Remote Data Mirroring Solutions for High Availability

David Arrigo

EMC Corporation

508-435-1000

arrigo_david@emc.com

Remote Data Mirroring Solutions

- Agenda
  - Why remote data mirroring?
  - Physical and logical mirroring
  - Integration with clustered solutions
  - Other remote mirroring Options

Remote Data Mirroring Solutions

- Why remote data mirroring?
  - Disaster readiness for unplanned events
    - Natural disasters
      - Hurricanes, earthquakes, locust (Oracle)
    - Human error
  - Data availability for planned events
    - Upgrades
      - Operating systems and applications
    - Disaster readiness testing
      - Internal site or outsourced to service providers

Being prepared means performing readiness testing

Remote Data Mirroring Solutions

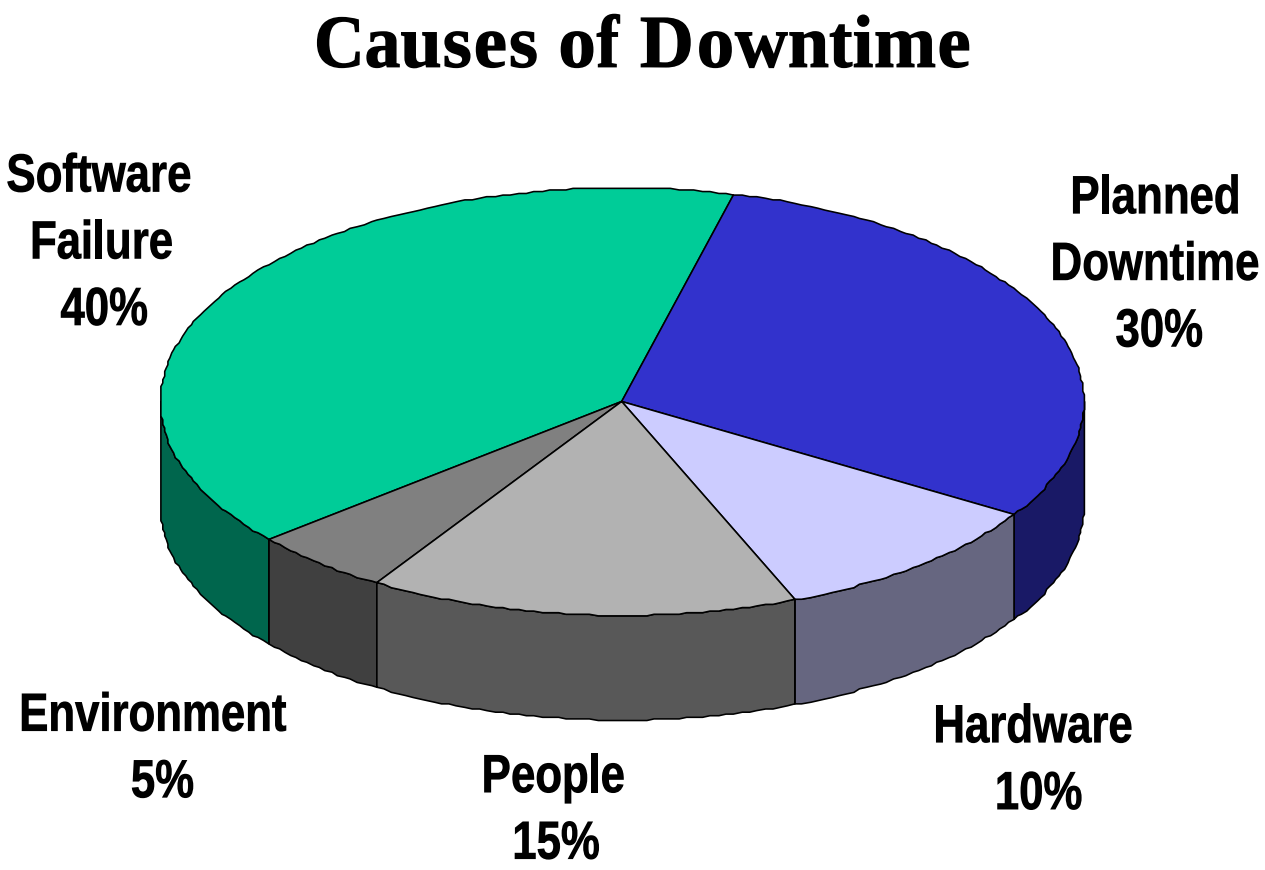

Source : IEEE Computer

Remote Data Mirroring Solutions

- Lost Revenue
- Loss of Productivity
- Damaged Reputation
- Financial Performance
- Other Expenses
  - Litigation
  - Cost of temporary employees of overtime
  - Equipment rental
  - Additional shipping costs

Remote Data Mirroring Solutions

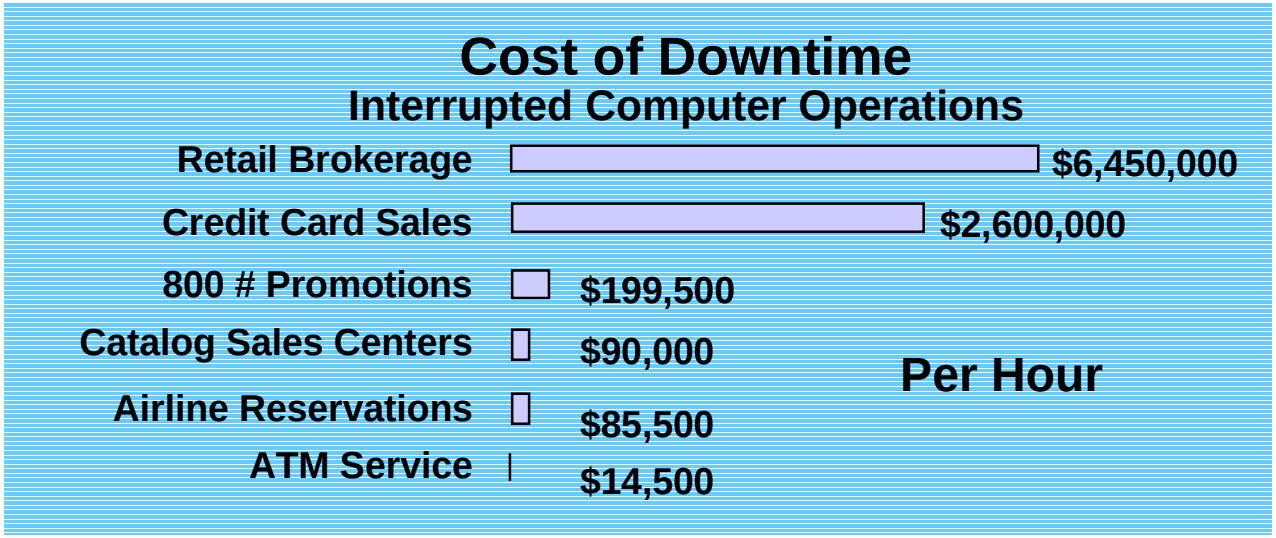

Source: Gartner Group and Contingency Planning Research

Remote Data Mirroring Solutions

Downtime Costs Add Up

- **America Online**
  August 1996 Outage: 24 hours
  Maintenance / Human Error
  Cost: $3 million in rebates
- **E*Trade**
  February 1999 through 3 March 1999 Four outages
  Cost: 22 percent stock price hit on 5 February 1999
- **eBay**
  June 1999 outage: 22 hours OS Failure
  Cost: $3 million to $5 million revenue hit
  26% decline in stock price

Remote Data Mirroring Solutions

Measuring Availability

Unavailability (minutes / year)	System Availability
50,000 (about 5 weeks)	90.0%
5,000 (About 3.5 days)	99.0%
500 (About 8 hours)	99.9%
50	99.99%
5	99.999%

Remote Data Mirroring Solutions

- Data Mirroring Solutions
  - Physical Mirroring
    - Hardware
      - » Example: EMC Symmetrix Remote Data Facility
    - Software
      - » Example: HP MirrorDisk/UX
  - Logical Mirroring
    - File System
      - » Example: Quest Shareplex/UX
    - Database
      - » Example: Oracle Advanced Replication

Each has advantages and disadvantages with respect to one another

Remote Data Mirroring Solutions
Physical Mirroring with Hardware

- Disk mirror in real time issuing a single I/O without host CPU's
- Resynchronization is performed independent of host
- Mirror operation is at the block level

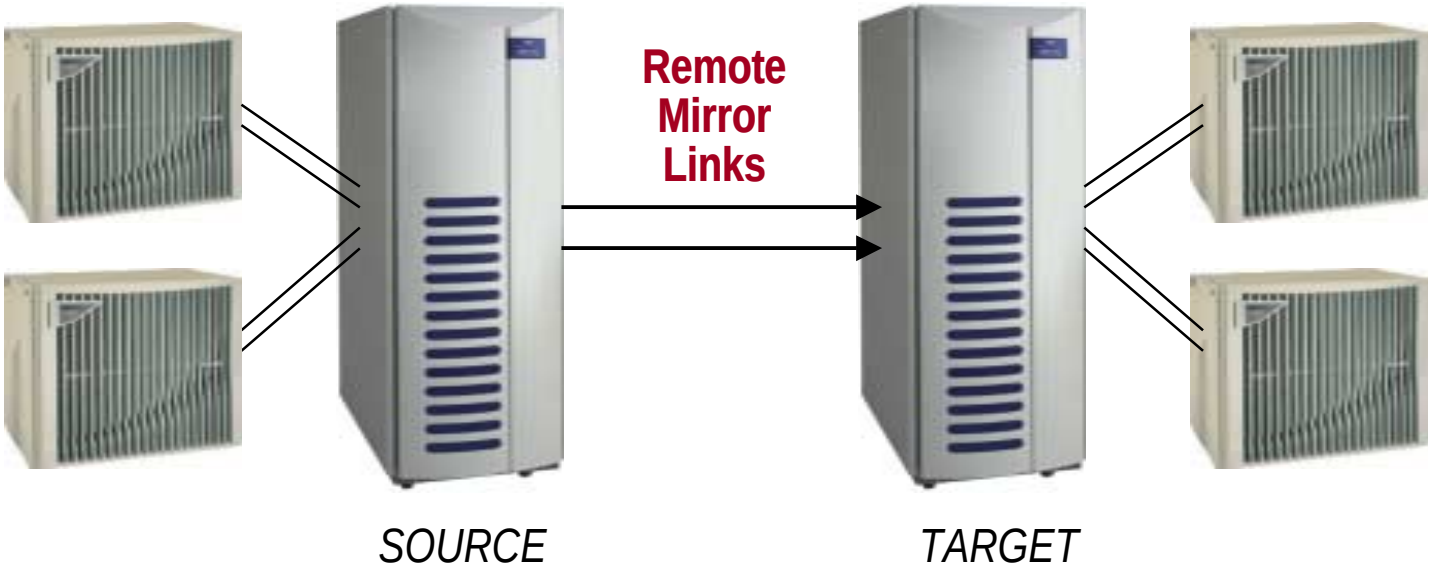

Major advantage is mirroring is not specific to a database or file system

Remote Data Mirroring Solutions
Physical Mirroring with Software

- Host CPU's required to perform mirroring operation issuing multiple I/O's
- Resynchronization requires host CPU's
- Mirror operation is at the block level

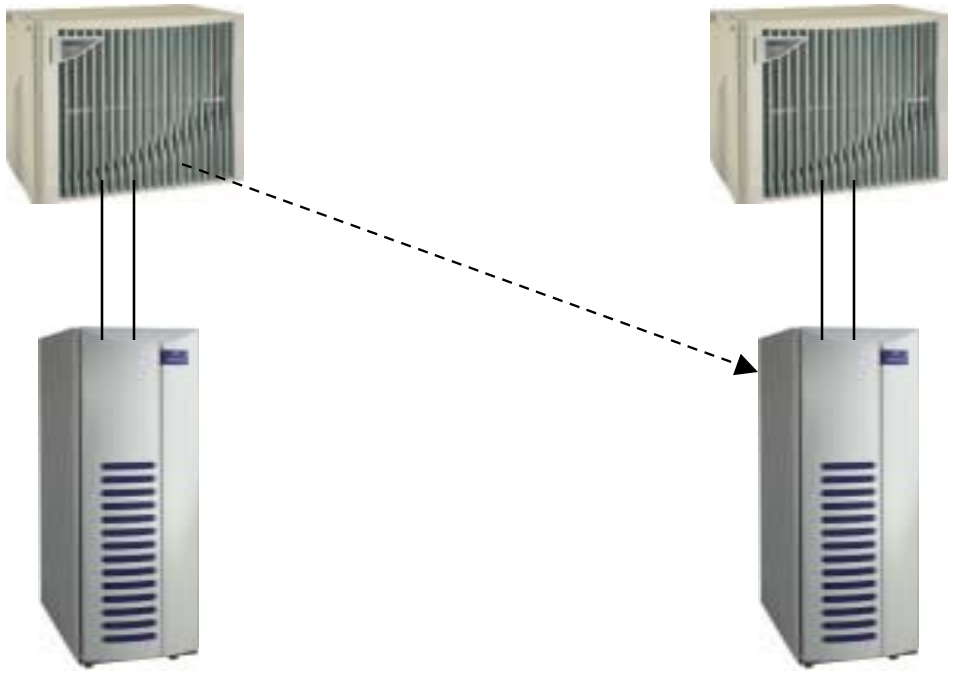

Major advantage is independence of any one vendors disk technology

Remote Data Mirroring Solutions

- Comparison of Physical Mirroring options:
  - Disk based system do not consume host CPU's
    - Single I/O issued for mirroring operation
  - Resynchronization does not consume host CPU's
    - Bit map tables maintained in storage cache vs. host memory
  - Software mirroring independent of disk technology
    - EMC or HP storage in the case of HP MirrorDisk/UX
  - Data copies are peers with software mirroring
    - May improve read performance with multiple read devices

Physical mirroring when performance, data currency, and ease of management are most important

Remote Data Mirroring Solutions
Logical Mirroring

- File system or database specific mirroring operation issuing multiple I/O's
- Host CPU's required to perform mirroring operation
- Resynchronization may require manual intervention to accomplish

Uses network to perform mirror operation

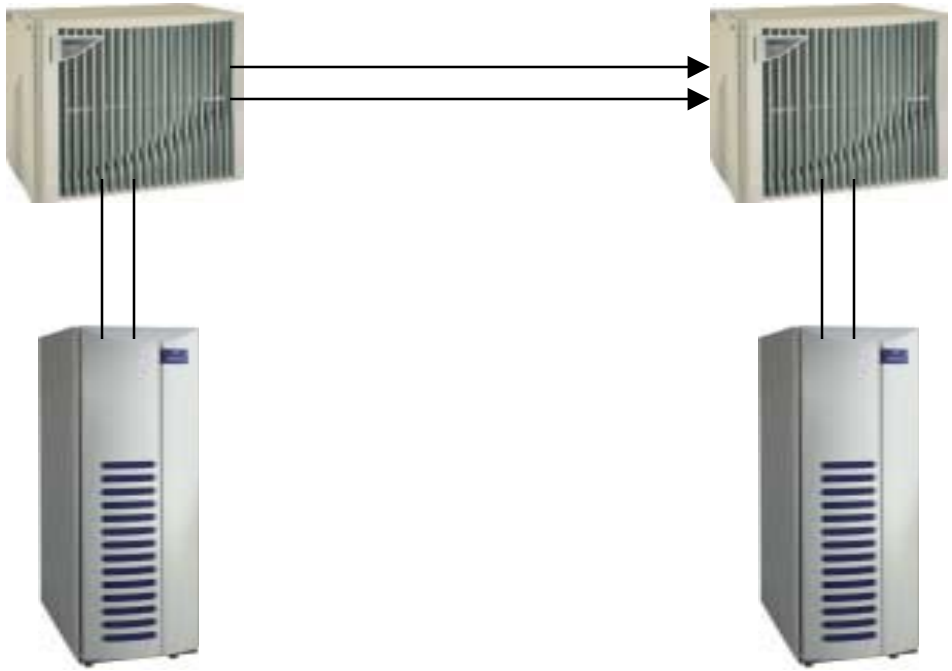

Major advantage is data corruption at remote site unlikely since transactions are mirrored

Remote Data Mirroring Solutions

- Comparison of Logical to Physical Mirroring
  - Remote data corruption less likely to occur
    - Remotely mirror transactions and not data blocks
  - Resynchronization may require manual intervention
    - Fail back usually requires manual process
  - Usually specific to a file system, database, or application
    - File System / Database mirroring or Transaction Monitor
  - Mirrors transactions and not data blocks
    - Results in lower performance

Logical mirroring when transactional consistency is most important

Remote Data Mirroring Solutions

Integrated Cluster Solutions for unattended failover

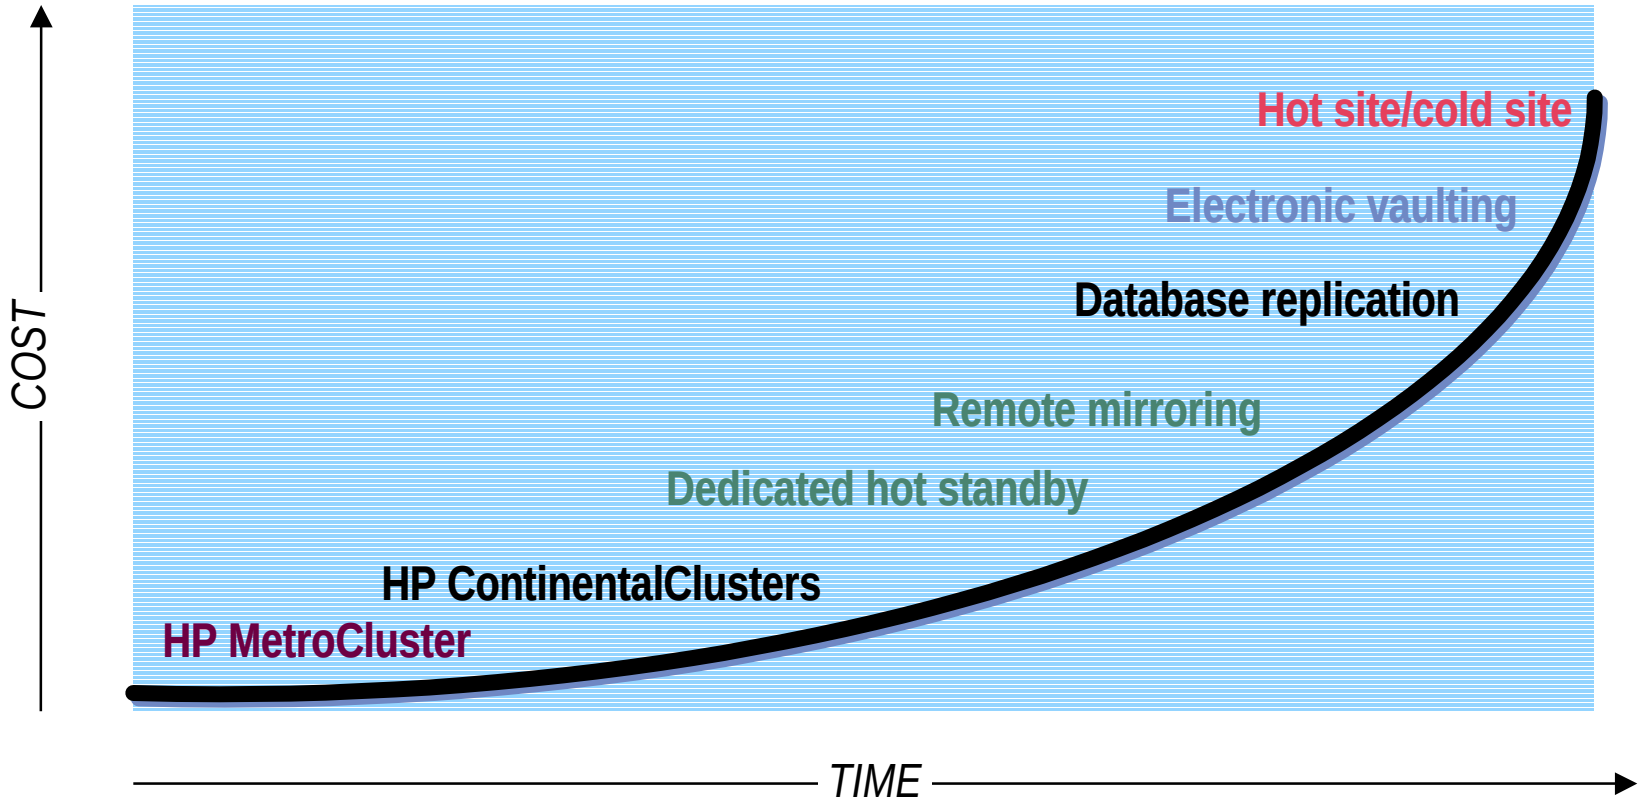

Cost of inaccessibility escalates quickly over time

Remote Data Mirroring Solutions

HP MetroCluster with EMC SRDF

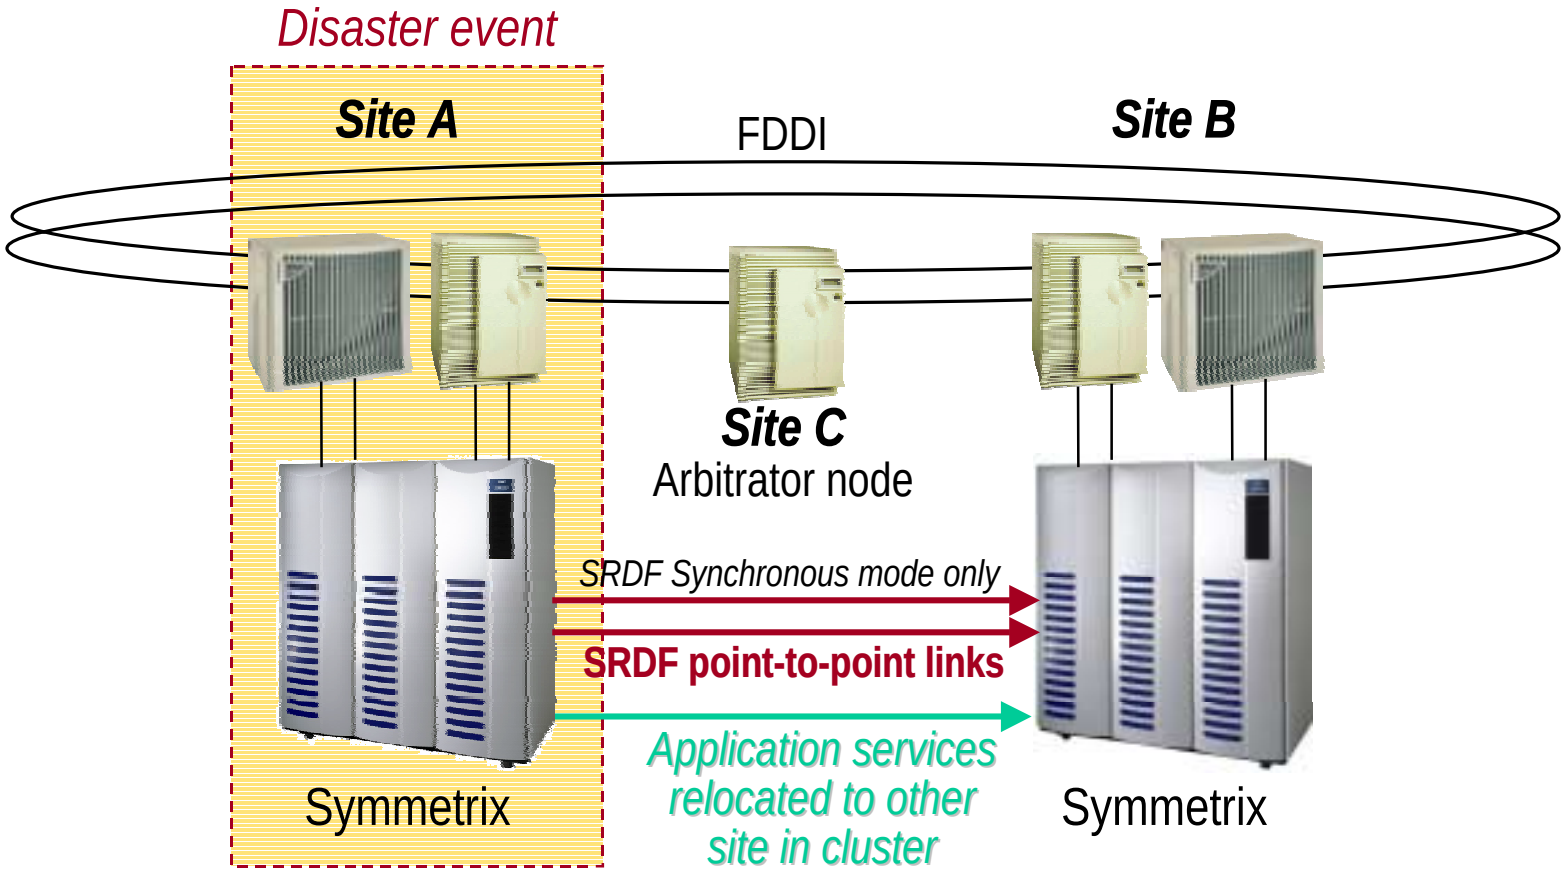

Example of an integrated cluster solution for automatic site failover

Remote Data Mirroring Solutions

HP ContinentalClusters with EMC SRDF

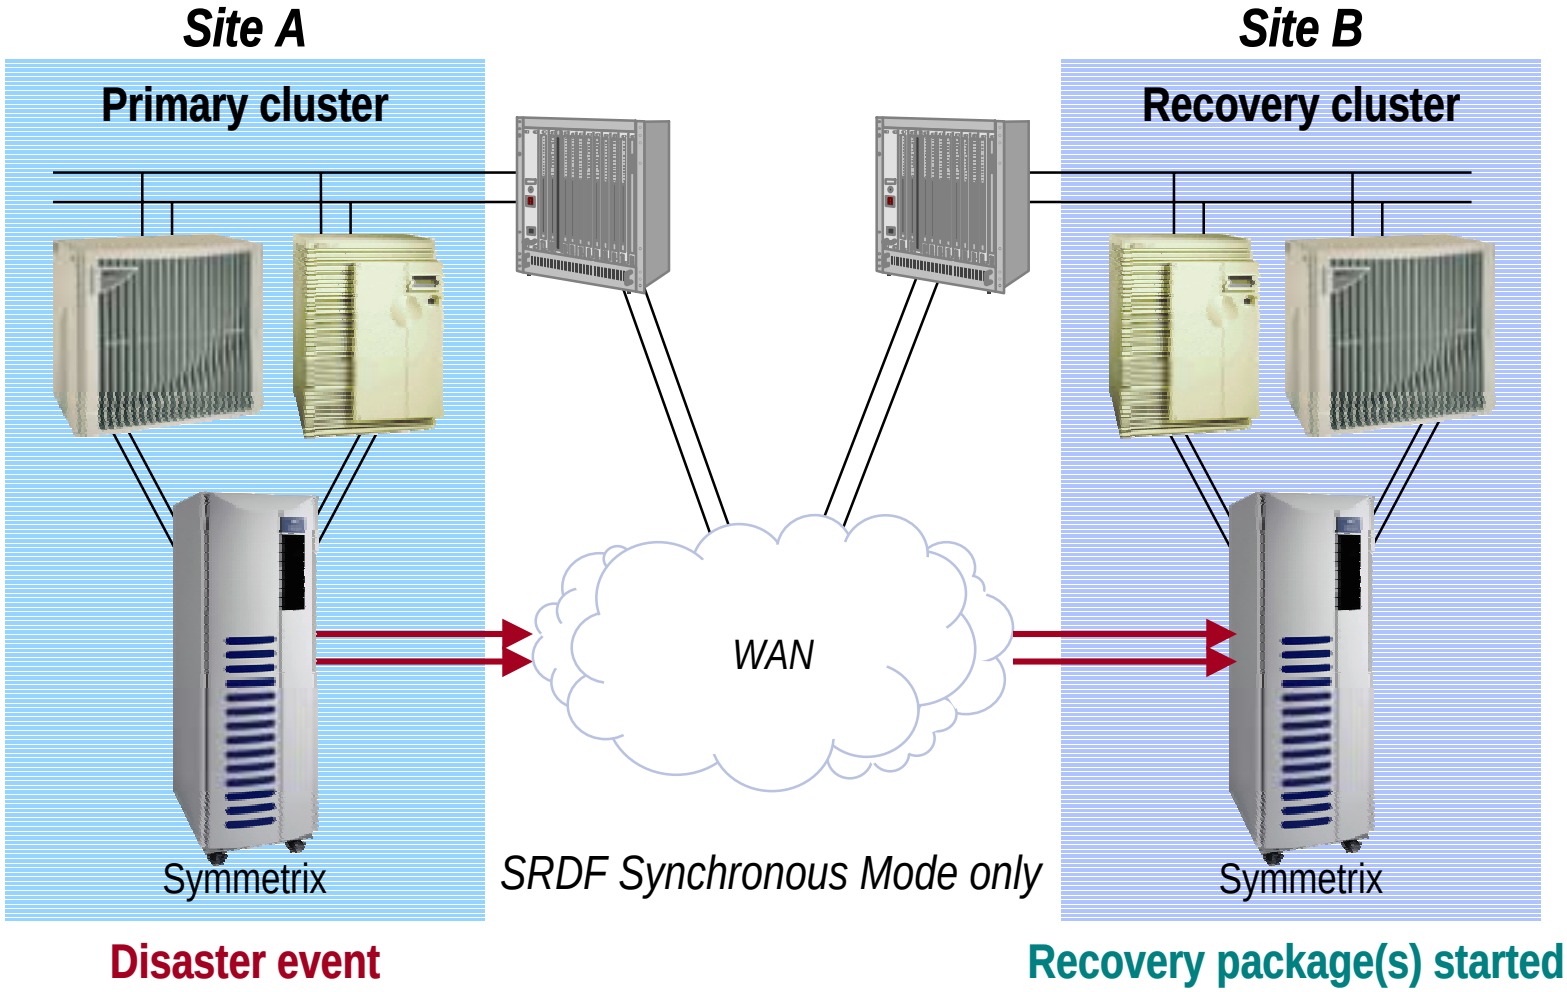

Example of an integrated cluster solution for semi-automatic site failover

Remote Data Mirroring Solutions

HP Campus Cluster using HP MirrorDisk/UX

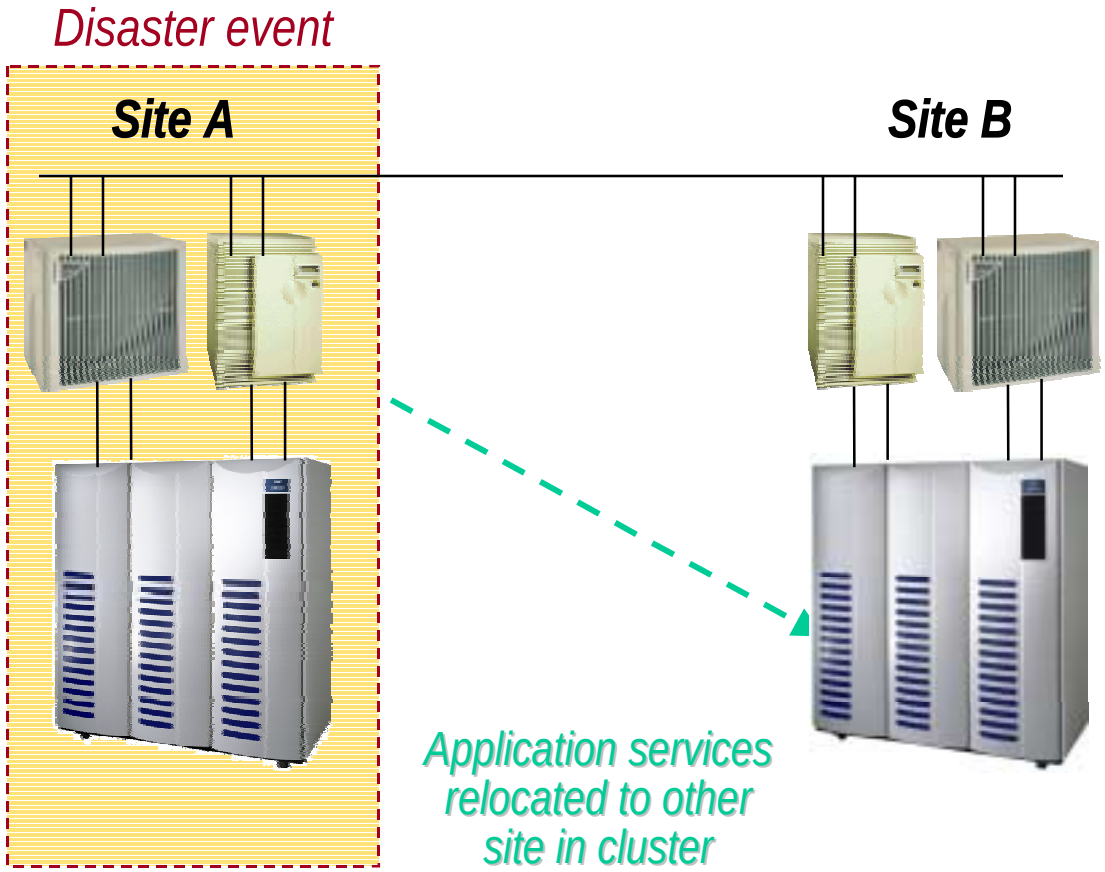

Example of an integrated cluster solution for automatic site failover

Remote Data Mirroring Solutions

- Advantages of automatic and semi-automatic site failover solutions
  - Rapid site recovery with no manual intervention
    - Not prone to human error during recovery process
  - Downtime avoided during off-hours periods
    - Middle of the night events in which there is minimal staffing
  - Integrated, tested, and supported solution
    - Engineered for end-user environment
  - Distances beyond that of a single datacenter
    - Tolerances beyond a single site or campus environment

Disaster tolerant solution to meet minimal downtime requirements

Remote Data Mirroring Solutions

- Important considerations when choosing a remote mirroring solution:
  - Synchronous or asynchronous operation
    - Importance of data currency
      - Requires currency up to the last committed transaction?
      - Tolerances to some data loss?
  - Support for failback process
    - Manageable resynchronization process
      - Full-copy or changed tracks / blocks
      - Ability to maintain changed data information if second fault event occurs
  - Recoverability of data at the remote location
    - Ability to roll forward committed and rollback uncommitted transactions
      - Available with physical and logical mirroring solutions
      - Use of non-synchronous mirroring may result in data loss

Remote Data Mirroring Solutions

- Other remote data mirroring options:
  - Point-in-time copies
    - Remotely mirror copies of point-in-time data
      - Addresses network costs since mirror is point-in-time
        - » Requires less network bandwidth since not real-time
      - Addresses I/O latency issues
        - » Extended distance environments
  - Database Redo-Log Mirroring
    - Remotely mirror redo-log files only
      - Addresses network costs as it requires less bandwidth

This can be most cost-effective approach for extended distance environments

Remote Data Mirroring Solutions

HP ContinentalClusters with EMC SRDF

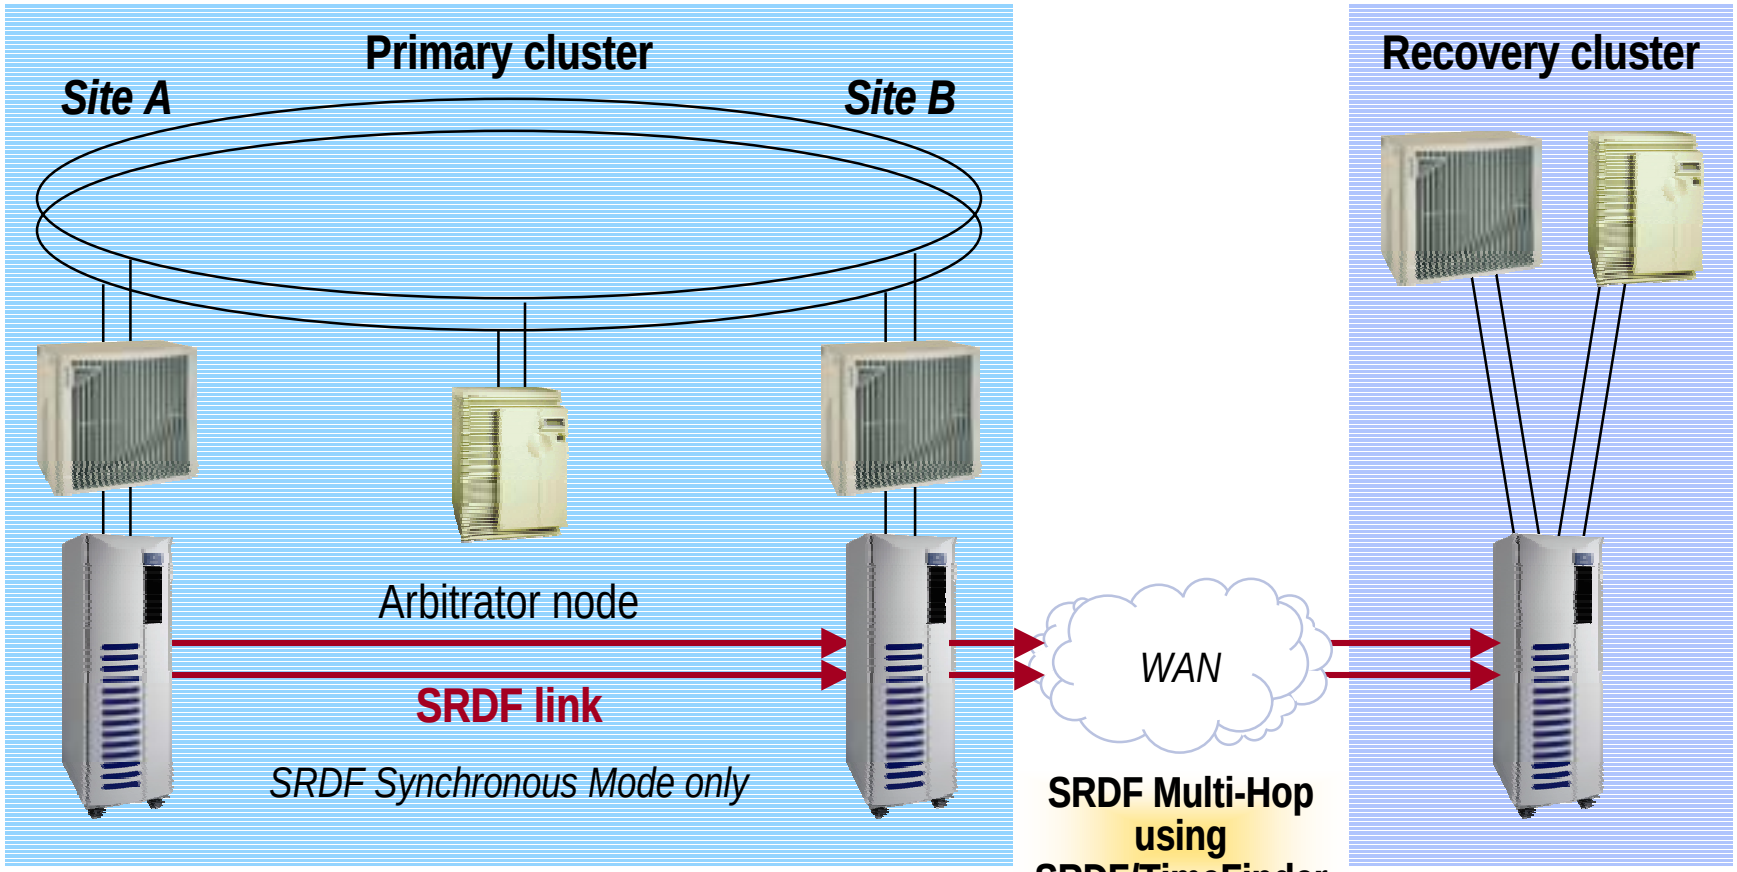

Failover between HP MetroCluster nodes

Example of using point-in-time mirroring to address network costs and mirror I/O delay

Remote Data Mirroring Solutions

How Point-in-Time Remote Mirroring Works

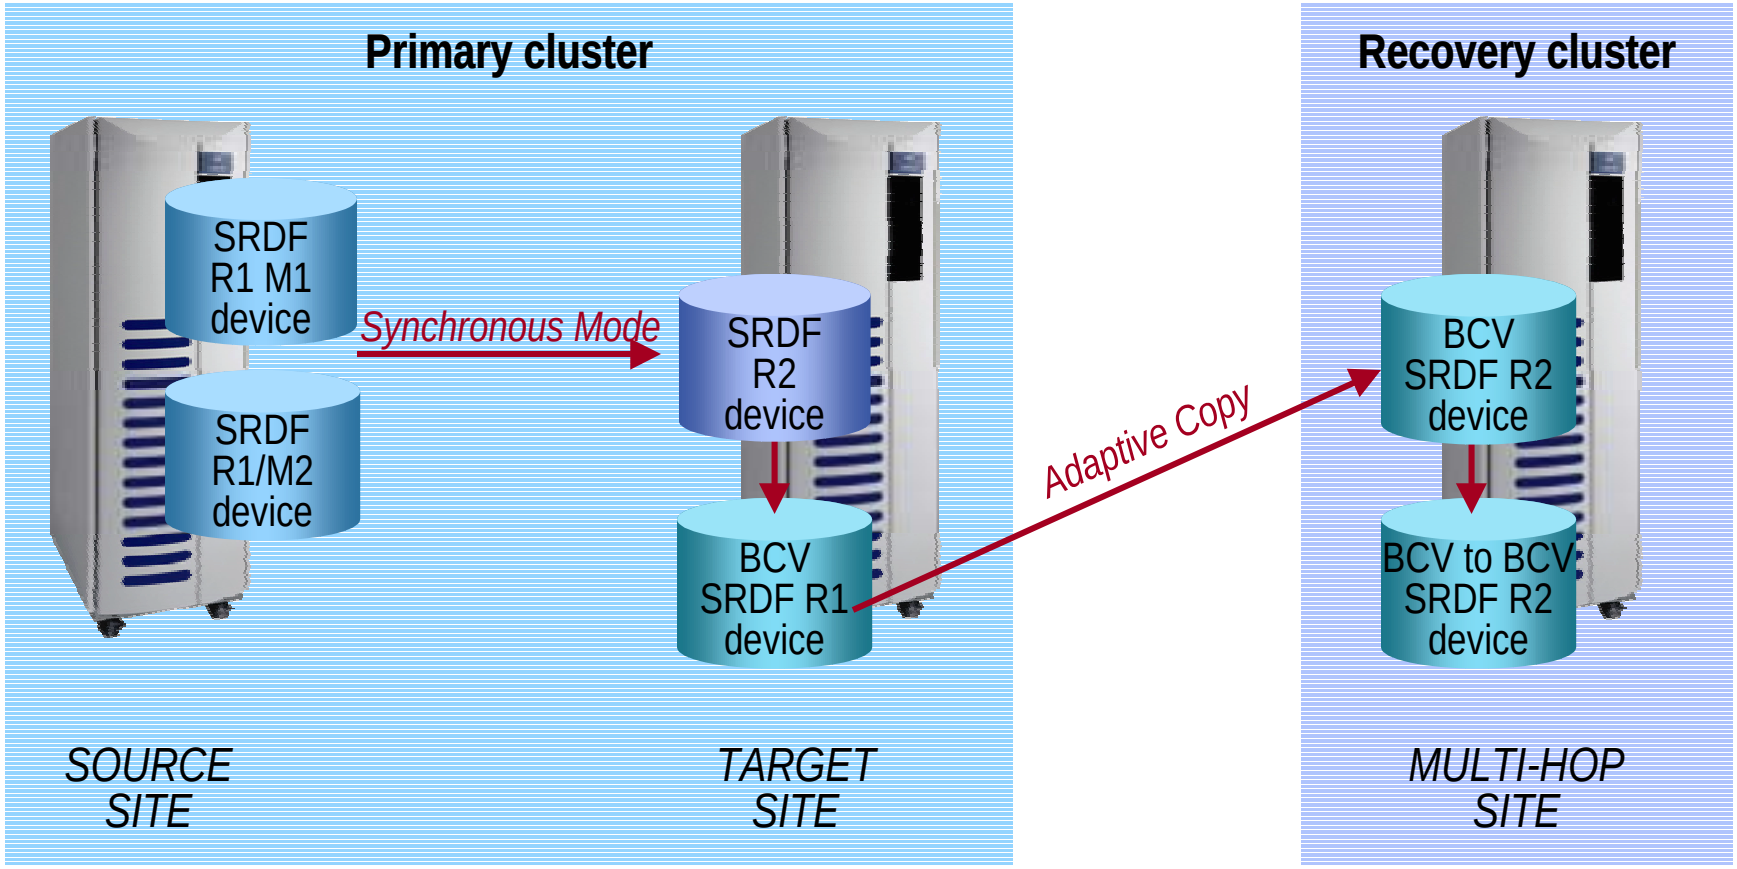

Minimize extended distance network requirements

Remote Data Mirroring Solutions

How to resynchronize using Point-in-Time for Remote Mirroring

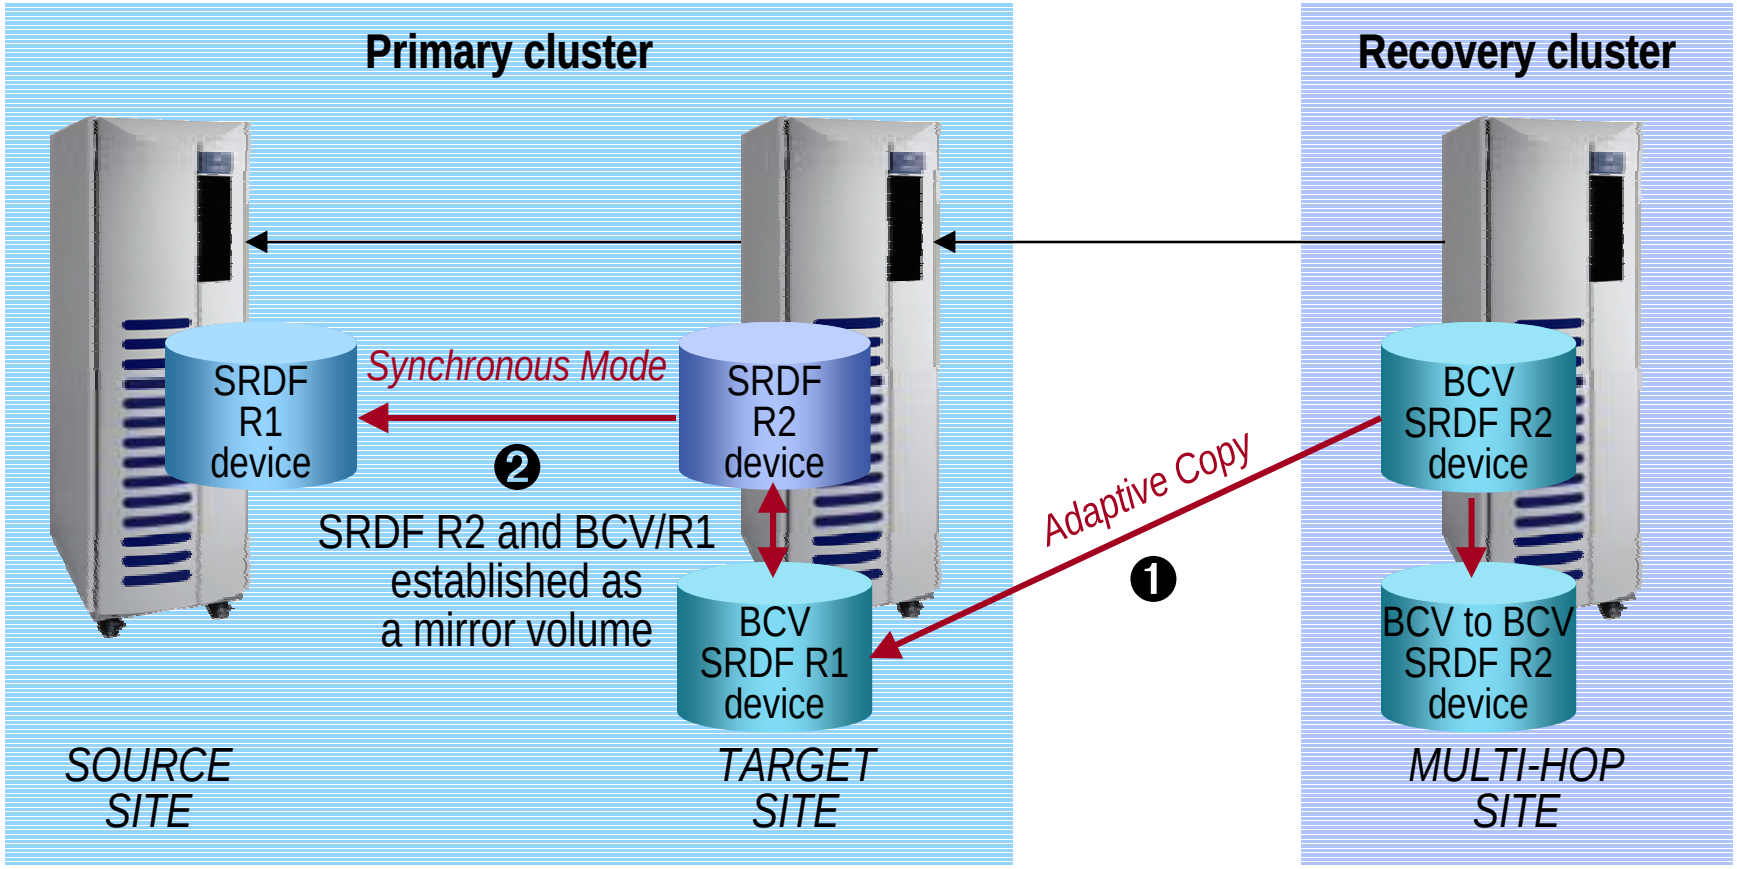

Requires a failback process that minimizes impact to ongoing operations

Remote Data Mirroring Solutions

Benefits of using Point-in-Time Remote Mirroring

- The primary cluster provides automatic site failover locally
  - Consists of an HP MetroCluster with EMC SRDF
- Extended distance mirroring with no application latency
  - Multi-Hop (Point-in-Time) mirroring operation performed independent of real-time processing
- Multi-Hop mirroring operation for changed tracks only
  - Symmetrix maintains invalid track information reducing resynchronization time
    - Also reduces switched network bandwidth requirements
- Allows for intercontinental mirroring of data
  - Can be used for data warehousing and DSS applications
    - Additional copies can be configured to support these requirements

The major benefit is reduced network costs and mirrored I/O performance

Remote Data Mirroring Solutions

Standby database enables the creation and maintenance of a duplicate, or standby copy of your production database

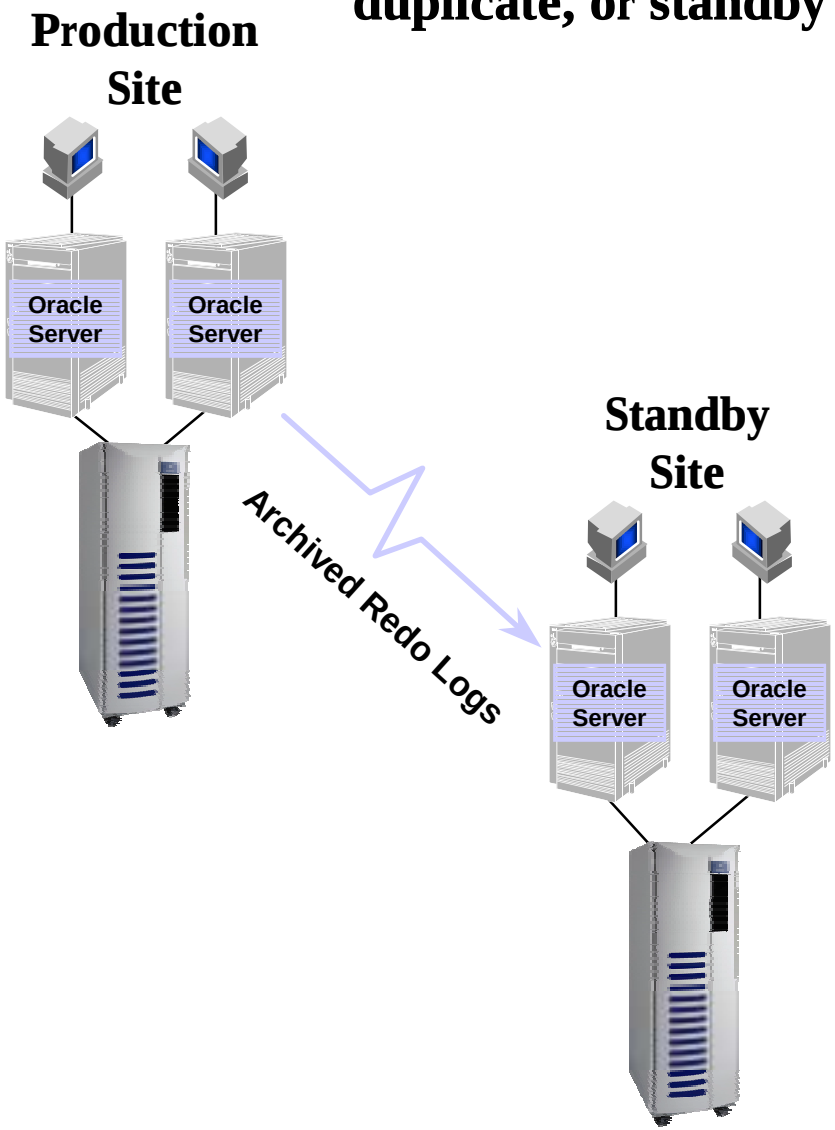

- Streamlined management of standby databases and elimination of human error
- Automatic log shipping and application
- Rules to enforce consistency between production and standby database and correct failures
- Standby database can be opened read-only and used as a reporting system

Thank You

Questions?