

Optimizing EMC RepliStor SMB Edition in an EMC Retrospect Environment

1/16/2006

Abstract:

EMC[®] Retrospect[®] software provides comprehensive, point-in-time backup and recovery for Windows, Linux, Macintosh, and Solaris computers at a company's central or branch office. EMC RepliStor SMB Edition provides data replication to protect a critical Windows server. It also replicates data from a computer at remote branch offices to a central location for backup. Utilizing RepliStor SMB Edition in tandem with Retrospect, a small or medium business (SMB) or distributed enterprise can guard against data loss from accidental deletion, virus, or hardware failure as well as free applications from the impact of backup operations.

Copyright © 2006 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.

EMC², EMC, RepliStor, and Retrospect are trademarks of EMC Corporation. All other trademarks used herein are the property of their respective owners. All other brand names are trademarks or registered trademarks of their respective owners.

[S70120106V1]

Table of Contents

- Overview4**
 - EMC Retrospect 4
 - EMC RepliStor SMB Edition 4
- Protecting Critical Data—Backup or Replicate?4**
- Supported Implementations of RepliStor SMB Edition and Retrospect.....5**
 - Replicating Data to Protect a Critical Server 5
 - Protecting a Critical Server by Replicating Data over a LAN..... 5
 - Protecting a Critical Server by Replicating Data over a MAN or WAN 7
 - Replicating Branch Office Data to a Central Location for Backup 7
- Sizing Data Replication Needs8**
 - Replicating Initial Amount of Data 8
 - Rate of Data Change..... 9
- Unsupported Scenario9**
- Summary9**

Overview

Retrospect and RepliStor SMB Edition can be used together to eliminate data loss on a critical Windows server or to protect data on local and remote computers automatically by replicating to a central office for backup. These capabilities enable a small and midsize business (SMB) or branch office environment to prevent the loss of critical information in the event of hardware failure and to reliably protect data at branch offices that do not have the resources needed to perform onsite backups. .

EMC Retrospect

Retrospect backup and recovery software provides automated, reliable data protection for servers, desktops, notebooks, and business-critical applications at an SMB or branch office. Retrospect has received widespread industry acclaim for its ease of use, fast backups and accurate restores, simplified disk, tape, and optical media management, and best-in-class disaster recovery.

The Retrospect application is installed on a Windows backup server. Retrospect client software is installed on each Windows, Linux, Macintosh, and Solaris computer being protected. Retrospect pulls backups from the client computers to the backup server. Backup data is then written to direct or network-attached disk or tape devices via SCSI, iSCSI, or Fibre Channel. Backups can also be transferred automatically to additional disk or tape resources for archiving or offsite storage for disaster recovery.

EMC RepliStor SMB Edition

EMC RepliStor SMB Edition delivers easy to manage, host-based Windows data replication. RepliStor SMB Edition replicates individual files, directories and the registry asynchronously from one Windows computer to another over a LAN or WAN. Businesses can replicate critical Windows data automatically for greater protection and copy data from a remote computer for centralized backup.

Application software is installed on each computer involved in the replication process. A computer containing data that is being replicated to another site is configured as a *source* computer. A computer receiving data is configured as a *target* computer. Replication can be set to occur continuously or according to a schedule.

Protecting Critical Data—Backup or Replicate?

Retrospect's and RepliStor SMB Edition's respective capabilities complement each other to achieve appropriate levels of data protection. **Table 1** illustrates how you can use Retrospect to protect your entire computer environment while adding RepliStor SMB Edition to increase protection for critical Windows systems.

Table 1 Retrospect and RepliStor SMB Edition perform complementary functions

Feature	Retrospect	RepliStor SMB Edition
Restore points	Multiple points captured hours, days, weeks, or months ago	Single point captured seconds or minutes ago
Recovery time	Minutes – Hours - Day	Seconds - Minutes
Protect against hardware failure	●	●
Protect against data corruption	●	
Bare Metal Restore	●	

Offsite data transfer	Tape, WAN	WAN
-----------------------	-----------	-----

Restore points – This is the point in time data can be restored to. RepliStor SMB Edition restores data to the time of last replication and limits potential data loss to within seconds of when the loss occurred. Retrospect can restore data to multiple points in time. Data loss is limited to the time between when the loss occurred and the last backup.

Recovery time – This is the time it will take to restore data after a loss. RepliStor SMB Edition can recover data within seconds to minutes. Using Retrospect, data recovery from disk can take minutes to hours, depending upon data volume. Restoring data from tape may take a few hours to a day, depending upon the volume of data and if tapes are onsite or offsite.

Protect against hardware failure – Protect against data loss due to computer hardware failure as well as from fire, flood, or other disaster.

Protect against data corruption – Protect against data loss due to viruses, user deletions, or corrupted data. Because these types of events are replicated, Retrospect should be used to provide secondary protection (through multiple restore points) for servers being replicated with RepliStor SMB Edition.

Bare metal restore – During the course of backups to tape or disk, Retrospect proactively gathers all information to completely restore a Windows server, desktop or notebook, even if it has become unresponsive. Selected restore points include Windows OS, applications, patches, settings and preferences.

Offsite data transfer – Copying to tape is the most common method for transferring Retrospect backups to offsite locations for disaster recovery (DR) and archiving, especially for large data volumes. Telecommunications lines can be used to transport data using RepliStor SMB Edition if data volumes are moderate. For instance a dedicated T1 line, costing approximately \$10,000 per year, can typically move 1 GB of data each hour. A dedicated T3 line is 30 times faster, but cost approximately \$100,000 per year. If data volumes are large, Retrospect tape backups offer the best method for recovery.

Supported Implementations of RepliStor SMB Edition and Retrospect

By combining RepliStor SMB Edition’s data replication capabilities and Retrospect’s comprehensive backup and recovery capabilities, SMBs and branch offices can achieve optimum levels of data protection, as well as automate the geographical consolidation or distribution of Windows data, as outlined in these four scenarios:

- RepliStor SMB Edition replicates data over the LAN, MAN or WAN for Windows computers that are part of a Retrospect client environment. This scenario provides increased protection for business critical Windows systems.
- RepliStor SMB Edition replicates Windows data on a branch office computer over a WAN to a central/main office computer that is backed up by Retrospect. This scenario provides hands-free data protection for one computer at a branch office with no onsite IT/admin personal.

Replicating Data to Protect a Critical Server

In these two scenarios, Retrospect performs traditional backups of up a file server, application server, desktop computer, or notebook computer on the network. RepliStor SMB Edition has been added to replicate data from one business-critical file or application server (source) to a second server (target) to protect against data loss if hardware failure or another disaster struck the source server.

Protecting a Critical Server by Replicating Data over a LAN

In this scenario, data is replicated over a LAN. Because large amounts of data can be moved quickly over a LAN connection, RepliStor SMB Edition can easily replicate selected data (a folder or group of important files) or the entire contents of the computer.

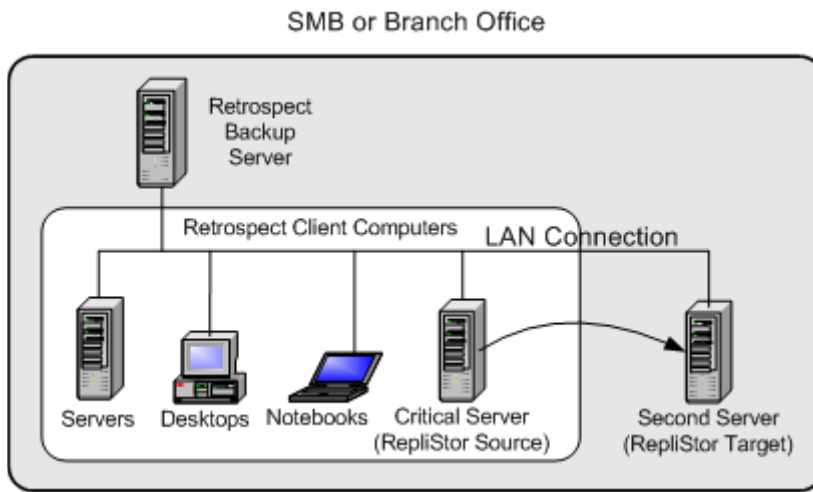


Figure 1 Replicating data over a LAN

This scenario eliminates data loss in the event a critical server fails. The target computer is configured as a file server or with identical applications as the source server. If data loss is a result of hardware failure, fire, flood, or other event causing physical damage to storage, servers, or connections, a replicated copy of data is available. If the source computer crashes, potential data loss is reduced to seconds and minutes. If data loss is caused by user error, inadvertent deletion, virus, or another event that results in data loss on the replicated volume, Retrospect is utilized to restore data to a point in time. Potential data loss depends on the amount of time between the last backup and the data loss. Restore times are contingent on the speed with which the backup media can be accessed.

During normal operations, the target computer is in a passive state and only receives replicated data. If the source computer fails, the target computer is placed in an active state and traffic is rerouted to the target computer. While recovery operations are being performed at the source computer, access to files, folders, and applications continues without any loss of data or productivity. When the source computer is ready to be brought back online, RepliStor SMB Edition can be used to sync data between the source and target computers. User connections can be restored to the source computer and normal replication operations can resume.

If all the contents of the RepliStor SMB Edition source server are replicated to the target server, there is no need for Retrospect to back up both computers (in case of user error, data corruption, virus, etc), because they contain identical data. Which server to back up depends on whether the source computer is a database or non-database server.

Replicating File and Non-Database Application Servers

If the source computer is a file server or non-database application server, either the source or target servers can be set up as Retrospect clients. For daily backups, Retrospect backs up the target server only. This frees applications and users from any performance impact during backup operations. In the event data needs to be recovered from a Retrospect backup, the data is written back to the *source* computer. This ensures data integrity on the source and target computers when replication resumes.

Replicating Microsoft Exchange, SQL Server, and Other Database Servers

If the source computer is running a database application, set up just the source computer as a Retrospect client, because the application is actively running only on the source computer. For most database applications, Retrospect's Open File Backup add-on protects all Retrospect clients and provides a true point-in-time backup while applications remain online.

However, backing up Exchange Server and SQL Server is more complex. Open file backups of Exchange Server or SQL Server while the application is running could result in data loss, because the application might not have time to

flush its tables to disk. Exchange Server, SQL Server, and other well-architected database applications deal with this kind of situation by recovering to a *crash-consistent state*. The restored application searches its logs and files until it finds a previous point in time when the log files are consistent with the data. Then it returns the application to a state that is consistent with that point in time. As a result, some of the more recent data in the application might not be backed up.

To ensure that all data is recoverable, use Retrospect's Exchange Server Agent or SQL Server Agent add-on to These add-ons utilize application-specific APIs from Microsoft to provide automated, online protection and restores without data loss.

Protecting a Critical Server by Replicating Data over a MAN or WAN

In this variation of the previous scenario, RepliStor SMB Edition has been added to replicate data from one business-critical server within a Retrospect client environment over a MAN or WAN to a second server (target) located at a remote site.

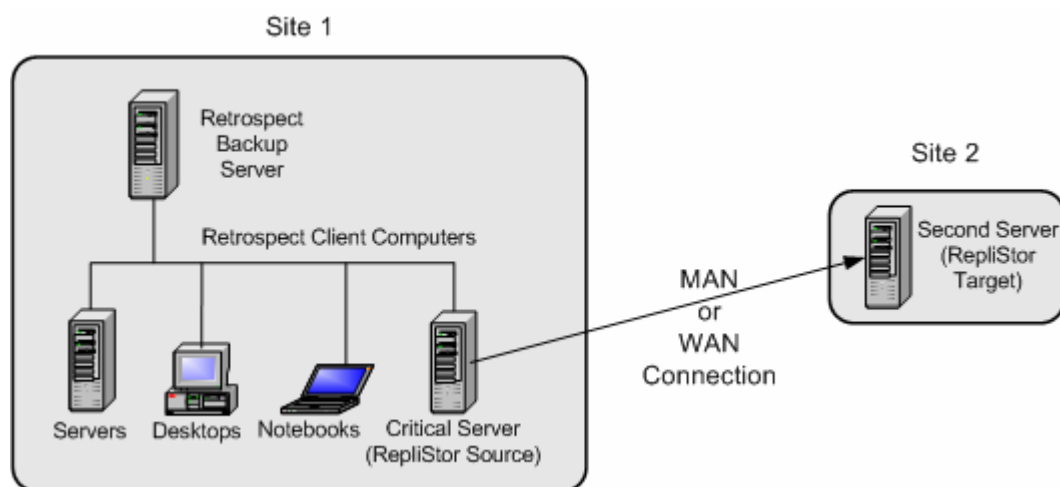


Figure 2 Replicating data over a MAN or WAN

The guidelines noted above for replicating and protecting critical servers over a LAN also apply when replicating data over a WAN or MAN. Unless a VPN connection is established with the target server, however, Retrospect backups take place on the *source* server only.

You should also keep in mind the limited throughput of telecommunication lines. A dedicated T1 line, which costs approximately \$10,000 per year, can typically move a maximum of 1 GB per hour. A dedicated T3 line, which costs approximately \$100,000 per year, can typically move 30 times that amount of data per hour. RepliStor SMB Edition can use all or part of the available telecommunications bandwidth at the option of the administrator. Data is also encrypted over the network for security.

This variation of scenario one is useful for a small business owner who wants to replicate information from one site to another, such as from a commercial office to a home computer, where data can be accessed during off hours or stored for an additional level of back up security. And because RepliStor SMB Edition can replicate data across town, across country, or across the globe subject to network transfer speed, it is quite useful for businesses and branch offices that want to eliminate the risk of a large scale disaster affecting source and target computers.

Replicating Branch Office Data to a Central Location for Backup

This scenario is ideal for small businesses that have multiple offices or stores in a geographical area, one or two computers at each location, and no IT resources to support them on a day-to-day basis.




Figure 3 Replicating data over a MAN or WAN for backup

RepliStor SMB Edition replicates data from each of the remote (source) computers over a WAN to a single target computer at the main office. The main office computer is then backed up by Retrospect. If a remote computer incurred data loss resulting from human error or corruption and a VPN connection was not set up, Retrospect restores would be performed on the target computer at the main office. An incremental sync between the source and target mirrors would then be performed with RepliStor SMB Edition at the main office to ensure data integrity before resuming replication operations.

By combining RepliStor and Retrospect, businesses can regularly protect remote computers without any onsite assistance.

Sizing Data Replication Needs

To size replication needs for a given scenario two questions should be considered first:

- How much data will be replicated initially?
- What is the rate at which data will change?

When copying data over a MAN or WAN, carefully evaluate the speed of the data link and the amount of data being replicated, because it directly affects the amount of data that can be replicated within a given amount of time.

RepliStor SMB Edition transmission speeds vary considerably over a MAN and should be evaluated on a case-by-case basis to determine the amount of data that can be moved over a given period. To determine the amount of data RepliStor SMB Edition can move over a WAN, use these figures as a guideline for planning:

- T1 link @ 50% utilization — 500 MB/hour
- T1 link @ 100% utilization — 1 GB/hour
- T3 link @ 50% utilization — 15 GB/hour

To evaluate RepliStor SMB Edition and measure its transmission speeds in a specific environment, download a free trial version of RepliStor SMB Edition at http://www.emcinsignia.com/replistor_trial.

Replicating Initial Amount of Data

If the amount of data being replicated is small, the initial transfer of data can move over a MAN or WAN without causing undue delays. It will take time to synchronize data between the source and target computers, but after the initial data has been copied to the target disk, only new or changed data needs to be replicated across the data link.

If the amount of data is large, the source and target systems must be pre-seeded by co-locating them or by physically transporting the data (on disk or tape) to the target site for the initial synchronization. Thereafter, only new or changed data needs to be replicated across the data link.

Rate of Data Change

If data changes at a rate that is dramatically faster than the throughput of the data link, information cannot be replicated fast enough to become available in a reasonable amount of time. For replication to be practical, the rate of data change must be within the range of the data link that connects the source and target sites.

For example, a branch office connected to a central office via a dedicated T1 link has a total of 1 GB of data changes each day. The T1 link replicates data at 1 GB per hour. Replication should proceed quickly and the replication should never experience backlogs.

However, if a branch office has a total of 25 GB of data changes per day moving over a T1 link that replicates data at 1 GB per hour (24 GB a day), the replication process will experience continuous backlogs, even if replication takes place 24 hours a day. Replication will fall behind by 1 GB after the first day, 2 GB after the second day, and so on. The data at the two offices will never synchronize.

Unsupported Scenario

Although it may be tempting to use Retrospect to back up computers to a Disk Backup Set and then use RepliStor SMB Edition to replicate the Disk Backup Set to a remote site, this scenario is currently not supported, because Retrospect allows only one Retrospect backup server to utilize a Disk Backup Set and its associated catalog file. To copy disk backups to additional disk resources, utilize Retrospect's Backup Set Transfer and synthetic full capabilities.

Summary

The demand for data is expanding at an astronomical rate in almost every business environment. Companies must look beyond traditional backup methods to protect critical data and applications, and to ensure that computers at remote branch sites are reliably protected. These tasks are especially challenging for SMBs with limited IT resources. By combining RepliStor SMB Edition's data replication capabilities and Retrospect's comprehensive backup and recovery technology, SMBs and the distributed enterprise can achieve optimum levels of data protection easily, reliably, and cost effectively.

About EMC

EMC Corporation (NYSE: EMC) is the world leader in products, services and solutions for information storage and management that help organizations extract the maximum value from their information, at the lowest total cost, across every point in the information lifecycle. Information about EMC's products and services can be found at www.EMC.com.

About EMC Retrospect and EMC RepliStor SMB Edition

EMC Retrospect and EMC RepliStor SMB Edition are part of the EMC Insignia line of software and hardware products, which enables small and medium businesses (SMBs) to store, protect, manage, and share vital business information. To learn more about EMC Insignia, contact your authorized EMC Velocity SMB channel partner or visit www.emcinsignia.com.