

# Do you want Business Continuity or Disaster Recovery in a Virtualized, Tiered Storage Solution?

Irene Adler  
Oracle Corporation

Monday, August 12, 2013, 3:00pm  
Session 14133

# Agenda



- Business Continuation (BC) Planning
- Business Resumption (BR) Planning
- VSM BC/BR Disaster Recovery Features
- Control Data Set and its Importance
- Clustered VTSS (Virtual Tape Subsystem)
- VSM Cross-TapePlex Replication
- How to safely plan a DR Test in a Production environment

# Business Continuity or Disaster Recovery?

## The answer is . . . both, of course!



- Business considerations in BC/DR planning
  - Different types of data
  - Compliance regulations
  - Budget – need to minimize storage costs
  - Data Growth – how long will my solution last
  - Service level agreements - how soon do I need to recover
  - Data protection – how secure is my data
  - Centralized data management control
- This session will explore some options that are available with Oracle StorageTek's Virtual Storage Manager (VSM) and Virtual Library Extension (VLE) products

# Types of Outages


- DR planning typically addresses one or more of these types of outages, which can be small in nature or large
  - Failure of a storage device
  - Loss of access to virtual volume
    - bad media
    - bad drive
  - Loss of CDS (tape environment metadata repository)
  - Loss of IP communication link to remote site(s)
  - Loss of FICON connection to DR site tape drive
  - Loss of power, cooling or network access
  - Data center facility damage due to:
    - natural disasters, such as earthquake, storm, flooding, etc.
    - other causes, such as fire, vandalism, theft, etc.

# Planning for Success



- Customers must analyze their business requirements
  - Determine data workloads
    - critical – must have high availability
    - important, but not as critical
  - Determine Recovery Time Objective (RTO) for each workload – this is the point in time, after a disruption, that business has to be resumed to full capabilities
  - Determine Recovery Point Objective (RPO) for each workload – this is the point in time, prior to the disruption, that data is known to be complete and consistent, or how far back in time it will take us to recover

# Data Classification



- Current
  - ✓ Frequent changes
  - ✓ Immediate access
  - ✓ Instantaneous protection

- Recent
  - ✓ Infrequent changes
  - Any modifications change the classification to "Current"
  - ✓ Slight access delay acceptable

- Archival
  - ✓ Very infrequent/no changes
  - ✓ Offsite/offline/Nearline protected

# Planning for Success



- Once you categorize your data and determine how long you can allow for recovery in each category, then you can plan how to configure your environment to achieve your goals
- The planning process identifies high-level business continuance (BC) and business resumption (BR) requirements
- Solutions vary in cost, but generally the higher cost solutions allow faster recovery times
- While it is imperative that a catastrophic disaster recovery plan must be in place for long-term outages
- However, just as important is to plan for brief outages

# Temporary Loss of Data – Plan for High Availability



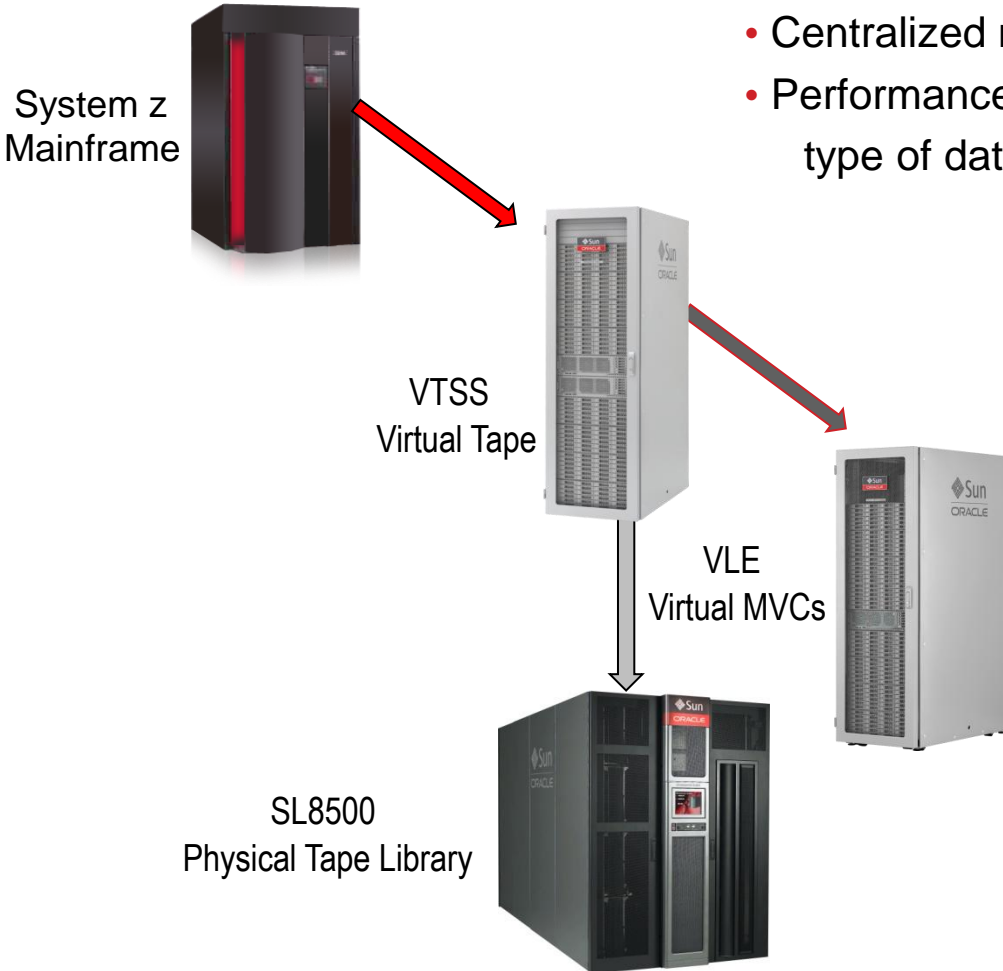
- Requirements:
  - No single point of failure
  - Non-disruptive access to customer data
- Tape virtualization makes high availability possible
- Oracle StorageTek's Virtual Storage Manager (VSM) offers solutions to meet customer requirements
- VSM is a global solution:
  - All Hosts can access all VSM subsystems
  - All VSM subsystems can access all back-end VTVs
- Centrally managed solution via VSM Policies
- Many different VSM options to choose from



# Oracle's StorageTek Virtual Storage Manager (VSM) Tiered Storage



- Provides seamless integration of hardware and software
- Centralized management for all parts of storage system
- Performance and cost of storage closely matched to type of data



- **Tier 1 – High Performance VTSS Disk 10TB – 1.2PB (300PB effective capacity)**
- **Tier 2 – Economical VLE Disk 330TB – 338PB VLE-to-VLE Dedup'd copy**
- **Tier 3 – Tape Up to 1000PB lowest cost, high capacity long-term archiving**

# Hierarchy of Performance, Capacity & Cost



Storage Type	Tier	Perform	Capacity	Cost
VTSS	High Performance Disk	Faster Access ↑	More Capacity ↓	Lower \$\$ Cost ↓
VLE	Economical Disk			
ACS	Archival Tape			

# Temporary Loss of Data – Plan for High Availability



- Requirements:
  - No single point of failure
  - Non-disruptive access to customer data
- Tape virtualization makes high availability possible
- Oracle StorageTek's Virtual Storage Manager (VSM) offers solutions to meet customer requirements
- VSM is a global solution:
  - All Hosts can access all VSM subsystems
  - All VSM subsystems can access all back-end VTVs
- Centrally managed solution via VSM Policies
- Many different VSM options to choose from

# VSM's High Availability Options



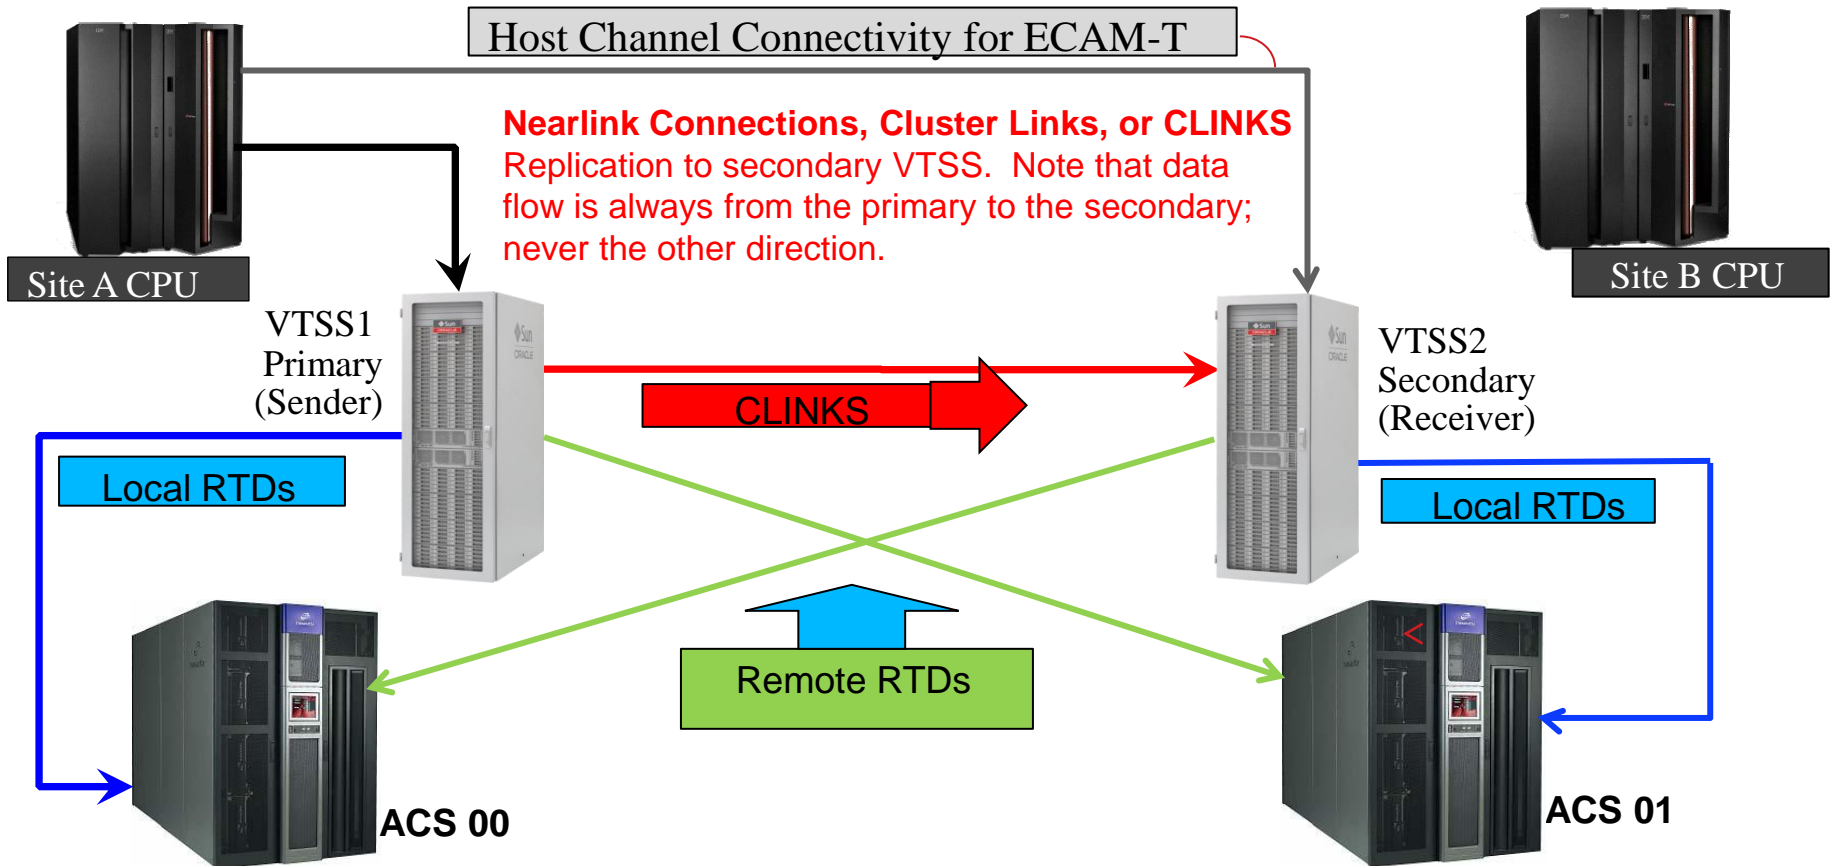
- VTV copies
  - 1 to 4 VTV copies
    - On different MVCs
    - On MVCs in up to 4 different ACS locations
    - On different media
      - for maximum performance
      - high capacity for archive
  - This allows customers to have VTV copies at their production site, DR site, and other site(s) automatically
  - VTSS migrates to VLE and/or real tape via Policy definitions
  - In the event that a VTV cannot be read on one MVC, due to media failure or drive failure, the host software that manages VSM (VTCS - Virtual Tape Control Software) will automatically select the VTV from an another MVC

# VSM's High Availability Option - Clustered VTSS

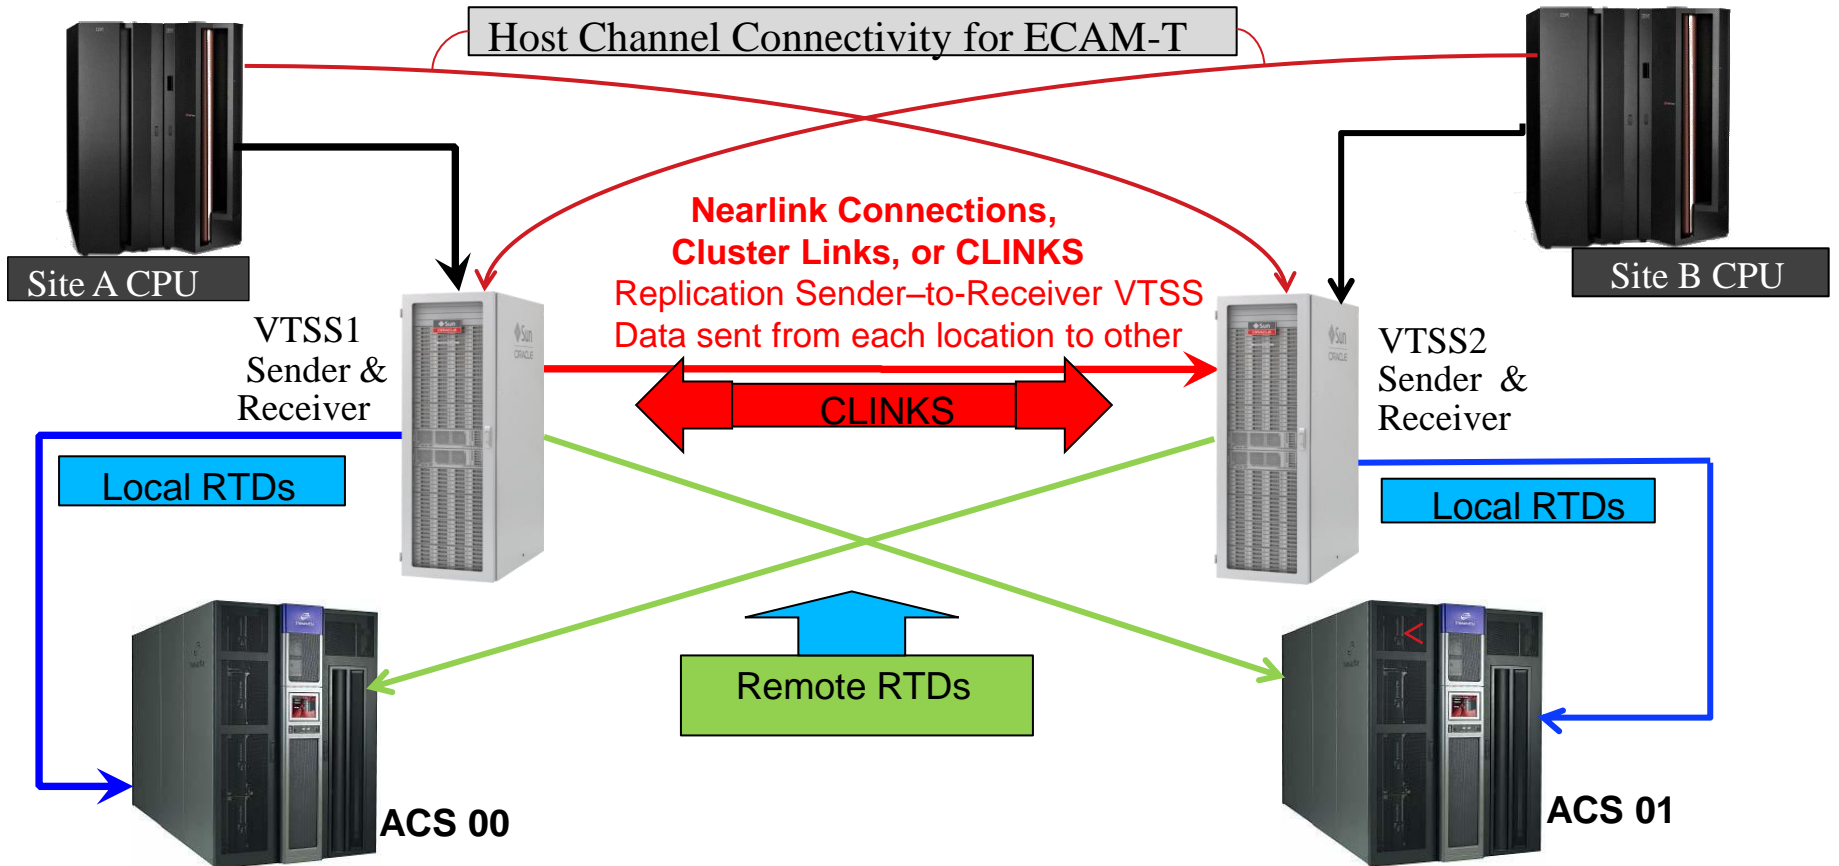


- Clustered VTSS offers the highest level of high availability
  - Virtual volumes are written to a Primary VTSS and then upon dismount are replicated to a Secondary VTSS – either in the same location or at a remote site.
  - Uni-directional – one site replicating to a second site
  - Bi-directional – each site replicating to the other
  - Replication can be:
    - via FICON or IP for VSM5 and IP for VSM6
    - Synchronous or non-synchronous
    - requires no host involvement
    - runs in background
  - Secondary VTSS acts as a warm standby
  - Provides immediate Business Continuance
  - Eliminates single point of failure

# Uni-Directional Clustered VTSS



# Bi-Directional Clustered VTSS



# VSM's High Availability Options (cont'd)



- VTV copies for HA relies on certain other factors to consider:
  - Having a minimum of a duplexed VTV copy
  - Having VTV copy (or copies) offsite electronically via:
    - remote vRTDs in remote VLEs
    - remote RTDs in remote ACSs
    - in same city; preferably in different city or state
- To mitigate possible Cluster connections (CLINKs) failures, consider if you should increase # of FICON or IP connections in Cluster
- To mitigate possible drive failures, consider if you should increase # of RTDs and/or FICON paths to remote site Real Tape Drives (RTDs)
- To allow secondary VTSS to takeover workload from Primary during an outage, ensure sufficient front-end Host channel connectivity is available



# Catastrophic Event Plan for Disaster Recovery



- Plan for long outage or permanent disruption to IT facility
- Ensure VSM data, CDS metadata and other critical system datasets (tape catalog, system/user catalogs) are offsite
- Other VSM features available for DR and DR Test
- Export/Import
- Cross-TapePlex Replication
- Physical Vaulting
- LCM Vaulting Feature
- Concurrent DR Test Software

# How Secure is the Control Data Set (CDS)



- In the event that a catastrophic failure would cause the CDS to be destroyed or become corrupted, the time required to restore the CDS and to bring the transaction activity current would be a monumental and a time-consuming task
- The CDS contains all the metadata for the tape environment, both physical and virtual
  - Must have all 3 copies of CDS available at all times
    - Primary, Secondary and Standby
    - If a problem occurs, switch to secondary is automatic
- With VSM, the high activity of updates to the CDS requires it be backed up at frequent intervals
  - Some customers back it up every 8 hours, some every 6 hours, others more often

# Synchronize Your Data Capture



- How to ensure data synchronization matches CDS metadata:
  - Ensure all Disaster Recovery critical data has been migrated and is secured at the customer's recovery location
  - Establish system checkpoints in order to establish a recovery point
  - Ensure that DR data remains consistent with metadata until the next checkpoint
  - The metadata copy must be valid and complete when a disaster is declared, either real or test
  - Allow the VSM customer to create CDS backups that reflect checkpoints of the VSM environment with a consistent set of DR data at a given point in time
  - The CDS backup, plus other backups (e.g. Tape Management Catalog, System Catalogs) then form a checkpoint of the entire system
  - VTCS DR Synchronization is a feature that can provide all of these

# VTCS DR Synchronization Feature



- VTCS DR Synchronization is a feature which will:
  - Monitor and wait for DR data to reach its recovery location
  - Then sets a Checkpoint in the CDS
  - Then a CDS backup should be taken
  - Meanwhile, don't allow overwrites on MVCs until the next DR baseline is established – this is protected by the utility software
  - In the event of a major outage, you now have a definite point to restore back to and then recovery can begin knowing your data is synchronized with the CDS metadata

# 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



- ELS External Vaulting Feature replaces the VSM Offsite Vault Feature
- Uses the CDS to store vault and vaulted volume information instead of using the customer's tape management system
- Off-site physical vaulting of DR tapes can be accomplished using LCM (Library Content Manager) to manage the vaulting process
- 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

# Cross-TapePlex Replication (CTR)



- Cross-TapePlex Replication operates as an electronic export and electronic import of VTV data into a separate TapePlex, with its own CDS configuration, via FICON or IP CLINKs for VSM5, or IP CLINKs for VSM6
- Metadata is transmitted from sending to receiving TapePlex over IP using SMC HTTP
- Originating TapePlex retains ownership of the VTV
- Receiving TapePlex cannot update or scratch VTV
- When VTV is scratched at originating TapePlex, all copies will be scratched
- **No VTSS connectivity to remote RTDs are required or allowed**
- **No Front-end Host channel connectivity to the receiving TapePlex**

# Cross-TapePlex Replication (Cont'd)



- Management Class is defined at data creation time and is part of the metadata that is exported with the data.
- Provides additional data copies for cascading unlimited number of data copies within TapePlexes and across TapePlexes by the receiving TapePlex's migration policies
- Creating TapePlex can replicate the data to up to two separate TapePlex systems; one copy may be synchronous
- VTV migration and buffer management are done independently by each TapePlex
- DR tests can be run without impacting or disrupting the production environment

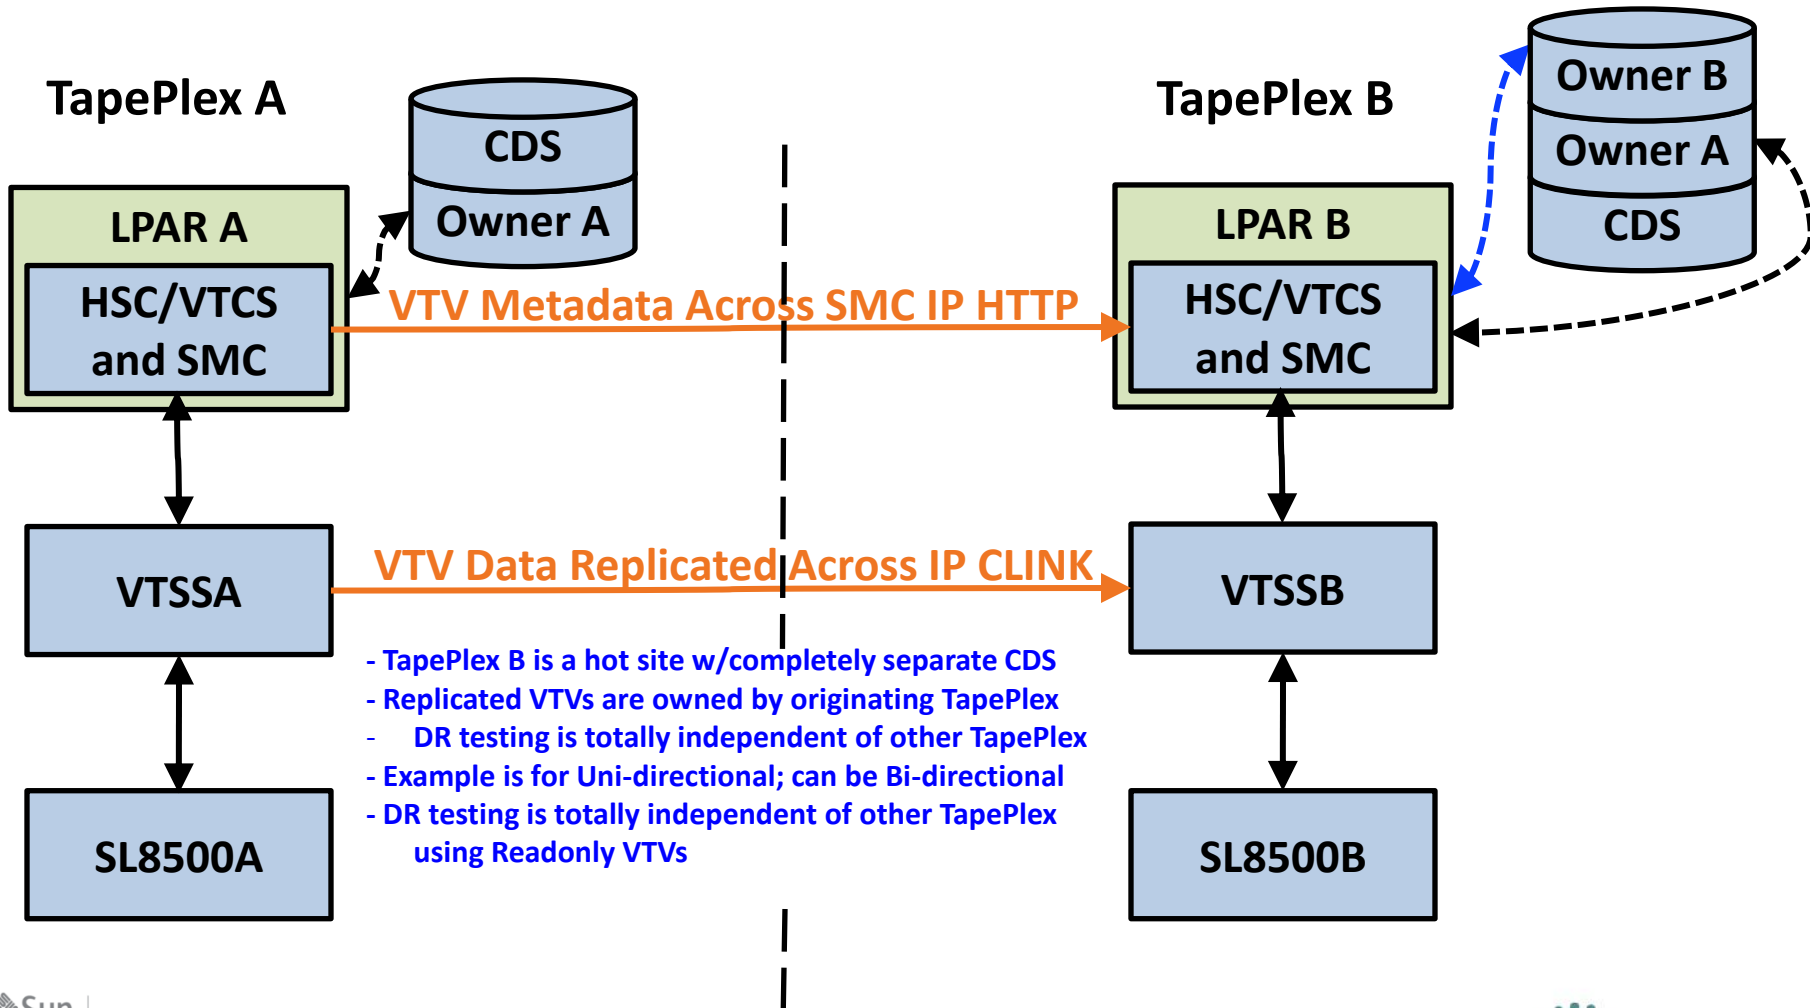


# Cross-TapePlex Replication (Cont'd)



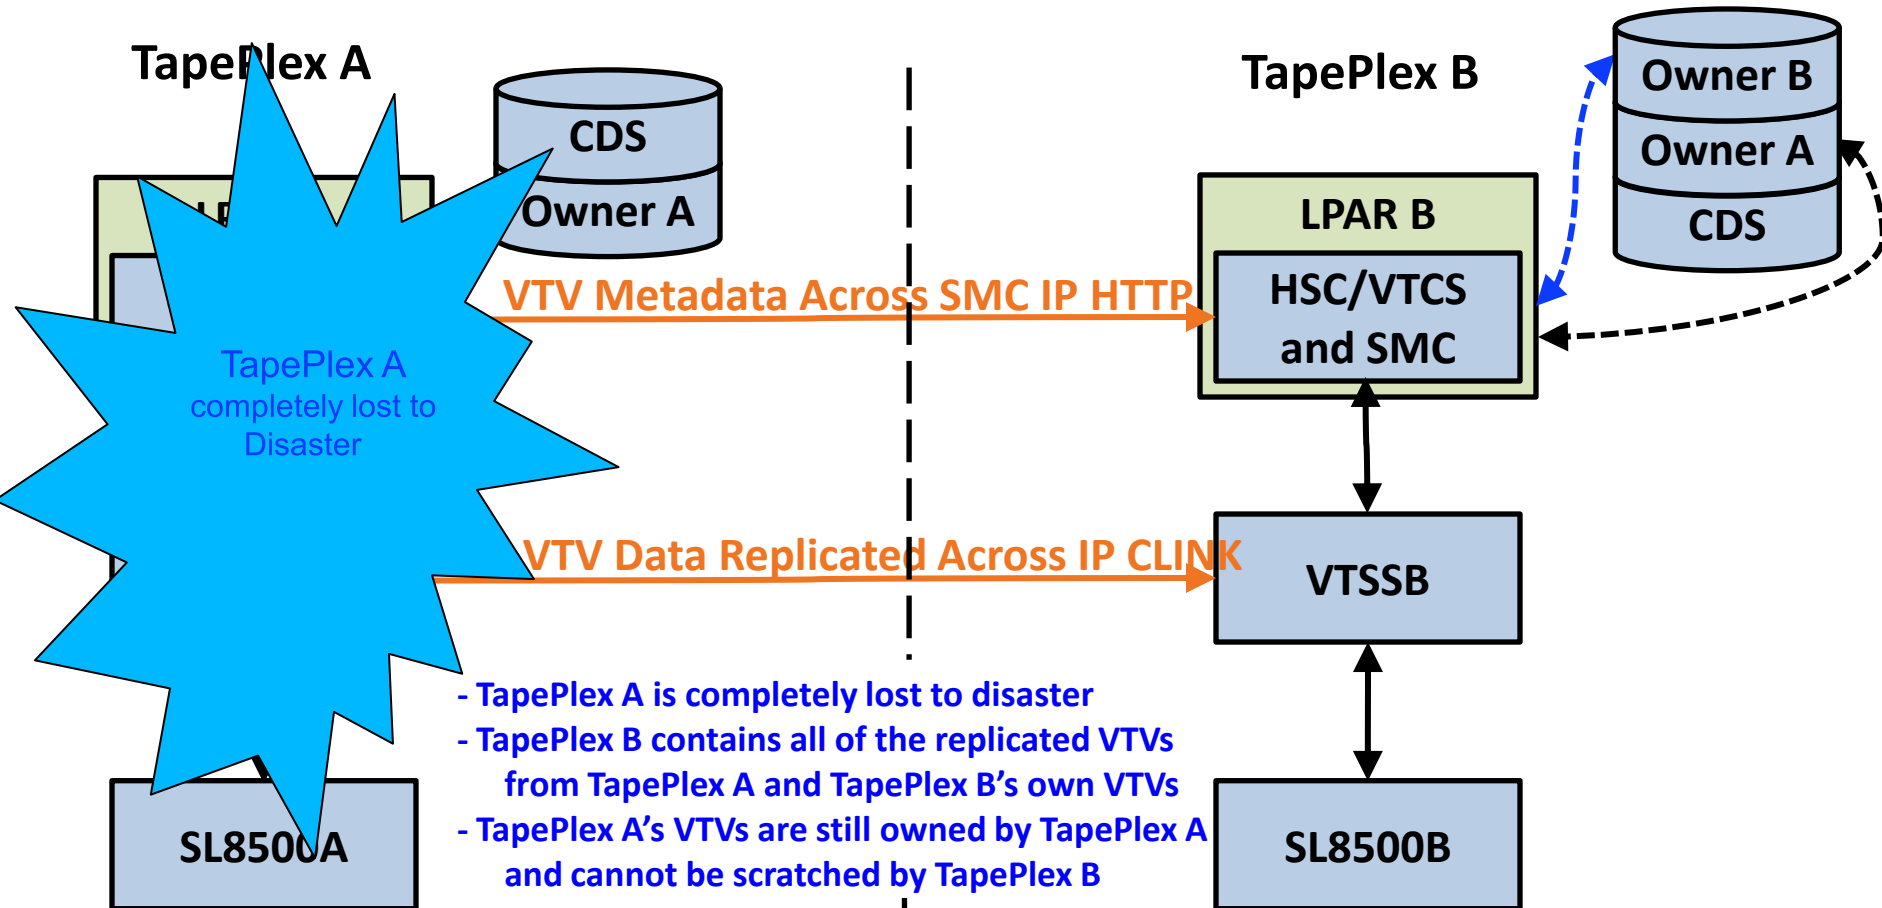
- If a VTV is needed to be sent back to the owning TapePlex, the electronic export *automatically* recalls the VTV back to the originating TapePlex
- The EEXPORT Utility can also be used to bring the VTV back to the originating TapePlex manually
- CTR Requirements:
  - VTV ranges for the replicated VTVs in both TapePlexes must be the same
  - No overlapping volsers allowed
  - Management Class names for the replicated VTVs must be identical in both TapePlexes

# CTR Uni-directional Configuration Normal Operations



# CTR Uni-directional Configuration


## Complete Loss of TapePlex A



# CTR Uni-directional Configuration

## Business Resumption after Loss of TapePlex A

1. Initially, VTVs owned by LPAR A are protected against being scratched
  - This is because they are still owned by TapePlex A
2. During an extended outage, VTVs owned by LPAR A can be changed to be owned by TapePlex B



# Concurrent Disaster Recovery Test (CDRT) Utility



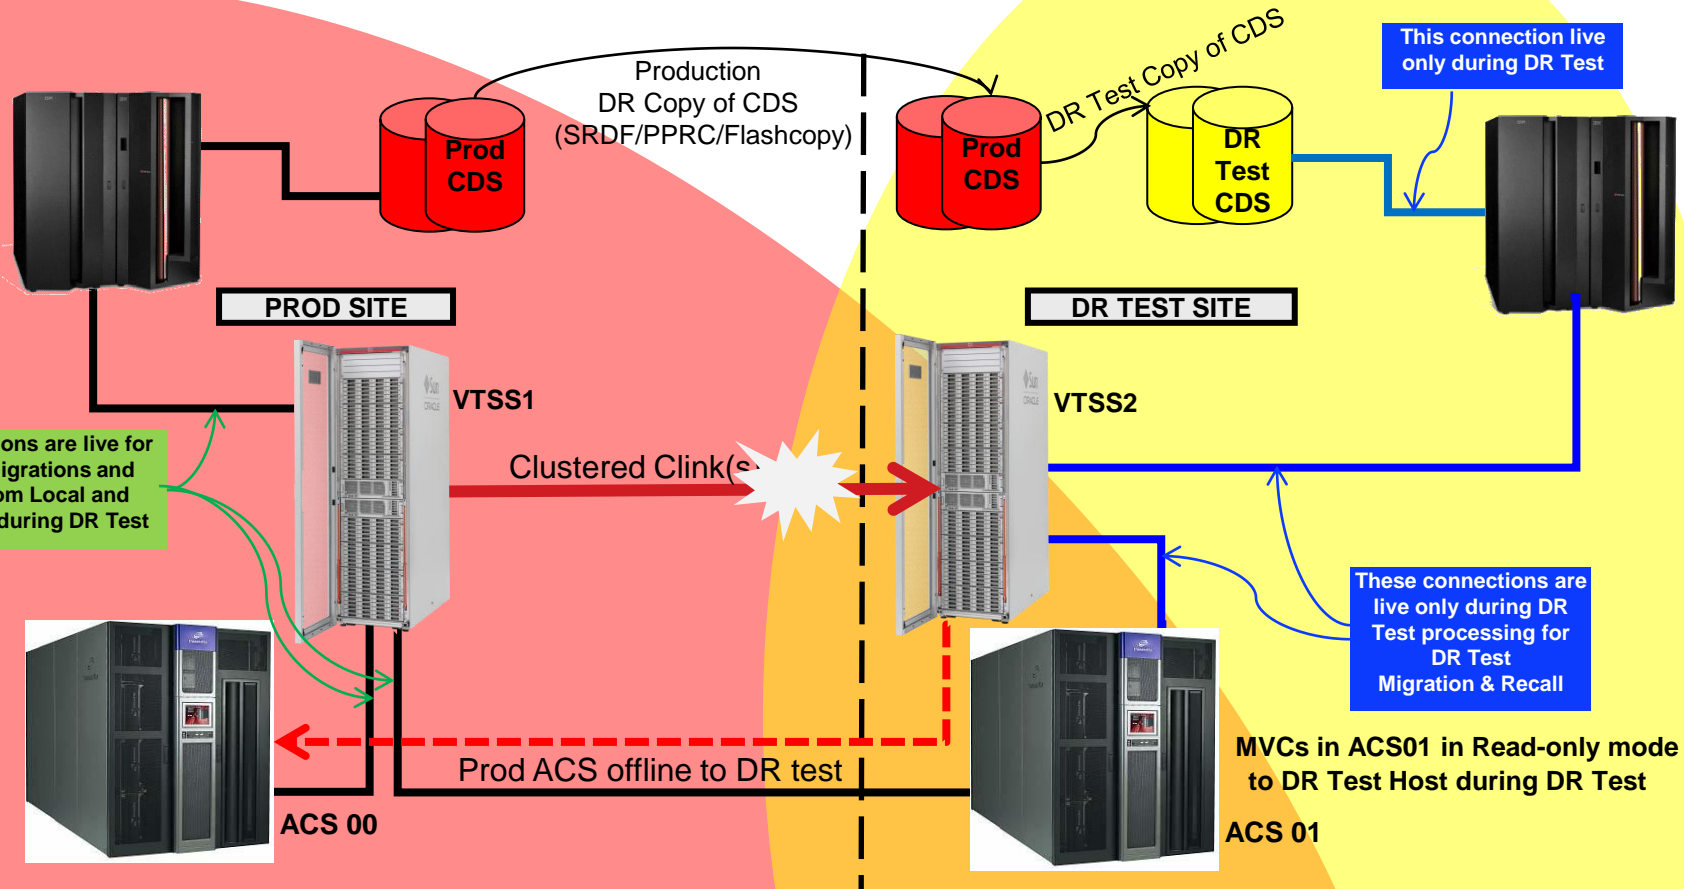
- CDRT Utility allows the Virtual Storage Manager system to run with 2 sets of CDSs
- Can be used in single TapePlex Clustered or non-Clustered configurations
- Production and DR Test can run concurrently
- Production CDS is primed and copied whenever utility is STARTed
- Restrictions & controls are put in place when the START keyword is used
- Restrictions and controls are removed when the STOP keyword is used
- CDRT Utility and Commands provided:
  - DRTEST Utility – “Primes”, Starts, Stops and “RESETs”
  - DRTEST Command – to Start and Stop
- Production & DR Test run concurrently
- Production VTVs & MVCs are frozen from the DRTEST at Start
  - DR Test runs from a “frozen” pool of VTVs and MVCs that are in readonly mode

# Concurrent Disaster Recovery Test (CDRT) (Cont'd)



- DR Test Scratch Subpool(s) must be defined to be used during the test
- Typically two Output MVCPools are used
- Certain restrictions are set and enforced by the CDRT software during the test
  - DR Test ACS is disconnected to Production Host(s)
  - Production VTSS is offline to DR Test Host(s)
  - Production ACS is disconnected from DR Test Host(s)
  - Set in both Production and DR Test
    - Float(off) in ACSs
    - No Enters, Ejects, Moves, Audits or Scratch Redistributions

# CDRT w/Clustered Configuration and Mirrored CDS



# In Conclusion . . .



- Careful planning is essential for Business Continuity and Disaster Recovery
- Explore the various options
- Consider what DR features VSM has available
- Pick the solution that meets your business requirements
- Ensure critical data is backed up often and resides off-site
- Test your DR plan
- Make sure your plan includes having a detailed, written document to be followed by your IT staff in the event of an outage – don't wait until it happens



*Thank you . . .*

*Questions ?*