

Introduction to EMC RecoverPoint 3.3: New Features and Functions

Applied Technology

Abstract

This white paper discusses EMC® RecoverPoint version 3.3 software. This software locally and/or remotely replicates one or more SAN or iSCSI volumes in one or more storage systems and maintains an online history of all changes to the volumes. RecoverPoint enables you to bring databases and applications back online quickly and easily after an event like data corruption. RecoverPoint technology enables you to dramatically reduce recovery time and increase the number of recovery points for your application environments.

June 2010

Copyright © 2007, 2008, 2009, 2010 EMC Corporation. All rights reserved.

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.

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

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

Part Number H2781.9

Table of Contents

Executive summary	4
Introduction	4
Audience	4
Introduction to RecoverPoint 3.3.....	4
New features in RecoverPoint 3.3	4
Distributed consistency groups.....	5
Overview.....	5
Distributed consistency groups in RecoverPoint 3.3.....	5
CLARiiON-based splitter sharing	7
Automated VMware vCenter Site Recovery Manager failback.....	7
Cluster Enabler 4.0	8
Support for new platforms and applications.....	8
Microsoft Windows Server 2008 R2 support.....	8
Microsoft Exchange Server 2010 support	8
Conclusion	9
References	9

Executive summary

The EMC® RecoverPoint family provides cost-effective, local continuous data protection (CDP) and continuous remote replication (CRR) solutions that allow for any-point-in-time data recovery. RecoverPoint/SE is the offering that simplifies continuous data protection and replication for the CLARiiON® CX, CX3 UltraScale™, and CX4 series of arrays. RecoverPoint is the full-featured offering that adds support for concurrent local and remote (CLR) replication of the same volumes, intelligent fabrics, heterogeneous server, and heterogeneous storage platforms. Both products are appliance-based, out-of-band data protection solutions designed to ensure the integrity of production data at local and/or remote sites. These products enable customers to centralize and simplify their data protection management and allow for the recovery of data to nearly any point in time.

RecoverPoint CDP provides local continuous replication of every write between a pair of local volumes residing on one or more arrays. RecoverPoint CRR provides remote replication between pairs of volumes residing in two different sites. For asynchronous remote replication multiple writes are collected at the local site, compressed, and sent across periodically to the remote site where they are uncompressed, written to a journal, and then distributed to the target volumes. For local replication and synchronous remote replication every write is collected and written to a local or remote journal and then distributed to the target volumes. In both cases the journal data provides the ability to roll back the target volume to any point in time.

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.

EMC RecoverPoint 3.3 introduces several new capabilities for replication including sharing of CLARiiON splitters, automated SRM failback, and distributed consistency groups.

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 and continuous remote replication. This white paper introduces the new features and functions in RecoverPoint version 3.3 and explains how they can be used to improve overall data protection and recovery.

Audience

This white paper is intended for systems integrators, systems administrators, and the members of the EMC and partners professional services community. This paper serves as an overview to the new features and functions available in RecoverPoint version 3.3. The "References" section provides links to more information on RecoverPoint.

Introduction to RecoverPoint 3.3

RecoverPoint 3.3 is a major release of the RecoverPoint software for RecoverPoint and RecoverPoint/SE. Users with RecoverPoint or RecoverPoint/SE who want the new features should upgrade to RecoverPoint or RecoverPoint/SE version 3.3.

New features in RecoverPoint 3.3

RecoverPoint 3.3 introduces new features and capabilities for local and remote replication. Some of the new features include:

- Distributed consistency groups
- CLARiiON splitter sharing
- Automated VMware vCenter Site Recovery Manager failback
- Cluster Enabler 4.0
- Support for Microsoft Windows Server 2008 R2

- Support for Microsoft Exchange Server 2010

Distributed consistency groups

Overview

RecoverPoint 3.3 supports a new replication capability called distributed consistency groups. To understand distributed consistency groups it is necessary to provide some details on the RecoverPoint building blocks.

Consistency groups

Replication by RecoverPoint is based on a logical entity called a consistency group. SAN-attached storage volumes at the primary and secondary sides — called replication volumes by RecoverPoint — are assigned to a consistency group to define the set of data to be replicated. Data consistency and dependent write-order fidelity are maintained across all volumes assigned to a consistency group, including volumes on different storage systems. The system supports one to 128 consistency groups.

RecoverPoint appliance

The RecoverPoint appliance (RPA) is RecoverPoint's intelligent hardware platform. The RPA is based on an industry-standard Intel server platform and runs RecoverPoint software on top of a custom-built 64-bit Linux kernel environment. The RPA manages all aspects of reliable data replication at all sites.

During replication for a given consistency group, an RPA at the source side makes intelligent decisions regarding when and what data to transfer to the target side. It bases these decisions on its continuous analysis of replication load and resource availability, balanced against the need to prevent degradation of host-application performance and to deliver maximum adherence to the specified replication policy. The target-side RPA distributes the data to the target-side storage.

Distributed consistency groups in RecoverPoint 3.3

The replication of a consistency group is handled by a single RPA, which means that the maximum throughput of a consistency group will be capped by the performance of the RPA. To overcome this limitation, some customers artificially split a consistency group into two or more consistency groups and manually assigned these new consistency groups to different RPAs. This introduced complexities in management and monitoring — for example, deciding which LUNs should go into which new consistency group required detailed performance analysis. Additionally each of the new consistency groups could only be cross-consistent if a Group Set operation or a *parallel_bookmark* command was performed. Also, it was very complex to ensure that all of these sub-consistency groups were operated upon as a single entity, such as would be required for recovery or failback.



Figure 1. Distributed consistency group architecture

These problems are solved when you use a distributed consistency group. A distributed consistency group is managed just like a regular consistency group, but the write load will be distributed across up to four RPAs¹. This enables a much higher write load, and the nature of the distribution helps ensure that no single RPA will perform the bulk of the replication.

¹ EMC recommends that either two or four RPAs be used for a distributed consistency group.

Name	Status	Transfer	Active RPA
ERP	Enabled	Active	RPA 2
ERP Prod	Production Source	N/A	
ERP Copy	Local Replica	Active	
Exchange	Enabled	Active	RPA 1
Exchange	Production Source	N/A	
Exchange NY	Remote Replica	Active	
SQL	Enabled	Active	RPA 1 (primary), 2, 5, 8
Production	Production Source	N/A	
CDP	Local Replica	Active	
CRR	Remote Replica	Active	
VMWare	Enabled	Active	RPA 1 (primary), 6, 7, 8
VMware	Production Source	N/A	
VMware@NY	Local Replica	Active	
VMware@UK	Remote Replica	Active	

Figure 2. Distributed consistency groups shown on the RecoverPoint GUI

A consistency group can be defined as a distributed consistency group, or can be converted into a distributed consistency group. The figure above shows four consistency groups, with two of them being distributed consistency groups, which can be seen by the display showing more than one RPA in the Active RPA column. Up to eight distributed consistency groups can be defined for RecoverPoint and RecoverPoint/SE.

The next figure also shows how a consistency group can be converted into a distributed consistency group. The primary RPA is selected in the General section — here RPA 1 is used — and the secondary RPA(s) are defined in the Advanced section of the Policy tab. The Advanced section will gray out the RPA that is not eligible to be used as a secondary RPA. In this example RPA 1 is not selectable since it is the primary RPA, and RPAs 6, 7, and 8 are selected as the secondary RPAs.

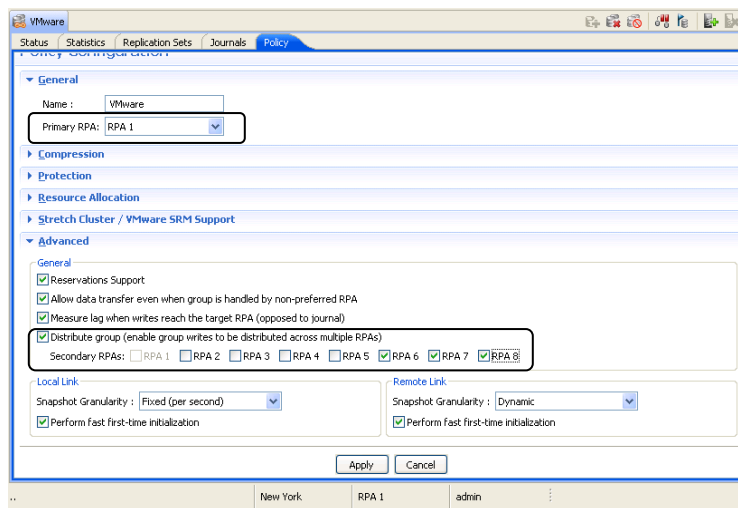


Figure 3. Distributed consistency group policy settings

Up to three RPAs can be selected as secondary RPAs; selecting more than three RPAs will result in an error. It is best to initially create a distributed consistency group if it is required since converting an existing consistency group into a distributed consistency group will result in the journal being lost. Additionally, converting from a distributed consistency group to a normal consistency group will also result in the loss of a journal. However, changing the number or selection of secondary RPAs will not result in a journal loss.

The write is split by the write-splitter and a copy is sent to the primary RPA. The RPA examines the target location of the write to determine which one or more RPAs will process the write. It does this by dividing the LUN into equal-size sections (shown in Figure 1 as the colored blocks), with the primary RPA handling writes for its portion distributing the remaining sections, using the Fibre Channel links, to the secondary RPAs. This results in different sections being distributed across a LUN and helps avoid any issues with a single RPA handling all the hot spots. In the event of a failure of the primary RPA, one of the secondary RPAs assumes the role of primary and performs the role of primary as well as its existing role as secondary. The combined throughput of a distributed consistency group is more than three times that of a

normal consistency group. The difference is due to the role of primary in handling its portion of the writes and distributing the remaining data to the other RPAs.

CLARiiON-based splitter sharing

Using the CLARiiON splitter with RecoverPoint adds multiple operating systems, including Linux and VMware ESX, as supported platforms that don't require intelligent fabric splitting. Without the CLARiiON splitter, support for Linux and VMware requires fabric-splitter technology from either Brocade or Cisco. CLARiiON splitter sharing enables you to utilize the same splitter resident in a CLARiiON CX3 or CX4 storage array for simultaneous attachment to multiple RecoverPoint or RecoverPoint/SE clusters.

The CLARiiON splitter can be used by up to four RecoverPoint or RecoverPoint/SE clusters. This enables the sharing of a single CLARiiON array without the need to use host-based or fabric-based write splitting. Multiple RecoverPoint clusters cannot use the same FC or iSCSI LUN; however, multiple RecoverPoint clusters can use the same array(s). While the CLARiiON-based splitter can be used by RecoverPoint versions earlier than RecoverPoint 3.3, splitter sharing is only available with RecoverPoint or RecoverPoint/SE 3.3.

Automated VMware vCenter Site Recovery Manager failback

The EMC RecoverPoint failback tool is a plug-in to VMware vCenter Server Virtual Center, just like VMware vCenter Site Recovery Manager. Once installed, users have access to the features and steps necessary to automate the failback process after VMware vCenter Site Recovery Manager has been used to mitigate a disaster. Similar to VMware vCenter Site Recovery Manager, users configure the relationship between the disaster recovery site and the production site. They then identify those LUNs that have been failed over, and simply press the failback button – EMC handles the rest.

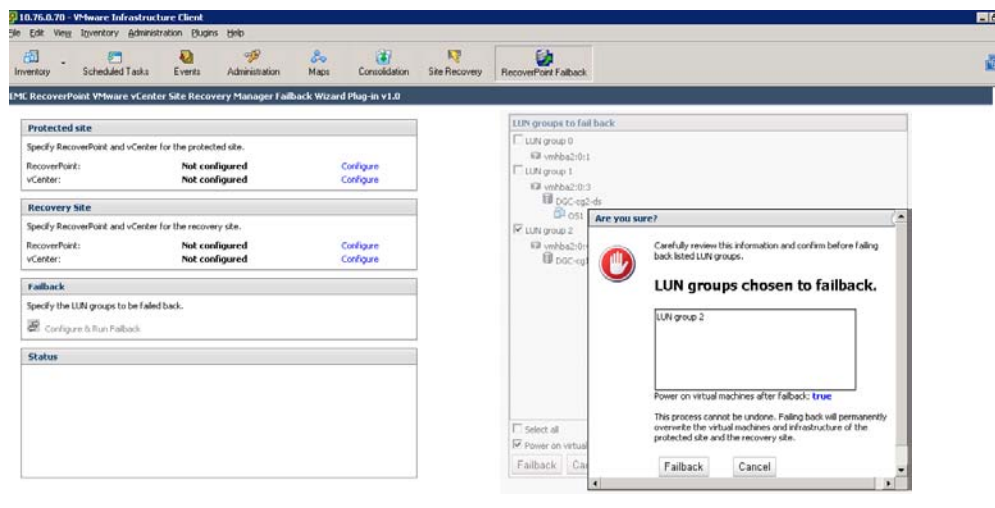


Figure 4. RecoverPoint vCenter Site Recovery Manager failback tool

The RecoverPoint vCenter Site Recovery Manager failback tool is not a replacement for VMware vCenter Site Recovery Manager. This tool does not provide users with a mechanism to execute custom script or business logic that is specific to their organization. This tool does not provide an ordering mechanism for virtual machine dependencies and does not reset the IP addresses on each virtual machine. This tool does not support a complete data center loss when and if the disaster recovery site fails before failing back to the production site. This is because the failback tool is not a replacement for VMware vCenter Site Recovery Manager. This tool requires access to the original RecoverPoint configuration to resolve replication configuration and settings to automate failback.

Cluster Enabler 4.0

Cluster Enabler has been updated to support Microsoft Windows Server 2008 R2 and Hyper-V for Windows Server 2008 R2. The upgrade to Cluster Enabler 4.0 is nondisruptive. Aside from support for Windows Server 2008 R2, there are no other feature updates for Cluster Enabler.

Cluster Enabler 4.0 supports RecoverPoint 3.3 and RecoverPoint/SE 3.3 for replication of configurations built from Microsoft Cluster Server on Windows Server 2003 and Microsoft Failover Cluster on Windows Server 2008 and Windows Server 2008 R2. Cluster Enabler 4.0 does not support Microsoft Cluster Shared Volumes.

Support for new platforms and applications

Microsoft Windows Server 2008 R2 support

The KDriver is system software installed on all hosts that access protected volumes, that is, volumes being locally replicated using continuous data protection and/or remotely replicated using continuous remote replication. The primary function of the KDriver is to “split” application writes so that they are sent not only to their normally designated storage volumes, but also to the RPA. The KDriver carries out this activity efficiently, with little perceptible impact on host performance, since all CPU-intensive processing necessary for replication is performed by the RPA.

Windows Server 2008 R2 is one of Microsoft’s server operating systems. Released to manufacturing on July 22, 2009, it is the successor to Windows Server 2003, which was released nearly six years earlier. Like Windows Vista and Windows 7, Windows Server 2008 R2 is built on Windows NT 6.x. RecoverPoint 3.3 has been qualified to support Windows Server 2008 R2 for local and remote replication of LUNs using the CLARiiON-based or intelligent fabric-based splitter drivers. RecoverPoint 3.3 also includes a 64-bit KDriver for Windows Server 2008 R2 hosts. Additionally, the utilities for RecoverPoint, called kutils, are installed automatically with the Windows Server 2008 R2 KDriver.

Microsoft Exchange Server 2010 support

Microsoft Exchange Server is a messaging and collaborative software product developed by Microsoft. It is part of Microsoft’s line of server products and is widely used by enterprises using Microsoft infrastructure solutions. Exchange’s major features consist of electronic mail, calendaring, contacts and tasks; support for mobile and web-based access to information; and support for data storage. Microsoft Exchange Server 2010 was officially launched on November 9, 2009, and it is expected that existing Exchange Server 2003 and 2007 customers will start migrating to Exchange Server 2010 in the first half of 2010.

Exchange 2010 is only available in 64-bit and is intended to be installed and run on a Windows Server 2008 R2 platform. RecoverPoint 3.3 has been qualified to support Microsoft Exchange Server 2010 for local and remote replication of LUNs using the Windows Server 2008 R2-based, CLARiiON-based, or intelligent fabric-based splitter drivers. RecoverPoint 3.3 also includes a utility called KVSS that invokes the Microsoft VSS framework to prepare and take a transportable Shadow Copy, using normal Microsoft VSSAPI routines. This utility has been qualified with Microsoft Exchange Server 2010.

RecoverPoint supports synchronous replication of Exchange Server 2010 Database Availability Groups (DAGs) using the EMC replication enabler for Microsoft Exchange Server 2010. The EMC replication enabler is a free software utility that integrates RecoverPoint and RecoverPoint/SE synchronous remote replication with the Exchange Server 2010 DAG architecture. RecoverPoint also supports replication of Exchange Server 2010 in a virtualized environment. For additional information, please contact your EMC representative.

Conclusion

EMC RecoverPoint offers continuous data protection, continuous remote replication, and concurrent local and remote data protection functionalities. With its customer-defined recovery point and recovery time objectives, RecoverPoint allows critical business processes to be available locally for operational recovery, or remotely at a disaster recovery site hundreds or thousands of miles away from the primary site. With support for consistency groups, RecoverPoint is a no-data-loss model with full write order consistency for replicated volumes that can span multiple heterogeneous storage systems and servers.

References

Visit the [EMC RecoverPoint page](#) for more information. White paper titles include:

- *EMC RecoverPoint Family Overview*
- *Improving VMware Disaster Recovery with EMC RecoverPoint*
- *Improving Microsoft Exchange Server Recovery with EMC RecoverPoint*
- *Using EMC RecoverPoint Concurrent Local and Remote for Operational and Disaster Recovery*
- *Solving Data Protection Challenges with EMC RecoverPoint*