

# Affordable Data Protection without the Compromise

*Leveraging Midrange Storage Platforms for Tiered Data  
Protection and Business Continuity*

---

## **Abstract**

Continuing storage technology developments have enabled tiered business continuity capabilities on midrange storage platforms. As a result comprehensive business continuity solutions are now more affordable and can be deployed in more areas in the enterprise. This white paper outlines the growing list of cost effective options that are available to address organizations' increasing emphasis on the protection of critical data assets.

3/23/2007

---

---

Copyright © 2007 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on [EMC.com](http://EMC.com)

EMC<sup>2</sup>, EMC, CLARiiON, RepliStor, Symmetrix, and where information lives are registered trademarks of EMC Corporation. MirrorView, SAN Copy, and SnapView are trademarks of EMC Corporation. All other trademarks used herein are the property of their respective owners.

S12700307V1 / H2703

---

## Table of Contents

|   |           |
|---|-----------|
| <b>The growth of affordable data protection solutions .....</b>                   | <b>4</b>  |
| Many factors driving demand for business continuity .....                         | 4         |
| <b>Choosing the right data protection solution .....</b>                          | <b>5</b>  |
| Evaluating business requirements .....  | 5         |
| Matching protection level to business value .....                                 | 5         |
| Compromise no longer necessary .....  | 8         |
| <b>Midrange storage platforms for cost-effective tiered data protection .....</b> | <b>9</b>  |
| <b>Taking advantage of affordable data protection technologies now .....</b>      | <b>11</b> |

---

## **The growth of affordable data protection solutions**

Until recently, only for the largest organizations has it been financially feasible to deploy comprehensive business continuity solutions to protect critical information systems throughout the enterprise. But the requirement to mitigate the risks resulting from a disruption to technology infrastructure has steadily grown from minimal to urgent over the past few years for most companies. Fortunately, recent technology developments have made a growing number of affordable solutions available to companies as they have begun to address these requirements.

### ***Many factors driving demand for business continuity***

Even as organizations continue to focus on the potential for high-profile disaster events like 9/11 and Hurricane Katrina, many other forces have contributed to an increase in the demand for enterprise-wide data protection and business continuity solutions.

#### **Availability expectations**

Compressed business cycles brought about by online processes and customer expectations for 365x24 business availability have led to an intolerance for any data unavailability and have driven this demand. As a result, companies cannot tolerate data being unavailable at any time, for any reason. While traditional disaster recovery addresses disaster events that create a site-wide interruption and require restoration of business functions in another location, organizations increasingly need to avoid downtime resulting from planned operational procedures such as backups and upgrades. Many of the recent developments in data protection and business continuity have addressed these requirements.

#### **Regulatory requirements**

For other organizations, it is the need to meet growing regulatory demands that dictates greater data protection capabilities. Over the past few years, as companies have prepared to comply with government and industry regulations, they have taken a closer look at how their business processes and infrastructure ensure business continuity under all circumstances.

#### **Centralized storage networks**

The move to centralized storage networks has also created a demand for affordable enterprise-wide data protection. While companies have centralized storage in order to achieve greater efficiencies in storage resource management, they are realizing that the risks of disruption of the centralized storage network are much greater with data concentrated on the network rather than spread out among many servers with direct-attached storage. So the move to centralized storage networks requires that business continuity considerations be addressed as part of the implementation.

These factors have forced organizations to reassess data protection levels and to explore more flexible business continuity solutions, even as pressure on IT budgets continues to grow. According to a recent survey of medium-sized organizations, nearly half are protecting data across buildings or across sites using replication technologies, while a majority of the remaining companies plan on doing so in the near future.

---

# Choosing the right data protection solution

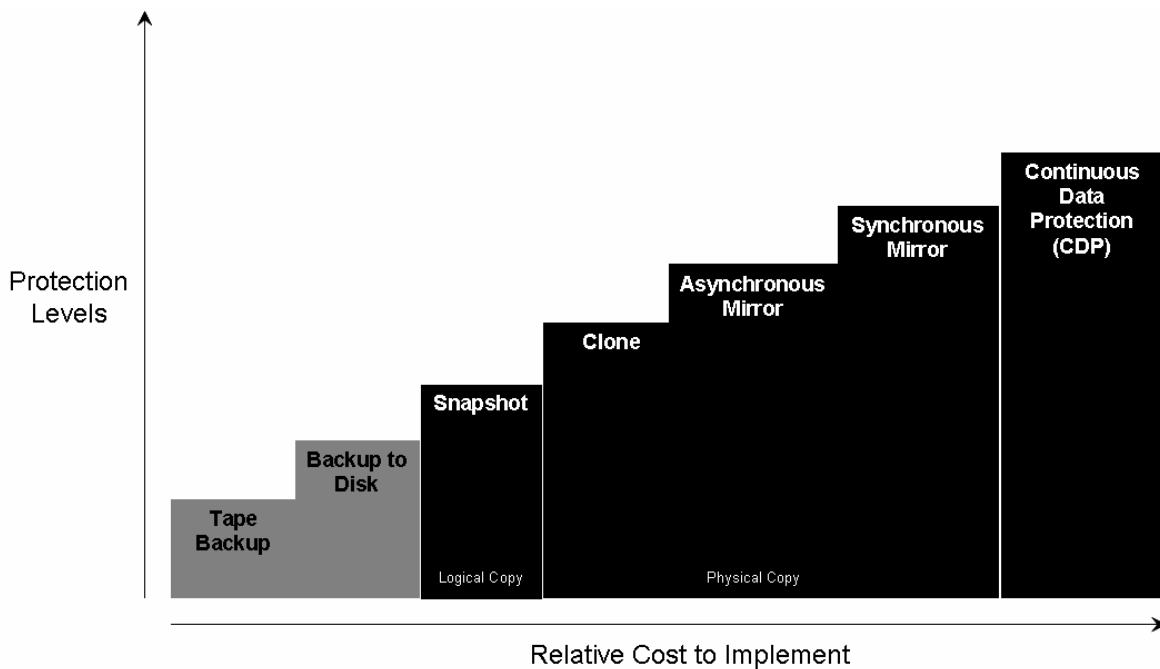
## *Evaluating business requirements*

IT organizations that have experienced the most success with data protection and business continuity initiatives are those that have worked with end users to establish the true business requirements first before making any technology decisions. Business users should be asked two questions regarding specific data and systems—“How much data can we afford to lose if our applications went down?” and “What is the longest outage we could tolerate?” The answers to these questions are referred to respectively as the Recovery Point Objective (RPO) and the Recovery Time Objective (RTO).

Before jumping into a technology solution, the business user should carefully consider two factors in order to objectively determine the criticality of specific data and applications—the monetary value a specific application and its associated data holds for the business and the costs that would be incurred should this system be unavailable. Once this business value is established, the IT organization and the business unit can make an objective assessment of value versus costs in determining the appropriate level of protection. With the emotional element removed, the technology decisions that are made using this approach tend to support the business requirements more cost effectively.

## *Matching protection level to business value*

With an understanding of value and the maximum acceptable data loss and outage duration, the RTO and RPO requirements can be mapped to one of a range of data protection solutions by considering the costs for each and making sure the costs are consistent with the value. If the IT organization has already established a catalog of data protection service levels, users can quickly determine if the cost for the RPO and RTO they believe they need is justified by the value of the data being protected.



**Figure 1: Tiered protection**

---

At a very basic level data is protected by making copies of the production data. The method used to make these copies determines the level at which the data is protected as well as the costs required to reach that level. Figure 1 illustrates the range of methods available to deliver increasing levels of data protection. How quickly a copy can be made determines how often the copies occur, which dictates the maximum potential data loss. More specifically, all transaction records that are captured after the last copy was made are potentially lost if a disruption occurs.

## Tape backup

Tape backup at the left of the figure provides the lowest level of protection and results in the greatest potential data loss and longest time to restore in the event of a disruption. Many medium-sized enterprises use tape backup for both disaster and operational recovery, and most large enterprises use tape for a portion of their applications. According to a recent report from Ovum<sup>1</sup>, tape solutions continue to meet over 50 percent of enterprise disaster recovery requirements.

As companies put more processes online, the volume of data captured and stored continues to grow rapidly, producing greater challenges for every organization executing these backup procedures. Increased data volumes also mean that backup procedures take increasingly longer to execute. Traditionally, this meant taking production systems offline sometime in the middle of the night. With the continued proliferation of online processes and 24x365 operations, this is no longer acceptable. The window to shut down production systems and backup data has disappeared.

## Backup-to-disk

Today, many companies use high capacity, low cost disks, such as LCFC drives, in storage arrays to do backup-to-disk (B2D). B2D is a faster and more reliable backup process so it alleviates some of these challenges and provides protection levels that are higher than tape backup at a slightly higher price point. Virtual tape technologies are another option for customers who would like to leverage the benefits of disk without modifying their existing backup procedures.

## Snapshot and cloning technologies

Another way customers are using disk as a key element of their operational recovery strategy is by deploying point-in-time replicas. These replicas can be snapshots or clones. A snapshot is not an actual copy of data but a pointer to the original data as the data was at the time the snapshot was taken. A clone is a separate physical copy of production data at a specific point in time. Snapshot technology allows a logical copy of data to be made with much greater frequency since a snapshot is instantly accessible, unlike a full copy which needs to be fully synchronized before it can be accessed. A snapshot also consumes less disk space than a clone. One thing that is important to consider is that with a clone, if there is a failure of the original dataset, you can restore the data set from that volume without significant data loss. Snapshot technology does not enable that level of protection.

Because of this, snapshots are usually implemented in combination with some form of actual physical copy such as tape or with clones. Once again, the decision regarding what combination is best should be based on the service level required for the environment. The primary benefit that snapshots or clones provide over B2D or virtual tape is the ability to make copies more frequently which decreases both the RPO and RTO.

---

<sup>1</sup> “Continuous Data Protection: Addressing timely data protection and much more”, Carl Greiner, October 2005

---

## Mirroring

A third replica category is a mirror. A mirror is another type of physical disk-based copy that continually tracks or mirrors the changes made to a production dataset and is usually stored on a second storage system. A mirror can be either asynchronous or synchronous. An asynchronous mirror accumulates all changes made to the production dataset and then applies these changes at a specified interval. Any transactions made between the time the accumulated changes are applied may be lost if the production dataset is disrupted.

A synchronous mirror continually applies each change made to the production dataset to the copy as well before committing the transaction back to the server. Since the copy is always in synch with the original, no transactions are lost in the event of a disruption to the original, bringing the RPO to zero. Since the server has to wait for both the production and mirrored datasets to change, there is a slight performance penalty that occurs that increases as the physical distance separating the production and mirrored datasets increases.

## Continuous data protection

While replicas can provide RPOs up to zero, they do not protect against data corruption errors because the replication process propagates the error to the replicated dataset. Continuous data protection (CDP) is a relatively new technology designed to continuously capture or track data modifications and store changes independently of the primary data, enabling recovery points from any point in the past. In effect, CDP creates an electronic journal of complete storage snapshots, one storage snapshot for every instant in time that data modification occurs. A major advantage of CDP is the fact that it preserves a record of every transaction that takes place in the enterprise. CDP can support RPOs of zero as well but also protect against data corruption errors because data can be rolled back to the exact instant before the error occurred.

The replication or CDP process can be executed using the processing power of the application server, an appliance, or the storage processor inside the storage array itself. Server-based replication may be a lower cost solution but it reduces the number of valuable server cycles that can be dedicated to production workloads. Executing replication procedures using the processing power of the storage arrays can offload servers from this task and can also provide a consistent way to replicate data across different server types. The decision that is going to make the most sense will vary based on the business and technical requirements of the organization.

## Extending protection with replication management tools

Availability can be enhanced further through the use of common replication management tools designed to simplify the automation and management of disk-based replicas. This is particularly valuable for environments with multiple storage arrays, hosts, and applications. Companies can execute replication jobs via a single interface, allowing them to add storage groups, interact with databases, applications, hosts, storage arrays, and virtual operating systems, and to offload backup cycles from production systems.

With the growth of web-based applications and web-services, more environments involve multiple databases and interrelated applications. For instance, sales, manufacturing, e-commerce, and customer service records all may share common databases and therefore need to be consistent with each other at the transaction level. Therefore it is important that a copy being made of one database must represent the same point in time as the copy being made of another interrelated database before any attempt is made to restart a portfolio of applications that span these databases. If the two copies are not consistent and a restart is attempted, significant data integrity issues may result. Advanced consistency software technologies ensure that replicas made across inter-related applications are consistent at the transaction level. This protects the integrity of all data in the event of a disruption.

---

### **Protecting critical records across a CLARiiON SAN at a local government organization**

The Information Technology department of a local government organization in Kansas had worked to serve the online information needs of its 170,000 residents, but was finding it difficult to manage the growing pool of direct-attached storage. In order to better manage and protect its data, the organization replaced its direct-attached storage for the 45 servers scattered across its facilities with an EMC CLARiiON CX-based SAN.

The CLARiiON SAN consolidates the most critical information and helps the organization address its storage provisioning challenges. EMC SnapView provides cloning capabilities and EMC SAN Copy software provides replication to another CLARiiON three miles away at the courthouse facility for the SQL server databases that hold registered deeds and GIS system records. These remote copies enable the local government organization to test upgrades, fixes, and configuration changes without disrupting the production system. The result has been significantly improved reliability and accuracy of these processes.

“EMC CLARiiON has provided the bullet-proof platform we needed,” said the organization’s network administrator. “We not only get a more manageable, straightforward storage environment that eliminates the provisioning headaches, but also get data that is more protected than ever. And we can recover from server failures much more quickly as well. If a server fails we can restart by simply booting a second server from the SAN. We’re just beginning to tap into the operational capabilities afforded by EMC replication technologies, like non-disruptive backup to disk.”

### ***Compromise no longer necessary***

No organization can cost-effectively protect all of its data assets with just one of these technologies. Less critical data may require a simple tape backup for operational recovery while the most critical data may require synchronous replication to a secondary system for disaster recovery coupled with the use of clones for operational recovery. Companies that use a single technology to meet the protection needs of multiple data types will likely see excessive exposure to data loss or excessive costs. The most effective approach combines these technologies—from backup, snapshots, and clones to mirrors and CDP systems—into a tiered protection infrastructure that delivers the most appropriate levels of protection to data based on its value to the organization. It is no longer necessary for organizations to compromise—either accepting greater risks than they need to, or implementing a solution that overprotects data and costs more to implement than is justified by the value of the data.

Taken together, these technologies allow an organization to not only protect its critical data assets but also use the copies of production data to support parallel processing activities such as backup, application development and testing, and data warehouse refreshes while increasing the availability of production systems. Decision support activities are more effective because the data warehouse can be refreshed with current data more frequently without bringing user query activity to a halt. Backups can be taken more frequently which reduces the potential for data loss. System upgrades, maintenance fixes, and configuration changes can be thoroughly tested offline using the most current copies of production data, which reduces the chances of failure or data integrity issues when the changes go live.

---

## Midrange storage platforms for cost-effective tiered data protection

Midrange storage platforms like EMC® CLARiiON® CX3 are helping large and medium-sized enterprises implement cost-effective tiered data protection that only the largest organizations could afford previously. CLARiiON offers a full range of end-to-end business continuity solutions that, when coupled with EMC consulting services, are designed to meet every business continuity need.

EMC CLARiiON provides the ability to deploy a full range of tiered storage options within a single platform. High performance disk drive technologies support mission-critical production volumes. Less critical environments like development, test, or reporting can utilize high capacity, low cost drives, and tape-emulating disk libraries can be used for backup and recovery operations.

In addition to working transparently with host-based replication software like EMC RepliStor®, CLARiiON also works with a full range of array-based replication software that delivers snapshots, clones, or local or remote mirrors without consuming valuable server cycles or LAN bandwidth.

### Protecting critical medical records at Willis-Knighton

Willis-Knighton is one of the fastest growing health systems in the United States with an aggressive plan to eliminate paper throughout its network of four hospitals in the Shreveport, Louisiana area. As paper disappears, however, digital records grow and must be properly protected—not only to ensure effective patient care in the event of data loss, but also to comply with strict federal regulations, such as HIPAA, that govern information security.

Among the most critical data at Willis-Knighton Health System are Picture Archiving and Communication System (PACS) images, each of which can be several gigabytes in size. Willis-Knighton stores these images on a high-performance EMC CLARiiON Fibre Channel storage area network (SAN) in a primary data center at the Willis-Knighton Medical Center. To further protect these vital medical records, Willis-Knighton uses EMC MirrorView Synchronous (MirrorView/S) software to continuously replicate the PACS images to a second CLARiiON SAN at a disaster recovery site seven miles away.

Jonathan Lee, Network Coordinator at Willis-Knighton Health System, said, “Loss of our PACS images could have a very serious impact on patient care because our physicians reference them frequently in the course of treating a patient. CLARiiON and MirrorView provide us with rapid access to actively used images, and the peace of mind of knowing that we could recover these images quickly in the event of a disaster. Using MirrorView is important because we need to verify that the data is intact when it reaches the disaster recovery site.

“Previously, our PACS images were stored on optical disk, which would have required hours to restore. If we ran into mechanical problems, it could take longer – possibly even days. Today, we can fail over to our disaster recovery site in a matter of minutes. So our physicians can continue to access the PACS images with virtually no disruption to patient care.”

EMC SnapView™ runs within the CLARiiON unit and delivers either snapshots or full-volume clones that can be used as the source data for fast, frequent, and non-disruptive backups to improve the restore process and to use for development, testing, or data warehouse procedures.

EMC SAN Copy™ also runs within the CLARiiON array and is designed to replicate data between CLARiiON and other storage platforms including EMC Symmetrix®, Hitachi, HP, IBM and Sun storage arrays. The use of SAN Copy and SnapView can be integrated to meet operational and disaster recovery requirements.

---

EMC MirrorView™ is advanced, array-based replication software that provides synchronous or asynchronous mirrors between two or more CLARiiON systems. By integrating MirrorView with SnapView snapshots, copies of production data can be used for parallel operational procedures from a secondary location.

#### **Keeping Tradecard's customers running 24x7**

TradeCard is a leader in supply chain automation services that manage global trade for thousands of organizations worldwide like Columbia Sportswear, Rite Aid, and Wolverine Worldwide. Because TradeCard's customers use its web-based platform to help buyers, sellers, and their trading partners manage trade transactions around the clock, the availability requirements for the system are very high. After several years of business growth, TradeCard began to outgrow its data protection strategy. "Previously, using database log replication, our quarterly DR drills would take a couple of hours to perform, require downtime to the production system, and upon completion, would require additional time to restore the DR system," explained Anthony Ercolino, Tradecard vice president of data center operations. "As a provider of 24x7 supply chain solutions, we sought to improve this."

In addition, Ercolino knew that the active logs could not be replicated until all work was completed and any data that had changed since the last replication was at risk until a new replication could be sent. As a result, TradeCard set out to implement a cost-effective data protection mechanism that would utilize its existing EMC CLARiiON SAN to provide non-disruptive protection of its customers' mission critical supply chain application data, allow the company to test its DR plans without affecting production system availability, and reduce the amount of data that could be lost in a disaster.

TradeCard was able to meet these goals when it implemented EMC RecoverPoint to keep the production images of its customers' supply-chain DB2 databases synchronized with a remote DR site 100 miles away using a VPN link through the public Internet. With RecoverPoint, TradeCard has seen its RPOs shrink to less than a minute and can now conduct DR tests without affecting production system availability or performance.

"RecoverPoint has more than met our expectations," said Ercolino. "It provides the availability and RPO levels we need, and allows us to test as often as needed, but at a level of cost that is appropriate for an organization our size. It seamlessly integrated with our EMC CLARiiON SAN and helped us save 50 percent on bandwidth costs. Our production application is hosted and we pay for the bandwidth we use, so the RecoverPoint data compression capabilities were critical. And the RecoverPoint interface makes it flexible and easy to manage. We didn't add any staff to support it and have been able to automate many of the procedures, which has reduced errors and improved productivity."

EMC RecoverPoint is an intelligent data protection and recovery solution that runs on out-of-band appliances located on the network. RecoverPoint provides comprehensive data recovery for SAN-attached devices across the entire data center including heterogeneous, continuous remote replication (CRR), and local, continuous data protection.

SnapView, MirrorView, and RecoverPoint provide application consistency technology to ensure that copies made in environments with multiple databases and interrelated applications represent a consistent point in time, which ensures the integrity of any process that uses those copies.

EMC Replication Manager provides automation and management of disk based replicas for operational and disaster recovery. Replication Manager reduces the need for manual scripting and offers simple management of disk-based replicas so that IT personnel can be more productive. Replication Manager's ability to manage multiple replication technologies with deep application integration offers additional value to Microsoft Exchange, SQL, and Oracle environments.

---

## Taking advantage of affordable data protection technologies now

Mid-range platforms like EMC CLARiiON are helping organizations implement cost-effective, tiered data protection solutions that only large organizations could afford previously. Many IT executives are still under the impression that an investment in data protection is essentially an insurance policy that sits idly until a disaster strikes. Initial deployment costs along with the lack of business benefits to justify the costs are cited most often in industry surveys as the reasons for not implementing business continuity solutions to protect data.

### Protecting critical customer data at a large community bank

When a large, northwest, privately-held, local community bank wanted to ensure the protection of 500 GB of critical customer account data without exorbitant bandwidth costs and availability interruptions, it replaced its tape-based backup and database scripts with replication solutions from EMC. The bank uses EMC RecoverPoint to replicate its database snapshots non-disruptively from its headquarters site in Vancouver to CLARiiON arrays at a remote site in Portland using just a pair of T1s.

The EMC solution has easily met the bank's 15 minute RPO objective—tests conducted mid-day during peak processing show only a three minute lag between the production and replicated datasets. The system has also proven to be very straightforward to maintain and has enabled the bank to eliminate complicated, third-party database scripts, which is important for a mid-sized organization with limited staff and financial resources. The bank is now looking to extend the RecoverPoint solution to e-mail information and loan document images with an eye towards growing its technical availability plan to more of a complete business continuity plan.

"When we saw what RecoverPoint could do we stepped back and said, 'Wow,'" said the bank's vice president of technical operations. "We're minimizing the bandwidth and replicating our mission-critical data without affecting availability. The ease of identifying the snapshot to load into the recovery system and then executing it is impressive. My tech spent only a week with it and felt very comfortable. Most banks use tape and a hot site to recover critical operations. RecoverPoint allows us to get our key branches up in an hour—and do it ourselves."

In reality, the benefits go well beyond an insurance policy. Technology that was originally intended solely for disaster recovery purposes has evolved to the point where it is now being used to make operational procedures—such as development, testing, and data warehouse refreshes—much more efficient. As a result, both IT staff and even end users utilizing decision support systems are much more productive. These ongoing productivity benefits make it much easier for medium-sized enterprises to justify the investments required.

EMC has helped thousands of companies successfully implement comprehensive data protection and business continuity solutions on its CLARiiON storage platforms. Its experienced team has extensive knowledge of replication technologies and storage deployment best practices along with the proven methodologies to help medium-sized organizations quickly assess both the risks and consequences of data loss, develop optimal data protection implementation plans, and maintain the availability of critical applications.

For more information, please contact your EMC Account Manager or visit <http://www.emc.com/continuity>.