Внутренняя архитектура систем корпоративного уровня

Sergey Pospelov

Engineering Manager, Mainframe
Sergey.Pospelov@emc.com
July 2013

EMC²

Mainframe at EMC World

2012 Mainframe Highlights

- Market
  - IBM announces zEnterprise EC12 8/2012
    - IBM System Z Q412 revenue +56%
- EMC MF Product Releases / Features
  - Virtual Provisioning for Mainframe
  - FAST VP for Mainframe
  - SRDF/SQAR
  - DLm 8000 – DLm/GDDR integration
  - EzSM 4.1 support for DLm
  - 5 MFE maintenance sets across 3 releases
- ESD Mainframe Organization Established
  - VP hired
  - Disparate organizations consolidated
    - QE, PM, Dev, DVT, MF labs

IBM and EMC Partnership for Mainframe

IBM LICENSED PARTNER

EMC MEMBER OF IBM ESP

IBM MEMBER OF EMC VMAX EAP

COOPERATIVE SUPPORT AGREEMENT

FIRST AND ONLY GDPS LAB

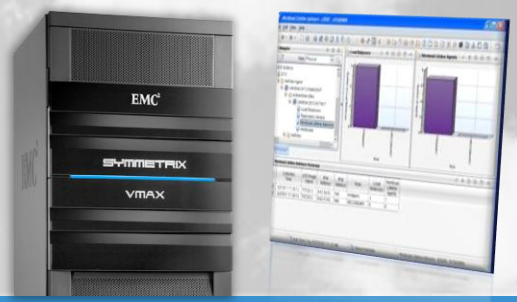

FIRST TO SUPPORT NEW Z PLATFORMS

EMC²

Customer Voice

- Information growth = 50%+ a year
  - Headcount growth = not even close
- Don’t want to be a plumber
- Need to use much less power
- Reduce my total cost
- Need to do more with fewer people
- Protect and secure my information
- My business is 24 x Forever
- Reduce the complexity

Key Requirements of Storage Systems

Storage Consolidation

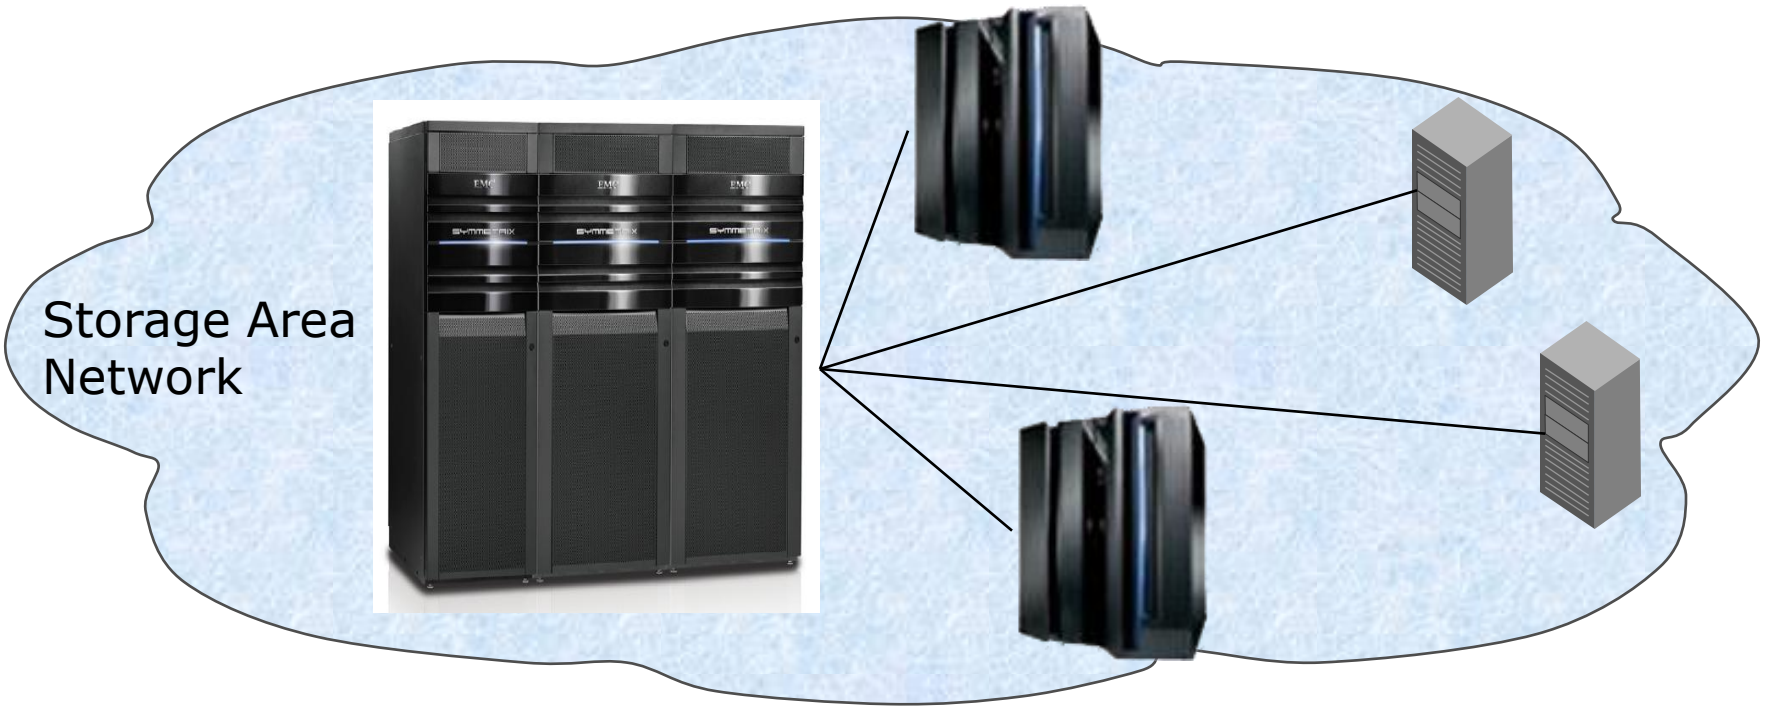

- ❑ High data availability
- ❑ Data protection
- ❑ Business continuity
- ❑ Manageability
- ❑ Reduce cost
- ❑ Data deduplication

The World's Most Trusted Storage System

25 Years Running the World's Most Critical Applications

1988	1990	1994	2000	2003	2005	2009	2011	2012	2013
 Symmetrix 4000	 Symmetrix 5500	 Symmetrix 3000/5000	 Symmetrix 8000	 Symmetrix DMX-1, -2	 Symmetrix DMX-3, -4	  Symmetrix VMAX 20K	 Symmetrix VMAX 10K	 Symmetrix VMAX 40K	New
ICDA, RAID, NDU	SRDF Consistency Groups	TimeFinder, PowerPath	AutoSwap, SRDF over Fibre Channel	Concurrent SRDF, SRDF/Star, GDDR	Secure Erase, Audit, Service	z/OS integration, DARE,FAST VP, Federated Live Migration	FAST VP, FLM, 100% VP, RP Splitter	3X scale Dense Configuration Option System Bay Dispersion Federated Tiered Storage FAST VP for System z, IBM i	

POWERFUL TRUSTED SMART

Storage Scalability

System Bay

Symmetrix V-Max 40 000

V-Max maximum specification:

- 32 six-core 2.8 GHZ Xeon ® processors
- 2 TB of memory
- Virtual matrix bandwidth 400 GB/s
- 128 Front End ports
- 128 Back End ports
- 4 PB usable storage capacity
- 3200 disk drives

Host connectivity:

- Fibre Channel
- iSCSI
- 10 Gigabit Ethernet
- FICON connectivity

Disk drive connectivity:

- 3.5″ Fibre channel drives (2 TB)
- Enterprise Flash drives (400 GB)
- SATA drives (2 TB)
- 2.5″ SAS drives (600 GB)

V-Max Director

FBDIMM Modules

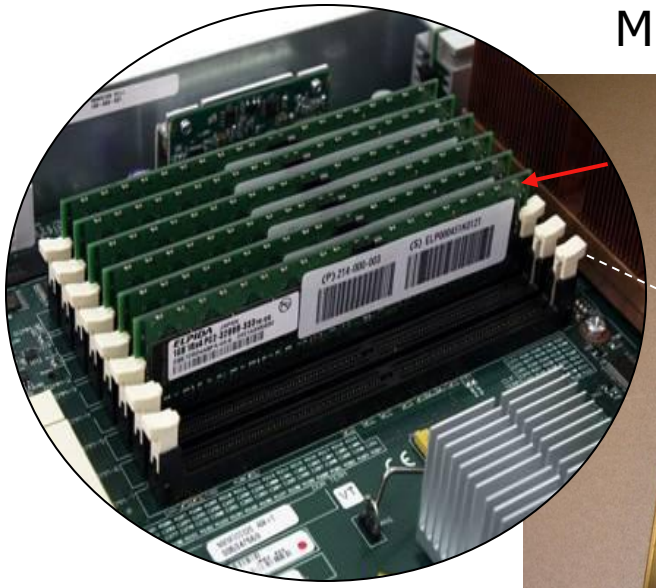

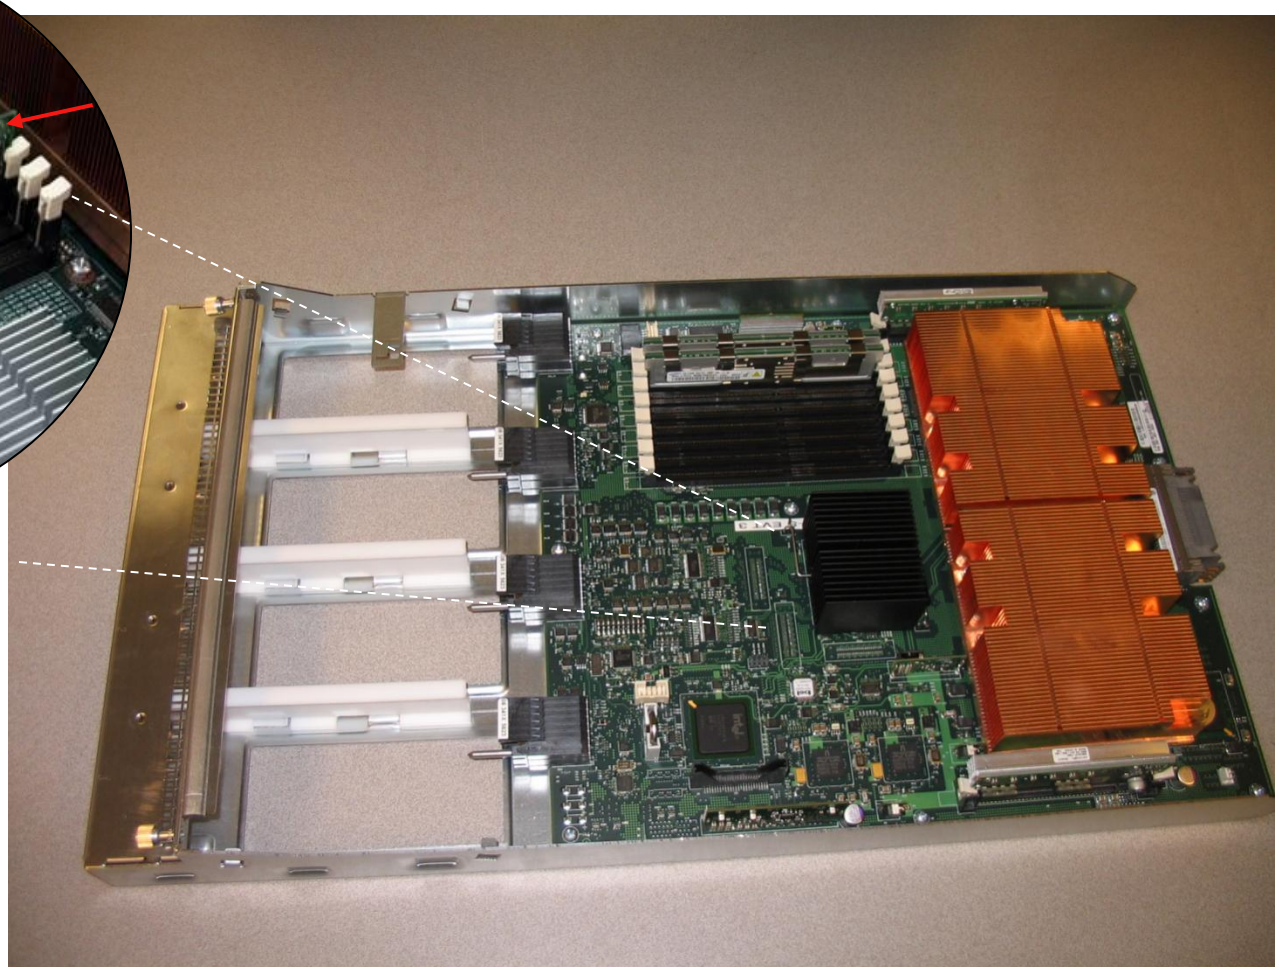

Feature	Description
Dimensions	2-¼”H x 12-½”W x 22-5/8”D
Weight (excluding SLiC’s)	13.8 lbs

V-Max Engine

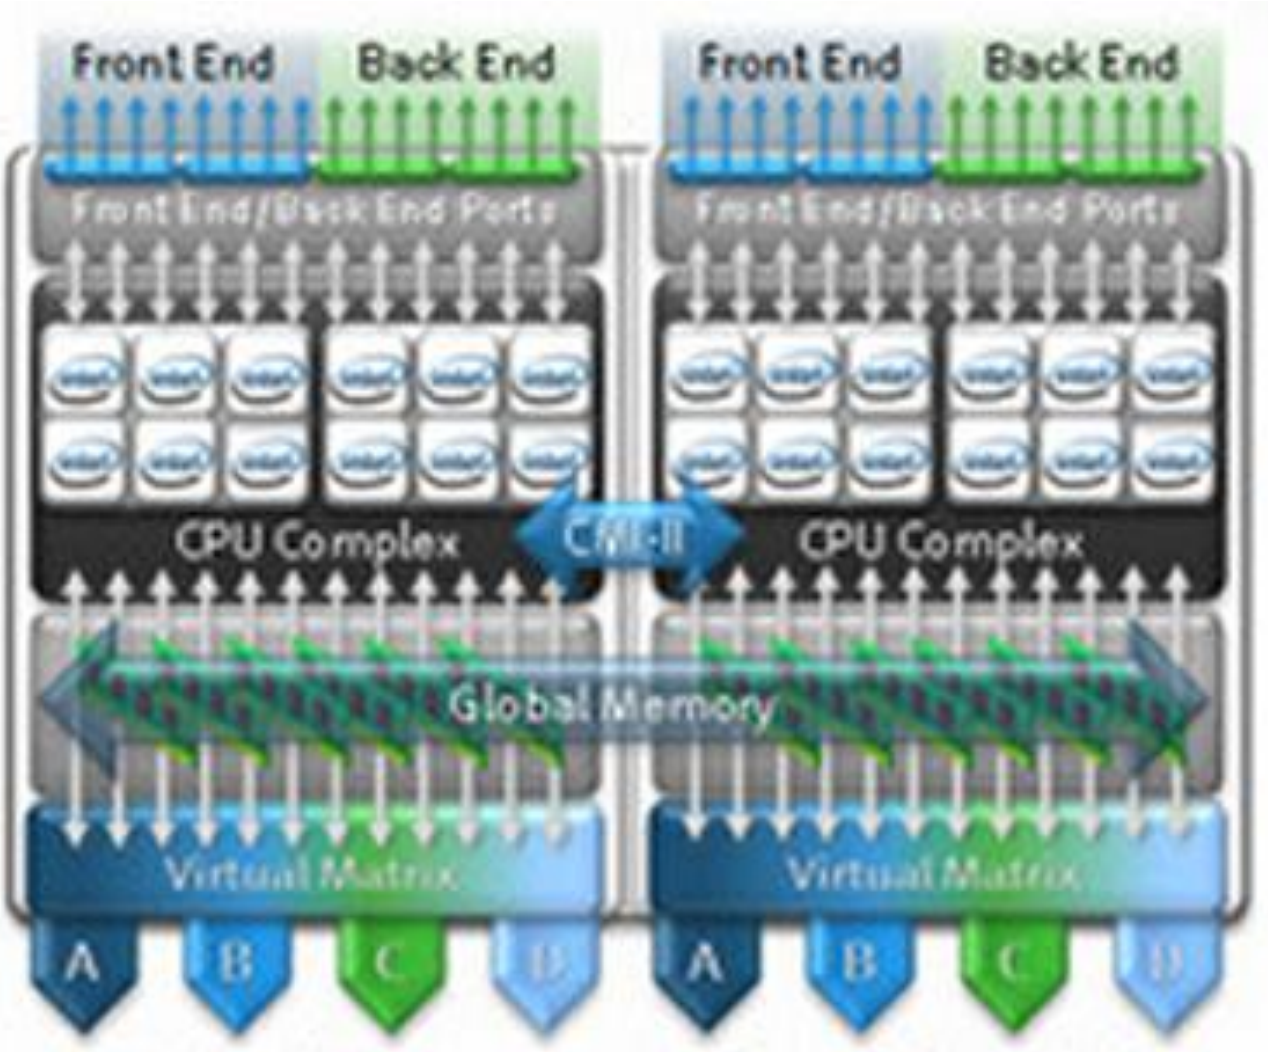

Virtual Matrix Architecture

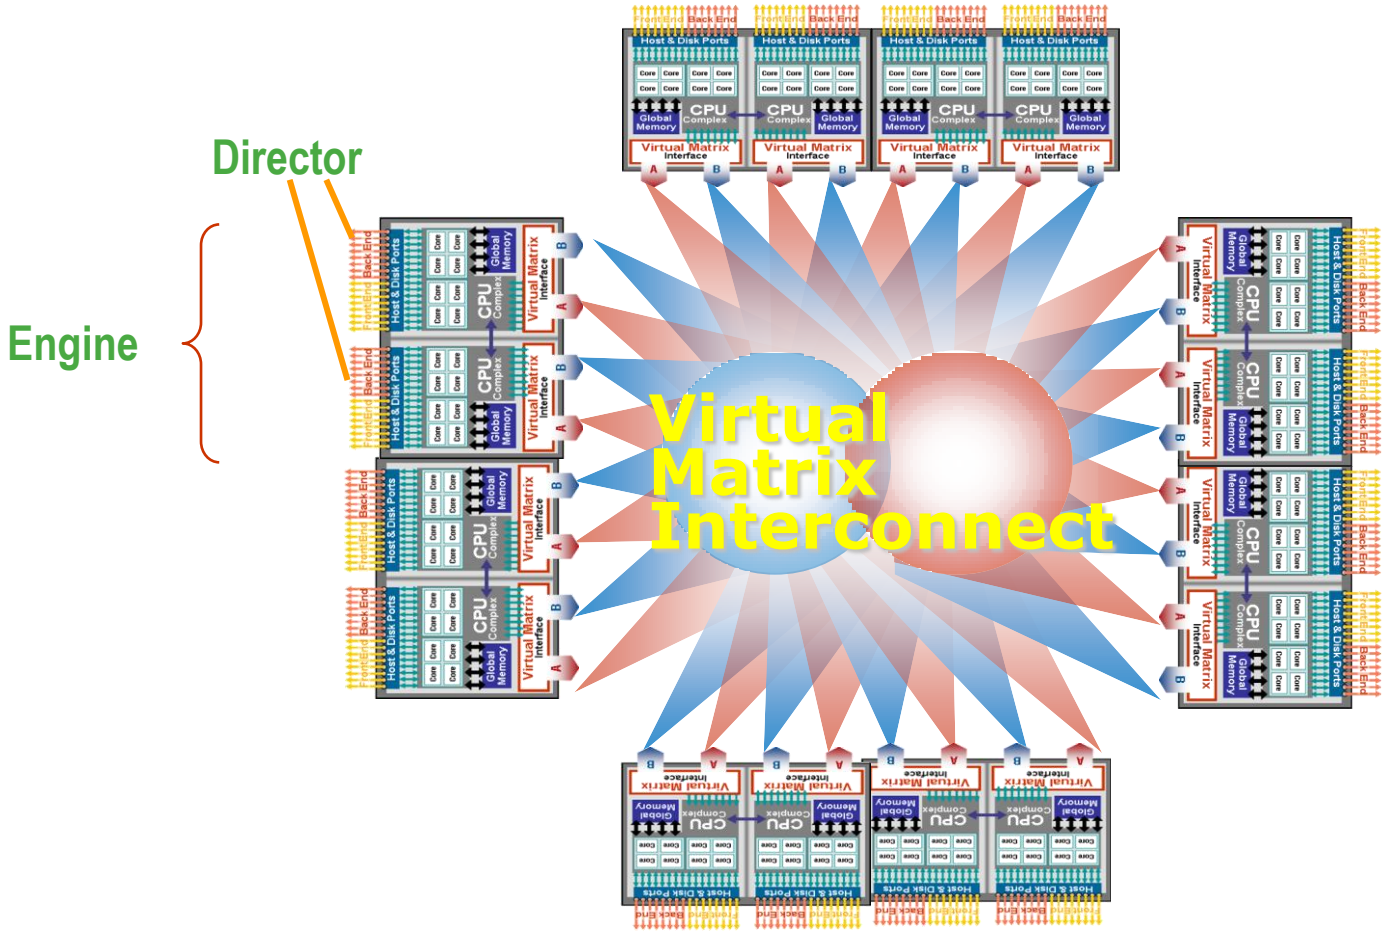

VMAX Interconnect: RapidIO fabric®, InfiniBand®

Basic Symmetrix Architecture

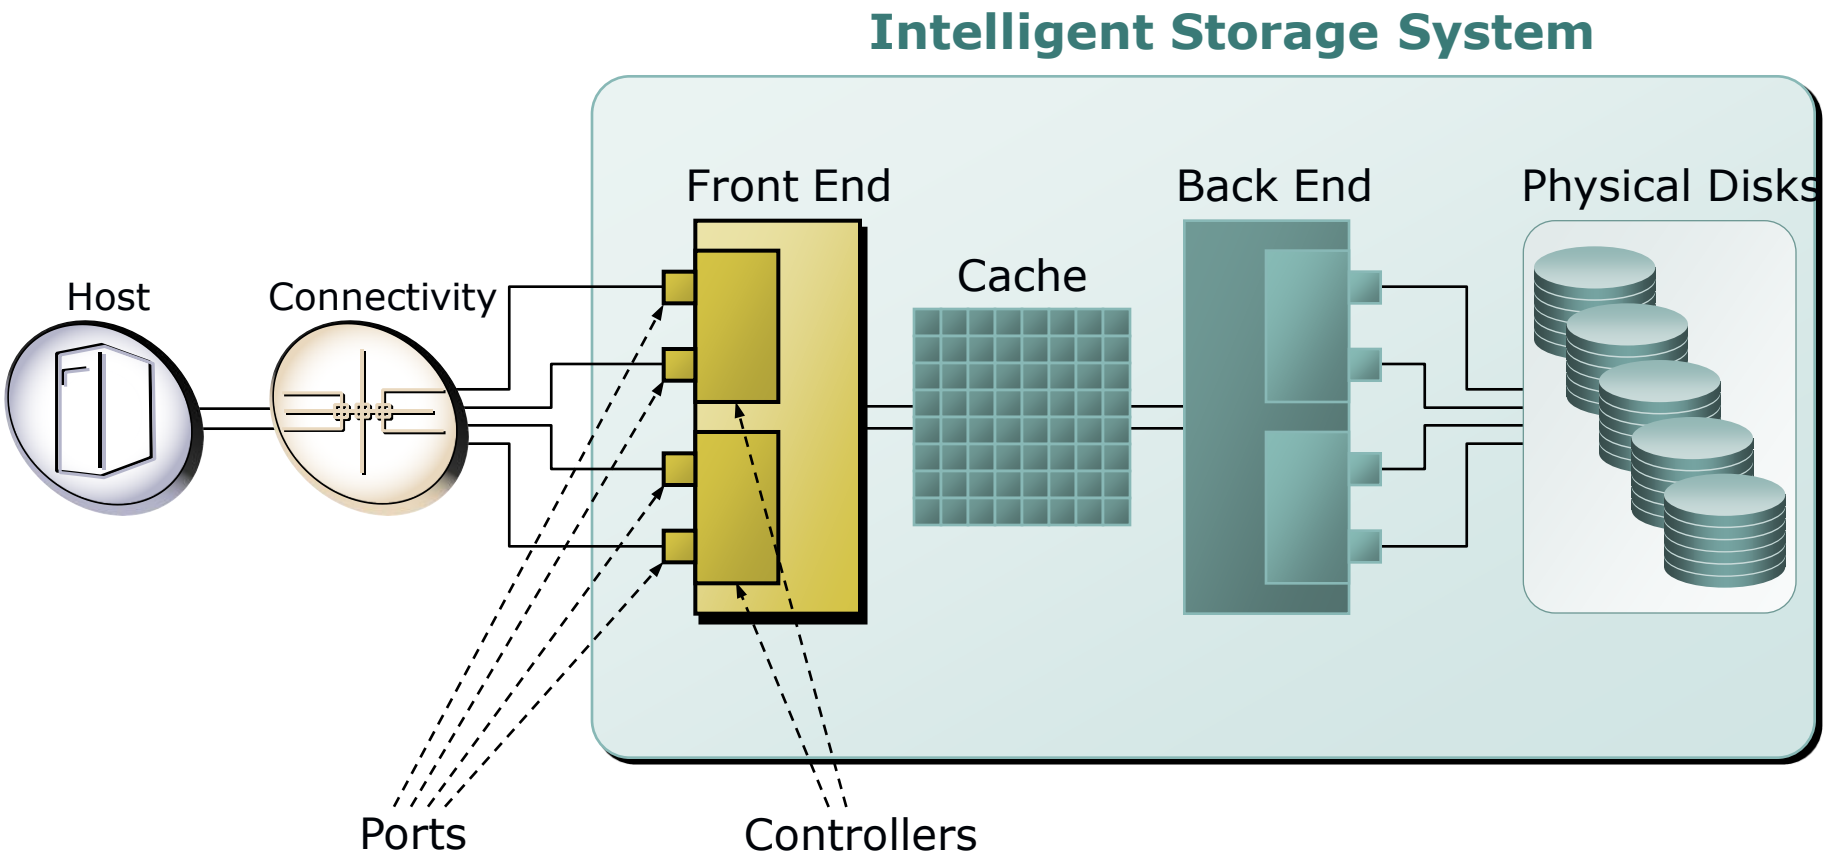

The Enginuity storage operating environment provides the intelligence that controls all components in Symmetrix.

Cache Read, Cache Write

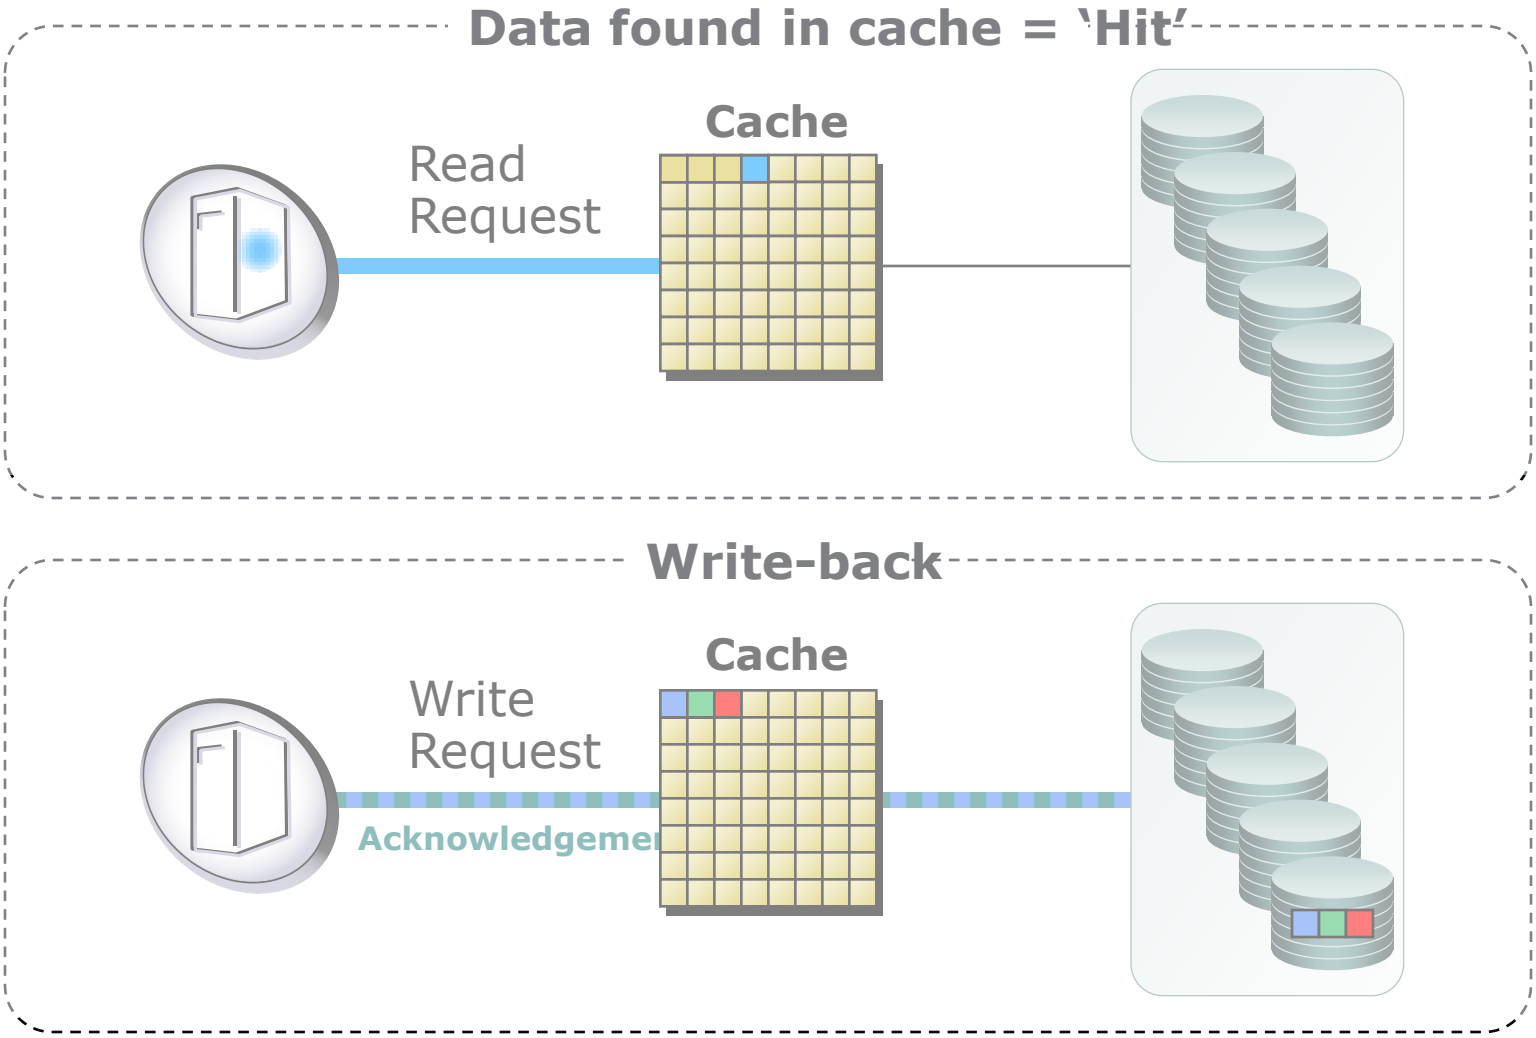

Power Vault

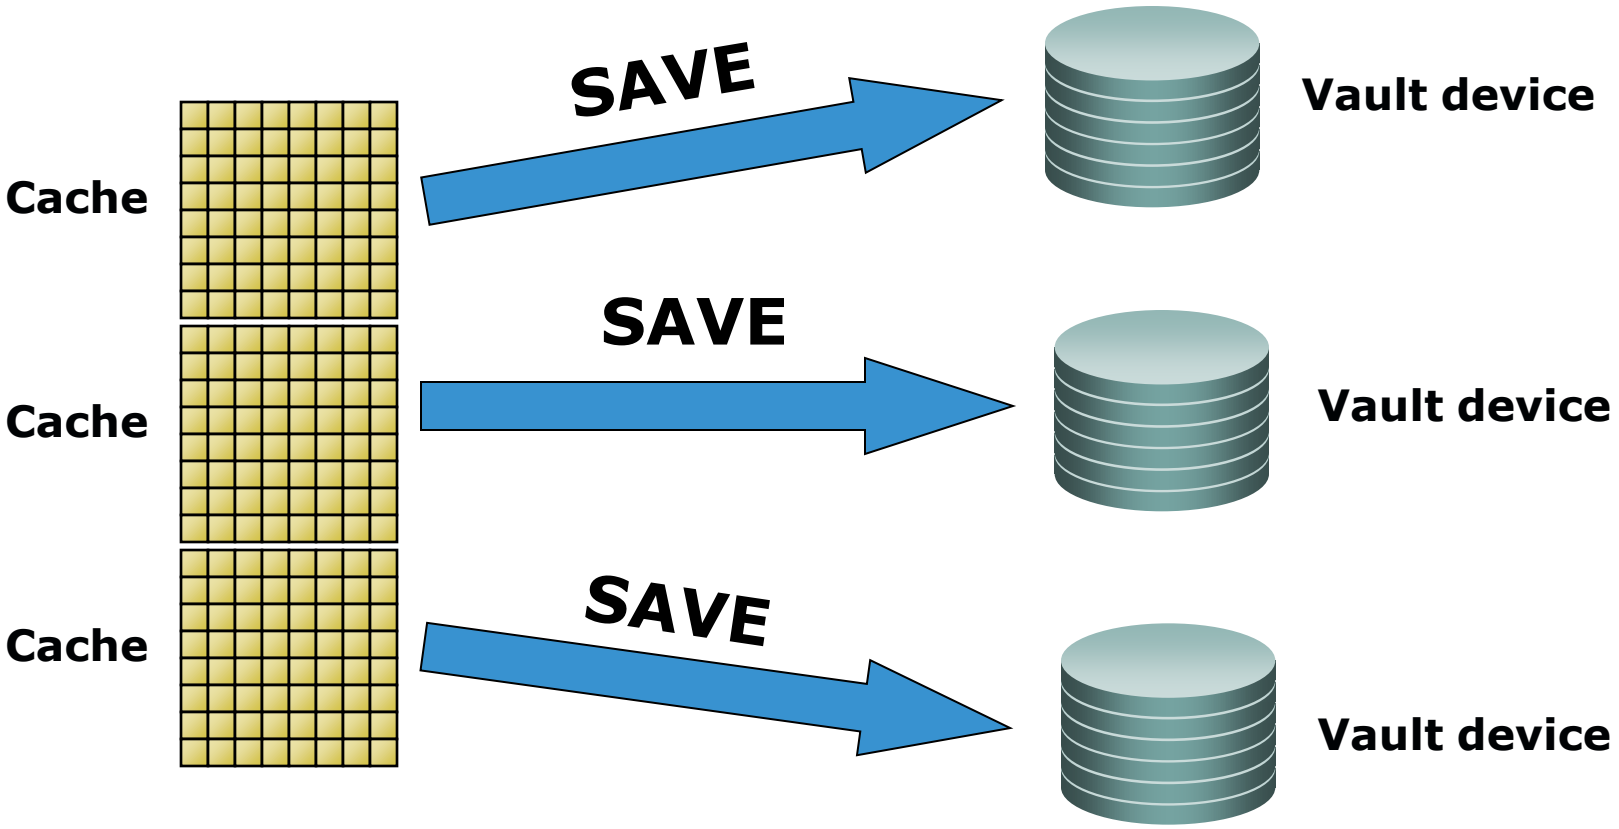

- ✓ Allows the contents of Global Memory to be saved in less than 5 minutes.
- ✓ 2 copies of Global Memory are written to the Vault Devices.

Symmetrix Logical Volumes

- RAID1 SLV:
  - Data is written to two hyper volumes on two different physical disks which are accessed via two different disk directors.
- Host is unaware of data protection being applied.

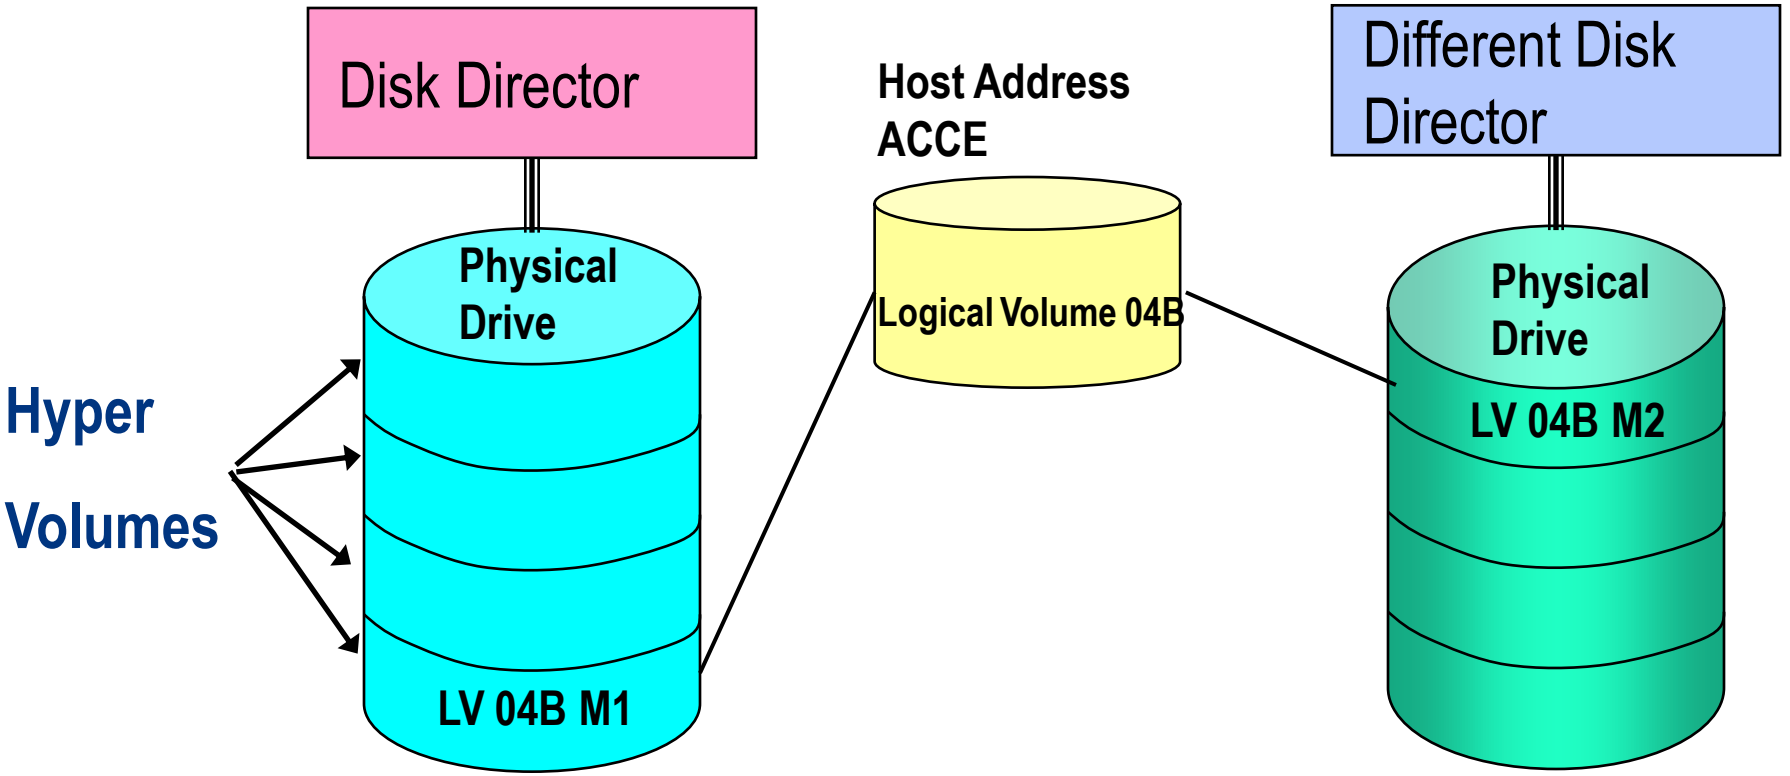

Hot Spares

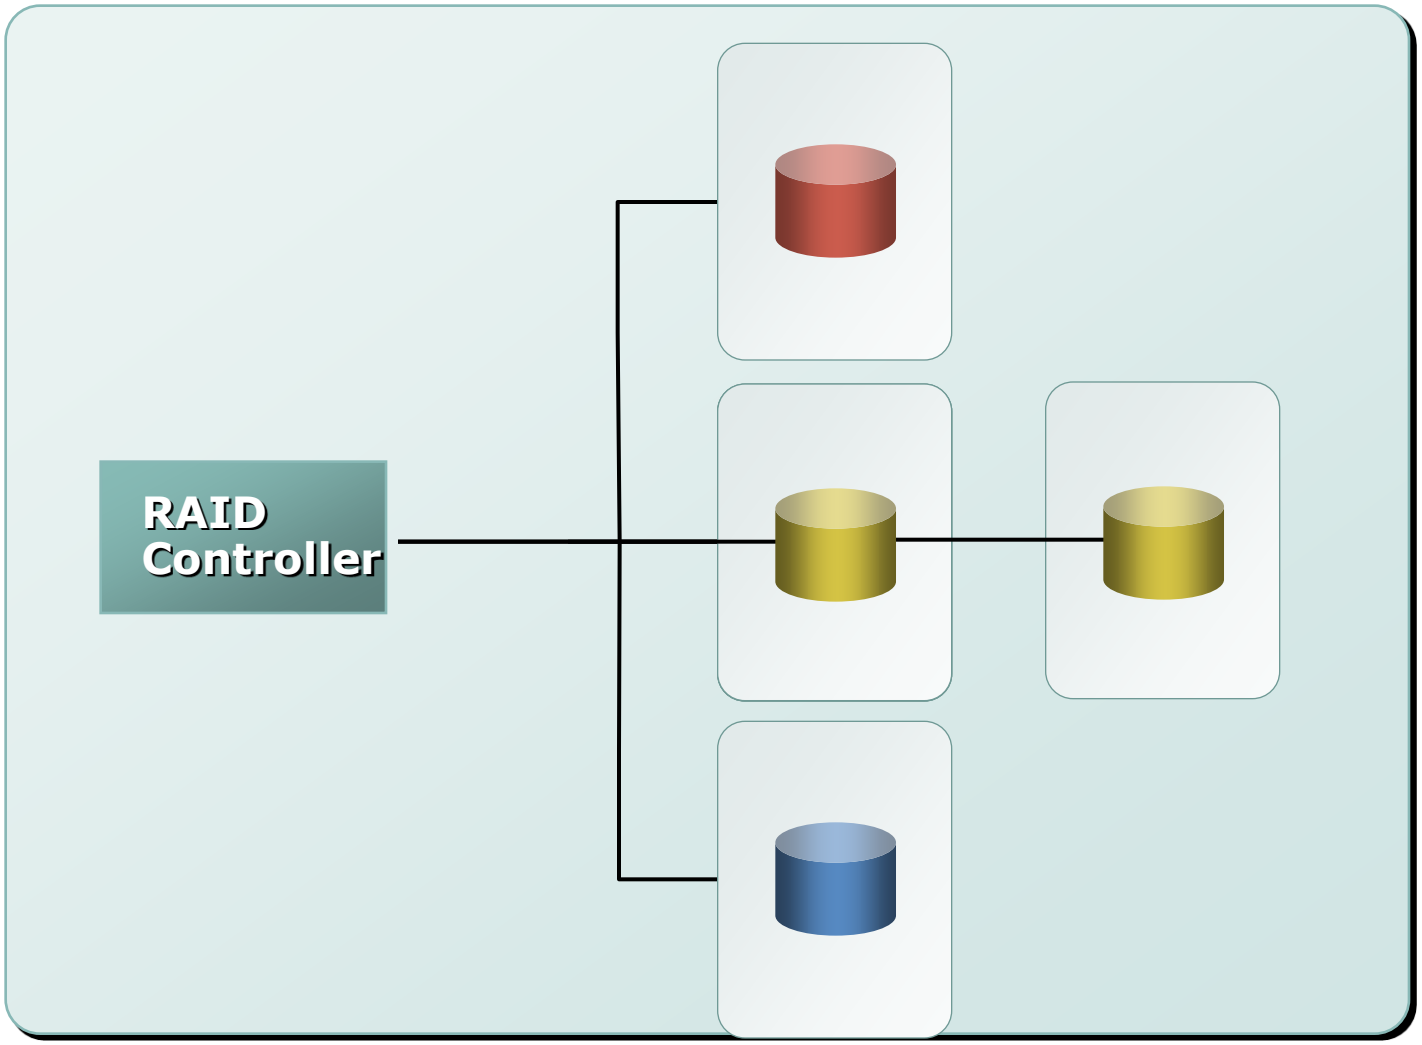

Shadow Partner

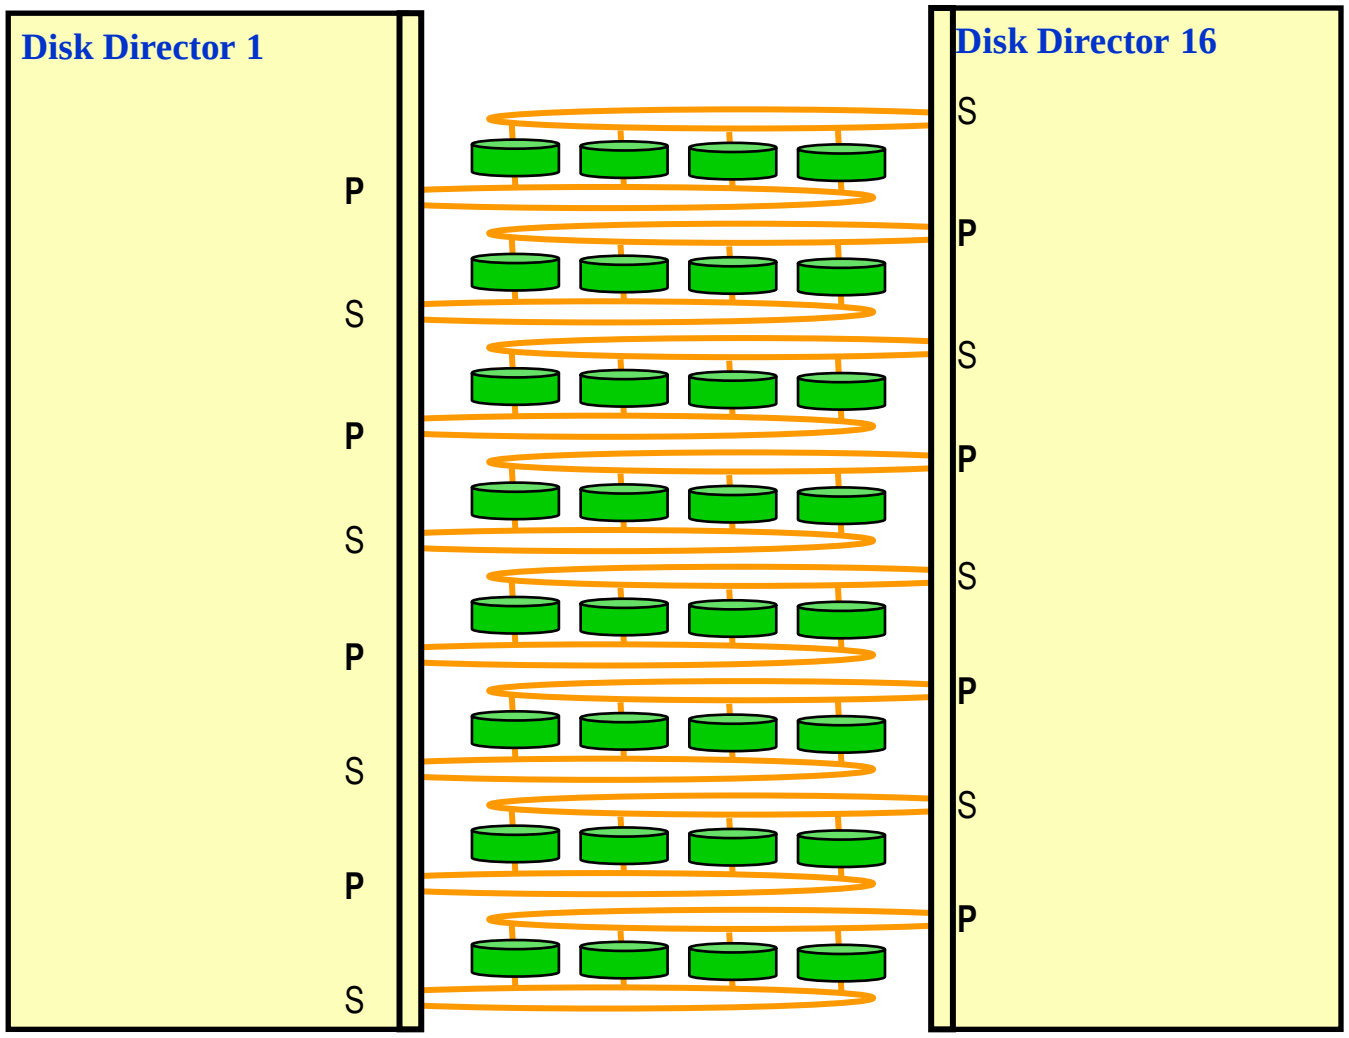

Global Memory : Design Considerations

- Each Director can be configured with 16 GB, 32 GB or 64 GB of memory.
- Directors of a given Engine *must* have the same memory configuration.
- In a single Engine system, memory is mirrored within the same Engine.
- In multiple Engine systems (2 thru 8 Engines), memory is mirrored across Engines.
- This does not require that all Engines have identical memory configurations.

Example: A

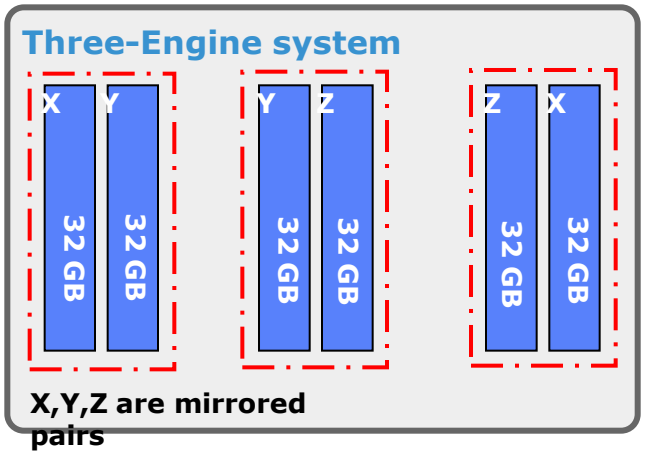

Example: B

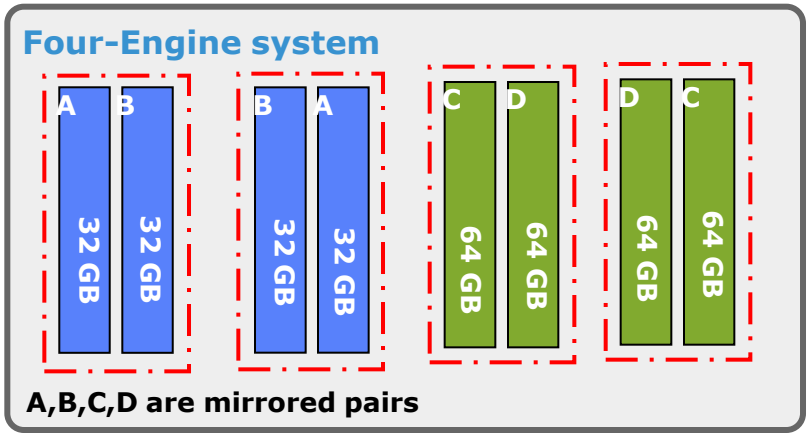

Thin Devices. Data Devices

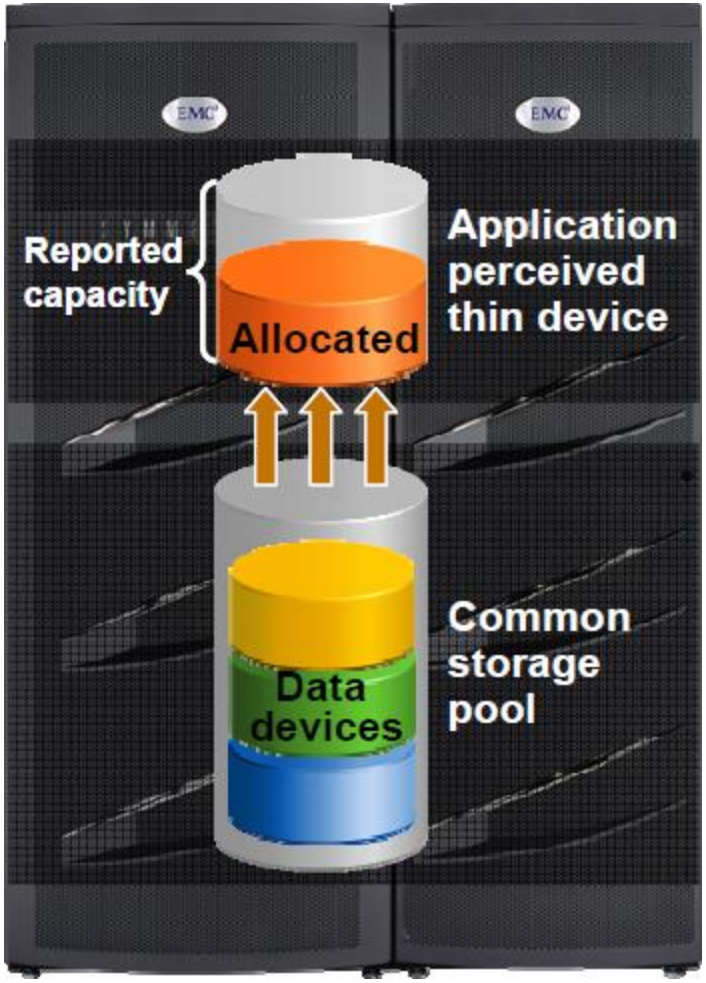

Thin devices do not need to have physical storage completely allocated at the time the device is created.

Thin Storage Pool is comprised of devices called data devices that provide the actual physical storage.

Writing to a portion of thin device allocates a minimum allotment of physical storage from the pool.

Virtual Provisioning

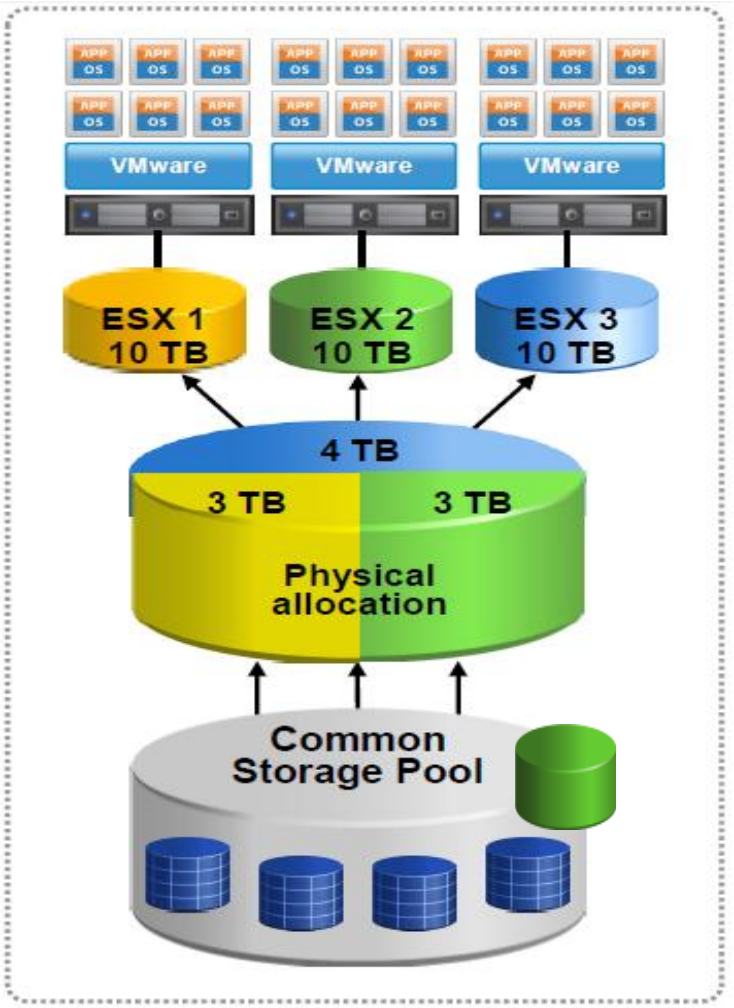

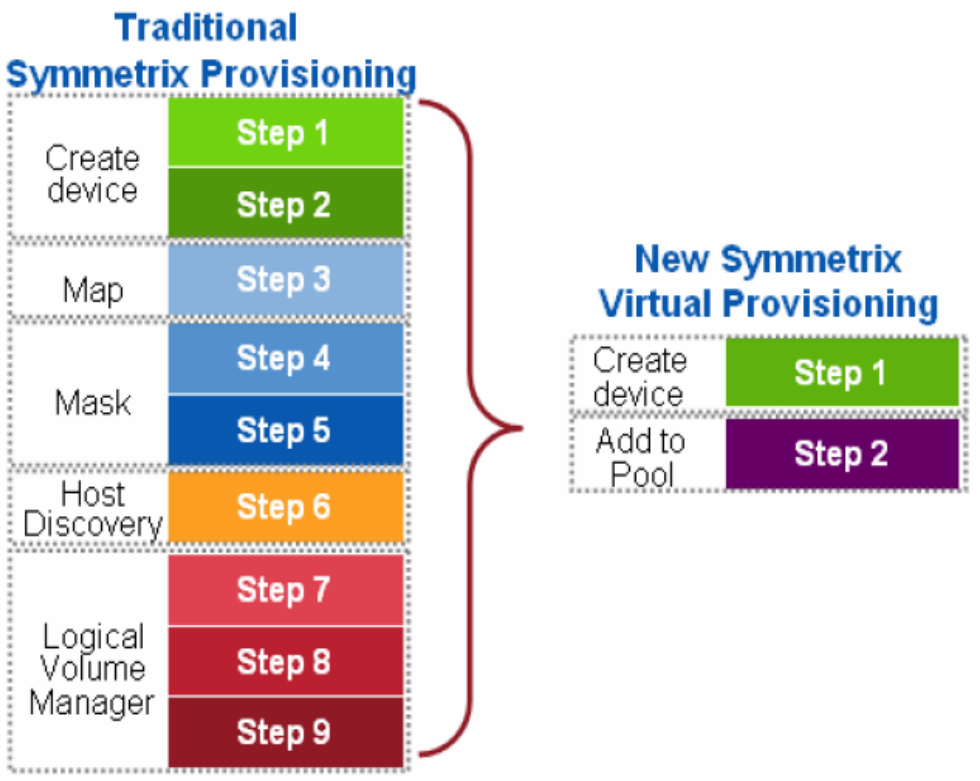

Fully Automated Storage Tiering (FAST)

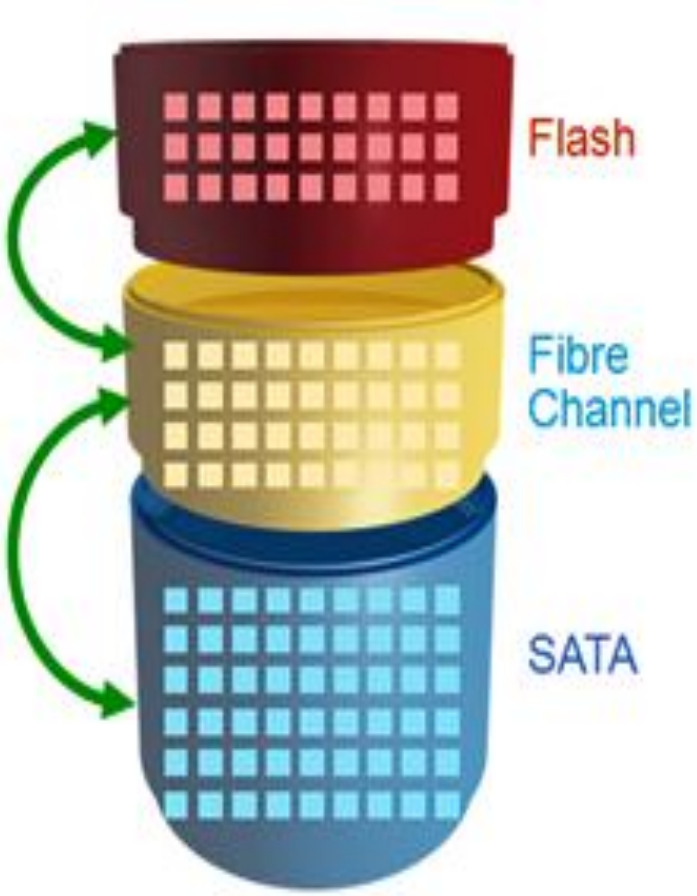

Puts the right information in the right place at the right time at the right cost

- Performance-sensitive applications go on mirrored 15K rpm drives.
- Bulk data and backups go on large SATA drives, ideally protected by RAID 6.
- Anything we aren't sure of goes on 10K rpm drives, possibly using RAID 5.

4% Flash Drives - 2.5X Faster System Response Time

EMC²

Dynamic data mobility: FAST

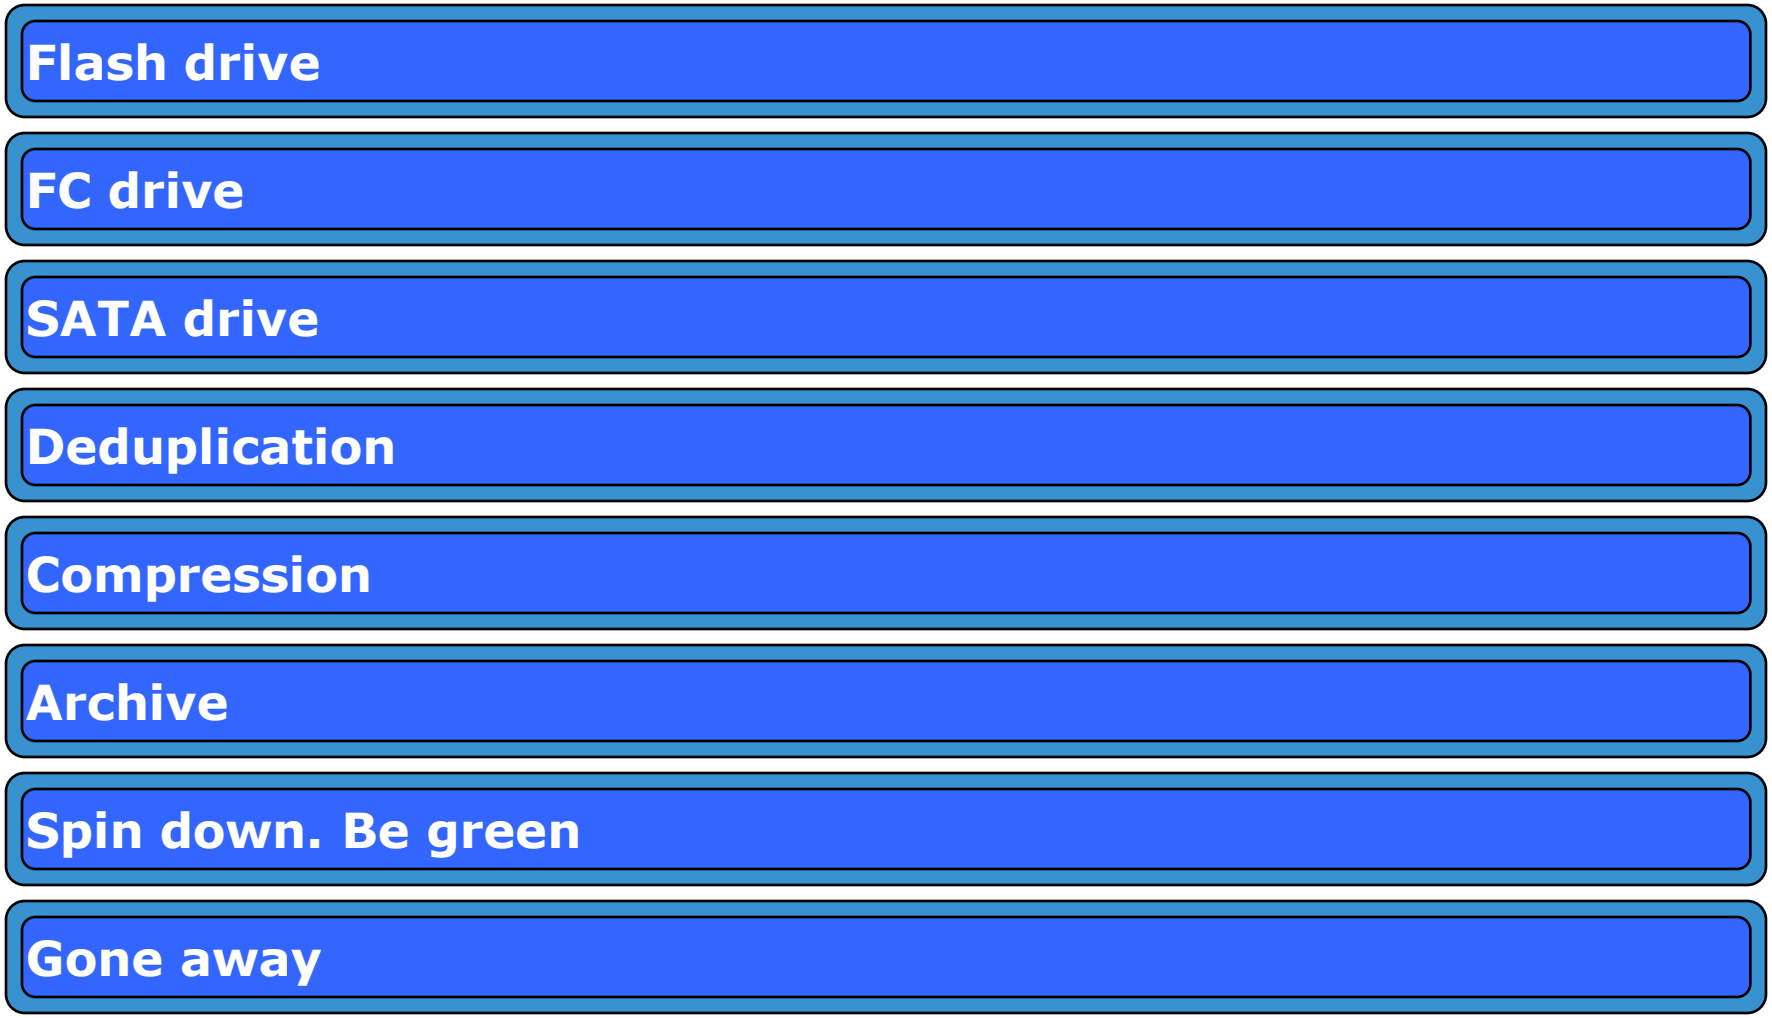

Building Mission Critical Systems

1. Do Not Sacrifice Quality for Cost

Make sound decisions, resilient systems require investments

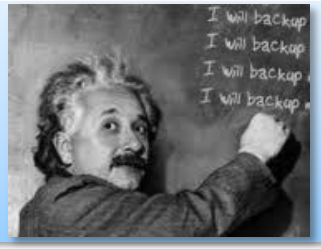

2. Ensure all critical data is replicated

Use Snapshots and Clones for local recovery

Define RTO and RPO to match business requirements

3. Integrate storage and system resiliency

Design redundant systems and automated restart technology

Deploy consistency technology for federated applications

4. Build site failover into regular prod/test process

Plan the flight, test the plan, fly the plane, review the flight

Fly often – build failover into production processes

EMC²

Local Replication

TimeFinder/Clone

- High-performance logical copies.
- Full volume and dataset level.
- Ideal for higher-tier applications.

TimeFinder/Snap

- Space-saving snapshot images.
- Typically requires less than 30% additional capacity.
- Ideal for lower-tier applications.

TimeFinder/Mirror

- Ultra-high-performance mirrors.
- Highly availability full-volume mirror.
- Ideal for higher-tier applications.

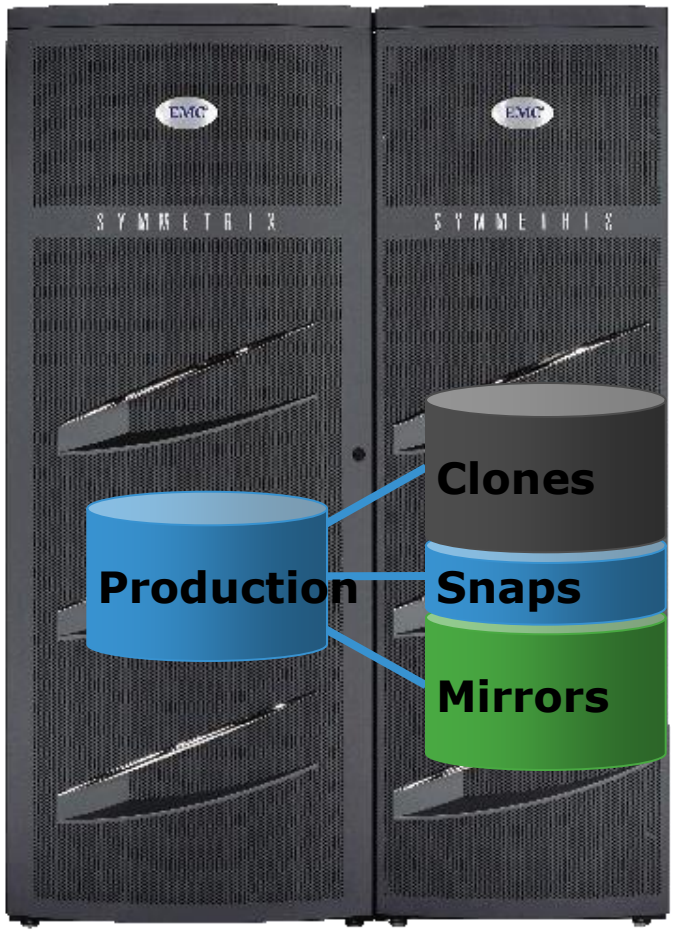

Clone. STD Devices

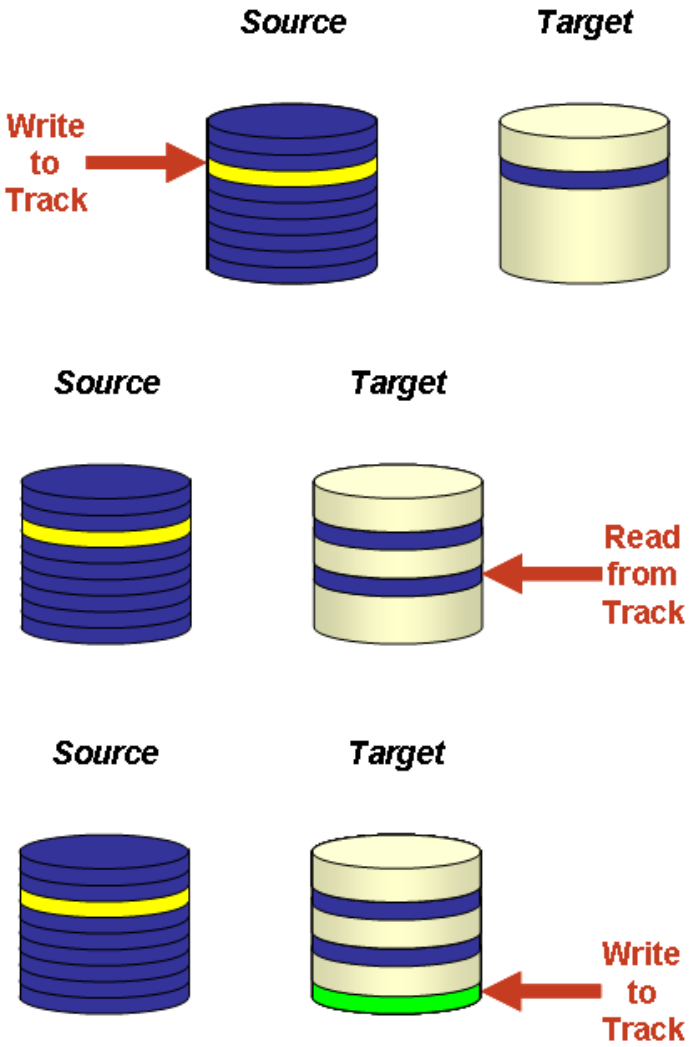

TimeFinder Clone is an instant point-in-time copy of a standard device.

- Copying can be immediate or deferred (copy on access).
- Copy sessions are maintained on the Symmetrix unit.
- Data can be copied from a single source device to as many as sixteen target devices, with automatic background copying of up to four target devices simultaneously.

SNAP. Virtual Devices

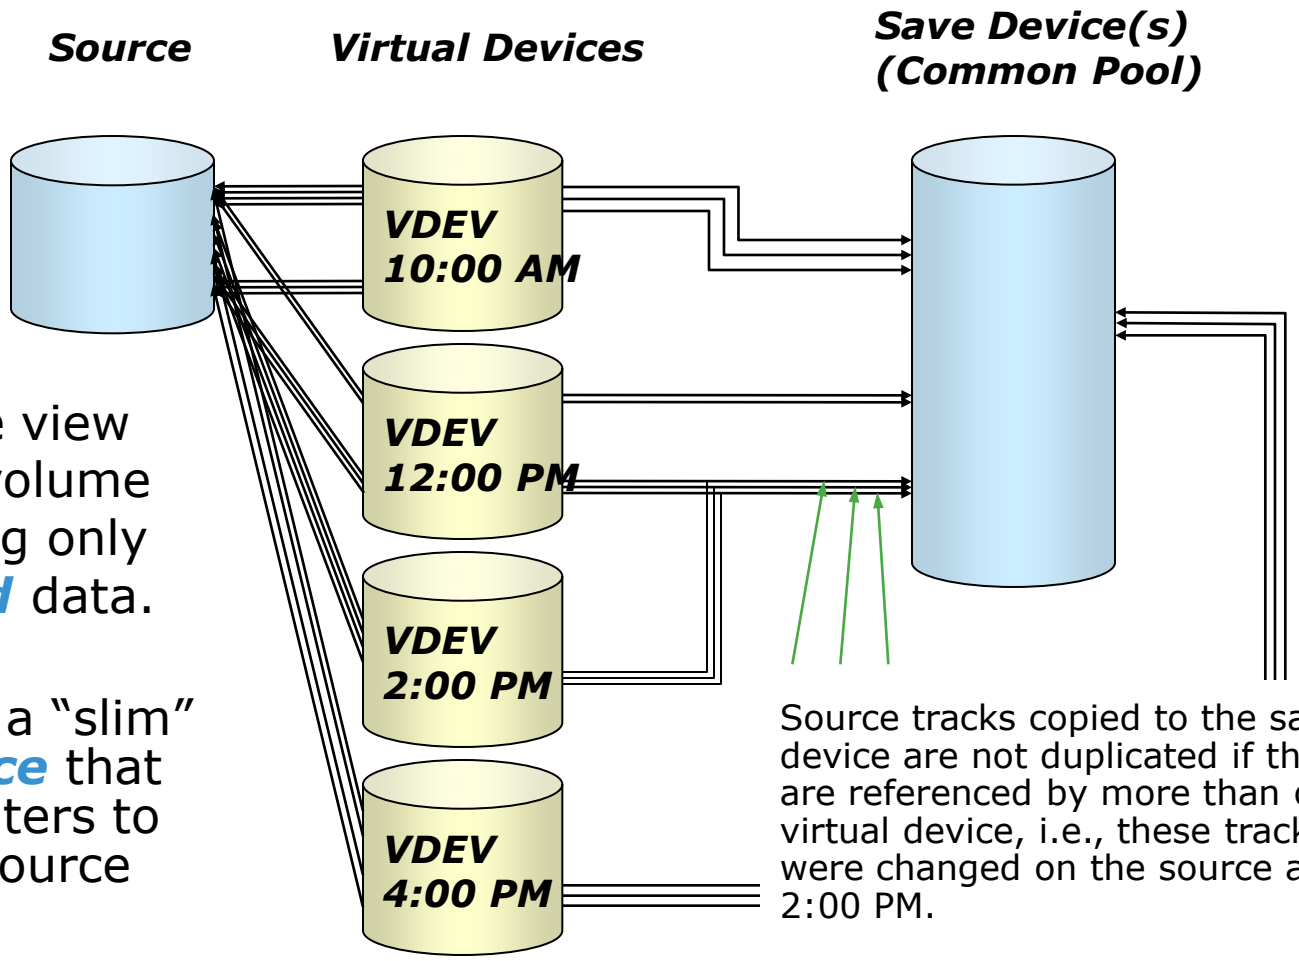

Point-in-time view of a source volume by duplicating only the ***changed*** data.

The target is a "slim" ***virtual device*** that contains pointers to the original source data.

Source tracks copied to the save device are not duplicated if they are referenced by more than one virtual device, i.e., these tracks were changed on the source after 2:00 PM.

Remote Replication: Synchronous mode

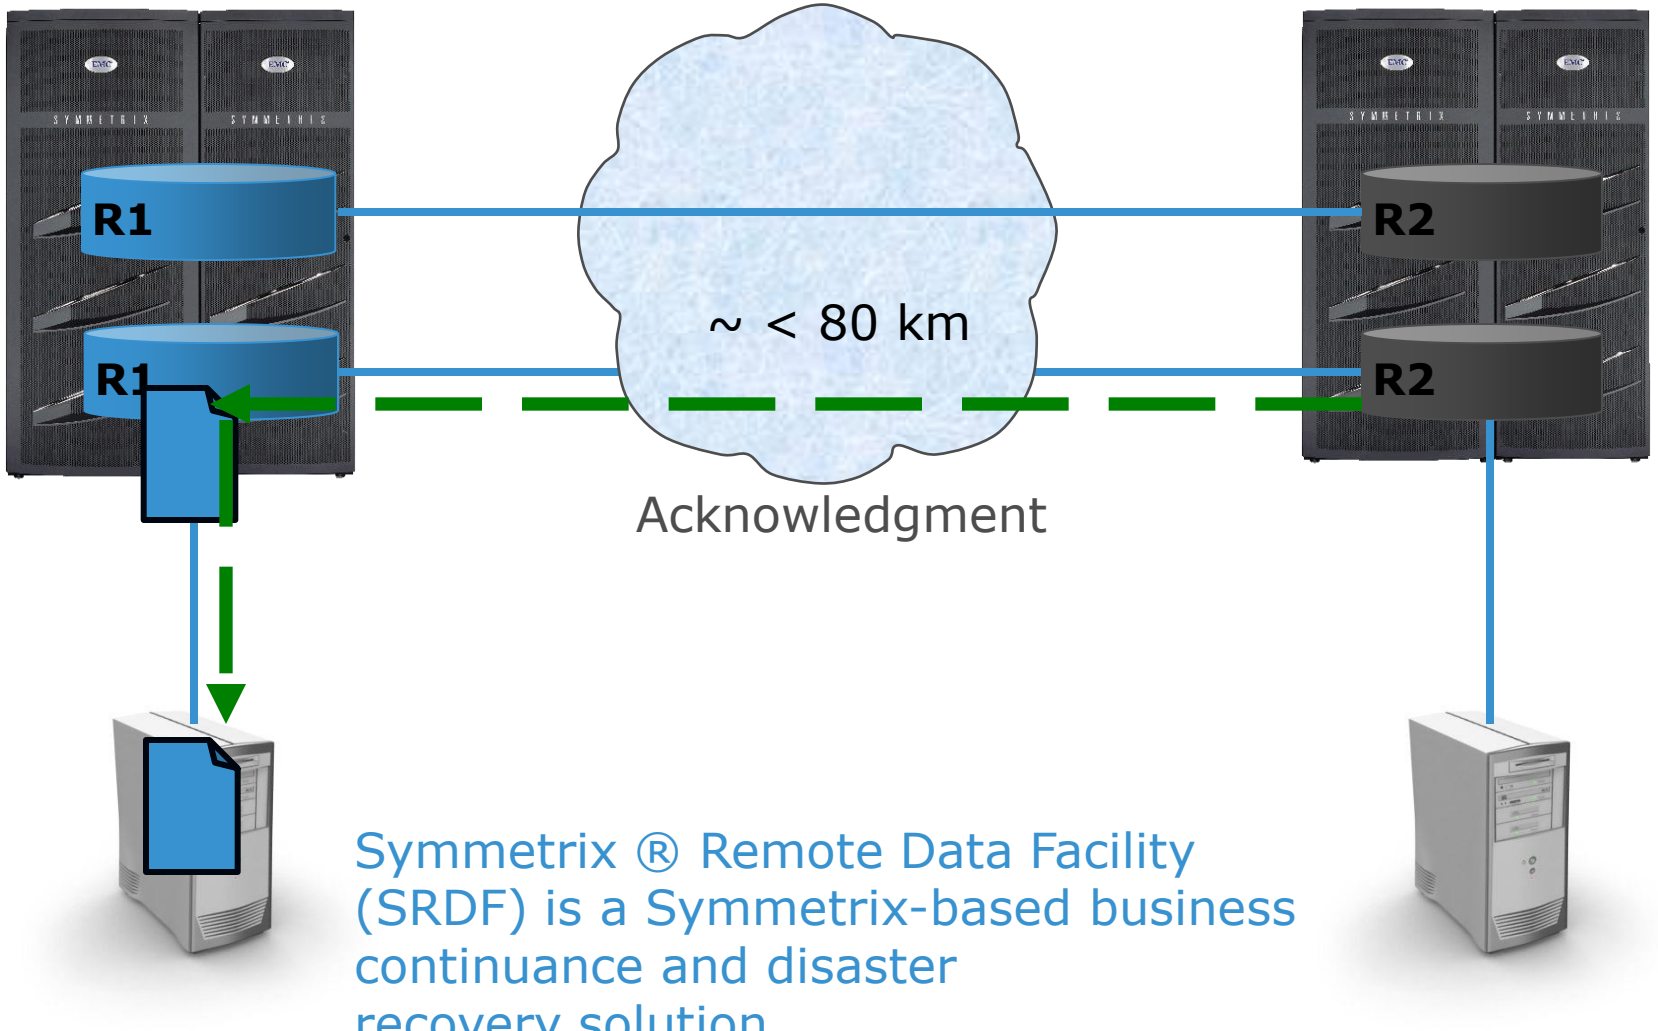

Symmetrix ® Remote Data Facility (SRDF) is a Symmetrix-based business continuance and disaster recovery solution.

Remote Replication: Asynchronous mode

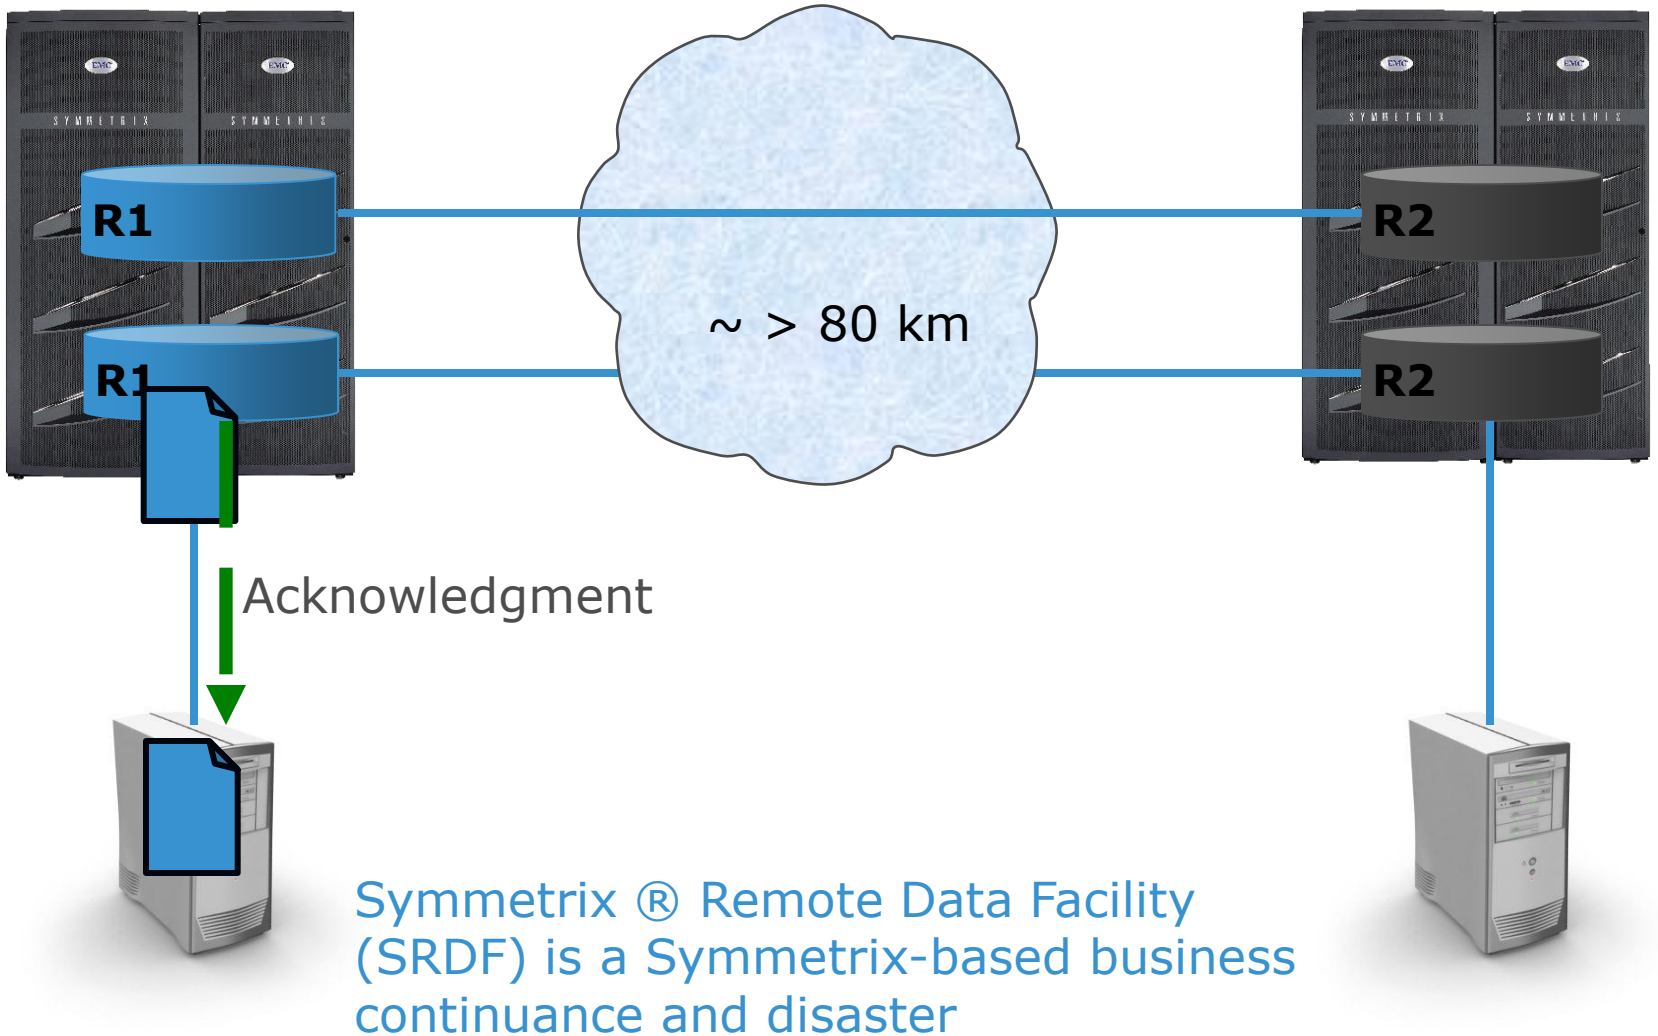

Symmetrix ® Remote Data Facility (SRDF) is a Symmetrix-based business continuance and disaster recovery solution.

Remote Replication: Cascaded

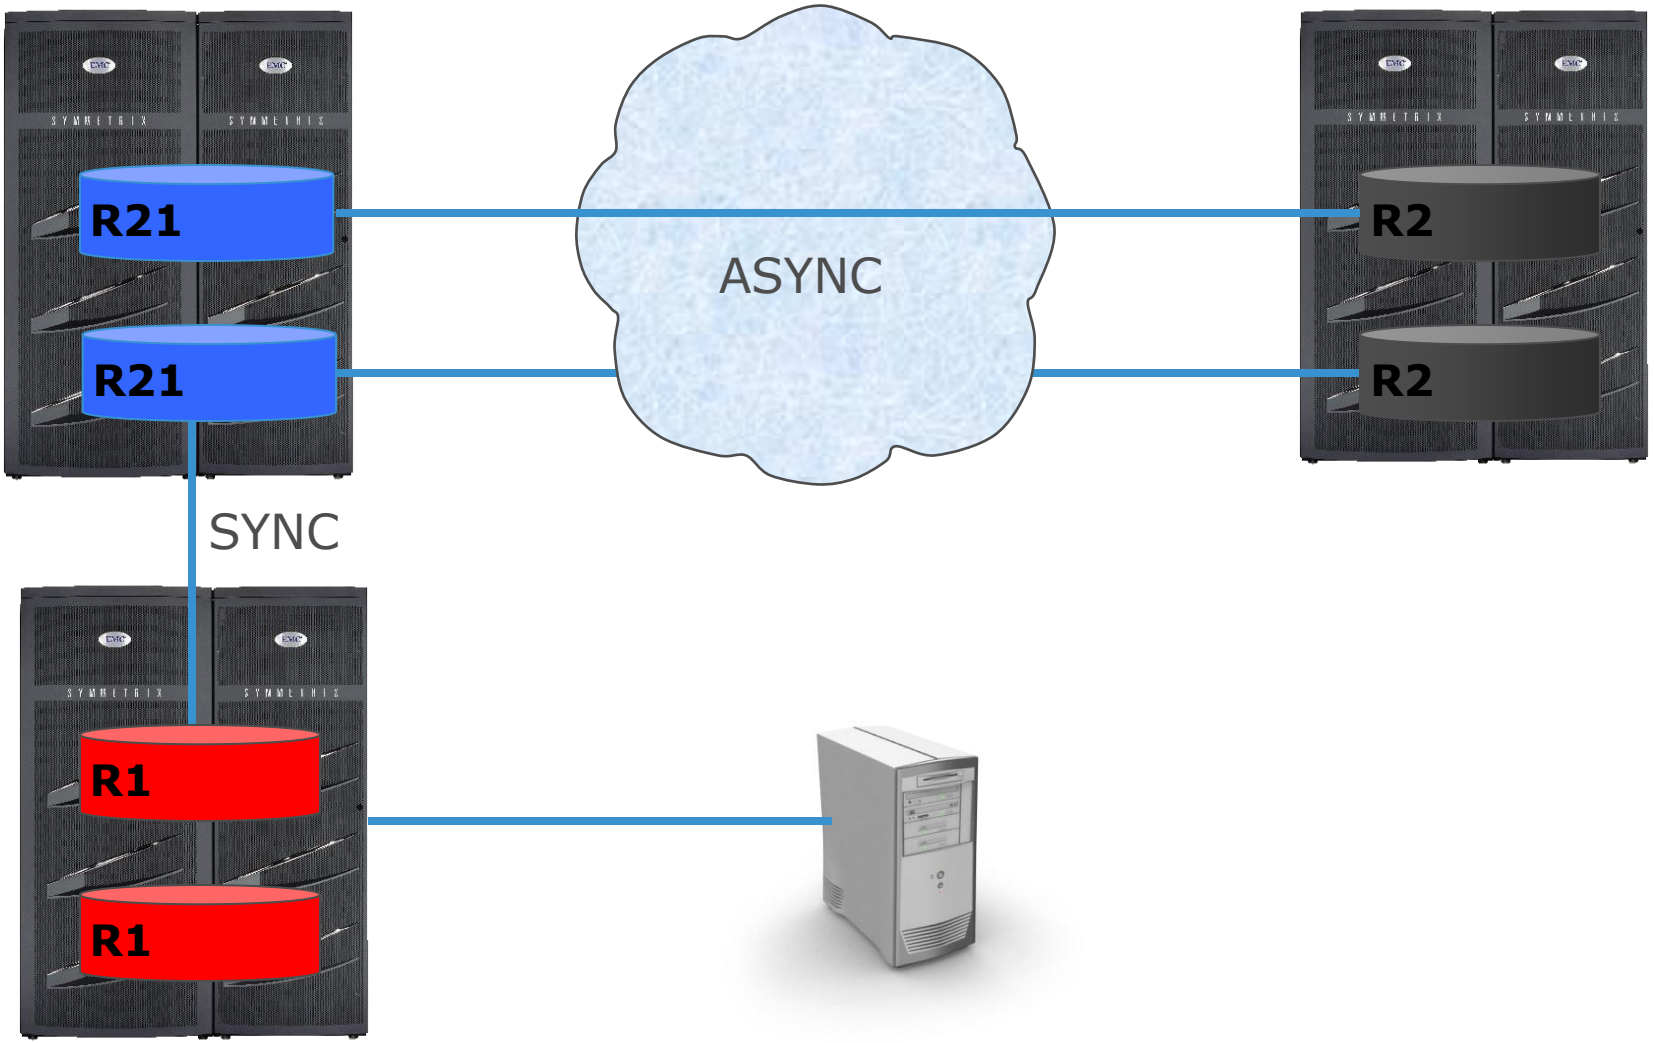

Remote Replication: SRDF/EDP. Extended Distance Protection

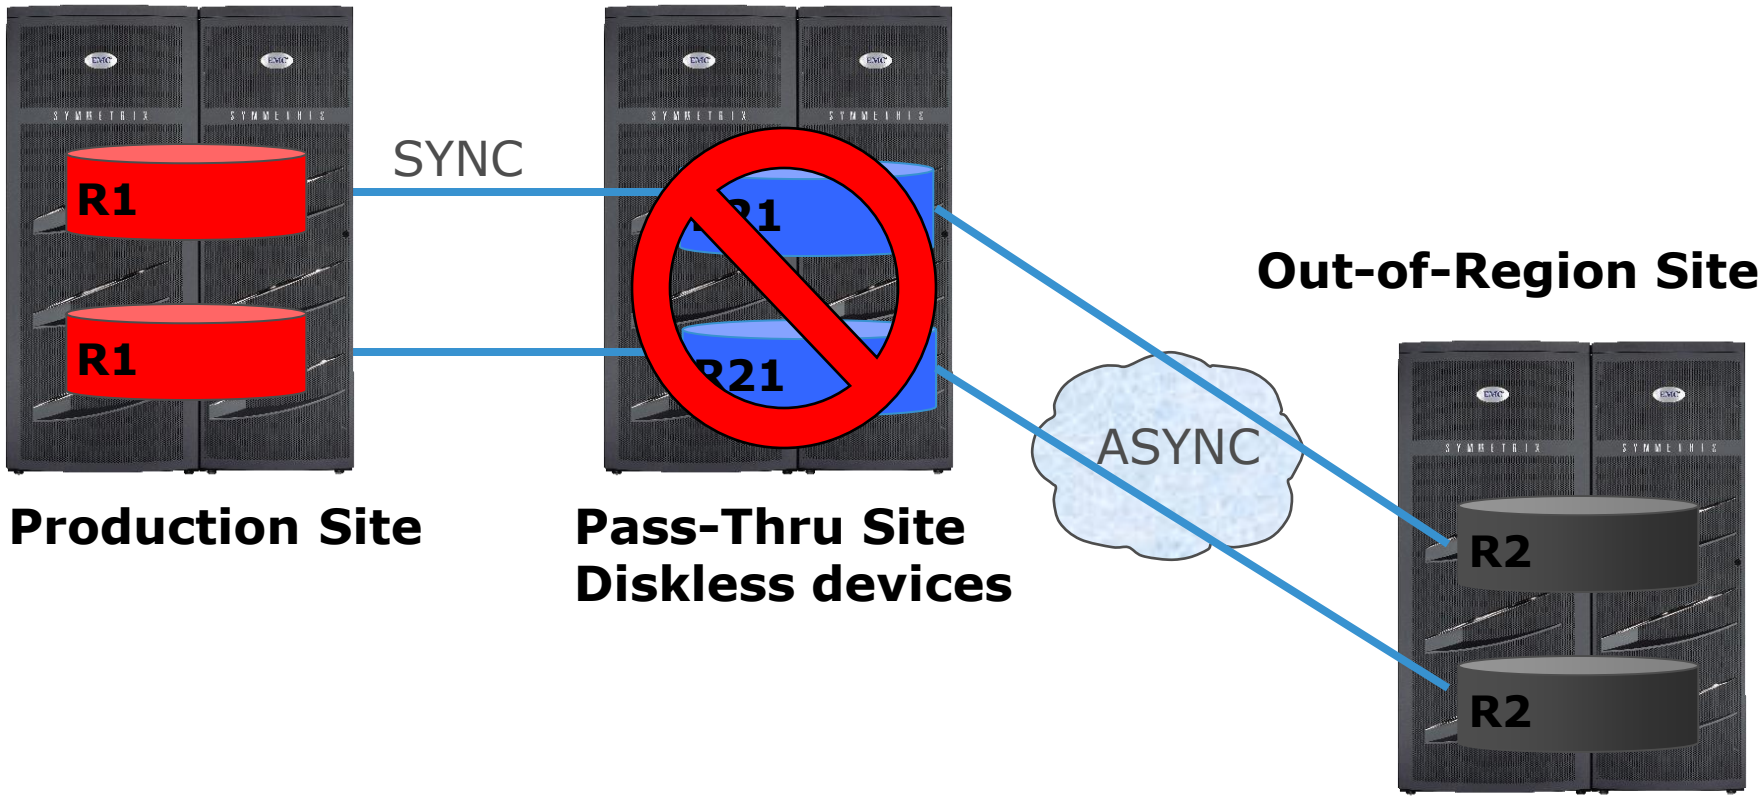

Ability to establish remote replication between production and out-of-region site and achieve no data loss at lower cost

Building Mission Critical Systems

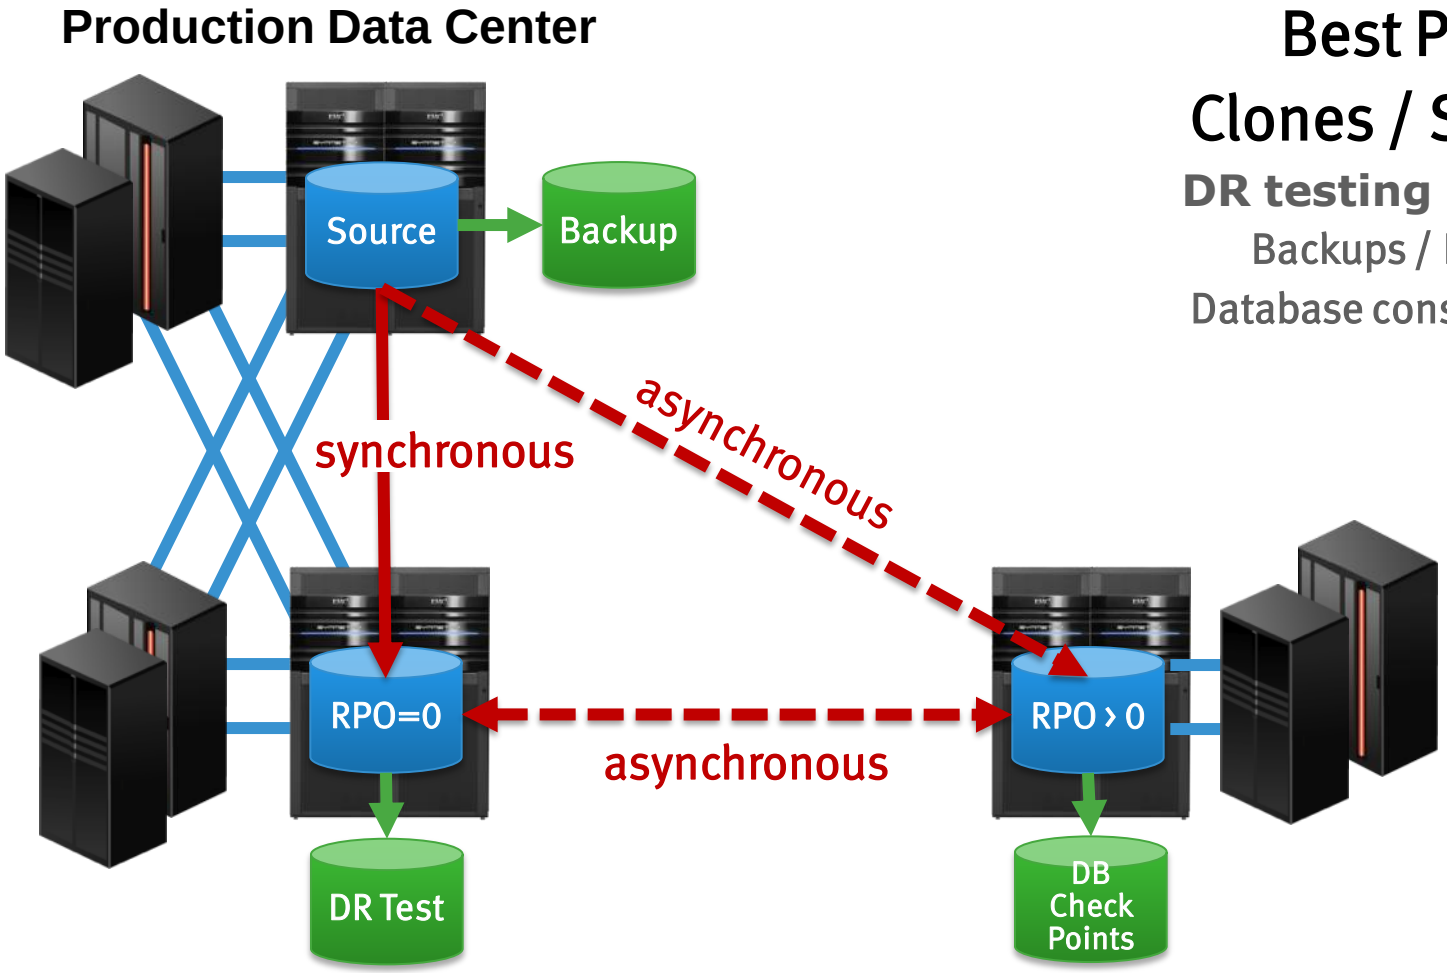

Best Practice

Clones / Snapshots

DR testing / validation

Backups / PIT recovery

Database consistency checks

Q&A

Thank you!

Sergey.Pospelov@emc.com

EMC²