

# USING EMC RECOVERPOINT CONCURRENT LOCAL AND REMOTE FOR OPERATIONAL AND DISASTER RECOVERY

Applied Technology

## Abstract

This white paper discusses EMC® RecoverPoint concurrent local and remote data protection. This solution locally replicates SAN volumes in one or more storage arrays and maintains an online journal of all changes to the volumes. It also replicates data to a remote storage system maintaining an online remote journal of changes. This solution is designed as a local operational recovery solution and an extended-distance disaster recovery solution. RecoverPoint enables databases and applications to be brought back online quickly and easily after an event like data corruption —recover to any point in time, going back to the exact, best, and most recent recovery point.

May 2012

Copyright © 2008, 2012 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate 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.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

VMware and ESX are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number h4175.5

## Table of Contents

<b>Executive summary</b> .....	<b>4</b>
<b>Introduction</b> .....	<b>4</b>
Audience.....	5
<b>Information protection challenges</b> .....	<b>5</b>
<b>Data protection technologies</b> .....	<b>5</b>
Tape backup .....	6
Snapshot .....	6
Replication.....	7
Continuous data protection .....	8
Benefits of RecoverPoint .....	9
<b>EMC RecoverPoint operations</b> .....	<b>10</b>
Managing RecoverPoint.....	11
System architecture .....	12
RecoverPoint appliance.....	13
Write splitter .....	13
Repository volume .....	13
Journal volume .....	13
Consistency groups .....	14
Consistency group operations .....	14
Application servers .....	16
Heterogeneous storage arrays.....	16
SAN considerations.....	16
Preserving data integrity during failure scenarios .....	17
Managing RecoverPoint appliance failure.....	17
Managing lost connectivity to source volumes.....	18
Managing lost connectivity to local or remote copy volumes.....	18
Managing host or write-splitter failures .....	18
Utilizing VNX series, CLARiiON, and Symmetrix operations with RecoverPoint.....	18
<b>RecoverPoint use-case scenarios</b> .....	<b>20</b>
Case 1: Storage migration .....	20
Case 2: Repurposing for development and test support.....	21
Case 3: Simplifying backup, restore, and disaster recovery.....	22
Case 4: Individual file or folder level recovery .....	23
Case 5: Federated application recovery using group sets.....	24
Case 6: Performing recovery testing fire drills .....	27
<b>Conclusion</b> .....	<b>29</b>
<b>References</b> .....	<b>29</b>

## Executive summary

The EMC® RecoverPoint family provides cost-effective, local continuous data protection (CDP), continuous remote replication (CRR), and concurrent local and remote (CLR) data protection solutions that include point-in-time data recovery. RecoverPoint/SE is the entry offering that simplifies replication and continuous data protection for VNX™ series, CLARiiON® CX™, CX3 UltraScale™, or CX4 arrays. RecoverPoint/EX supports local and remote replication for Symmetrix VMAX 10K, Symmetrix VMAX 20K, Symmetrix VMAX 40K, VPLEX Local, VPLEX Metro, VNX series, and CLARiiON CX3 and CX4 arrays. RecoverPoint/CL is the full-featured offering that adds support for intelligent fabrics, heterogeneous servers, and heterogeneous storage platforms. Both products are appliance-based, out-of-band data protection solutions designed to protect production data to local and/or remote sites. These products enable customers to centralize and simplify their data protection management, and they allow for recovery of the data to nearly any point in time.

RecoverPoint CDP provides block-level local replication between LUNs in the same SAN using technology that journals every write for later recovery to any point in time. RecoverPoint CRR provides block-level remote replication between LUNs in two different SANs using technology that journals groups of writes for later recovery to significant points in time. RecoverPoint CLR provides simultaneous block-level local and remote replication for LUNs with one copy residing locally in the same SAN with every write journaled, and the second copy residing remotely in a different SAN with significant groups of writes journaled. Recovery of the local or remote copy can occur without affecting the other copy.

RecoverPoint utilizes policies that map the recovery time objectives (RTO) and recovery point objectives (RPO) by consistency groups, allowing for flexibility in protecting multiple applications. RecoverPoint 3.5 includes new features and functions that improve usability, performance, scalability, and support for RecoverPoint and RecoverPoint/SE.

## Introduction

Today's businesses are faced with an ever-increasing amount of data that threatens to undermine their existing storage management solutions. Data protection is no longer the simple copying of yesterday's changed files to tape. Critical data changes occur throughout the day, and to protect this data customers are frequently turning to new technology such as continuous data protection (CDP).

This white paper is designed to give the reader an introduction to data protection concepts, explore how EMC RecoverPoint operates, and give some use-case scenarios for RecoverPoint. This paper serves as an overview to the RecoverPoint software and discusses key software functions and features. The [References](#) section provides related information, including an administrator guide and white papers.

## Audience

This white paper is intended for systems integrators, systems administrators, and members of the EMC and partners professional services community.

## Information protection challenges

Protecting information is the key to a successful businesses operation. Implementing remote site protection for critical business information is not a simple proposition. The first step, even before analyzing technology, is to understand the current business processes and develop a clear set of objectives and plans that reflect what is required to safeguard against any disaster that could make the data at the primary site unavailable.

An evaluation of the current business applications and their information must be completed as part of designing a protection solution. If the production site goes offline due to a disaster, and the business processes must be transferred to the recovery site, how much data loss can be tolerated before the business is deemed unable to restart production? This is termed the recovery point objective (RPO). For a few critical business applications, such as real-time financial transactions, businesses cannot afford to lose any data in the event of a disaster. In this case their RPO must be zero. For most business applications, the loss of a few minutes to a few hours of data can be easily tolerated, and their RPO is much more flexible. It is fair to say that RPO is dependent upon business rules and processes more than anything else.

Once the RPO requirements are well understood, the second challenge is how long it takes to restart the business applications at the recovery site with the data at the remote site. This measurement is termed the recovery time objective (RTO). In the case of real-time financial transactions, it may be very important that the application comes back online in a matter of seconds without any noticeable impact to the end users. For other applications, a delay of a few minutes or hours may be tolerable.

It is fair to say that the shorter the RTO and the RPO, the more difficult or costly it may be to implement a disaster recovery process successfully. A perfect configuration with no data loss guaranteed and data instantly available may come at a complexity and price that are not practical. It is important to distinguish between absolutely critical business applications (Tier 1) and other applications. Any disaster recovery solution chosen must have the flexibility to support applications that have differing RPO and RTO values without compromising protection for the Tier 1 applications.

## Data protection technologies

Over the past 30 years, data protection has evolved as business needs changed. At first, most computer systems were stand-alone, with their entire data residing on a single system. Networking and interconnectivity between systems were expensive and limited. Yet, there was a need to protect the data that resided on these systems.

Out of this need arose the capability to back up data to tape. Tape was the prevalent interchange media for data, and every major system had one or more tape drives. As these systems evolved, so did the backup technology. As systems became interconnected, it became more common to share a single tape drive or tape library between multiple systems. In these systems, one server owned the tape library, and the other systems became clients and sent their data across the network.

## Tape backup

The advantages of tape backup are that the tapes are very durable and tend to be written in a format that is portable across systems. Tapes became a common mode of interchanging data between systems. However, the disadvantages of backup slowly became a problem as computer systems evolved. To back up an application and all of its data requires that the application be shut down, so that its data files are in a consistent state.

Shutting down an application is naturally disruptive to the users of these applications, so backups were usually performed during a quiet time, typically midnight to 6 A.M. This timeframe when the applications were shut down came to be known as the backup window. Many customers could not afford to have their applications shut down for more than once a day, so these customers had to find some other way to protect their data.

Additionally, backup operations were rarely trouble-free – sometimes a server would be shut down or an application may not have shut down cleanly, resulting in a missed or corrupted backup. Additionally, as data volumes grew, the time to back up this data increased, adding pressure to the backup window.

Backup vendors responded by creating alternate methods of backing up the data that reduced the backup window. These included technologies such as open-file backup, implementing incremental or differential backups, using SAN or NAS APIs to access the data directly, and supporting faster, streaming tape drives. Regardless of how many changes were made, backups tended to remain an inherently slow, cumbersome activity.

## Snapshot

On the heels of backup came the concept of snapshots. A snapshot is a copy of a file system, volume, or LUN that contains an image of the data as it appeared at the point in time at which the copy was initiated. The snapshot may either duplicate or replicate the data it represents.

Snapshot technology can be implemented on the host, in the storage network, or at the array level. Host-based snapshots may be performed at the volume level as in the Veritas Volume Manager Snapshot facility, or at the file system level as in Microsoft's Volume Shadow Copy Services (VSS). When implemented inside of any array, most snapshots are at the physical, or block level, with the exception of products such as Network Appliance's WAFS file system, where the snapshots can be extended to a file level.

A snapshot may be a full copy of a volume or LUN (for example, the EMC Symmetrix® TimeFinder® business continuance volumes), or it may be a replicate snapshot, which just contains the changes necessary to apply to the current version of the LUN to re-create the image at a specific point in time (for example, the EMC VNX SnapView™ snapshot).

Snapshots tend to be less disruptive to applications and environment. In fact, many applications have specific interfaces that can be invoked to create a snapshot-ready image of their data. For example, VSS has an operation to quiesce an application prior to a snapshot being created that writes in-memory data to disk to ensure the applications on disk files are in a consistent state prior to the snapshot. This helps improve the recoverability of the application's state from a snapshot image.

Snapshot technology can be host-based, network-based, or array-based. Host- and network-based technologies tend to be more generic, and less dependent upon a specific array vendor's storage, while array-based technology is usually tied to the vendor's storage product and may have limitations, such as it can only support snapshots using resources available inside the array.

Host- or network-based products tend to have fewer of these limitations, as they build on resources presented to them from the underlying storage infrastructure. For example, the Veritas Volume Manager Snapshot facility creates an exact copy of a primary volume at a particular instance in time. After a snapshot is taken, it can be accessed independently of the volume from which it was taken.

Regardless of the implementation, snapshots are less disruptive, more reliable, and faster than traditional tape backup. However, unless they are integrated with the application, they may miss important data residing in-memory at the time of the snapshot. Array-based technologies are limited to the proprietary storage, and a customer with storage arrays from different vendors will find that the snapshots are not compatible across the arrays. Finally, snapshots can consume significant resources, which results in many customers limiting the number of consecutive snapshots that are maintained at any time. This means the customer is still at risk for data loss that occurs between the snapshots.

## Replication

Some customers have moved to replicating their data as a form of protection. This could be done locally, through mirroring, or remotely using separate replication technology that copies data from the local host to a remote data facility.

Regardless of the implementation, replication is inherently nondisruptive to the applications, highly reliable, and fast. It is designed to protect against physical failures, either of the storage array, for local replication, or of a site failure in the case of remote replication. However, it will not protect against logical data failures or corruption, as the corrupted data will be quickly replicated to the target. Additionally, most replication products do not provide the user with the ability to recover a previous version of the image, and would require users to implement snapshot capabilities on top of the replication if they want this feature.

## Continuous data protection

A continuous data protection product is designed to monitor changes to one or more data objects and store a copy of these changes in a separate location called a journal. This journal can then be used to re-create the object as it existed at any previous point in time. A CDP product is either file system-centric, where the object is a file, or storage block-centric, where the object is the volume or LUN.

Historically, the CDP file system approach started in the Windows community, where most applications utilized files to hold their data. Block-based CDP, on the other hand, started in the UNIX (and now Linux) community, where database applications traditionally bypassed the file system and operated directly at the disk/block level.

File system CDP products are typically found in Microsoft Windows environments, and usually offer a file system or a graphical explorer interface for their configuration and recovery operations. Recovery is end-user-driven and is usually provided through an extension to the file system (for example, when you right-click on a file or folder then click **Show previous versions**) or as a shared file system based on NFS or CIFS that can be mounted to a recovery server.

Block-based CDP operates as a layered feature of the underlying storage infrastructure, and usually operates independent of the host's file system and volume manager. Recovery is typically storage- or database administrator-driven, is provided through capabilities outside of the platform being protected, and is managed by the CDP implementation.

A CDP system also allows the user to establish write consistency between two or more objects that reside on different systems. For example, a database has two different objects, the files that maintain the database's data, and the files that maintain the database's logs. All databases will write to the log files before they commit the write to the data files. If a CDP product did not enforce write consistency between the two, a restoration of previous versions of the data and log files could result in a corrupted database. In this example, the administrator would identify the data and log files as part of the same consistency grouping to ensure that write order between the data and log files is maintained.

It is important to note that CDP systems deliver what is known as an "atomic" view of the data. All the data across all the disks is shown at exactly the same moment in time. It is as if time stopped at that exact moment. This atomic view provides consistency and stability across databases, applications, federations, and even entire data centers. CDP can dynamically re-create entire application environments without application involvement. In fact, the alternate view staging can be done on a completely different SAN or even in a separate geographic location.

CDP systems can re-create copies of the same time point repeatedly because the process is nondestructive to the underlying data. These instantly created time-based versions of data can easily be backed up using traditional, scheduled backup systems to removable media, replicated over distance or exported for alternative uses, all without impacting the ongoing online production applications.

Figure 1 displays the various data protection modes.

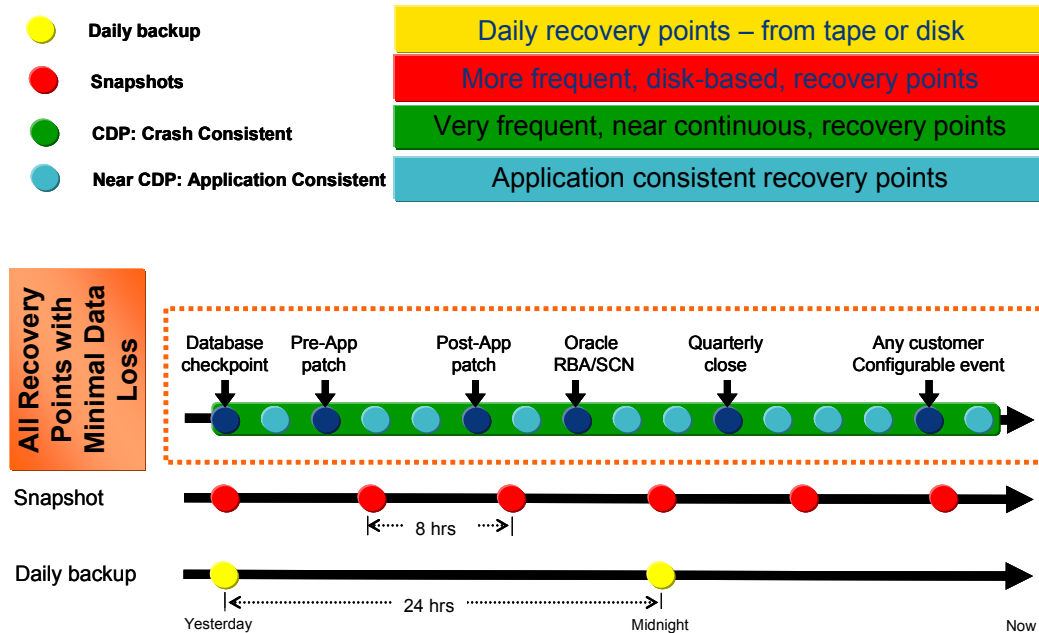


Figure 1. Data protection modes

There are a wide range of additional uses for time-stamped views. For instance, time-stamped views can be used as the source for loading a data warehouse, for cloning a production environment for use by a development and test organization, or for noninvasive backup. As an example, the customer can perform an audit on the data before a backup is performed – improving confidence in an effective recovery should the backup copy ever be needed. Obviously, if the backup failed, it would be a simple matter to re-create the image and restart the backup.

### Benefits of RecoverPoint

There are several features that RecoverPoint brings to the market. These include:

- **Support of heterogeneous storage and server environments**  
RecoverPoint supports a wide variety of possible storage and server environments to meet the needs of those customers that do not utilize a single vendor for their storage and server solution.
- **Awareness of applications and their environments**  
Application recovery is becoming more complex and time-consuming. Users should choose a product that integrates application consistency into the recovery process.
- **Simultaneous support for application recovery and disaster recovery**  
RecoverPoint allows customers to choose where and how their data is protected. Using RecoverPoint, data can be protected locally where it would be used for

application recovery, and it can also be protected remotely, where it would be used for disaster recovery. Unlike other products, RecoverPoint will protect data both locally and remotely with each copy having its own recovery point objectives.

- **Noninvasive to the application or server being protected**

RecoverPoint is designed to minimize the impact to a production application's I/O throughput or CPU load. This is done by keeping the RecoverPoint footprint on the application server to a minimum, and moving the replication processing to the external RecoverPoint appliance.

- **Out-of-band solution**

RecoverPoint is an out-of-band product. It resides in the SAN, but it is not involved in the flow of production data from the server to storage. Therefore, an unlikely failure in the RecoverPoint product will not affect application access to storage.

- **Scalable, reliable platform**

RecoverPoint is built on top of an industry-standard server platform. Each RecoverPoint appliance cooperates in managing the data protection processes for the site. The appliance is built on a high-availability architecture that ensures that the failure of a single appliance will not affect ongoing data protection operations. Additionally, RecoverPoint data protection and recovery performance can be scaled by adding additional appliances to the configuration when needed.

- **Supports business policies and service level agreements**

Companies assign different values to their different applications. RecoverPoint supports policies that allow differing RPOs and RTOs on a per-application basis for both the local copy and remote copy of the application data.

- **Can be extended by use of APIs/CLIs**

RecoverPoint has both an interactive and programming command line interface (CLI) that can be extended by customers. Additionally, RecoverPoint provides sample scripts that allow integration with traditional business applications such as Oracle Database Server, Microsoft Exchange Server, and Microsoft SQL Server.

- **Tightly integrated with business continuity technologies**

RecoverPoint provides both local and remote protection using the same software and appliance architecture. A customer that starts with RecoverPoint CDP can extend protection by adding RecoverPoint CRR for remote or concurrent local and remote replication.

## EMC RecoverPoint operations

RecoverPoint is designed to replicate changes at a block level on one or more SAN volumes (*source volumes*) residing in one or more storage arrays. It allows the

replicated targets (*target volumes*) to reside in one or more heterogeneous storage arrays.

RecoverPoint maintains transactional-consistent journals for each application (*consistency group*) defined within a RecoverPoint system. The journal allows convenient rollback to any point in time, enabling instantaneous recovery for application environments. Additionally, RecoverPoint also supports application integration that provides intelligent application-consistent recovery points for multiple third-party applications such as Microsoft Exchange Server or SQL Server.

## Managing RecoverPoint

RecoverPoint is managed using a Java-based GUI, which is called the RecoverPoint Management Application, or through an interactive SSH session using the RecoverPoint CLI. Figure 2 shows the main management application. This GUI is platform-independent and can be launched from any Internet browser running on a variety of Windows systems. The health and status of RecoverPoint can be viewed at a single glance and the user can quickly obtain details by clicking on an item of interest.

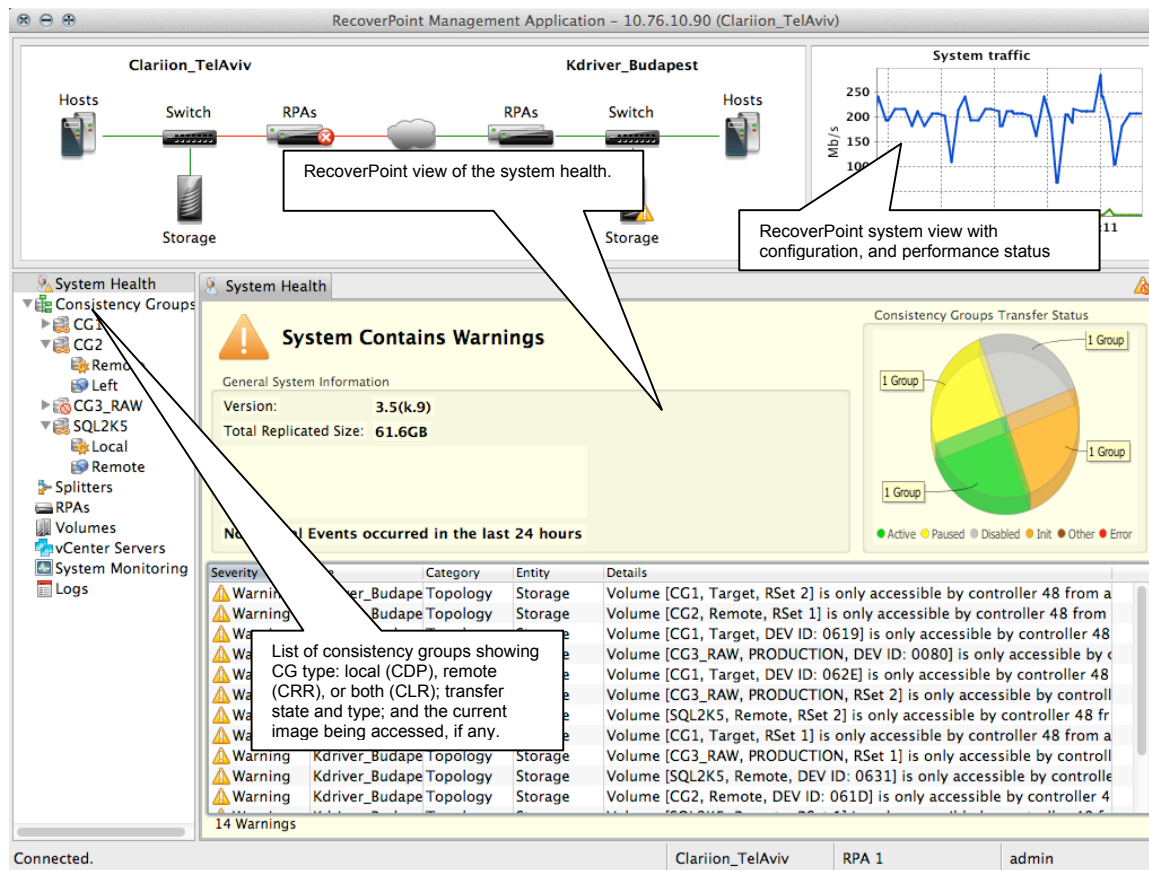


Figure 2. Using the RecoverPoint Management Application GUI

To automate RecoverPoint operations, use the RecoverPoint CLI. This CLI can be invoked in an *interactive* mode, or in a *programmatic* mode. Interactive mode offers automatic help, command option prompting, and automatic command completion.

Programmatic mode is designed to be used inside scripts where the inputs may be generated and outputs may be further processed by the script. The example in Figure 3 shows an interactive mode login using ssh.

```

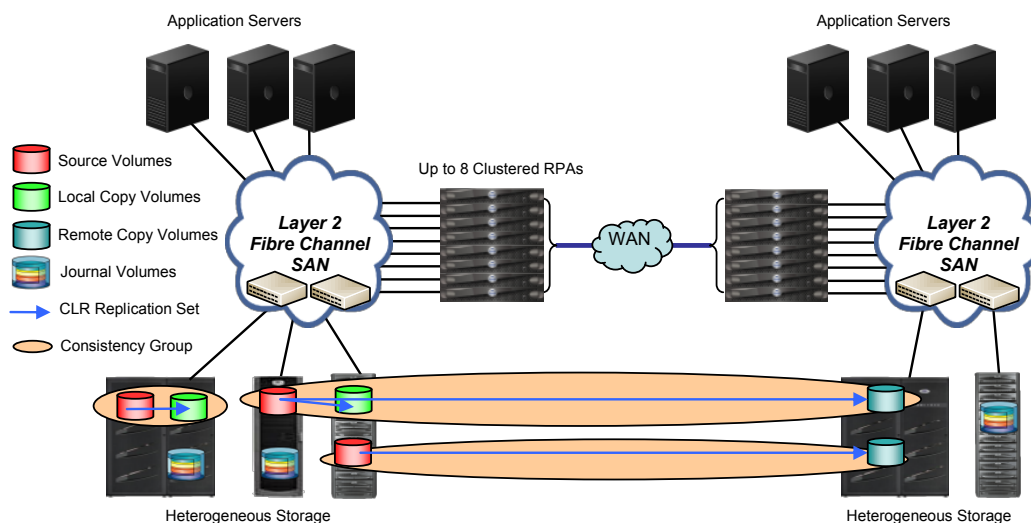
$ ssh admin@137.69.186.184
Password:
Last login: Thu Dec  3 16:03:15 2009 from 10.127.41.10
Site:
London:
RPAs: OK
Volumes: OK
Splitters: OK
New York:
RPAs: OK
Volumes: OK
Splitters: OK
WAN: OK
System: OK
[Test mode] New York>|

```

**Figure 3. Using CLI to manage RecoverPoint**

### System architecture

The specific components of EMC RecoverPoint are shown in Figure 4. This configuration shows the source, local copy, and remote copy volumes residing across multiple storage arrays in two different sites. There are three types of replication shown: local replication using CDP, remote replication using CRR, and both local and remote replication using the combination of CDP and CRR called concurrent local and remote. Details on the components are described in the next section.



**Figure 4. EMC RecoverPoint architecture**

## RecoverPoint appliance

EMC RecoverPoint is mainly software, however the RecoverPoint appliance (RPA) runs the RecoverPoint software on top of a custom 64-bit Linux kernel inside a secure environment built from an industry-standard server platform. An RPA manages all aspects of data protection for a storage group, including capturing changes, maintaining the images in the journal volumes, and performing image recovery. Moreover, one appliance can manage multiple storage groups, each with differing policies.

There are at least two active RPAs per site that constitute a RecoverPoint cluster. Physically, the RecoverPoint cluster is located in the same facility as the host and storage subsystems. All RPAs in a cluster have identical functionality. In normal operation, all RPAs are active all of the time. Consequently, if one of the RPAs in a cluster goes down, EMC RecoverPoint supports immediate *switchover* of the functions of that appliance to one or more of the remaining RPAs.

## Write splitter

The host-based write-splitter driver (KDriver) is system software installed on all hosts that access the volumes that contain the protected data. The primary function of the KDriver is to “split” or “mirror” an application’s write traffic, so that the written data is sent to both the production storage volumes and to the RPA. The KDriver carries out this activity efficiently, with little perceptible impact on host performance, since all CPU-intensive processing necessary for data protection is performed by the RPA.

The write-splitter functionality is included in the Symmetrix VMAX series, VPLEX Local, VPLEX Metro, VNX series, CLARiiON CX3 UltraScale, and CX4 UltraFlex™ line of storage arrays, where it carries out the write splitting inside each storage processor. Alternatively, the *splitter* function can be carried out inside an intelligent fabric, such as provided by the Connectrix® AP-7600B switch using the Brocade SAS API. Intelligent fabric is also supported with the Connectrix MDS 9200 switch family or the Connectrix MDS 9500 director family with the SSM or 18/4 module and the SANTap protocol. When the splitter function is carried out by an array-based write splitter or by intelligent fabric-based write splitter, a host-based KDriver is not required.

## Repository volume

The repository volume is a SAN-attached volume used only by RecoverPoint. The repository volume is used to maintain the configuration and communication between RPAs in a cluster. Similar to a cluster server’s quorum volume, the repository volume contains the status of the overall RecoverPoint system and acts as a resource arbitrator during RPA failover and takeover operations. There is no user-accessible information stored on the repository volume.

## Journal volume

Journal volumes hold data waiting to be distributed to target volumes and also retain copies of the data previously distributed to the target volumes. Each consistency group has its own journal volumes that allow for differing retention periods across

consistency groups. Each consistency group has two or three journal volumes, one assigned to the local copy volumes (if present), one assigned to the remote copy volumes (if present), and one assigned to the production or source volumes. Journal volumes are used for the source and both copies in order to support production failover from the current active source volume to either the local or remote copy volume.

Storage efficiency is realized in the history journal by retaining only the changes between journal entries. Additionally, the journal volume is also compressed, resulting in even more storage savings. Source and target images on the CDP volumes are always consistent upon completing the distribution of each write change.

### Consistency groups

Figure 5 shows the architecture of RecoverPoint consistency groups. Volumes are grouped into consistency groups, and replicated to target volumes in the local and/or remote site. Writes to source volumes will be replicated in the same order on the local and/or remote volumes, which ensures that transactional consistency is maintained.

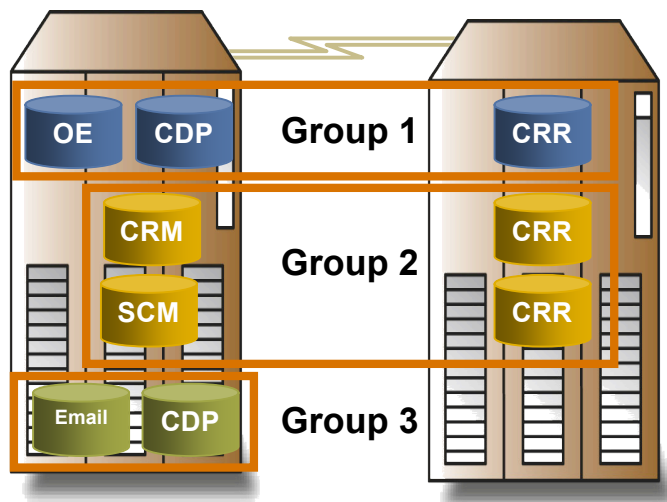


Figure 5. RecoverPoint consistency groups

If one of the source volumes in a consistency group goes offline, such as what may happen in a *rolling* disaster where a partial loss affects only a few of the source volumes, then the replication for all volumes in the consistency group will be paused until the source volume comes back online. In this way, consistency groups prevent dependent writes from getting out of sync, thus ensuring the integrity and consistency of the data at the remote site.

### Consistency group operations

A consistency group is made up of one or more groups of volumes. These volumes, called replication sets, identify the production copy, local copy, and/or remote copy volumes for replication operations. The production copy and local copy (if present) reside in the same array or SAN and are protected using true-CDP technology. The

remote copy (if present) resides in a different SAN and is protected using near-CDP technologies. The only requirement for a replication set is that the member volumes are the same size;<sup>1</sup> however, they can span different arrays from different manufacturers and may have different geometries or performance characteristics.

When creating a consistency group, the user creates one or more named replication sets and assigns volumes to the set from the pool of available volumes. Either two or three journal volumes are also defined: one for the production copy, one for the local copy (if present), and one for the remote copy (if present). These journals are used to maintain in-flight replicated data and also operate as a history volume, allowing the local and/or remote copy volumes in a consistency group to be rolled back to a previous point in time.

Each consistency group also has a set of policies (shown in Figure 6) used to optimize the storage and network resources by managing replication lag, data compression, and WAN bandwidth prioritization.

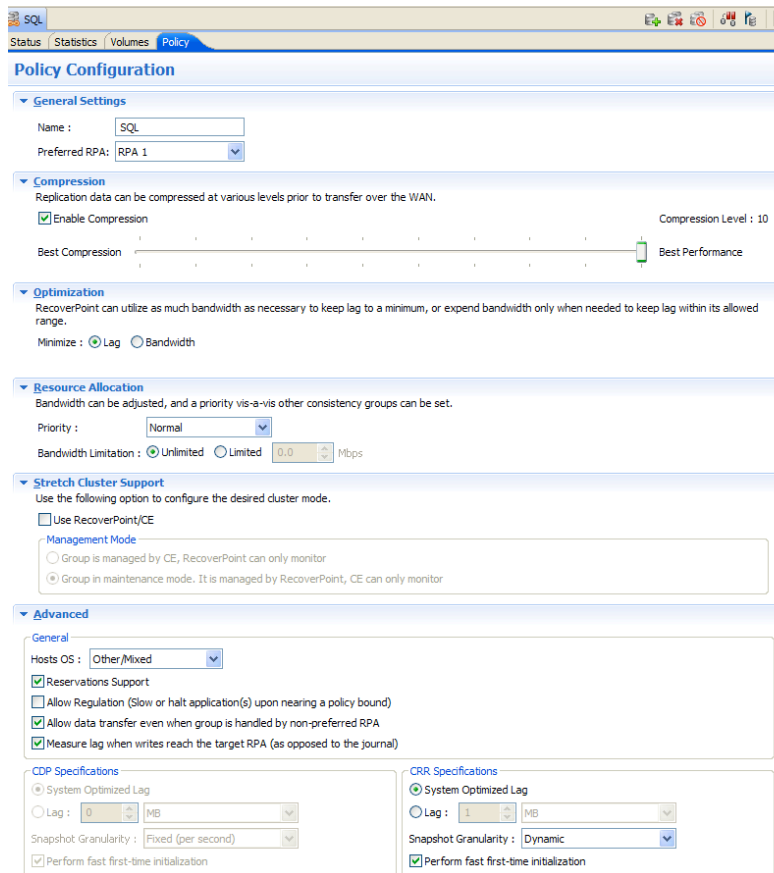


Figure 6. RecoverPoint consistency group policy settings

<sup>1</sup> A target replication volume can be slightly larger and/or have a different geometry; however, replication will be limited to the block count on the smallest volume of the pair. Some operating systems, such as Solaris, require identical geometry. Consult the *EMC Support Matrix* for further details.

## Application servers

RecoverPoint is an out-of-band, block-level local and remote replication solution. As such, it is agnostic about the server operating system, file system, and volume management environment. RecoverPoint is designed to support heterogeneous server platforms, including many of the common Linux, UNIX, and Windows operating systems. Additionally, RecoverPoint also supports VMware® ESX® Server for protection of both the ESX Server physical volumes and all virtual machines with their data.

## Heterogeneous storage arrays

RecoverPoint replicates data between one or more arrays in the same data center or between two data centers as shown in Figure 7. RecoverPoint provides data protection with guaranteed consistency across a data center's heterogeneous server, network, and storage environment, protecting and optimizing an organization's valuable infrastructure environment. Additionally, IT organizations can implement a tiered storage infrastructure built from lower-cost storage. This infrastructure can be used to protect critical application environments running on more expensive storage.

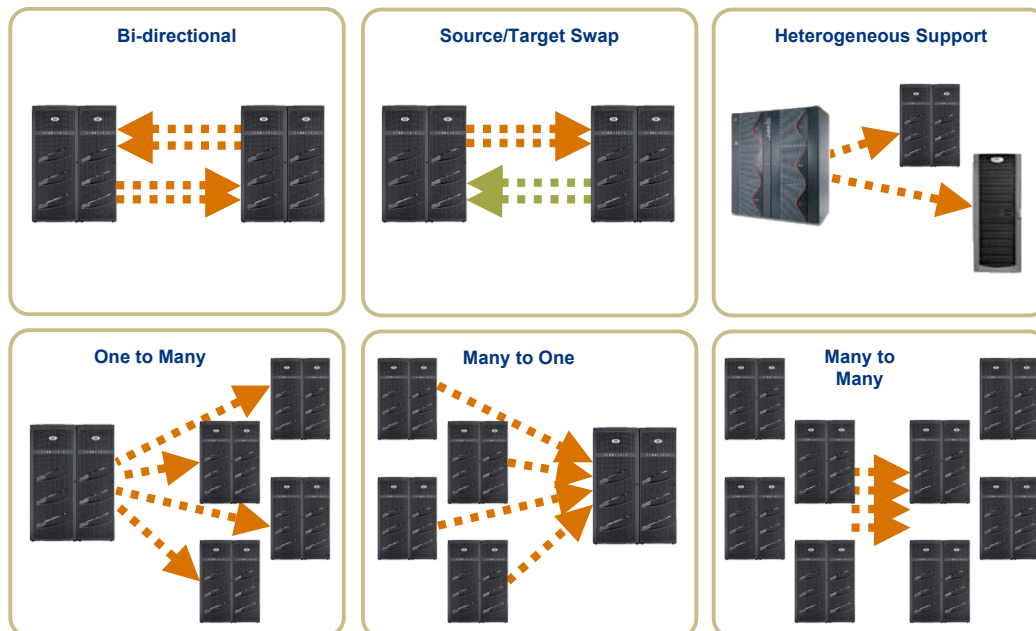


Figure 7. Current RecoverPoint deployment options for array-to-array replication

## SAN considerations

RecoverPoint configurations are required to have each appliance connected to the host and storage through separate (A/B) SANs or fabrics. Each RecoverPoint appliance has four Fibre Channel ports: half of the ports are connected to the A-fabric, and the other half are connected to the B-fabric. For a minimum two-node RecoverPoint cluster, two Fibre Channel ports from both appliances will be connected

to the A-fabric, and the other two FC ports from both appliances will connect to the B-fabric. Hosts should have multipathing capability, which can be supplied by EMC PowerPath® or other qualified multipathing products, and any storage array should have its LUNs presented through at least two ports on the array.

RecoverPoint supports any of the Layer-2 SAN environments present on the *EMC Support Matrix*. This includes SAN infrastructures based on Brocade and Cisco configurations. For SANs based on Brocade or Cisco, RecoverPoint also supports intelligent-fabric protocols for mirroring the writes between protected volumes. For Brocade, RecoverPoint supports the SAS API provided in the Connectrix AP-7600B switch. For Cisco, RecoverPoint supports SANTap in the Connectrix MDS 9200 series switch or MDS 9500 series director.

## Preserving data integrity during failure scenarios

RecoverPoint is designed to preserve data integrity in the event of a variety of failure scenarios. RecoverPoint leverages consistency groups that enable it to:

- Preserve the data integrity and consistency across all members of the consistency group
- Minimize the amount of time that data is unavailable to the end users

To maintain data consistency across replication pairs in a consistency group, RecoverPoint software protects data integrity in the event of the following failures:

- RecoverPoint appliance failure
- Lost connectivity to one or more source volumes
- Lost connectivity to one or more local or remote copy volumes
- Host connectivity failure or host splitter failure

## Managing RecoverPoint appliance failure

A minimum of four RecoverPoint appliances are required for RecoverPoint CLR, two in each site. EMC does not support single-appliance RecoverPoint configurations. RecoverPoint appliances operate as a highly coupled, active/active, shared nothing cluster. The failure of a RecoverPoint appliance will be detected and its workload will be automatically distributed among the remaining RecoverPoint appliances without data loss.

When the appliance fault is corrected or a replacement appliance is brought online, the workload will be automatically returned to the original appliance. This capability allows for a nondisruptive upgrade of the RecoverPoint appliance. When an appliance is taken offline for maintenance, the remaining appliance will take over its operations. With maintenance completed, the recently serviced appliance is added back into the RecoverPoint cluster to resume operations.

## Managing lost connectivity to source volumes

A RecoverPoint consistency group will maintain data consistency across multiple volumes that reside on one or more local storage systems. If one of the volumes on the local system becomes unresponsive, this will be detected by the host splitter, the CLARiiON splitter, or the fabric splitter. The RecoverPoint appliance will halt the local and remote replication processes for the affected consistency group, which leaves the local and remote copy volumes in a fully consistent and usable state. The user can either correct the fault or remove the affected replication set from the consistency group so that other replication sets can update successfully. Once the failure is resolved, the administrator can restart the replication process.

## Managing lost connectivity to local or remote copy volumes

A RecoverPoint consistency group will maintain data consistency across multiple volumes that reside on one or more local storage systems. If one of the local target volumes fails, then CDP protection for all of the volumes in the storage group will be paused until the faulty volume is repaired. The CRR protection for all of the volumes will not be affected. If one of the remote copy volumes fails, then the CRR protection will be paused until the faulty volume is repaired, however, the CDP protection will continue. If connectivity to the SAN for the production host fails, and a host-based splitter is being used, then the host's volumes cannot be monitored for changes. In both cases, the consistency of the mirrored volumes will be unaffected, and CLR operations will continue for the remaining replication sets in the consistency group. When connectivity is restored, the affected local or remote copy volumes will be automatically resynchronized with the host volumes.

## Managing host or write-splitter failures

If connectivity to a local host fails when using a host-based write splitter, then the volumes accessible by the failed host will not be monitored for changes. However, the consistency of all RecoverPoint local and remote copy volumes will be unaffected, and RecoverPoint operations will continue for the remaining volumes in the consistency group. If the host splitter crashes, then the host goes into an unmanageable state, and the RecoverPoint appliance can be configured to pause the protection process for all members of the consistency group or continue protection for the unaffected volumes.

## Utilizing VNX series, CLARiiON, and Symmetrix operations with RecoverPoint

### Layered array operations overview

Both CLARiiON and Symmetrix provide array-specific capabilities for array-based volume copies (SnapView and TimeFinder) as well as array-to-array replication (VNX MirrorView™ and SRDF®). In general RecoverPoint production volumes can be the source for VNX, CLARiiON, and Symmetrix/DMX layered array snapshot and replication operations with a few considerations detailed in the following sections. In addition, the RecoverPoint local and/or remote replica copies can be the source for

array-based snapshot operations performed by TimeFinder<sup>2</sup> and SnapView. RecoverPoint's heterogeneous local and remote data protection with CDP-based rollback capabilities are very much complementary to array-based snapshot and replication facilities.

### **Using EMC VNX, CLARiiON, and Symmetrix array features with RecoverPoint**

MirrorView, SnapView, SRDF, and TimeFinder along with RecoverPoint can use the same production volumes as both solutions may fully coexist with each other. The one caveat is that the user must ensure to never “reverse” the replication direction in one of the products (for example, performing a reverse synchronization with SRDF) without disabling the other product. If you use SRDF or TimeFinder then you will require an RPQ.

#### ***Performing operations on RecoverPoint production copies***

There are many ways that RecoverPoint can be combined with native array capabilities. A common request from customers is to combine array-based replication with local protection using RecoverPoint CDP. For example, the production volumes for an Oracle application can be protected locally using RecoverPoint CDP and synchronously replicated using MirrorView/S. All writes by the Oracle application to the protected volumes will be intercepted by the RecoverPoint splitter and a copy will be sent to the RecoverPoint appliance. Additionally, the write will be intercepted by MirrorView/S and synchronously replicated to the MirrorView/S target.

If the production volume needs to be re-created by MirrorView or another array-based operation, it is necessary to disable the specific RecoverPoint consistency group that contains the affected volumes before re-creating those volumes. For example, if the Oracle production volume needs to be re-created by MirrorView, it will be necessary to disable the specific RecoverPoint consistency group that contains the affected volumes before re-creating those volumes. When this is completed and MirrorView resumes normal operations the user enables the RecoverPoint consistency group. At this point RecoverPoint will perform a full sweep of the production and replica copies to bring them back into sync.

#### ***Performing operations on RecoverPoint replica copies***

RecoverPoint supports the use of TimeFinder and SnapView on local and remote replica copies. To use one of these layered array operations, it is first necessary to access the point-in-time image as a physical volume hosted by the appropriate array. The user accomplishes this by selecting “enable image access” for the specific consistency group copy that contains the selected volumes. The user then selects an appropriate point-in-time image or consistent bookmark from the RecoverPoint journal requesting logged access as the image access mode. This causes RecoverPoint to re-create the physical volumes as they existed at the requested point in time.

---

<sup>2</sup> If you are using a RecoverPoint LUN as the source for a TimeFinder or RDF operation then an RPQ will be required.

Once the volume is rolled back to the point in time and is visible in the SAN, it becomes possible to use a layered array operation, such as TimeFinder or SnapView, to create a BCV or fracture a clone. Once the BCV or clone is created or fractured, the RecoverPoint image is no longer needed. At this point the user may select “disable image access” for the consistency group copy, which causes RecoverPoint to resynchronize the replica copies. The BCV or clone LUNs can now be used for any purpose without any impact to RecoverPoint. If the BCVs or clones need to be resynchronized, the user repeats the “image access” steps, selecting a different point-in-time copy. Once the image is available, the user can resynchronize the BCV or clone and then split or fracture it off again.

## RecoverPoint use-case scenarios

This section covers multiple use-case scenarios for using EMC RecoverPoint for local and remote operations and disaster recovery.

### Case 1: Storage migration

RecoverPoint can be used to provide an application and host independent data migration between storage arrays, to redeploy storage volumes, or to consolidate from heterogeneous storage arrays. RecoverPoint consistency groups ensure that the data at the target volumes represents an identical copy of the production volumes. Additionally, RecoverPoint will perform both the initial copy as well as maintain the synchronization between the source and target volumes using SAN resources without impacting the production application.

By keeping the source and target volumes continuously synchronized, the migration of the application server to the new storage can be performed during prescheduled application maintenance windows. Finally, RecoverPoint adds a second level of protection in case of a migration or storage failure by having the target volume and the journal available as a real-time backup of the production data. Using RecoverPoint to migrate data between arrays will speed the migration process and maximize application availability when compared to other solutions.

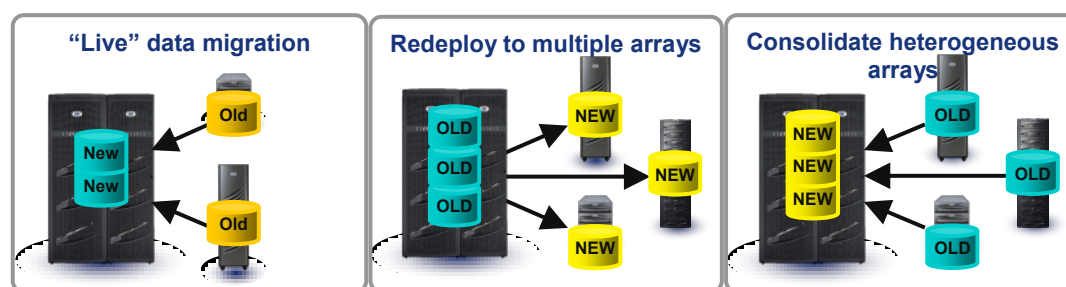
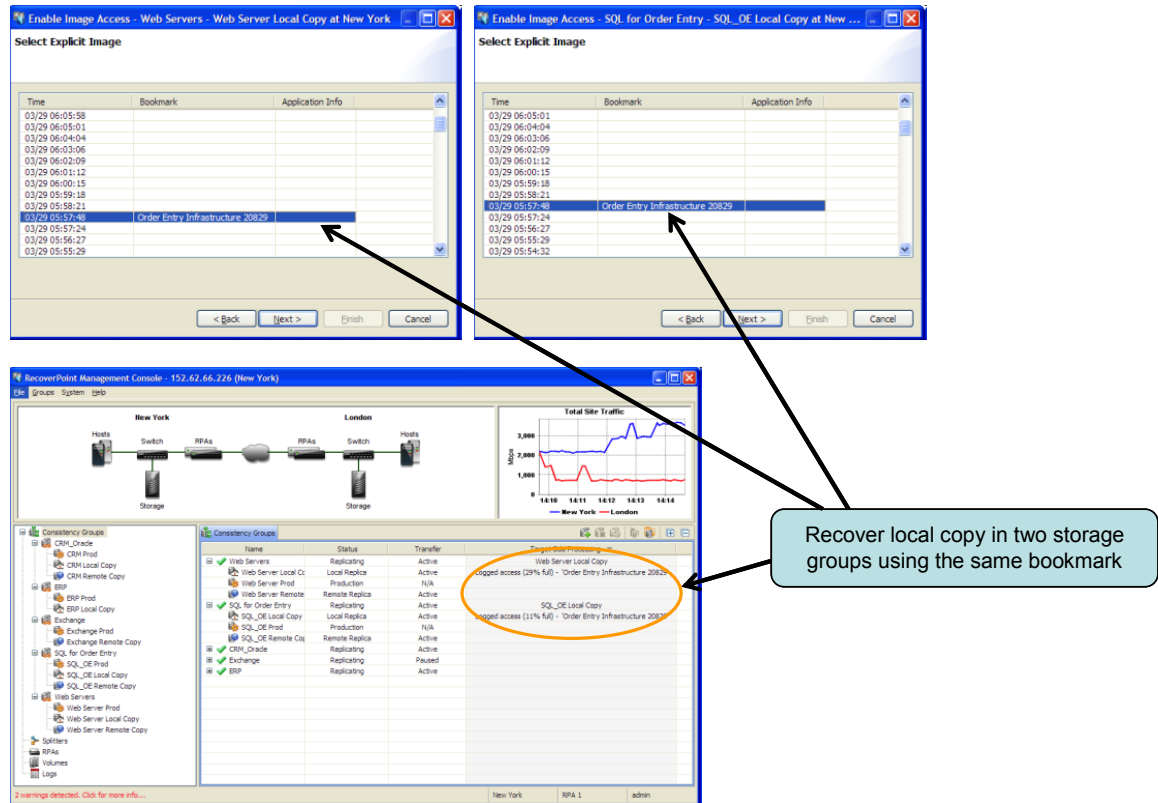


Figure 8. Deployment scenarios

## Case 2: Repurposing for development and test support

RecoverPoint capabilities can be used to quickly repurpose a point-in-time mirror image by taking a snapshot of the application's data while the application remains online. The snapshot of the image can be used for a variety of purposes, such as development and testing of product updates, data mining, or compliance.



For multiple consistency groups, a group set would be created that periodically tags each consistency group's image with the same bookmark. When taking the snapshot, each consistency group would be recovered back to the same exact point in time using the RecoverPoint management GUI and selecting the appropriate local or remote copy and the identical bookmark for each storage group, shown below.

### Figure 9. Using group sets for application data repurposing

Once the images have been recovered back to the appropriate point in time, the application data would be recovered with database consistency and integrity ensured. At this point a snapshot can be taken. The snapshot can be mounted on the development and test systems and can be used for reads, writes, and changes, or they could be copied to a different system using traditional SAN- or WAN-based tools.

Application data repurposing is a natural extension of the RecoverPoint data protection capabilities. Application production will not be affected by the repurposing process, which enables the recovery and repurposing of a concurrent, consistent copy of the production data. Additionally, RecoverPoint's ability to recover to any point in time means that the copy can be refreshed at any point with data from earlier or later points in time, without impact to the production application.

Finally, when using the RecoverPoint's logged access to clone volumes, the physical cloned volumes can be used by array-specific applications such as EMC's TimeFinder or SnapView to further manage the images as shown in Figure 10.

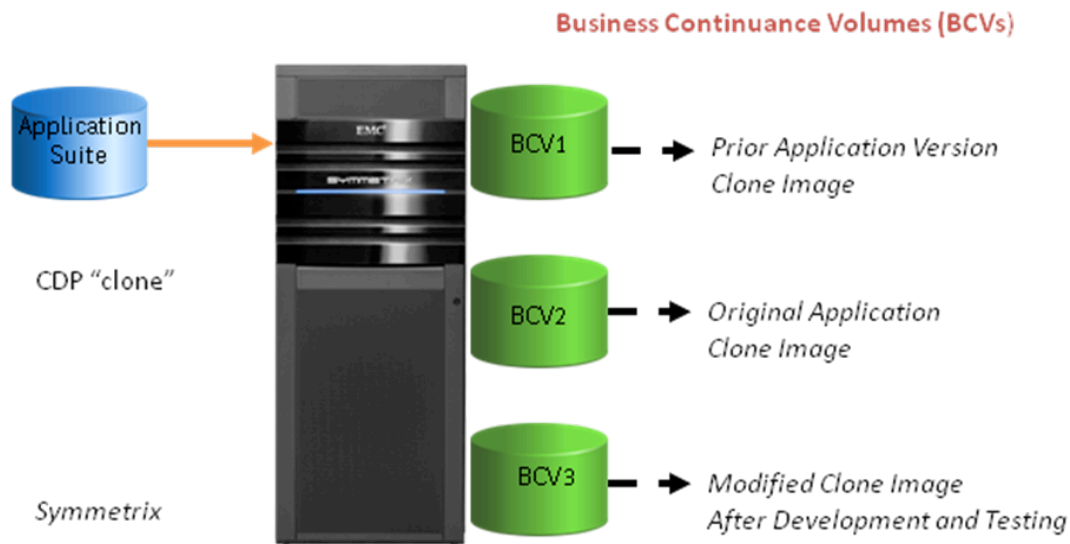


Figure 10. Using TimeFinder clones with RecoverPoint

### Case 3: Simplifying backup, restore, and disaster recovery

RecoverPoint enables a simplification of the backup, restoration, and disaster recovery for many applications, such as Microsoft Exchange. By utilizing RecoverPoint consistency groups, all of the volumes associated with Exchange, including the Exchange Storage Group Volumes and Exchange Log Volumes, are managed with the same policy and their recovery information is stored locally and/or remotely in the history journal. Using the Microsoft VSS API through the RecoverPoint KVSS utility ensures that recovery of Exchange is fully supported by Microsoft.

Using RecoverPoint ensures no disruption to the production database when making replica copies of the production volumes. Additionally, all volumes are fully consistent, with write-order fidelity maintained at all times. This ensures that the replica volumes can be used for a variety of purposes, such as fast mailbox recovery, backup to tape with no impact to production, and compliance readiness testing.

The RecoverPoint history journal, along with RecoverPoint's virtual access, enables near instantaneous recovery of critical data locally at any point in time, and/or remotely at significant previous points in time. This ensures that a company will be able to maintain their applications in case of a local data corruption or a regional disaster that impacts the production site.

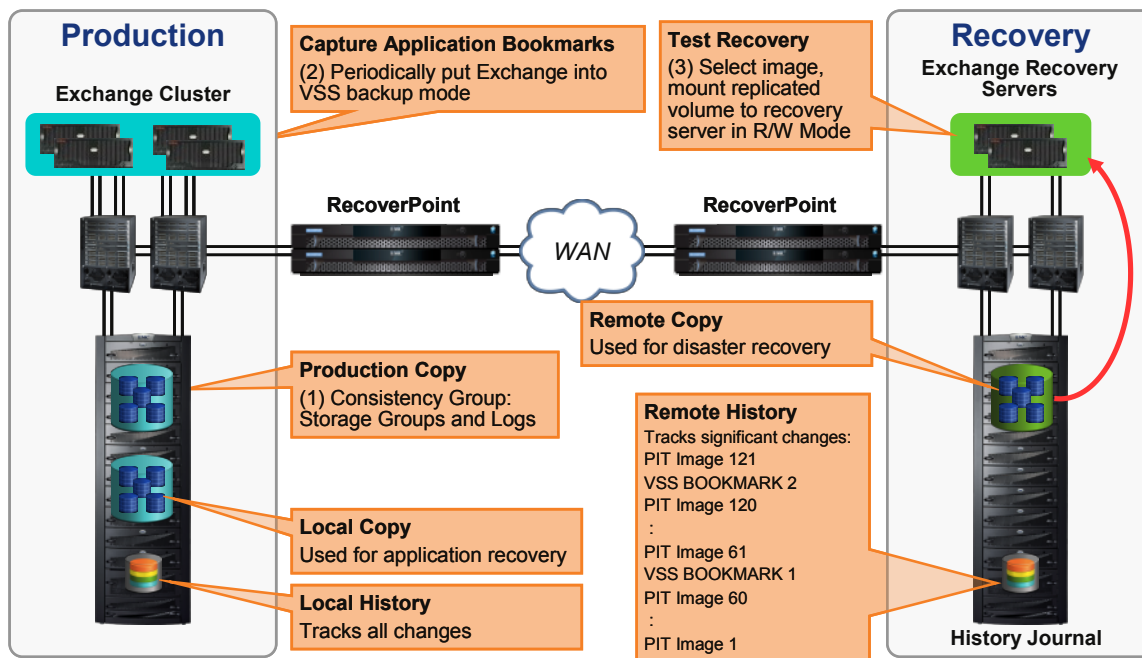


Figure 11. Exchange protection on a CLARiiON with RecoverPoint

#### Case 4: Individual file or folder level recovery

Data corruption of a file or folder at the production site may occur as a result of human error, a rolling disaster, or a machine failure. RecoverPoint CDP protects data at the block level, recording every write into the history journal. Since RecoverPoint is a block-level product, it is unaware of the file system or volume contents so the administrator must use a recovery server to identify the appropriate recovery point and extract the files from the recovered image and move them back to the production server.

It may be difficult to determine the exact moment at which data corruption began. To help find when the trouble started, the RecoverPoint logged image access feature facilitates testing to identify the recovery point created just prior to the data corruption. Once identified, the recovered file or folder can be copied back to the source site using external means (such as FTP, copying it to an external media device, or any other means available).

One of the benefits of recovery using logged image access is that the production data continues to be protected during the recovery, and normal application operations are not affected by the recovery process. To recover from data corruption the following steps would be followed:

1. Access an image near the time at which data corruption is believed to have first occurred, as shown in Figure 12.

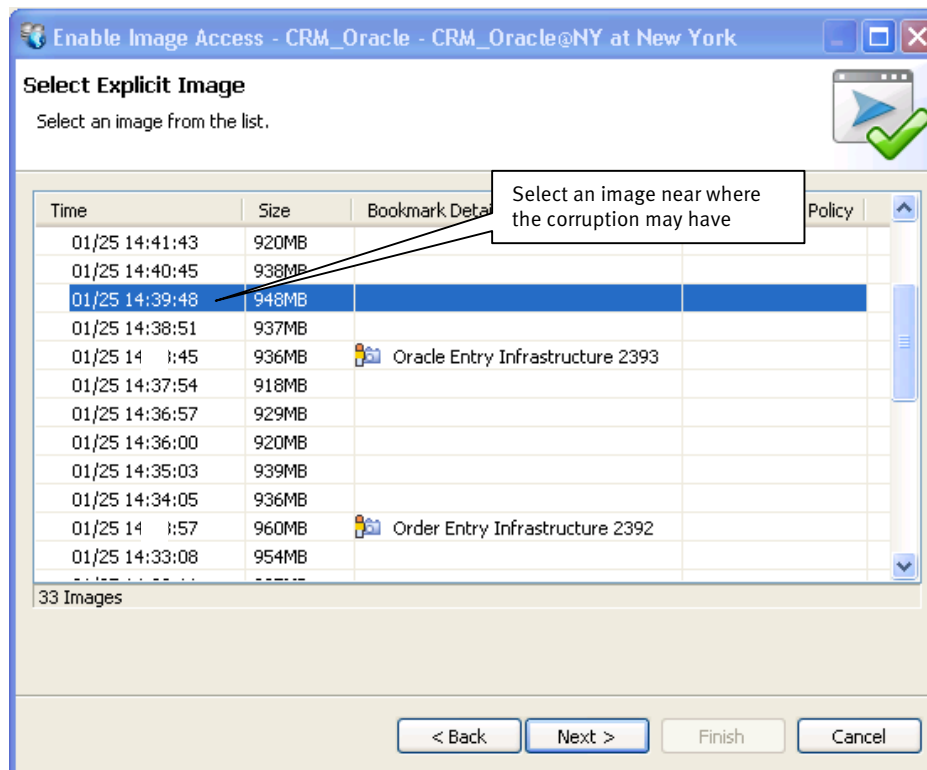


Figure 12. Selecting the image for recovery

2. Mount each of the volumes in the consistency group onto a recovery server. Using RecoverPoint virtual access, the volume(s) will be instantly available and visible in the SAN.
3. On the recovery server, test and verify that the data is the correct version and that it is not corrupted on this image. If corrupted, recover and test an earlier version until a valid image is identified that occurred shortly before the beginning of the corruption on the production volume.
4. After recovering a valid image, copy the recovered data to the production system. This can be done by mounting a share from the production server and copying the files and folders to the share. Alternatively, a removable media device, such as a USB key, can be used to transport the files to the production server.
5. Once the files are copied over, shut down the recovery server and then restart the mirroring between the source and target volumes. Since virtual access was used, both the source and target volumes remained synchronized during the recovery process.

### Case 5: Federated application recovery using group sets

One of the key components of data consistency is the ability to restart applications. As configurations become more complex, it becomes harder to protect the consistency of data. However, with RecoverPoint, data consistency and the ability to restart applications are assured.

In this scenario, there are three interrelated applications running across different servers and different storage systems that provide Internet services, including order entry and customer records management (CRM). The system is comprised of a front-end set of web servers running IIS on Microsoft Windows, a midtier server running Microsoft SQL Server, and a back-end server that runs the Oracle CRM database. Each of these systems is set up in a different consistency group as shown in Figure 13.

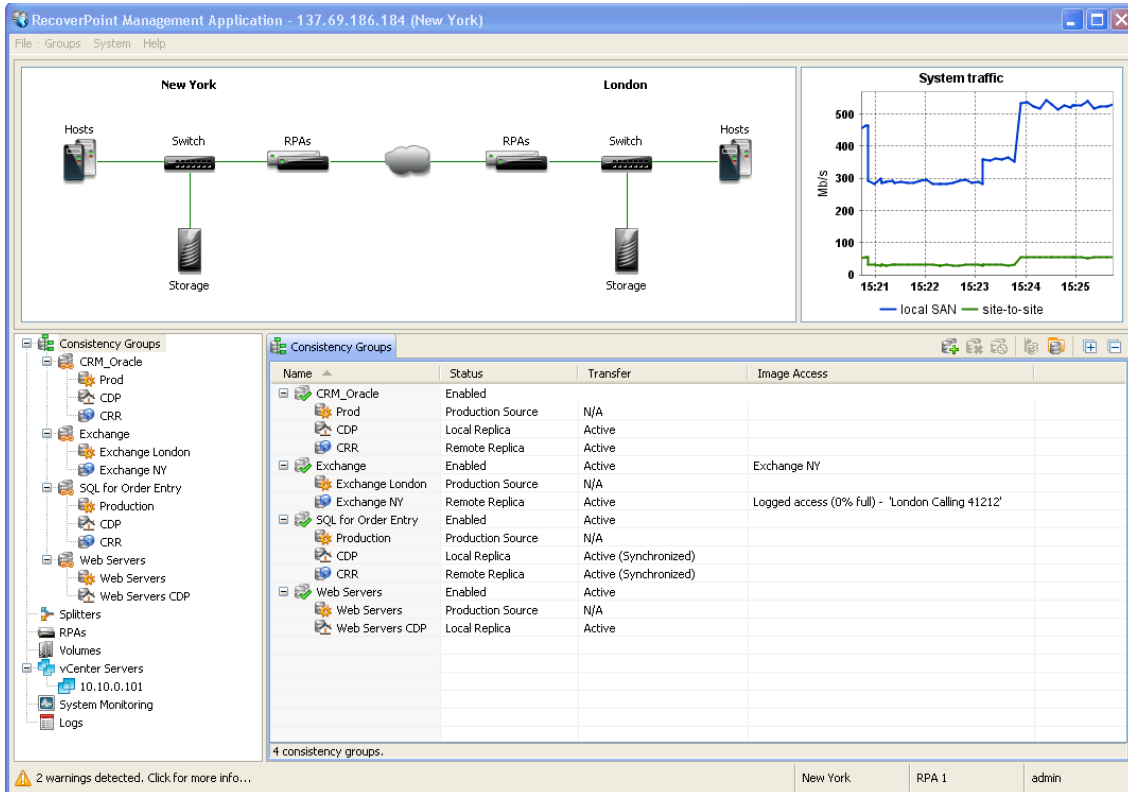


Figure 13. Management Application showing the active groups

Each application is protected using policies that support the differing RPO and RTO requirements for the individual applications. However, the RecoverPoint group sets capability can be used to periodically synchronize the recovery points across all three applications.

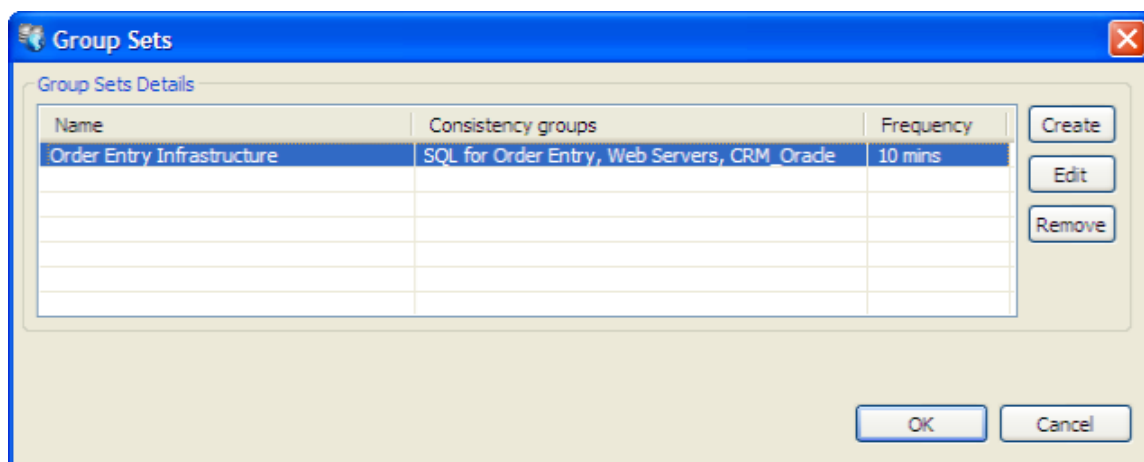


Figure 14. Group Sets window

A group set ensures that each of the included consistency groups (here Web Servers, SQL for Order Entry, and CRM\_Oracle) will be brought into consistency every few minutes and an appropriately named bookmark entry will be created for each group. The production copies for each consistency group in a group set must reside on the same site. The group set members can have differing target copies, such as local, remote, or both.

### Federated recovery

Using group sets makes it easy to identify the correct image to access across all members. In the example in Figure 15, the bookmark **“Order Entry Infrastructure 2393”** is being selected as the image to access for each of the three groups.

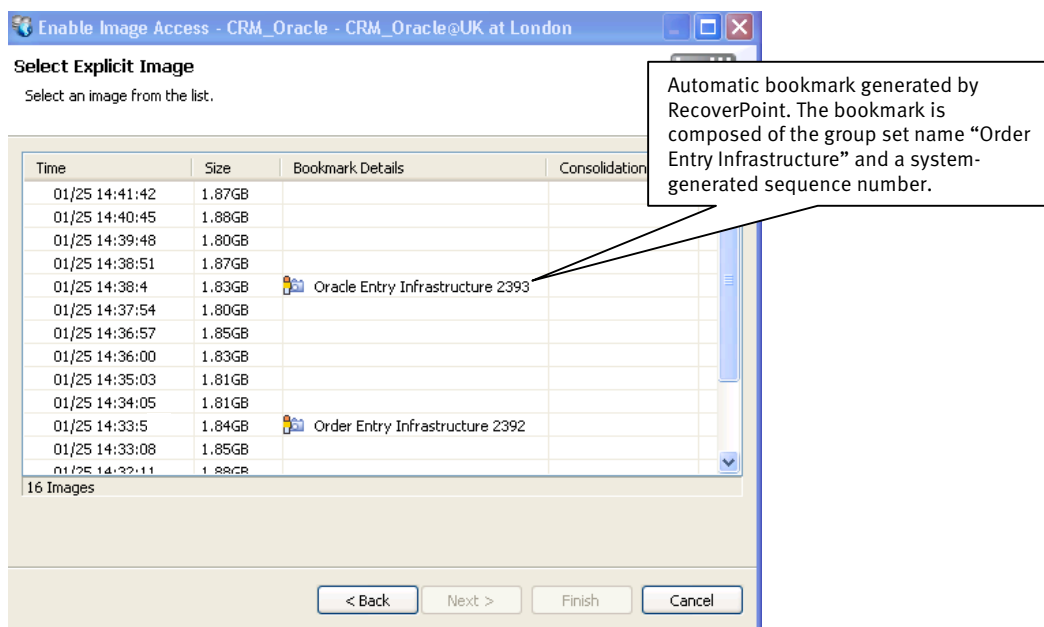


Figure 15. Selecting the appropriate image by a group set bookmark

When completed for all three groups, the RecoverPoint Management Console screen (Figure 16) will show that all three have been rolled back to the same target image “Order Entry Infrastructure 20823”. At this point, all volumes at the disaster recover site required for these three applications are fully consistent, and the application servers can be restarted without any data loss or inconsistencies among the three applications.

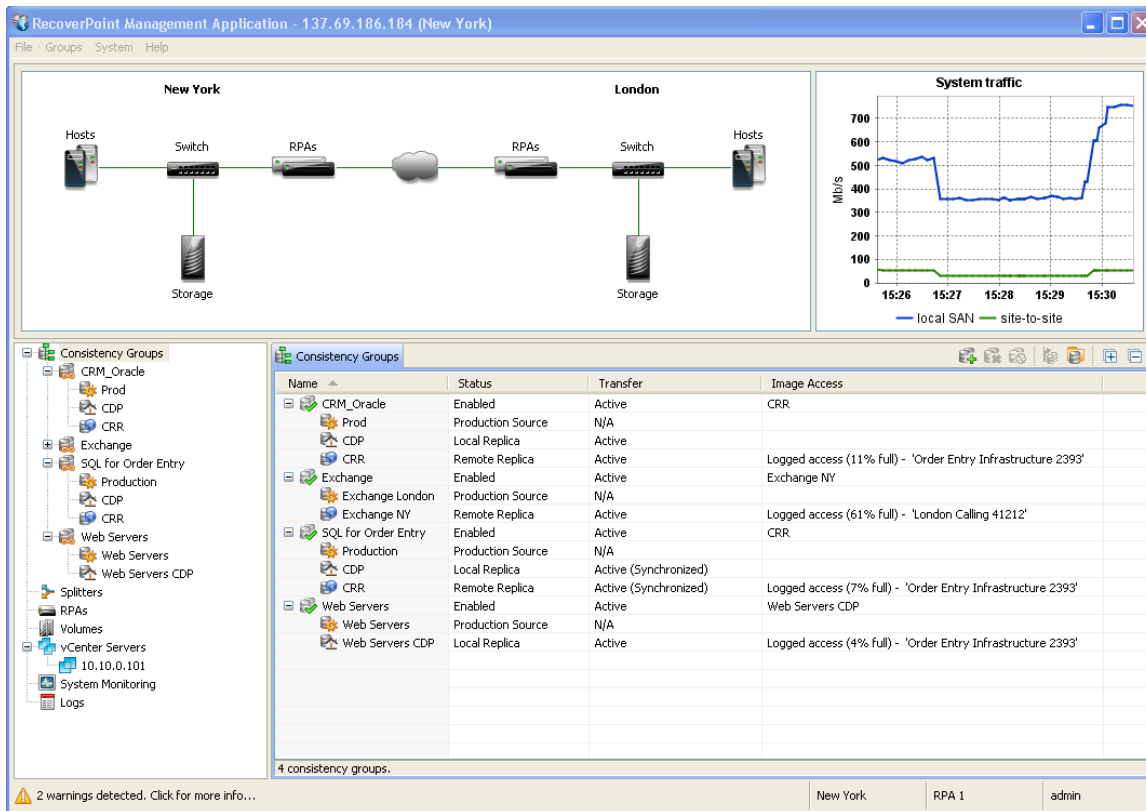


Figure 16. The RecoverPoint Management Application showing logged access of the same point in time for three applications

### Case 6: Performing recovery testing fire drills

RecoverPoint enables you to conduct comprehensive and regular tests of your application or disaster recovery data without any disruption to your production servers or any interruption to your data recovery protection. These tests, called fire drills, are important to ensure the integrity and effectiveness of recovery capabilities in light of the frequent server, system, and network changes that occur over time.

RecoverPoint’s physical recovery capabilities along with target-side processing allow administrators to perform a real test of the recovery data by actually running the applications and modifying the recovered data, all without impacting the production environment. Additionally, RecoverPoint enables testing without leaving production systems vulnerable during testing, without disrupting production application

availability in any way, and without the need for expensive data resynchronization after testing.

Using the RecoverPoint GUI or programmable CLI, the user selects the storage group to be tested and picks the copy to be accessed, either the local copy or the remote copy. The user selects an appropriate point-in-time image using the time and date stamp, bookmark, or other application-specific information. Logged access (physical) is then selected, which instructs RecoverPoint to rewind the data on the physical local or remote copy volumes back to the appropriate point in time. Finally, any data currently being replicated from the production side during this process is logged in to the history journal for the appropriate copy volumes but not written to the volumes.

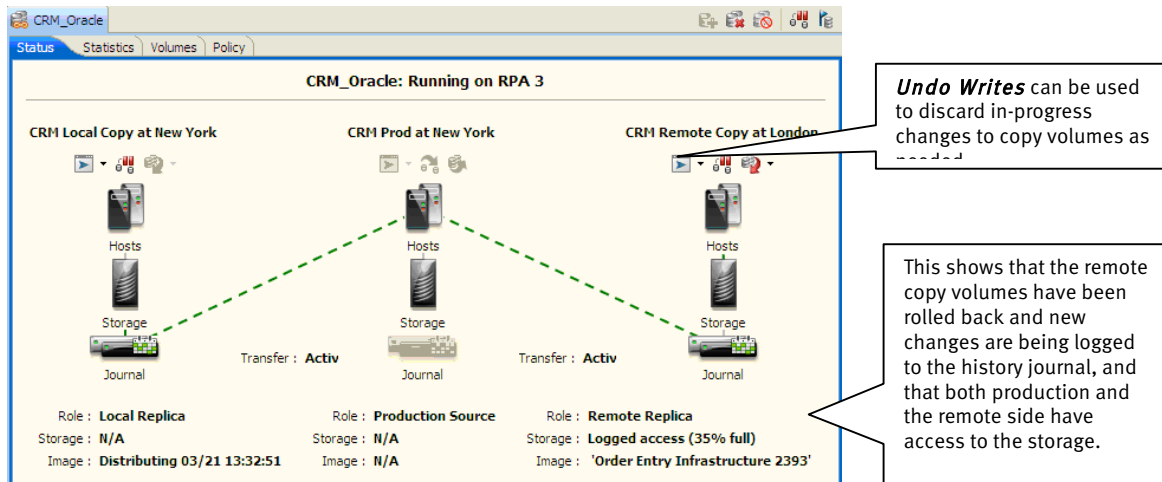


Figure 17. Selected image for fire drill testing

When RecoverPoint completes the rewinding, the local or remote copy volumes are unmasked and presented to the SAN where they can be mounted as read/write and used by one of the recovery application servers. All reads and writes operate directly against the copy volumes and can employ any advanced features of the array, such as mirroring, snapshot cloning, and SAN backup. During the testing process, it may be necessary to reset the target volumes back to their original state – this can easily be done by selecting the undo writes option, which will remove all writes to the volumes by applications at the disaster recovery site.

After disaster recovery validation testing, the testing images can be discarded without requiring a costly resynchronization with the production volumes. Instead, RecoverPoint will undo all changes to the target volume and will write all the pending changes that were logged in to the history journal.

RecoverPoint enables more frequent disaster recovery testing using images of the production data in different disaster recovery scenarios. Users can be confident of a successful recovery in the case of a production site disaster.

## Conclusion

EMC RecoverPoint offers both continuous data protection and continuous remote replication technology for heterogeneous SAN-attached storage. With its support for any-point-in-time recovery, RecoverPoint is designed to protect critical business processes and improve operational recovery with minimal impact to production environments. Through the use of consistency groups, RecoverPoint provides local and remote with write-order consistency for protected volumes that can span multiple heterogeneous storage systems and servers.

## References

More information on EMC RecoverPoint can be found at the [RecoverPoint page](#) on EMC.com and in the following documents on the EMC [Powerlink®](#) website:

- Introduction to EMC RecoverPoint 3.5: New Features and Functions — Applied Technology
- Improving Microsoft Exchange Server Recovery with EMC RecoverPoint — Applied Technology
- EMC RecoverPoint Family Overview — A Detailed Review
- Improving VMware Disaster Recovery with EMC RecoverPoint – Applied Technology
- Solving Data Protection Challenges with EMC RecoverPoint – Best Practices Planning
- *EMC RecoverPoint Administrator's Guide* (Powerlink only)