



Leveraging HP StorageWorks XP enterprise storage on HP NonStop servers

Ron LaPedis, CBCP, CISSP, ISSAP,
ISSMP

June 2005

© 2004 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice



Agenda

- Terminology
- What is enterprise storage?
- How enterprise storage is used on an HP NonStop server
- How HP NonStop Remote Database Facility (NonStop RDF) software leverages enterprise storage



Cluster terminology

- HP Metrocluster for NonStop servers
 - Metrocluster allows NonStop servers to be geographically separated for disaster tolerance
 - ServerNet connections up to 15 kilometers (9 miles) and NonStop RDF/ZLT connections up to 100 kilometers (62 miles) are supported configurations
- HP Continentalclusters for NonStop servers
 - Continentalclusters allows systems to be geographically separated over unlimited distances through the use of NonStop RDF over Expand networking for the ultimate in disaster tolerance

Disaster tolerant solutions for the NonStop server



Solution	Benefit	Data replication	Comms	Storage	Distance	Time
Basic disaster recovery	Low-cost entry level	Tape backup or electronic vaulting	LAN, WAN	Any storage	Unlimited	1–7 days
Metrocluster with ServerNet	Local failover	NonStop RDF Host-based disk replication	ServerNet, DWDM, Fibre Channel	Internal disk StorageWorks XP	15 km	1–60 minutes
Metrocluster with remote disk	Regional failover	NonStop RDF Host-based disk replication	DWDM, Fibre Channel, FC over WAN	Internal disk Remote StorageWorks XP	StorageWorks XP—100 km	1–60 minutes
Continentalcluster S	Unlimited distance failover	NonStop RDF	WAN	Any storage	Unlimited	1–60 minutes

Disk terminology

- NonStop disk volume—This is also called a “disk,” and it can be mirrored or nonmirrored
 - If nonmirrored, this is only \$Disk-P
 - If mirrored, this is made up of \$Disk-P and \$Disk-M
- LDEV (logical device)—This is a single *virtual disk image* within the XP that maps to what has traditionally been called a “disk” on the NonStop server and is given a name like \$Disk-P or \$Disk-M
- Logical unit number (LUN)—The path to a host group on an XP port
 - Two LUNs per XP LDEV, the primary and the backup (for example, \$Disk-M and \$Disk-MB)

RAID—Redundant Array of Independent Disks



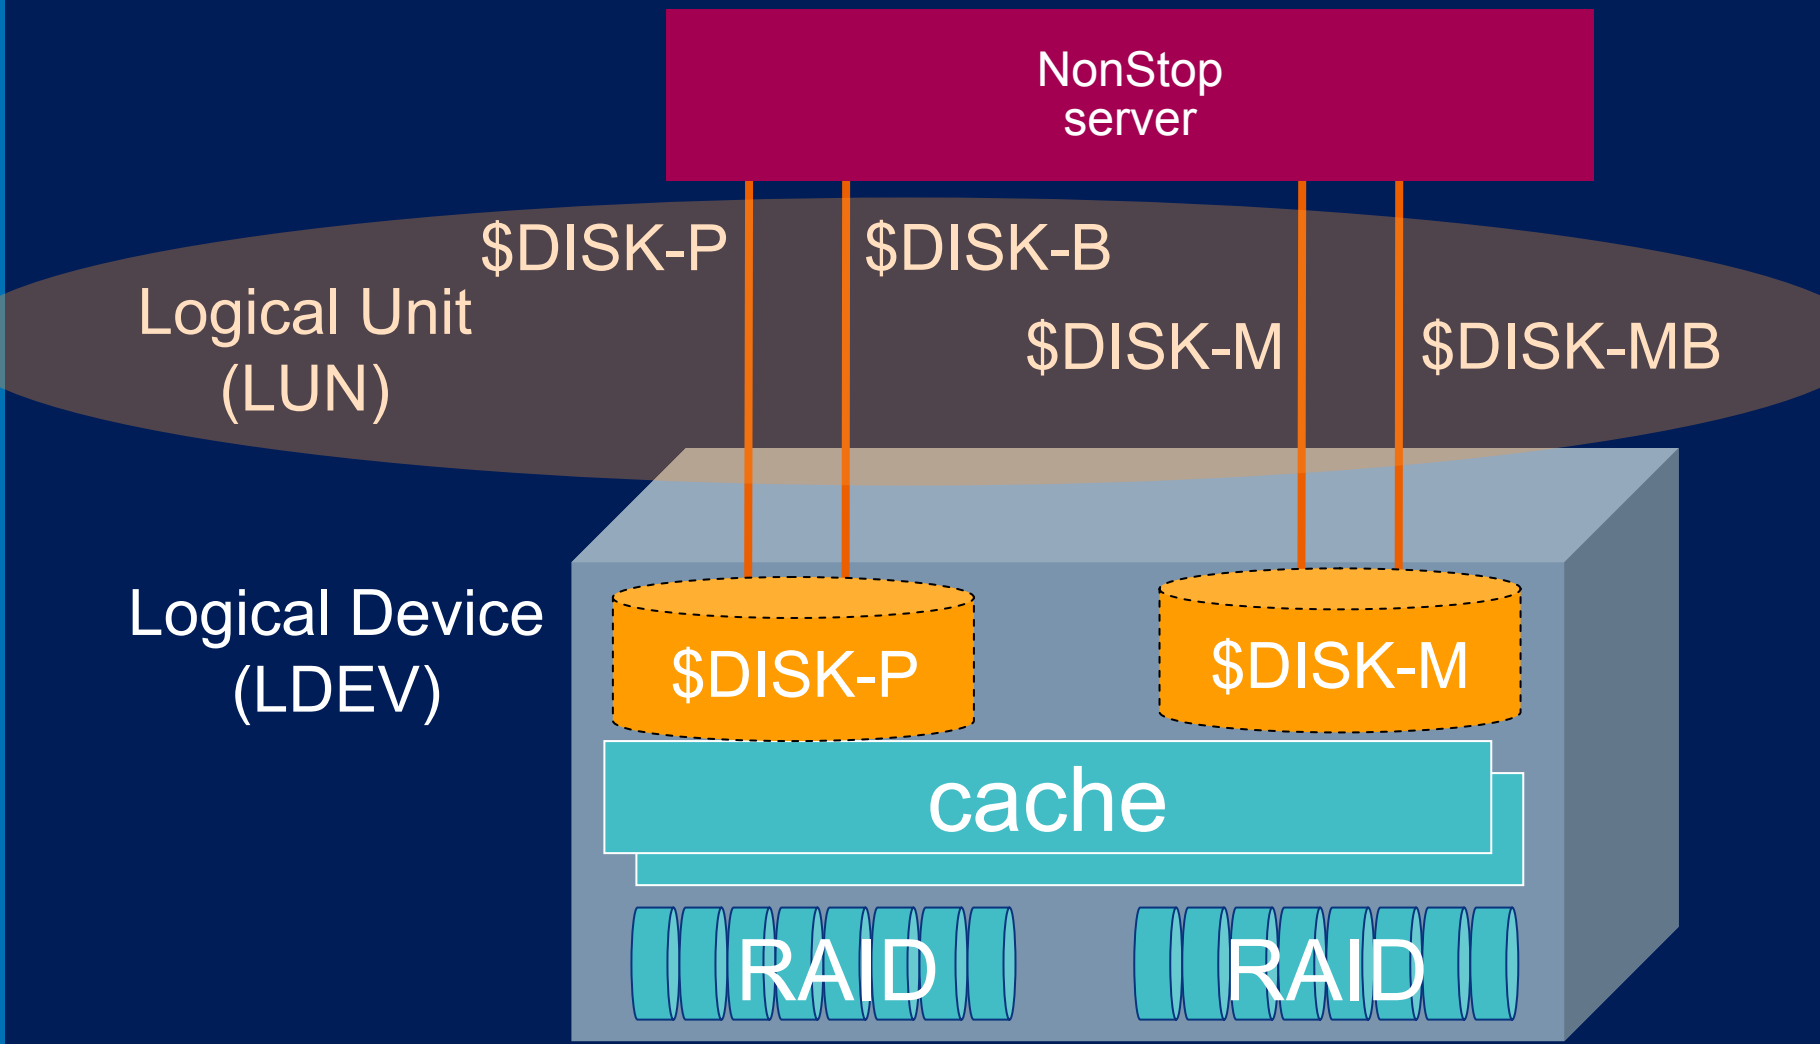
- RAID 1—Traditional mirroring
 - Two drives for each LDEV
- RAID 5—Level 5 provides data striping at the byte level and also stripe error correction information, resulting in excellent performance and good fault tolerance
 - Level 5 is one of the most popular implementations of RAID

Enterprise storage system



- Storage system with intelligent firmware
 - Highly available, with terabytes of storage
 - Multiplatform support
- Usually no data interchange between platforms


LUN versus LDEV




How much mirroring is enough?

- Host-based mirroring (\$Disk-P, \$Disk-M) on enterprise storage is strongly recommended
 - Host-based mirroring provides better protection against connection path and enterprise storage problems, taking full advantage of error recovery of NonStop system-based mirroring, whether using one or two enterprise storage disk arrays.
- One StorageWorks XP may be sufficient per site because it is fully fault tolerant
 - Customers need to understand their level of risk acceptance

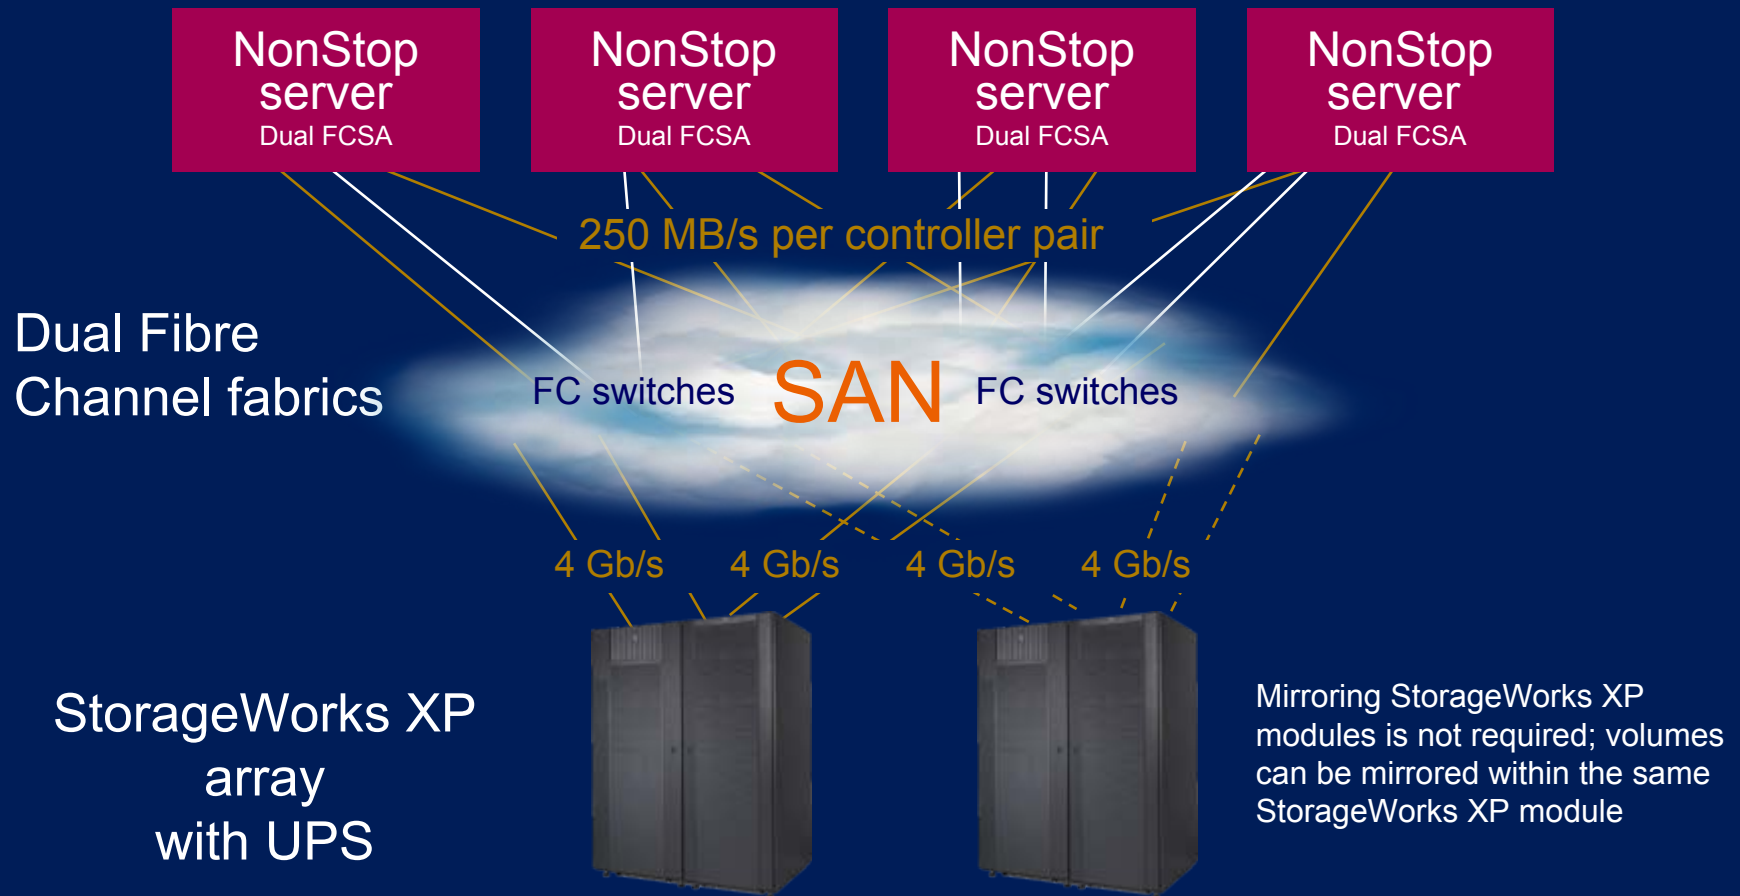
Host-based mirroring – one disk array



Host-based mirroring – dual disk arrays



NonStop SAN architecture



HP StorageWorks XP enterprise storage

How enterprise storage is used on a NonStop server




- Only one server can mount an XP LDEV at a time
 - No loss of functionality; two NonStop servers could never mount the same internal disk at the same time either
 - Two systems still can open the same file using remote access over HP Expand networking software
- NonStop RDF/Zero Lost Transactions software leverages enterprise storage for disaster tolerance

How enterprise storage is used on a NonStop server



- As with a traditional (internal) disk, a volume must be cleanly dismounted from one system before it can be mounted on another system
 - All files on the volume closed, volume disabled in NonStop TMF, disk cache flushed, and volume removed through Subsystem Command Facility (SCF)
- Business copy is minimally supported—the application must be quiesced just as if you are using the “backup” program
 - NonStop TMF cannot make a disk snapshot transactionally consistent

Continuous Access is not supported for NonStop server volumes



How enterprise storage is used on a NonStop server



- When a system crashes, for whatever reason, the disks mounted on it **cannot** be switched to a different system—most of the time
 - When a system is restarted after a crash, NonStop TMF automatically recovers disk volumes and the audited files on them using the audit trails
 - NonStop SQL keeps system-specific information in its files
 - There is an SQL/MP fix-up utility to adjust the embedded names if you want to move a disk to a system with a different node number

How enterprise storage is used on a NonStop server



- You could switch the system's entire disk configuration—including the NonStop TMF audit volumes—to another system with the same node name/number and bring up the original NonStop TMF configuration to start volume recovery.

Can I mirror my entire disk configuration remotely?

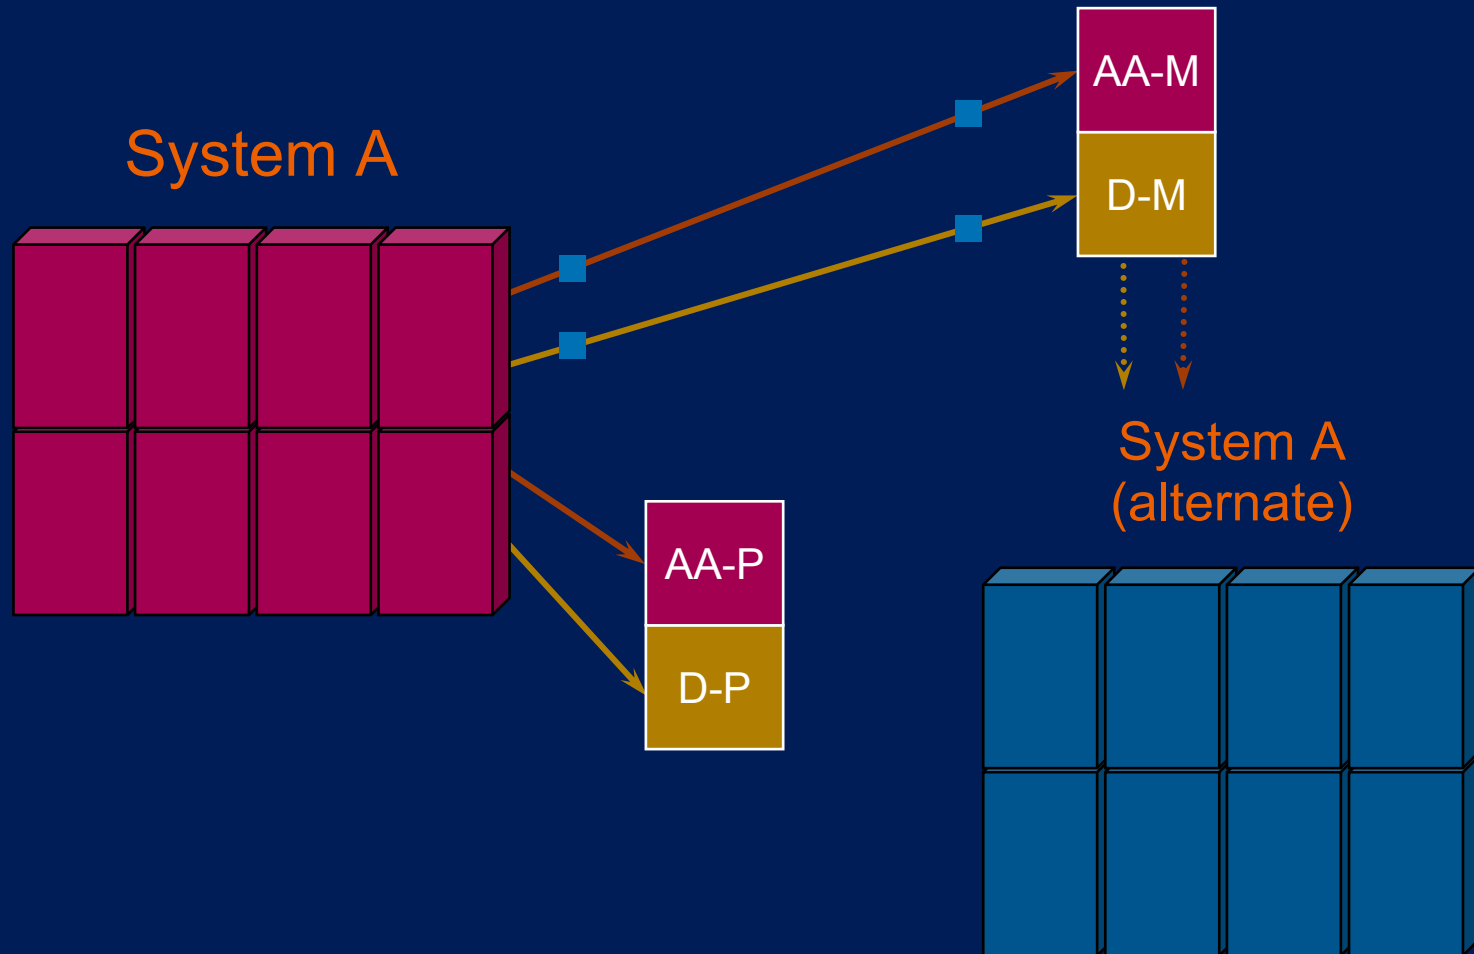


Remember Nomadic Disk?

February 11, 2005

- HP announces the end of sales of the Nomadic Disk technology for HP NonStop S-series servers. Nomadic Disk technology was first introduced in 1995 for use with NonStop K-series servers.
- The functionality of Nomadic Disk is available through the use of HP StorageWorks XP Disk Array.

Can I mirror my entire disk configuration remotely?



Can I mirror my entire disk configuration remotely?



- If you mirror your entire disk configuration remotely, it could be switched to another system if the primary system fails; however...
 - The target system must have same node name/number
 - Substantial communications bandwidth is required
 - Information is sent twice: database files and NonStop TMF audit trail

Can I mirror my entire disk configuration remotely?



- Even a 10-second communications outage could have catastrophic consequences because you have to revive all your mirrors (smart revive can mitigate this somewhat)
- Losing one remote mirror before a total failure can make your entire database unrecoverable
- NonStop TMF will perform volume recovery on system startup
 - If NonStop TMF needs a dump tape to complete volume recovery, where is it?
- How often would you exercise the failover process?

Steps to switch an entire disk configuration

1. Stop applications on the backup system
2. Stop NonStop TMF
3. Stop operating system
4. Rename system
5. Start operating system
6. Switch disks to the backup system
7. Start NonStop TMF with the configuration from the failed system
8. Volume recovery by NonStop TMF
9. Integrate existing disks into the moved NonStop TMF configuration
10. Integrate both systems' tables into the same NonStop SQL catalog
11. Reconfigure NonStop TMF tape catalog if necessary
12. Take new online dumps
13. Bring up applications

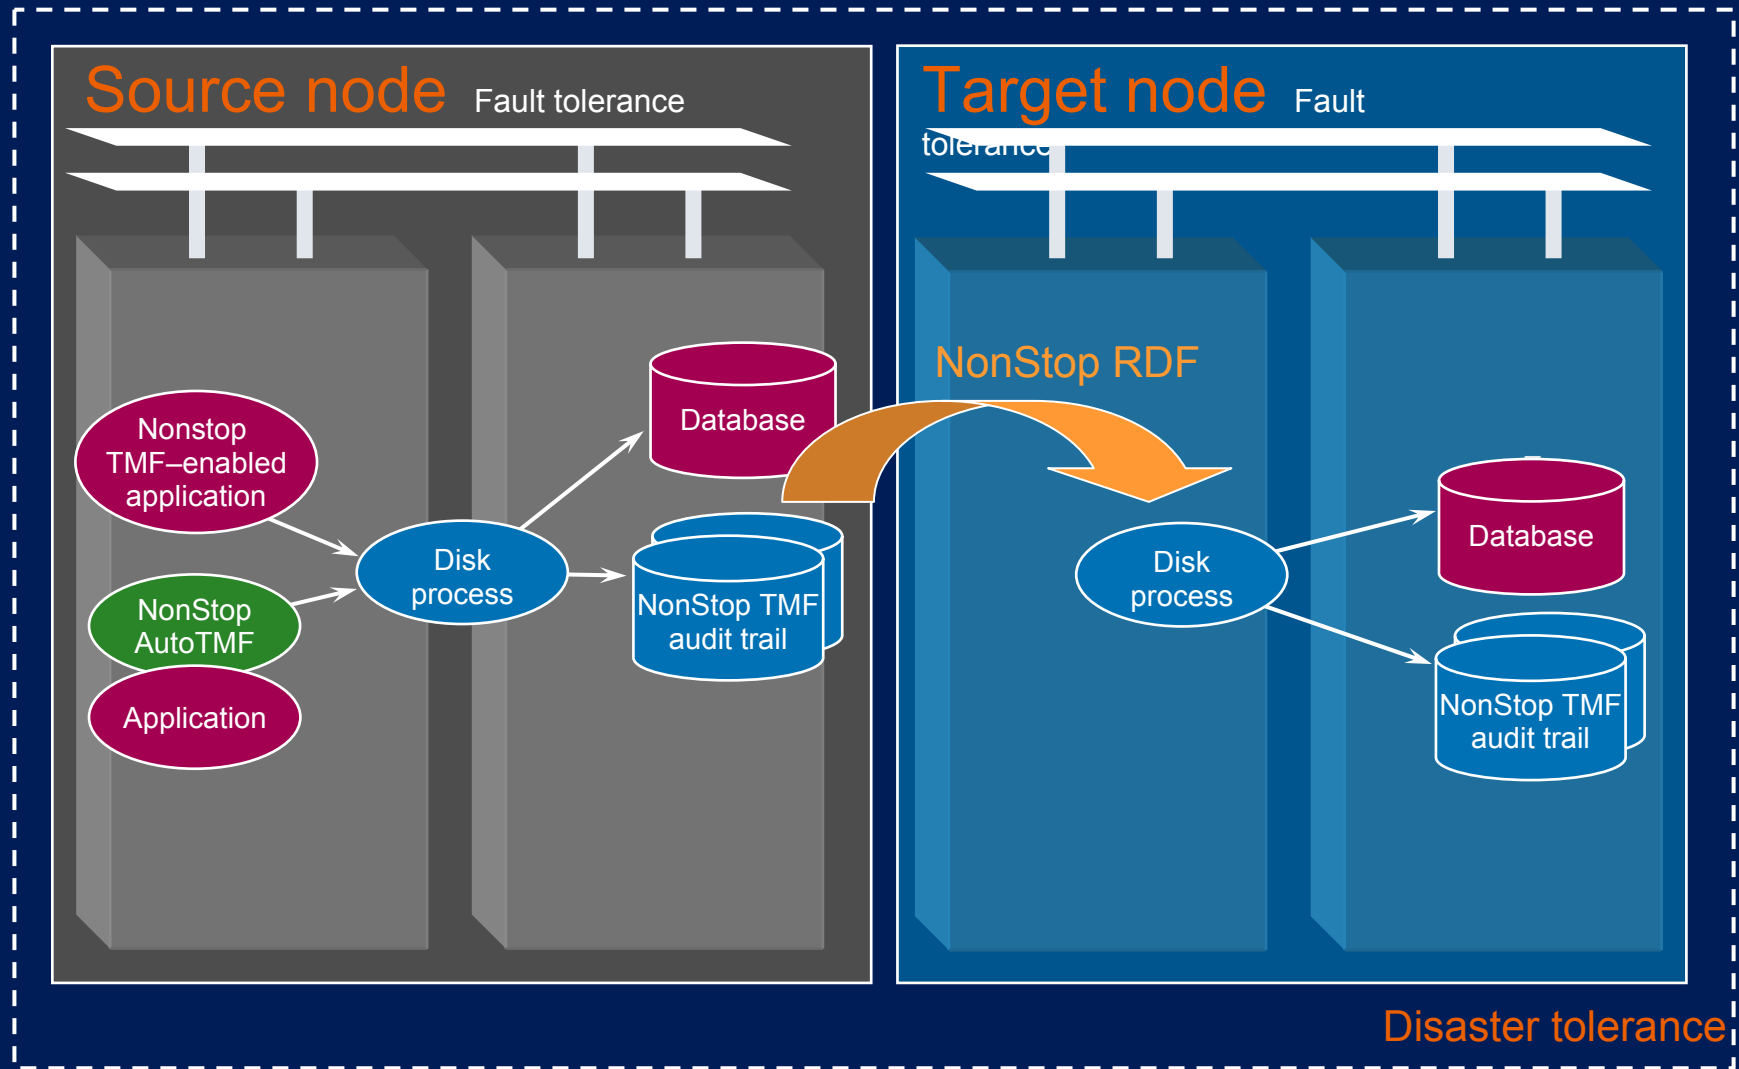


There must be an
easier way for
another system to
take control of my
database after a
system or site failure

There is: NonStop RDF software


- High-speed, low-latency database replication software
- Peer-to-peer replication for NonStop servers only
- Focused on disaster tolerance
- Uses low-level system interfaces

NonStop TMF and NonStop RDF software: Unbeatable for NonStop system-to-system replication




Some NonStop RDF topologies


Centralized




Simplex




Multiple duplicate sites




Reciprocal/
split workload



Ring




Triple contingency



NonStop RDF failover

1. Issue NonStop RDF takeover command
2. Start applications

Backup system takeover timeline





NonStop RDF/ZLT
software leverages
enterprise storage.

The issue of “lost” transactions

- All asynchronous replication products (including NonStop RDF) are vulnerable to losing some transactions if there is an unplanned outage of the primary system
 - Some have lower latency than others
 - NonStop RDF has the lowest latency in the NonStop system environment because parts of it are built into the operating system
- Synchronous or lockstep replication
 - Guarantees no loss of transactions
 - Increases commit response times
 - Transaction cannot commit until data is committed in backup database
 - Greater distance = greater latency = lower TPS

Our solution: NonStop RDF/ZLT

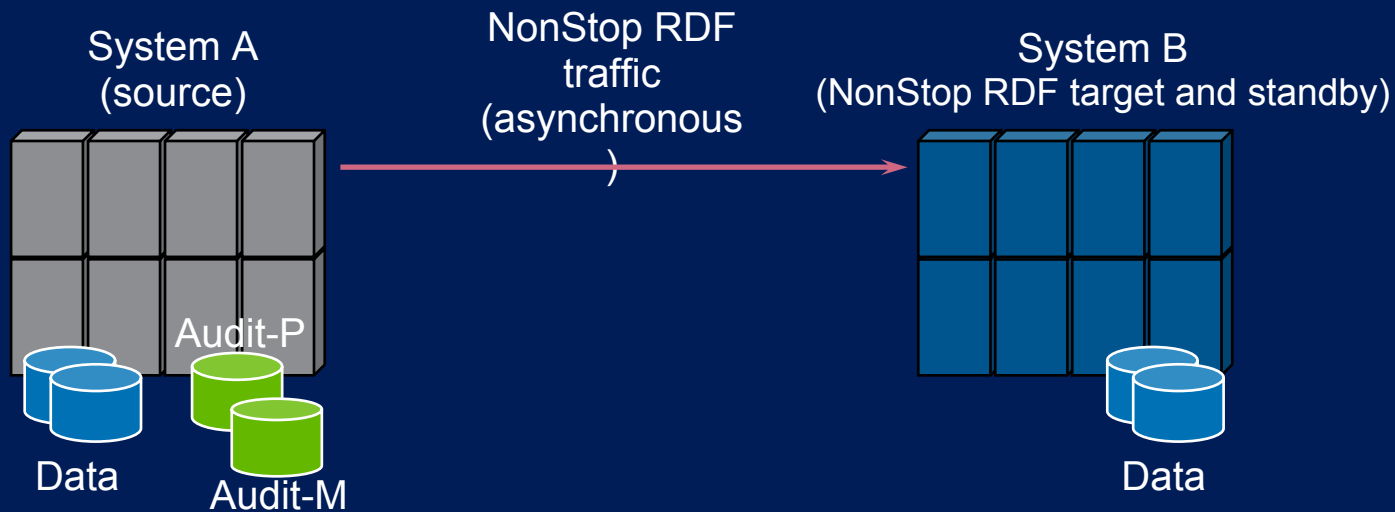
- A new member of the NonStop RDF family, NonStop RDF/Zero Lost Transactions (NonStop RDF/ZLT [SA88V1]) can eliminate the loss of committed data even after an unplanned outage—without any application changes
- NonStop RDF/ZLT software was shipped on June 1, 2004
 - HP Services must be involved in planning and configuration

How NonStop RDF/ZLT leverages StorageWorks XP

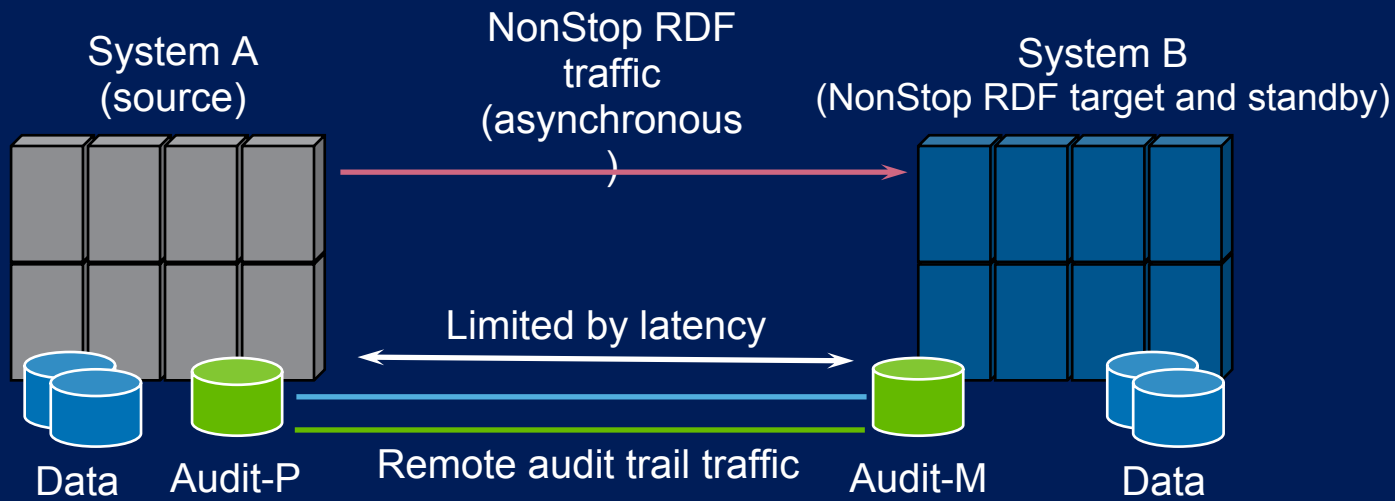


- For each audit trail volume in the NonStop RDF configuration, one of the two disk mirrors making up the volume is located remotely from the system it is protecting
- NonStop RDF standby system
 - Can be a third system
 - Can be the target system itself
 - Caveat: The audit trail volume names must be unique across the source and standby systems
- The distance from system to remote volume is dictated by the customer's latency and throughput requirements—
There is no change up to 100 kilometers (62 miles)
- If the standby system is not the target system, there is no distance limit for location of the target


NonStop RDF without ZLT configuration




NonStop RDF/ZLT moves the audit mirror to the backup site




NonStop RDF/ZLI moves the audit mirror to a location in between the sites



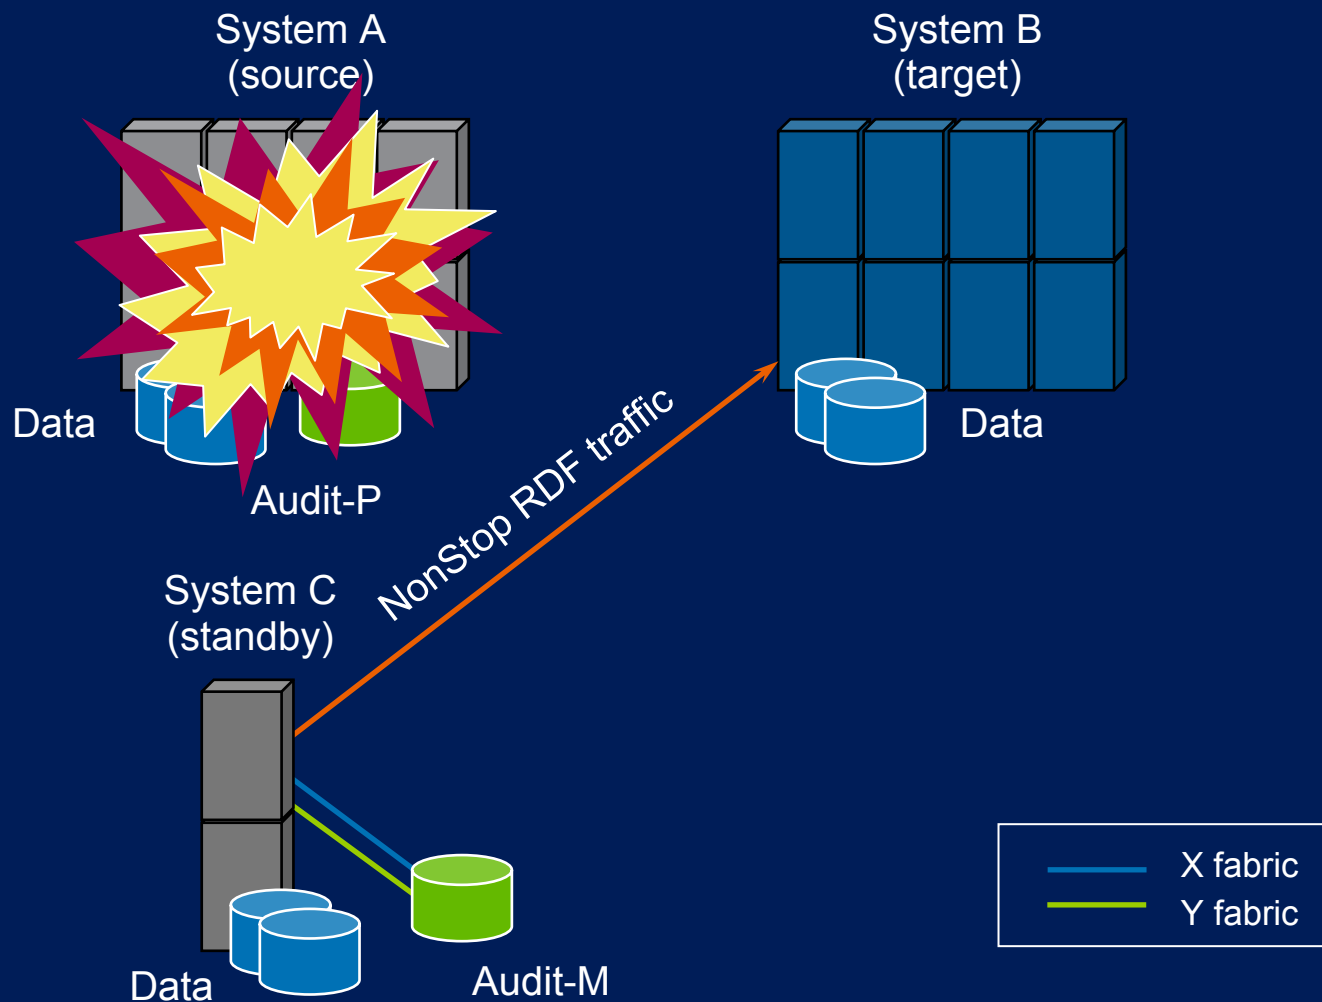
NonStop RDF/ZLI remote mirror with third system as standby: Normal processing



NonStop RDF/ZLI remote mirror with third system as standby: **Source system fails**



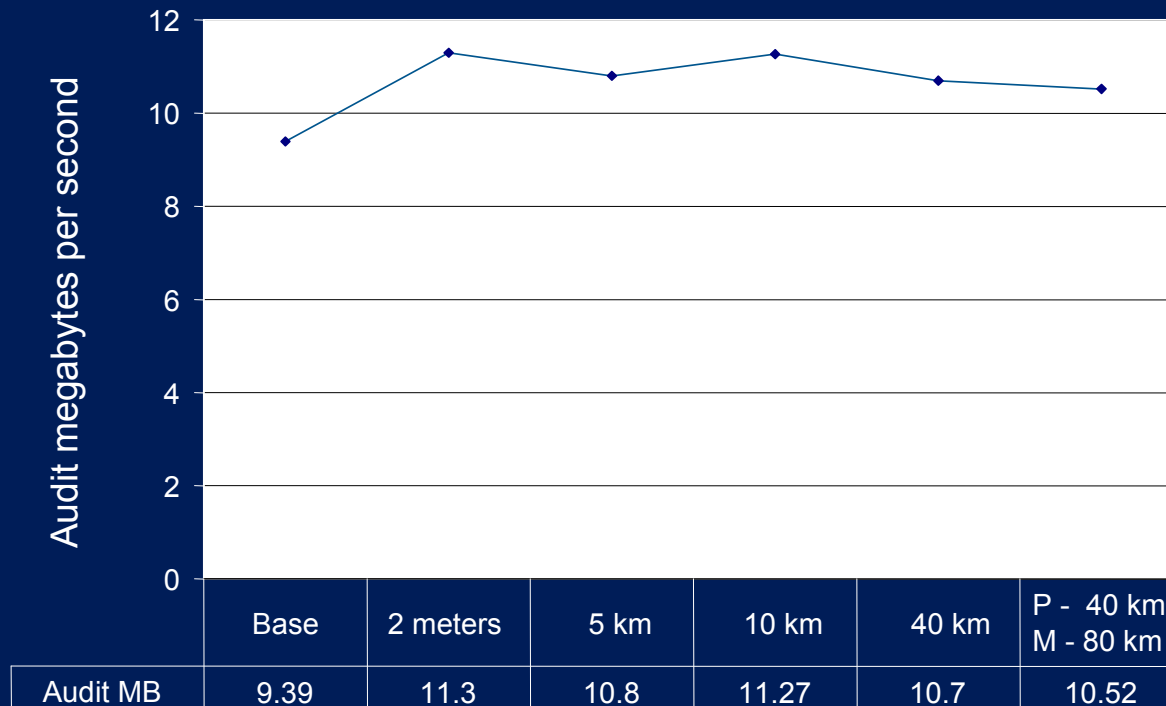
NonStop RDF/ZLT remote mirror with third system as standby: NonStop RDF recovers in-flight transactions



NonStop TMF audit trail performance

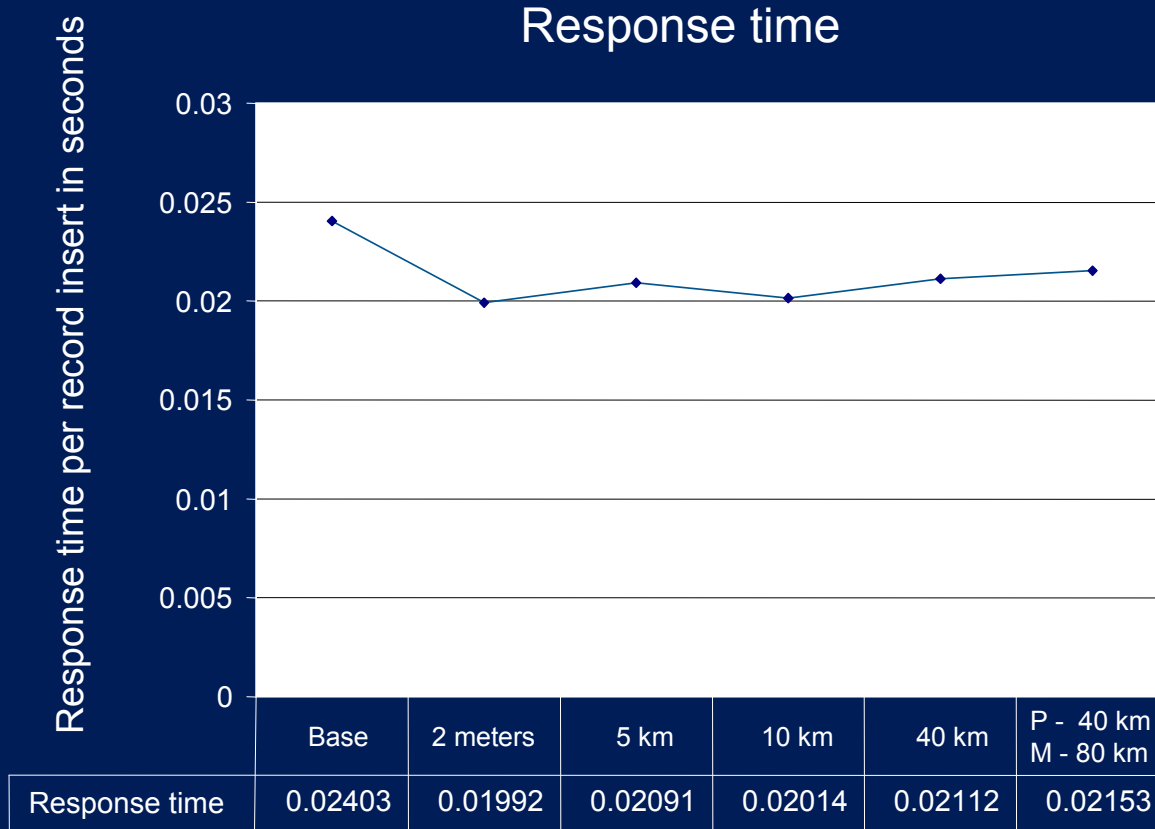


Audit generation in MB/s



Disclaimer: Future product plans, dates, and functionality are subject to change without notice.

NonStop TMF audit trail performance



Disclaimer: Future product plans, dates, and functionality are subject to change without notice.

NonStop RDF/ZLT remote mirror summary



- With access to the remote mirror, no transactions are lost
 - NonStop RDF reads all remaining audit prior to actual takeover processing
 - Lower overhead method delivers the same result as lockstep
- Distance is limited by disk technology
 - Fibre Channel is estimated at 100 kilometers
- What if an audit mirror goes down prior to unplanned outage?
 - If local mirror, no transactions are lost during NonStop RDF takeover
 - If remote mirror, some final transactions may be lost
 - In the future, the user will be allowed to determine whether to continue or to stop transaction processing

Why do I need NonStop RDF if I have enterprise storage?

- More flexible
- Much lower communications bandwidth
- Faster failover
- Easier to exercise—and switch back
- Both systems can be processing transactions in reciprocal or split workload configuration
- Backup system does not need to be brought down during NonStop RDF failover processing

Contrast RDF with Enterprise Storage



NonStop RDF failover

1. Issue NonStop RDF takeover command
2. Bring up applications

Entire disk switch

1. Stop applications on the backup system
2. Stop NonStop TMF
3. Stop operating system
4. Rename system
5. Start operating system
6. Switch disks to the backup system
7. Start NonStop TMF with the configuration from the failed system
8. Volume recovery by NonStop TMF
9. Integrate existing disks into moved NonStop TMF configuration
10. Integrate both systems' tables into the same NonStop SQL catalog
11. Reconfigure NonStop TMF tape catalog if necessary
12. Take new online dumps
13. Bring up applications

For more information

- hp.com/go/nonstopcontinuity
- hp.com/go/continuity
- ron.lapedis@hp.com

- <http://nedmigration.cac.cpqcorp.net/domains/EnterpriseStorage.htm>
- nomi.trapnell@hp.com

Q & A



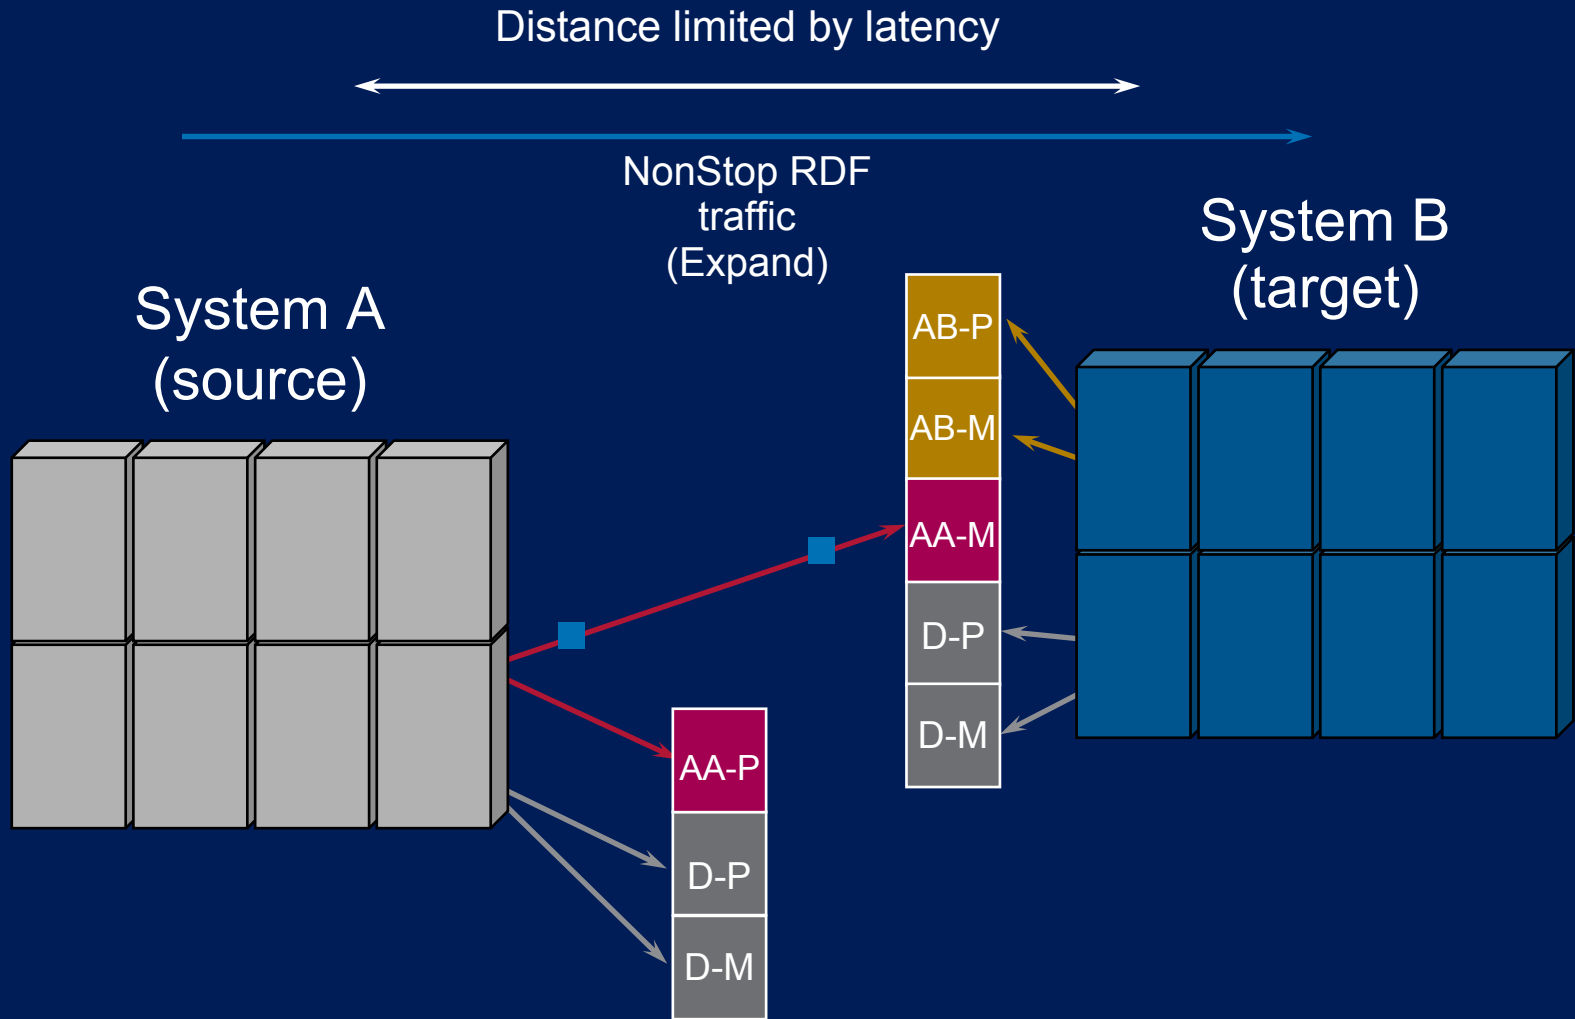
i n v e n t




Presenter notes

- The next slides show some StorageWorks XP configuration examples and can be inserted after slide 43.

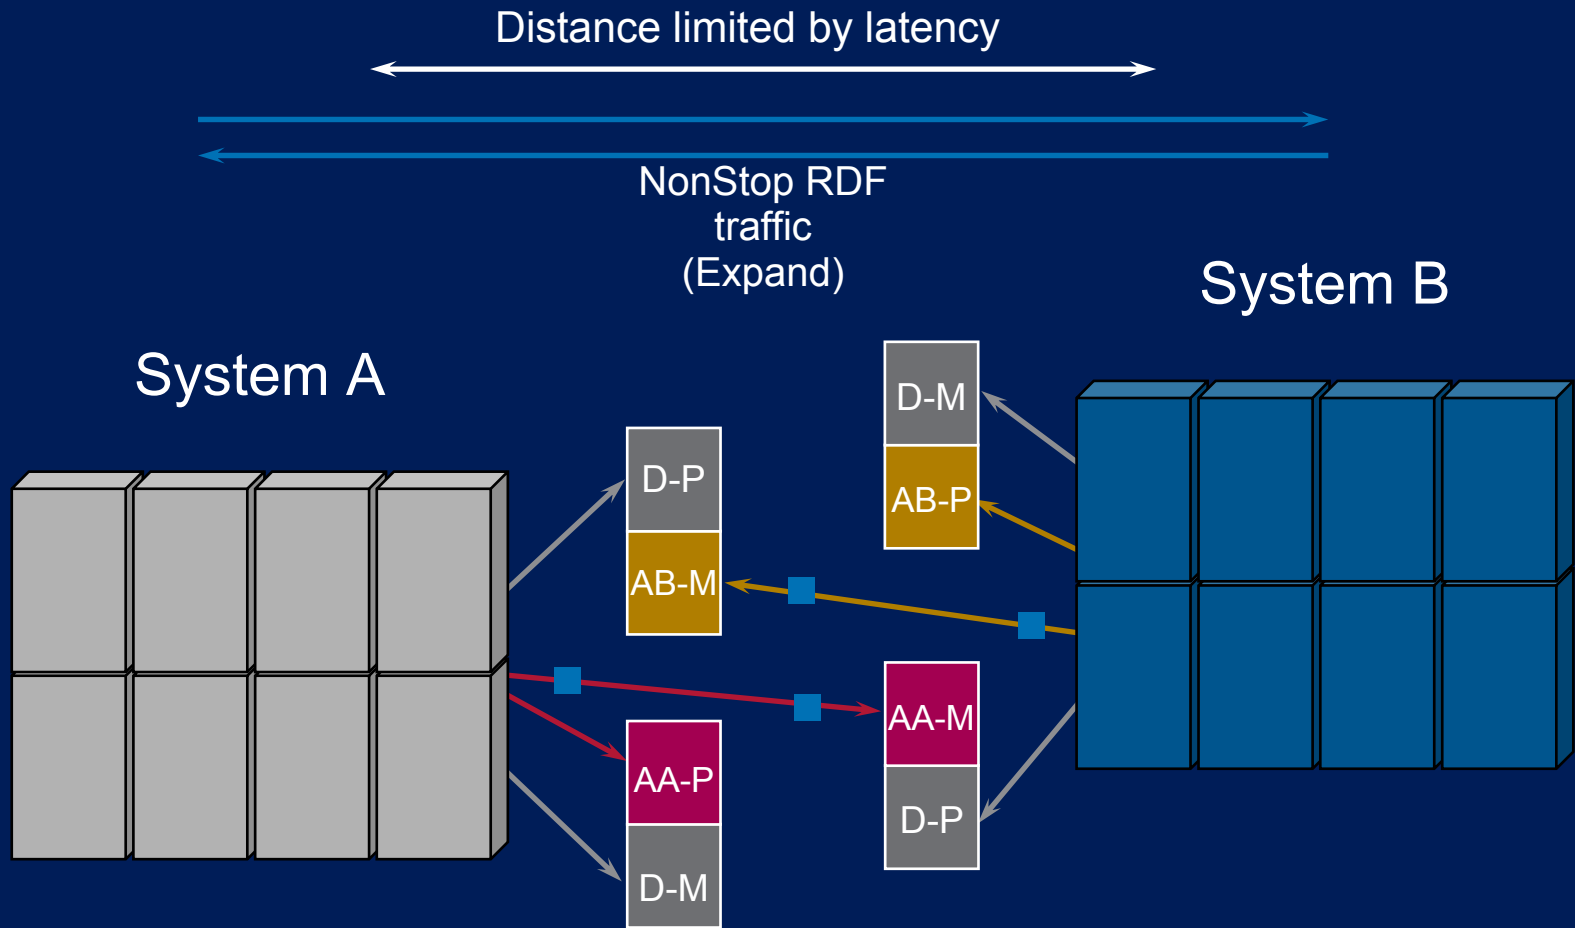
StorageWorks XP with IP switches: Single direction



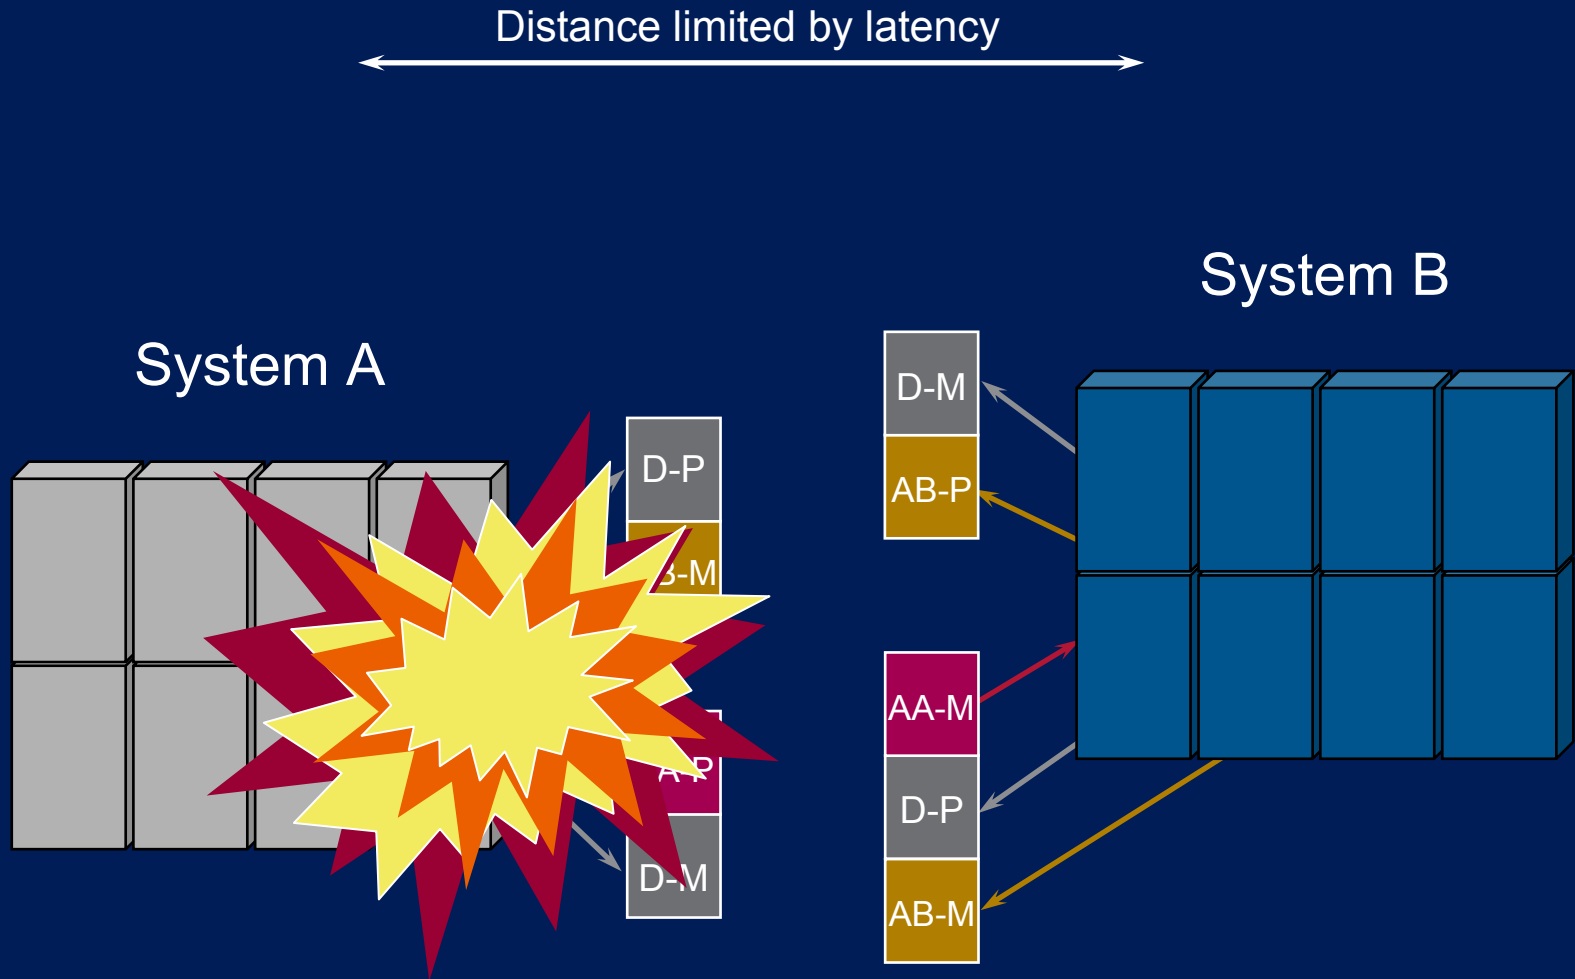
StorageWorks XP with IP switches: Single direction



Dual StorageWorks XP with IP switches: Reciprocal direction



Dual StorageWorks XP with IP switches: Reciprocal direction





Presenter notes

- These slides describe HP's Business Critical Services (BCS) and can be inserted where you like for emphasis

economics

It's all about ~~IT features~~

Are you being served?

improving / service-



ITe = cost ↓ + quality ↑ + risk ↓ +

HP Services are required for installation of both the StorageWorks XP disk array and NonStop RDF/ZLT

customers need to mitigate risk!

enhanced business agility


MUS- cost ↑

Why services for the NonStop server?


Ensure that the rest of your infrastructure is as robust as your NonStop server.

Achieve the high levels of IT availability and performance required by your adaptive enterprise through people and processes:


- Comprehensive processes and tools
- Proactive services
- Reactive services
- Support partnerships
- Account management



HP services which help your journey to adaptability



NonStop Server Support Portfolio



Services for the NonStop server value proposition

Enable agility in a business-critical environment by providing a highly available, secure, manageable, scalable IT infrastructure.

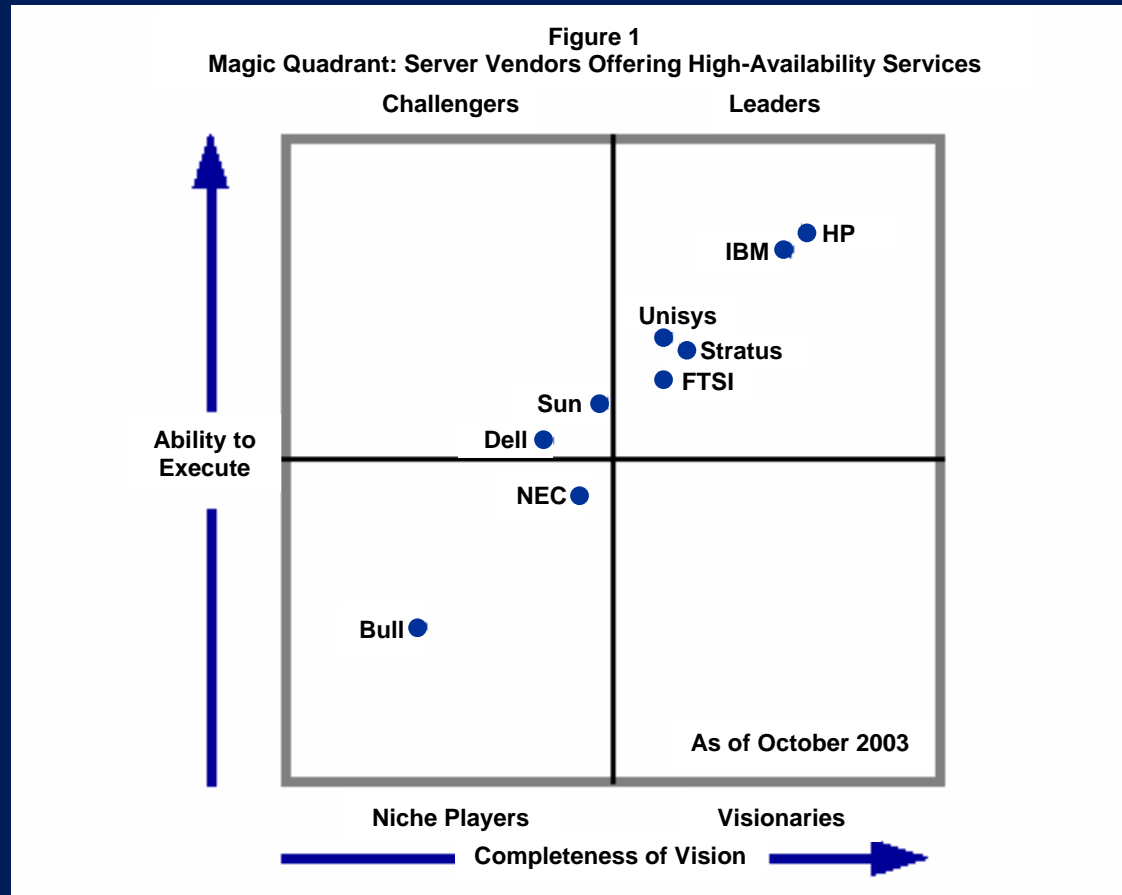


**Track
record**

Gartner Magic Quadrant

Evaluation of selected server vendors for their High Availability services

The Magic Quadrant is copyrighted October, 2003, by Gartner, Inc. and is reused with permission, which permission should not be deemed to be an endorsement of any company or product depicted in the quadrant. The Magic Quadrant is Gartner, Inc.'s opinion and is an analytical representation of a marketplace at and for a specific time period. It measures vendors against Gartner defined criteria for a marketplace. The positioning of vendors within a Magic Quadrant is based on the complex interplay of many factors. Gartner does not advise enterprises to select only those firms in the "Leaders" quadrant. In some situations, firms in the Visionary, Challenger, or Niche Player quadrants may be the right matches for an enterprise's requirements. Well-informed vendor selection decisions should rely on more than a Magic Quadrant. Gartner research is intended to be one of many information sources including other published information and direct analyst interaction. Gartner, Inc. expressly disclaims all warranties, express or implied, of fitness of this research for a particular purpose.



Source: Gartner Research Note, October, 2003, Server Vendor's High-Availability Services: Magic Quadrant

Can the rest of your infrastructure support your NonStop server?



Industry average MC availability

- 27 hours per year downtime*
- 99.69% availability

• \$108K/hour mean*

• \$2,916K cost of downtime

HP Mission Critical availability

- 3.06 hours per year downtime
- 99.97% availability

• \$108K/hour mean*

• \$330.5K cost of downtime

*Source: Gartner : North American Customers Reveal Preferred Services for Mission –Critical Systems 9/5/03

Why HP mission critical support?

Superior customer service experience

- >8X less downtime than industry average
- 95% of calls direct connected to technical resource within 5 minutes
- 98% of call-to-resolution time commitments met
- 20-30% of reactive calls avoided due to enabling technology's preventive analysis capabilities
- Experience and expertise
 - 5,000 high availability experts
 - thousands of mission-critical customers worldwide

HP mission critical leadership



- #1 support company
- Leader in mission-critical support for open environments
- #1 in storage consulting and support services
- Gartner Group High Availability Magic Quadrant leader
- SAP award of excellence every year since 1995
- Leader in high availability for support for Windows environments

Business value enabled by Mission Critical Services



Complete, accurate
back office settlements
Investment Banking

Reliable delivery of
multimedia services
Wireless telecom carriers

Accurate, timely
customer billings
Wireline telecom carriers

Ensure quality and
safety of food
processing
Food manufacturing



24 x 7 electronic trading
capability
Investment banking

Reduce manufacturing
design cycle times
Consumer electronics
manufacturing

Improve supply chain
efficiency and accuracy
Retailing and
manufacturing

Maintain retail shelf
inventory
Retailing

Summary

- Delivering and measuring IT business value is critical to IT success
- Proactive, comprehensive approach is needed to:
 - Minimize the business risk of downtime
 - Increase IT quality of service
 - Manage IT costs
 - Gain agility
- HP Mission Critical Services assist IT in achieving these objectives



HP Services: committed to your success



- Continued support of legacy systems
- Unsurpassed experience and expertise
- Lifecycle services using a holistic approach
- Designing, building, managing, and evolving customer infrastructures to meet business challenges
- Proactive service partnership
- Comprehensive consistent, global support
- Best-in-class tools, processes, and remote support through Global Customer Support Center (Support Center Practices Certified)
- Investment in continuous improvement

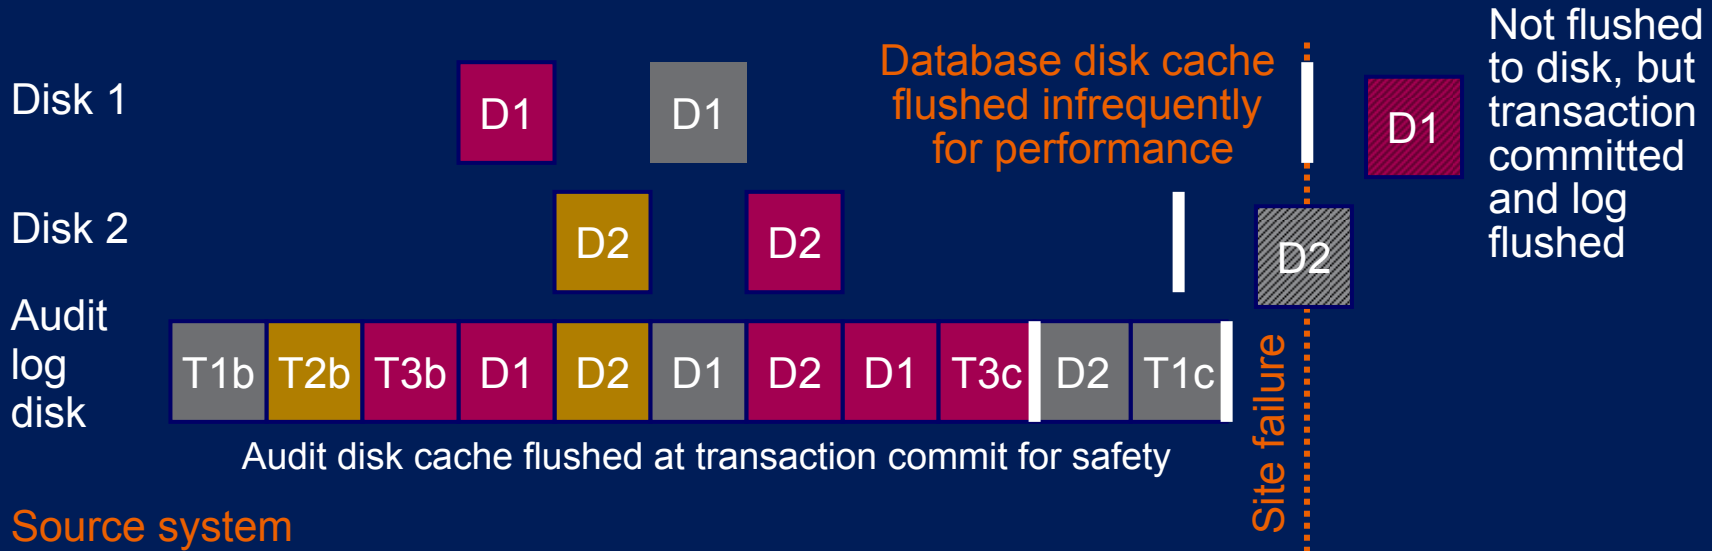


Your satisfaction is the measure of our success!

Presenter notes

- This slide is an animated illustration of why replicated Enterprise Storage (Continuous Access) does not show a consistent image of the database and needs software support from the OS or database management system to make it so.

Physical disk does not equal logical database



Target system

