

Overview

New Features in HP Tru64 UNIX Logical Storage Manager Version 5.1B-3

LSM now offers a configurable policy for mirrored volume I/O that will limit retries on some device errors.

This new policy is turned off by default, but when manually turned on, will improve application response time on some device errors, allowing LSM to kick out a plex faster, using another mirror to service the I/O request.

Costly delays can occur on some device errors as retries are being serviced. Consider older I/O storage subsystems, for example, where devices are growing more error-prone. Although the retries may eventually result in successful I/O completion, there is no guarantee that will happen. This policy will help kick out potentially faulty plexes faster and keep I/O moving via a serviceable plex. As with any plex detach, administrator action will be required to restore redundancy.

This policy is targeted for those environments that have strict application response time requirements and active administration.

Performance improvements were also made to the startup of LSM and disk rescan operations along with additional performance enhancements to further improve plex detach processing time in a TruCluster environment.

Product Overview

HP Tru64 UNIX® Logical Storage Manager (LSM) is a cluster-integrated, host-based solution for storage management and a layered product for Tru64 UNIX systems. LSM reduces downtime by using host based software RAID technology to create redundancy, reconstruct data from failed disks, and allow changes to the system's storage configuration without taking the system down.

LSM provides the following features:

- Data spanning across physical disks (concatenation) allows the creation of large file systems that span multiple disks.
- Data mirroring or shadowing (RAID 1) mirrors critical data protecting against loss of data and to improving read I/O performance, including support for root file system and swap space mirroring.
- Data striping (RAID 0) improves I/O performance by interleaving data across multiple disks.
- Combined mirroring of a striped set (RAID 0+1) protects data while improving system performance.
- Data striping with parity (RAID 5) provides data redundancy through the use of parity that reconstructs data in the event of a disk failure. RAID 5 logging also prevents the corruption of data by logging new data writes before committing them to the storage device.
- Comprehensive online management functions include:
 - Add and remove disks to and from LSM control
 - Create and change concatenated, mirrored, RAID 5 and striped configurations
 - Access mirrored data after multiple disk failures, if at least one copy of the data is available
 - Monitor performance
 - Migrate data for disk repair or load balancing
 - Detach a mirror for backup with minimal interruption
 - Backup and restore utilities for LSM configuration records
- Hot sparing automatically relocates mirrored and RAID 5 objects to a spare disk after an I/O failure. You can designate those disks to use as spares or let the system find free space.
- LSM is cluster integrated so that LSM operations continue despite the loss of cluster members - as long as the cluster itself continues to operate and a physical path to the storage is available. LSM manages the shared storage in a cluster as a single system, so that LSM disk groups can be used by all cluster members and the LSM configuration can be changed by any member.
- LSM is supported in an HP TruCluster Server cluster with the following restrictions:
 - LSM volumes cannot be used for the boot partitions of individual cluster members.
 - LSM cannot be used to mirror a quorum disk or any partitions on that disk.
 - LSM RAID 5 volumes are not supported in clusters.
- UNIX File System (UFS) partitions and Advanced File System (AdvFS) storage domains can be migrated to LSM volumes.
- Dirty Region Logging significantly speeds resynchronization of volumes after a system crash by resynchronizing only those

Overview

mirrored regions that are marked in the log as active at the time of the crash.

- Fast Plex Attach permits making a temporary copy of volume data available for backup. When completed, the mirrors resynchronize in much less time than usually required.
- Multiple management interfaces are provided:
 - The command line interface consists of a set of basic, but powerful shell-level commands that require input from the user. Many of these commands require a thorough understanding of LSM concepts and existing volume configuration.
 - The voldiskadm menu interface provides easy-to-use menus with prompting to do disk and disk group operations.
 - The Storage Administrator Graphical User Interface (lsmsa) is a graphical interface based on Java[®] that simplifies the management of LSM volumes and deftly handles the management of larger storage configurations. An object tree displays the hierarchical relationship between LSM objects, and the Command Launcher displays a list of tasks that you can perform on LSM objects.
 - The Visual Administrator Graphical User Interface (dxlsm) is a graphical interface based on Motif[®] that provides an intuitive approach to volume management. Using color, icons, windows, and pop-up menus, it portrays the system volume configuration and volume I/O activities and allows you to use mouse-based operations to do online storage management, reconfiguration, and performance monitoring and tuning. Storage Administrator offers much of the same functionality and scales better.

LSM functionally resides below the local file system layer and above the device driver layer and transparently supports data managed by:

- The UNIX File System (UFS)
- The Advanced File System (AdvFS)
- Other applications that use UNIX physical disk partition devices (bdevsw/cdevsw interface)

A device entry is created for each volume that provides the read/write data interface and the policies for managing data in the volume.

Hardware RAID devices are supported as standard disks. Each RAID device logical unit is viewed as a physical disk.

Configuration Limits

The following definitions and limits define the maximum configuration supported by LSM:

- Maximum of 8192 volumes per system of which three volumes are reserved for system use and two are dedicated to root and swap volumes.
- Maximum of 32 plexes per volume. A plex is a mirror (or copy) of data; non-mirrored volumes have one plex and mirrored volumes may have 2 to 32 plexes.
- Maximum of 8189 plexes per system.

Standard Features

Hardware Requirements

LSM Version 5.1B-3 operates on any AlphaServer supported by Tru64 UNIX Version 5.1B-3 or later. There are no restrictions on the disk devices supported beyond the valid configurations defined in the Tru64 UNIX Operating System QuickSpecs.

Software Requirements

Operating System

LSM Version 5.1B-3 requires Tru64 UNIX operating system Version 5.1B-3 or later.

The Storage Administrator Graphical User Interface (lsmsa) requires that Java Version 1.1.7B be installed on the system. The Java subsets are provided on the Tru64 UNIX Operating System CD-ROMs.

The Visual Administrator Graphical User Interface (dxlsm) requires that Motif be installed on the system.

Migration Tools

Migration tools are provided to migrate:

- Tru64 UNIX UFS partitions and AdvFS storage domains to LSM volumes.
 - Tru64 UNIX Version 4.* disk groups to Version 5.* disk groups (**NOTE: this will migrate Version 4.* Block Change Logging (BCL) to Version 5.* Dirty Region Logging (DRL).**)
-

Optional Software

Advanced File System Utilities

The Advanced File System Utilities (AdvFS Utilities) extend the high availability and flexibility of the Advanced File System. The Advanced File System Utilities provide a GUI to simplify management tasks and utilities to dynamically resize file systems, defragment files, balance the percentage of space used on volumes, and clone filesets. The Advanced File System Utilities is a separately licensed software product for Tru64 UNIX. Reference the QuickSpecs at <http://www.compaq.com/products/quickspecs/Division/11024.html>.

HP TruCluster Server

HP TruCluster Server provides highly available and scalable solutions for users in mission-critical environments. It delivers easier, more sophisticated UNIX clustering capabilities by adding a fully clustered shared file system to the rich functionality already found in the HP cluster products. Because TruCluster Server appears as a single system, it requires less effort to manage an entire cluster than separate systems, significantly reducing your management costs.

By combining the advantages of symmetric multiprocessing, distributed computing, and fault resilience, a cluster running TruCluster Server offers its users high availability while providing scalability beyond the limits of a single system. TruCluster Server also significantly reduces, but does not eliminate, the impact of hardware and software failures.

Refer to the TruCluster Server QuickSpecs at http://www.compaq.com/products/quickspecs/North_America/10657.html.

Performance Manager

Performance Manager for Tru64 UNIX is a management application that centralizes the monitoring of UNIX system performance information for a network of systems. A single instance of Performance Manager for Tru64 UNIX can monitor resources for many systems and alert the administrator when problems arise. Performance Manager is a separately licensed software product for Tru64 UNIX.

Standard Features

Software Licensing

Basic LSM functionality, including disk spanning and concatenation, is licensed with the Tru64 UNIX operating system. Additional features, including disk striping (RAID 0), mirroring (RAID 1), RAID 5, and the graphical user interfaces, are available under a separate license. For more information about HP's licensing terms and policies, contact your local HP office or your HP authorized reseller.

Licensing Management Facility Support

This software product supports the Tru64 UNIX License Management Facility. License units for this product are allocated on an Unlimited System Use basis. For more information on the License Management Facility, see the Tru64 UNIX Operating System QuickSpec.

Growth Considerations

The minimum hardware and software requirements for any future version of this product may be different from the requirements for the current version.

Distribution Media

This product is distributed with the Tru64 UNIX Operating System CD-ROM. Tru64 UNIX comes with a basic LSM license which allows users to create data that spans across physical disks (concatenation). A software license is required to use the software graphical user interfaces and the RAID capabilities; data mirroring (RAID 1), data striping (RAID 0), and data striping with parity (RAID 5).

Software Warranty

This software is provided by HP with a 90-day conformance warranty in accordance with the HP warranty terms applicable to the license purchase.

Ordering Information

LSM Software Licenses: QL-2GVA*-AA
Software Product Services: QT-2GVA*-**
Software Documentation: AA-Q3NC*-TE

Software Product Services

A variety of service options are available from HP. For more information, contact your local HP office or your HP authorized reseller.

© Copyright 2005 Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice.

UNIX is a registered trademark of The Open Group.

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained.