



# EMC<sup>®</sup> Symmetrix<sup>®</sup> TimeFinder<sup>®</sup> for VMAX<sup>™</sup> 40K, VMAX<sup>™</sup> 20K/VMAX<sup>™</sup> Series

## Product Guide

P/N 300-011-719

REV A04

Copyright © 2010- 2012 EMC Corporation. All rights reserved. Published in the USA.

Published May, 2012

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC<sup>2</sup>, EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to the technical documentation and advisories section on the EMC online support website.

# CONTENTS

<b>Preface</b>	
<b>Chapter 1</b>	<b>TimeFinder Overview</b>
	TimeFinder family of products overview..... 14
	TimeFinder, Symmetrix, and Engenuity concepts ..... 15
	Engenuity..... 15
	Track..... 15
	Symmetrix cache..... 15
	Symmetrix session ..... 16
	Symmetrix device (logical volume) ..... 16
	TimeFinder source and target devices..... 17
	Open systems metadevices..... 17
	Virtual Provisioning..... 18
<b>Chapter 2</b>	<b>TimeFinder/Clone Operations</b>
	TimeFinder/Clone overview ..... 22
	Session types ..... 23
	Protection session ..... 23
	SDDF session ..... 24
	Differential TimeFinder/Clone session ..... 25
	Full-device TimeFinder/Clone session life cycle ..... 25
	Copy options..... 26
	Precopy..... 27
	Copy ..... 27
	Nocopy ..... 28
	I/O optimization ..... 28
	Asynchronous Copy on First Write..... 28
	Avoid Copy on Full Track Write..... 29
	QOS ..... 29
	Restore operations..... 29
	Multi-device operations ..... 30
	TimeFinder/CG ..... 30
	Cascaded TimeFinder/Clone..... 31
	Extent-level TimeFinder/Clone..... 31
	TimeFinder/Clone examples..... 32
	TimeFinder/Clone sessions with the copy option..... 32
	TimeFinder/Clone session with the nocopy option..... 33
	TimeFinder/Clone differential session ..... 33
	TimeFinder/Clone with Virtual Provisioning ..... 35
<b>Chapter 3</b>	<b>TimeFinder/Snap Operations</b>
	TimeFinder/Snap overview..... 40
	TimeFinder/Snap concepts ..... 40
	Virtual device (VDEV)..... 40
	SAVE device ..... 41
	Virtual sessions ..... 41
	Copy modes ..... 41

	Monitoring SAVE device usage .....	42
	Operational overview .....	42
	Create - identifying the source and the virtual device.....	42
	Activate - create a snapshot .....	43
	Recreate - reuse a virtual device .....	46
	Restore data from virtual devices .....	47
	Incrementally restoring back to the source device .....	47
	Fully restoring to anywhere.....	49
	Terminating a virtual copy session .....	50
	Copying a virtual device to another virtual device (duplicate snap) .....	50
<b>Chapter 4</b>	<b>TimeFinder VP Snap Operations</b>	
	TimeFinder VP Snap overview.....	52
	VP Snap operations.....	52
	Restore operations.....	53
	VP Snap example .....	53
<b>Chapter 5</b>	<b>TimeFinder Integration</b>	
	TimeFinder and SRDF .....	56
	R1 and R2 devices in TimeFinder operations.....	57
	SRDF/S and TimeFinder operations .....	58
	SRDF/A and TimeFinder.....	60
	SRDF adaptive copy operations and TimeFinder.....	61
	TimeFinder and Virtual Provisioning .....	64
<b>Appendix A</b>	<b>TimeFinder/Mirror</b>	
	TimeFinder/Mirror.....	68
	Standard and BCV devices .....	68
	TimeFinder/Mirror operations.....	68
	Clone Emulation mode .....	69

# FIGURES

	<b>Title</b>	<b>Page</b>
1	Data tracks on physical devices (drives) and in cache .....	16
2	Striped metadevices in TimeFinder/Clone operations.....	18
3	Concatenated metadevices in TimeFinder/Clone operations.....	18
4	Host write I/O to thin devices .....	19
5	Target device keeping point-in-time copies .....	22
6	Source device and target device used simultaneously by different hosts.....	23
7	TimeFinder/Clone protection sessions .....	24
8	SDDF session tracking changes.....	24
9	TimeFinder/Clone session.....	25
10	Clone from clone target (both sessions are cascaded clone).....	31
11	Full device clone sessions using copy and copy with precopy.....	32
12	Device clone session using the nocopy option .....	33
13	Full device clone session followed by an incremental clone copy.....	34
14	TimeFinder/Clone with Virtual Provisioning: 9:00 session before host write I/O.....	35
15	TimeFinder/Clone with Virtual Provisioning: 9:00 session after host write I/O .....	36
16	TimeFinder/Clone with Virtual Provisioning: 10:00 session before host write I/O.....	37
17	TimeFinder/Clone with Virtual Provisioning: 10:00 session after host write I/O .....	38
18	Protection bits and pointers .....	43
19	Protection bits after write I/Os to the source device.....	44
20	Pointers after data has been copied to the SAVE devices.....	45
21	Protection bits and pointers after write I/Os to the virtual device.....	46
22	Incrementally restoring back to the source device .....	48
23	Fully restoring to a target device .....	49
24	VP Snap session .....	53
25	SRDF basic configurations.....	56
26	R1 and R2 devices participating in TimeFinder operations. ....	58
27	Simultaneous TimeFinder/Clone with Enginuity version 5875 .....	59
28	SRDF/A device-level pacing activated for R1 volumes .....	60
29	SRDF/AR single-hop data flow .....	62
30	SRDF/AR multi-hop data flow .....	62
31	Standard device with an established BCV mirror .....	69



# TABLES

	<b>Title</b>	<b>Page</b>
1	Thin devices and full-device TimeFinder/Clone.....	64
2	Thin devices and TimeFinder/Snap.....	64
3	Thin SRDF devices and TimeFinder/Clone.....	65
4	Thin SRDF devices and TimeFinder/Snap.....	65





# PREFACE

*As part of an effort to improve its product lines, EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.*

*Contact your EMC representative if a product does not function properly or does not function as described in this document.*

---

**Note:** This document was accurate at publication time. New versions of this document might be released on the EMC online support website. Check the EMC online support website to ensure that you are using the latest version of this document.

---

## Audience

Users requiring basic knowledge of TimeFinder concepts and operation.

## Related documentation

The following EMC publications provide additional information:

- ◆ *EMC Solutions Enabler TimeFinder Family CLI Product Guide*
- ◆ *EMC Mainframe Enablers TimeFinder/Clone Mainframe SNAP Facility Product Guide*
- ◆ *EMC Mainframe Enablers TimeFinder/Mirror Product Guide*
- ◆ *EMC Symmetrix Remote Data Facility (SRDF) for VMAX 40K, VMAX 20K/VMAX, DMX Series Product Guide*
- ◆ EMC Support Matrix

## Typographical conventions

EMC uses the following type style conventions in this document:

<b>Normal</b>	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> <li>Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus</li> <li>Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, functions, and utilities</li> <li>URLs, pathnames, filenames, directory names, computer names, links, groups, service keys, file systems, and notifications</li> </ul>
<b>Bold</b>	Used in running (nonprocedural) text for names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system calls, and man pages  Used in procedures for: <ul style="list-style-type: none"> <li>Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus</li> <li>What the user specifically selects, clicks, presses, or types</li> </ul>
<i>Italic</i>	Used in all text (including procedures) for: <ul style="list-style-type: none"> <li>Full titles of publications referenced in text</li> <li>Emphasis, for example, a new term</li> <li>Variables</li> </ul>
<code>Courier</code>	Used for: <ul style="list-style-type: none"> <li>System output, such as an error message or script</li> <li>URLs, complete paths, filenames, prompts, and syntax when shown outside of running text</li> </ul>
<b>Courier bold</b>	Used for specific user input, such as commands
<i>Courier italic</i>	Used in procedures for: <ul style="list-style-type: none"> <li>Variables on the command line</li> <li>User input variables</li> </ul>
< >	Angle brackets enclose parameter or variable values supplied by the user
[ ]	Square brackets enclose optional values
	Vertical bar indicates alternate selections — the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

## Where to get help

For additional information on all Symmetrix-related publications, refer to EMC online support (<http://Powerlink.EMC.com>).

**Product information** — For documentation, release notes, software updates, or for information about EMC products, licensing, and service, refer to EMC online support.

**Technical support** — For technical support, go to EMC online support. On the Support page, you will see several options, including one for making a service request. Note that to open a service request, you must have a valid support agreement. Please contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

## Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to:

`techpubcomments@emc.com`



# CHAPTER 1

## TimeFinder Overview

This chapter introduces the TimeFinder family of local replication software and describes TimeFinder, Symmetrix, and Engenuity concepts. Topics include:

- ◆ TimeFinder family of products overview ..... 14
- ◆ TimeFinder, Symmetrix, and Engenuity concepts ..... 15

## TimeFinder family of products overview

This document describes the EMC® TimeFinder® family of products for VMAX™ 20K/VMAX™ and VMAX™ 40K arrays.<sup>1</sup>

The TimeFinder family of products are Symmetrix local replication solutions designed to nondisruptively create point-in-time copies of critical data. You can configure backup sessions, initiate copies, and terminate TimeFinder operations from mainframe and open systems controlling hosts using EMC Symmetrix host-based control software.

The TimeFinder local replication solutions include TimeFinder/Clone, TimeFinder/Snap, and TimeFinder VP Snap. TimeFinder/Clone creates full-device and extent-level point-in-time copies. TimeFinder/Snap creates pointer-based logical copies that consume less storage space on physical drives. TimeFinder VP Snap provides the efficiency of Snap technology with improved cache utilization and simplified pool management.

Each solution guarantees high data availability. The source device is always available to production applications. The target device becomes read/write enabled as soon as you initiate the point-in-time copy. Host applications can therefore immediately access the point-in-time image of critical data from the target device while TimeFinder copies data in the background.

TimeFinder/Clone is suitable if:

- ◆ Full-volume copies are intended for recovery scenarios
- ◆ Full-volume or extent-level point-in-time copies of production data have to be immediately available to applications for activities such as reporting and testing
- ◆ The majority of data on the production volumes changes between subsequent backup sessions
- ◆ Multiple copies of production data are needed, and you want to reduce disk contention and improve data access speed to the production data.

TimeFinder/Snap is suitable if:

- ◆ Only a fraction of data on the production volumes changes between subsequent backup sessions
- ◆ Only a fraction of data on the production volumes frequently changes during the peak I/O activity window when multiple point-in-time copies are required.

TimeFinder VP Snap is suitable if:

- ◆ You want to create space-efficient snaps for thin devices
- ◆ You want multiple sessions to share capacity allocations within a thin pool, thus reducing the storage required for saved tracks.

---

1. Native TimeFinder/Mirror is no longer supported on Symmetrix VMAX Series arrays and Enginity levels 5874 and higher. If you want to use TimeFinder/Mirror scripts with Symmetrix VMAX Series arrays, you can do so by running TimeFinder/Clone in emulation mode. Note that emulation mode uses TimeFinder/Clone and not the TimeFinder/Mirror infrastructure. “Clone Emulation mode” on [page 69](#) provides more details.

TimeFinder includes the following features:

- ◆ Supports RAID 1, RAID 5, RAID 6, and RAID 10 protection schemes.
- ◆ Provides restore capabilities.
- ◆ Provides incremental re-synchronization between the source and the target.
- ◆ Supports Virtual Provisioning.
- ◆ Tightly integrated with SRDF.

## TimeFinder, Symmetrix, and Enginuity concepts

This section introduces TimeFinder, Symmetrix, and Enginuity™ concepts.

### Enginuity

Enginuity is the operating environment for all Symmetrix arrays. Enginuity provides the following services:

- ◆ Manages system resources to optimize performance across a wide range of I/O profiles
- ◆ Manages and monitors internal data flow through Symmetrix hardware components
- ◆ Ensures data integrity and data availability through advanced fault monitoring, detection, and correction capabilities
- ◆ Provides concurrent maintenance and serviceability
- ◆ Offers the foundation for disaster recovery and business continuity solutions such as TimeFinder and Symmetrix Remote Data Facility (SRDF)
- ◆ Provides uniform access through APIs for internal calls and offers an external interface to allow integration with other software providers.

### Track

Enginuity manages data within Symmetrix arrays by breaking the data into tracks. A track is the basic data unit handled in Symmetrix arrays. The track size is 64KB on Symmetrix VMAX Series arrays running Enginuity levels 5874 and higher. It varies among Symmetrix hardware models and Enginuity versions.

Enginuity is capable of handling smaller data units, but for most practical purposes, tracks are the smallest data chunks manipulated in Symmetrix arrays.

### Symmetrix cache

The Symmetrix cache is at the heart of Symmetrix I/O operations. As data moves within the Symmetrix array, between arrays, or between the array and the host applications, it always flows through the Symmetrix cache from the source to the target location.

In addition to keeping recently used data readily available for the next application's request, the Symmetrix cache contains I/O-specific metadata for each device in the array. Device-level and track-level metadata resides in cache tables and in dynamically allocated

cache slots. Recently used application data resides in ordinary cache data slots and on physical drives. In the event of a power outage, the power vault mechanism protects the cache.

TimeFinder uses dynamically allocated cache slots to allow special I/O handling during TimeFinder background operations when data is copied from the source to the target device.

## Symmetrix session

Engenuity uses the Symmetrix session mechanism to deliver Engenuity features without interrupting production applications. TimeFinder/Clone, TimeFinder/Snap, Open Systems Replicator, IBM FlashCopy, and SRDF/Star all use the session mechanism.

Symmetrix sessions have the following common characteristics:

- ◆ Users configure, initiate, poll, and terminate Symmetrix sessions via Symmetrix host-based control software that typically runs on a dedicated controlling host.
- ◆ Engenuity provides all session-specific functionality without requiring any resources from the application hosts.
- ◆ Engenuity limits the number of concurrent sessions per Symmetrix device to 16. Symmetrix host-based control software may impose additional restrictions.

---

**Note:** TimeFinder/Snap multi-virtual sessions allow up to 128 virtual point-in-time copies. [“Virtual sessions” on page 41](#) provides additional details. TimeFinder/Snap sessions on Symmetrix VMAX 40K systems always use multi-virtual mode.

---

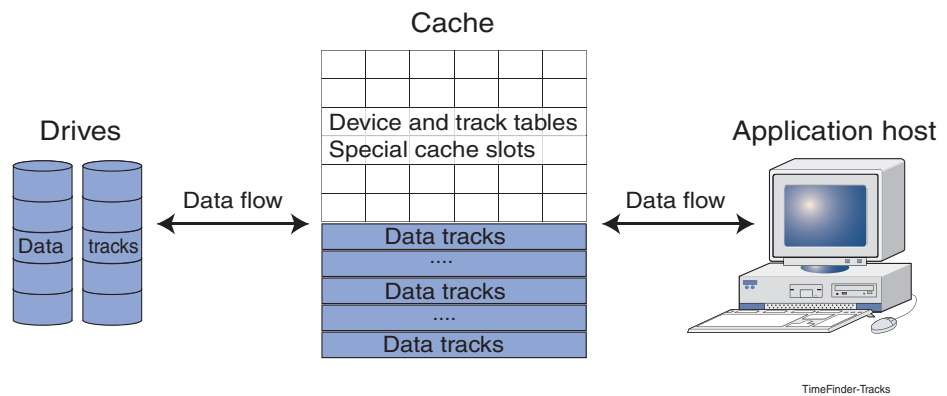


---

**Note:** TimeFinder VP Snap allows an additional 32 sessions per Symmetrix device.

---

Figure 1 on page 16 shows data flow through Symmetrix cache.



**Figure 1** Data tracks on physical devices (drives) and in cache

## Symmetrix device (logical volume)

A Symmetrix device (Symmetrix logical volume) is a host-addressable storage device in a Symmetrix array. Engenuity presents this storage capacity to the application hosts by assigning a unique Symmetrix device number to each host-addressable storage device.



Internally, Enginuity maintains a collection of data tracks, which are written to physical drives, and the corresponding metadata, which is written to physical drives and to the Symmetrix cache to manage data flow within a Symmetrix array.

## TimeFinder source and target devices

In TimeFinder operations, the Symmetrix device that keeps the original data is the source device and the device that keeps a point-in-time copy is the target device. TimeFinder supports the following device configuration options:

- ◆ The source and the target device can have different RAID schemes. TimeFinder supports all of the RAID schemes that are supported with Symmetrix systems.
- ◆ The target device can be larger than the source device. In this case, you cannot restore data from the target to the original source device.
- ◆ Metadevices are supported. Both the source device and the target device must be metadevices.

[“Open systems metadevices” on page 17](#) provides greater detail.

- ◆ The source, the target device, or both can be thin devices.

[“Thin devices and data devices” on page 18](#) and [“TimeFinder and Virtual Provisioning” on page 64](#) provide greater detail.

- ◆ TimeFinder source and target devices must have the same device emulation type (both FBA or both CKD).

## Open systems metadevices

Open systems metadevices are comprised of multiple Symmetrix devices, the meta head device, a number of meta members, and the meta tail device. The application host addresses a metadevice through the meta head device and perceives the Symmetrix devices grouped in the metadevice as a single addressable storage device. Metadevices are used to create storage devices larger than the maximum Symmetrix logical device size and to configure open systems RAID schemes similar to mainframe RAID 10 schemes.

Open systems metadevices can be used in TimeFinder operations, provided that they meet the following requirements:

- ◆ In TimeFinder/Clone, both the source and the target device must be metadevices of the same size, with the same number of meta members.
- ◆ In TimeFinder/Snap, both the source and the virtual device must be metadevices of the same size, with the same number of meta members.
- ◆ If you use striped metadevices, the source and the target device in TimeFinder/Clone operations or the source and the virtual device in TimeFinder/Snap operations must have the same stripe size.

Figure 2 on page 18 shows striped metadevices configured as TimeFinder/Clone device pairs.

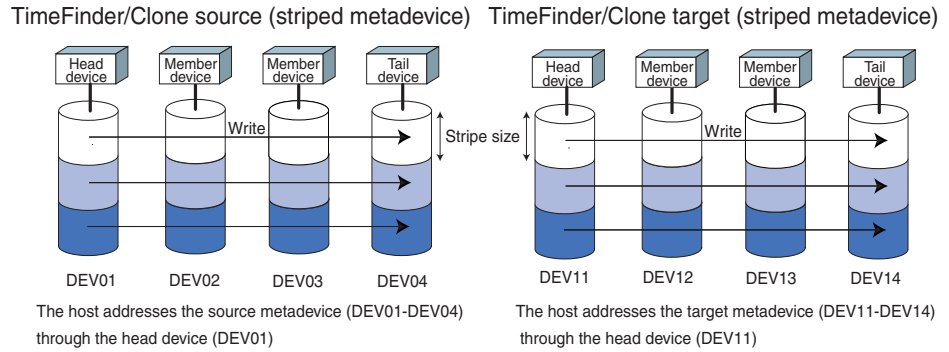


Figure 2 Striped metadevices in TimeFinder/Clone operations

Figure 3 on page 18 shows concatenated metadevices configured as TimeFinder/Clone device pairs.

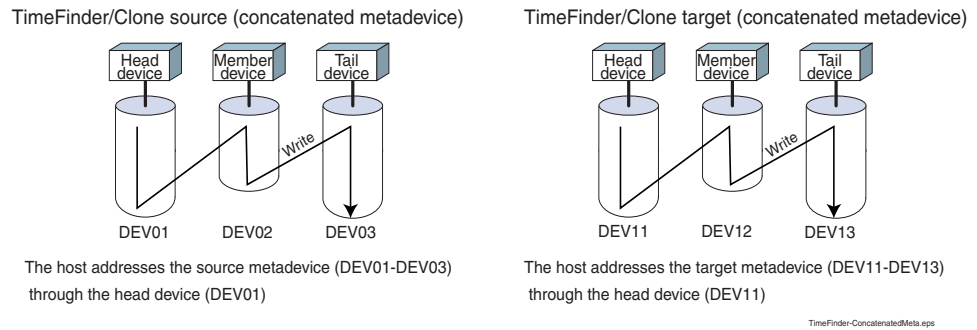


Figure 3 Concatenated metadevices in TimeFinder/Clone operations

## Virtual Provisioning

Symmetrix Virtual Provisioning™ is an Engenuity feature that allows you to create Symmetrix devices of a specified size without requiring that the entire specified physical capacity be available when you create them.

In a virtually provisioned storage environment, Engenuity allocates physical storage space (12 tracks at a time) on-demand for any host write to previously unallocated space. This functionality allows you to create devices with host-reported capacities far exceeding the actual physical capacities available in the Symmetrix array.

Virtual Provisioning supports the FBA and CKD device emulations, and requires special Symmetrix device types discussed in the following sections.

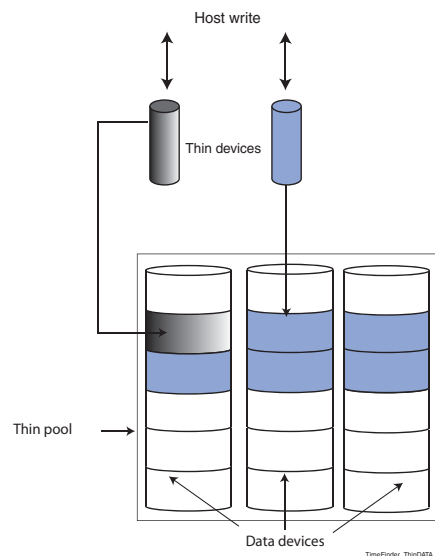
### Thin devices and data devices

A virtually provisioned device known as a thin device is a special Symmetrix device that has a minimal amount of storage space associated with it until the first write I/O to that device occurs. Engenuity allocates only as much physical storage capacity as is required to service the host write request.

A thin device must be associated (bound) to a pool of devices (known as data devices) that have RAID attributes associated with them. Data devices are similar to the SAVE devices used in TimeFinder/Snap operations.

When the host issues the first write request to a thin device, Enginuity allocates physical storage space from a pool of data devices bound to the thin device, and then writes data to the data devices.

Figure 4 on page 19 shows write I/O operations to thin devices.



**Figure 4** Host write I/O to thin devices

Unlike thin devices, data devices cannot be mapped to the host. They can be accessed only through the associated thin devices. In addition, data devices have physical storage and RAID attributes, while thin devices do not.

## Fully-provisioned (thick) device

In contrast to thin devices, “thick” devices are ordinary Symmetrix devices whose storage capacity is fully allocated when devices are created. Thick devices are also known as fully-provisioned or standard devices.

With Enginuity version 5874 and higher, TimeFinder supports thick-to-thin and thin-to-thick source-target device pairs, which can be used for local thick-to-thin migration. [Table 1, “Thin devices and full-device TimeFinder/Clone,” on page 64](#) provides greater detail.



# CHAPTER 2

## TimeFinder/Clone Operations

This chapter provides a high-level overview of TimeFinder/Clone operations. Topics include:

- ◆ TimeFinder/Clone overview ..... 22
- ◆ Session types ..... 23
- ◆ Full-device TimeFinder/Clone session life cycle ..... 25
- ◆ Copy options..... 26
- ◆ I/O optimization ..... 28
- ◆ Restore operations..... 29
- ◆ Extent-level TimeFinder/Clone..... 31
- ◆ TimeFinder/Clone examples..... 32

## TimeFinder/Clone overview

Full-device TimeFinder/Clone is a licensed Symmetrix local replication solution that creates full-device 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 data. In mainframe host environments, TimeFinder/Clone also supports extent-based point-in-time copies in addition to full-device copies.

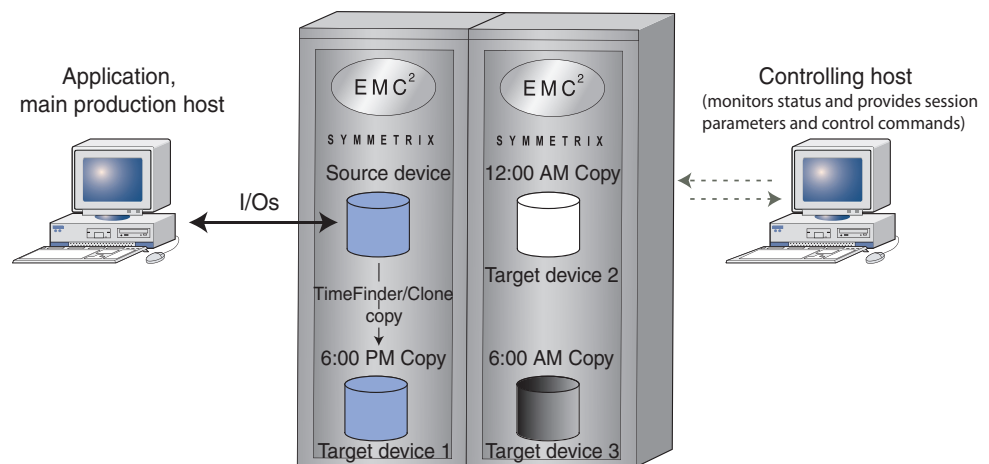
[“Extent-level TimeFinder/Clone” on page 31](#) describes extent-level specifics.

TimeFinder/Clone allows you to make copies of data simultaneously on multiple target devices from a single source device. The data is copied from the source device, creating a physical backup copy called a clone. The data is available to a target device’s host immediately upon activation of the TimeFinder copy session, even if the copy process has not completed.

TimeFinder/Clone copies are appropriate in situations where multiple copies of production data are needed for recovery, testing, backup, or report generation. Using clone copies can reduce disk contention and improve data access speed by assigning users to clones rather than to the production data.

TimeFinder/Clone does not require that the source device or the target device (see [“TimeFinder source and target devices” on page 17](#)) be mapped to an application host. The source device is typically mapped to the production host, but the target device may never be mapped to a host (depending on the intended use of the point-in-time copy).

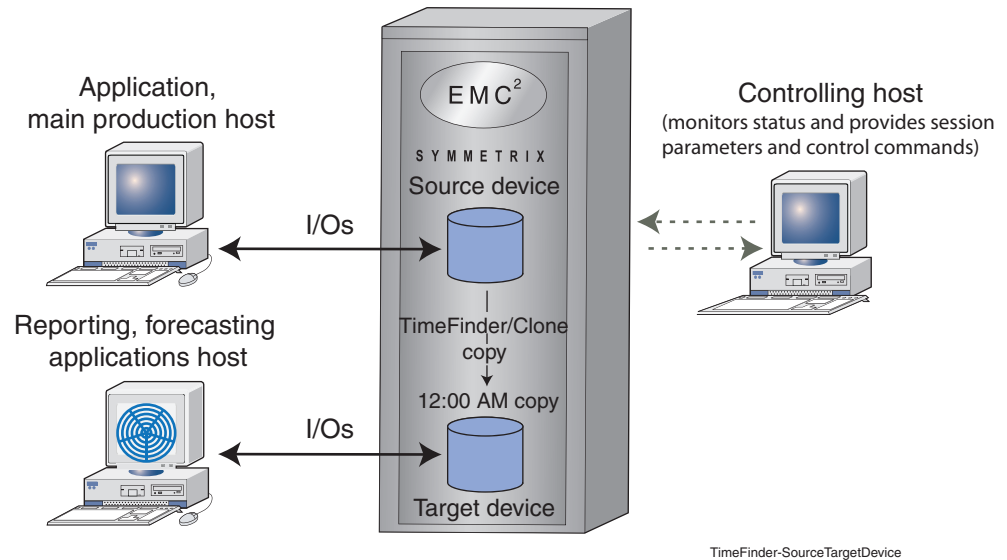
[Figure 5 on page 22](#) shows a TimeFinder/Clone session creating the 6:00 PM copy of the source device on a dedicated target device. In this example, TimeFinder/Clone creates three point-in-time copies of the source device to three different target devices every day. Target devices are not mapped to any host to preserve the point-in-time data for potential recovery operations.



**Figure 5** Target device keeping point-in-time copies

In another example, the reporting and forecasting applications require access to the 12:00 AM copy of the source device data while the source device keeps servicing I/Os from the main production application.

Figure 6 on page 23 shows a TimeFinder/Clone session copying the 12:00 AM data from the source device to the target device, which is mapped to the host running reporting and forecasting applications. Although the TimeFinder/Clone copy session is still copying data, the main production host and the reporting host are unaware of the background copy process. The production host continues without contention while the reporting host can access the 12:00 AM data on the target device as soon as the 12:00 AM copy session begins.



**Figure 6** Source device and target device used simultaneously by different hosts

## Session types

TimeFinder/Clone creates a protection session and, optionally, a Symmetrix Differential Data Facility (SDDF) session for the source device and the target device. Protection and SDDF sessions are the tools that TimeFinder/Clone uses to monitor copy operations.

## Protection session

Protection sessions maintain a protection bit for every source device and every target device track. Engenuity sets protection bits when you create the TimeFinder/Clone session and maintains them until you terminate the session. Protection bits are stored in cache metadata tables.

A protection bit set for a source device track indicates that Engenuity has yet to copy the track to the target device. Similarly, a protection bit set for a target device track indicates that the track has not yet received data from the source device. When Engenuity copies a track from the source device to the target device, it clears the corresponding protection bits.

Figure 7 on page 24 shows an example of TimeFinder/Clone protection sessions.

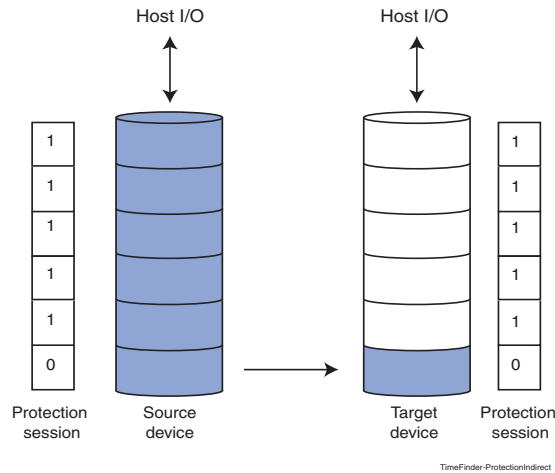


Figure 7 TimeFinder/Clone protection sessions

### SDDF session

A Symmetrix Differential Data Facility (SDDF) session is a Symmetrix session used to monitor changes on a Symmetrix device. An SDDF session maintains its own bitmap with a bit for every device track.

When you initiate the SDDF session, Engenuity initially sets all of the SDDF bits, then clears an SDDF bit when the corresponding track changes due to a host write I/O. Any subsequent host write I/O to the same track has no impact on the SDDF bit.

**Note:** The TimeFinder/Clone copy process does not impact the SDDF bits.

Figure 8 on page 24 illustrates an SDDF session tracking changes on a device over a period of time.

- 1) SDDF session initiated at 12:00 AM
- 2) SDDF session at 6:00 PM




Figure 8 SDDF session tracking changes



## Differential TimeFinder/Clone session


A TimeFinder/Clone session that uses protection and SDDF sessions is known as a differential session. When you create a differential TimeFinder/Clone session, you create two protection sessions and two SDDF sessions, one of each for the source device and the target device.

When you initiate the differential TimeFinder/Clone session, the SDDF sessions start tracking changes on the source device and the target device. When you initiate the next point-in-time copy of the source device, Enginuity looks up the SDDF protection bitmap to determine which tracks have changed on either the source device or the target device since the previous TimeFinder/Clone session. Enginuity then marks the tracks to be copied in both protection session bitmaps.

## Full-device TimeFinder/Clone session life cycle

TimeFinder/Clone allows you to fully control its operations through session management commands from the controlling host that runs Symmetrix host-based control software.

[Figure 9 on page 25](#) illustrates a TimeFinder/Clone session. In this example, the controlling host is a dedicated system (as is typically the case). The main application host is running I/Os to the source device while the user-controlled TimeFinder/Clone copy session is creating a point-in-time copy of the source device on the target device.



**Figure 9** TimeFinder/Clone session

The following steps describe how to control a full-device TimeFinder/Clone session. [“Extent-level TimeFinder/Clone” on page 31](#) explains the differences that apply to extent-level TimeFinder/Clone operations.

1. Create a clone copy session by specifying the source device and the target device and set session parameters, such as copy options and session type. [“Multi-device operations” on page 30](#) and [“Copy options” on page 26](#) provide more information.

Once created, the TimeFinder/Clone session is inactive until you activate the session in [step 2](#). When the session is inactive, the target device is not available to its host. You control the duration of the inactive phase by choosing the activation time.

2. Initiate (activate) the point-in-time copy. The TimeFinder/Clone session transitions from the inactive phase to the active phase.
3. Optionally, query for the status of the TimeFinder/Clone copy process.
4. When data from the source device is copied to the target device, proceed as follows:
  - If the session you created in step 1 is not differential, you can remove/terminate it. TimeFinder/Clone provides the option to automatically terminate the session once the copy process completes.
  - If the session you created in step 1 is differential, initiate a subsequent point-in-time copy at a later time. This TimeFinder/Clone session will copy only data that has changed on the source device or the target device since the previous TimeFinder/Clone session. You can initiate subsequent differential TimeFinder/Clone sessions as long as you keep the SDDF sessions running. [“SDDF session” on page 24](#) and [“Differential TimeFinder/Clone session” on page 25](#) provide more details.

[“TimeFinder/Clone examples” on page 32](#) shows TimeFinder/Clone full-device sessions using different copy options.

## Copy options

This section explains the mechanisms that TimeFinder/Clone uses to:

- ◆ Copy data from the source device to the target device
- ◆ Guarantee a consistent point-in-time image on the target device while the source device and the target device are processing host I/Os.

TimeFinder/Clone uses different copy options to make a point-in-time copy of source data available on the target device. Copy options determine:

- ◆ What kind of point-in-time copy is created on the target device
- ◆ When in the session’s life cycle the TimeFinder/Clone session copies data
- ◆ What triggers the TimeFinder/Clone copy process.

---

**Note:** Regardless of which copy option or I/O optimization you are using, a host read I/O to a protected source device track is not affected by TimeFinder operations. It is an ordinary read I/O. Engenuity does not copy the track to the target and does not change the protection bits.

---

## Precopy

The precopy option allows you to synchronize data from the source device to the target device during the inactive phase, that is before you initiate the point-in-time copy process. Once you activate the session, the precopy phase completes and the active copy phase starts.

When you copy data using the precopy option during the inactive phase, the copy process (also known as the background copy process) does not impact host I/Os since it does not require special I/O handling. As the source tracks are copied to the target device, TimeFinder/Clone clears the corresponding protection bits on the source device and the target device, and sets them again if a subsequent I/O arrives for the source device track. No I/O can arrive for the target device since the device is not available to the host during the inactive phase.

The source tracks copied to the target device that remain unchanged on the source device between the time they were copied and the time you activate the session do not have to be copied during the active phase. The active phase will then take less time to complete.

The precopy option works best when the I/O rate on the source device is low during the inactive phase so that not many tracks change. Note the following about using the precopy option:

- ◆ If you use the precopy option and the target device is mapped to a host, the target device is not available to the host until you activate the session.
- ◆ The precopy option implies that you allow a lapse of time for the precopy phase to be able to copy some data from the source device to the target device.

---

**Note:** TimeFinder/Clone emulation mode uses the precopy option by default. [“Clone Emulation mode” on page 69](#) provides more information. Extent level TimeFinder/Clone does not support the precopy option.

---

## Copy

When you use the copy option, the TimeFinder/Clone session copies all of the source device tracks to the target device. The copy process starts when you activate the session (initiate the point-in-time copy) and completes when all of the source device tracks are copied to the target device and no protected tracks are left.

You will typically use the copy option to create a full source device copy on the target device (gold copy) and keep data on the target device intact for recovery purposes. You can also use the copy option if you need to make the full-device point-in-time copy available to another host. The other host can then access the target device while the production host continues to access the source device.

With the copy option, you usually want to minimize the duration of the inactive phase. Typically, you create and immediately initiate the session so that the target device is promptly available to its host.

When you use the copy option, two mechanisms drive the copy process from the source device to the target device:

- ◆ The background process unique to the copy and precopy option that scans the source device tracks and copies them to the target device if the protection bit is set.
- ◆ The copy process triggered by host write I/Os to protected tracks on the source device or partial track write I/Os to protected tracks on the target device.

## Nocopy

The difference between the copy option and the nocopy option is that the nocopy option does not allow the background process. Like the copy option, the nocopy process starts when you activate the TimeFinder/Clone session and completes when no protected source tracks exist on the source device. However, with the nocopy option, the copy process copies data from the source device to the target device only if triggered by a host I/O.

A write I/O to a protected track on the source device or a partial track write I/O to the target device always results in full tracks being copied to the target device. In open systems environments, by default, a read I/O to a protected track on the target device results in data being copied to the target device (CopyOnAccess). This default behavior can be changed so that a read I/O to a protected track on the target device does not result in data being copied (CopyOnWrite). In mainframe environments, a read I/O to a target device never results in data being copied.

Unlike the copy option, the time required to copy protected source tracks to the target device depends entirely on the host I/O rate. When host I/Os cease, so does the TimeFinder/Clone copy process.

Because the copy process is entirely driven by host I/Os to either device, the nocopy option is suitable when the host I/O rate is high and you want to make a point-in-time image instantly available to the host accessing the target device.

Like with the copy option, you usually want to minimize the duration of the inactive phase when using the nocopy option.

## I/O optimization

This section describes the techniques that Enginuity uses to optimize I/O performance.

### Asynchronous Copy on First Write

Since the point-in-time copy of the source device has to be immediately available to the host accessing the target device and both the source device and the target device must be able to service host I/Os, TimeFinder must ensure that the protected source track is copied to the target device before the new write I/O is allowed to proceed.

If a host write I/O arrives to a protected source/target device track, Enginuity does not accept it until TimeFinder copies data from the source device to the target device. This mechanism is also known as Copy on First Write (COFW). Compared to an ordinary host write I/O, a write I/O to a protected track causes a delay in the host response time since the incoming write I/O is not accepted until the source track is copied to the target device.

Asynchronous Copy on First Write is a performance improvement mechanism that provides a faster host response time than COFW. When Asynchronous Copy on First Write is enabled, Enginuity accepts an incoming host write I/O to a protected source track and immediately returns a status to the host, but intercepts the write I/O in the Symmetrix cache. Enginuity maintains multiple versions of the same track. One version holds the new write data and the other holds the old source track data, which is also stored on the physical drive. The track version that contains the new write data cannot be written to the source device until the old track version is copied to the target device.

The main quality of the Asynchronous Copy on First Write is that the write intercept does not cause any delays in host response time. Because Enginuity maintains multiple track versions in cache, Asynchronous Copy on First Write requires more cache resources to handle write I/Os to protected source tracks.

Asynchronous Copy on First Write is a system-wide setting and is enabled by default. The Asynchronous Copy on First Write feature provides the most benefit in situations where the protected tracks are not in cache when the new write I/O arrives.

## Avoid Copy on Full Track Write

Avoid Copy on Full Track Write is a system wide setting. When enabled, if a write I/O arrives to a protected target track, Enginuity does not intercept it. It allows the write I/O and does not copy the track from the source device. Only when a partial track write I/O arrives to the target device does Enginuity intercept the I/O and copy the full track from the source device. After copying the full track, Enginuity accepts the partial track write I/O to the protected target track. This is necessary to guarantee that a consistent point-in-time image of the source device track is available on the target device before the host sends a partial track write I/O to the target device.

## QOS

QOS allows you to assign priorities to copy jobs to optimize Symmetrix cache utilization and prevent cache overflows. With Enginuity 5875 and higher, QOS supports request based copies. This feature optimizes copy performance and provides finer QOS control by allowing target devices to initiate the copy process from the source.

## Restore operations

Full-device TimeFinder/Clone supports restore operations. All concepts explained in [“Multi-device operations” on page 30](#) and [“Copy options” on page 26](#) apply to restore operations. Restore sessions are ordinary TimeFinder/Clone sessions where the roles of source devices and target devices are swapped. Extent-level TimeFinder/Clone does not support restore operations. [“Extent-level TimeFinder/Clone” on page 31](#) provides more details about differences that apply to extent-level TimeFinder/Clone.

Restore operations can be used to copy target data to another device (full restore), or back to the original source device (incremental restore). In the case of a full restore, the original session terminates and a copy session to the target of the restore starts. In the case of an incremental restore, the original session copy direction is reversed and changed data is copied from the target device to the source device. Restore operations require that the original session is differential or the source device is fully copied.

## Multi-device operations

EMC host-based control software allows you to send TimeFinder/Clone control commands to multiple devices and perform large-scale operations. Instead of issuing a control command to a single device pair, you issue the control command to multiple device pairs.

TimeFinder/Clone supports the following multi-device operations:

- ◆ Create
- ◆ Activate
- ◆ Restore (differential sessions only)
- ◆ Terminate

## TimeFinder/CG

TimeFinder/Consistency Group (TimeFinder/CG) is an option available with TimeFinder/Clone and TimeFinder/Snap product offerings. With Enginuity 5874 and higher, TimeFinder/Clone and TimeFinder/Snap both include the TimeFinder/CG option. You do not need a separate license for TimeFinder/CG.

TimeFinder/CG is based on the Enginuity Consistency Assist (ECA) feature that enables you to create a consistent point-in-time image across multiple devices. This is necessary for tasks such as backing up all of the devices that belong to a particular application or backing up multiple devices distributed across multiple sites.

When you use TimeFinder/Clone to create a consistent point-in-time copy across multiple devices, you need to ensure that the point-in-time copy starts simultaneously across a range of devices that participate in the process. To accomplish this, you need to block host I/Os across multiple devices when you activate multiple TimeFinder/Clone sessions. The time interval during which the host I/Os are blocked is also referred to as the ECA window.

To create a consistent point-in-time copy across multiple devices:

1. Create a consistency group including all of the devices that will participate in the TimeFinder/Clone copy process as source devices.
2. Create TimeFinder/Clone sessions for the devices that belong to the consistency group.
3. Activate multiple TimeFinder/Clone sessions. TimeFinder/CG guarantees that the activation process is consistent by blocking host I/Os to all the devices that belong to the consistency group. Until all TimeFinder/Clone sessions are activated, no I/O to any device in the consistency group is allowed.


For more information about using Symmetrix host-based control software, refer to:

- ◆ *EMC Solutions Enabler TimeFinder Family CLI Product Guide*
- ◆ *EMC Mainframe Enablers TimeFinder/Clone Mainframe SNAP Facility Product Guide*

## Cascaded TimeFinder/Clone

Full-device TimeFinder/Clone supports cascaded operations. The cascaded clone capability allows you to create two full-device TimeFinder/Clone sessions using the same device as a source device and as a target device. You can then run both sessions concurrently as long as you activated the copy process sequentially (to preserve the copy direction).

Figure 10 on page 31 shows how cascaded sessions are accepted from left to right. You can activate session 1 to copy from device A to device B. Then, you can activate session 2 to clone device B to device C. If your copy direction is  $A \Rightarrow B \Rightarrow C$ , you can activate session  $B \Rightarrow C$  only after you activated session 1 from  $A \Rightarrow B$ .



**Figure 10** Clone from clone target (both sessions are cascaded clone)

Cascaded clone on thin devices is supported in environments running Enginuity version 5875 and higher. The only exception is that if device B has been restored to device A, when A or C is a thin device, the creation of a new session  $A \Rightarrow C$  is not permitted.

Mixed thick and thin devices are not supported with cascaded operations. All devices in cascaded sessions must be thick or thin.

Restore to Target (RTT) allows you to perform an incremental restore to a cascaded clone target. For example, devices in an  $A \Rightarrow B \Rightarrow C$  cascaded clone session can copy data from device C to device A (via device B). RTT is supported in environments running Enginuity 5875 and higher.

Cascaded clone supports restore operations and all of the copy options listed in “Copy options” on page 26 and “Multi-device operations” on page 30.

Cascaded clone is supported with VP Snap with some restrictions. Refer to the *EMC Solutions Enabler TimeFinder Family CLI Product Guide* for additional details.

## Extent-level TimeFinder/Clone

TimeFinder/Clone supports extent-level operations in mainframe host environments. Extent-level TimeFinder/Clone is used to implement dataset-level snaps using the TimeFinder/Clone Mainframe Snap Facility.

An extent is a sequential set of tracks that contains all or part of a dataset. You can define multiple extents on the source device and use TimeFinder/Clone to replicate them to one or more target extents. The extents do not have to be at the same location on the target device as they were on the source device. TimeFinder/Clone can also replicate datasets that are broken into multiple extents that span devices. The “Performing a SNAP DATASET copy” section of the *EMC TimeFinder/Clone Mainframe Snap Facility Product Guide* provides additional details.

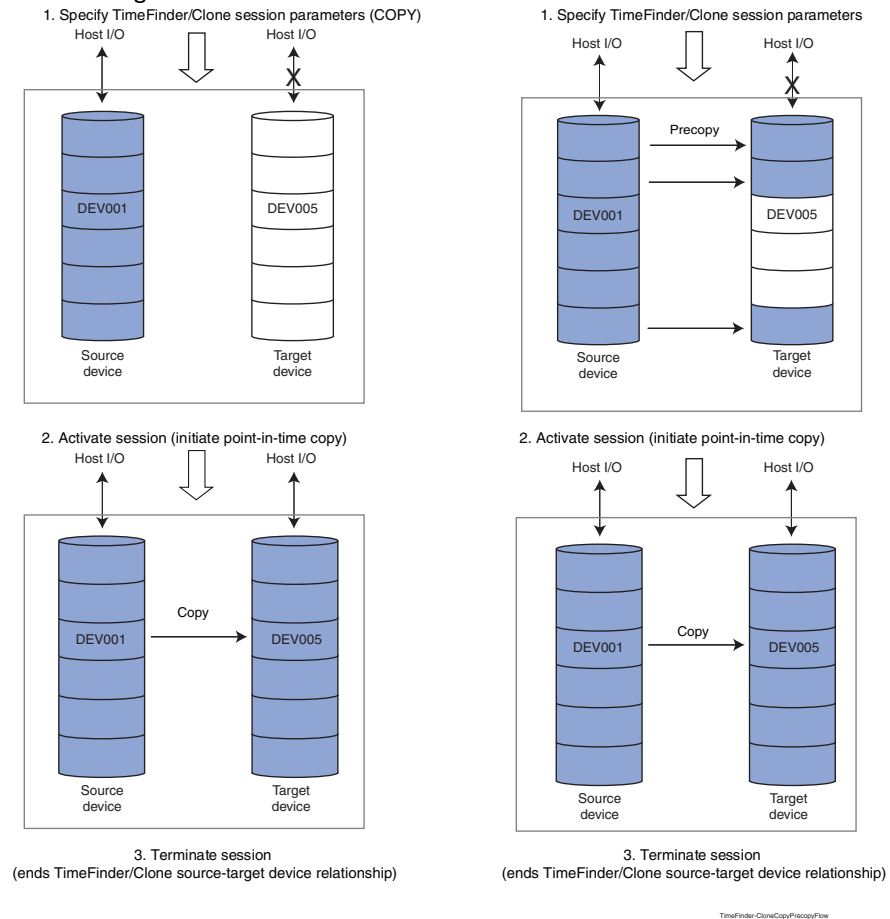
**Note:** Extent-level TimeFinder/Clone supports all copy options described in “Copy options” on page 26 except the precopy type. Extent-level TimeFinder/Clone does not support restore operations.

## TimeFinder/Clone examples

This section provides examples of several common TimeFinder/Clone operations.

### TimeFinder/Clone sessions with the copy option

Figure 11 on page 32 shows a TimeFinder/Clone session using the copy option (left) and another session using the copy option with precopy (right). Note that the precopy option allows TimeFinder to start copying data to the target device before the session is activated and the target device becomes available to the host.



**Figure 11** Full device clone sessions using copy and copy with precopy



## TimeFinder/Clone session with the nocopy option

Figure 12 on page 33 illustrates a TimeFinder/Clone session using the nocopy option. Note that TimeFinder does not copy data until it is triggered by host I/O to the source device.

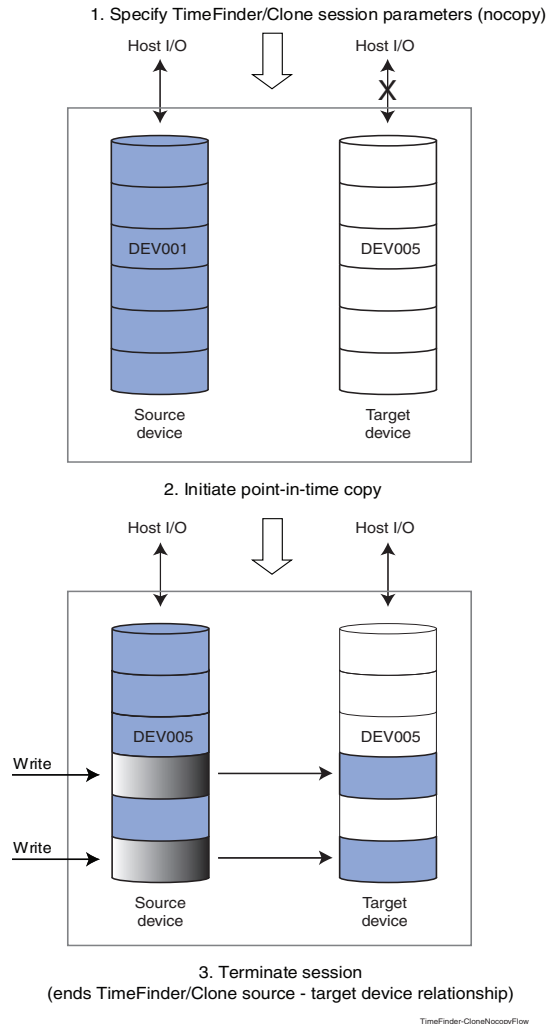


Figure 12 Device clone session using the nocopy option


## TimeFinder/Clone differential session

Figure 13 on page 34 illustrates a differential TimeFinder/Clone session created with the copy option. Note that any subsequent session copies the following data from the source device to the target device:


- ◆ Source device tracks that have been changed by host I/Os since the last TimeFinder/Clone session.
- ◆ Source device tracks whose target counterparts have been changed by host I/Os since the last TimeFinder/Clone session.

Therefore, when you activate the next differential session, the target host sees the corresponding changes to the source device data available on the target device and changes to the target device data overwritten by data from the source device.


1. Specify TimeFinder/Clone session parameters (-copy, differential)



2. Initiate a full-device point-in-time copy



3. Initiate another differential point-in-time copy




TimeFinder-CloneDifferentialFlow

Figure 13 Full device clone session followed by an incremental clone copy

## TimeFinder/Clone with Virtual Provisioning

TimeFinder/Clone works differently in virtually provisioned storage environments. In these environments, thin source and target devices contain pointers to data devices in thin pools.

Figure 14 on page 35 illustrates a TimeFinder/Clone nocopy session that will create a point-in-time copy of the data as it existed at 9:00 am. The time shown in this illustration is 9:05. The session has been activated, but the source device has not received any write I/Os.



**Figure 14** TimeFinder/Clone with Virtual Provisioning: 9:00 session before host write I/O

Figure 15 on page 36 illustrates the same session, but now the time is 9:10. Because this is a nocopy session, data only moves if it is triggered by a host I/O. There has been a write I/O to the first track (data “a”), and the data has been replaced with data “x”. Enginuity has copied data “a” to a new allocation in the target thin pool, the protection bit for the first track has been cleared, and the target device now points to the new location of data “a”.




Figure 15 TimeFinder/Clone with Virtual Provisioning: 9:00 session after host write I/O

Figure 16 on page 37 illustrates a TimeFinder/Clone nocopy session with the same source device that will create a point-in-time copy of the data as it existed at 10:00 am. The time shown in this illustration is 10:05. The session has been activated, but the source device has not received any write I/Os. Note that the 9:00 session is still active.




Figure 16 TimeFinder/Clone with Virtual Provisioning: 10:00 session before host write I/O

Figure 17 on page 38 illustrates the same session, but now the time is 10:15. There has been a write I/O to the second track (data “b”), and the data has been replaced with data “y”. This data is new to both the 9:00 and the 10:00 point-in-time’s targets. Engenuity has copied data “b” to an allocation in the target thin pool for the 9:00 point-in-time copy, and a different allocation for the 10:00 point-in-time copy. The protection bits for the second track have been cleared, and the target device for each point-in-time copy points to a copy of data “b” in the target thin pool.




Figure 17 TimeFinder/Clone with Virtual Provisioning: 10:00 session after host write I/O

# CHAPTER 3

## TimeFinder/Snap Operations

This chapter provides a high-level overview of TimeFinder/Snap operations. Topics include:

- ◆ TimeFinder/Snap overview..... 40
- ◆ TimeFinder/Snap concepts ..... 40
- ◆ Operational overview ..... 42
- ◆ Copying a virtual device to another virtual device (duplicate snap) ..... 50

## TimeFinder/Snap overview

TimeFinder/Snap creates space-saving, logical point-in-time images called snapshots. You can create multiple snapshots simultaneously on multiple target devices from a single source device. Snapshots are not full copies of data; they are logical images of the original information, based on the time the snapshot was created.

TimeFinder/Snap logical copies are particularly useful when you take multiple backups or recovery copies of a source device throughout the day, with only a small percentage of data changing on the device. Through the use of device pointers to the original data, virtual devices (VDEVs) allow you to allocate space based on the expected amount of changes to a device. Rather than replicating a full-device, you backup only data that has changed.

A set of pointers to the source device data tracks is instantly created upon activation of the snapshot. This set of pointers is addressed as a logical volume and is made accessible to a secondary host that uses the point-in-time image of the underlying data. 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. This means that snapshots are created in seconds and can be retired when no longer needed.

In contrast to a full device copy, a snapshot only needs to store the data that has changed on the source device. Multiple snapshots can be created to suit the needs of multiple business processes. Secondary servers see the snapshot as an additional device and have full read/write capabilities to that device.

---

**Note:** Snapshots cannot be used for recovery scenarios because the snapshots are lost if the source device becomes unavailable.

---

## TimeFinder/Snap concepts

Like TimeFinder/Clone, TimeFinder/Snap uses source and target devices, but the target device is a special Symmetrix device known as a virtual device (VDEV).

### Virtual device (VDEV)

A virtual device (VDEV) is a Symmetrix host-addressable cache device used in TimeFinder/Snap operations to keep pointers to point-in-time copies of the source device. Virtual devices are space-efficient since they contain only address pointers to the actual data tracks stored on the source device or in a pool of SAVE devices. Virtual devices can be metadevices as long as the source device is also a metadevice.

During the TimeFinder/Snap session, a virtual device pointer points either to the source track or the SAVE device track depending on whether or not the track has been copied to the SAVE device.



## SAVE device

A SAVE device is a Symmetrix device that is not accessible to the host and can only be accessed through virtual devices that point to it. SAVE devices provide pooled physical storage and are configured with any supported RAID scheme. SAVE devices cannot be metadevices. They store either source data copied to the SAVE pool during the TimeFinder/Snap session or updates from the host mapped to the virtual device.

Since TimeFinder/Snap operations are designed to create point-in-time copies of the source device when only a fraction of the source device changes over time, the SAVE device pool storage capacity can be much smaller than the capacity of the source device.

---

**Note:** Source, virtual, and SAVE devices have to be configured with the same device emulation type (FBA or CKD).

---

## Virtual sessions

TimeFinder/Snap sessions are also referred to as virtual sessions. Like TimeFinder/Clone sessions, virtual sessions use a protection session on the source device.

[“Protection session” on page 23](#) provides more information.

You can run virtual sessions as:

- ◆ Basic virtual sessions where each virtual session is associated with a single virtual device and a virtual point-in-time copy. A protection session runs on the source device for each point-in-time copy. Therefore, you can create a maximum of 16 virtual point-in-time copies of a single source device in open systems environments. In mainframe environments, you can create a maximum of 8 virtual point-in-time copies.
- ◆ Multi-virtual sessions (multi-virtual mode) where a single protection session runs on the source device for up to 128 virtual point-in-time copies. An additional protection session runs on each virtual device associated with a point-in-time copy.

---

**Note:** TimeFinder/Snap sessions on Symmetrix VMAX 40K systems always use multi-virtual mode.

---

You can run only one type of a virtual device session on a single source device.

---

**Note:** Since a virtual session is associated with a point-in-time copy and a particular virtual device, you need to keep a virtual session active as long as you need its snapshot copy. By terminating a virtual session, you remove the associated point-in-time copy because Enginuity automatically releases the storage space in the associated SAVE device pool.

---

## Copy modes

TimeFinder/Snap operates in nocopy mode and uses the Asynchronous Copy on First Write technique.

[“Nocopy” on page 28](#) and [“Asynchronous Copy on First Write” on page 28](#) provide more details.

## Monitoring SAVE device usage

Using virtual copies requires proper planning to prevent the SAVE devices from filling up with pre-updated data. If the SAVE devices fill up, you will begin to lose the pre-update images of the newly changed tracks in the virtual copy session, the virtual device will be set to Not Ready, and the session will fail. (Only sessions with I/O activity will be in a failed state, sessions without I/O activity will continue to operate normally). If this happens, you must terminate the failed sessions to clear the tracks on the SAVE devices. Once a session is terminated, the virtual data is lost and the SAVE device space associated with the session is freed and returned to the SAVE device pool for TimeFinder/Snap use. In addition, you should examine how you are using the SAVE device pools and consider adding more SAVE devices.

## Operational overview

TimeFinder/Snap operations enable you to create and manage virtual copy sessions between a source device and multiple virtual target devices. When you activate a virtual copy session, a point-in-time copy of the source device is immediately available to its host through the corresponding virtual device.

The attached host views the point-in-time copy through virtual device pointers to both the source device and SAVE device, for as long as the session remains active. If you terminate the virtual copy session, the copy is lost and the space associated with the session is freed and returned to the SAVE device pool for future use.

You control the TimeFinder/Snap session life cycle via Symmetrix host-based control software as described in [“Full-device TimeFinder/Clone session life cycle” on page 25](#). These are the basic session control commands:

1. Create (identify the source and the virtual device)
2. Activate (create a virtual point-in-time copy)
3. Terminate (remove the point-in-time copy and release disk space in SAVE device pool).

## Create - identifying the source and the virtual device

To create a snap copy, you must establish a device pair relationship by identifying the source device and the virtual device. When you identify the devices, Enginuity sets protection bits on the source tracks and pointers from the virtual device to tracks on the source device.

Figure 18 on page 43 illustrates a source device with protection bits applied and a virtual device with pointers from the virtual device to tracks on the source device. Note that the virtual device is unavailable to the host until you activate a point-in-time copy.




Figure 18 Protection bits and pointers

## Activate - create a snapshot

After you identify the source and target devices, the host activates the point-in-time copy. The virtual device becomes available to the host when the copy is activated.

Once the host has activated the point-in-time copy, read and write I/Os to the source and virtual devices are handled by Engenuity.

### Write I/O to source device

When there is a write I/O to the protected source track, Engenuity uses the Asynchronous Copy on First Write technique ([“Asynchronous Copy on First Write” on page 28](#)). Engenuity copies the track to the SAVE device before the write I/O occurs and points the virtual device pointer to a SAVE device track. When the track is copied to the SAVE device, the protection bit is cleared if the session is a basic snap session. If the session is running in multi-virtual mode, the protection bit is cleared only if no other virtual device points to this source track.

Figure 19 on page 44 illustrates a source device that has received write I/Os to tracks 2 and 4. Before the write I/Os are allowed to the source device, Engenuity copies the data to the SAVE device and clears the protection bits for the source tracks (example shows a basic virtual session).





Figure 19 Protection bits after write I/Os to the source device

Figure 20 on page 45 illustrates a source device with some tracks that Engenuity has copied due to write I/Os to the source device. Once Engenuity has copied the data, it overwrites the tracks on the source device and changes the pointers from the virtual device to point to the tracks on the SAVE devices.



**Figure 20** Pointers after data has been copied to the SAVE devices

### Read I/O to source device


A host read I/O to the source device is not affected by TimeFinder operations. It is an ordinary read I/O. Engenuity does not copy the track to the target and does not change the protection bits.

### Write I/O to virtual device

If the virtual device points to the protected source device track, Engenuity clears the protection bit, copies the source data to the SAVE device, and updates the virtual device pointers to point to the SAVE device.

Since a track on a SAVE device can be shared among multiple virtual devices, a write I/O operation to a virtual device requires that Engenuity allocate a new SAVE device track and update the corresponding virtual device pointer, while keeping the original SAVE device track with other virtual device pointers intact.

Figure 21 on page 46 illustrates a source device with some tracks that Engenuity has copied (and then overwritten) due to write I/Os to the virtual device. In this example, Engenuity has copied tracks 1 and 3 and cleared their protection bits. Tracks 2 and 4 still have their protection bits cleared due to the earlier write I/Os to the source device.



**Figure 21** Protection bits and pointers after write I/Os to the virtual device

## Read I/O to virtual device

A read I/O to a virtual device is an ordinary read I/O serviced by the device that the virtual device pointer points to (the source device or the SAVE device).

## Recreate - reuse a virtual device

You can recreate virtual sessions to activate a new point-in-time image. When you recreate a virtual copy session, Engenuity resets all of the virtual pointers to point at the source device. To recreate a virtual session, you have to issue the command against an active session.

Since the recreate operation replaces the previous point-in-time image with a new one, the used tracks in the SAVE devices that were associated with the previous session are freed during the processing of this command.

## Restore data from virtual devices

Two types of restore operations can be performed for virtual copy sessions:

- ◆ Incremental restore back to the original source device.
- ◆ Full restore to any target device outside of the existing virtual session. The target device must be of the same size and device emulation type as the source device.

---

**Note:** Restoring from a thick device to a thin device (or from a thin device to a thick device) is not supported.

---

A restore operation requires a new protection session between the source and the restore target device, so a restore can only be performed if an additional session is available for use.

By default, any existing virtual copy sessions persist until manually terminated. After the virtual device has been restored to another device, you must terminate the restore session before another restore operation is allowed from that virtual device.

## Incrementally restoring back to the source device

Through the use of device pointers to the SAVE device, Enginuity can incrementally restore the virtual device back to the source device.

Once the restore completes, Enginuity maintains both the original virtual copy and restore copy sessions and you must manually terminate them if they are not needed for future use. When you terminate the original session, Enginuity deletes the pointers from the virtual device and frees the SAVE device space for future use.

---

**Note:** You must terminate the restored copy session before you terminate the original virtual copy session.

---

In open systems environments, Enginuity sets the target device and the virtual device to the Not Ready state while protection bits are set. The target device automatically becomes available for use (Ready) as soon as protection bits are set. Enginuity then begins copying the tracks until all protected tracks have been restored. To continue using the original virtual session, you must manually set the virtual device to the Ready state after the restore operation has completed.

If there are multiple virtual sessions for the same source device, restoring to the source device from any of the virtual sessions other than the latest one will result in data being written to the SAVE pool. This is because the original data on the source device must be moved to the SAVE pool to preserve the point-in-time copies of the other virtual sessions. The amount of data that is written to the SAVE pool is equal to the difference in changed tracks between the latest virtual session and the restored virtual session.

Figure 22 on page 48 illustrates an incremental restore back to the source device. In this example, track 2 and track 4 were changed and are copied back to the source device.

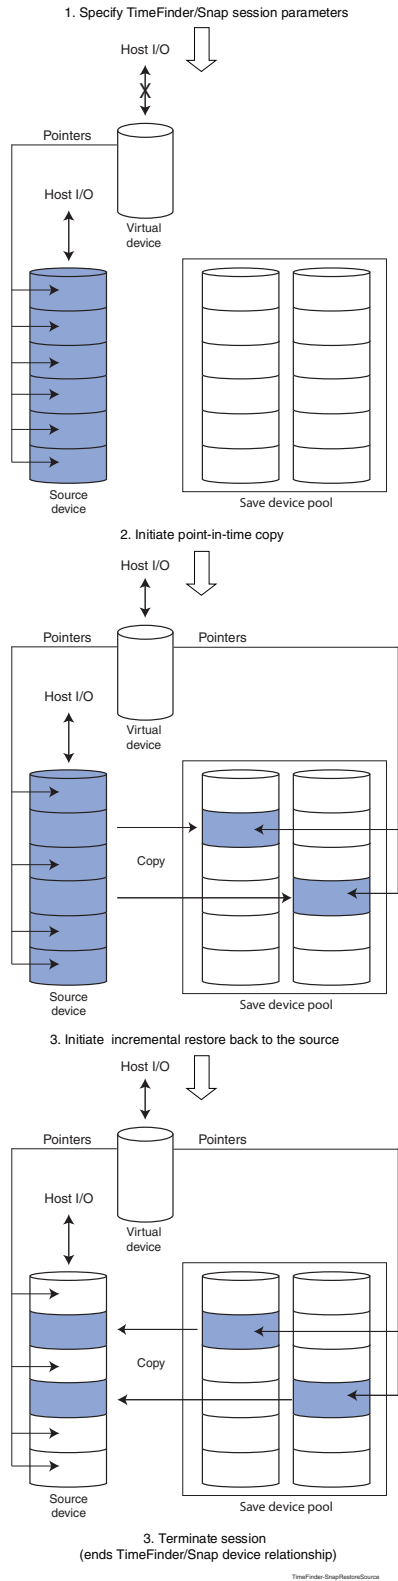


Figure 22 Incrementally restoring back to the source device




## Fully restoring to anywhere

A full restore operation from the virtual device is allowed to any device of the same size and emulation type as the source device. Any changes made to the virtual device during the active virtual copy session will be restored to the specified target device. The restore target inherits the virtual device pointers. A new restore copy session begins between the source and target device and any changed tracks pointed to by the pointers are then copied to the target device from the source device and SAVE devices. Upon completion of the restore operation, Engenuity maintains both the original virtual and the restore sessions until you manually terminate them.

In open systems environments, Engenuity sets the target device and the virtual device to the Not Ready state while protection bits are set. The target device automatically becomes available for use (Ready) as soon as protection bits are set. Engenuity then begins copying the tracks until all protected tracks have been restored. To continue using the original virtual session, you must manually set the virtual device to the Ready state after the restore operation has completed.

**Note:** You must terminate the restored session before you terminate the original virtual session.

Figure 23 on page 49 illustrates a full restore to a target device.



**Figure 23** Fully restoring to a target device

## Terminating a virtual copy session

Terminating a virtual copy session deletes the pairing information between the source and virtual device and resets virtual device pointers (to point back to the virtual device).

Terminating the session causes the host accessing the virtual device to lose access to data pointed to by the virtual device.

## Copying a virtual device to another virtual device (duplicate snap)

Duplicate snap functionality allows you to duplicate a point-in-time copy of a virtual device, which is paired in a previously activated snap session, to another virtual device. This second point-in-time copy session resides with the source device of the original snap session.

Use the duplicate option to begin a virtual copy session that will take one or more source virtual devices and copy them to one or more target virtual devices. The newly created session has its own time stamp indicating the time of its creation.

The following restrictions apply to the duplicate snap functionality:

- ◆ Snap create and activate operations cannot be mixed between normal snap sessions and duplicate snap sessions within the same operation.
- ◆ Snap create and activate operations cannot be mixed between normal snap sessions and duplicate snap sessions within the same operation because the duplicate operations require the `-duplicate` option.
- ◆ The maximum number of duplicated sessions in the Created state is two.
- ◆ When a duplicate snap session is in the Created state, the original session cannot be terminated or recreated until the duplicate session is activated.

Once you activate a duplicate snap session, the copied session works like a normal snap session between the original source device and the duplicated virtual device. This session may be used for restore, recreate, and terminate operations.

# CHAPTER 4

## TimeFinder VP Snap Operations

This chapter provides a high-level overview of TimeFinder VP Snap operations. Topics include:

- ◆ TimeFinder VP Snap overview..... 52
- ◆ VP Snap operations..... 52
- ◆ VP Snap example ..... 53

## TimeFinder VP Snap overview

VP Snap leverages TimeFinder/Clone technology to create space-efficient snaps for thin devices by allowing multiple sessions to share allocations within a thin pool. VP Snap provides the efficiency of Snap technology with improved cache utilization and simplified pool management. With VP Snap, tracks can be stored in the same thin pool as the source, or in another pool of your choice.

---

**Note:** VP Snap on FBA devices is supported by the Mainframe Enablers TimeFinder/Clone Mainframe Snap Facility version 7.4.

---



---

**Note:** VP Snap source and target devices are optimized by FAST VP, but the shared target allocations are not moved.

---

## VP Snap operations

VP Snap sessions copy data from the source device to the target device only if triggered by a host write operation. Read I/Os to protected tracks on the target device do not result in data being copied.

For a single activated VP Snap session on a source device, the target represents a single point-in-time copy of the source. Copied data resides on allocations in the thin pool. For example, if tracks 100, 200, and 300 are written on the source device, the point-in-time data for each track resides in a unique allocation in the thin pool.

When there is a second VP Snap session from the same source device to a different target, the allocations can be shared. For example, if there is a write I/O to tracks 1100, 1200, and 1300 on the source device, the data is new to both targets' point-in-time, and the point-in-time data can be saved in a single set of allocations that is shared by both target devices.

However, if there is another write I/O to tracks 100, 200, or 300 on the source device, since the data is new for only the second session's point-in-time, when the tracks are copied to the second target device, the point-in-time data is put into a set of allocations that are uniquely owned by the second target device. In other words, the allocations for these tracks cannot be shared.

If more VP Snap sessions are added to the same source device, data is copied to the targets based on whether the source data is new with respect to the point-in-time of each copy. When data is copied to more than one target, only a single shared copy resides in the thin pool.

If there is a write I/O to one or more of the tracks stored in a shared allocation, the affected allocation for that target will be split off from the shared group because the data is now different than the data for the other targets that are using that allocation. The new data of the written target will be stored in a separate allocation while the shared allocation will still contain the data of the other targets.

When VP Snap sessions are terminated, the target device is removed from any shared allocations that were part of the session, and any non-shared allocations for that device are deallocated. When all but one of the VP Snap sessions are terminated, the last remaining session uses the same space in the thin pool, but it is no longer a shared allocation. Upon termination of the last session, the space is deallocated.


## Restore operations

VP Snap supports incremental restore operations back to the original source device. The source device does not have to be fully copied and the original session between the source device and the target device is maintained.

## VP Snap example

VP Snap allows multiple target devices to share allocations within the storage pool. In the example below, instead of copying data “b” to one allocation for the 9:00 point-in-time copy, and a different allocation for the 10:00 point-in-time copy, VP Snap allows the targets for both point-in-time copies to share the allocation that contains data “b”.

[Figure 24 on page 53](#) shows the target for the 10:00 point-in-time sharing the allocation for data “b” with the 9:00 point-in-time copy.



**Figure 24** VP Snap session



# CHAPTER 5

## TimeFinder Integration

This chapter describes TimeFinder integration with other Enginuity features and products. Topics include:

- ◆ [TimeFinder and SRDF .....](#) 56
- ◆ [TimeFinder and Virtual Provisioning .....](#) 64

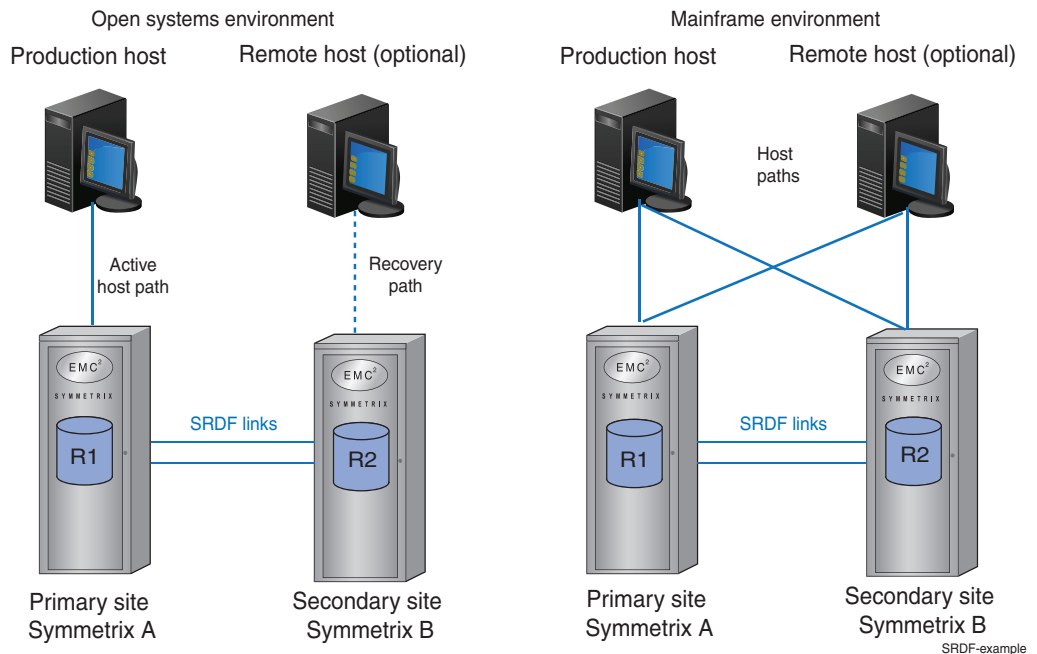
# TimeFinder and SRDF

TimeFinder is closely integrated with the EMC® Symmetrix® Remote Data Facility (SRDF®). The SRDF family of products are Symmetrix-based disaster recovery, parallel processing, and data migration solutions.

The SRDF family of products are based on active remote mirroring and dependent-write consistent copies of data maintained at one or more remote locations. A dependent-write is a write I/O request that cannot be issued by an application until a prior, related write I/O request is completed. Dependent-write consistency is required to ensure transactional consistency once the applications are restarted at the remote location.

SRDF solutions require at least two arrays (the primary array and the secondary array). The arrays can be located at sites in the same room, in different buildings within the same campus, or many kilometers apart.

Figure 25 on page 56 shows a basic SRDF configuration in an open systems (left) and in a mainframe environment (right). The production host is connected to Symmetrix A (the primary Symmetrix array). The primary device (R1) in Symmetrix A contains production data that is remotely mirrored to the secondary device (R2) in Symmetrix B, while the production host is in operation. Symmetrix A and Symmetrix B are connected to each other through the SRDF links. Each R1 device is paired with its remote R2 partner device to constitute an SRDF pair.



**Figure 25** SRDF basic configurations

During ordinary SRDF remote mirroring operations, the production host connected to Symmetrix A is issuing I/Os to the R1 devices. Although receiving data from their R1 partner devices, the R2 devices in Symmetrix B are made Read Only or Not Available to the host connected to Symmetrix B. The R2 devices cannot become Read/Write enabled to the host as long as the SRDF relationship between the R1 and R2 device exists.



Depending on the type of remote mirroring, the SRDF solutions operate in synchronous mode (SRDF/S), asynchronous mode (SRDF/A), or adaptive copy mode.

The *EMC Symmetrix Remote Data Facility (SRDF) for VMAX 40K, VMAX 20K/VMAX, DMX Series Product Guide* provides more information about SRDF operations.

The following are the key benefits of TimeFinder integration with SRDF:

- ◆ **Remote command support for simplified automation** — EMC host-based control software commands can be transferred across the SRDF links, which allows you to issue a single command to the primary array to initiate TimeFinder operations on both the primary and secondary array.
- ◆ **Consistent data images across multiple devices and Symmetrix arrays** — The SRDF/Consistency Group (SRDF/CG) feature is used in SRDF/S solutions to guarantee that a dependent-write consistent image of production data on the R1 devices is created across the SRDF links.

The TimeFinder/Consistency Group (TimeFinder/CG) feature guarantees that a consistent point-in-time image of data written across multiple local devices (TimeFinder source devices) is created on another set of local devices (TimeFinder target devices).

SRDF/CG and TimeFinder/CG both use the Enginuity Consistency Assist infrastructure. By using TimeFinder/CG in an SRDF configuration, you can create dependent-write consistent local and remote images of production data across multiple devices and Symmetrix arrays.


With Enginuity version 5874 or higher, the SRDF/CG feature does not require a separate license. It is offered with the SRDF/S product license. Similarly, the TimeFinder/CG feature does not require a separate license with Enginuity version 5874 or higher. TimeFinder/CG is offered with the TimeFinder/Clone and the TimeFinder/Snap product licenses.

## R1 and R2 devices in TimeFinder operations

You can use TimeFinder and SRDF products to complement each other in environments that require both local replication and remote mirroring. R1 and R2 devices can participate in TimeFinder/Clone and TimeFinder/Snap operations to create local replicas (typically gold copies required for disaster recovery scenarios). The following rules apply:

- ◆ You can use R1 and R2 devices as source devices in TimeFinder/Clone and TimeFinder/Snap.
- ◆ You can use R1 devices as target devices in TimeFinder/Clone (except for VP Snap).
- ◆ You can use R2 devices as target devices in TimeFinder/Clone (except for VP Snap) unless SRDF remote mirroring is active. This restriction applies because the R2 device cannot receive data from its R1 partner device across the SRDF links and from the local TimeFinder source device at the same time. If you want to use R2 devices as TimeFinder/Clone target devices, you must first suspend the SRDF remote mirroring.

Figure 26 on page 58 shows R1 and R2 devices participating in TimeFinder operations.



**Figure 26** R1 and R2 devices participating in TimeFinder operations.

**Note:** Special considerations apply to TimeFinder operations running from R2 devices that participate in an active SRDF/A session. “SRDF/A and TimeFinder” on page 60 provides more information.

## SRDF/S and TimeFinder operations

SRDF/S solutions support any type of TimeFinder copy sessions running on R1 and R2 devices as long as the conditions described in “R1 and R2 devices in TimeFinder operations” on page 57 are met.

The following sections explain simultaneous TimeFinder/Clone solutions that employ SRDF/S for remote mirroring and TimeFinder/Clone for local replication.

### Simultaneous TimeFinder/Clone and failover capability with Engenuity 5875

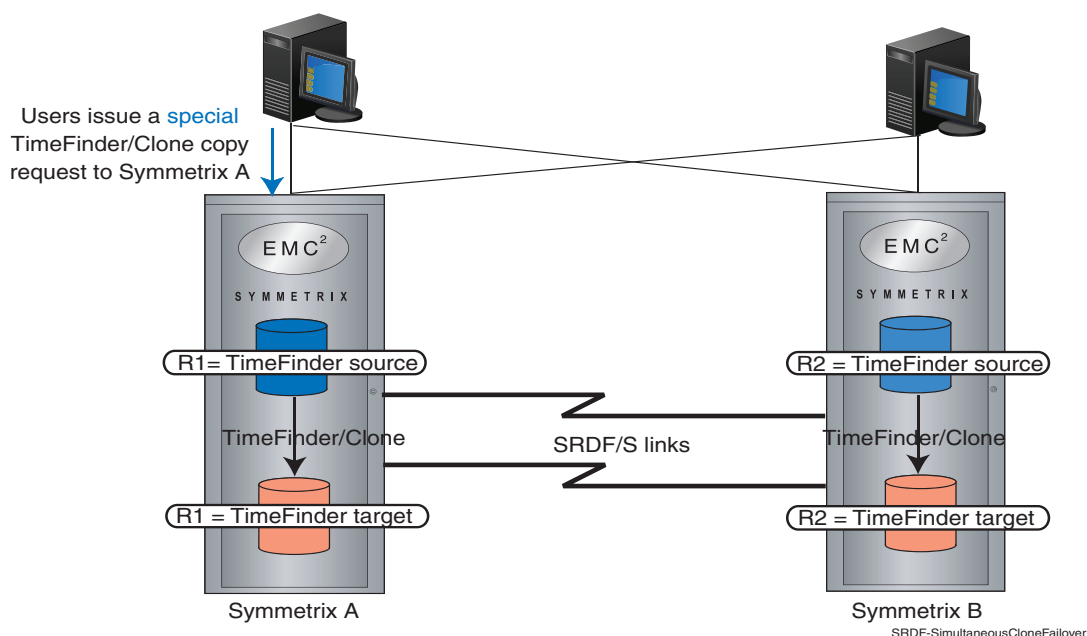
With Engenuity version 5875, simultaneous TimeFinder/Clone operations provide failover capabilities in mainframe host environments using full device or dataset level TimeFinder/Clone operations.

With Engenuity version 5875 running on both Symmetrix arrays, simultaneous TimeFinder/Clone operations no longer move data across the SRDF/S links to create remote images on the R2 TimeFinder/Clone target devices in Symmetrix B. Instead, you can issue a special TimeFinder/Clone copy request to Symmetrix A. This single request to Symmetrix A simultaneously creates TimeFinder/Clone copies on both sides of the SRDF/S relationship.

These simultaneous TimeFinder/Clone operations involve the following:

1. The R1 devices (TimeFinder/Clone source devices) in Symmetrix A must be synchronized with their SRDF partners in Symmetrix B.
2. Once the SRDF/S pairs are synchronized, you issue a special TimeFinder/Clone copy request to Symmetrix A. This request is transferred across the SRDF/S links to Symmetrix B to simultaneously initiate TimeFinder/Clone copy sessions from TimeFinder/Clone source to TimeFinder/Clone target devices on both sides of the SRDF/S relationship.

Figure 27 on page 59 shows an example.



**Figure 27** Simultaneous TimeFinder/Clone with Engenuity version 5875

When simultaneous TimeFinder/Clone copy sessions start in Symmetrix A and in Symmetrix B, both arrays hold the same data image:

- ◆ The TimeFinder/Clone source devices in Symmetrix B capture current data since they are synchronized with their SRDF partners in Symmetrix A.
- ◆ The TimeFinder/Clone target devices in Symmetrix B capture the same point-in-time copy as their SRDF partners in Symmetrix A because of simultaneous TimeFinder/Clone operations that take place on both sides of the SRDF/S relationship.

By preserving the same data image as Symmetrix A, including in-process clone copies, Symmetrix B can become the failover site and can assure that all data made available to the host after the failover operation is current data.

This feature can be used with Concurrent SRDF on the SRDF/S branch of a concurrent SRDF configuration. It is not supported in SRDF/A, Cascaded SRDF, or in SRDF/Star configurations and does not support SRDF pairs in which the R2 devices are larger than their R1 partners.

The *EMC Mainframe Enablers TimeFinder/Clone SNAP Facility Product Guide* provides details about controlling simultaneous TimeFinder/Clone operations using mainframe host-based control software.

## SRDF/A and TimeFinder

SRDF/A supports any type of TimeFinder copy session on the R1 devices. On R2 devices that participate in an active SRDF/A session, prior to Engenuity version 5875, only TimeFinder/Clone full-device copy sessions with the precopy option are supported.

With Engenuity 5875, you can run TimeFinder/Clone full-device and dataset (extent) level copy sessions using any copy option and TimeFinder/Snap regular and multi-virtual sessions on the R2 devices as long as you enable the SRDF/A device-level pacing option. You can still run the TimeFinder/Clone sessions using the precopy option on the R2 devices without enabling the SRDF/A device-level pacing option.

### SRDF/A device-level pacing option

With Engenuity version 5875, the SRDF/A write pacing functionality features a new option known as the SRDF/A device-level pacing.

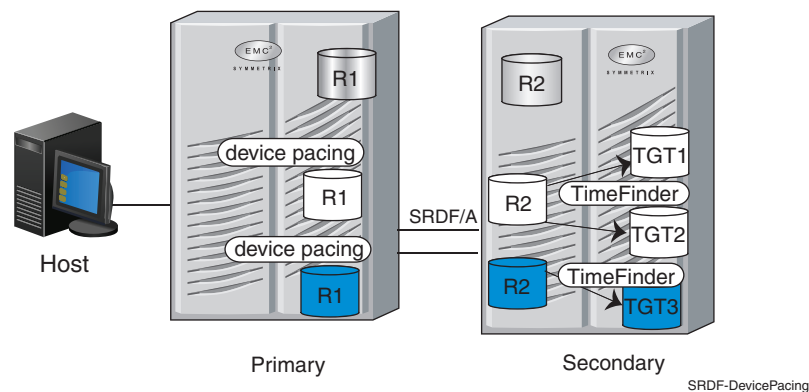
The original SRDF/A write pacing feature extends host write I/O response time by adding a write pacing delay to the overall host write I/O response time. This prevents SRDF/A operational interruptions caused by high levels of Symmetrix cache utilization. If no corrective actions are taken when the SRDF/A maximum cache utilization thresholds are reached, the SRDF/A session is dropped.

If you enable this feature and the SRDF/A maximum cache utilization thresholds are met, the write pacing feature adds an appropriate write pacing delay to the overall host response time to keep the SRDF/A session active. The write pacing feature therefore temporarily slows down host I/O rates to match the slower SRDF/A service rates.

The SRDF/A device-level pacing feature is a write pacing option that may apply a write pacing delay for individual SRDF/A R1 devices if:

- ◆ TimeFinder copy sessions run on the R2 partners of the R1 devices
- ◆ The SRDF/A maximum cache utilization thresholds are met.

[Figure 28 on page 60](#) shows how device-level pacing is activated for individual R1 devices.



**Figure 28** SRDF/A device-level pacing activated for R1 volumes

Since the device-level pacing option applies the write pacing delay only for individual R1 devices whose R2 partners participate in TimeFinder operations, it indirectly paces the SRDF/A cycle requests issued to the appropriate R2 devices.

If device-level pacing is active, the host write I/Os are paced as follows:

- ◆ If you do not specify the maximum write pacing delay, the device-level pacing feature always applies the appropriate write pacing delay to the overall host write response time to keep the SRDF/A session active. By default, the maximum write pacing delay is 50 milliseconds.
- ◆ If you have specified the maximum write pacing delay, the device-level pacing feature applies the appropriate write pacing delay to keep the SRDF/A session active up to the specified maximum write pacing delay. You can specify a maximum write pacing delay of up to one second.

The key benefit of the SRDF/A write pacing feature and its device-level pacing option is the dynamic, self-paced mechanism. Once you enable the device-level pacing option, the device-level pacing technique is employed only when required and applies the appropriate write pacing delay to the host write I/O response time to keep the SRDF/A session active. When write pacing is enabled but pacing is not required to keep the SRDF/A session running, host write I/Os are not paced.

The *EMC Symmetrix Remote Data Facility (SRDF) for VMAX 40K, VMAX 20K/VMAX, DMX Series Product Guide* provides more information about the SRDF/A write pacing.

## SRDF adaptive copy operations and TimeFinder


Like SRDF/S, SRDF adaptive copy operations support any type of TimeFinder copy sessions running on R1 and R2 devices as long as the conditions described in [“R1 and R2 devices in TimeFinder operations” on page 57](#) are met.

### SRDF/AR

SRDF/Automated Replication (SRDF/AR) is the most widely used integrated solution using the SRDF adaptive copy for remote integration and TimeFinder for local replication. SRDF/AR provides a long-distance disaster restart solution as a:

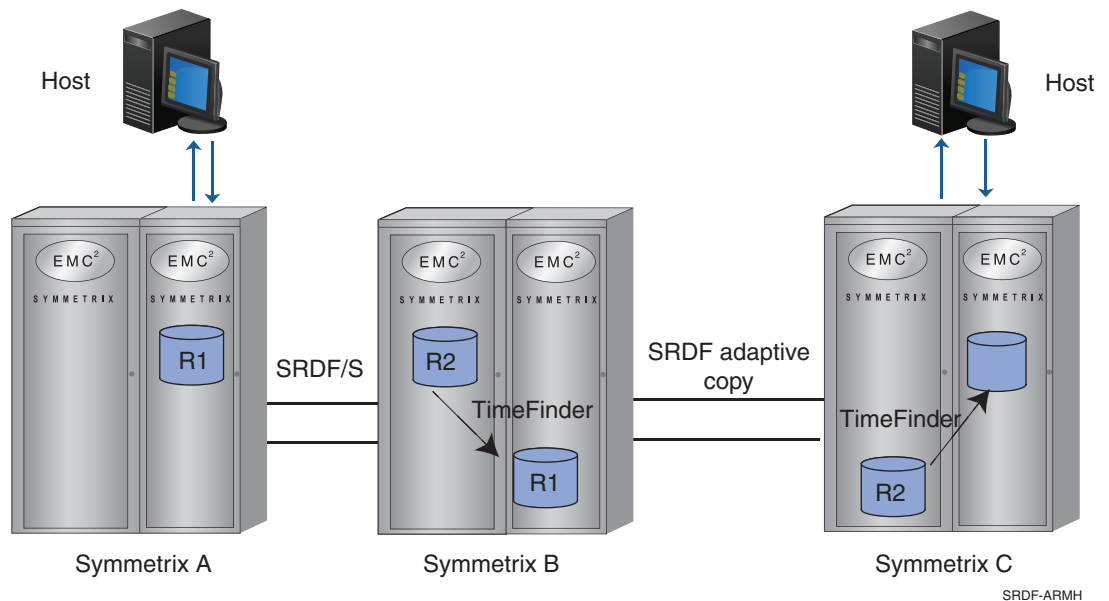
- ◆ Two-site configuration operating in adaptive copy mode using SRDF/DM in combination with TimeFinder.
- ◆ Three-site configuration using a combination of SRDF/S, SRDF/DM, and TimeFinder. This solution operates in synchronous mode between Symmetrix A and Symmetrix B and in adaptive copy mode between Symmetrix B and Symmetrix C.

[Figure 29 on page 62](#) shows a single-hop SRDF/AR configuration. The R1 device is also the TimeFinder target device, replicated across the SRDF links. TimeFinder thus separates the remote replication and production processing at the production site in Symmetrix A. The R2 device holds an asynchronous restartable copy of the R1 data. The R2 device can be configured as a TimeFinder source device to enable independent asynchronous processing in Symmetrix B.



**Figure 29** SRDF/AR single-hop data flow

Figure 30 on page 62 shows a multi-hop SRDF/AR configuration. This disaster restart solution operates in synchronous mode between Symmetrix A and Symmetrix B and in adaptive copy mode between Symmetrix B and Symmetrix C. The three-site solution provides a remote data copy in Symmetrix B (to guarantee a disaster restart with zero data loss if Symmetrix A fails) and an asynchronous restartable copy in Symmetrix C. TimeFinder is required in Symmetrix B to separate synchronous remote mirroring from remote replication between Symmetrix B and Symmetrix C.



**Figure 30** SRDF/AR multi-hop data flow

SRDF/AR is a native Symmetrix solution that does not require a Symmetrix controlling host attached to Symmetrix B or Symmetrix C because its design fully supports automation from the controlling host connected to Symmetrix A. The controlling host attached to Symmetrix A can issue SRDF and TimeFinder commands to R1 and R2 devices. SRDF/AR can be used in mainframe and open systems host environments.

---

**Note:** SRDF/AR uses TimeFinder/Clone Emulation mode with Enginuity versions 5874 and higher and TimeFinder/Clone Emulation mode or TimeFinder/Mirror with Enginuity versions 5773 and lower. [“TimeFinder/Mirror operations” on page 68](#) and [“Clone Emulation mode” on page 69](#).

---

## TimeFinder and Virtual Provisioning

Full-device TimeFinder/Clone and TimeFinder/Snap both support Virtual Provisioning and allow thin source and target devices in TimeFinder operations. Extent-level TimeFinder/Clone operations do not support Virtual Provisioning and thin devices.

Enginuity versions 5773 and higher support multi-virtual TimeFinder/Snap sessions and the Asynchronous Copy on First Write optimization technique for thin devices used as source devices in TimeFinder/Clone or TimeFinder/Snap sessions.

Mixed thick and thin devices are not supported with cascading operations. All devices in cascading sessions must be thick or thin.

[Table 1 on page 64](#) lists TimeFinder/Clone support for Virtual Provisioning.

**Table 1 Thin devices and full-device TimeFinder/Clone**

TimeFinder/Clone source	TimeFinder/Clone target	Enginuity version		
		5773 <sup>1</sup>	5874 <sup>1</sup>	5875 and higher <sup>2</sup>
Thin	Thin	Yes	Yes	Yes
Thin	Thick	No	Yes	Yes
Thick	Thin	No	Yes	Yes

1. Data devices can be configured as RAID 1 and RAID 6 but not RAID 5 groups.
2. Data devices support RAID 1, RAID 5, and RAID 6 groups.

Full-device TimeFinder/Clone support for Virtual Provisioning provides you with a mechanism for migrating to a virtually provisioned storage environment within the same Symmetrix array. By pairing thick devices with thin devices in a TimeFinder/Clone source-target relationship, you can copy data and then move production operations to thin devices.

[Table 3 on page 65](#) lists TimeFinder/Snap support for Virtual Provisioning.

**Table 2 Thin devices and TimeFinder/Snap**

TimeFinder/Snap source	TimeFinder/Snap target	Enginuity version		
		5773 <sup>1</sup>	5874 <sup>1</sup>	5875 and higher <sup>2</sup>
Thin	Virtual	Yes	Yes	Yes
Virtual	Thin <sup>3</sup>	Yes	Yes	Yes

1. Data devices can be configured as RAID 1 and RAID 6 but not RAID 5 groups.
2. Data devices support RAID 1, RAID 5, and RAID 6 groups.
3. Restore to the same device or to a different device.

## SRDF and TimeFinder in virtually provisioned environments

SRDF supports Virtual Provisioning and allows TimeFinder/Clone and TimeFinder/Snap operations using thin SRDF devices. The rules for SRDF device roles in TimeFinder operations described in [“R1 and R2 devices in TimeFinder operations” on page 57](#) still apply.



Table 3 on page 65 and Table 4 on page 65 list SRDF and TimeFinder support for Virtual Provisioning.

**Table 3** Thin SRDF devices and TimeFinder/Clone

Thin SRDF volume	SRDF mode	TimeFinder/Clone device role	Enginuity version			
			5773	5874	5875	5876
R1	SRDF/A	source or target <sup>4</sup>	No	Yes	Yes	Yes
R2	SRDF/A	source	No	Yes <sup>1</sup>	Yes <sup>2</sup>	Yes <sup>2</sup>
R1	SRDF/S <sup>3</sup>	source or target <sup>4</sup>	Yes	Yes	Yes	Yes
R2	SRDF/S <sup>3</sup>	source	Yes	Yes	Yes	Yes

1. Only full-device TimeFinder/Clone with the precopy option is supported.
2. Full-device TimeFinder/Clone and extent level TimeFinder/Clone copy sessions using any copy option are supported as long as the SRDF/A device-level pacing is enabled. TimeFinder/Clone with the precopy option is allowed without the SRDF/A device-level pacing feature enabled.
3. SRDF/S or SRDF adaptive copy.
4. VP Snap cannot use SRDF devices as target devices.

**Table 4** Thin SRDF devices and TimeFinder/Snap

Thin SRDF volume	SRDF mode	TimeFinder/Snap device role	Enginuity version		
			5773	5874	5875 and higher
R1	SRDF/A	source or virtual	No	Yes	Yes
R2	SRDF/A	source	No	No	Yes <sup>1</sup>
R1	SRDF/S <sup>2</sup>	source or virtual	No	Yes	Yes
R2	SRDF/S <sup>2</sup>	source	No	Yes	Yes

1. The SRDF/A device-level pacing feature must be enabled.
2. SRDF/S or SRDF adaptive copy.

Similar to the TimeFinder/Clone support for Virtual Provisioning, the SRDF support for virtual provisioning allows you to mirror thick R1 devices to thin R2 devices. Once all of the devices are synchronized, you can perform a failover operation to migrate to the virtually provisioned remote array and resume production operations.



# APPENDIX A

## TimeFinder/Mirror

This chapter provides a high-level overview of TimeFinder/Mirror operations. Topics include:

- ◆ TimeFinder/Mirror ..... 68
- ◆ Clone Emulation mode ..... 69

## TimeFinder/Mirror

On Symmetrix DMX series arrays running Engenuity versions prior to 5874, TimeFinder /Mirror is supported as one of the TimeFinder local replication products. TimeFinder/Mirror is not supported on Symmetrix VMAX Series arrays, but Clone Emulation mode allows you to use TimeFinder/Mirror commands to run TimeFinder/Clone operations.

[“Clone Emulation mode” on page 69](#) provides more details.

### Standard and BCV devices

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

A standard device can have any supported RAID configuration with at least one BCV mirror. The BCV device can be configured as a RAID 0, RAID 1, and RAID 5 device.

---

**Note:** Symmetrix host-based control software uses the terms standard device (STD) and BCV device outside the scope of TimeFinder/Mirror to refer to ordinary Symmetrix devices.


---

### TimeFinder/Mirror operations

TimeFinder/Mirror involves associating and establishing the BCV device as a mirror of a specific standard Symmetrix device. The BCV device gets synchronized data from the standard device while it is established as a BCV mirror to the standard device. As a result, the BCV device becomes inaccessible using its original device address while it is established to the standard device. Once the BCV device is synchronized with its source, you can split it from the standard device, thereby making it available again to its host for backup or other host processes through its original device address.

After host processing on the BCV device is complete, the BCV may again be established and synchronized with the standard device.

Figure 31 on page 69 illustrates a standard device with an established and split BCV mirror.



**Figure 31** Standard device with an established BCV mirror

## Clone Emulation mode

Clone Emulation mode transparently converts TimeFinder/Mirror commands to TimeFinder/Clone commands. If you want to run your existing scripts that use TimeFinder/Mirror commands, they will still work, but the underlying mechanism is TimeFinder/Clone.

Use Clone Emulation mode only if you need to continue using TimeFinder/Mirror scripts. Otherwise, TimeFinder/Clone is the recommended solution.

The following operational characteristics apply to Clone Emulation:

- ◆ Forces precopy mode
- ◆ Supports multi-device commands.
- ◆ With Engenuity version 5875 and higher, Clone Emulation supports thin SRDF BCV devices.

