Disaster Tolerant OpenVMS Clusters

Laurence Fossey
Keith Parris

Who are we?

- Laurence Fossey
- Keith Parris

Agenda

- Introduction and background
- Foundation topics
- Management, monitoring & control, DTCS
- Quorum
- Storage and Volume Shadowing
- Long Distance Issues

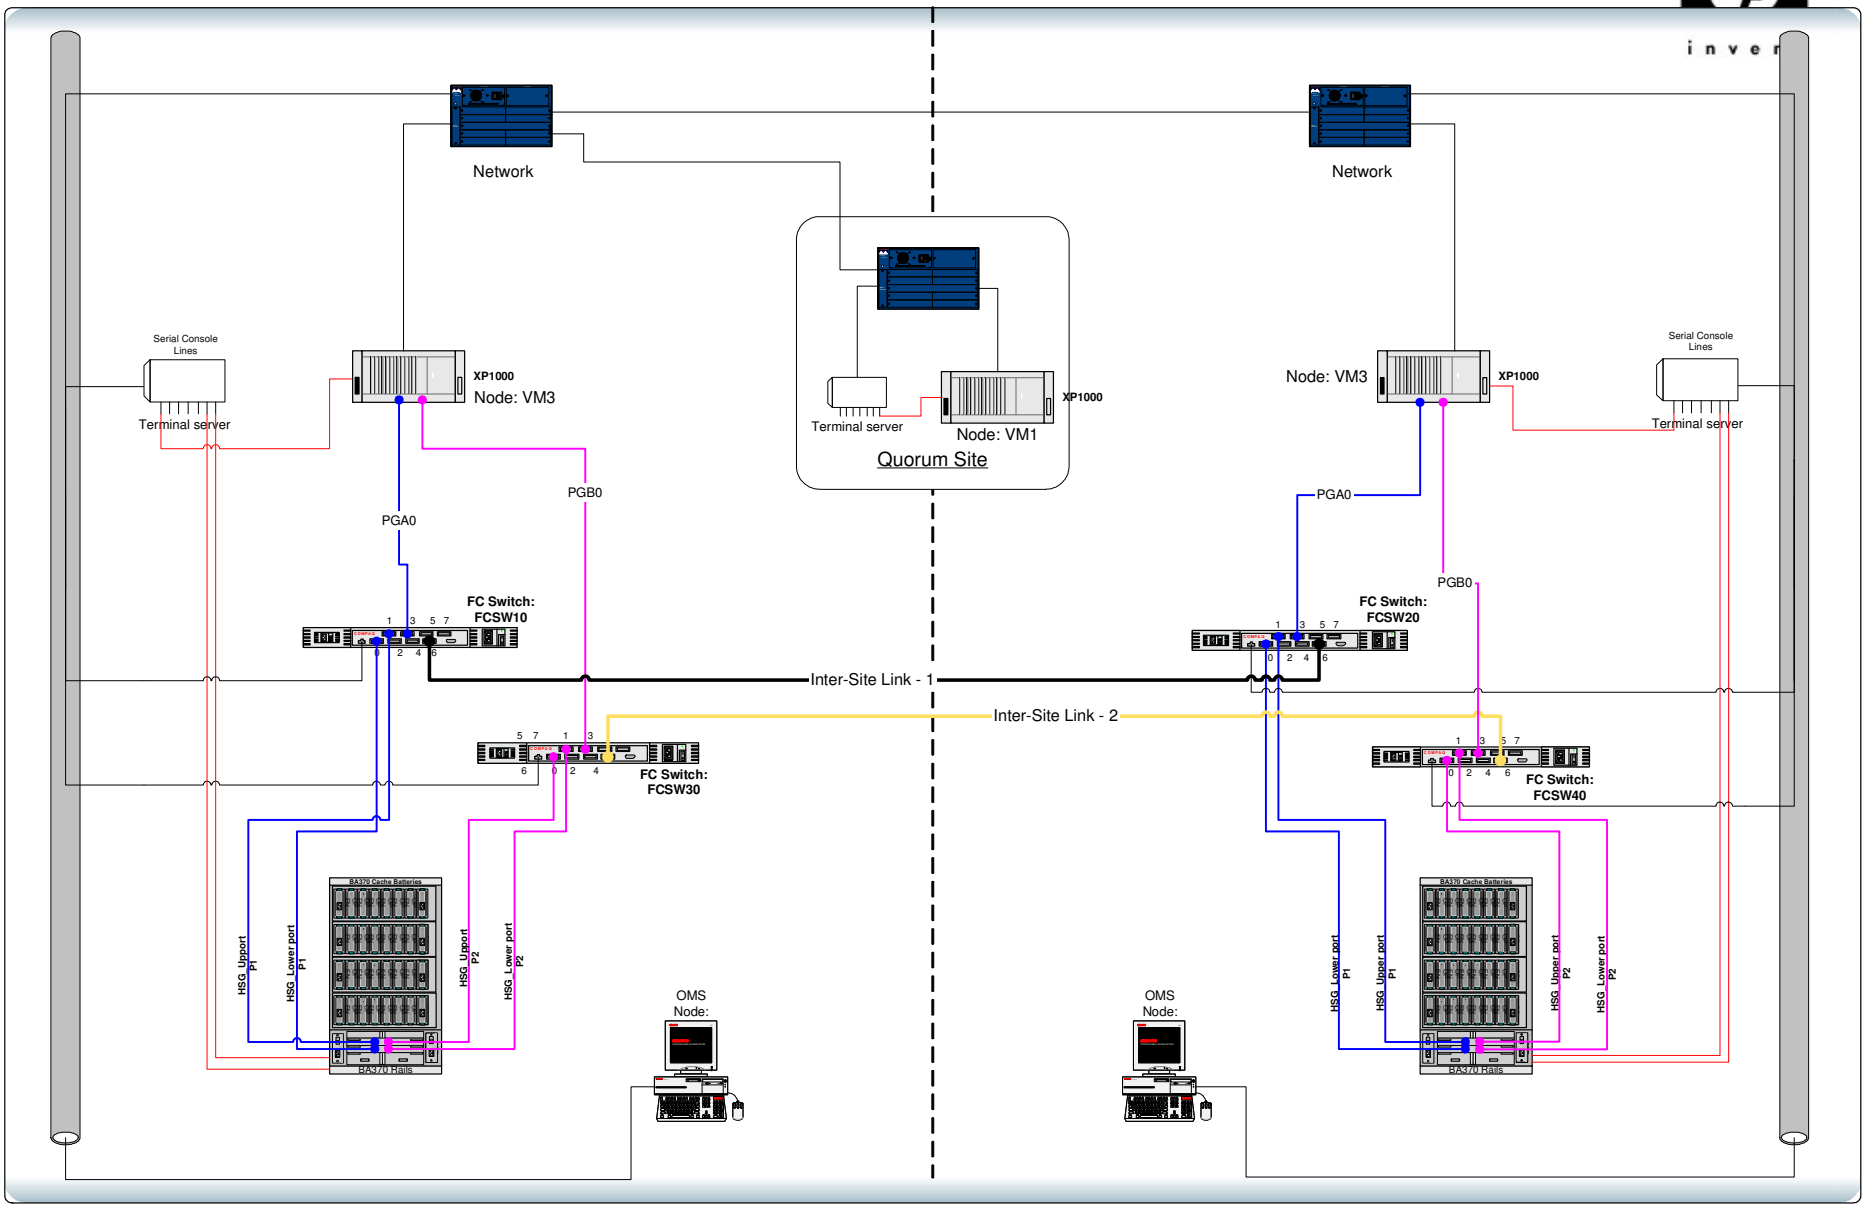
Network
Network
Serial Console Lines
Terminal server
XP1000
Node: VM3
Terminal server
Node: VM1
XP1000
Quorum Site
Node: VM3
XP1000
Serial Console Lines
Terminal server
PGA0
PGB0
PGA0
PGB0
FC Switch: FCSW10
FC Switch: FCSW20
Inter-Site Link - 1
Inter-Site Link - 2
FC Switch: FCSW30
FC Switch: FCSW40
HSG Upport P1
HSG Lower port P1
HSG Upport P2
HSG Lower port P2
BA370 Cache Batteries
BA370 Rails
OMS Node:
OMS Node:
HSG Lower port P1
HSG Upper port P1
HSG Upper port P2
HSG Lower port P2
BA370 Cache Batteries
BA370 Rails

Introduction

Positioning Disaster Recovery Solutions

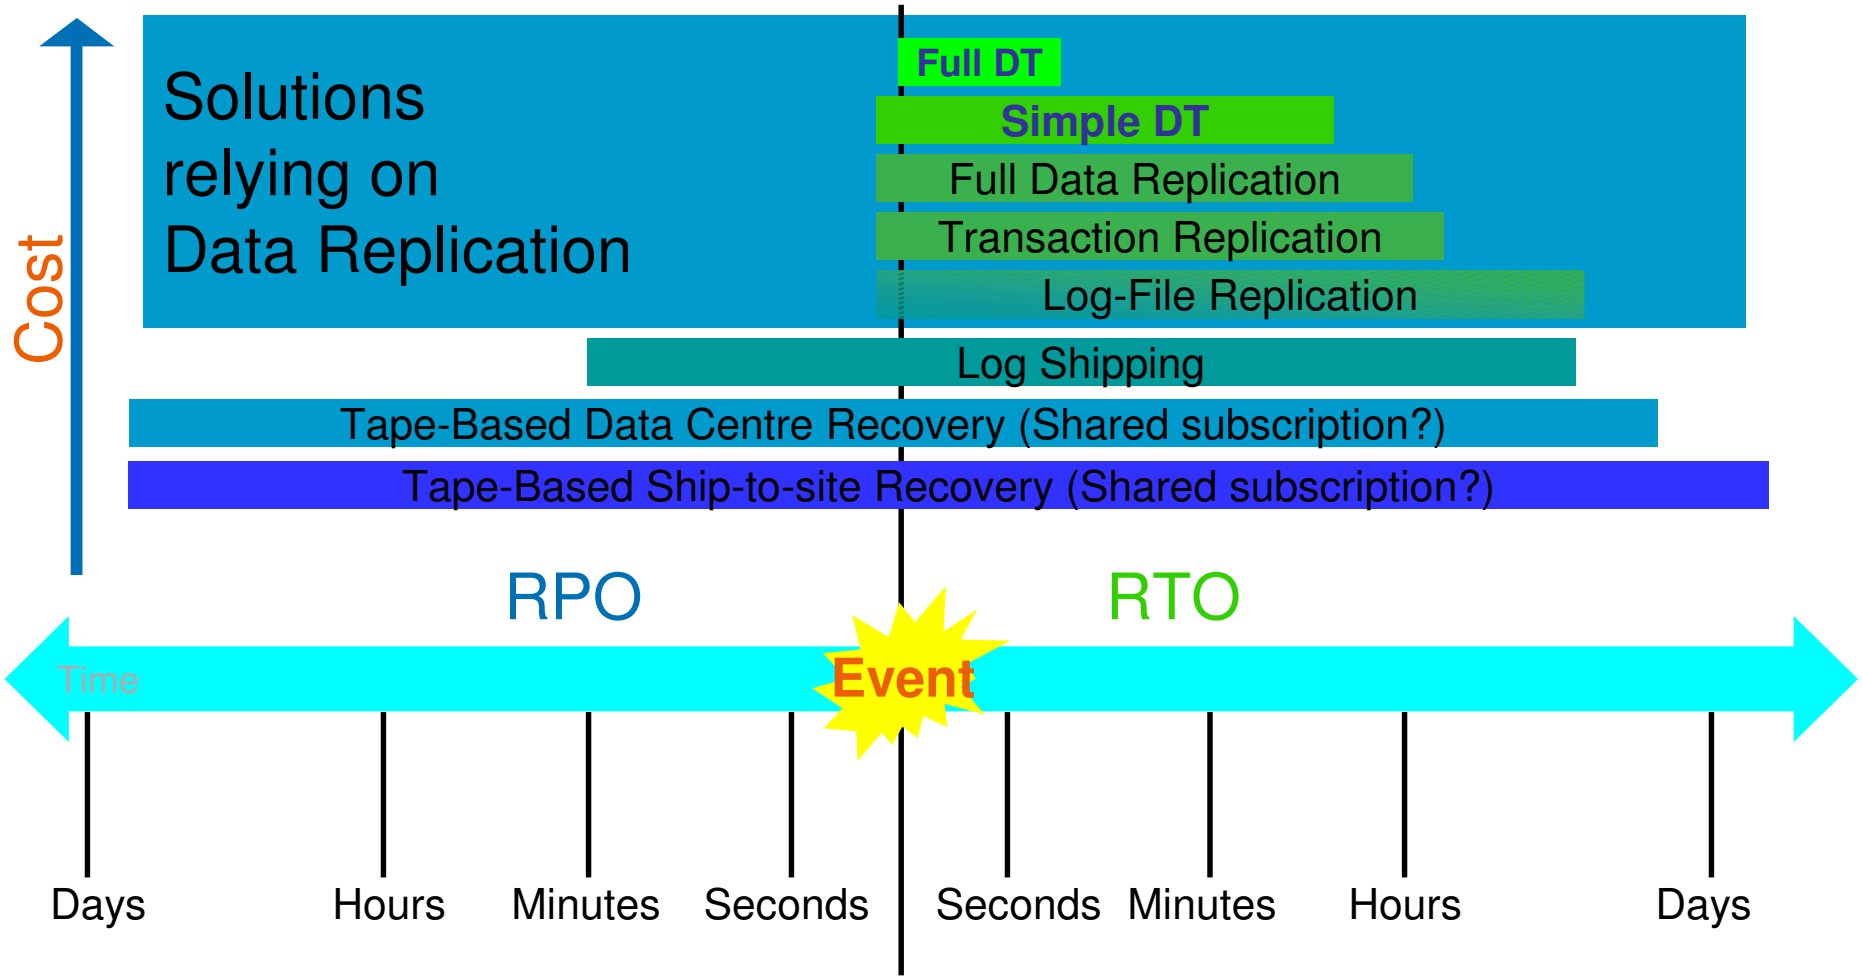

Problems or disasters?

The basic design model

Real-world problems are not that simple

Recovery completion is extended

And what about the most common problems?

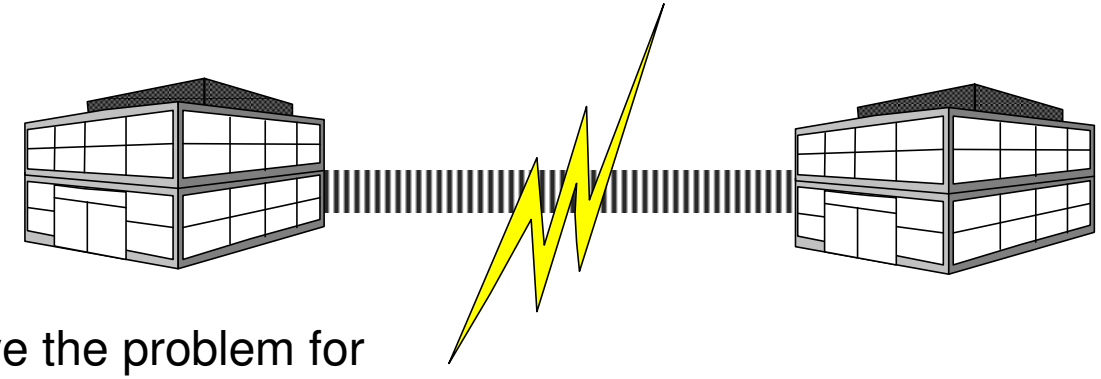

Can your technology solve the problem for you?

Foundation for OpenVMS Disaster-Tolerant Clusters

Disaster-Tolerant Clusters: Foundation

- Two or more datacenters a “safe” distance apart
- OpenVMS Cluster Software for coordination
- Inter-site link for cluster interconnect
- Data replication, to provide 2 or more identical copies of data at different sites
  - e.g. Host-Based Volume Shadowing
- Management and monitoring tools
  - Remote system console access or KVM system
  - Failure detection and alerting
  - Quorum recovery tool (especially for 2-site clusters)

Disaster-Tolerant Clusters: Foundation

- Configuration planning and implementation assistance, and staff training
  - HP recommends Disaster Tolerant Cluster Services (DTCS) package

Disaster-Tolerant Clusters: Foundation

- Carefully-planned procedures for:
  - Normal operations
  - Scheduled downtime and outages
  - Detailed diagnostic and recovery action plans for various failure scenarios

Planning for Disaster Tolerance

- Remembering that the goal is to continue operating despite loss of an entire datacenter
  - All the pieces must be in place to allow that:
    - User access to both sites
    - Network connections to both sites
    - Operations staff at both sites
  - Business can't depend on *anything* that is *only* at either site

Multi-Site Clusters

- Consist of multiple sites in different locations, with one or more OpenVMS systems at each site
- Systems at each site are all part of the same OpenVMS cluster, and share resources
- Sites must connected by bridges (or bridge-routers); pure routers don't pass the SCS protocol used within OpenVMS Clusters

Bandwidth of Inter-Site Link(s)

Link Type	Bandwidth
DS-3 (a.k.a. T3)	45 mb
ATM	155 mb (OC-3) or 622 mb (OC-12)
Ethernet	Regular: 10 mb Fast: 100 mb Gigabit: 1 gb
Fibre Channel	1 or 2 mb
Memory Channel	100 MB
[D]WDM	Multiples of ATM, GbE, FC, etc.

System Management Issues

- System disks
- System parameters
- Cluster-common disk
- Shadowed data disks

Management, Monitoring, Control and DTCS

The full complexity of DT

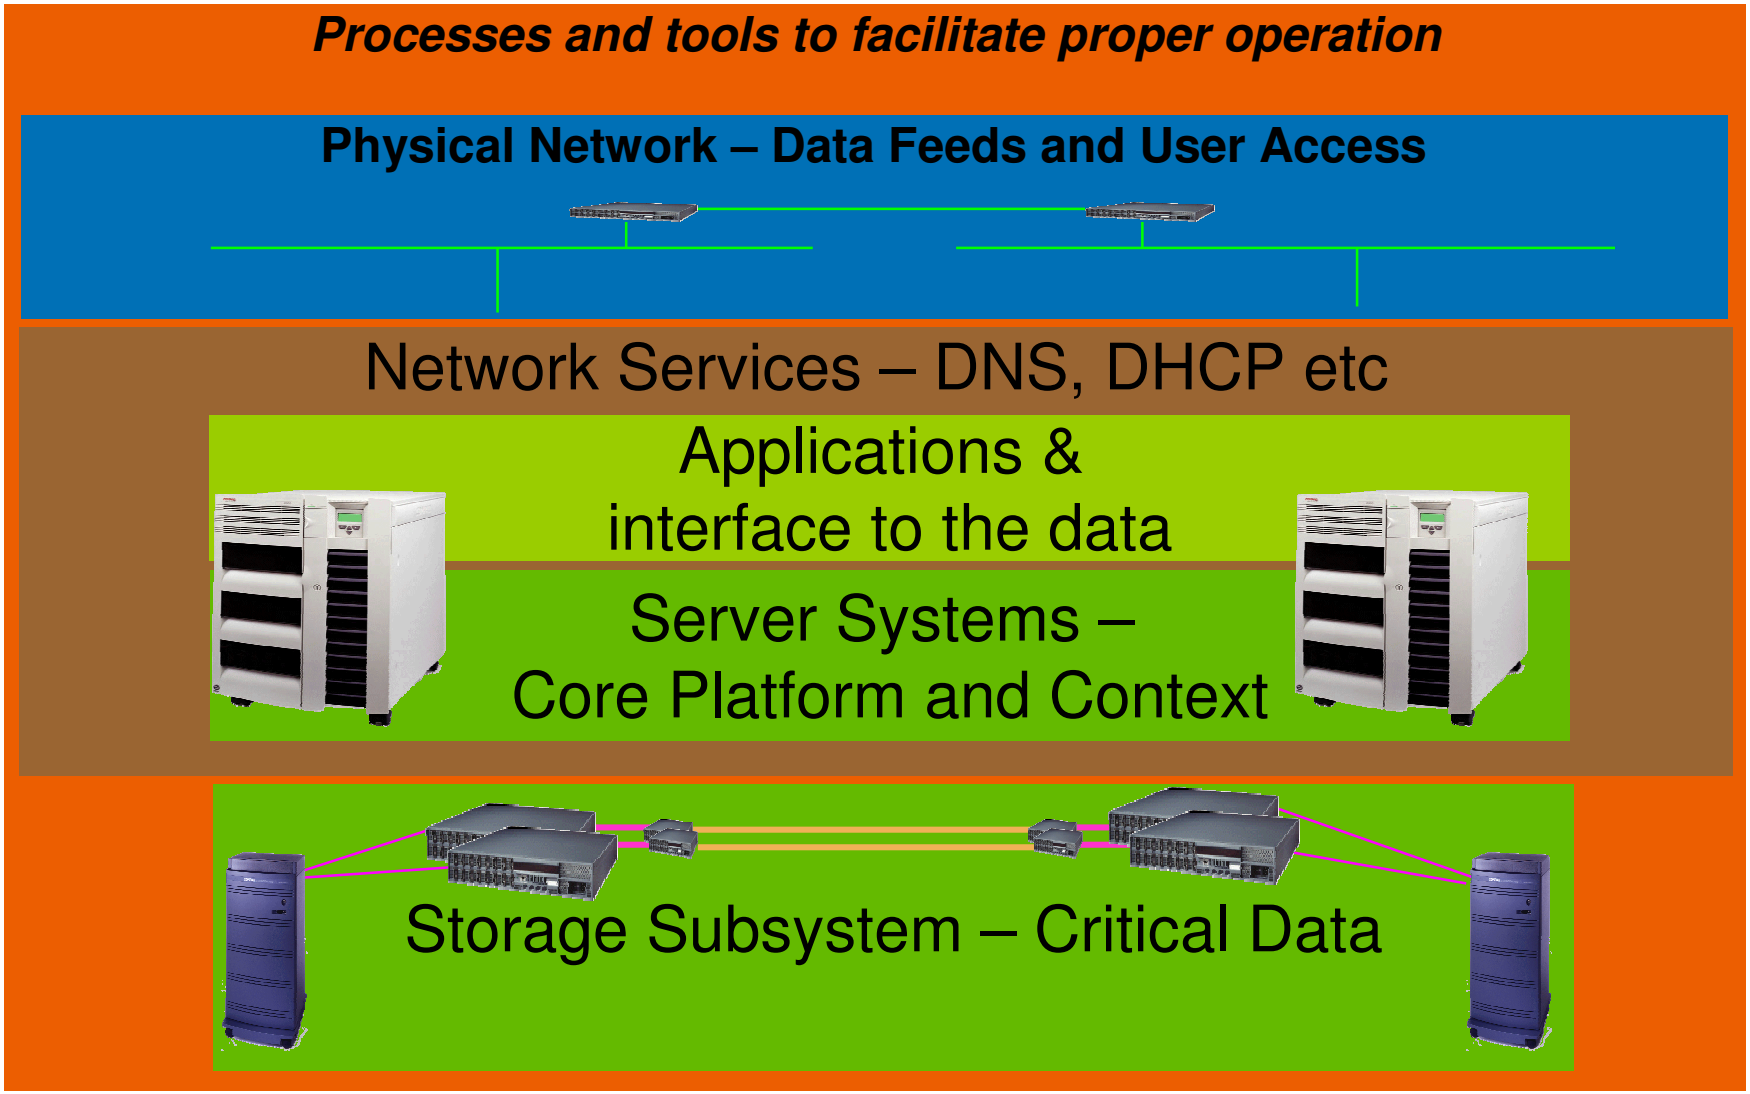

How Can You Improve The Recovery Time?

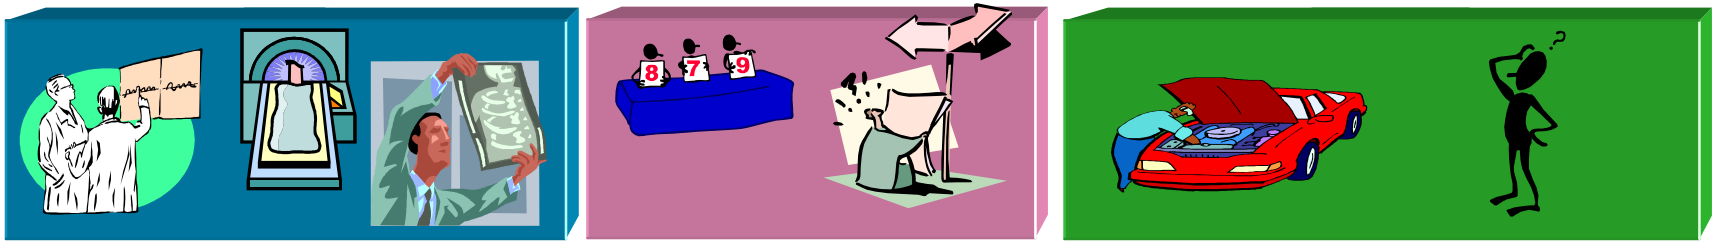

Diagnose Problem? → Decide what to do? → Implement Recovery?

Time

Recovery in disaster tolerant systems

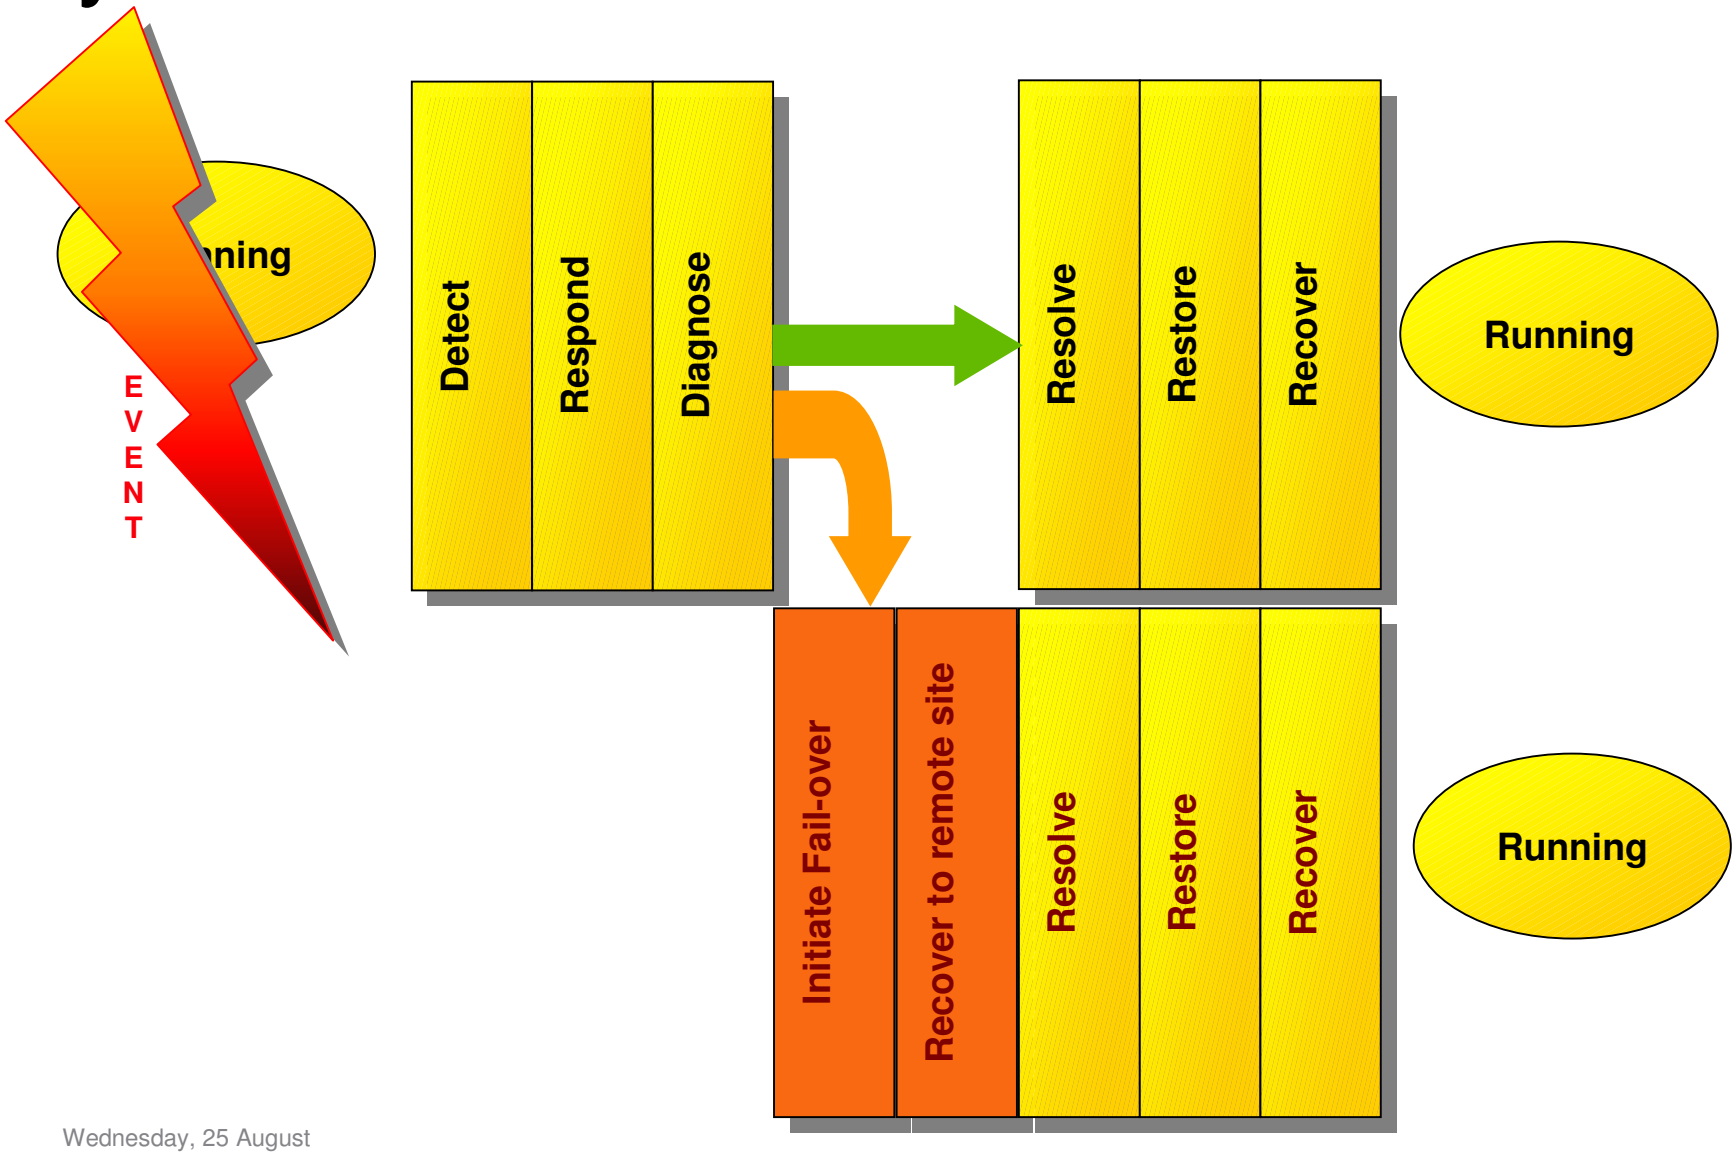

Effective Management

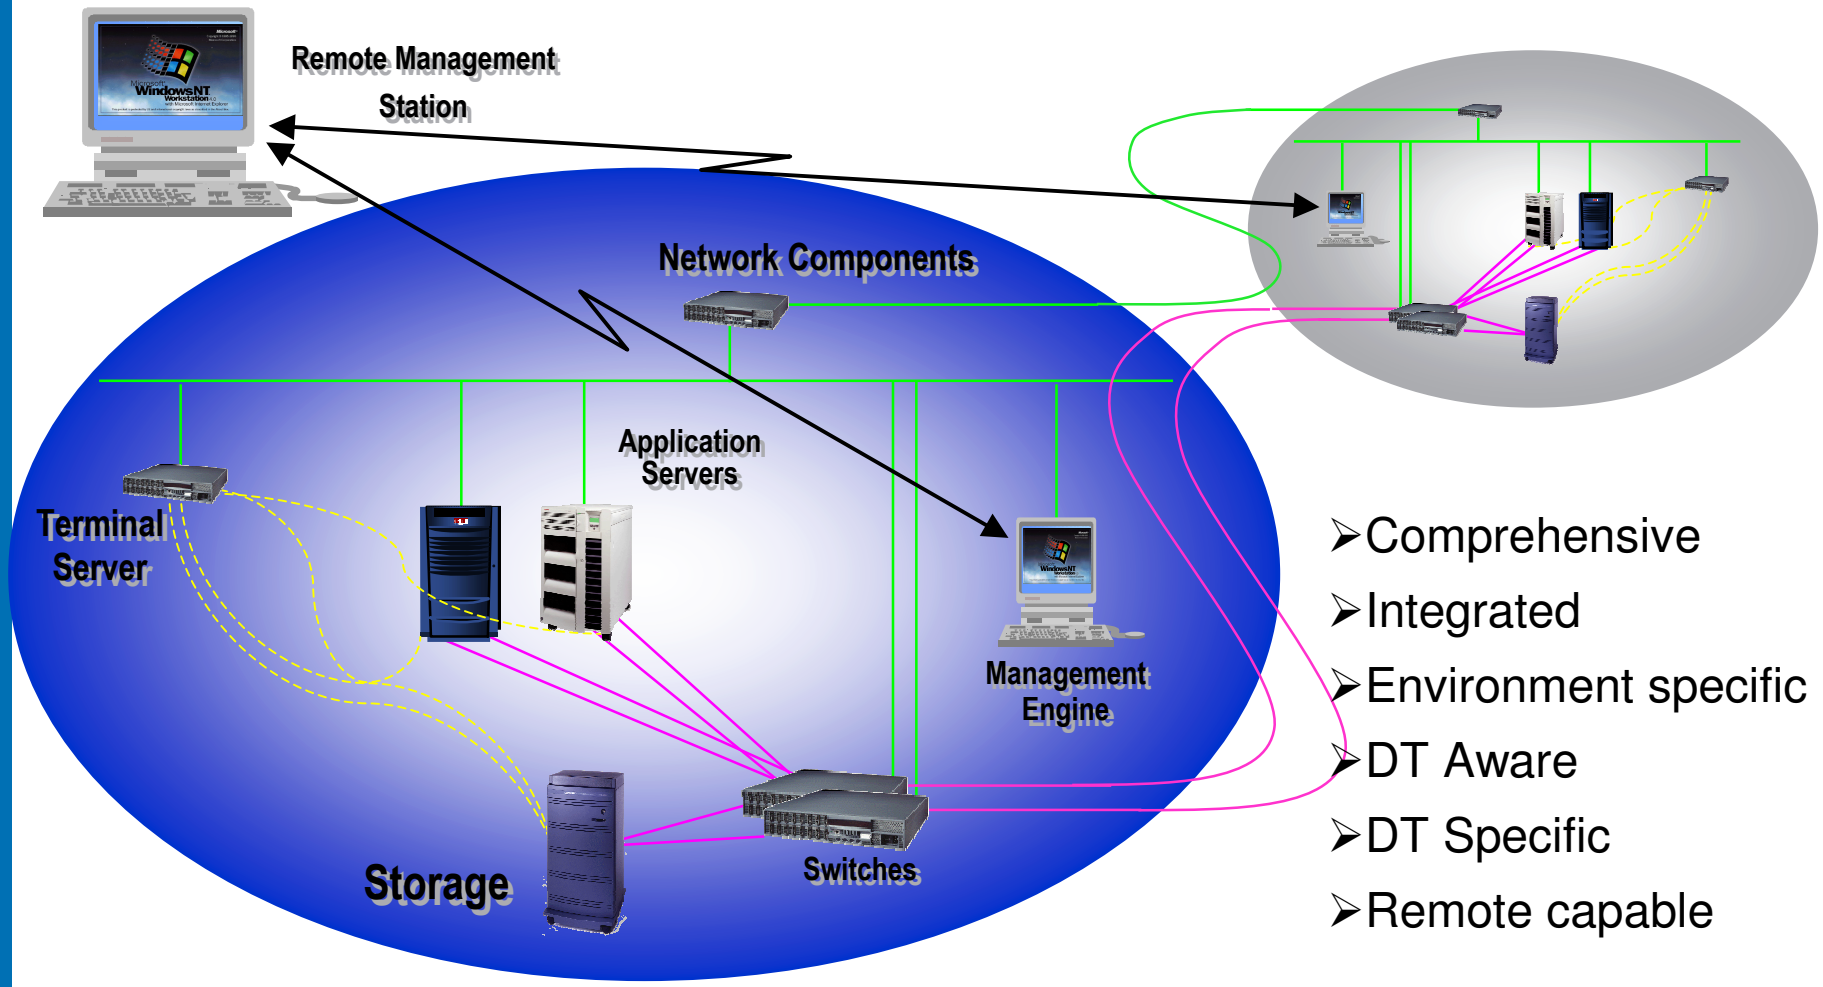

- Comprehensive
- Integrated
- Environment specific
- DT Aware
- DT Specific
- Remote capable

Effective Management

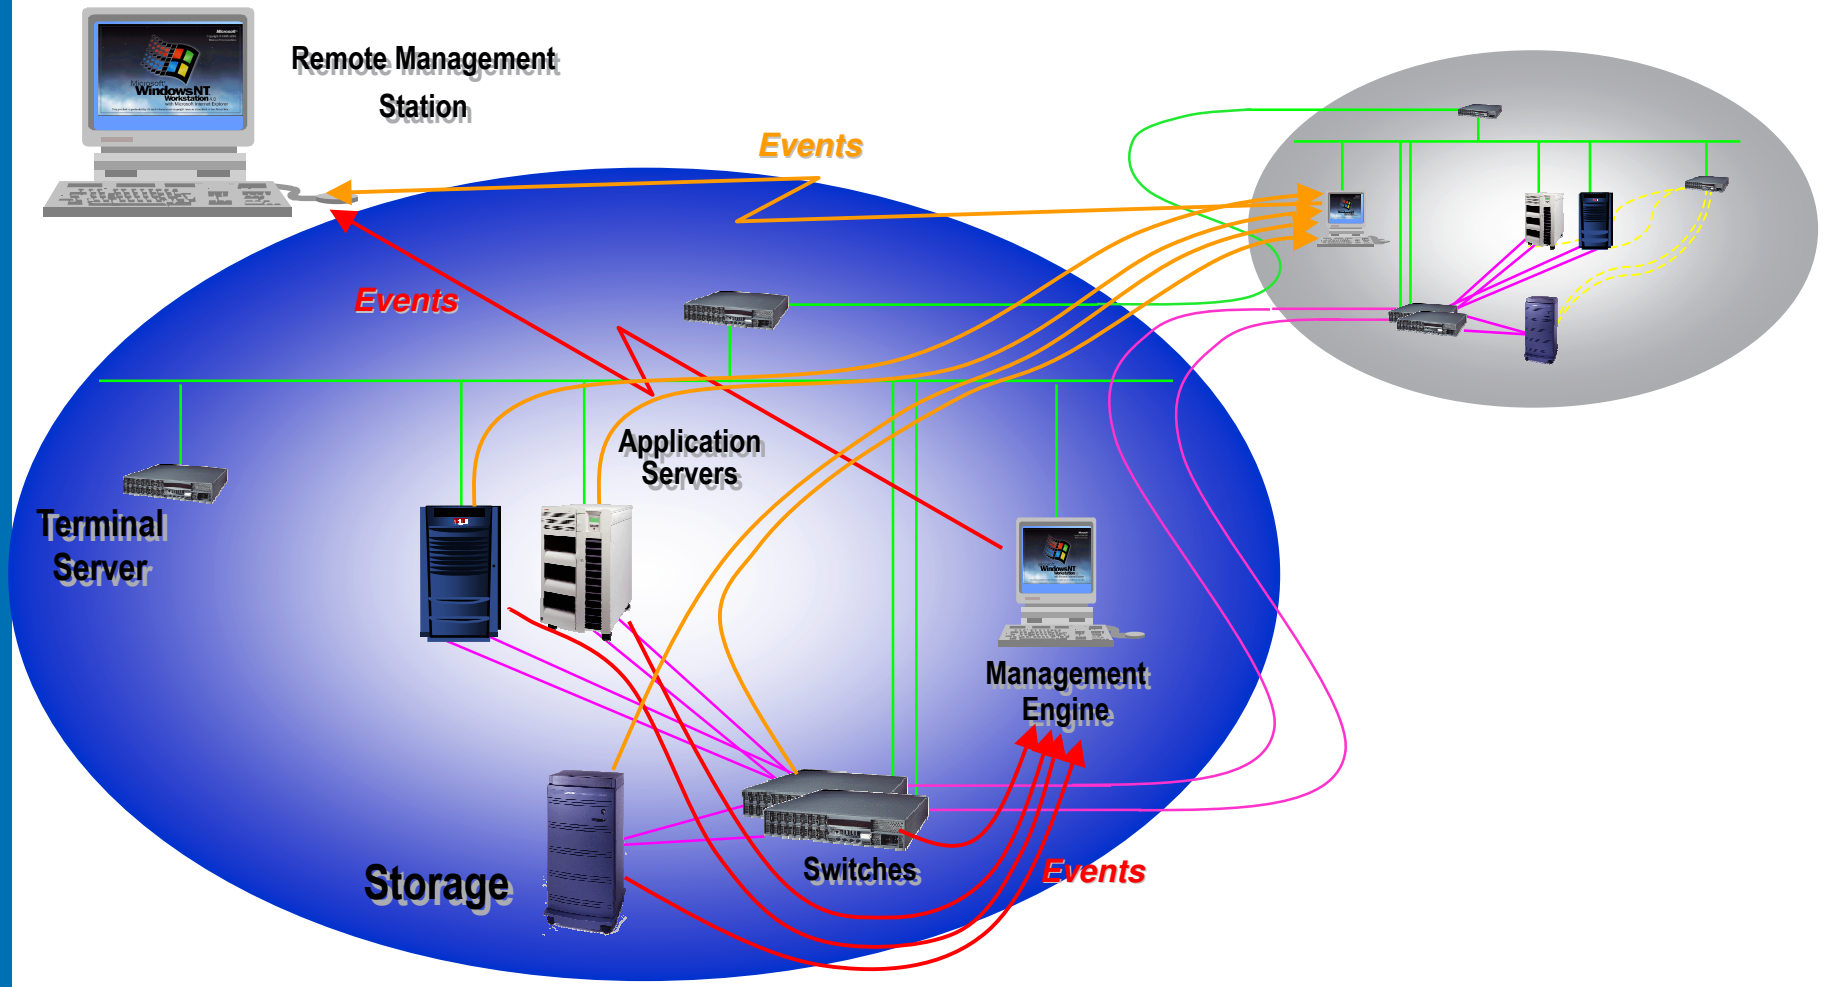

Robocentral™ Connectivity

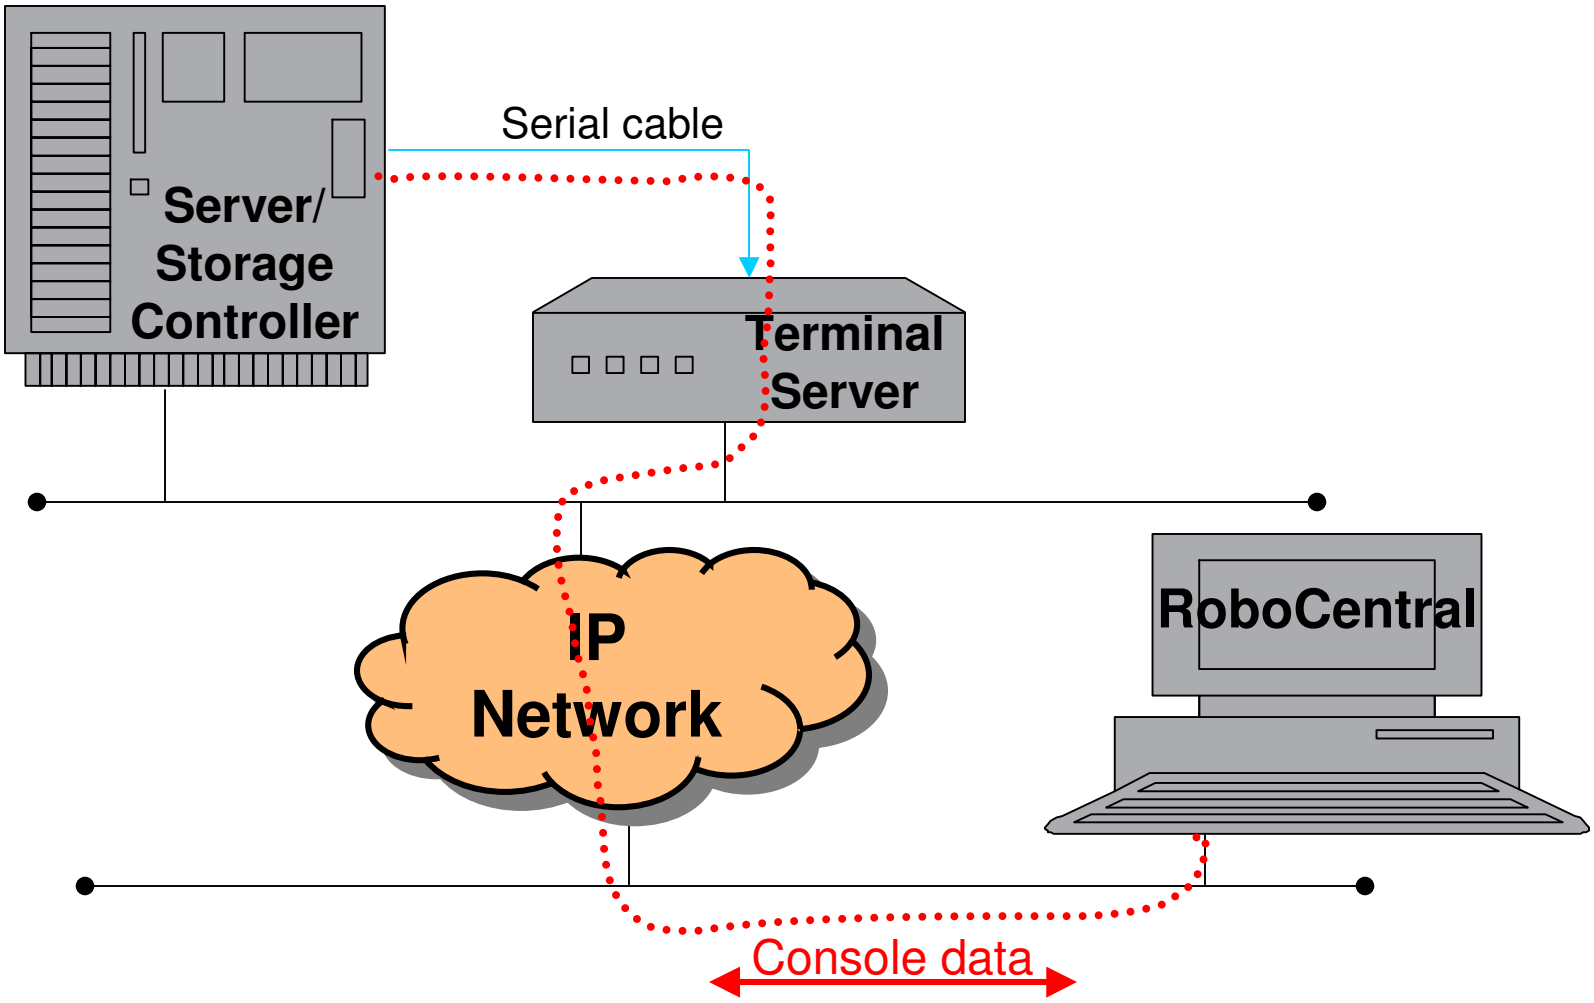

OMS Optional Software Components

- OpenVMS cluster quorum adjustment tool
  - DTCS tool allowing operator to recover an OpenVMS cluster once it has lost Quorum
- DTCS recovery manager
  - DTCS specific GUI allowing management of various types of dual site storage failures which can occur in a Storage Area Network (SAN) environment when running DRM
    - Can also be used to good effect in non-Windows environments such as OpenVMS to support “site swapping”

DT SAN Recovery Man

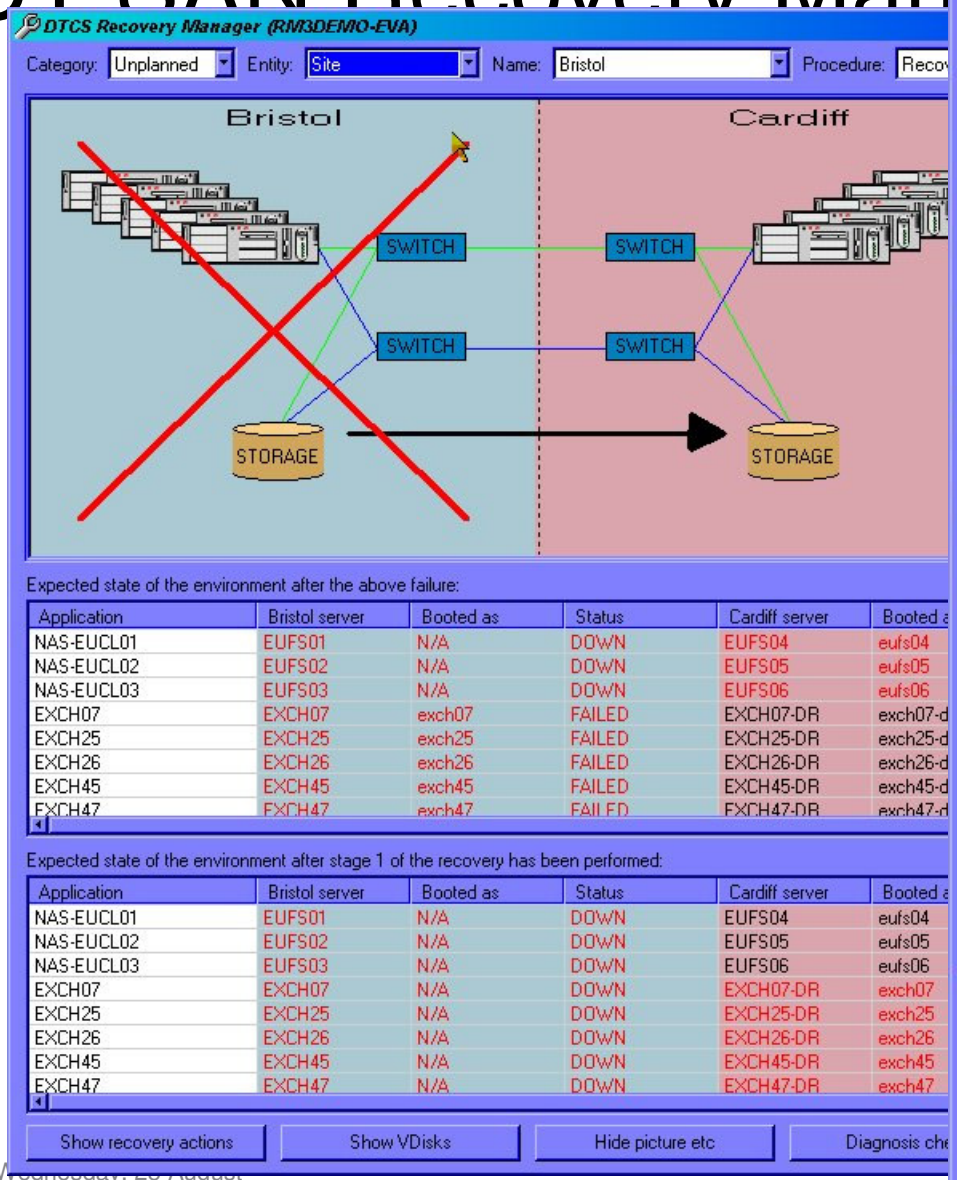

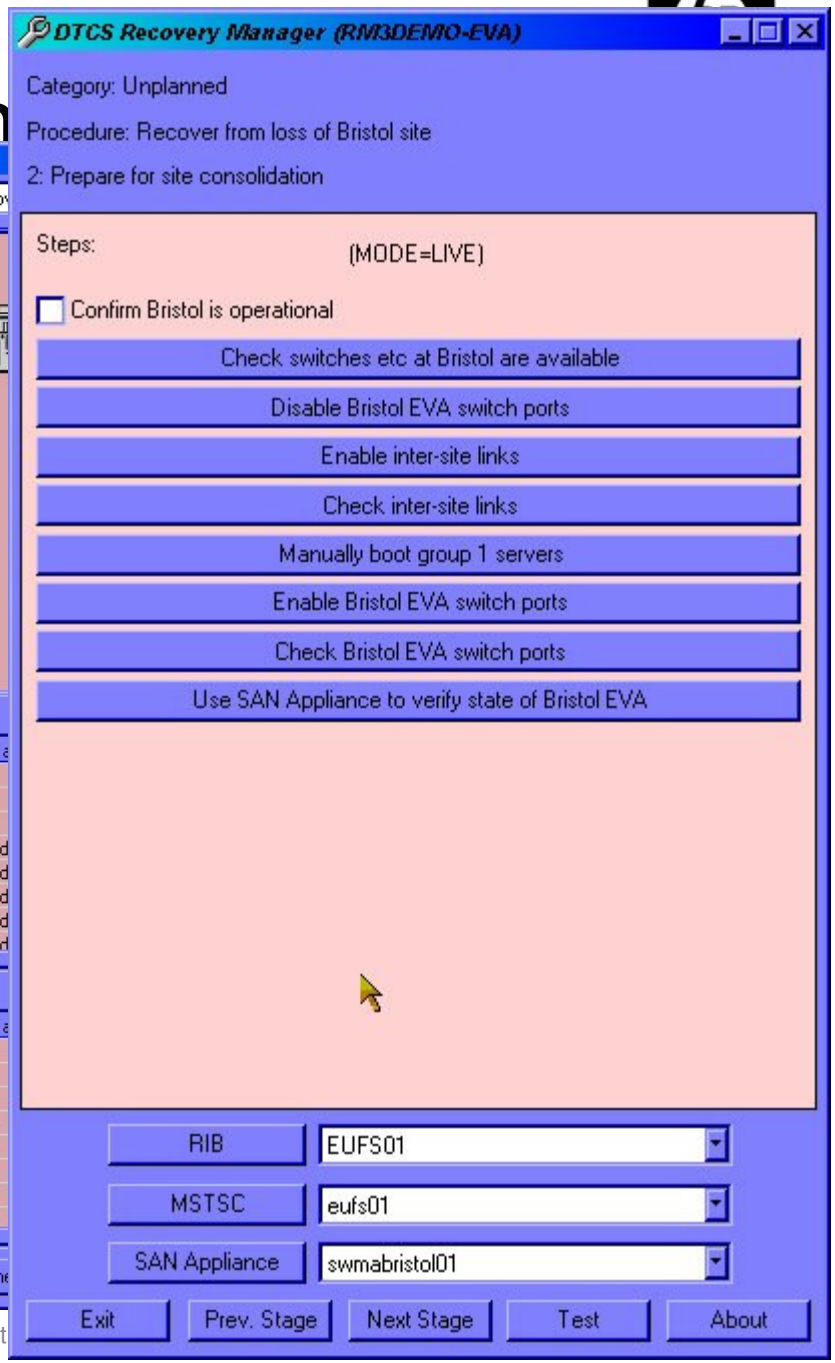

Quorum Schemes

Quorum configurations in Multi-Site Clusters

- 3 sites, equal votes in 2 sites
  - Intuitively ideal; easiest to manage & operate
  - 3rd site contains a “quorum node”, and serves as tie-breaker

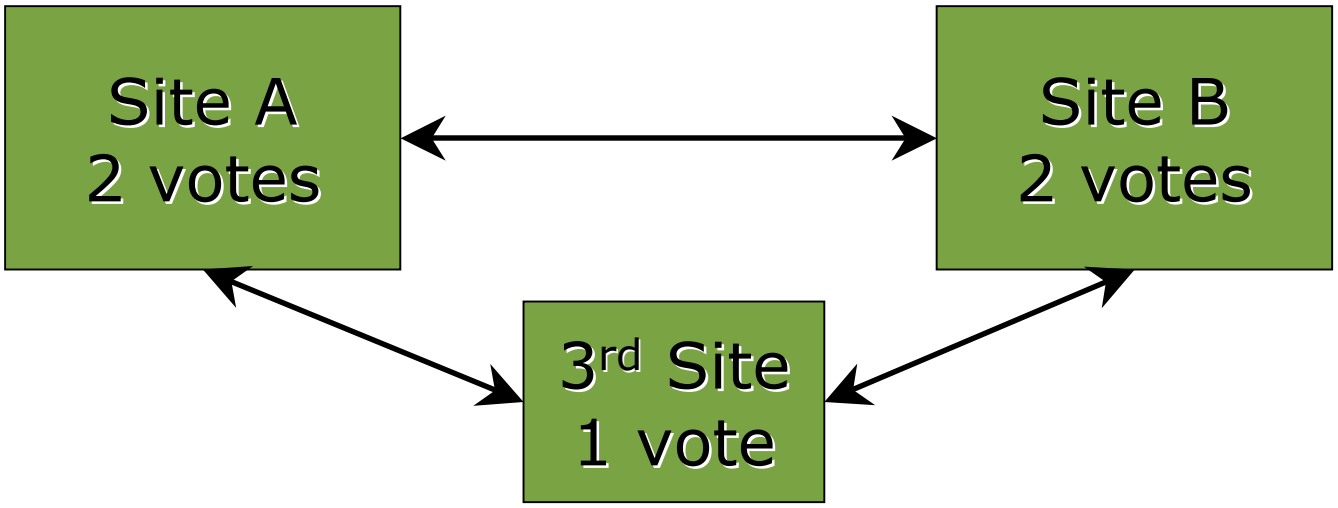

Quorum configurations in Multi-Site Clusters

- 2 sites:
  - Most common & most problematic:
    - How do you arrange votes? Balanced? Unbalanced?
    - If votes are balanced, how do you recover from loss of quorum which will result when either site or the inter-site link fails?

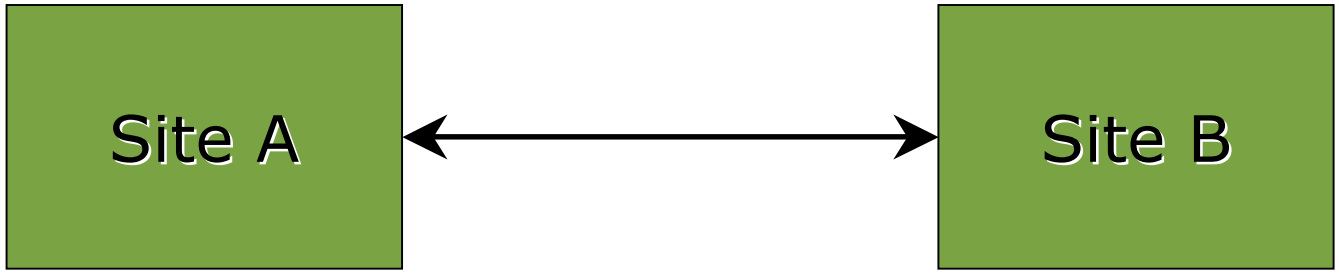

Quorum configurations in Two-Site Clusters

- One solution: Unbalanced Votes
  - More votes at one site
  - Site with more votes can continue without human intervention in the event of loss of either the other site or the inter-site link
  - Site with fewer votes pauses on a failure and requires manual action to continue after loss of the other site

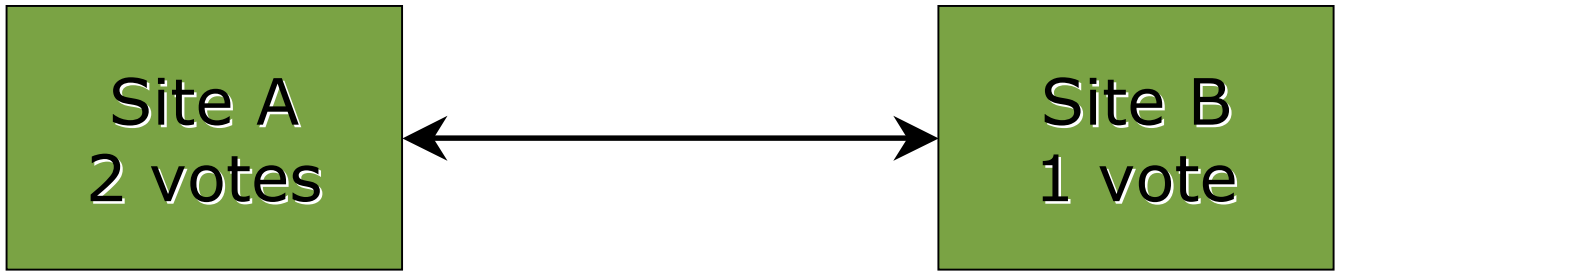

Can continue automatically

Requires manual intervention to continue alone

Quorum configurations in Two-Site Clusters

- Unbalanced Votes
  - Common mistake:
    - Give more votes to Primary site
    - Leave Standby site unmanned
  - Result: cluster can't run without Primary site, unless there is human intervention, which is unavailable at the (unmanned) Standby site

Can continue automatically

Requires manual intervention to continue alone

Quorum configurations in Two-Site Clusters

- Unbalanced Votes
  - Also very common in remote-shadowing-only clusters (for data vaulting -- not full disaster-tolerant clusters)
    - 0 votes is a common choice for the remote site in this case
      - But that has its dangers

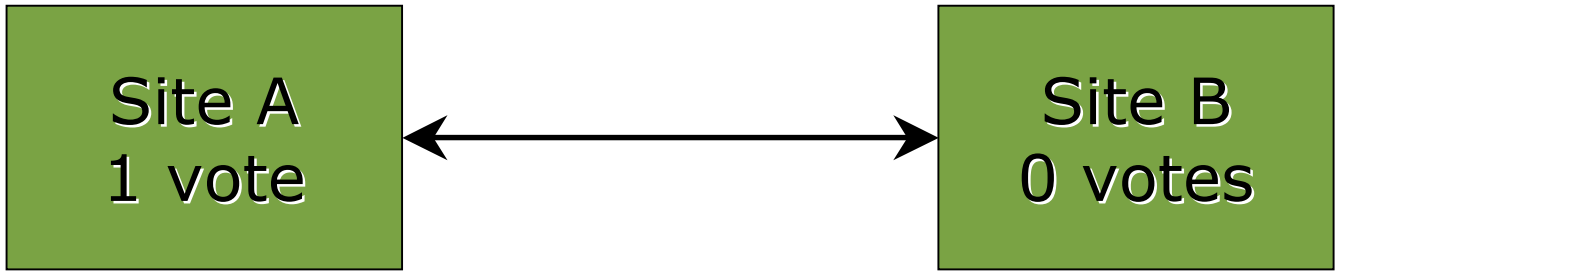

Can continue automatically

Requires manual intervention to continue alone

Optimal Sub-cluster Selection

Connection manager compares potential node subsets that could make up surviving portion of the cluster

1) Pick sub-cluster with the most votes
2) If votes are tied, pick sub-cluster with the most nodes
3) If nodes are tied, arbitrarily pick a winner
   - based on comparing SCSSYSTEMID values of set of nodes with most-recent cluster software revision

Optimal Sub-cluster Selection: Two-Site Cluster with Unbalanced Votes

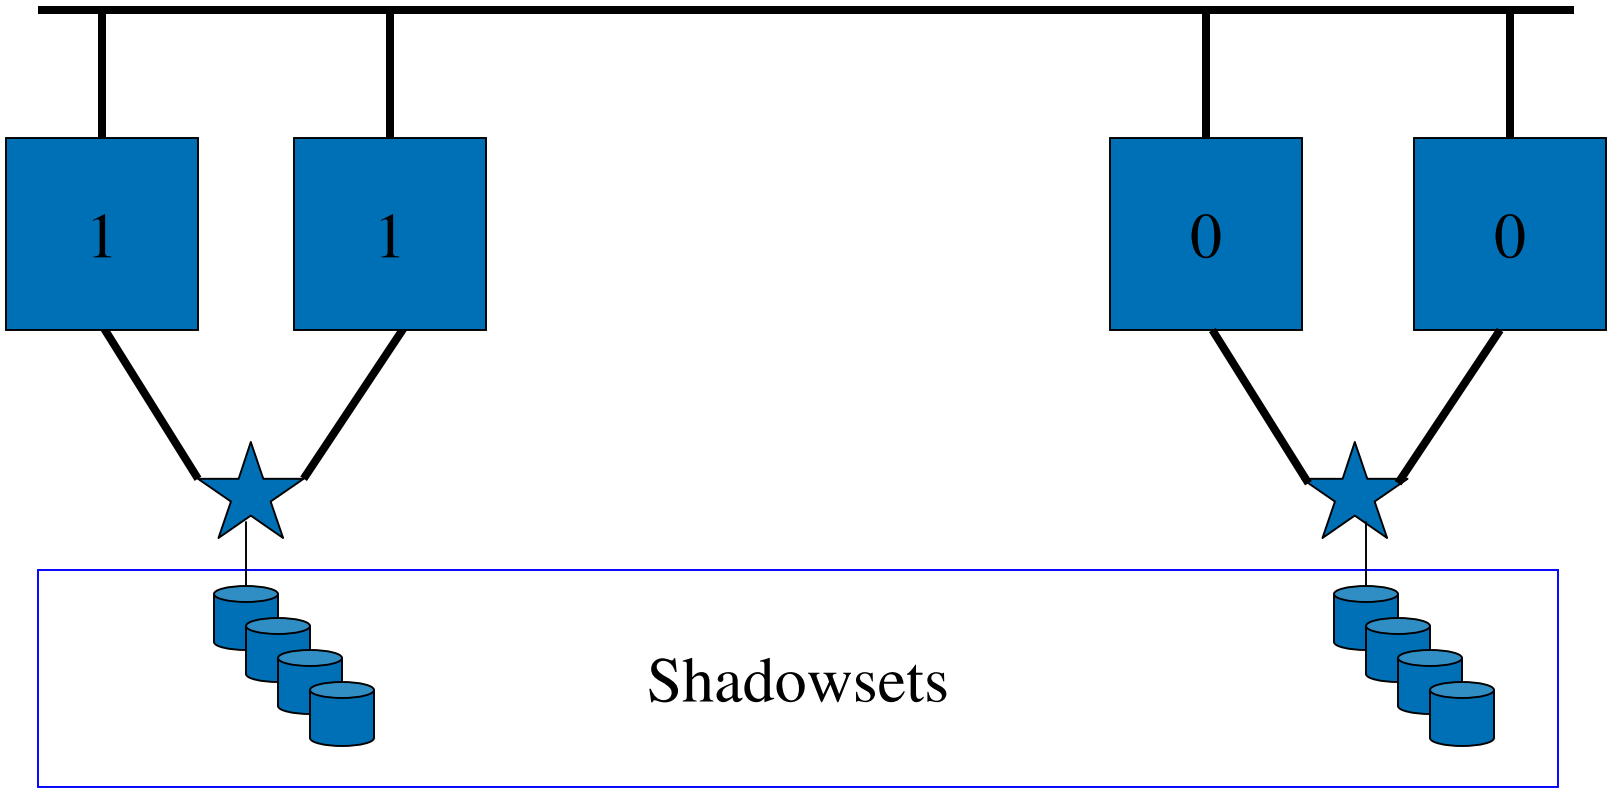

Optimal Sub-cluster Selection: Two-Site Cluster with Unbalanced Votes

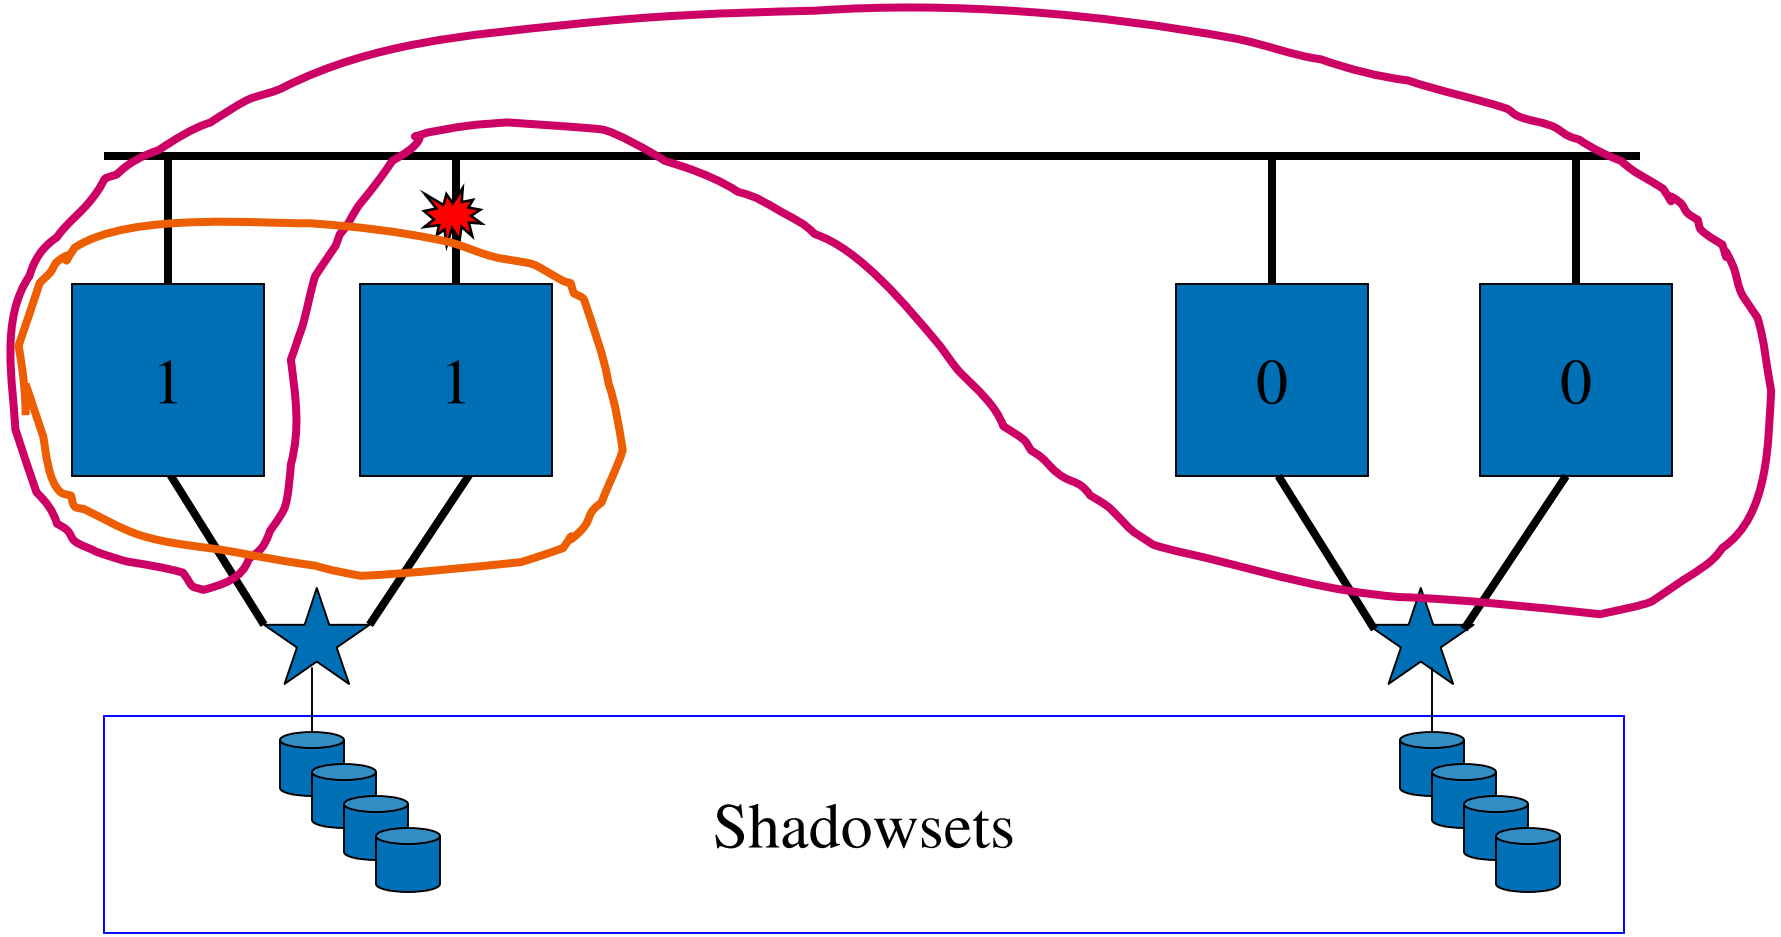

Which subset of nodes is selected as the optimal sub-cluster?

Optimal Sub-cluster Selection: Two-Site Cluster with Unbalanced Votes

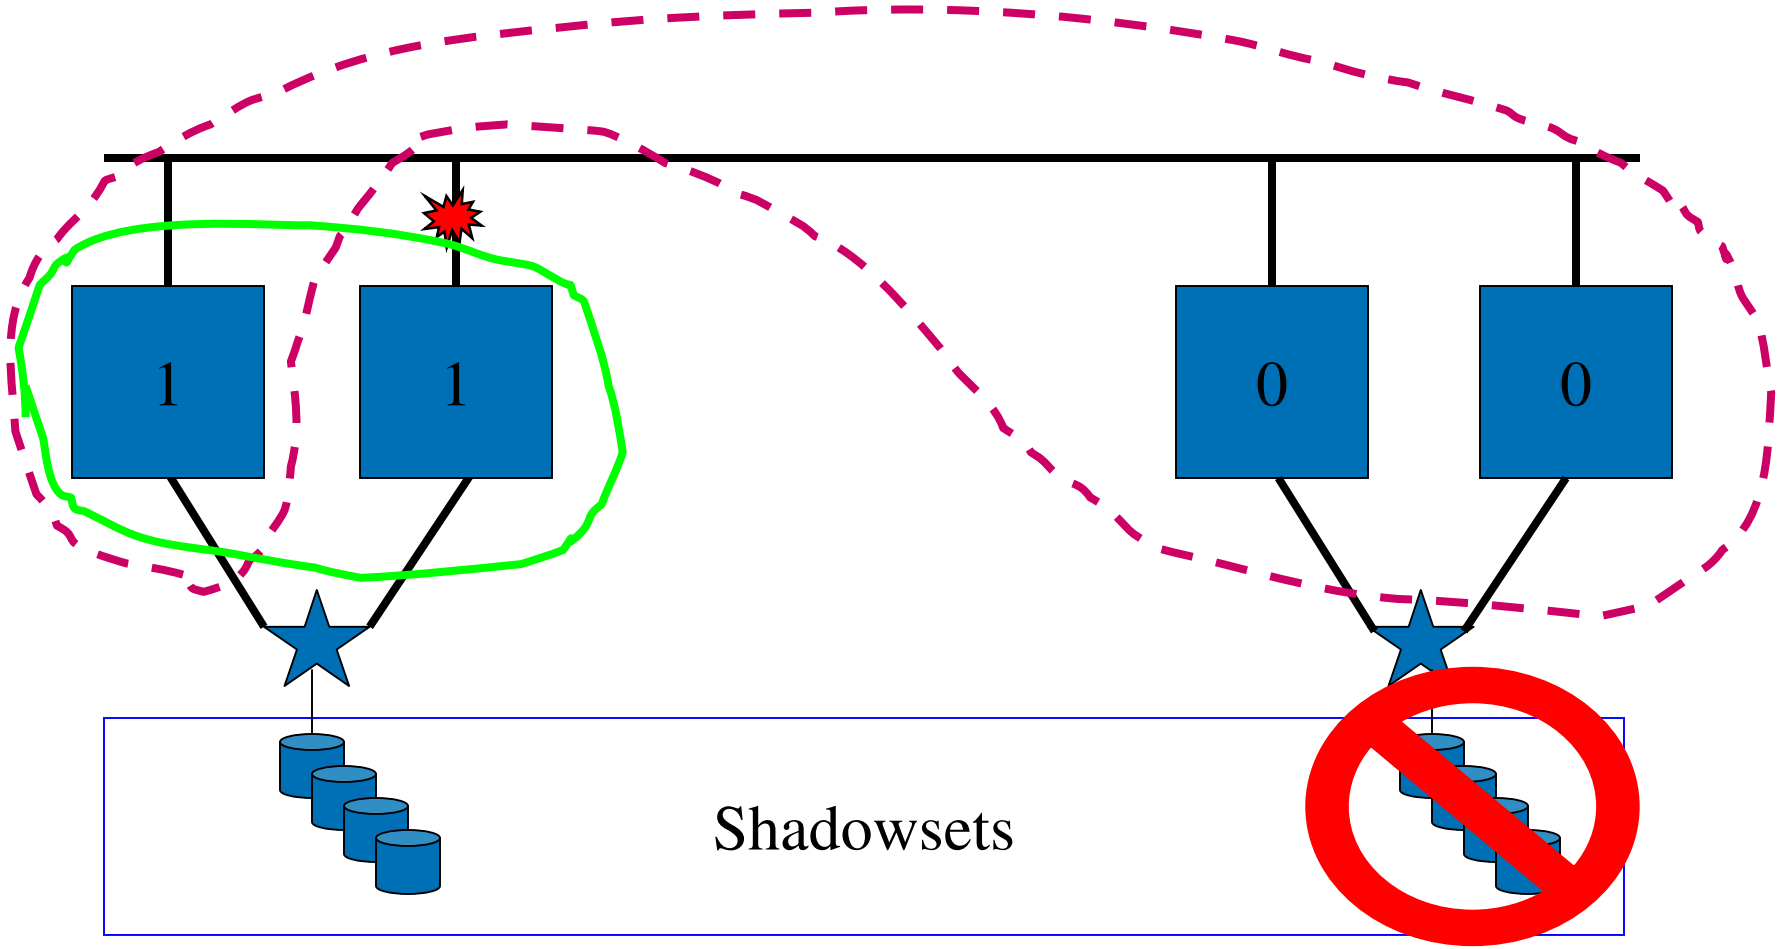

Nodes at this site continue

Nodes at this site CLUEXIT

Optimal Sub-cluster Selection: Two-Site Cluster with Unbalanced Votes

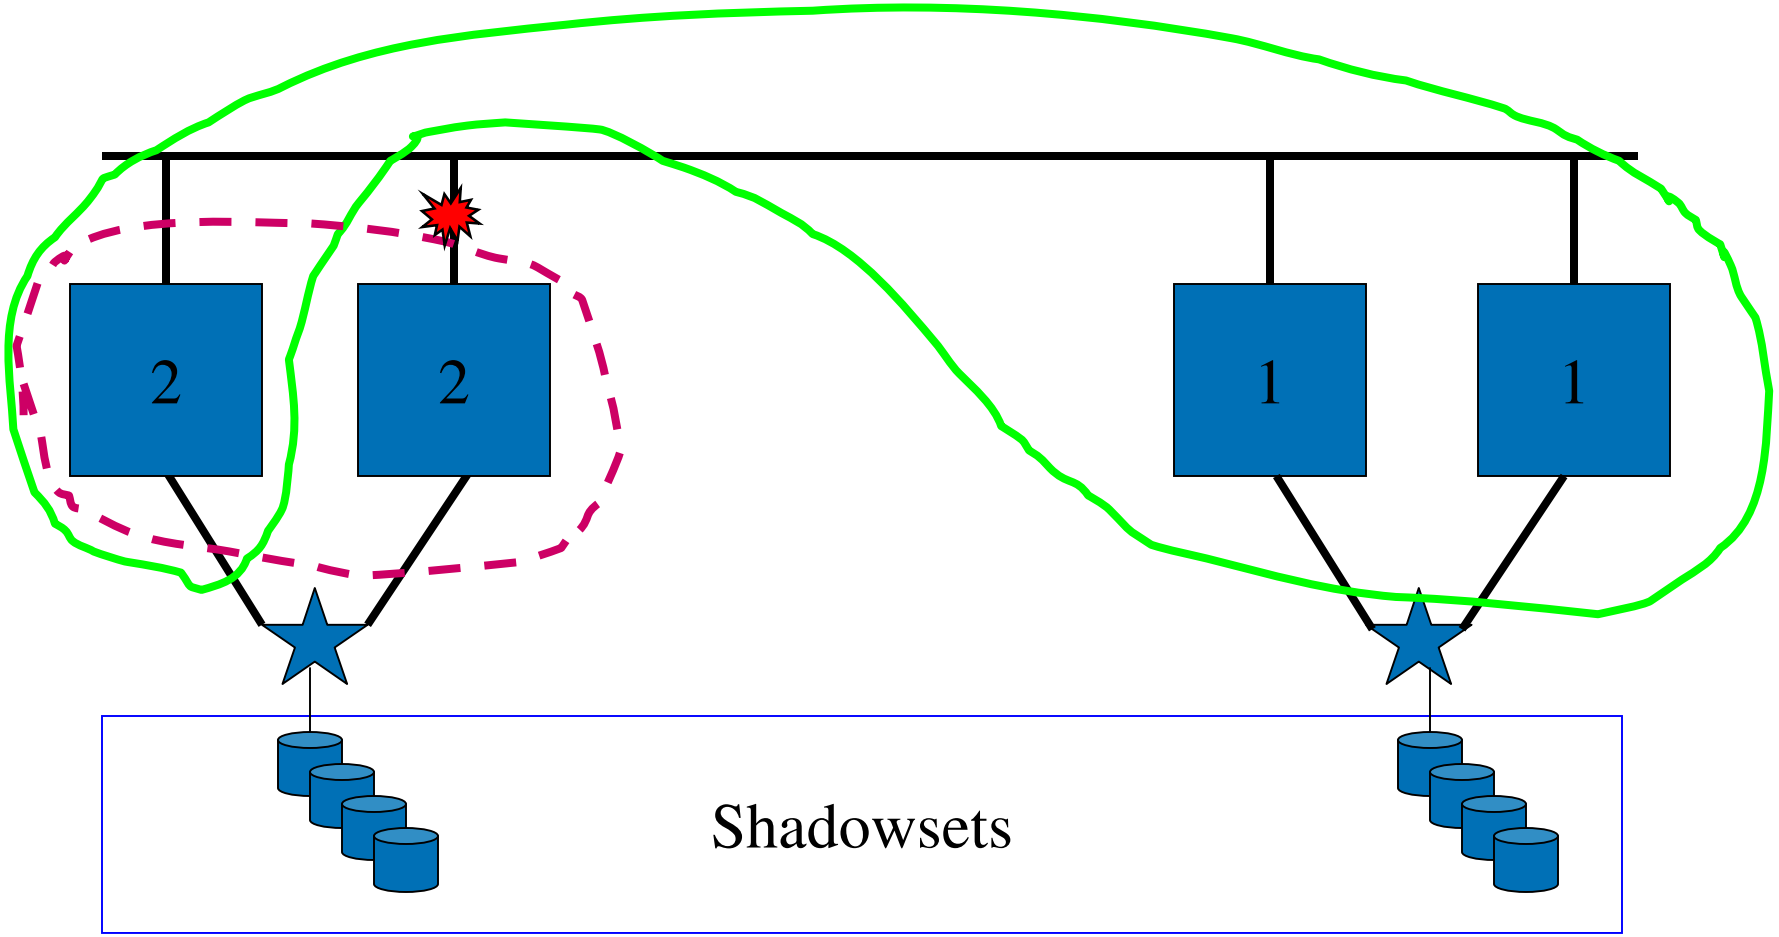

One possible solution

Quorum configurations in Two-Site Clusters

- Balanced Votes
  - Equal votes at each site
  - Manual action required to restore quorum and continue processing in the event of either:
    - Site failure, or
    - Inter-site link failure

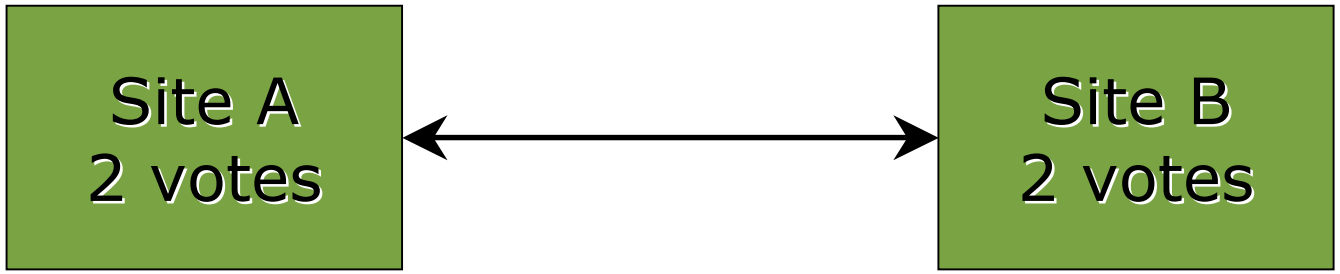

Requires manual intervention to continue alone Requires manual intervention to continue alone

Quorum Recovery Methods

- Software interrupt at IPL 12 from console
  - IPC> Q
- Availability Manager (or DECamds):
  - System Fix; Adjust Quorum
- DTCS (or BRS) integrated tool, using same RMDRIVER interface as AM / DECamds

Quorum configurations in Two-Site Clusters

- Balanced Votes
  - Note: Using REMOVE_NODE option with SHUTDOWN.COM (post V6.2) when taking down a node effectively "unbalances" votes:

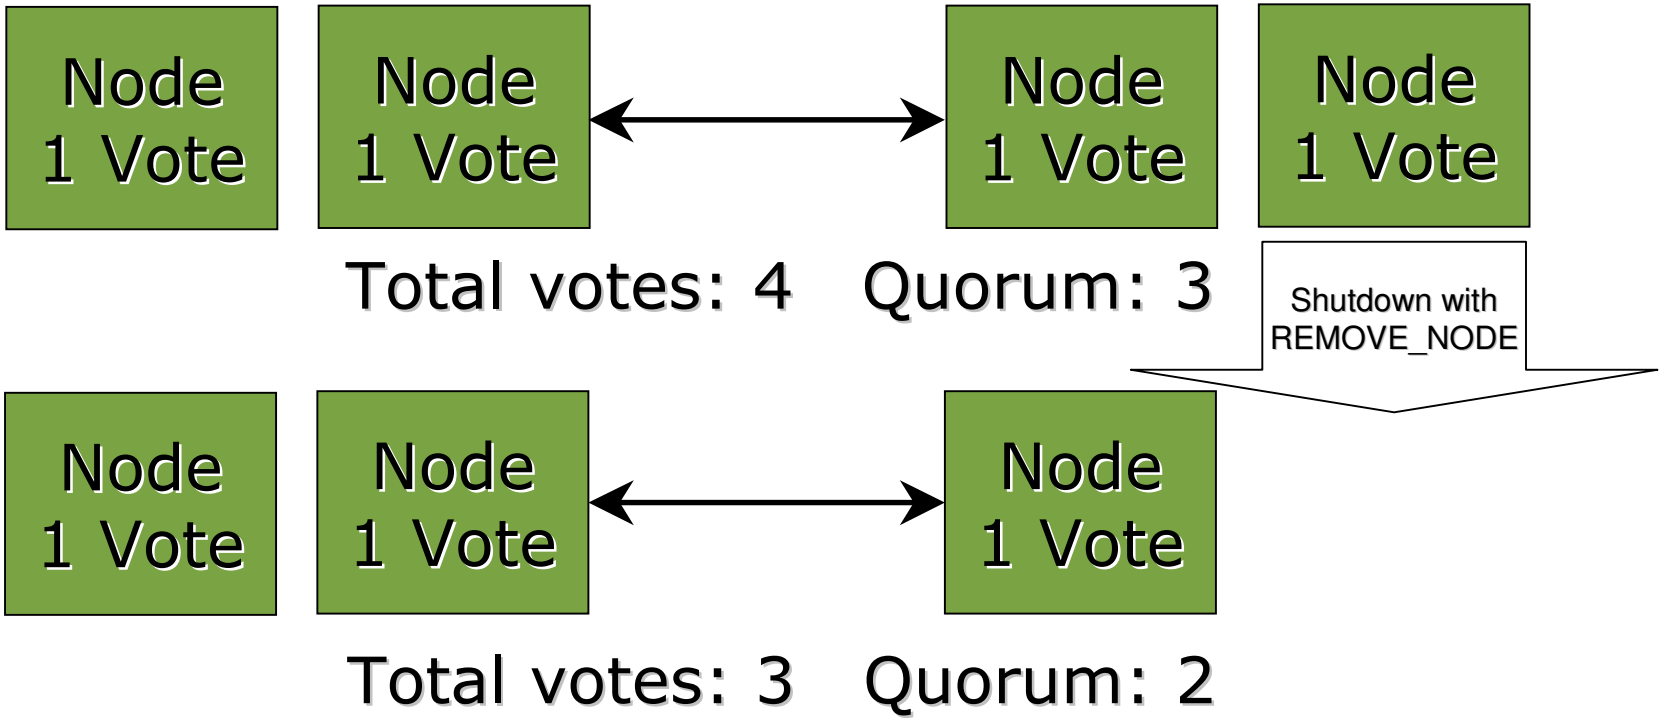

Storage and Volume Shadowing

Conventional OpenVMS Cluster design

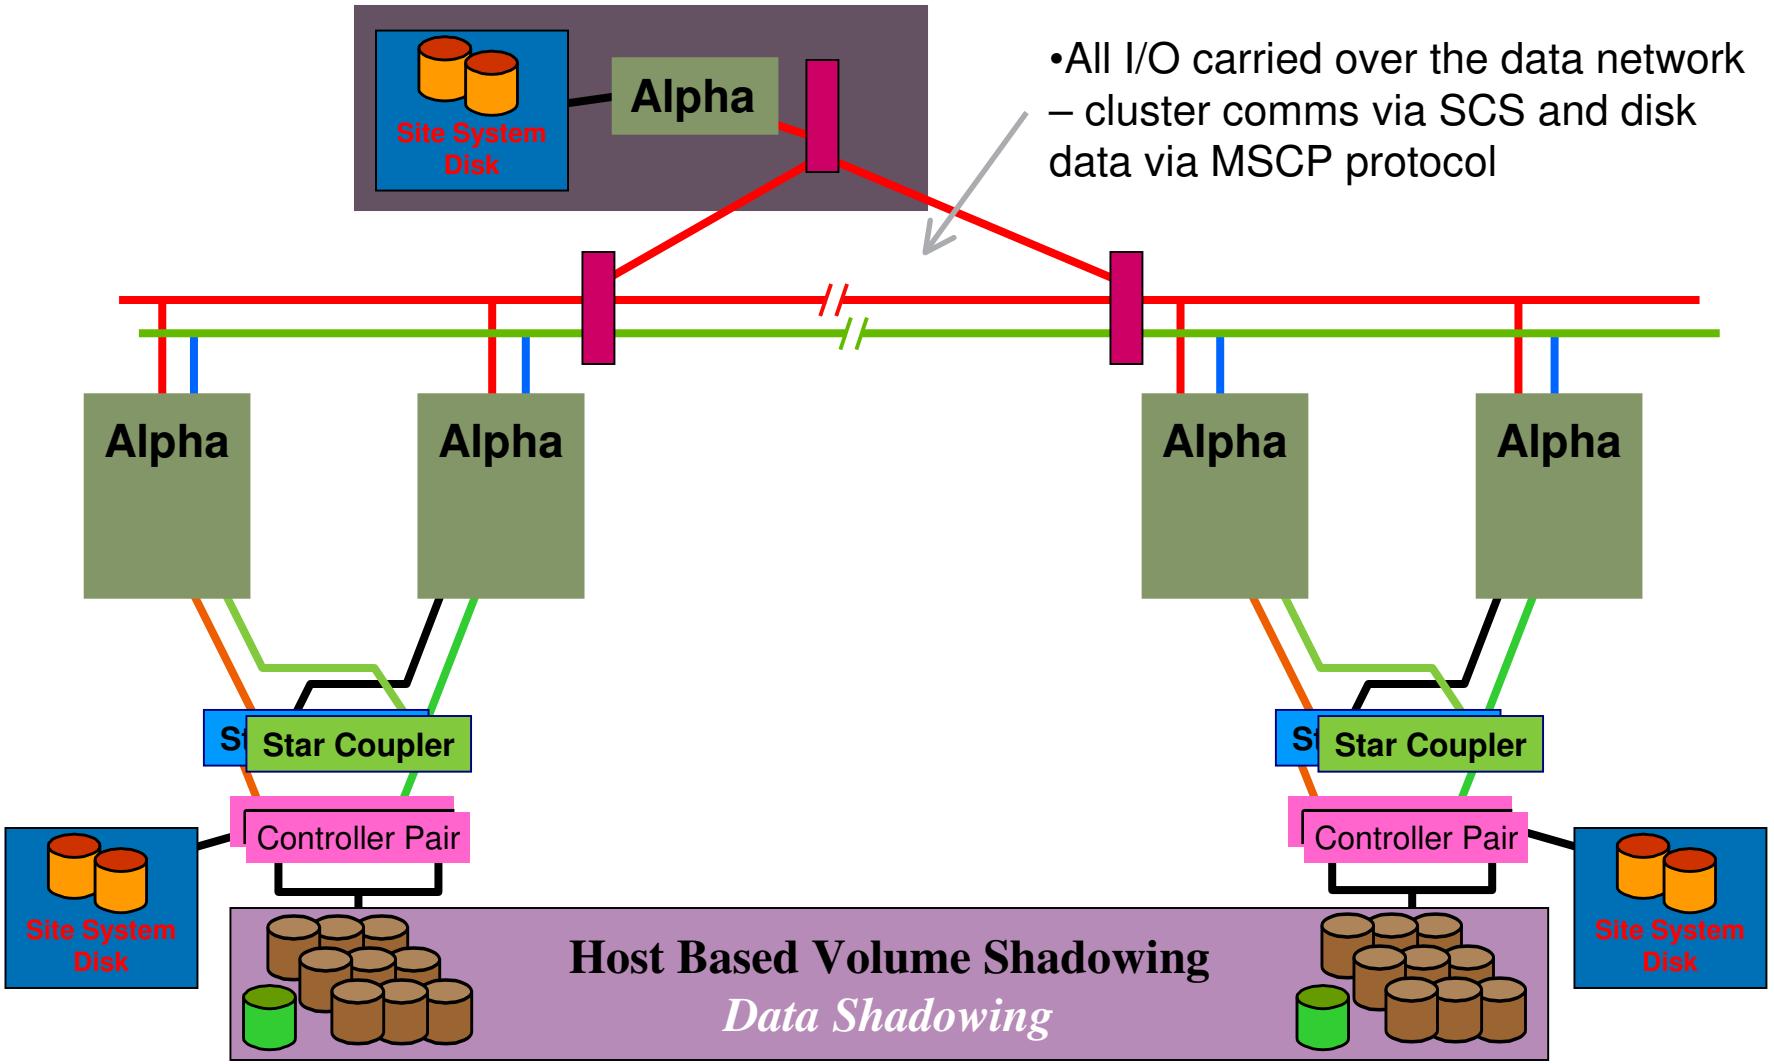

FibreChannel DT Cluster

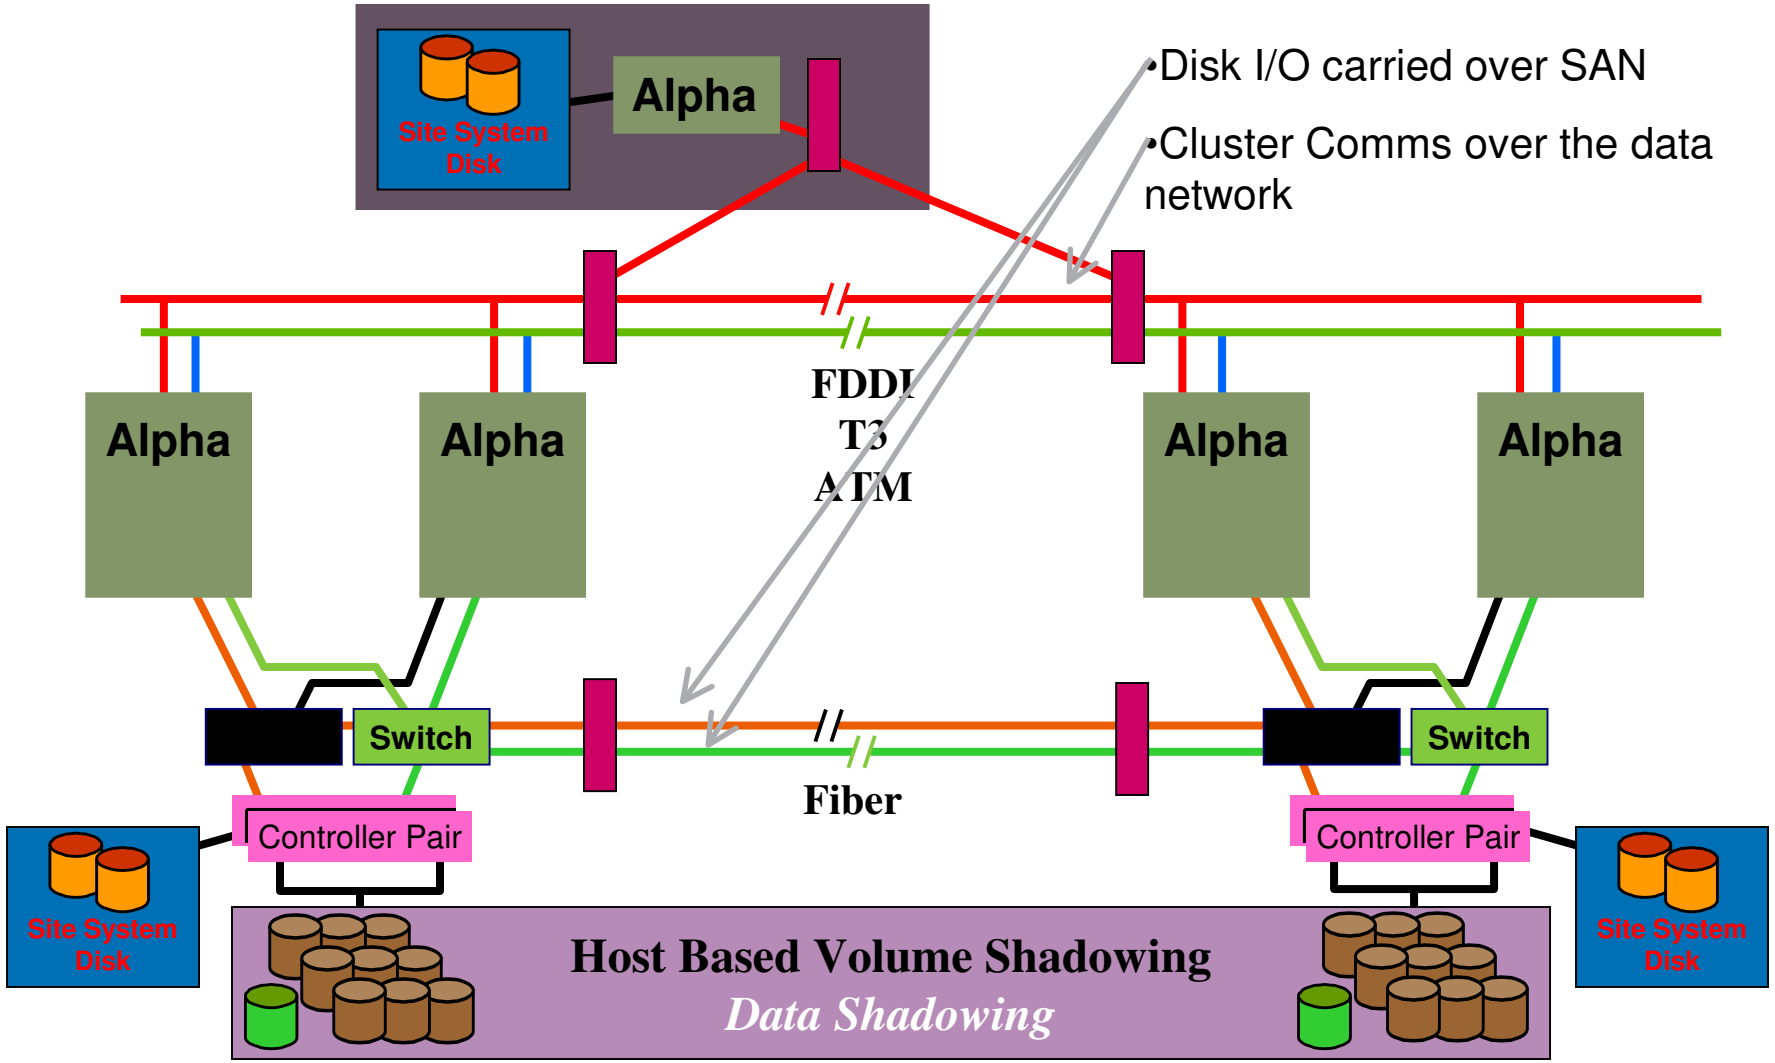

Split-Site OpenVMS with FibreChannel

- Advantages over conventional cluster (MSCP)
  - Much better performance
    - No MSCP server overhead / latency
    - Not limited by network protocol performance
    - 100 Mbyte/Second bandwidth throughout (200 Mbyte with 2 Gbit FC)
  - “Direct” paths to remote storage
    - Easier to maintain sites (all nodes in a site can go down without removing access to storage)
- Disadvantages
  - Issues with failure modes
  - Configuration must be handled carefully
  - Difficult to distinguish between local and remote disks
- Option of using DRM for data replication

CA & DRM - Reading

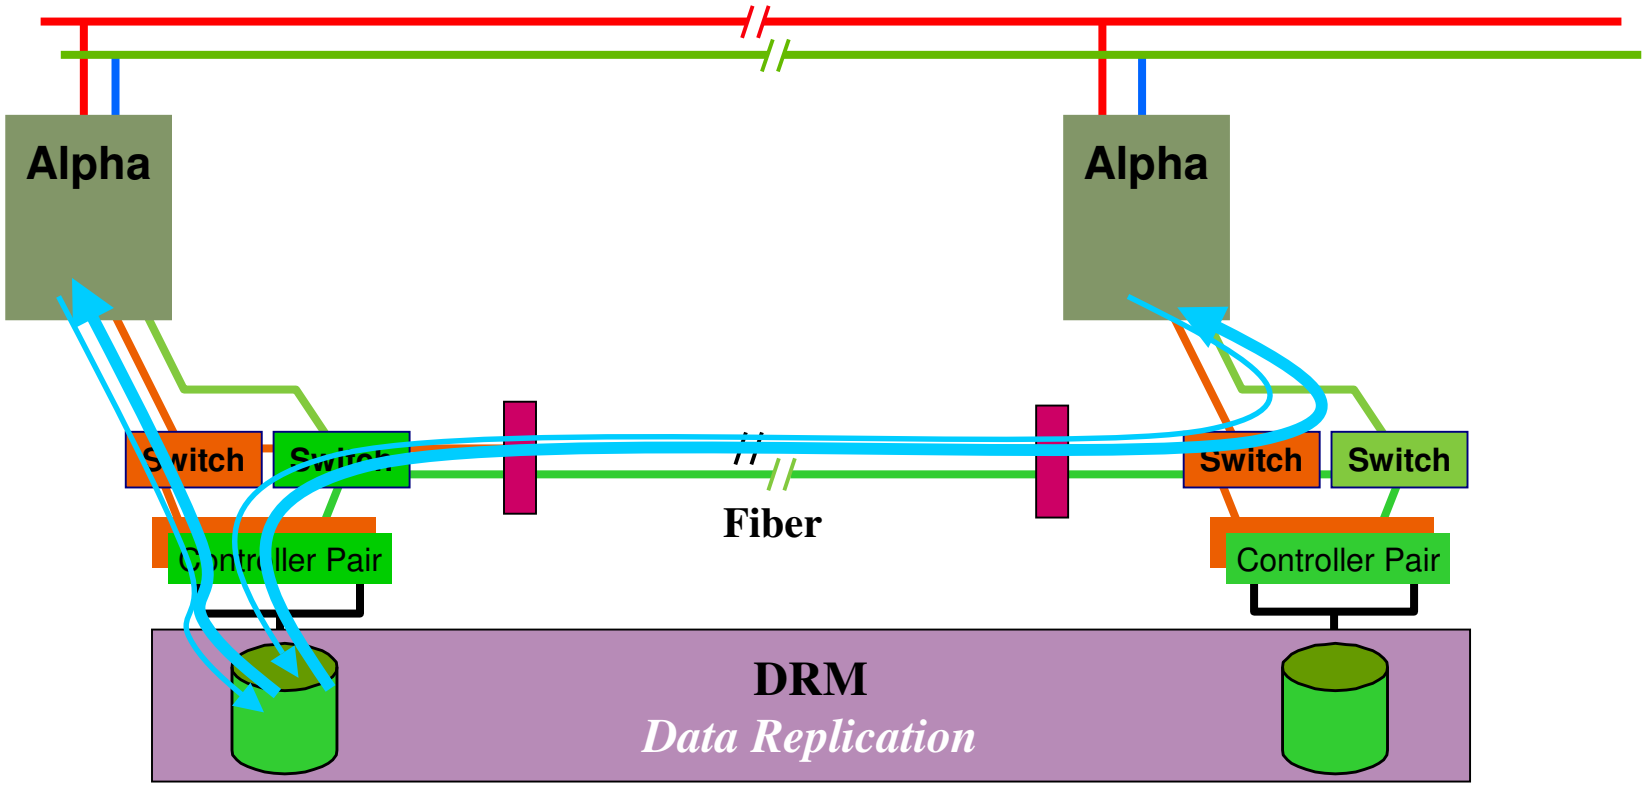

Host-Based Shadowing - Reading

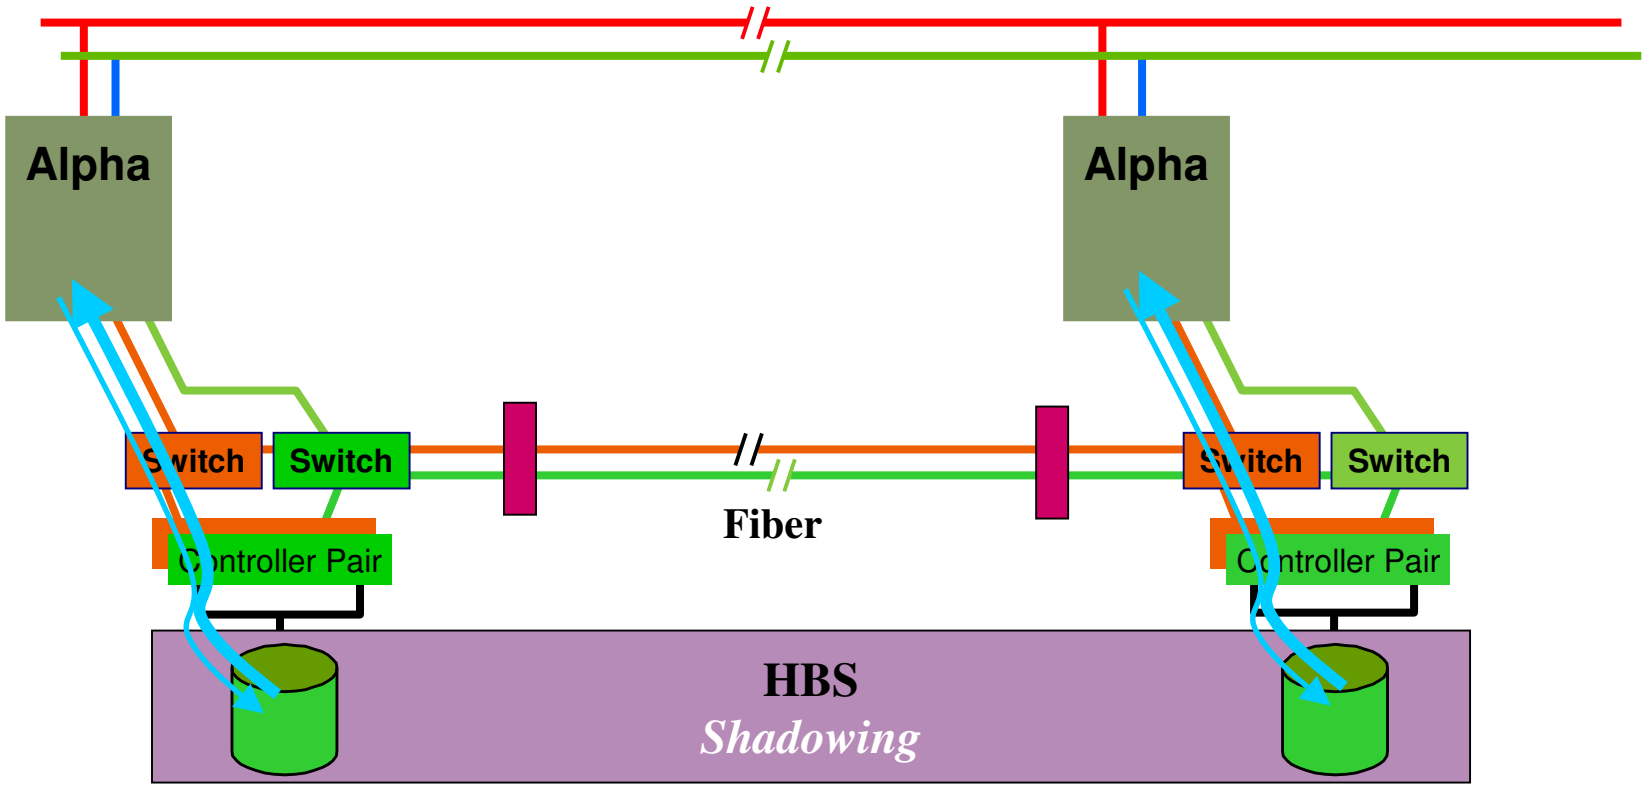

CA & DRM - Writing

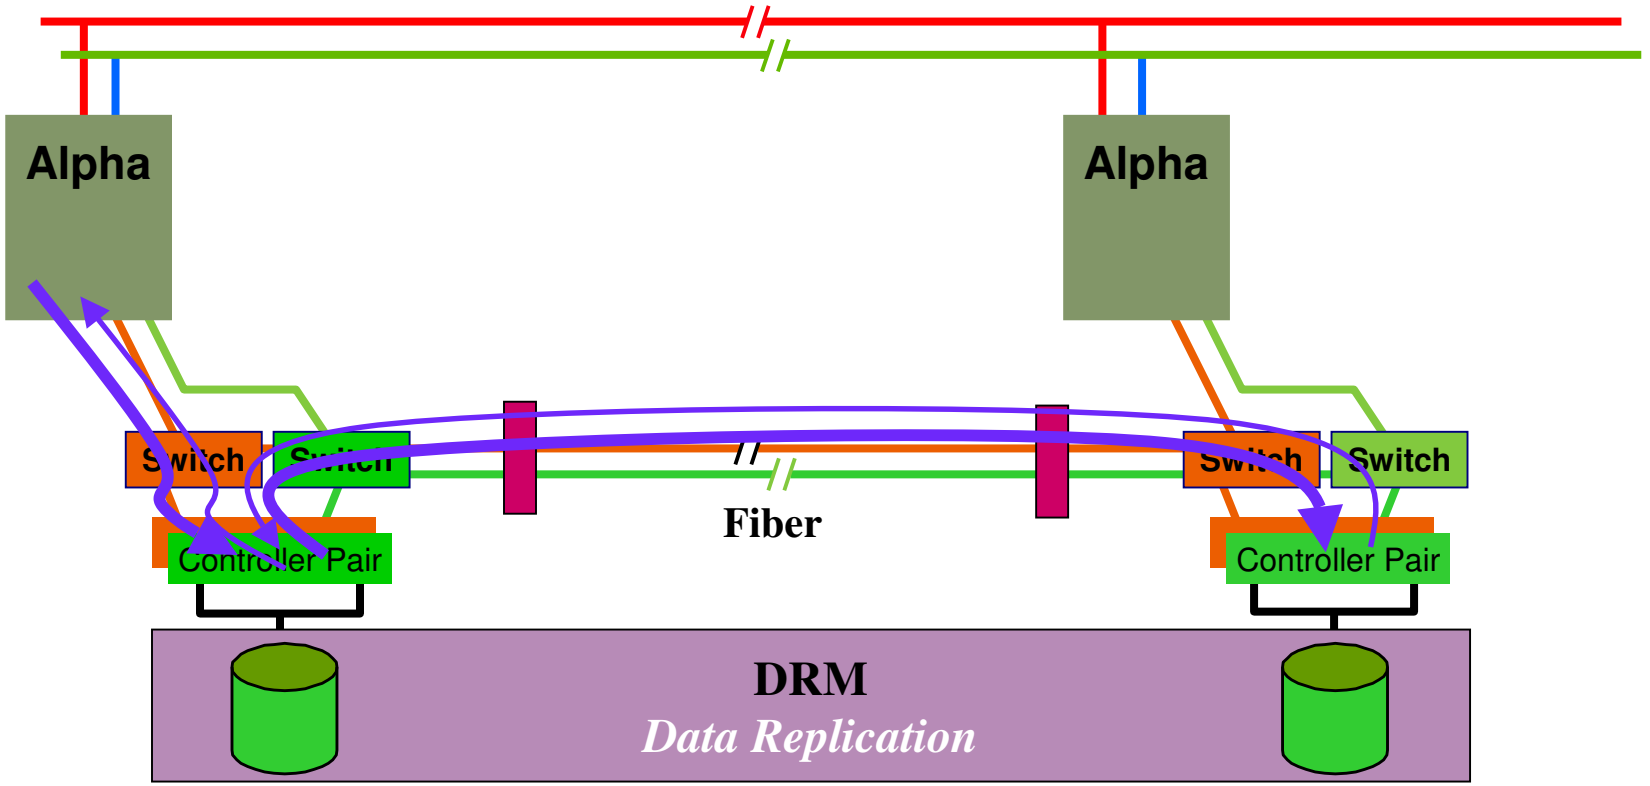

Host Based Shadowing - Writing

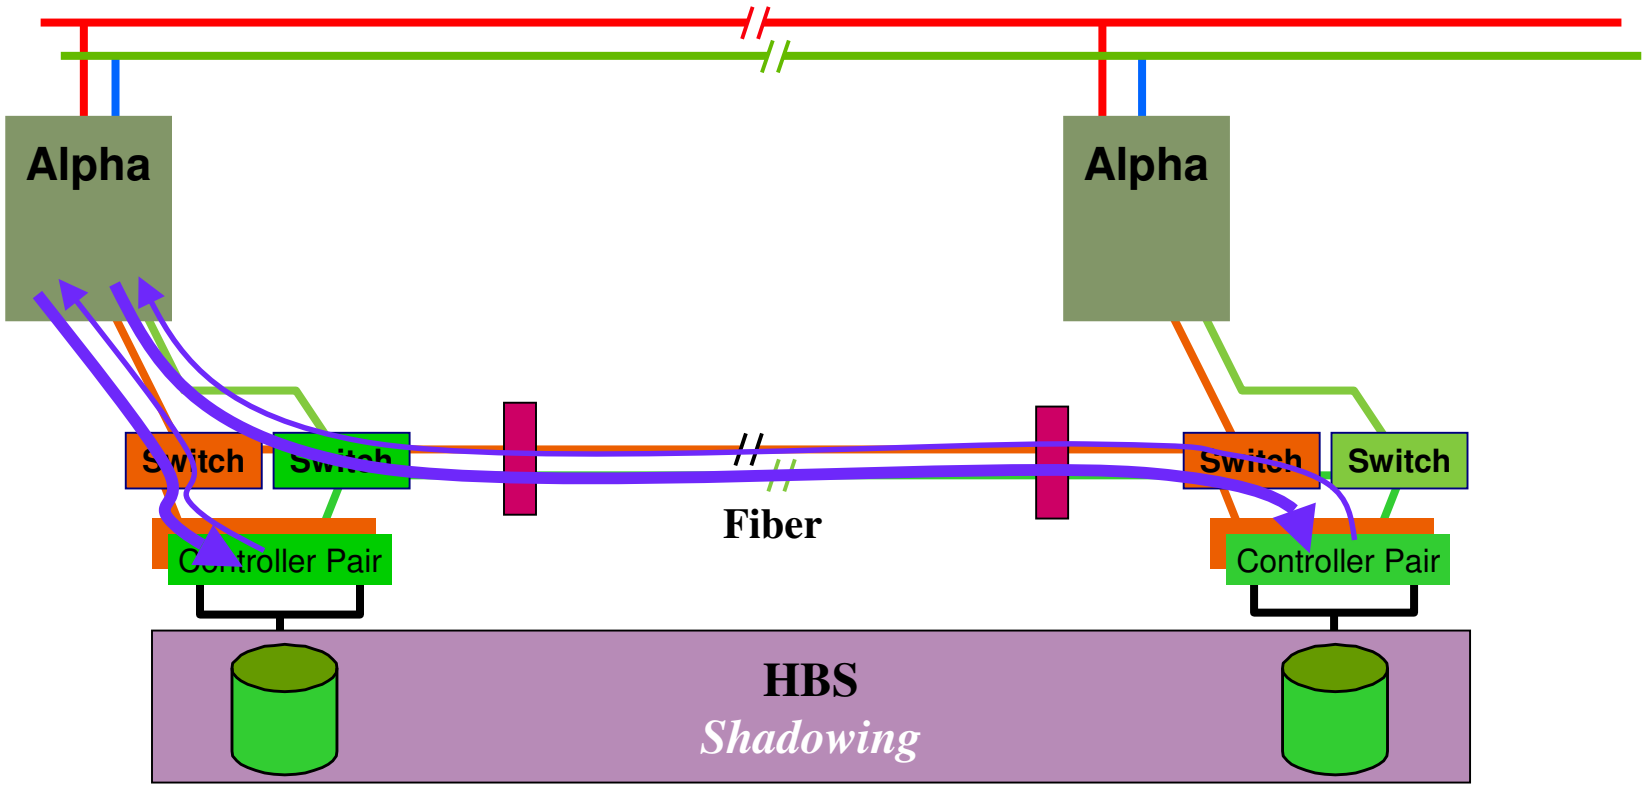

FibreChannel DT Cluster with HBS

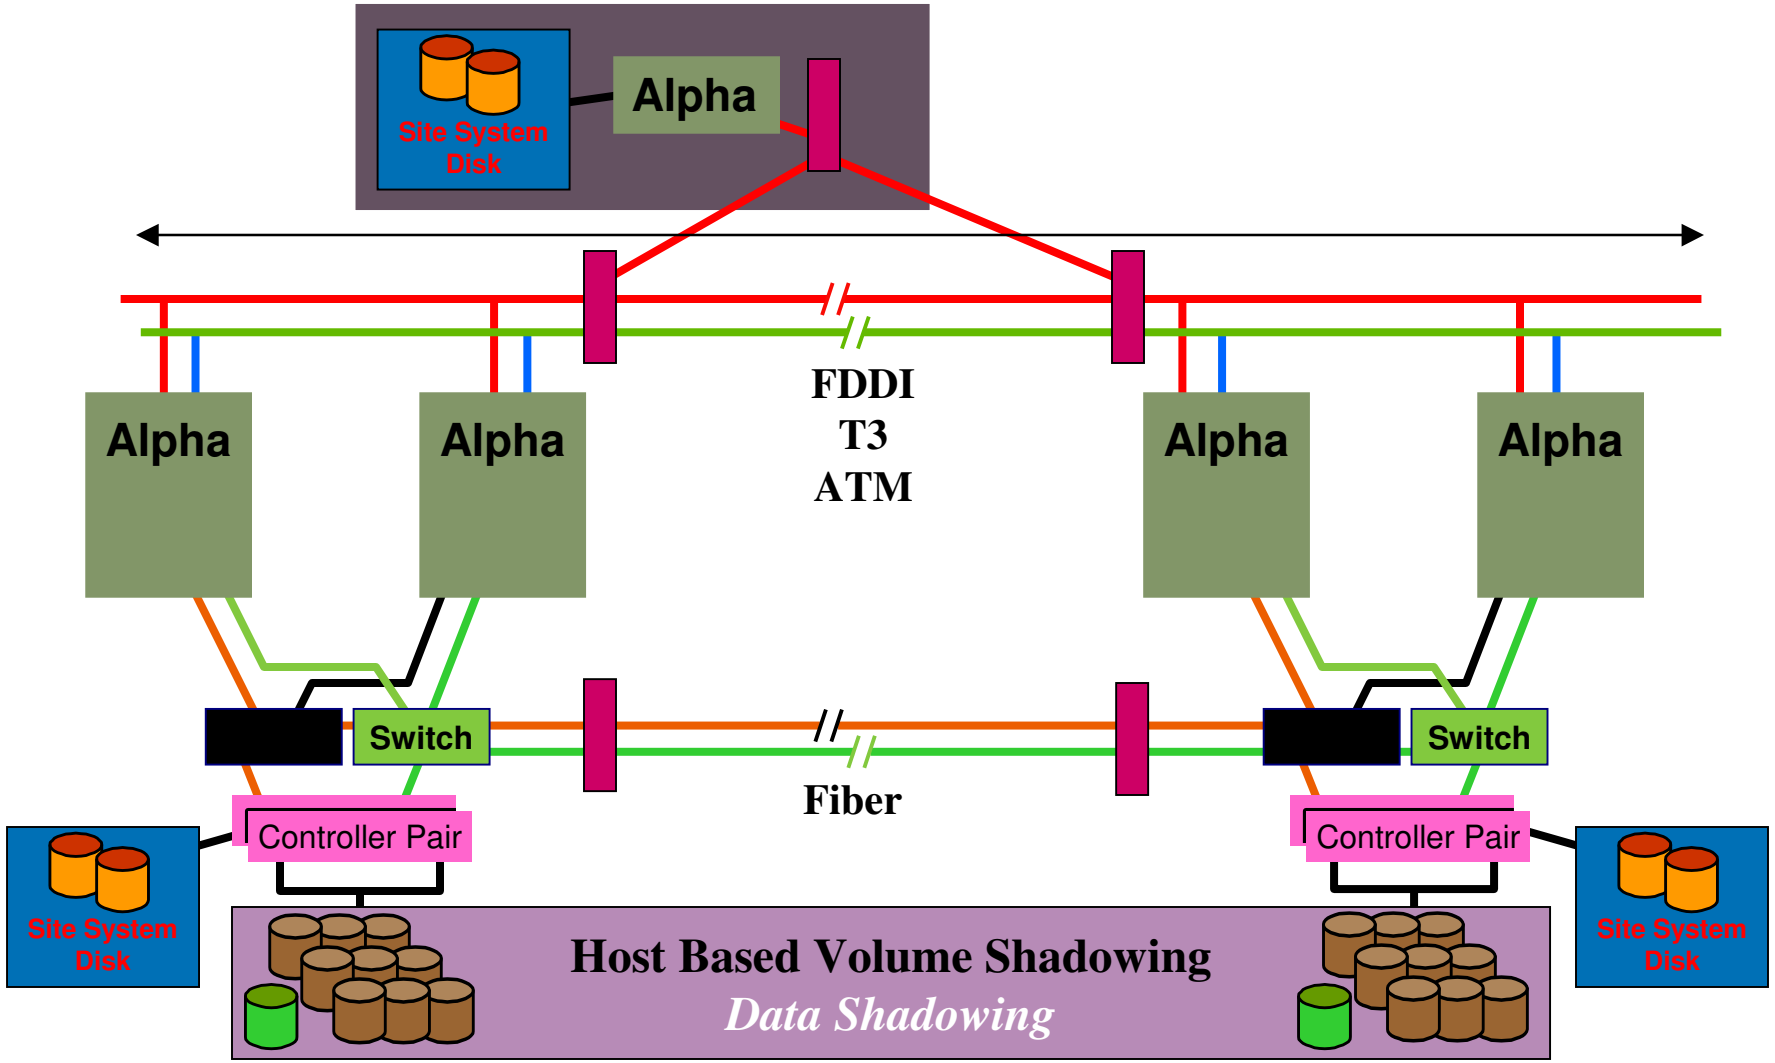

FibreChannel DT Cluster with CA

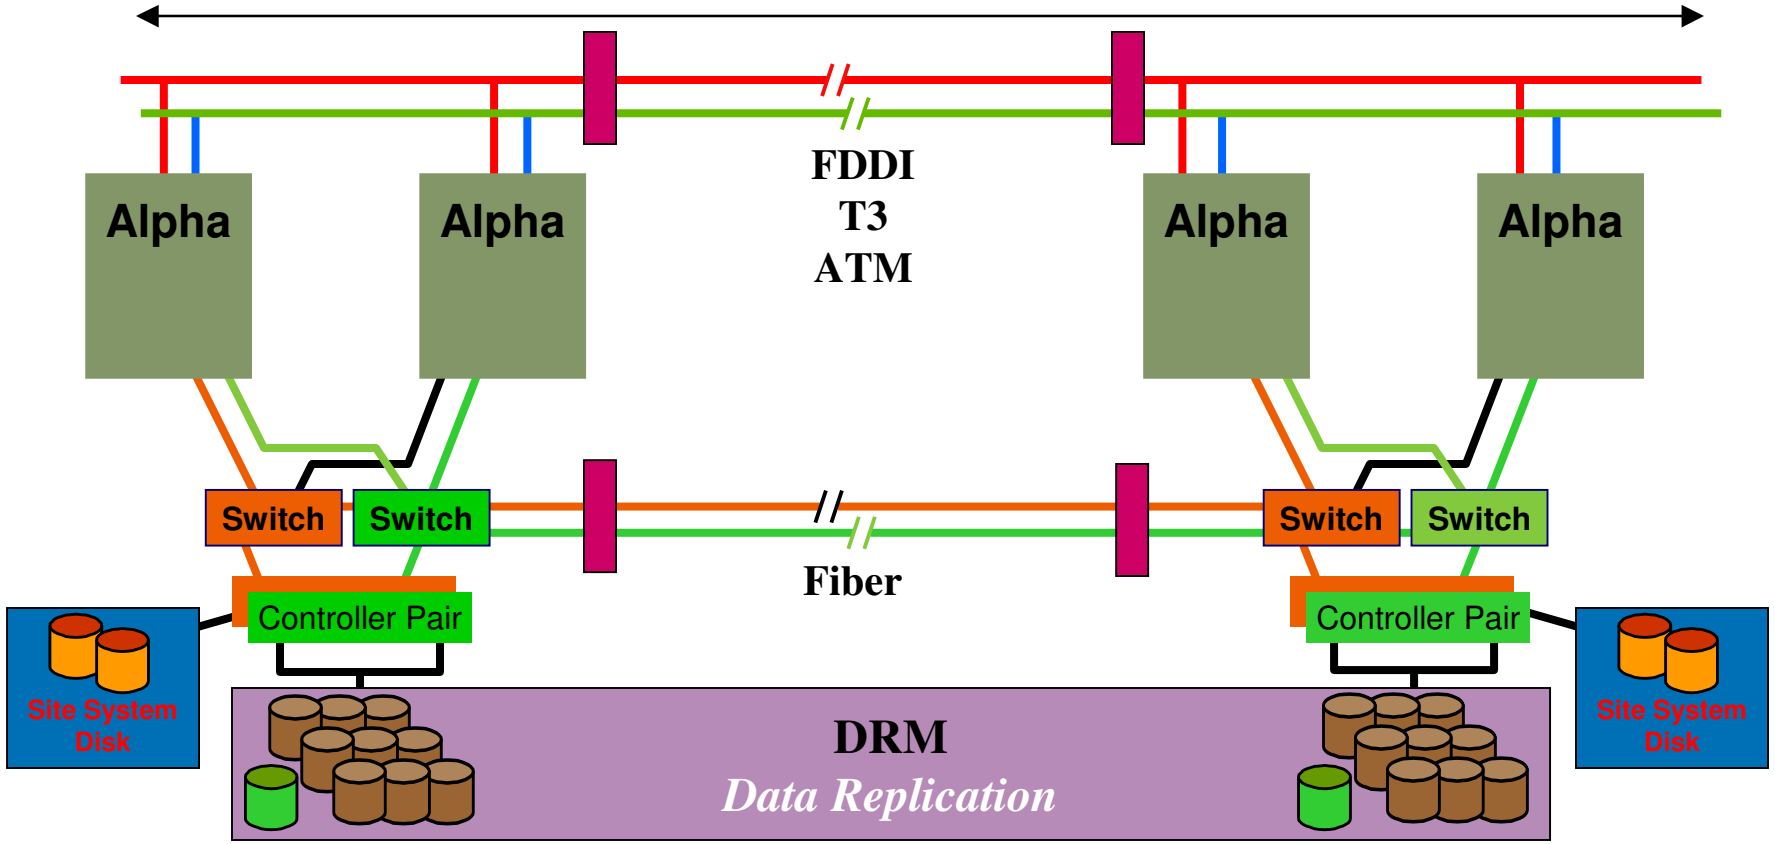

Split-site Replication Comparison

	OpenVMS Host Based Shadowing	DRM / CA
Shadow / Mirror Model	Peer-peer	Master-slave
Shadow / Mirror Features	Fully synchronous	Synchronous or Asynchronous
Configuration Data	Simple metadata on each volume	Metadata in database in NVRAM
Automatic Continuation after Failure	Yes	If secondary fails and Failsafe Mode Disabled – Yes If secondary fails and failsafe mode enabled – No If primary fails – No
Down-time after site failure	Short period for cluster quorum recalculation, lock adjustment and application recovery. No reboot required	10-20 Minutes Reboot and storage reconfig may be required
Recovery after cluster / site repair	Few seconds No reboot required	10-20 Minutes Reboot and storage reconfig may be required

Split-site Replication Comparison (cont)

	OpenVMS Host Based Shadowing	DRM / CA
Shadow catchup performance	Host systems provide processing power for shadow copies – overall performance will reduce as applications will have reduced available CPU power and I/O bandwidth	Storage controllers provide processing power for copies – additional load on inter-site link but otherwise no additional overhead
Merge copies	As shadow-catchup but more processing required *(Mini-merge will improve this)*	As shadow catchup. Logging also available to improve re-synch times.
Steady-state read performance	OpenVMS can read from either volume – read throughput doubles	OpenVMS only sees a single volume – no performance change
Steady-state write performance	OpenVMS must write to both volumes – write throughput halves	OpenVMS only sees a single volume – no performance impact
Monitoring	Fully integrated into OpenVMS – full alert / problem reporting through normal OpenVMS mechanisms	Operators must actively check subsystem for many error conditions Normal reporting difficult
Manageability	Fully integrated into OpenVMS – simple commands to reconfigure / recover	Mixed environment. SAN Appliance for set-up and performance management DTCS GUI or menus for recovery

Storage and Volume Shadowing

Host-Based Volume Shadowing

- Host software keeps multiple disks identical
- All writes go to all shadowset members
- Reads can be directed to any one member
  - Different read operations can go to different members at once, helping throughput
- Synchronization (or Re-synchronization after a failure) is done with a Copy operation
- Re-synchronization after a node failure is done with a Merge operation

Shadowing Full-Copy Algorithm

- Host-Based Volume Shadowing full-copy algorithm is non-intuitive.
  1. Read from source disk
  2. Do Compare operation with target disk
  3. If data is identical, we're done with this segment. If data is different, write source data to target disk, then go back to Step 1.
- Shadow_Server process does copy I/Os
  - Does one 127-block segment at a time, from the beginning of the disk to the end, with no double-buffering or other speed-up tricks

Speeding Shadow Copies

- Implications:
  - Shadow copy completes fastest if data is identical beforehand
    - Fortunately, this is the most-common case – re-adding a shadow member into shadowset again after it was a member before
- If you know that a member will be removed and later re-added, the Mini-Copy capability can be a great time-saver.

Data Protection Mechanisms and Scenarios

- Protection of the data is obviously extremely important in a disaster-tolerant cluster
- We'll examine the mechanisms the Volume Shadowing design uses to protect the data
- We'll look at one scenario that has happened in real life and resulted in data loss:
  - "Wrong-way shadow copy"

Protecting Shadowed Data

- Shadowing keeps a “Generation Number” in the SCB on shadow member disks
- Shadowing “Bumps” the Generation number at the time of various shadowset events, such as mounting, or membership changes

Protecting Shadowed Data

- Generation number is designed to monotonically increase over time, never decrease
- Implementation is based on OpenVMS timestamp value
  - During a “Bump” operation it is increased
    - to the current time value, or
      - if the generation number already represents a time in the future for some reason (such as time skew among cluster member clocks), then it is simply incremented
  - The new value is stored on all shadowset members at the time of the Bump operation

Protecting Shadowed Data

- Generation number in SCB on removed members will thus gradually fall farther and farther behind that of current members
- In comparing two disks, a later generation number should always be on the more up-to-date member, under normal circumstances

System Management





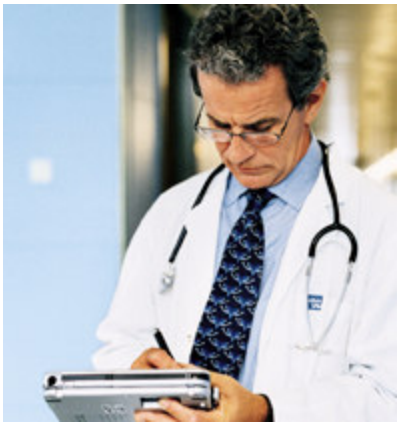

Documentation

- Comprehensive and accurate documentation essential for FAST recovery.
- DTCS is a business solution, TIME IS MONEY…..
- Many approaches which should include:-
  - Site Documentation
    - include Network and Cluster
  - Disaster Recovery Plan, components critical to running of the applications on the DTC

Topics for Site Documentation - Cluster/Network/SAN

- Cluster Configuration
  - Disk Controller configuration
  - Fibre Channel Switches
  - Node Configurations
  - System disks
  - Application disks
  - SYSGEN parameters
  - Logical names
  - Start-up files - map
  - Page/Swap Files Configuration
  - License management information
- Network Configuration
  - Network device configuration details (priorities, filters etc..)
  - Port Usage details
  - FDDI ring map
  - Software/firmware rev levels
  - Changing management & port modules.
  - Console port connections
- SAN Configuration
  - Switch configuration details (address, ports, ISL etc…)
  - Switch Zoning

System disk design - 1

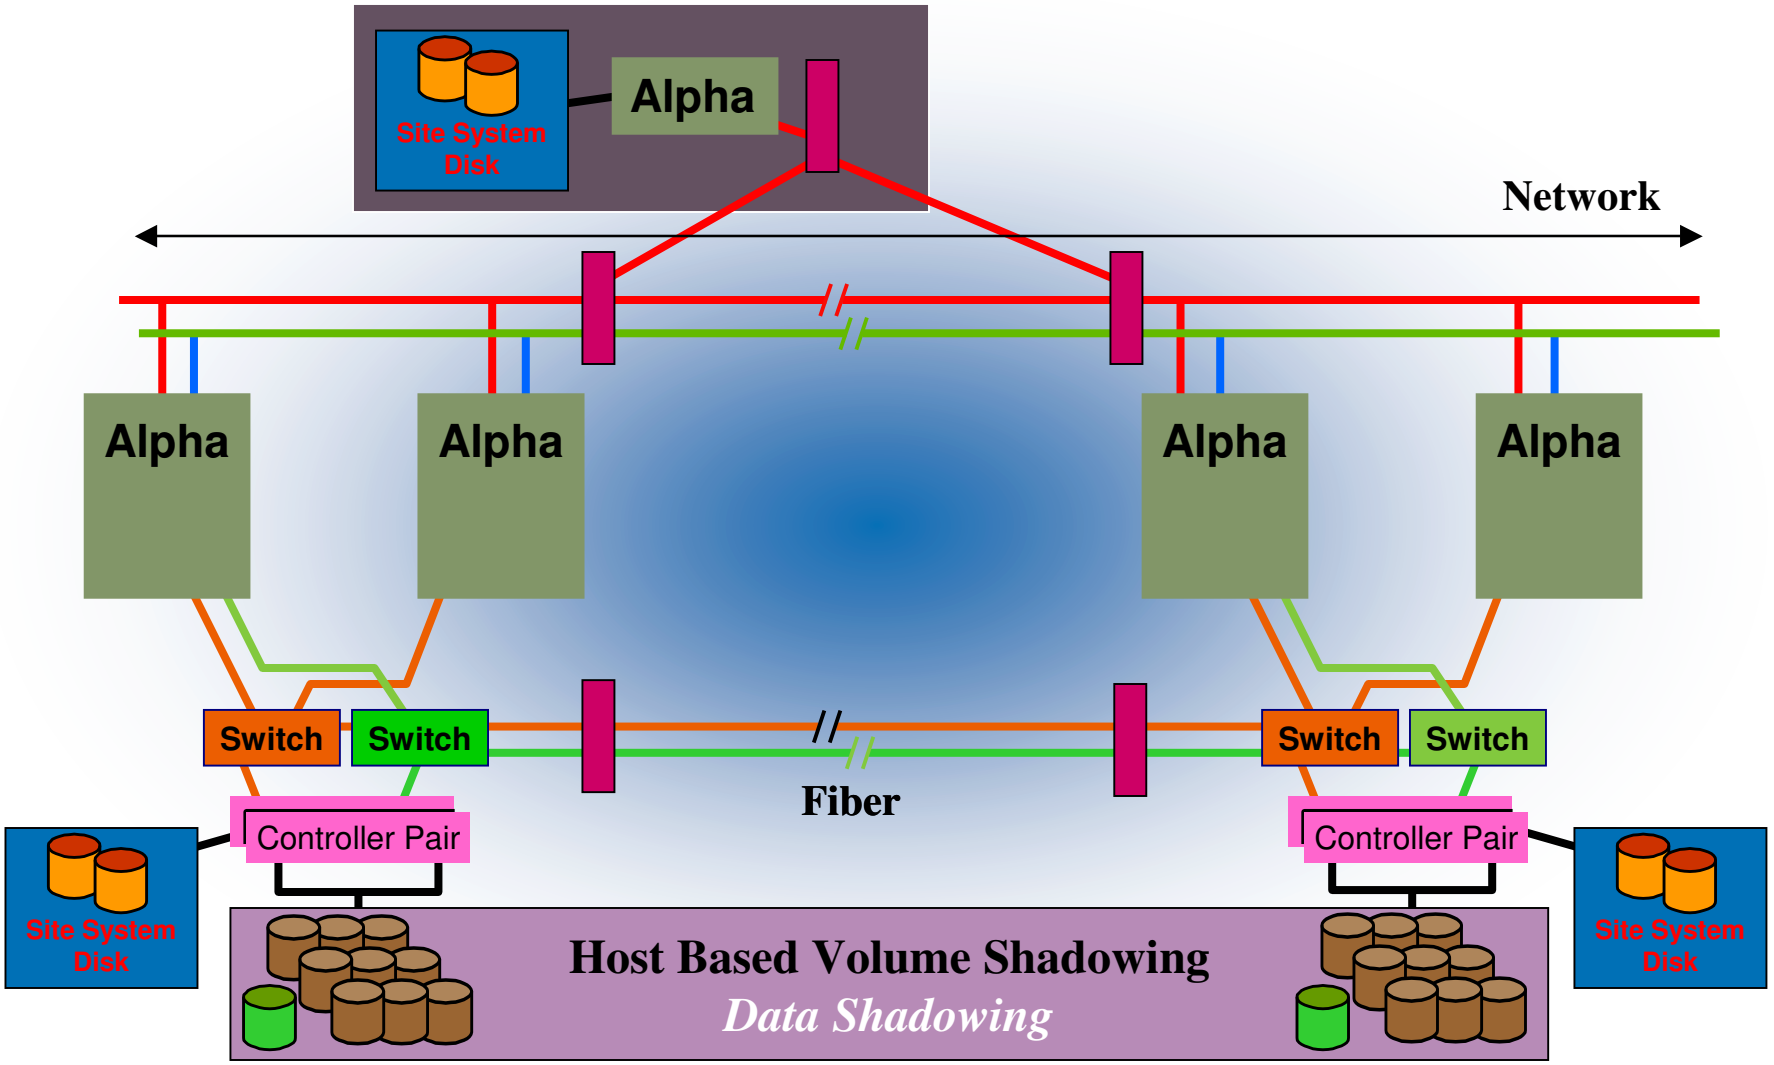

System disk design - 2

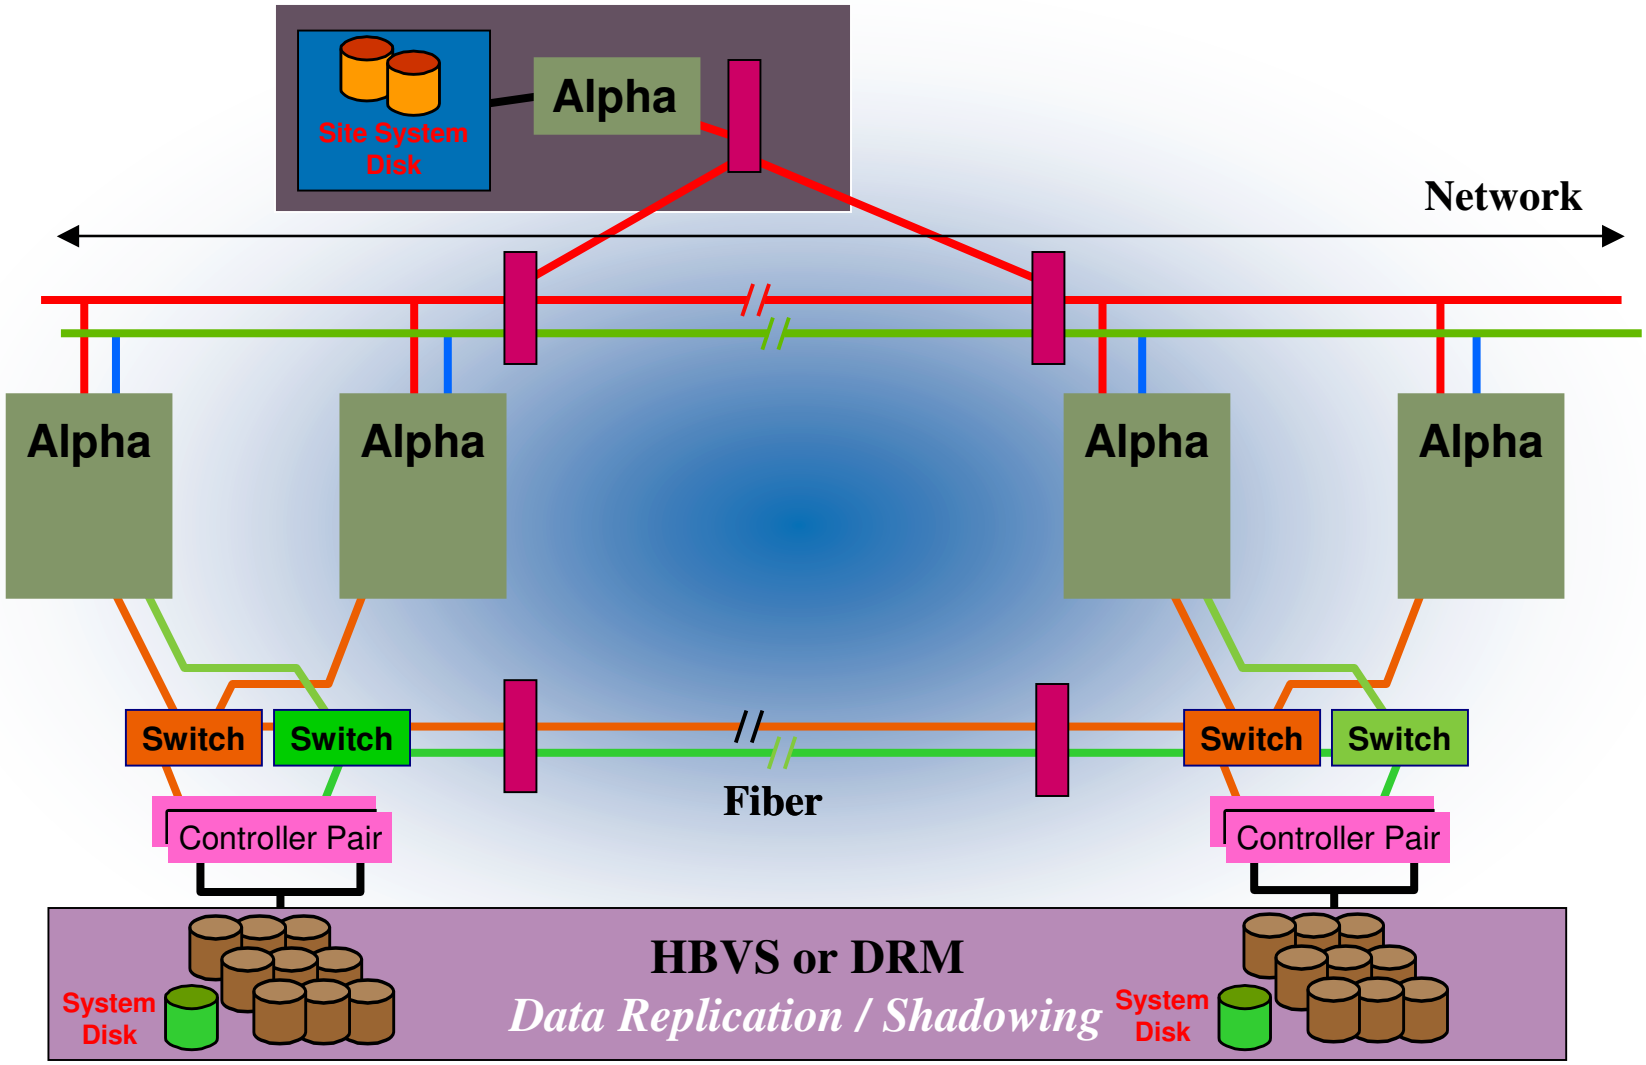

Beware of too much flexibility….

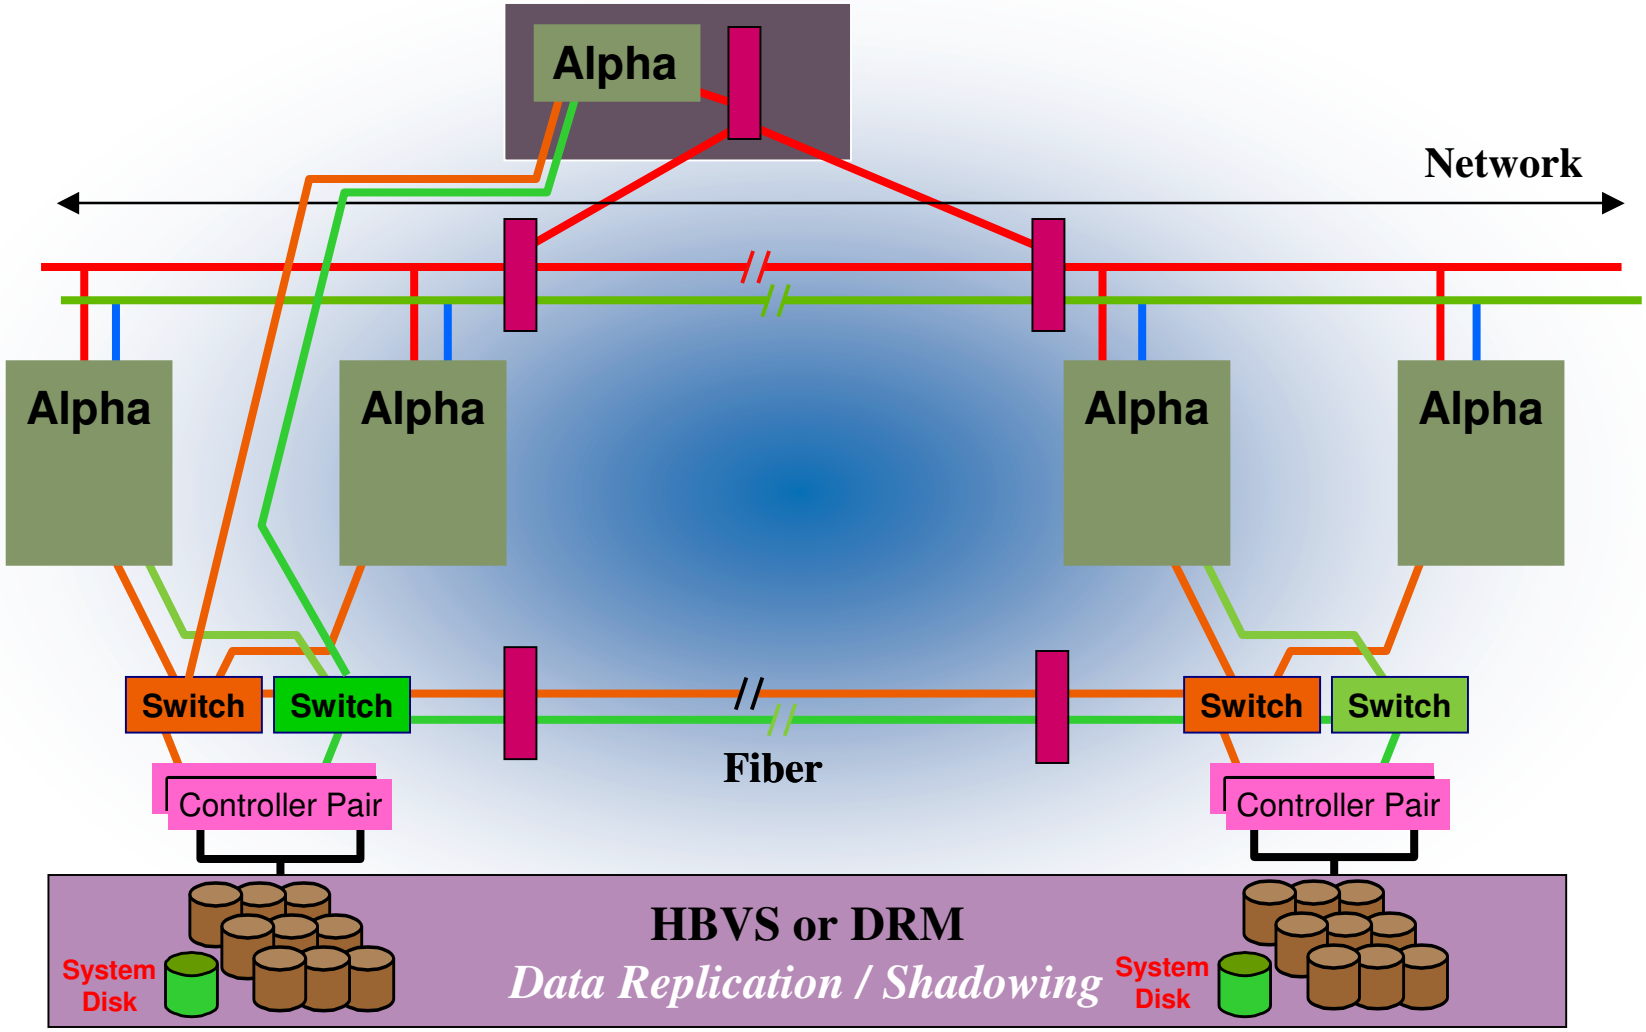

System disk issues for DT

- separate system disk
  - upgrade-ability
  - must have ability to boot from the same physical member
  - some “insurance” in having 2 system disks
  - paging?
  - performance (inter-switch link) esp. latency
  - must maintain all copies of OpenVMS properly
- single system disk
  - one copy of VMS to maintain
  - Cluster will hang
    - if inter-site link breaks
    - If there is a “locking” problem on the disk

Example System start-up flow

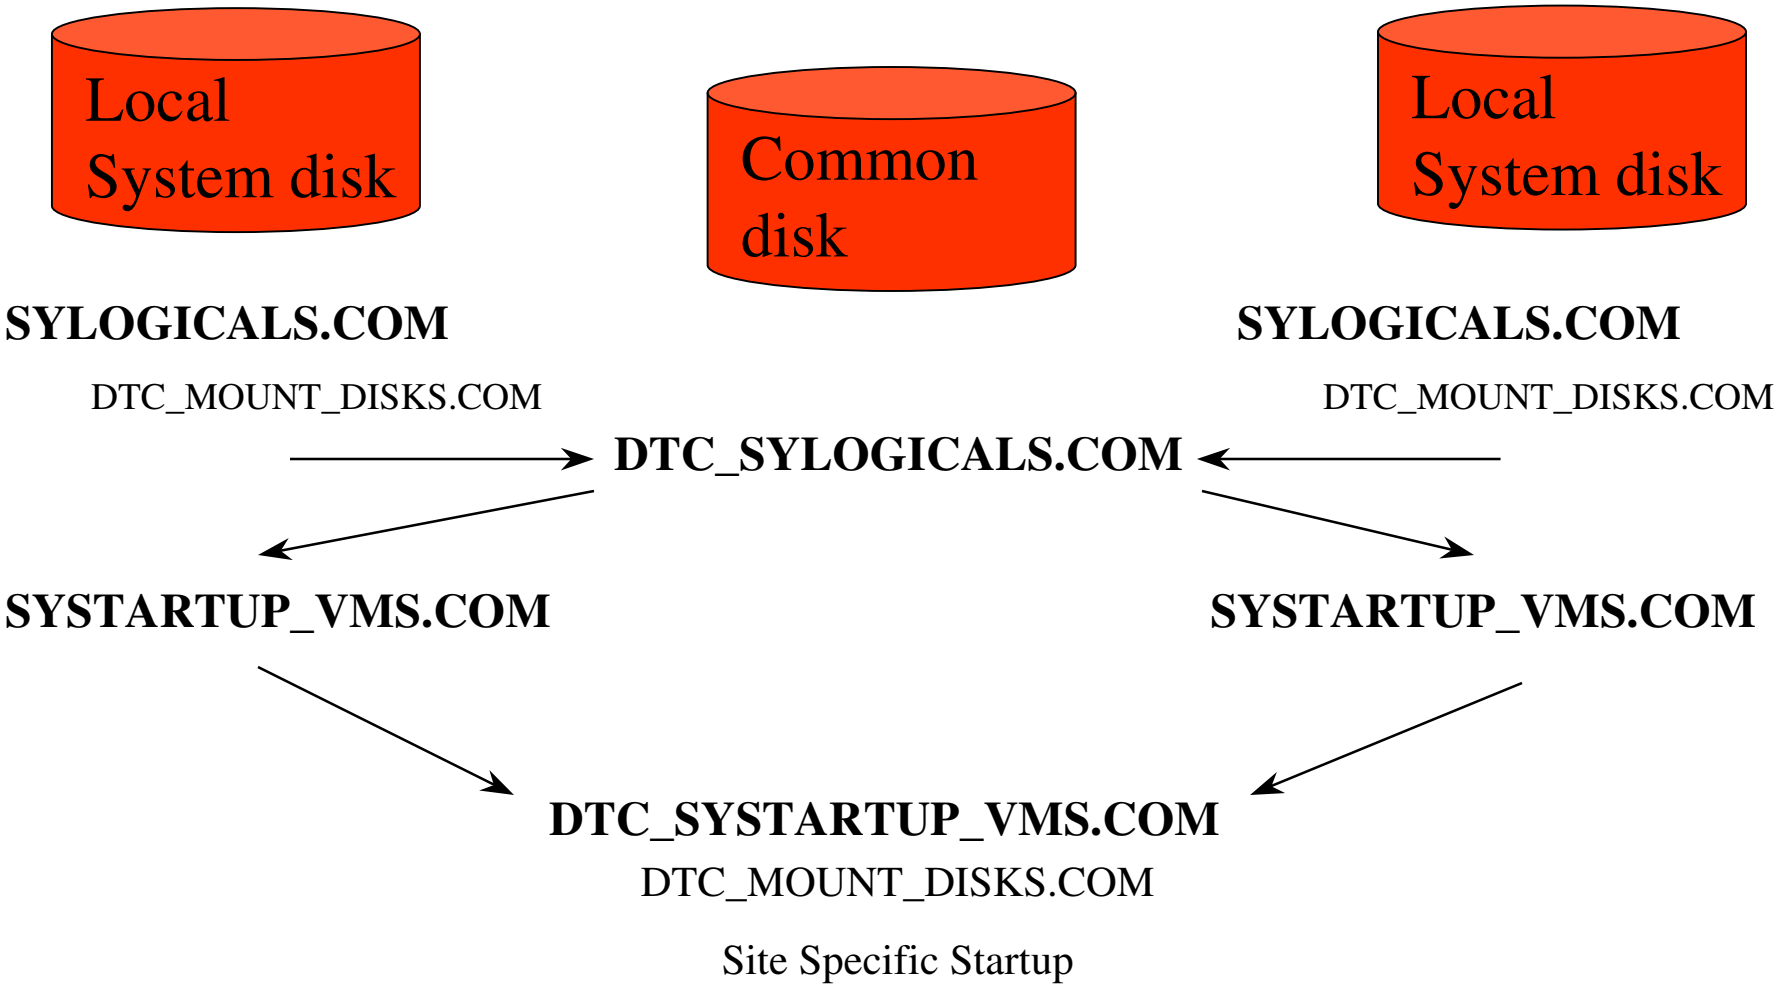

Disk mounting / dismounting

- mounting
  - mount sequence
    - which of the disks were the last on-line?

- beware disk on-line detection

Dismounting

- Clean dismounts minimises merges

.....and more

- So where is your continuity now?

Understanding what is critical

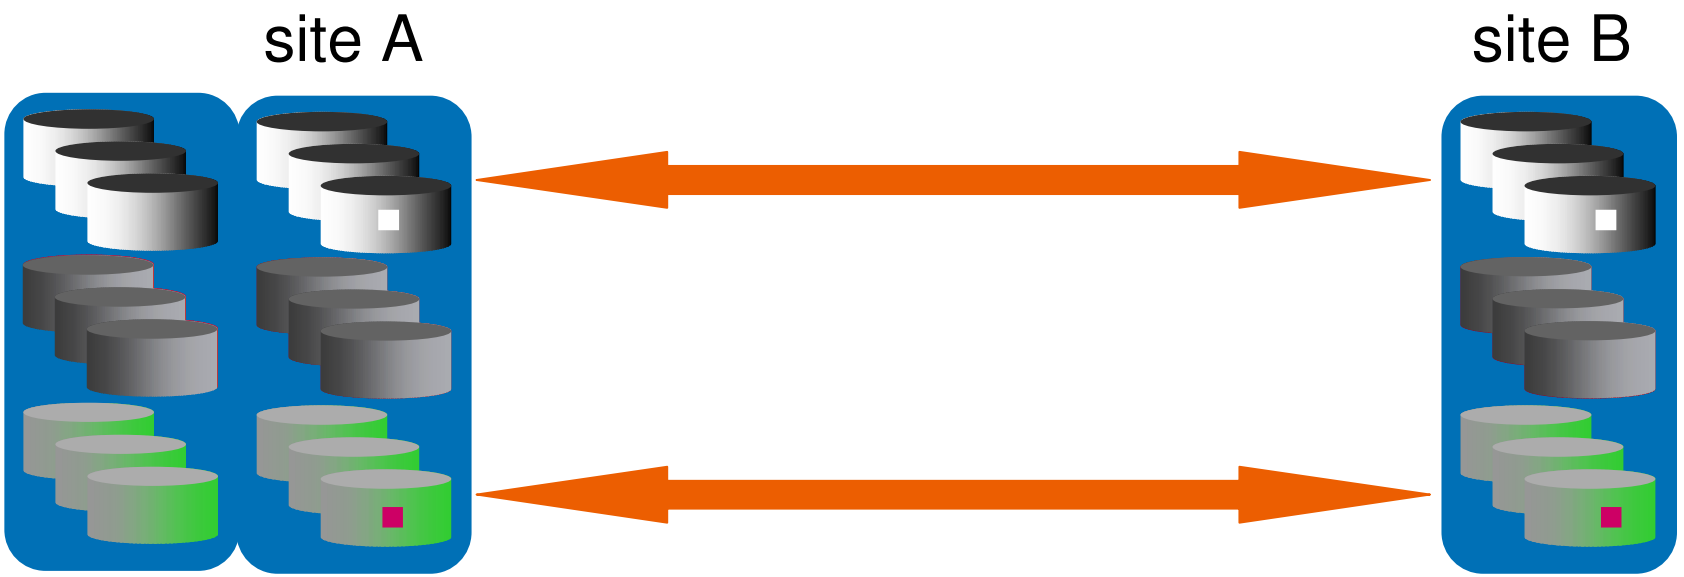

SYSGEN Parameter Tuning for DTCS - 1

- RECNXINTERVAL = Big Enough
- SHADOW_MBR_TMO = Bigger than RECNXINTERVAL
- SHADOW_MAX_COPY = ?
  - Dynamic parameter
  - Setting is configuration dependent, 4 is the default value

RCNXINTERVAL – too low

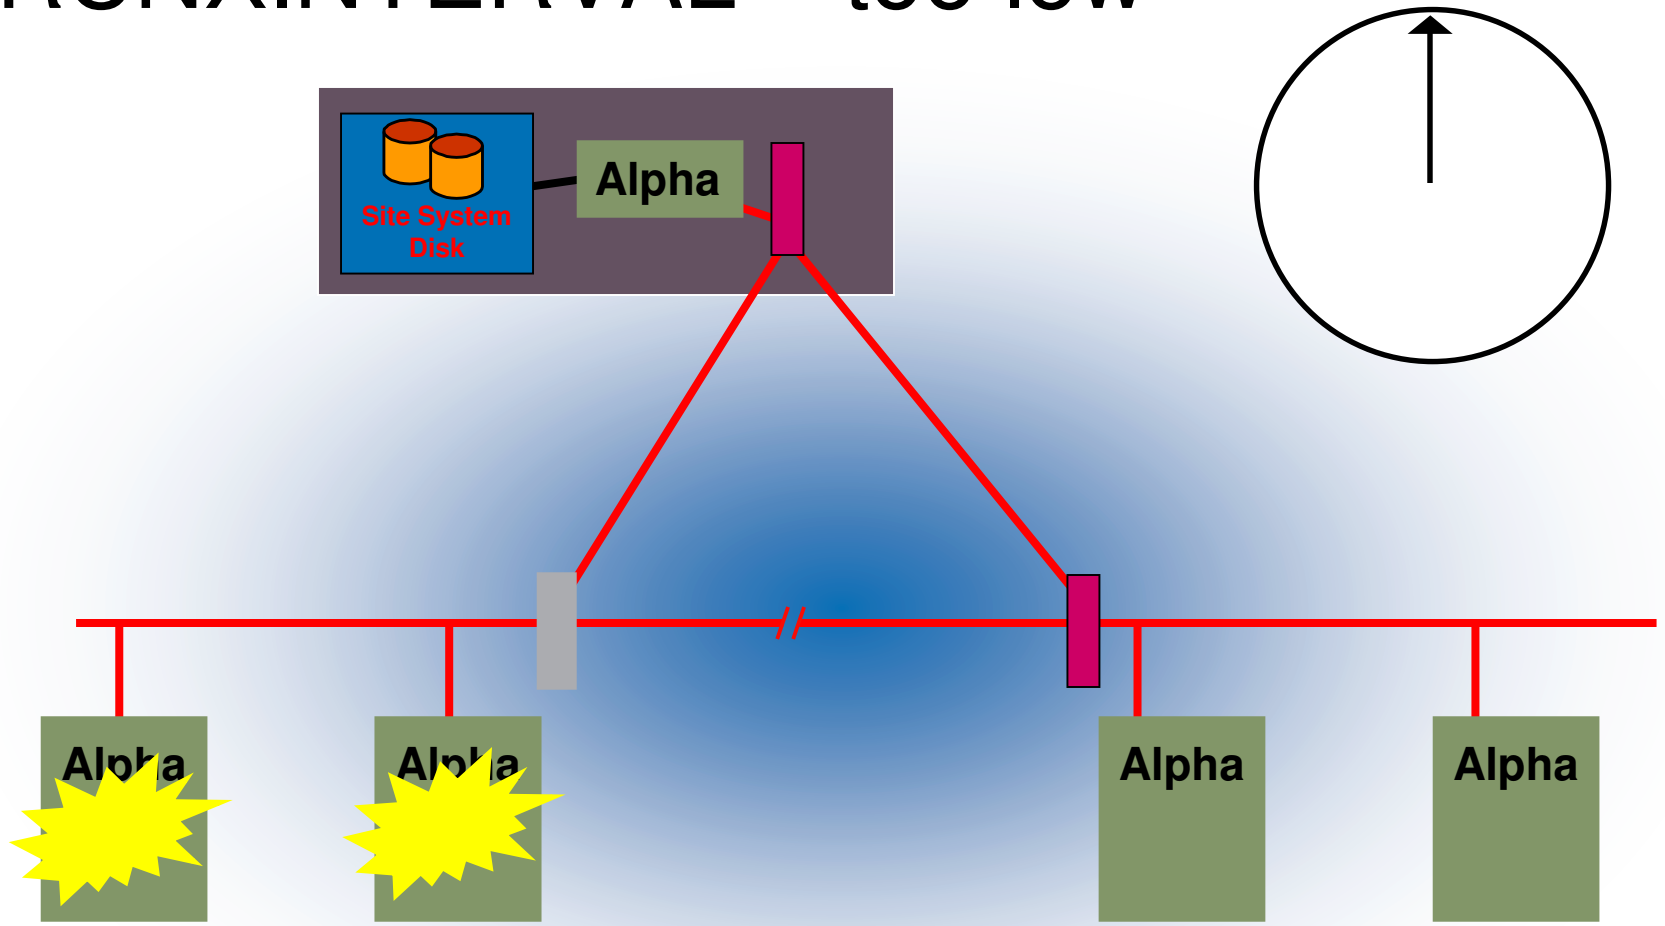

RCNXINTERVAL – set correctly

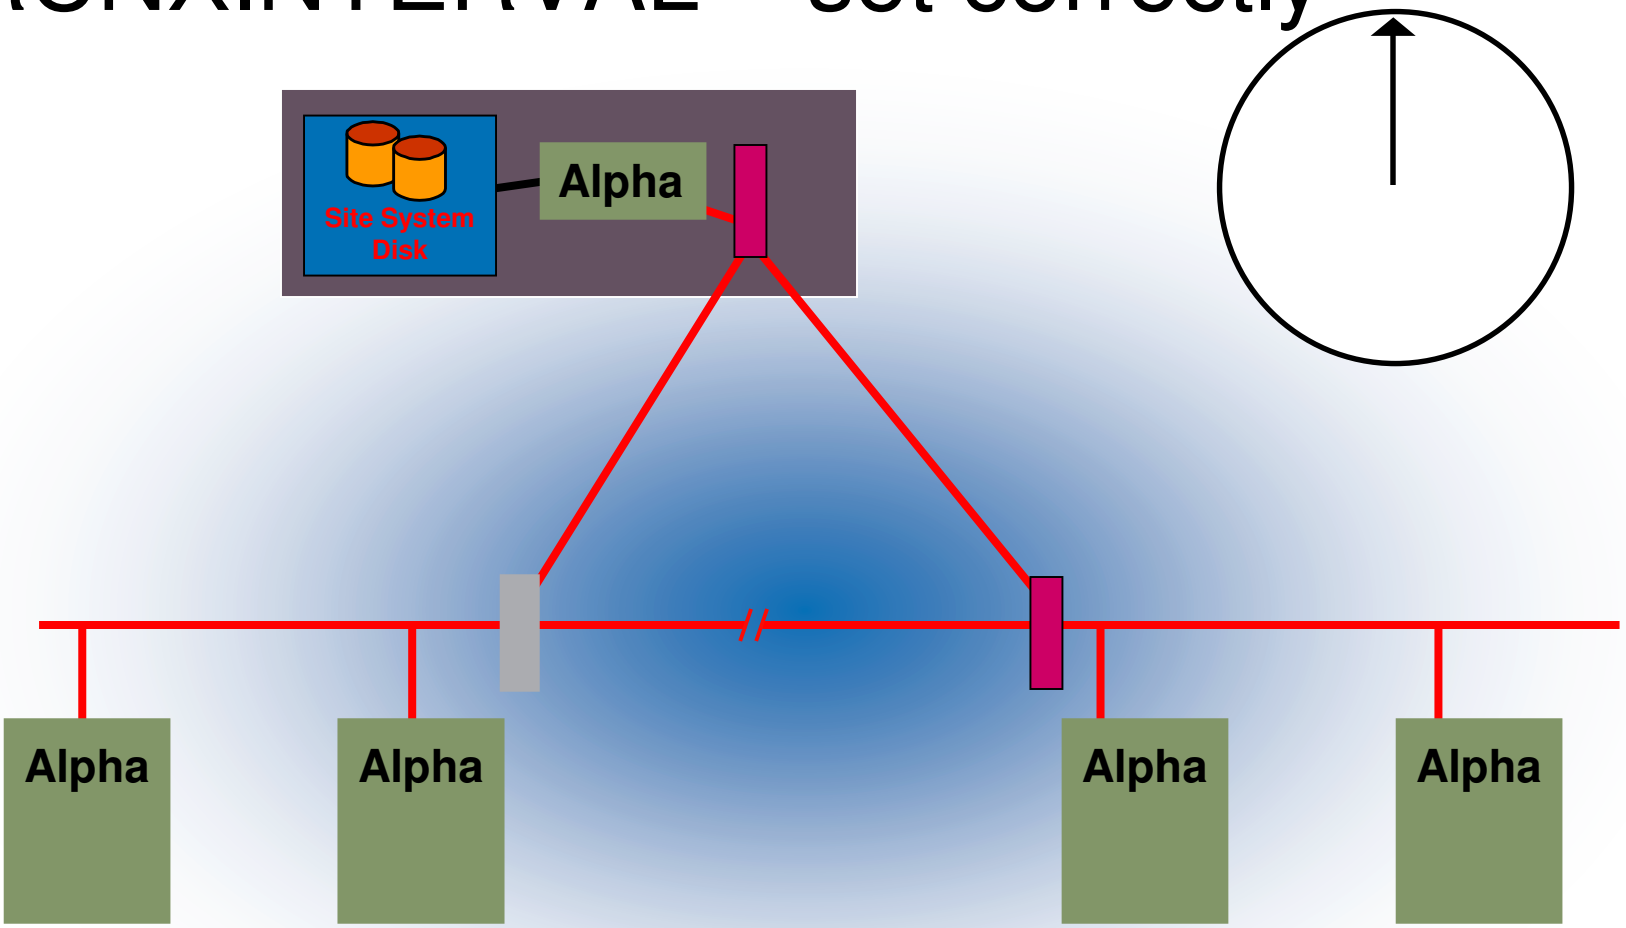

SYSGEN Parameter Tuning for DTCS - 2

- MSCP parameters
  - MSCP_CREDITS=128
  - MSCP_BUFFERS=2048 or larger

SYSGEN Parameter Tuning for DTCS - 3

- PEDRIVER parameters
  - NISCS_LOAD_PEA0 =1
    - Enables SCS on all LAN adapters
  - NISCS_PORT_SERVE = 3
    - Enables CRC checking on SCS packets
  - NISCS_MAX_PKTSZ=4468
    - (OpenVMS 6.1 & above for FDDI. Not required in OpenVMS V7.3 or above)
  - LAN_FLAGS = %x40 (Bit 6)
    - (OpenVMS V7.3 ATM Jumbo Packet Support)

SYSGEN Parameter Tuning for DTCS – 4

- Parameters/Systems tuned to cope with additional load after data centre failure

- Consider:
  - GBLPAGES
  - GBLSECTIONS
  - MAXPROCESSCNT
  - NPAGEDYN
  - Page and Swapfiles

- AUTOGEN does not always do a good job in the split-site cluster case.

Long-Distance Cluster Issues

Long-Distance Clusters

- OpenVMS officially supports distance of up to 500 miles (833 km) between nodes
- Why the limit?
  - Inter-site latency

Long-distance Cluster Issues

- Latency due to speed of light becomes significant at higher distances. Rules of thumb:
  - About 1 ms per 100 miles, one-way or
  - About 1 ms per 50 miles, round-trip latency
- Actual circuit path length can be longer than highway mileage between sites
- Latency primarily affects performance of:
  1. Remote lock operations
  2. Remote I/Os

Lock Request Latencies

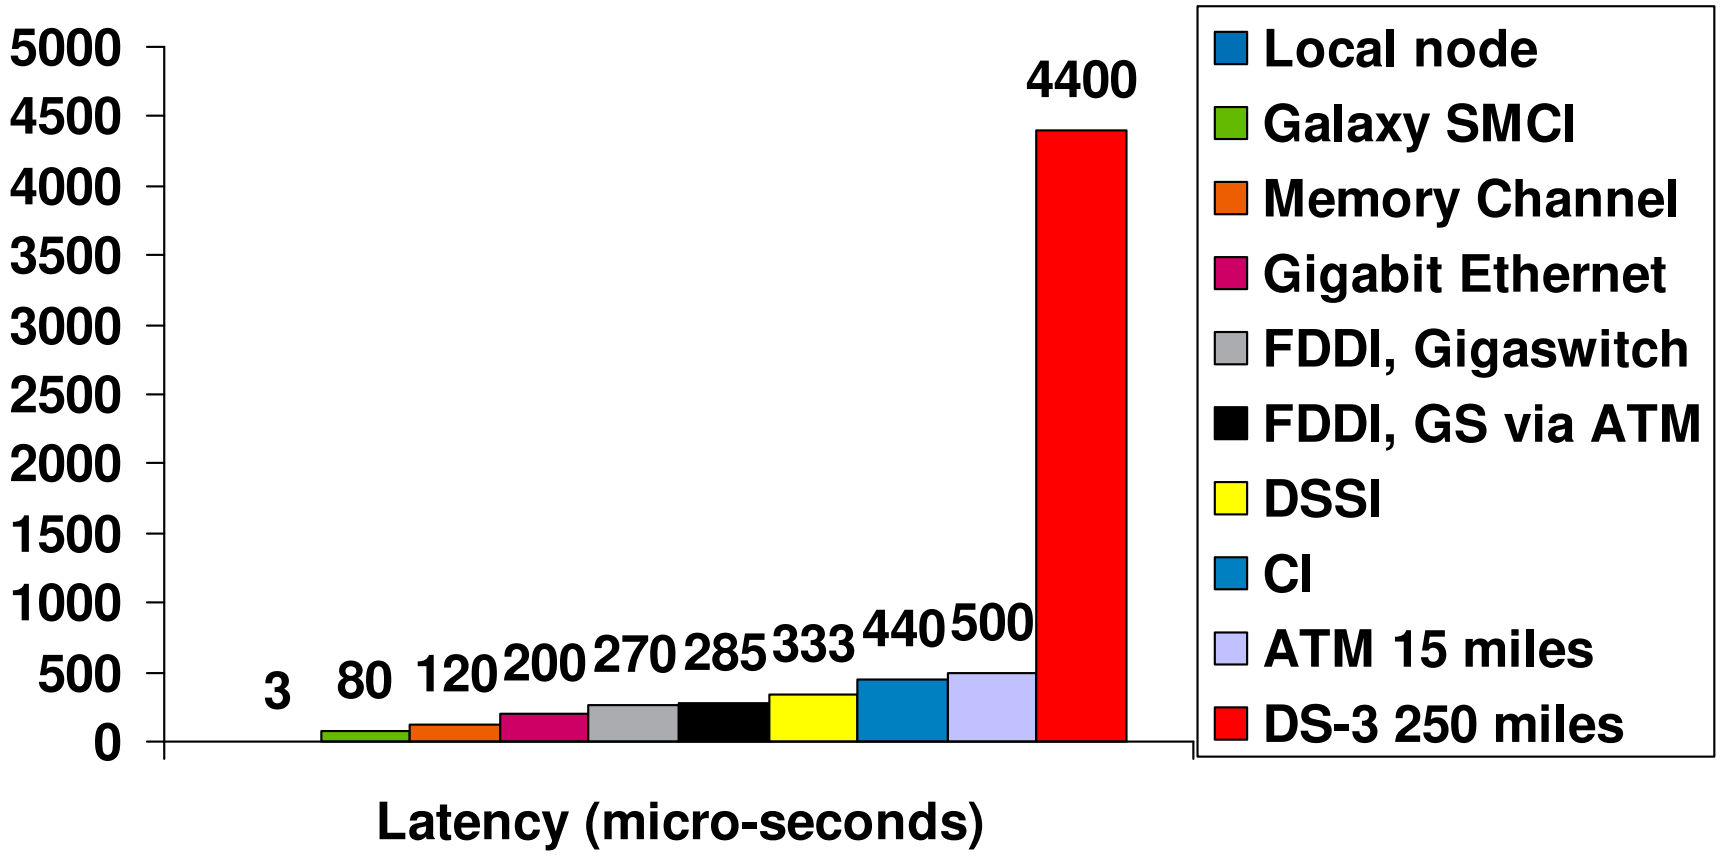

Differentiate between Latency and Bandwidth

- Can’t get around the speed of light and its latency effects over long distances
  - Higher-bandwidth link doesn’t mean lower latency

Latency of Inter-Site Link

- Latency affects performance of:
  - Lock operations that cross the inter-site link
    - Lock requests
    - Directory lookups, deadlock searches
  - Write I/Os to remote shadowset members, either:
    - Over SCS link through the OpenVMS MSCP Server on a node at the opposite site, or
    - Direct via Fibre Channel (with an inter-site FC link)

  Both MSCP and the SCSI-3 protocol used over FC take a minimum of two round trips for writes

Long Distance Issues





Synchronous replication - reading

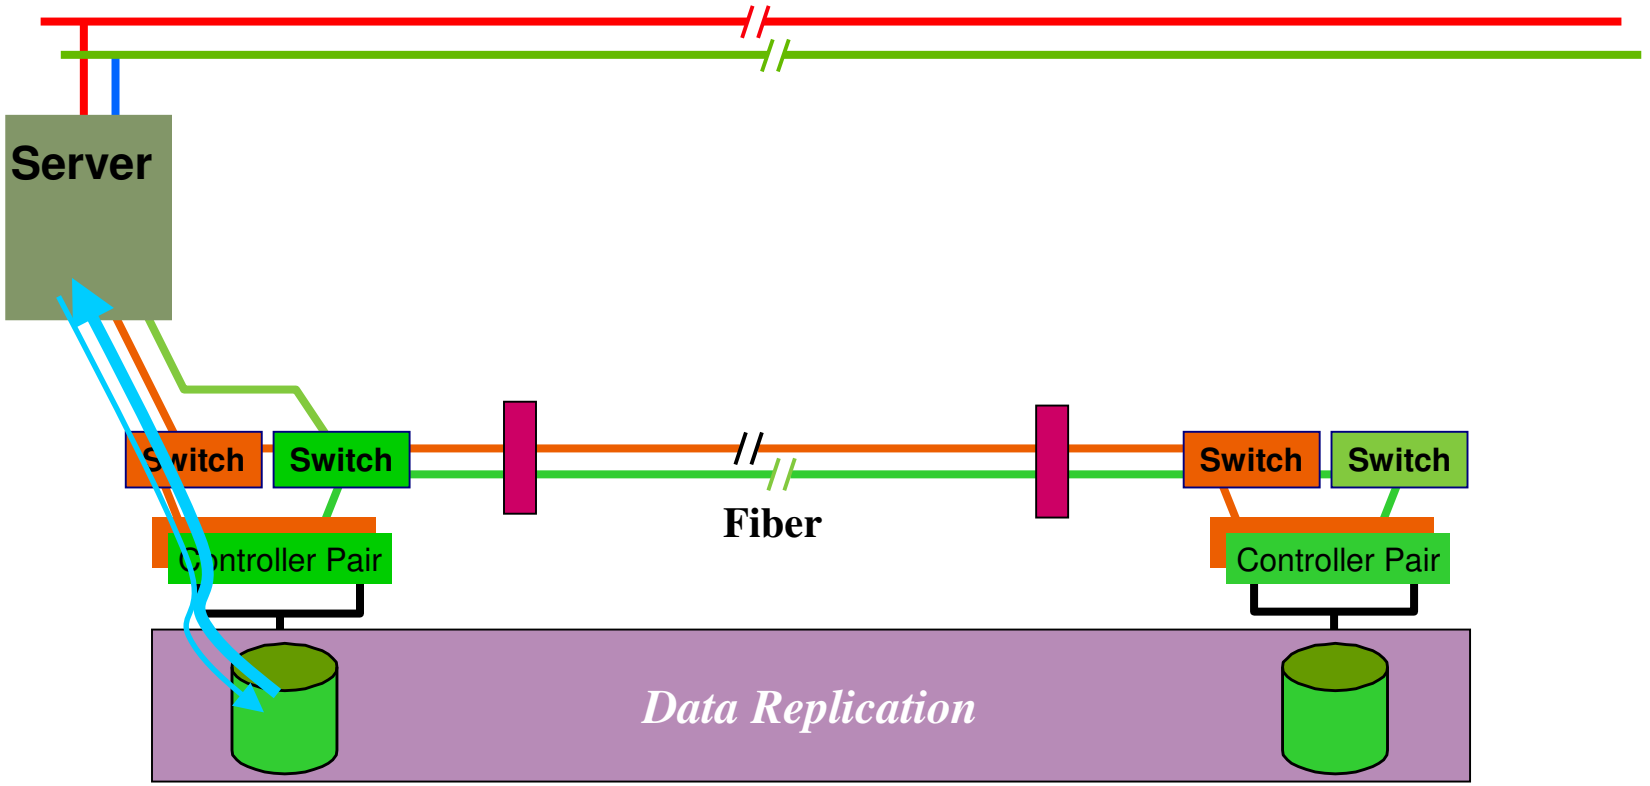

Asynchronous replication - reading

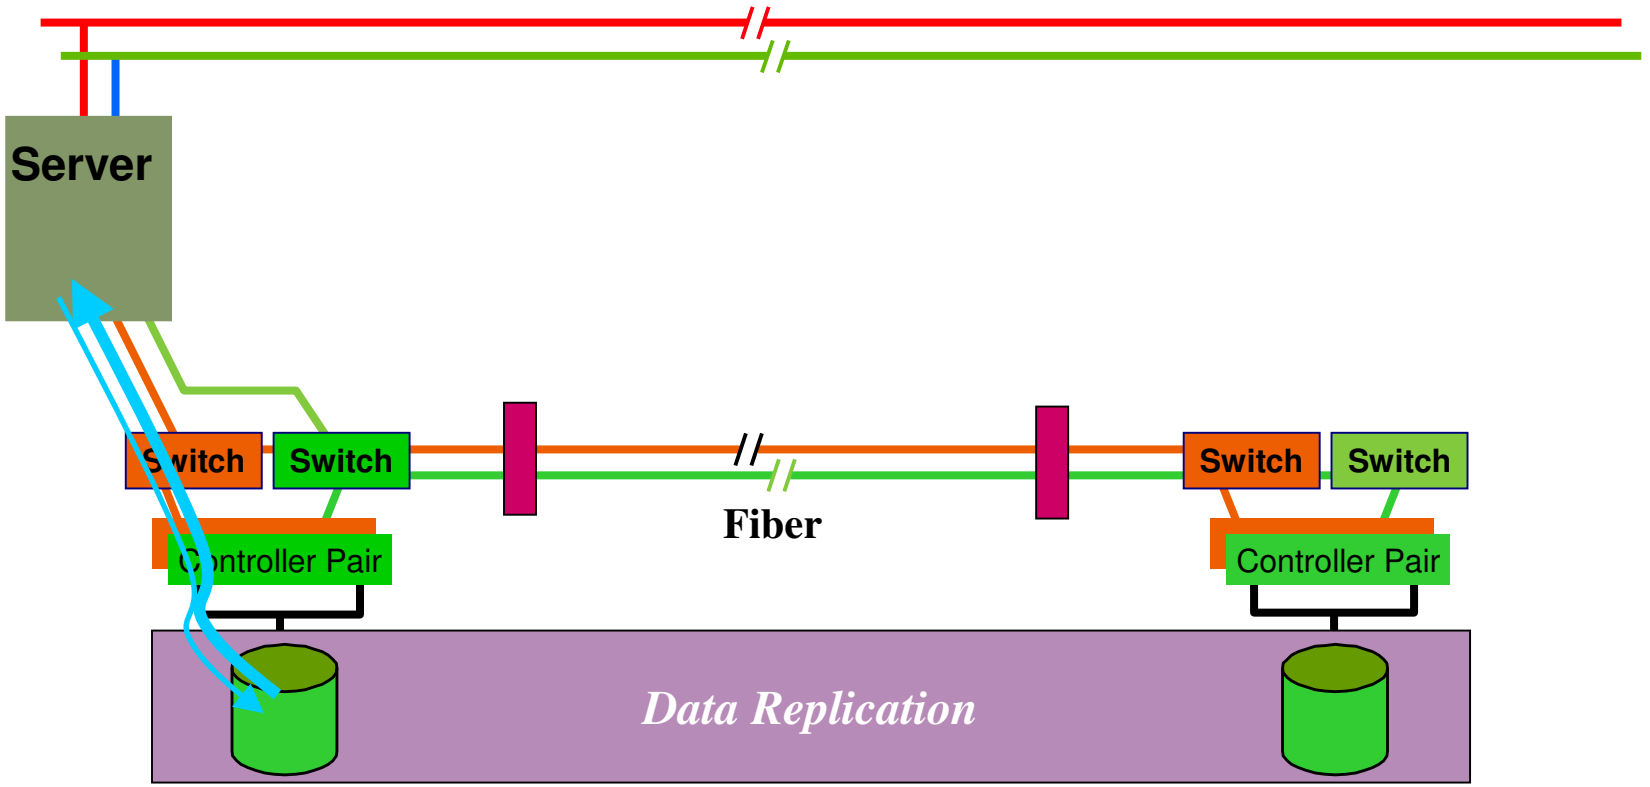

Synchronous replication - writing

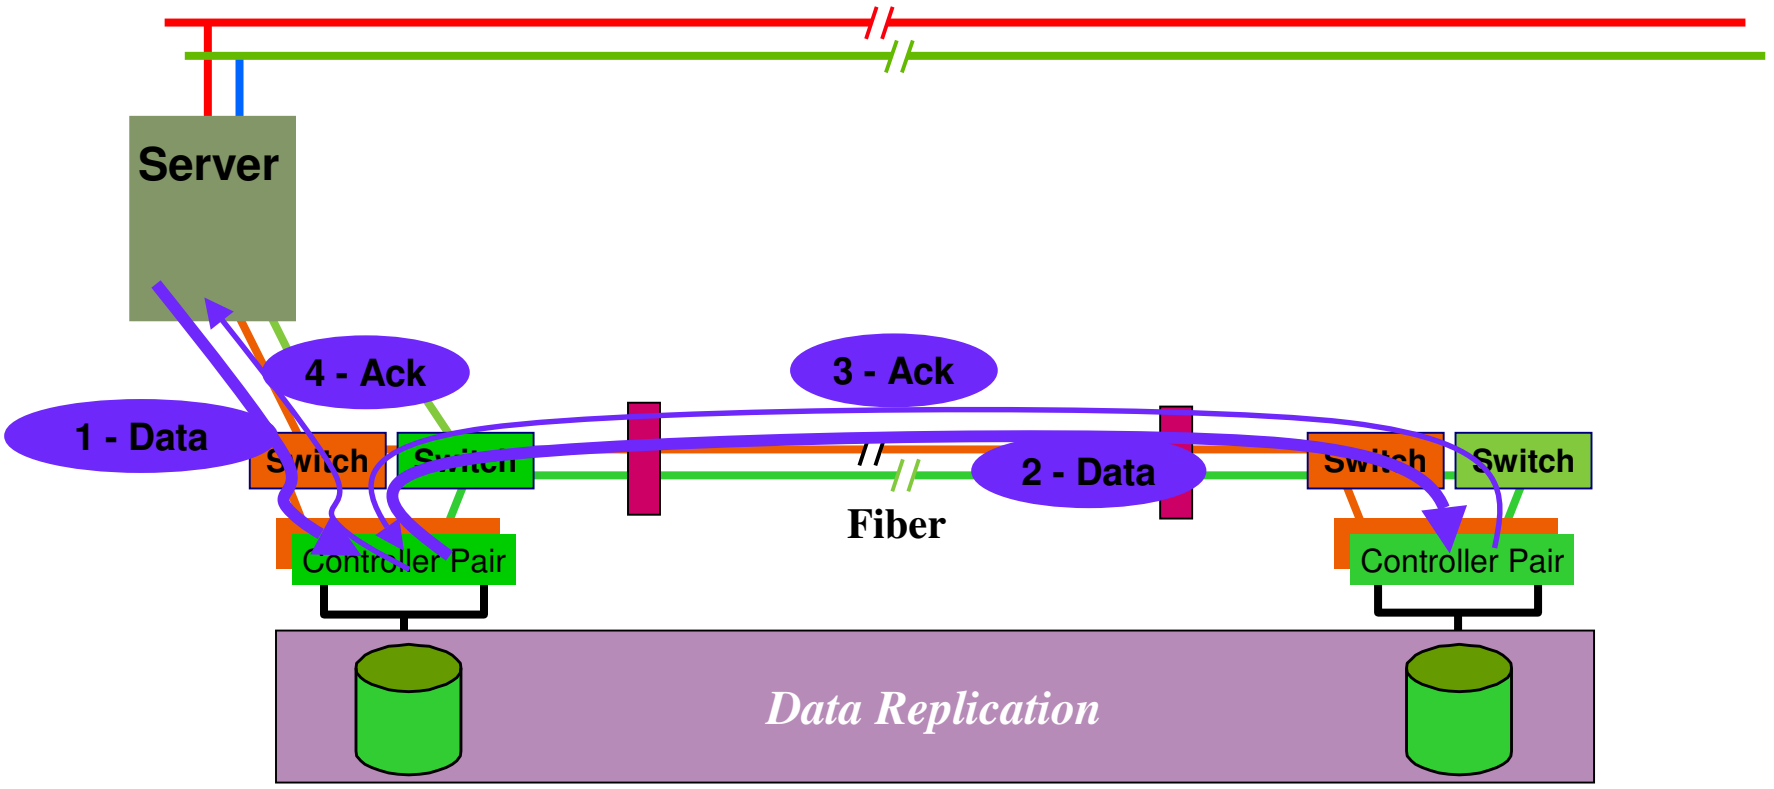

Asynchronous replication - writing

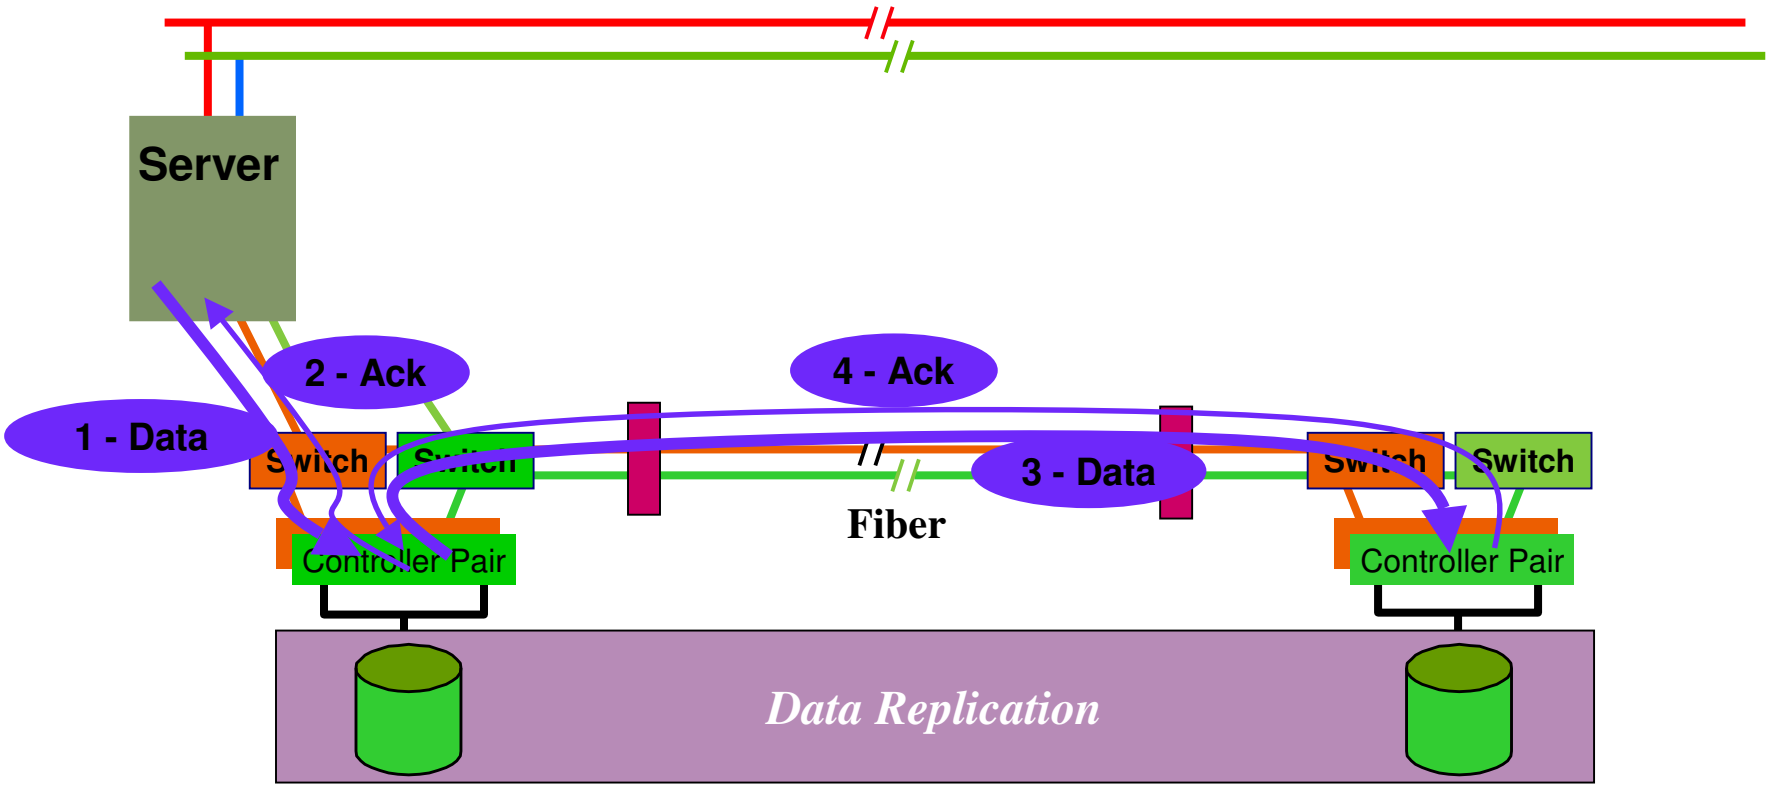

Synchronous vs Asynchronous - Perceptions

- Synchronous
  - ☺ Will guarantee no loss of data
  - ☹ More expensive to implement than asynchronous
  - ☹ Needs high bandwidth
  - ☹ Doesn't work over distance
- Asynchronous
  - ☺ Will only lose minimal amount of data
  - ☺ Cheaper than synchronous
  - ☺ Needs lower bandwidth
  - ☺ Allows unlimited distance

So what do we expect asynch to do for us

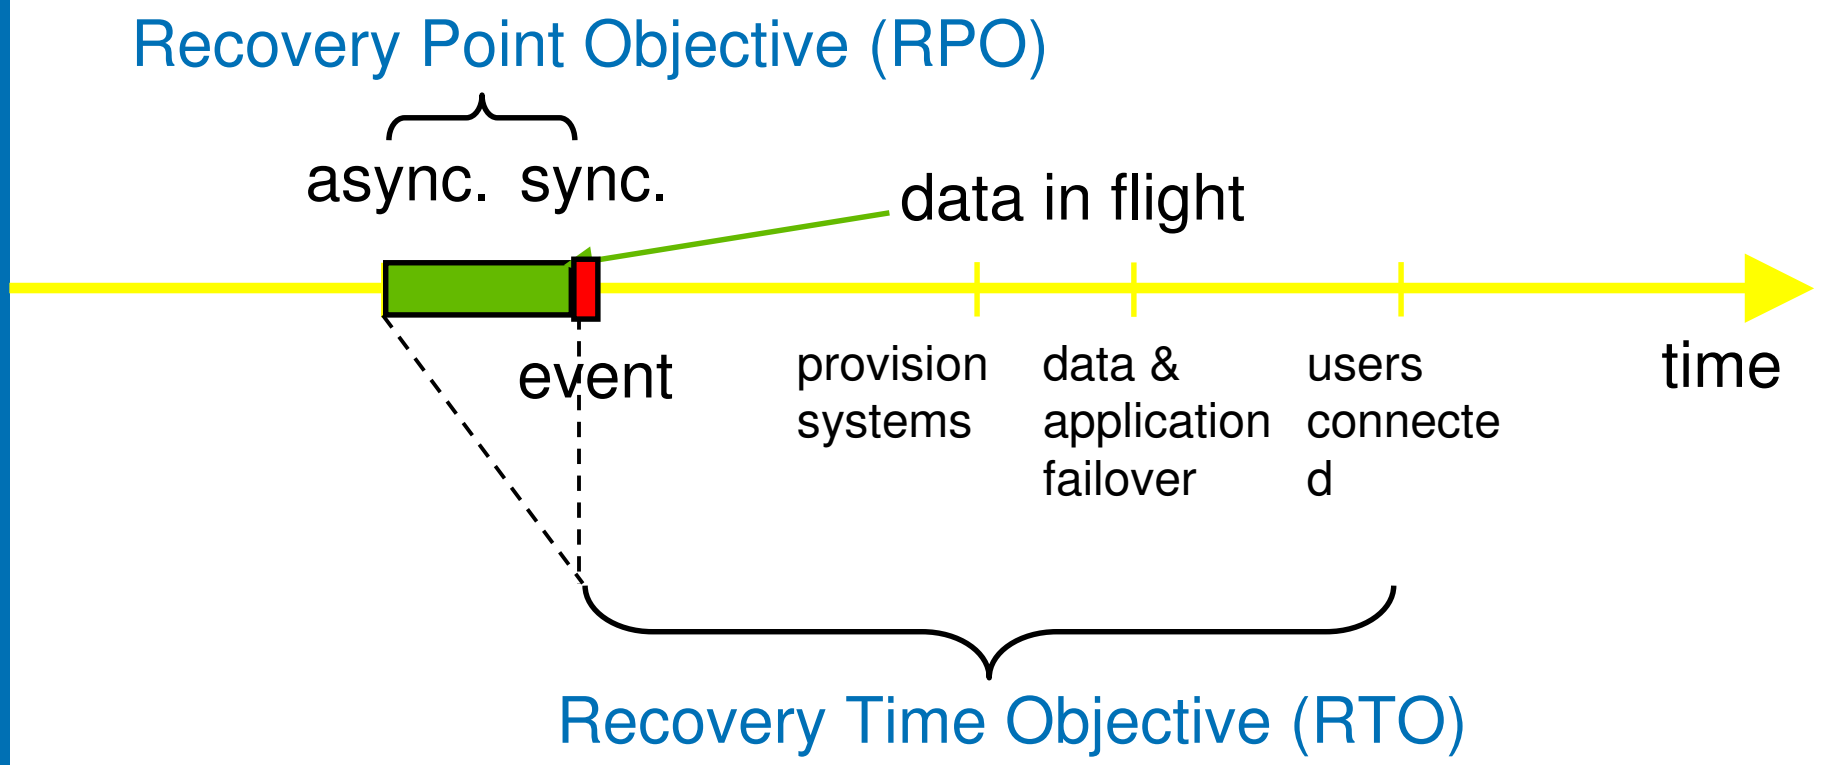

Note: not to scale

The buffering effect…..

- For a given size data block – the first bit reaches the second site at the same time irrespective of the link “speed”
- But the final bit of the packet will be later on the slower links

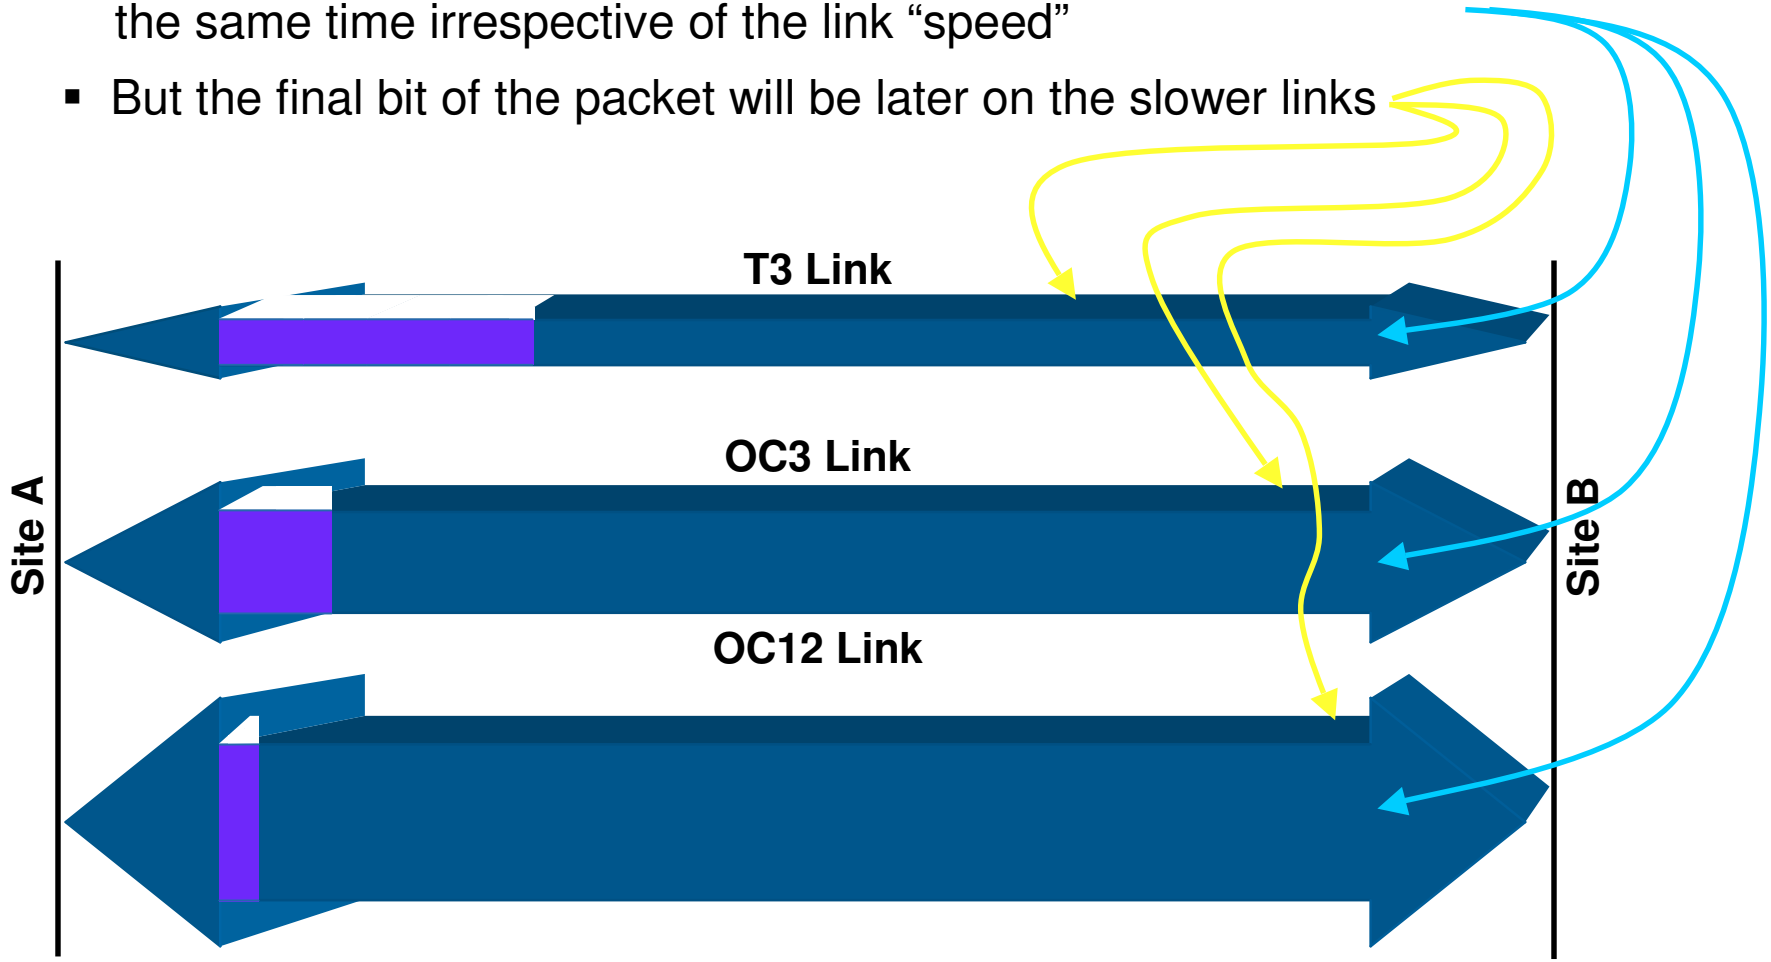

The buffering effect…..

- Consider a link with 8 buffers …….

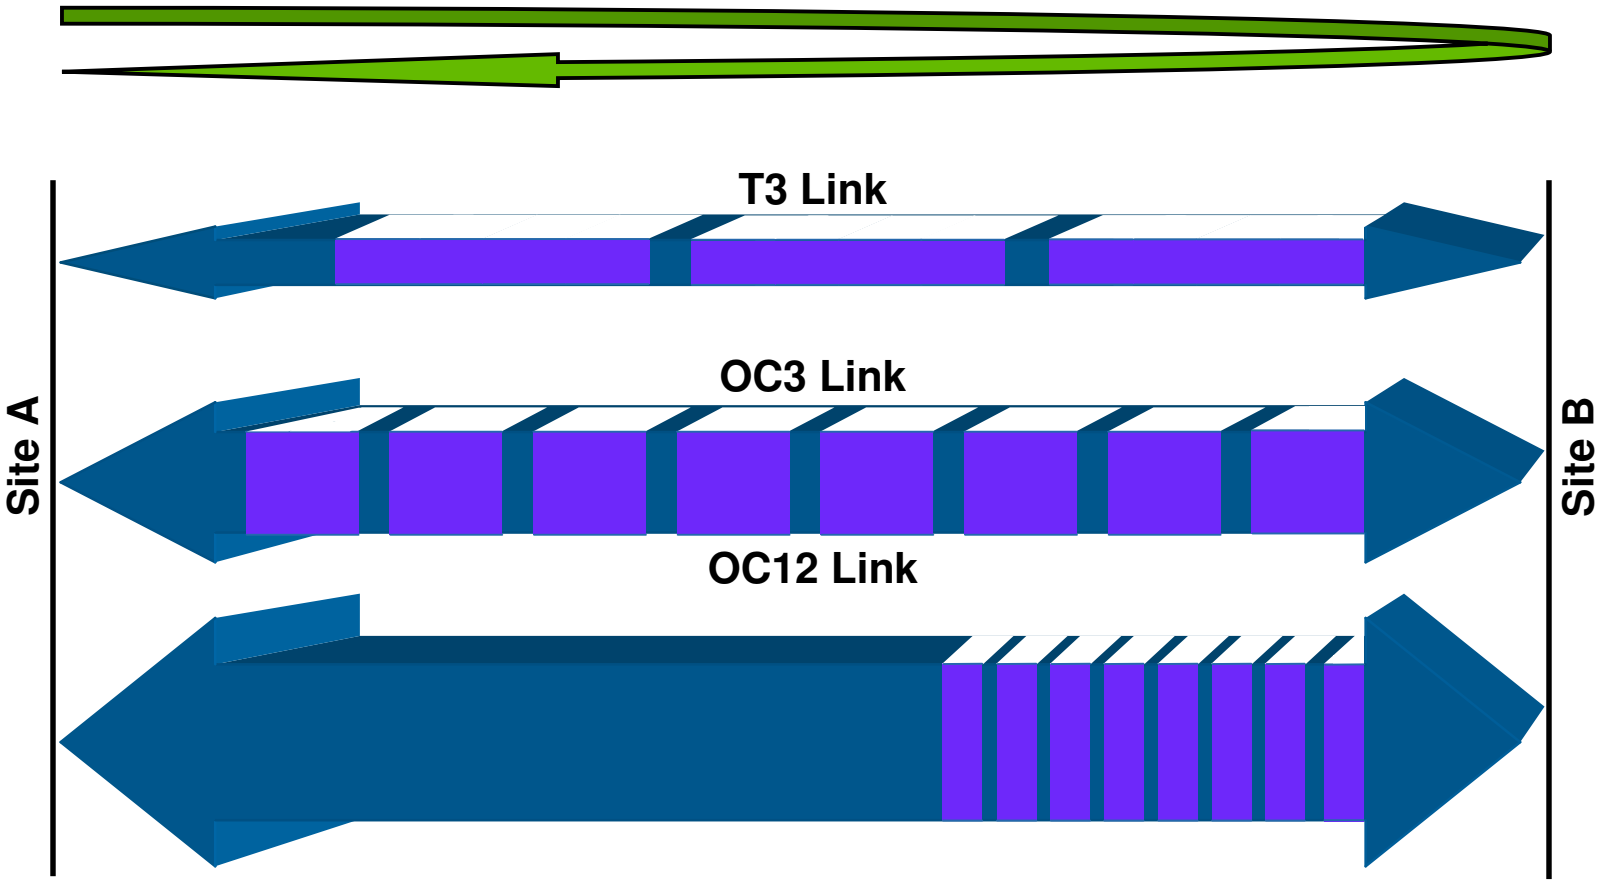

Synchronous data replication

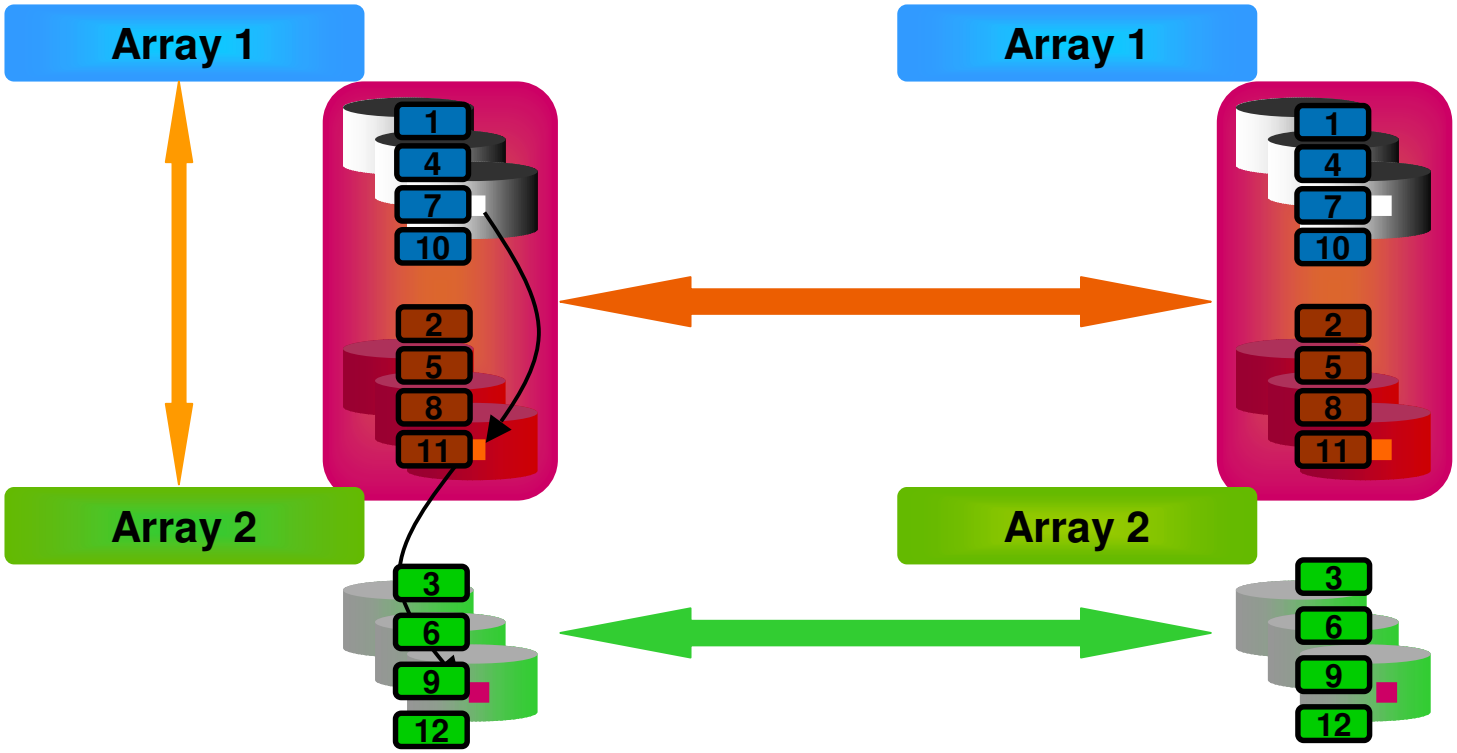

Asynchronous data replication - ideal

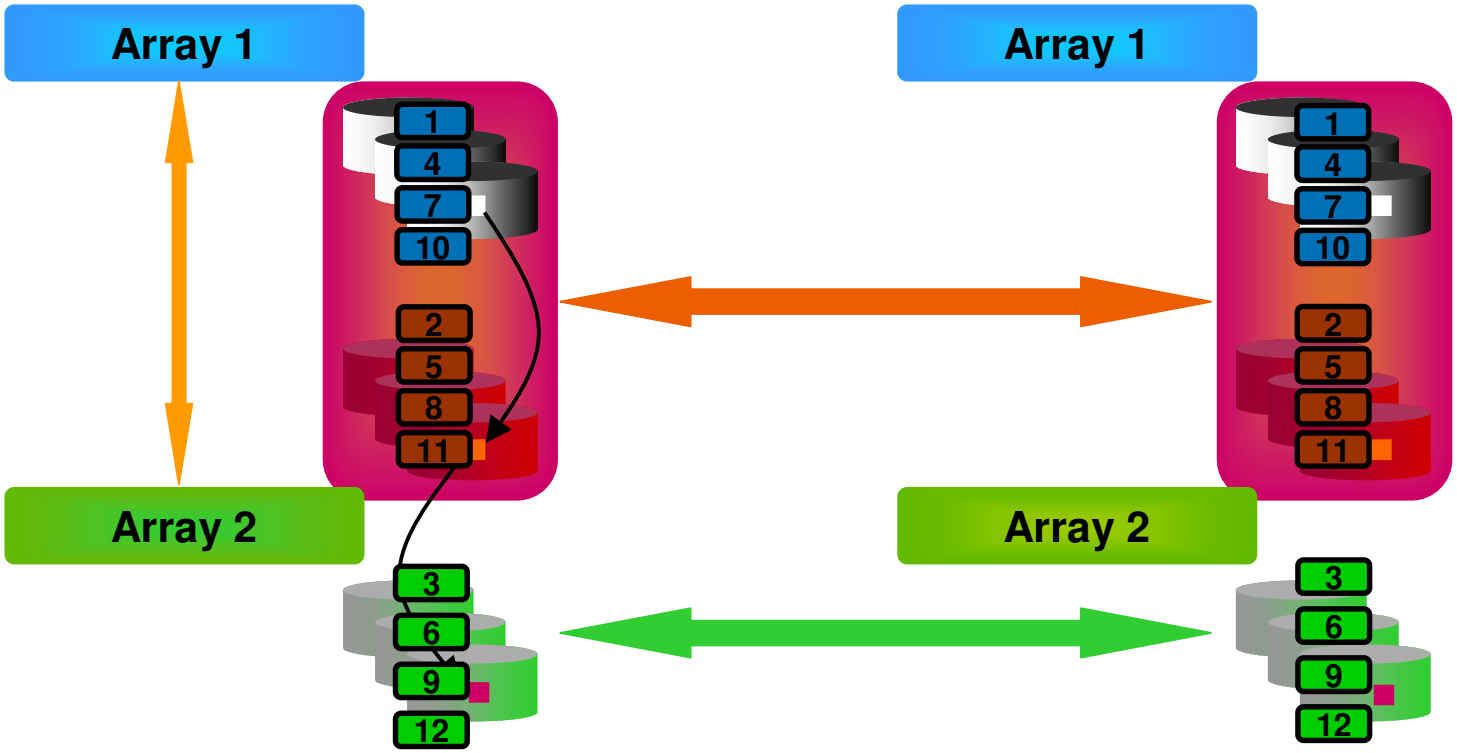

Asynchronous replication under stress