



Tru64 UNIX and HP-UX File Systems and Volume Managers

Thomas Aussmann
Consultant Proactive Services
Hewlett-Packard GmbH
thomas.aussmann@hp.com


© 2004 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice



Agenda

- Volume Manager Overview
- Logical Volume Manager (LVM)
- Logical Storage Manager (LSM)
- VxVM – LSM Feature Comparison
- File Systems
- Comparison File Systems and Volume Managers
- Resources


17.04.2007 www.it-symposium2007.de 2



Volume Manager Overview

- Comparison may fail
 - No LVM within Tru64 UNIX
 - Only abstraction Layer within Tru64 UNIX
 - Tru64 UNIX LSM similar to VERITAS VxVM
 - Based on VxVM V2.3 code base


17.04.2007 www.it-symposium2007.de 3



Volume Manager Overview

- Tru64 UNIX
 - abstraction layer between physical disk and file system
 - Logical volume is implemented through
 - Hardware (RAID)
 - Device driver (partition, LSM)
 - File system (AdvFS)
 - no management for logical volumes needed
 - No LVM commands available
 - Must configure RAID, AdvFS or partitions, though


17.04.2007 www.it-symposium2007.de 4



Volume Manager Overview

- HP-UX
 - More than abstraction layer
 - Used for managing disk partitions
 - Limitations:
 - No striped mirrors
 - No hotspare
 - No RAID5
 - „whole disk approach“ (no LVM)
 - File system cannot span multiple disks
 - Only one file system partition
 - Difficult to extend


17.04.2007 www.it-symposium2007.de 5



Logical Volume Manager (LVM)

- LVM naming
 - Physical volume
 - Volume group
 - Logical volume
 - Associated to volume group
 - LVM extent
 - Smallest allocatable unit of space
 - Associated to logical volume


17.04.2007 www.it-symposium2007.de 6



Logical Volume Manager (LVM)

- LVM commands
 - „pvcreate“, initialize a physical volume
 - Write physical volume reserved area
 - Write volume group reserved area
 - „vgcreate“, create volume group
 - Physical volume is broken into physical extents (PE)
 - Groups logical volumes
 - „lvcreate“, create logical volumes
 - Creates both block and raw LV
 - Consists of sequentially numbered logical extents (LE)
 - LE points to PE on physical volume
 - Contains file system

17.04.2007 www.it-symposium2007.de 7



Logical Volume Manager (LVM)

- Look at LVM configurations with
 - pvdisplay, vgdisplay, lvdisplay
- Delete LVM configurations with
 - pvremove, vgremove, lvremove
- Extend LVM configurations with
 - vgextend (add physical volumes to group)
 - lvextend (increase space or mirror to volume)
- Reduce LVM configurations with
 - vgreduce (remove physical volumes from group)
 - lvreduce (reduce space or number of mirror copies)

17.04.2007 www.it-symposium2007.de 8

```

wallace:~# lvs -v /dev/vg00/lvol2
--- Logical volumes ---
LV Name                /dev/vg00/lvol2
LV VG                  /dev/vg00
LV Permission          read/write
LV Status              available/syncd
Mirror copies         1
---snip---snip---snip---
--- Distribution of logical volume ---
/dsk PV Name      LE on PV  PE on PV
/dsk /dev/dsk/c1t15d0 750      750
/dsk /dev/dsk/c3t15d0 750      750
---
LE --- Logical extents ---
LE PV1 PE1 Status 1 PV2 PE2 Status 2
000 /dev/dsk/c1t15d0 00075 current /dev/dsk/c3t15d0 00075 current
001 /dev/dsk/c1t15d0 00076 current /dev/dsk/c3t15d0 00076 current
---snip---snip---snip---
007 /dev/dsk/c1t15d0 00823 current /dev/dsk/c3t15d0 00823 current
008 /dev/dsk/c1t15d0 00824 current /dev/dsk/c3t15d0 00824 current
wallace:~#
wallace:~#

```

17.04.2007

www.it-symposium2007.de

9


Logical Volume Manager (LVM)

- Additional commands
 - vgchange (activate/deactivate, control membership)
 - Related to ServiceGuard
 - vgcfgbackup (create or update group configuration)
 - vgcfgrestore (display or restore group configuration)
 - Displays only vgcfgbackup info, no life configuration
 - lvsplit, lvmerge (split or merge LVM mirror)
 - vgsync, lvsync (LVM mirror synchronization)
 - Even more commands for:
 - Boot disk preparation, export/import of volume groups, ...
 - HP-UX LVM Reference
 - <http://docs.hp.com/hpux/online/docs/B2355-60103/00/42/4255-con.html>

17.04.2007

www.it-symposium2007.de


10



Logical Volume Manager (LVM)

- Things to take care off
 - lvmextend adds space, does not notify file system
 - FS superblock and metadata structures need notification
 - umount, extendfs and mount file system
 - Or use fsadm utility for online notification/expansion (license)
 - lvmreduce may corrupt data
 - Backup, umount, newfs, mount and restore FS
 - fsadm utility works sometimes
 - Checks for used blocks at end of FS
 - Defragment may help
 - Reducing FS does not reduce logical volume


17.04.2007 www.it-symposium2007.de 11



Logical Volume Manager (LVM)

- LVM and VERITAS VxVM can co-exist
 - On same system but not on same disk
 - LVM and VxVM configuration data located on disks
- VERITAS VxVM is mandatory for VERITAS CFS
- No need to use LVM for root anymore
 - VERITAS VxVM has rootability
- LVM will continue to be enhanced
 - LVM will remain HP-UX 11i default volume manager


17.04.2007 www.it-symposium2007.de 12



Logical Volume Manager (LVM)

- LVM Example
 - Initialize disk
 - `pvcreate /dev/rdisk/c0t0d0`
 - Create pseudo device for LVM subsystem
 - `mkdir /dev/vg01`
 - `mknod /dev/vg01/group c 64 0x030000`
 - Check owner/permission (root/640)


17.04.2007 www.it-symposium2007.de 13



Logical Volume Manager (LVM)

- LVM Example cont.
 - Create volume group
 - `vgcreate -s <pe_size> /dev/vg01 <pv_path> <pv_path> ...`
 - Create logical volume
 - `lvcreate -L <size> -n <name> /dev/vg01`
 - Create and mount file system
 - `newfs <special>`
 - `mount ...`

17.04.2007 www.it-symposium2007.de 14




Logical Volume Manager (LVM)

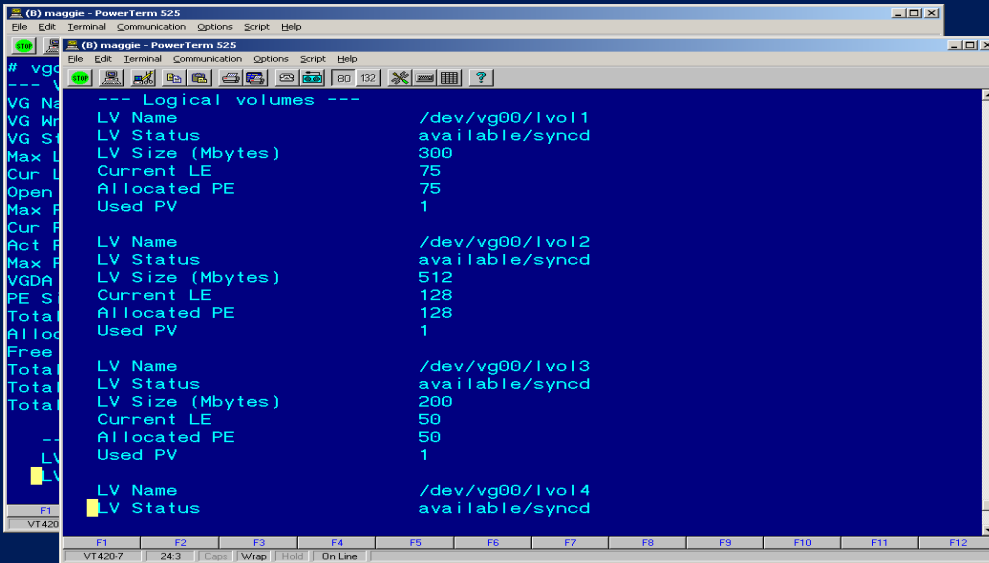
- Default system disk layout

/dev/vg00/lvol3	/	vxfs
/dev/vg00/lvol1	/stand	hfs
/dev/vg00/lvol4	/tmp	vxfs
/dev/vg00/lvol5	/home	vxfs
/dev/vg00/lvol6	/opt	vxfs
/dev/vg00/lvol7	/usr	vxfs
/dev/vg00/lvol8	/var	vxfs

17.04.2007
www.it-symposium2007.de
15



Logical Volume Manager (LVM)



```

--- Logical volumes ---
LV Name          /dev/vg00/lvol1
LV Status        available/syncd
LV Size (Mbytes) 300
Current LE       75
Allocated PE     75
Used PV          1
LV Name          /dev/vg00/lvol2
LV Status        available/syncd
LV Size (Mbytes) 512
Current LE       128
Allocated PE     128
Used PV          1
LV Name          /dev/vg00/lvol3
LV Status        available/syncd
LV Size (Mbytes) 200
Current LE       50
Allocated PE     50
Used PV          1
LV Name          /dev/vg00/lvol4
LV Status        available/syncd
    
```


17.04.2007
www.it-symposium2007.de
16



Logical Storage Manager (LSM)

- Tru64 UNIX LSM based on V2.3 VxVM code
 - VxVM commands and outputs almost equal to LSM
 - Just replace „vol“ with „vx“
- Various shell-level LSM commands available
 - volassist, volclonedg, vold, voldctl, voldg, voldisk, voldiskadd, voldiskadm, voldisksetup, voledit, volencap, volevac, volinfo, volinstall, voliod, vollogcnvt, volmake, volmend, volmigrate, volmirror, volnotify, volplex, volprint, volreattach, volreconfig, volrecover, volrestore, volrootmir, volsave, volsd, volsetup, volstat, voltrace, volume, volunmigrate, volunroot, volwatch
- GUI based commands „lsmsa“, „dxlsm“

17.04.2007 www.it-symposium2007.de 17



Logical Storage Manager (LSM)

- LSM naming
 - Subdisk
 - Logical representation of contiguous disk blocks
 - Compare with LVM PE's
 - Plex
 - Made up of one or more subdisks
 - Instance of volume data
 - Volume
 - Contains at least one plex; two or more for mirror volume
 - Virtual disk device
 - Contains file system
 - Compare with LVM LV's


17.04.2007 www.it-symposium2007.de 18



Logical Storage Manager (LSM)

- LSM naming cont.
 - Diskgroup
 - Collection of disks belonging to named group
 - Compare with LVM volume group
 - Disk
 - physical disk
 - Equal to LVM physical volume
 - Private region
 - Contains LSM configuration data, compare with VGRA/PVRA
 - Public region
 - Contains either free space or subdisks

17.04.2007 www.it-symposium2007.de 19




LVM – LSM Comparison

LSM Term	LVM Term
•Physical Disk	•Physical volume
•Subdisk	•Physical extent
•Volume	•Logical volume
•Disk group	•Volume group
•Private region	•PVRA/BDRA/VGRA
•Free space	•Unused physical extent
•Plexes	•Mirrors
•Dirty Region Logging (DRL)	•Mirror Write Cache (MWC)
•Dynamic multipathing*	•PVlinks

* VxVM term, no need for that within Tru64 UNIX V5.*


17.04.2007 www.it-symposium2007.de 20



VxVM - LSM Feature Comparison

	HP-UX 11i v2 VxVM 3.5	Tru64 5.1B LSM (VxVM 2.3 Code Base)
Layered Volumes	Yes	No
RAID 5	Yes	Yes
VMSA GUI, VEA GUI	Yes	lmsa
Online Relayout (vxrelayout)	Yes	No
Dirty Region Logging	Yes	Yes
Disassociating a plex for backup (volpex displex)	Yes	Yes
Snapshot volumes (volassist snapshot ...)	Yes	Yes
Non-Persistent Fast Resync (Fast Mirror Resynchronization (FMR))	Yes	No
Persistent Fast Resync (Data Change Object (DCO))	Yes	Wildcat 5.1B (Smash and Resilvering Log (SRL))
SmartSync Recovery Accelerator for Databases	Yes	No
Hot Spares / Relocation	Yes	Yes
Dynamic MultiPathing (DMP)	Yes	NA - base OS handles in SCSI/CAM


17.04.2007 www.it-symposium2007.de 21



VxVM - LSM Feature Comparison

Power Fail Timeout vxpft	Yes	NA - handled in SCSI/CAM
Task Management vxtask	Yes	No
Autoconfig	Yes	Yes
Config load balancing	Yes	Yes
Dynamic IO Sizes	Yes	Yes
Configuration Saving	Yes (dgcfgbackup, dgcfgrestore)	Yes (volsave, volrestore)
Limits: Volume Size	2TB	1TB
Limits: Number of Volumes	None	8189
Limits: Plexes per Volume	32	32
Limits: Plexes	None	8189
Limits: Number of Disks	None	None
Limits: Number of diskgroups	None	None


17.04.2007 www.it-symposium2007.de 22



File Systems

- Tru64 UNIX supported file systems
 - UFS, AdvFS, NFS, cdfs, dvdfs,
 - mfs, procfs
 - dfs, efs (DCE filesystems)
 - fdfs, ffm (used by streams)
 - pcfs, sysv
- HP-UX supported file systems
 - HFS (UFS), JFS (VxFS), cdfs, NFS


17.04.2007 www.it-symposium2007.de 23



File Systems

- Large File Support
 - Support for files larger 2^{32}
 - HP-UX 11i V2/VxFS 2TB
 - HP-UX 11i V3/VxFS 16TB
 - Tru64 UNIX AdvFS 16TB
 - AdvFS does not require special flags or options
 - VxFS requires option „largefiles“ for files larger 2 GB
 - „mkfs -o largefiles“, „mount -o largefiles“
 - „fsadm -o largefiles“ converts nolargefiles FS


17.04.2007 www.it-symposium2007.de 24



File Systems

- Direct I/O
 - HP-UX supports async I/O for raw devices only
 - Sybase/Oracle DB on raw volumes
 - OnlineJFS comes with direct I/O
 - bypass buffer cache on normal file systems
 - controlled through extended mount options
 - no kernel level Oracle asynchronous I/O
 - Comes with Storage Foundation 4.1
 - No limitations on Tru64 UNIX
 - Either „single server“ or direct I/O within cluster
 - Based on drdmgr, default direct I/O
 - Tapes always are „single server“


17.04.2007 www.it-symposium2007.de 25



File Systems

- Tru64 UNIX
 - different commands for UFS and AdvFS
 - newfs, fsck, dump, rdump, restore ... (UFS)
 - mkfdmn, mkfset, chfset, chfile, rmdmn, rmfset, showfdmn, ... , vdump, ... , addvol, rmvol, defragment, ... (AdvFS)
 - see „man advfs“ for all related commands
- HP-UX
 - single command for specific actions
 - mkfs, newfs, fsck, fsadm (licensed)
 - use option „fstyp“ for file system type
 - default file system defined at /etc/default/fs


17.04.2007 www.it-symposium2007.de 26



File Systems

	Tru64 UNIX® V5.1B (AdvFS)	HP-UX 11iv2 JFS (VxFS)
storage model	multi-volume	single volume (V3) multi-volume (V4)
journals	meta-data optional – user file data, Atomic Data Logging	meta-data
allocation abstraction	extents	extents
Recovery	automatic on mount	external tool fsck, run in bcheckrc
on-line resize	addvol, rmvol, or mount -o expand	(volume mgr cmds), fsadm, VEA
read-only file system copies	clones (clonefset, mount)	Storage chkpts (fsckptadm, mount)
on-line defragmentation	defragment, vfast	fsadm


17.04.2007 www.it-symposium2007.de 27



File Systems

	AdvFS	VxFS
create a file system	# mkfdmn vol dom # mkfset dom fset # mount dom#fset dir	# mkfs vol # mount vol dir
increase file system size	# addvol vol dom	<i>increase volume size</i> # fsadm
list the storage of a file system	# showfdmn dom	# df
list mounted file systems	# mount	# mount
determine unmounted file systems	# more /etc/fstab # ls -R /etc/fdmns	# more /etc/fstab # fsck /...vol...


17.04.2007 www.it-symposium2007.de 28



File Systems

	AdvFS Tru64 UNIX	VxFS HP-UX
multi-volume model in file system	yes "domain" abstraction representing a pool of volumes to be used for a file system (mkfdmn, addvol, rmvol)	yes (V4) for VxFS 3.5 every file system is associated with a single volume (mkfs) with VxVM V4 can use vsets for multi-volume
multiple mountable rooted trees per file system	yes AdvFS "fileset" abstraction representing an individual mountable tree within a domain (mkfset)	no every file system has one mountable tree with exception of the special case of a snapshot


17.04.2007
www.it-symposium2007.de
29



Comparison File Systems and Volume Managers on Tru64 and HP-UX

	Tru64	HP-UX
Base File System	AdvFS Base	JFS
Full Featured File System	AdvFS Utilities	OnlineJFS
Cluster File System	TruCluster CFS	VERITAS CFS
Base Volume Manager	LSM Base	VxVM Base, LVM
Full Featured Volume Manager	LSM Advanced	VxVM, MirrorDisk/UX
Enhanced Volume Manager		CVM, SLVM

17.04.2007
www.it-symposium2007.de
30



Resources

- HP-UX System and Network Administration
- HP-UX online reference
 - <http://docs.hp.com/en/B2355-60103/>
- HP-UX LVM online reference
 - <http://docs.hp.com/hpux/onlinedocs/B2355-60103/00/42/4255-con.html>
- Tru64 UNIX online reference
 - http://h30097.www3.hp.com/docs/pub_page/V51B_DOCS/V51B_DOCLIST.HTM
- Tru64 UNIX Logical Storage Manager
 - http://h30097.www3.hp.com/docs/base_doc/DOCUMENTATION/V51B_HTML/A_RH9BDTE/TITLE.HTM

17.04.2007 www.it-symposium2007.de 31

