BUSINESS CONTINUITY AND DISASTER RECOVERY STRATEGIES

Roberto Lotti
Sr. Account Technology Consultant

EMC² where information lives

Agenda

- EMC Introduction
- Workshop Objectives
- Business drivers for BC/DR
- Terminology
- BC/DR technologies
- EMC BC/DR Solutions
- Case Studies

Workshop Objectives

- Explain Why Customers need a BC/DR Strategy
- Explain Capabilities, Complexity, and Choice
- Understand BC and DR from a technological standpoint
- Describe the main EMC Solutions for BC/DR

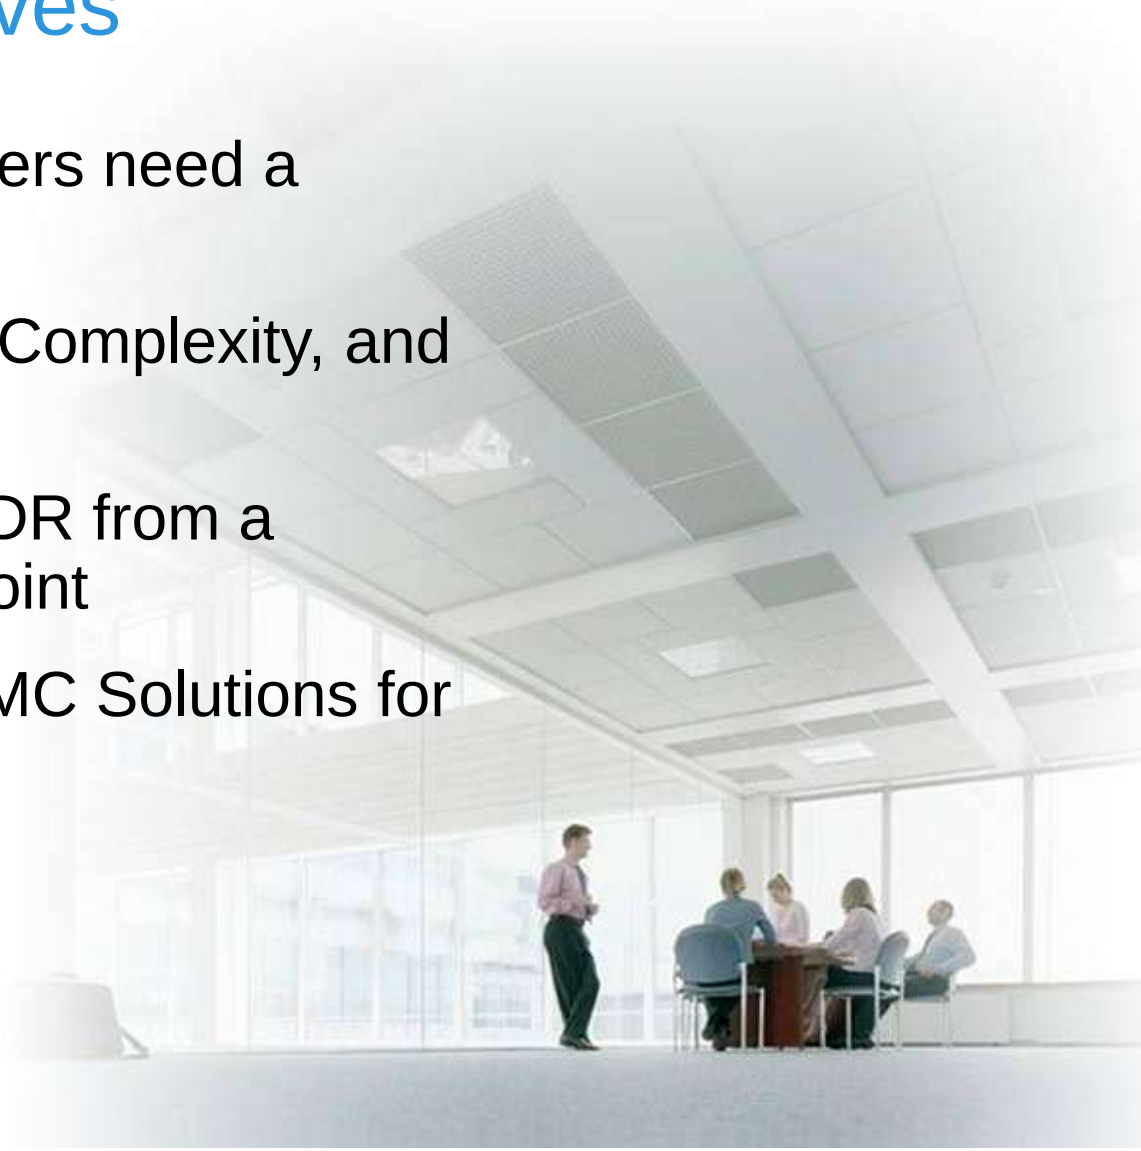

EMC² where information lives®

BC & DR Drivers

Why You Should Care

A study from research firm Frost & Sullivan estimates that North American Business Continuity and Disaster Recovery spending will reach $23.3 billion by 2012.

That is up more than 50 percent from $15.1 billion in 2006.

"We are seeing increased concern from small and mid-sized enterprises about how they protect their data,"

ByteandSwitch
Storage Networking and Beyond

Business Continuity and Disaster Recovery Decision Drivers

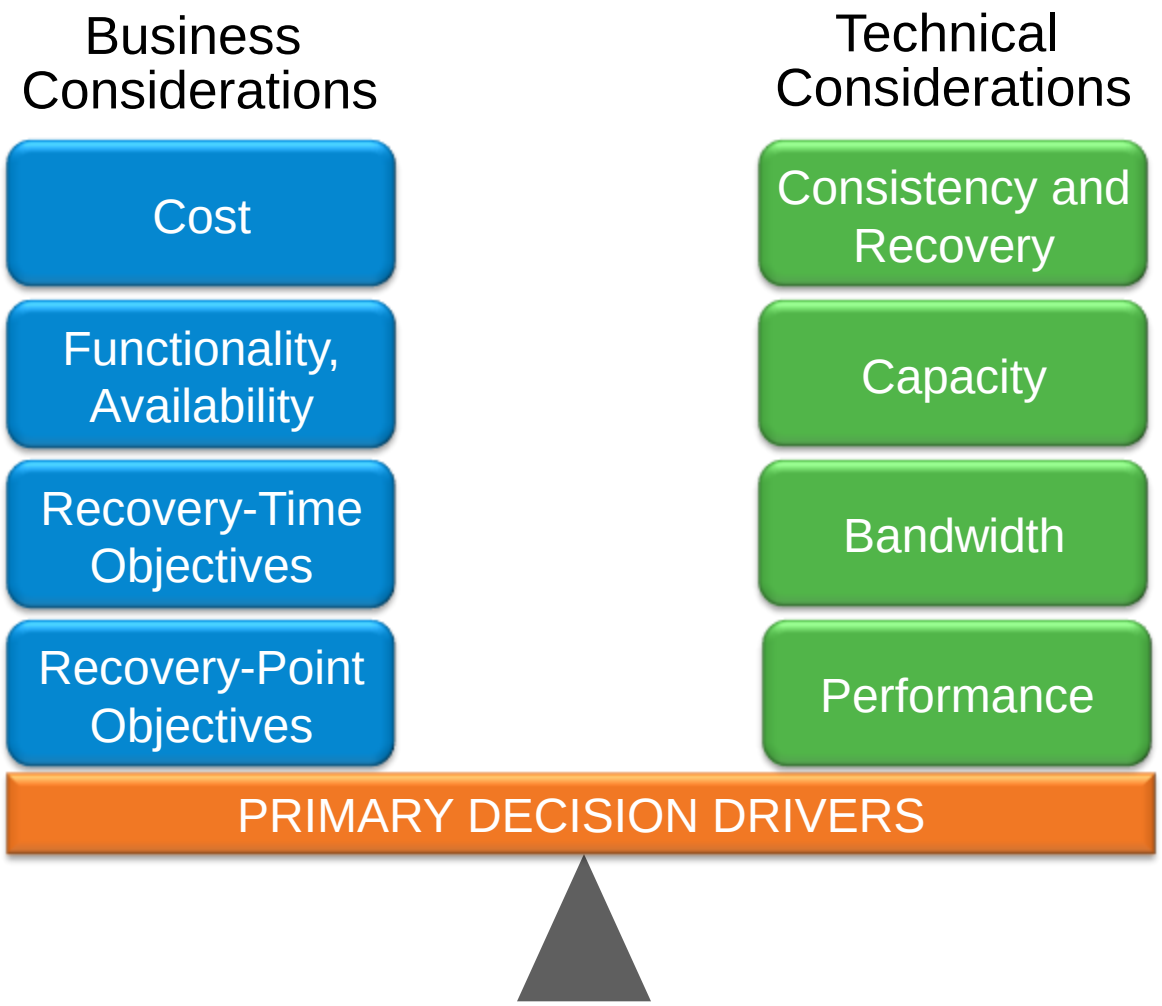

The Cost of Downtime Per Hour By Industry

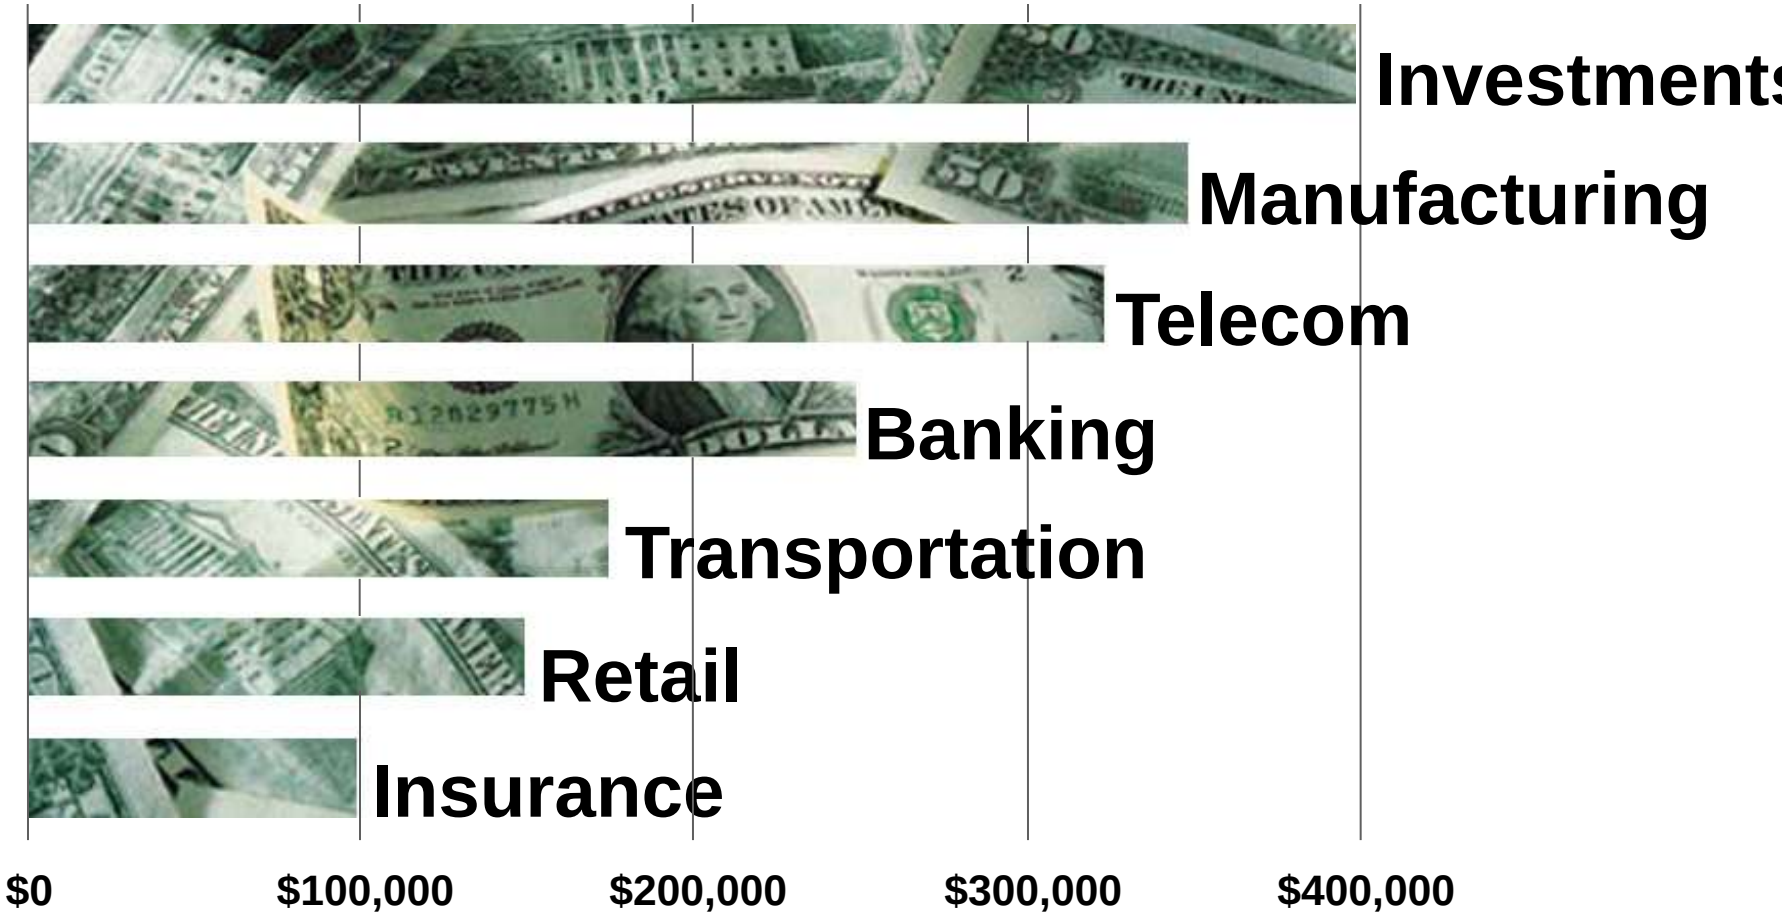

Source: AMR Research

The Impact of Business Continuity

Productivity Impact

- Employees affected
  - Email !
  - Systems

Revenue Impact

- Direct + Indirect losses
- Compensatory payments
- Lost future revenue

Brand Impact

- Customers
- Suppliers
- Financial markets
- Banks
- Business partners
- The Media

Financial Impact

- Revenue recognition
- Cash flow

Business Continuity Considerations

- What are your company's most critical processes and data needs?
- How much data can you afford to lose?
- How quickly do you need to restore your critical processes?
- How vulnerable are your operations to disasters?

Events that Impact Information Availability

Events that require a data center move: <1% of occurrences

- Examples?

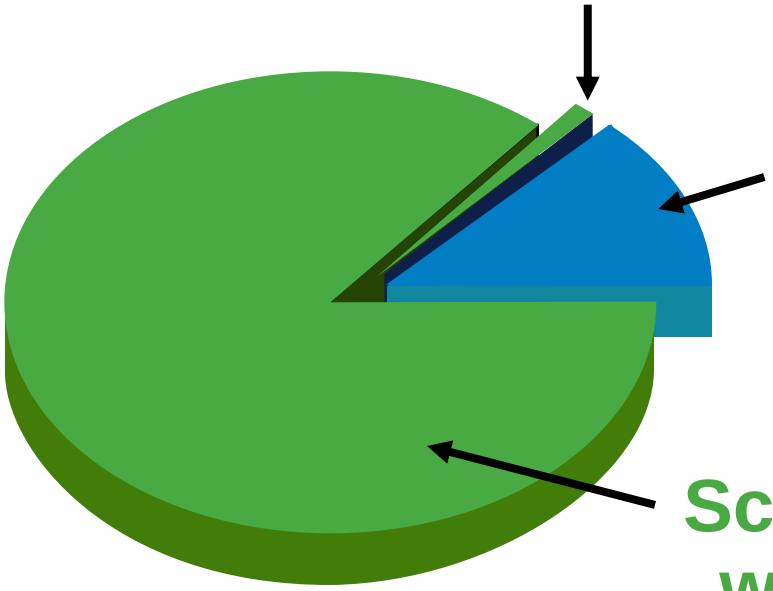

Unscheduled events/failures: 15%

- Examples?

Scheduled events/competing workloads: 85%

- Examples?

Events that Impact Information Availability

Events that require a data center move: <1% of occurrences

- Disaster events
  - Fire, flood, storms, etc.
- Data center move or relocation
- Workload relocation

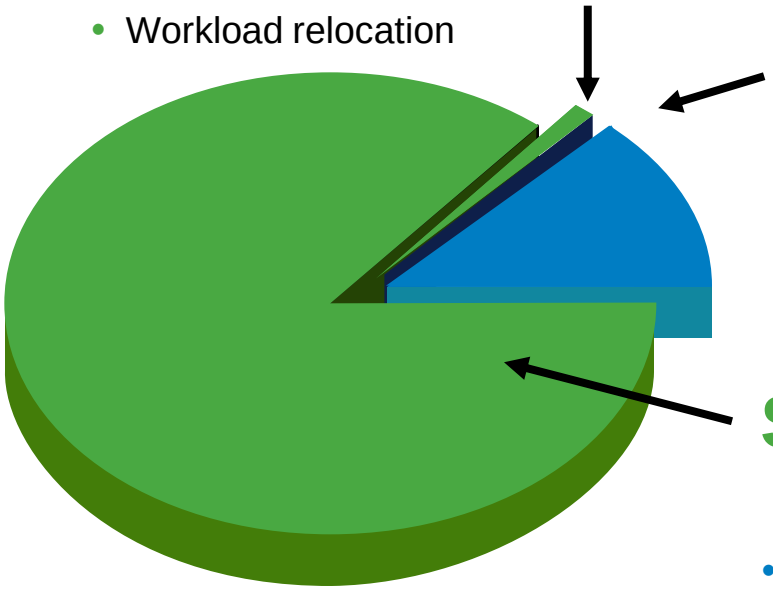

Unscheduled events/failures: 15%

- Server failure
- Application failure
- Network / storage failure
- Processing or operator error

Scheduled events/competing workloads: 85%

- Maintenance and migrations
- Backup and restore
- Batch processing
- Reporting
- Data warehouse extract, build, and load

Terminology

A Key Differentiation

Understand the difference between
Disaster Recovery (DR) and Business Continuity (BC)

- **Disaster Recovery:** Restoring IT operations following a site failure
- **Business Continuity:** Reducing or eliminating application downtime

Protecting Information is a Business Decision

Recovery-point objective (RPO):

How much data can you afford to lose, can you determine a sync point

Recovery-time objective (RTO):

How long can you afford to idle your business and survive?

- Fast recovery times enable continuous business operations
- Slow recovery times—or data loss—translates into *Business Recovery*

Business Decision: ***(RPO + RTO) < Acceptable Business Risk***

Balancing Business Requirements and Cost

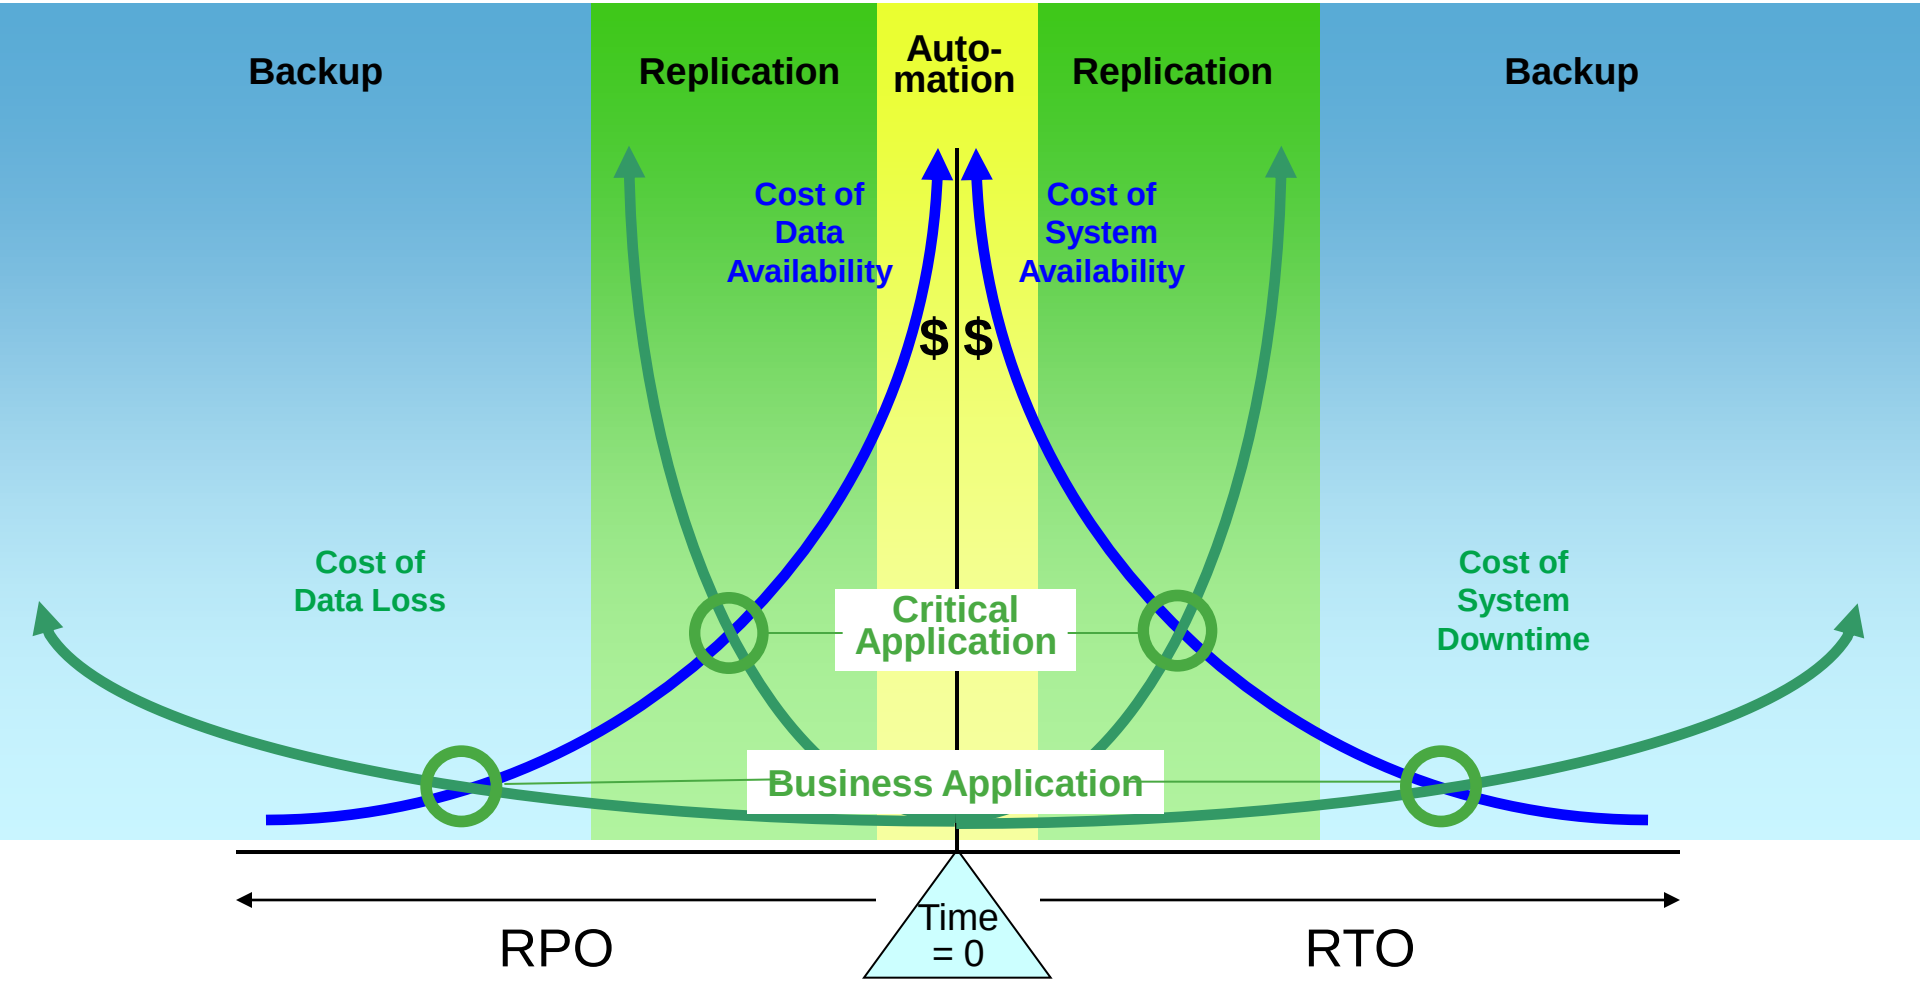

The point in time to which critical data must be restored to following an interruption before its loss severely impacts the organization

The maximum acceptable length of time that can elapse following an interruption to the operations of a business function before its absence severely impacts the organization

Typical Service Levels by Replication Technology

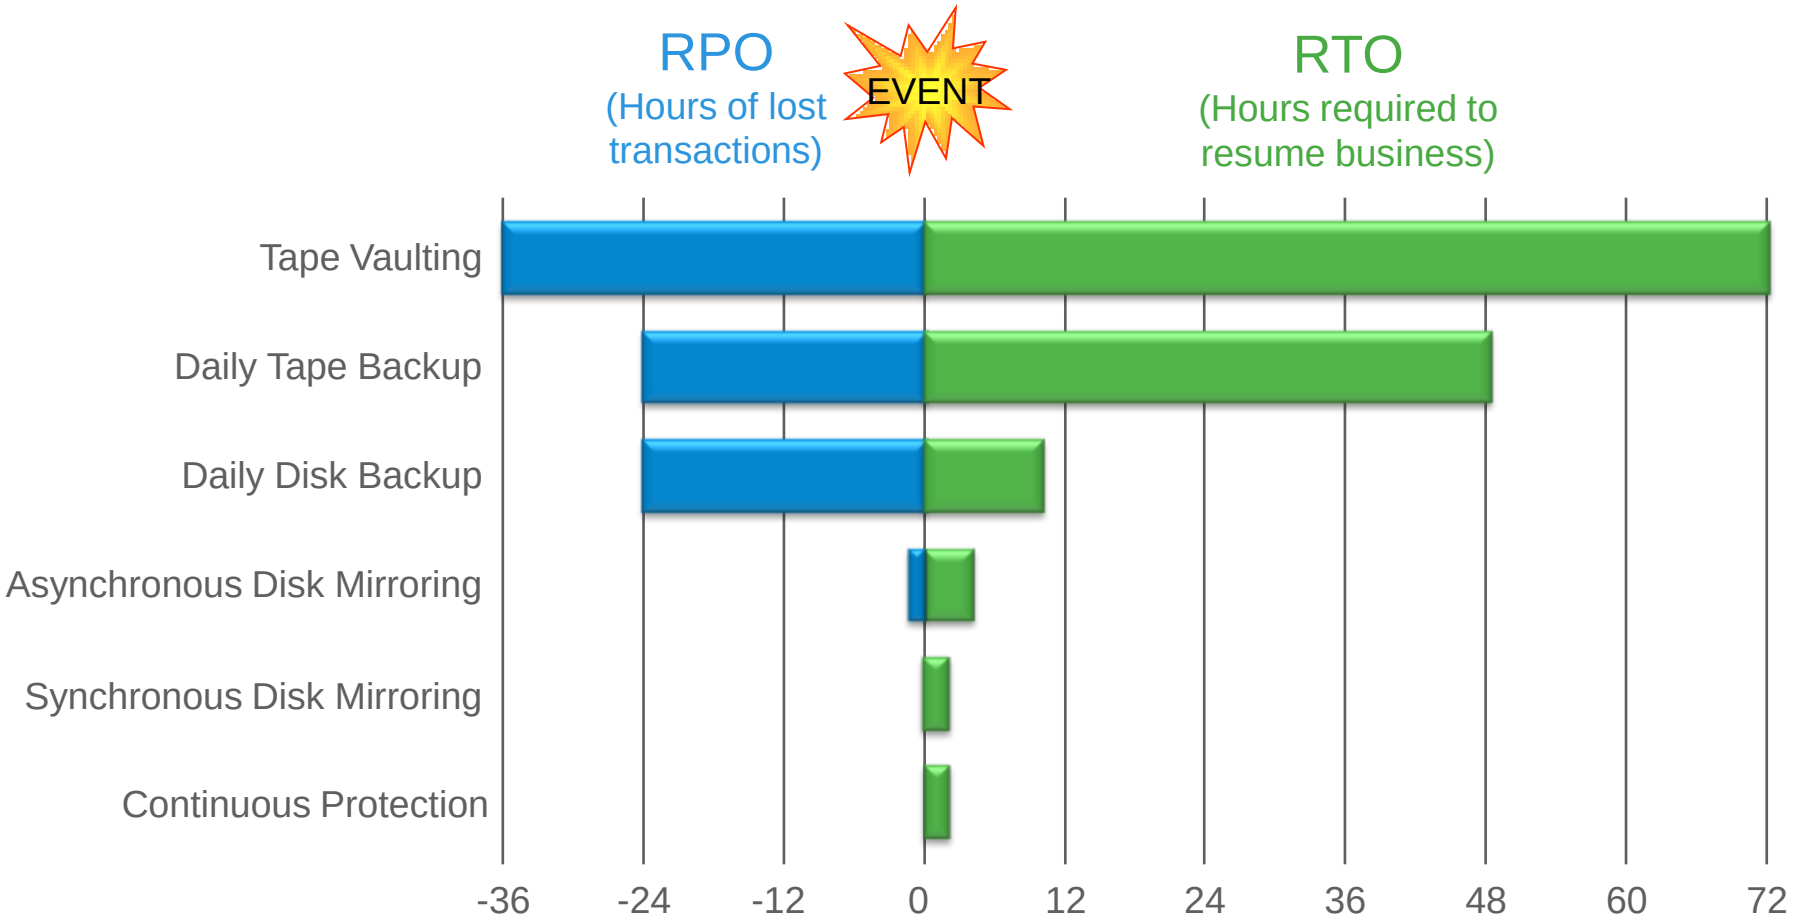

Basic Replica Tipologies

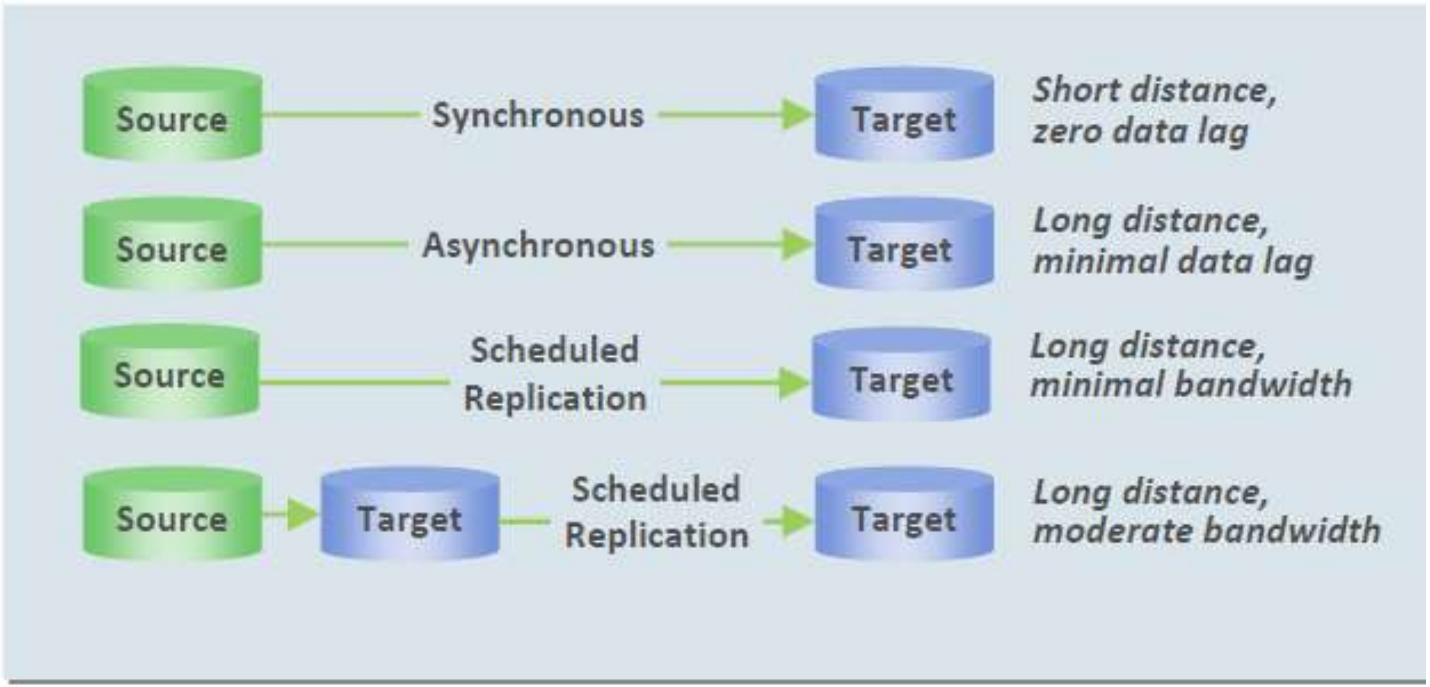

Replica Sizing Parameters

Characterization of I/O Workload	Channel Characterization
• Characteristic sizes [bytes] • I/O rate at each characteristic size [1/s] • % read vs. write at each size [%] • Probability of cache hit [%] • Network activity peaks [distribution of request deltas]	• Throughput (better for large sequential access) [bytes/s] • IOPS (better for small random access) [1/s] • Latency (Service Time) [s] • Queue Length [1]

Replica protocols

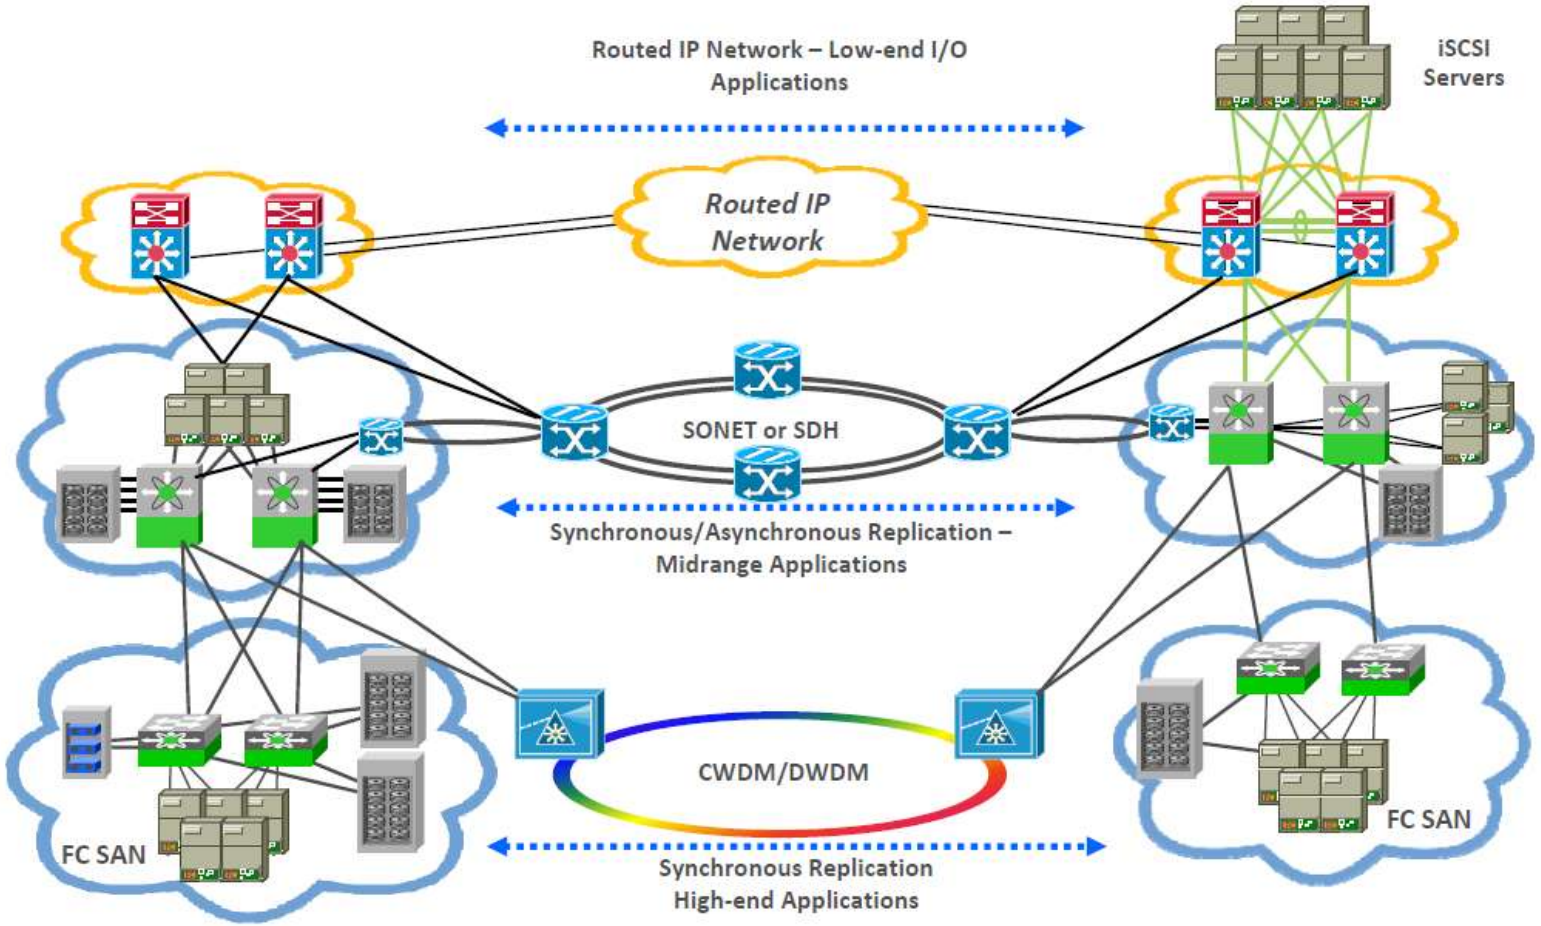

Latency

- Latency in dark fiber is ~ 5ns/m or 5us/km (One 10km link can have 50us latency)
- Worst ….. A round-trip time (RTT) can be 100us
- Latency over SONET/SDH is higher
- Latency over IP networks is generally much higher
- Latency directly impacts application performance:
  - Increased idle-time while application is waiting for read data
  - Increased idle-time while application is waiting for write acknowledgement
  - Reduces I/Os per second (IOPS)

	Dark Fiber	CWDM	DWDM	SONET / SDH
Distance	Up to 10km	Up to 100km	Up to 200km	1000s of km
Latency	5us/km	5us/km	Low	Medium
Bandwidth	1000Mbps on 10GFC	8 channels x 2Gbps	>32 channels x 10Gbps	OC-192 ~ 10Gbps
Reliability	$1x10^{12}$	$1x10^{12}$	High	High
Cost	Low	Low	High	Medium

Line Type	Line Speed / Overall Capacity
T1	1.5 Mbits
T3 (DS3)	45 Mbits
OC3	155 Mbits
OC12	655 Mbits
OC48	2.5 Gbits
OC192	10 Gbits

Application Consistency

Which type of consistency will be required for the applications that you're going to protect ?

Crash Consistency

This is the equivalent of pulling the power from a server while the applications are running, and then powering up the server again. Replication solutions that have limited knowledge of the applications are easier to put together. During recovery you are reliant on the application's capability to start up on its own merits, or possibly with some intervention. Following a fail-over, the data will not have transactional consistence, if transactions were in-flight at the time of the failure. In most cases what occurs is that once the application or database is restarted, the incomplete transactions are identified and the updates relating to these transactions are "backed-out" or some extra procedures or tools may be required.

Application Consistency

There are ways of ensuring that if a copy is taken, or if a system is shut down, all necessary transactions within a database are complete and caches are flushed inorder to maintain consistency. Scripts can be written, following best practice for each application to ensure processes take place in a certain order, or there are applications which can automate these procedures for each application. Some technologies use agents which are application specific. The choice is again down to importance of data, RPOs, RTO's and the available budgets within the organisation.

Deployment Options

Bi-directional

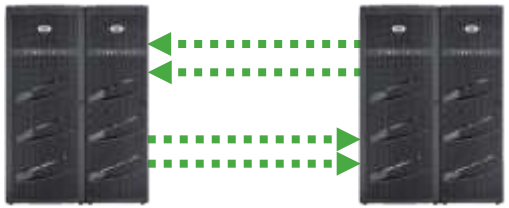

Source/Target Swap

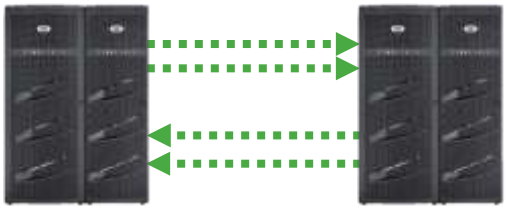

Concurrent and Dynamic

One to Many

Many to One

Many to Many

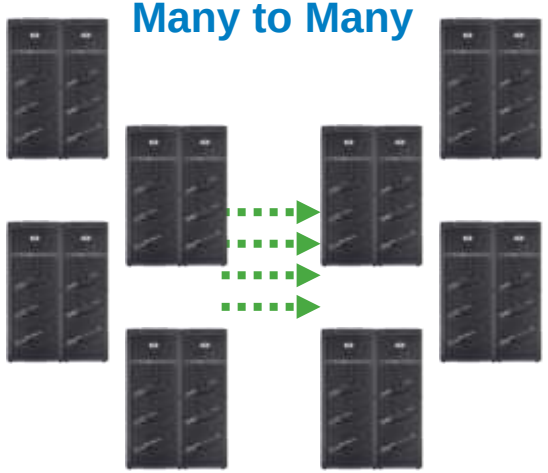

BC/DR: Replica Deployment Options

Synchronous Replication

- No data exposure
- Some performance impact
- Limited distance

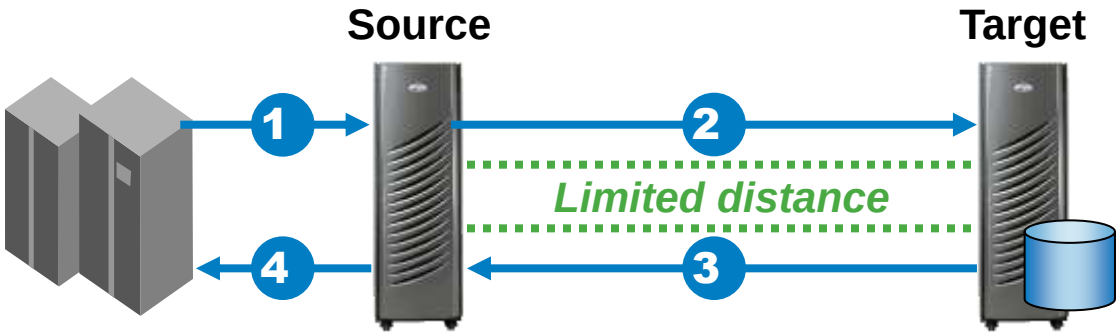

Asynchronous Replication

- Seconds of data exposure
- No performance impact
- Unlimited distance

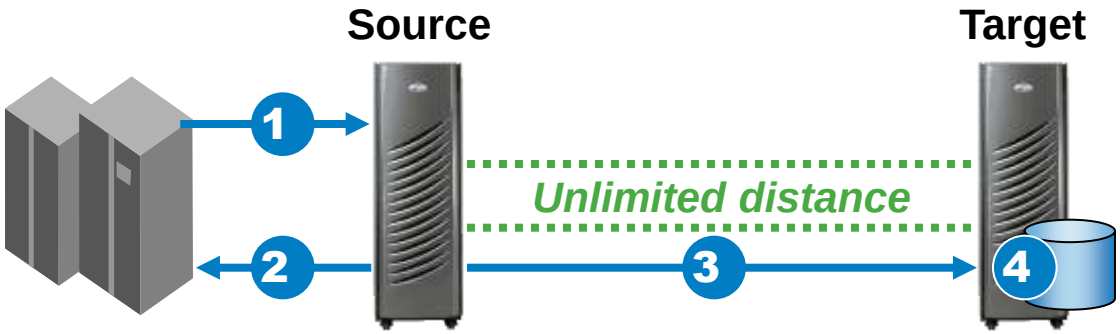

Deployment Options to Fit Your Needs

Mode Change feature

- Change between S and A with consistent remote copies
- Reduces bandwidth requirements
- Increases performance of host applications during peak periods

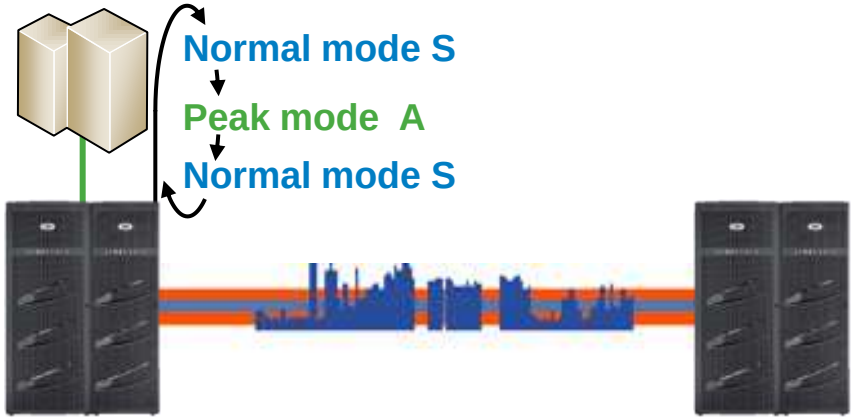

Concurrent

- No data loss on single point of failure
- Geographically dispersed protection
- Unlimited distance

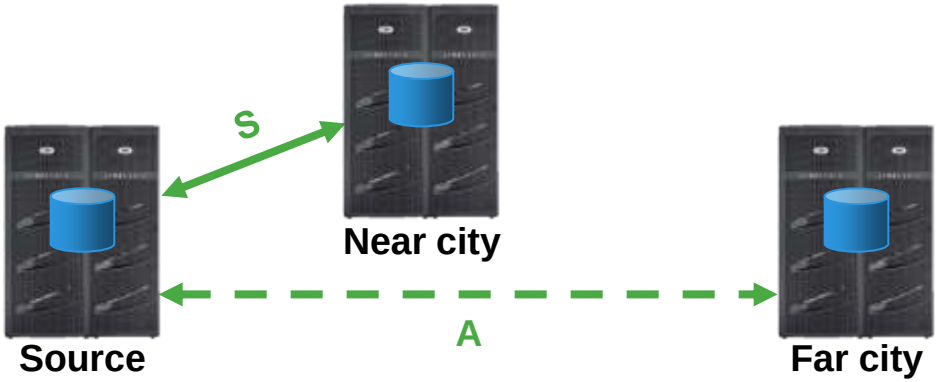

BC/DR: Site Deployment Options

Multisite Replication

- No data loss on single point of failure
- Geographically dispersed protection
- Unlimited distance
- Ability to enable new Async between remote locations

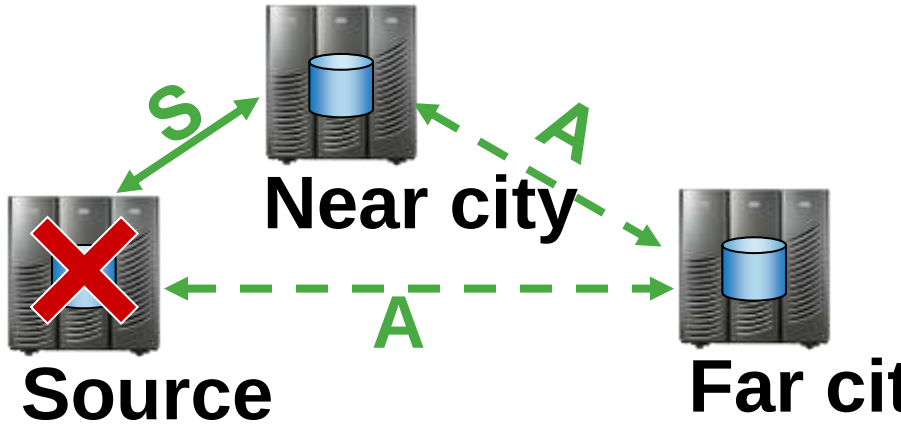

Multisite Replication

- Zero or hours of data exposure
- No performance impact
- Unlimited distance
- Requires BCVs

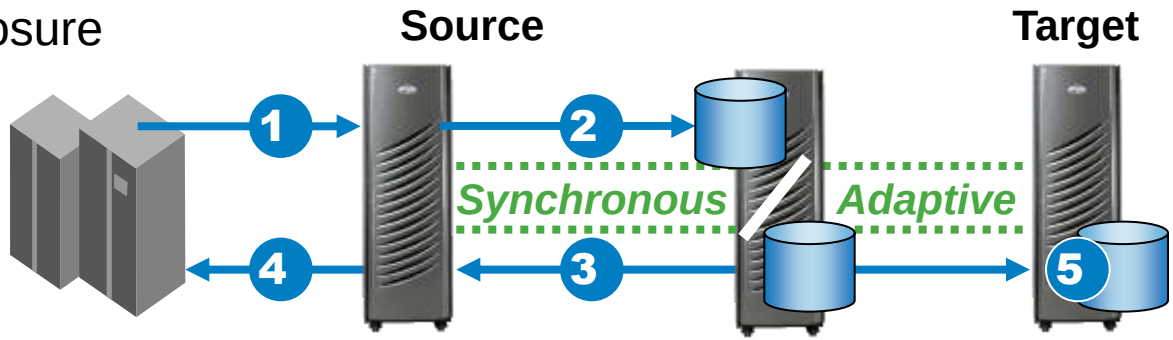

Overview on Data Consistency

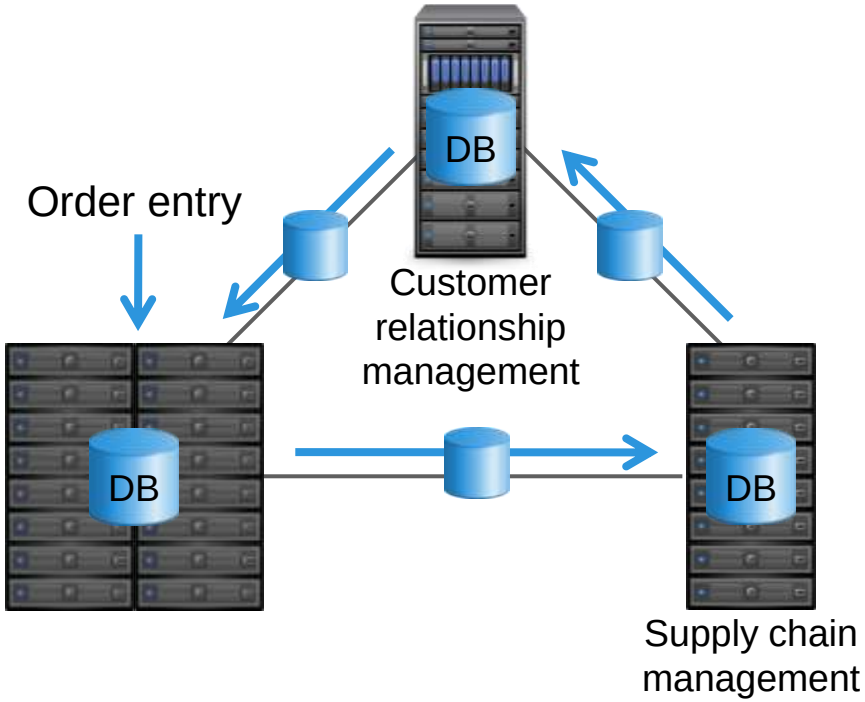

- Applications and data are interrelated (federated)
- All data movement must be stopped/
  started at the same point in time
- To restart applications, you must have *all* the data—not parts of it
- Recovery requires dependent-write consistency across *all* volumes and systems

Systems share information…

How do you get a consistent view of it all?

Data Consistency: The Solution

Getting All the Data at the Same Time

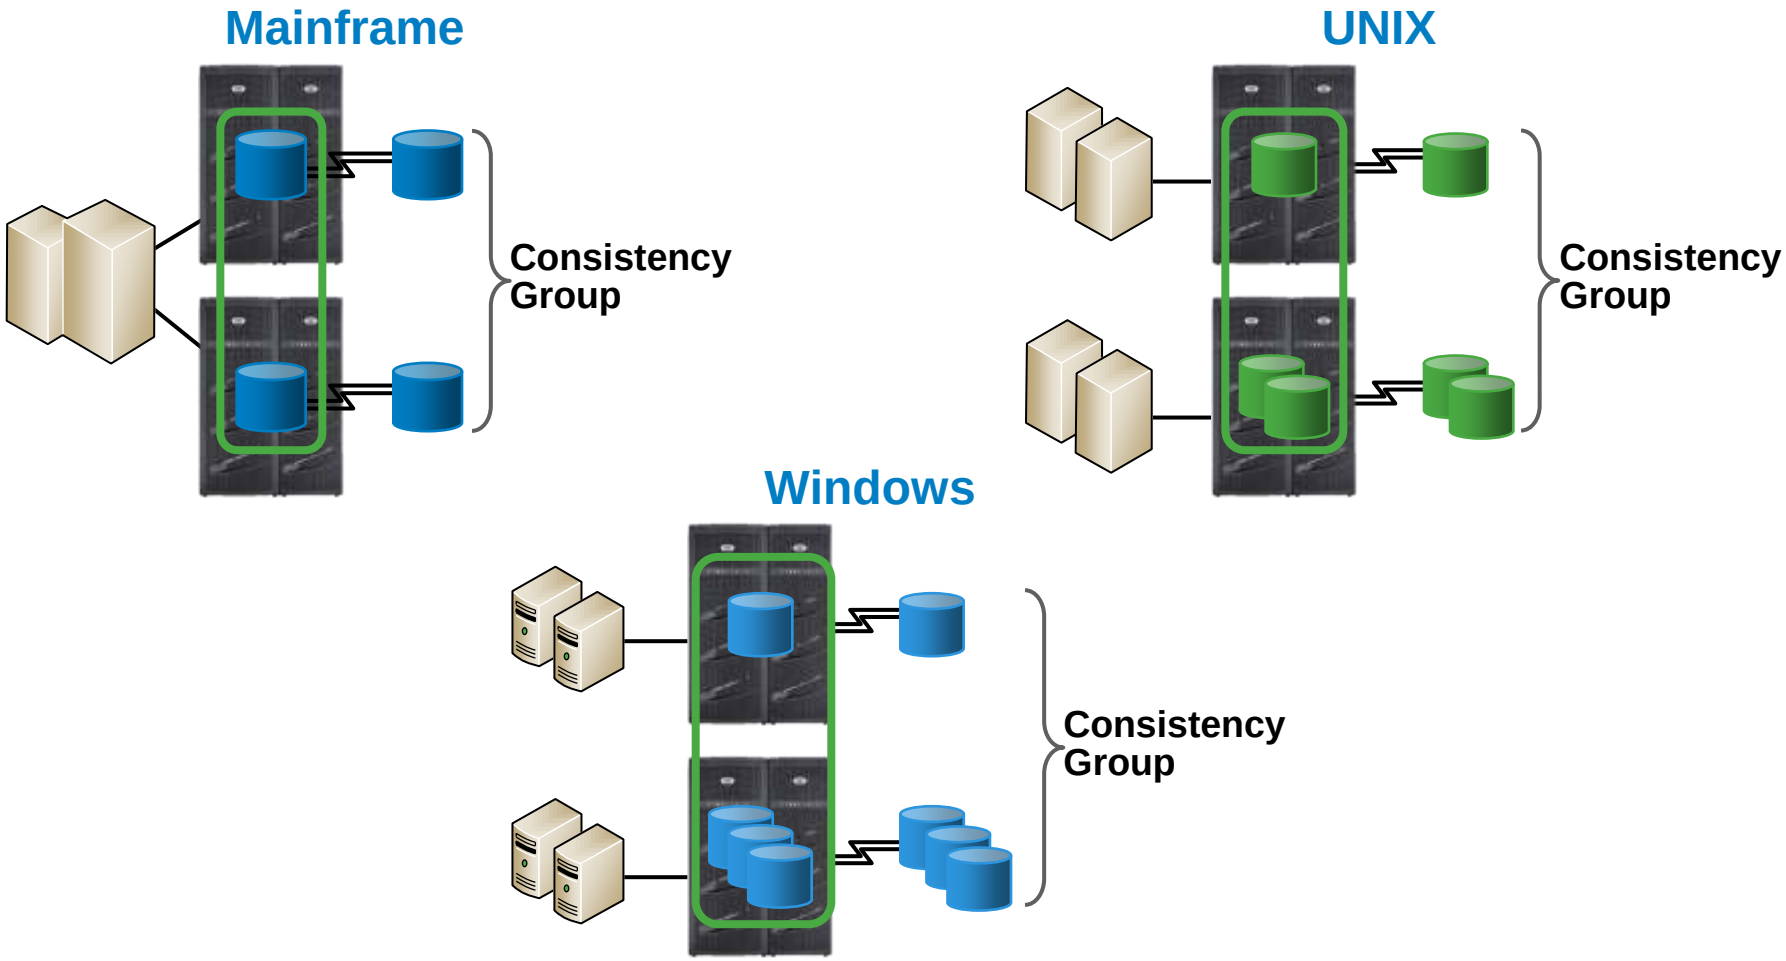

Data Consistency: The Enterprise Solution

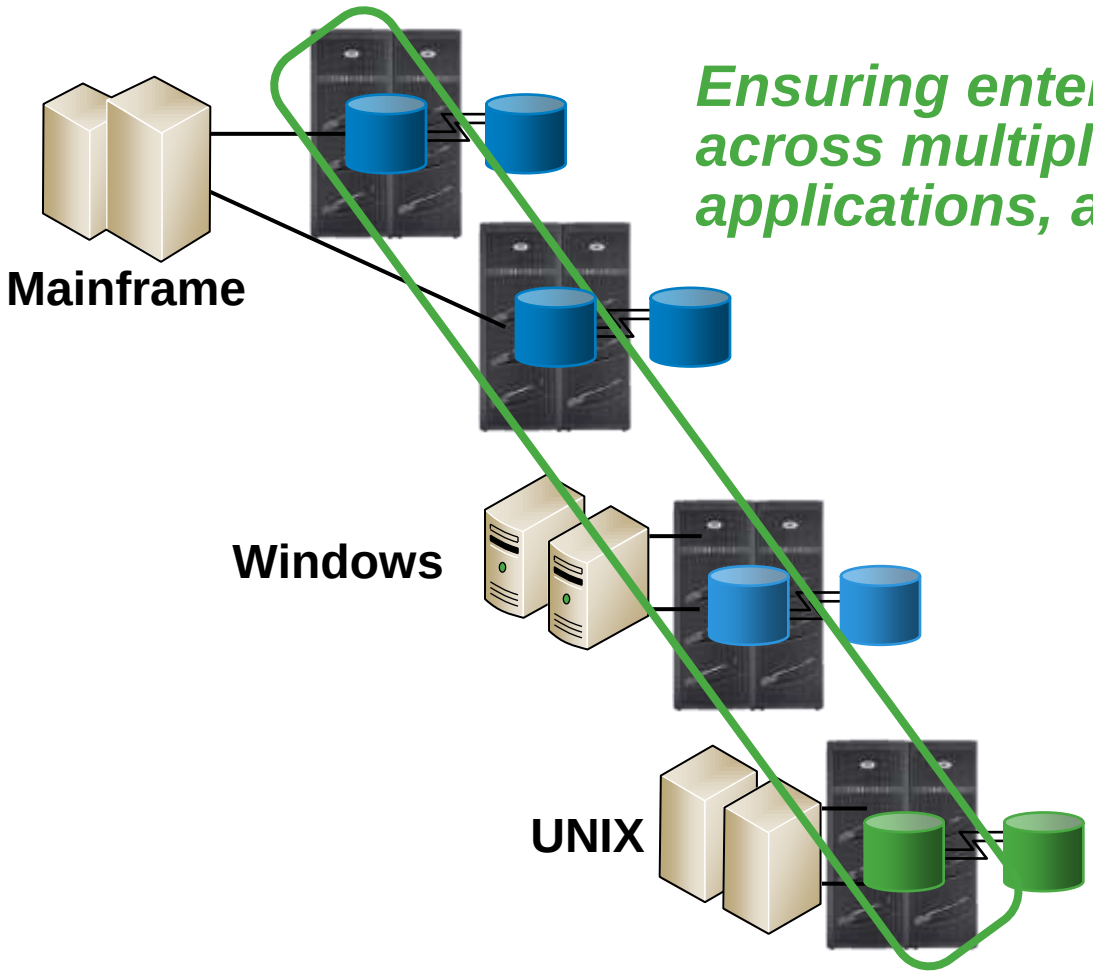

Ensuring enterprise-wide consistency across multiple databases, applications, and platforms

Data Consistency - BC/DR Last Hope Gold Copy

- Gold Copy
  - copy of secondary image to protect against failure during synchronization

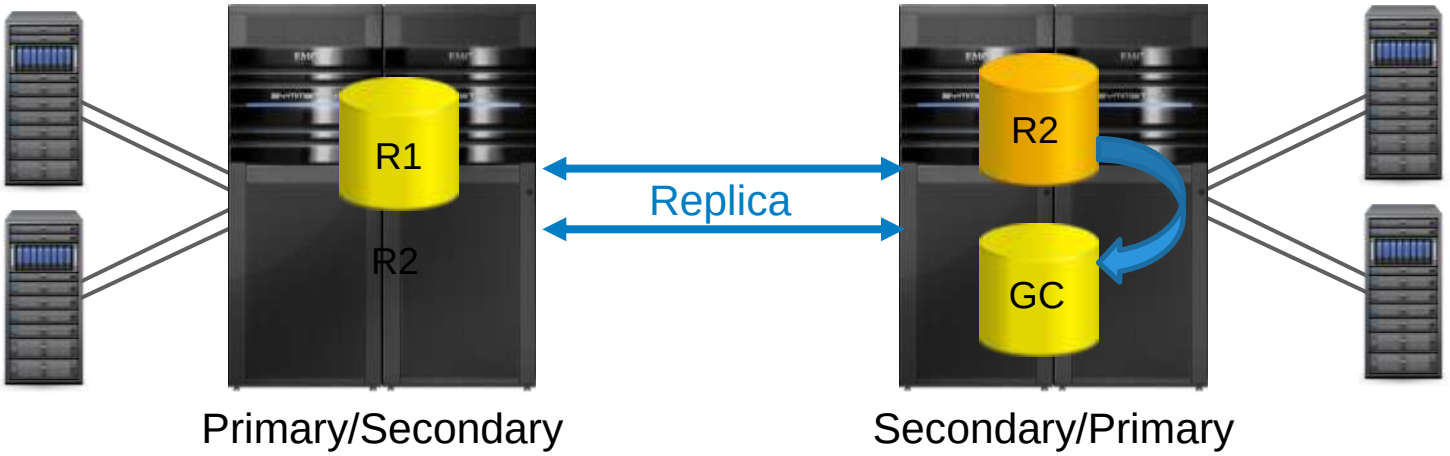

BC & DR Technologies

Portfolio: EMC BC/DR Technologies

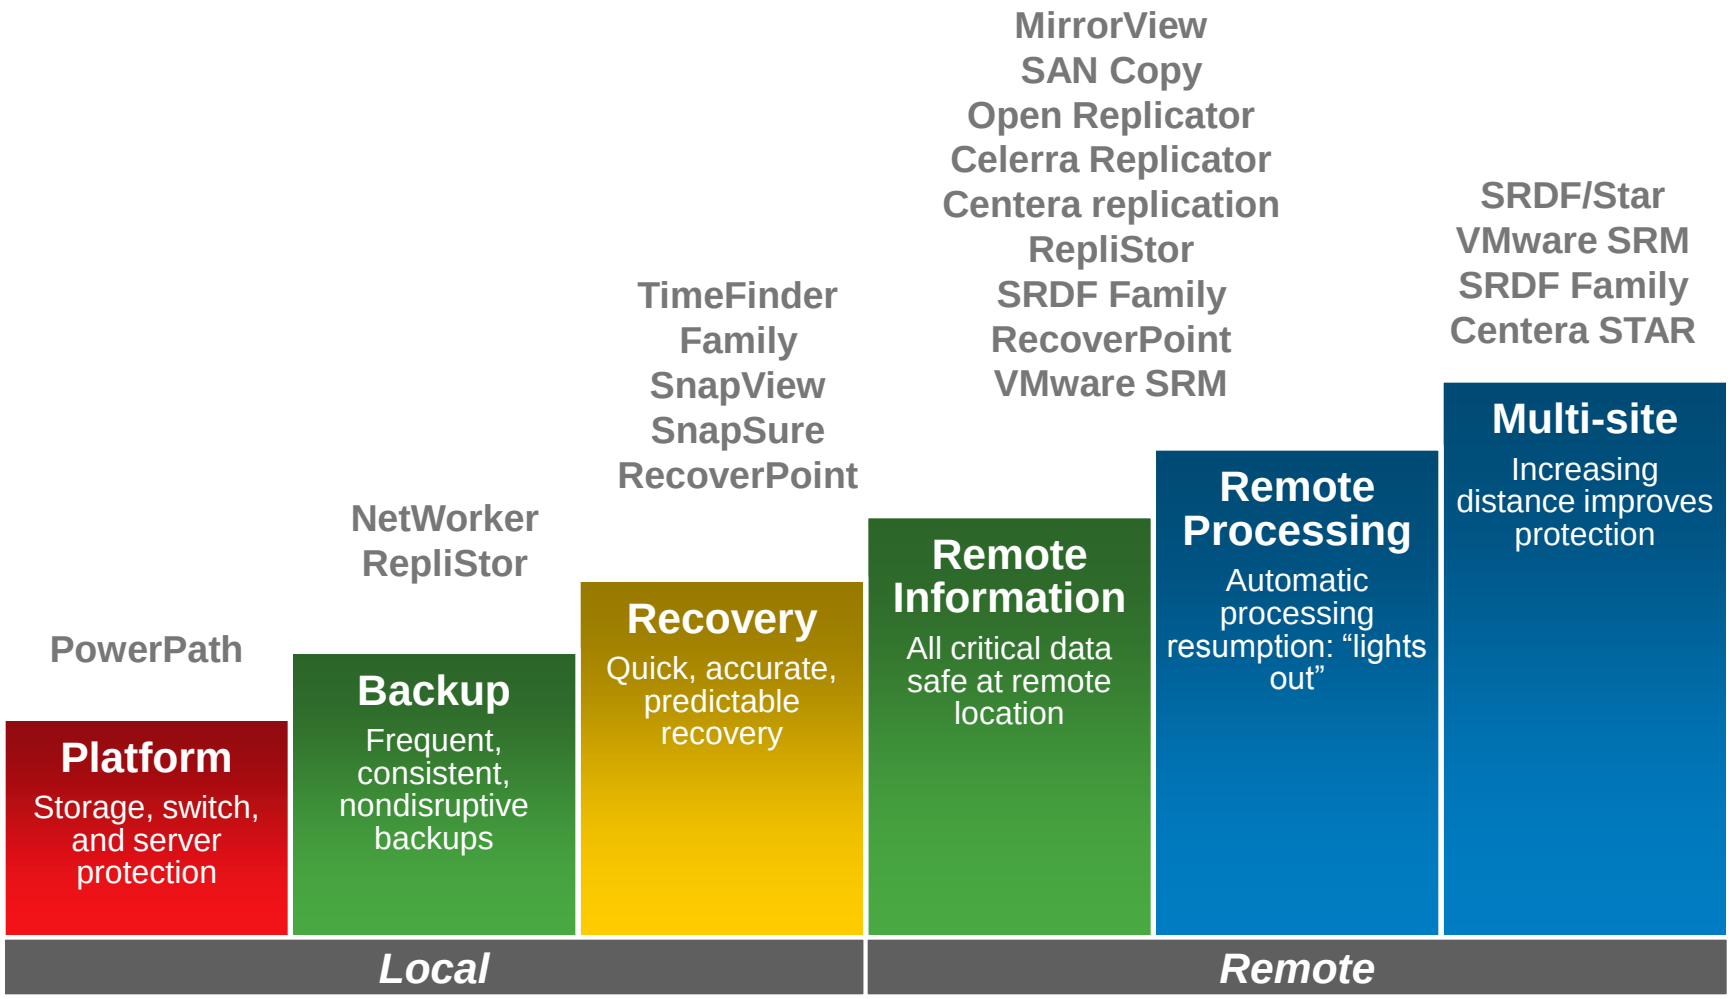

Enterprise Dedicated Replica

SW

Symmetrix Remote Data Facility (SRDF)

Time Finder (TF)

EMC SRDF

Industry-leading remote replication for Symmetrix

SRDF FAMILY

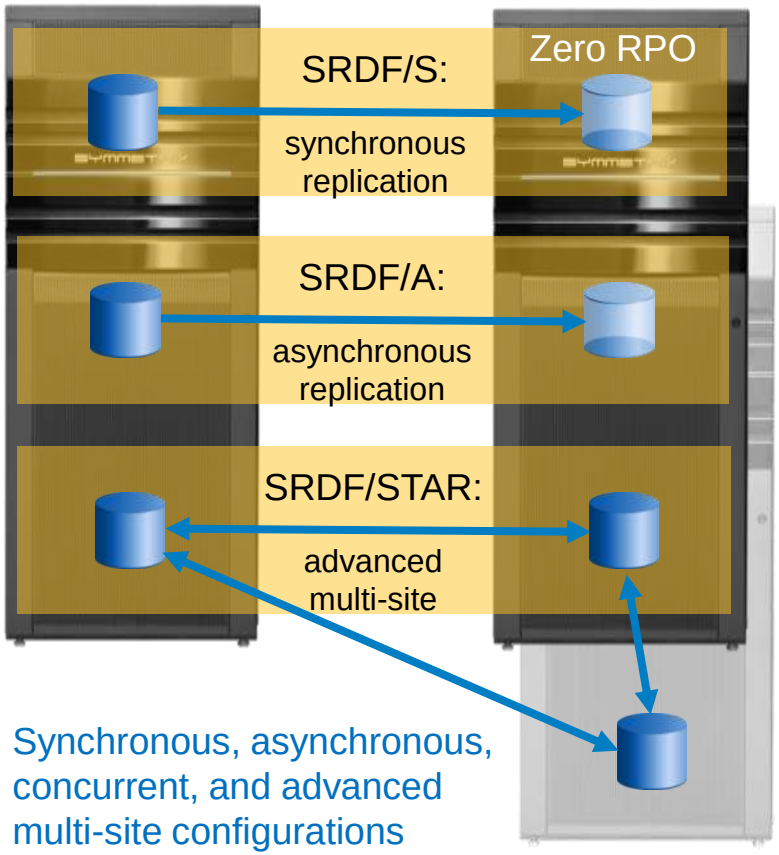

Synchronous, asynchronous, concurrent, and advanced multi-site configurations

- Provides advanced multi-site deployments
- Protects open systems and mainframes
- Provides coordinated and automated failover/restart of applications, servers, storage

SRDF/Synchronous Mode Operations

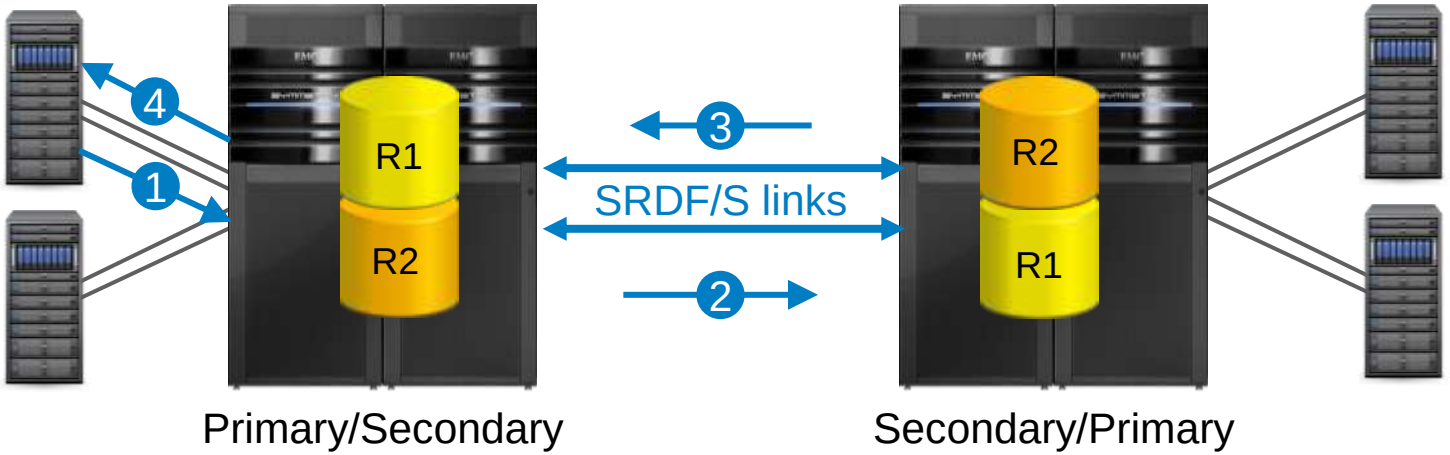

1. I/O write received from host/server into source cache
2. I/O is transmitted to target cache
3. Receipt acknowledgment is provided by target back to cache of source
4. Ending status is presented to host/server

SRDF/Asynchronous Remote Replication and Local or Regional Disasters

- Primary requirement: Provide for minimal data loss (RPO = less than one minute) at an out-of-region site in the event of a local or regional disaster
- Objective achieved: Yes, with Site B available for disaster restart

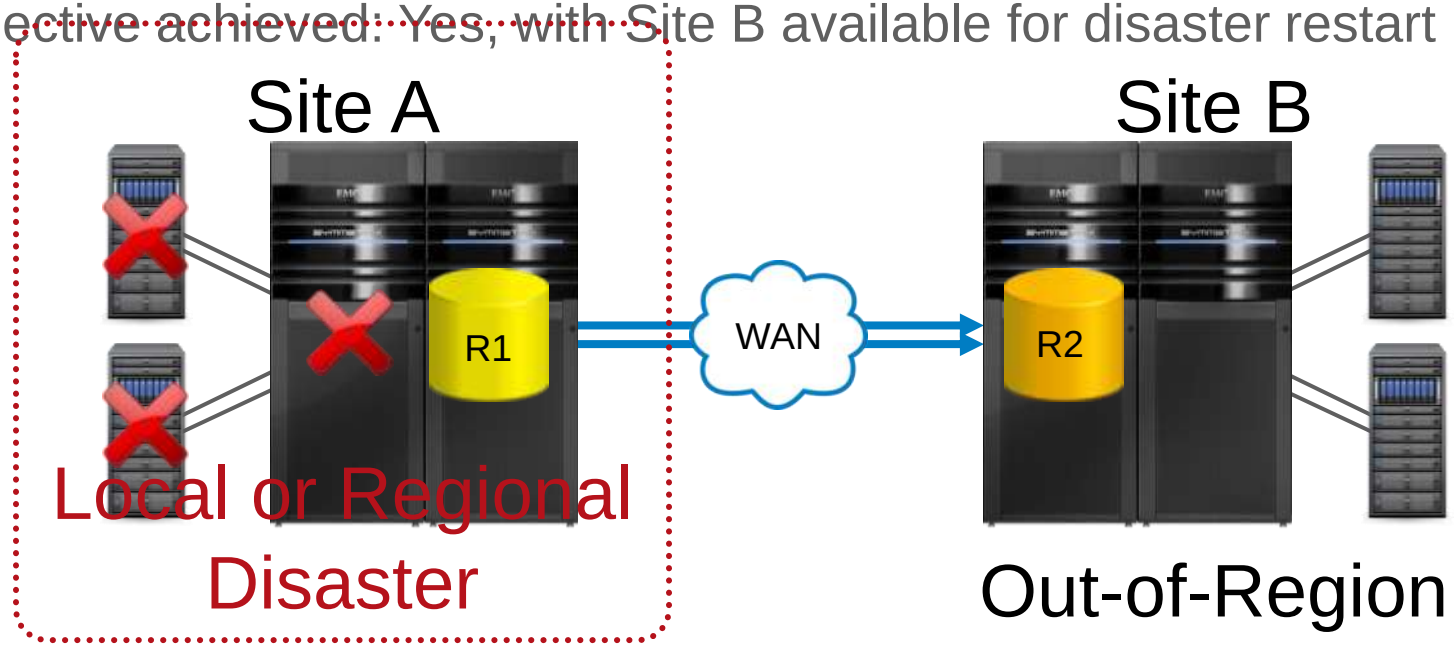

Provides disaster restart in the event of a regional disaster

SRDF/AR Multi-Hop

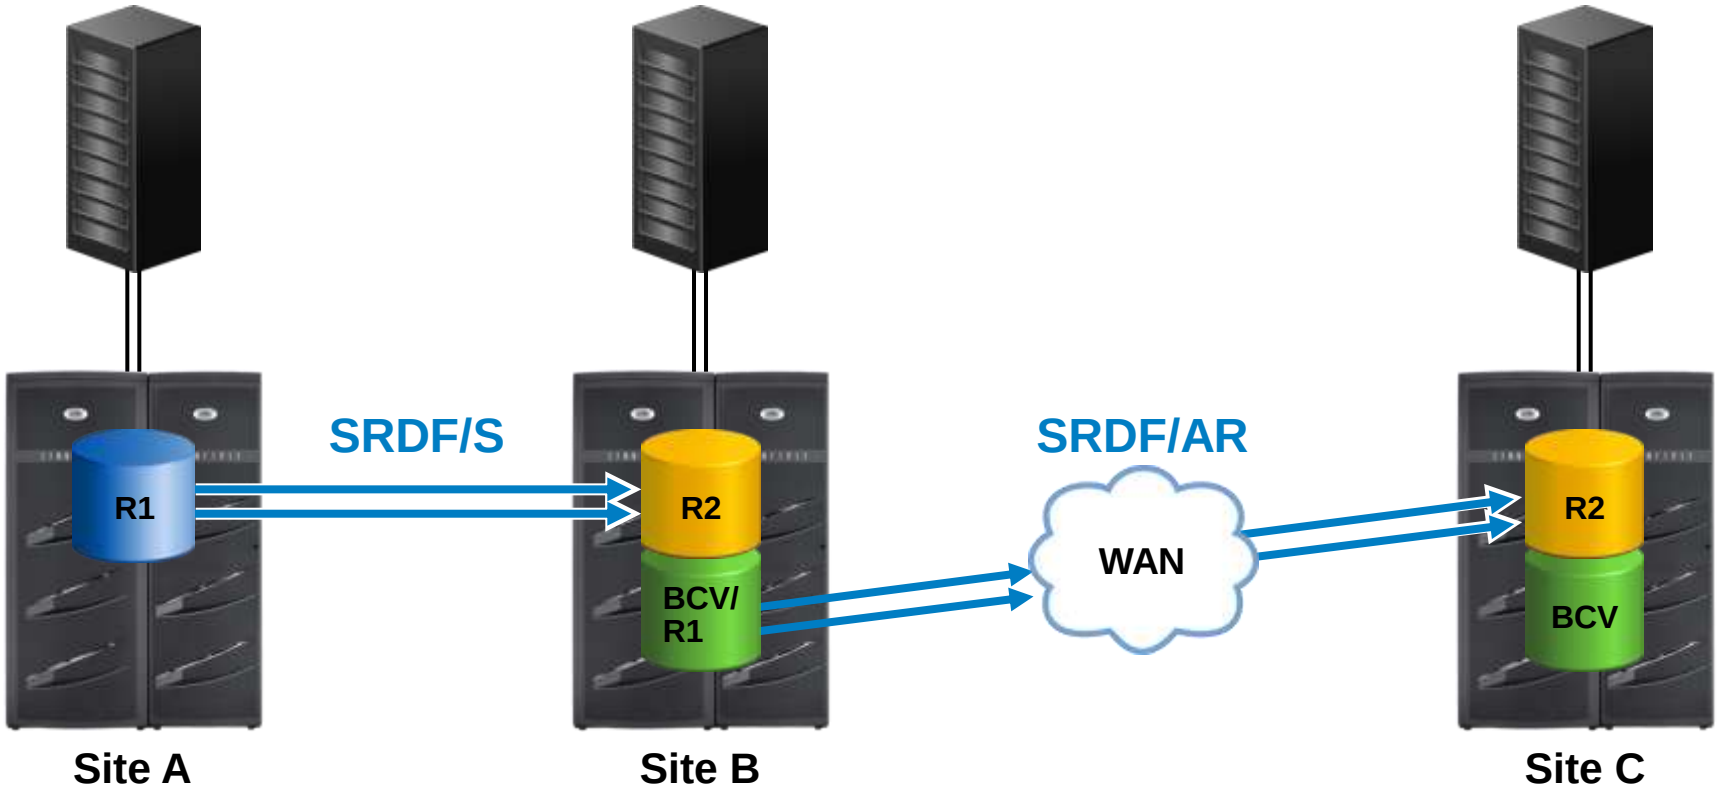

Possible no data loss solution while reducing extended distance intersite network bandwidth requirements resulting in lower costs

Concurrent Remote Replication (1 of 3)

Primary requirement with this example: Provide for a no-data-loss solution (RPO = 0) in event of local disaster

Provide for minimal data loss (RPO = less than two minutes) in event of regional disaster with a three-site solution

Objective achieved: Yes - Site B remains available

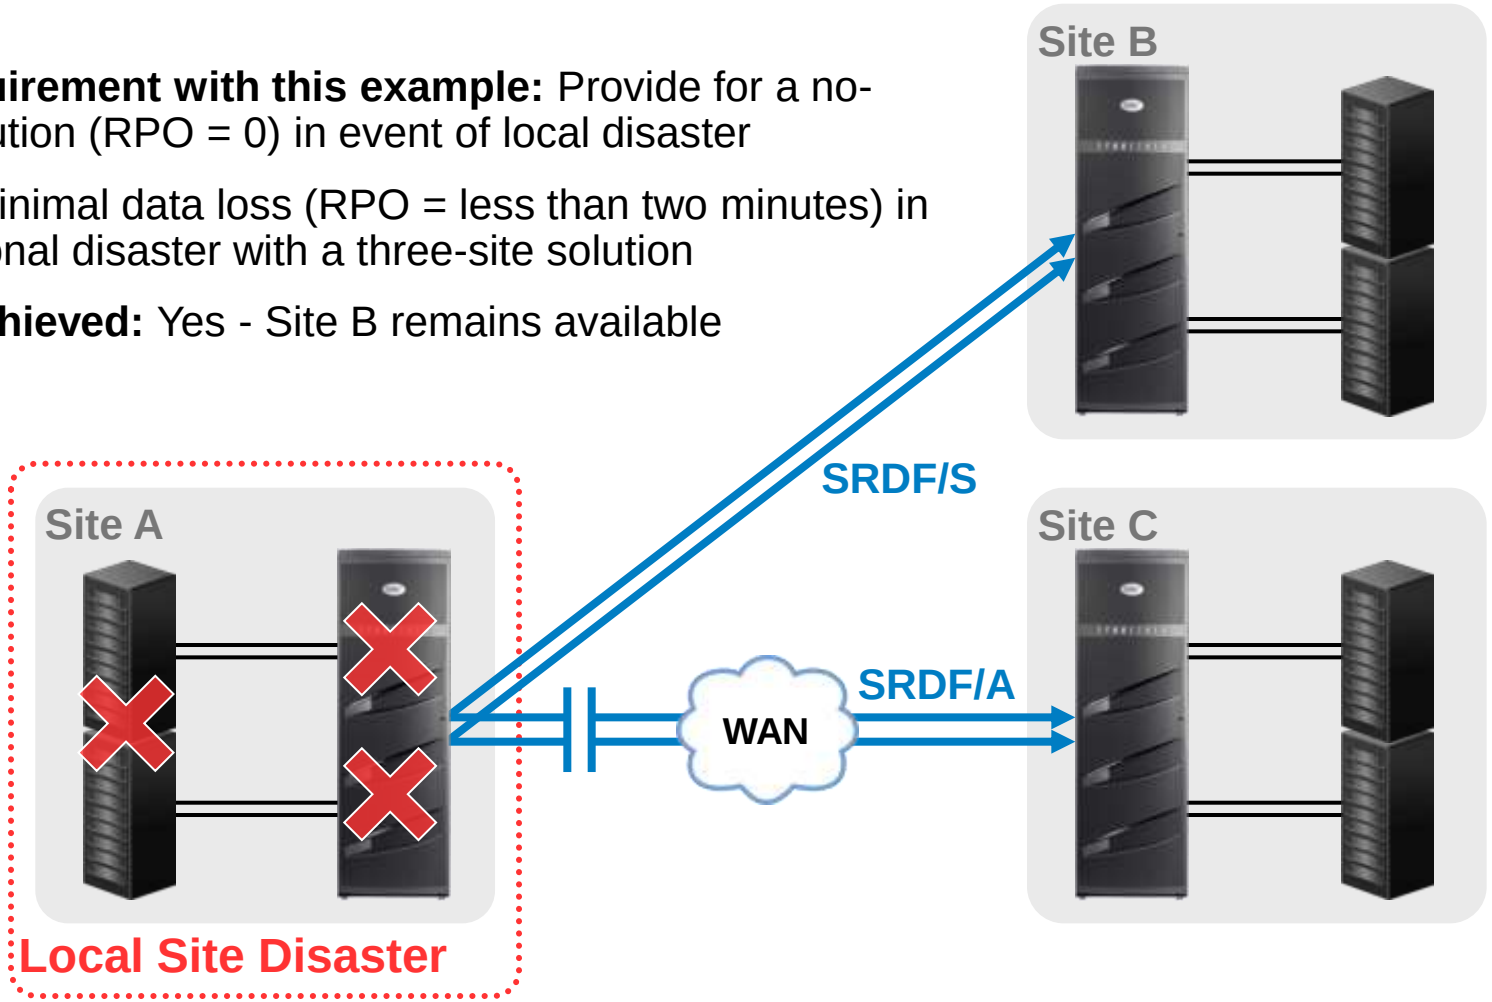

Provides continuous data protection between Sites A and B with no data-loss exposure

Concurrent Remote Replication (2 of 3)

Primary requirement with this example: Provide for a no-data-loss solution (RPO = 0) in event of local disaster

Provide for minimal data loss (RPO = less than two minutes) in event of regional disaster with a three-site solution

Objective achieved: Yes, since Site A and C remain available

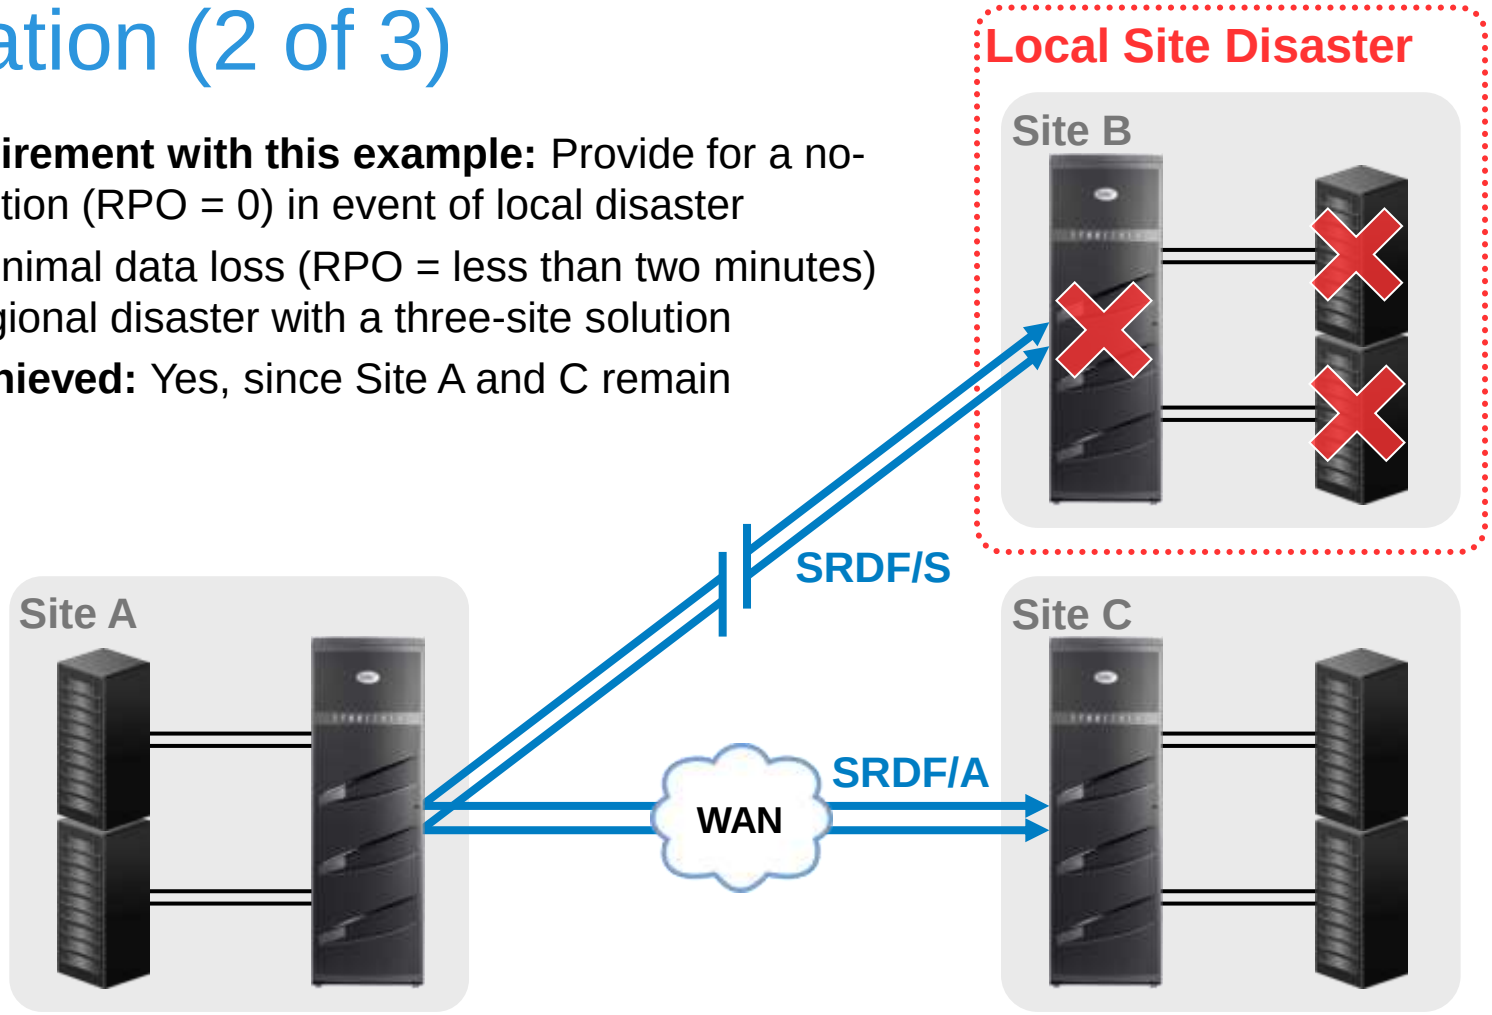

Provides continuous data protection between Sites A and C with no data loss—data-loss exposure will now depend on data lag between Sites A and C

Concurrent Remote Replication (3 of 3)

Primary requirement with this example: Provide for a no-data-loss solution (RPO = 0) in event of local disaster

Provide for minimal data loss (RPO = less than two minutes) in event of regional disaster with a three-site solution

Objective achieved: Yes - Site C remains available

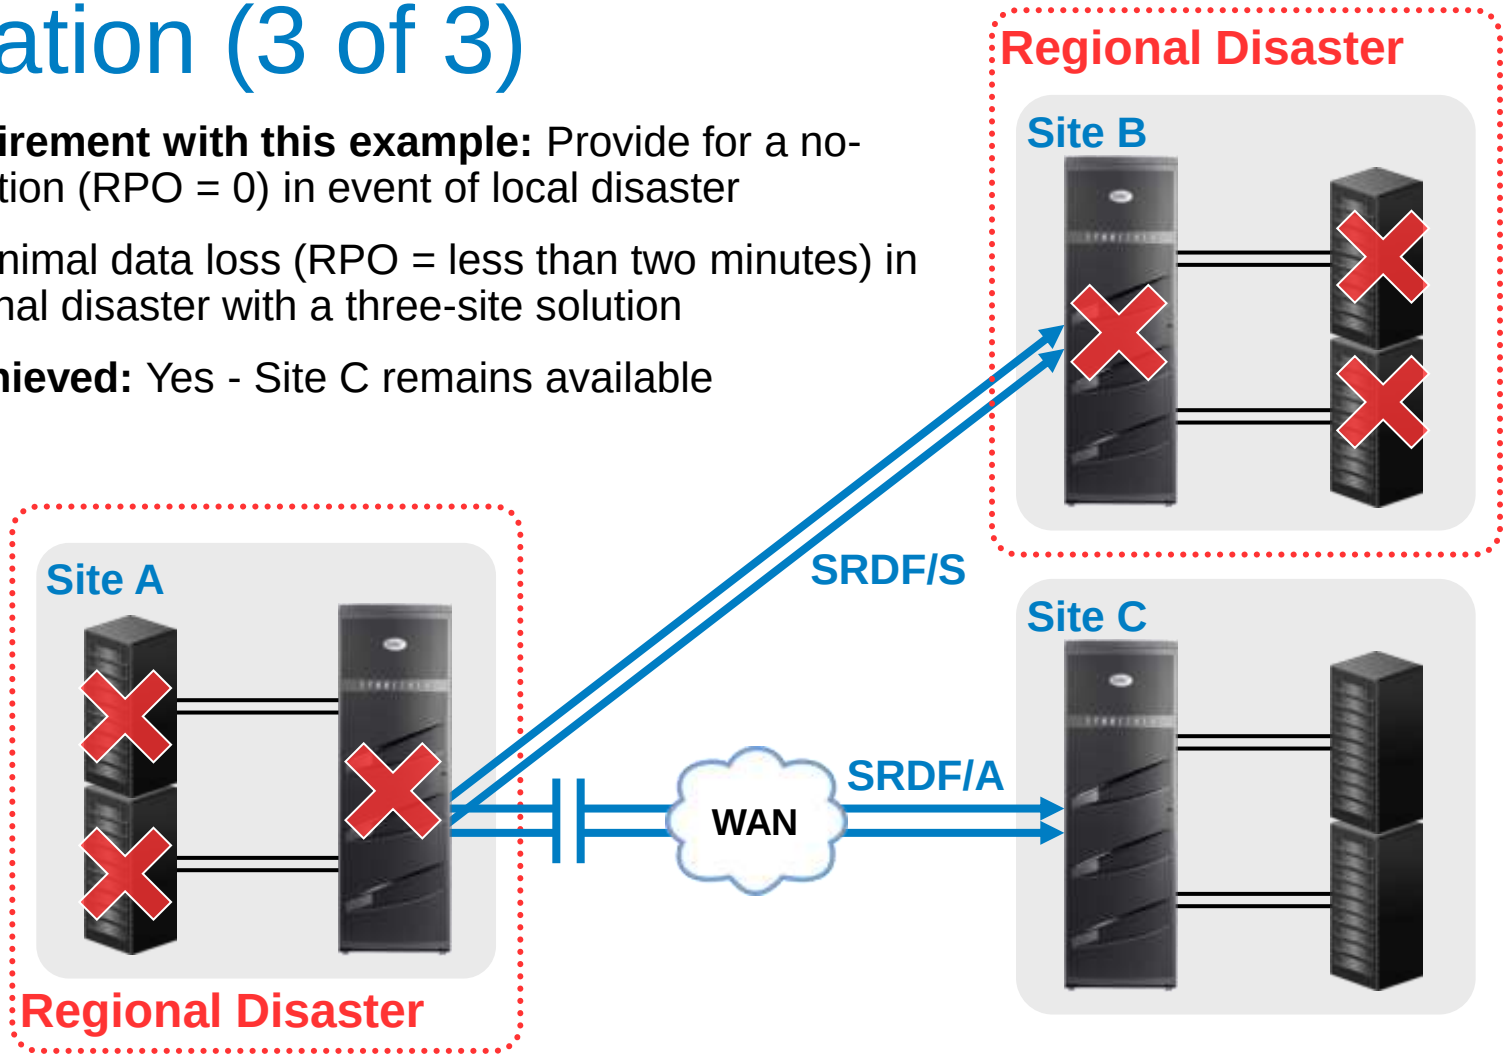

Provides continuous data protection between Sites A and B with no data-loss exposure

Cascaded SRDF

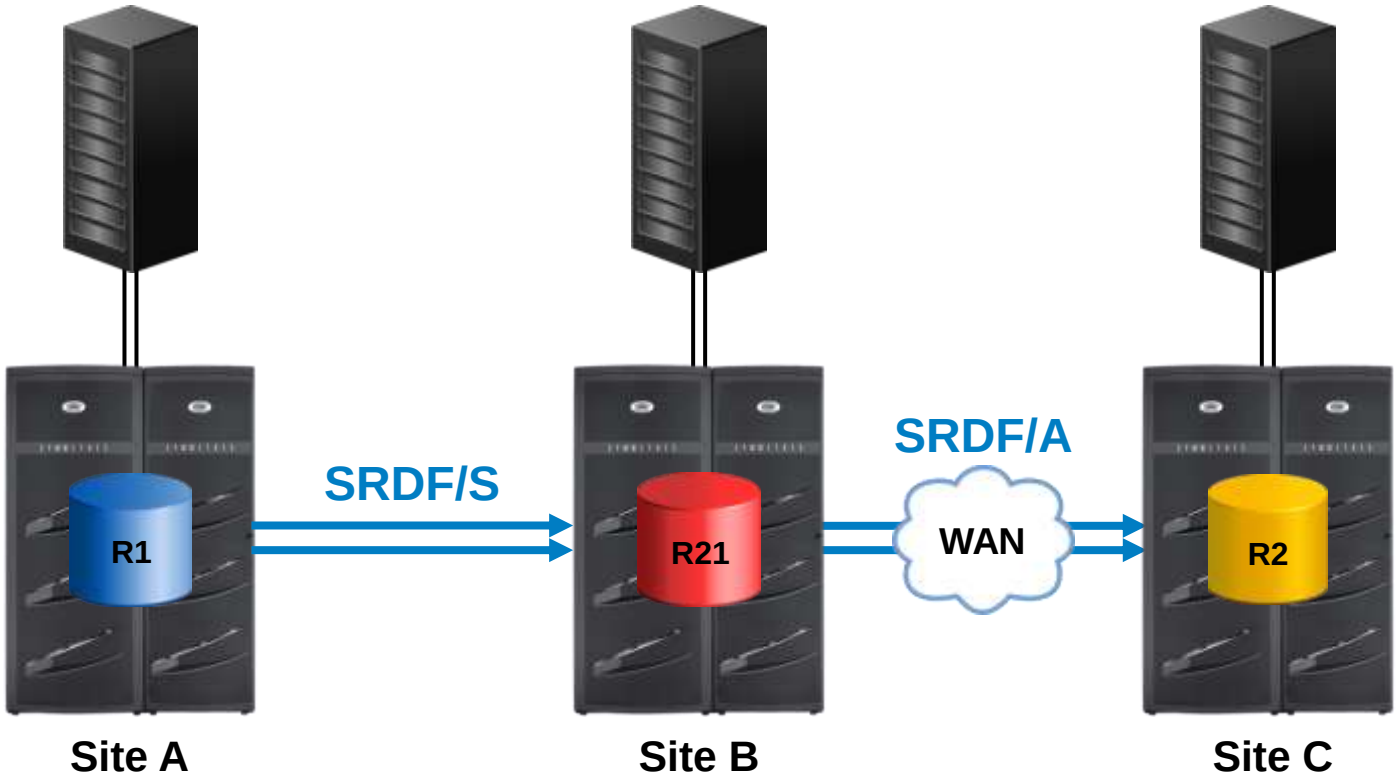

Achievable no data loss solution with improved R

SRDF/Extended Distance Protection—Zero Data Loss Asynchronous Replication

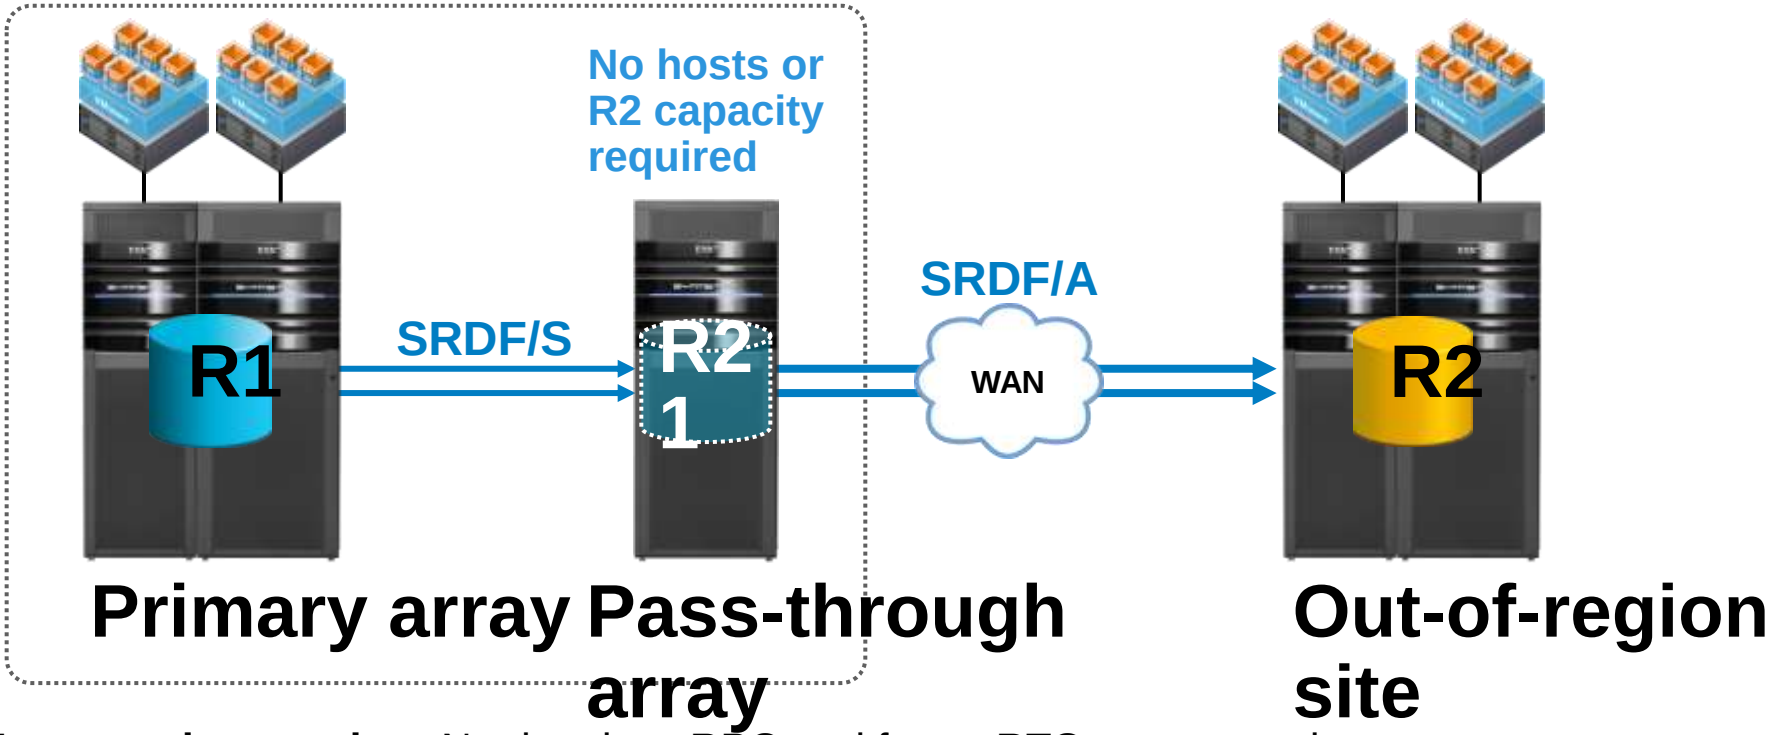

- **Increased protection:** No data loss RPO and faster RTO at remote site
- **More affordable:** Pass-through system reduces server and capacity costs
- **Investment protection:** Symmetrix DMX-3 and DMX-4 supported as source and target
- **Competitive differentiation:** SRDF-exclusive capability

SRDF/Star with Concurrent

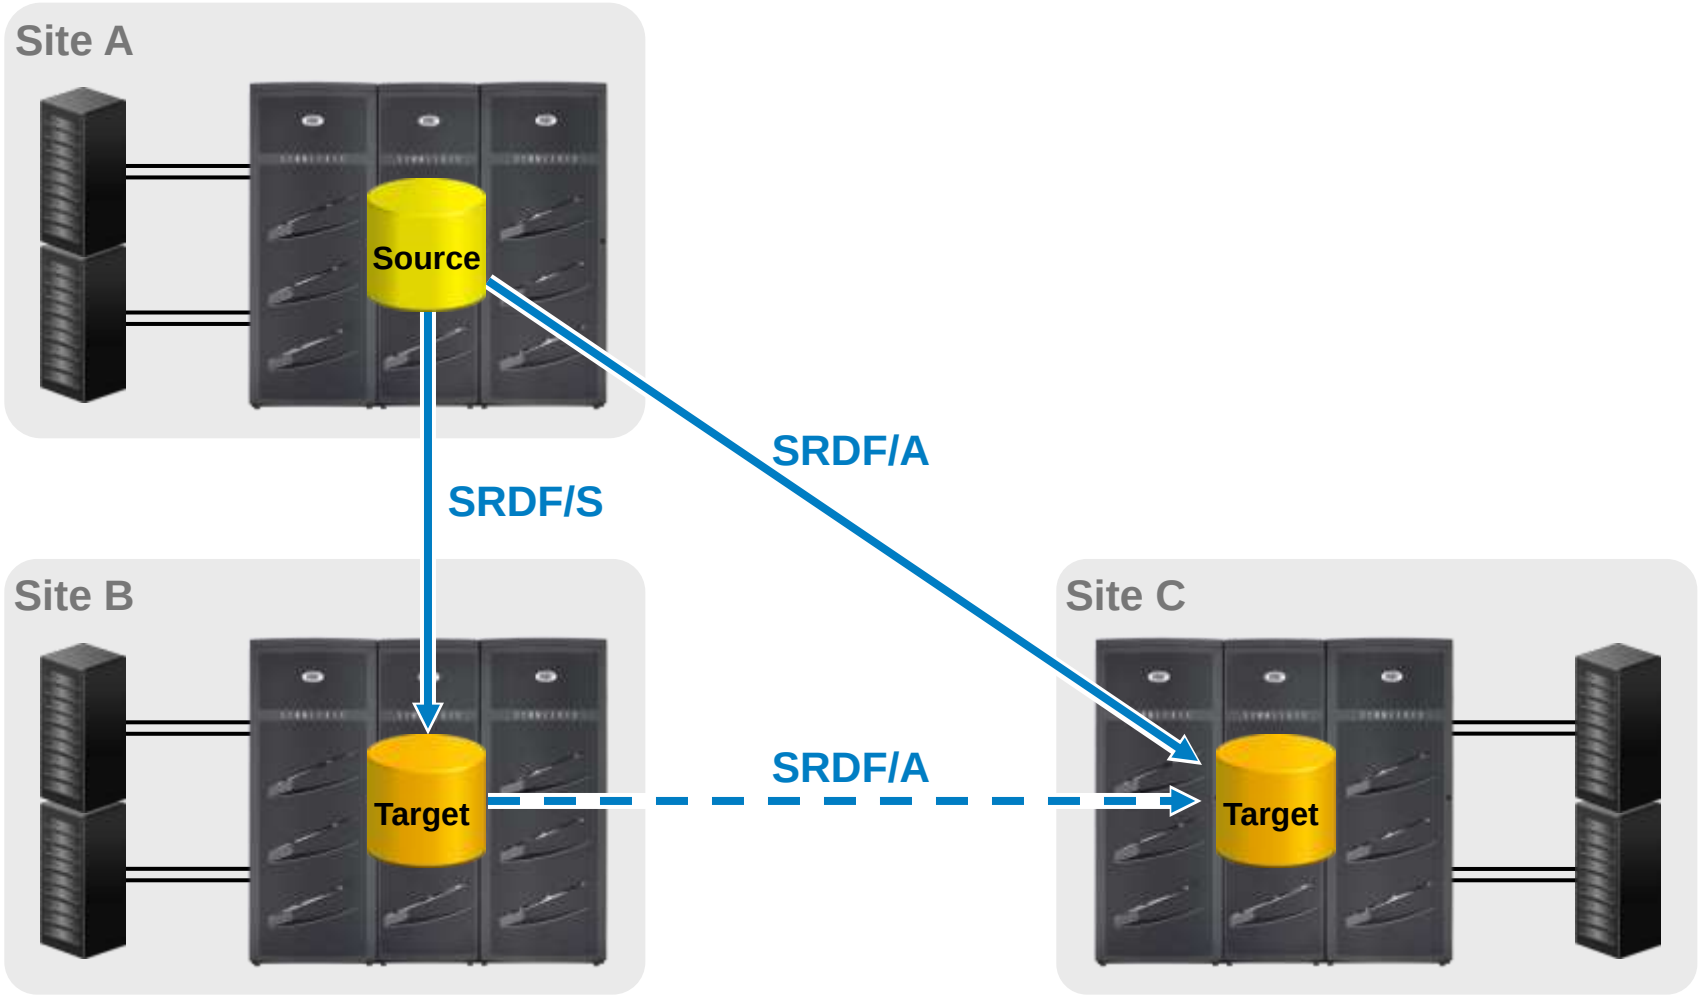

SRDF/Star—Continuous Protection Due to Site-Failure Event and User Failover to Site B (1 of 2)

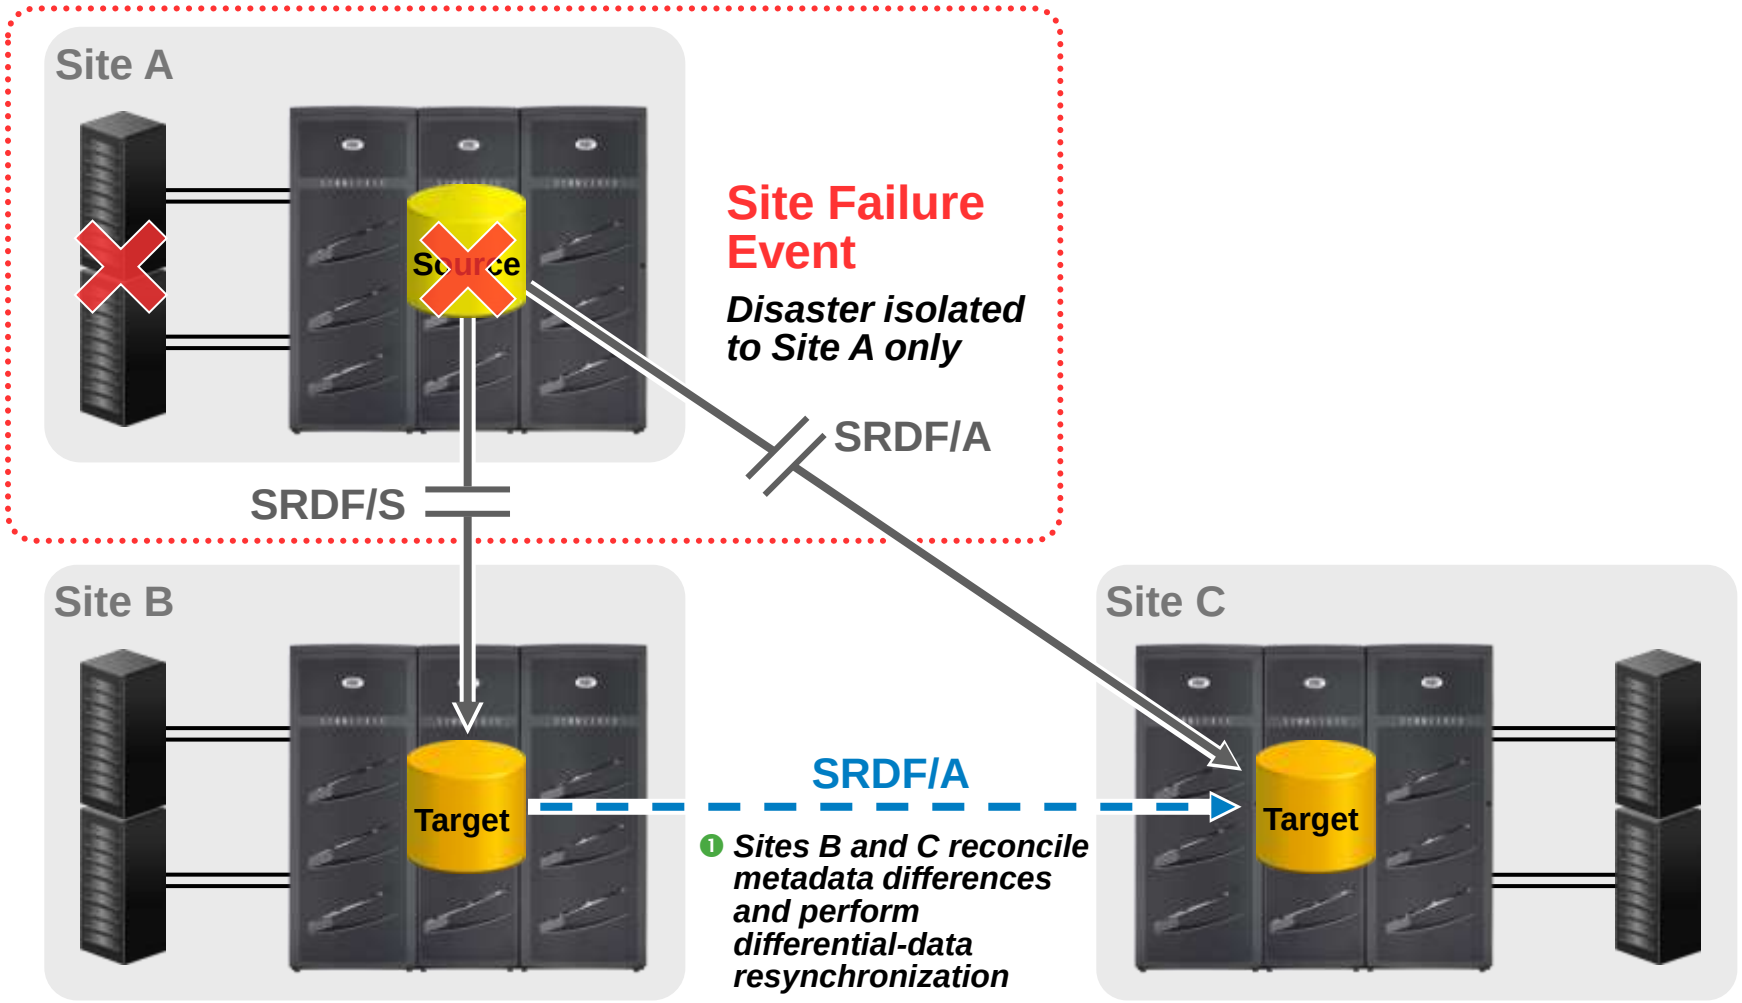

SRDF/Star—Continuous Protection Due to Site-Failure Event and User Failover to Site B (2 of 2)

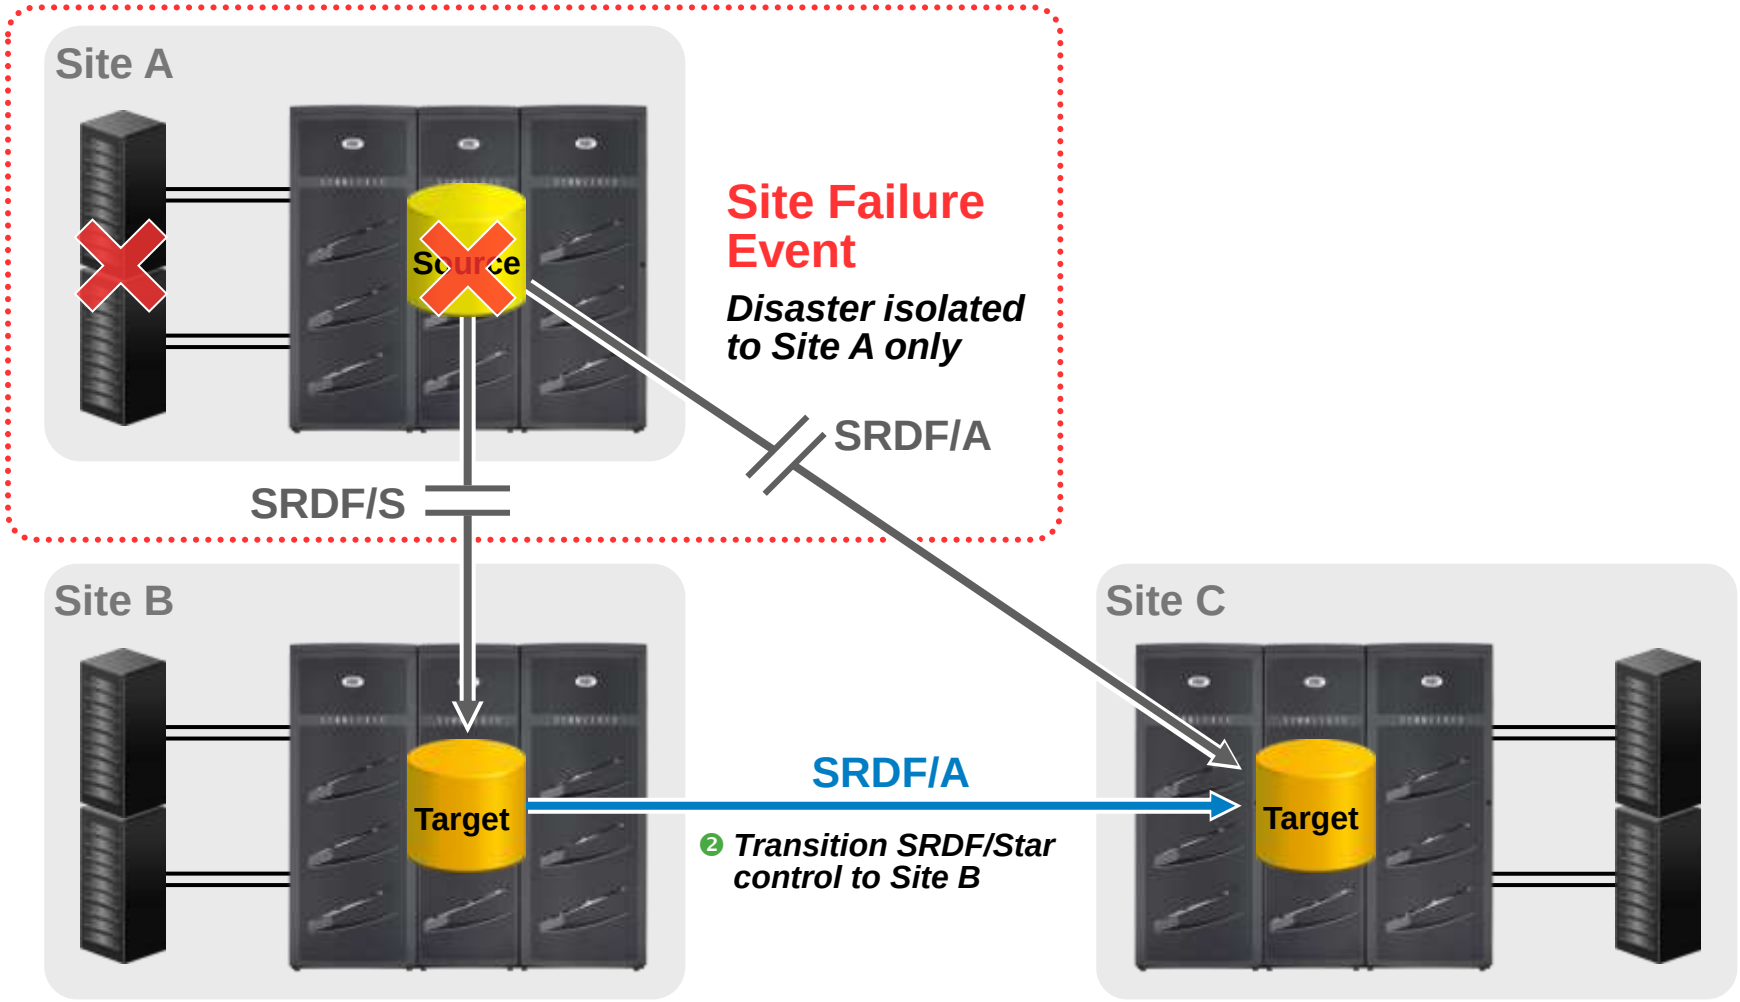

SRDF/Star—Continuous Protection Due to Site-Failure Event and User Failover to Site C (1 of 2)

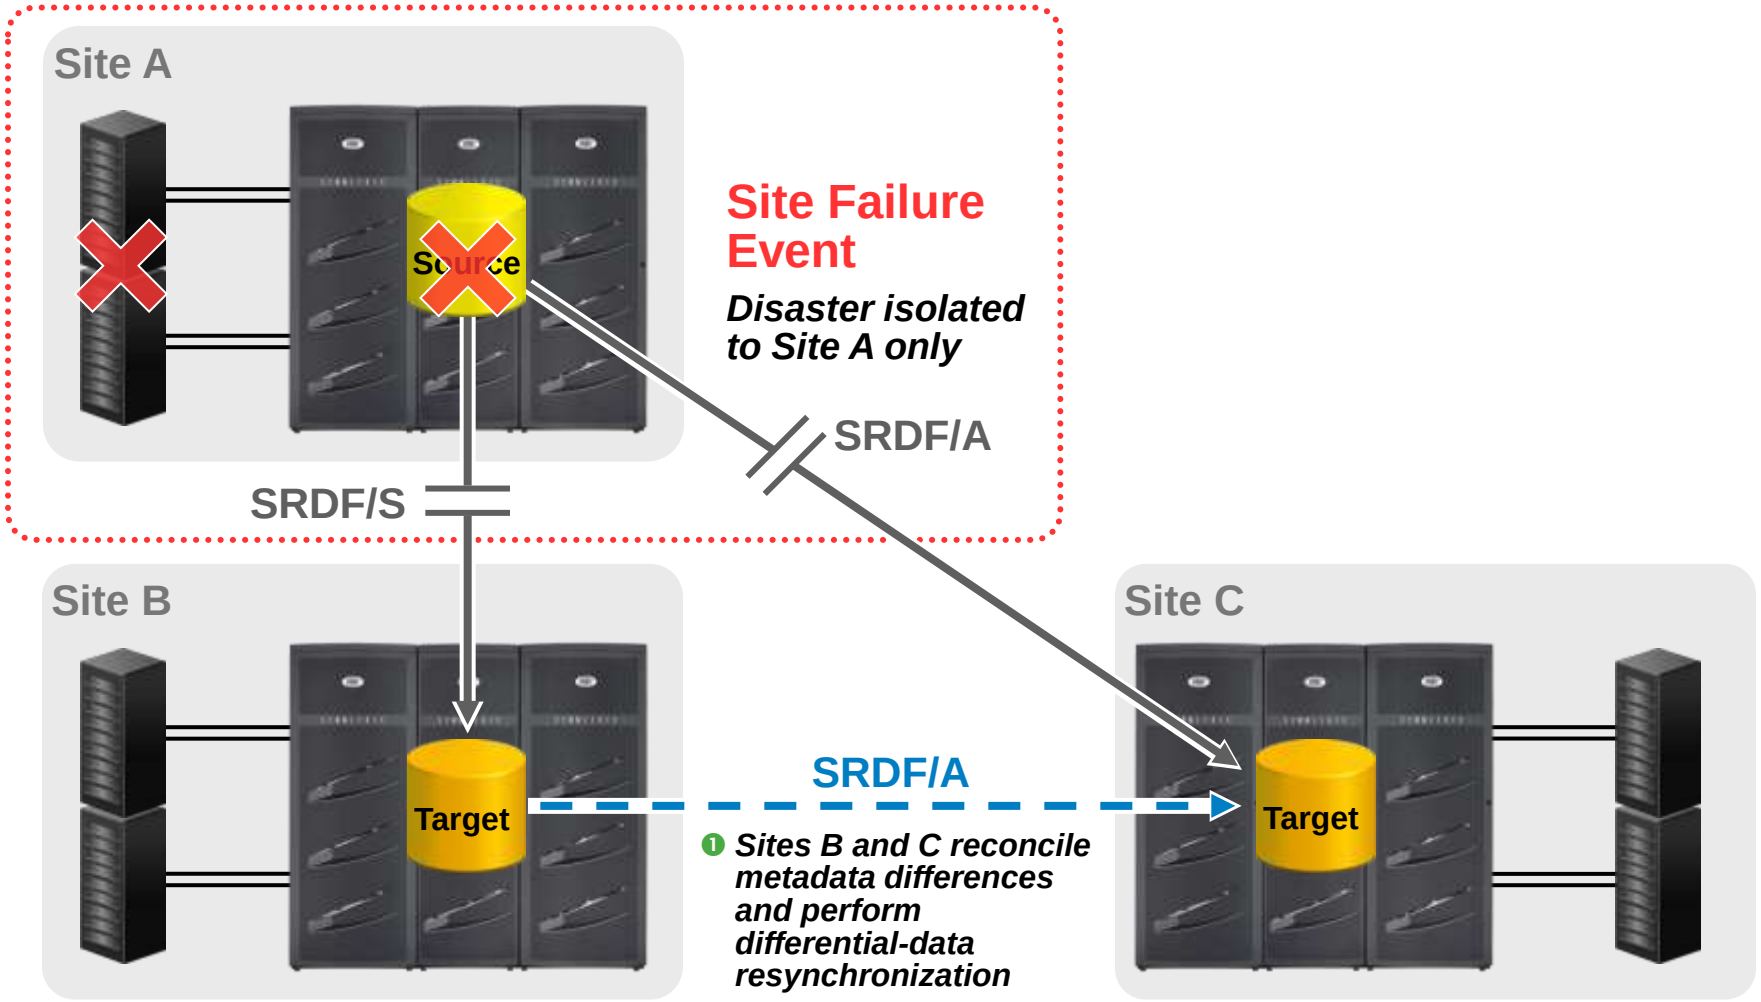

SRDF/Star—Continuous Protection Due to Site-Failure Event and User Failover to Site C (2 of 2)

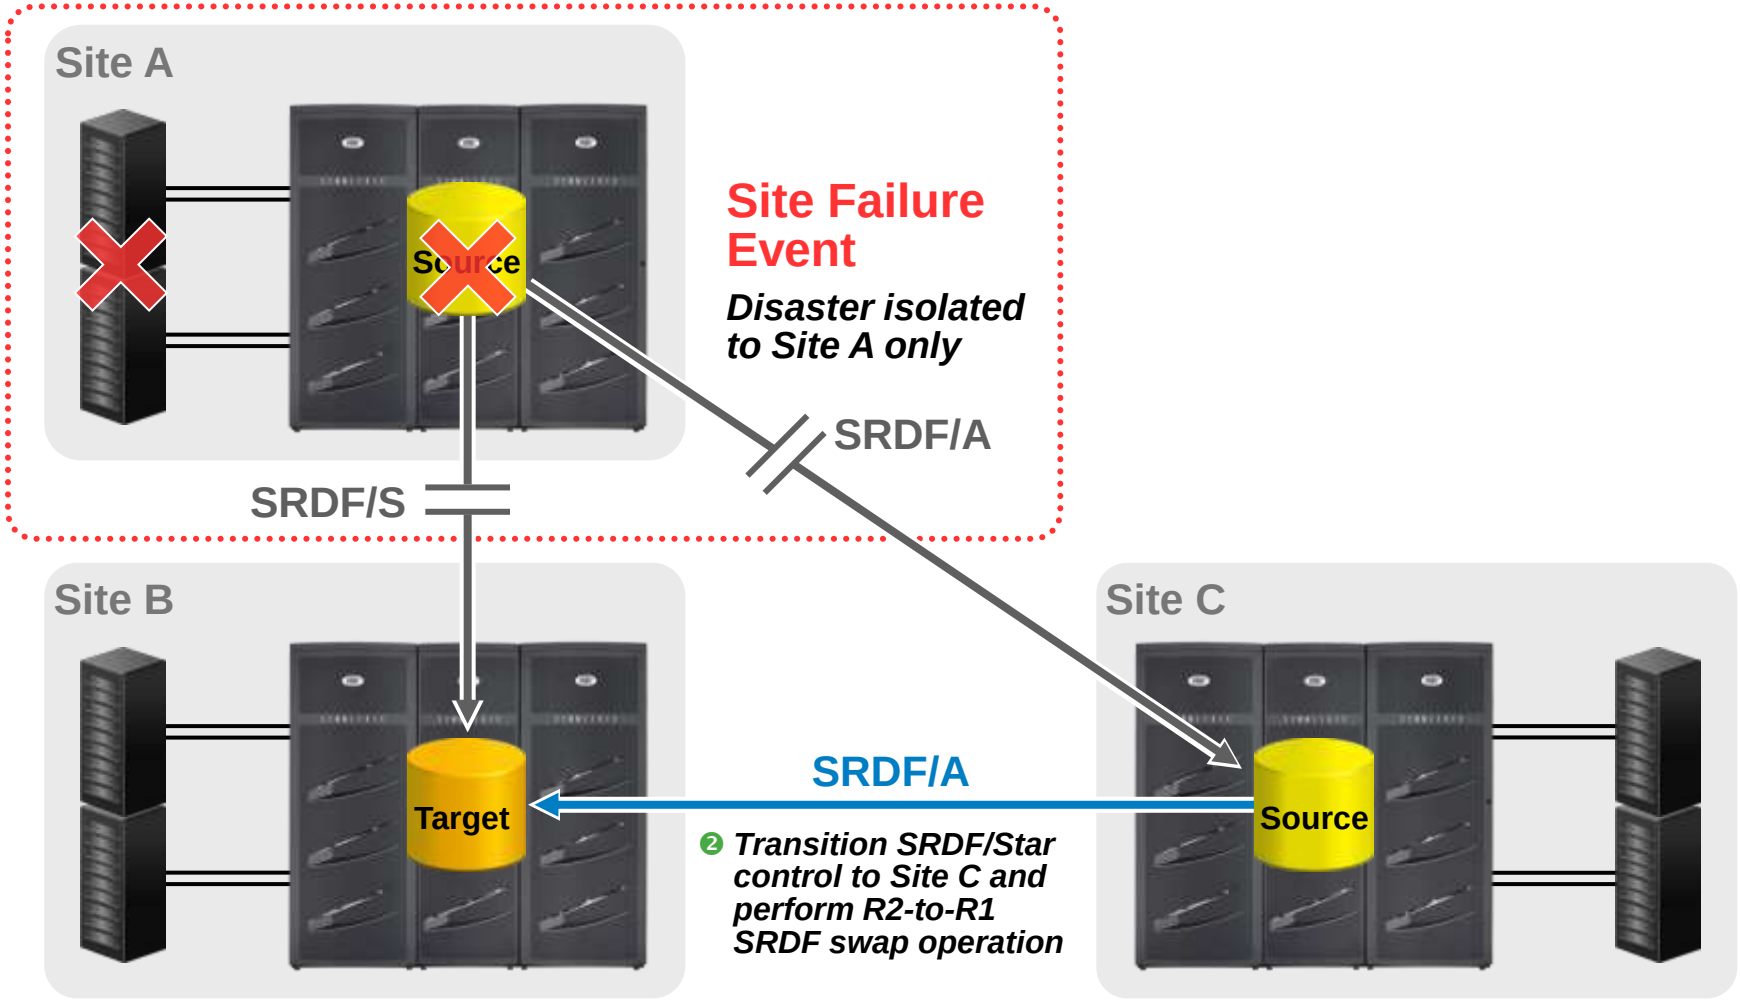

EMC TimeFinder

Industry-leading local replication for Symmetrix

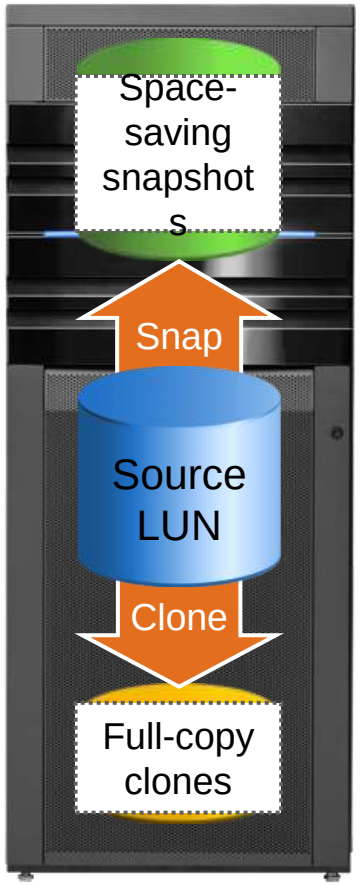

- Array-based software
  - Snapshots: Space-saving views of production information
  - Clones: Full copies of data
- Enhances productivity
  - Tens of thousands of replicas for greater parallel processing
- Delivers co-developed application-integration
  - For fast deployment, faster restore/restart

TimeFinder/CG (Consistency Groups)

- Used to create a consistent point-in-time copy of related data
- Multiple volumes activated or “snapped” as one logical unit
  - Ensures consistency across multiple volumes
- Consistency Groups can span:
  - An associated set of volumes
  - An entire database
  - An entire system
  - Across systems

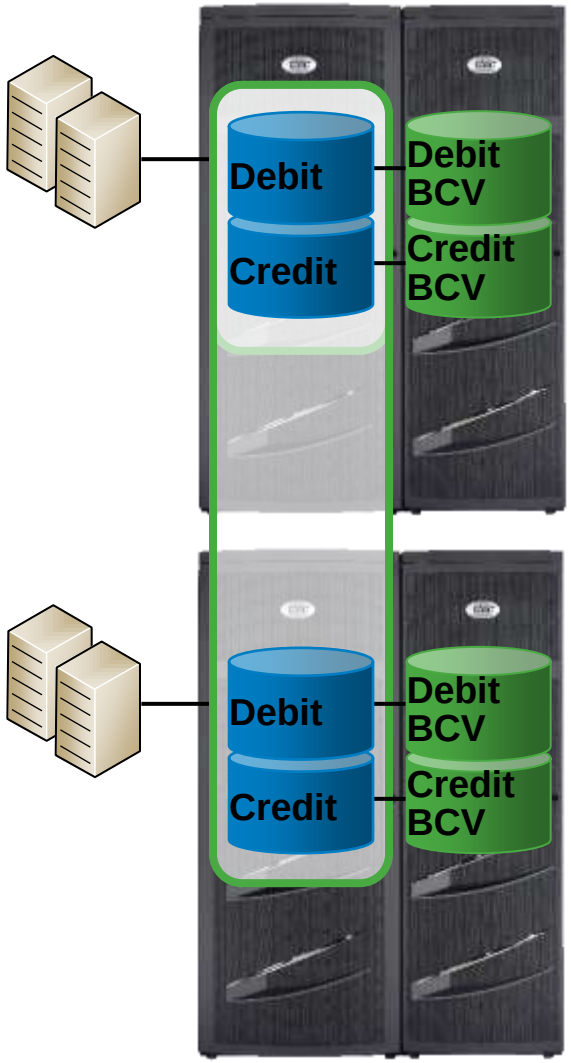

Modular Dedicated Replica sw

MirrorView(MV) – IP Replicator (IPR)

SnapView(TF) – SnapSure (SS)

EMC MirrorView

Cost-effective multi-site remote replication for CLARiiON

MIRRORVIEW

Synchronous or Asynchronous

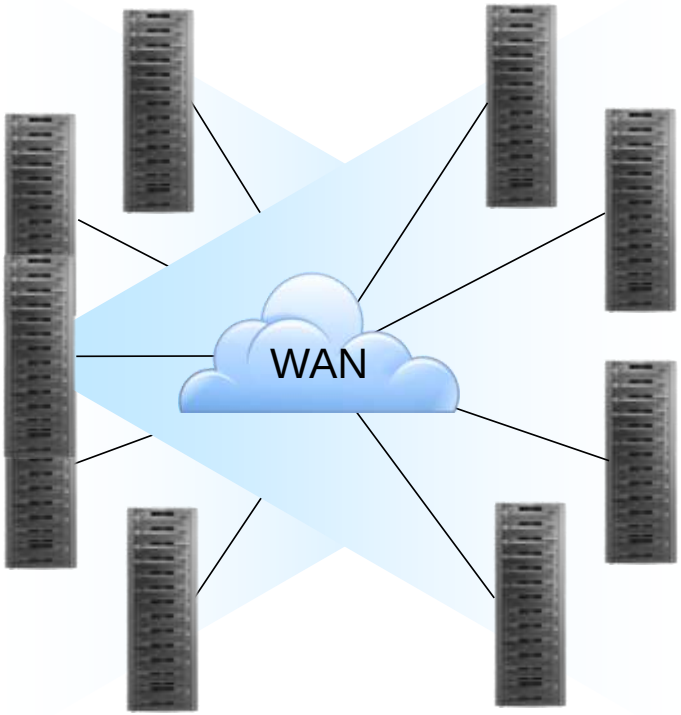

Fan-in (n:1) and Fan-out (1:n)
Multi-site Configurations

- Integrates with Unisphere for setup and management
- Integrates with SnapView for local replication

MirrorView

CLARiiON's Remote Replication Software

MirrorView/Synchronous

RPO: Zero seconds

- Both images identical
- Limited distance
- High network bandwidth
- One primary to one or two secondaries

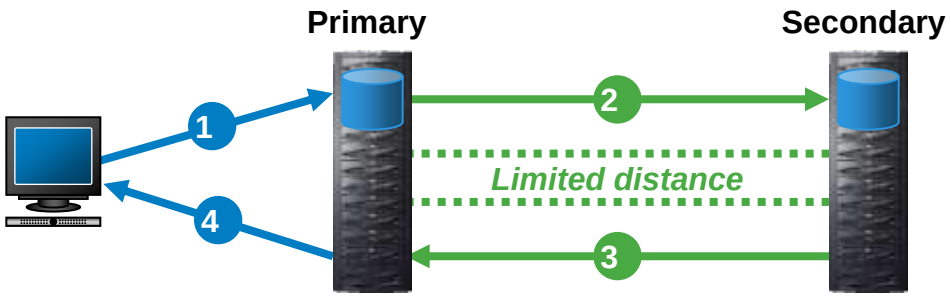

MirrorView/Asynchronous

RPO: 30 minutes to hours

- Target updated periodically
- Unlimited distance
- Restartable copy on secondary if session fails
- Optimized for low network bandwidth (consumes 100 Mb/s maximum)
- One primary to one secondary

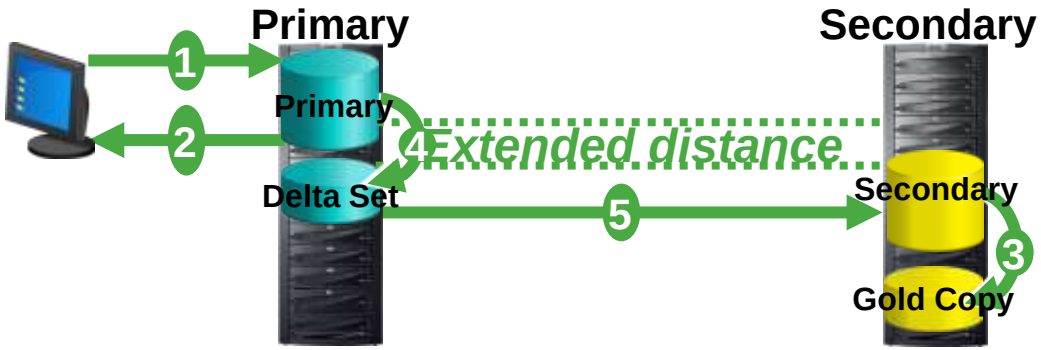

52

MirrorView

P = Primary; S = Secondary

Disaster-Recovery Consolidation with 4:1 Fan-in

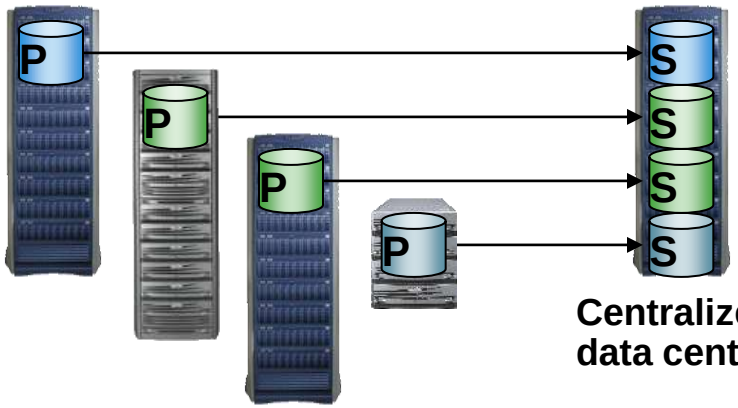

Concurrent Mirroring for Parallel Processing (*MirrorView/Synchronous only*)

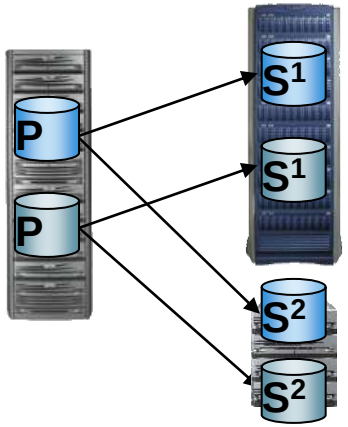

Bi-Directional Support for Multiple Sites

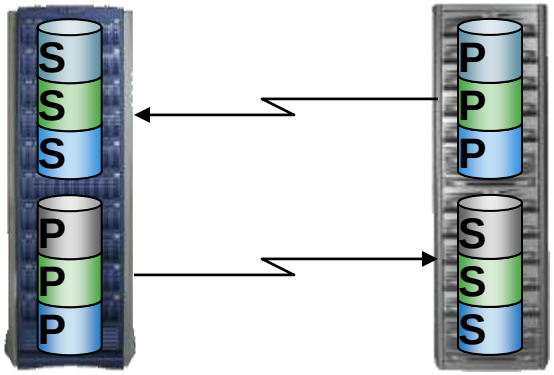

MirrorView / SnapView Integration to Eliminate Planned and Unplanned Outages

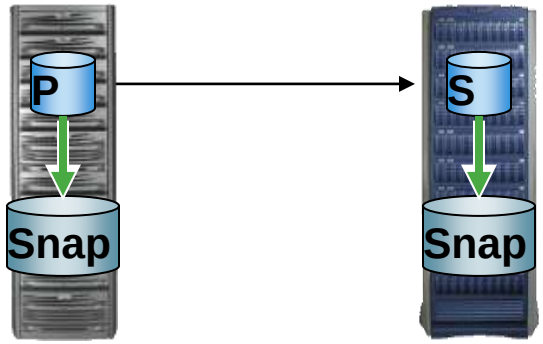

EMC SnapView

Point-in-time snapshots and full-image clones for CLARiiON

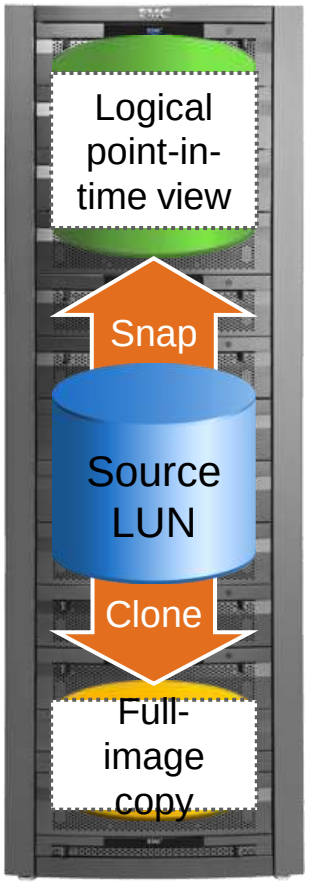

- Array-based software consumes zero host cycles
  - Snapshots: Logical views of production information
  - Clones: Copies of data
- Readable, writeable replicas
  - Provide concurrent data access
  - Enables increased operational efficiencies

EMC Celerra Replicator

Point-in-time file system and iSCSI LUN replication for Celerra

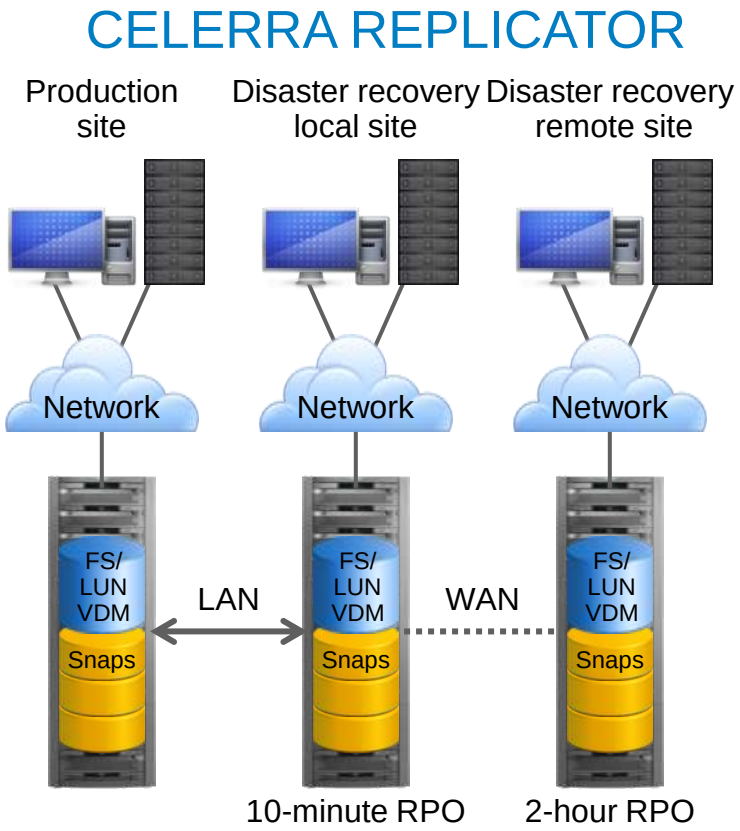

- Includes ‘set and forget’ policy definitions
- Ensures quality of service with scheduled bandwidth throttling
- Enables 1-to-N data replication for multi-site disaster recovery

Heterogeneous Replica sw

RecoverPoint - Continous Remote Replication (CRR)

RecoverPoint - Continous data processing (CDP)

RecoverPoint – Continous local and remote protection (CLR)

EMC RecoverPoint

One way to protect everything you have better

RECOVERPOINT FAMILY

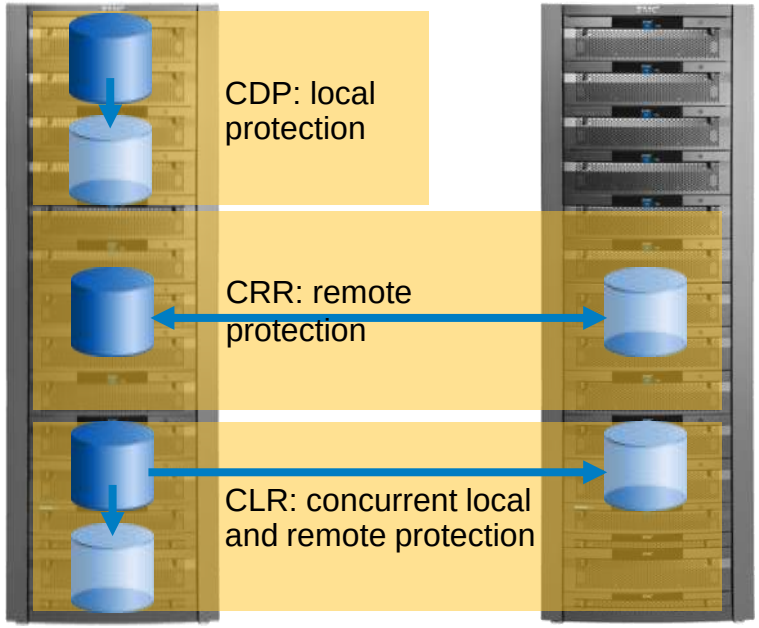

- Federated, clustered, cloud applications
- Dynamic synchronous and asynchronous replication
- Concurrent local and remote protection
- Flexible deployment options
- Integrated with VMware

- Protects any host, any application, any array
- Provides affordable data protection
  - Use existing network bandwidth
  - Save on storage costs
- Enables any point-in-time recovery
  - Continuous local and remote protection
  - Application-consistent bookmarks

Local Protection Process: Continuous Data Protection (CDP)

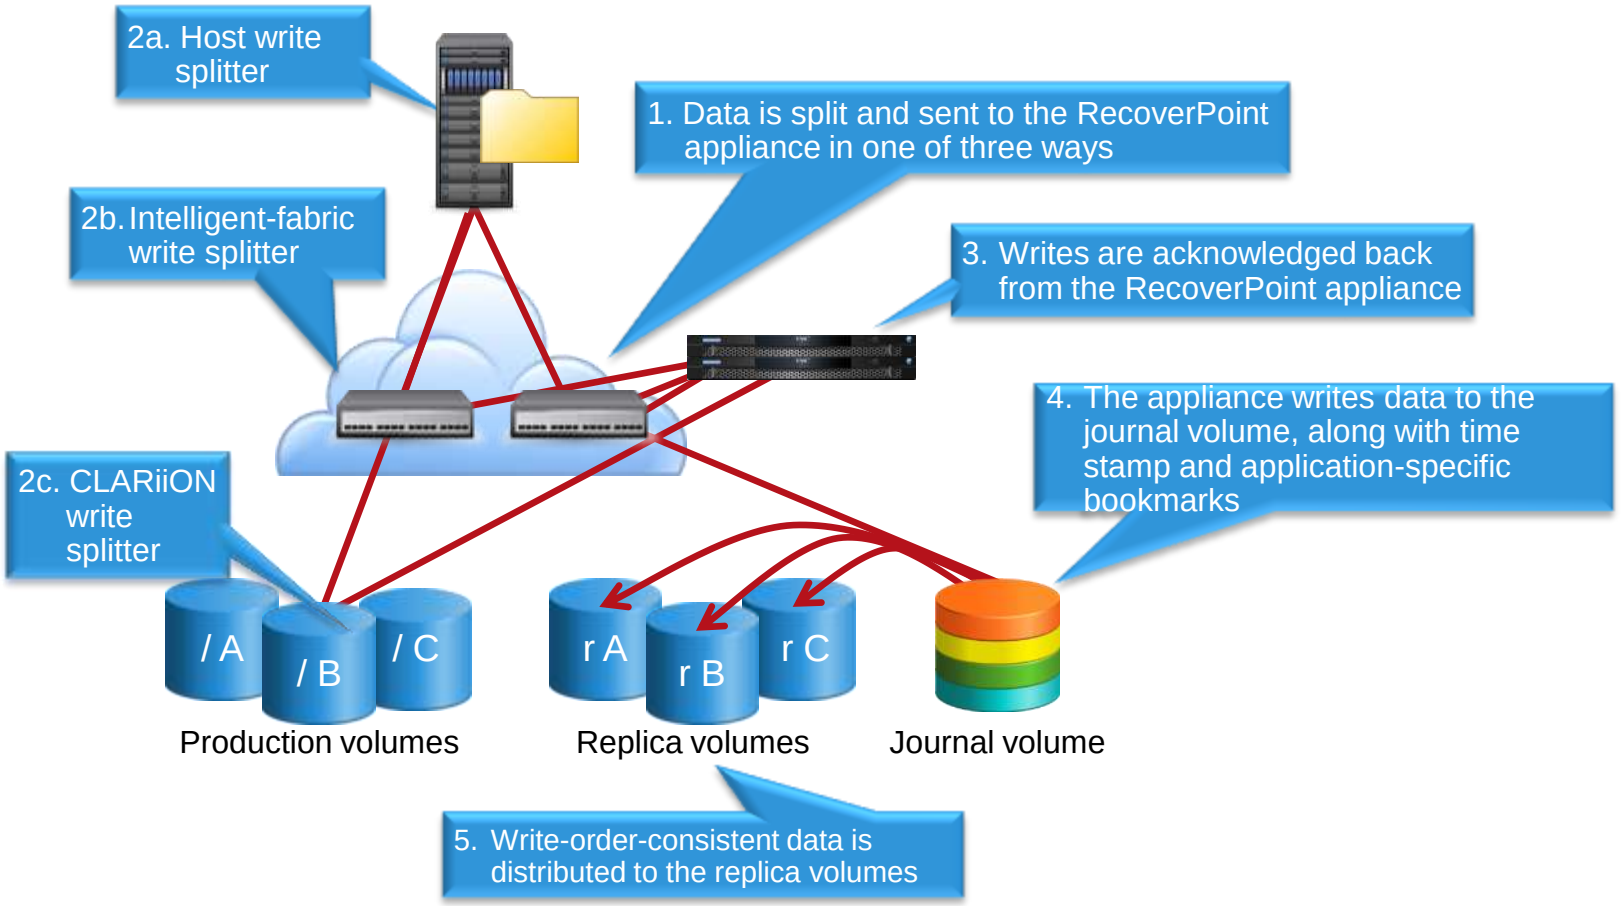

Remote Protection Process: Continuous Remote Replication (CRR)

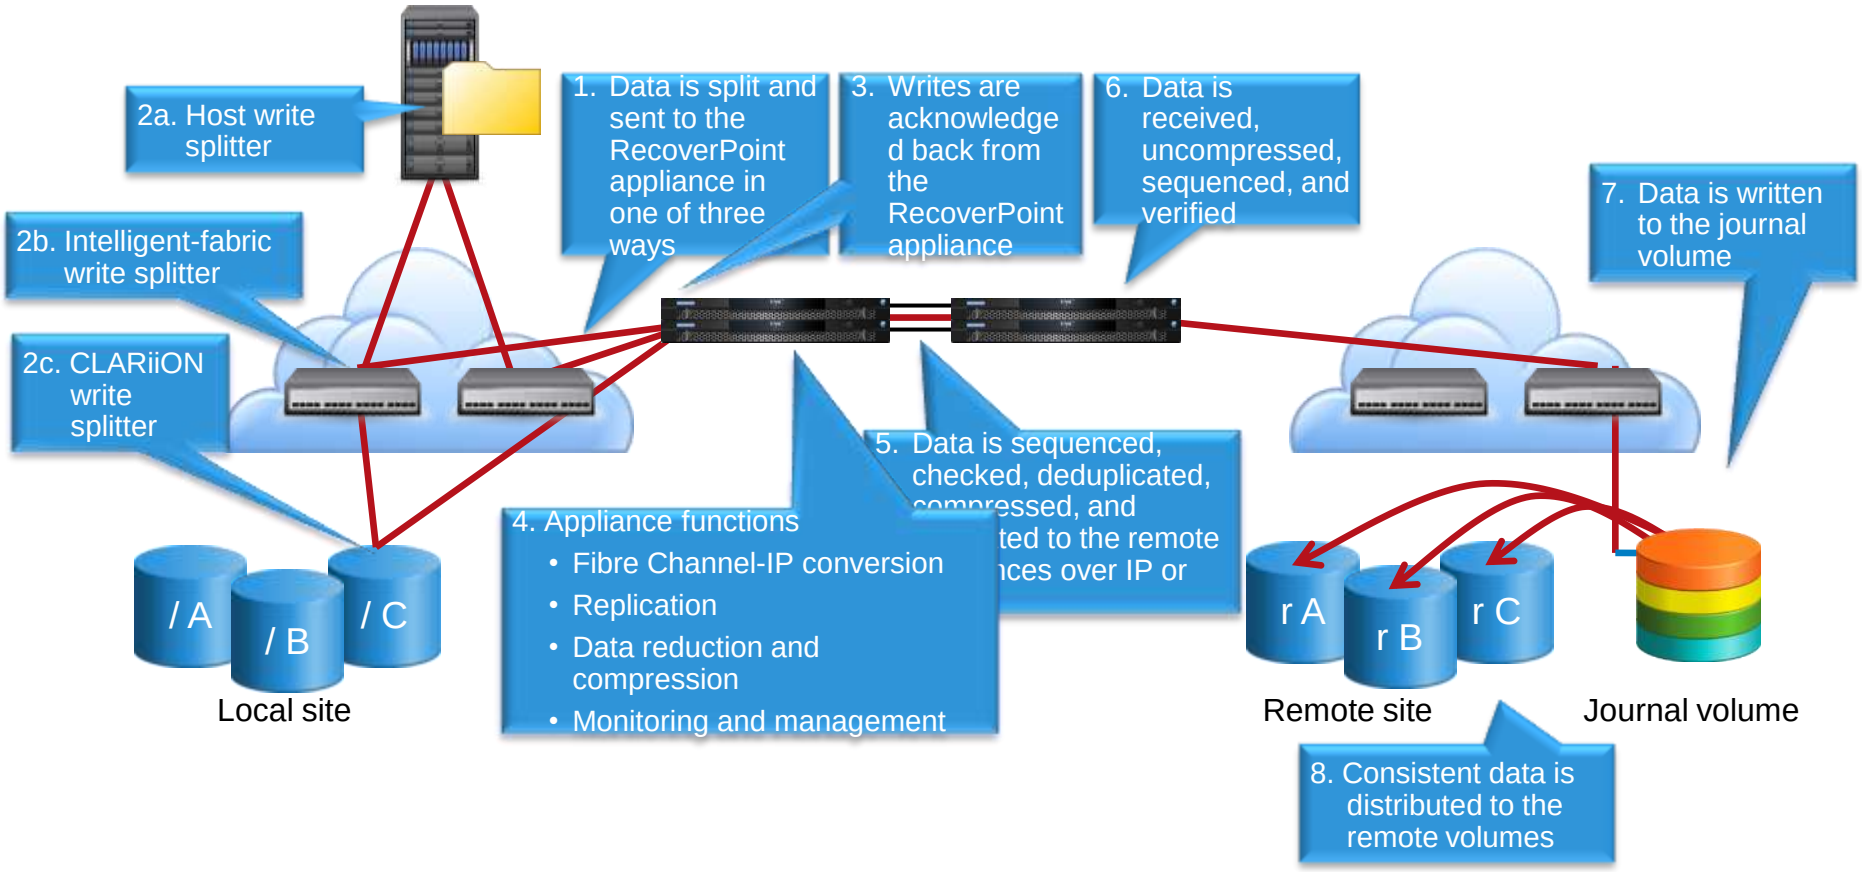

Synchronous Remote Replication: Zero RPO

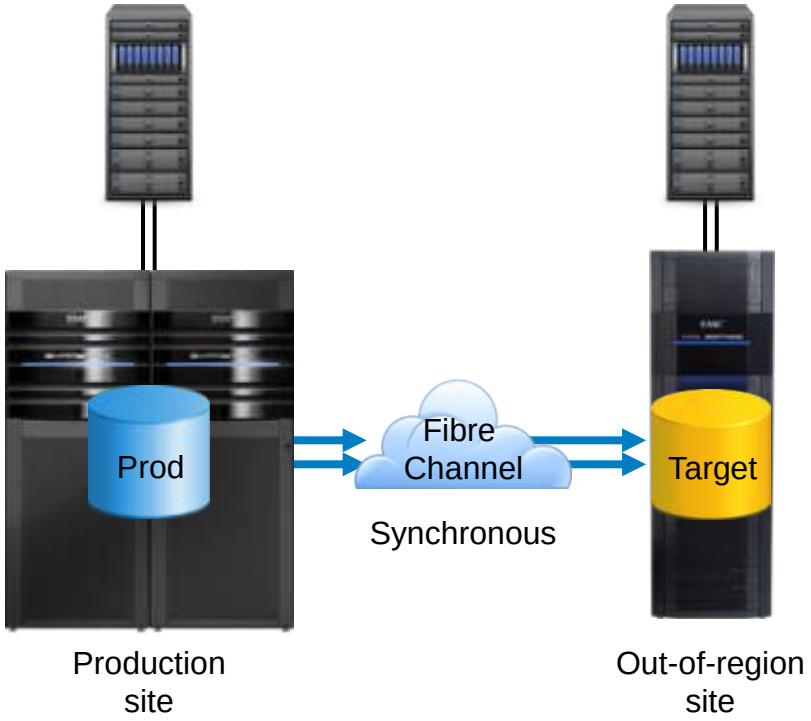

- Replication only supported over Fibre Channel network
- Sites separated by up to 200 km
- Waits for acknowledgement from the remote site
- Can dynamically switch to and from asynchronous
- All synchronization will be asynchronous
- Supports RecoverPoint and RecoverPoint/SE

Stretched CDP over Fibre Channel

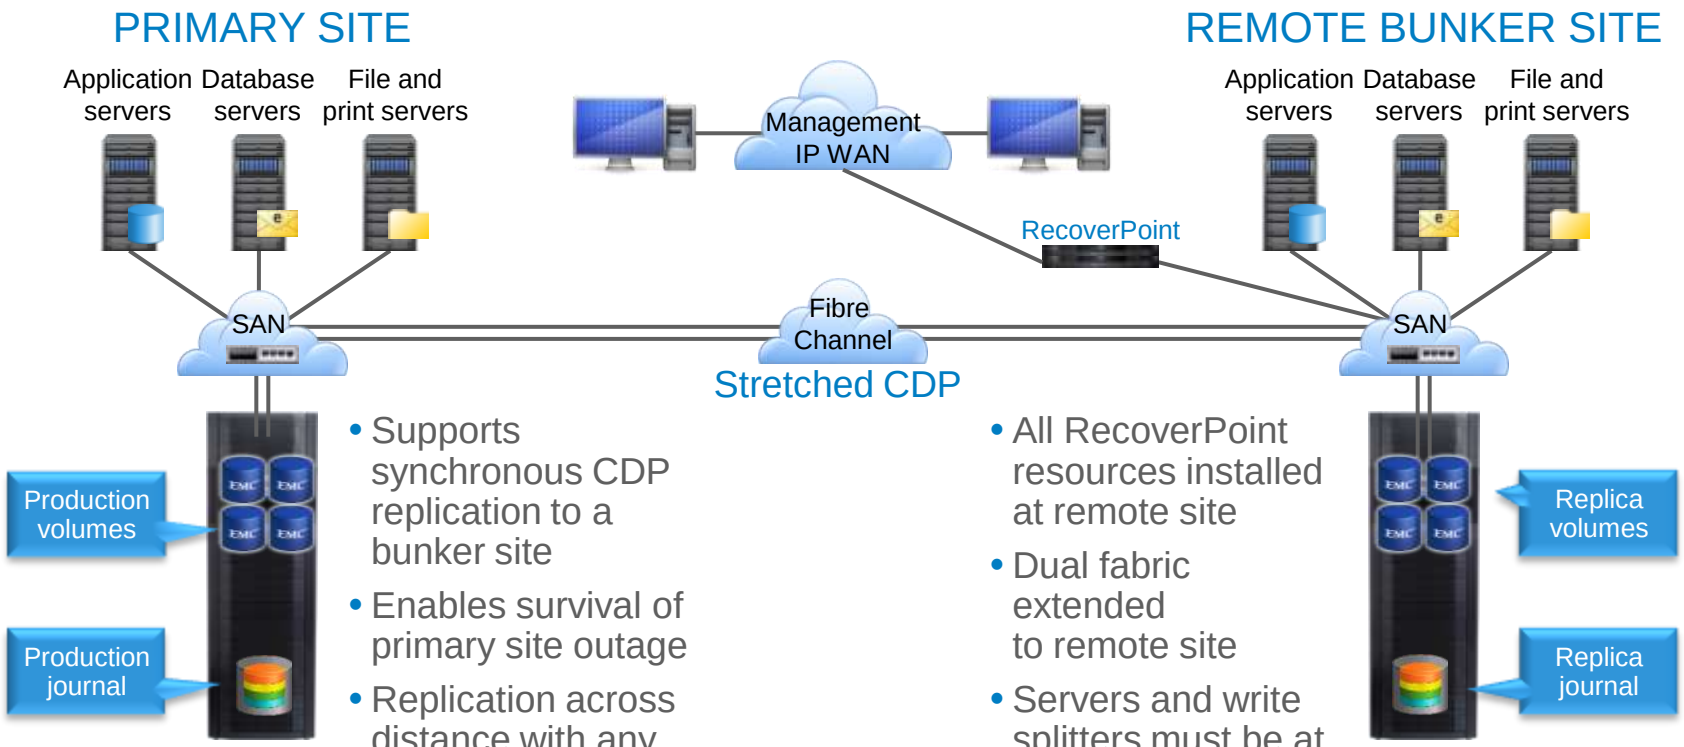

- Supports synchronous CDP replication to a bunker site
- Enables survival of primary site outage
- Replication across distance with any point-in-time recovery
- Applications can access any point-in-time image as read/write without impacting production

- All RecoverPoint resources installed at remote site
- Dual fabric extended to remote site
- Servers and write splitters must be at both sites
- Promoting CDP replica to production is not supported

Not supported by RecoverPoint/SE

Cascaded Replication Topology

Local and remote data protection minimizes data loss

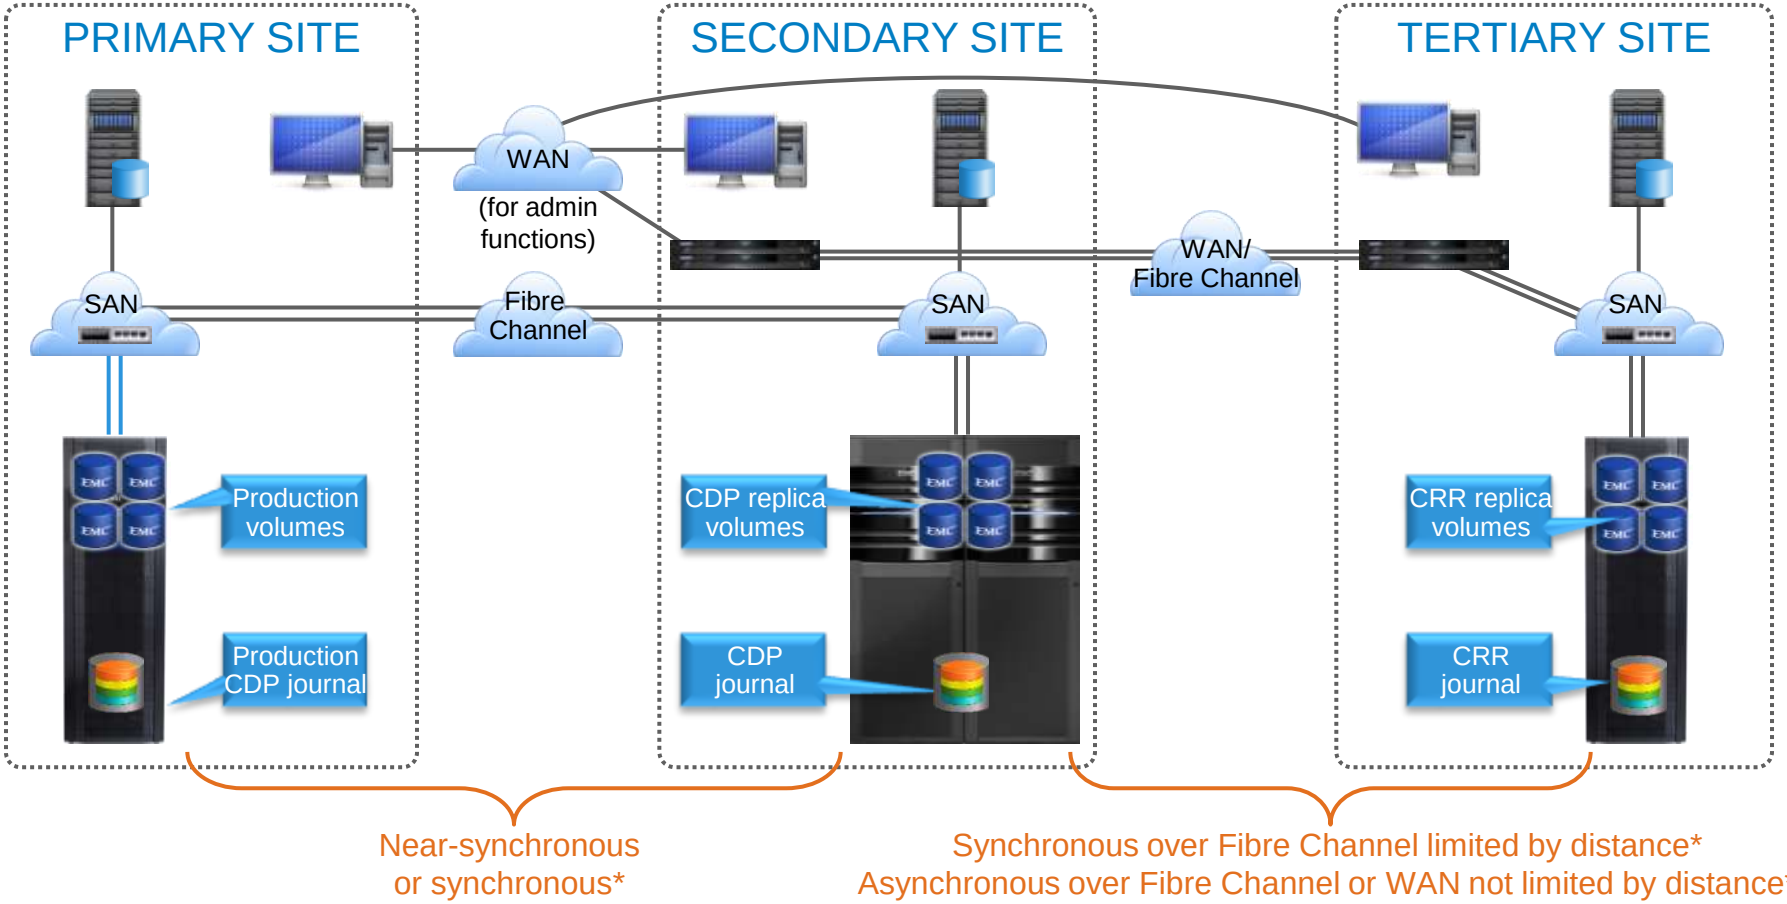

** Refer to the EMC Support Matrix for the maximum distance for your configuration*

Not supported by RecoverPoint/SE

Continuous Recovery Points

- Daily backup → Daily recovery points—from tape or disk
- Snapshots → More frequent disk-based recovery points
- Any point in time
- Significant point in time

} All recovery points

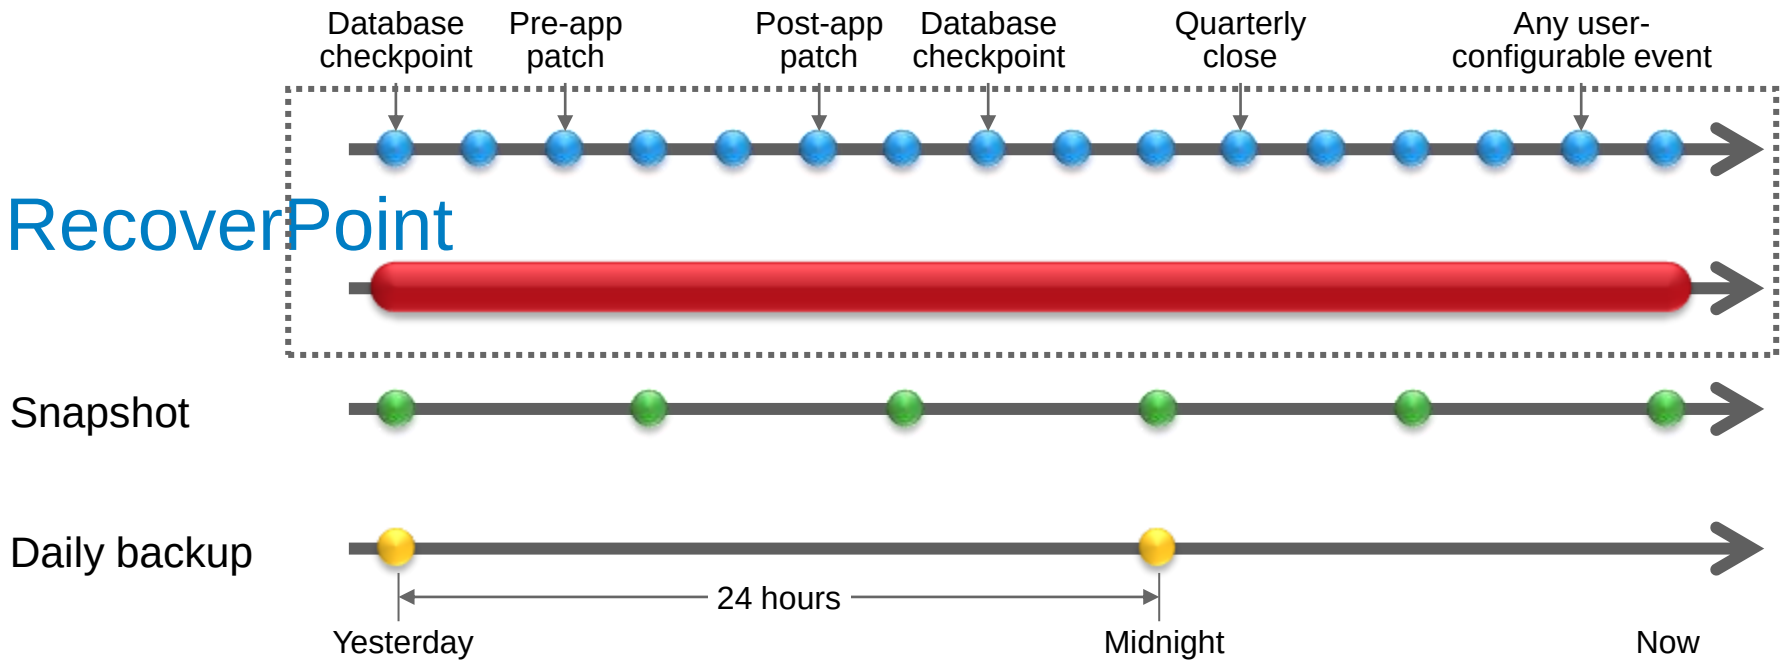

Journaling for Application-Aware Recovery

Journal includes replicated data plus metadata

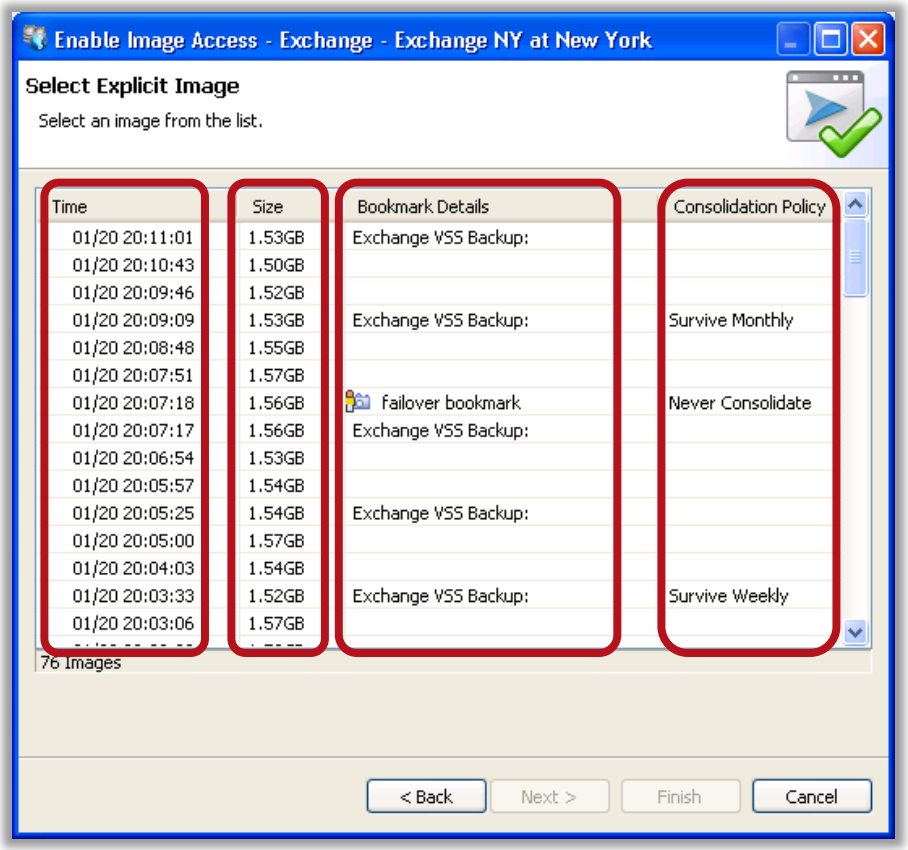

- Time/date:
  - Exact time (down to milliseconds) that the image was stored
- Size:
  - Size of image
- Image bookmark details:
  - System-, EMC-, or user- generated bookmarks
- Bookmark consolidation options
  - System decides (blank)
  - Survive a daily, weekly, or monthly consolidation
  - Never consolidate

Disaster Recovery for Virtualized Environments

Integration with VMware vCenter Site Recovery Manager

Site Recovery Manager requires replication technology

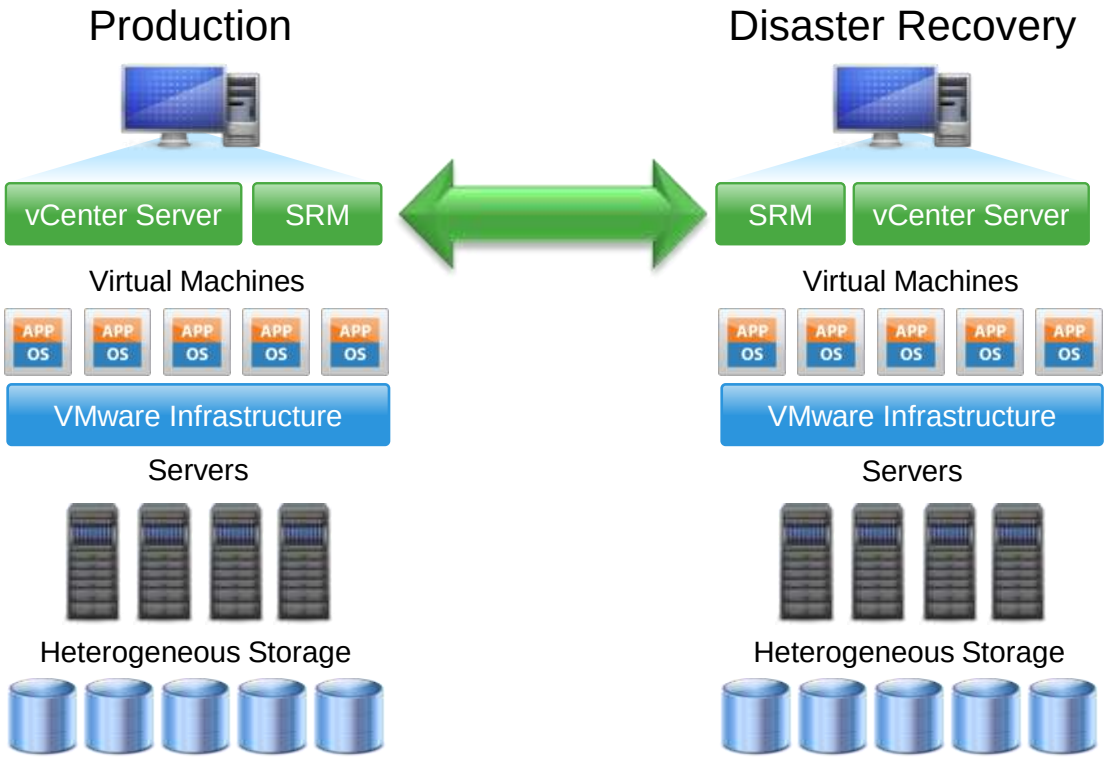

- Leverages EMC advanced replication software
- Removes manual recovery complexity
- Centralizes management of recovery plans
- Simplifies and automates disaster recovery workflows:
  - Configuration
  - Testing
  - Failover
  - Failback

EMC² where information lives

THANK YOU