

Oracle StorageTek's VTCS DR Synchronization Feature

Irene Adler
Oracle Corporation


Thursday, August 9, 2012: 1:30pm-2:30pm
Session Number 11984

Agenda

- Tiered Storage Solutions with VSM's VTSS/VLE/Tape
- Some Recovery Options with Virtual Storage Manager
- Disaster Recovery Considerations
- What DR problem are we trying to fix
- What is VTCS DR Synchronization?
 - How does it work?
 - How do you implement this feature?
- How does Synchronization fit in with current DR strategies:
 - Remote RTDs
 - Clustered VTSS
 - Cross-TapePlex Replication
 - Tapeless


Modern, Large Capacity Data Storage – Integrated, Highly-Scalable Solution

- 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 Disk 0.8TB – 23PB**
- **Tier 2 – Economical Disk 0.22TB–338PB**
- **Tier 3 – Tape Up to 1000PB**

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

Hierarchy of Performance and Capacity

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

Some VSM Disaster Recovery Options



- VTV Copies
- Export/Import
- Auto Archiving
- Physical Vaulting (with LCM)
- Clustered VTSS (synchronous or asynch replication)
- Extended Clustering (also synch or asynch replicat)
- Cross-TapePlex Replication (electronic export)
- Concurrent Disaster Recovery Test

Disaster Recovery Considerations



- Recovery Time Objective (RTO)
 - The time required to restore business to an acceptable level of operational capability.
- Recovery Point Objective (RPO)
 - A point of time from which to begin recovery.
- There is no mechanism to ensure a CDS backup is valid when a DR is declared.
- How is data protected between CDS backups?
- How can you ensure a consistent set of data can be used as a DR baseline?
- How can you prevent VTV overwrites when an MVC is Drained or Reclaimed?

What's the problem we want to fix?

- VTCS does not establish system checkpoints to ensure that all Disaster Recovery critical data has been migrated and is secured at the customer's recovery location.
- 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.

What does this mean?

- Wait for DR data to reach its recovery location.
- Set checkpoint.
- Take CDS backup.
- Don't allow overwrites on MVCs until the next DR baseline is established.
- And the next DR baseline is established by repeating the above three steps.
- DR Synchronization is a new feature of VTCS which will provide the capability to do this.

VTCS DR Synchronization Overview



- Enhancements to VTCS to support new feature:
 - New Utility to monitor DR critical VTVs until they have reached the desired DR location.
 - New Utility to set a date/timestamp (synch point) in the CDS.
 - Modify VTCS CONFIG RECLAIM processing to specify the number of hours to protect an MVC/VMVC from re-use after Drain/Reclaim processing.
 - Modify MVC (real MVC tape) and VLE VMVC (virtual MVC) re-use based on checkpoint and/or set protection time.
 - Modify VTCS to defer specific VLE cleanup tasks, done at the end of Drain/Reclaim, until time of first migration to that VMVC.

VTCS Migration Monitor Utility



- The purpose of the DRMONitr Utility is to monitor a set of migrates and/or electronic exports.
- This utility allows a job stream to stall awaiting a specific set of data to reach the desired DR location. Once completed, the job stream can continue.
- The utility assumes that the customer has a set of VTVs that require migration to a particular Storage Class.
- At that time, a CDS checkpoint should be set and then the CDS backup taken.
- In addition, backups of Tape Management Catalog, System/User Catalogs and other types of metadata should be taken at this point.

DRMONitr Utility

- Parameter Syntax for VTVs to be monitored:

DRMONitr ——— MGMTclas(*mgmtclas* or *mgmtclas_list*)
 ├── VTV(*vtv_list* or *vtv_range*)
 └── DSN(*dataset_list*)

These three keyword parameters are mutually exclusive and one must be specified.

┌_ STORclas(*storclas* or *storclas_list*) _┐

┌_ REPLICAT _┐

One of the above two parameters are required to indicate target location.

┌_ MAXAGE(*nnn*) _┐

(in hours)

┌_ TIMEOUT(*nnn*) _┐

(in minutes)

These two parameters are optional.

More about DRMONitr Parameters

- Three, mutually exclusive, parameters identify the data to be monitored until it has reached its target destination. One must be specified:
 - MGMTclas – list of management classes; **or**,
 - VTV – list or range of VTV volsers; **or**,
 - DSN – list of datasets.
- STORclas **or** REPLICAT – two parameters which designate target location for DR VTVs. One must be specified.
 - STORclas – specifies the storage class location(s) where designated, critical VTV data is to be migrated – the DR destination location; **or**,
 - REPLICAT - specifies that VTV replication is to be monitored in a VSM Tapeless configuration.

More about DRMONitr Optional Parameters



- MAXAGE(*nnn*)
 - Optional parameter which specifies the maximum VTV age, in hours, to be considered when monitoring migrations by management class (when MGMTclas has been specified).
 - If the VTV's age is less than MAXAGE, the VTV is monitored; if greater than MAXAGE, the VTV will not be monitored.
 - If MAXAGE is not specified, any VTV belonging to the specified MGMTclas will be monitored.
- TIMEOUT(*nnn*)
 - Optional parameter which specifies the allowable time, in minutes, for the monitor utility to run.
 - When the amount of time the monitor has been active exceeds the TIMEOUT value, the utility aborts with return code 8.
 - If TIMEOUT is not specified, there is no limit on the time the utility can run (wait).


VTCS Checkpoint Utility



- The purpose of the DRCHKPT Utility is to set (or to clear) a date/timestamp checkpoint in the CDS, which represents a synch point, which establishes the system recovery point from which MVC content can be recovered for a period of time.
- The timestamp reflects a point in time where the physical layout of the MVCs are preserved.
- All VTV backup copies, up to this synchronization point, can be referenced **even** if the MVC has been Drained or Reclaimed since creation.
- The Monitor Utility should run every time a CDS backup is scheduled to ensure every VTV makes it to its destination, followed by the Checkpoint Utility to set the recovery point, immediately followed by a backup of the CDS, TMC and other critical DR metadata.

DRCHKPT Utility


- Parameter Syntax for date/timestamp Checkpoint recovery point in the CDS:



- The SET parameter sets the current date / timestamp in the CDS as the Checkpoint.
- The CLEAR parameter clears or removes the Checkpoint date / timestamp from the CDS.

New Protection Parameter for VTCS Reclaim

- Change in Parameter Syntax for VTCS CONFIG RECLAIM:



This parameter is optional.

- The PROTECT parameter is optional.
- Specifies the time, in hours, that VTCS will prevent an MVC from being re-used after being Drained or Reclaimed.
- Thus, the MVC is protected for that duration of time.
- Valid values are 1-99.
- The DECOM utility will generate the PROTECT parameter if it is set in the CDS.

How Does this affect MVC Re-use?

- Without the VTCS DR Synchronization feature, MVCs were eligible for reuse immediately after a Drain or Reclaim was run.
- With this feature enabled, VTCS will prevent MVC reuse after a Drain or Reclaim based on any Checkpoint timestamp set in the CDS and/or any post Drain/Reclaim reuse protection policy set in the configuration via the CONFIG RECLAIM PROTECT parameter.
- This may mean that more MVCs may be required due to MVCs that will not be eligible for new migrations for a protected period of time.

How Does this affect VMVCs?

- For VLE-based VMVCs, VTV delete processing currently occurs at the end of Drain/Reclaim processing.
- With VTCS DR Synchronization feature enabled, VMVCs will “look and feel” more like real tape.
- VTV delete requests will be deferred until the first Migrate.
How this will work is:
 - the VMVC will be inventoried;
 - then VTCS will process the VTV list and associated timestamp information from the CDS;
 - valid VTVs will be retained and all others will be deleted.

Requirements for VTCS DR Synchronization



- Dependencies: ELS 7.0 and above.
- Enabling PTFs:
 - for ELS 7.0 L1H16VZ
 - for ELS 7.1 L1H170E
 - For either release, the PTF must be applied to all Hosts.
- New DR Monitoring and DR Checkpoint Utilities must be established to run immediately before every CDS backup is taken.

Reporting Modifications



- With DR Synchronization, the MVC report will identify all DR-protected MVCs with the letter 'P' under the Usage column. This indicates the MVC is protected and cannot be selected for migration.
- A protected MVC remains available for recalls.
- MVCRPT example:

TIME 13:24:52 VTCS MVC Summary Report DATE 2012-05-09

MVC Number	Size (GB)	Mount Count	Status	Last Mounted	Location	Owner / Comment / Consolidate				
Volser	VTVs	%Used	%Avail	%Frag	IBLDRUTM	Date	Time	VTSS	/ ACS ID	Consolidate Time
VMC000	43	63.10	36.90	0.00	244.2	258	I - - - - - P - -	2012May09 13:23:28	DVTSS16	VLELIB0 S0
VMC001	1	1.43	98.57	0.00	250.0	257	I - - - - - U - -	2012May09 13:24:42	DVTSS16	VLELIB0 S0

2 Initialized MVCs processed

Display Modifications

- Display MVCPOOL
 - Will now display the checkpoint timestamp in the CDS, and/or,
 - any configuration PROTECT policy specified.
 - These will be displayed at the end of the Display content as:
 - DR checkpoint set: 2012Feb27 11:50:54** → showing ckpt date and time
 - MVC protection period set: 30 hrs** → showing hours set

- Display CDS

- Will now display the checkpoint timestamp in the CDS:

PRIMARY = VTCSDEV.VLE1442.DVTTEST.CDS

VOLSER = VTCS00 ACTIVE

CDS LEVEL = 060100 VER = I6746446

VSM LEVEL = G

VSM DR CHECKPOINT = 20120525 09:52:14

.

Operational Scenario using Checkpoint



- Once the customer has determined which data is DR critical:
 - The DRMONitr Utility is run against the data. This will stall a job stream until the VTV copies have reached their MVC target destination(s).
 - Then the DRCHKPT Utility is run, which establishes a checkpoint time. This serves to protect MVC content until some point in the future (i.e. the next DRCHKPT run).
 - After DRMONitr and DRCHKPT have completed, then the CDS backup should be taken, along with other DR critical datasets, such as the TMC and MVS Catalogs.
- In the event a disaster is declared (real or test), the CDS is restored from the current backup. Because the backup's MVC contents are protected, they can be used without the need for audits because they will be ***synchronized*** with the CDS.

How would this work without PROTECT set?



- Consider the following example:
 - Tuesday @ 3:59am - DRMONitr Utility completes
 - all critical VTV data has been successfully migrated to target storage classes
 - Tuesday @ 4:00am – DRCHKPT Utility runs and CDS backup taken
 - at this point, both the CDS and the CDS backup checkpoint reflects Tuesday @ 4:00am.
 - Tuesday @ 18:00pm – MVC1 is reclaimed
 - MVC1 drain/reclaim time reflects Tuesday @ 18:00pm
 - Tuesday @ 22:00pm – Migration runs and is not allowed to select MVC1 because the MVC1 drain/reclaim timestamp is greater than the time of the CDS checkpoint.
 - Wednesday @ 3:59am – DRMONitr Utility completes
 - Wednesday @ 4:00am – DRCHKPT Utility runs and CDS backup taken
 - MVC1 is no longer protected at this point.
 - Wednesday @ 22:00am – Migration selects MVC1 and is allowed because MVC1 is no longer protected because MVC1's drain/reclaim timestamp is less than the new CDS checkpoint time.

With CONFIG PROTECT? How does it differ?

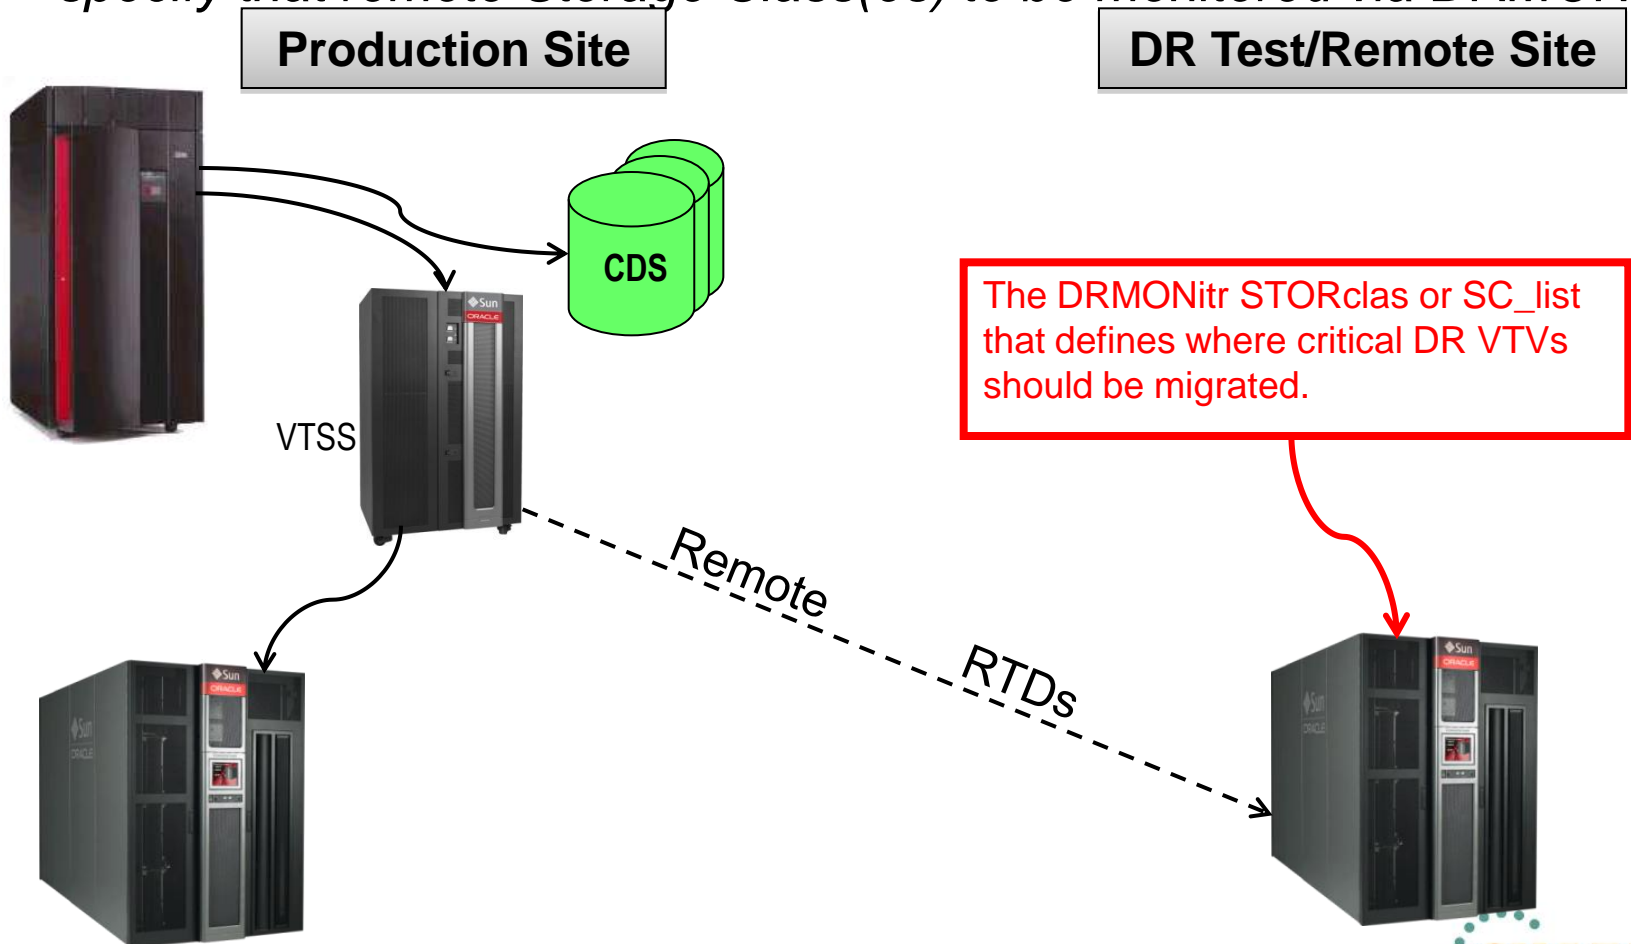


- Using the same operational scenario, but now using CONFIG PROTECT parameter:
 - Tuesday @ 3:59am – DRMONitr Utility completes
 - Tuesday @ 4:00am – DRCHKPNT Utility is run and CDS backup taken
 - Tuesday @ 4:00am – MVC1 is reclaimed
 - MVC drain/reclaim time reflects Tuesday @ 4:00am
 - Note – using a 30-hour PROTECT policy, MVC1 cannot be re-used until Wednesday @ 10:00am.
 - Tuesday @ 22:00pm – Migration runs and is not allowed to select MVC1 because time of MVC selection (Tuesday @ 22:00pm) is less than MVC1 drain/reclaim timestamp **plus** 30 hours (Wednesday @ 10:00am).
 - Wednesday @ 22:00pm – Migration selects MVC1 and is allowed because the time of MVC selection is past MVC1 drain/reclaim timestamp plus 30 hours.
- Care must be taken to establish the optimum protection period for adequate MVC content protection between CDS backups.

How does this work with your DR solution?

- Remote RTD Configuration:

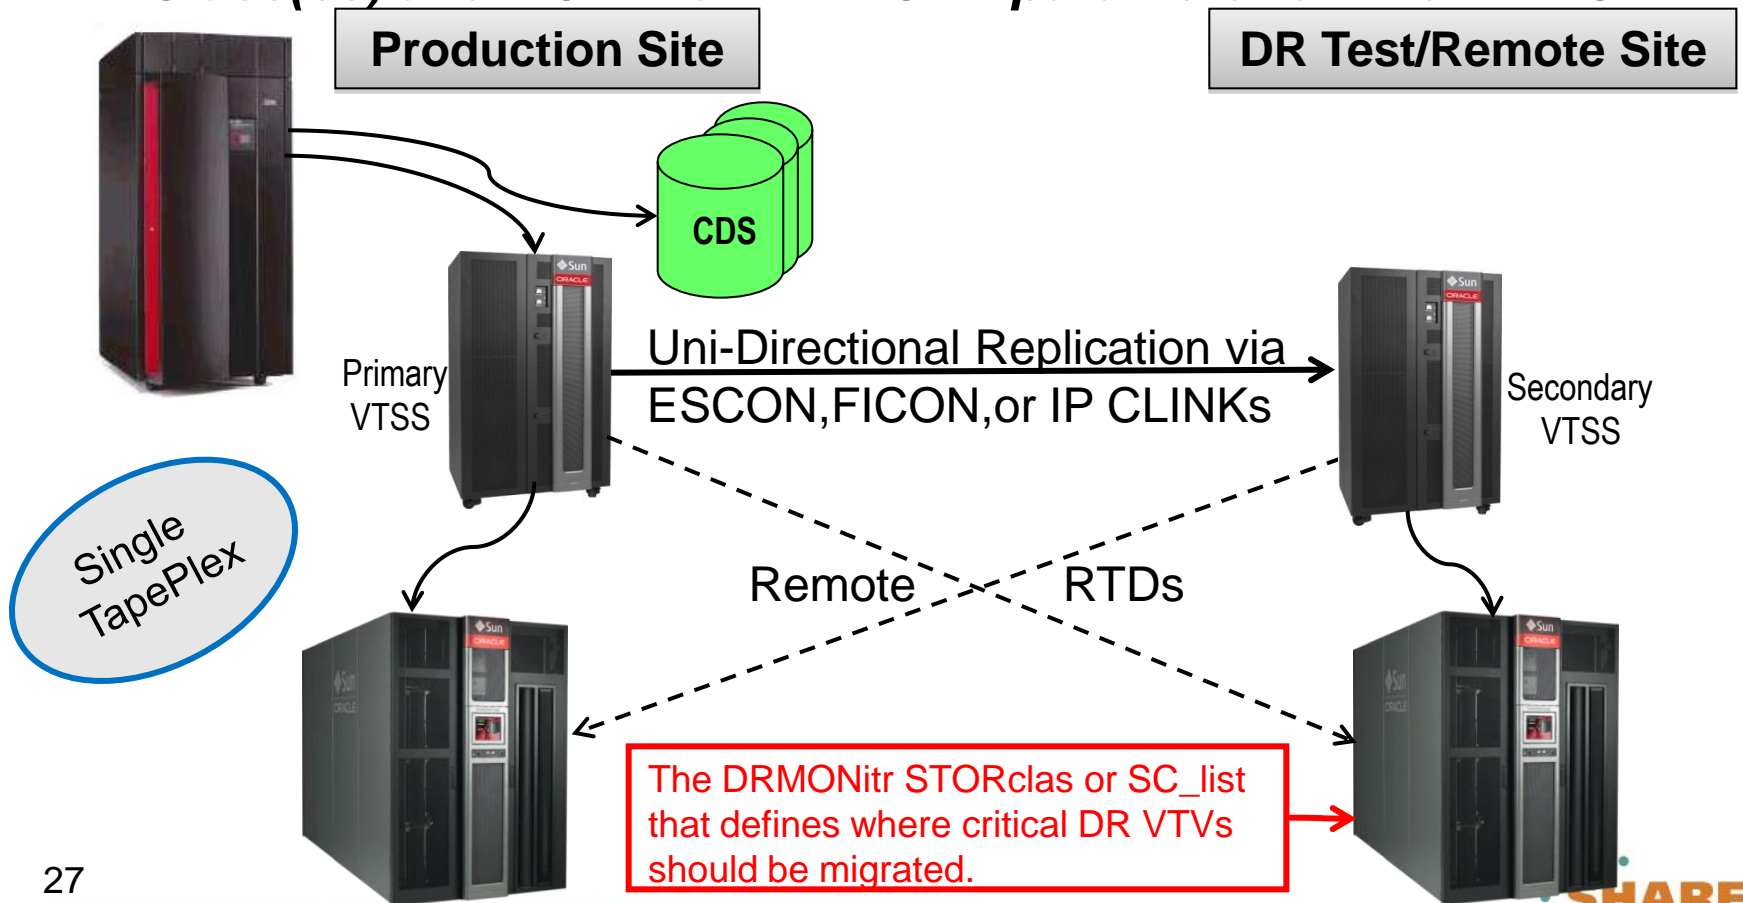
- *If you are creating your DR VTV copy via remote RTDs, then you would specify that remote Storage Class(es) to be monitored via DRMONitr.*



How does this work with your DR solution?

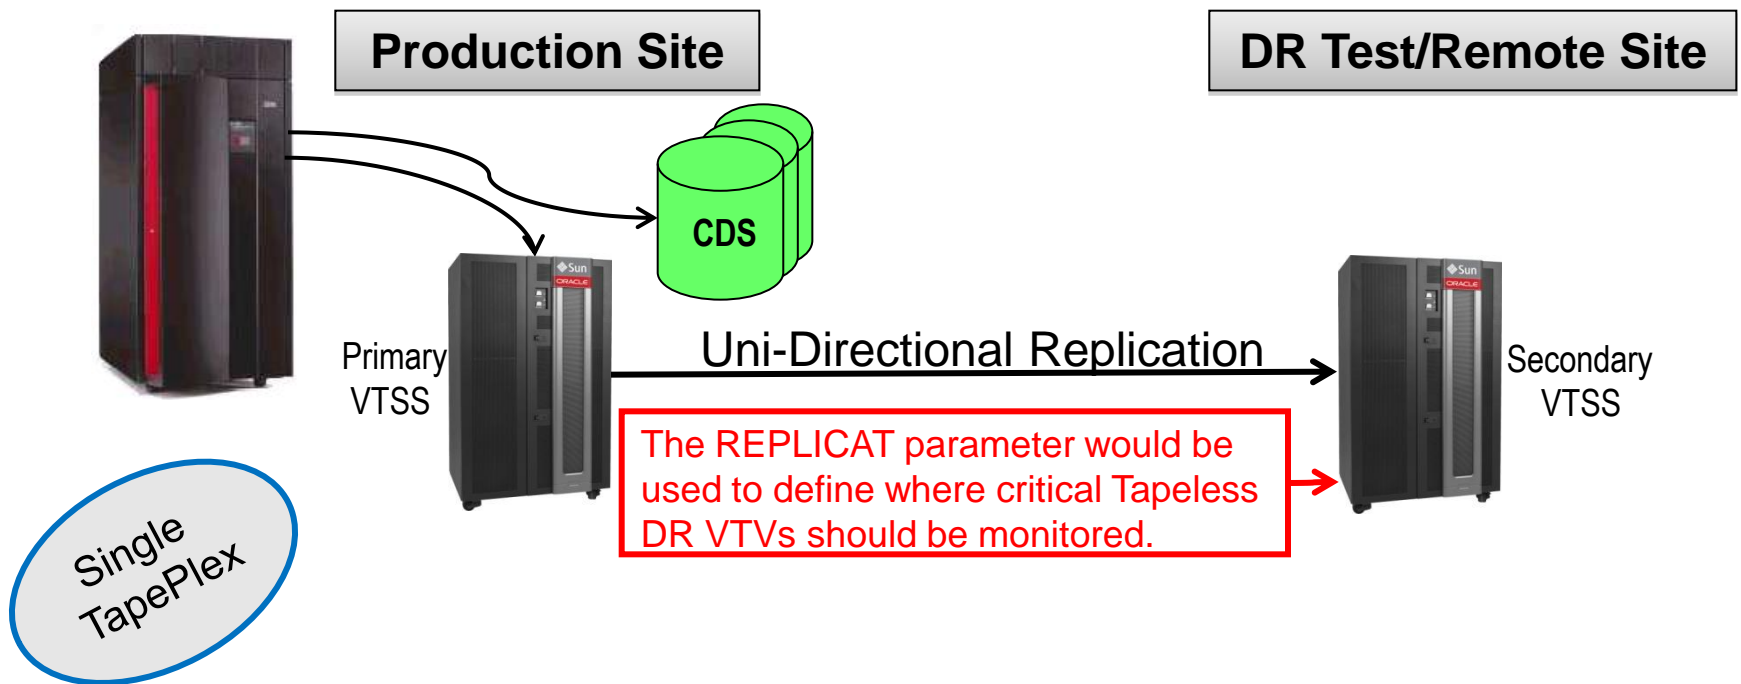
- Uni-Directional Clustered VTSS Configuration:

- In a Clustered configuration, if the secondary VTSS migrates the DR VTV copy to an MVC, then you should specify the **remote DR Storage Class(es)** and **NOT the REPLICAT parameter on the DRMONitr.**



How does this work with your DR solution?

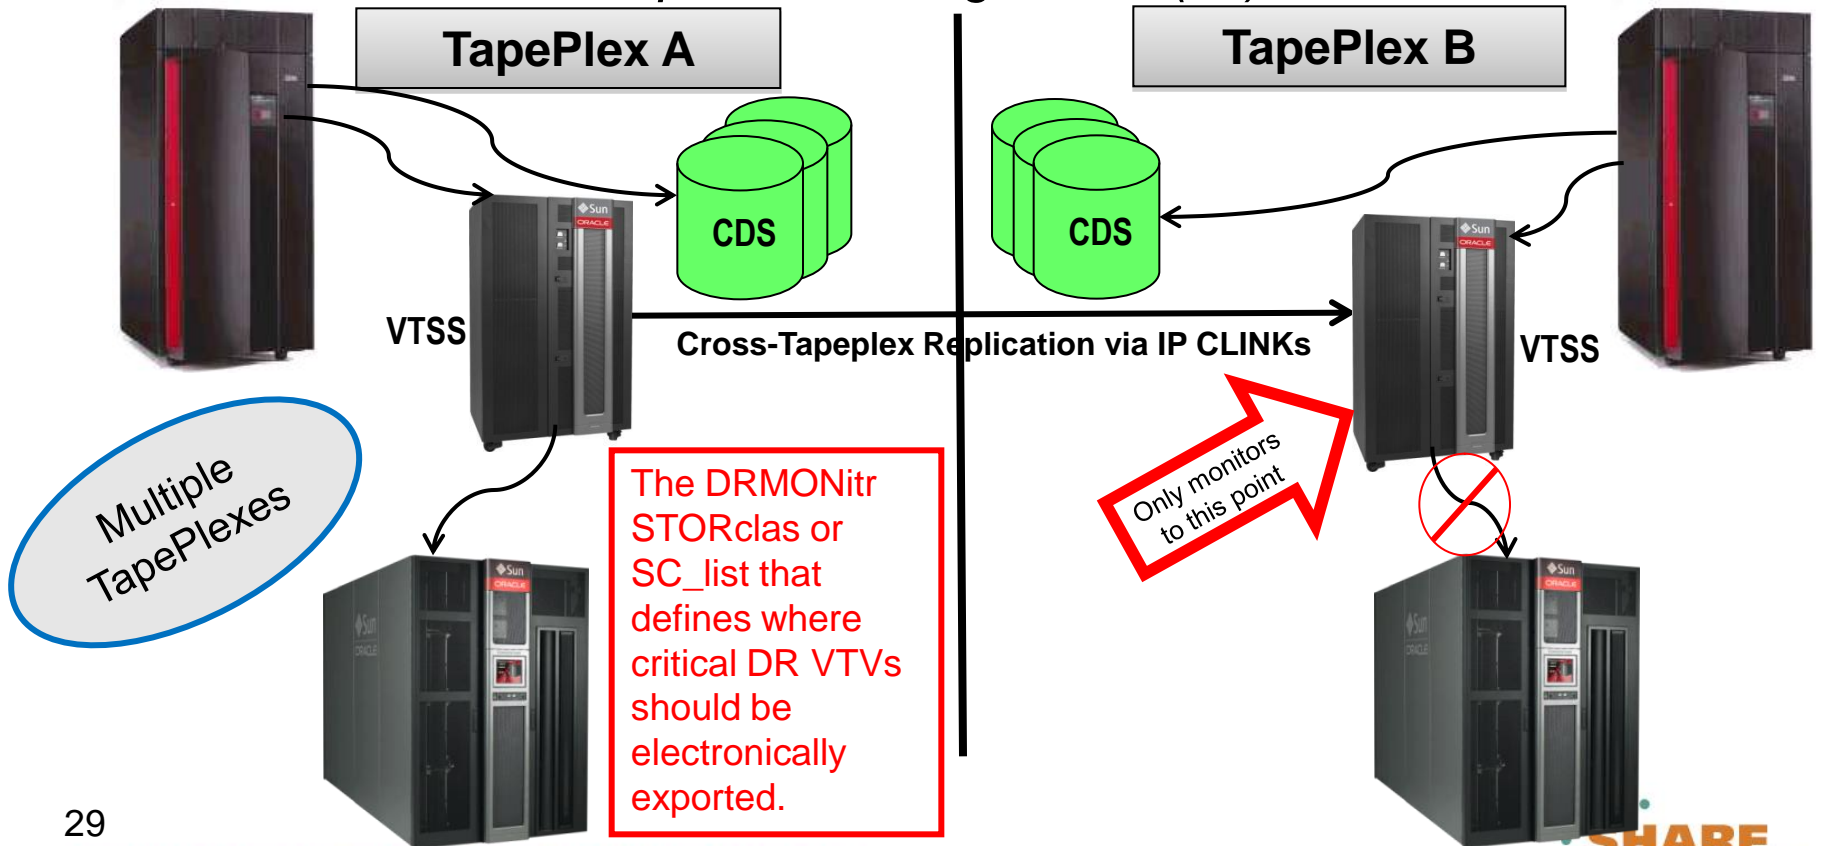
- Tapeless VTSS Configuration:
 - In a Tapeless Clustered VTSS configuration, you should specify the REPLICAT parameter on the DRMONitr Utility.



How does this work with your DR solution?

- Cross-TapePlex Configuration:

- *In a CTR configuration, where CTR copies are monitored, the Monitor only waits for the **electronic export** to occur to the remote TapePlex and not for the VTV to be migrated to an MVC in the remote TapePlex. Specify the remote TapePlex Storage Class(es) and NOT REPLICAT.*



What to be aware of . . .

- Determine which Storage Class(es) are mission critical for DR. Those will be the ones you will want monitored.
- The Monitor Utility will wait until all VTVs associated with those Storage Class(es) have reached their desired location.
- DRMONitr should be scheduled at the place in the daily cycle where all VTVs are believed to be created (if at all possible) and before the CDS backup is taken.
- The DRMONitr, DRCHKPT and CDS Backup utilities should be set up in a single jobstream, executed in that order, and condition coded to execute contingent on the success of the previous steps.
- Establish the optimum protection period for adequate MVC content protection between CDS backups.
- Additional MVCs may be required in the environment to allow for the ones that are being protected and temporarily not allowed for migration.
- Use the REPLICAT parameter only with Tapeless configurations.

Thank You

Questions ?