

Improve Recall Performance with Oracle StorageTek VSM's Concurrent Tape Recall/Mount Feature

Irene Adler
Oracle Corporation

Session 10872
Thursday, March 15, 2012 4:30pm – 5:30pm



Agenda

- What problems are we trying to address?
 - VSM customers want to use larger sizes for their virtual volumes, such as 4GB
 - However, because of the recall time it takes to read the dataset back into the VTSS buffer, customers intentionally use smaller VTV sizes – 400MB instead of 4GB
 - How to use larger VTV sizes and yet improve recall time?
- The solution: VSM's Concurrent Tape Recall/Mount Feature
- How is it different from Standard Recall
- How does Concurrent Tape Recall/Mount work?
- Software/Hardware Requirements
- How to Enable the Feature

The Problem

- When a virtual volume no longer resides in the VTSS (Virtual Tape SubSystem), VTCS (Virtual Tape Control Software) must recall the volume back into the buffer from a tape residing in a library by performing the following steps:
 - Mount a MVC (Multi-Volume Cartridge) on an RTD (Real Tape Drive). This includes the robotic time for the mount and the threading of the tape in the drive.
 - Locating to the start of the VTV on the MVC. On average this is half way down the MVC.
 - Recalling the VTV (reading it off the MVC) back into the VTSS buffer.
 - Mounting the VTV on the requested Virtual Tape Drive.

Recall Timing Factors

- The speed of the tape drive is a factor in recall time.
 - Oracle StorageTek T9840x family of tape drives are higher speed drives; whereas,
 - Oracle StorageTek T10000x family of tape drives are higher capacity (T10KB 1 terabyte; T10KC 5 terabyte native capacity)
- Thread time is another recall factor in mount and locate time
 - Mount and locate times are significantly faster on a T9840x drive than on a T10Kx drive.
 - The library mount time is approximately the same, but the T9840x family threads the tape much faster (4 seconds vs. 18 seconds).

Recall Timing Factors (cont'd)



- Cartridge design method is yet another recall factor in mount and locate time.
 - T9840 cartridges are center mounting, so that the average locate is only one fourth of the media length instead of one half (11 seconds vs 41 seconds).

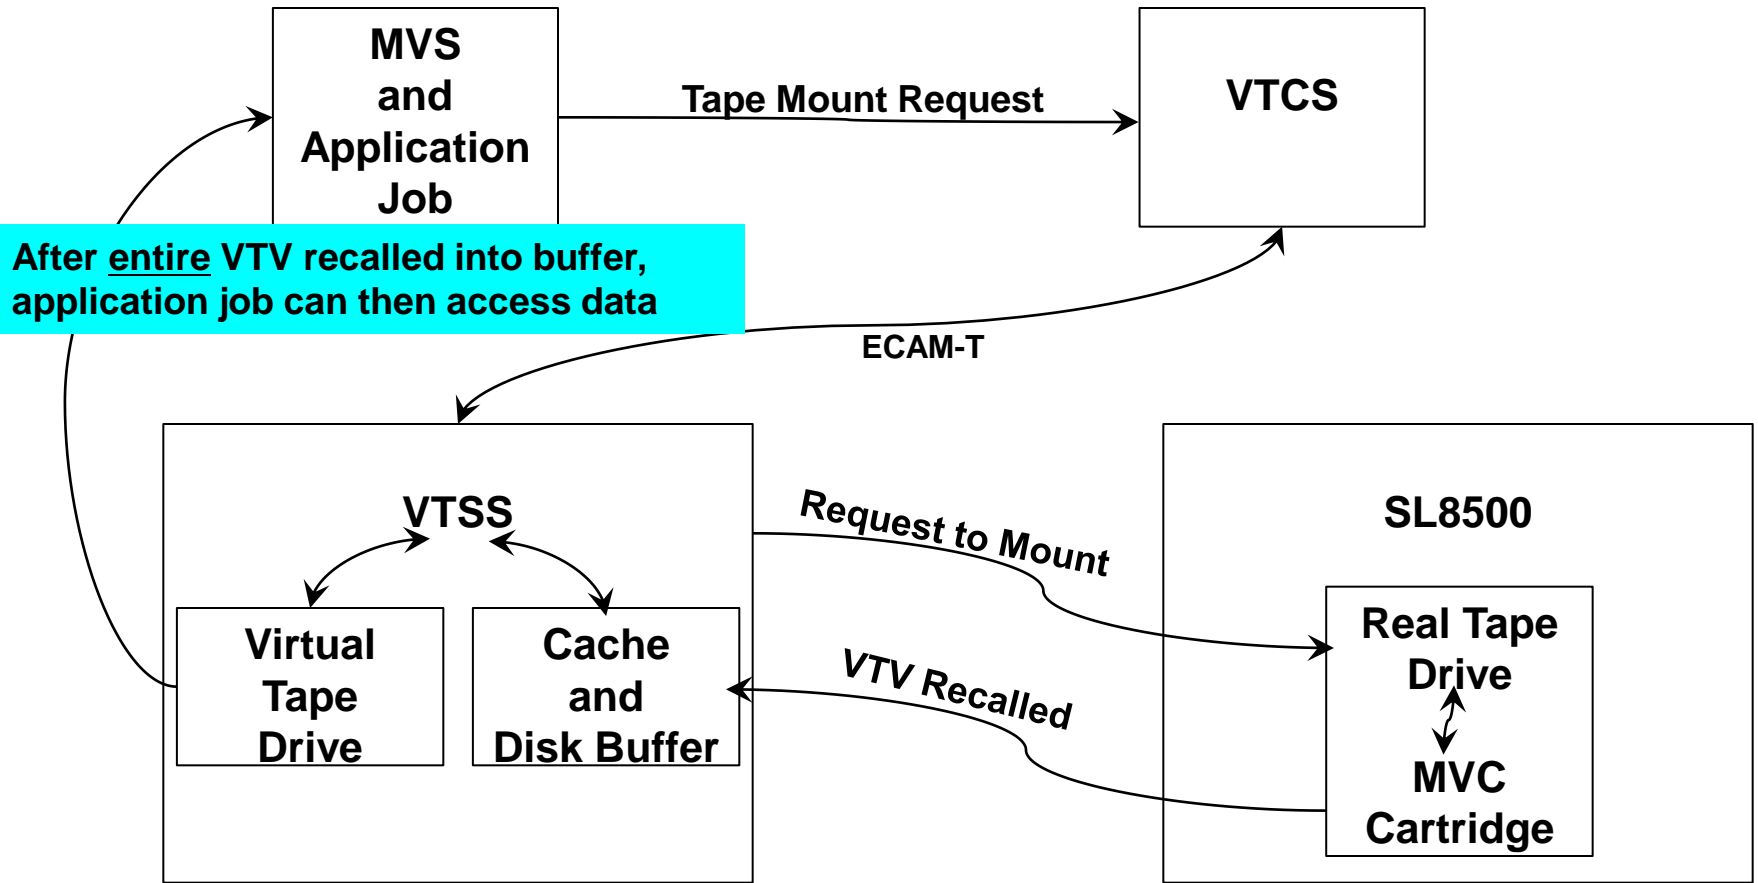
What Does All This Mean?

- There was a time when VTVs could only be 400MB or 800MB in size. Customers wanted larger VTVs.
- Now 2GB or 4GB sizes can be defined by the customer.
- So now, what was the problem?
- Customers had a strong desire to move to using larger VTV sizes – especially those who have chosen to go with the higher capacity T10000x tape drives. This, however, is not possible for those applications that are highly impacted by the longer recall times.
- Many customers decided to stay with the smaller VTV sizes rather than wait for the VTV to be recalled in its entirety before the data can be accessed by the application.

Regular Tape Mount/Recall

- Recalls of VTVs from tape has always meant that the entire VTV had to be recalled back into the buffer before the application could begin to access the data.
- The time to recall the VTV is determined by the recall data transfer rate *times* the size of the VTV. The VTSS FICON recall data transfer rate is 18 MB/sec from real tape. Depending on the size of the VTV, this can be the dominant portion of the overall recall time. Because of this, many customers will intentionally use smaller VTV sizes, 400 MB vs. 4GB. This reduces the recall time, but significantly increases the number of VTVs that the customer, MVS, and VTCS have to manage.

Regular Mount/Recall Request



The answer . . .

- VSM Concurrent Tape Recall/Mount is a new feature in VSM which will drastically improve the mount time for a Virtual Tape Volume (VTV) that is not resident on the Virtual Tape Subsystem (VTSS) buffer.

Concurrent Tape Recall/Mount Feature

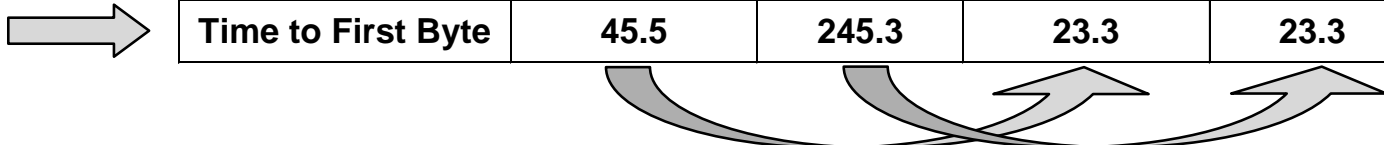


- This feature was introduced in January 2012 and is also referred to as “Early Time To First Byte” or by the acronym ETTFB.
- Mounting a VTV on a VTD when the VTV is resident in a VTSS buffer can be very quick. However, when the VTV is not resident in the VTSS, the mount will be delayed and, thus be slower.
- Delays in mount processing will also result in delaying the time before the data can be accessed as well as increasing the overall run times for jobs.

Concurrent Tape Recall/Mount Feature

- The following table is a summary of **approximate** recall performance. The interesting number is the time until the application can read the first byte of data from the requested VTV (Time to First Byte).

All Times in Seconds	Time Without ETTFB		With ETTFB	
	SL3000/T9840D		SL3000/T9840D	
	400MB VTV	4GB VTV	400MB VTV	4GB VTV
Mount Time	12.3	12.3	12.3	12.3
Locate Time	11	11	11	11
Recall Time	22.2	222	22.2	222
VTV Mount Time	0	0	0	0
Time to First Byte	45.5	245.3	23.3	23.3



How Does ETTFB Work?

- The concept of the ETTFB feature is to reduce the time to first byte as seen by the application job.
- This is achieved by overlapping the VTV recall and mount.
- Instead of waiting for the entire recall to complete, after the recall has been initiated, the VTV is concurrently mounted on the VTD as well.
- This allows the application to access the VTV while it is still being recalled.
- VTSS is responsible for pacing the VTD during the recall process. If the application attempts to read part of the VTV that has not yet been recalled, the application's I/O request will be blocked until the required VTV data has been recalled.

More on the ETTFB Feature

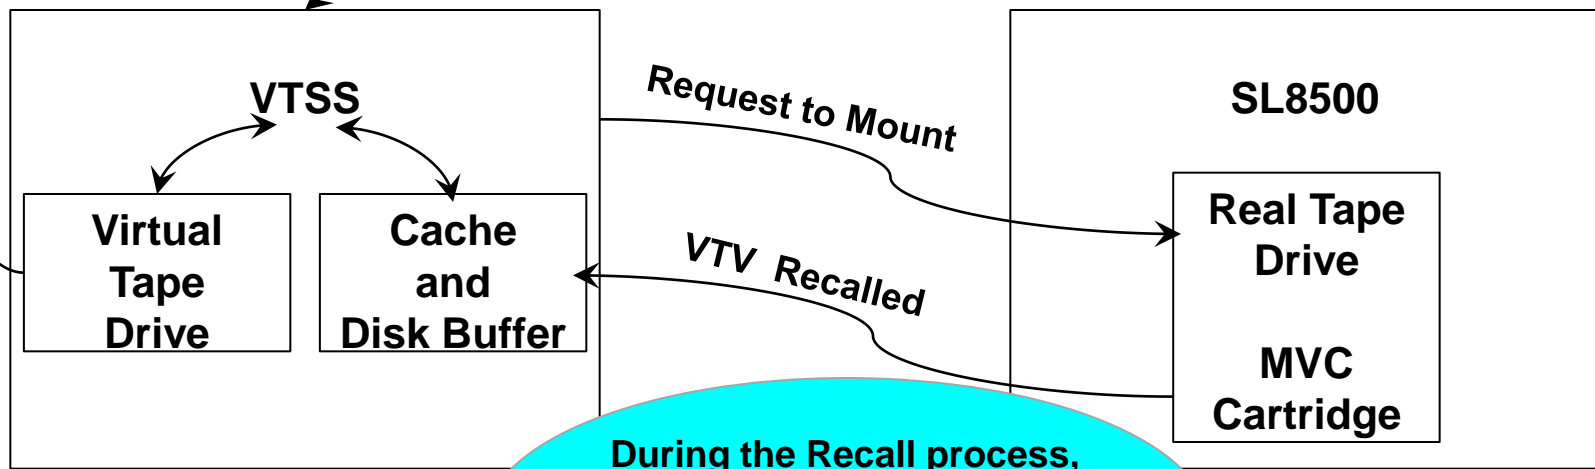
- For applications that access the VTV data in a mostly serial manner, the recall will typically proceed faster than the application can read the VTV. At 4:1 compression, an 18MB/sec recall rate can provide data to the application at 72MB/sec.
- Unfortunately, the ETTFB feature provides less of a benefit for those applications that stack multiple files on a single VTV, such as HSM and image management applications. In these types of applications, the desired data is usually not at the beginning of the VTV, but rather at some random location on the VTV. So, on the average, half of the recall will have to complete before the desired data is available.

Concurrent Tape Mount/Recall Request



As VTV blocks are recalled into buffer, the application job can begin to access the data

ECAM-T



During the Recall process, the VTSS paces the VTD, not allowing application to read data that has not yet been recalled into buffer

What ETTFB is *NOT* Used For

- ETTFB is intended for recalls resulting from a specific mount request and where the VTV is not already resident in the VTSS buffer. Therefore, the ETTFB feature is **not** used for:
 - MVC Reclaim, or
 - MVC Drain

Software and Hardware Requirements

- The ETTFB feature requires both VTSS microcode and ELS software to be at specific levels.

Hardware

- VTSS Microcode D/H02.14.24.00 or higher will support ETTFB

Software

- This feature is available on ELS 7.0 and ELS 7.1 with the appropriate PTFs installed:
 - ELS7.0 (FMID SSEA700) L1H16MY (with pre-req L1H1672)
 - ELS7.1 (FMID SSEA710) L1H16MZ (with pre-req L1H1674)

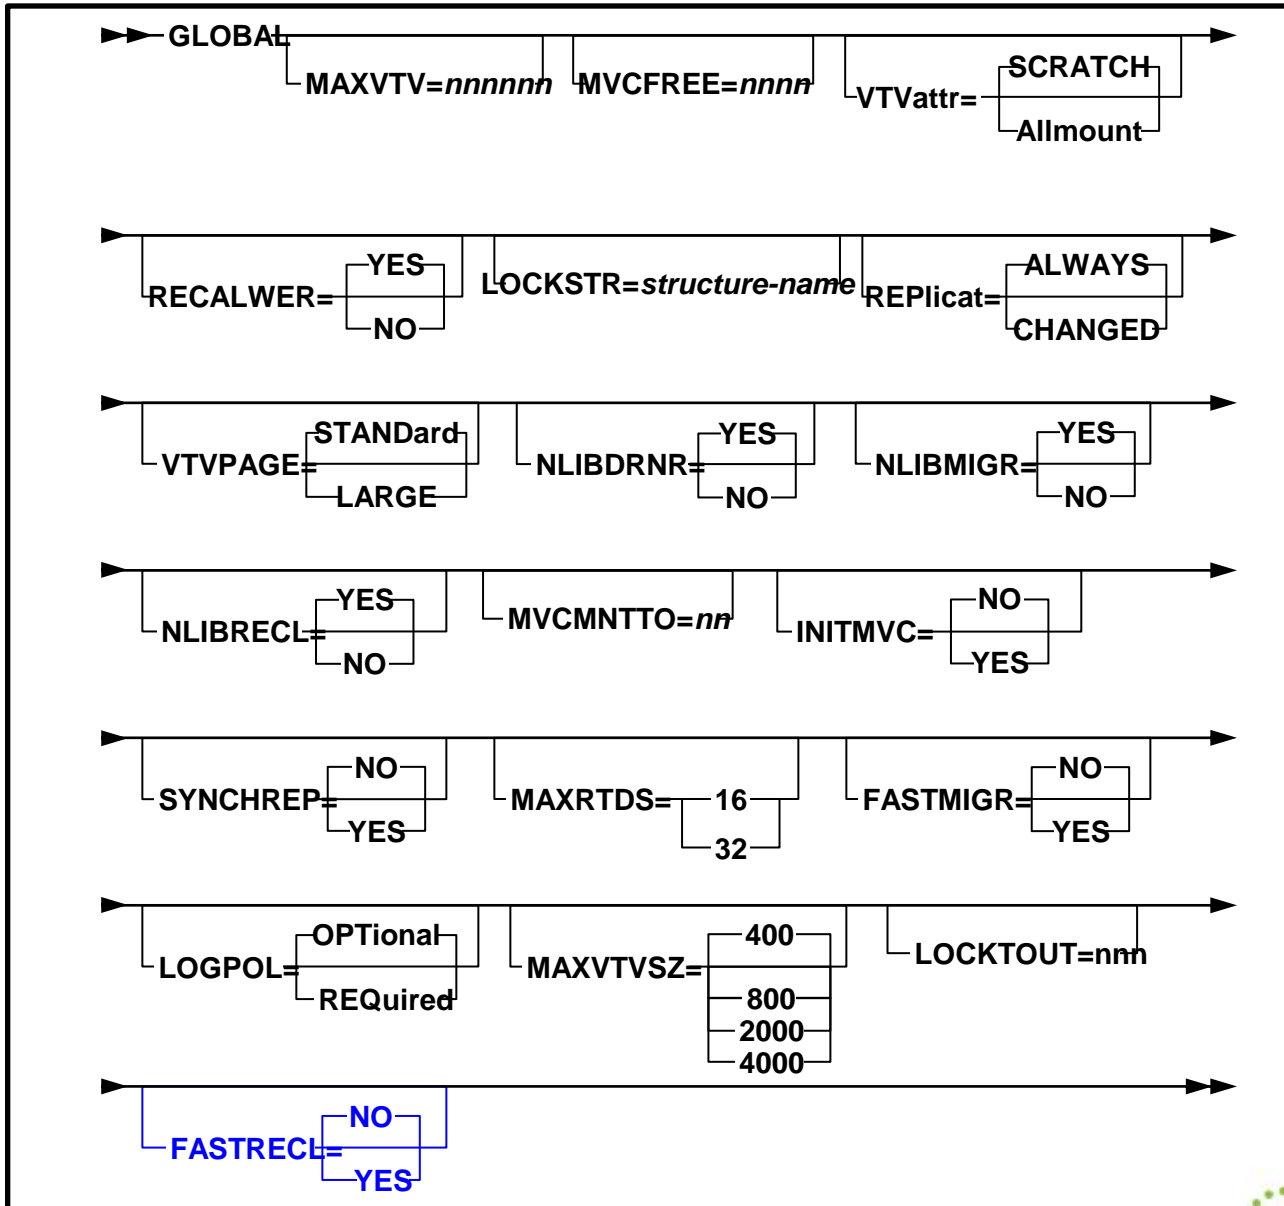
Note: VTCS 6.2 cannot use this feature. However, it is allowed to co-exist in an environment with ETTFB enabled with following PTF:

- VTCS6.2 (FMID SWS6200) L1H16NN (supersedes L1H16JV)

How to Enable ETTFB

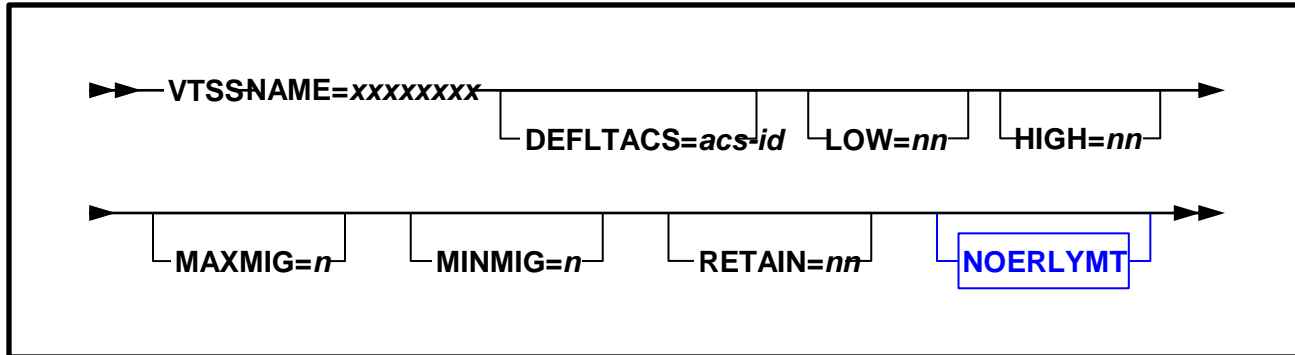
- Two new VTCS CONFIG parameters are being provided:
 - A new VTCS CONFIG GLOBAL statement parameter FASTRECL=YES/NO will enable or disable the ETTFB feature for all VTSSs in the configuration.
 - Note: This feature is Disabled by default.
 - A new VTCS CONFIG VTSS statement parameter NOERLYMT will allow the customer to disable the ETTFB feature for an individual VTSS. This parameter is only relevant if the FASTRECL=YES is globally specified.

VTCS CONFIG GLOBAL STATEMENT TO ENABLE



VTCS CONFIG VTSS STATEMENT TO DISABLE

- The ETTFB feature is enabled by the FASTRECL=YES on the GLOBAL Statement. To disable the feature for a specific VTSS, code the NOERLYMT parameter on the CONFIG VTSS Statement as shown below:



Error Handling During Recall

- The ETTFB feature interfaces with MVS's DDR Swap to assist with the recovery in the case of a recall failure. This causes the VTV to be dismounted and then remounted in the case of a Read Data Check being reported back to MVS.
- The VTSS will report a Read Data Check on the VTD if the recall fails for any reason.
- If the recall fails after the VTD has been unloaded, then there is no need to notify MVS or the application. The recall is failed back to VTCS and the partial VTV is deleted from the VTSS.
- If the RECALWER(YES) parameter is specified, this will not work while ETTFB is enabled. If a failed ETTFB recall occurs, VTCS will not leave a partial VTV on the VTSS.

Error Handling During Recall (continued)



- There are two cases of error handling:
 - 1) If the mount of the VTD was satisfied, then MVS and DDR swap will re-drive the mount request. VTCS treats this as a new mount, with the caveat that VTCS must remember that the first attempt failed. In order to do this, the VTV is flagged in the CDS as no longer being eligible for Concurrent Tape Recall/Mount.
 - On the retry, if there is one, VTCS will perform the recall as a normal, non-ETTFB recall.
 - 2) If the mount of the VTD was not satisfied, then VTCS is still responsible for completing the recall and then mounting the VTV on the requested VTD. However, VTCS will then perform the recall as a normal, non-ETTFB recall.
 - If this occurs, then the VTV is again flagged in the CDS as no longer being eligible for ETTFB.

VTV ETTFB Error Flag Status

- VTVs that have been through either of the two previously mentioned processes will have had the CDS flagged that the VTV is no longer being eligible for ETTFB. This error condition remains in place for this VTV until one of the following occurs:
 - VTV is migrated to a new copy
 - VTV is scratched
 - A new version of the VTV is created
 - VTVMaint SCRATCH(ON) is run against the VTV
 - VTV is Imported
- When any of the above changes occur, the flag in the CDS is reset and the VTV again becomes eligible for ETTFB handling.

To Summarize

- Customers want to use larger VTV sizes – 4GB vs. 400MB.
- In order for customers to define large VTV sizes, VTV recall times need to be faster.
- Oracle StorageTek tape drive offerings have increased speeds and feeds significantly in recent years.
- Regardless, it still takes a certain amount of time to recall a VTV from a tape back into the VTSS buffer if it is not already resident there.
- The desired objective is to improve on the recall time by allowing the application to begin accessing the data as it is still being recalled from the tape into the buffer.
- The feature that will allow that to happen is called VTCS Concurrent Tape Recall/Mount.

Questions ??