



An Oracle White Paper
September 2010

StorageTek Virtual Library Extension Overview

Introduction: Why would a customer want an Oracle VLE?

Oracle's StorageTek Virtual Library Extension (VLE) provides customers the flexibility to keep more active data sets within disk for longer periods of time and at the same time complement tape storage. This enables a wider range of applications, e.g. HSM, SAR, or image-type data such as billing statements, to be utilized more effectively within Oracle's StorageTek Virtual Storage Manager (VSM) subsystem.

These data types typically need to be stored within the disk longer because of their usage patterns, which could be from 45 to 90 days, before the probability of re-use diminishes. Without VLE there is more reliance upon tape storage, which often creates situations where data is being constantly migrated and recalled back and forth from tape. This is not an effective use of tape resources.

VLE also allows customers to not utilize tape or not have tape directly attached, if they desire. However, this is not recommended, since tape is the most efficient means to store data for longer periods of time. A combination of StorageTek VSM 5, StorageTek VLE and tape storage (StorageTek 9840 and T10000 drives) provides the most reliable and economical storage system for data.

It is important to note that the VLE is not just commodity disk but was designed specifically to be implemented within the VSM architecture as a second tier of disk storage. VLE provides virtualization functionality as well as policy management - a differentiation from commodity disk.


What is StorageTek Virtual Library Extension (VLE)?

The VSM5/VLE solution consists of:

- VSM5 Virtual Tape Storage Subsystem (VTSS) hardware and microcode (VSM5s with IFF3 cards).
- Virtual Tape Control Subsystem (VTCS) software and its co-requisite software.
- VLE Hardware

The VLE provides tier 2 disk storage for VSM5 only. To Virtual Tape Control Software (VTCS), a VLE looks like a single tape library (closely equivalent to a SL3000 library) but instead of tape, the Virtual Tape Volumes (VTVs) are stored on disk. With VLE, you can configure either a disk and tape or a disk-only (for example, Tapeless VSM configurations) back-end VTV storage solution. A VSM5 VTSS can migrate VTVs to and recall them from a VLE, just as is done with a real tape library. Because VTCS sees the VLE disk storage as VMVCs, VTCS can control the VLE with existing tape policies and processing.

Figure 1. StorageTek Virtual Library Extension Architecture



What are Oracle VLE Capacity Options?

The VLE will provide four JBOD capacity options: 55TB, 110TB, 165TB, and 220TB native capacity options all in a ZFS RAID Z3 configuration for high reliability. Since the data flowing into the Virtual Library Extension is compressed, effective capacities of 220TB, 440TB, 660TB and 880TB (respectively) can be achieved, assuming a device-standard 4:1 compression ratio. Each VMVC is equal to 250GB (native). The VLE uses 2TB SAS disk drives. The following table summarizes the VLE offered configurations:

VIRTUAL LIBRARY EXTENSION SPECIFICATIONS				
CUSTOMER USABLE NATIVE CAPACITY	CUSTOMER USABLE EFFECTIVE CAPACITY (AT 4:1 COMPRESSION)	VMVC Count (250GB native, 1TB effective ea.)	Number of JBODS	Number of Disk Drives
55TB	220TB	220	2	48
110TB	440TB	440	4	96
165TB	660TB	660	6	144
220TB	880TB	880	8	192

An X-Option will be provided that will allow upgrades to the largest (220TB native) from any of the smaller configurations.

What differentiates VLE as not just commodity disk?

The data is protected using ZFS RAID Z3. RAID groups are formed in an 8 + 3 configuration. Notice that 3 parity drives are used (traditionally 2 are used in RAID) for higher reliability. Second, VTSS-VLE provides end-to-end data integrity. The VTSS appends a parity (CRC) on each page of data. CRC checks are performed by both the VTSS and the VLE during migrate and recall operations to the VLE.

How is the VLE connected?

VLEs are connected to the VSM5 VTSS via multiple TCP/IP connections. The following connectivity options are supported:

- A single VLE can connect up to 8 VTSSs, so multiple VTSSes can share a single VLE. VLE supports 16 physical connections, but having at least 4 connections per VTSS is recommended.
- A single VSM5 VTSS can connect to up to 4 VLEs to increase effective buffer space.

- A single VSM5 VTSS can attach to (in any combination): VLEs, Real Tape Drives, or other VSM5s.

How is the VLE managed?

The Virtual Library Extension is managed via already existing policy management. Utilizing a management class with a storage class of virtual library configures the product. When you set up a VLE, as with a normal library you load it with volumes (define VMVCs on the VLE). With archive, existing archive policies can be exploited to migrate data first to a VLE and the later after some period time data can move to real tape.

The following are simple examples of typical VLE configurations:

Figure 2.0. Two copies of data; one copy on VLE and one copy on Tape.

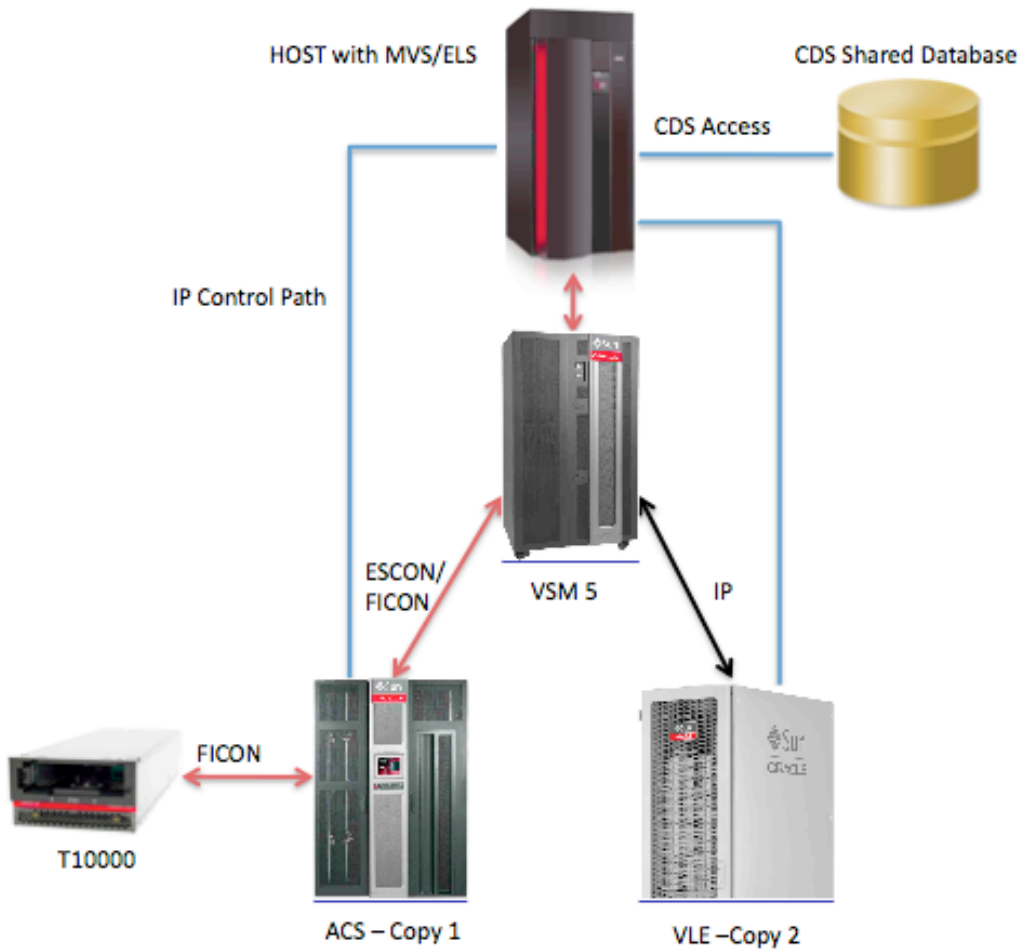


Figure 2.1. One copy of data on VLE with spillover to Tape.

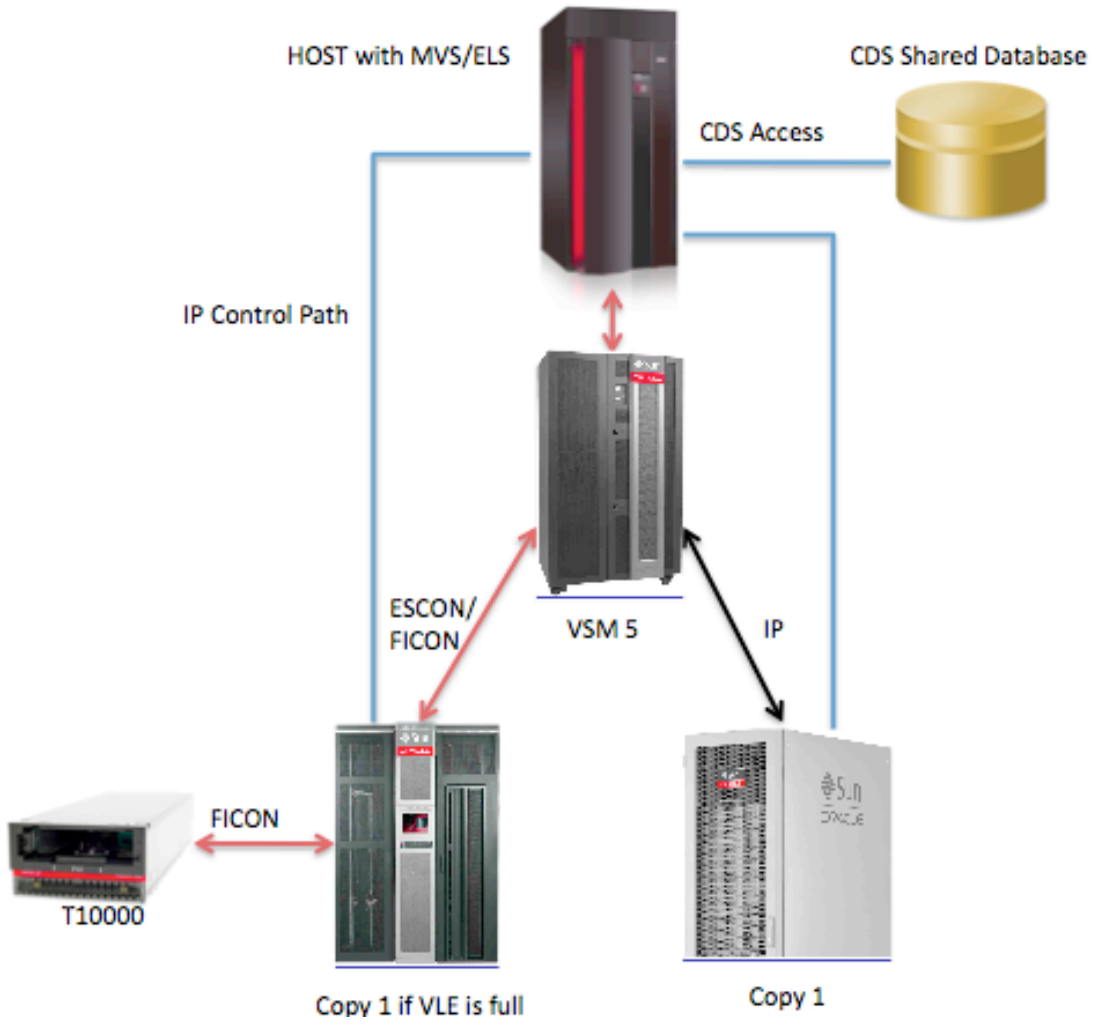
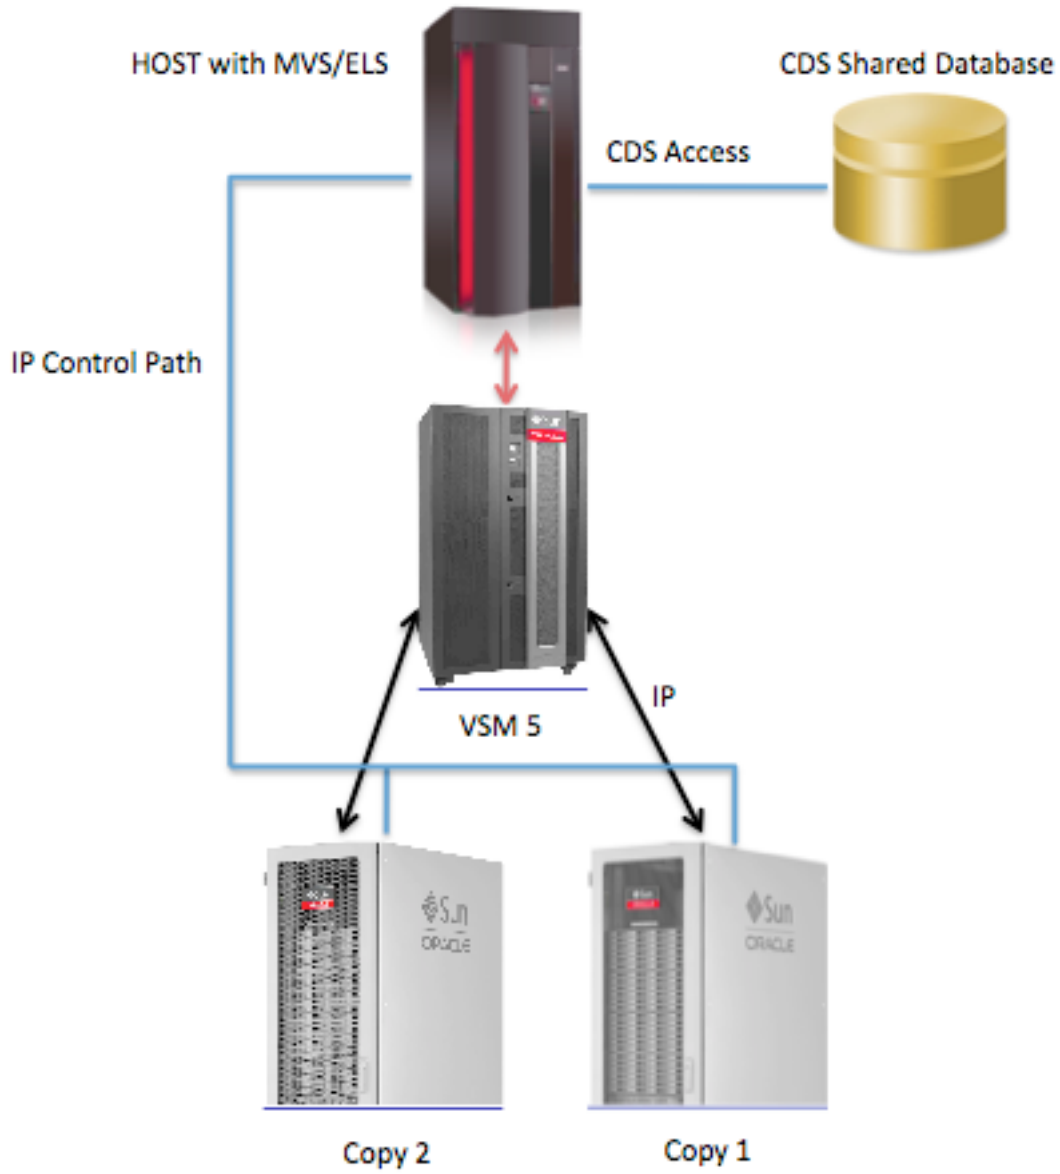


Figure 2.2. Two copies of data - one copy on VLE1 and one copy on VLE2.



What can we expect in terms of performance?

Each VTSS IP port is rated at 40MB/Sec, which is native performance. If a 4:1 compression ratio is attained the effective data rate would be 160MB/Sec of customer data. The throughput between VSM 5 and VLE in aggregate can reach up to 640MB/s (across 16 ports).

COMPARISONS BETWEEN ENTERPRISE TAPE DRIVES AND VLE

COLUMN HEAD	T10000B	9840D	VIRTUAL LIBRARY EXTENSION
Cartridge Capacity (Native)	1000GB	75GB	250GB
Migrate Rates (Native)	120MB/Sec	30MB/Sec	40MB/Sec
Recall Rates (Native)	120MB/Sec	20MB/Sec	40 MB/Sec
Reclaim	Yes	Yes	Yes*
Average MVC Locate Time for a Recall	41 Seconds	11 Seconds	0 seconds
SL3000 Robotic Mount and Drive Load Time	26 Seconds	12 Seconds	0 Seconds

*On real tape, expired VTVs are freed from an MVC by moving current VTVs back into a VTSS buffer (recall) and then re-writing into tape (migration). On VLE reclaim operates by deleting the expired VTVs from disk – no data transfer.

Best Practices

A single VLE is NOT a high availability (HA) solution. It is expected to be 99.99% available. Like the existing Real Tape VSM solution, it is highly recommended that at least two copies of each VTV be made.

A mixed VLE and real tape environment would require at least two storage classes to be defined. One storage class would be for the VLE virtual MVCs (VMVC) and one storage class for the real MVCs.

A tapeless environment would require at least two VLEs. Storage classes then need to be defined so that the two or more MVC copies are on different VSM VLEs. Unlike real tape, it is not sufficient ensuring that the multiple copies are on different MVCs.

Each VSM5 with IFF3 cards has four IP connectors, one per IFF3 card. Each IP connector is capable of supporting four concurrent connections for a total of sixteen connections.

For redundancy and throughput, it is recommended that at least one connection per connector is configured for each attached VLE or clustered VSM5. Thus, it is recommended that the VSM5 (via IP) is only connected to four target boxes (four connections per connector). This implies that up to four VLEs may be attached to a single VSM5.

The VLE can support up to 32 connections across 16 physical IP connectors. Therefore, up to 8 VSM5s (8 vsm5 * 4 connections each = 32 total connections) can be connected to a single VLE.

To avoid “islands” of data, it is recommended that all VSM5s in the tapeplex are connected to all VSM VLEs in the tapeplex.



Overview of StorageTek Virtual Library
Extension
September 2010

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com



| Oracle is committed to developing practices and products that help protect the environment

Copyright © 2010, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0810

SOFTWARE. HARDWARE. COMPLETE.