

Combining CDP and Replication for Comprehensive HA/DR Solutions: EMC RecoverPoint 3.1

December 2008



All enterprises, regardless of size, must provide recovery options for both applications and data. Regular backups combined with manual application recovery provide the minimum, but as recovery requirements become more stringent, enterprises look for faster, more granular data recovery capabilities coupled with automated application recovery. Many of these enterprises, whether they are driven by business or regulatory mandates, must also provide multi-site disaster recovery (DR) solutions. Candidate technologies to address these issues include continuous data protection (CDP), replication, and high availability (HA) clustering.

Enterprises can deploy multiple products, attempting to cobble together the solutions they need to provide comprehensive HA/DR capabilities, or they can turn to integrated solutions that combine multiple technologies. The new breed of “CDP” appliances includes replication options because the two technologies are so synergistic when it comes to improving data protection, but to date these offerings have left cluster integration to the customer. With EMC’s RecoverPoint 3.1 release, we see a considered combination of all three that supports off-the-shelf capabilities unique in the industry. In this Product Profile, we’ll take a close look at how well matched CDP and replication technologies are to addressing critical data protection issues such as backup window, recovery point objective (RPO), recovery time objective (RTO), recovery reliability, and DR, and then explore how RecoverPoint’s features map to solving problems in these areas. With this latest RecoverPoint offering, EMC is bringing not only complementary technologies together to address critical customer needs, but is continuing to deepen RecoverPoint’s integration with EMC’s product portfolio and other critical enterprise products in ways which improve flexibility, speed deployments, and extend EMC’s centralized management strategies.

The Need for Combined HA/DR Solutions

In today’s new regulatory environment, an enterprise’s HA needs go beyond just local solutions. The use of a multi-site DR capability is mandated by the government in certain industries, but even when it is not driven by regulatory issues there are few enterprises that would not benefit from more

comprehensive availability management solutions that can cover everything from the loss of a single file all the way up to the loss of an entire site. Local site solutions can handle issues such as data loss and server failure, but do not provide the depth to deal with manmade disasters like flooded data centers and widespread power outages or natural disasters like hurricanes, tornadoes, and earthquakes. Best practices recommend

P R O D U C T P R O F I L E

that secondary copies of critical data for recovery be kept at remote sites that are at least 100 miles away so that large-scale disasters at one site do not affect secondary sites.

To maximize overall availability, enterprises must deploy strategies that provide not only for data recovery but also for automated application recovery. While there are a number of different approaches to provide data recovery, replication technologies allow secondary copies of data to be continuously maintained at remote sites, with the most recent data state always available for recovery. But replication by itself offers only the most recent recovery point, and if for some reason that copy is corrupted, enterprises may not be able to recover. For this reason, replication is often combined with snapshot technologies to ensure the availability of multiple recovery points. While snapshots generally provide more recovery granularity than the legacy “once daily” backups, they do not provide as good a recovery granularity as continuous data protection (CDP). CDP transparently splits application writes as they occur, sending one copy to primary storage and a second copy to a local, disk-based journal. This journal allows administrators to select any previous point in time, and perform very rapid restores directly from disk. Given shrinking backup windows and increasingly stringent RPO and RTO, CDP offers significant value for data protection in that it supports near real-time RPOs, enables very short RTOs, and imposes no backup windows. When combined with CDP, replication creates a very granular, very resilient data recovery solution that covers both local (HA) and remote (DR) requirements.

Clusters have historically been the mainstream technology of choice to implement automated application recovery solutions. Clusters back “primary” servers with “secondary” servers that can automatically take over an application workload, including the application’s disks and client-side facing IP addresses, for automated recoveries that do not require any human intervention or new client-side access processes. Built around industry-standard hardware and software components, clusters can provide a good availability management solution for almost any application environment. Manual and automated “failover” choices provide flexible options to deal with either planned or unplanned downtime. While clusters are an effective solution, they have been limited to local configurations.

Most clustering solutions allow cluster nodes to be placed up to 30 kilometers (km) apart, which limits their effectiveness in a disaster scenario. “Geo-cluster” solutions on the other hand, allow the remote nodes to be located at geographically dispersed locations, hundreds or thousands of miles away. By combining CDP, asynchronous replication, and clustering, HA/DR solutions can be crafted that provide for rapid recovery directly from disk without data loss for most failures, and recovery from a very remote (i.e. unlimited distances) site with very minimal data loss to handle that rare but very extensive DR situation.

Enter EMC RecoverPoint

EMC has a strong pedigree in open systems replication, having introduced the first synchronous replication solutions in this

P R O D U C T P R O F I L E

space in the mid 1990s. As DR requirements evolved, EMC introduced solutions that supported heterogeneous, asynchronous replication capabilities to enable long distance solutions that flexibly supported additional storage array options. RecoverPoint, EMC's entry in this latter space, is an out-of-band, network-based appliance that combines CDP and replication into a single solution that provides extremely granular recovery to application-consistent recovery points to meet both local and remote recovery requirements. The RecoverPoint product line includes RecoverPoint, targeted for use in heterogeneous storage environments, and RecoverPoint/SE, a lower cost and slightly less feature-rich version targeted for use in CLARiiON-only environments.

Early CDP appliances were bare bones solutions supporting only CDP – in effect technologies looking for problems to solve. EMC's product strategy with RecoverPoint, however, has leveraged a strong CDP foundation but added critical complementary technologies, like snapshots, replication, and compression, to create a very efficient, true application recovery oriented solution. Multiple "source" volumes can be combined into consistency groups within which write order fidelity is maintained, and each consistency group can accept writes from multiple servers concurrently to reliably support federated applications. Local data capture can be performed at the host (through operating system-specific filter drivers), in the network (through documented APIs such as Cisco's SANtap and Brocade's Fabric Application Platform), and in storage arrays (EMC's CLARiiON CX3 or CX4 arrays). This flexible approach to data

capture allows a wide variety of heterogeneous operating system platforms to be supported, including Windows, Linux, and Unix, as well as both physical and virtual machine environments. Standalone and clustered server data sources are supported. Integrated snapshot capabilities allow snapshot backups to be automatically created from the CDP journal per administrator-defined policies. Support for synchronous and asynchronous replication of SAN-attached storage provides flexible options for distributing data to remote locations near and far, while the simultaneous use of CDP and replication enables very granular recovery capabilities in all sites. Support for replication between dissimilar storage arrays enables the creation of lower cost configurations. In addition to supporting "any point in time" (APIT) recovery, the use of a policy engine allows RecoverPoint to leverage application-specific APIs to "bookmark" transaction-consistent recovery points within applications which are readily available to administrators through the RecoverPoint management GUI. Built-in bandwidth reduction and compression enables sites to be kept in sync through replication even in very low bandwidth environments. RecoverPoint configurations feature automatic local failover for high availability using RecoverPoint appliance (RPA) pairs, and can support bi-directional replication of different consistency groups between appliances.

EMC RecoverPoint 3.1: Raising the Bar

In November 2008, EMC shipped RecoverPoint 3.1, including important new

P R O D U C T P R O F I L E

features to improve application recovery and overall availability as well as performance and scalability in single and/or multi-site configurations. Key new features include “stretched CDP” configurations, snapshot consolidation, virtual provisioning support, an “asynchronous replication over Fibre Channel” (FC) option, and longer distance options for Microsoft Cluster Services (MSCS) configurations.

Stretched CDP. Stretched CDP uses synchronous replication to give RecoverPoint customers two additional configurations to improve deployment flexibility. What stretched CDP support does is allow the CDP journal and target volumes to be located in a second site that can be up to 30 km away, offering increased resiliency based on a two site configuration. In this configuration, the RPA (or appliance pair) is located at site 2, and the write splitters at the primary site send data directly to it; there is no RPA at the primary site. This configuration offers additional resiliency to primary site failures without the additional expense of CRR (RecoverPoint replication).

Stretched CDP provides another option though, allowing RecoverPoint customers to create what are called “cascaded” configurations. In creating DR configurations, there has been a technology dilemma. Synchronous replication, which will maintain the data states at two sites in perfect lockstep, imposes latencies which limit configuration distances to generally no more than 30 km (roughly 18 miles). While asynchronous replication overcomes the distance limitation, there is generally a lag time in applying changes to the secondary

site which can result in data loss on recovery. In two site configurations, administrators have been forced to choose between limited distances or possible data loss. RecoverPoint’s support for stretched CDP, combined with CRR, provides an innovative way to solve both problems in a single solution called a “cascaded” configuration.

Cascaded configurations are 3 site configurations which include a primary, a secondary, and a tertiary site. Stretched CDP is used to create the primary and secondary sites, and data at these sites is kept in sync with the tertiary site using CRR (which replicates between the RPAs at the secondary and tertiary sites). CRR’s asynchronous replication allows the tertiary site to be hundreds or thousands of miles away. In this cascaded replication configuration, if the primary site fails, recovery occurs to the secondary site and replication between the secondary and tertiary sites continues to operate without disruption, eliminating any chance for data loss.

Cascaded configurations offer recovery without data loss from events which affect the primary site, 99% of which will not be extensive enough to affect both the primary and secondary sites. However, in the event that a catastrophic disaster that affects both sites occurs, you still have the tertiary site to provide recovery, albeit with some possible data loss. Relative to two-site configurations, cascaded configurations offer the advantages of recovery without data loss from the great majority of primary site failures with the safety net of recovery from widespread catastrophes that solutions using only synchronous replication could not offer.

P R O D U C T P R O F I L E

Snapshot consolidation. Several key aspects of how administrators want to use RecoverPoint have emerged over the last several years. First, while very granular recovery capabilities can be extremely important for very recent data, the value of extremely granular recovery declines over time. From several hours up to a day or two, the ability to generate recovery points from any previous point in time has value in fine-tuning RPOs and performing root cause analysis of data corruption problems. Eventually, this level of granularity is no longer needed, and the amount of CDP journal space required to support this APIT recovery capability can become an expensive proposition. Second, when actually performing recoveries (as opposed to root cause analyses), administrators have a strong preference for recovering from application-consistent (rather than crash-consistent) recovery points. To give customers what they need in the most cost-effective manner, EMC has introduced policy-driven snapshot consolidation.

Snapshot consolidation allows customers to define separate time periods and apply different data retention policies to each. For example, for data less than 24 hours old, all data, including any inserted bookmarks, can be retained to support APIT recovery. For data that is older than 24 hours but less than 72 hours old, an administrator could specify a policy that “collapses” retained CDP journal data into hourly recovery points (which may or may not be application-consistent). For data older than 72 hours, administrators can define snapshot consolidation policies to retain weekly, monthly, or quarterly recovery points.

Snapshot consolidation policies can be set at the consistency group level. Snapshot consolidation saves significant CDP journal space while still providing extremely granular recovery capabilities where they are most needed (with very recent data) along with very easily accessible older recovery points.

Virtual provisioning support. Virtual provisioning enables you to present a large amount of storage capacity to a server, but then actually consume physical space only as it is needed from a shared pool. This improves TCO by reducing initial over-allocation of storage capacity, reducing costs and simplifying ongoing management. Virtual provisioning is supported in EMC’s CLARiiON and Symmetrix storage arrays, and with the 3.1 release RecoverPoint is now “virtual provisioning aware”. It recognizes “thin” LUNs and will not automatically convert “thin” LUNs to “thick” LUNs during the replication process, preserving the physical capacity space savings that virtual provisioning offers. RecoverPoint supports “thin-to-thin” LUN replication, but can also allow mixing of “thin” and “thick” LUNs where this is desirable.

Asynchronous replication over FC. Historically, RecoverPoint CRR has previously supported asynchronous replication modes over IP. With this release, asynchronous replication can now also be performed across FC, subject to FC distance limitations. FC infrastructures can often support higher throughput than IP infrastructures, and in certain cases this additional performance is highly desirable. There are certain requirements to support

PRODUCT PROFILE

this configuration – dual fabrics and write splitters at both sites – but this option provides the flexibility to support higher performance asynchronous replication configurations.

RecoverPoint/CE for MSCS Geo-clusters. On Windows, MSCS (Microsoft Cluster Services) is the most widely deployed HA cluster solution by far. RecoverPoint/Cluster Enabler (CE) provides a long distance replication option for improved disaster tolerance. Targeted for use with Windows Server 2003 and Windows Server 2008, Enterprise and Data Center Editions, RecoverPoint/CE is an off-the-shelf solution that supports “geo-cluster” configurations with automated failover/failback between MSCS nodes at sites that are geographically distributed without distance limitations. In use, administrators are likely to reserve the decision to fail over remotely to themselves, but leverage RecoverPoint/CE’s automation to make that failover a simple, one-step administrative process.

In providing this solution, RecoverPoint/CE has resolved a key historical stumbling block to geo-cluster configurations: split brain syndrome. Split brain syndrome is a condition which can arise when communication between cluster nodes fails and two or more nodes view themselves as the new “primary”, a situation which can lead to data corruption. To address split-brain syndrome, local MSCS configurations use a quorum disk, but as implemented by Microsoft the quorum disk cannot be used to address this same issue for geographically distributed clusters. RecoverPoint/CE implements a geographically-aware “file

share witness” to address this same problem for these geographically distributed configurations to ensure reliable computing even in the event of failures that cause split brain syndrome.

Improved performance and scalability. RecoverPoint is targeted for use in enterprise-class environments so high performance and scalability are critical requirements. With this release, EMC has made some significant extensions in this area:

- Improved CDP and replication throughput, although actual improvements vary by configuration
- Each appliance can now support up to 64 consistency groups, 32 of which can be active
- Compression performance has been improved by up to 20%, resulting in more efficient use of available network bandwidth
- Up to 6,144 LUNs are now supported in a single RecoverPoint configuration

Taneja Group Opinion

CDP and replication are both valuable data protection technologies, providing options for addressing backup window, RPO, RTO, recovery reliability, and DR issues. When deployed together, however, their whole is greater than the sum of their parts. And when deployed together in an integrated solution where each technology is cognizant of the other, they become indispensable in meeting stringent application recovery requirements both locally and remotely.

P R O D U C T P R O F I L E

RecoverPoint combines these technologies with a clear focus on rapid and reliable application – not just data – recovery. And, as with many solutions from EMC, integration points with EMC portfolio products and other widely deployed third party vendor technologies help to make it a comprehensive solution that is easy to deploy. Case in point: RecoverPoint provides the highest granularity data capture in EMC's data protection product line, and integrates with other key enterprise products, like Replication Manager, NetWorker, Symmetrix, and CLARiiON to simplify management, and other third-party applications, like Microsoft MSCS, Exchange, and SQL Server (through support for VSS-based snapshot backups available prior to 3.1), and Oracle and other databases (through application recovery kits also available prior to 3.1), so that this recovery granularity can be harnessed easily (without custom scripting) and in meaningful ways to improve both local and remote data protection processes. Write-splitters all

support CDP/replication for both virtual and physical machine environments, regardless of whether they are deployed on the host, in the fabric, or in an array. In the latest release, support for stretched CDP (including cascaded 3 site replication configurations), snapshot consolidation, and longer distance options for MSCS (RecoverPoint/CE) in particular give RecoverPoint customers and prospects unique and valuable options just not available as off-the-shelf products from other vendors. If backup windows, recovery granularity, and long-distance DR are areas where you could improve your existing data protection infrastructure, you'll want to review what the RecoverPoint product line has to offer.

***NOTICE:** The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.*