

# Invista Disaster Recovery with RecoverPoint

*Applied Technology*

---

**Abstract**

This document covers the utilization of EMC<sup>®</sup> RecoverPoint to provide disaster recovery for EMC Invista<sup>™</sup>, leveraging both the RecoverPoint splitter driver and the Cisco MDS 9000 series SAN Tap service by RecoverPoint.

January 2007

---

---

Copyright © 2007 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 H2557

---

# Table of Contents

<b>Executive summary .....</b>	<b>4</b>
<b>Introduction .....</b>	<b>4</b>
Audience .....	5
Business challenges .....	5
<b>Invista overview .....</b>	<b>6</b>
<b>RecoverPoint overview.....</b>	<b>8</b>
CRR and CDP in one common platform.....	8
RecoverPoint capabilities.....	8
RecoverPoint fabric-based replication benefits.....	9
<b>SANTap: Optimized architecture for data replication.....</b>	<b>10</b>
SANTap benefits for data replication.....	11
<b>Replicating Invista Virtual Volumes with RecoverPoint.....</b>	<b>12</b>
Virtual-to-virtual remote replication.....	12
Virtual-to-physical remote replication.....	13
Replicating Invista volumes with RecoverPoint on the host.....	13
Replicating Invista volumes with RecoverPoint and SANTap.....	14
<b>Conclusion.....</b>	<b>15</b>
<b>References.....</b>	<b>15</b>

---

## Executive summary

EMC® RecoverPoint is an enterprise scale appliance designed to protect application data on heterogeneous SAN-attached servers and storage arrays. RecoverPoint combines industry-leading continuous data protection technology with a bandwidth-efficient, no-data-loss, replication technology. This allows RecoverPoint to protect data both locally and remotely. The Cisco MDS series storage services module (SSM) coupled with RecoverPoint creates a superior data-protection solution for the heterogeneous enterprise that offers an unprecedented level of functionality, manageability, flexibility, and performance.

EMC Invista™ enables the virtualization of storage in networked storage environments that dramatically reduces the amount of downtime associated with the movement of data across storage tiers in support of Information Lifecycle Management (ILM) strategies.

Unlike other offerings, Invista places the virtualization intelligence in the network, where that intelligence is applied—and where it does not impact server or application performance. The results are higher application availability, reduced administrative overhead, more-effective and efficient use of storage resources, and reduced cost; all leading to greater business value from the advantages of a tiered-storage ILM strategy.

The integration of Invista and RecoverPoint provides a unique, network-based virtualization solution with extensive and flexible remote disaster-recovery capabilities. Both Invista and RecoverPoint leverage intelligent switch technology to provide a truly heterogeneous replication environment, enabling more-complex virtual-to-virtual and less-costly virtual-to-physical disaster recovery options.

## Introduction

Intelligent storage networks promise a new heterogeneous, flexible, high-performance platform for hosting storage applications. RecoverPoint has delivered on this promise by utilizing the Cisco MDS 9000 series SANTap service, offering a fully heterogeneous fabric-based data-replication solution with no need for a host agent. RecoverPoint replicates data in any SAN-attached array, across any distance, or to any other array, with none of the traditional downsides of deploying devices, in-band (within the data path) or out-of-band, in conjunction with host-based software agents.

RecoverPoint provides continuous remote replication (CRR) and continuous data protection (CDP) for enterprise data center applications. RecoverPoint reduces the total cost of data protection and management by considering the value of different types of data and the associated requirements for data accessibility, performance, availability, and protection.

RecoverPoint's intelligent features specifically address key customer concerns related to data migration and replication, recovery, storage utilization, and the high cost of bandwidth and storage management. Utilizing the SANTap, RecoverPoint can be deployed out-of-band (outside the data path) and without a host agent, since SANTap provides a reliable copy of storage *write* operations. SANTap functionality is enabled through the SSM, which can be inserted into any modular switch within the Cisco MDS 9000 family of multilayer intelligent storage switches.

Invista leverages intelligent SAN switches for deployment of network-based storage virtualization. Invista takes advantage of additional specialized processing power in the switch to perform the basic operation of virtualization at wire speed. The intelligent switch is a high-performance platform that can be incorporated into existing SAN infrastructures.

Implementing virtualization at the storage-network level allows customers to mobilize their data nondisruptively across diverse, heterogeneous back-end storage platforms.

Together, Invista and RecoverPoint provide an enterprise-class storage-virtualization solution with comprehensive data protection and disaster recovery capabilities for heterogeneous host and storage environments.

---

## **Audience**

This white paper is targeted to corporate management and business decision-makers, including storage and server administrators, IT managers, and application engineers, as well as storage integrators, consultants, and distributors.

## **Business challenges**

The following are some business challenges customers face today:

- Planning downtime to support changes in the storage infrastructure
- Complex reporting of storage capacity in multivendor environments
- Managing data availability across heterogeneous storage arrays
- A competitive environment that can not withstand any data loss in the event of server failures, catastrophic events, or data corruption
- Corporate service-level agreements (SLAs) that require near-instant point-in-time recovery and minimal downtime for mission-critical applications
- Increasing capital and operational expenses in maintaining multiple solutions across heterogeneous environments

In today's environment, you are likely managing hundreds of servers and applications, and gigabytes to terabytes of storage. With this growth comes complexity and inefficiency; for example, upgrading to new storage technologies takes careful planning and—ultimately—a migration requiring planned downtime that is disruptive to the business.

EMC recognizes that many customer environments have storage from EMC as well as from other vendors. Typically, storage allocation takes place on a device-by-device basis, which is time-consuming and costly.

Another major inefficiency is the notion of *cloning* (copying) data across multivendor storage arrays. The methods that many companies use today to create point-in-time copies across different storage arrays are costly and cause downtime. The challenge is leveraging a tiered storage environment for such operations as local replication without downtime.

Customers must ensure that their environments are protected from data loss due to common problems in the event of failures and catastrophic events, which can bring data centers to a standstill.

Customers also need guaranteed instant access to any single point-in-time copy of their data.

Lastly, costs are increasing in several areas, including:

- Bandwidth
- Creating extra copies of data
- Training and skill sets for maintaining disparate point solutions

Customers are looking for solutions that ease their fears of losing data at a lower total cost of ownership. Invista and RecoverPoint address these challenges by providing:

- Disaster-recovery capabilities for Invista-virtualized environments
- Virtual-to-virtual volume replication, which allows IT managers to employ virtualization technologies across multiple sites for disaster tolerance and enhanced availability
- Virtual-to-physical heterogeneous remote replication of Invista environments to increase cost savings and flexibility
- Synchronous or asynchronous bidirectional replication between sites for more flexible disaster recovery and increased application availability

---

## Invista overview

Enterprise computing environments are larger and more complex than ever. As organizations grow, the demand on IT to deliver on the information requirements of the business grows accordingly. For most organizations, that means constructing complex IT infrastructures that can deliver the right information to the right place at the right time, 24 hours a day, seven days a week. Businesses require continuous information availability, and are unwilling to endure infrastructure downtime when delivering on their needs.


In addition to such unplanned events as natural or man-made disasters, planned downtime is an IT reality, accounting for 60 percent to 75 percent of all downtime. Some key reasons for planned downtime are:

- Scheduled maintenance
- Lease rollovers
- Technology refreshes
- Data center migrations/additions
- Service-level optimizations
- Reconfigurations

The cost of downtime to the enterprise can be staggering.

Enterprise-class infrastructures typically include multivendor server environments, diverse connectivity technologies, and multivendor tiered-storage environments. Organizations require the ability to allocate any storage to any application based on the needs of the business, and to do so nondisruptively. Networked storage virtualization enables organizations to deliver the right information at the right performance level, and the right functionality to the business at the lowest total cost.

EMC Invista (represented in [Figure 1 on page 7](#)) enables the virtualization of storage in networked storage environments that dramatically reduces the amount of downtime associated with the movement of data across storage tiers in support of ILM strategies.



IWA-000022

**Figure 1. Invista architecture**

Invista provides the following benefits:

- High-performance architecture:
  - Leverages next-generation "intelligent" SAN switches for high performance, even with intensive random I/O applications
  - Works in enterprise-class customer environments
- Advanced functionality:
  - Dynamic volume mobility
  - Network-based volume management
  - Heterogeneous point-in-time copies (clones)
- Support for heterogeneous environments:
  - Works with EMC and third-party storage
  - Runs on intelligent switches from Brocade or Cisco, leveraging standard APIs

---

## RecoverPoint overview

EMC RecoverPoint provides a comprehensive solution for the entire data center, providing local CDP and CRR, with CDP at the remote site. RecoverPoint protects companies from data loss due to such common problems as server failures, data corruption, software errors, viruses, and end-user errors. RecoverPoint also protects against catastrophic events that can bring an entire data center to a standstill.

RecoverPoint customers have realized dramatic cost savings by eliminating complex, nonperforming data-protection schemes, and deploying a single, easy-to-manage solution. RecoverPoint supports intelligent application bookmarks, instantaneous recovery, and bidirectional local and remote replication, providing true any-point-in-time protection. The unique clustered architecture of RecoverPoint provides linear scalability to support the most demanding environments. With support for heterogeneous storage, hosts, networks, and SANs, RecoverPoint protects storage investments, enhances business continuity, and facilitates storage consolidation.

### ***CRR and CDP in one common platform***

The RecoverPoint data protection solution employs a continuous remote replication (CRR) module and a local continuous data protection (CDP) module. These modules run on the RecoverPoint appliance to reduce cost and ease management, and can be deployed separately or together based on your unique needs for end-to-end data protection.

The RecoverPoint CRR module is a third-generation, enterprise-class data replication solution that offers any-point-in-time recovery with complete protection from site disasters. Enterprise performance, scalability, and instant recovery are combined to guarantee data consistency, with recovery that takes seconds or minutes instead of hours or days. RecoverPoint offers bidirectional replication with no distance limitation, guaranteed data consistency, and advanced bandwidth reduction technology designed to dramatically reduce WAN bandwidth requirements and associated costs.

The RecoverPoint CDP module tracks changes to data and journals the changes along with application-specific bookmarks. This capability is used to protect critical data from physical or logical failures, such as server outages, data corruption, software errors, viruses, and common user errors.

RecoverPoint leverages either host-based splitter to split the I/Os; also, Cisco SANTap can split the I/Os at the intelligent network level.

### ***RecoverPoint capabilities***

RecoverPoint capabilities include:

- End-to-end data protection — RecoverPoint provides an end-to-end data protection solution for the entire data center. RecoverPoint is the first solution to deliver a comprehensive data protection suite that can be deployed to protect data continuously, both locally and remotely. This means that the multiple products currently required for local and remote data protection can be replaced by one centrally managed solution, dramatically lowering complexity and cost.

RecoverPoint offers multiple management features, including automated discovery of LUNs, centralized management through a web-based graphical user interface (GUI), and a programmatic command line interface (CLI). RecoverPoint is also SNMP-enabled for integration with standard system-management applications, and it adheres to business objectives and SLAs tied to policies. RecoverPoint determines the optimized data-protection process intelligently, dramatically reducing complexity and maximizing the use of IT and system resources.

- Long-distance replication — Unique bandwidth compression and bidirectional replication, combined with guaranteed write-order consistency, shatter today's distance/latency limitations. RecoverPoint enables up-to-date protection from regional disasters, with no impact on application performance or reliability.

- 
- Immediate recovery — RecoverPoint CDP stores changed data and time together with application-specific bookmarks, enabling immediate recovery to any point in time with no data loss. RecoverPoint CDP can recover applications quickly to a selected time or event with guaranteed consistency, even when data spans multiple heterogeneous storage volumes and servers.

## ***RecoverPoint fabric-based replication benefits***

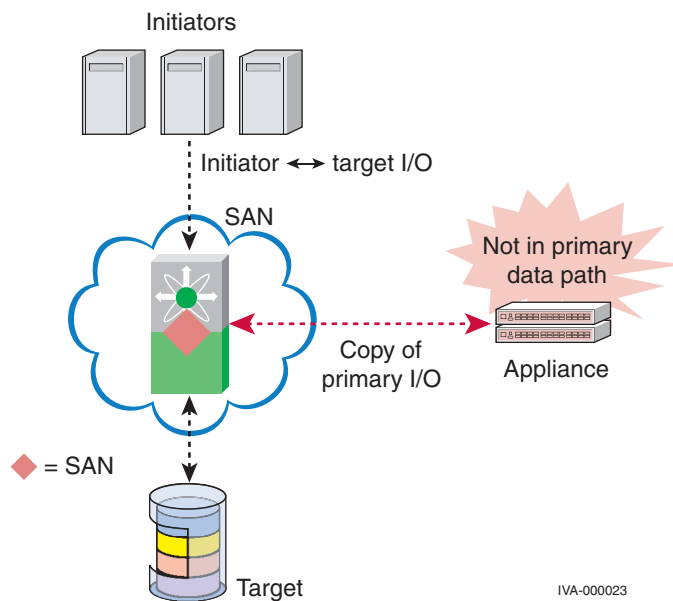
RecoverPoint benefits include:

- Universal enterprise data protection — RecoverPoint with SANTap is an end-to-end solution for CRR and CDP across heterogeneous server and storage platforms, enabling complete data protection for the entire enterprise. Because storage systems at the primary and secondary sites do not have to be the same, there is flexibility to deploy lower-cost storage (at the secondary site, for example) or to leverage existing storage. With the SSM SANTap architecture, RecoverPoint can be used with virtually every open systems-based environment. No changes are required on the storage arrays, hosts, or SAN configuration.
- Guaranteed data consistency — RecoverPoint guarantees a consistent replica of business-critical data in the event of a possible failure or disaster. With RecoverPoint, consistency is maintained at all times, even through rolling disasters or during resynchronization.
- Intelligent use of bandwidth — RecoverPoint employs intelligent bandwidth-reduction technologies to deliver unprecedented bandwidth savings. As a result, the system always provides superior protection for the available bandwidth, while dramatically reducing WAN costs, particularly over long distances. Data reduction is achieved through application-aware and storage-aware algorithmic techniques that conserve bandwidth to an extent not possible with traditional compression technologies.
- Policy-based data replication — RecoverPoint offers a full spectrum of replication modes, such as synchronous, asynchronous, small-aperture snapshot, and point in time. The replication process is managed automatically, with strict adherence to user-defined policies that are tied to desired business objectives. The system adapts its replication dynamically for each application according to these policies, based on the available bandwidth and the application workload. This greatly simplifies data and disaster-recovery management for complex and heterogeneous environments. For example, within the same storage array, replication policies can be set for minimal lag for critical applications, and policies can be set to minimize bandwidth for less-critical data.
- Recovery to any point in time — RecoverPoint efficiently maintains a journal to allow convenient rollback to any point in time, enabling fast and effective recovery from any disaster. RecoverPoint supports multiple transactional-consistent snapshots at the remote site, allowing reliable recovery in database environments. Frequent, small-aperture snapshots (seconds apart) are utilized to minimize the risk of data loss due to data corruption.
- Data processing at the secondary site — RecoverPoint enables direct read/write access to the replicated image at the secondary site. There is no requirement to make an additional copy first. The system supports robust failover and failback capabilities, reducing management and operational costs.
- Intuitive management — RecoverPoint offers multiple management features that make implementation and operation extremely easy. The system, consisting of any number of appliances, is centrally managed through a single intuitive GUI or CLI, both of which are accessible from a secure Internet connection. RecoverPoint provides automated discovery of LUNs, is SNMP-enabled for integration with standard enterprise network- and system-management applications, and offers CallHome capability that proactively reports system status in the event of any failure.

## SANTap: Optimized architecture for data replication

Storage applications can be delivered either on the host or initiator, on the storage or target, or inside a SAN-attached appliance. A purpose-built appliance such as EMC RecoverPoint, which is designed to provide a specific storage application such as replication, offers the benefit of supporting heterogeneous storage while minimizing the impact to hosts, applications, or array performance. Long-distance data replication is uniquely suited for this architecture, since host and storage are relieved from the burden of managing the data movements across the WAN and dealing with complex data-consistency issues inherent in cross-site data replication.

As [Figure 2 on page 10](#) shows, an appliance placed in the SAN/WAN junction must be connected either in-band (between host and storage) or out-of-band (with the aid of a host agent). Both approaches can have obvious drawbacks: The in-band approach can compromise SAN performance, integrity, and availability, and can be very disruptive to deploy; out-of-band deployment removes these drawbacks, but requires installation of a driver in the host, adding potential complexity to the solution's deployment.



**Figure 2. SANTap-based connectivity for RecoverPoint appliance**

The Cisco MDS 9000 series SANTap service eliminates the need for host agents, enabling a simplified and agentless implementation of the RecoverPoint data replication appliance. This out-of-band appliance approach offers an optimized architecture for data replication that protects existing investments in storage arrays with no host footprint, and delivers high-performance bidirectional data replication between SANs across any distance, with no impact to host and array performance. With ample processing, and utilizing native SAN and WAN interfaces, the RecoverPoint appliance replicates data efficiently across multiple homogeneous or heterogeneous arrays without the need for protocol converters or edge-connect devices.

---

## ***SANTap benefits for data replication***

SANTap provides:


- **Seamless deployment of the EMC RecoverPoint data replication appliance**  
The RecoverPoint appliance can be deployed seamlessly with no disruption. The SANTap architecture eliminates the need for host-side agents, further simplifying the deployment.
- **No disruption of the primary I/O from the server to the storage array**  
The RecoverPoint design takes advantage of SANTap to eliminate the risk of impacting the availability and performance of deployed applications.
- **Deployment flexibility and investment protection**  
RecoverPoint data replication services can be provided to all existing servers and storage in the SAN, irrespective of the operating systems or array types. This allows customers to get more out of the existing storage and server infrastructure investments.
- **On-demand storage services**  
RecoverPoint data replication services can be provisioned on demand with no application downtime for any server or storage connected to any port of a storage network.
- **Unlimited scalability, with no performance bottlenecks**  
A single SSM provides 320,000 IO/s and 20 Gb/s throughput. Moreover, SANTap has been implemented in a distributed architecture that enables multiple SSMs on a storage network to provide the SANTap services. Coupled with the linear scalability of data replication appliances, this means that customers are no longer constrained by the performance limitations of using host CPU cycles and in-band appliances.

## Replicating Invista Virtual Volumes with RecoverPoint

RecoverPoint enables volume-level replication in an Invista-virtualized environment. Since the replication is done at the volume level, the integration of the two technologies provides a powerful disaster recovery option for both virtual-to-virtual and virtual-to-physical environments. Both Invista and RecoverPoint function at the storage network-level, leveraging intelligent switch technology. RecoverPoint also offers a host-based splitter function that enables the splitting at the host level.

### *Virtual-to-virtual remote replication*

As [Figure 3](#) demonstrates, RecoverPoint performs remote replication at the network level, enabling virtual-to-virtual volume replication over distance, which in turn enables the use of virtualization at both the primary site and the disaster recovery site. This solution gives the customer ultimate flexibility at both locations.

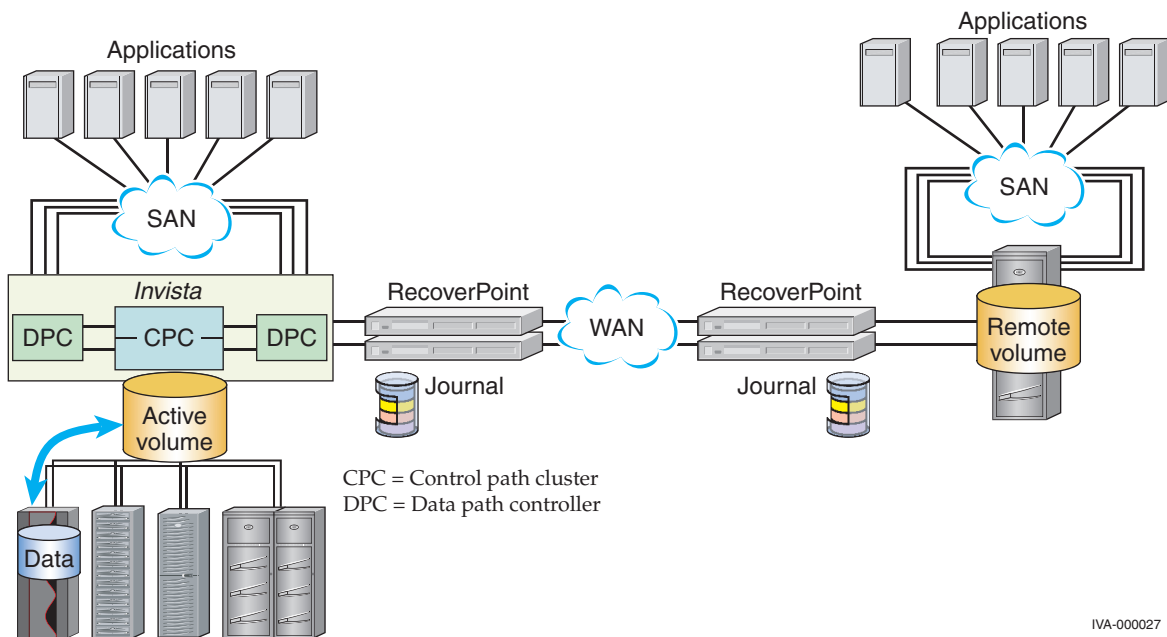


**Figure 3. Virtual-to-virtual remote replication**

## Virtual-to-physical remote replication

As [Figure 4](#) shows, Invista and RecoverPoint can be deployed together to replicate a virtualized environment to a physical environment. Both RecoverPoint and Invista support heterogeneous host and storage environments, so you can remotely replicate the more expensive tiers of storage to lower-cost storage on the target side. This configuration allows the deployment of a lower-cost replication solution for Invista environments.

Virtual-to-physical replication also provides a much-simpler solution for customers who do not require virtualization in the disaster recover site.



**Figure 4. Virtual-to-physical remote replication**

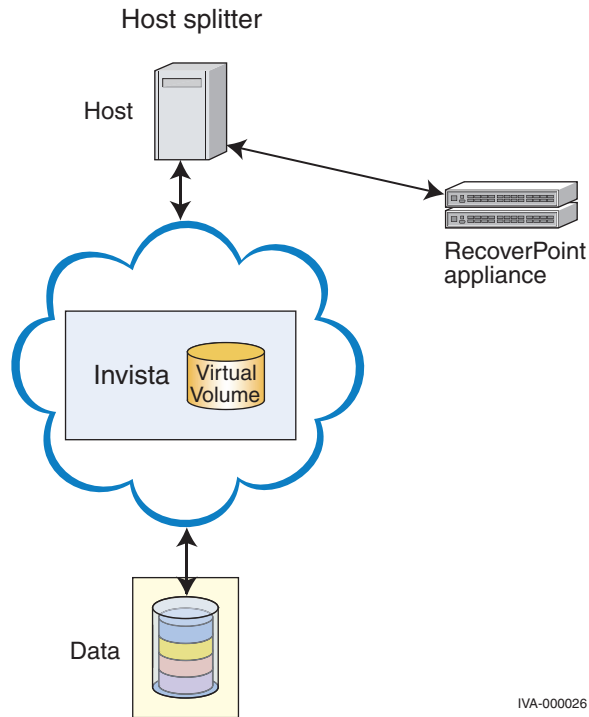
## Replicating Invista volumes with RecoverPoint on the host

The RecoverPoint driver (RPDriver) is system software that you install on all hosts that access replication volumes (volumes that contain data to be replicated). The primary function of the RPDriver is to *split* application writes so that they are sent not only to their normally designated storage volumes but also to the RecoverPoint appliance (RPA). The RPDriver 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.

The RPDriver can be used to replicate an Invista Virtual Volume in the same way it is used to replicate a physical volume. Since the Invista Virtual Volume looks just like any other SAN-attached volume to the host, the RPDriver splits the write, sending one write to the Virtual Volume and the other to the RPA.

The following steps (represented in [Figure 5 on page 14](#)) occur when using the RPDriver with Invista:

1. The host sends a write to an Invista Virtual Volume.
2. The write is intercepted by the RPDriver and passed through to Invista. A copy is sent to the RPA.
3. Invista translates the Virtual Volume into the physical location and performs the write while the RPA handles the write appropriately.
4. The host receives acknowledgment from both Invista and the RPA.



IVA-000026

**Figure 5. Using the RPDriver with Invista**

## ***Replicating Invista volumes with RecoverPoint and SANTap***


Like Invista, RecoverPoint leverages the Cisco SSM, a blade that supports the Cisco MDS 9000 series of intelligent fabric switches. Invista uses the SSM to assist in the virtualization of storage. In the case of RecoverPoint, the SSM provides a SANTap service that can be used to intercept and redirect a copy of a write between an initiator and target. Currently, Invista and RecoverPoint require separate SSMs for each function.

RecoverPoint uses the SANTap services to eliminate the need for a host splitter on each application server. SANTap monitors the writes to specific LUN targets, and sends a copy of the write to the RecoverPoint appliance for further CDP or CRR processing.

In a virtualized environment, SANTap sits in front of Invista, intercepting and splitting writes to the appropriate Virtual Volumes. Instead of SANTap forwarding a write to a physical storage target, it forwards the write to the Invista Virtual Target behind which the Virtual Volume is accessible.

The steps (represented in [Figure 6 on page 15](#)) are as follows:

1. The host sends a write to the Invista Virtual Volume.
2. The write is intercepted by SANTap and is forwarded to both the Invista Virtual Target and the RPA.
3. Invista translates the Virtual Volume into the physical location and performs the write, while the RPA handles the write appropriately.
4. Both Invista and the RPA send an acknowledgment to the switch, and the switch sends an acknowledgment to the host.



**Figure 6. Replicating Invista volumes with RecoverPoint and SANTap**

## Conclusion

EMC Invista coupled with EMC RecoverPoint creates a superior virtualization solution for the heterogeneous enterprise, offering an unprecedented level of functionality, flexibility, and performance. By leveraging intelligent network technology, RecoverPoint can replicate Invista Virtual Volumes in heterogeneous storage and server environments across any distance, with no impact to the primary I/O path or applications. The combined solution provides both virtual-to-virtual and virtual-to-physical replication for Invista-virtualized environments, providing a powerful, flexible, cost-effective solution.

## References

You can review EMC's proven solutions for virtualization, data replication, data lifecycle management, disaster recovery, and continuous data protection at <http://www.EMC.com>, or you can contact an EMC representative at 1-800-607-9546 (1.508.435.1000 outside the U.S.) for more information.

