



DMF Extends Active Archiving to Lustre

*Direct Archiving Feature Leverages
the Power of DMF for Any
POSIX-Compliant Filesystem*

For 20 years SGI® DMF (Data Migration Facility) has been the standard for advanced tier virtualization, being used by hundreds of clients to manage some of the world's largest active archives. By presenting to the user space a unified "online" view of data that is distributed across multiple tiers of disk and tape, DMF provides users with automated access to all data all the time. In this way, very large data stores can be maintained with much less cost by keeping high-performance production disk at much smaller levels. Users only know that the data is online and available, without needing to manage manual copying to offline silos.

With the introduction of DMF 5.2 and a new feature called DMF Direct Archiving, the power of DMF is now expanded to include any POSIX-compliant filesystem such as Lustre.

The DMF Direct Archive method is more efficient because it requires less time, bandwidth, and managed filesystem capacity. This provides a secure, easily scalable active archive solution for high speed scratch filesystems like Lustre. Thus, instead of investing in more expensive high speed storage for scratch space, DMF Direct Archiving enables IT managers to quickly, easily, and cheaply add capacity to their storage infrastructure while maintaining full interoperability with Lustre or other POSIX filesystems.

Introducing DMF Direct Archiving

Available with DMF 5.2, SGI is introducing a new feature called DMF Direct Archiving that allows users to implement a DMF active archive solution for any POSIX-compliant filesystem. With DMF Direct Archiving, users are now able to manually archive files directly to secondary or tertiary storage via DMF from any POSIX-compliant filesystem, thus allowing far greater scope for leveraging the power of DMF. By opening the tiered archive environment to POSIX filesystems, DMF can be used as an archive environment for high speed scratch filesystems leveraging the cost benefits of DMF while easing the management requirements associated with traditional archive and backup.

In a traditional DMF implementation, all data in the tiered storage pools is managed within XFS or CXFS filesystems. In this way, users with high performance disk requirements are able to keep such infrastructure small even while online access is available to massive volumes of data distributed throughout the virtualized tiers. Data is automatically migrated off of high speed, expensive storage to less expensive storage including tape based upon a variety of user-driven policies, such as last access, type of file, file tags or custom policies written to DMF's API that can be assigned via the DMF policy engine. By presenting all the files as online and available to users, but virtualizing the storage tiers such that the bulk of the data can live in low-cost MAID or tape systems, IT administrators are able to provide the quality of service that their users demand for large volumes of data, while also keeping the costs dramatically lower to satisfy budget constraints.

With the addition of DMF Direct Archiving, any POSIX compliant filesystem can be identified as an *unmanaged filesystem* to DMF and can benefit from the same virtualized tiering. An example of an unmanaged filesystem would be a Lustre file system. With DMF Direct Archiving files in the unmanaged filesystem can be easily and efficiently copied into a DMF managed archive environment, which extends the cost and management benefits of DMF to any POSIX compliant filesystem like Lustre.

This new feature includes the new *dmarchive* command and associated library subroutines within DMF. When the *dmarchive* command is issued to a file, DMF copies the file data to either Tier 2 or Tier 3 Storage in the DMF archive environment while placing the metadata in a visible DMF-managed filesystem, as shown in Figure 1.

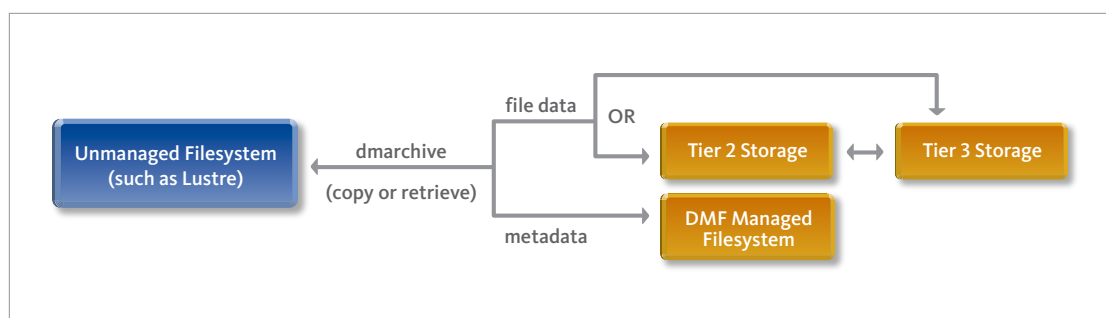


Figure 1: Archiving Files from an Unmanaged Filesystem to Secondary Storage

By streamlining this process into a single more efficient step, DMF Direct Archiving saves users from a two-step manual procedure to move files into a DMF-managed environment before it can be migrated to low-cost tier 2 or tier 3 storage. Without DMF Direct Archiving, a user would have to first copy the file to a DMF managed filesystem (such as /dmf) and then migrate the files. For example:

```
% cp -a /lustrefs/work /dmf
% dmput /dmf/work/*
```

However, using DMF Direct Archiving, you can achieve the same results with a single command:

```
% dmarchive -a /lustrefs/work/dmf /dmf/work/*
```

In this way, the file data will be copied directly to less expensive DMF-managed secondary or tertiary storage like a SGI® COPAN™ MAID array or tape library, and only the file metadata will be copied to the DMF-managed filesystem (/dmf). This feature enables the DMF-managed filesystem to be much smaller, cost-optimized storage that is always online.

During the archive process, files from the unmanaged filesystem can be in one of two DMF file states: 1) *archiving files*, which are files where the original resides on an unmanaged filesystem and whose offline copies are in the process of being copied to the DMF archive, and 2) *offline files*, which is what the files will be when the process completes. Once files are offline and reside on the DMF archive, then free space on the unmanaged filesystem can be reclaimed. When the files need to be retrieved again into primary storage, the data is copied directly from DMF-managed lower tiers to the unmanaged filesystem.

Figure 2 below shows an example of a DMF Direct Archive environment for a Lustre filesystem. The Lustre server is serving the /lustrefs/work filesystem, which is mounted on both the DMF server and the DMF client where user applications are running. This allows users to run the *dmarchive* command from the DMF client when they want to archive their data from the Lustre filesystem. The DMF server is managing the /dmf filesystem, which is NFS-mounted at /mnt/dmfusr1 on the DMF client. User's archived files appear always available and online in the NFS-mounted DMF-managed filesystem.

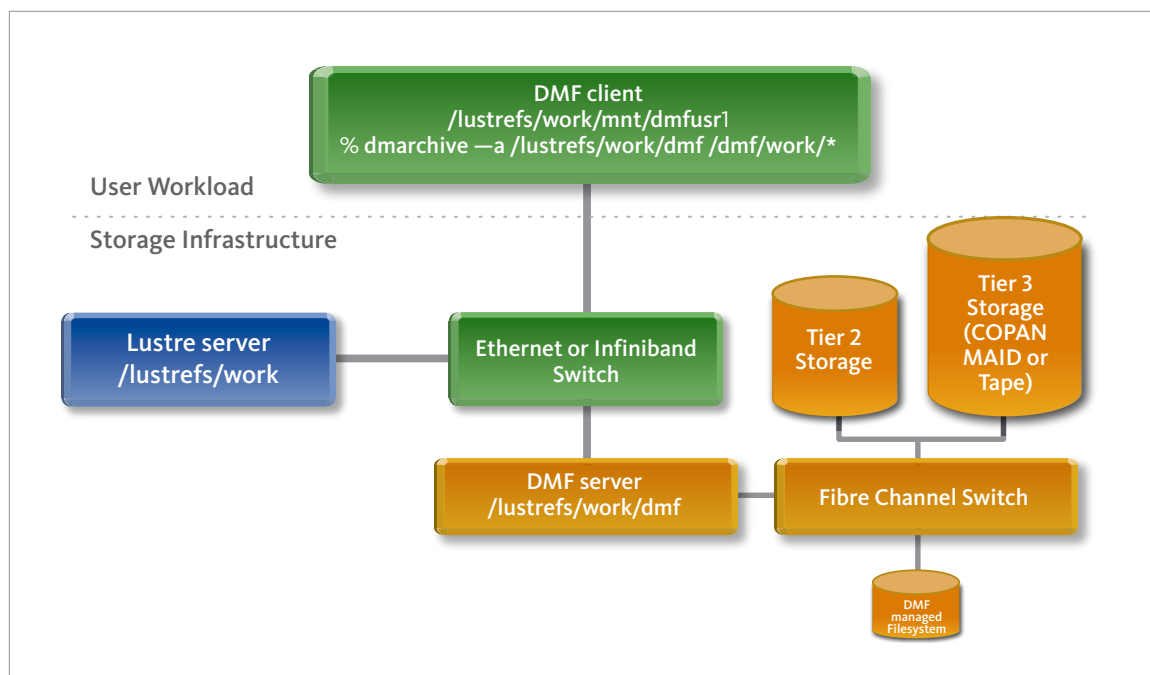


Figure 2: Lustre DMF Direct Archiving Environment Example

The DMF Direct Archive method is therefore more efficient because it requires less time, bandwidth, and managed filesystem capacity. This provides a secure, easily scalable archive solution for high-speed scratch filesystems like Lustre. Instead of investing in more expensive high-speed storage for your scratch space, leverage DMF with DMF Direct Archiving to quickly, easily, and cost-effectively add capacity to your storage infrastructure.

Corporate Office
 46600 Landing Parkway
 Fremont, CA 94538
 tel 510.933.8300
 fax 408.321.0293
 www.sgi.com

North America +1 800.800.7441
 Latin America +55 11.5185.2860
 Europe +44 118.912.7500
 Asia Pacific +61 2.9448.1463

© 2010 SGI. SGI and COPAN are registered trademarks or trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries. All other trademarks are property of their respective holders. 14122010 4272