Using Oracle's Virtual Storage Manager (VSM) for Cost-Effective Disaster Recovery

Piotr Polanowski
Nancy Rassbach
Oracle

August 5, 2010
Session Number 8070

SHARE in Boston

Agenda

- Overview - Disaster Recovery
- Real Tape and Disaster Recovery
- Virtual Tape and Disaster Recovery
- Managing multiple sites with VSM

Disaster Recovery Overview

- A disaster recovery strategy has three basic components:
  - Getting the data to the disaster recovery site
  - Testing the disaster recovery plan
  - Deploying the disaster recovery plan in case of a real disaster
- And don’t forget:
  - Returning to normal operations after a disaster

Disaster Recovery Considerations

- Recovery Time Objective (RTO) and Recovery Point Objective (RPO) are critical part of disaster recovery planning
  - RTO is influenced by availability of data at the DR site
  - RPO is determined by how data is transferred to the DR site
- Just as for other parts of the DR solution, better RTO and RPO generally translate to more cost
- Tape is a part of a complete Disaster Recovery planning strategy.
  - Relies on Catalogs, Tape Management Systems and other infrastructure for the overall DR solution

Disaster Recovery Testing

- Testing the disaster recovery plan is critical
  - Without the ability to test, there is no DR solution
  - The closer that the DR test mimics what would happen in a real disaster, the better the test
- DR testing can be non-disruptive or disruptive, depending on the resources available to the test

Real Tape DR Configurations

Physical Vaulting

Electronic Vaulting

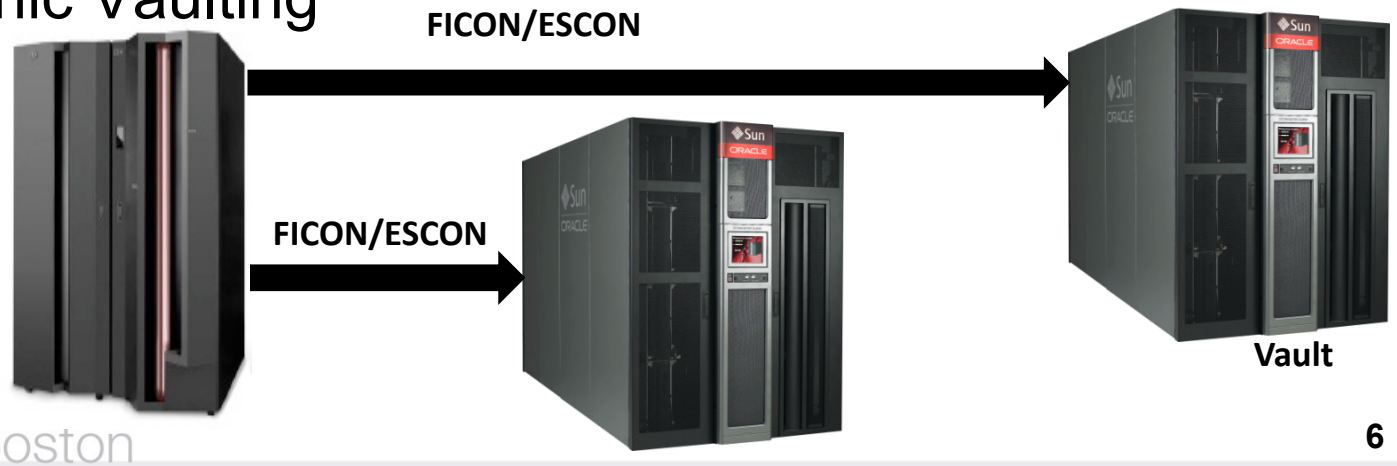

Real Tape Disaster Recovery

- Requires customer applications to perform tape copies of files needed for DR
- Physical vaulting - tape copies shipped offsite periodically
  - Directly to the DR site
  - To a vault site to be transported to the DR site if needed
- Electronic vaulting - tape copies are sent to the DR site using channel extension from the primary site
  - Could be within the same TapePlex (but a different physical location) as the production copy
  - Could be a separate TapePlex, managed by a separate SMC at the DR site using client/server

Real Tape DR Testing

- Physical vaulting
  - If library available, bring up HSC to access remote ACS; otherwise, mount tapes manually
- Electronic vaulting, same TapePlex
  - Can use Concurrent Disaster Recovery Test utility to access tapes in remote ACS from the DR site
- Electronic vaulting, different TapePlex
  - Can run a DR test at any time using tapes copied from production

Real Tape DR Testing Restrictions

- DR testing with real tape assumes no modification to the tape volumes
- If tapes are modified during the DR test, then they are invalid for use in a “real” disaster
  - This restriction implies that the DR test must not MOD to the end of the production tapes, and must not overwrite the production tapes
- The restriction is the same regardless of the method for getting the tapes to the remote site

DISP=MOD and Tape Data

- In general, the use of DISP=MOD for critical tape data sets is not recommended
  - If a job using DISP=MOD fails to complete, the status of the tape is unknown. The job cannot be restarted without manual intervention
- Alternative recommendations
  - Use GDGALL or concatenate input
  - For HSM volumes, mark as full on dismount
- DISP=MOD is not an issue provided that restart procedures always start with the files in a known state

Real Tape Real Disaster Recovery

- Recovery from a real disaster with real tape is very close to the DR test procedure
- Tapes are moved to the DR site, if necessary
- Processing continues at the DR site using the copied tapes

Virtual Tape Disaster Recovery

- VTCS supports equivalent processes to real tapes
  - Offsite MVCs – Shipping physical MVCs to a remote location
  - Multiple VTV copies to channel-connected remote ACS(s)
  - Clustered VTSSs - send VTVs directly to a second site
  - Cross TapePlex Replication - send VTVs directly to a different TapePlex at a second site
- Virtual tape eliminates the need for the application to perform tape copy functions
  - VTCS policies control the number of copies on real tape or replicated to another VTSS (no mainframe cycles needed)

Offsite MVCs With Vaulting

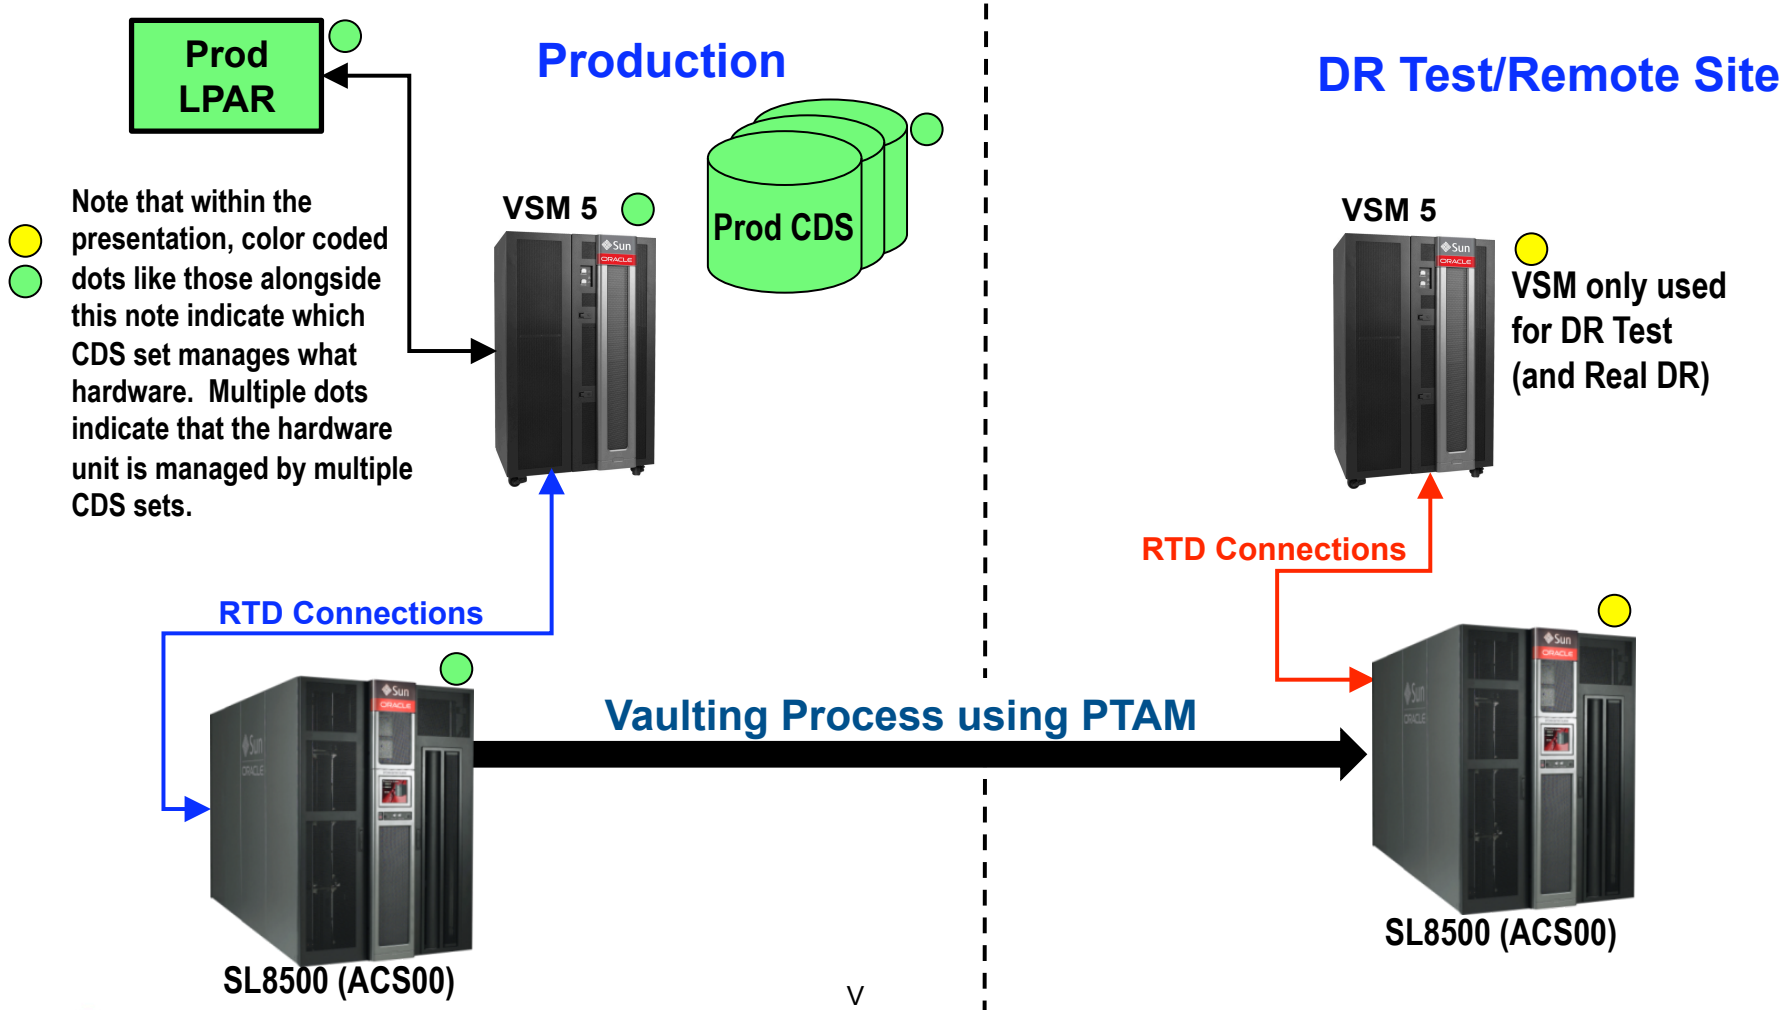

Offsite MVCs With Vaulting – DR Test

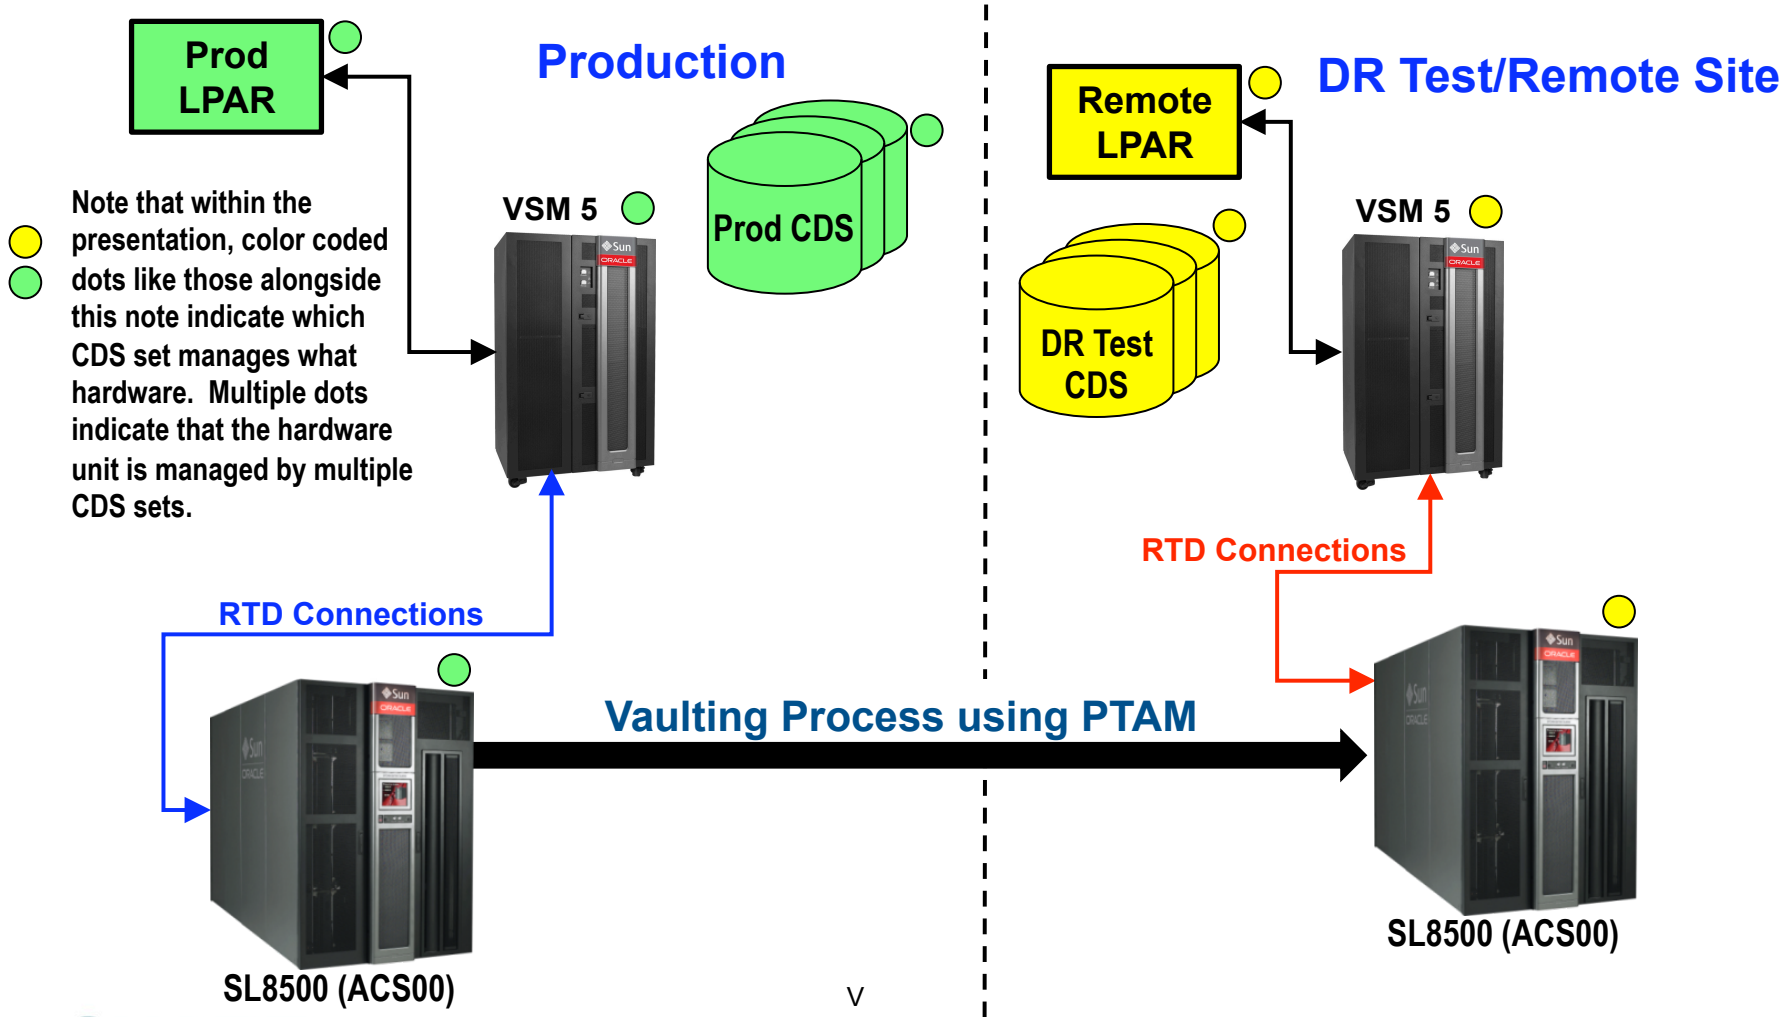

Offsite MVCs
DR Testing

- DR testing with offsite MVCs requires bringing up HSC/VTCS at a remote site with VSM and library hardware
  - No special software requirements
- Restrictions:
  - No restriction on changing VTVs, as VTV copies in the buffer will be deleted after the test
  - Changing MVCs is allowed as long as reclaim is not done - the production CDS will not know about any data added to the MVC
  - DR hardware must not be defined to production system
- Recovery from a real disaster is the same process as DR test

Multiple VTV Copies to Remote ACS Non Clustered VSM

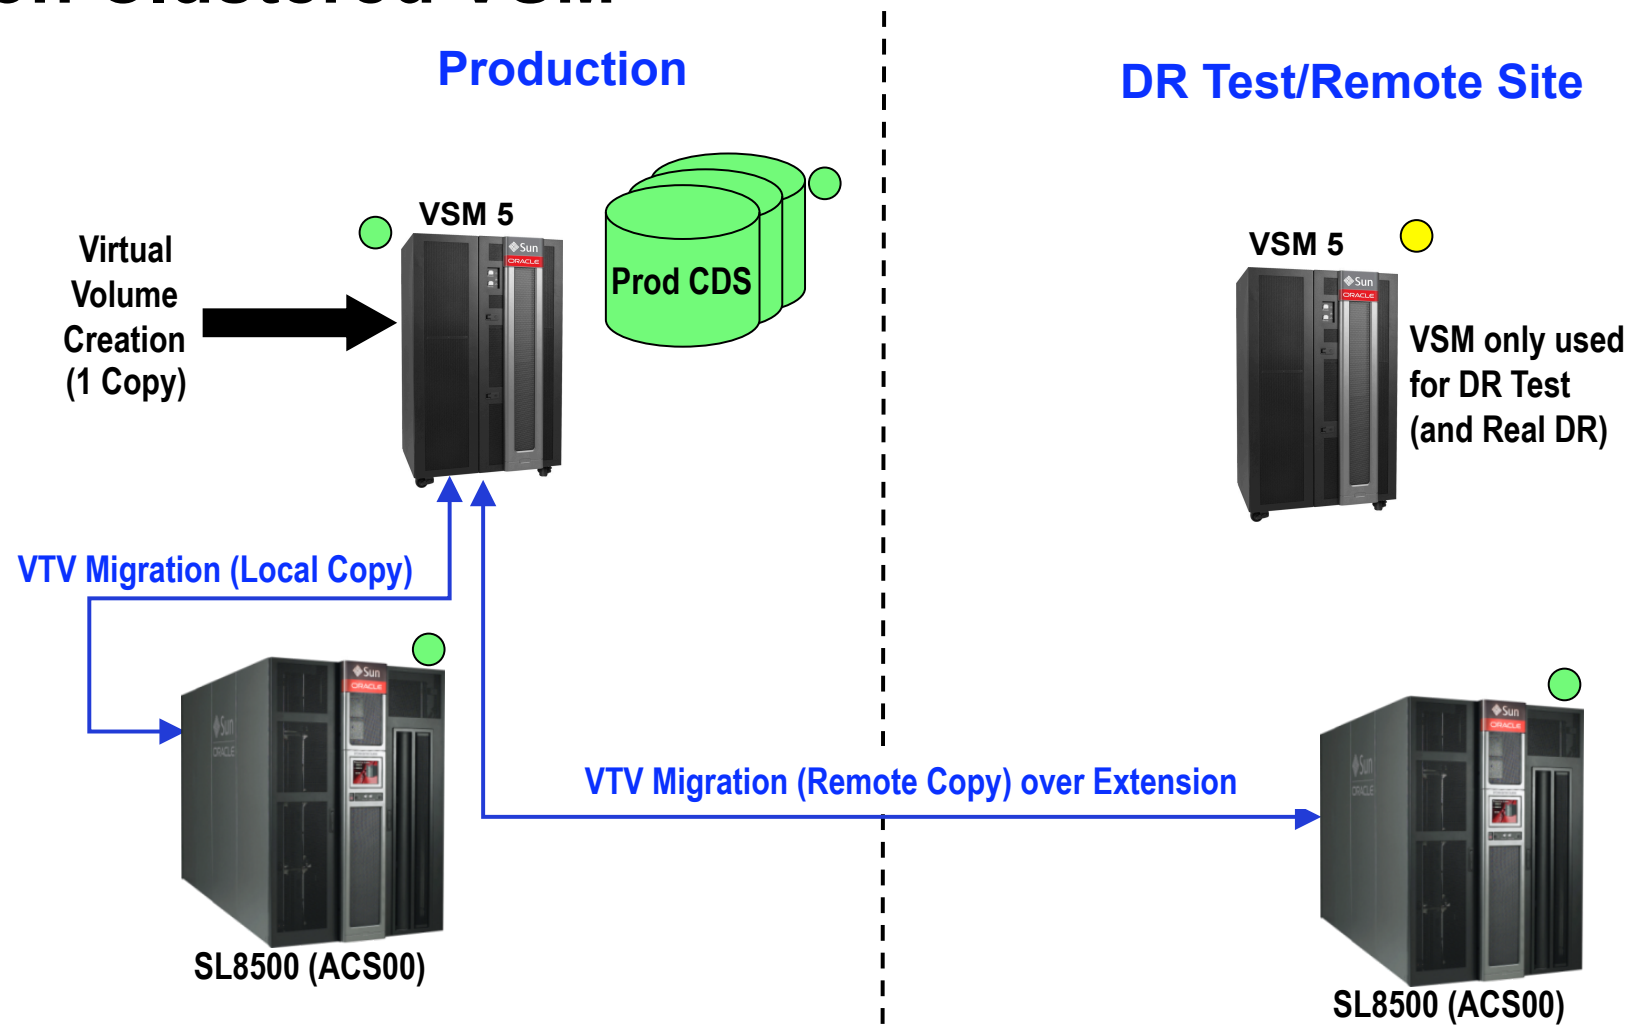

Multiple VTV Copies to Remote ACS Clustered VSM

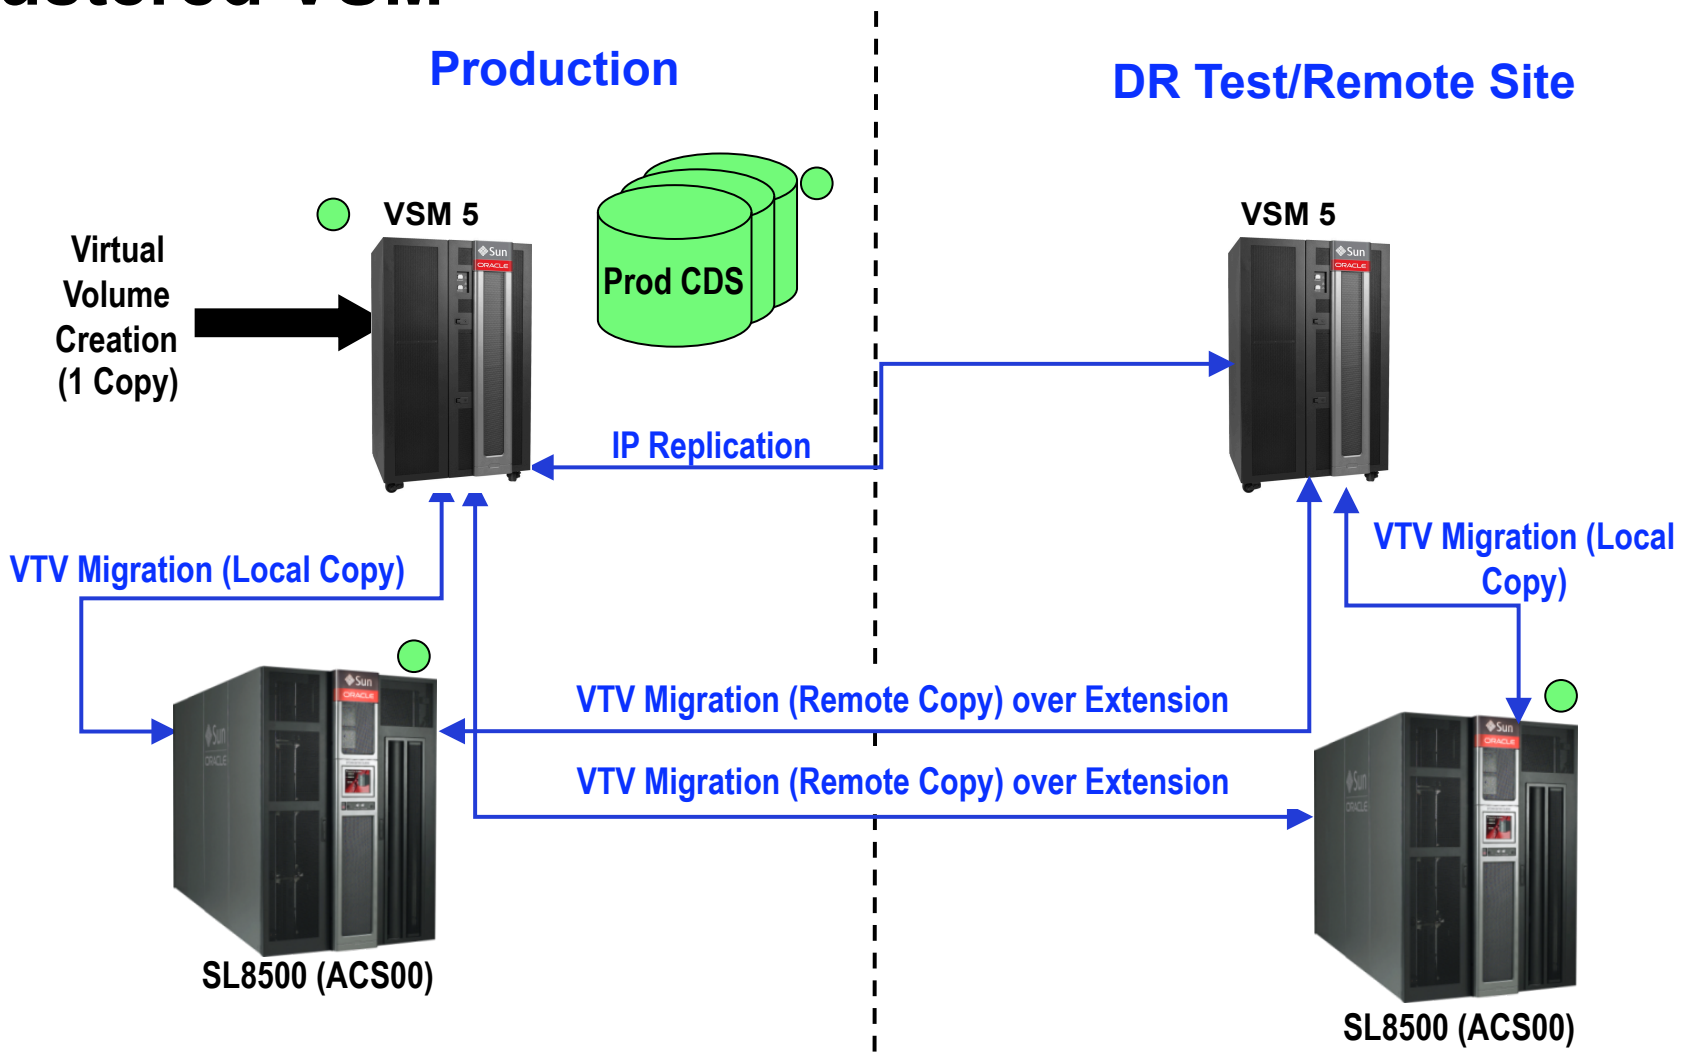

Multiple VTV Copies to Remote ACS & Clustered VSM - DR Test

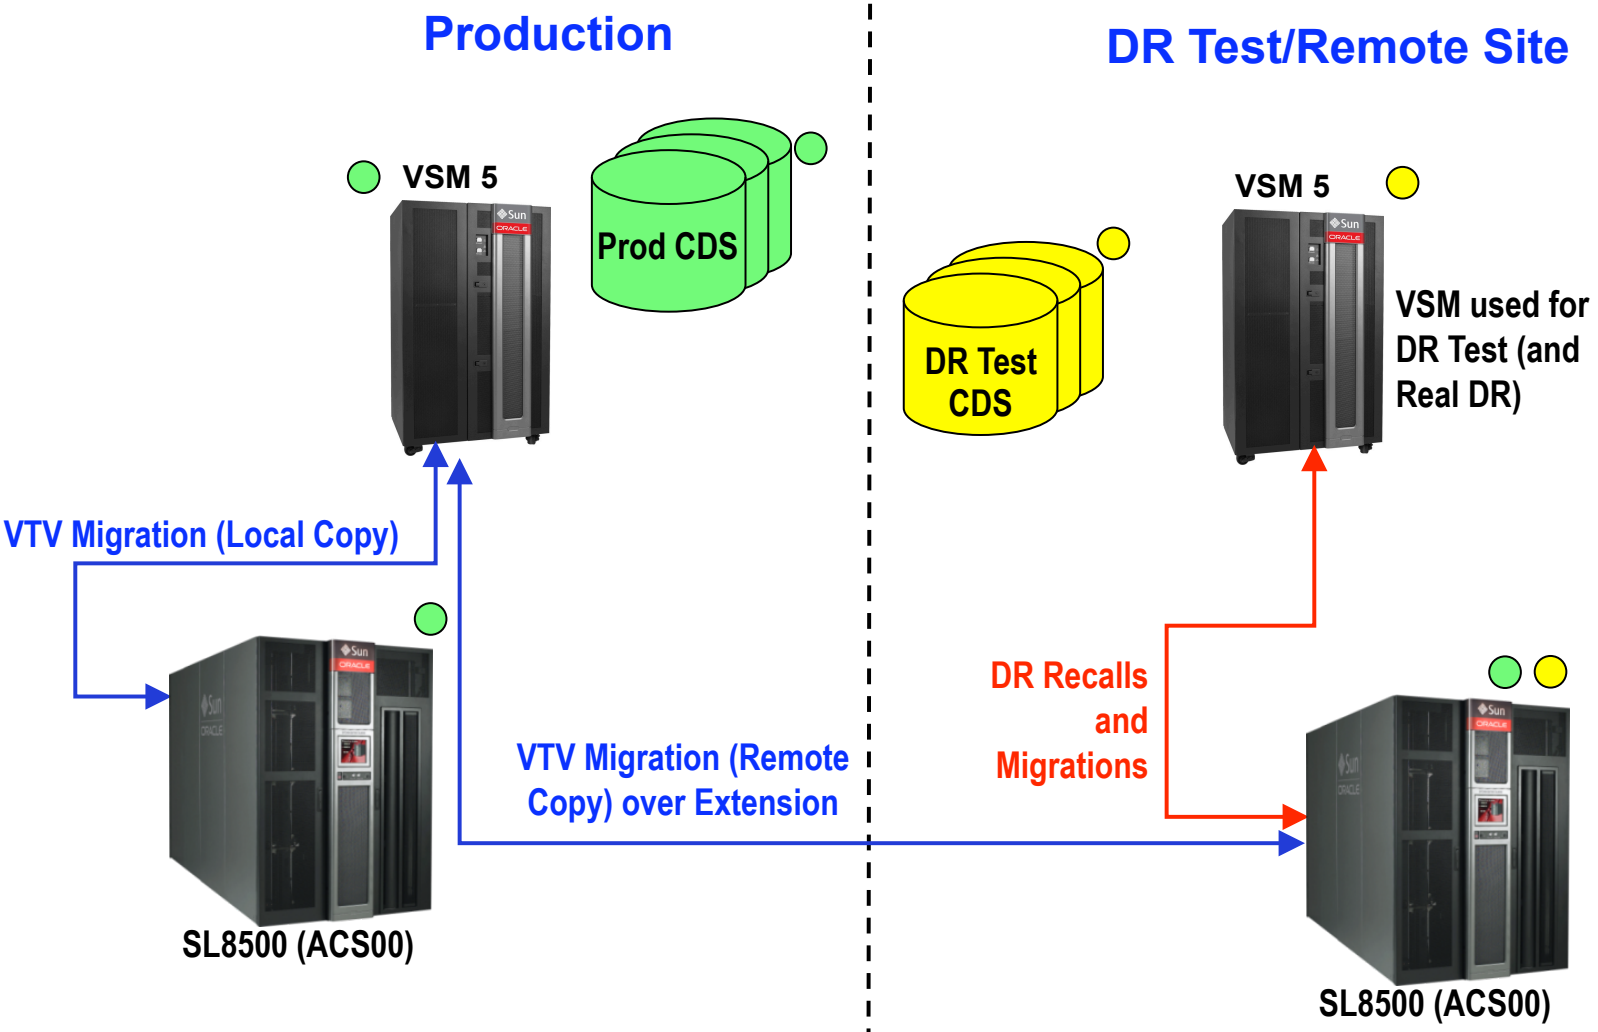

Remote ACS or Clustered VTSS
DR Testing

- DR testing with offsite ACS (or VTSS to which VTV copies are directed) can use the CDRT utility for DR testing
- CDRT imposes several restrictions on DR site hardware :
  - The DR VTSS in clustered configuration must be offline to production during the test
    - Replication is temporarily stopped, or a separate VTSS must be available to devote to DR testing.
  - Changing VTVs during a DR test is allowed but not recommended. If VTVs are changed during the test, then cleanup after the test must be done carefully to avoid data loss.
  - Production MVCs cannot be altered during the test (new ones are created if necessary)

Cross TapePlex Replication (CTR) All Connections

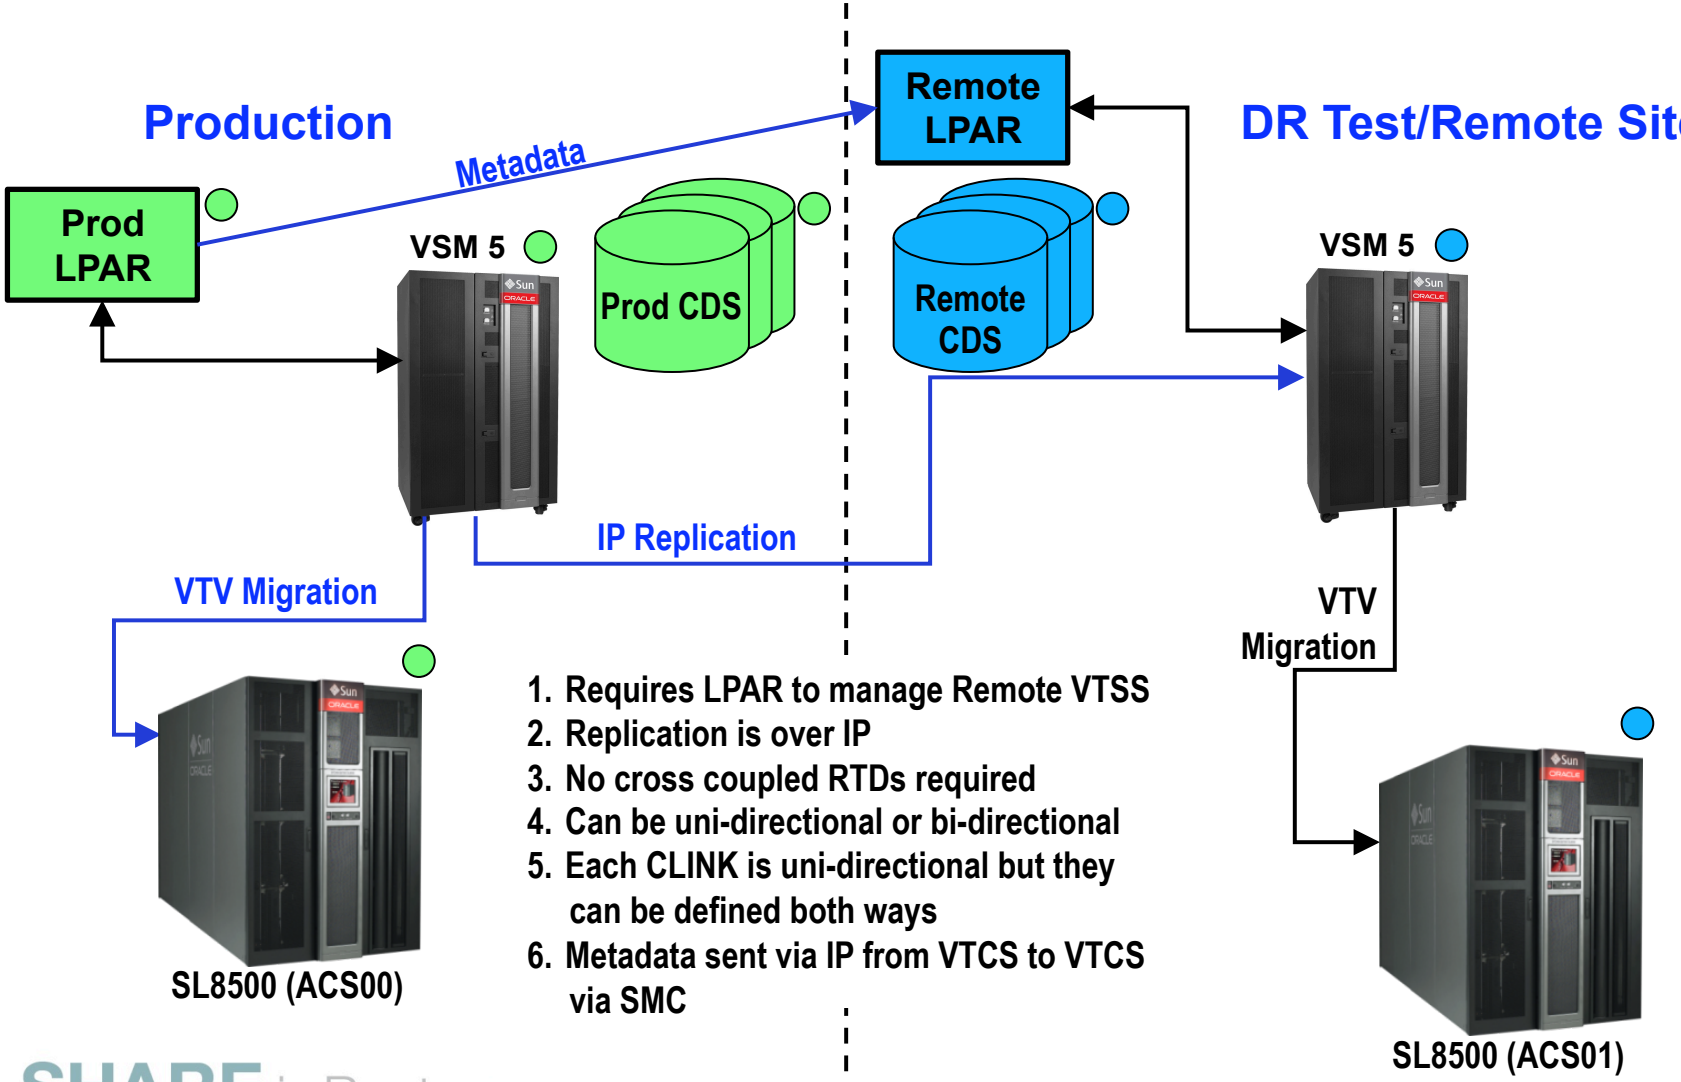

1. Requires LPAR to manage Remote VTSS
2. Replication is over IP
3. No cross coupled RTDs required
4. Can be uni-directional or bi-directional
5. Each CLINK is uni-directional but they can be defined both ways
6. Metadata sent via IP from VTCS to VTCS via SMC

Cross TapePlex Replication (CTR or Electronic Export)

- Cross TapePlex Replication lets you manage multiple copies of data in multiple locations
- Data is replicated between two VTSSs using a FICON or IP CLINK
- Metadata about the replicated VTVs is sent from sending TapePlex to receiving TapePlex using SMC TCP/IP feature
- VTV metadata includes the MGMTclas
  - The receiving TapePlex uses its local MGMTclas definitions to manage VTV copies of the replicated data
- VTVs replicated to another TapePlex are flagged as “owned” by the sending TapePlex, and can only be scratched or updated by the owner

Cross TapePlex Replication (CTR)
DR Test

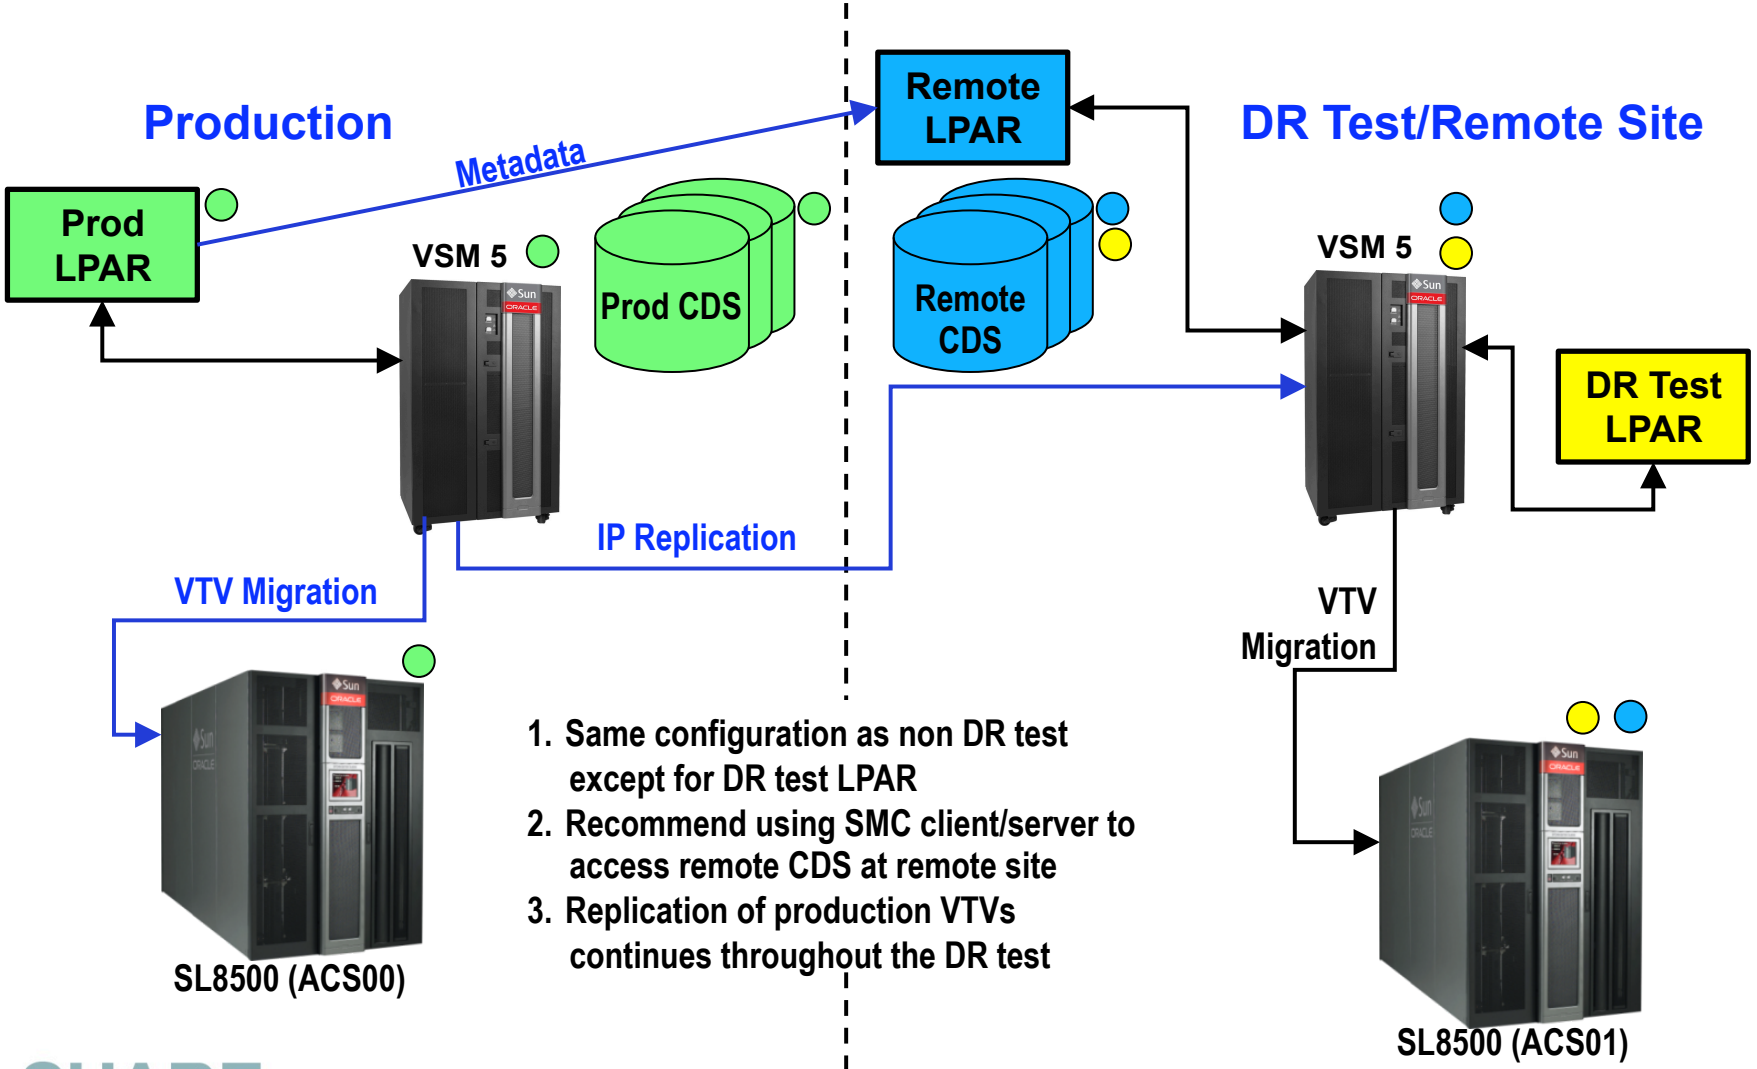

1. Same configuration as non DR test except for DR test LPAR
2. Recommend using SMC client/server to access remote CDS at remote site
3. Replication of production VTVs continues throughout the DR test

CTR: DR Testing

- Parallel with scenario above using real, customer-copied volumes instead of virtual volumes
- No special setup for DR testing
- Access replicated VTVs using the DR TapePlex
- Restrictions:
  - Cannot modify VTVs replicated from another TapePlex
  - Must ensure no scratch or update of production volumes needed for the DR test
  - Separate scratch and MVC pools for DR testing

Managing Multiple Sites

- The Cross TapePlex Replication is “many-to-many”
- Data can be received from any number of other TapePlexes
- The receiving site can replicate to additional TapePlexes

Multiple Site Cross TapePlex Replication

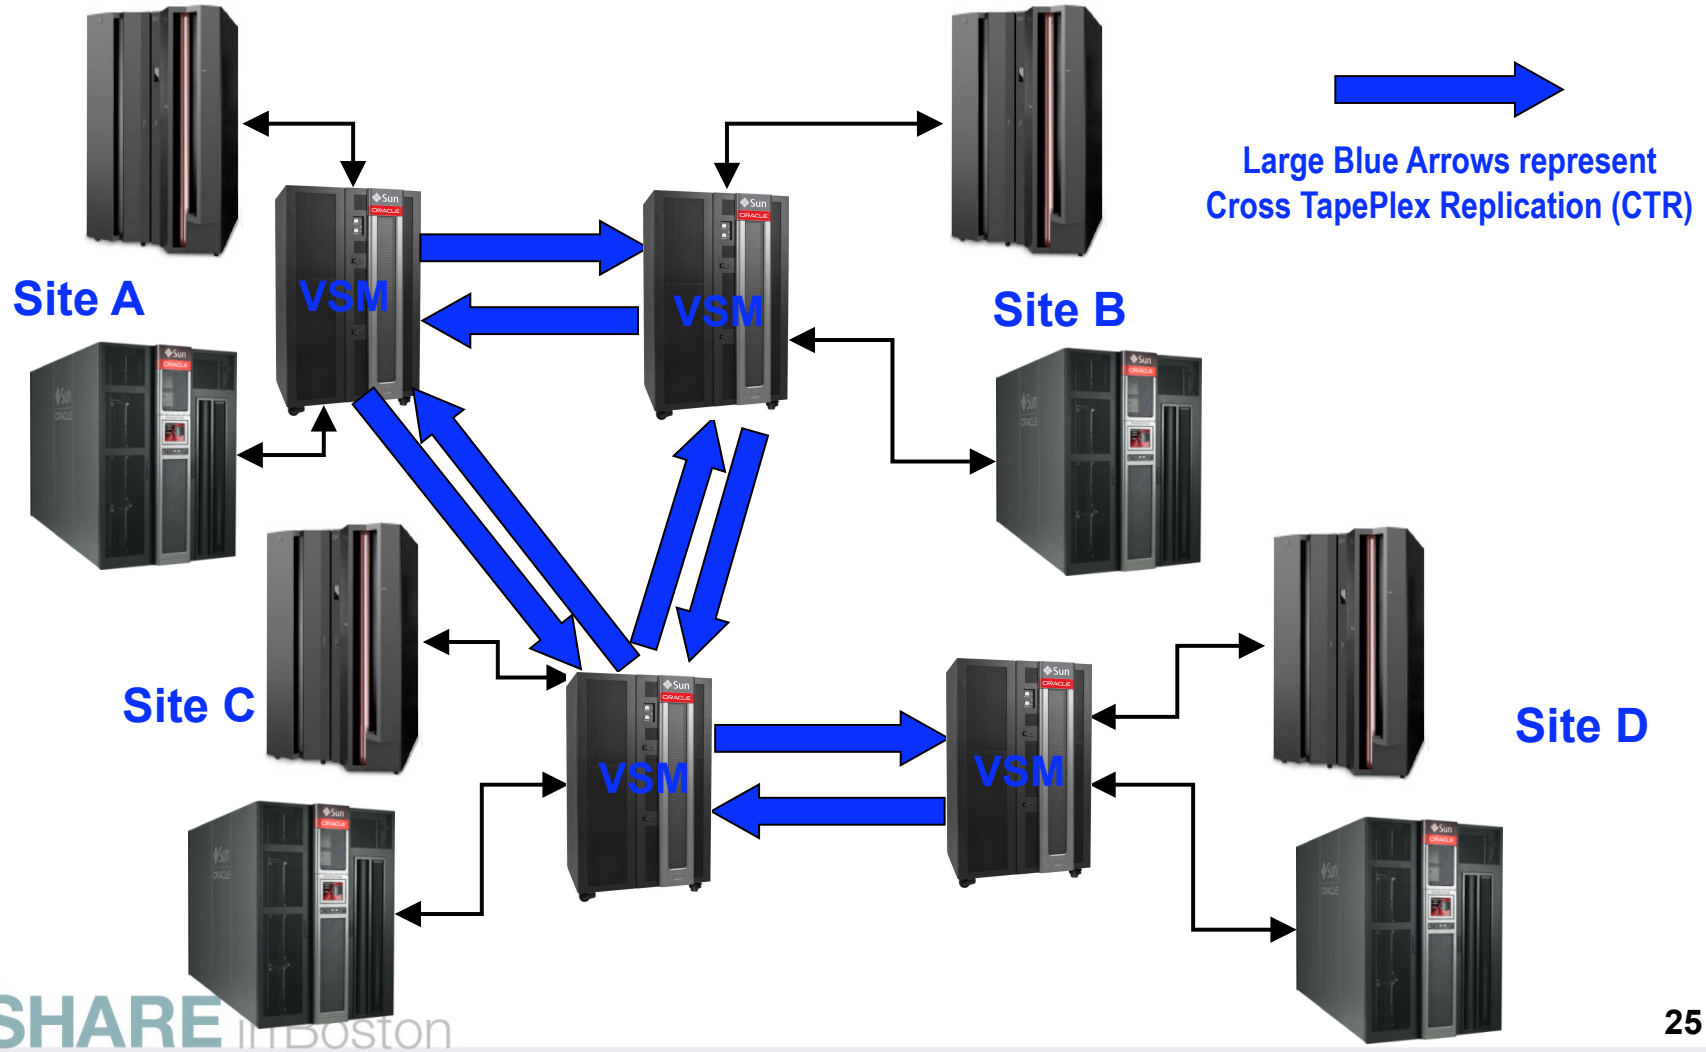

Cross TapePlex Replication (CTR) "Cold" DR Site

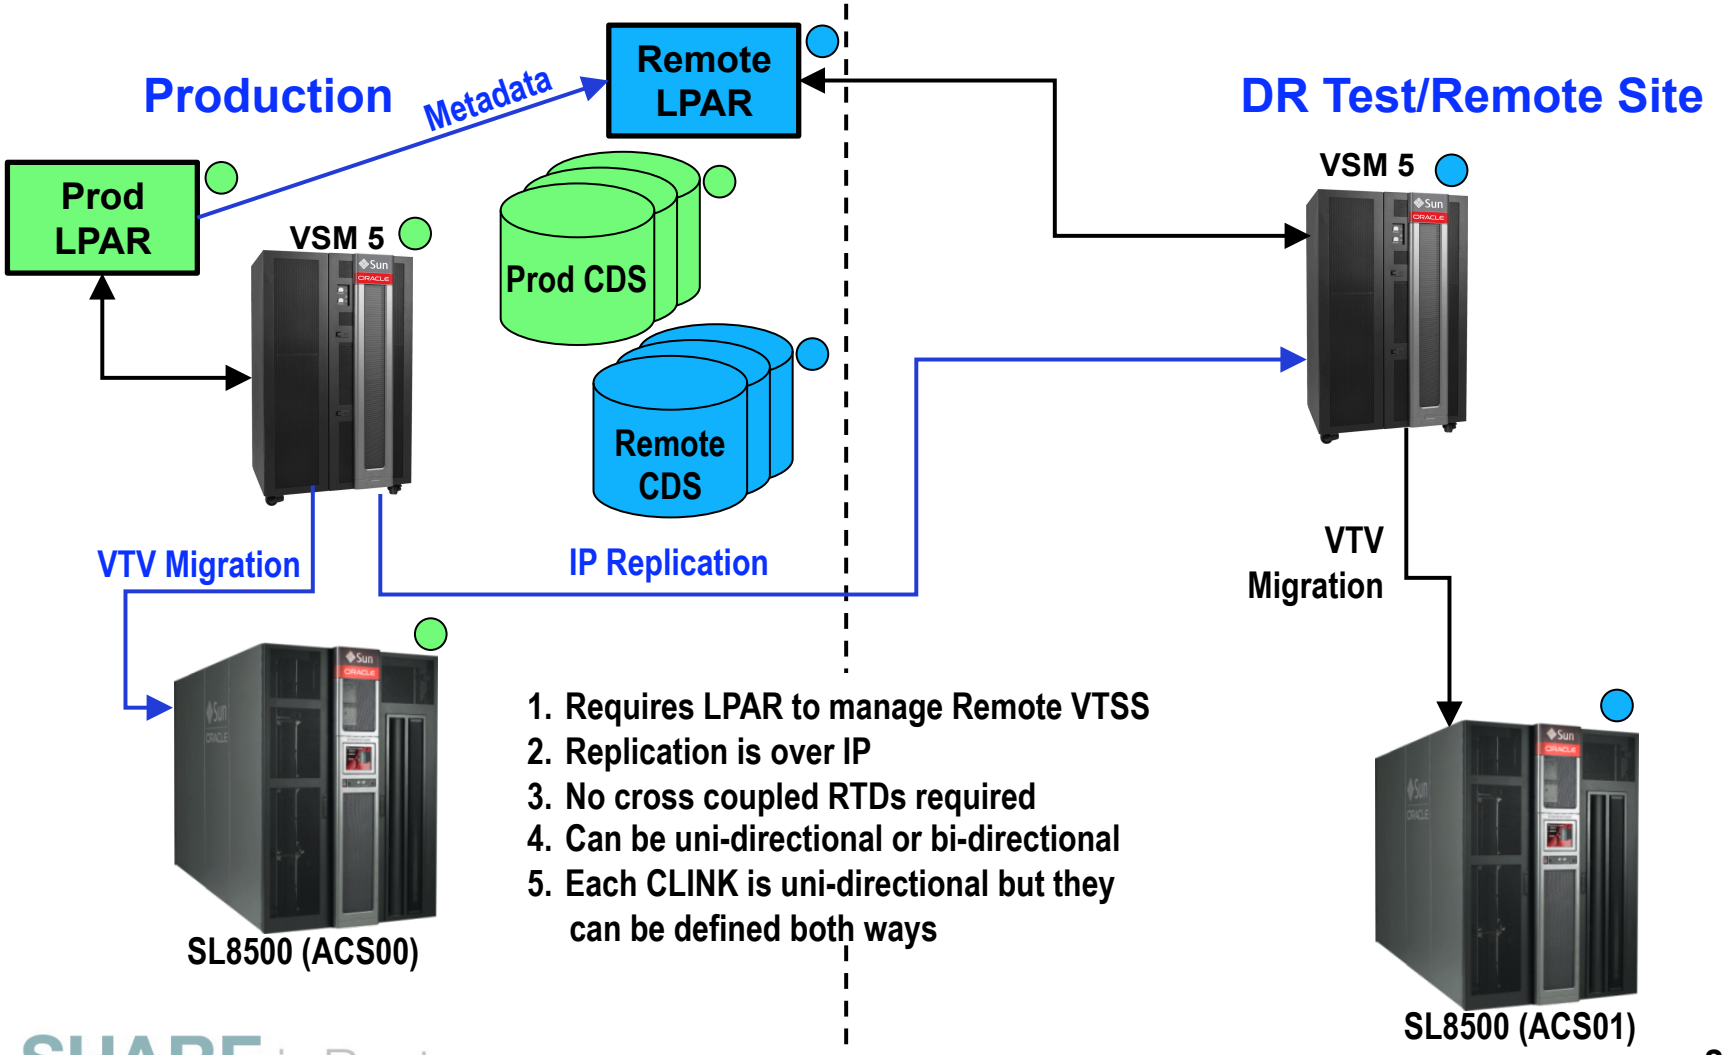

1. Requires LPAR to manage Remote VTSS
2. Replication is over IP
3. No cross coupled RTDs required
4. Can be uni-directional or bi-directional
5. Each CLINK is uni-directional but they can be defined both ways

CTR With a “Cold” DR Site

- If the DR site does not have an active LPAR, how can I run CTR?
  - Run a second TapePlex at the production site
  - If using client/server, the second TapePlex can be run on an LPAR currently running only SMC
  - New feature allows running multiple TapePlexes (multiple copies of HSC/VTCS) on one LPAR
  - DR CDS in remote site, or mirrored to remote site

Cross TapePlex Replication (CTR)
"Cold" DR Site - DR Test

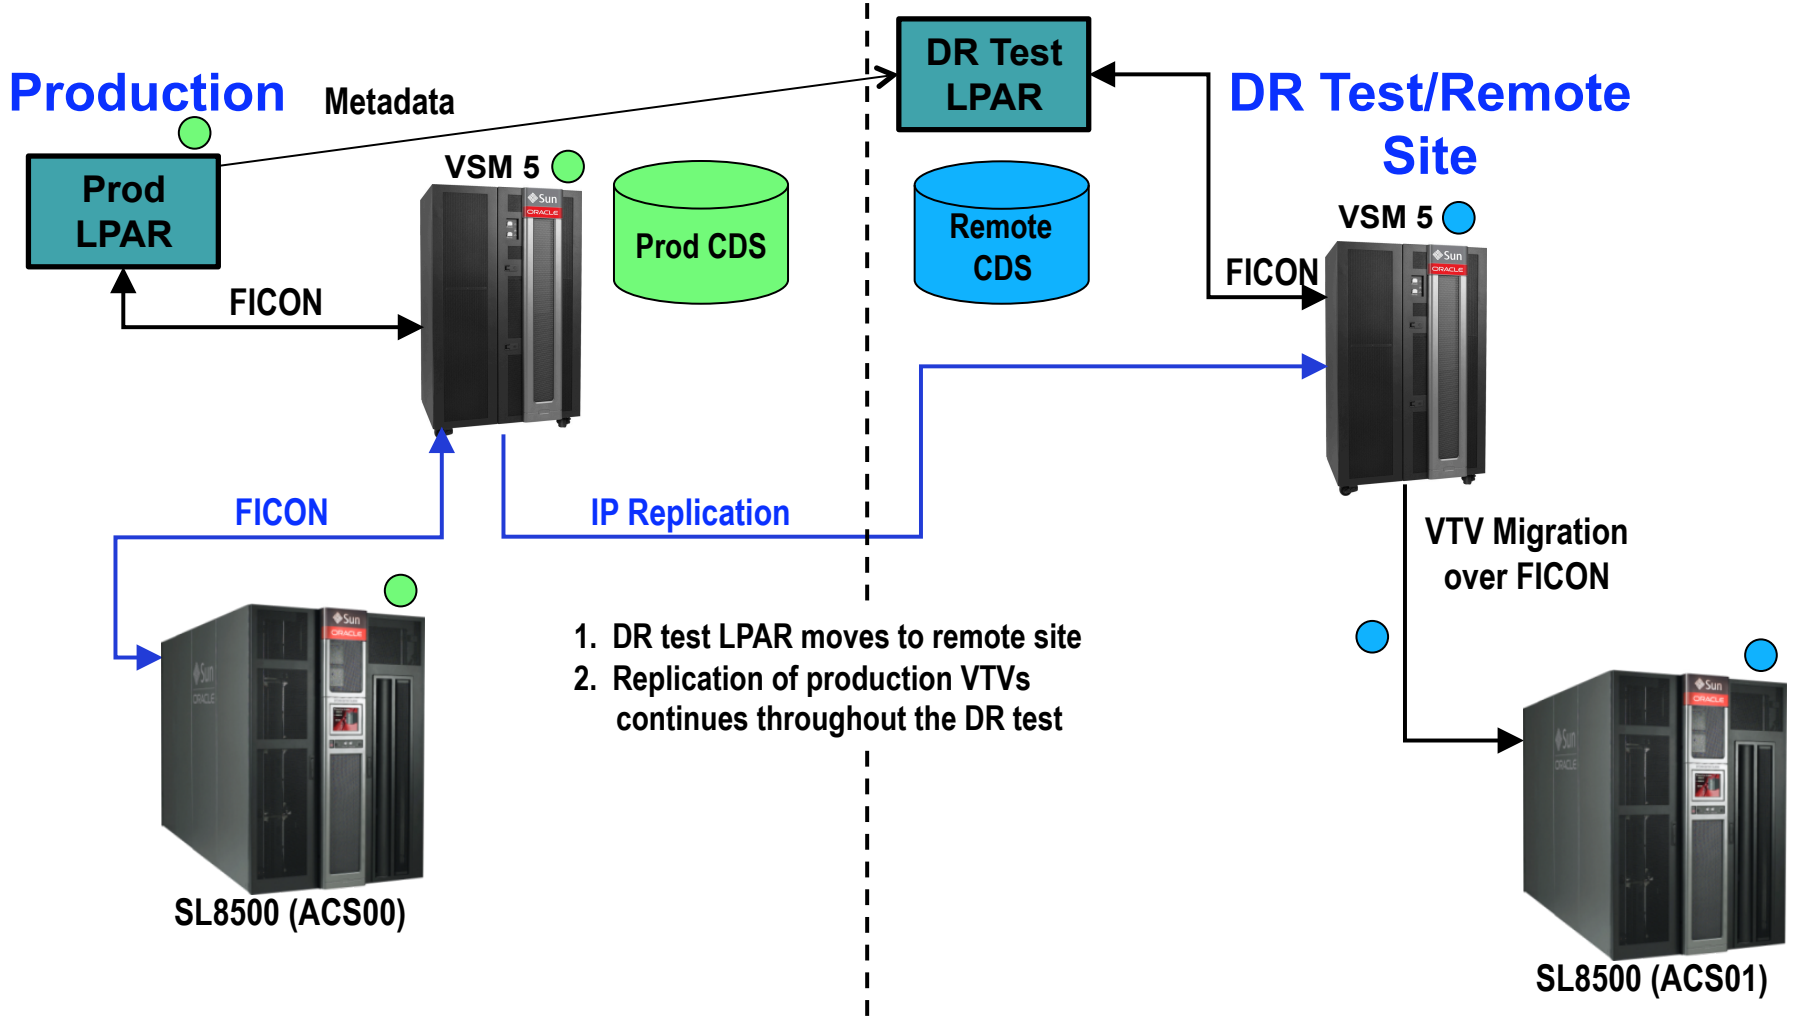

CTR: “Cold” Site - Starting DR Test

- Special considerations apply when running a DR test when the DR site normally has no LPARs
  - Shut down the HSC/VTCS for the DR TapePlex that normally runs at the production site
  - Ensure that the SMC systems at the production site include a SERVER statement that points to the LPAR that will be running HSC/VTCS during the DR test at the DR site
  - These two steps ensure that replication of VTVs to the DR TapePlex can continue during the test

CTR: “Cold” Site - DR Test Cleanup

- At the conclusion of the test, clean up DR test data
  - Scratch VTVs created during the test
  - Drain MVCs used during the test
- Shut down the HSC/VTCS and the test LPAR.
- Re-establish the normal production configuration using the now more current CDS from the DR site
  - Recommended method: FTP the CDS from the DR site to the production site
  - Alternatives: reverse mirror, copy to tape and physically transport
  - Restart the HSC/VTCS at the production site on the updated CDS

CTR With a Real Disaster

- What about a real disaster – how do I manage the CTR tapes?
  - VTVs continue in read-only status to protect data after a disaster (just like the DR test)
  - Temporarily, run only the DR (or second host) TapePlex (just like the DR test)
  - Once you are confident that your recovery has been successful, you can change your POOL definitions to allow the VTVs received via CTR to be scratched
- For a longer term disaster, re-establish the original “owning” TapePlex, and set up to replicate the VTVs into its hardware

VSM DR Capabilities Overview
Concurrent Multi-tiered Solution Set

- Tools
  - Recovery Utility
  - Concurrent DR Test Utility
  - Physical Vaulting Utilities
  - Virtual Volume Multiplexing
  - Multiple ACS Support
  - Remote Vaulting
  - Export/Import
  - Enhanced Clustered VTSS
  - Electronic Export/Import
- Enablers
  - Testing Options
  - Concurrent DR Test Utility
  - Electronic Export/Import
  - Channel Extension
    - Brocade
    - Cisco
  - Professional Service Implementation Offerings

VSM – New DR Product Features

- Several new product features (ELS 7.0) improve Disaster Recovery functionality
- LCM 7.0 Vault function improves management of DR vaulting, especially for MVCs
- ELS 7.0 VOLPARM feature allows allocation of restricted ranges for DR scratch and MVC output during a Concurrent Disaster Recovery Test (CDRT)
- VTCS Cross TapePlex Replication (CTR) feature allows a remote TapePlex to manage copies of virtual volumes replicated from another TapePlex

Summary

- VSM offers the ability to DR in many ways
- Have had over 10 years of experience
- DR capabilities continue to evolve
- Cross Tape Replication provides “real time” data replication and testing that closely matches a real disaster situation