

TimeFinder Foundations

Upon completion of this module, you will be able to:

- Identify the basic features of TimeFinder
- Describe the technical positioning of TimeFinder
- Describe the architectural components and theory of operations of TimeFinder
- List several TimeFinder management tools
- Describe several examples of how TimeFinder can be integrated into a Symmetrix environment

The objectives for this module are shown here. Please take a moment to read them.

TimeFinder – A Point in Time Business Tool

- TimeFinder is a family of product provide multiple point in time copies of data in order to distribute business tasks. This allows for simultaneous action of previously sequential business tasks
- An example of this is the ability to backup a point in time copy of critical data while production operations continue on the source data
- TimeFinder provides flexibility, match the right product to your service-level requirements
- Understanding the features and functions of each of the TimeFinder product family options is imperative in order to identify the appropriate solution


TimeFinder is a family of products enabling businesses to provide multiple point in time copies of data in order to distribute business tasks. This allows for simultaneous action of previously sequential business tasks. An example of this is the ability to backup a point in time copy of critical data while production operations continue on the source data.

In today's IT environments, it is critical to have choice and flexibility when building local replication solutions in order to match the right product to your service-level requirements. The TimeFinder family of product offers the broadest choice of alternatives, from full-volume mirrors and clone copies to space-saving snapshot copies. Understanding the features and functions of each of the TimeFinder product family options is imperative in order to identify the appropriate solution.



The TimeFinder Product Family

- Offers solutions to several business continuance scenarios
- Consists of two base solutions:
 - TimeFinder/Clone
 - Mainframe SNAP Facility
 - TimeFinder/Snap
- Consists of four add-on solutions:
 - TimeFinder/Mirror
 - TimeFinder/Consistency Groups
 - TimeFinder/Exchange Integration Module
 - TimeFinder/SQL Integration Module



© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 3 

The TimeFinder family offers solutions to several business continuance scenarios. The TimeFinder family consists of two base solutions: TimeFinder/Clone and TimeFinder/Snap. TimeFinder/Clone includes Clones and Mainframe SNAP.


The add-on solutions include: TimeFinder/Mirror, TimeFinder/Consistency Groups, TimeFinder/Exchange Integration Module, and TimeFinder/SQL Integration Module. This module covers TimeFinder/Clone, TimeFinder/Snap, and TimeFinder/Mirror.





The TimeFinder Product Family (continued)

- Offers solutions to several business continuance scenarios
- Consists of two base solutions:
 - TimeFinder/Clone 
 - Mainframe SNAP Facility
 - TimeFinder/Snap
- Consists of four add-on solutions:
 - TimeFinder/Mirror
 - TimeFinder/Consistency Groups
 - TimeFinder/Exchange Integration Module
 - TimeFinder/SQL Integration Module



The diagram illustrates the TimeFinder Family product structure. A central vertical bar labeled 'TimeFinder Family' contains two main product boxes: 'TimeFinder/Clone' (Ultra-functional, High Performance Copies) and 'TimeFinder/Snap' (Economical Space Saving Copies). To the right of these boxes are four add-on options: 'TF/Mirror Classic High Performance Option', 'TF/CG Consistency Group Option', 'TF/EIM Exchange Integration Option', and 'TF/SIM SQL Integration Option'.

© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 4



TimeFinder/Clone is a full-copy product for open systems that enables the creation of point-in-time copies that can be used for backups, decision support, data warehouse refreshes, or any other process requiring parallel access to production information. Clone copies can be used to reduce disk contention and improve data access speed by assigning users to copies of data, rather than accessing the one production copy. TimeFinder/Clone also includes a mainframe SNAP facility that enables the creation of full-volume copies or dataset-level copies and is fully compatible with STK Snapshot functionality.

TimeFinder/Snap is a pointer-based, space-saving snapshot solution for Symmetrix DMX systems. TimeFinder/Snap, for virtual device copy operations, provides multiple copies of production data for testing, backups, or report generation. Snapping to virtual devices uses a copy-on-first-write technique as a way to conserve disk space when making copies. Only writes to tracks on the source device or target virtual device cause any incremental storage to be consumed. The space savings using virtual devices can be significant when you consider that most applications change only a small percentage of data on a volume. However, as with any pointer-based implementation, you can expect some performance degradation that will vary according to application characteristics and I/O profile.



The TimeFinder Product Family (continued)

- Offers solutions to several business continuance scenarios
- Consists of two base solutions:
 - TimeFinder/Clone
 - Mainframe SNAP Facility
 - TimeFinder/Snap
- Consists of four add-on solutions:
 - TimeFinder/Mirror 
 - TimeFinder/Consistency Groups
 - TimeFinder/Exchange Integration Module
 - TimeFinder/SQL Integration Module

TimeFinder Family

TimeFinder/Clone
Ultra-functional, High Performance Copies

TimeFinder/Snap
Economical Space Saving Copies

TF/Mirror
Classic High Performance Option

TF/CG
Consistency Group Option

TF/EIM
Exchange Integration Option

TF/SIM
SQL Integration Option

© 2006 EMC Corporation. All rights reserved. TimeFinder Foundations - 5 

TimeFinder/Mirror is a business continuance solution that allows you to use special Business Continuance Volume (BCV) Symmetrix devices. Copies of data from a Standard Symmetrix device (which are online for regular I/O operations from the host) are sent to and stored on BCV devices to mirror the primary data. Uses for the BCV copies can include backup, restore, decision support, and applications testing. Each BCV has its own host address, and is configured as a stand-alone Symmetrix device.



The TimeFinder Product Family (continued)

- Offers solutions to several business continuance scenarios
- Consists of two base solutions:
 - TimeFinder/Clone
 - TimeFinder/Snap, Mainframe SNAP Facility
- Consists of four add-on solutions:
 - TimeFinder/Mirror
 -  – TimeFinder/Consistency Groups
 - TimeFinder/Exchange Integration Module
 - TimeFinder/SQL Integration Module

TimeFinder Family

TimeFinder/Clone
Ultra-functional, High Performance Copies
Open Systems & MF

TimeFinder/Snap
Economical Space Saving Copies
Open Systems & MF

TF/Mirror
Classic High Performance Option

TF/CG
Consistency Group Option

TF/EIM
Exchange Integration Option

TF/SIM
SQL Integration Option

© 2006 EMC Corporation. All rights reserved. TimeFinder Foundations - 6 

TimeFinder/Consistency Groups (TimeFinder/CG) is available as an option to any of the base solutions, enabling Engenuity Consistency Assist (ECA) to maintain data coherency across TimeFinder-protected volumes, and to ensure data consistency and restartability.

TimeFinder/Exchange Integration Module (TimeFinder/EIM) is available as an option to any of the base TimeFinder solutions, for those customers needing to quickly and easily integrate TimeFinder and Microsoft Exchange. TimeFinder/EIM also ships with Symmetrix Integration Utilities (SIU), formerly available in ResourcePak for Windows.

TimeFinder/SQL Integration Module (TimeFinder/SIM) is available as an option to any of the base TimeFinder solutions for those customers needing to quickly and easily integrate TimeFinder and SQL Server. TimeFinder/SIM also ships with Symmetrix Integration Utilities (SIU).

 			
TimeFinder Family Service – Level Comparison			
	TimeFinder/Clone (includes MF SNAP)	TimeFinder/Snap	TimeFinder/Mirror Option
Performance	• Supports high I/O workloads and functionality requirements	• Supports moderate I/O workloads and functionality	• Supports extremely high I/O workloads and availability needs
Availability	• Full logical copy • Data is immediately available	• Space-saving virtual copy • Data is immediately available	• Full physical copy • Data is available following Split
Functionality	• Does not require mirror position • Can be RAID 5 protected • Copy is read/write enabled • Up to 16 concurrent copies of single volume in an Open System • Up to 4 concurrent copies in a mainframe environment • Protected restore • Incremental re-sync • Can be source for SRDF Family • MF SNAP offers volume and dataset-level copies	• Does not require mirror position • Can be RAID 5 protected • Copy is read/write enabled • Up to 16 copies of single volume • Incremental re-sync • Cannot be source for SRDF Family	• Requires mirror position • Can be RAID 1 protected • Copy is read/write enabled • Read-through capability • Up to 16 copies of single volume • Protected restore • Incremental re-sync • Can be source for SRDF Family
Economics	• Requires 100% of the capacity of the source volume or dataset • Can be RAID 5 protected	• Requires fraction of capacity of source volume • Can be RAID 5 protected	• Requires 100% of capacity of the source volume • Support for RAID 1 protection
© 2006 EMC Corporation. All rights reserved. TimeFinder Foundations - 7 			

So, with all the choice and flexibility in making TimeFinder family replicas, how do you know which type to use, and when? It's a matter of balancing customer needs and mapping them back to a few key areas—performance, availability, functionality, and economics.

Pointer-based images are not a physical copy of your information; they are logical views of the original information based on the time the image was created. Snapshots and snaps are created in seconds and can be retired when no longer needed.

In contrast to a full data copy, a snapshot uses only a fraction of the original space. Multiple snapshots can be created to suit the need of multiple business processes. Secondary servers see the snapshot as an additional mountable disk volume. Servers mounting a snapshot have full read/write capabilities on that data.


TimeFinder/Mirror and TimeFinder/Clone copies are mirror image full copies of the data implemented. As soon as the information has been copied to its mirror and split from its source, it is a fully complete and independent physical copy of the data.

TimeFinder/Mirror, TimeFinder/Clone, and TimeFinder/Snap can be used together to provide the best solution for mixed production environments, providing the right solution for the right application and enhancing recovery options.



TimeFinder – Benefits

- Repurposing
- Non-disruptive incremental resynchronization
- Consistent methodology for all platforms
- Maximize business operations by enabling parallel processing



Non-disruptively copies of an active application volume, or set of volumes, within the same Symmetrix

© 2006 EMC Corporation. All rights reserved. TimeFinder Foundations - 8 

TimeFinder allows companies to make more effective use of their most valuable resources by enabling parallel information access, as opposed to traditional sequential information access, thus eliminating the need to do things like quiesce an application for backups.

This provides tangible benefits to business including:

- Accelerated upgrades — EMC TimeFinder enabled a rehearsal of application upgrades to quickly identify and remedy problems, minimize risk, and reduce production downtime
- High system availability — EMC TimeFinder can eliminate backup window
- Cost savings—The EMC TimeFinder can shorten the maintenance window, minimizing infrastructure costs and improving service levels



TimeFinder Foundations

THEORY OF OPERATION



© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 9 

Having briefly defined what TimeFinder is, we will now discuss the internal mechanisms of the tool.

TimeFinder Architectural Components

- **TimeFinder/Mirror**
 - Standard and BCVs
 - Copying and Restoring data
- **TimeFinder/Clone**
 - Standard and Clones
 - Copying and Restoring data
- **TimeFinder/Snap**
 - Standard, VDEVs, and Save devices
 - Snapping

Each TimeFinder product has its own distinct architecture and operational characteristics. These attributes are essential when addressing the wide scope of customer data replication requirements. In the course of the next few slides, we will be reviewing the components that make up each TimeFinder implementation.



TimeFinder/Mirror – Simplest Implementation

- Establish BCV with Standard
- Stop I/O for point-in-time copy
- Split BCV (BCV now uniquely addressable)
- Start I/O against Standard
- Execute operations against BCVs in parallel
- Re-establish BCV as required with Standard (resynchronization)



© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 11 


To understand the fundamental working of TimeFinder/Mirror, here are some details on establishing a BCV pair.

- The user needs to first establish a logical connection between the Standard and the BCV. This will cause the Symmetrix system to synchronize these volumes, or to make them exact mirrors.
- Once the system indicates to the host that this has been accomplished, all writes to the Standard will also be posted to the BCV.
- Stop I/O or quiesce the application. Essentially you are waiting for the application and access method buffers to write all data to the Symmetrix system.
- At this point, split the BCV from the Standard, thereby stopping the synchronization between volumes. Processing can then take place against the BCV as well as the Standard.

After the volumes have been split, the system maintains a table of “invalid tracks”. If the BCV is re-established back to the original Standard, the system uses the “invalid tracks” information to minimize the time required to “resynch” the BCV pair.

Business Continence Volumes

- A BCV is a Symmetrix volume with special attributes that allow it to be attached to another Symmetrix Logical Volume, within the same Symmetrix, as the next available mirror
- Each BCV device has its own host address Target ID, Logical Unit Number etc., Symmetrix device number
- BCVs must be the same size (cylinders), type (FBA/CKD) and emulation (Mainframe 3380/3390) as the Standard Device




© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 12


Mercer Road
EMC Education Services

The most fundamental element of TimeFinder/Mirror is a specially defined volume called a Business Continuity Volume. A BCV is a Symmetrix volume with special attributes that allows it to be attached to another Symmetrix Logical Volume within the same Symmetrix as the next available mirror. It must be of the same size, type, and emulation (for mainframe 3380/3390) as the device which it will mirror. Each BCV has its own host address and Symmetrix device number.



A BCV can be configured with a mirror device, as shown in this slide. When a BCV is paired with a Standard device, the relationship with its mirror is broken, except in the case of a protected establish. Immediately after the split of the BCV from the Standard, the split BCV and its mirror resynchronize in the direction specified on the Split action.

SRDF Protected BCV Devices




- When an SRDF Protected BCV is established, data is not synchronized to the SRDF Target. The SRDF link between the R1/BCV and its R2 pair is logically suspended
- When an SRDF Protected BCV is split, changed tracks are synchronized to the SRDF Target

BCVs can be SRDF protected. When an SRDF Protected BCV is established, data is not synchronized to the SRDF target. The SRDF link between the SRDF pair is logically suspended. When an SRDF protected BCV is split, changed tracks are synchronized to the SRDF target.

Multi-BCVs

- Symmetrix Differential Data Facility keeps track of changes to multiple BCVs, one after the other
- Incremental establish or restore now possible with up to eight (default) BCVs



© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 15


Mercer Road
EMC Education Services

Different BCVs can be established and then split from a Standard at different times of the day. It used to be that the Symmetrix did not keep track tables for the previous BCVs, and the track table was only for the last, most recent BCV. Thus, a BCV that was split at 3:00 a.m. could not be re-established if another BCV had been established and split at 4:00 a.m. The 3:00 a.m. BCV would have to be established (full volume copy) again.

Symmetrix Differential Data Facility (SDDF), answers the need for a more efficient way to process changes that occur on the data stored in Symmetrix systems. It is used in TimeFinder so only the tracks on the BCVs that have changed since the last backup are backed up, thereby reducing time and resource use.

Multi-BCVs (continued)

- Symmetrix Differential Data Facility keeps track of changes to multiple BCVs, one after the other
- Incremental establish or restore possible with default value of 8 BCVs, this can be increased to as many as 16



© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 16

Mercer Road
EMC Education Services


With the Multi-BCV function, up to eight (default) BCVs will be maintained by the Symmetrix. In other words, a BCV that was split at 4:00 a.m. can be re-established, even though another BCV was established and split at 5:00 a.m. In this way, a user can split and incrementally re-establish volumes throughout the day or night, and still keep re-establish times to a minimum, thus allowing multiple BCVs to maintain an incremental relationship with the same Standard.

Incremental information can be retained between a STD (Standard) device and multiple BCV devices, provided the BCV devices have not been paired with different STD devices.

The maximum number of BCVs that can be paired with the same STD device can be increased to as many as 16 by using the TimeFinder Solutions Enabler command variable.

Concurrent BCVs

- Establish two BCVs with the same Standard simultaneously, or one after the other
- The BCVs can be split individually or simultaneously




- Each BCV will occupy a mirror position - locally protected R1 Volumes cannot have two BCVs established with them concurrently
- Concurrent restores are not allowed

Concurrent BCVs is a TimeFinder/Mirror - BCV feature that allows two BCVs to be simultaneously established with a Standard. The BCV pair can be split, providing customers with two copies of the data. Each BCV can be mounted online and made available for processing.



TimeFinder/Clone

- High performance full-volume copies
 - Volume level
 - Dataset level (mainframe)
- Any-to-any protection type
 - Including RAID 5
- Highly flexible configuration options
 - Symmetrix mirror position not required on the source
- Up to 16 copies of a source
 - Immediately read and write-able
 - NOCOPY option
- Support for TimeFinder/mirror commands
 - TimeFinder/Mirror emulation support



© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 18 


TimeFinder/Clone is a full-copy product that enables the creation of point-in-time copies that can be used for backups, decision support, data warehouse refreshes, or any other process that requires parallel access to production information. TimeFinder/Clone also includes a Mainframe SNAP facility that enables the creation of full-volume copies or dataset-level copies.

Clones can be protected using any type of supported protection scheme, including RAID 5, and do not require a Symmetrix mirror position on the source. You can create up to 16 active clones of a single source volume (4 in a mainframe environment) which are available immediately for both read and write access. TimeFinder/Clone also includes a “no copy” option which enables you to perform a copy process, only when the actual data is requested. And, as with all TimeFinder Family products, TimeFinder/Clone supports the TimeFinder/Consistency Groups option to ensure data consistency between volumes and even across Symmetrix systems.

TimeFinder/Clone is ideal when high performance, RAID 5 protection and highly-functional point-in-time copies are required. If ultra-performance and availability is needed, the TimeFinder/Mirror option can be added to TimeFinder/Clone.



TimeFinder/Clone “Copy on First Access”



- No data moved initially
- Appears to host like source device
- Source and Target must be the same size and type (CKD/FBA)
- Source and Target devices can be either standard devices or BCV devices

- Clone command to activate point-in-time
- Data copied on first access
- Optional setting copies all tracks, no waiting for access

© 2006 EMC Corporation. All rights reserved.TimeFinder Foundations - 19

A clone is a point-in-time image of the full source volume in open systems. The source and target devices can be either standard devices or BCV devices, as long as they are all of the same size and emulation type (FBA or CKD).


Copying to a clone allows you to copy data from a single source device to as many as sixteen target devices, with automatic background copying of up to four target devices simultaneously. The target of a copy operation is a Symmetrix clone. Copying to a clone can be either immediate or deferred.

Deferred copies those tracks on a full device that contains the data of interest. Copying occurs when there are writes to the source device or reads/writes to the target device. The clone pair state remains CopyOnAccess until you terminate the copy session or all tracks have been accessed. This is the default option in Open Systems.

Immediate copies the full source device to the target device. This requires the “-copy” option for Open Systems. The clone pair state is CopyInProgress until the clone pair reaches the Copied state. This is the default option in a mainframe environment.



Copying Data to the Clone



- **Copy On Access**
 - The first time a track on the source device is written to, or a track on the target is accessed (read or write)
 - Default setting for Open Systems, option for Mainframe
- **Full Background copy**
 - The `-copy` option specifies complete copy as background operation
 - Default setting for Mainframe, option for Open Systems
- **Pre-copy option**
 - Start copying tracks in the background, before the clone session is activated
 - Default setting for Mainframe, option for Open Systems

© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 20 


Using the copy on access method when activating a Clone means that the first time a track is either written to or read from, the track is copied to the Clone. In this case, both the source and the Clone are affected by the I/O.

In open systems, this functionality is enabled by default. When a copy session is finally activated, the device pair state will be CopyOnAccess.

In the mainframe environment, there is an option rather than a default setting for CopyOnAccess. It is invoked by adding the `MODE(NOCOPY)` parameter to the command.



Copying Data to the Clone (continued)



- **Copy On Access**
 - The first time a track on the source device is written to, or a track on the target is accessed (read or write)
 - Default setting for Open Systems, option for Mainframe
- **Full Background copy**
 - The `-copy` option specifies complete copy as background operation
 - Default setting for Mainframe, option for Open Systems
- **Pre-copy option**
 - Start copying tracks in the background, before the clone session is activated
 - Default setting for Mainframe, option for Open Systems

© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 21 

In addition to the CopyOnAccess method, a full volume copy can be created by copying all source tracks as a background operation. In Open Systems, this is executed by the create argument with the “-copy” option. In the mainframe environment, this is the default operation. When the copy session is eventually activated, data begins background copying so that a full copy of the data will become available on the target device. While background copying, the state of the device pair is CopyInProgress; when the operation completes, the state goes to Copied. The copy session must be activated before the target host can access the data.

The -precopy function starts copying tracks in the background, before the clone session is activated. In open systems, this can be used with the create argument using -precopy option. In the mainframe environment, this is the default setting. While in the Created state, the pre-copy process never actually completes; instead, the process keeps checking for new writes to be pre-copied to the target device until the session is activated. Once activated, the normal background copy mechanism takes over and the pre-copy operation ceases.

Restoring Data from a Clone



Restoring Original Data


- Full restore
 - Copy target data to original source or another device
- Incremental restore
 - Copy target data back to the original source device
- In a Mainframe environments, restore can be achieved by snapping the target volume back to the original source or another device

You can use the restore command to copy target data to another device (full restore), or back to the original source device. In the case of a full restore, the original session will terminate and a copy session to the target of the restore will start. In the case of an incremental restore, the original session will terminate and an incremental copy session back to the original source device will start. To support this operation, the session must have been created with the differential option and the device must be in a fully copied state.

In a mainframe environment, the target volume can be copied to another device or back to the original source device by performing another SNAP VOLUME operation. To support a differential SNAP back to the original source volume, the original SNAP VOLUME command must have been issued with the DIFFERENTIAL(YES) parameter.

TimeFinder/Snap for Symmetrix DMX

- Snapshots create logical point-in-time images of a source volume
- Requires only a fraction of the source volume's capacity
- Multiple snapshots can be created from a source volume and are available immediately
- Snapshots support both read-only and read/write processing




TimeFinder/Snap creates space-saving, logical point-in-time images or “snapshots”. The snapshots are not full copies of data; they are logical images of the original information, based on the time the snapshot was created. It’s simply a view into the data.

A set of pointers to the source volume data tracks is instantly created upon activation of the snapshot. This set of pointers is addressed as a logical volume (VDEV) and is made accessible to a secondary host that uses the point-in-time image of the underlying data.

TimeFinder/Snap for Symmetrix DMX (continued)

- Supports mainframe and open systems host environments
 - 15 concurrent snaps in open systems
 - 8 concurrent snaps on mainframe
 - Full function access to the snap
- Complements TimeFinder and provides unmatched replication flexibility




TimeFinder/Snap supports:

- Open and mainframe data volumes
- Up to 15 concurrent snaps of a single source volume in open systems and eight on mainframe
- Full function access to the snap; that is, the snap may also be updated with the updates residing in the Save Area


Save Devices contain any original tracks that were changed as a result of a first copy on write to a source device or a new write to a virtual device (VDEV) during a virtual device copy session. Symmetrix supports the creation of multiple named SAVE device pools, allowing commands to use a specific pool.

In many situations, you find yourself trying to support multiple service levels with a single solution. EMC Snap now allows you to complement existing TimeFinder environments, fulfilling multiple service level requirements while balancing the economics of the solutions.



The Production view is a standard view into an active production volume from the host's perspective. EMC Snap creates the same type of view into your production data at a specific point in time. The cache-resident pointers maintain the point-in-time nature of the snapshot as unchanged data is “shared” with the Production view and changed data is temporarily collected in the pre-defined Save Area.


Best practice for Snap would not require more than 20% of the source volumes' capacity in the save area.



TimeFinder/Snap uses a process called “Copy on First Write” when handling writes to the production data when a snapshot is running.


For example, let’s say a snapshot is active on the production volume. When a host attempts to write to the data on the production volume, the original Track C is first copied to the Save Area, then the write is processed against the production volume. This process of pointers maintains the consistent, point-in-time copy of the data for the ongoing snapshot.

How to Perform a Backup with TimeFinder/Snap



- Backup host will see an image of the volume as it existed at the time of the snapshot
- Will read from both the shared source data and the copy of first write data that resides in the Save Area
- Some I/O contention on the shared source volume

Here is an example of a backup operation from a snapshot. The snapshot gets created and the backup host mounts the snapshot and performs the backup against the snapshot. It is important to note that the backup runs a sequential read process against the snapshot, so the production application may encounter some performance contention during the backup, due to the fact that the snapshot and the production host are both looking at many of the same spindles for the unchanged data.



Here is an example of an incremental restore operation back to the original Standard. When the restore is initiated, the host application must be offline so as not to create the potential for data corruption during the restore process. The restore completes after all the collected changed data in the Save Area as that particular snapshot is copied back to the Standard.



You can also restore to a BCV of the Standard when it is in split mode, or to a completely separate volume.

FUNCTIONAL OPERATION CONCEPTS


TimeFinder Foundations - 29 Mercer Road
EMC Education Services

TimeFinder/Mirror – BCV Operations

- Establish BCV pairs
 - Synchronize the Standard to the BCV
 - Options
 - Full or Incremental Establish
- Split BCV pairs
 - Terminate mirror relationship between BCV and Standard
 - Options
 - Differential Split
 - Reverse Split
 - Instant Split
 - Consistent Split
- Restore BCV pairs
 - Synchronize contents of BCV to the Standard
 - Options
 - Full or Incremental Restore
- Verify all BCV device
 - Provide current status of BCV/Standard pairs



TimeFinder/Mirror - BCV allows the user to create several host-based business continuity operations and control these operations through host commands. The user can establish a BCV, then split it to make the BCV available to a host/server. After completing the business continuity processes on the BCV device, the user can re-establish the BCV pair. Typical BCV operations include establish, split, re-establish, restore, incremental restore, and verify.




BCVs must first be “established”. For example, the BCV is first made Not Ready (unavailable to the host), then the entire contents of the Standard is copied to the BCV.

When this is completed, the BCV can be “split” from the Standard and mounted to the host or another host. A split is “instant”.

BCVs can also be “re-established”, thereby updating the BCV with changes that have occurred on the Standard volume.

Splitting a BCV



- Activity is quiesced on the production volume; create a point-in-time version of the data
- The split command ends the mirroring between the Standard (production volume) and the BCV
- Activity is restarted on production
- BCV is now addressable as a point-in-time copy of the original production volume
- As updates occur, Symmetrix maintains a log of updates to both volumes

In this example, the production volume and the BCV were established (see upper diagram), and then split (lower diagram). Updates are now being maintained in track tables for both the production volume and the BCV until they are re-established or restored.

Instant Split enables a volume that was previously established to be split almost instantaneously—in less than a second from the viewpoint of the host or server. Completion of the split process runs as a background task. This is the default Split operation for TimeFinder.

Default Split operations for TimeFinder occur as follows:


- BCV devices are set to Read/Write status immediately after a split command is executed. The BC applications have immediate access to the BCV devices.
- Requests (from the BCV or the STD) for tracks that have not been split are handled on a priority basis.




Point-in-Time Consistency with Consistency Assist

- Engenuity Consistency Assist enables multiple volumes to be captured as one logical unit
 - Ensures consistency across multiple volumes
- I/O from host is temporarily held for all devices in the group
 - Dependent write ordering is maintained
 - Timefinder, BCV pairs, snap pairs, clone pairs are activated
 - I/O resumes without application impact
 - Devices are consistent and restartable as a set

Production Volume



Production Volume



Restartable Application

Consistent activation enables multiple volumes to be created as "one", ensuring consistency restartable image



Host Application


© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 33 

When you create a point-in-time image across multiple devices, it is imperative that the entire set of logical volumes be captured at the exact same time.

One way to achieve this is to shut down, or totally quiesce an application so no I/O occurs while you create the sessions. This is obviously a problem in today's application environments.

EMC has a solution to this problem. It's called Engenuity Consistency Assist. When you create a Timefinder BCV Establish, Clone, or Snapshot and invoke Engenuity Consistency Assist, the Symmetrix will align the I/Os of those devices and halt all I/O from the host systems very briefly—much faster than the applications can detect—while it creates the sessions and copies the tracks. It then resumes normal operation without any application impact. This ensures write dependency across all devices, creating a restartable image.



The device groups and the composite group are entities that you can create and use to manage and control BCV pairs. You can create a composite group to control a set of BCV pairs that spans multiple Symmetrix arrays. A composite group provides greater flexibility than a device group, which can define devices only on a single Symmetrix array.

Symmetrix device groups or composite groups must be created on the controlling host for the targets to be consistently split.


Composite groups are used in conjunction with Consistent split to create a re-startable database image. This will be discussed in up-coming slides.

In a mainframe environment, device ranges can be specified in the SPLIT command, or GNS Groups can be created and used to control multiple device pairs and ensure consistency.




Host and Enginuity Consistency Assist (ECA) Consistent Splits

Host-based Split



- Open Systems:
 - PowerPath holds I/O during split
 - Can affect multiple hosts
- Mainframe:
 - IOS holds I/O during split
 - Can affect multiple hosts
- Performed from host doing I/O

ECA-based Split



- Symmetrix holds I/O during split
- Performed by any Symmetrix-attached host

© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 35




When databases become very large (larger than one terabyte), the time and resources required for host based software to protect or run decision-support queries on these databases becomes a critical factor. The time required to quiesce or shut down the database for offline operations is no longer acceptable for most database applications. When you need to perform such database operations, the use of TimeFinder consistent split allows you to split off a consistent, DBMS-restartable copy of the database in seconds with no interruption to online service.

TimeFinder software with PowerPath support provides a consistent-split implementation of instant split that allows you to split off a consistent, DBMS-restartable BCV copy of your database without having to quiesce or shut down the database first. The impact on your production environment is negligible.




Host and Enginuity Consistency Assist (ECA) Consistent Splits (continued)

Host-based Split



- Open Systems:
 - PowerPath holds I/O during split
 - Can affect multiple hosts
- Mainframe:
 - IOS holds I/O during split
 - Can affect multiple hosts
- Performed from host doing I/O

ECA-based Split



- Symmetrix holds I/O during split
- Performed by any Symmetrix-attached host

© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 36



PowerPath or Enginuity Consistency Assist allows hot splits of re-startable database images, without quiescing or placing DB in hot backup mode. It provides application support for Oracle, IBM UDB, and SQL Server.

PowerPath will suspend I/O to the devices at the host level, instant split is executed, PowerPath then resumes I/O to the devices.


Enginuity Consistency Assist (ECA) will hold I/O at the Symmetrix level, which eliminates the need for PowerPath. Consistency is maintained across multiple servers.

In a mainframe environment, PowerPath is not required. Instead, the I/Os subsystem is used to hold I/Os at the host level across a single or multiple LPARs.



Consistent Split

- Performs an Instant Split across a group of devices using a single Consistent Split command
- Composite Group Identifies the BCVs to split
- All the BCVs in the group are thus consistent point-in-time copies
- Used to create a consistent point-in-time copy of an entire system, an entire database, or any associated set of volumes
- Can create consistent splits on remote BCVs through SRDF




© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 37 

Consistent split performs an Instant Split across a group Composite Group of Devices using a single Consistent Split command. This creates a consistent point-in-time copy of all the BCVs in the Composite group. The consistent split functionality can also be used on remote BCVs through SRDF.

Consistent split is an implementation of instant split that avoids inconsistencies and restart problems that can occur if you split a database-related BCV without first quiescing the database. The difference between a normal instant split and a consistent instant split is the addition of PowerPath or Enginuity Consistency Assist (ECA) to suspend the database device writes at the host level for a very short time while the foreground split occurs. This feature can be used to create a consistent point-in-time copy of an entire system, an entire database, or any associated set of volumes.


Reverse Split




- A reverse split will overwrite the newly synchronized production data on the BCV with the original data from its mirror

In an established state, any logical corruption on the STD will be propagated to its BCV. The reverse split option allows the recovery back to a prior point in time. The reverse split will revert the BCV to a prior point in time image of the standard from when the establish command was issued. This prior point in time copy can then be used to recover the standard from a failed condition. A Reverse Split is only available for Locally Mirrored BCVs.

Re-establishing a BCV Pair



- Volume pair (Production and BCV) have been split
- Tracks on production updated after split
- Tracks on BCV have also been updated after split
- Symmetrix keeps table of changed tracks after split
- At re-establish, changed tracks are written from production to BCV (incremental re-establish)
- Synchronization then completes



© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 39


Mercer Road
EMC Education Services

In this case, the production volume and the BCV were split, and updates have occurred on both volumes (top). A “re-establish” will join the two volumes back as a BCV pair (bottom).

All updates that have occurred on the production volume must be transferred to the BCV.

All updates that occurred on the BCV also must be replaced. This is accomplished by overwriting the tracks on the BCV with the corresponding tracks from the production volume. This destroys the updates that have occurred on the BCV.

Remember, in a re-establishment activity, the direction of the data flow is from the production volume to the BCV.



The “Restore” is very powerful and can serve many uses. However, caution must be taken before issuing the command. Using the incremental restore option, the changes that have occurred on the production volume will be destroyed (lower diagram). Corresponding tracks from the BCV will be written to the production volume. Changes that have occurred on the BCV will be written to the production volume.

Using the Full Restore option, all the data on the BCV will be written to the production volume. This requires the standard volume to be made unavailable to the host first (offline/unmounted).

Remember, in a Restore function, the direction of the data flow is from the BCV to the production volume. During a Restore operation, the BCV is first made Not Ready (unavailable to the host); the contents of the BCV are then copied to the Standard.



TimeFinder/Clone: Operations

Step 1

Create

- Establish relationship between Standard and Clone

Step 2

Activate

- Clone is now active and available immediately for read/write access
- Production I/O is processed against standard

Step 3

Recreate and Activate

- Clone is re-attached to Standard for new point-in-time copy (incremental)

Step 4

Restore

- Re-attached to Standard and incremental, or full restore performed

Step 5

Terminate

- Causes the target host to lose access to Clone, Removes Clone

© 2006 EMC Corporation. All rights reserved.TimeFinder Foundations - 41

The TimeFinder/Clone creation and usage process consists of five main steps:

- Create
- Activate
- Recreate and Activate (or “Establish”, which combines the two procedures into one)
- Restore
- Terminate

Each step is a critical component of a disk-based local protection and recovery solution. Let’s spend some more time on each of the steps.

TimeFinder/Clone: Operations Details

Step 1 **Create**

- Establish initial relationship between Standard and Clone
- Only required the first time
- Two states:
 - Create in progress
 - Created



- In a mainframe environment, Create and Activate can be performed as a single step

In order to activate a Clone, you must first create the pairing between the Standard and the yet-to-be-created clone.

Open Systems: “Create” sets up the pairing between the devices in the device group with one or more target devices that are associated with the group.

While the operation is in progress, the state of the device pair is “Create In Progress”. When the operation completes, the state changes to “Created”. This “Create” process is done once, and subsequent use of the clone pair only requires an activate, re-activate, or restore command.

Mainframe: In a mainframe environment, Create and Activate can be performed as a single step.

TimeFinder/Clone: Operations Details

Step 2 **Activate**

- Clone is active and available immediately for read/write access
- Optional settings:
 - Base setting performs background copy
 - “No-copy” optional setting uses copy-on-first-access process from background copy as needed
- Production I/O continues against Standard
- In a mainframe environment, Create and Activate can be performed as a single step



- In a mainframe environment, Create and Activate can be performed as a single step

Once the clone pairing has been created, it can be activated so the copy process begins and the clone can be mounted for use.


Open Systems: “Activate” enables an internal copy session with the devices in the device group with one or more target devices that are associated with the group. While the operation is in progress, the state of the device pair is either “CopyInProgress” or “CopyOnAccess”, depending on which option is selected. When the operation completes, the state changes to “Copied”.

Mainframe: In a mainframe environment, Create and Activate can be performed as a single step.

TimeFinder/Clone: Operations Details

Step 3 **Recreate and Activate**

- Clone is re-attached to Standard to create new point-in-time copy
 - Clone is marked “not ready”
- Only changes since original clone activation are copied
 - Uses a “differential” option so only incremental update required



- In mainframe, use SNAP VOLUME with parameter DIFF(YES)

Open Systems: When the time comes to update the existing clone with any new changes since the last time it was created, simply issue a “Recreate” command. Recreate is the ability to incrementally copy changed data to the clone device. The clone session needs to be created with the –copy (or –precopy) and –differential options. Once the clone session is fully copied, changes to the source device (made after the point-in-time copy initiated) can be propagated to the target device.

This feature uses the industry-proven Symmetrix Differential Data Facility (SDDF) functionality that has been part of the TimeFinder Family for years.

An “Activate” command is necessary to establish the new point-in-time copy and initiate the copy of the changed data. “Recreate” can then be used again on the same session to copy over the new data, if desired.

Mainframe: In a mainframe environment, the clone can be re-attached, updating the existing clone to the Standard by using the SNAP VOLUME command with the DIFF(YES) parameter.

TimeFinder/Clone: Operations Details

Step 4 **Restore**

- Attach Clone back to Standard for restore operation
- Restore options:
 - Full restore back to Standard or independent volume
 - Incremental restore back to Standard
 - Standard is available immediately for processing activities



- In mainframe, use SNAP VOLUME command from the Target to the Source with the parameter DIFF(YES)

Open Systems: If a disaster or data corruption occurs and you need to restore from the Clone, you can use the “Restore” operation. The Restore operation enables you to copy the Clone data back to the Standard (original source), either fully or incrementally. Or, you can restore to a completely different device. Like the “Recreate” command, the original clone session needs to be created with the “-differential” flag and in a fully copied state in order to issue a restore command.

During a restore process, the original Standard is available for application processing as soon as the restore is initiated.

Mainframe: In the mainframe environment, use the SNAP VOLUME command from the Target to the Source with the parameter DIFF(YES).

TimeFinder/Clone: Operations Details

Step 5 Terminate

- Causes the target host to lose access to Clone
 - Deletes the pairing information in the Symmetrix array
 - Removes any hold on the target device
- Once the Clone session is terminated, the information is no longer available on the Clone



- In mainframe environment, termination is executed using the STOP SNAP DATASET or STOP SNAP VOLUME

Terminating a copy session deletes the pairing information in the Symmetrix array and removes any hold on the target device. Terminating a session while the device pairs are in the CopyOnAccess or CopyInProgress state will cause the session to end.

In the mainframe environment, termination is executed using the STOP SNAP DATASET command or STOP SNAP VOLUME command.



TimeFinder/Snap: Operations

Step 1

Create

- Establish relationship between Standard and Virtual Device.

Step 2

Activate

- Activating the copy session starts the copy on first write mechanism
- Production I/O is processed against standard

Step 3

Restore

- Copy tracks from Virtual Device to standard or other device

Step 4

Terminate

- Causes the target host to lose access to data pointed to by the virtual device

© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 47 

The TimeFinder/Snap creation and usage process consists of four main steps:

- Create
- Activate
- Restore
- Terminate

Each step is a critical component of a disk-based local protection and recovery solution. Let's spend some more time on each of the steps.

TimeFinder/Snap: Operations Details

Step 1 **Create**

- Establish initial relationship between Standard and the Virtual Device, Only required the first time
- Two states:
 - PAIRING of Source and Target Devices:
 - CREATE operation



- In mainframe, you can snap volumes or datasets, Create and Activate operation can be performed as a single step

In order to activate a Snap session, you must first create the pairing between the Standard and the Virtual Device.

“Create” sets up the pairing between the devices in the device group with one or more target devices that are associated with the group.

While the operation is in progress, the state of the device pair is “Create In Progress”. When the operation completes, the state changes to “Created”. This “Create” process is done once. The virtual data becomes accessible to its host when the copy session is activated.

In the mainframe environment, you can snap volumes or datasets with the SNAP VOLUME or SNAP DATASET. The Create and Activate operation can be performed as a single step.

TimeFinder/Snap: Operations Details

Step 2 **Activate**

- Snap is active and available immediately for read/write access
- Starts the virtual copy session using the “copy on first write” mechanism
- Production I/O continues against Standard
- Target host can access the Snap and has access to data on the source host until the Snap session is terminated



- In mainframe, Create and Activate operation can be performed as a single step

Activating the Snap session starts the copy on first write mechanism and places the target device in the Read/Write state. The target host can access the Snap virtual copy. It has access to data on the source host until the Snap session is terminated.

In the mainframe environment, the Create and Activate operation can be performed as a single step.

TimeFinder/Snap: Operations Details

Step 3 **Restore**

- Restore back to Standard or Other Volume
- Restore options:
 - Full restore back to Standard or independent volume
 - Incremental restore back to Standard
- Standard is available immediately for processing activities



© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 50  Mercer Road
EMC Education Services

If a disaster or data corruption occurs and you need to restore from the Virtual Device, you can use the “Restore” operation. The Restore operation enables you to copy the Snap data back to the Standard (original source), either fully or incrementally.

An increment restore copies only changes made to the virtual device tracks during the active copy session back to the original source device.

A full restore copies all the virtual device tracks back to the original source device or an other volume.

Beginning with Solutions Enabler version 5.4, TimeFinder/Snap performs a persistent snap restore. Persistent means you can restore to a source device that has other virtual device snap sessions, and these sessions along with the restore session continue during and after the restore operation. In this way, both the restore session and other virtual device snap sessions remain intact after the restore operation.

Once you restore from a virtual device, you need to terminate this restore session before performing another restore operation from the same device.

TimeFinder/Snap: Operations Details

Step 4 Terminate

- Causes the target host to lose access to data pointed to by the virtual device
 - Deletes the pairing information in the Symmetrix array
 - Removes any hold on the target device
- Once the virtual copy session is terminated, the information is no longer available on the virtual device.



Remove
Relationship



- In mainframe environment, termination is executed using the STOP SNAP DATASET or STOP SNAP VOLUME

Terminating a copy session deletes the pairing information in the Symmetrix array and removes any hold on the target device. Terminating the session causes the target host to lose access to data pointed to by the virtual device. Terminating a session while the device pairs are in the CopyOnWrite state will cause the session to end. Once the virtual copy session is terminated, the information is no longer available on the virtual device.

In the mainframe environment, termination is executed using the STOP SNAP DATASET command or STOP SNAP VOLUME command.



TimeFinder Foundations

MANAGEMENT OPTIONS



© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 52 

Now let's explore several ways to manage the TimeFinder family of products.

TimeFinder Management Tools

- **EMC Control Center**
 - Easy, point-and-click access
 - Excellent for ad hoc TimeFinder operations
- **EMC Replication Manager**
 - Discovers replication environments
 - Automates replication process
 - Integrates replication technologies at the application level

TimeFinder can be managed by a variety of management tools such as:

EMC Control Center - Through the graphical user interface (GUI) within EMC ControlCenter software related devices are grouped together in device groups. TimeFinder operations may be performed on all devices in a device group using a single command, since group information is maintained in the SYMAPI database.

EMC Replication Manager - Replication Manager is simple, automated, and open. It is simple because Replication Manager is one application that automates, simplifies, and manages disk-based replications. It insulates the user from the storage environment, replication software being used, the type of server, the operating system, disk volumes, etc.

TimeFinder Management Tools (continued)

- **TimeFinder/Exchange Integration Module (TimeFinder/EIM)**
 - Provides a comprehensive backup management interface specifically for Windows servers that support Microsoft Exchange databases residing in Symmetrix storage
 - Produces exact copies of the production volumes that hold the Exchange server information stores and logs
 - Full and single mailbox restores in a fraction of the usual time
- **TimeFinder/SQL Integration Module (TimeFinder/SIM)**
 - Provides a comprehensive backup and recovery management interface specifically for Windows servers that support Microsoft SQL Server databases
 - Integrates and collectively automates the command actions and behavioral features

TimeFinder/Exchange Integration Module (TimeFinder/EIM) - The TimeFinder/EIM software produces exact copies of the production volumes that hold the Exchange server information stores and logs. This enables Exchange administrators to perform both full- and single- mailbox restores in a fraction of the usual time.

TimeFinder/SQL Integration Module (TimeFinder/SIM) - TimeFinder/SIM provides a comprehensive backup and recovery management interface specifically for Windows servers that support Microsoft SQL Server databases. TimeFinder/SIM integrates and collectively automates the command actions and behavioral features of EMC TimeFinder, SRDF, and the tools supplied with Microsoft SQL Server 2000 products.

TimeFinder Management Tools (continued)

- **Solutions Enabler TimeFinder Feature**
 - TimeFinder/Mirror - General monitor and control using BCVs
 - TimeFinder/Snap - Snap copy sessions
 - TimeFinder/Clone - Clone copy sessions
 - TimeFinder/CG – Composite / GNS groups
- **Perform control operations on device pairs within your TimeFinder environment**
- **Automation for Open Systems scripting occur through scripting ,using Unix Shells, or Windows Command Line Interface**
- **Mainframe Management Tools**
 - Single command operations
 - Operator Commands through SRDF-Host Component
 - Through ISPF panels
 - Batch operations for scripting through JCL

The SYMCLI TimeFinder component extends the basic SYMCLI command set to include TimeFinder or business continuance commands that allow you to perform control operations on device pairs within your TimeFinder environment. TimeFinder Family CLI provides functionality for general monitor and control operations using BCVs, Snap copy sessions, Clone copy sessions and Consistency groups.

Commands such as symmir, symbcv, and symioctl allow you to perform a wide spectrum of monitor and control operations on standard/BCV device pairs within your TimeFinder/Mirror environment. With the use of the symioctl command, you can send I/O control commands to a specified database server. The TimeFinder/Clone command, symclone, allows you to create point-in-time copies of a source device on multiple target devices. The TimeFinder/Snap command, symsnap, allows you to create virtual device copy sessions between a source device and multiple virtual (VDEV) target devices. Automation for open systems scripting occurs through scripting.

Mainframe Management Tools

TimeFinder operations in a mainframe environment can be performed using several tools:

- Single command operations
- Operator Commands through SRDF-Host Component.
- Using ISPF panels, individual commands can be issued or batch JCL jobs built. Batch operations for scripting is executed through JCL



TimeFinder Foundations

INTEGRATED IMPLEMENTATIONS



© 2006 EMC Corporation. All rights reserved.


TimeFinder Foundations - 56 

Now let's explore how TimeFinder can be integrated as an enterprise solution.



TimeFinder/Clone: TimeFinder/Mirror Emulation Mode Capability

- Leverage existing TimeFinder/Mirror commands with TimeFinder/Clone
- Support for all current TimeFinder/Clone capabilities in open systems environments
 - Symmetrix mirror position not required
 - Immediately read- and write-able
 - Up to 8 copies of a source
 - NOCOPY option
 - TimeFinder/Consistency Group support
- Simplifies conversion to RAID 5-protected BCVs
- Faster introduction of TimeFinder/Clone into existing TimeFinder/Mirror environments



** Open systems environments*

© 2006 EMC Corporation. All rights reserved.

TimeFinder Foundations - 57 

TimeFinder/Clone includes a “TimeFinder/Mirror Emulation Mode” feature for open systems environments that enable users with existing TimeFinder/Mirror scripts to leverage those exact same scripts with TimeFinder/Clone.


This capability enables you to deploy TimeFinder/Clone more quickly and easily into existing TimeFinder/Mirror environments by leveraging the same scripts.

For instance, if a user is using unprotected or mirrored TimeFinder/Mirror BCVs and wants to convert to RAID 5 protected clone BCVs, simply point to the TimeFinder/Clone application rather than TimeFinder/Mirror. The script remains the same while TimeFinder/Clone simply does a command-to-command conversion during the process.

TimeFinder/Clone’s Emulation Mode supports conversion of TimeFinder/Mirror scripts when previously used in either open systems or mainframe environments.

TimeFinder Integration with SRDF

- Best of both worlds—remote disaster restart and creating time
- Use SRDF Family-based solution to mirror between Symmetrix
- Create copies on the remote Symmetrix's SRDF Family-based target volume
- Run processes against remote snapshots
 - Data warehouse refreshes
 - Decision support
 - Application testing
 - Software updates




Using local replication technology within the same Symmetrix system, TimeFinder allows the creation of independently addressable copies of active production volumes for mainframe and open systems information storage. These copies can be used to run simultaneous tasks in parallel with one another, such as data warehouse loads, decision support systems, application testing, third-party software updates, and data backups.

You can gain the same benefits of TimeFinder at a remote location that is part of an SRDF family-based process as well. You can now perform many of those same tasks at the remote location, greatly reducing the amount of activity at the production site.

Using TimeFinder and SRDF family-based solutions together allows you to maintain a complete disaster restart program, while making productive use of the remote Symmetrix and the data it is protecting.

TimeFinder/Exchange Integration Module

- Automates the creation of full volume copies and snaps in an Exchange environment for backup and restore
 - Multiple Exchange servers; one backup server configuration
 - Hot splits for Exchange
 - Supports MSCS (Microsoft Cluster Service)
 - SRDF Family support
- Maximize collaboration and communication
 - Full-volume copy or snapshot
 - BCVs can be used to reduce downtime caused by Exchange restore and recovery




The TimeFinder/Exchange Integration Module reduces the downtime associated with Exchange restore and recovery activities. TimeFinder/EIM automates the process of shutting down Exchange properly, splitting of BCVs, restarting Exchange, and running a consistency check on the BCVs. It is highly recommended that the BCVs then be backed up to tape. TimeFinder/EIM is ideally suited for configuration where one backup server is supporting multiple production Exchange servers and can be used in an MSCS (Microsoft Cluster Service) environment.

The BCV data can then be used to do an Information Store, Directory, or Single Mailbox Recovery. By not having to go back to a copy on tape, companies can get their mission-critical Exchange servers up and running faster.

TimeFinder/SQL Server Integration Module

- Backups of SQL 2000 via Virtual Device Interface (VDI) database snapshot functionality with no impact on the production server
- Single or multiple database backups/restores in one command
- Supports SQL filegroup backup and restore
- Enables rapid restore
- Standby database creation is near instantaneous
- Data can be used for decision support or database maintenance



TimeFinder/SIM performs full database and filegroup backups using SQL's virtual device interface. The snapshot backup is made online with no disruptions to the production database. The backup copies are then available for a variety of operations such as tape backup and presentation to other servers for decision support, testing, and database maintenance.

TimeFinder/SIM backups are also commonly used to create Standby SQL Servers. Once initialized, they can be kept current by applying periodic transaction log backups. TimeFinder/SIM provides a robust alternative to the traditional log shipping method of creating and maintaining a standby server, using the SAN and eliminating network traffic.

Rapid recovery is probably the primary benefit. As with other TimeFinder family integrated solutions, recovery from the BCV tremendously reduces that window.

TimeFinder/SIM supports the SRDF family for the creation of remote copies to be used in the event of a disaster.

TimeFinder Foundations

BUSINESS JUSTIFICATION

This final section describes a typical business justification for the TimeFinder family.




As EMC's customers work to support their overall business, they have to be particularly focused on information storage and management. The fact that information continues to grow at a robust rate, shows there's still dramatic growth in data and in our customers' information in general.

There are new requirements around compliance; new levels of protection and recovery. It's not just about backup - it's actually how customers can restore and get back into business at the different service levels based on the application need.

All of this needs to be done at a lower cost. Every customer continues to be asked to do either more with less or more with the same. This is not going to change. You will continue to need to extract the most value out of every investment you make. You're going to look for fast payback periods. You're going to look for a measurable and demonstrable return on investment. The right remote replication solution will help you overcome these challenges.

Local Replication Benefits

- Increase application availability while reducing downtime
 - Improve Recovery Time Objective (RTO)
 - Improve Recovery Point Objective (RPO)
- Provide near-instant recovery
- Reduce backup windows
 - Minimize/eliminate impact on application




Local replication can significantly enhance current business and technical operations by providing access points to production data for parallel processing activities like backups, disk-based recovery after logical corruptions, and creating test environments for faster application time-to-revenue.

The Recovery Point Objective, RPO, is the point in time to which systems and data must be recovered after an outage, for example, the end of the previous day's processing. RPOs are often used as the basis for the development of replication and backup strategies and as a determinant of the amount of data that may need to be recreated after the systems or functions have been recovered.


Recovery Time Objective, RTO, is the period of time within which systems, applications, and functions must be recovered after an outage, for example, one business day. RTOs are often used as the basis for the development of recovery strategies, and as a determinant as to whether or not to implement the recovery strategies during a disaster situation. Similar terms include Maximum Allowable Downtime.

Local Replication Benefits (continued)

- Enable parallel processing for:
 - Data warehouse refreshes
 - Decision support
 - Database recovery “checkpoints”
 - Application development and testing



Every business strives to increase productivity and utilization of its most important resource—information. The information asset is key to finding the right customers, building the right products, and offering the best service. The greater the extent that this corporate information can be shared, re-used, and exploited, the greater competitive advantage a company can gain.




Once the key considerations are fully understood, it is time to weigh them against the business and technical considerations of the proposed projects. In most cases, it will not be a one-size-fits-all proposition. Each project may require you to balance your business and technical considerations to achieve the appropriate service level that you desire for a given project.

It is a careful balancing act that requires the ability to choose from more than one product in order to meet mixed service levels. Fortunately, the TimeFinder family offers you unparalleled choice and flexibility.


TimeFinder Technical Positioning

- When to use
 - TimeFinder Clone
 - TimeFinder SNAP
 - TimeFinder Mirror

Depending on the customer's environment, replication needs, and performance requirements, TimeFinder Clone, TimeFinder SNAP, and TimeFinder Mirror complement one another to insure maximum replication coverage for the entire data processing environment.




Today, the complexity and time critical aspects of a typical customer environment present an unprecedented data replication challenge. Thanks to the feature rich capability of the TimeFinder family of products, a solution to these challenges is at hand. For instance, customers who require RAID 5 protected copies, mainframe dataset level copies, or are required to work within a specific number of available Symmetrix mirror positions, the TimeFinder/Clone offering would be an excellent solution. For environments that require an advanced combination of established mirrors and snapshot copies, ultra-high performance BCVs, and ultra-high BCV availability, the TimeFinder/Mirror offering would be the right choice.



For those environments that require a more economical solution, fast, easily accessible, short term copies, and extremely small recovery point objectives without the cost of full copies, TimeFinder/SNAP proves beneficial. But for those situations where advanced replication is a requirement, including such technologies as SRDF/AR or Remote SRVC, full copies for high performance and availability, mainframe dataset level copies, and ultra-high BCV availability, the integration of TimeFinder/Clone and TimeFinder/Mirror would be appropriate.

TimeFinder – Benefits Review

- Repurposing
- Non-disruptive incremental resynchronization
- Consistent methodology for all platforms
- Maximize business operations by enabling parallel processing



Non-disruptively copies of an active application volume, or set of volumes, within the same Symmetrix

Please take a moment to review the benefits of this very important business tool.

Module Summary

Key points covered in this module:

- The TimeFinder family of products consists of:
 - TimeFinder/Mirror - Monitor and Control using BCVs
 - TimeFinder/Snap - Snap copy sessions
 - TimeFinder/Clone - Clone copy sessions
 - TimeFinder/CG - Consistency groups
 - TimeFinder Exchange Integration Option
 - TimeFinder SQL Integration Option
- TimeFinder functionality offers solutions to a vast array of business continuance scenarios
- EMC ControlCenter and Solutions Enabler are key tools used to manage a TimeFinder implementation
- Several examples of how TimeFinder can be integrated into a Symmetrix environment

These are the key points covered in this training. Please take a moment to review them.