

Reclaiming Space on High-Capacity Tapes with VTCS Dynamic Reclaim

Irene Adler
Oracle Corporation

Session 9929
Wednesday, August 10, 2011: 6:00 PM-7:00 PM

Agenda

- What problems are we trying to address?
 - With high-capacity tape drives, tapes now hold an enormous amount of customer data
 - How to deal with expired space on tapes fast and efficiently
- What is VTCS Dynamic Reclaim?
- How is it different from Standard Reclaim
- How does Dynamic Reclaim work?
- Software/Hardware Requirements
- VTCS implementation parameters

Oracle's StorageTek T10000B and T10000C High Capacity Tape Drives with VSM



- With the introduction of higher-capacity tape drives, MVCs (Multi-Volume Cartridges) now hold thousands of gigabytes of data, which are individual datasets, or VTVs (Virtual Tape Volumes) stacked on the tape, often with different expiration dates
- Capacity of Oracle's StorageTek T10000x Tape Drives
 - T10000B has a native capacity of 1 terabyte
 - T10000C has a native capacity of 5 terabytes
- Over time, tapes become fragmented with non-current data caused by expiring VTVs interspersed throughout the tape
- How can unused space be reclaimed fast and efficiently?

Reclaiming Space on Tapes

- Traditional Reclaim means that the MVC has to be mounted on an RTD (Real Tape Drive) and the valid VTVs remaining on the tape past the first non-current VTV recalled back into the VTSS (Virtual Tape Subsystem) buffer and re-migrated out to a new MVC using a second RTD
- This reclamation method costs the customer:
 - RTD and VTSS buffer resources
 - amount of time required to run reclaims
- Oracle's StorageTek T10000B and T10000C tape drives introduced a new tape technology called tape partitioning
- This provided an opportunity to *dynamically* reclaim unused space on tapes

Standard MVC Reclaim Overview



- VTVs are migrated in consecutive block-ids on an MVC
- When fragmentation occurs on these tapes, it requires the user to run a Reclaim job, which recalls VTVs back into the VTCS buffer, and re-migrates them back out to a different MVC
- This recall and re-migration continues until there are no more non-current VTVs on the tape
- Then the logical end of tape (EOT) is set immediately behind the last current VTV on the tape
- Then the MVC becomes eligible to accept new migrations
- This process ties up RTD tape drives and VTSS buffer resources, in addition to precious time in a busy tape environment

What is VTCS Dynamic Reclaim?

- VTCS Dynamic Reclaim is VSM's implementation of reclamation for partitioned MVCs
- Dynamic Reclaim is a VTCS feature that was introduced with Enterprise Library Software (ELS) 7.1 to work with Oracle StorageTek's T10000B and T10000C tape drives to allow dynamic reclamation of non-current space on MVCs
- Dynamic Reclaim requires no RTD mounts and no VTSS buffer resources
- Dynamic Reclaim works by creating automatically linked partitions on the tape whereby VTVs can span across partitions and be non-contiguous on the tape
- VTCS Dynamic Reclaim is both a hardware enabling feature and a VTCS software feature

Dynamic Reclaim Overview



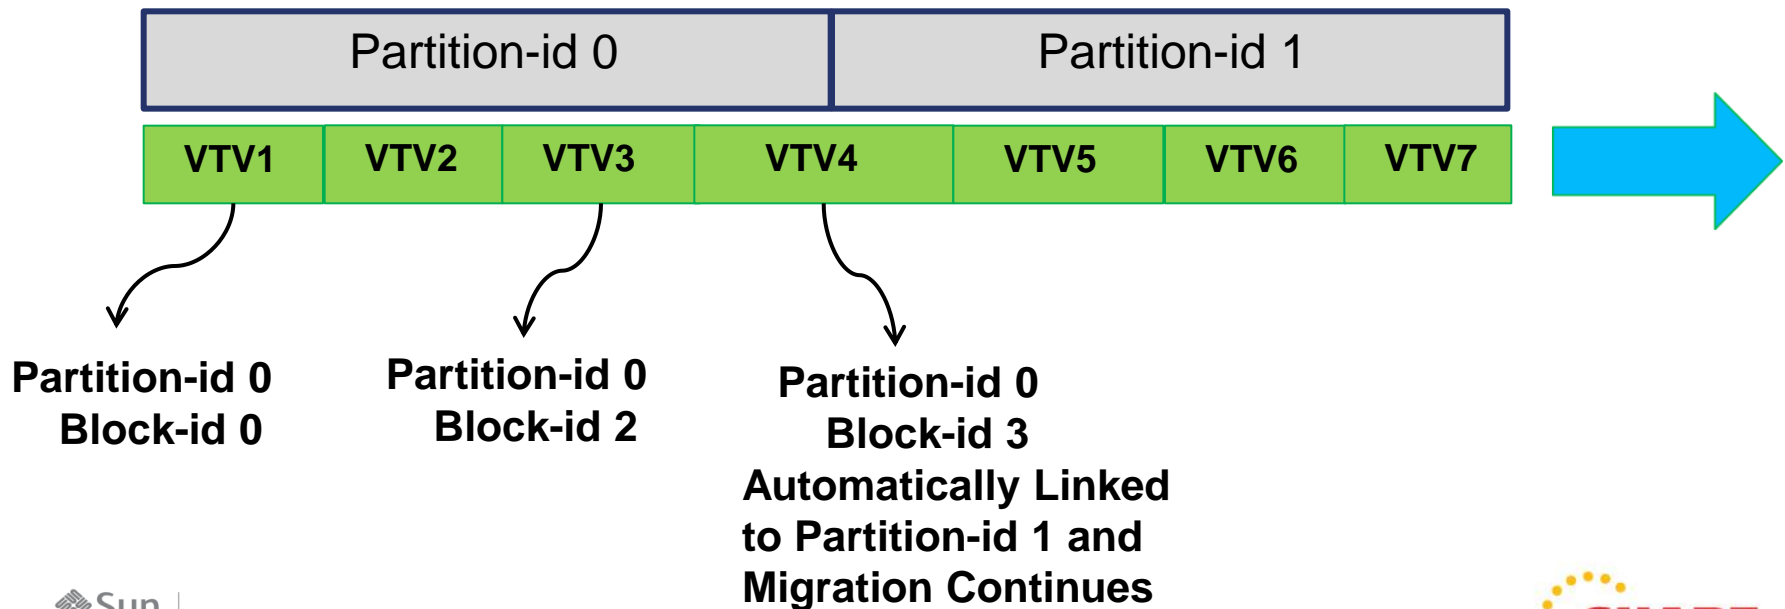
- Dynamic Reclaim uses new, Oracle StorageTek tape technology called tape partitioning
- The first time a partition-format MVC is mounted for migration, VTVs are migrated in consecutive block-ids into partitions, unlike standard-format MVCs where VTVs are migrated to consecutive block-ids
- VTVs can span over to a second partition
- Tape Partition sizes:
 - T10000B tape has 192 partitions (0-191) at 4.84GB per partition
 - T10000C tape has 480 partitions (0-479) at 9.37GB per partition
- Without Dynamic Reclaim enabled, the block-id of the next block on the tape eligible to be written to is kept in the CDS; with Dynamic Reclaim, partition-id and block-id are preserved

MVC Partition Map

- Dynamic Reclaim requires an H-level CDS to support a new format record type used for partition mapping
- The MVC Partition Map (MPM) record carries partition status information
- The MPM record indicates if partitions can or cannot be used for migration

How Dynamic Reclaim Works

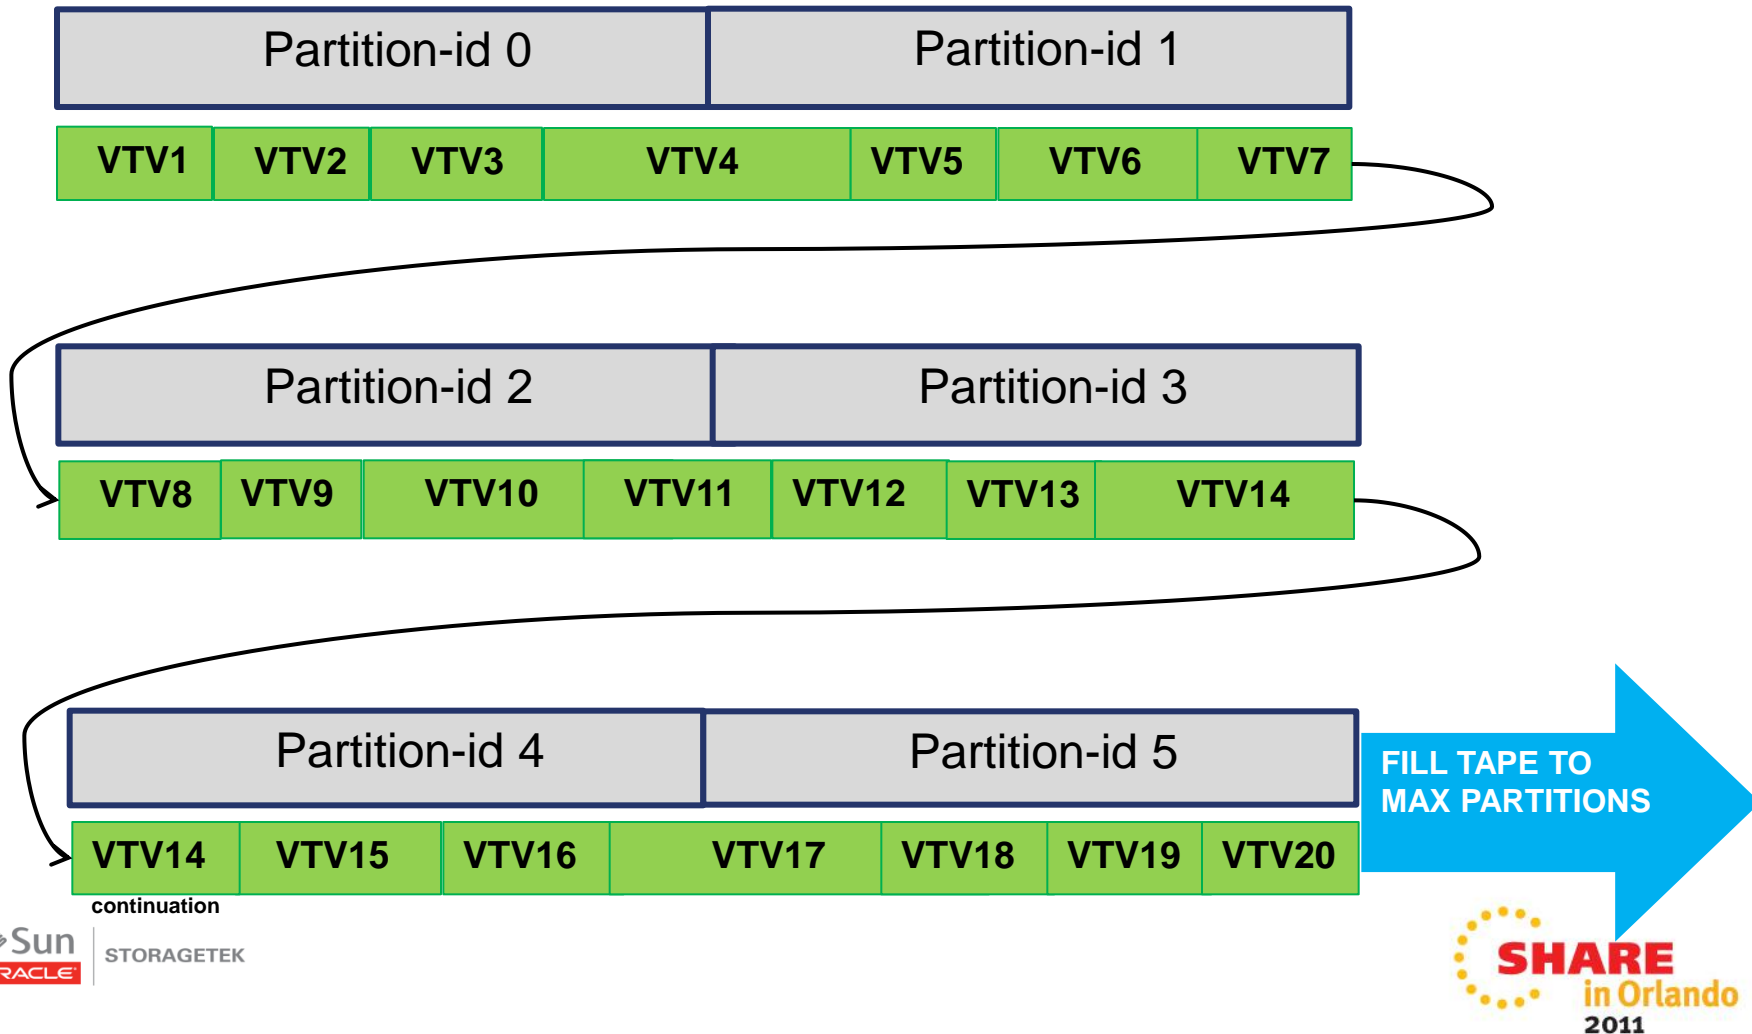
- The first time that a Dynamic Reclaim, or partition-format MVC is used, VTVs are migrated consecutively to MVCs, without regard for partition boundaries, and the partition-ids and block-ids are updated in the CDS



How Dynamic Reclaim Works (cont'd)


- VTV Migration to MVC starts

Note: VTVs 4, 11, 14 and 17 span Partitions




How Dynamic Reclaim Works (cont'd)

- Virtual Volumes have been migrated to MVC
 - CDS updated to reflect VTV Partition-id and Block-id




VTV4 spans Partitions 0 and 1



VTV11 spans Partitions 2 and 3

VTV14 spans Partitions 3 and 4



continuation



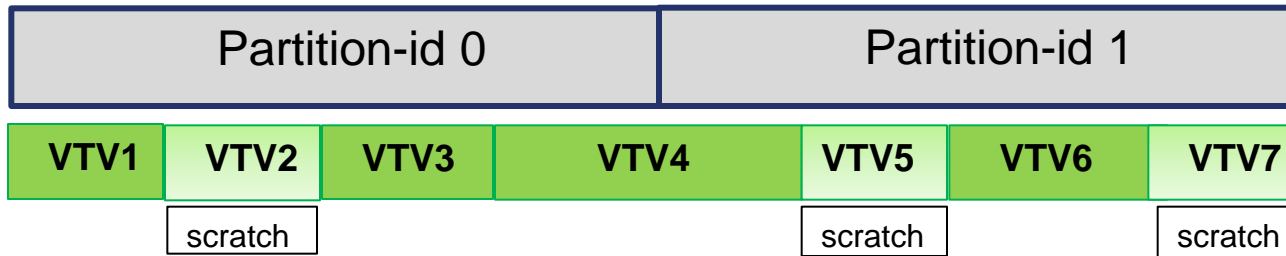
VTV17 spans Partitions 4 and 5

More
Partitions
Remain

How Dynamic Reclaim Works (cont'd)

- Tape starts to become fragmented as VTVs become Scratch

Note: VTVs 2, 5, 7, 8, 9, 13, 15, 16 and 19 are Scratch



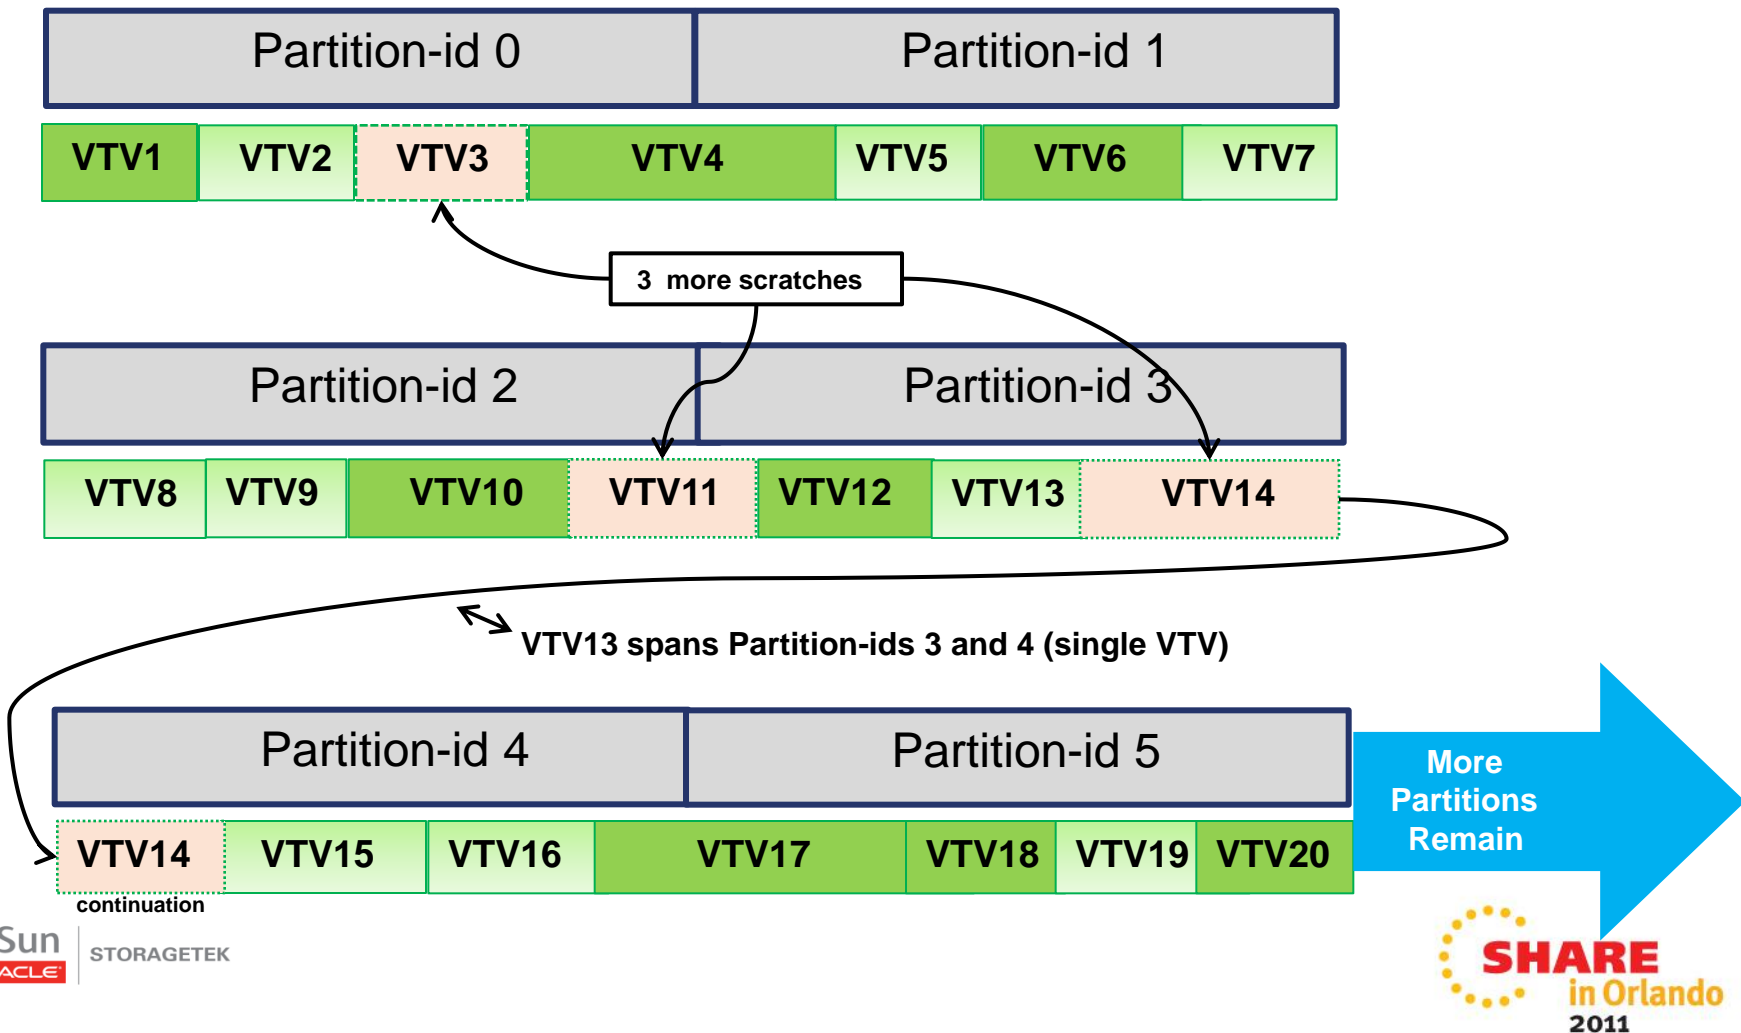
More
Partitions
Remain



How Dynamic Reclaim Works (cont'd)

- Tape continues to fragment – 3 more VTVs Scratched


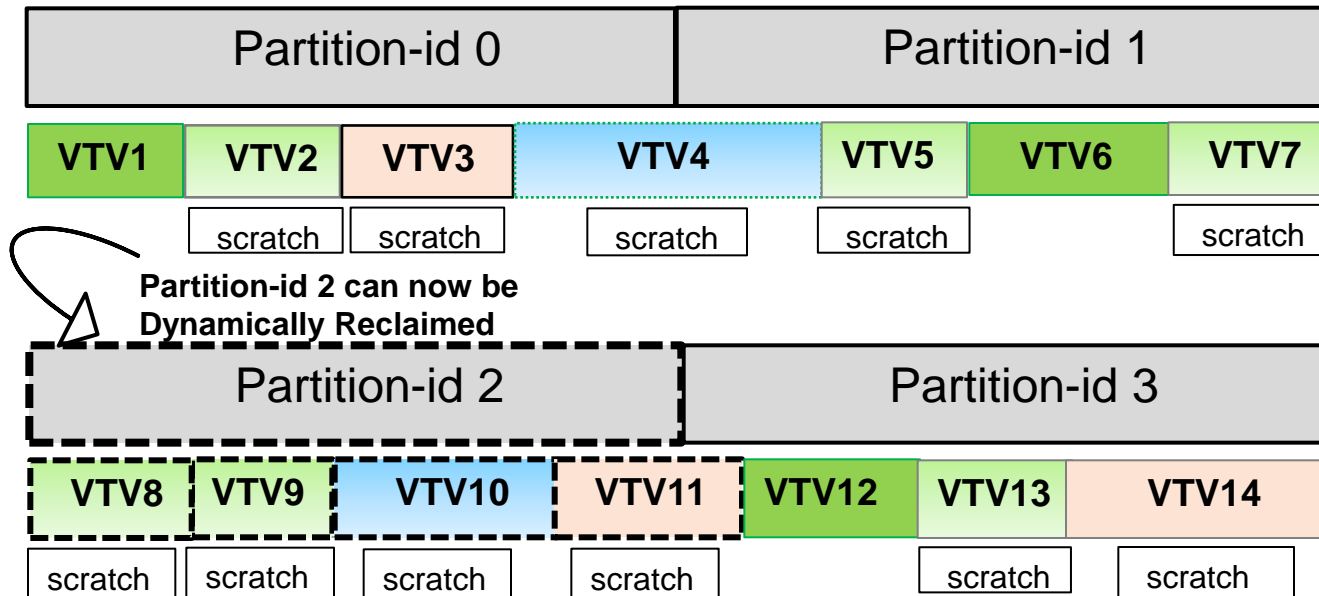
Note: All Partitions still have active VTVs and are not eligible to be Dynamically Reclaimed



How Dynamic Reclaim Works (cont'd)

- Tape fragmentation continues - four more VTVs Scratched

Note: Partition-id 2 has no more active VTVs and becomes eligible for Dynamic Reclaim




Dynamic Reclaim Executes

- So now we have an MVC that has been filled with VTVs
- The tape has become fragmented
- Once the tape becomes fragmented over and above the customer-defined threshold for Dynamic Reclaim, VTCS Dynamic Reclaim runs automatically
- During dynamic reclamation, each partition is examined to see if there are any active VTVs
- Only partitions with no active VTVs will be reclaimed during this execution process

Dynamic Reclaim Executes (cont'd)


- When VTCS Dynamic Reclaim runs, Partition-id 2 is reclaimed
Each Partition is examined during the reclamation process to check for valid VTVs



Partition-id 2 had no valid VTVs, has now been reclaimed and is eligible to accept new migrations when this MVC is again selected by VTCS



After Dynamic Reclaim runs, EOT is reset to Partition-id 2/Block-id 0



continuation

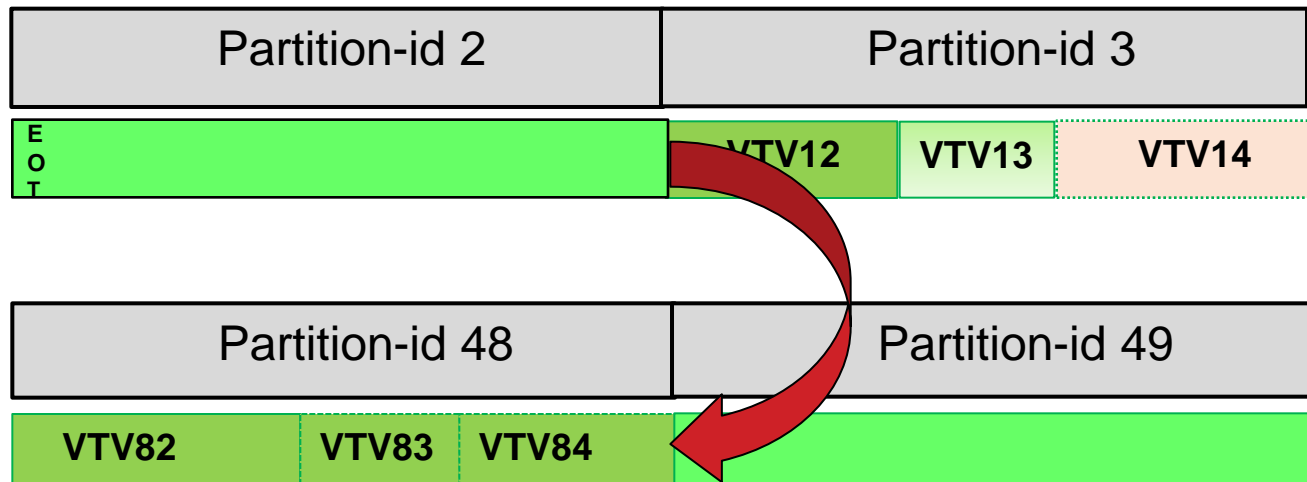
More
Partitions
Remain

After VTCS Dynamic Reclaim Runs

- Partitions-ids 0, 1, 3, 4 and 5 cannot be freed even though they contain non-current VTVs
- Partition-id 2 is the only partition freed in previous example
- The first block in partition-id 2 becomes block-id 0
- The EOT is marked at partition-id 2, block-id 0
- At this point, there are actually two partitions with block-ids of 0, partition-ids 0 and 2
- CDS is updated with new partition/block statuses

After Dynamic Reclaim Runs (cont'd)

- When VTCS again selects this MVC for migration, VTVs will start being written to the partition-id/block-id marked EOT
- When migrations start occurring to partition-id 2, the first block-id on that partition is changed to block-id 0 and then the block-ids are numbered consecutively from that point on and automatically linked to the next empty partition



System Requirements for Dynamic Reclaim



- HSC/VTCS software must be ELS 7.1 on all Hosts.
- 'H' level CDS is required.
- Microcode to support T10000B tape drives:
 - VTSS Microcode must be at D02.11.00.00 or higher
 - T10000B Microcode must be at 1.46.209 or higher (for RTDs)
- Microcode to support T10000C tape drives:
 - VTSS Microcode must be at D02.12.00.00 or higher
 - T10000C Microcode must be at 151.313 or higher (for RTDs)
- VTSSs and all attached RTDs of that device type must be at the proper microcode level to support Dynamic Reclaim.
- Supported on Oracle's StorageTek VSM5 and VSM4.

VTCS Dynamic Reclaim Enablement

- The first time an MVC cartridge is used, the format is determined by whether Dynamic Reclaim is enabled, by defining the VTCS CONFIG INPLACE parameter, which is a new parameter on the VTCS CONFIG RECLAIM Statement
The settings for this parameter are:
 - INPLACE(YES) – this specifies that Dynamic Reclaim is enabled at the global level for all Storage Classes and that the fragmentation percentage to start Dynamic Reclaim (INPTHRSH) will also be specified globally
 - INPLACE(NO) – this specifies that Dynamic Reclaim is enabled, but the parameter will be specified for certain Storage Classes on each STORCLAS statement
- It is important to point out that, regardless of whether INPLACE(YES) or INPLACE(NO) is coded, if the INPLACE parameter is present at all, then Dynamic Reclaim is enabled. There is no default value for this parameter

Dynamic Reclaim Threshold Parameters

- Along with the INPLACE parameter, Dynamic Reclaim requires that the in-place fragmentation threshold, INPTHRSH, also be defined on the VTCS CONFIG RECLAIM statement
- INPTHRSH – This fragmentation percentage is defined by the customer and it determines when a partitioned-format MVC becomes eligible for Dynamic Reclaim or Demand Reclaim. This percentage must be a value that is lower than the THRESHLD percentage value
- INPTHRSH values are between 3 and 97. The default is half the THRESHLD value

Other Considerations

- CDS should be re-sized to accommodate new 'H' level records that support partitioned MVCs
 - The latest CDS Size Calculator (H Format V7.1, Nov 2010) is available on the MOS (My Oracle Support) VSM Community
- What should the INPTHRSH percentage setting be?
 - Varies from customer-to-customer -and- depends on:
 - the customer's average VTV sizes –and–
 - the mix of expiration dates in each Storage Class
 - The default setting for INPTHRSH is half of THRESHLD
 - set lower if using larger VTV sizes -or–
 - set higher if using smaller VTV sizes
 - Customers tend to have a mixture of VTV sizes. Therefore, the recommendation would be to use the default setting in the beginning and monitor for the environment
- RTV utility supports partitioned MVCs
 - requires PTFs L1H161X and L1H1632

Thank You

Questions ?