

SGI LTFS

Olga Weiss

MTS, Storage Software

Al Stipek

Tech Lead, Storage Software

© 2013 SGI



Overview

- What is LTFS?
- Why LTFS?
- How LTFS Works
- SGI LTFS Features
- Future Directions

What is LTFS?

LTFS = Linear Tape File System

Refers to:

Tape format standard developed by IBM in 2010

<https://en.wikipedia.org/wiki/LTFS>

<http://www-03.ibm.com/systems/storage/tape/lufs/>

File system interface to data stored on magnetic tape

Open Source implementations

IBM wrote and open-sourced

HP, and Quantum have open source forks of that code

All the vendors also have \$commercial versions

Drive Compatibility

LTFS takes advantage of a feature in newer tape drives which support two partitions on a tape
In effect, like having two separate tapes on one cartridge

Works with

LTO5 and 6 drives

Oracle T10000C and T10000D drives

LTFS Editions

LTFS can come in two “editions”:

Single Drive Edition

It's the user's job to get a tape into a drive and then LTFS-mount that tape

The LTFS file system consists of that single tape

Typically vendors release this as Open Source

Library Edition

An entire tape library is mounted as LTFS filesystem

The 1st level subdirectories correspond to the VSNs of each tape of the library

Typically this assumes the whole library belongs to LTFS

Why LTFS?

Use tape as drive

See files and directories as a tree – fast convenient access to data

Search / open / copy files as they are on a drive

...but don't forget that this is a tape

And tape doesn't allow random access

Transfer large files between systems

Vendor-neutral data retrieval from tape, independent of any hardware or software platform

How LTFS Works

LTFS partitions a tape into **two segments**

The smaller first partition is used to store index information and metadata about the contents of the second part

Metadata is the descriptive information about user data stored on the tape

This information enables searching and accessing of files in the other, data partition (using UI)

The second part contains data

SGI LTFS Features

SGI LTFS is:

Single Drive Edition (handles one tape at a time with manual mount)

But more than one LTFS tape can be mounted at the same time

Linux support only

The CLI script works on the DMF server only

In general intended to work with **DMF**, i.e. share drives

...but can share LTFS drives with **any specified application**

Supports v2.0.0 of the LTFS standard

Available in ISSP 3.2

OpenVault

Tapes are mounted using the **OpenVault** mounting service

No dedicated drives or library partition is required

The LTFS cartridges would be in the same library with DMF (or another application) cartridges, just in a separate OpenVault cartridge group

Similarly, the LTFS drives would be in a separate OV drive group

Drives can be moved to and from the LTFS drive group, so no dedicated drives are required

This will work with all tape libraries SGI supports

Command Line Interface

The `sgi-ltfs` RPM provides a program **ov_ltfs** with commands:

show information about LTFS tapes/drives/config setting

import new tapes into the OpenVault system, LTFS drive group

format/mount a tape; can do that on a specified drive; otherwise the script will try to find a drive for the operation – see next slide

unmount LTFS on the specified tape

reserve/give back the borrowed drive to DMF

eject a tape to be physically taken from the library

purge the OpenVault record for the tape if it will never be seen again

Borrow A Drive

In a busy DMF environment, a drive can be "borrowed" for LTFS use and returned to DMF immediately after

If a **mount** or **format** operation is issued and no drives are available in the LTFS group, the CLI will find a compatible drive for the request

Can search for a drive in "allowed drive group" list specified

Options --temp and/or --wait can be specified

-- temp return the drive to DMF immediately after this operation

-- wait wait for the drive (can be specified one) if it is busy

Mounting

The **ov_ltfs mount** operation does two things

1. Arrange to physically mount the LTFS tape in a drive
2. Logically mount the tape as a FUSE filesystem

The default mountpoint is */dmf/ltfs/\$vsn*

After mounting, you ssh to the DMF server and use standard Linux commands to move data to and from the tape

cp, tar, etc

DMF Manager

We added LTFS support in dmfman.

Reserve a drive

Check status

Mount a tape

Unmount a tape

[Stop OpenVault](#)

OpenVault libraries

C11

C11d00

C11d01

C11d02

C11d03

C11d04

C11d05

C11d06

lto6_lib

LTO6_2

LTO6_1

OpenVault libraries:

Library C11:

<input type="checkbox"/>	Drive Name	Drive Group	Status	Mounted Volume
	No filter applied			
<input type="checkbox"/>	C11d00	dg_c11	Enabled	No volume mounted
<input type="checkbox"/>	C11d01	dg_c11	Enabled	No volume mounted
<input type="checkbox"/>	C11d02	dg_c11	Enabled	No volume mounted
<input type="checkbox"/>	C11d03	dg_c11	Enabled	No volume mounted
<input type="checkbox"/>	C11d04	dg_c11	Enabled	No volume mounted
<input type="checkbox"/>	C11d05	dg_c11	Enabled	No volume mounted
<input type="checkbox"/>	C11d06	dg_c11	Enabled	No volume mounted

Library lto6_lib:

<input type="checkbox"/>	Drive Name	Drive Group	Status	Mounted Volume
	No filter applied			
<input type="checkbox"/>	LTO6_1	ltfs_dg	Enabled	No volume mounted
<input type="checkbox"/>	LTO6_2	ltfs_dg	Enabled	No volume mounted

- Enable drives
- Disable drives
- Reserve Drive for LTFS...
- Release LTFS Drive...

These two drives have
already been reserved
for LTFS use

DMF Manager - Mozilla Firefox

Configuration Storage Messages Statistics Backups Help timeToExpire resetExpire Change Password Log Out

Overview Libraries

Library Servers

- LS500
 - vg_lto2
 - maid_ls
 - vg_c11
- Hold Flag Repo
 - hoa
 - hfull
 - hextern
- Unassigned Tapes
 - 000029
 - DMF604
 - DMF605
 - DMF606
 - DMF607

LTFS Information

Drives

Drive Name	Library Name	Status	VSN
LTO6_2	Ito6_lib	unloaded	
LTO6_1	Ito6_lib	unloaded	

Tapes

VSN	Library Name	Mounted	Drive Name	Mnt Point

Unassigned Tapes

Cartridge PCL	Type	Library Name
000029	Ultrium6-2500	Ito6_lib
DMF604	Ultrium6-2500	Ito6_lib
DMF605	Ultrium6-2500	Ito6_lib
DMF606	Ultrium6-2500	Ito6_lib
DMF607	Ultrium6-2500	Ito6_lib

LTFS Configuration

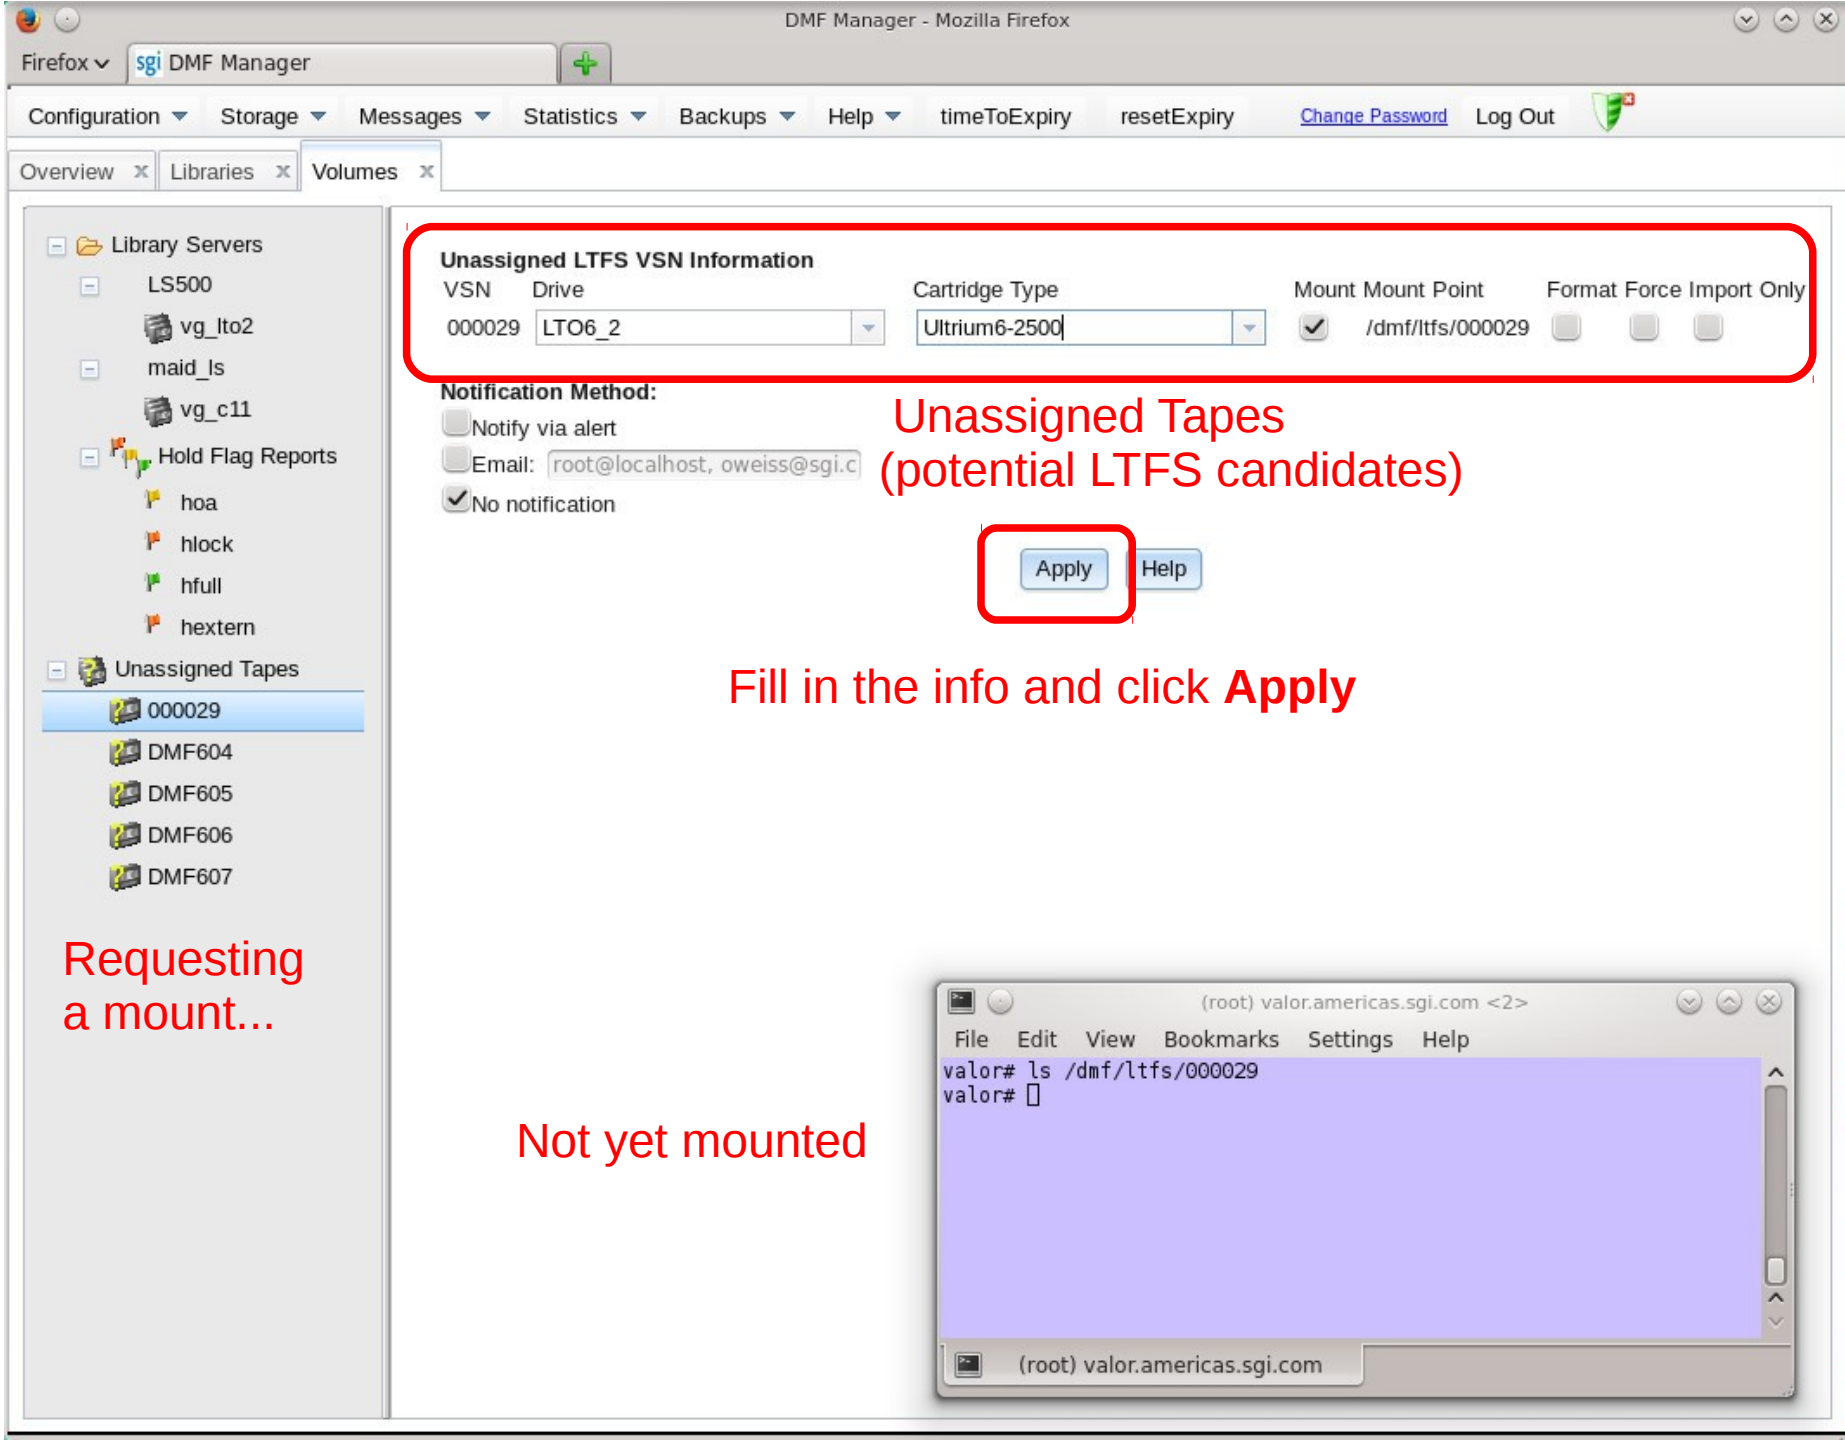
Parameter	Value
Config file path:	/usr/local/etc/ltfs.conf.local
ov_server	valor
ltfs_drive_group	ltfs_dg
ltfs_cart_group	ltfs_carts
mnt_root	/dmf/ltfs
max_drives	2
ltfs_keyfile	/dmf/home/ltfs_keys
Allowed Drive Groups	Max To Reserve

Help

Menu:

- Add Volumes...
- Display Dump Volumes...
- Show LTFS Information...
- Set Volumes Preferences...
- Refresh Volumes
- Show Key
- What is 'Volumes'?
- Close

The current state of LTFS can be displayed via a menu option off the Volumes panel.



- Library Servers
 - LS500
 - vg_lto2
 - maid_ls
 - vg_c11
 - Hold Flag Reports
 - hoa
 - hlock
 - hfull
 - hextern
 - Unassigned Tapes
 - 000029
 - DMF604
 - DMF605
 - DMF606
 - DMF607

Unassigned LTFS VSN Information

VSN	Drive	Cartridge Type	Mount	Mount Point	Format	Force	Import Only
000029	LTO6_2	Ultrium6-2500	<input checked="" type="checkbox"/>	/dmf/ltfs/000029	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Notification Method:

Notify via alert

Email:

No notification

Unassigned Tapes
(potential LTFS candidates)

Fill in the info and click **Apply**

Requesting
a mount...

Not yet mounted

```
(root) valor.americas.sgi.com <2>
File Edit View Bookmarks Settings Help
valor# ls /dmf/ltfs/000029
valor#
```


- Library Servers
 - LS500
 - vg_lto2
 - maid_ls
 - vg_c11
 - Hold Flag Reports
 - hoa
 - hlock
 - hfull
 - hextern
- LTFS Assigned Tapes
 - 000029
- Unassigned Tapes
 - DMF604
 - DMF605
 - DMF606
 - DMF607

LTFS Information

Drives

Drive Name	Library Name	Status	VSN
LTO6_2	lto6_lib	loaded	000029
LTO6_1	lto6_lib	unloaded	

Tapes

VSN	Library Name	Mounted	Drive Name	Mnt Point
000029	000029	yes	LTO6_2	/dmf/lufs/000029

Unassigned Tapes

Cartridge PCL	Type	Library Name
DMF604	Ultrium6-2500	lto6_lib
DMF605	Ultrium6-2500	lto6_lib
DMF606	Ultrium6-2500	lto6_lib
DMF607	Ultrium6-2500	lto6_lib

LTFS Configuration

Parameter	Value
Config file path:	/usr/local/etc/lufs.conf.local
ov_server	
lufs_drive_group	
lufs_cart_group	
mnt_root	
max_drives	
lufs_keyfile	
Allowed Drive Groups	

...and mounted!

```
(root) valor.americas.sgi.com <2>
File Edit View Bookmarks Settings Help
valor# ls /dmf/lufs/000029
valor# ls /dmf/lufs/000029
README
valor# cat /dmf/lufs/000029/README
Hello from 000029.
valor#
```

DMF Manager - Mozilla Firefox

Firefox v Sgi DMF Manager

Configuration Storage Messages Statistics Backups Help timeToExpiry resetExpiry [Change Password](#) Log Out

Overview Libraries Volumes

Library Servers

- LS500
 - vg_lto2
 - maid_ls
 - vg_c11
- Hold Flag Reports
 - hoa
 - hlock
 - hfull
 - hextern
- LTFS Assigned Tapes
 - 000029
- Unassigned Tapes
 - DMF604
 - DMF605
 - DMF606
 - DMF607

Mounted LTFS VSN Information

VSN	Drive	Unmount	Mount Point	Eject	Remember
000029	LTO6_2	<input checked="" type="checkbox"/>	/dmf/lufs/000029	<input type="checkbox"/>	<input type="checkbox"/>

Notification Method:

- Notify via alert
- Email:
- No notification

Now let's unmount.

Future Directions

Planned features for SGI LTFS version 2

export operation: copy files to LTFS tape directly from the DMF storage skipping "bring files online" step

Request the mount of an LTFS tape from an NFS client

Support all types of tape/drive that are LTFS-compatible / supported by SGI

Automatically detect cartridge type

Maybe Later??

Interop testing

Sparse file / xattr support

sgi