Enhanced Disaster Recovery Options

Irene Adler and Mitch Mackrory
Oracle Corporation

Thursday, August 11, 2011: 1:30 PM-2:30 PM
Session 09894

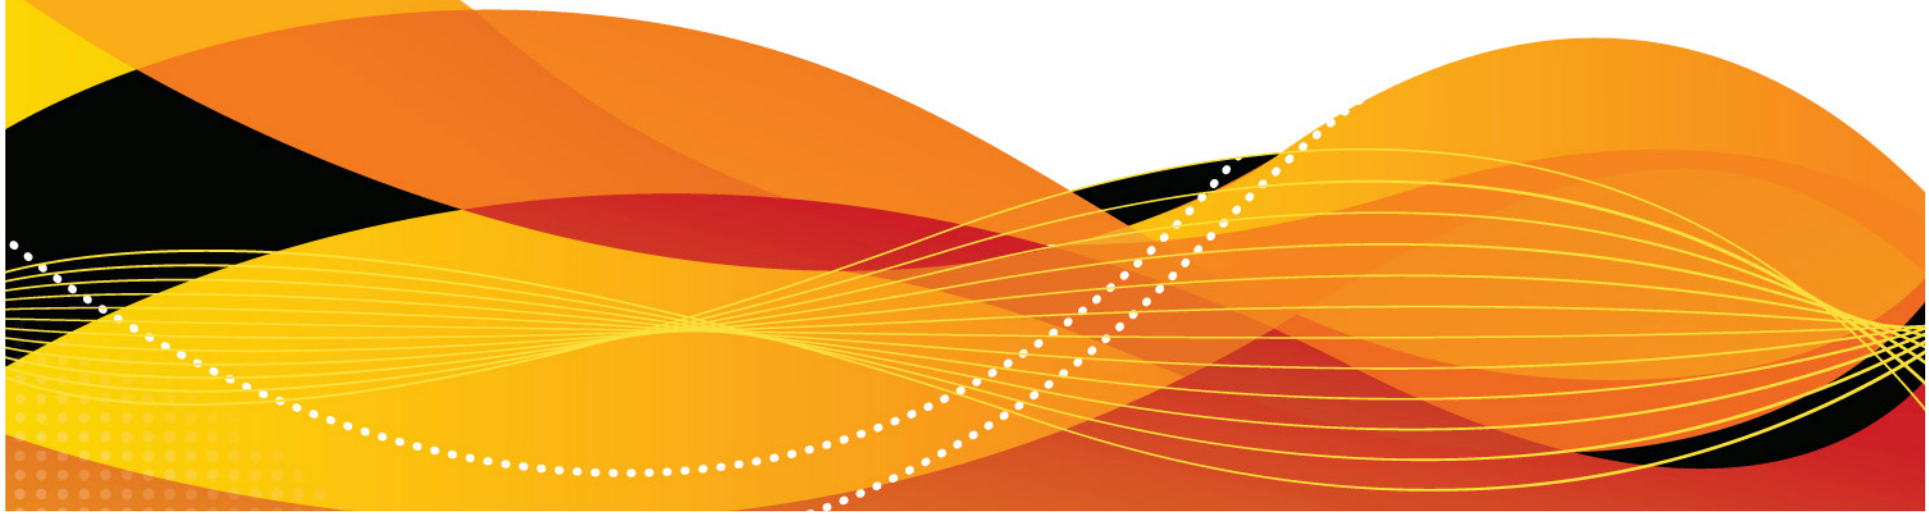

Agenda

- Some Definitions
  - (BC/DR/RPO/RTO)
- Tapeless, Tiered Tapeless, Touchless and Break Link
- Hierarchy of Performance & Capacity
- Some VSM Recovery Options
- VTV Copies
- Auto Archiving
- Export/Import and Physical Vaulting
- RTV Utility
- Clustered VTSS
  - Standard Clustered
  - Extended Clustered
- Cross-TapePlex Replication
- Concurrent DR Test
- Q&A

Some Definitions: (BC/DR/RPO/RTO)

- Business Continuity
  - The ability to continue operations smoothly after a failure or outage
- Disaster Recovery
  - The ability to recover from a disaster
- RPO
  - The Recovery Point Objective is how far back in time we need to go to be able to recover
- RTO
  - The Recovery Time Objective is how long it will take us to recover

Tapeless, Tiered Tapeless, Touchless and Break Link

- These are terms loosely used within Oracle and/or outside
- **Tapeless** means no tape is attached
  - This can be used for part or all of a configuration
- **Tiered Tapeless** means we have a tiered structure of capacity and performance in a tapeless configuration
  - Includes VLE (Virtual Library Extension) as a second tier
- **Touchless** is when we have any of the above and we add automated tape
- **Break Link** is an important concept where we break the automated link between disk and tape
  - This can be a critical feature of a disaster recovery design

Hierarchy of Performance and Capacity

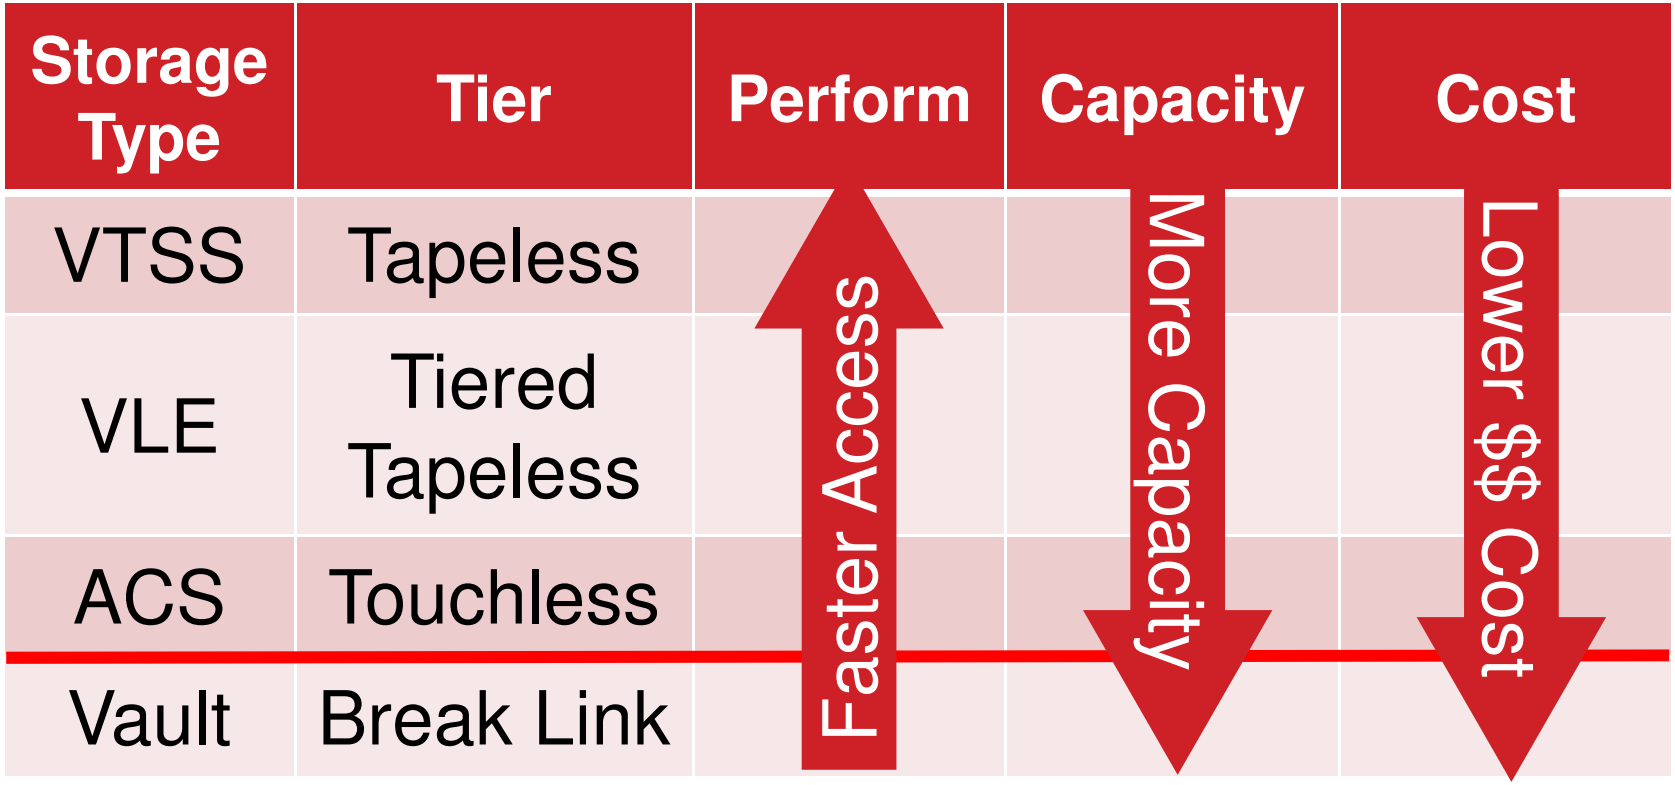

Some VSM Recovery Options

- VTV copies
- Export/Import
- Auto Archiving
- Physical Vaulting
- RTV Utility
- Clustered VTSS
- Extended Clustering
- Cross-TapePlex Replication
- Concurrent Disaster Recovery Test

VTV Copies

- VTV copies
  - 1 to 4 VTV copies
    - On different MVCs
    - On MVCs in up to 4 different locations
    - On different media
      - for maximum performance
      - high capacity for archive
  - This allows customers to have VTV copies at their production site, DR site, and other site(s) automatically via parameter definitions
  - VTV copies can be synchronous or asynchronous

Automatic Archiving

- VSM Archiving features allow customers ability to automatically archive virtual volumes from high performance to high capacity media based on customer-defined policies

- ARCHAge(nnn) specifies the age (in days) from creation of a VTV before it is archived automatically as specified by ARCHPol

- ARCHPol specifies up to four Storage Classes that specify:
  - ACS and media type of the archive MVCs
  - VTCS can archive multiple Storage Classes and can:
    - *archive VTVs to different MVCs in different ACSs –or-*
    - *different MVCs in the same ACS*

- If ARCHPol is defined, ARCHAge must also be specified on the Management Class statements

Automatic Archiving (continued)

- Usage would be to migrate VTVs to high performance media for a period of time, say 90 days, when VTVs are considered to be ready for deep archive on high capacity media
- When VTVs meet the ARCHAge criteria, they would be automatically recalled back into the VTSS buffer and re-migrated out according to the ARCHPol parameters
- Some customers choose to make all copies up front and then when Archive executes, it will just delete the high performance copies
  - this eliminates the need to recall and migrate again later
  - saves on VTSS, VLE, and RTD resources
  - but, it does mean additional copies at creation time

EXPORT / IMPORT

- The VTCS EXPORT/IMPORT features allow the customer to EXPORT VTVs to a different site for DR testing or physical vaulting

- Volumes can be EXPORTed:
  - by VTV or VTV-range (VTVs are consolidated to new MVCs)
  - by Management Class(es) (VTVs are consolidated to new MVCs)
  - by MVC or MVC-range (additional MVC copy not made)
  - by Storage Class(es) (additional MVC copy not made)

- EXPORT creates a Manifest File, which contains all of the metadata that pertains to the volumes being EXPORTed

- IMPORT will import volumes using the Manifest File into a separate CDS and VSM system at another TapePlex

Physical Vaulting

- Off-site physical vaulting of DR tapes can be accomplished using LCM (Library Content Manager)
- LCM is a software product that was developed specifically for Oracle that interfaces with HSC and VTCS to provide tape management services for Nearline and VSM
- LCM 7.0 replaces the ExLM (Expert Library Manager) product
- LCM 7.0 controls off-site vaulting in concert with the customer's tape management system, in addition to other tape utility functions
- LCM 7.0 is a replacement for the standalone VSM Offsite Vaulting Utilities

Physical Vaulting (continued)

- LCM's off-site vaulting services interfaces with the following tape management systems:
  - CA-1
  - DFSMSrmm
  - CA-TLMS
  - Control-T

- LCM uses the VAULT statement to define to LCM an ELS-controlled Vault and allows assignment of attributes via parameters for volumes assigned to the Vault
  - LCM assigns vault slot numbers to each tape being vaulted
  - LCM assigns a time period that the tape should remain in vault
  - LCM controls library ejects of cartridges to be vaulted off-site
  - LCM determines when tapes should be returned to the production environment from the vault and produces a pick list
  - Returning tapes are entered back into Library

RTV (Real Tape Volume) Utility

- RTV is a standalone VTCS utility that is designed to read VTV data directly from an MVC, without any assistance from VTCS, mounted on a native tape drive
- Can be run when VTCS is down
- Used to recover VTV(s) in the event VTCS cannot be brought up due to an outage or disaster of some nature
- Used to read VTVs in TapePlex that has no VSM systems
- RTV works by reading a single VTV directly from an MVC, decompressing the VTV, then writing the data to a single output tape

Clustered VTSS

- With Clustered VTSS, virtual volumes are written to a primary VTSS and then upon dismount are replicated to a secondary VTSS
- Uni-directional - one site replicating to a second site
- Bi-directional - each site replicates to the other
- Replication can be:
  - via ESCON, FICON or IP
  - synchronous or non-synchronous
  - requires no host involvement
  - runs in background
- Secondary acts as a warm standby
- Provides immediate Business Continuance
- Eliminates single point of failure

Uni-Directional Clustered VTSS

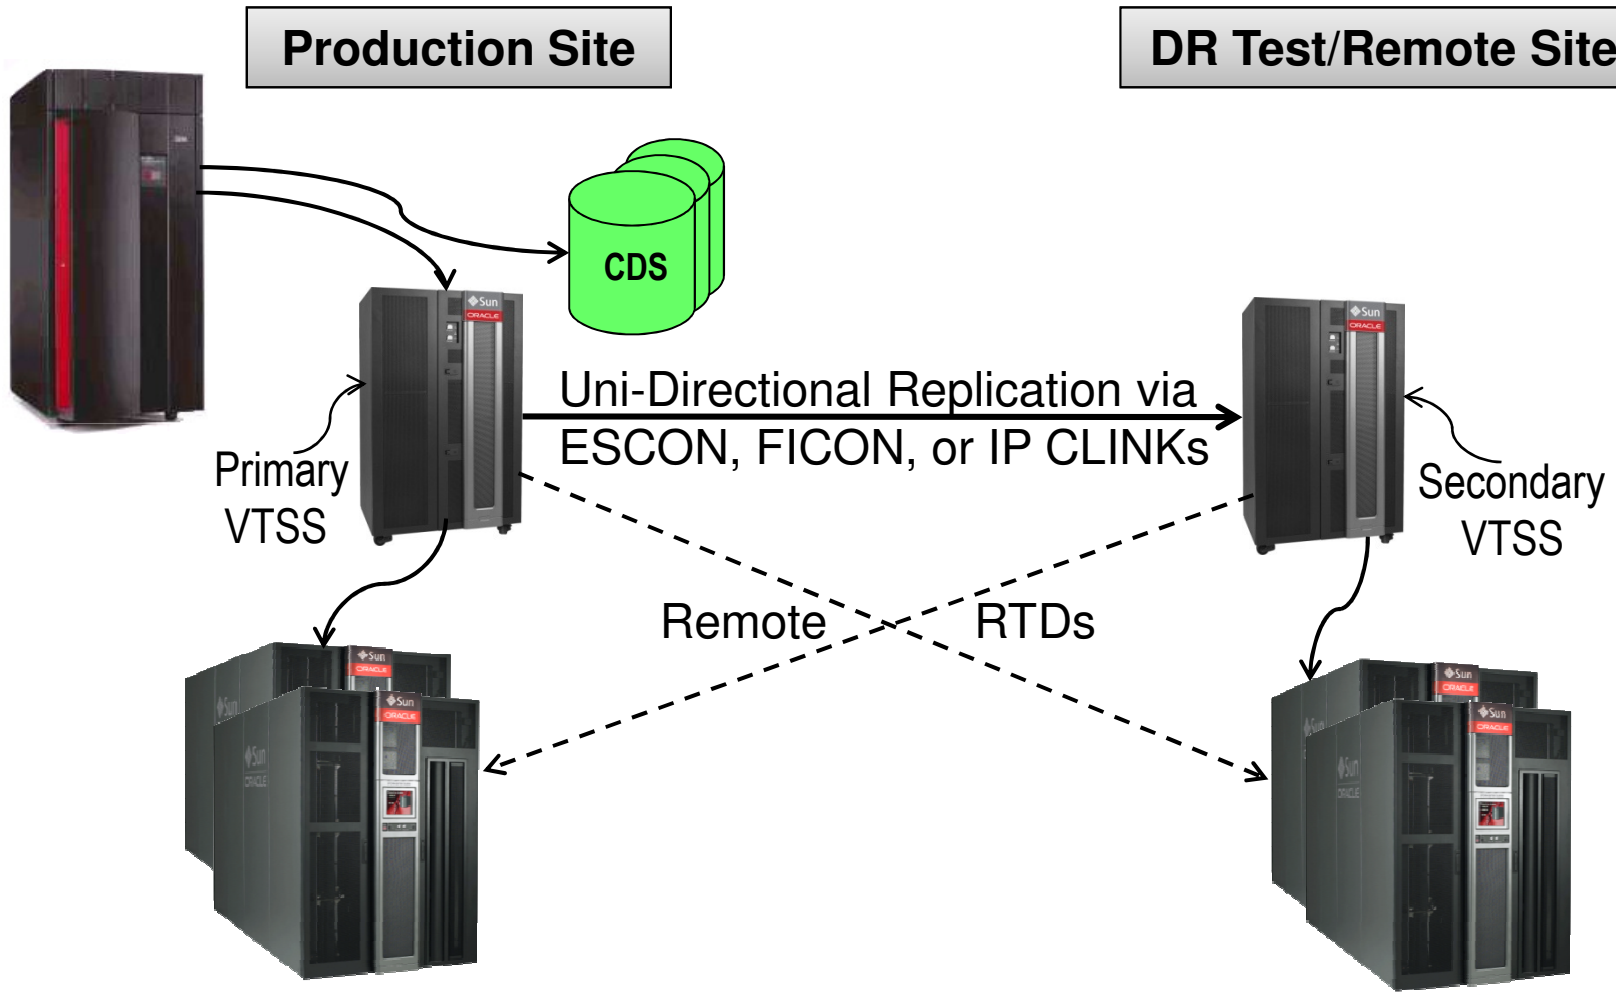

Bi-Directional Clustered VTSS

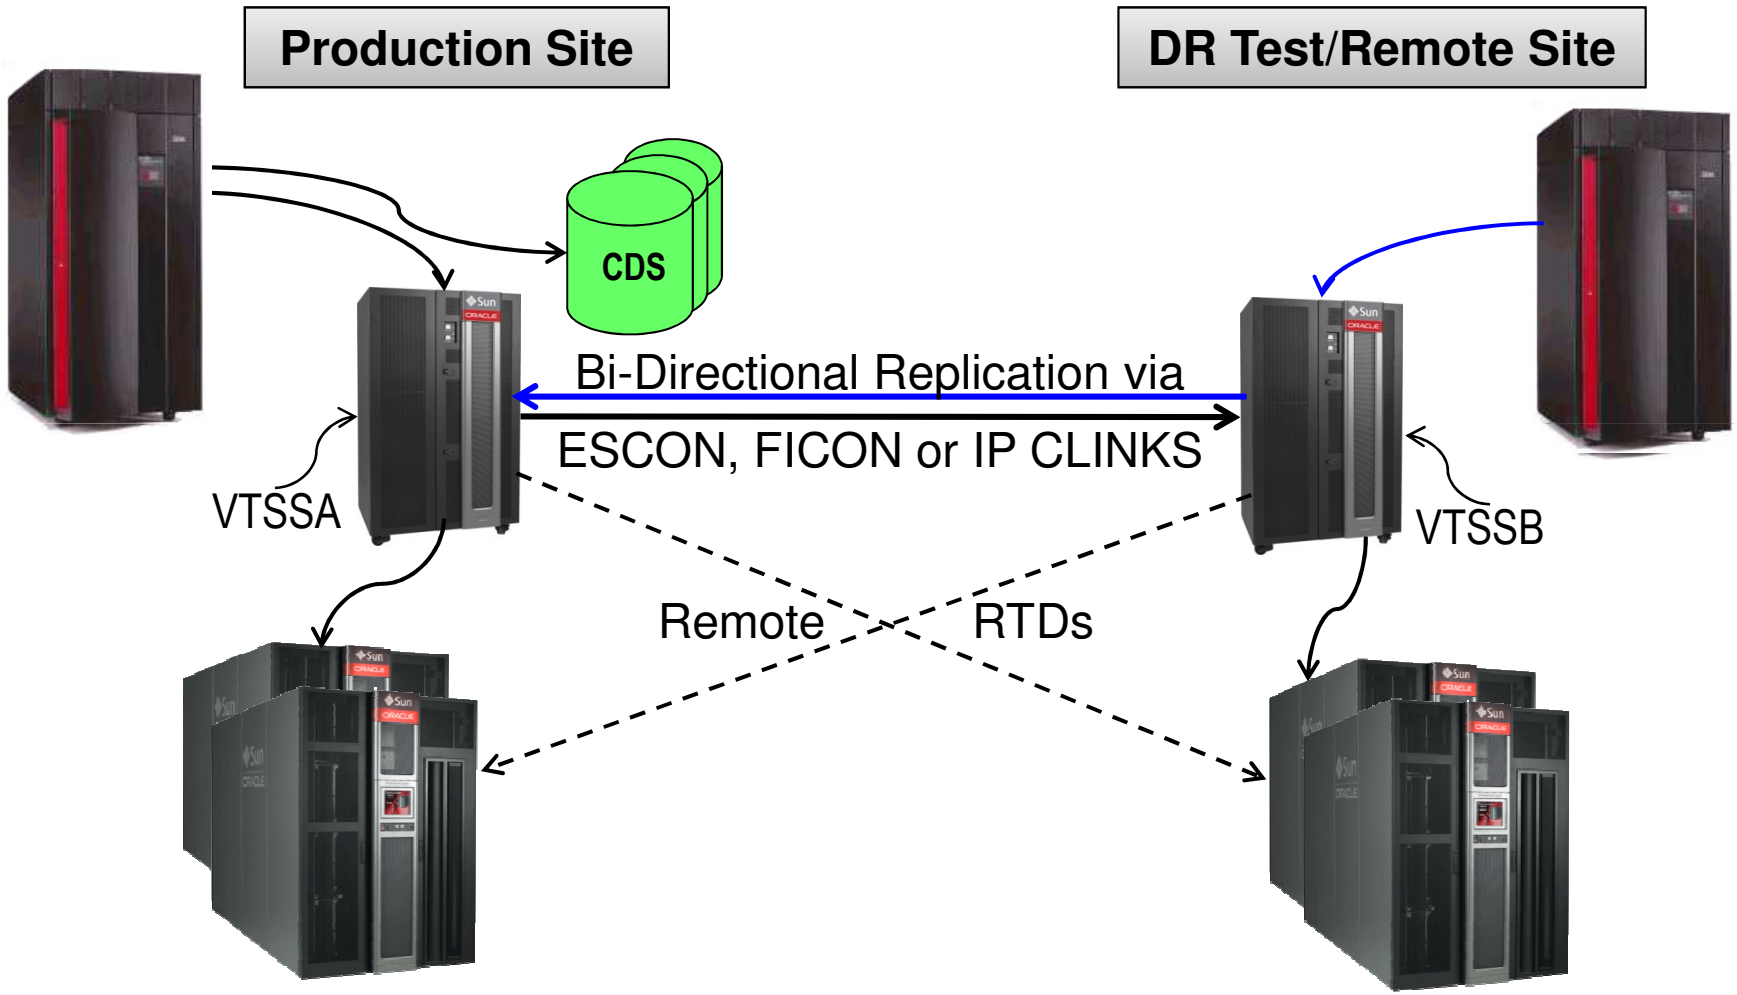

Extended Clustered VTSS

- **Many-to-one** VTSS Clustering - introduced with VTCS 7.0
- Allows many sites to replicate to a single DR repository
  - Customers could have multiple sites with a requirement to store a DR/BC copy at a common site

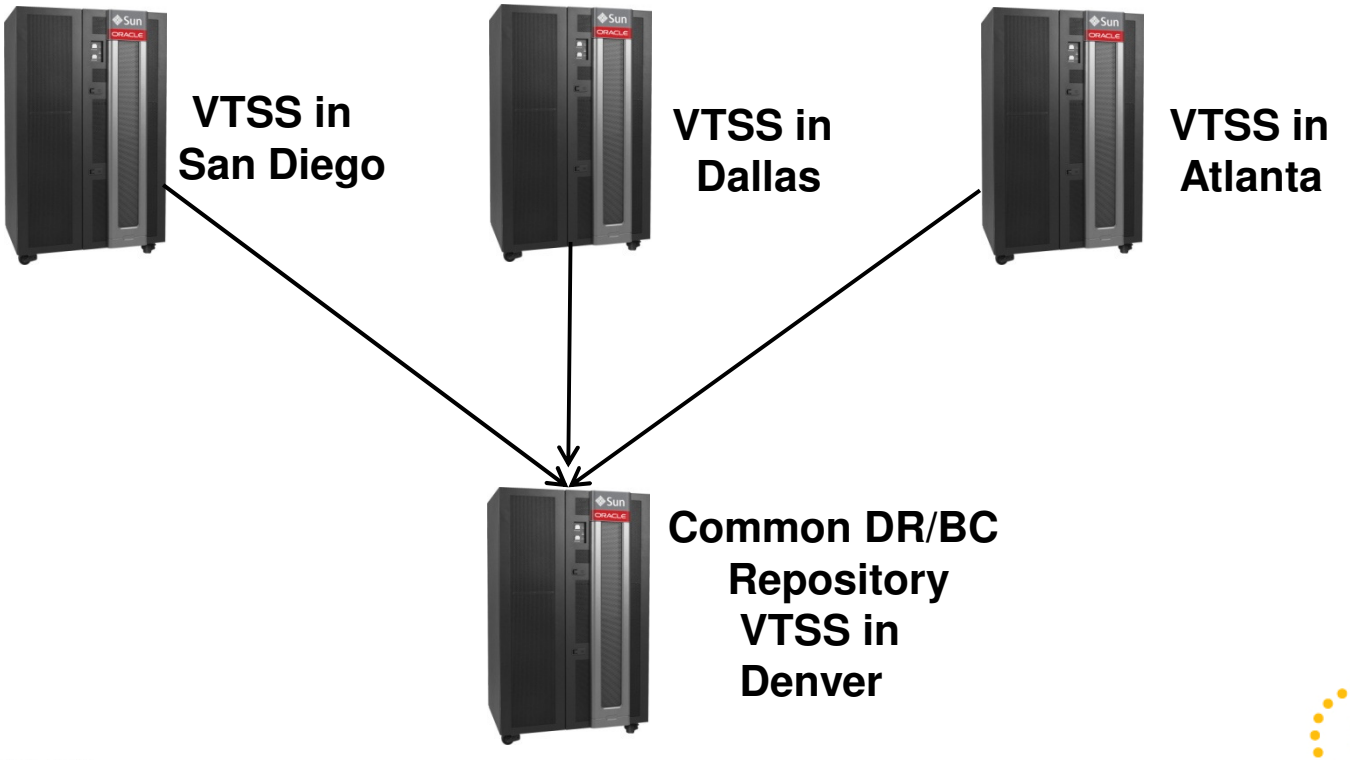

Extended Clustered VTSS (continued)

- **One-to-many** VTSS Clustering
- Allows one site to replicate to multiple sites
  - Customers may have requirement to have some workload replicated to one site and another workload to a second or third site

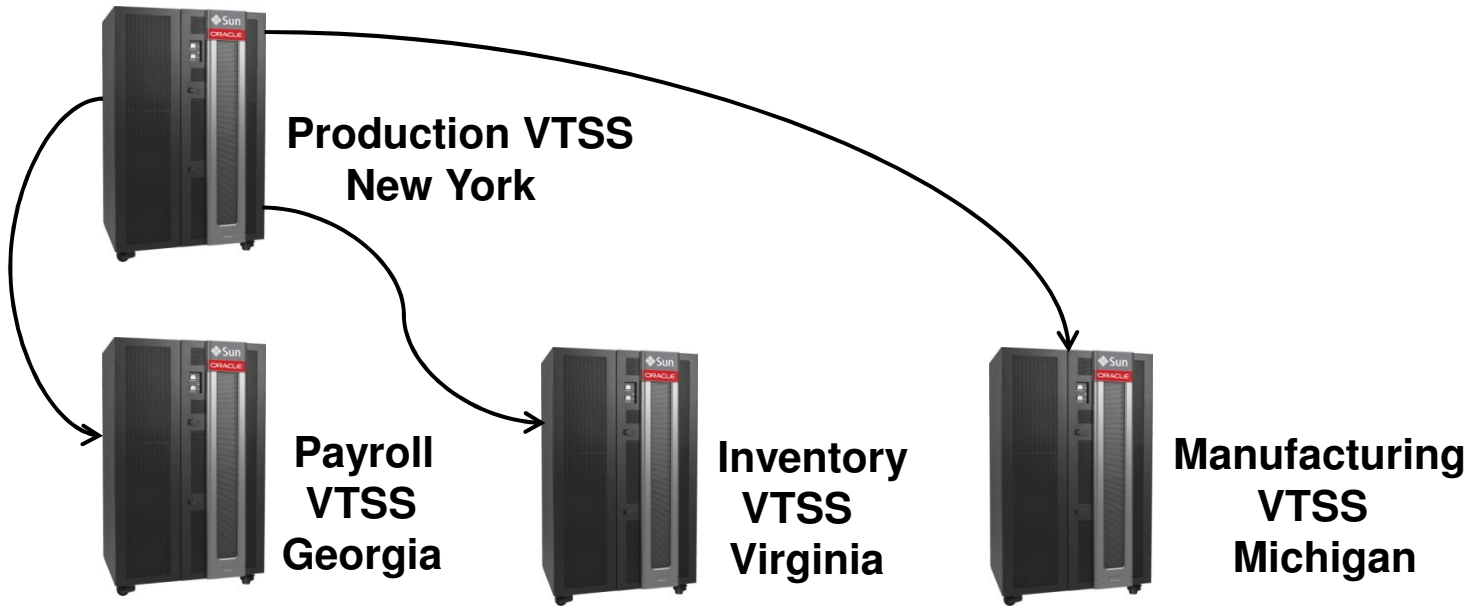

Cross-TapePlex Replication

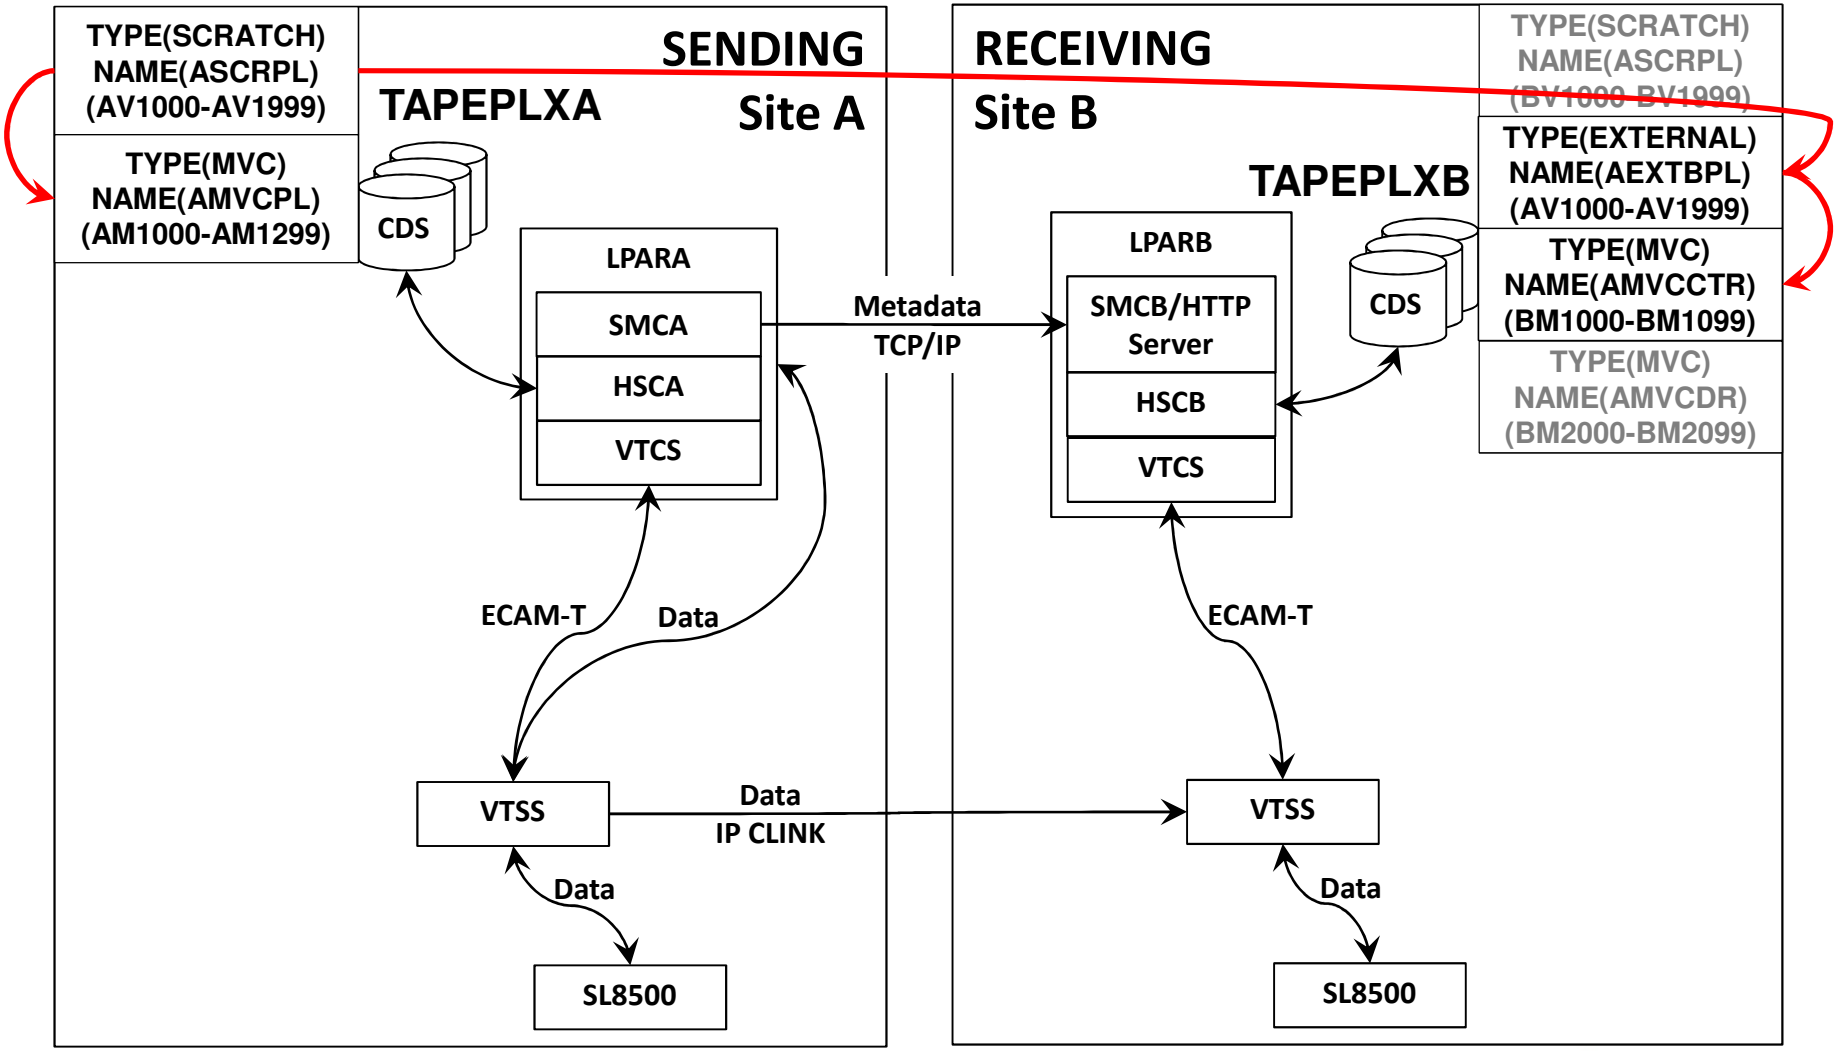

Cross-TapePlex Replication with DR Test

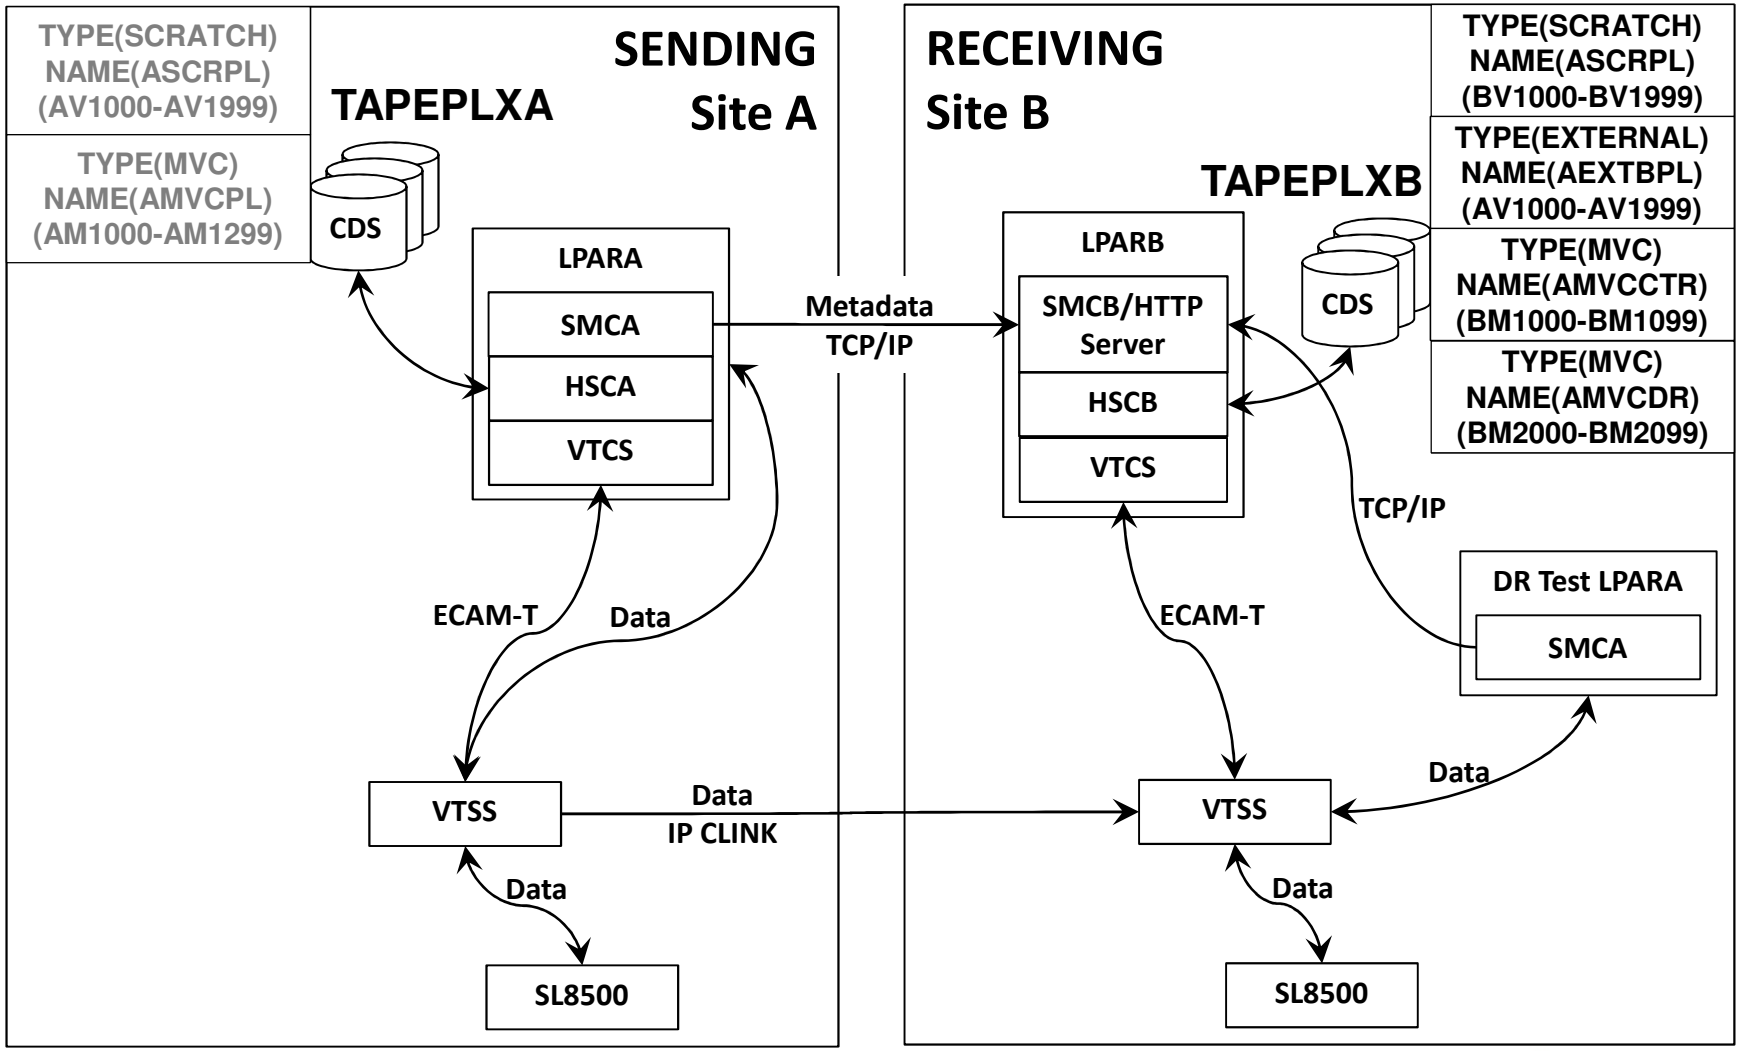

Concurrent Disaster Recovery Test

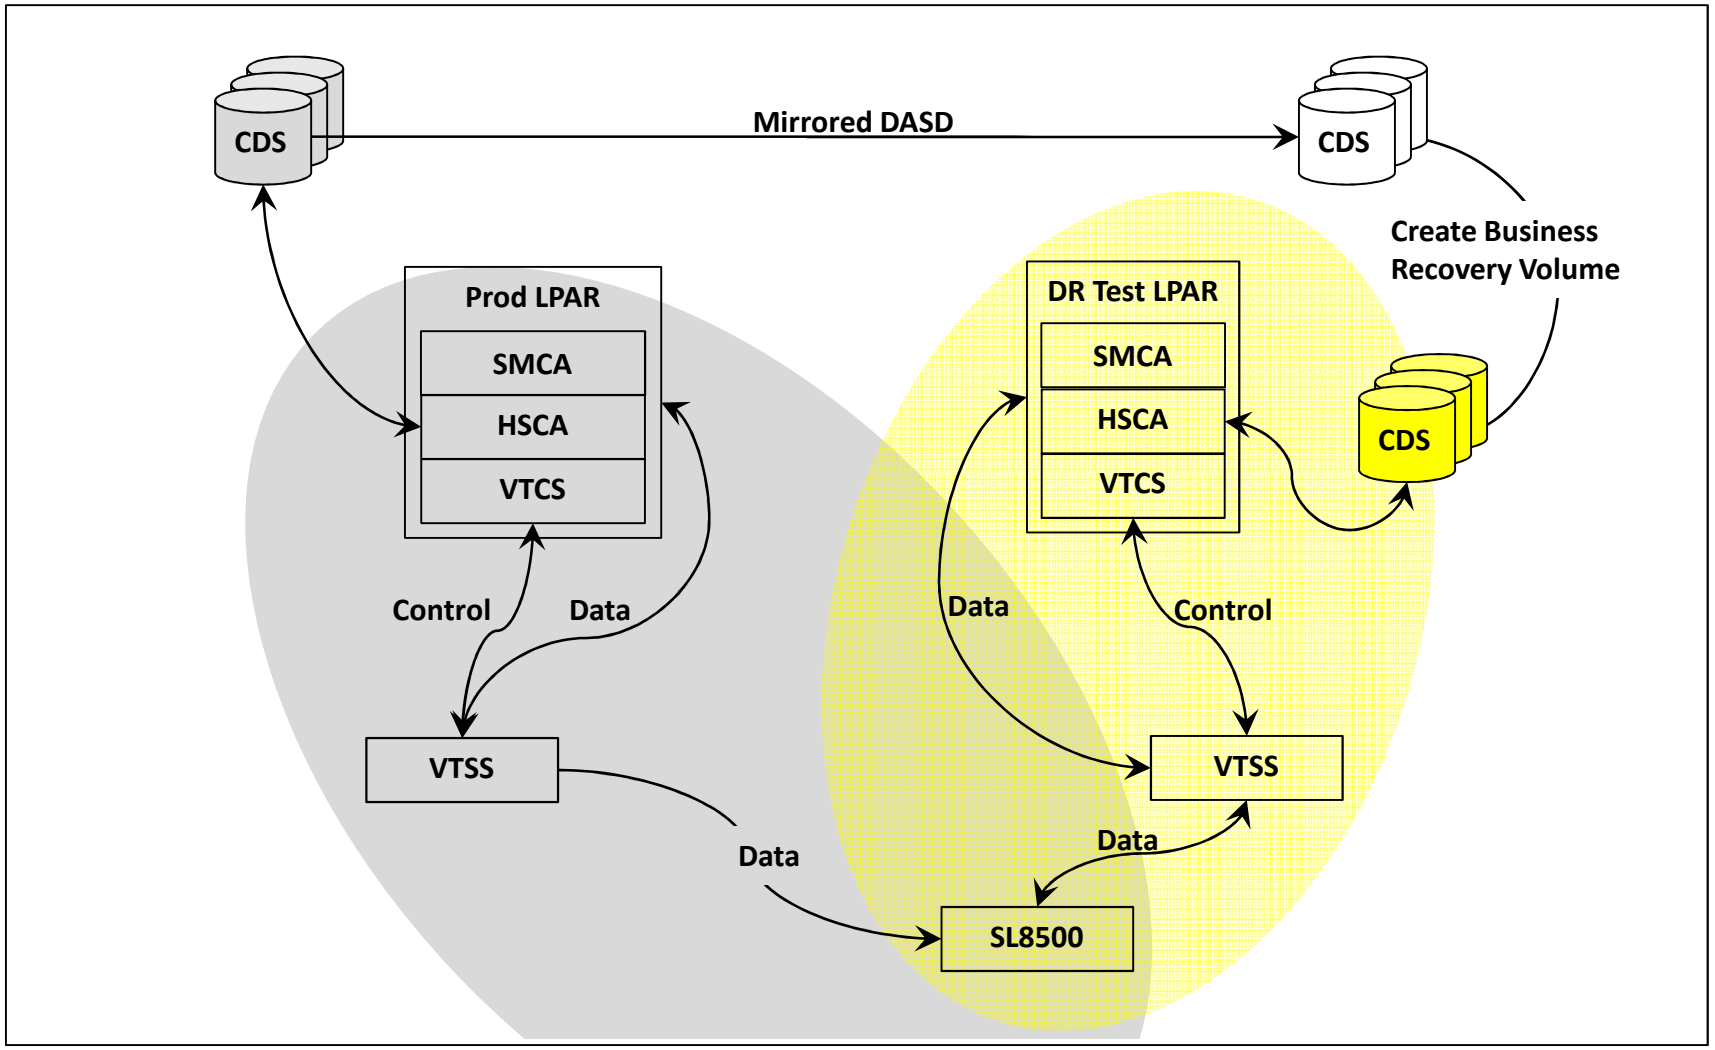

QUESTION TIME