

EMC Fastrax™
(Online Intra-storage
Data Movement
Technology)
Product Design
P/N VH50996
REV 3

EMC Corporation 171 South Street, Hopkinton, MA 01748-9103
Corporate Headquarters: (508) 435-1000, (800) 424-EMC2
Fax: (508) 435-5374, Service: (800) SVC-4EMC

Copyright © 1999, 2000 EMC Corporation. All rights reserved.

Printed June, 2000

No part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written consent of EMC Corporation.

The information contained in this document is subject to change without notice. EMC Corporation assumes no responsibility for any errors that may appear.

All computer software programs, including but not limited to microcode, described in this document are furnished under a license, and may be used or copied only in accordance with the terms of such license. EMC either owns or has the right to license the computer software programs described in this document. EMC Corporation retains all rights, title and interest in the computer software programs.

EMC Corporation makes no warranties, express or implied, by operation of law or otherwise, relating to this document, the products or the computer software programs described herein. EMC CORPORATION DISCLAIMS ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. In no event will EMC Corporation be liable for (a) incidental, indirect, special or consequential damages or (b) any damages whatsoever resulting from the loss of use, data or profits, arising out of this document, even if advised of the possibility of such damages.

Trademark Information

Informix® and INFORMIX-OnLine® are registered trademarks of Informix Software, Inc.

Oracle® is a registered trademark and Oracle Parallel Server™ [or Oracle7 Parallel Server™] are trademarks of Oracle Corporation.

PDB™, SPARC™, SPARCcenter™, SPARCcluster™, and SPARCserver™, and UltraSPARC™, and are trademarks of SPARC International, Inc. licensed exclusively to Sun Microsystems.

DiskSuite™, Network File System™, NFS™, Solaris™, Solstice™, Sun™, Sun Microsystems™ and SunOS™ are trademarks of Sun Microsystems, Inc.

Sybase4 MPP™ is a trademark and SYBASE® is a registered trademark of Sybase, Inc.

UNIX® is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd.

VERITAS®, VERITAS FirstWatch®, VxVM®, and VxFS®, are registered trademarks of VERITAS Software Corporation.

ICDA® (Integrated Cached Disk Array), EMC® (the EMC logo), and Symmetrix® are registered trademarks and EMC™ and SRDF™ are trademarks of EMC Corporation.

All other trademarks used herein are the property of their respective owners.

EMC Fastrax

Product Design

EMC Engineering

CONTENTS

<i>Introduction</i>	2
<i>Fastrax Concepts</i>	2
<i>Fastrax Base Technology</i>	4
<i>Enhanced Management for Backup/Restore</i>	8

FIGURES

<i>Figure 1 - Enterprise Storage Architecture</i>	3
<i>Figure 2 - Fastrax Data Engine Front Internal View</i>	4
<i>Figure 3 - Fastrax Data Engine Rear Internal View</i>	5
<i>Figure 4 - Fastrax Concurrent Backup Algorithm</i>	7
<i>Figure 5 - External Management of the Enterprise Storage Domain</i>	8
<i>Figure 6 - Logical Backup of one LBO</i>	10
<i>Figure 7 - Logical Restore of one LBO</i>	10

Introduction

ABOUT THIS DOCUMENT

This document is intended for non-EMC engineering personnel and customers interested in the technical aspects of Fastrax. It contains information on the Fastrax product in the form of a technical innovation roadmap, discussing the product evolution, and architectural benefits. This document exists to assist non-EMC personnel in gaining a broad, comprehensive understanding of the Fastrax architecture for the primary purpose of beginning discussions with EMC representatives regarding Beta testing, product purchase, and integration partnerships.

RELATED READING

For more information on the Symmetrix enterprise product line, features, and benefits, please refer to Symmetrix Model 8xxx, 5xxx and 3xxx Product Guides, available on request.

Fastrax Concepts

Fastrax is a natural evolution of the Enterprise Storage definition. It is where the Symmetrix attributes of enterprise connectivity, cascadability, and information-centricity are extended all the way to secondary storage without impacting the host and/or network between storage and tape resources. It is an important piece of the storage fabric and architecture that creates the Enterprise Storage Network (ESN) Domain.

Fastrax is a unique implementation that off-loads the data movement (between disk and secondary storage) from the network and servers into the Symmetrix Enterprise Storage Domain. The first solution/product built on this intelligent storage architecture is a backup and restore solution to support VLDB Symmetrix customers. In the future this same architecture can support HSM, remote data vaulting, archiving, tape copying, differential merging, virtual tape operations, and intra-storage data migration.

Fastrax is an intelligent platform built to handle data movement from disk to secondary storage and back. It is a technology enabler and architecture from EMC that, combined with leading storage management software, becomes a unique backup/restore solution for customers with VLDBs in both mainframe and open systems environments.

Figure 1 displays a possible enterprise network surrounding Fastrax. Fastrax is connected to the Symmetrix Remote Data Facility (RDF) port of all attached Symmetrix units through ESCON and/or Fibre Channel links. An extended distance solution is also possible due to the functional benefits of SRDF. The data to back up is transferred directly from the Symmetrix units to the Fastrax Data Engine where it is written to the attached tape library.

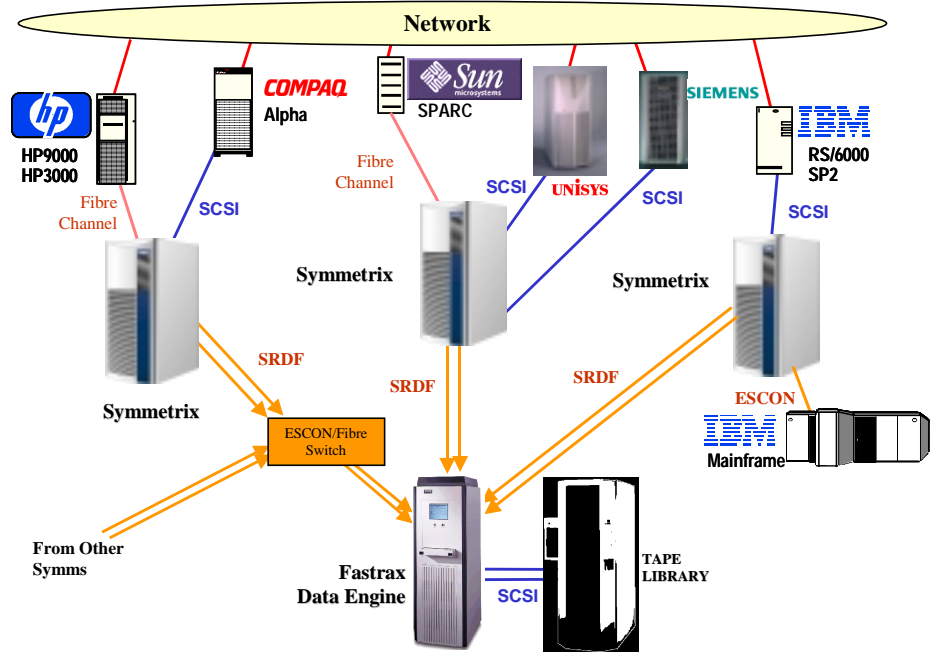


Figure 1 - Enterprise Storage Architecture

FEATURES

The Fastrax solution offers many benefits:

- **Enterprise-wide**, Fastrax can back up and restore all information residing on Symmetrix regardless of host or operating system dependency.
- Eliminates the need for all network traffic and almost all host and application processing cycles. Off-loads host and network cycles while retaining **point-in-time** capability. It enables **"instant" backups** while off-loading all the backup activities from the mission-critical application/database side (i.e. data movement) to the Symmetrix domain.
- Provides **hardware-level performance** while retaining **logical-level flexibility**.
- Provides **faster restore times** as the Fastrax Data Engine allocates space on the data-owning server and writes all the SCSI blocks directly to the Symmetrix.
- **Centralizes** the management of storage resources and protection of data.
- Provides a **scalable architecture** for future growth. Scalability applies to performance and capacity.
- Fastrax's architecture allows the customer to **protect their current investments** in tape, back-up solutions and storage management.
- Provides an electronic, real-time remote vaulting solution through its **extended distance data transfer** capability
- Ability to integrate with the Symmetrix DeltaMark (known previously as Differential Data Facility (SDDF)) in order to handle **differential, incremental backups** (i.e. the capability of backing up only changed data and merging the changes during the restore process).
- Integration mechanisms for the leading backup/restore products to offer **logical backup** capabilities and **sophisticated tape library and media management**.

Fastrax Base Technology

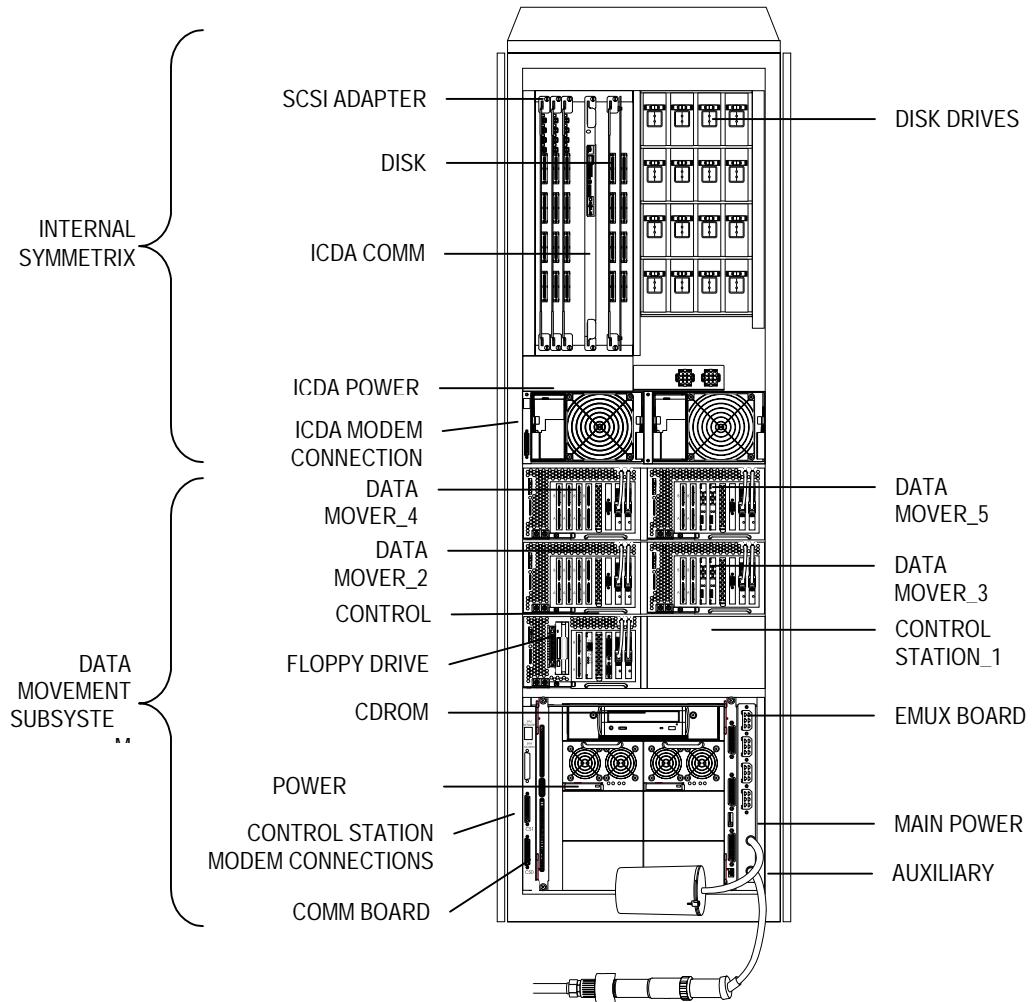


Figure 2 – Fastrax Data Engine Front Internal View

HARDWARE

The Fastrax system hardware consists of a core cabinet containing a cluster from one to four Data Mover processing units and a single Control Station unit. The cabinet also features an external PC keyboard and monitor (console) and houses two redundant power supply systems shared by all processing units and the internal Symmetrix. Additional hardware in the Fastrax system is a Fastrax-dedicated Symmetrix Integrated Cached Disk Array (ICDA). The ICDA is housed in the same core cabinet.

All processing units (Data Movers and Control Station) are currently Intel-based and plug into a common backplane. The backplane features two redundant on-board 10-Mbit Ethernets enabling TCP/IP communication between the units, and two redundant on-board serial buses used by the Control Station to monitor the Data Mover heartbeat and remotely reboot the Data Movers in case of emergency. Each processing unit connects to the internal backplane through two Ethernet and two serial adapters. Additionally, the Control Station features a CD-ROM, a floppy drive, an additional (external) Ethernet adapter, a modem enabling remote servicing (dial-in and call-back), and a connection to the PC keyboard and monitor on the outside of the core cabinet.

Each Data Mover unit runs an EMC proprietary, real-time, embedded operating system specialized in efficient high performance data moving (Data Access in Real Time, or DART). The Control Station currently runs the SCO® UnixWare® operating system.

Each Data Mover unit can have both Front End connections to one or more Symmetrixes and Back End connections to one or more tape drives. All Data Mover units connect to the internal disk storage via fast/ wide SCSI and share access to all disk drives. The Control Station also connects to the internal disk storage via SCSI, but does not have access to all configured disk drives.

Data Movers with Front End connections feature an ESCON/PCI adapter or a Fibre Channel adapter, enabling them to connect to Symmetrix Remote Data Facility Adapters (RAs) through one to six links (each Data Mover) and emulate the target side of an SRDF configuration. Front End Data Movers in the core cabinet need not all connect to RAs in the same Symmetrix. Currently, the connections between the RA and the FE Data Mover are point-to-point.

Each Data Mover has a Back End connection which interfaces with tape drives outside the core cabinet via one to two fast/ wide SCSI interfaces. These tape drives are typically located in a Tape Library Unit (TLU). With a SCSI-attached TLU, the robotics device of the TLU is directly attached to the SCSI interface of one of the Data Movers. Future releases may feature a network-attached TLU as well as a Fibre connection between the Fastrax Data Movers and the TLU. Alternatively, the SCSI- or network-attached TLU may be directly SCSI-connected to the backup server, which controls it directly through backup software. This configuration does present a distance limitation due to the SCSI protocol of 25 meters. SCSI pass-through via Fastrax eliminates this SCSI limitation.

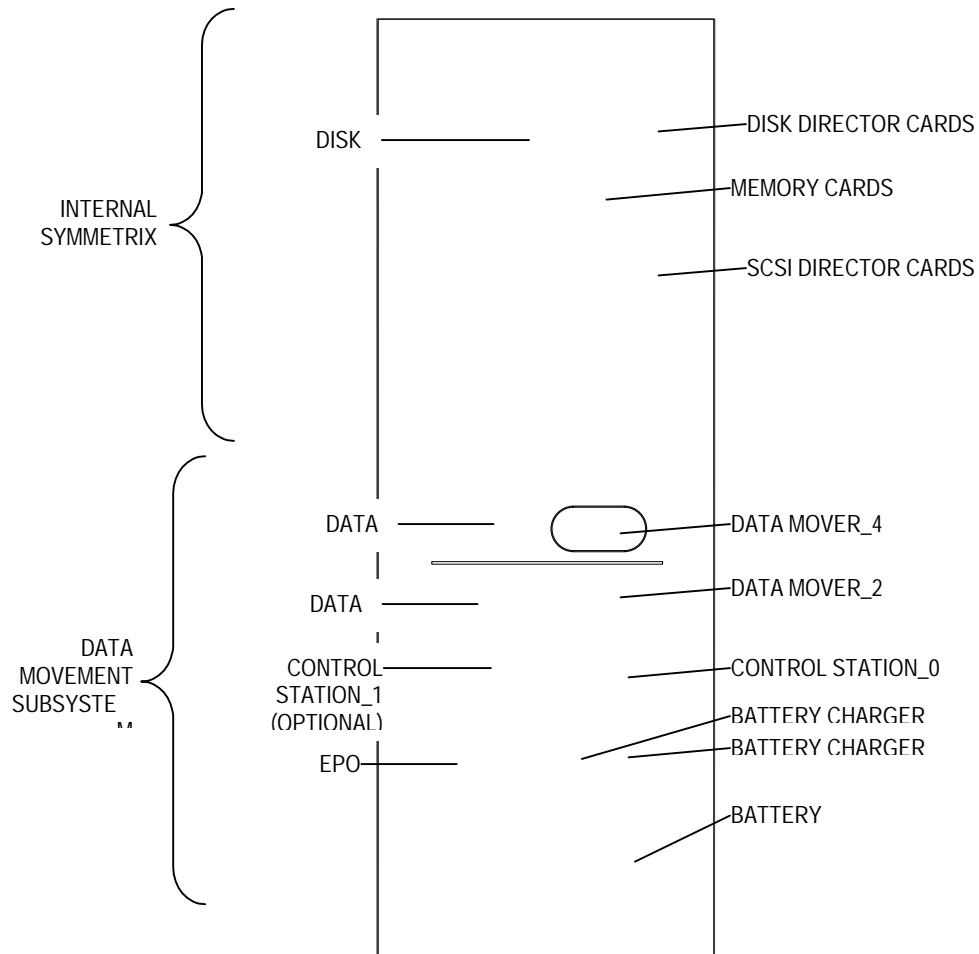


Figure 3 – Fastrax Data Engine Rear Internal View

SOFTWARE

Fastrax software enables "point-in-time" data backups to be made from multiple Symmetrix units to multiple tape drives through one or multiple Fastrax Data Engines for optimal data movement. This means that Fastrax can back up live data stored on multiple Symmetrix even as it is being modified, using minimal host resources. Fastrax's scalability allows for multiple backups to occur simultaneously from multiple Symmetrixes.

Fastrax software has the following main components:

1. EMC's Symmetrix API. Fastrax controls are the latest capability within the SymmAPI. Fastrax features include existing mapping, session and object specification and data movement controls. Using the SymmAPI a backup application interfaces with Fastrax to determine the set of blocks on the Symmetrix that make up the object(s) to be backed up (file, tablespace, VSAM dataset, etc). The backup application defines this set of blocks as one or more Logical Backup Objects (LBO) through the SymmAPI and directs the Symmetrix to back them up via Fastrax. A Logical Backup Object can be any host-level logical entity as defined by the backup application.
2. Interface software running on the host machine. Today this is a Fastrax command line interface or a third party backup application interfacing with EMC's Symmetrix API (SymmAPI). The interface software is responsible for starting, specifying, and closing backup and restore sessions as well as managing the tape media and the TLU robotics. It can also query/monitor Fastrax for backup/restore progress and various TLU information.
3. Fastrax features in the EMC microcode running in the attached Symmetrix unit(s). The Fastrax backup concept is based on the Concurrent Backup functionality implemented in the Symmetrix microcode that enables data to be backed up without burdening the host. Concurrent Backup provides point-in-time data capture functionality while allowing the host to continue with normal I/O operations.

Concurrent Backup works as follows: at backup time, the Symmetrix back end sequentially scans the Symmetrix internal data structures for tracks previously marked by a host backup application integrated with Fastrax. The back end then schedules the appropriate disk track data to be transferred across the SRDF links to the Fastrax unit. Refer to Figure 4.

In order to ensure a point-in-time coherency across the set of marked/protected tracks, each incoming front end write which affects a marked track is temporarily suspended while the back end immediately copies the contents of the affected target track to the sidefile. Only after this is complete can the incoming write request be serviced and the target track overwritten. The back end can then resume its sequential scans of the Symmetrix internal tracks. High performance is assured by this mechanism and the fact that several scans occur in parallel, so the dynamics and order of data transfer from the Symmetrix to Fastrax cannot be predicted. This is why Fastrax does not store all data in order on tape.

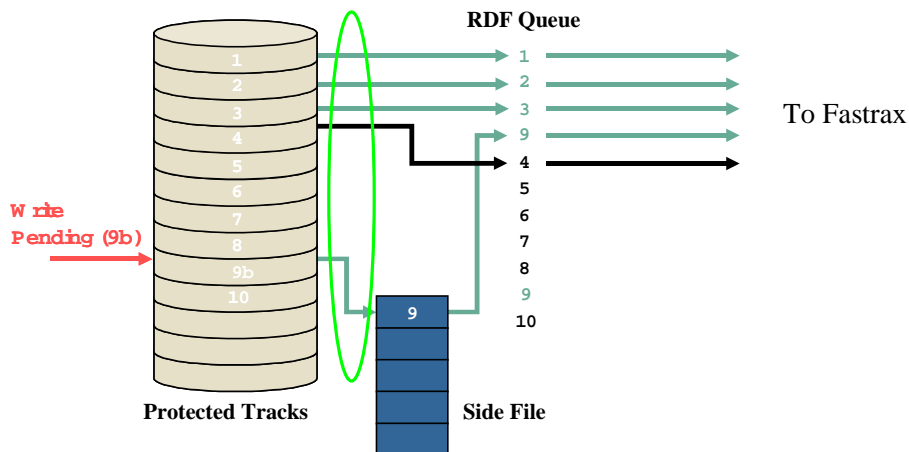


Figure 4 - Fastrax Concurrent Backup Algorithm

4. EMC code running in the Fastrax core cabinet: DART operating system plus Fastrax agents running on the Data Movers. During a backup session, a stream of incoming tracks (for a given backup definition) is accepted by a Front End Fastrax agent running on a Data Mover attached to the corresponding Symmetrix RA through the SRDF links. The agent is responsible for identifying incoming tracks and arranging them into distinct track queues. The track queues are physically located in the shared storage area or internal memory of the internal Symmetrix, visible to all Data Movers. Storage space for track queues is allocated dynamically in 10MB units called disk records.

As an individual disk record of a track queue fills with incoming tracks (which may be out-of-order, as pointed out previously), the agent 'closes' the disk record and passes its 'ownership' (pointer to the shared storage area containing the disk record) to the Back End Fastrax agent executing on the same Data Mover or a cooperating Data Mover. This agent receives the 'ownership' notification through a token passed over the on-board Ethernet. By default, the agent delays processing individual disk records until it accumulates 10 disk records per track queue. This is a tunable parameter. It then reads several disk records from the shared storage pool and writes them sequentially to an attached tape drive. A sequence of ten disk records belonging to the same incoming track queue thus forms a tape segment of 100MB. When the cooperating Front End and Back End agents are running in the same Data Mover, the "disk" records are actually passed through memory, increasing performance.

5. Data Recovery. The backup application reads the tape catalog and metadata to understand the object(s) to be restored. This includes the size of the LBO data to be retrieved. Using the SymmAPI, the backup application pre-allocates space for the backup object. A virtual circuit is created from the Tape Library Unit through the Fastrax Data Engine to the host Symmetrix. A Data Mover reads the tape files that contain the required tracks. The information from the tape file is written to dedicated cache and a Data Mover requests the Symmetrix to retrieve the specified tracks and place them into their pre-allocated space in the correct order. This algorithm provides for more efficient restore windows as the backup data is written directly to the Symmetrix from the Tape Library Unit.

Enhanced Management for Backup/Restore

Fastrax users have several alternatives for backup/restore process management. The basic level is Fastrax's own command line interface (see *SymmBack*, page 8 below). Advanced backup configuration and process management is made possible through partner applications (primarily from backup/restore Independent Software Vendors (ISVs)) that interface with Fastrax through the Symmetrix API (SYMAPI) framework.

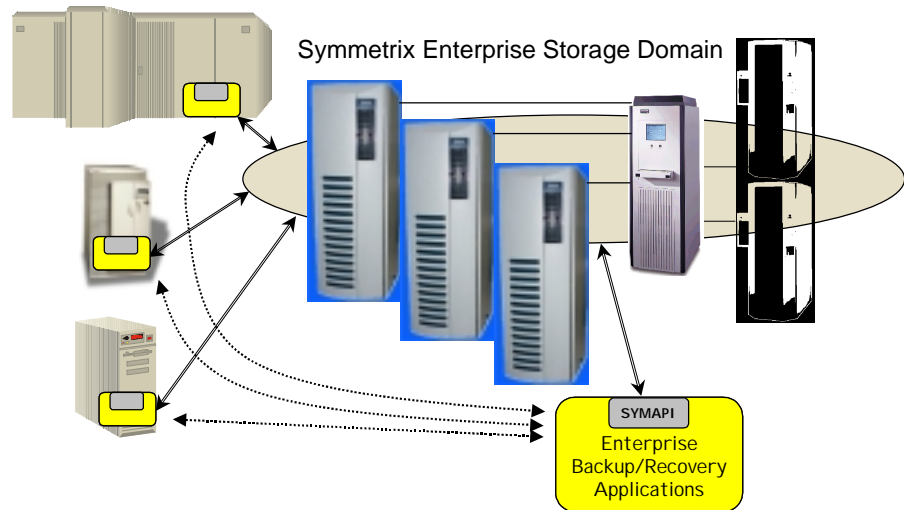
SYMAPI

SYMAPI is a linkable library of functions that can interface with Symmetrix units attached to hosts that are running in a mainframe or open systems environment. The SYMAPI software modules reside on a host system and allow system integrators and ISVs access to Symmetrix configuration, status, and performance data, as well as key Symmetrix features. SYMAPI retrieves data from a Symmetrix unit through the use of standard I/O or channel commands.

SYMAPI allows ISVs to interface with various Symmetrix solutions, facilities, and features, such as TimeFinder, File-level TimeFinder, DeltaMark, logical-to-physical mapping, and Fastrax.

SYSTEM ARCHITECTURE WITH SYMAPI

One or more host-based ISV Storage Management and backup Applications (SMAPPs) interface with the Symmetrix/Fastrax storage domain through the Fastrax API, which is an integral part of the SYMAPI. It is this Fastrax API which allows the SMAPPs to determine the mapping between the object to be backed-up and the set of blocks on the Symmetrix which make up the object. These SMAPPs can also leverage the benefits of the Enterprise Storage Domain through various other components of the SYMAPI. The SMAPPs and their software agents execute on host computers connected to the client Symmetrix units that can be (and typically would be) physically remote to the Fastrax server system. The SYMAPI communicates with Symmetrix via a set of Symmetrix system calls. Symmetrix then propagates these system calls to Fastrax through the



SRDF links. See Figure 5.

Figure 5 - External Management of the Enterprise Storage Domain

EMC'S SMAPP OFFERING

Based on the SYMAPI, EMC has developed a command-line based utility (for open systems) called *SymmBack* to demonstrate the SYMAPI usage model to SMAPP development partners. *SymmBack* will complement the Fastrax product in environments where a partner SMAPP is not available and the deployment of Fastrax is desired. *SymmBack* is intended as a testing and SYMAPI demonstration utility and can potentially be used to provide customers with a basic set of commands that allow Fastrax to be managed and used by the customer if no other means are available. Used for off-line restore for certain applications not supported by the backup application.

FASTRAX/SMAPP INTERACTION The interaction between Fastrax and a partner SMAPP falls into one of two categories:

1. TLU and Media Management.
2. Data storage and retrieval

The SMAPP/Fastrax model of communication is based on the concept of streams. A stream is a virtual, run-time control handle from the SMAPP through the Symmetrix, SRDF link, and Fastrax to a remote tape drive attached to one of Fastrax's BE Data Movers. The SMAPP establishes, controls, and terminates these streams during their lifetime using SYMAPI primitives that are translated into messages by SYMAPI for transfer to Symmetrix and Fastrax. Several streams can be active concurrently, but the maximum number of streams that can be open at any time is limited to the number of tape drives attached to the Fastrax system.

Once a SMAPP establishes a Fastrax stream, it defines a set of objects to be backed up or restored as part of that particular stream, and then passes their definitions to the SYMAPI. In the case of a backup process/stream, the SYMAPI engages the Symmetrix Concurrent Backup facility. In the case of a restore stream, the SYMAPI engages Fastrax's software that transfers the required tracks from Fastrax's tapes via the SRDF link back to Symmetrix. In either case, the SMAPP (or any other instance within the host) is excluded from the direct manipulation of the bulk data, and the data transfer occurs exclusively between Symmetrix and Fastrax. However, the SMAPP is responsible for identifying and mounting the required tapes into Fastrax's tape drives before the Symmetrix/Fastrax data transfer can occur; essentially, all TLU and media management.

The object of data storage/retrieval in the Fastrax architecture is an abstract SMAPP-defined entity called a Logical Backup Object (LBO). It is stored/retrieved by Symmetrix/Fastrax as a whole during the lifetime of a stream according to the detailed specification provided by the SMAPP. Each LBO may or may not have associated with it SMAPP-specific metadata. An LBO can be a complete Symmetrix Logical Volume (or a contiguous extent thereof), a generic file system file or a portion of a file, a database data file, a raw disk partition, a complete database tablespace, single database table, or a set of records of a database table. A fundamental principle is that within Fastrax the LBOs are platform-independent; Fastrax operates on LBOs from both mainframe and open systems environments. LBOs are typically hundreds of MB in size; a single stream can handle several such LBOs. Due to the manner in which Symmetrix receives the backup definition and transfers data to Fastrax (see "Software" on page 6), individual LBOs from a particular stream are usually not stored contiguously on tape; tape segments for the LBOs in a backup definition are typically interleaved.

LOGICAL

Fastrax's LBO storage/retrieval mechanism is sufficiently flexible to allow sophisticated backup/restore of logical entities that have a complex mapping (at a possible granularity of 512 bytes, or one SCSI block) to Symmetrix physical storage. The key to this functionality is the metadata attached to each LBO, which can be used by the SMAPP to store the description of an LBO's logical-to-physical storage mapping (source signature) on tape at backup time. At backup time, the following steps need to occur:

1. SMAPP resolves LBO into a set of Symmetrix tracks plus metadata.
2. SMAPP defines the LBO trackset to Symmetrix.
3. Symmetrix transfers this set of tracks to Fastrax.
4. Symmetrix transfers the LBO's metadata to Fastrax to be appended to the LBO trackset on tape.

Figure 6 illustrates this logical backup process.

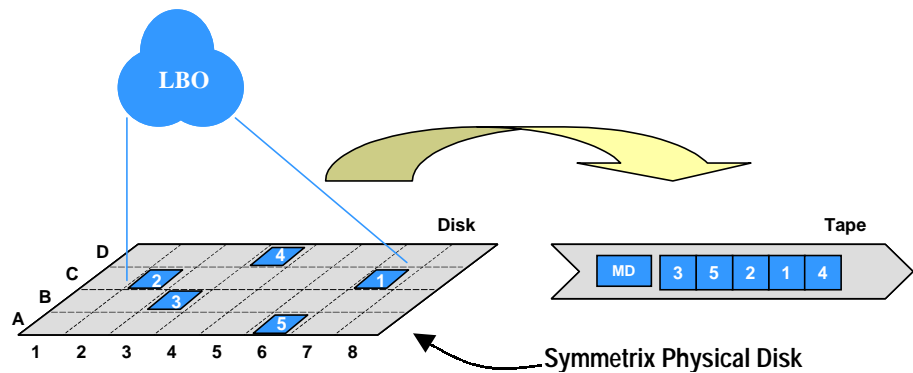


Figure 6 - Logical Backup of one LBO

LOGICAL RESTORE

At restore time, the following needs to occur per LBO:

1. SMAPP retrieves source LBO' metadata from Fastrax.
2. SMAPP preallocates a "dummy" physical target space for LBO".
3. SMAPP resolves the target LBO" mapping.
4. SMAPP calculates the restore mapping matrix (from LBO' to the dummy LBO").
5. SMAPP defines the restore-mapping matrix to Symmetrix/Fastrax.
6. Fastrax transfers the trackset to Symmetrix.

Figure 7 illustrates this logical restore process.

Figure 7 - Logical Restore of one LBO

