

Disaster Recovery for Windows Using EMC RecoverPoint/Cluster Enabler

Applied Technology

Abstract

EMC® RecoverPoint/Cluster Enabler software integrates with Microsoft Failover Cluster software that uses Node Majority and Node and File Share Majority quorum modes to allow geographically dispersed cluster nodes to replicate their data using RecoverPoint/SE and RecoverPoint continuous remote replication. EMC RecoverPoint/CE addresses common cost bottlenecks in implementing a disaster recovery solution for Microsoft Windows Server environments.

October 2011

Copyright © 2008-2011 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.

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

Part Number h5935.4

Table of Contents

Executive summary	4
Audience.....	4
Introduction	5
Cluster Enabler.....	5
Cluster introduction.....	6
The concept of a cluster	6
Cluster benefits and limitations.....	6
Solving problems by clustering.....	6
Limitations.....	7
EMC RecoverPoint concepts	8
EMC RecoverPoint/Cluster Enabler concepts	9
Conclusion	10
References	11

Executive summary

EMC® RecoverPoint/Cluster Enabler (RecoverPoint/CE) software enables geographically dispersed Microsoft Failover Clusters to replicate their data using EMC RecoverPoint, MirrorView™, and SRDF®. When Cluster Enabler is used with RecoverPoint, then continuous remote replication (CRR) can be used for replication between cluster nodes. Geographically dispersed clusters offer increased levels of high availability, disaster recovery, and automation over non-clustered solutions.

RecoverPoint/CE works seamlessly with applications designed to take advantage of Failover Clusters, including Exchange Server 2003, Exchange Server 2007 and SQL Server, in Microsoft Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server Core for Windows Server 2008, and in Hyper-V environments. RecoverPoint/Cluster Enabler does not support Microsoft Exchange Server 2010. RecoverPoint/Cluster Enabler (RecoverPoint/CE) is the version of Cluster Enabler that supports RecoverPoint. Cluster Enabler version 4 has a new base and plug-in configuration that requires both components to be at the same version level. Additionally, RecoverPoint/CE has been updated to version 4.1.1 and supports RecoverPoint version 3.4.2 and later.

With Windows operating environments proving suitable for mission-critical application deployments, even in enterprise data centers, demands for business continuity solutions in Windows environments have flourished. However, the affordable heritage of Windows is often challenged by other costly components in the overall solution. In terms of disaster recovery solutions for Windows environments these high-cost components can encompass enterprise-class storage platforms, storage software licenses, network equipment and communication links, and integration services.

Data availability as part of disaster recovery is performed by remote data replication (RecoverPoint/CE) and application failover (Microsoft Failover Clusters). Both of these components are necessary to satisfy the recovery point objective (RPO) and recovery time objective (RTO) for the application.

Audience

This white paper is intended for any customers, partners, or EMC personnel wishing to gain a general understanding of how RecoverPoint/CE improves disaster recovery for Windows-based applications. It is intended to serve as a technical overview document as opposed to a step-by-step guide. For more in-depth information, please see the product guide and software release notes on EMC Powerlink®. Readers are presumed to have a basic understanding of RecoverPoint and Microsoft Failover Clustering technology.

Introduction

This white paper first covers RecoverPoint/CE, then the concept of a cluster and finally it discusses the EMC RecoverPoint family and its operation with RecoverPoint/CE

Cluster Enabler

Cluster Enabler allows existing Microsoft Failover Cluster customers to extend the protection of their current solution to include site disaster recovery. It reduces RTO for existing Microsoft customers by automating resource and application failover between sites.

By deploying Cluster Enabler, companies can:

- Realize significant cost savings compared to traditional array-based disaster recovery solutions
- Mitigate risks with a complete, integrated, Microsoft-certified solution
- Obtain greater investment protection and return on investment from RecoverPoint/CE's deployment flexibility and functional optimization
- Achieve competitive advantage through superior RTO and RPO compared to competitive solutions




Figure 1 RecoverPoint/CE logical view

RecoverPoint/CE 4.1.2 is an add-on software component to RecoverPoint. It supports Microsoft Failover Cluster nodes that use Node Majority and Node and File Share Majority quorum modes. It supports nodes that have the storage remotely replicated synchronously or asynchronously by RecoverPoint. It supports RecoverPoint CRR consistency groups and tolerates CLR consistency groups where Cluster Enabler

manages the remote replica and the local replica is manually managed. It does not support nodes with storage resources replicated locally with CDP.

Cluster introduction

This section introduces the concept of server clustering and its benefits and limitations. It then goes on to discuss cluster organization, infrastructure scaling, and cluster operating modes and how clustering is used on multiple, geographically dispersed sites.

Concept of a cluster

The concept of server clustering involves taking two or more computers and organizing them to work together to provide higher availability, reliability, and scalability than can be obtained by using a single server. When failure occurs in a server cluster, resources can be redirected and the workload can be redistributed. Typically, the end user experiences a limited failure, and may only have to refresh the browser or reconnect to an application to begin working again.

Cluster benefits and limitations

A server cluster provides high availability by making application software and data available on several servers linked together in a cluster configuration. If one server stops functioning, a process called failover automatically shifts the workload of the failed server to another server in the cluster. The failover process is designed to ensure continuous availability of critical applications and data.

While clusters can be designed to handle failure, they are not fault tolerant with regard to user data. The cluster by itself does not guard against loss of a user's work. Typically, the application software handles the recovery of lost work; the application software must be designed to recover the user's work, or it must be designed in such a way that the user session state can be maintained in the event of failure.

Solving problems by clustering

Clusters can be used to solve three typical problems in a data center environment:

- **High Availability** refers to the ability to provide end-user access to a service for a high percentage of scheduled time while attempting to reduce unscheduled outages. A solution is highly available if it meets the organization's scheduled uptime goals. Reducing unplanned downtime and then working to improve total hours of service operation achieve availability goals.
- **High Reliability** refers to the ability to reduce the frequency of system failure, while attempting to provide fault tolerance in case of failure. A solution is highly reliable if it minimizes the number of single points of failure and reduces the risk that failure of a single component/system will result in the outage of the entire service offering. Reliability goals are achieved using redundant, fault-tolerant hardware components, application software, and systems.

- **High Scalability** refers to the ability to add resources such as computers while attempting to improve performance. A solution is highly scalable if it can be scaled up and out. Adding additional resources (for example, CPUs, memory, disks) can scale the service up. Adding additional computers can scale the service out.

A well-designed service solution uses redundant systems and components so that the failure of an individual server does not affect the availability of the entire service.

Limitations

While a well-designed solution can guard against application failure, system failure, and site failure, cluster technologies do have limitations. Cluster technologies depend on compatible applications and services to operate properly. The software must respond appropriately when failure occurs. Cluster technology cannot protect against failures caused by viruses, software corruption, or human error. To protect against these types of problems, organizations need solid data protection and recovery plans. The RecoverPoint family provides the data protection and recovery necessary to protect against these types of failures and EMC RecoverPoint/CE integrates with Microsoft Failover Clusters to automate the application and data failover.

RecoverPoint/CE is a software extension of Failover Cluster functionality. RecoverPoint/CE allows Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2 and Windows Server Core for Windows Server 2008 running Microsoft Failover Clusters to operate across multiple connected storage arrays in geographically distributed clusters. Formerly Microsoft Failover Clusters were called server clusters, which used Microsoft Cluster Services (MSCS) for the Failover Cluster functionality.

A geographically dispersed cluster (GDC) can be implemented with Microsoft Failover Clusters to enable automated application failover in a DR solution. A GDC stretches the nodes of a Microsoft high-availability cluster over extended (or geographic) distances. Note that excessive latency on the communication link between the GDC nodes can delay the transmission of the “heartbeats” of the production node to the point that the delay causes the secondary site to falsely activate and try to take ownership of the operation – a condition known as the split-brain syndrome. The latency constraints of GDC implementations often considerably limit the maximum distances between sites compared to the maximum distances supported by the underlying replication technologies.


Microsoft has addressed this limitation with Failover Clusters in Windows Server 2008 and Windows Server 2008 R2. However, with MSCS on Windows 2003, the maximum latency between nodes in a geographically dispersed cluster must be less than 500 ms to avoid the split-brain syndrome. Due to this latency constraint, Windows MSCS GDC solutions in the market today can be separated by only a few hundred kilometers.

EMC RecoverPoint concepts

The EMC RecoverPoint family provides local and remote data protection, enabling reliable replication of data over any distance; that is, locally within the same site, and/or remotely to another site—even halfway around the globe. RecoverPoint CRR uses either Fibre Channel or an existing IP network to send the data snapshots to the remote site using techniques that preserve write-order. Replicated point-in-time images are retained in a journal where they can be accessed in the future such as for an application restart from a historic image.

The EMC RecoverPoint family utilizes continuous data protection (CDP) technologies as part of its data replication. Using RecoverPoint CRR, users can achieve lower RPO and RTO over traditional host- or array-based replication. RecoverPoint CRR accomplishes this through a combination of small-aperture snapshots, synchronous replication, dynamic synchronous replication (which dynamically changes between synchronous and asynchronous replication), dynamic bandwidth management, write-order preservation, and an on-disk journal that records all the changes between point-in-time images.

RecoverPoint supports consistency groups that ensure write-order consistency across related mirrored volumes. This approach is common practice for transactional databases such as Microsoft SQL Server, where log volumes and database volumes must be kept logically consistent with one another. It also applies to federated applications, such as SAP, where applications and databases operating across multiple server and/or storage arrays must be kept logically consistent with each other. CRR provides bi-directional replication with no distance limitation and guaranteed data consistency offering complete protection from site disasters. RecoverPoint's CDP technology ensures that applications can be recovered to any



specified point in time with no data loss.

Figure 2. RecoverPoint CRR architecture

RecoverPoint appliances can be inserted into the network with minimal disruption to applications and without incurring overhead on either servers or storage. RecoverPoint software will automatically discover the data volumes available at both sites and provide wizards to configure the consistency groups and replication sets necessary to replicate the LUNs between the sites.

Once the LUNs are replicated by RecoverPoint and available to the cluster nodes, the EMC Cluster Enabler Manager GUI is used to configure the cluster and storage system environment. The GUI is a Microsoft Management Console (MMC) snap-in. The EMC Cluster Enabler Manager GUI provides wizards that validate that the proper software components exist on the cluster nodes, and discovers existing cluster nodes, their corresponding storage systems, and the device assignments. It then converts the cluster to a RecoverPoint/CE cluster and converts any existing cluster groups to RecoverPoint/CE groups. The RecoverPoint/CE configuration wizard can also be used to add additional nodes to the cluster as part of the conversion process. This feature is most useful for adding the geographically dispersed nodes to the cluster if they are not already present, such as when converting a cluster to a geographically dispersed cluster.

EMC RecoverPoint/Cluster Enabler concepts

RecoverPoint/CE controls the CRR aspects of moving cluster resources between storage systems. RecoverPoint/CE only supports LUNs that belong to CRR consistency groups; it will not support LUNs that belong to a CDP consistency group. It will tolerate CLR consistency groups by failing over the CRR component. When using CRR without RecoverPoint/CE, users must configure separate clusters at the primary and secondary sites. In this case, moving volumes between the clusters involves unmounting LUNs from the cluster at the primary location, making the replica available to the secondary cluster, and then redirecting application clients to the new cluster. By integrating with Microsoft Failover Clusters, RecoverPoint/CE enables the cluster to automatically manage volume failover in addition to application failover, which greatly improves recovery time objective.

Figure 1 on page 5 is a high-level depiction of the RecoverPoint/CE environment. Cluster nodes are geographically dispersed over two sites. CRR is used to replicate data volumes P1 and P2 at the primary site, Site 1, to data volumes S1 and S2 at the secondary site, Site 2. In the event of a local node failure, the Failover Cluster software will move the application to Node 2. In the event of a site failure, the cluster group can be moved manually to Site 2 at the user's discretion or automatically by Failover Cluster software in the event of a failure. In either case the latest version of the data volume replicas S1 and S2 will be presented to the nodes in Site 2.

Geographically dispersed cluster (GDC) technologies can be implemented with RecoverPoint/CE to enable automated application failover in a DR solution. GDC stretches the nodes of a Microsoft high-availability cluster over extended (or geographic) distances. RecoverPoint/CE software is the only EMC software element that you must install on each cluster node. RecoverPoint/CE uses the RecoverPoint software APIs over a secure communications path (using SSH sessions) to control the replication environment. RecoverPoint/CE manages the replicas at the RecoverPoint consistency group level.

RecoverPoint/CE enables any node to be a member of a GDC, and with its integration with Microsoft Failover Cluster services it ensures that the replicated data is presented to the correct nodes in the event of a failover event. RecoverPoint/CE

requires a RecoverPoint consistency group for each Failover Cluster group. Therefore, as cluster resources are moved between sites, the movement occurs in a consistent manner across all of the volumes in the cluster group.

Note that excessive latency on the communication link between the GDC nodes can delay the transmission of the heartbeats of the production node to the point that the delay causes the secondary site to falsely activate and try to take ownership of the operation. The latency constraints of GDC implementations often considerably limit the maximum distances between sites compared to the maximum distances supported by the underlying replication technologies.

Microsoft has addressed this limitation with Failover Clusters in Windows Server 2008. However, with Microsoft Cluster Services (MSCS) on Windows 2003, the maximum latency between nodes in a geographically dispersed cluster must be less than 500 ms. Due to this latency constraint, Windows MSCS GDC solutions in the market today can only be separated by a few hundred kilometers. However, as a result of its advanced replication mechanism, RecoverPoint/CE is capable of supporting replication across any distance involved in a GDC configuration.

Conclusion

Businesses can achieve significant competitive advantage by adopting best-of-breed IT solutions based on cutting-edge technologies. Ultimately, an enterprise should seek a solution that leverages the many different advances in data protection and apply these technologies in a way that makes sense for their budgets and their environments. By integrating with Microsoft Failover Clusters, EMC RecoverPoint/CE enables customers to secure competitive advantage. Together with RecoverPoint CRR, it is a comprehensive data protection solution for the entire data center, addressing a complete range of data protection challenges while significantly reducing the cost and complexity of protecting business-critical data. By deploying RecoverPoint/SE CRR or RecoverPoint CRR and RecoverPoint/CE customers can:

- Replace multiple existing data protection, replication, and recovery solutions with a single solution, therefore reducing cost and management complexity
- Reduce their dependence on storage-consuming snapshots, lowering overall storage costs
- Allow replication to and from any SAN-based storage array, allowing more cost-effective storage arrays to be used at the remote site or the reuse of existing storage
- Reduce the data transferred between sites due to replication by up to 90% by using RecoverPoint's data compression and deduplication technology
- Perform replication with write-order consistency across standard networks
- Reduce acquisition and operational costs compared to traditional array-based replication solutions

- Utilize bi-directional replication and recovery between sites regardless of distance, allowing the use of existing data centers as DR sites, lessening infrastructure and operational costs
- Secure competitive advantages by having a lower RPO and RTO for business applications in the event of a local or regional disaster

References

For more information, please see the following documents on EMC Powerlink and EMC.com:

- *EMC RecoverPoint/Cluster Enabler Version 4.1.2 Product Guide* (Powerlink only)
- *EMC RecoverPoint/Cluster Enabler Version 4.1.2 Release Notes* (Powerlink only)
- *EMC RecoverPoint Family Overview* white paper
- *EMC RecoverPoint/Cluster Enabler — A Detailed Review* white paper