

EMC Replication Manager

August 2007



Without any fanfare, data replicas have quietly become the lifeblood of enterprise storage infrastructures. In any given data center, a myriad of data replicas exist for every mission critical process imaginable. For example, data protection, business continuity, testing, and environment upgrades now all depend entirely on sophisticated data replication technologies. While it is easy enough to *create* a replica, it is a much harder task to *manage and integrate* replicas within complex IT operations.

To that end, we find in our ongoing interviews that many IT teams have yet to make an appropriate investment in the management methodologies or solutions that could automate and improve their data replication. Stuck in their highly manual, error-prone processes and customized scripts, many enterprises are setting themselves up for non-optimal performance due to the absence of a centralized replication management approach. But with the right replication management products, it is possible to transform replication into a powerful, smooth running, policy-driven process with very few manual interventions.

Such technologies do exist, and one example is the EMC Replication Manager offering. Designed to orchestrate a wide range of heterogeneous replication activities, EMC Replication Manager is the kind of product that any EMC customer can leverage to establish ongoing strategic control over replicas. In this profile, we will take a detailed look at replication management and what EMC Replication Manager can do for customers and their businesses. In our opinion, this kind of product is non-optional for any data center that leverages replication technologies today.

Why Orchestrate Replicas?

Creating data replicas has become an absolutely essential component of any data center environment. We have very quickly evolved over the past decade from a state where relatively simple local replication technologies and straightforward data control scripts have blossomed into incredibly complex, multi-faceted processes. Replication now routinely spans multiple applications, storage platforms, and data replication software products.

In our interviews with IT decision-makers, local data replication is consistently mentioned in the same breath as the mission critical application environments they support. In short, *the local replica of information is now as important as the production copy itself*. Why is this so? Because a usable, consistent local data replica creates a pivot point for all kinds of advances across the entire infrastructure, freeing up resources from the application,

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

server, network and storage layers to handle other production workloads. Not surprisingly, several key business and IT processes are now entirely enmeshed with data replicas:

- Disk-based recovery for instant restore of databases
- Disk-based back-up to eliminate data protection windows
- Data repurposing for instant refresh of development, testing or reporting
- Production data access assurance
- All forms of environment upgrades
- Information Lifecycle Management to transport replicas of application data to lower cost infrastructures

These ubiquitous data center processes are 100% dependent on local disk-based data replicas in today's enterprise. Yet each one also represents its own dimension of planning and management complexity. As complexity grows, local replication procedures are prone to become equally complex and unwieldy. Applications, servers, storage, protection schema, and service levels change over time and require corresponding changes in how local replicas must be created, managed, mounted, and accessed.

The historical answer to this complexity has been to "automate" *portions* of the replication process through increasingly intricate customized scripts. However, since most IT teams lack adequately deep expertise in such areas, the opportunity to mismanage local replication becomes quite real. As the number of applications grow, the number of replicas grow and scripting becomes incapable of handling the complexities of

scheduling. As a result, it becomes difficult to keep business goals aligned around changing protection practices and service goals and increasingly difficult to deeply integrate with a broadening and complex pool of multi-tier applications. Furthermore, hardware or software changes can render replication scripts obsolete and require significant re-engineering long after the knowledge and expertise supporting the original scripts have vanished. In short, a scripted environment may slowly spiral out of control and become more difficult and expensive to manage as an infrastructure grows and changes over time.

For this reason, we believe that *orchestration and pervasive automation* of local data replication processes through advanced software is a new requirement. Users can gain entirely new levels of control by placing the creation, management and usage of local data replication in the hands of a centralized software product.

So, why should customers invest in a new technology to orchestrate data replication? Simply put: replication management software provides many technological advantages over traditional approaches - all of which add up to business-level efficiencies. We outline these technological advantages as follows:

Advantage: Pervasive Automation

An advanced replication management product enables IT teams to centrally automate their entire range of replication activities, not merely individual replica practices around an array, host or application. By removing and eliminating

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

future needs for traditional scripting approaches, data replication processes will become streamlined, adaptive, and highly reliable.

Advantage: Ease of Use

An advanced replication management solution can deliver ease of use that cuts through complex practices surrounding different technologies and applications. Whether an application requires multiple storage volumes to be quiesced, the use of different replication technologies on different volumes, or any variety of different complex steps, an advanced replication product can use underlying policies and structured best practices to simplify decision making and planning for replicas. Furthermore, an advanced replication product can summarize and convey information about the on-going state of replication, identify potential storage changes that have not been taken into account, and allow administrators to easily confirm the configured state of protection as well as the success of replication activities.

Advantage: Deep Integration

An advanced replication management solution integrates directly with mission critical operations systems, applications and databases so that replicas can be created, managed and mounted for instant restores, repurposing, or ILM without manual intervention. This is a key distinction that separates the “old way” of manually managing data replication integration on an application-by-application basis from the new way of managing and automating replicas. Interestingly, in this new paradigm of a centralized point-and-click replication

management platform, all application integration issues are already resolved; the administrator simply defines replication scheduling and service parameters. This effectively masks all of the former complexity of application-level integration.

Advantage: Adaptation and Scalability

An advanced replication management solution can scale seamlessly to meet the needs of the largest enterprise while being flexible enough to address every scattered storage pool or any type of storage architecture in a single unified solution. This means that as the needs of the data center change, the same common replication management platform will scale easily and without disruption. Whether managing a few terabytes or many petabytes, the same replication management solution will be seamlessly available to the mission critical applications that it services. This stands in sharp contrast to manual processes and custom scripts that are notoriously non-scalable and failure-prone as the infrastructure changes.

Delivering True Efficiency

And so, what is the ultimate impact of these numerous advantages? Specifically, we see that backup windows, recovery management, business continuity, and overall storage management practices gain higher levels of efficiency when a replication management platform is deployed. This adds up to real ROI gains for the storage environment, but *only* if end users pick a replication management solution that is well suited to their infrastructure. As such, we will now take a look at how EMC Replication Manager

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

provides the opportunity for EMC end users to markedly improve data replication, streamline business processes and achieve real economic returns.

Meet EMC Replication Manager

EMC Replication Manager is a host-based software application that orchestrates all aspects of heterogeneous data replication for the IT administrator. Leveraging a very simple to use point-and-click GUI, Replication Manager takes control of the entire end-to-end process of discovering, creating, managing, and executing automated replicas and recovery images.

Replication Manager leverages its client-server architecture to scale across the entire enterprise. It enables the user to conduct automated discovery of applications, hosts, and storage resources across the infrastructure (each Replication Manager Server can support up to 30 agents.) This topology then becomes the basis against which the administrator can assign and manage all data replication activities from a single centralized interface. To be sure, this is no small task to integrate heterogeneous replication capabilities so deeply into a storage infrastructure.

EMC Replication Manager's interoperability list is long, including all of today's major enterprise operating systems (e.g. Windows, Solaris, HP-UX, IBM AIX, VMWare ESX Server, and Linux), major enterprise file systems and volume managers (e.g. NTFS, Ext3, VxFS, VxVM, ASM, UFS, HP-HFS, JFS/JFS2, IBM AIX LVM and HP LVM), and the most popular clustering technologies (e.g.

MSCS, VCS, IBM HACMP, Sun Cluster, and HP/Service Guard, Oracle ASM with Oracle RAC).

Replication Manager broadly integrates with mission-critical applications to ensure consistent replicas of production data. Replication Manager delivers total automation within a sophisticated framework that can keep replica practices aligned with industry-wide best practices as well as enterprise policies and practices surrounding service levels, protection practices, and storage utilization. To better understand how the EMC Replication Manager product has achieved this goal, we have identified and grouped its key differentiators into the following descriptive categories.

Differentiator: Deep Automation

At the highest level, this technology's most powerful differentiator is its automation capabilities for all aspects of data replication. Whether EMC customers are utilizing SnapView, SAN Copy, the TimeFinder family, Celerra SnapSure, Celerra Replicator, Invista clones, RecoverPoint CDP, or advanced features on top of replication such as federated environments or protected restore, Replication Manager orchestrates all of these storage technologies effectively.

Replication Manager considers replication software products in an end-to-end context with respect to top-level applications, host operating systems, file systems, drivers, primary disk volumes, and replication disk volumes. This comprehensive view means that Replication Manager can ultimately mask significant complexity and automate

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

the entire range of activities required for data replication from quiesce to best practice execution. In fact, Replication Manager comes with well developed templates and practices that organizations can leverage to improve their replica practices, optimize storage utilization and replica rotations, and provide optimal protection when combining multiple complex replica and protection technologies.

Furthermore, with application intelligence built on top of this comprehensive view, Replication Manager can maintain the environment in the face of changes by discovering alterations to storage and hosts. Replication Manager then adapts to maintain the same level of service prior to the changes. This is particularly valuable during upgrades that might otherwise throw a manually managed replica environment into a significant re-engineering cycle, or in storage environments that grow through the addition of new volumes to complex storage groups. With manual replication management, these volumes might be left out of critical replicas, but with Replication Manager, they are automatically added to existing configurations.

The practical implication of EMC Replication Manager's broad interoperability and deep integration with the infrastructure is that the vast majority of end users can deploy Replication Manager with confidence and no integration concerns.

Differentiator: Ease of Use

Replication Manager simplifies on-going management by scheduling recurring events

or ad hoc events from its centralized GUI, and integrating pre and post replication processes that were previously manual tasks. For example, Replication Manager can schedule replication, provide mounts to an authorized host, and execute a post-replication backup process, making the reuse of replicas for backup or development simple and interaction free. Such simplification makes it possible to structure replicas in this fashion with a few mouse clicks.

Along with this simplification, Replication Manager incorporates five levels of access for granular user control that can allow non-storage administrators, such as DBAs and application developers, to manage data replicas with ease and make changes to jobs and schedules at will.

Replication Manager rolls up all replication into a master console for a 10,000 foot view of all replication activities throughout the enterprise. Users can drill down on specific items, easily viewing the overall state of protection and status of individual activities. As a result, Replication Manager provides a single comprehensive framework for replication management that turns what used to be multi-person in-depth engineering projects into a single person, GUI managed administrative activity. Furthermore, the product provides the flexibility to delegate specific tasks to application or database staff.

Differentiator: Sophisticated Management

Replication Manager's capabilities deliver not only simple, automated replication, but also a sophisticated management framework

PRODUCT PROFILE

that increases the effectiveness of enterprise replica use. Replication Manager contains built-in best practices that can be used to control and optimize the implementation of replicas. These policies walk administrators through vendor recommended approaches to replicas, such as technology used and the number of replicas maintained.

This comprehensive policy and best practices engine encompasses different combinations of protection technologies. For example, Replication Manager supports the simultaneous use of Exchange snapshots with RecoverPoint CDP. Using both Exchange snapshots and RecoverPoint CDP, Replication Manager helps administrators choose the optimal number of data copies to protect Exchange and minimize overall storage utilization. Moreover, once a replica is created, Replication Manager automatically determines whether or not a desired replication schema is indeed supportable, and alerts the administrator to any out of bound conditions or lack of best practice conformance.

Equally important, Replication Manager uses a sophisticated set of policies to create alignment with agreed upon service and protection levels. This allows enterprise-wide policies to be created that mandate the frequency and execution timing of these replication activities, along with optimizing storage tier utilization, and mandating retention, expiration and reclamation of replicas.

Altogether, the sophisticated management in Replication Manager allows organizations to guarantee best practices and maintain high

service levels. To this end, Replication Manager can simultaneously enforce best practices around data protection, reduce operational costs, and improve infrastructure utilization by synchronizing and coordinating these critical protection activities. Replication Manager's sophisticated management framework manages these activities in lock-step according to pre-defined enterprise policies.

Differentiator: Application Centric Execution

A critical differentiator for Replication Manager is the technology's direct integration with mission critical applications. Replication Manager integrates tightly at the application level (e.g. Oracle, MS Exchange, MS SQL Server, and IBM DB2 UDB) to provide intelligent automation of replicas. Specifically, Replication Manager ties in with various consistency state tools to manage quiescing or hot backup modes and captures application consistent replicas. The offering comes ready to automate all aspects of replication planning and execution for these applications. As a result, Replication Manager can manage everything from scheduling to recovery consistency for each application instance. It should be noted that the offering can support a full range of consistency states on a per application instance basis, including fully consistent recovery, crash consistent recovery, or application consistent recovery.

For example, Replication Manager integrates directly with Microsoft's VSS technology, to create, mount and restore VSS compliant replicas. Additionally, for Oracle

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

environments, Replication Manager integrates with management APIs enabling it to bring Oracle databases in and out of hot back-up mode. On top of any type of replica management, Replication Manager can use Federated Storage Groups to create consistent replicas for multiple subsystems that make up an application.

The implications of this level of application-centric replication integration are huge: IT teams can literally set service-level driven schedules and policies across heterogeneous mission critical applications. Ultimately, from a common centralized point-and-click GUI, administrators can guarantee compliance with downtime windows, for an entire infrastructure. This completely eliminates the historical need for diving into application-specific management for every replication activity. That alone constitutes a massive step forward in efficiency.

Differentiator: Storage Integration

Replication Manager also provides a very high level of integration with the EMC CLARiiON CX series, Symmetrix-DMX, and iSCSI Celerra. In particular, it merits noting that Replication Manager has been optimized to inter-operate with the new CX3 UltraScale series CLARiiON running SnapView and SANCopy replication software.

As an example of the importance of storage integration and replica management in the context of applications, let's review the triad of CLARiiON CX plus SnapView/SANCopy and Replication Manager. For example, for any CX deployment with SnapView, Replication Manager provides the

intelligence to automate the entire replication lifecycle and eliminate scripting. The orchestration moves from (1) discovery of resources to (2) storage allocation for SnapView replicas, (3) retention policy creation, (4) placement of the top-level application into a consistent state, (5) execution of the SnapView replication function, to (6) engaging third party backup applications, as required. This level of automated offload translates into zero backup window data protection for EMC CX series customers using SnapView. In our opinion, the amount of human management and manual effort required to equal that workflow is non-trivial. This is a big win for CLARiiON shops.

Replication Manager's Impact

There is no doubt that deploying a technology like EMC Replication Manager can bring about significant benefits to the enterprise. In our assessment of this offering, we have identified 4 general areas of impact that merit consideration by any potential end user of this technology. We encourage prospective buyers to contextualize these impact areas for their own infrastructure to assist in the evaluation process.

Impact: Operational Efficiency

One of the biggest business level impacts that an IT team will notice when they deploy this kind of replication management solution is how streamlined their operations become. With no more manual configurations, application owners can now drive their own replication requirements, and recovery operations can take place via a simple point-and-click. That shaves untold weekly hours

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

off of most team schedules. We know from experience that there are hard dollars in team efficiency attached to that kind of a transformation.

Impact: Streamlines Protection

By making replicas easy to use and more broadly applied within an organization, Replication Manager can reinvent data protection. Replicas have always been useful as a data protection instrument for large mission critical systems by creating a replica for mounting and backup via a dedicated system that does not impact operations, as well as keeping the replica in place after backup as a “gold copy” that can provide almost instant data restore. But the usefulness of that approach was limited by management complexity, especially in the face of trying to guarantee recovery consistent replicas instead of crash consistent replicas.

Replication Manager makes achieving a zero backup window environment throughout an organization very straightforward. Replication Manager automates and centralizes the myriad of discrete steps required to bring the application into a recovery consistent state, create a replica, and execute any required 3rd party backup application processes. Replication Manager transforms a formerly complex process into what is literally a point-and-click exercise, while providing deep application-aware integration that assures known-good, recover-consistent replicas are created. We anticipate that the ROI gains resulting from simple, fast and efficient disk-based recovery will likely be where many teams look to

justify their investment in Replication Manager.

Impact: Easy Data Reuse

Replicas have also long supported production data duplication for testing and development, but have similarly been limited in usefulness by management complexity. Replication Manager enables organizations to address their development and testing needs pro-actively and comprehensively. Rather than typically reactive fulfillment of development needs with crash consistent data copies, enterprises can easily manage automated data refreshes that provide clean, recovery consistent data sets. As a result, organizations are able to make much broader use of data copies, and increase the quality of their development and testing efforts.

Impact: Enables Service Level Focus

Perhaps the most important impact is that this enables the storage team to establish a service-level orientation for replication. Because the offering provides for an entire range of scheduling and policy-based controls integrated with mission critical applications and storage platforms, it becomes a key product for closing the gap between business-level service objectives and their replication environments. Without the ability to automate across all relevant levels of the replication infrastructure, IT faces a difficult task in achieving availability, RTO and RPO metrics. Without masking the complexity of replication components, this kind of efficiency is not possible. As such, we believe that Replication Manager can play a key role in enabling service level-focused

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

management, and the resulting business efficiencies in enterprise data centers.

Taneja Group Opinion

The old days of relying on scripts to manage the replication schema for an enterprise storage environment are now over. The complexity that we see in terms of growth rates, supported applications, and variety of storage management products all point to the need for advanced approaches and increasing levels of automation. The new class of software product that we refer to as “replication management software” directly addresses this challenge with a range of new functionality specifically designed to abstract out the tactical challenges of point replication technologies and automate the end-to-end workflow.

Simply put, replication management software will become an essential part of any storage infrastructure. In fact, given the complexity trajectory that we see in data centers today, we will go so far as to say that the notion of managing data replicas *without* a replication management product soon will be unthinkable for most companies.

Replication Manager is EMC’s offering in this very useful new category. We believe that virtually any EMC customer who critically evaluates the potential benefits of this offering will come away nodding their heads in agreement. Replication Manager solves a longstanding challenge through smart integration and automation across a wide

range of mission critical elements (*please note our chart at the end of this report that itemizes EMC Replication Manager application coverage areas versus usage cases discussed in this profile.*)

From an ROI perspective, we expect that most customers will justify the investment in Replication Manager based on its ability to demonstrably improve disk-based recovery operations *and* the management efficiency gains that result from reducing or eliminating the cost of developing, maintaining and documenting custom scripts. Administrators can use Replication Manager to create replicas for a variety of purposes and applications, at will. Both factors are directly quantifiable and we encourage customers to evaluate this technology along those lines.

Replication is now a mission critical activity. As such, it now demands powerful management technologies. By eliminating the anxiety of failure-prone custom-built scripts and manual integration, Replication Manager provides a way for IT to meet increased service levels and enforce protection policies across the enterprise in an automated and centralized approach. We are very encouraged to see that products, like Replication Manager, are empowering end users to address these replication challenges. For any EMC customers, we believe that EMC Replication Manager is the clear choice.

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

EMC Replication Manager Applications and Use Cases

		Use Cases			
		Recovery	Back-up	Repurposing	ILM
Applications	Exchange	✓	✓	N/A	✓
	SQL Server	✓	✓	✓	✓
	Oracle	✓	✓	✓	✓
	DB2 UDB	✓	✓	✓	✓

***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.*