

# WHAT IS THE POWER OF RECOVERPOINT?

Roy Mikes  
Infrastructure/Datacenter Architect  
roy@mikes.eu

## Table of Contents

About This Article	3
Who Should Read This Article?	3
Introduction	4
Why we need replication	6
Data Protection Technologies	8
Backup	8
Snapshots	8
Continuous data protection	9
What is RecoverPoint?	11
Splitter	13
Management	15
VMware Site Recovery Manager	18
Storage Replication Adapter	18
Developing a Recovery Strategy	20
How it works	21
Consistency Groups	21
Policy	21
VMware SRM Protection Groups	23
VMware SRM Recovery Plan	25
Run a Recovery Plan	26
AppSync	28
Conclusion	30
References	32

Disclaimer: The views, processes, or methodologies published in this article are those of the author. They do not necessarily reflect EMC Corporation's views, processes, or methodologies.

## **About This Article**

Today's businesses face 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 changes. Critical data changes instantly, and to protect this data we are frequently reaching to new technologies. The EMC RecoverPoint product family provides a comprehensive data protection solution for mainly enterprise customers, providing integrated continuous data protection and continuous remote replication to recover applications to any point in time.

This article will help you understand the need for replication performed by RecoverPoint. Because it is not all technical, this article covers also the non-technical discussions.

As such, this material is probably most useful to those with little or no familiarity with this topic. Readers who fall into this category would be well served to read this article.

## **Who Should Read This Article?**

This article is written for IT professionals who are responsible for managing and defining the direction of protecting data in their data center(s).

These include:

- Storage Administrators
- Operational, middle level managers
- IT managers (CIO, Chief information officer)

Organizations and individuals who have the same interests will benefit from this article as well. My goal is to give a general guideline to provide insight into Replication and Recovery, which should not be too difficult to read.

## Introduction

Though certainly not a new product, I recently finished an EMC RecoverPoint configuration. After installation, configuration, and a recovery test I was amazed by the great potential of this appliance/software. RecoverPoint makes it easier to protect applications that grow, providing a wizard that allows you to modify the applications' protection configuration to add new storage volumes. According to EMC, you'll never have to worry about data protection again. As far as I can judge now, this is 100% true.

### What Is the Power of RecoverPoint?

RecoverPoint systems enable reliable data replication over any distance within the same site (CDP), to another distant site (CRR), or both concurrently (CLR). Specifically, RecoverPoint systems support data replication that applications are writing over Fibre Channel to local SAN-attached storage. The systems use existing Fiber Channel infrastructure to integrate seamlessly with existing host applications and data storage subsystems. For remote replication, the systems use existing IP connections to send the replicated data over a WAN, or use Fiber Channel infrastructure to replicate data asynchronously or synchronously. The systems provide failover of operations to a secondary site in the event of a disaster at the primary site.

Similar to other continuous data protection products, and unlike backup products, RecoverPoint needs to obtain a copy of every write in order to track data changes. RecoverPoint supports three methods of write splitting: host-based, fabric-based, and in the storage array. EMC advertises RecoverPoint as heterogeneous due to its support of multi-vendor server, network, and storage environments.

Each site requires installation of a cluster that holds a minimum of two RecoverPoint appliances for redundancy. Each appliance is connected via Fiber Channel to the SAN, and must be zoned together with both the server and the storage. Each appliance must also be connected to an IP network for management. All replication takes place over standard IP for asynchronous replication and Fiber Channel for synchronous replication.

Beyond integration with EMC products such as the CLARiiON<sup>®</sup>, VNX<sup>®</sup>, VMAX<sup>®</sup> storage arrays, Replication Manager<sup>®</sup> and EMC Control Center<sup>®</sup>, RecoverPoint integrates with VMware vCenter and Microsoft Hyper-V, enabling protection to be specified per virtual machine instead of per volumes that are available. It also integrates with Microsoft Shadow Copy, Exchange, SQL Server, and Oracle Database Server which enables RecoverPoint to temporarily stop writes by the host in order to take consistent application-specific snapshots.

EMC RecoverPoint's concurrent local and remote (CLR) data protection technology eliminates the need for separate solutions as it provides CDP and CRR of the same data. The solution now provides more flexibility to replicate and protect data in many local and remote-site combinations with less storage footprint whether for production applications or for test and development.

Despite the simple looks and the lots of 'sounds goods', you really have to know what you are doing with this application. It can be confusing because of the many possibilities. Therefore, you should be careful what you do.

## Why we need replication

Data replication is the same data stored on multiple storage arrays. Besides people, the most valuable asset is your business data. Without people to maintain the equipment, even the most sophisticated and powerful machinery would cease to function. Without people doing the day-to-day operations, the organization would stop functioning. On the one hand, we need people. On the other side, technique is needed to protect against disasters such as data loss. My previous<sup>[1]</sup> Proven Professional Knowledge Sharing article discussed Disaster Recovery at a high level. I encourage you to read it as it is a good complement to this topic.

Data replication is an increasingly important topic these days as more and more databases are deployed. One of the challenges in database replication is to introduce replication without restricting performance. This can be very difficult in large environments.


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 disasters that could make data at the primary site unavailable. Before applications are transferred to production, design of a protection solution must be completed. In fact, this should be mandatory. Ask yourself what will be left, if the production site goes offline due to a disaster and the business processes on the primary site becomes unavailable or even lost. To prevent this kind of scenario, data has to be transferred to a recovery site.

Another interesting approach is how much data loss can be tolerated before the business is deemed unable to restart production? This is termed the recovery point objective (RPO), which is "... the maximum tolerable period in which data might be lost". For 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.

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 reasonable to assume that the shorter the RTO and the RPO, the more difficult or costly it may be to successfully implement a disaster recovery process. A perfect configuration that guarantees no data loss and data that is instantly available may come at a complexity and price that often are not practical. Therefore, it is important to distinguish between absolutely critical business applications and other applications.

*Impact versus investment*




## Data Protection Technologies

Data protection has evolved over time and became more important. In the early years, most computer systems were stand-alone, with their entire data residing on one single system. Networking and interconnectivity between systems were expensive and limited. Yet, there was a need to protect this 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 applications evolved, so did the backup technology. This is not altered by time. Applications become more complex and data is growing. It's logical that these technologies evolve too, and are interrelated.

A reasonable question is: How do you choose the right replication method for optimal data protection? However, the right question is: How important is your data? Answering this question isn't easy because how do you measure importance?

### Backup

The simplest form of replication is Backup. A backup is a copy of data from your files or database that can be used to reconstruct that data. Backups can be divided into physical backups and logical backups. I assume no explanation is needed on the idea behind backup via tape or disk. However, there are several specific backup methods for backup solutions for Microsoft, such as SQL, Exchange, or SharePoint. Or backup solutions for Oracle. Let's also not forget VMware. What about Desktops, Laptops, or Remote Office Solutions or Disaster Recovery, and so on? As you can see, there's a lot to consider and each application has its peculiarities.




### Snapshots

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 time when 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.

A snapshot may be a full copy of a LUN, 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.

Snapshots tend to be less disruptive to applications and environment. Remember that when a snapshot is taken there still can be data in memory. Make sure you flush this data to disk when you need an absolute consistent snapshot.



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. Meanwhile 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 independent of the volume from which it was taken.

Regardless of the implementation, snapshots are less disruptive, more reliable, and faster than traditional backup. It's worth noting that snapshots can consume significant resources.

### **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 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 LUN.

File system CDP products are typically found in Microsoft Windows environments, and usually offer a file system. 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.

A CDP system also enables 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.

## What is RecoverPoint?

EMC RecoverPoint<sup>[2]</sup> is a single solution that provides the advantages of host-based and array-based solutions while replicating data from any SAN-based array to any other SAN-based array over existing Fibre Channel or IP networks using any combination of host-based, VNX<sup>®</sup>-based, or intelligent fabric-based write-splitting options. RecoverPoint is a product variant that is optimized for VNX series storage arrays.

Both RecoverPoint and RecoverPoint/SE provide synchronous local replication using continuous data protection (CDP), synchronous and asynchronous continuous remote replication (CRR), and concurrent local and remote (CLR) data protection. The RecoverPoint family protects companies from data loss due to common problems such as server failures, data corruption, software errors, viruses, and end-user errors, while also protecting against catastrophic events that can bring an entire data center to a standstill.

The RecoverPoint family supports application bookmarks, instantaneous recovery, and bi-directional local and remote replication. RecoverPoint provides point-in-time DVR-like recovery, and its unique clustered architecture provides linear scalability to support the most-demanding environments. RecoverPoint support for heterogeneous storage, hosts, networks, and SANs enables storage investment protection, enhances business continuity, and facilitates storage consolidation.

RecoverPoint application software runs on an EMC-provided and -supported appliance that provides the core functionality and management for the system. The RecoverPoint appliance is built from a standard Dell 1 $\mu$  high-availability server running a customized 64-bit Linux 2.6 kernel. Appliances are sold and deployed in a two- to eight-node cluster configuration per site. A RecoverPoint cluster enables active-active failover between the nodes.

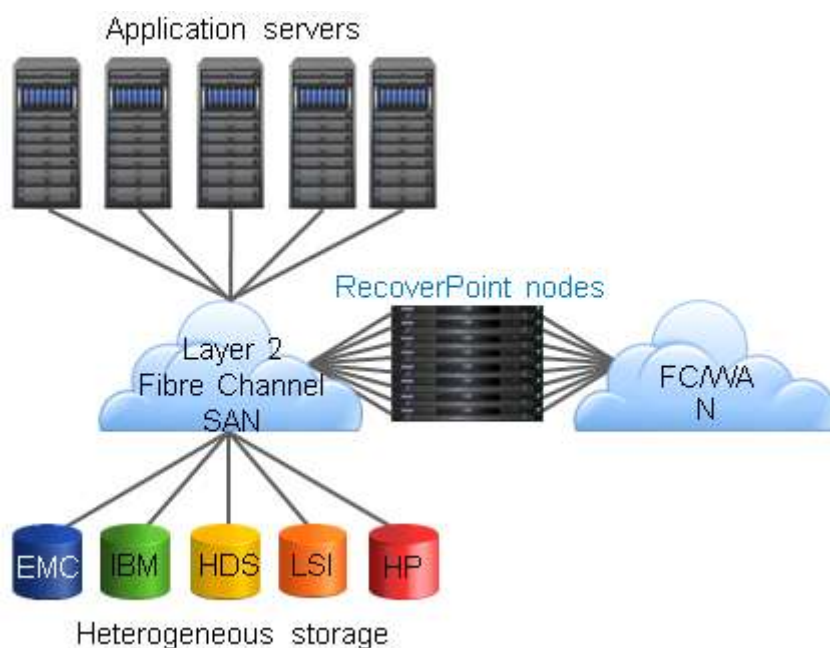
Each RecoverPoint appliance provides four physical Fiber Channel ports that are auto-sensing and support 2, 4, and 8 Gbps/s connections. Each RecoverPoint appliance provides two 1 Gbit Ethernet ports which are used for management and monitoring. The other is used for remote replication over the WAN. Hosts and storage arrays are connected to the RecoverPoint appliance using standard Fiber Channel SANs enabling host fan-in and array fan-out.

Each RecoverPoint cluster can define up to 64 consistency groups per RecoverPoint appliance, with all consistency groups transferring data at any one time. A RecoverPoint cluster can support up to 128 consistency groups. If one of the appliances fails, the consistency groups defined to the failed appliance will be temporarily transferred to the

remaining appliances, and data transfer will continue. Once the appliance is repaired, the consistency groups will be transferred back to their original appliance.

RecoverPoint can support up to 2,048 replication sets, with each replication set containing the production LUN, a local replica LUN, and/or a remote replica LUN. The maximum number of LUNs that can be managed is 2,048 production LUNs with 4,096 local and remote replicas for a total of 6,144 LUNs.

RecoverPoint supports EMC and third-party storage by using write-splitting technology. The function of the write splitter is to mirror writes to protected LUNs to the RecoverPoint appliance. The host driver is a host-resident lightweight driver residing at the bottom of the I/O stack, just above any existing multi-path software (such as PowerPath®). The EMC VNX splitter runs on the EMC VNX storage processor and supports write splitting for all of the VNX Fiber Channel and iSCSI volumes.



Refer to the EMC Support Matrix for a full list of the storage supported by RecoverPoint.

## Splitter

The magic in RecoverPoint, the Splitter makes a copy of every write I/O and sends it to RecoverPoint (RPAs) for replication, local or remote. To split these writes you need a write splitter.

There are three types of splitters.

- **Host-based** that is installed on the host itself just above the multi-path software
- **Fabric-based** that is installed within your Fiber Channel fabric switches (Brocade or Cisco)
- **Array-based** that is installed in FLARE on your array (VNX Only)

**Note:** *RecoverPoint provides out-of-band replication and therefore is not involved in the I/O process. This is important because often people suggest this impacts the I/O process but that is NOT true. Instead, a separate component of RecoverPoint, called the splitter (or KDriver), is involved.*

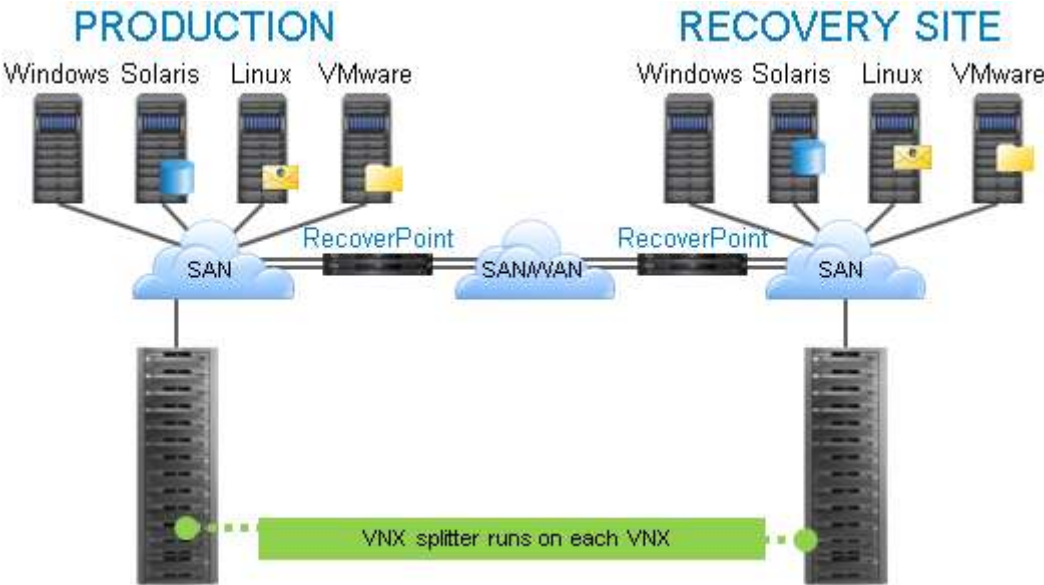
The primary function of a splitter is to split or “duplicate” application writes so that they are sent to their normally designated storage volumes and the RPA simultaneously. The splitter carries out this activity efficiently, with little perceptible impact on host performance, since all CPU-intensive processing necessary for replication is performed by the RPA.

A splitter is proprietary software that is installed on either host operating systems, intelligent fabric switches, or storage subsystems (see three types above). A host-based splitter resides on a host server that accesses a volume being protected by RecoverPoint. This splitter resides in the I/O stack, below the file system and volume manager layer, and just above the multi-path layer. This splitter operates as a device driver and inspects each write sent down the I/O stack and determines if the write is destined for one of the volumes that RecoverPoint is protecting. If the write is destined for a protected LUN, then the splitter sends the write downward and will rewrite the address packet in the write so that a copy of the write is sent to the RecoverPoint appliance. When the ACK (acknowledged back) from the original write is received, the splitter will wait until a matching ACK is received from the RecoverPoint appliance before sending an ACK up the I/O stack.

A fabric-based splitter is part of the storage services on intelligent SAN switches from Brocade or Cisco. These intelligent fabric-based write splitters operate at wire speeds and split writes, with the original sent on to the target LUN and a copy of the original sent to the RecoverPoint appliance.

A VNX storage processor also has a write splitter. When a write enters the VNX array (either through a Gigabit Ethernet port or a Fiber Channel port), its destination is examined. If it is destined for one of the LUNs being replicated by RecoverPoint, a copy of that write is sent back out one of the Fiber Channel ports of the storage processor to the RecoverPoint appliance. Since the splitter resides in the VNX array, RecoverPoint can support any open systems server that is qualified for attachment to the VNX array.

*Example of a VNX-based Write Splitter*



## Management

The RecoverPoint Management Application allows you to manage the RecoverPoint system. Site management provides access to all boxes in the local RPA cluster, as well as the RPA cluster at the other site. Almost all of the information necessary for routine monitoring of the RecoverPoint system appears on the RecoverPoint Management Application.

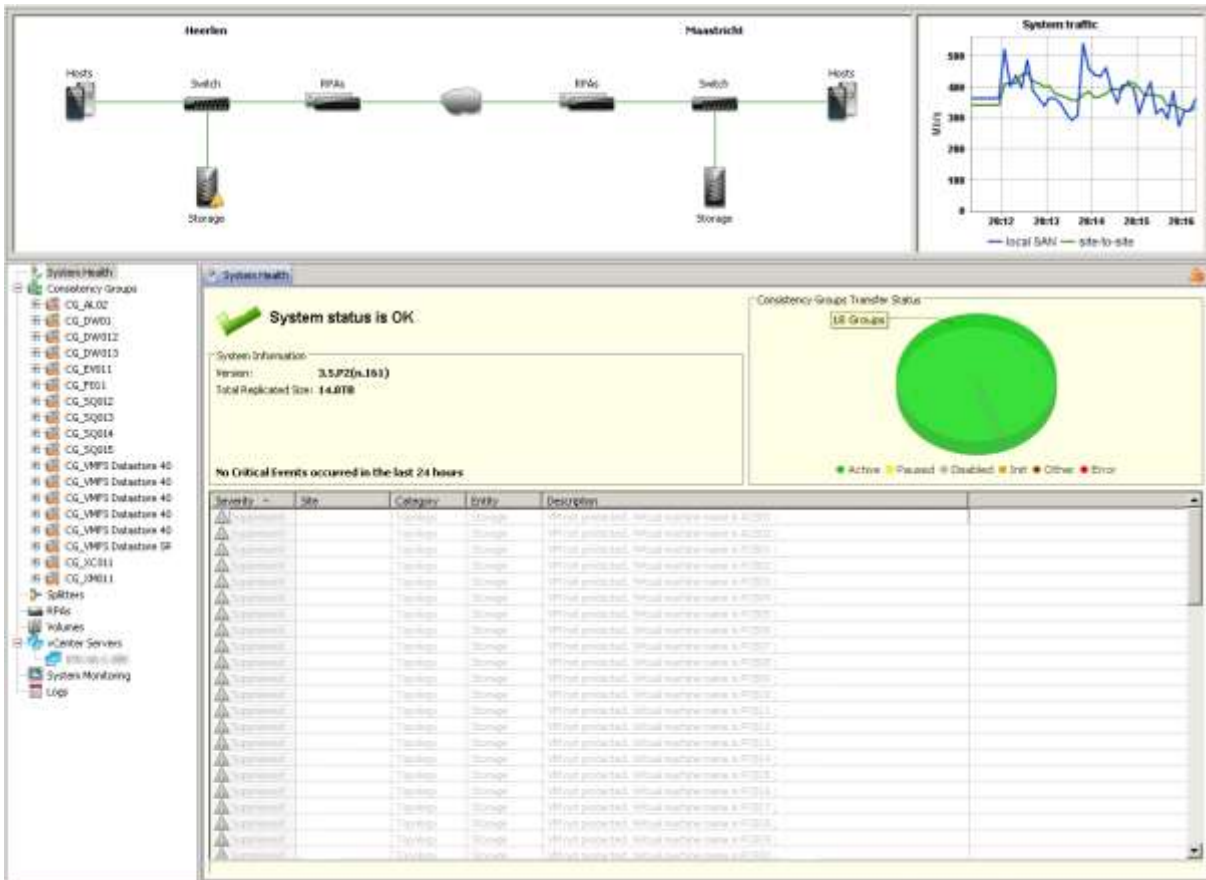
The System pane provides an overview of system health at a glance. The pane shows the status of major components of the RecoverPoint system environment, including the hosts, switches, storage devices, RPAs at two sites, and WAN connection. The Traffic pane displays the amount of SAN and WAN traffic passing through the RPAs.

The Navigation pane allows you to navigate to the different views available in the Component pane. In the Navigation pane, click the component on which you wish to focus. The corresponding view appears in the Components pane.

The RecoverPoint management application is a Java-based GUI that runs on multiple operating system platforms, including Windows, Linux, and UNIX. It is used for the initial installation of the RecoverPoint environment as well as for administration, configuration, monitoring, and recovery processes.

RecoverPoint is managed through the management IP interface on each RecoverPoint appliance. For high-availability support, a virtualized management IP address can be assigned for remote use. Primary management is through a web-based management application that supports HTTP or HTTPS. In addition to a web-based management interface, a command-line access (CLI) is available via SSH over SSL v3.

Example of the Management Application: System pane (top), Navigation pane (right), and Component pane (left).



How is the RecoverPoint management GUI positioned? As shown below, the management application GUI can run on any workstation running the OS as mentioned above or server that has the Java runtime environment (JRE) installed. Up to four simultaneous GUI sessions can be initiated on the same or across different platforms. In addition to managing RecoverPoint, users can monitor the health of RecoverPoint and its configuration as well as request and view RecoverPoint logs.



RecoverPoint includes integration with Unisphere, providing the ability to monitor and manage RecoverPoint replication from the same management console used to provision VNX storage. This is supported with RecoverPoint 3.3 SP1 or later and requires the VNX to be running FLARE 30. In the image above, I am running RecoverPoint version 3.5 SP1 P1 and VNX FLARE 32. Additionally, RecoverPoint must have been installed with the RecoverPoint Deployment Manager for this capability to be activated. If RecoverPoint was installed without Deployment Manager, it must be reinstalled with Deployment Manager to enable this capability.

## VMware Site Recovery Manager

Downtime is expensive! Disaster preparedness and recovery planning is an iterative process, not a one-time event. Traditional disaster recovery plans are complex and quickly get out of sync with evolving IT configurations. It's a good idea to replace your manual written plan with centralized recovery plans managed directly from VMware vCenter Server. You need to continually revisit disaster recovery plans to ensure they remain aligned with current business goals and test those plans regularly to ensure that they perform as planned.

I was one of the few who played with version 1.0 at VMworld in 2008 in Cannes. The latest version is 5.x and a lot has changed and improved. Because 'everything' is virtualized, recovery is very easy, be it to a second data center or the cloud.

VMware Site Recovery Manager (SRM)<sup>[3]</sup> provides business continuity and disaster recovery protection for virtual environments. Protection can extend from individual replicated data stores to an entire virtual site to ensure the simplest and most reliable disaster protection for all virtualized machines. SRM leverages cost-efficient vSphere Replication and supports a broad set of high-performance storage-replication products to replicate virtual machines to a secondary site. Of course, RecoverPoint is one of them. SRM provides a simple interface for setting up recovery plans that are coordinated across all infrastructure layers, replacing traditional error-prone plans. Recovery plans can be tested non-disruptively as frequently as required to ensure that they meet business objectives. How cool is that!

### Storage Replication Adapter

A storage replication adapter (SRA) is software provided by storage vendors that ensures integration of storage devices and replication with VMware Site Recovery Manager. Actually, it is nothing more than some scripts wrapped in a piece of software. However, it is good that it is there; many would not know where to start.



Basically, all the power is located in the adapter because storage vendors know their product better than VMware.

Nevertheless, VMware works closely with each of its storage partners to drive toward mutual support of SRM. Due to different product release cycles, levels of testing, and partner agreements, not all storage devices will be supported at the general availability date of a new version of SRM. I recommend contacting the storage vendor for the best information on when their device is planned to be certified with SRM.

Also note that the SRM compatibility matrix for a specific release only lists those arrays that are supported on the VMware SAN compatibility list (<http://www.vmware.com/resources/compatibility/search.php?deviceCategory=san>) with the corresponding version of ESXi. For example, storage array entries for SRM 5.0 will show only those arrays that are supported with ESXi 5.0.

For the full list of storage replication adapters supported by SRM 5.x, see <http://www.vmware.com/resources/compatibility/search.php?deviceCategory=sra>.

## Developing a Recovery Strategy

The primary task of this step is to determine how you will achieve your disaster recovery goals for each of the systems and system components that were identified. For most organizations, the design of a recovery strategy solution is a fairly custom process. While the design principles and considerations are mainly common, designers typically have to make a number of compromises.

Backup and recovery are components of business continuity, the term that covers all efforts to keep critical data and applications running despite any type of interruption (including both planned and unplanned). Planned interruptions include regular maintenance or upgrades. Unplanned interruptions could include hardware or software failures, data corruption, natural or man-made disasters, viruses, or human error. Backup and recovery is essential for operational recovery; that is, recovery from errors that can occur on a regular basis but are not catastrophic, i.e. data corruption or accidentally deleted files. Meanwhile, disaster recovery is concerned with catastrophic failures. Believe me, nothing is as interesting as a big failure because it's the moment you actually learn something.

When planning for a recovery, the time it takes to completely restore data and for business applications to become available is called the Recovery Time Objective (RTO). So, after determining your recovery time on the basis of the data to restore, you can determine how much time you actually have to perform your recovery.

I intentionally do not discuss Recovery Point Objective (RPO) because this is near zero with RecoverPoint. In fact, it can return to several points in time.

Only questions left are: What, When, Who, and How. To figure this out, I recommend reading another EMC Proven Professional Knowledge Sharing article I wrote:

[http://mikes.eu/download/2011KS\\_Mikes-Disater\\_Recovery\\_in\\_a\\_Cloudy\\_Landscape.pdf](http://mikes.eu/download/2011KS_Mikes-Disater_Recovery_in_a_Cloudy_Landscape.pdf) Or  
<https://education.emc.com/guest/certification/benefits/ks.aspx>

## How it works

I am not going to detail how to install and configure VMware Site Recovery Manager. There is a good administrator guide<sup>[3]</sup> which explains everything. My focus here is RecoverPoint. After you have installed SRM at the protected and recovery sites, you must connect the two sites to create a site pair. Install the appropriate storage replication adapters on the SRM server hosts at both sites.

After you have connected the protected site and recovery site, you must configure the array managers so that SRM can discover replicated devices, compute datastore groups, and initiate storage operations.

The array manager configuration wizard leads you through a number of steps. When the configuration process is complete, the wizard presents a list of replicated datastore groups. You typically configure array managers only once, after you have connected the protected and recovery sites. You do not need to reconfigure them unless array manager connection information or credentials have changed, or you want to use a different set of arrays.

## Consistency Groups

RecoverPoint Consistency Groups enable the user to associate the LUNs or volumes in their replication sets together so that all operations work consistently across all the replicated LUNs or volumes.

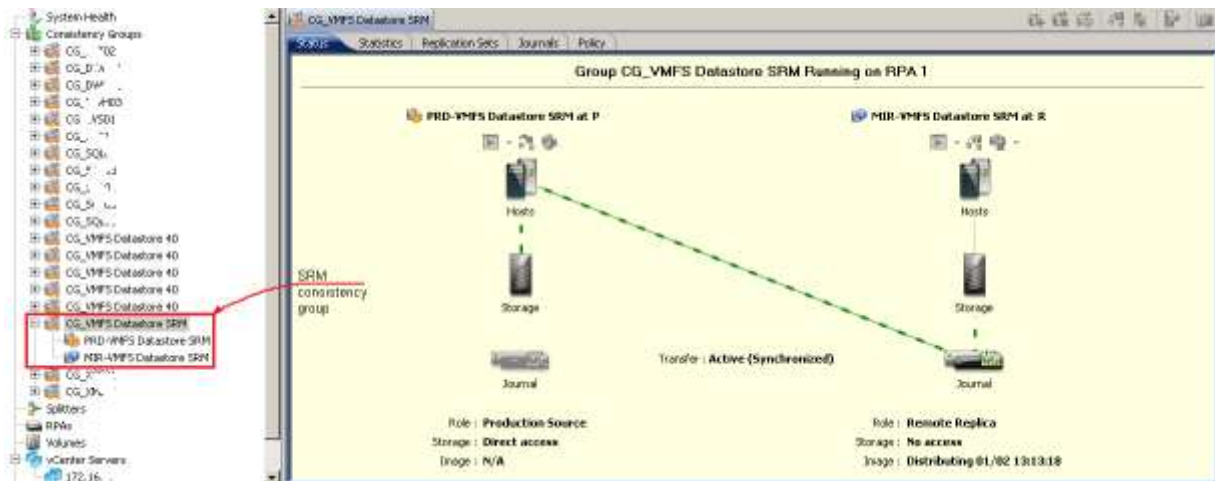
Another use for Consistency Groups is when different service levels need to be applied to different applications. The user can ensure that applications can be protected independently by creating a consistency group for each application and then managing the consistency group independent of other groups, such as performing local and/or remote recovery or starting and stopping individual consistency groups. Additionally, the user can set a specific SLA or policy optimizations such as application importance, use of RecoverPoint resources, and recommended RPOs on a per-consistency group basis.

Note: Be careful when virtual machines use RAW LUNs

## Policy

Replication with RecoverPoint is policy-driven. A replication policy, based on the particular business needs of your company, is uniquely specified for each consistency group.

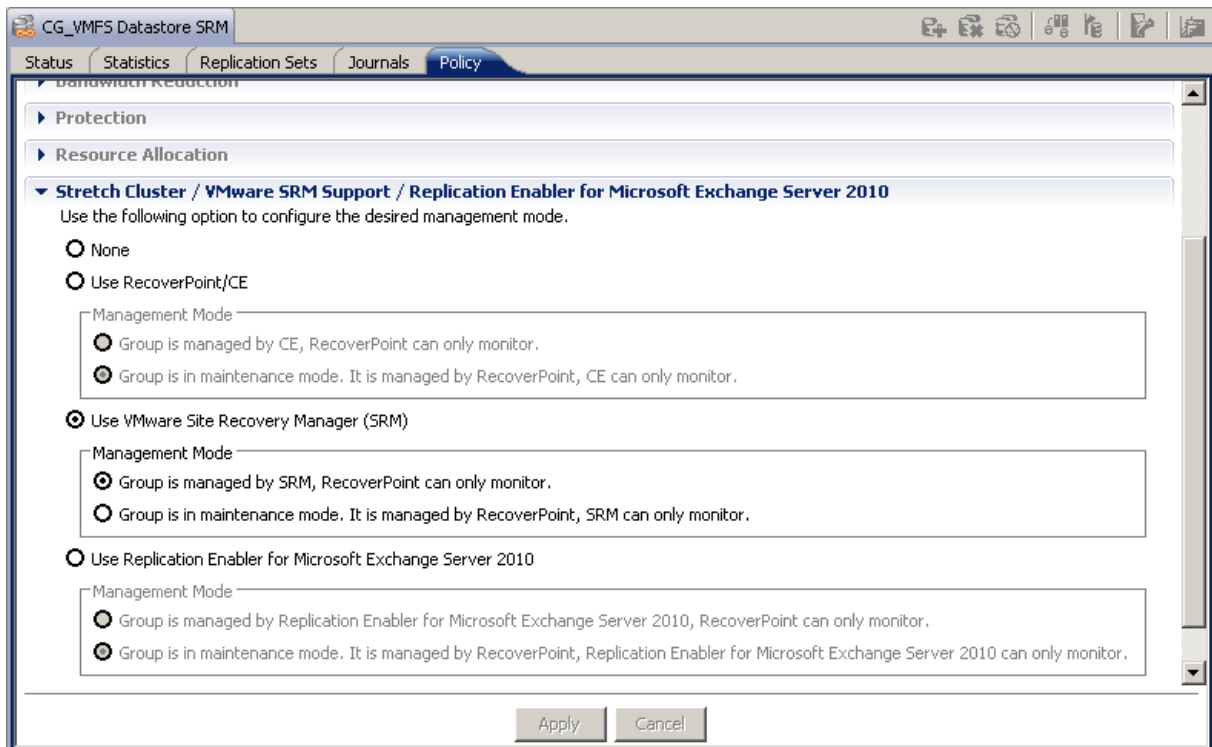
Replication behavior changes dynamically during system operation in light of the policy, the level of system activity, and the availability of network resources.



The left pane in the diagram above shows several Consistency Groups configured. The one we focus on is *CG\_VMFS Datastore SRM*.

This Consistency Group consists of a VMFS LUN of 2TB. This LUN contains all virtual machines that we are going to Protect with VMware Site Recovery Manager. The Replication Set contains two LUNs; Production (PRD) and Replica (MIR) LUN.

Use the following option to configure the desired management mode in Consistency Group Stretch Cluster / SRM Support Policy Settings



In this case, there are several settings to consider using SRM.

1. Use VMware Site Recovery Manager (SRM)

Check this option to enable VMware SRM support. This option is valid when a RecoverPoint Storage Replication Adapter for VMware Site Recovery Manager is installed on the vCenter Servers.

a. Group is managed by SRM. RecoverPoint can only monitor.

Check this option to activate VMware SRM support. When selected, VMware SRM manages the group and can perform failover and test failover from one site to the other. When activated, all RecoverPoint user-initiated capabilities are disabled. The user cannot access images, change policies, or change volumes. Bookmarks cannot be created in the RecoverPoint Management Application, but they can be created using the RecoverPoint command line interface bookmark commands.

b. Group is in maintenance mode. It is managed by RecoverPoint. SRM can only monitor.

Check this option for planned or unplanned maintenance of the RecoverPoint system. When activated, VMware SRM support is disabled and user-initiated RecoverPoint capabilities are enabled. When activated, all RecoverPoint user-initiated capabilities, such as image access, image testing, changing policies, and creating bookmarks are available.

Obviously, 1a should be selected to manage as much as possible from SRM and not RecoverPoint. From a RecoverPoint perspective you are good to go!

### **VMware SRM Protection Groups**

Once you are satisfied with your RecoverPoint array manager configuration you're ready for a next major step: configuring Protection Groups, which are used to protect virtual machines. After protection is established, placeholders are created and inventory mappings applied for each virtual machine in the group. If a virtual machine cannot be mapped to a folder, network, and resource pool on the recovery site, it is listed with a status of Mapping Missing, and a placeholder is not created for it. Wait for the operations to complete as expected to ensure that the protection group was created and virtual machines were protected.



One Protection Group can contain or point to one ESXi datastore. Alternatively, it is possible for one Protection Group to contain many datastores. This happens when virtual machines are spread across many datastores for disk performance optimization reasons or when a virtual machine has a mix of virtual disks and RDM mappings. An SRM Protection Group could be loosely compared to the storage groups or consistency groups you may create in your storage array.

In my example, I use three virtual machines, DC01 (a domain controller), VCS01 (the vCenter Server Machine), and TSRM01 (a test virtual machine). In a later stage, I use these three for a failover.

## VMware SRM Recovery Plan

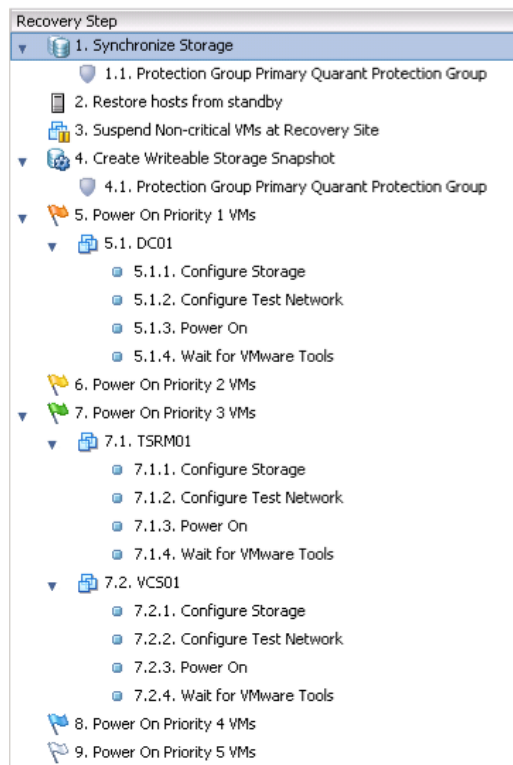
Finally, we reach the Recovery Plan. At this stage we have properly configured everything needed to carry out a recovery test and to establish how virtual machines are recovered. A basic recovery plan includes steps that use default values to control how virtual machines in a protection group are recovered at the recovery site. You can customize the plan to meet your needs. Recovery plans are different from protection groups in that recovery plans indicate how virtual machines in one or more protection groups are restored at the recovery site.

During tests, keep the virtual machine that is recovered during the test isolated from other machines in your environment. Errors can occur if duplicate machines are brought online and begin interacting with other machines in your production network. You can isolate virtual machines restored during test recoveries in an isolated network.

When we look deeper into this Recovery Plan deeper, we see the following tabs.



Let's focus on the Recovery Tab '**Recovery Steps**'. It displays the progress of individual steps. Of course, these steps are editable once made. Using Site Recovery Manager, organizations can execute automated tests of their recovery plans without disrupting their



environment. How cool is that! Site Recovery Manager makes it easy to create an isolated environment for testing while leveraging the recovery plan that would be used in an actual failover.

When performing a Non-Disruptive test SRM:

- Leverage storage snapshot capabilities to perform recovery tests without losing replicated data.
- Connect virtual machines to an existing isolated network for testing purposes.
- Automate cleanup of testing environments after completing failover tests.
- Automate execution of recovery plans.

## Run a Recovery Plan

Traditional recovery plans are often difficult to test and keep up to date, and depend on exact execution of complex, manual processes. Also, there is a real threat of causing damage when testing your DR plan. In a virtualized environment, testing is simpler because non-disruptive tests can be executed using existing resources. Hardware independence eliminates the complexity of maintaining the recovery site by eliminating failures due to hardware differences.

When testing your disaster plan, note anything that's not going according plan, and pass the plan back to the people who designed the plan so they can update it. This process improves the quality and accuracy of the disaster plan.

Clearly, periodic, realistic testing of the recovery plan is highly recommended to succeed in your mission.

Ready to perform a recovery? First, conduct a test. Because as said, a recovery plan makes significant alterations in the configurations of the protected and recovery sites and it stops replication. Do not run any recovery plan that is not tested. In the case of array-based replication, recovered virtual machines and services might need to be supported at the recovery site for a period of time. Reversing these changes might cost significant time and effort and can result in prolonged service downtime.

SRM will automate the entire failover process and bring your site online in a matter of seconds or minutes depending on the size of your virtual site. All virtual machines in the recovery plan are migrated to the recovery site and corresponding virtual machines in the protected site are normally shut down. Since we are testing in an isolated environment, servers remain on.

### **STEPS**

In the VMware client, click **Recovery Plans** in the left pane and click the recovery plan to run. In the command area, click **Test**.




Site Recovery Manager is running. At point 4, RecoverPoint kicks in.

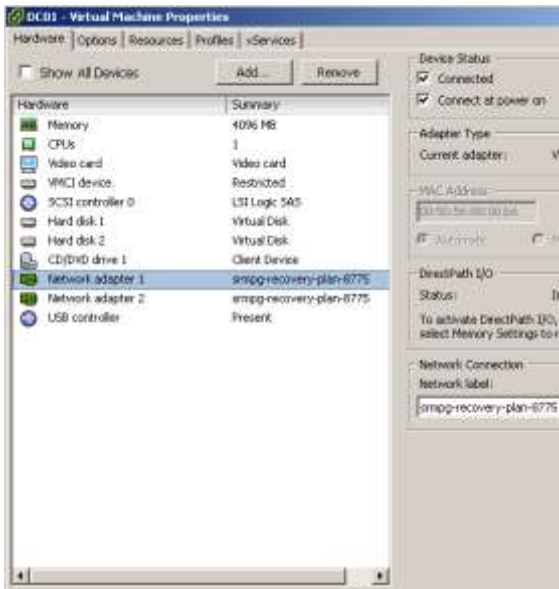
*Before Test*



*When Test is running*



When the test is running, RecoverPoint creates a writable snapshot. All writes are still replicated and saved in the Journal. The Snapshot will be presented to the VMware ESXi host at the Recovery site.



Recovery Step	Status	Step Started	Step Completed
1. Snapshot Storage	Success	1/7/2012 12:55:13 PM	7-1-2012 12:55:49
1.1. Protection Group Primary Quietest Protection Group	Success	1/7/2012 12:55:13 PM	7-1-2012 12:55:49
2. Restore hosts from standby	Success	1/7/2012 12:55:49 PM	7-1-2012 12:55:49
3. Suspend Non-critical VMs at Recovery Site	Success	1/7/2012 12:55:49 PM	7-1-2012 12:57:42
4.1. Protection Group Primary Quietest Protection Group	Success	1/7/2012 12:55:49 PM	7-1-2012 12:57:42
5. Power On Priority 1 VMs	Running	1/7/2012 12:57:42 PM	12%
5.1. DC01	Running	1/7/2012 12:57:42 PM	52%
5.1.1. Configure Storage	Running	1/7/2012 12:57:42 PM	12%
5.1.2. Configure Test Network	Running	1/7/2012 12:57:42 PM	12%
5.1.3. Power On	Running	1/7/2012 12:57:42 PM	12%
5.1.4. Wait for VMware Tools	Running	1/7/2012 12:57:42 PM	52%
6. Power On Priority 2 VMs	Running	1/7/2012 12:57:42 PM	11%
7. Power On Priority 3 VMs	Running	1/7/2012 12:57:42 PM	11%
7.1. TSPM01	Running	1/7/2012 12:57:42 PM	52%
7.1.1. Configure Storage	Running	1/7/2012 12:57:42 PM	12%
7.1.2. Configure Test Network	Running	1/7/2012 12:57:42 PM	12%
7.1.3. Power On	Running	1/7/2012 12:57:42 PM	12%
7.1.4. Wait for VMware Tools	Running	1/7/2012 12:57:42 PM	52%
8. Power On Priority 4 VMs	Running	1/7/2012 12:57:42 PM	11%
9. Power On Priority 5 VMs	Running	1/7/2012 12:57:42 PM	11%

The picture above shows the isolated network label which prevents servers from becoming visible as double identity in your domain.

## AppSync

Another great functionality is AppSync<sup>[2]</sup> which offers a self-service approach for protecting virtualized Microsoft applications in EMC VNX deployments. AppSync is an advanced protection management software that offers a better way to manage the protection and replication for critical business applications and databases. Application owners can protect and recover their own data quickly and easily using this technology and enables application owners and database administrators to make their own copies and restore their own data. This is a fairly new product that has a great potential.

### SLA-Driven

Offers multiple service levels



Application	Service Plan	Measurement	Technology
 vmware	 <b>GOLD</b>	<b>ZERO DATA LOSS</b>	<b>SYNC REPLICATION</b>
 SQL Server	 <b>SILVER</b>	<b>RPO = MINUTES</b>	<b>ASYNC REPLICATION</b>
 Exchange Server	 <b>BRONZE</b>	<b>RPO = HOURS</b>	<b>HOURLY SNAPSHOTS</b>

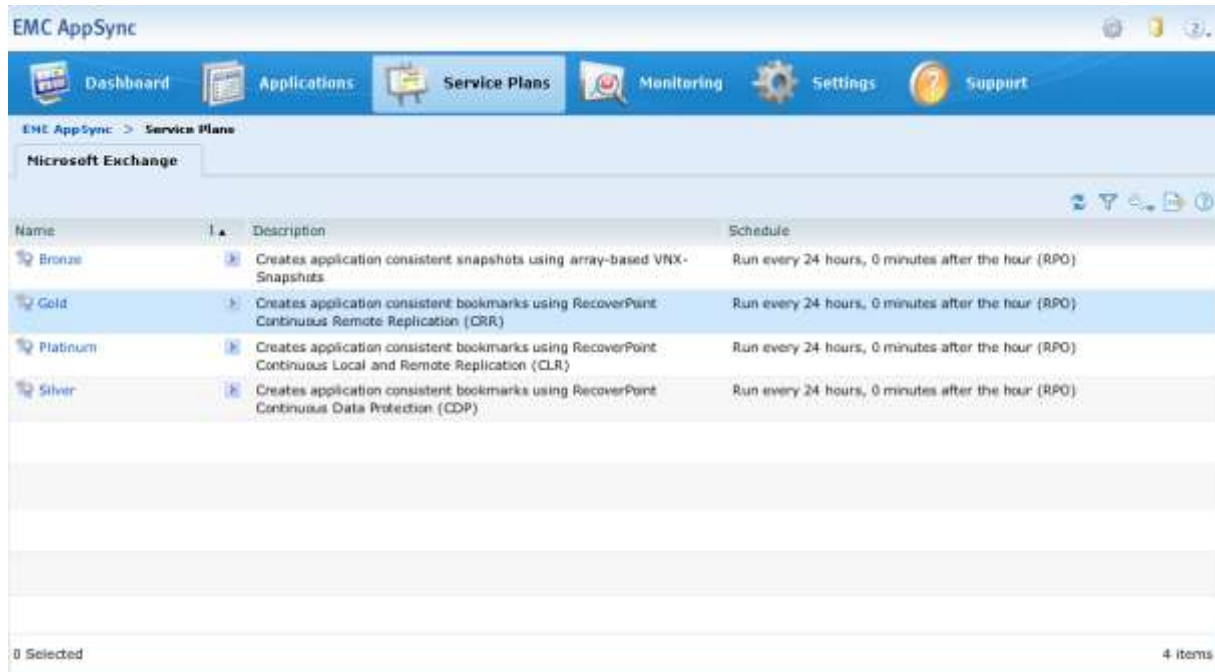
EMC AppSync offers a self-service SLA-driven approach for protecting virtualized Microsoft applications such as SQL Server, Exchange, and SharePoint in VNX deployments. The service plan runs immediately. A panel displays progress as application storage is discovered and mapped. Application protection begins according to service plan settings.

AppSync<sup>[4]</sup> creates application-consistent copies of Exchange, SQL, and SharePoint using VNX Snapshots or RecoverPoint bookmarks. Like Replication Manager, AppSync uses VSS to create application consistency on, for example, Exchange. Unlike Replication Manager, AppSync uses VNX Snapshots where Replication Manager uses SnapView Snapshots.

VMware administrators can use AppSync with EMC Virtual Storage Integrator (VSI) plug-in to VMware data stores directly from VMware vCenter. An EMC-developed plugin to the VMware vCenter management software, VSI It enables you to provision, monitor, and

manage VMware vSphere datastores on EMC storage arrays directly from vCenter, greatly simplifying management of virtualized environments.

AppSync has a Unisphere look-a-like interface (see below) that is very handy since many people are familiar with this.



## Conclusion

Besides people, the most valuable asset to a business is data. Without its people to maintain the equipment, even the most sophisticated and powerful machinery would cease to function. Without people performing day-to-day operations, the organization would stop functioning. This statement makes data replication important.

Systems became smarter and we were overtaken by new technology. As applications evolved, so did the backup technology. This is not altered by time. Applications become more complex and data is growing. It's logical that these technologies evolve too.

A different form of backup is the concept of snapshots. A snapshot is a copy of a file system just as backup is. However, the technology is different. Snapshot technology can be implemented on the host, in the storage network, or at the array level. Snapshots tend to be less disruptive to applications and environment.

A continuous data protection (CDP) product is designed to monitor changes to one or more data objects and store a copy of these changes in a journal. A CDP product is either file system-centric, where the object is a file, or storage block-centric, where the object is the LUN. File system CDP products are typically found in Microsoft Windows environments, and usually offer a file system. 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. How do you choose the right replication method for optimal data protection? This question is important for further recovery strategies. How, what, and where to recover is key!

After choosing at least the how, what, and where, there's one bridge to take; Recovery Point Objective (RPO) and Recovery Time Objective (RTO). If you've found answers to these questions, you can proceed to a product selection. RecoverPoint is a single solution that provides the advantages of host-based and array-based solutions while replicating data from any SAN-based array to any other SAN-based array over existing Fibre Channel or IP networks using any combination of host-based, VNX based, CLARiiON®-based, or intelligent fabric-based write-splitting options. Both RecoverPoint and RecoverPoint/SE provide synchronous local replication using CDP, synchronous and asynchronous continuous remote replication (CRR), and concurrent local and remote (CLR) data protection. The RecoverPoint family protects companies from data loss due to common problems such as server failures, data corruption, software errors, viruses, and end-user errors, while also protecting against catastrophic events that can bring an entire data center to a standstill.

We agreed on the importance of Data Protection. In addition to data loss, downtime is expensive too. Disaster preparedness and recovery planning is an iterative process, not a one-time event. Traditional disaster recovery plans are complex and quickly get out of sync with evolving IT configurations. You need to continually revisit disaster-recovery plans to ensure they remain aligned with current business goals and test those plans regularly to ensure that they perform as planned. In my estimation, VMware Site Recovery Manager does this.

VMware Site Recovery Manager provides business continuity and disaster recovery protection for virtual environments, ensuring the simplest and most reliable disaster protection for all virtualized machines. Site Recovery Manager provides a simple interface for setting up recovery plans that are coordinated across all infrastructure layers, replacing traditional error-prone plans. Recovery plans can be tested non-disruptively as frequently as required to ensure that they meet business objectives.

AppSync is another great functionality, offering a self-service approach for protecting virtualized Microsoft applications in EMC deployments. AppSync, an advanced protection management software, offers a better way to manage protection and replication for critical business applications and databases. Application owners can protect and recover their own data quickly and easily using this technology. It enables application owners and database administrators to make their own copies and restore their own data.

There are many more functionalities of RecoverPoint than I discussed in this article. I suggest you read the documentation on [emc.com](http://emc.com). There are not many products that I know of with the same functionalities, simplicity, and versatility as RecoverPoint. Bottom-line, RecoverPoint is a great product.

## References

- [1] Proven Professional Knowledge Sharing article 2011, <http://www.mikes.eu>
- [2] <http://www.emc.com>
- [3] <http://www.vmware.com>
- [4] <http://convergingclouds.com/2012/09/07/emc-appsync/>

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.