

Data Migration Facility Development Update

Thomas W. Lanzatella, Cray Research, Inc.

ABSTRACT: *This paper reviews plans and status for all versions of DMF which are currently released as well as for all versions which are in development. A Roadmap for future releases and feature content is provided. We also review the differences between DMF for UNICOS and DMF for IRIX.*

Introduction

Several major events have occurred since our last status report on DMF.

1) DMF 2.5 was successfully field tested and released on 30 January, 1997.

2) DMF 2.4.7 was released on 28 March, 1997. This version supports the Cray T3E.

3) DMF 2.6 for IRIX has been executing on an Indy workstation since 14 April demonstrating manual and automatic retrieval.

4) We have adopted a 4-digit naming convention for DMF release names. This change was implemented in order that DMF release procedures conform with those used by other asynchronous products. This change will commence with the next releases of DMF. Watch the Cray Research Service Bulletin for detailed information.

DMF 2.4

DMF 2.4 is nearing the end of its useful life as a product. The development focus for the next version of DMF 2.4 will be support for security in UNICOS/mk 1.6 which is expected to be available during July of this year. Orderability of DMF 2.4 will end on 1 August for all platforms except Cray T3E. DMF 2.4 will continue to be available until DMF 2.5 is provided for Cray T3E. We anticipate that DMF 2.5 will be available for Cray T3E during 4Q97.

DMF 2.5

DMF 2.5.4 was released on 30 January, 1997 after a successful field test at Boeing. We would like to extend our sincere thanks to Boeing for accepting the risk and adopting a

partnership role with Cray/SGI in the release of this product. DMF 2.5 is now in production use at Boeing and on the production file server within Cray at Eagan as well as at a few other customer sites. Stability of the product is very high. No serious problems have been reported. We are still chasing a problem which can result in a corrupted database. This problem has not resulted in any data loss. Recovery from this problem is accomplished by replaying journal files from the time of the last good backup image of the DMF databases. DMF 2.5.4 is the current released version of the product. We recommend that sites upgrade to DMF 2.5.4.2 which will be available during June, 1997. Watch the Cray Research Service Bulletin for precise information.

No additional feature are planned for DMF 2.5.

DMF 2.5 - Content Review

The primary purpose of DMF 2.5 product is to improve DMF database technology. The database package used within DMF 2.5 is the Raima Data Manager (RDM), produced by the Raima Corporation (see <http://www.raima.com>). RDM provides a database transaction commitment mechanism that protects the DMF databases from damage after an unscheduled interrupt by ensuring that only whole transactions are applied to the databases. RDM also provides a transaction journaling capability. Transaction journaling is a method of tracking the history of database transactions. In the event of a disk failure or unscheduled interrupt, the system administrator can replay the transaction journals through a special utility which ensures that each DMF operation is reflected by the DMF databases.

DMF performance is not affected by the integration of RDM. Pilot studies show that performance of database operations in DMF 2.5 is the same as or better than under DMF 2.4.

Copyright © Cray Research Inc. All rights reserved.

Installing DMF release 2.5 requires that you convert your DMF databases with the new `dmdbcv` utility. This is a one-way conversion process. Once you have moved to DMF release 2.5, you cannot convert back to an earlier version.

You should schedule a block of time during which you can leave DMF down for the conversion. The `dmdbcv` utility runs at a rate of about .5 Million database entries per hour on a CRAY T90 machine. Be sure to consider the sum of database entries in the daemon database and in each media-specific process (MSP) database when you calculate the expected DMF down time for conversion.

Prior to DMF release 2.5, the database package had a problem with inefficient reuse of index file space, which required the DMF administrator to compress and rebuild the DMF database on a regular basis. The new database package does not have this problem. Therefore, once you have converted to DMF release 2.5, you will no longer have to compress and rebuild your DMF databases. This improvement simplifies the Cray Data Migration Facility (DMF) Administrator's Guide, publication SG-2135 2.5.

Journal application is now uniform between the tape MSP databases and the DMF daemon database. A single utility is used to apply journals to all DMF databases. Journal application capabilities which used to be offered through `dmvoladm` and `dmcataadm` have been moved to the new `dmdbrecover` utility. This improvement simplifies the Cray Data Migration Facility (DMF) Administrator's Guide, publication SG-2135 2.5.

Support for the old tape MSP and for the station MSP ends with DMF release 2.5. When you convert your DMF databases to the new format, the conversion utility, `dmdbcv`, will detect old tape MSP databases and convert them to the new format for the advanced tape MSP. You do not have to convert any of your DMF tapes. The advanced tape MSP reads old tape MSP media. Sites still running the station MSP should convert to the FTP MSP.

In DMF release 2.4 and earlier, the DMF administrator had the ability to change the name of database files from the released path names. In DMF release 2.5, the DMF administrator can change only the name of the directory in which databases are stored. Database names cannot be changed.

Source availability for DMF is ended with DMF release 2.5. For those few remaining DMF customers who have DMF source licenses, a source package will continue to be available.


DMF release 2.5 uses FLEXlm licensing to control access to the software. A license key is required in order to run the product. You will receive a license key from the Cray Research Distribution Center with installation instructions. If you do not receive the license key, contact the Distribution Center.

The DMF 3.0 Program

Data Migration Facility is established as a reliable and efficient tool for storage management services in the Cray/UNICOS environment. Development objectives for storage management are driven by customer requirements. The evolution of the

typical Cray environment to one of multiple-heterogeneous machines connected by high-speed networks require changes in DMF architecture. Distribution of computational and storage resources require that DMF be broken into distributable components. In this way, the critical functions of DMF, i.e., native file system migration and safe storage of data, can be accommodated across a wide area with consideration for the economic properties of storage devices.

Support for heterogeneous environments is a long-standing requirement for DMF. The acquisition of Cray by SGI fuels this requirement by making certain the likelihood of mixed configurations spanning a broad range of performance characteristics. The objectives of the DMF enhancement program are to provide customers with the ability to migrate from the platform of their choice and to store data on the platform or their choice.

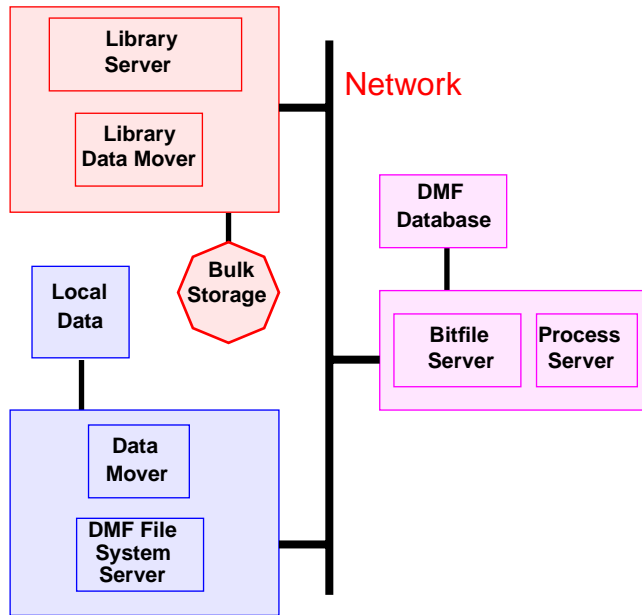


The Migration Domain for Cray/SGI

Management of native file system capacity continues to be a critical requirement for DMF. This process relies on UNIX kernel interfaces that support file state changes allowing an agent such as DMF to gain control of the data associated with a path in the native file system. To date, only UNICOS has had an efficient, reliable interface for this purpose. In the last couple of years the Data Management Interface Group (DMIG) has promoted an interface called the Data Management API or DMAPI which standardizes an interface supporting native file

system migration. The advent of DMAPI makes feasible the opportunity for DMF to manage file systems other than the UNICOS NC1 file system. In particular, since the IRIX file system, XFS, supports DMAPI, there is an immediate opportunity to manage XFS file systems with DMF.

The current architecture of DMF is monolithic in the sense that all DMF components must run on the same machine. The DMF 3.0 program will transform DMF into a storage management tool having the same external characteristics as DMF today but its components will be separable to facilitate remote execution and monitoring.



DMF 3.0 Architecture

During this process, we will adopt terminology from the IEEE Storage System Reference Model. The three major components of DMF 3.0 are the File System Server (replacing the DMF client), the DMF Bitfile Server (replacing the DMF server) and the Library Server (replacing the DMF media specific process).

This architecture is portable in the sense that all file system dependencies are localized in the File System Server while all media specific dependencies are localized in the Library Server. The Bitfile Server is considered to have no operating system dependencies.

To ensure that data can be moved from point to point with maximum efficiency, data movers will employ an authenticated, socket-based protocol. A data mover operates on behalf of each file system server and each library server. Multiple File System Servers and multiple Library Servers will be supported.

A single Bitfile Server will dispense bitfile identifiers (formerly called file handles) for the entire migration domain.

And finally, the three databases employed by DMF today will be replaced by a single database. We anticipate that DMF 3.0 will be available during 2H98.

DMF 2.6 for IRIX

The current status of DMF 2.6 is that the DMF daemon and the FTPMSP have been demonstrated with manual and automatic retrieval. Separately, tape-to-tape merging has been demonstrated with the advanced tape MSP. All of the daemon and MSP database utilities have been ported and demonstrated. Automatic space management utilities have also been ported and demonstrated.

We anticipate that DMF performance testing will commence in June, 1997 while field tests begin in early July. We are planning to release DMF 2.6 on 15 August.

The primary goal for this phase of the DMF 2.6 project is to offer a fully supported DMF product for the stand-alone IRIX environment. Our goal is defined to be a lateral port of DMF 2.5 into the IRIX 6.2/6.4 environment with as few changes as possible. Interoperability with Cray machines running UNICOS/DMF is a secondary objective as is DMF operation with the IRIX Fail Safe feature.

In the area of media support, SCSI versions of Redwood, Timberline and MagStar will be the first supported devices. In the area of robotics Grau, STK and IBM will be the first devices supported. IRIX supports many other manufacturers robots and transports. These will be supported as they are validated for use with DMF.

External differences between DMF 2.5 and DMF 2.6 have been kept to a minimum but there will be differences in areas affected by operating system dependencies. The following list details the major differences:

- 1) The kernel interface supporting file state transitions will be the DMIG/DMAPI interface.
- 2) The DMF Automatic Space Management utilities will be adapted to the sparse allocation capabilities of the XFS file system.
- 3) DMF 2.6 will employ a new SGI product named OpenVault for tape mounting services instead of the UNICOS Tape Subsystem. If OpenVault cannot be provided in time for a DMF 2.6 release in August, other arrangements for tape mounting services will be made.
- 4) XFS dump/restore will be adapted to DMF-migrated files as on UNICOS.
- 5) The Fast-flexible I/O library (ffio) will be ported to IRIX.
- 6) IRIX does not have the UNICOS User Database (UDB) feature. DMF features that rely on the UDB will be removed until alternate implementations are found. These features are .keep processing (ARCHLIM) and processes mediated by the ARCHMED field.
- 7) The fsmon utility will be eliminated and its function will be combined with dmmctl. A new utility named dmfsmon will provide the services of a combined fsmon/dmmctl.

8) The DMF tape MSP will be adapted to operate with character-special devices.

9) The `dmmode` command will be removed in order to avoid making changes to the IRIX kernel.

10) In UNICOS, the commands `ls` and `find` were modified to know about migrated files. Rather than modify the IRIX versions of these commands, we will offer new DMF commands `dmls` and `dmfind` with DMF 2.6. A new command, `dminfo`, will be provided that offers a replacement for the shell testing capabilities.

11) The client/server configuration option in DMF 2.5 will be removed from DMF 2.6. This change was made because there is no shared file system under IRIX and because the form and function of client/server configuration options will change in DMF 3.0.

12) The `dmhit` command is replaced by a command called `dmscanfs`. This command will return more information that was previously provided by `dmhit`.

13) The commands `dmdalter`, `dmdump` and `dmdbase` are replaced by a single command, `dmdadm`, which has an administrator interface similar to `dmcatadm` and `dmvoladm`.

14) The `dmmigall` command is renamed to `dmmigrate`.

15) The utilities `dmastat` and `dmmspuse` will not be part of the initial release of DMF 2.6 due to current uncertainty over the format of log files. The `dmdubug` command will be eliminated.

DMF Roadmap

The next couple of years will see several versions of DMF in use by a growing customer base. We see DMF 2.4 remaining in the field as long as there are Cray T3E machines. While it would be most prudent for all DMF customers to move to DMF 2.5 for the added safety afforded by the new database technology, this will not be immediately possible for T3E customers or for customers choosing not to upgrade to UNICOS 9.0. We expect DMF 2.5 to be available for T3E during 4Q97.

SGI customers will be able to run DMF 2.6 starting in August. DMF 2.5 and DMF 2.6 will be the primary products until DMF 3.0 arrives during 2H98.

