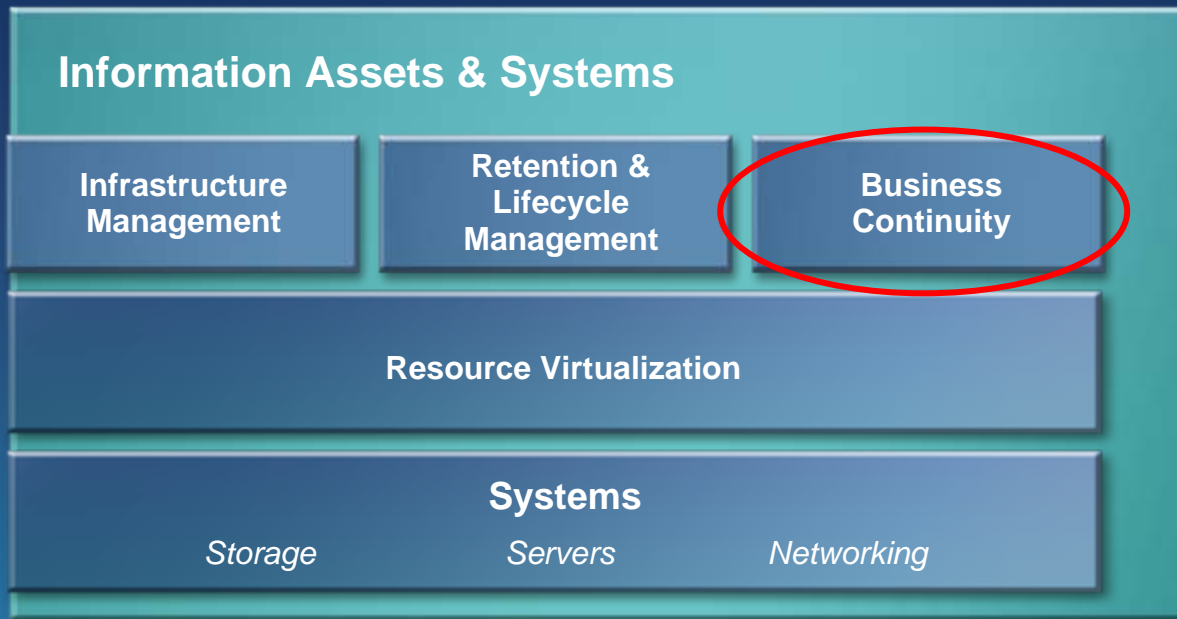


Il Mondo dei Partner
INNOVARE E CRESCERE. INSIEME **2006**
Technical World

Business Continuity – Funzionalità di replica
dati

Stefano Tirasso

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

Session Agenda

1. A technical update of new functions on IBM Disk

Mirroring:

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




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

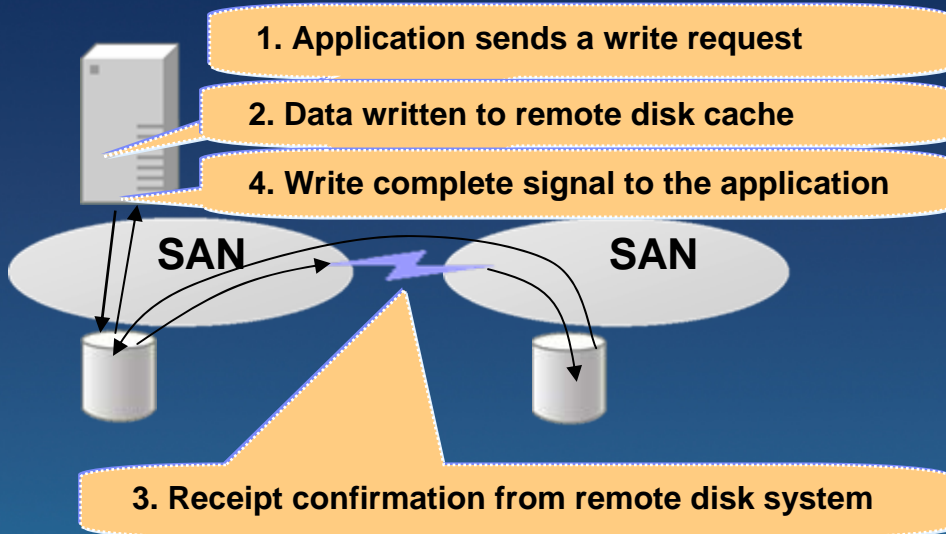
- Asynchronous mirroring (global)
 - DS8, DS6, ESS
 - DS4000
 - SVC (Ann'd 2Q06)
 - N Series

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?

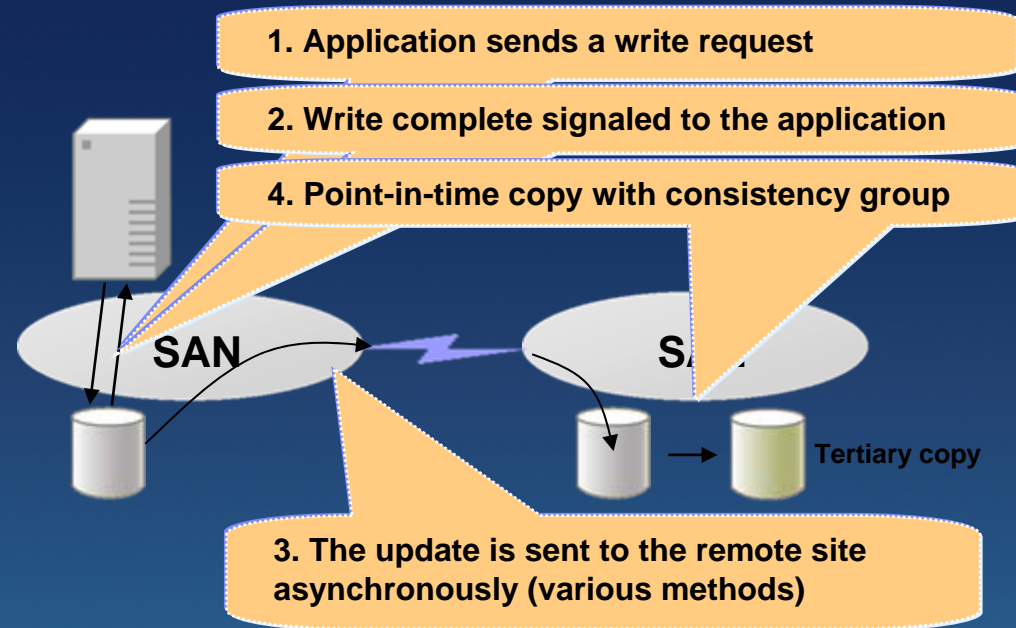
- For asynchronous mirror with data consistency, at any distance
- Creates data consistent copy

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 like disk

Metro Mirror Consistency Group

DS8000, DS6000, ESS 800

- Unlimited
- Multiple disk subsystems (or LSS's) require automation software

SAN Volume Controller

- Maximum of 1,024 LUNs

DS4000

- Does not support Metro Mirror Consistency Group

N series

- Done differently - application level

Global Mirror Consistency Group

DS8000, DS6000, ESS 800

- Up to 8 total disk systems (source + target)
- Any attaching platforms
- Up to 17 primary side systems (RPQ)

SAN Volume Controller

- Available now with SVC 4.1

DS4000

- DS4000 Global Mirror supports up to 64 LUNs in a Consistency Group

N Series

- Done differently – application level

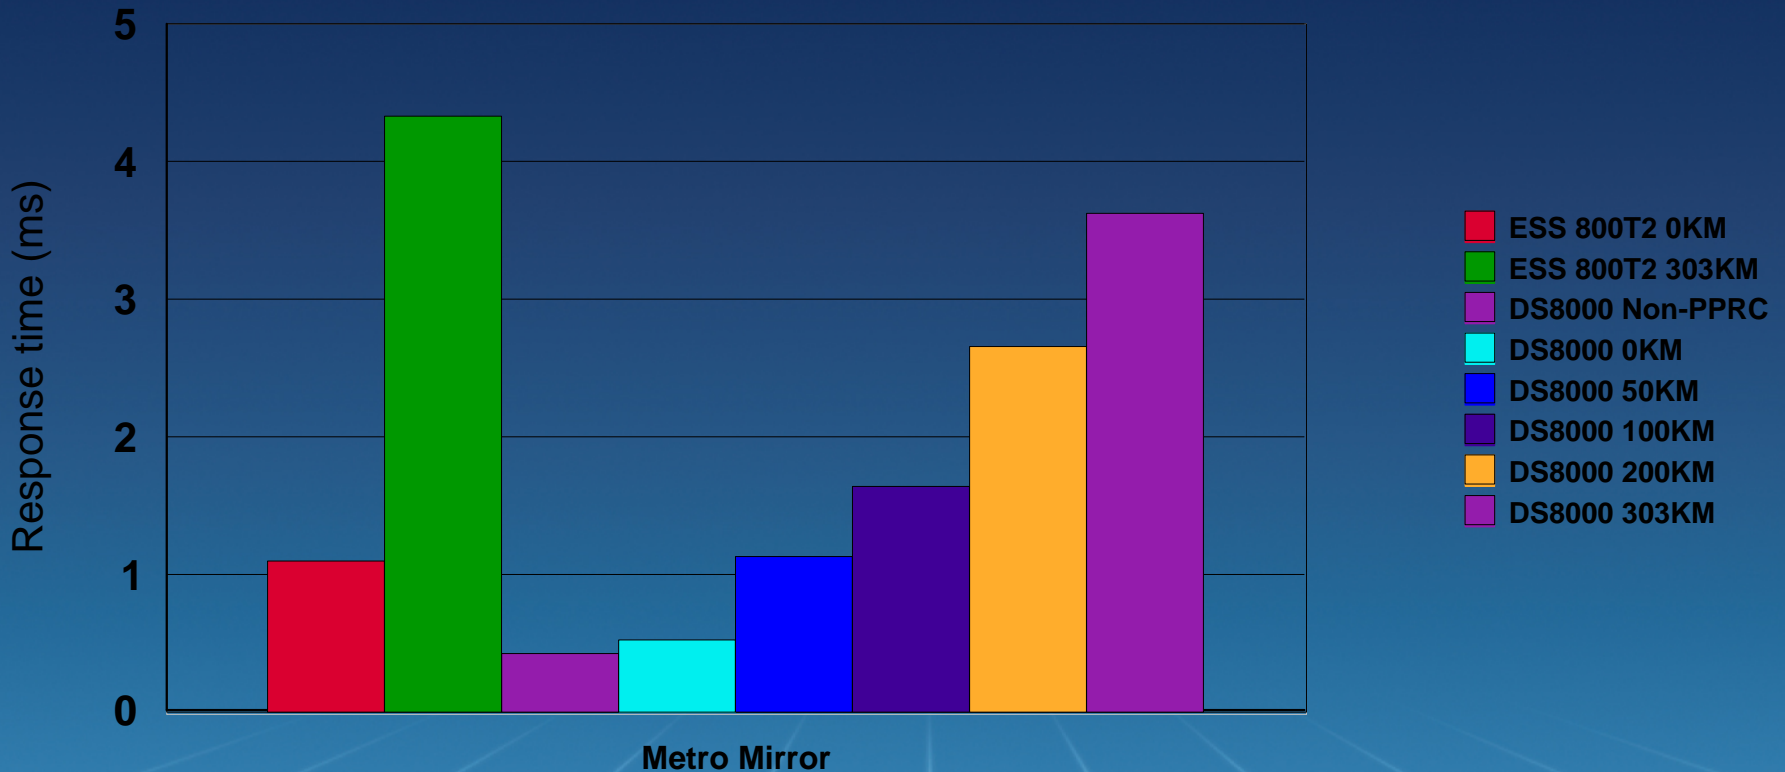
Il Mondo dei Partner
INNOVARE E CRESCERE. INSIEME **2006**
Technical World

DS6000, DS8000, Disk Mirroring



DS8000 Metro Mirror Service Time Improvements

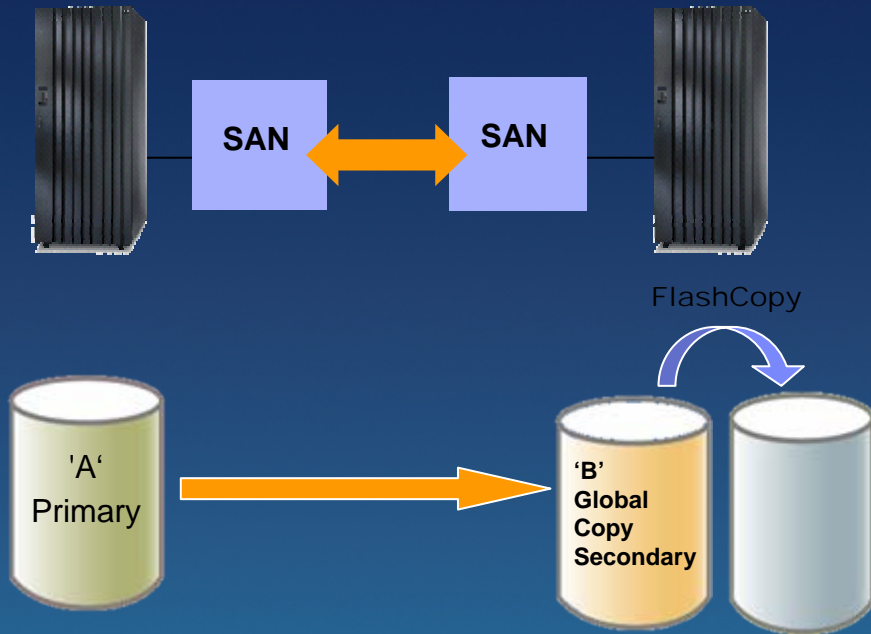
4KB Write Hit Response Time



Response time is for a low IO rate (<1000 io/sec)

303km distance via 1 FCP PPRC Port mapped to (1) CN 2000 Storage Extension Platform OC-12 NIC

IBM Global Mirror Update



Native performance

Performance Transmission

Consistent Data

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

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)

Enhancements:

- **GDPS/Global Mirror:** GA in Oct 2005
- Three site **Metro/Global Mirror:** available since Nov 2005
- Near term **Futures:**
 - More three site enhancements
 - Space Efficient FlashCopy
- **Statement of Direction :** Common Restart Point with z/OS Global Mirror (XRC)

IBM Enterprise Three Site Disk Mirroring

Metro/Global Mirror
A->B->C



Available today

- Enhancements planned for 2006

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

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



Il Mondo dei Partner
INNOVARE E CRESCERE. INSIEME **2006**
Technical World

DS4000 Disk Mirroring



DS4000 Enhanced Remote Mirroring – Overview

Storage-based data replication


- Ongoing, real-time replication of a logical drive from one FASiT or DS4000 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):

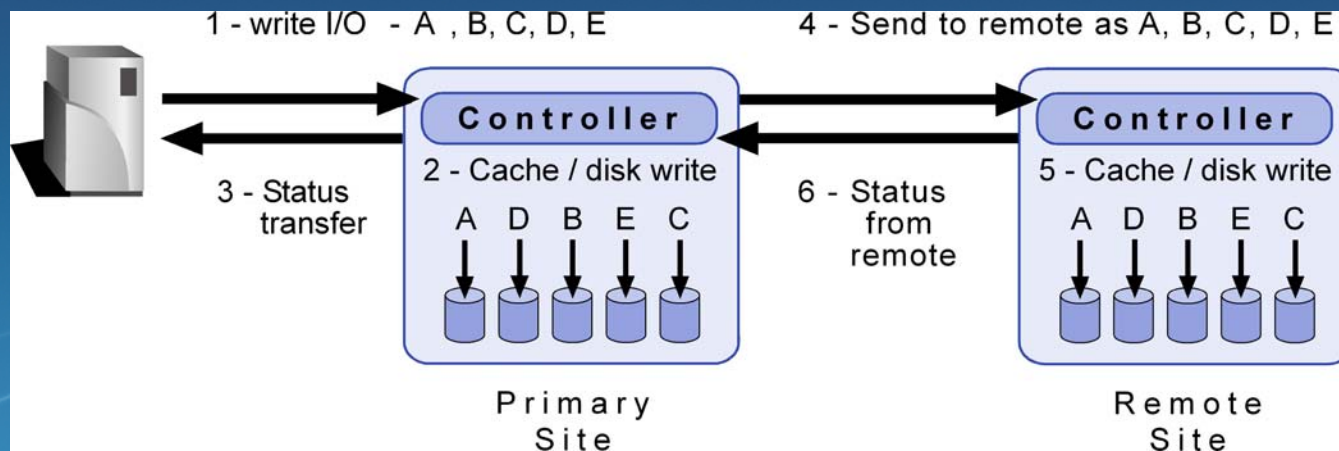
- **DS** Storage Manager 9.xx



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 depth of queue: 128 I/Os per logical drive
- Maximum size of Consistency Group: 64 LUNs





Il Mondo dei Partner
INNOVARE E CRESCERE. INSIEME **2006**
Technical World

**SAN Volume Controller
Disk Mirroring**



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

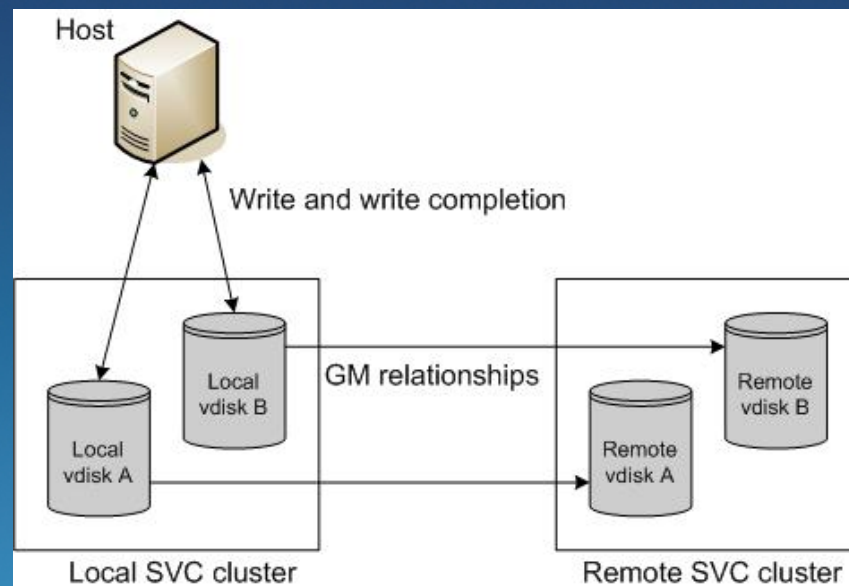
Designed to maintain consistent secondary copy at all times

- 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



SAN Volume Controller Global Mirror Overview

Announced as part of SVC Version 4.1

- Same ordering feature code as for Metro Mirror

Maintains consistent secondary vdisk at all time

- Once background copy has completed
- Does not wait for secondary IO before completing host IO

Secondary volume

- Same size as Primary volume
- No need for journal volume

Same configuration sizes as SVC 3.1.0.x

- 1024 relationships
- 256 Consistency Groups

Limits are aggregate for both SVC Metro Mirror and SVC Global Mirror Same configuration sizes as SVC 3.1.0.x

- i.e. 1024 Metro Mirror + Global Mirror relationships in total

Different protocol than DS8000, DS6000, ESS Global Mirror

- Cannot con-join CG with DS6000, DS8000, ESS

Il Mondo dei Partner

INNOVARE E CRESCERE. INSIEME **2006**

Technical World

N series Disk Mirroring

High Level Overview



N3000
N5000
N7000

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

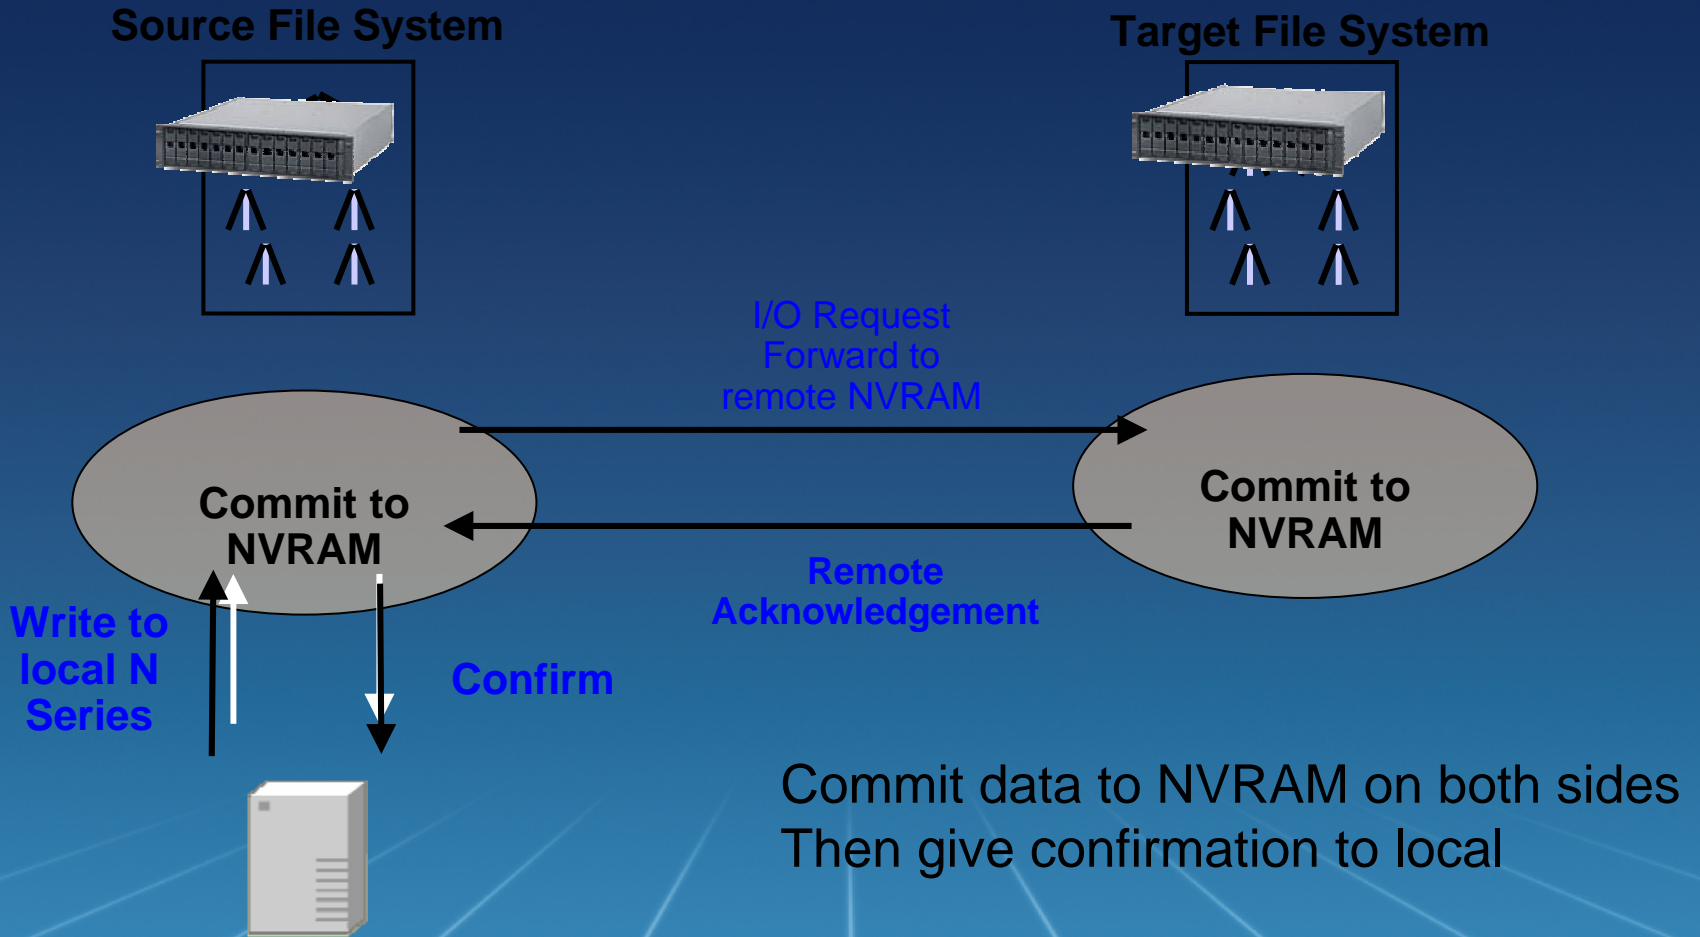
- Snapshots use no additional disk space when first taken, and Snapshots consume space when the file system changes

- Snapshots can be taken manually or automatically on a schedule



SnapMirror Synchronous


Based upon data forwarding to remote N Series NVRAM



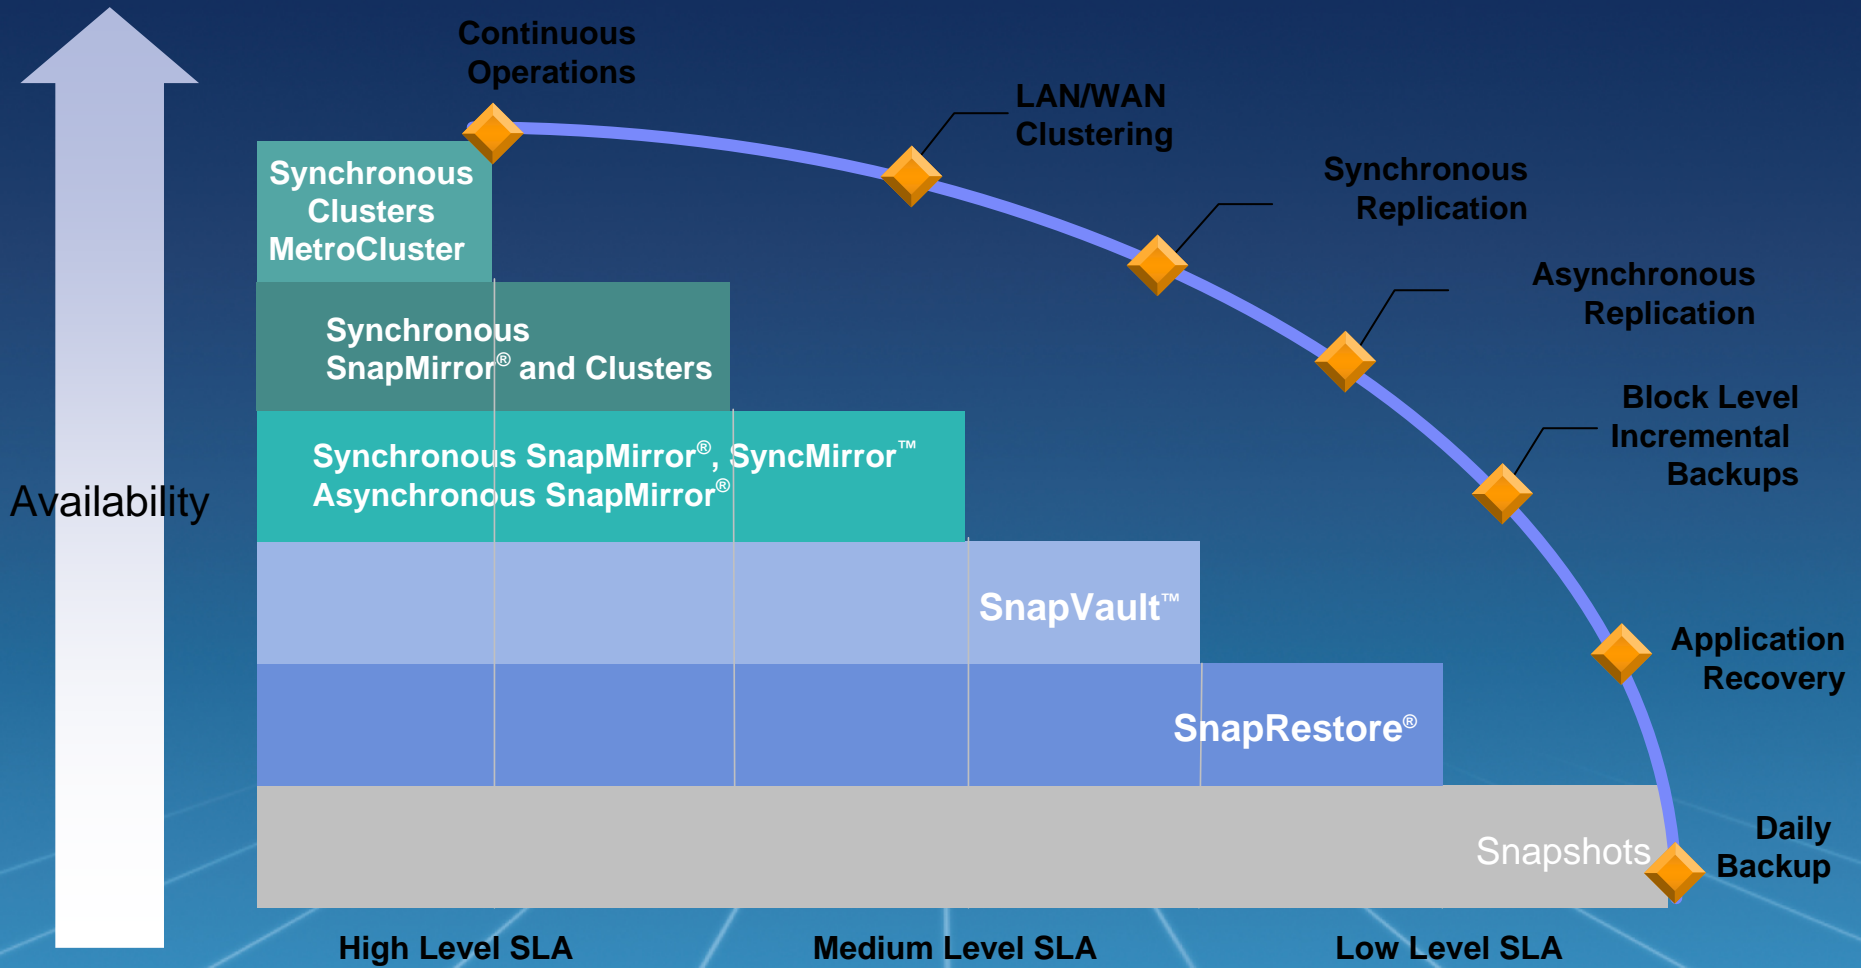
SnapMirror Asynchronous



Changes to source FS or qtree are shipped to **remote N Series NVRAM** on an independent schedule



N Series Disaster Protection features



IBM Disk Mirroring Replication Function Summary

	ESS / DS6000 / DS8000	DS4000	N Series	SVC (multi-vendor storage)
Point-in-Time Copy	FlashCopy®	FlashCopy or VolumeCopy	FlashCopy	FlashCopy
Synchronous Replication	Metro Mirror	Metro Mirror	SnapMirror SyncMirror	Metro Mirror for SVC
Asynchronous Replication	Global Mirror	Global Mirror	SnapMirror	Global Mirror for SVC (2Q06)
Three Site Mirroring	Metro/Global Mirror	n/a	SnapMirror	n/a