

Survey Of Volume Managers

Nasser M. Abbasi

May 24, 2000

Compiled on July 20, 2022 at 10:37pm

Contents

1	Advantages of Volume Managers	1
2	Terminology used in LVM software	1
3	Survey of Volume Managers	2
3.1	SGI XLV. First generation volume manager	2
3.2	SGI XVM Volume Manager. Second generation, enhanced XLV	4
3.3	Linux LVM (Logical Volume Manager)	4
3.4	Veritas Volumn Manager. (Sometimes called VxVM)	5
3.5	The Vinum Volume Manager	7
3.6	Compaq SANworks Enterprise Volume manager	7
3.7	Sun StorEdge Volume Manager (SSVM). (repackaged version of Veritas Volume Manager.)	8
3.8	AIX Logical Volume Manager	9
3.9	Sequent Volume Manager (ptx/SVM)	9
3.10	HP Shared Logical Volume Manager (SLVM)	11
4	Other links to Volume Managers	12

1 Advantages of Volume Managers

From <http://www.uwsg.iu.edu/usail/peripherals/disks/logical/>

Some of the advantages of using a logical volume manager:

1. They offer greater flexibility for disk partitioning.
2. The size of logical volumes can be modified according to need, while the operating system is running
3. Logical volumes can span multiple disks.
4. Disk mirroring is often supported, for greater data reliability

2 Terminology used in LVM software

Figure below Shows the main terms used in Logical Volume Manager software.

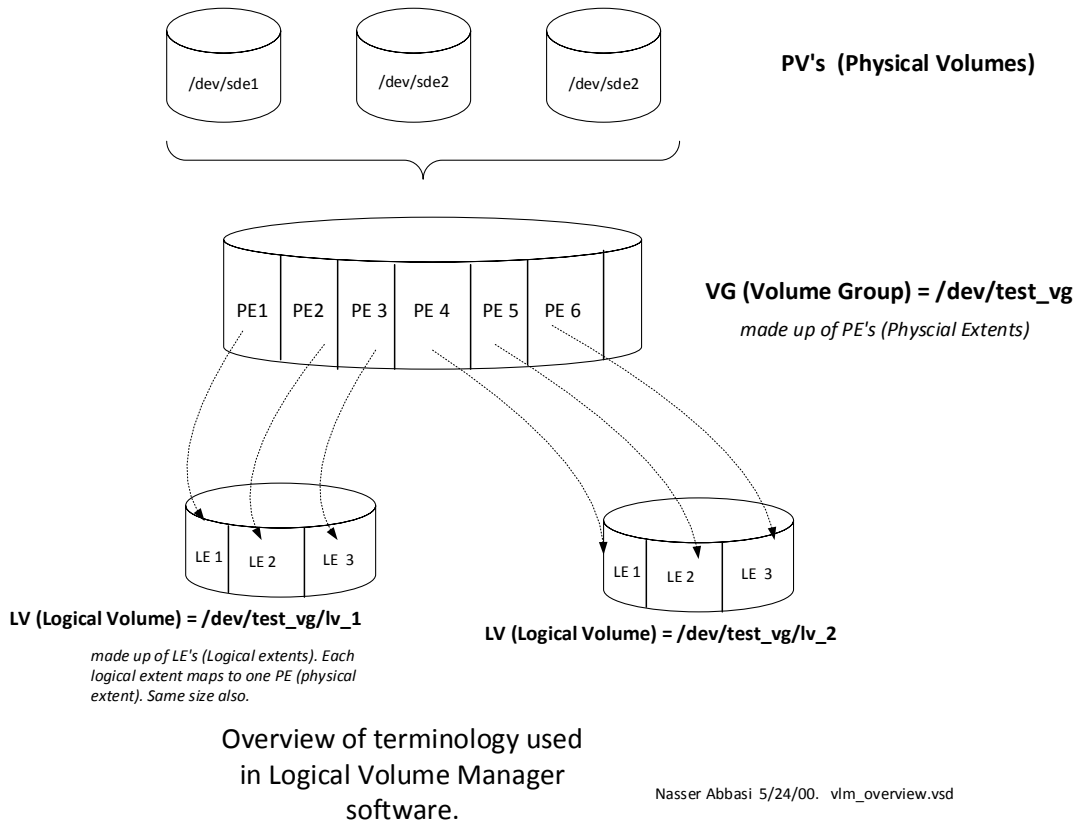


Figure 1: Terms used in LVM software.

3 Survey of Volume Managers

This a list of volume managers that I collected searching the web:

3.1 SGI XLV. First generation volume manager

XLV is SGI volume manager that is integrated in XFS (SGI file system). see <http://www.sgi.com/Technology/xfs-whitepaper.html>

From above URL:

3.5.1 Basics

The xlv volume manager (XLV) is an integral part of the XFS filesystem(1). The volume manager provides an operational interface to the system's disks and isolates the higher layers of the filesystem and applications from the details of the hardware. Essentially, higher-level software "sees" the logical volumes created by XLV exactly like disks.

Yet, a logical volume is a faster, more reliable "disk" made from many physical disks providing important features such as the following (discussed in detail later):

- concatenating volumes for a larger disk
- striping volumes for a larger disk with more bandwidth
- plexing (mirroring) volumes for a more reliable disk

The use of volumes enables XFS to create filesystems or raw devices that span more than one disk partition. These volumes behave like regular disk partitions and appear as block and character devices in the /dev directory. Filesystems, databases, and other applications access the volumes rather than the partitions. Each volume can be used as a single filesystem or as a raw partition. A logical volume might include partitions from several physical disks and, thus, be larger than any of the physical disks.

Filesystems built on these volumes can be created, mounted, and used in the normal way.

The volume manager stores all configuration data in the disk's labels. These labels are stored on each disk and will be replicated so that a logical volume can be assembled even if some pieces are missing. There is a negligible performance penalty for using XLV when compared to accessing the disk directly; although plexing (mirroring data) will mildly degrade write performance.

3.2 SGI XVM Volume Manager. Second generation, enhanced XLV

See <http://www.sgi.com/developers/devnews/2000/xvm.html>

3.3 Linux LVM (Logical Volume Manager)

From <http://linux.msede.com/lvm/> by Heinz Mauelshagen. Above URL contains good presentations and more documentation on LVM. This is based originally on the OSF LVM.

From above URL

The Logical Volume Manager (LVM) is a subsystem for on-line disk storage management which has become a de-facto standard for storage management across UNIX implementations.

The Logical Volume Manager adds an additional layer between the physical devices and the block I/O interface in the kernel to allow a logical view on storage. Unlike current partition schemes where disks are divided into fixed-sized continuous partitions, LVM allows the user to consider disks, also known as physical volumes (PV), as a pool (or volume) of data storage, consisting of equal-sized extents.

A LVM system consists of arbitrary groups of physical volumes, organized into volume groups (VG). A volume group can consist of one or more physical volumes. There can be more than one volume group in the system. Once created, the volume group, and not the disk, is the basic unit of data storage (think of it as a virtual disk consisting of one or more physical disks).

The pool of disk space that is represented by a volume group can be apportioned into virtual partitions, called logical volumes (LV) of various sizes. A logical volume can span a number of physical volumes or represent only a portion of one physical volume.

The size of a logical volume is determined by its number of extents. Once created, logical volumes can be used like regular disk partitions - to create a file system or as a swap device.

LVM was initially developed by IBM and subsequently adopted by the OSF (now OpenGroup) for their OSF/1 operating system. The OSF version was then used as a base for the HP-UX and Digital UNIX operating system LVM implementations. Another LVM implementation is available from Veritas which works differently. The Linux implementation is similar to the HP-UX LVM implementation.

3.4 Veritas Volume Manager. (Sometimes called VxVM)

From <http://www.sun.com/storage/software/volumemgr.html>

VERITAS Volume Manager software is an advanced, system-level disk and storage array solution that alleviates downtime during system maintenance by enabling easy, online disk administration and configuration. The product also helps ensure data integrity and high availability by offering fast failure recovery and fault tolerant features. VERITAS Volume Manager software provides easy-to-use, online storage management for enterprise computing and emerging Storage Area Network (SAN) environments.

Through the support of RAID redundancy techniques, VERITAS Volume Manager software helps protect against disk and hardware failures, while providing the flexibility to extend the capabilities of existing hardware. By providing a logical volume management layer, VERITAS Volume Manager overcomes the physical restriction imposed by hardware disk devices.

Highlights:

Data redundancy (RAID 0, 1, 0+1, 1+0, 5).

Dynamic multipathing (DMP) support.

Intuitive Java technology-based platform-independent graphical user interface.

Enables easy movement of data between nodes in a SAN environment.

Features	Benefits
HIGH AVAILABILITY -----	
On-line storage reconfiguration and database table space growth	Uninterrupted data access during planned system maintenance
Mirroring of storage (RAID-1)	Higher availability with fault tolerance and faster data access
Data redundancy (RAID-5, RAID-1+0, RAID-0+1)	Reduces storage costs for data redundancy or by maintaining access to data even when multiple disks fail
Hot relocation of failed redundant storage	Automatic restoration of data redundancy when disks fail
DRL (Dirty Region Logging)	Provides fast recovery after a system failure
SCALABILITY/PERFORMANCE -----	

On-line performance monitoring and tuning tools	Tools to assist in identification and minimization of I/O bottlenecks
Striping (RAID-0) and selective disk mirroring	Increases throughput and bandwidth while providing scalable performance and balancing of application data loads
Dynamic Multipathing (DMP) support	Increases performance to multi-controller disk arrays by spreading I/O between the multiple paths into the array. Easier access to data and automatic path recovery across redundant fibre channel loops
Spanning of multiple disks	Eliminates physical storage limitations
On-line relayout	On-line tuning ability enables reconfiguration of storage layout
Task monitor	Ability to throttle volume recovery speed and monitor VM I/O tasks
CENTRALIZED MANAGEMENT -----	
Intuitive Java-based platform-independent graphical user interface	Displays logical view of storage devices, providing easy monitoring of disk I/O configuration and management that improves productivity
Device driver, file system, and database independence	Easy integration with disk subsystems and arrays including Hardware RAID systems
Free space pool management for automatic or directed allocations	Simplifies administration and provides flexible use of available hardware
Logical diskgroups	Enables easy movement of data between nodes in a SAN environment
Maximum volume size button	Automatically calculates the maximum free space available for building volumes on chosen disks
Unique disk ids	Eliminates disk misidentification allowing disks to be added and removed and configured automatically and on-line
HETEROGENEOUS SUPPORT -----	
Ability to move diskgroups between different versions of	Easier migration with reduced downtime volume Manager

Available on numerous Sun hardware platforms	Reduces training costs
Supports multiple Sun disk arrays	Provides maximum flexibility by allowing businesses to select the storage hardware solution that best meets their needs
INTEGRATION -----	
Data Snapshot	Provides mirrored data snapshot which enables on-line backup of data via backup application such as VERITAS NetBackup

More on Veritas volume manager are here <http://www.veritas.com/us/products/volumemanager/> which contains number of white papers and data sheet on the product.

3.5 The Vinum Volume Manager

From <http://www.lemis.com/vinum.html>

FreeBSD volume manager. From above URL

The Vinum Volume Manager is a block device driver which implements virtual disk drives. It isolates disk hardware from the block device interface and maps data in ways which result in an increase in flexibility, performance and reliability compared to the traditional slice view of disk storage. Vinum implements the RAID-0, RAID-1 and RAID-5 models, both individually and in combination.

See also <http://www.shub-internet.org/brad/FreeBSD/vinum.html>

3.6 Compaq SANworks Enterprise Volume manager

From <http://www.compaq.com/products/storageworks/Storage-Management-Software/evmindex.html>

Product Description

Key Features

Cloning and Snapshot - SANworks Enterprise Volume Manager use can be optimized by selecting either cloning or snapshot, depending on the application. Snapshots are virtual copies and clones are physical copies. Snapshot is ideal for quick recovery. Both are ideal for non-disruptive backup.

Web-based application - SANworks Enterprise Volume Manager is accessible from any system that has a web browser.

Multi-platform support - SANworks Enterprise Volume Manager operates on the RAID Array 8000 (RA8000) and the Enterprise Storage Array 12000 (ESA12000) using HSG80 controllers in switch or hub configurations. SANworks Enterprise Volume Manager supports Windows NT V4.0, Windows 2000, Sun Solaris V2.6, 7 and Tru64 Unix V4.0F/G. Other platform support is to follow. SANworks Enterprise Volume Manager provides consistent storage management regardless of the platform.

Plug and Play with existing applications - Microsoft Exchange, StorageWorks Enterprise Backup Solution, VERITAS NetBackup, VERITAS Backup Exec, Legato NetWorker, CA ARCserveIT, Oracle, and Microsoft SQL and with plans to support other applications in the future.

Supports LAN-less backup - Backup data is isolated from the general purpose LAN, so there is no network performance degradation during backup. All volume movement is on the SAN.

FC-AL or Switched Fibre Channel topologies supported - Customers with either technology can take advantage of the features of EVM. Snapshots are available in switch configurations only.

3.7 Sun StorEdge Volume Manager (SSVM). (repackaged version of Veritas Volume Manager.)

From <http://www.sun.co.uk/services/educational/catalog/courses/UK-ES-310.html> and <http://www.carumba.com/talk/veritas/volumemanager.shtml>

3.8 AIX Logical Volume Manager

See <http://www.redbooks.ibm.com/abstracts/sg245432.html>

3.9 Sequent Volume Manager (ptx/SVM)

From <http://www.numaq.com/products/software/layered/svm.html>

Sequent Volume Manager (ptx/SVM[tm]) provides enhanced disk management facilities and increased system availability for Sequent's NUMA-Q[tm] servers running DYNIX/ptx, Sequent's enhanced implementation of the UNIX operating system. Based on an emerging standard, VxVM[tm] from Veritas[tm] Software Corporation, ptx/SVM offers significant enhancements over traditional mirroring products. These include features that significantly improve data integrity, system availability, and system performance.

In addition to its sophisticated mirroring capabilities, ptx/SVM also provides disk striping, disk concatenation, hot sparing, and on-line disk management. With on-line disk management, a system administrator can optimize disk performance by moving data between disks while the system is running.

Drives under ptx/SVM control can be dynamically resynchronized with one or more mirrored partners, independent of the disk controller, without taking the system off-line. ptx/SVM can also control the resynchronization rate, which can be set to minimize impact on performance or to minimize the time required to perform a resync operation.

ptx/SVM offers users the advantages of open systems with its access to powerful and dynamic volume management tools.

ptx/SVM Highlights Disk Mirroring Data availability and integrity are increased with the continuous maintenance of up to 32 copies of critical data. ptx/SVM automatically uses these data mirrors in the event of a disk failure.

ptx/SVM adds greatly to system availability by allowing system administrators to dynamically create, remove, and allocate mirrors, as well as perform on-line resynchronization and snap-shot backups with a minimum impact on users.

If mirroring is used, depending on the layout, ptx/SVM may automatically divide the read load among all the mirrors, creating multiple read paths which can enhance system performance.

Disk Concatenation Disk concatenation allows a user to create logical volumes that can span multiple disks. Two or more physical disks or disk segments can be viewed as a single entity.

Disk Striping Striping allows portions of multiple disks to be viewed as a single logical entity. Striping improves performance by distributing the data of a heavily-used partition over several disks, thus increasing the number of heads available to read and write data.

Hot Sparing ptx/SVM allows the designation of dedicated disks as hot spares that are used to replicate mirrored data from disks that have failed, thus increasing the availability of mirrored data.

Disk groups ptx/SVM allows for the segregation of disks into logical groups called disk groups. Disk groups improve access to data objects by maintaining separate databases of data objects and allows for the ability to create up to 100,000 data objects providing the ability for a system to

It enables an administrator to perform ptx/SVM administrative actions (e.g., creating, modifying, and removing ptx/SVM objects) using GUI dialogs instead of the command line.

It provides top-down and detailed views of the ptx/SVM configuration.

It reports many ptx/SVM error conditions, such as I/O errors and failure of the volume configuration daemon.

Dirty Region Logging (DRL) DRL is a fast resynchronization mechanism for private storage. If a mirrored volume needs to be resynchronized due to a system crash, only the addresses with outstanding writes stored in the log need to be resynchronized.

Sequent Support Sequent offers full product support for ptx/SVM, including a manual for system administrators and training classes specifically for ptx/SVM. Consulting services are also available to assist with particular system configurations and implementations.

3.10 HP Shared Logical Volume Manager (SLVM)

From <http://www.software.hp.com/STK/partner/relnotes/10.20/chp3.html>

For 10.10, LVM has been modified to support volume groups that are connected to multiple systems. This feature is called Shared LVM (SLVM) and is supported only on S800 machines.

SLVM is a mechanism that permits multiple systems in an MC/LockManager cluster to share (read/write) disk resources in the form of volume groups. The objective is a highly available system by providing direct access to disks from multiple nodes and by supporting mirrored disks, thereby eliminating single points of failure.

SLVM permits a two system cluster to have read/write access to a volume group by activating the volume group in shared mode.

SLVM is designed to be used only by specialized distributed applications (such as Oracle Parallel Server) that use raw access to disks, rather than going through a file system. The applications must provide their own concurrency control for their data, as well as transaction logging and recovery facilities, as appropriate. Applications that are not network aware, such as file systems, will not be supported on volume groups activated in shared mode.

SLVM requires services provided by MC/LockManager and thus only clusters that have MC/LockManager will be able to use shared activation.

4 Other links to Volume Managers

- Logical Volume Manager HOWTO.
See <http://www.ds9a.nl/lvm-howto/HOWTO/cvs/lvm-howto/lvm-howto.html>
- AIX logical volume manager Overview.
see http://nscp.upenn.edu/aix4.3html/aixbman/baseadm/lvm_overview.htm
- news group for Veritas volume manager: veritas.volume_manager.english.
See http://support.veritas.com/nntp/rules_ddProduct_VOLUMEMAN_crumb_on.htm
- veritas volume manager white papers.
see <http://www2.veritas.com/us/products/volumemanagernt/whitepapers.shtml>
- Heinz Mauelshagen's LVM (Logical Volume Manager) howto.21/02/2000
see <http://linux.msede.com/lvm/doc/LVM-HOWTO>
- SGI XFS file system
see <http://www.sgi.com/Technology/xfs-whitepaper.html>
- IBM paper. Open systems, storage, and performance for the mainframe mind.
see <http://www.storage.ibm.com/hardsoft/diskdrils/technolo/open.htm>
- A Walkthrough of the Logical Volume Manager for Linux
see <http://www.gweep.net/~sfoskett/linux/lvmlinux.html>