



Branch Office Data Consolidation and Information Protection

Implementation Notes and Best Practices

*Using LEGATO RepliStor and NetWorker with
EMC NetWin Storage*

Abstract

Branch Office Data Consolidation and Information Protection is a solution for consolidating servers as well as replicating data from branch offices to a central regional data center. This document discusses the architecture, components, installation, and configuration for the solution, and includes guidance for replication and recovery.

Published 6/17/2004

Copyright © 2004 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, Symmetrix, Celerra, CLARiiON, CLARAlert, Documentum, HighRoad, LEGATO, Navisphere, PowerPath, ResourcePak, SnapView/IP, SRDF, TimeFinder, VisualSAN, and where information lives are registered trademarks and EMC Automated Networked Storage, EMC ControlCenter, EMC Developers Program, EMC OnCourse, EMC Proven, EMC Snap, Access Logix, AutoAdvice, Automated Resource Manager, AutoSwap, AVALONidm, C-Clip, Celerra Replicator, Centera, CentraStar, CLARevent, Connectrix, CopyCross, CopyPoint, DatabaseXtender, Direct Matrix, Direct Matrix Architecture, EDM, E-Lab, Engenuity, FarPoint, FLARE, GeoSpan, InfoMover, MirrorView, NetWin, OnAlert, OpenScale, Powerlink, PowerVolume, RepliCare, SafeLine, SAN Architect, SAN Copy, SAN Manager, SDMS, SnapSure, SnapView, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix DMX, Universal Data Tone, and VisualSRM are trademarks of EMC Corporation. All other trademarks used herein are the property of their respective owners.

Part Number H1252

Table of Contents

Introduction	5
Solution Overview	5
Replication	6
Protection	7
Recovery	7
Automated, Rapid Recovery When Failures Happen	7
Low-Cost Data Management and Increased Recoverability of Systems and Branch Operations	7
Components	8
EMC NetWin 200.....	8
Windows Storage Server 2003.....	8
LEGATO RepliStor	9
LEGATO NetWorker Suite	10
Volume Shadow Copy Service (VSS).....	10
Setup and Installation	12
EMC NetWin 200 and Microsoft Windows Storage Server 2003	12
Network Overview	12
Site Planning Prerequisite	13
EMC NetWin 200 Setup.....	14
EMC NetWin 200 Deployment.....	15
NetWin 200 Administration	17
LEGATO RepliStor	17
Before You Install.....	17
RepliStor Software Use of Windows Active Directory	18
Obtaining a License Key	19
Setting Up Secure Access	19
RepliStor Remote Installation	19
LEGATO NetWorker	19
NetWorker Installation Roadmap.....	19
Configuration	21
Replication.....	21
Creating Specifications	21
Global Exclude Specifications	21
File/Directory Specifications	22
Share Specifications	22
Registry Specifications.....	22
Synchronizing Specifications	23
Updating and Forwarding Options	23
Throttling Options.....	23
Failover	23
Super Automatic Switch Over (ASO) with Aliases.....	23
Switchover Setup	24
Data Backup	26
How LEGATO NetWorker Software Works	26
NetWorker Setup	27
Data Backup Methods	28
Database Backup Considerations	28

Pause Updates and Turning On/Off Forwarding and Open Files.....28
Offline Database Backups29
Online Database Backups29
Data Recovery30
Recovering Data with RepliStor30
Recovering Data Backed Up with NetWorker to the Data Center.....31
Frequently Asked Deployment Questions32
Additional Resources34

Introduction

The Branch Office Data Consolidation and Information Protection Solution is a LEGATO®, EMC®, and Microsoft partner solution. This solution was in part qualified in the Microsoft Partner Solution Center (MPSC), and enjoys the support and serviceability advantages of being offered by Microsoft Gold Partners. This document provides an overview of the architecture, components, installation, and configuration for the solution, along with guidance for operating the replication and recovery.

The purpose of this white paper is to provide overview information for system integrators (SIs) and system administrators to perform hardware and software installation and configuration of a solution designed to consolidate data from branch offices to a central location. Operators who monitor the daily backups and maintain the servers and clients on a network may also find this document useful. This information will include a description of the architecture and components, configuration of those components, and guidance for replication and recovery.

This document uses as an example the architecture built at the MPSC. However, each customer's environment will be different and therefore the application of the information in this document will vary based on the actual environment.

Each component for the Branch Office Data Consolidation and Information Protection solution has its own set of prerequisites which can be found in the documentation provided with that component. Generally, this document assumes that the customer's environment already has a Windows Domain, standard TCP/IP-based networking is in place, and the environment is operational.

Solution Overview

Branch Office Data Consolidation and Information Protection is a solution for consolidating servers as well as replicating data from branch offices to a central regional data center. Multiple file servers at any one location can be consolidated to the NetWin™ 200, running Microsoft® Windows® Storage Server 2003. Once consolidated, information can be efficiently and cost-effectively replicated from multiple sites to a central location to enable easier, lower-cost management and ensured recoverability of the data. LEGATO® software components enable the data replication, backup, and recovery operation.

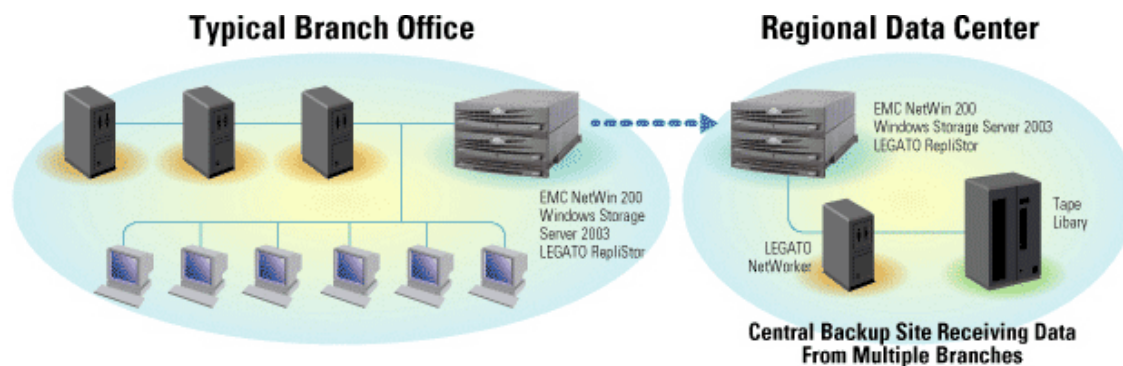


Figure 1. General Configuration and Data Flow for Data Consolidation

The combination of EMC NetWin™ 200 devices, LEGATO RepliStor® at the branch offices and regional data center, and LEGATO NetWorker® at the regional data center provide optimal branch office data protection and consolidation with improved business efficiencies and lower total cost of ownership for both the branch office and data center. When remote office data is replicated and consolidated to a central location on NetWin 200 it can be backed up with trained personnel while leveraging equipment and software already established and proven in the data center. The approach reduces risk because the entire process is managed by trained data center administrators and operations staff. Data protection is improved in the branch office because it is replicated offsite as soon as it is created. An up-to-the-moment second copy of the data is always available in the data center for disaster recovery. The data center can deliver rapid business resumption if a remote location goes down making data immediately available. Also, the approach is one that can greatly reduce the costs of hardware, software, support, and

administration for all remote locations as it eliminates the need to maintain backup and recovery infrastructure within branch offices. Over the long run, this data protection strategy ensures reliable branch office backups, quick recovery of data and business continuity when needed while increasing branch office and data center productivity, all with lower total cost of ownership (TCO). The combined EMC, LEGATO, and Microsoft solution enables the following.

Replication

- Realtime, asynchronous replication of Windows data including files, directories, volumes, shares and selected Windows registry keys.
- Replicates only changed Windows data as changes are written; does not replicate unnecessary data.
- Bidirectional replication enables distributed locations to share realtime business data, such as registry receipts and pricing changes.
- Support of existing TCP/IP network connections with throttling and compression, enables replication to work across networks shared with other types of business traffic and keeps costs low.
- Low CPU utilization minimizes impact on application performance and support for Windows Storage Server eliminates CPU impact entirely.
- Support of application servers with direct-attached storage and SAN storage for replication enables replication for entire branch office (see Figure 2).
- Capability to replicate data among Microsoft® Windows®, Microsoft® Windows® 2000, and Microsoft® Windows® Server 2003 systems enables these systems to co-exist in the same environment and share data.
- Replicates open files to ensure data gets replicated as quickly as possible after the initial Windows system writes to disk, and without needing to wait for the application to close the file.
- Write-order preservation ensures database integrity through the replication process.

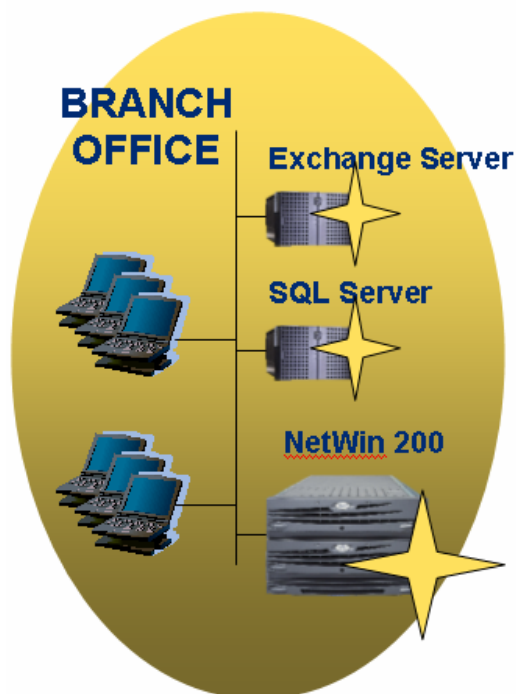


Figure 2. Typical Replication Points from a Branch Office

Protection

- Capability to continue replicating data while also performing backup operations on the centralized data set, to ensure continuous protection.
- Enables point-in-time backup of data to be performed at any time with no performance impact on business operations within the branch.
- High-speed backup and disk backup options for staging data enable easier management, lowest-cost management, and highest-speed recovery.

Recovery

- Centralized backup processes streamline the recovery operation with fewer locations to search for offline data (tapes).
- Centralized data is kept discreet for each branch system to enable easier management and faster recovery of just that data from that branch, rather than requiring recovery of all branches.
- Replicated data is immediately available online when failures occur, for easy replication back to the original systems to recover their operations.
- Rapid recovery operations from disk and from tape, along with disk-staging, ensure fastest-possible recovery from data backed up.
- Several levels of recovery options are provided depending on the level of failure that has occurred, and include:
 - Immediate recovery using online replicated data which can be used in-place or replicated back to the original system or site.
 - Staged recovery from backup media.
 - Recovery from backup media used in combination with re-synchronization of replicated data.

Automated, Rapid Recovery When Failures Happen

- RepliStor monitors Windows file and print services on Windows Storage Server, and can failover those services to a second NAS for rapid recovery when needed.
- RepliStor enables servers and network-attached storage to be configured in clusters for lowest TCO and highest return on investment, supporting both data and application failover.
- Certified on Microsoft Windows 2000 for Server, Advanced Server, and Datacenter Server.
- Certified on Microsoft Windows Server 2003 for Server, Enterprise Server, and Datacenter Server.

Low-Cost Data Management and Increased Recoverability of Systems and Branch Operations

Once the NetWin 200 platform is deployed with the data replication and recovery capabilities of RepliStor, information is protected. In the event of failure, this solution provides:

- Data replication across multiple sites
- Automated recovery of NetWin data from a second site
- Business continuity through the recovery process

For more information on this solution, visit the LEGATO website: <http://www.legato.com/nasalliance>

Components

The following sections describe the hardware and software included in the Branch Office Data Consolidation and Information Protection Solution.

EMC NetWin 200

EMC NetWin 200 is a NAS solution that combines standard Intel server technology running the Microsoft Windows Storage Server 2003 operating system and an EMC CLARiiON® storage array. By leveraging Microsoft's advanced Virtual Disk Service (VDS) and Volume Shadow Copy Service (VSS), EMC provides integrated storage management of the CLARiiON storage array from the Windows Storage Server 2003 interface.

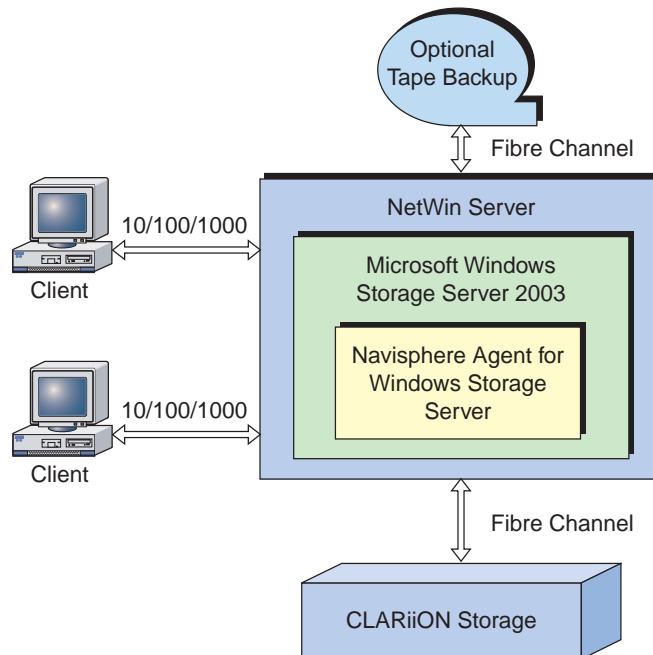


Figure 3. NetWin 200 System Components

Using NetWin with RepliStor, customers can consolidate data to a reliable, robust storage platform and replicate remote office data to a central or primary location where it can be managed by a backup and recovery professional using standard processes.

With NetWin 200, administrators can manage both server and storage resources from a single console using the familiar Windows interface. They will reduce time and resources by significantly simplifying monitoring, provisioning, and other storage management activities. NetWin 200 provides native Windows Storage Server 2003 management for the CLARiiON array or can be managed through EMC ControlCenter™.

Windows Storage Server 2003

Microsoft Windows Storage Server 2003, the NAS version of Microsoft's Windows Server 2003 operating system, provides standard system management tools for configuration of server and attached captive storage as well as file-level snapshot and path-failover capabilities. It enables file server consolidation and provides one of the best platforms for simplified file sharing, backup, and replication of business-essential data. Because Windows Storage Server 2003 integrates with existing infrastructure technology, enterprises can make full use of commonly used network environments and standard management software, as well as Active Directory service.

Windows Storage Server is a fixed-function device dedicated to file and print serving, as the operating system does not activate services unnecessary for file serving and storage tasks. Windows Storage Server 2003 NAS devices are

headless (without monitor, keyboard, or mouse) and are managed through a Web-based user interface or Terminal Server session.

By design, the ease of deployment and seamless integration into existing networks enables expanded storage capacity to users on the network in just minutes. The Windows Storage Server 2003 devices are especially appropriate for organizations that do not wish to put extensive IT resources into setting up multiple file servers and managing them, such as in a branch office scenario. They are designed to provide the storage administrator with the ability to keep information available to those who need it when they need it. The system is also designed to protect data from security risks, system failure, or data loss resulting from catastrophic events and is rapidly recoverable should the need arise.

Windows Storage Server 2003 enables organizations of all sizes to store and protect valuable data, from file servers in branch offices to fully integrated data centers. Windows Storage Server 2003 is optimal for server consolidation, backup, restore, replication and integration with storage area networks.

Combined with LEGATO RepliStor, it ensures high data availability through replication across multiple sites.

LEGATO RepliStor

At its core LEGATO RepliStor is a replication solution which replicates files, directories, volumes, shares, and registry keys from a source system, where data needs to be protected, to a target system or systems, where the data is then replicated. RepliStor delivers realtime, asynchronous file replication in one-to-one, one-to-many, and many-to-one configurations across local and wide area distances, using standard TCP/IP network connections. The ability to replicate data from geographically distributed sites to a centralized server over LAN or WAN makes RepliStor the optimal solution for centralizing data and can significantly reduce, or in most cases eliminate, IT requirements at remote locations, while leveraging administration and hardware resources found in the data center.

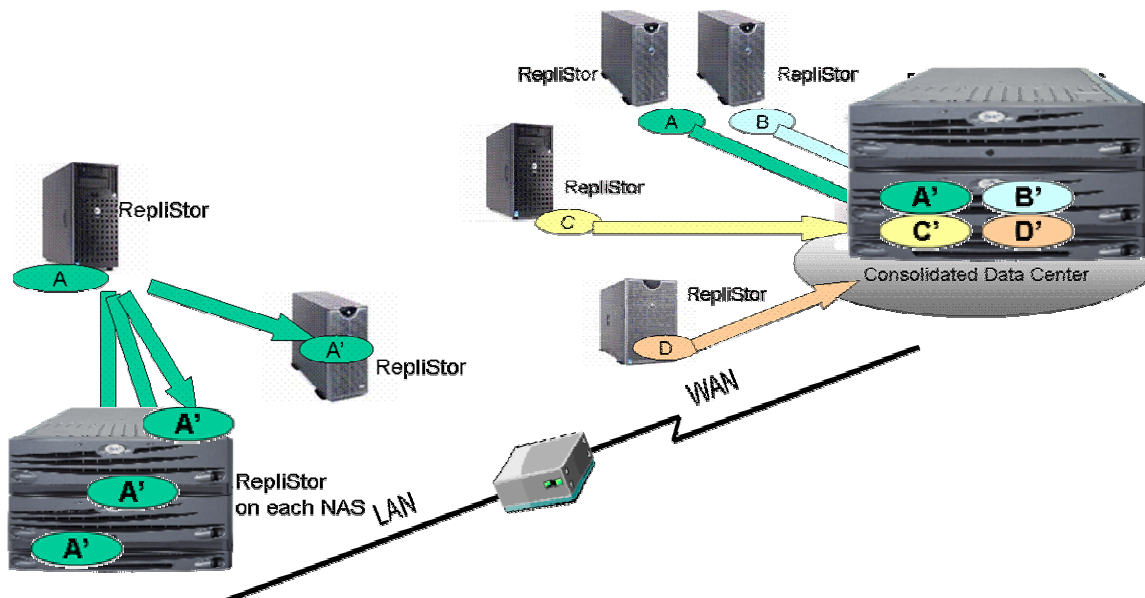


Figure 4. RepliStor One-to-Many and Many-to-One Data Replication

RepliStor enables easy remote management of replication operations from a centralized site. RepliStor does this, in part, by providing a single replication console to simplify management of replication across distributed NetWin 200 Windows servers and other Windows systems, and across any number of sites. This console enables a system administrator to monitor replication status constantly, to change replication direction, reconfigure replication, and to perform troubleshooting all without leaving the centralized site. RepliStor can also be easily deployed across the network, enabling new branches to be protected also without requiring the administrator to leave the centralized site.

After the initial synchronization of the data, RepliStor mirrors only the file operations for changes made to specific files, minimizing network traffic. It then replays those files' operations on the target system, ensuring both systems are in sync.

Advanced features of granular, rules-based replication such as filtering and scheduling allow complete flexibility in the selection of what information gets replicated and when the process takes place. Replication can be set to run automatically, or replication and synchronization can be scheduled to take place at a particular time or day. RepliStor's throttling options enable administrators to set limits on the percentage of available bandwidth that will be used for replication, preserving bandwidth for other types of network traffic to lower costs by using shared networks. RepliStor also provides compression of data, to use fewer networks than required by other replication products to move the same data. Administrators can also dedicate a network interface card to RepliStor traffic, making it easy to isolate replication from the rest of the network which helps to better manage replication impact on network bandwidth.

In the event of system failure or even if IT just needs to perform maintenance on a system, clients can be redirected to a current copy of data located on another system. That second copy of data on disk is also critical to a much more rapid recovery than restoring data from tape. RepliStor can also be configured for failover so that a source system moves network information to the target if the source system fails. This preserves the network connections from the RepliStor client to the RepliStor server as well as the data in the event of source system failures. Failover can take place over a LAN or a WAN without impact to application performance. Recovery from failover can take seconds, virtually eliminating downtime.

Finally, RepliStor's *save deletes* option provides an easy way to preserve files that might be deleted at the branch by accident or through intentional destruction. Administrators at the central site can easily and quickly return deleted files to the branch by reversing the direction of the replication.

LEGATO NetWorker Suite

LEGATO's information protection software suite protects the business-critical data of enterprise customers by simplifying and centralizing backup and recovery operations across all platforms. Built upon an open, highly scalable architecture, NetWorker increases IT staff productivity by automating the protection of storage assets for the largest corporate data centers to the smallest branch offices by enabling:

- Instant backup — recover and protect massive, increasing volumes of data faster
- Increased service levels and reduce downtime costs
- Increased IT staff productivity and lower TCO of storage resources
- Greater reliability, performance, and scalability of data protection operations across the enterprise
- Increase ROI of DAS, NAS, and SAN investments
- NetWorker PowerSnap Modules for complete management of instant backup and recoveries
- Microsoft Windows Server 2003 Volume Shadow Copy Services backup/restore and ASR recovery
- Improved NetWorker memory management for world-class backup/recovery performance to tape/disk
- Enhanced NAS and SAN protection
- Improved scalability, availability
- Similar policies and functionality for most supported operating systems
- Extensive reporting
- Access control on job functions

Volume Shadow Copy Service (VSS)

With the release of Windows Server 2003, Microsoft introduces VSS. VSS acts as a coordinator between all the components that create, archive, modify, back up, or restore data: the operating system, storage hardware, applications like Exchange 2003, and utility or backup programs like LEGATO NetWorker.

In NetWorker release 7.1, if a VSS client license is present, NetWorker software takes advantage of VSS technology to create snapshot backups of volumes and exact copies of files, including all open files. For example, databases and files that are open due to operator or system activity are backed up during a volume shadow copy. In this way, files that have changed during the backup process are copied correctly. Shadow copy backups ensure that:

- Applications can continue to write data to the volume during a backup.
- Open files are no longer omitted during a backup.
- Backups can be performed at any time, without locking out users.

Setup and Installation

This section provides an overview of the setup and installation processes required for this solution. The Information pertinent to this particular solution is provided along with references to product documentation for detailed installation and administration procedures. Setup and installation information is the same for the data center and branch office with the exception of LEGATO NetWorker which gets installed only at the data center.

EMC NetWin 200 and Microsoft Windows Storage Server 2003

The following sections provide planning, setup, and deployment information for the EMC NetWin 200 and Microsoft Windows Storage Server 2003.

Network Overview

Figure 5 is a detailed network diagram of the Branch Office Data Consolidation and Information Protection Solution and depicts two NetWin 200s in the main data center and one NetWin 200 in the branch office. Depending upon the volume of data being consolidated to the data center, more than one NetWin 200 might be required. Also, depending on the customer's IP configuration and other applications sensitive to the NetWin back-end IP addresses, different back-end IP addresses may be required for each NetWin when using more than one at one site. The diagram gives an example of using two NetWins at one site and how to configure the NetWin back-end IP addresses where this is necessary.

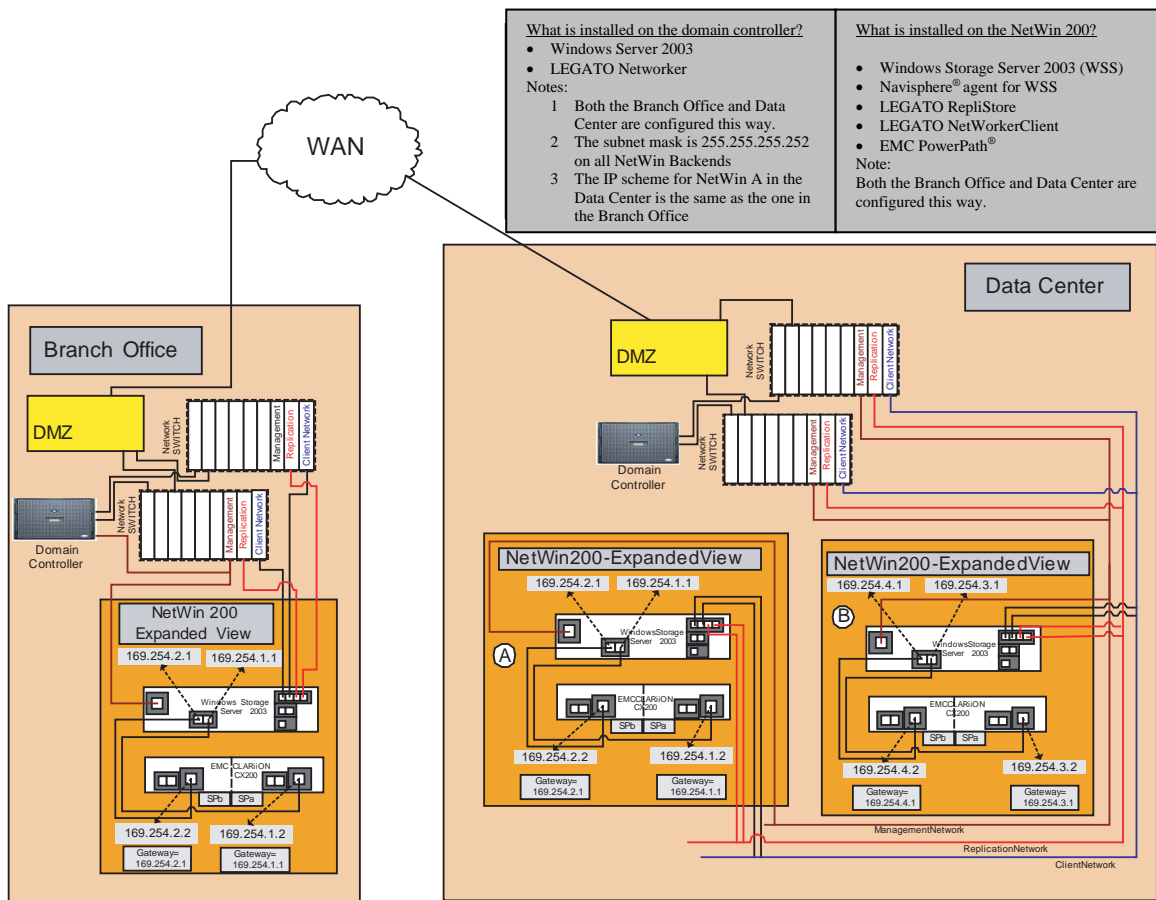


Figure 5. Branch Office Data Consolidation Network Diagram

Where any office uses a single NetWin 200, the IP addresses can be configured exactly as described in the *EMC NetWin 200 Installation Guide*, provided it does not conflict with the any existing IP addresses in the customer's network. All the NetWin 200 IP addressing schemes for all offices come down to being dependent upon the customer's network environment.

Note: The *EMC NetWin 200 Installation Guide* is not generally available and is used by the qualified EMC service technician during the NetWin planning and preliminary configuration process.

NetWin 200 Client Network

The server that comes with the NetWin 200 has six network ports, (not counting the built-in server management port found on the Dell server.) Two of these ports are built-in to the server, and the other four are in a quad-port PCI-based NIC.

As previously explained and depicted in Figure 5, the two built-in server network ports are used for the NetWin 200 back-end network and connect via cross-over cables to SP A and SP B on the CLARiiON storage array. The quad-port NIC is used for the client network and, at a minimum, uses a single cable that connects to that network. If the NetWin 200(s) in the data center will be used solely for storing data replicated from the Branch Offices, then all network traffic to the NetWin 200 (s) would be replication traffic (and backup traffic). These differences will be identified during the site planning process, which is outlined below. The diagram gives an example of how the quad-port NICS might be used. This example shows:

- Fault tolerance on the client network by adding a second client network port and attaching it to a second network switch.
- Fault tolerance and isolation of replication traffic from the client network by using the two remaining ports for a separate VLAN and allocating only replication traffic to them.

This example also has server management traffic isolated to a different VLAN, connected via the built-in port dedicated for server management. This may not be applicable if the customer is not implementing the server management features.

Any VLANs or the use of the other NIC ports will require planning and ultimately will be determined by the customer requirements and existing network infrastructure.

Site Planning Prerequisite

Site planning with the customer is necessary before attempting implementation of this solution. Some of the questions that should be addressed include:

For each Branch Office

- How much data will be stored on the NetWin?
- Of that data, how much needs to be replicated?
- Of the data to be replicated, how much changes on a daily basis? This is the key factor that you will need to determine the network bandwidth that you will require to replicate the data.

For the Data Center

- Will the NetWins in the data center be used only for branch office data replication, or will they also be file servers for local data center clients?
- What is the total amount of data to be replicated from all the branch offices? This is necessary to determine the amount of NetWin storage required at the data center as it is possible to exceed the capacity of one NetWin; therefore multiple NetWins might be required.

In addition to the areas above, the following also need to be addressed as part of the site planning process:

- **Determine bandwidth between each branch office and the data center**

TIP: Use Windows Performance Monitor to measure the rate of change of data occurring on your systems over a week. Use that data to then plan your network requirements. Remember that RepliStor replicates only the changes in data as they occur, and does not re-replicate the entire data set when one portion changes.

- **Server Consolidation Planning**

Besides data consolidation via data replication between offices for backup and recovery purposes, implementation of the NetWin presents the opportunity to consolidate the customer's file servers. Consolidation can take place in both the Branch Offices and Data Center where multiple file servers are currently in use. A consolidation plan around these servers needs to be part of the site planning process.

- **Directory Structure Planning**

Good directory structure planning for the Branch Offices and the Data Center needs to be completed before deploying this solution in order to understand what file shares need to be created on the NetWin 200 in both the branches and data center, and must be understood when setting up replication between offices using RepliStor. Where servers are being consolidated as described above, directory structure planning should be looped into that task.

- **High-Availability Planning**

In addition to the HA networking option demonstrated in the sample diagram, an additional HA consideration is the power connection. The NetWin 200 comes with one standby power supply (SPS). For greater HA, two SPSs should be used. Further each SPS can be connected to separate power grids. This will allow continued operations in the event one grid fails. Customers interested in this should contact their EMC Sales Representative and include it as part of the site planning process.

EMC PowerPath[®] is also noted on the diagram. EMC PowerPath is an option that enables I/O load balancing between the two FC (Fibre Channel) connections which enhance performance. For HA, it also enables failover, so if any failure occurs at any point-one fibre path, the other path will pickup the load and there will be no operational interruption (but there will be a likely performance degradation until the failed component is repaired or replaced). Customers interested in this should work with their EMC sales representative as well as the SI, and include it as part of the site planning process.

EMC NetWin 200 Setup

Once the site planning is completed and the NetWin 200(s) has been obtained, a qualified EMC service technician must perform the following initial setup procedure before the SI or customer can implement this solution.

1. Read the latest *EMC NetWin 200 Release Notes*.

Note: It is recommended that everyone involved in the project, including the customer, read the *EMC NetWin 200 Release Notes*.

2. The NetWin 200 Configuration Worksheet must be completed by the EMC service technician, the SI, and the customer together. This worksheet is supplied by the EMC service technician (see Appendix E of the *EMC NetWin 200 Installation Guide*).
3. Install and configure the EMC NetWin 200 systems by performing the following steps:
 - a. Unpack and install the NetWin 200 into the server rack.
 - b. Connect all cables and power on the system.
 - c. Perform the base CLARiiON software installation.
 - d. Configure the NetWin 200 server user name and password as specified on the NetWin 200 Configuration Worksheet.
 - e. Configure the service processor's IP configuration for the integrated network ports as specified on the NetWin 200 Configuration Worksheet.

Note: This step may or may not be necessary depending on whether the NetWin 200 default IP addresses conflict with customer's network environment, or if there is more than one NetWin 200 at each site.

- f. Configure the NetWin 200 default user IDs and passwords as specified on the NetWin 200 Configuration Worksheet.
- g. Configure the public network ports as specified on the NetWin 200 Configuration Worksheet.
- h. Run the NetWin 200 setup scripts to set default cache and LUN failover settings.
- i. Initialize global security using Navisphere®.
- j. Set initial properties for each SP including Date/Time, Fibre Speed, Statistics Logging (enabled), and the SPS battery test time.
- k. Create RAID groups and bind LUN 0.
- l. If the customer is using PowerPath, set the LUN failover values to reflect this on the array.
- m. Verify HBA settings.
- n. If the customer purchased PowerPath, install it on the server.
- o. Create the EMC reserved volume.
- p. Configure the EMC *e-mail home* feature as specified on the NetWin 200 Configuration Worksheet.
- q. Complete the installation verification tests.

After the preceding steps have been completed, the NetWin 200 can be turned over to the customer, and is ready for the *Branch Office Data Consolidation and Protection* solution deployment.

EMC NetWin 200 Deployment

This section outlines the deployment tasks and recommended reading for completing the deployment of the NetWin 200 in the *Branch Office Data Consolidation and Information Protection* solution. These tasks can be performed by the SI or the customer once an EMC service technician has completed the NetWin 200 preliminary setup steps.

Important

- The tasks in Table 1 are a subset of a complete EMC NetWin 200 task list that can be found beginning on page 3-8 of the *EMC NetWin 200, Ver. 1.2.0, System Administrators Guide, 300-001-185, Rev A03*.
- These steps reflect the specific document version noted above. As new documentation revisions are released, the steps and page numbers referenced are subject to change.
- The *Information Location* column below references pages in that document, as well as to the Microsoft Windows Storage Server 2003 online help, accessible through the Web UI.
- Though only the steps necessary to complete this solution are included below, it is recommended that you refer to and read the entire *EMC NetWin 200 System Administrators Guide*.

Table 1. Preliminary Tasks

Task	Information Location	Description
If not already done, power up the NetWin system.	Page 4-3	Provides the proper steps for power up. Note: It is critically important to always follow the power up and power down steps as documented.
Install the antivirus software.	Page 9-7	Provides information for installing antivirus software on your NetWin 200 system.
Join the NetWin 200 Server to your Windows domain.	Refer to Microsoft documentation for joining a server to a domain.	Join NetWin 200 Server to the customer's Windows domain.
Log in to the NetWin 200 Web UI.	Page 4-4	Describes how to login to the Web UI.
If not completed already, enter a DNS suffix for your NetWin 200 system.	Page 1-6, step 6	Provides the steps required for entering the DNS suffix.
If not completed already, add a hot-spare disk on the array.	Page 6-5	<i>Managing Hot Spare Disks.</i> Be sure to follow the recommendations on how many hot spares to add to the array.
Verify the latest Microsoft patches and updates are on your system.	Page 1-7, step 8	Assures your system is fully protected with the latest updates from both Microsoft as well as EMC.
Schedule backups of the NetWin 200 front-end server for server system state information and customer data.	Chapter 9	Describes backing up and restoring the NetWin system state and NetWin user.

At this point, all the initial configuration steps should be complete. It is recommended that you complete the reading outlined in Table 2 before continuing.

Table 2. Preliminary Reading

Read	Information Location	Description
<i>Introduction to the EMC NetWin System</i>	Page 1-2 and 1-3	Provides an overview of the NetWin 200 Architecture.
<i>High-Availability Features</i>	Page 1-4 and 1-5	Describes the HA features of your NetWin 200. Note: LEGATO RepliStor is used in place of EMC OnCourse™ (mentioned on page 1-5) for this solution. Besides the content of this solution document, page 9-8 also has a reference to RepliStor.
<i>Getting Started</i>	Page 1-6 and 1-7	The pertinent getting started steps have been covered in this table. This is recommended here only as a double-check.
<i>Understanding NetWin Terminology</i>	Chapter 2, Page 2-2 and 2-3	Describes the terms you should be familiar with while working with the NetWin 200.
<i>Contacting Customer Support</i>	Page 10-35	Overview of how to engage EMC support in the event you need assistance.
<i>Using the Web UI</i>	Microsoft online Help; select the Welcome tab.	Provides orientation to the Microsoft Windows Storage Server 2003 Web interface.
<i>Managing your NetWin System</i>	Chapter 4	Provides indepth information to help you understand array concepts, power up and power down procedures, using the NetWin Restore DVD and a number of other important steps.

NetWin 200 Administration

Prerequisite: Before proceeding with administration tasks such as creating and managing volumes and shares, all preliminary configuration steps should be complete and you should have a good understanding of what is required to manage the EMC NetWin 200.

Administering the NetWin 200 includes various tasks such as:

- Managing volumes
- Managing shares
- Managing the CLARiiON array
- Managing shadow copies (a new feature in Microsoft Windows Storage Server 2003)
- Managing NetWin local users and groups
- Backup, restore, antivirus monitoring, and replication
- Troubleshooting

Because it is important to understand why you perform these tasks, this table references entire chapters rather than specific sections within the chapters. After the solution is deployed, any remaining sections of the administrator guide that have not been read should be read. After deployment, it is also recommended to refer to the *Documentation Reading Path* which can be found in Chapter 3, pages 3-2 through 3-7, for further reading.

Table 3. Reading Path for NetWin Deployment in a Branch Office Data Consolidation Solution

Task	Information Location	Description
<i>Managing your NetWin Volumes and Shares</i>	Chapter 5	Among other administrative tasks, this chapter describes how to create and share volumes, as well as how to create and manage shares.
<i>Managing your CLARiiON Array</i>	Chapter 6	Describes various array management tasks including how to perform rescans of the array, how to manage hot spares, and how to change the Management Network Configuration.
<i>Managing your NetWin Shadow Copies</i>	Chapter 7	Describes how to prepare for, create, delete, access, and manage Shadow Copies.
<i>Managing your NetWin Local Users and Groups</i>	Chapter 8	Describes how to add, modify, and remove users and groups.
<i>Backup, Restore, Antivirus, and Replication Support</i>	Chapter 9	Among other administrative tasks, this chapter describes how to back up and restore the NetWin System State as well as User Data.

This completes the EMC NetWin 200 planning, setup, and deployment steps.

LEGATO RepliStor

LEGATO RepliStor needs to be installed on any system to be designated as a source or a target (for example, both on the branch systems and on the data center systems). It can be installed from a CD-ROM or a command line. Instructions are detailed on pages 15-20 of the *LEGATO RepliStor Installation Guide*.

Before You Install

Be aware that RepliStor software must be installed as a service that runs under the SYSTEM account or an authenticated account in the administrators group (or Domain Admins, if installed in a domain). It cannot run under a user's account because of Windows operating system limitations. Specifically, RepliStor software needs to communicate with the RepliStor kernel mode driver. Only SYSTEM or Administrator users are allowed to perform

this function. For information on securing remote servers, see *Setting Up Secure Access to Remote Servers* on page 22.

During installation, RepliStor software prompts you for the following information:

- The pathname to install RepliStor programs. The recommended path is:

```
%SystemDrive%\Program Files\Legato RepliStor
```

- The pathname to store RepliStor data files. The default path is:

```
<CommonAppDataFolder>\Legato RepliStor Data
```

Data files can be stored in the same location as the program files.

Note: To change the program directories after installing RepliStor software, you must reinstall it. To change the location of the DATA directory after installing RepliStor software, you must change it on the Directories tab in the Options dialog box (select Options from the Maintenance menu).

RepliStor Software Use of Windows Active Directory

When installing RepliStor software in a Windows 2000 or later domain, you can optionally register it in Active Directory. This has the following benefits:

- There is a common site list for all RepliStor installations in the domain. If any server makes a change to the site list, all servers see that change immediately.
- There is a common users list for all installations in the domain (on Users tab in Options dialog box).
- The site list is available immediately. Manually adding a site is only necessary if the site is outside the domain.

Note: In order to register RepliStor software in Active Directory, the user performing the installation must be in the Domain Admins group. Once RepliStor software is installed, a domain administrator can assign administrative rights to any other user or group.

RepliStor software uses the base Active Directory schema and does not extend the schema. The only portion of Active Directory that RepliStor software modifies is:

```
DC=domain
```

```
  CN=System
```

```
    CN=LEGATO
```

```
      CN=RepliStor
```

Under the RepliStor entry is an entry for each installed RepliStor server, as well as entries for each external RepliStor server (as added into the Site List). Use the following information to determine the additional disk space the system requires to replicate RepliStor Active Directory updates:

- Each RepliStor installation typically requires less than 1 KB of additional space in Active Directory, although this depends on the number of characters entered in the text fields in the Site Properties dialog box (for example, Description, Site Name, Account, Password).
- All attributes that RepliStor software uses in Active Directory are sent during both full and partial replications.
- RepliStor Active Directory updates are confined to user-initiated events, such as installing and removing RepliStor software, and modifying site properties.

Obtaining a License Key

When you first install RepliStor software, you can select **Request Evaluation License** for a trial version. Before you install a non-evaluation copy of RepliStor, however, you must obtain a license key from LEGATO for each computer on which the product will run. When you obtain a license key, you automatically register the software and activate support for RepliStor software.

If you are installing RepliStor from a CD-ROM, you must contact LEGATO and provide the serial number listed on the software package to obtain a license key. For additional information, see *Licensing and Registration* on pages 9 and 14 of the *LEGATO RepliStor Installation Guide*.

Setting Up Secure Access

The following section describes how to set up access on the RepliStor servers after installing the product.

Important: In order to attach to a remote RepliStor site, you must have an account on the remote system in the Windows administrator group. If the remote system is in your domain and you are logged into that domain, then you automatically have a valid account on the remote system because Windows maintains one database of users for the entire domain. However, verify that your account is in the administrator group on the remote system.

It is possible, though unlikely, for a user to create a program and run it under the SYSTEM account, thus gaining access to a remote RepliStor server that is thought to be secure. If you are concerned about this, make sure that the RepliStor server is not running under the SYSTEM account, but is running under an authenticated account in the administrator group.

RepliStor Remote Installation

The RepliStor Remote Install Utility is a command line utility that can be used in script files to perform a remote install, upgrade, or uninstall of RepliStor. With this utility, it is possible to create a batch file that will install RepliStor on any number of remote servers.

LEGATO NetWorker

There are four NetWorker components that can be installed dependent upon the environment including the NetWorker Client, Server, Storage Node, and License Manager.

The minimum components that must be installed are:

- NetWorker Client
- NetWorker Server

Note:

You can install LEGATO License Manager with any of these NetWorker software installation types.

NetWorker Server and Client should be installed on the designated NetWorker Server.

NetWorker Client should be installed on the designated data center NetWin device or any additional clients to designated to be backed up by NetWorker.

NetWorker Installation Roadmap

The following roadmap provides an overview of the installation tasks. For detailed step-by-step instructions, refer to the *NetWorker Installation Guide, Release 7.1, Microsoft Windows Version*.

- Ensure that all installation requirements are satisfied (see *Installation Requirements*).
- (Optional) Determine whether a reboot is required during the installation (see *Checking the Microsoft Windows Installer Version*).
- Access the installation files (see *Accessing the NetWorker Installation Files*).

- Install the NetWorker software. See the appropriate section:
 - *Installing NetWorker Software*
 - *Installing or Removing NetWorker Software by Using SMS*
 - *Installing NetWorker Software in an MSCS Environment*
- Test the NetWorker software (see *Verifying the Installation*).
- Configure one or more storage devices (see *Chapter 3: Configuring Storage Devices*).
- Enable and register the NetWorker software (see *Chapter 5: Evaluating and Licensing NetWorker Software*).

Configuration

This section outlines requirements for configuring replication and failover for the Branch Office Data Consolidation and Information Protection Solution.

Replication

This section outlines information and procedures for configuring the branch office and data center locations for data replication. For basic RepliStor operating instructions, refer to the *RepliStor Administrator's Guide, Release 5.3, Microsoft Windows Version*.

Before configuring or administering a remote RepliStor site, you must attach to and start that site.

Important: To connect to a remote RepliStor server, you must have an account on the remote system that is in the Windows administrators group.

Creating Specifications

After attaching to the local and remote sites, create specifications that tell RepliStor software which data to replicate (or exclude from replicating). A specification identifies the path and file information that RepliStor software requires to replicate data from a source system to a target system.

You can configure the following types of specifications with the RepliStor software:

- Global Exclude — Identify files, directories, and shared directories that you do not want replicated from the source system. A Global Exclude specification takes precedence over all other specifications created after it.
- File/Directory — Identify the files, directories, and shared directories to replicate to a target system.
- Share — Replicate all of the shares on the source system to the target system.
- Registry — Replicate specific keys within the HKEY_LOCAL_MACHINE Windows Registry hive.
- DFS — Replicate Windows NT Distributed File System (DFS) shares on Windows NT only.

Keep the following in mind when creating specifications:

- Create Global Exclude specifications first, since they take precedence over all other specifications.
- Do not replicate the following types of files:
 - Virtual drives (same as redirected drives)
 - RepliStor files
 - Windows system files
- RepliStor replicates compressed files on a source system as compressed files on the target system.
- If the source or target system uses NTFS file systems, there are special requirements for the user ID in addition to those described in the *LEGATO RepliStor Installation Guide*:
 - The user ID used to add a specification must have Full Control permission on the source system for the files, directories, and shares to be replicated.
 - The user ID used to add or modify a specification must have Full Control access on the target system to the files, directories, and shares to be replicated.

The account under which RepliStor software is run, typically the SYSTEM account, must have Full Control access to files on both the source and target systems.

For step-by-step instructions for creating the following specifications, refer to the corresponding section in the *RepliStor Administrator's Guide, Release 5.3, Microsoft Windows Version*.

Global Exclude Specifications

Global Exclude specifications identify files, directories, and shared directories that you do not want replicated from the source system. The RepliStor software comes with the following Global Exclude specifications:

- One for each swap file (c:\pagefile.sys)
- RepliStor program directory (usually c:\<Program FilesDirectory>\LEGATO RepliStor)
- RepliStor DATA directory (c:\Documents and Settings\All Users\Application Data\LEGATO RepliStor Data)
- All System Volume Information directories on all drives
- All Recycler directories (or Recycled directories on FAT volumes)
- All \$EXTEND directories on NTFS file systems.

Important: Configure Global Exclude specifications before configuring file/directory or share specifications. If you configure file/directory or share specifications first, RepliStor software performs the initial synchronization, and any subsequent Global Exclude specifications do not take effect retroactively.

File/Directory Specifications

File/directory specifications identify the files, directories, and shared directories to replicate to a target system.

Once you have configured a file/directory specification, you must synchronize the specification in order for RepliStor software to replicate the selected files.

Share Specifications

A share refers to a shared directory or folder. Shares can be replicated through a File/Directory specification, but if you want to replicate the shares separately, follow the steps in this section.

If you are creating a Share specification, then you must maintain identical file structures on both the source and target systems (for example, you cannot create a specification that replicates C:\SHARE\ on the source system to D:\BKSHARE\ on the target system). If you do not have identical file structures on both systems, then you must create a File/Directory specification to replicate that file structure from the source system to the target system and select the **From Shares** option in the **Options** dialog box to replicate and translate the shares. If you are replicating multiple shares located in several locations on the source system, it is recommended that you replicate the entire drive from the source system to the target system.

Keep the following in mind when creating a share specification:

- You cannot create a share specification if you are running RepliStor software in a Microsoft Cluster Server (MSCS) environment.
- Once you have configured a share specification, you must synchronize the specification before RepliStor software can begin replicating the selected files.
- You cannot replicate network printer shares.

Registry Specifications

When an application is installed on a system, configuration information is generally placed in the HKEY_LOCAL_MACHINE Registry hive. You can create specifications to replicate specific keys within the HKEY_LOCAL_MACHINE Windows Registry hive. This configuration information may be static, where the information does not change after installation, or it may be dynamic, where the application updates its Registry key during normal use. Thus, to switch over to the target an application with dynamically updated configuration information, and have that application run as it did on the source, you may need to replicate the updates to its registry key.

Note: Some applications also update the HKEY_USERS Registry hive with per-user information. RepliStor software cannot replicate these hives.

Important: You must be familiar with the exact information a specific application places in the Registry in order to know which keys to replicate. In addition, do not replicate machine-specific information from a source to target key, since this can cause serious problems on the target.

When creating a Registry specification, you can specify that permissions associated with an application's key be replicated from the source system to the target system. If you do replicate permissions, the target system must be within the source system's domain, or the target keys will be inaccessible. Note that the permissions are based on Security Identifiers (SIDs), not on text account names, which is an option with file specifications. A SID has no meaning outside of its domain.

If you do not replicate permissions to the target, the target uses the parent key's default permissions. As a rule, only replicate small increments of the Registry — for example, only the key for a critical application. In addition, to avoid writing over configuration information in the target Registry you may want to stage the replication and specify a different key path on the target system.

Synchronizing Specifications

Before RepliStor software can replicate data successfully, an exact copy of the data from the source system must reside on the target system. To do this, you must synchronize the specifications.

Important: If you accept the defaults for synchronization when creating specifications, RepliStor software does the initial synchronization for you automatically.

Updating and Forwarding Options

RepliStor software allows you to schedule updating and forwarding. Remember, updating is a function on the target system during which updates from the source are applied to data files on the target system. Forwarding is a function on the source system during which changes to data are forwarded from the source to the target system. Forwarding and updating are always on by default.

You can schedule forwarding or site forwarding. Turning off forwarding not only discontinues forwarding, but also discontinues the heartbeat signal that is sent from the source system to the target system.

Throttling Options

Throttling is an optional feature that can give you greater control over network bandwidth. With throttling, you can determine the maximum data rate, in kilobytes per second, for your replicated data. By setting this rate, you ensure that the amount of replicated data transferred across the network from the current source system will never exceed the maximum data rate as long as throttling is enabled.

Throttling can be permanently enabled, or turned on and off on a schedule. You can also configure the throttling rate for all target sites, each target site individually, or no target sites. RepliStor software can also protect against poorly designed throttle settings. Since it is possible for replicated data to accumulate faster than the throttle rate, the files on the source system could grow until all disk space has been exhausted. To avoid this, enable the RepliStor Disk Space Monitor options (see *Configuring RepliStor Options* on page 58). When these options are enabled, the throttle rate is automatically disabled if the free disk space goes below the warning level. It is re-enabled automatically when the free disk space goes above the warning level.

Failover

RepliStor software can be configured to switch over from the source system to a target system when the source system fails. This section outlines how to configure RepliStor switchover options in a Windows Active Directory environment using aliases. For additional options and detailed instructions on configuring switchover, refer to the *RepliStor Administrator's Guide, Release 5.3, Microsoft Windows Version*.

Super Automatic Switch Over (ASO) with Aliases

With Super ASO, you can define any number of source system aliases to use during switchover, along with IP address and DNS server information. Aliases are designed primarily for Windows 2000 and later configurations and allow you to switch over between two Windows 2000 or later servers with Active Directory configured. They also make it easier to recover from a switchover with minimal disruption of client connections. With an alias, the source

system can be brought up without affecting client connections to the alias on the target system. This allows the source system to be repaired and brought up, and the data synchronized back to the old source without having to bring down the target system and disconnect any clients using the target system.

Figure 6 shows that clients connected to the alias on the source before switchover are still connected to the alias on the target after a switchover. Specifications can be sent automatically to the target and continue to replicate back to the source system. When the source is restored, simply remove the source machine name from the target, and all aliases and specifications are removed from the target automatically. The aliases are added back to the source. You then re-enable and synchronize the specifications on the source; client connections are disrupted for only the time it takes to bring up the alias, only a few seconds.

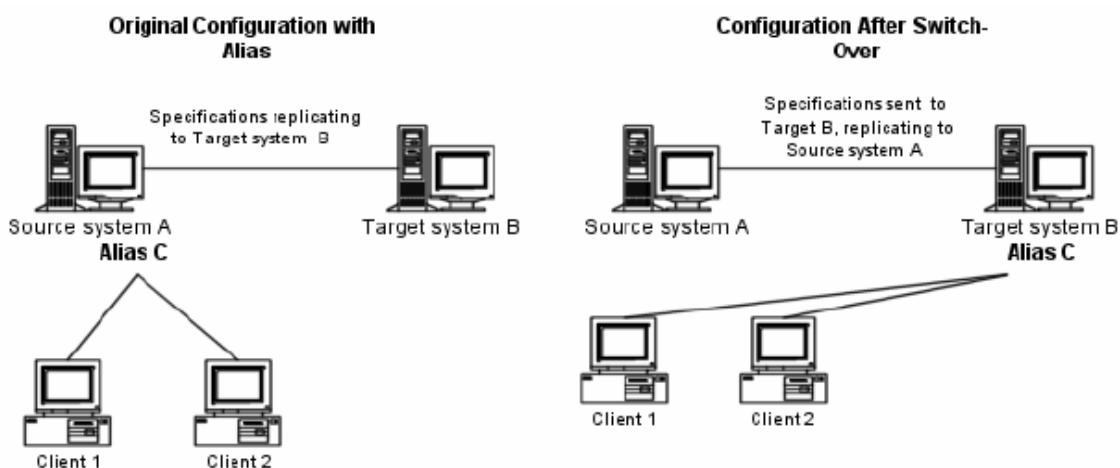


Figure 6. Using Aliases for Super ASO

System Configuration

With Super ASO, the target computer must be configured in one of the following ways:

- Windows Primary Domain Controller (PDC) to Windows Backup Domain Controller (BDC)
- Windows BDC to Windows PDC
- Windows BDC to Windows BDC
- Windows Stand-Alone Server to Windows Stand-Alone Server
- Windows Workstation to Windows Workstation
- Windows Member Server to Windows Member Server

Switchover Setup

The following task list is an overview of the tasks involved in setting up switchover:

Task 1: Start the RepliStor client and attach to the source system. Ensure that the RepliStor server is running.

Task 2: Configure switchover options for the source system:

- For Super ASO in a Windows 2000 or later domain-controller environment or for switchover across a WAN, define one or more source system aliases. You can also specify IP addresses (except for WANs) and optional DNS information to be used during switchover.
- In an environment where there are too many variables to allow an automatic switchover, such as a WAN failover, you may want to define the alias for manual activation on the target system. For more information about manual alias activation, see *Configuring Manual Alias Activation* in the *RepliStor Administrator's Guide 5.3, Microsoft Windows Version*.

- Determine how often the RepliStor heartbeat signal is sent from the source to the target system.
- If you are not defining aliases, select the IP addresses that you want to forward to the target system.
For more detailed information about source system options, see *Configuring Switchover on the Source System* in the *RepliStor Administrator's Guide 5.3, Microsoft Windows Version*.

Task 3: Attach to the target system, and ensure that the RepliStor server is running.

Task 4: Configure switchover options for the target system:

- Select Super ASO.
- If you created aliases on the source system, select Accept Aliases (Super ASO only).
- Specify whether you want specifications forwarded and automatically synchronized (Super ASO with aliases).
- Determine who will be notified when a switchover has occurred.

For detailed information about configuring ASO target options, see *Configuring Super ASO on the Target System* in the *RepliStor Administrator's Guide 5.3, Microsoft Windows Version*.

Data Backup

The following sections explain the configuration for backing up data at the data center using NetWorker.

How LEGATO NetWorker Software Works

NetWorker software works as a client/server system consisting of NetWorker servers, clients, storage nodes, storage devices, and storage volumes.

The NetWorker client is any computer that has the NetWorker client software installed, and that backs up the hosts data to be written to a backup device. Data is stored in savesets. A NetWorker client can be a workstation or a dedicated server. A NetWorker client sends its data over the TCP/IP network to another host that writes the data to a backup device.

The NetWorker server is a computer that orchestrates the back up of all hosts under its control. It holds the media database, the client file index, and the configuration database. The media database contains information about which savesets are on which backup media and where that backup media currently is. The Client File Index knows which files are in which savesets, and the configuration databases hold all the policy and hardware setup information.

A Storage Node is, essentially, a client that can backup directly to a backup device, without going across the TCP/IP network. There are two variants of Storage Node. The Dedicated Storage Node can only write data from the local host, whereas a full Storage Node can receive data from other clients.

The NetWorker server, Storage Node, and Clients can all be running a different platform or operating system from each other. NetWorker servers can also back up to other NetWorker servers as part of a hierarchical backup structure.

When you submit a request to view or browse backup data, the NetWorker server looks to the client file index for a record of each backed-up data item. In cases where the data item has been backed up several times over a specified period, the client file index provides this information, as well as the option of recovering different versions of the same data item.

When you request recovery of a selected data item, the NetWorker server searches the media database to locate the specific volume or volumes containing the data items you want.

Once you select the data you want to recover and the NetWorker server locates the appropriate volume, the data is written back to the original client location; or in certain cases, such as directed recovery, to other selected client locations. In cases where recovered data might overwrite data already on the client, the NetWorker server alerts you and provides several options for solving the conflict. You can also configure in advance the action to take when recovery conflicts occur.

Figure 7 shows the NetWorker data flow, starting with the NetWorker client and ending with the storage volume.

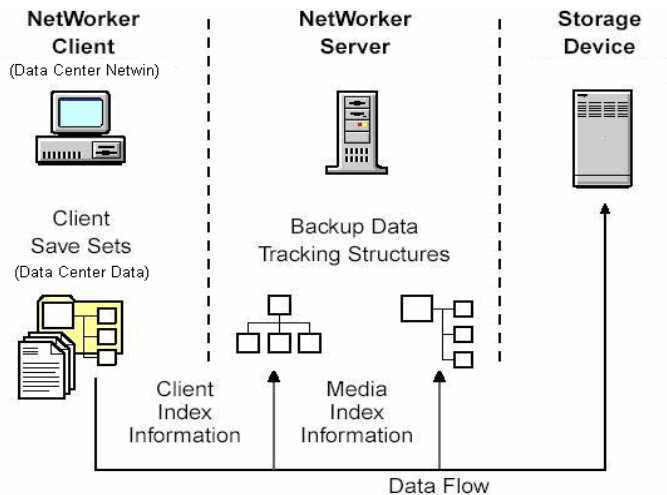


Figure 7. Data Backup Flow

For detailed step-by-step instructions or to customize NetWorker resources and attributes, and for information on how to configure the NetWorker software for your particular needs, refer to the *LEGATO NetWorker 7.1 Administrator's Guide for Microsoft Windows*.

NetWorker Setup

The following task list provides an overview of the tasks involved in setting up NetWorker. For detailed step-by-step setup instructions and to customize NetWorker resources and attributes for your particular needs, refer to the *LEGATO NetWorker 7.1 Administrator's Guide for Microsoft Windows*.

Task 1: Start the NetWorker Administrator client and connect to a NetWorker server.

The Server window opens (with the Configure tab active). The Server window displays the name of the server to which you are currently connected as a NetWorker client. Note that a server is by necessity a client of itself.

Task 2: Mount a Volume for Backups.

To perform backups, the backup media must be configured and available for the NetWorker server to write data to it. Once you have started the NetWorker Administrator program, label and mount a blank media volume.

If Device autodetection is enabled, the device should be listed in the Devices window. If it is not enabled, you must either perform an autodetect, or create the device manually.

Task 3: Configure NetWorker clients (Data Center NetWin device).

The LEGATO NetWorker product provides preconfigured settings for backing up clients, including the server, which is considered a client of itself. You can use these preconfigured settings, or create new settings.

Note: Each client backing up to a NetWorker server must be assigned to a backup group, either a preconfigured group or a group you create yourself. In many cases, the NetWorker default group is sufficient for the clients, and creating additional groups is not necessary.

You can now start this as backup.

Data Backup Methods

NetWorker software contains preconfigured resources you can start using immediately; therefore, when the volume is successfully labeled and clients are correctly configured, you are ready to perform a backup.

There are two types of backups in the NetWorker environment: *scheduled* and *manual*.

- Scheduled backups are configured in the NetWorker Administrator and are assigned a specific time for the backup to run, usually when network usage is low. For example, the preconfigured Default group is scheduled to begin its backup at 3:33 A.M. You can override scheduled backups to perform the backup immediately, if necessary.
- Manual backups are backups performed by the user of the client computer, the NetWorker User program (winworkr).

To run a scheduled backup, you need only enable the default schedule.

You can test the scheduled backup by overriding the scheduled time and starting the backup immediately.

For detailed instructions on performing a backup, refer to the *LEGATO NetWorker 7.1 Administrator's Guide for Microsoft Windows*.

Database Backup Considerations

Pause Updates and Turning On/Off Forwarding and Open Files

Whether you pause forwarding on the source or pause updating on the target, you have the same challenge. It may be that a file is open when the pause took place and subsequent updates will not be applied until the “*pause*” is turned off. If they are database files that are opened continuously, you may lose a few transactions.

This is a problem that you would have with any type of snapshot or backup scenario. To address open files, RepliStor allows the backup of the open file, but if an application is in the process of writing to that file, it obviously cannot back up data that has not been sent yet.

With this type of configuration it is preferable to turn off forwarding on all sources, then do a pause updates on the target, then do the backup, then turn off pause updates and turn on forwarding on the sources. In other words don't just pause updates. If you have lots of data coming in with updates paused, it may become unwieldy very quickly on the target, since all the Branch Offices would still be replicating to the Central Office and the data would need to be queued there. It is better to spread that data out amongst all the sources, for two reasons:

1. Each source needs to queue a portion of the data that the target would need to queue
2. The queuing mechanism on the source has less overhead.

Important: When turning on/off pause updates on the target server in the Central Office, this affects all of the specifications on the target server. No updates from any of the Branch Office RepliStor source servers can be made.

Pause updates is useful because it will perform a Close Application Files function that makes sure any future updates are saved in the data directory (block.rdf file) and the files are not reopened. We would need to pause the updates from RepliStor so the files are closed on the NetWin 200 in the Central Office. To keep this paused for any length of time would impact performance. Thus, we pause them only for the time needed to create a snapshot, then re-enable them and back up from the snapshot.

Whether you turn off forwarding on the source RepliStor server or pause updating on the target RepliStor server, the result is a point-in-time representation of the files, which is not the same as if all the files on the source were closed and then replicated. It may be that a file is open when the pause took place and subsequent updates will not be applied until the “*pause*” is turned off.

Offline Database Backups

The solution described in this document assumes backups will occur only at the Central Office, so the data base must first be quiesced at the Branch Office, then the snapshot can be taken in the Central Office. If data bases are running in the Branch Offices, the only way to back them up is to follow the procedure that the vendor provides. LEGATO NetWorker provides database modules exclusively for this purpose. The procedure usually involves quiescing a database (putting it into what is known as hot backup mode), taking a snapshot, then putting the database back into normal operating mode, then making a backup from the snapshot that was taken. Usually the snapshot will occur on the same host as the running database.

Additionally, if NetWorker is not running in the Branch Offices, there is no way to recover clients' data in the Branch Offices using only NetWorker. The snapshots will not be taken by NetWorker in the Branch Offices. One way to ensure databases get backed up properly is to use the tools that Windows Storage Server 2003 and NetWin provide. Quiesce the database, take a snapshot on the source RepliStor server in the Branch Office, and replicate the snapshot to the Central Office using RepliStor, where a backup can be made.

More information available at: <http://portal2.legato.com/products/replistor/BackupSolSheet.pdf>

Online Database Backups

A valid SQL Server online backup can be made in one of two ways:

1. The tablespace(s) are in hot backup mode when the snapshot is taken. This is the best practice.
- Or*
2. The restartable point-in-time image is backed up using a recovery manager. The backup produced using a method like this is valid; the point-in-time image on the snapshot is not a valid backup.

The following are the best practice steps undertaken using RepliStor and NetWorker to create a valid SQL Server backup in the central office.

1. Quiesce SQL Server running in the Branch Office, placing the tablespace(s) in hot backup mode.
2. Wait until all changes are sent to the Central Office. Turn on Pause Updates on the RepliStor Target server. (Additionally, turn off forwarding in the Branch Office.)
3. Create the snapshot on the RepliStor Target server in the Central Office.
4. Take the tablespace(s) out of hot backup mode.
5. Force the SQL to archive the current online redo log.
6. Make a copy of the Control File in the archive log file system.
7. Turn off pause updates on the RepliStor Target server and turn on forwarding on the RepliStor Source server.
8. Perform an SQL backup with NetWorker in the Central Office using the snapshot taken in the Branch Office.

Data Recovery

These sections outline using RepliStor and NetWorker capabilities together to recover data.

Recovering Data with RepliStor

Depending on the level of failure that occurs, different options may be needed to speed recovery. The combined solution from EMC, LEGATO, and Microsoft provides various options for recovering data and business operations, according to the level of failure that has occurred. Some examples of the levels of failure that can occur and the levels of recovery that may be needed are outlined in Table 4.

Table 4. Levels of Failure and Potential Recovery Scenarios

Level of Failure	Potential Recovery Scenarios
Single file is accidentally deleted	<ul style="list-style-type: none"> Use Windows Storage Server VSS capability onsite to rapidly recover the data. Restore from a replicated copy of the file, from the centralized site to the local site, through the RepliStor central management console
Entire Volume or Directory is lost	Restore from a replicated copy of the volume or directory, from the centralized site to the local site, through the RepliStor central management console
Entire system is lost	Data Recovery: <ul style="list-style-type: none"> Rebuild the system Restore the data from the replicated copy, from the centralized site to the local site, through the RepliStor central management console
	Automated System Recovery Option: Failover automatically to a second local system with replicated data, onsite within the branch, using RepliStor's automated switchover capabilities
Entire branch is lost	Data Recovery: <ul style="list-style-type: none"> Rebuild the branch systems Restore the data from the replicated copy, from the centralized site to the local site, through the RepliStor central management console If the size of the data is too large to send across the network OR if security is a concern: <ul style="list-style-type: none"> Rebuild the branch systems Restore the data locally from tape (or other removable media) shipped from the central office, using NetWorker Synchronize the data from the central site to the local site to recapture any new data written since the backup was made
	Automated Branch Recovery Option: Restart branch application services using the replicated dataset, at the central site or at an alternative branch office location, using RepliStor's automated switchover capabilities
Data corruption occurs at the branch	<ul style="list-style-type: none"> Use Windows Storage Server VSS capability onsite to rapidly recover the data Recover from backup at the central site using NetWorker, and then replicate the data down to the branch office using RepliStor

For detailed information on performing replication with RepliStor, consult the Operations Guide. Some details on NetWorker recovery operation are provided within this white paper; for more information, consult the NetWorker Operations Guide.

Recovering Data Backed Up with NetWorker to the Data Center

This section outlines the steps necessary to perform a basic recovery of data to the data center. For additional information, refer to the *NetWorker Administrators' Guide, Release 7.1, Microsoft Windows Version*.

You can designate the entire filesystem, specific directories, or individual files for recovery. Note that if you select a directory for recovery, all of its nested subdirectories and files are also selected for recovery. To exclude certain files from recovery after selecting an entire directory, you must unselect them.

Important: If VSS is licensed and enabled and you want to recover Microsoft Exchange Server or Microsoft SQL Server data, you must manually take the databases offline before beginning the recovery.

To recover data to the data center NetWin device:

1. In the NetWorker user program located on the data center NetWin device, click **Recover** to open the **Source Client** dialog box.
2. Select the source client whose data you want to recover (data center NetWin), and click **OK**. The local client is the default selection.
3. Select the destination client for the recovered data (data center NetWin), and click **OK**. The local client is the default selection.
4. In the left pane of the **Recover** window, click the appropriate directory folder.
5. Mark each directory and/or file to be recovered by performing one of the following:
Data Recovery
 - Select the directory or file, and click the **Mark** button. To clear an item, click the **Unmark** button.
 - Right-click the directory or file. To clear an item, right-click the directory or file again.

Note:

If you select **My Computer** for recovery, individual drives are not recovered. Only the SYSTEM or VSS SYSTEM savesets are recovered. You must select individual drives for recovery. When you mark a directory or file for recovery, a check mark appears next to that item.

To recover data from a scheduled backup other than the most recent one, you must change the browse time.

6. Click the **Start** button to begin the recovery.

It takes the NetWorker server a few moments to recover the files, depending on file size, network traffic, server load, disk properties, or tape positioning. During this time, messages appear in the NetWorker user program for each recovered file, so you can monitor the progress of the recovery.

Notes:

If an error occurs while recovering Microsoft Exchange Server or Microsoft SQL Server data, you must re-start the recovery process. If there is a naming conflict with the file you are recovering, the NetWorker server prompts you to choose one of several workarounds. If the recovery is successful, the NetWorker User program displays the following (or a similar) message: Received 1 file(S) from NSR server server Recover completion time: 5-30-03 2:03pm

If the recovery fails due to a problem with VSS or a writer, an error message will indicate this. Use the Windows Event Viewer to examine the event logs for additional information.

Frequently Asked Deployment Questions

This section provides answers to questions frequently asked during the planning process, as decisions are made about network connections, consolidation resources, and feasible technology options for each.

Q. I'm concerned that I might not have enough bandwidth from my branch offices to replicate my data. How can I know for sure?

The key metric that you need to measure to determine the network bandwidth that you require is the rate of change of the data that you need to replicate. Why? Because RepliStor replicates only the CHANGES in data as they occur. The easiest way to measure the rate of change of the data is to use Windows Performance Monitor.

On average, if you don't have time to measure, LEGATO has found that about 10 percent of data change is a good number to work with for planning purposes.

Q. If I have an existing network connection can I use it, or am I going to need a new network connection?

Depending on the size of data that you need to replicate, you'll need to plan the size of the network that you'll need in order to keep up with that replication. Most LEGATO customers find that they can use existing TCP/IP network connections, and there is no need to dedicate connections. LEGATO helps to enable shared use of network connections by providing key capabilities that include:

- Compression, which reduces the size of the data being replicated. Note that the compression ratio is dependent on the type of data being compressed and will vary.
- Throttling, which allows RepliStor to respond dynamically to changes in the available network bandwidth and preserve bandwidth for other types of business traffic.

Q. Can I use RepliStor's scheduling capability to only send replicated data at night using a synchronization, and what will this do to my level of network usage?

Yes, you can do this. However, this is generally NOT your most efficient option. Other options include:

- Schedule throttling during the day so it limits how much bandwidth it will use during the day, and then let it go full blast during the night.
- Schedule forwarding off during the day and forwarding on at night.

With the above schemes, the source drive must have enough disk space to hold the changed data during the day.

The main advantage is that it doesn't need to scan millions of files for changed data. It is typical for a large portion of a sync to go by without sending much data, because it must decide what files have changed. Note that depending on the file structure, the different levels of synchronization that RepliStor provides will offer different efficiencies:

- With millions of small files, it is usually more efficient to go ahead and do a full synchronization since comparing each file individually can easily add 25-50 percent overhead in time.
- With a smaller number of very large files, it is usually worthwhile to do an incremental synchronization since the time required to compare the smaller number of files is earned back by not having to send as much data over the network.

Q. Is there a maximum number of concurrent data streams, or "sessions," that I can replicate into a single target system?

No, there is no limit although there are some practical considerations to pay attention to during the planning process.

The first thing to know is that RepliStor is multi-threaded to handle multiple file operations at the same time, when replicating data at the centralized site. This enables it to better keep up with multiple data streams than alternative offerings. Do not confuse the maximum number of multi-threaded operations that RepliStor performs simultaneously (16) with being a limit on the number of data streams that RepliStor can receive (unlimited).

The target systems need to be powerful enough to handle the sum of what is being processed individually at each branch office. Multiple CPUs will help, as will adequate memory.

Q. What happens on the RepliStor receiving side, if the number of operation requests coming in at the same time exceeds the number that RepliStor can handle (16)?

RepliStor will receive them and queue them until RepliStor can perform the file operation. There is no problem.

Additional Resources

This section includes a list of product guides and white papers related to this solution. These guides and papers can be obtained from your EMC, LEGATO, or Microsoft representative.

<i>EMC NetWin 200 Version 1.2.0, System Administrator's Guide, 300-001-185, Rev A03</i>	This guide is intended for system administrators and contains information on managing and troubleshooting the EMC NetWin 200 storage system.
<i>LEGATO NetWorker Release 7.1 Multiplatform Version Disaster Recovery Guide</i>	This guide contains information on disaster preparation and recovery using the LEGATO NetWorker® software. This information is intended for system administrators who perform backup and recovery procedures, managers who want to learn how to implement a disaster-recovery program, and users who are responsible for implementing disaster-recovery plans and procedures.
<i>LEGATO License Manager 2nd Edition UNIX, Linux, and Microsoft Windows Version Installation and Administrator's Guide</i>	This guide provides information on how to configure and manage the LEGATO® License Manager software and is intended for system administrators and operators.
<i>LEGATO NetWorker Release 7.1 Microsoft Windows Version Administrator's Guide</i>	This guide provides information on how to configure and manage the LEGATO® NetWorker software on Microsoft Windows computers in a typical environment and is intended for system administrators and operators.
<i>LEGATO NetWorker, Microsoft Windows Version, Release Supplement</i>	This supplement provides additional information on new features, known limitations, and configuration tips obtained after the release of the software.
<i>LEGATO NetWorker Release 7.1 Microsoft Windows Version Installation Guide</i>	This guide provides instructions on how to install the LEGATO NetWorker® software.
<i>LEGATO RepliStor Release 5.3 Microsoft Windows Version Administrator's Guide</i>	This guide provides information for system administrators and operators on how to configure and manage LEGATO RepliStor software.
<i>LEGATO RepliStor Release 5.3 Microsoft Windows Version Installation Guide</i>	This guide provides information for system administrators and operators on how to install LEGATO RepliStor software.
<i>LEGATO RepliStor Release 5.3 Microsoft Windows Version Release Supplement</i>	This supplement provides additional information on new features, known limitations, and configuration tips obtained after the release of the software.
<i>LEGATO RepliStor Technical Manual November 7, 2003 Version 5.0.7</i>	This manual is for technical support personnel, SEs, OEMs, and customers. It supplies a good basic understanding of RepliStor, including configuration issues and troubleshooting techniques.
<i>Storage Management Using Windows Server 2003 and Windows Storage Server 2003 Virtual Disk Service and Volume Shadow Copy Service (Microsoft white paper available at http://www.microsoft.com)</i>	This white paper describes both the business benefits of VSS and VDS, and the technological details of the services. Section 1 introduces VDS and VSS services to the technological decision maker, and explains how the services help to reduce storage management complexity. Section 2 offers the technical details of the services, and is intended for storage administrators who need to understand how the services work to correctly deploy storage solutions.

