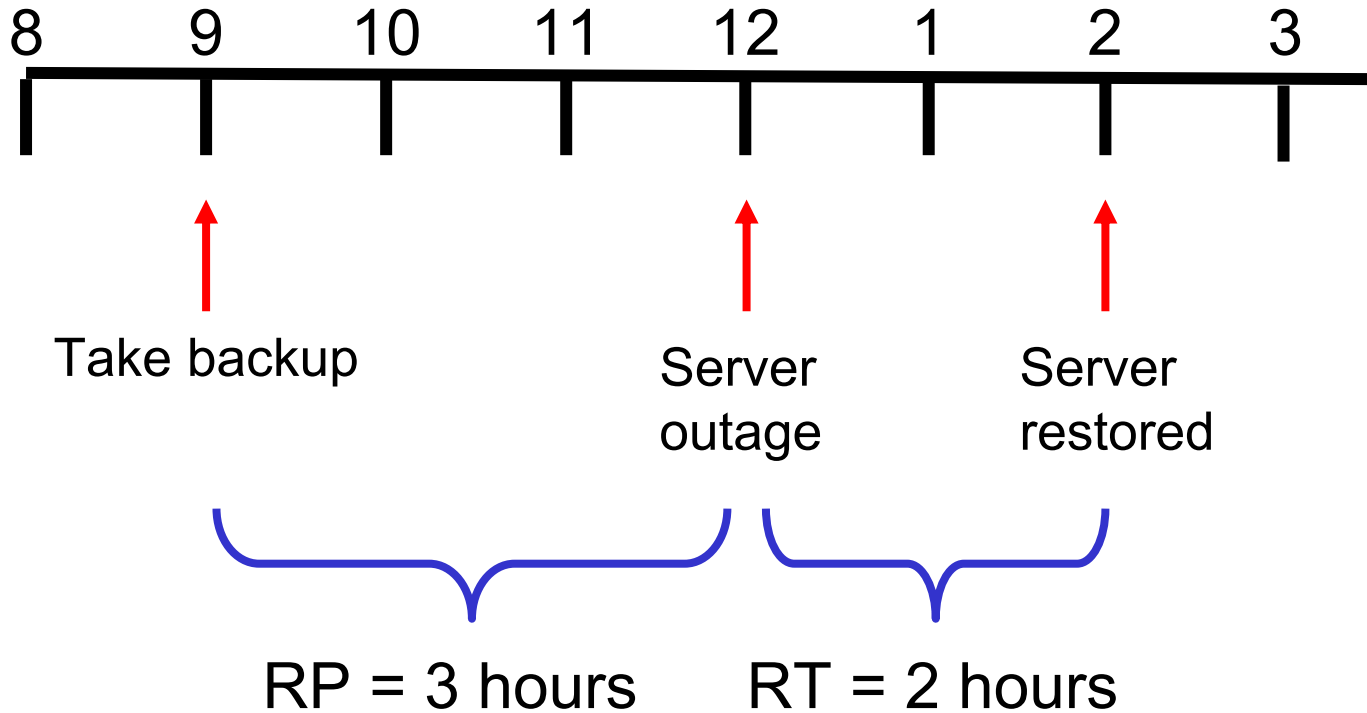
Amount of data that is lost  
(Age of data used for restart)



# Recovery Metrics



# R<sub>ecovery</sub> T<sub>ime</sub> and R<sub>ecovery</sub> P<sub>oint</sub>



# Recovery Time Objective (RTO)

What is an appropriate RTO?



# Recovery Time Objective (RTO)

What is an appropriate RTO?



BOARD OF GOVERNORS OF THE FEDERAL RESERVE SYSTEM  
DIVISION OF BANKING SUPERVISION AND REGULATION  
**SR 03-9** dated **May 28, 2003**

“Core clearing and settlement organizations” = two hours

“Firms that play significant roles” = four hours

<http://www.federalreserve.gov/boarddocs/srletters/2003/sr0309.htm>

# Recovery Time Objective (RTO)

What is an appropriate RTO?



**One of our customers had an objective of two days**



# IBM System Storage™ Resiliency Family



# IBM System Storage™ Resiliency Family

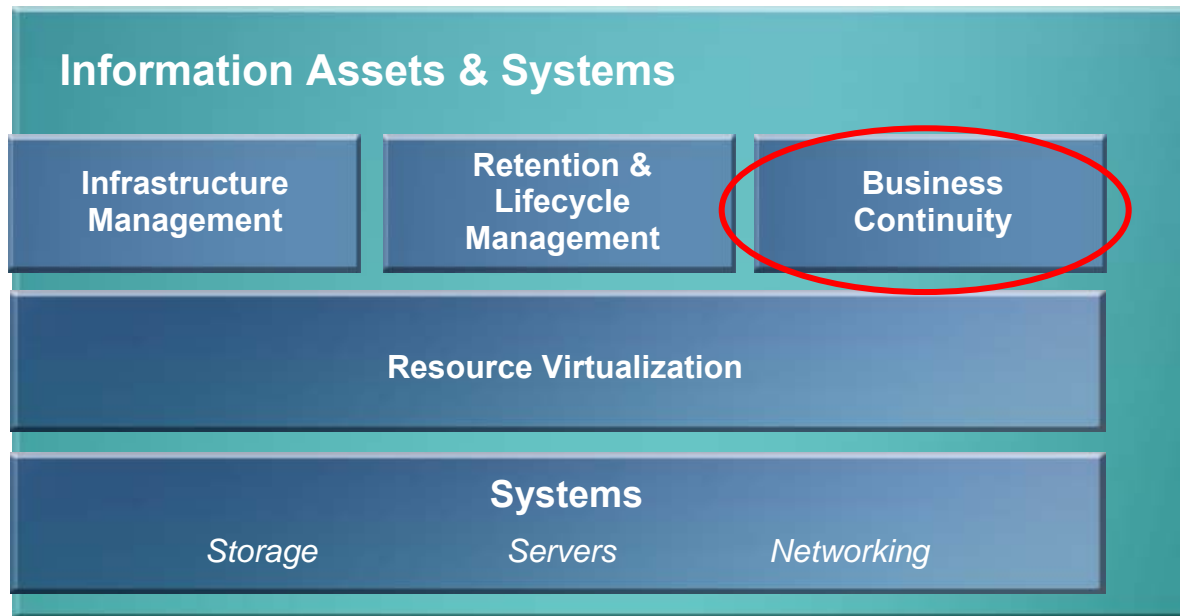
**An overview of the functions on IBM Disk Mirroring:**

- ▶ **DS8000, DS6000**
- ▶ **DS4000**
- ▶ **SAN Volume Controller**
- ▶ **N Series**

**And how to use them in business continuity and disaster recovery planning**



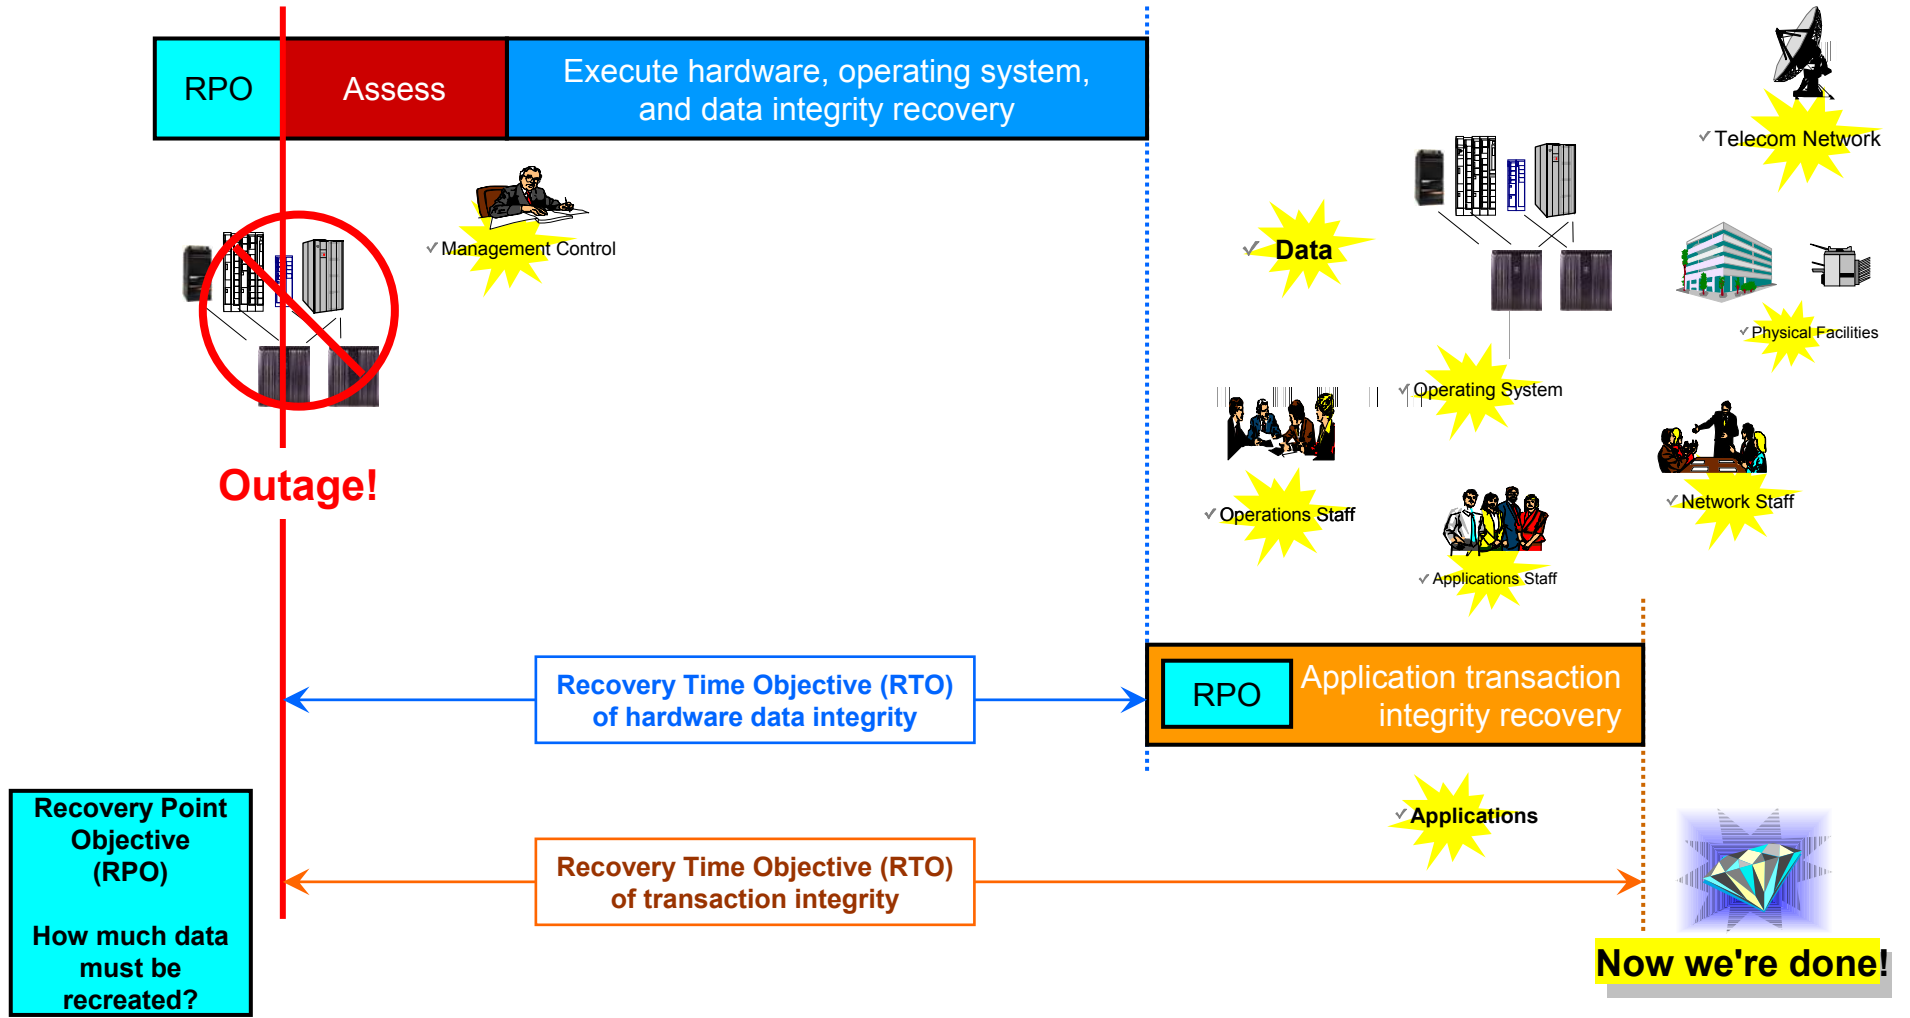
# Information Assets & Systems



- Clustering
  - ▶ Mainframe
  - ▶ UNIX
  - ▶ Linux
  - ▶ Windows
- Replication
  - ▶ FlashCopy
  - ▶ Metro Mirror
  - ▶ Global Mirror
- Continuous Data Protection
  - ▶ for Files
- Backup / Recovery
  - ▶ for enterprises
  - ▶ for departments and SMB


Timeline of an IT Recovery ==>

Storage disk mirroring provides power-outage crash-consistent image at this point in recovery




## General IBM Disk Replication Selection Guidance

- Do initial disk storage platform selection first based on these factors:
  1. The nature of the existing infrastructure
    - Open servers? zSeries? iSeries?
    - Storage virtualization required?
  2. Cost requirements
    - RPO? RTO? Distance to target? Telecom links?
  3. Application requirements
    - Database restart? Write performance?
    - Size and scalability? Number of copies?
  
- Refine if necessary based on these factors:
  1. Consistency Group scalability
  2. Performance and distance characteristics

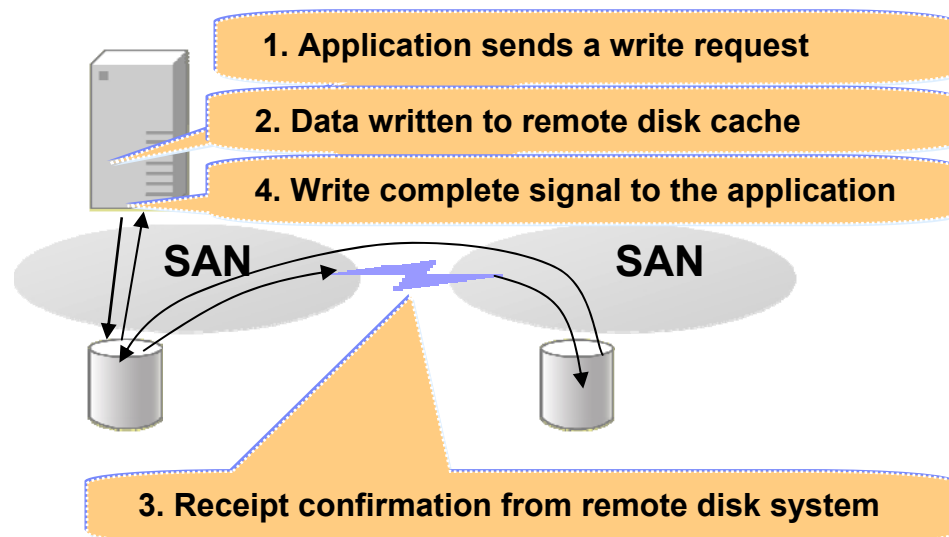


# IBM Disk Mirroring Technologies

- **Metro Mirror**
  - ▶ **Synchronous mirroring (campus)**
    - DS8, DS6, ESS
    - DS4000
    - SAN Volume Controller
    - N series
- **Global Mirror**
  - ▶ **Asynchronous mirroring (global)**
    - DS8, DS6, ESS
    - DS4000
    - SVC
    - N Series
- **Metro / Global Mirror**
  - ▶ **Three site synchronous and asynchronous mirroring**
    - DS8000, ESS
    - N Series



# Review: IBM Metro Mirror

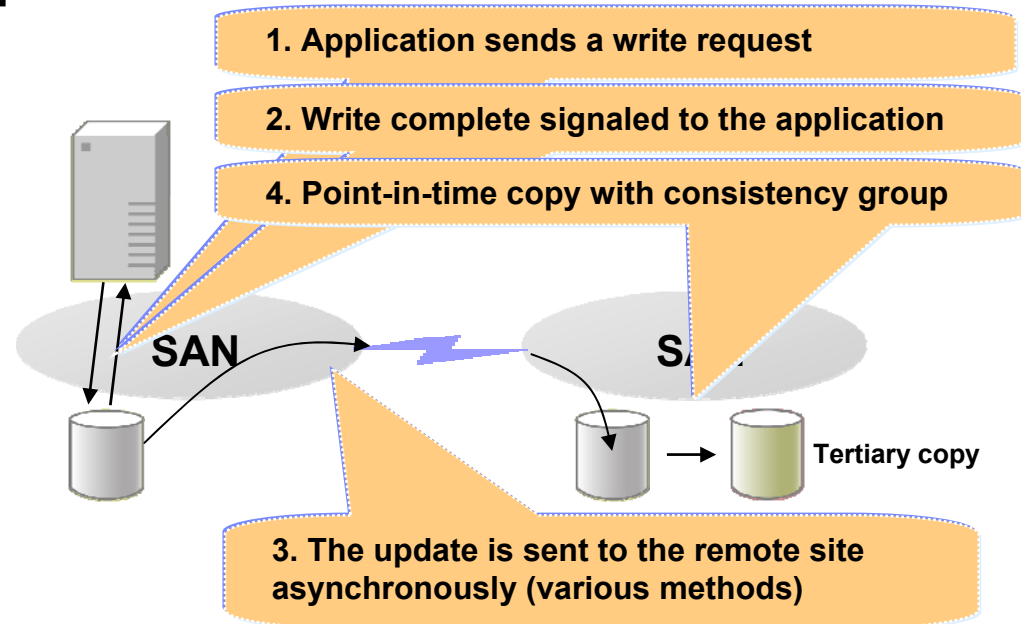


Remote disk is exactly in “lock step”  
at the record write level

- What is Metro Mirror?
  - ▶ Synchronous mirror of data volumes to a remote location
  - ▶ Typical suspend / resume capability
  - ▶ Available on:
    - ESS, DS6000, DS8000
    - DS4000
      - Note: no Metro Mirror Consistency Group – for DS4000, use Global Mirror if you need Consistency Group
    - SAN Volume Controller
    - N series
- What is Metro Mirror used for?
  - ▶ For applications that need zero data loss, Metro Mirror provides the tightest Recovery Point Objective
  - ▶ About 1 ms delay per 100km
  - ▶ Distance depends on platform
    - DS8/6/ESS = 300 km
    - SVC = 100 km
    - DS4000 recommends 10-50 km

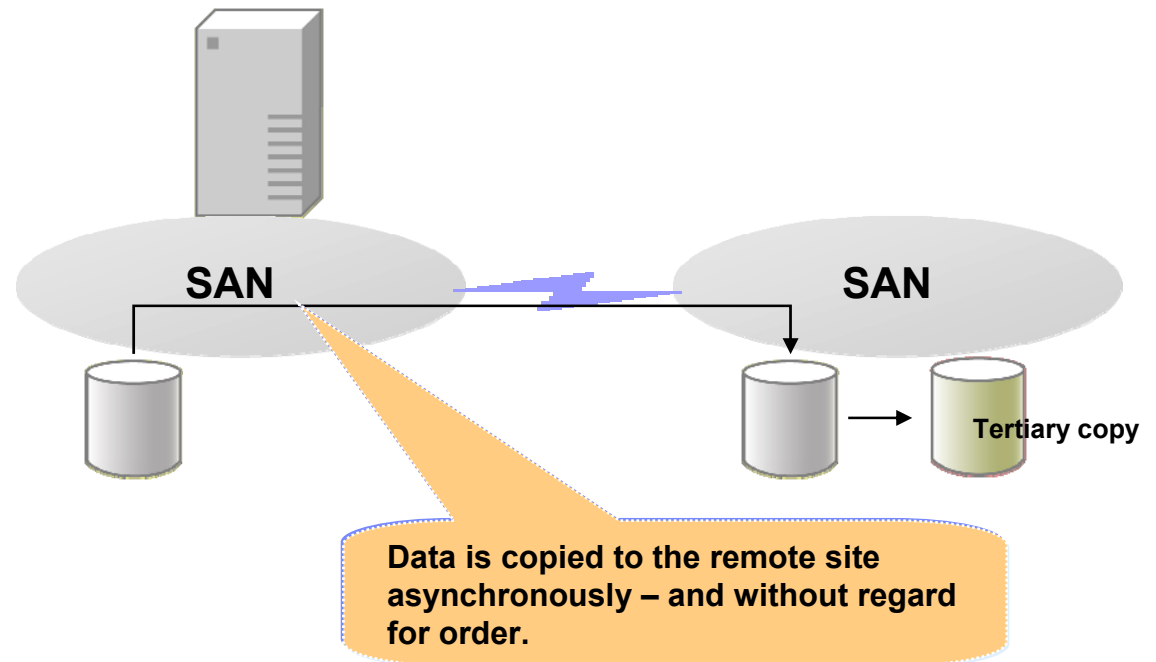
# Review: IBM Global Mirror

- **What is Global Mirror?**
  - ▶ Asynchronous disk mirror
  - ▶ With continuous data integrity
  - ▶ Performed at a volume level
- **What is Global Mirror typically used for?**
  - ▶ For asynchronous mirror with data consistency, at any distance
  - ▶ Creates data consistent copy
- **Objective:**
  - ▶ Continuously create a data consistent copy at remote site, suitable for database restart
- **Available on:**
  - ▶ DS8000, DS6000, ESS; DS4000; N Series
  - ▶ SAN Volume Controller (announced 2Q06)




## Review: IBM Global Copy

- What is Global Copy?
  - ▶ Non-synchronous volume level disk mirror to remote location
    - Data is *not* sent in *data consistent manner*
  - ▶ Data integrity only *at the end of the copy*
    - 'Go-to-Sync', then FlashCopy
  - ▶ Available on:
    - ESS, DS6000, DS8000
    - DS4000



- What is Global Copy typically used for?
  - ▶ Data migration
    - Data center or application moves
  - ▶ Data broadcast
    - For example, transmitting database log files
  - ▶ Typically *\*not\** for D/R by itself

# What is Data Consistency and Consistency Group?



L = Log   J = Journal   D = Data

## ■ Metro Mirror

- Consistency Groups typically use a Freeze/Thaw technique
- Should any Mirroring Link fail, all links are “frozen” keeping all volumes consistent
- Provided for all Enterprise Disks and SVC

## ■ Global Mirror

- Enterprise Disks can create volume consistency in up to eight disk systems (source and target)
- DS4000 provides consistency group up to 64 LUNs
- SVC announced Global Mirror on May 23, 2006

IBM® System Storage™



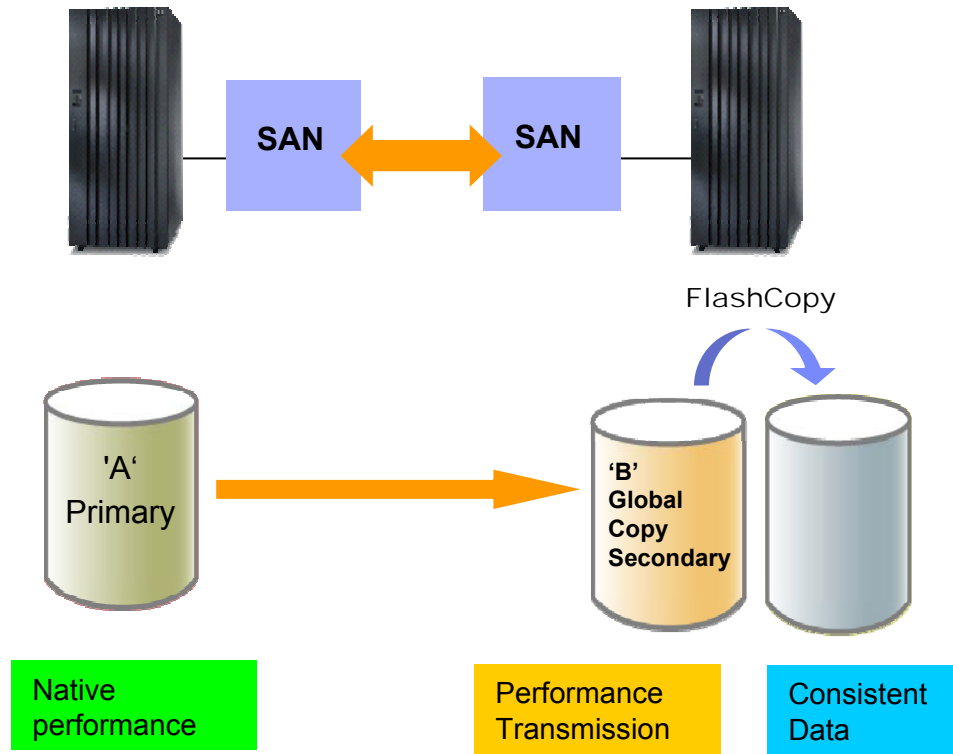
# DS6000, DS8000, ESS Disk Mirroring



© 2007 IBM Corporation

IBM Systems

## IBM Global Mirror Update



- **Continuing worldwide acceptance since GA in May 2004**
  - ▶ Supported on DS8000, DS6000, ESS
- **Performance:**
  - ▶ Maintain production application performance even when bandwidth limited
- **Data Currency**
  - ▶ Maintaining 3-5 seconds (bandwidth permitting)
- **Scalability of Consistency Group:**
  - ▶ Up to 8 subsystems (17 with IBM RPQ)


### European bank

Largest current Global Mirror user: →


- 40 TB, > 15 ESSs in mirrored config
- > 27,000 IO/sec
- Data lag: 4 to 7 seconds (200 KM)
- Heterogeneous zSeries and Open

## IBM Enterprise Three Site Disk Mirroring

Metro/Global Mirror  
A->B->C



z/OS Metro/Global Mirror  
A->B, A->C



- Available today
  - ▶ Enhancements planned for 2006 and 2007
- Two versions:
  - ▶ Cascaded Metro/Global Mirror
  - ▶ Multi-Target z/OS Metro/Global Mirror
- Design objectives:
  - ▶ 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**

IBM® System Storage™



# DS4000 Disk Mirroring




© 2007 IBM Corporation

IBM Systems

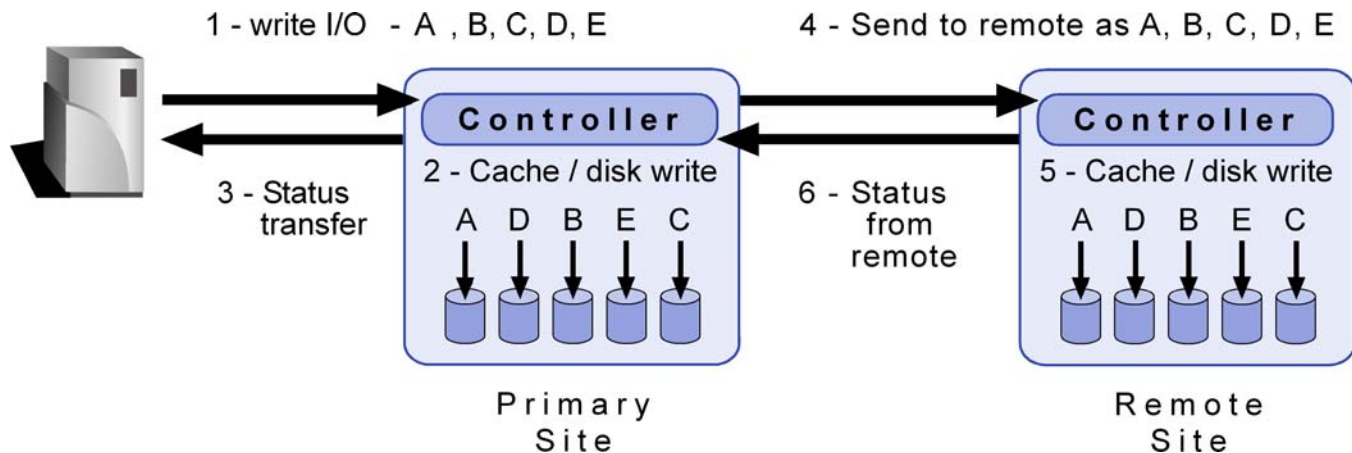
## DS4000 Enhanced Remote Mirroring – Overview

- **Storage-based data replication**
  - ▶ Ongoing, real-time replication of a logical drive from one DS4000or FAStT\* storage subsystem to another
  
- **Comprised of :**
  - ▶ **Metro Mirror** (synchronous writes)
    - *DS4000 Metro Mirror does \*not\* have a Consistency Group option*
  - ▶ **Global Copy** (non-synchronous writes)
  - ▶ **Global Mirror**
    - Supports Consistency Group
    - Maximum CG size: 64 LUNs
  
- **Control software (current release):**
  - ▶ FAStT Storage Manager 9.60



## DS4000 Global Mirror with Consistency Group

- Write operations to the secondary subsystem matches I/O completion order on the local subsystem for all volumes in the consistency group
  - ▶ Method: FIFO (First In, First Out queue)
  - ▶ Maximum size of Consistency Group: 64 LUNs



IBM® System Storage™



# SAN Volume Controller Disk Mirroring




© 2007 IBM Corporation

IBM Systems

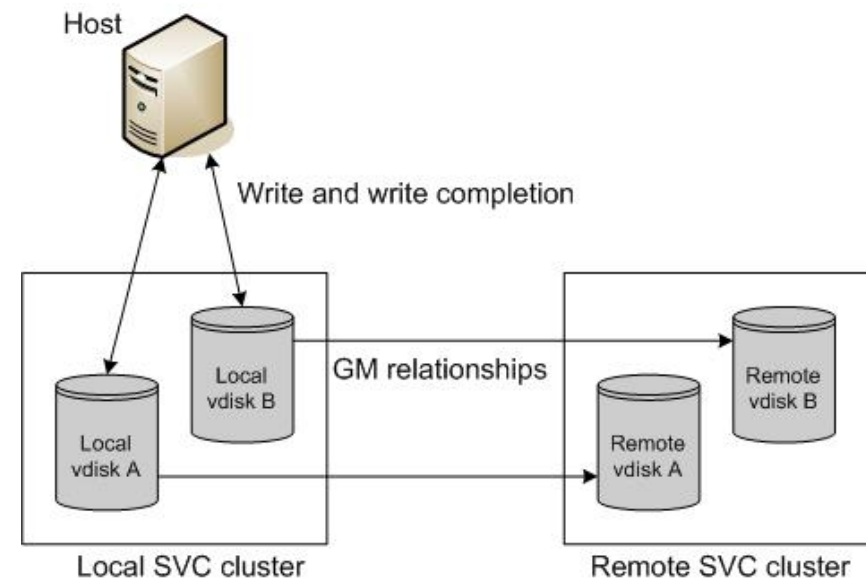
# SAN Volume Controller Disk Mirroring

- SVC disk mirroring
  - Virtual LUNs are copied to virtual LUNs
  - CG spans *nodes* and *I/O Groups*
  - Up to 256 CGs per cluster with SVC 3.1 (was 32 CGs before)
  - SVC microcode supports Consistency Group at the level of the *cluster*
- Inter-*cluster* remote copy is supported
  - One virtual disk comes from each of two clusters
  - One cluster to one cluster only
  - Intra-*cluster* remote copy is supported (“*loop-back*”)



## SAN Volume Controller Global Mirror

- Long distance asynchronous remote mirroring function
- Practically unlimited distances for business continuity
- Does not wait for secondary I/O before completing host I/O
  - ▶ Minimizes performance impact to applications
- Designed to maintain consistent secondary copy at all times
  - ▶ Once initial copy has completed
- Built on Metro Mirror code base
- Metro and Global Mirror delivered as single feature
  - ▶ Offers great implementation flexibility
  - ▶ No additional license charge for existing MM users




# N series Disk Mirroring

## High Level Overview

N3000 Series  
N5000 Series  
N7000 Series

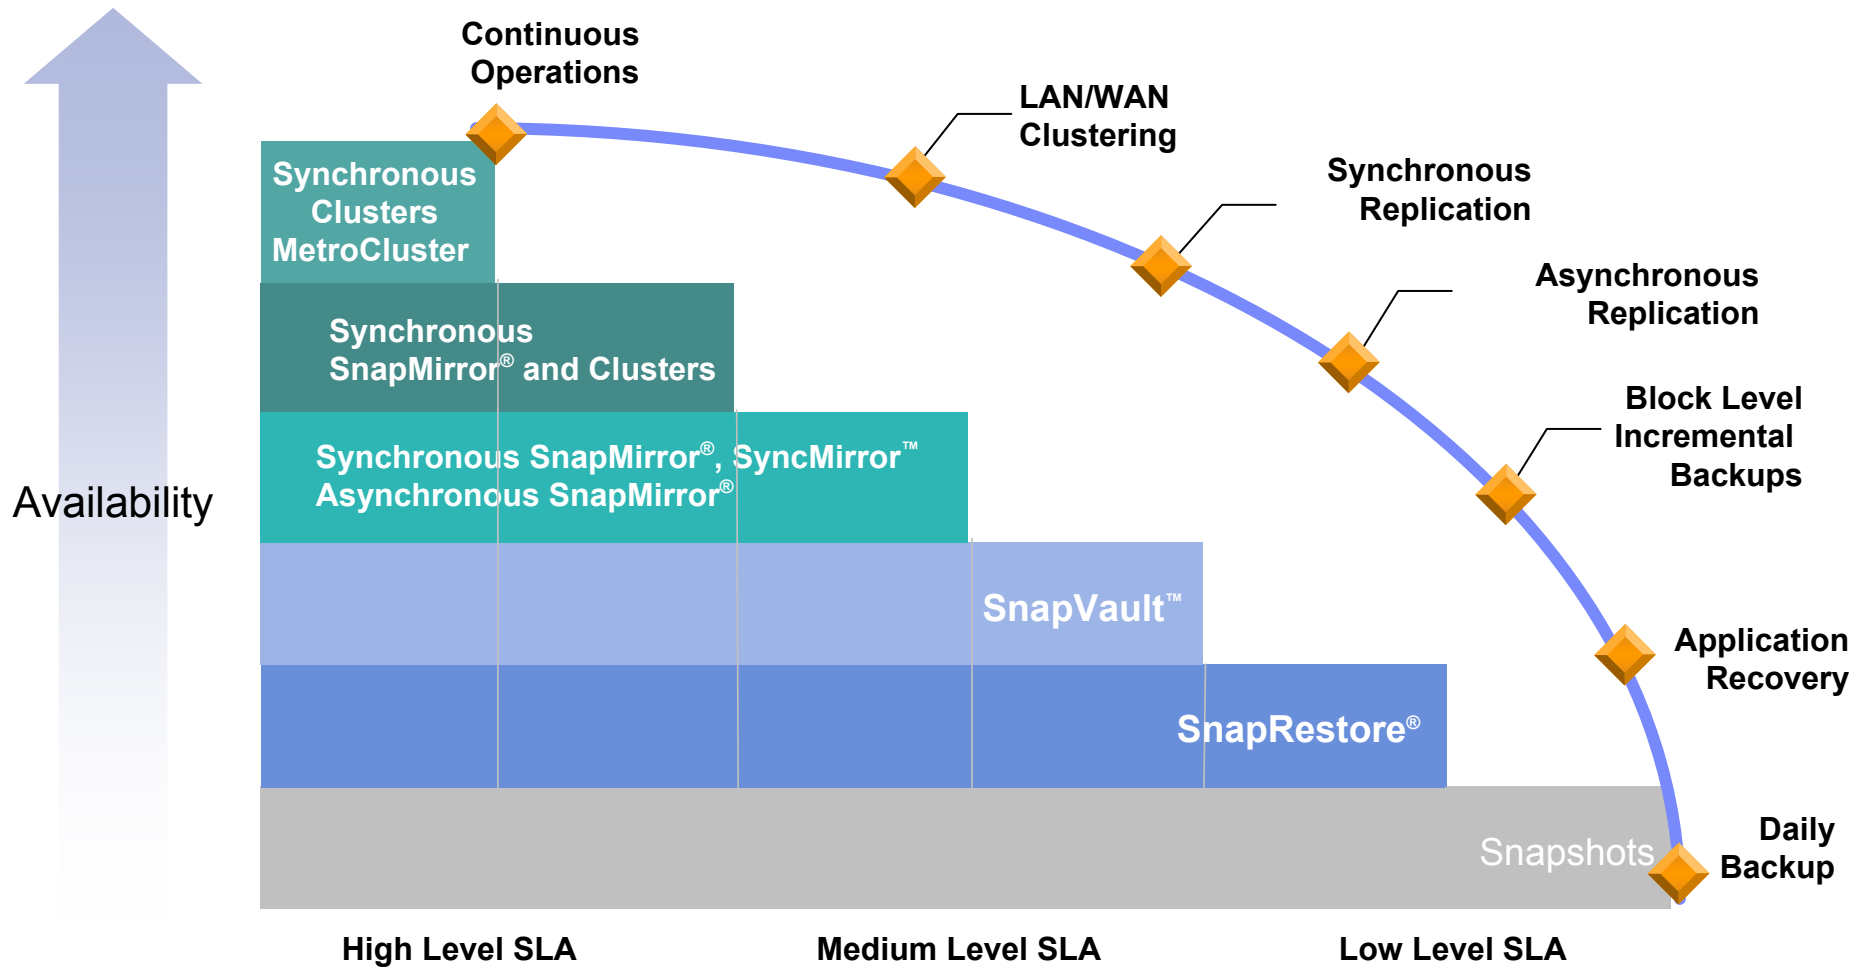


## N Series SnapShot



- Is a read-only, “freeze framed” version of a N Series filer’s file system, frozen at a point in time
  - Only block level changes are stored
  - A volume can maintain up to 255 snapshots concurrently
- Snapshots are readily accessible via “special” subdirectories that appear in the current or active file system
  - Snapshots use no additional disk space when first taken, and Snapshots consume space when the file system changes
- Snapshots can be taken manually or automatically on a schedule

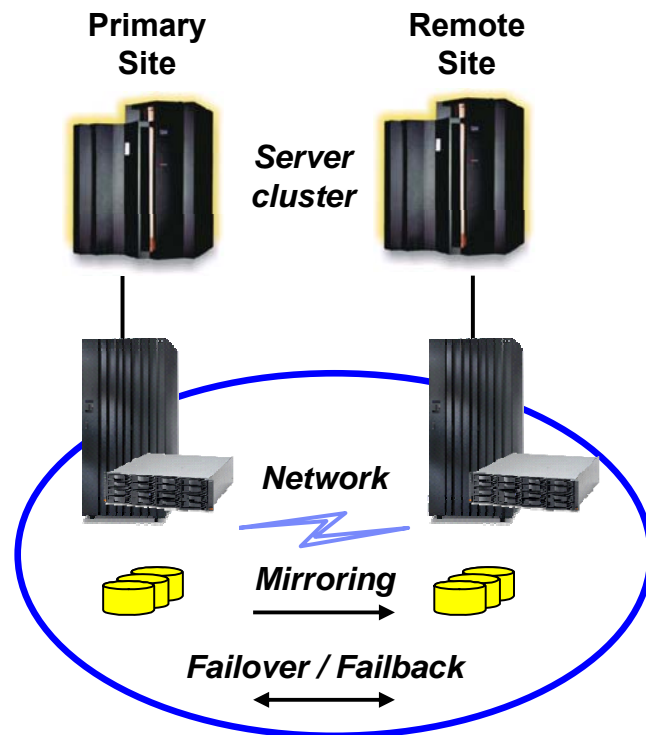
# N Series Disaster Protection features





# IBM Disk Mirroring Control Software

# IBM Metro Mirror -- Open Systems Automated Recovery



**Enterprise Storage**  
**Up to 300 km distance**

## Features

- ▶ **Robust, mission critical high availability for Open Systems environments**
- ▶ **Scalability to thousands of volumes**
- ▶ **Simplified implementation and powerful management of failover / failback**
- ▶ **Worldwide 24X7 support by IBM Global Services**

## Metro Mirror Automation

- ▶ **pSeries: TotalStorage Continuous Availability for AIX (AIX HACMP/XD)**
- ▶ **Heterogeneous: TPC for Replication 3.2**


# IBM Global Mirror -- Open Systems Automated Recovery

## Features


- ▶ Same as Metro Mirror

## TPC for Replication 3.2

- ▶ Robust, mission-critical high availability automation for open server configurations
- ▶ Two-site Metro Mirror, Global Copy, Global Mirror
- ▶ DS8000, DS6000
- ▶ ESS
- ▶ SAN Volume Controller
- ▶ Three site config support for DS8000



# IBM zSeries, zSeries + Open Replication Automated Recovery



**Software Automation**

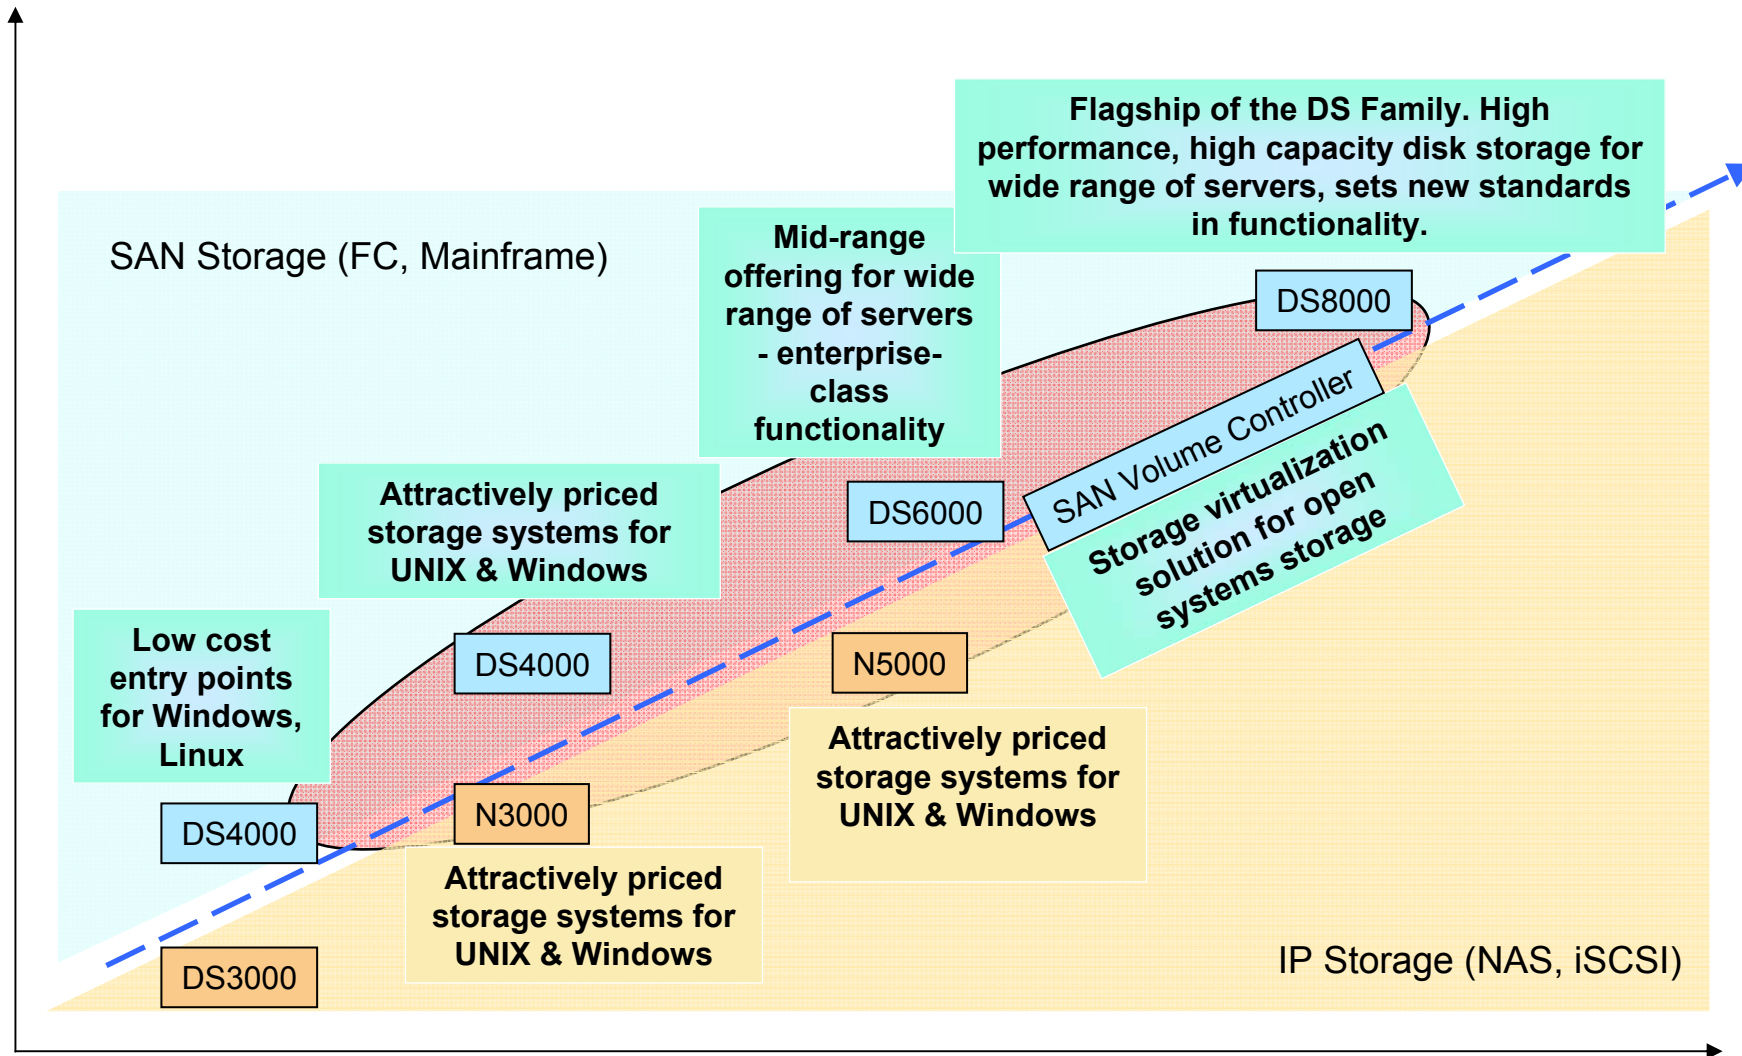
## zSeries and Open Server Solutions

- ▶ GDPS (Geographically Dispersed Parallel Sysplex)
  - IBM Global Services
- ▶ GDPS HyperSwap
- ▶ Geographically Dispersed Parallel Sysplex (GDPS) Open LUN Management -- Metro Mirror
- ▶ GDPS / Global Mirror

## Basic zSeries Server Solutions

- ▶ PPRC Migration Manager
  - IBM Storage Services

# IBM System Storage Disk Positioning



# End to End High Availability IT Business Continuity Solution



## This is....

....the end of Storage System part and the start of

- a. coffee break
- b. meal
- c. dinner
- d. Klemen / 2 long, go home.....
- e. next session

**Ps.**

**Klemen Bačak, IBM Slovenija d.o.o.**

**GSM: +386 40 456607**

[klemen.bacak@si.ibm.com](mailto:klemen.bacak@si.ibm.com)

धन्यवाद

Hindi

謝謝

Traditional Chinese

תודה רבה

Hebrew

*Tesekkurler*

Turkish

Спасибо

Russian

شكراً

Arabic

*Thank You*

English

Gracias

Spanish

Grazie

Italian

ありがとうございました

Japanese

Obrigado

Brazilian Portuguese

ขอบพระคุณ

Thai

谢谢

Simplified Chinese

**Merci**

French

Danke

German

நன்றி

Tamil

감사합니다

Korean

HVALA