IBM TotalStorage N series systems with SnapMirror Software

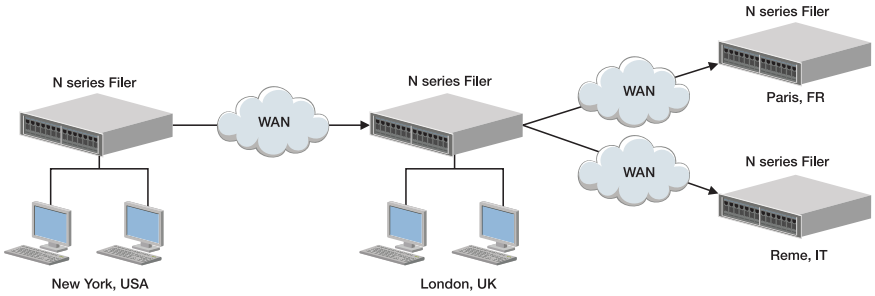

Figure 1. IBM TotalStorage N series with SnapMirror: cascaded configuration

Highlights

- ***Fast data replication and failover—can help reduce downtime in case of a failure at the primary site***
- ***Access to mirrored data—enables offloading tape backup, potentially increasing the value of your disaster recovery investment***
- ***Volume or Qtree replication—mirrors selected data sets, helping to reduce networking infrastructure requirements***

The challenge: Rapid access to mission-critical data

Today, global enterprises need to protect and quickly recover data in the event of natural or man-made disasters, operator errors, or technology and application failures. They also need an efficient way to distribute data to remote locations. Without an effective data protection and distribution strategy, operations can be brought to a standstill, resulting in significant lost revenue.

The solution: IBM TotalStorage N series systems with SnapMirror software

IBM TotalStorage® N series systems with SnapMirror® technology can help you implement the disaster recovery and data distribution solution that your enterprise needs. By replicating data at high speeds over a LAN or a WAN, SnapMirror technology can help support high data availability and quick recovery for mission-critical applications.

SnapMirror technology is designed to mirror data to one or more N series filers. It supports constant updates to the mirrored data to keep it current and available for disaster recovery, offloading tape backup, read-only data distribution, testing on non-production filers, online data migration, and more. If your enterprise is geographically dispersed and all locations need access to the same data set—such as training videos, CAD tools, and the like—SnapMirror can distribute the same data to all locations. By automatically updating this data and allowing local access to mirrored data, SnapMirror can help improve employee productivity and efficiency.

Preserves valuable network bandwidth

IBM TotalStorage N series systems with SnapMirror technology have many bandwidth-saving features that can help you lower the infrastructure cost of data replication and disaster recovery. You can perform an initial full-volume transfer using tapes, and then use the tapes to populate data in remote locations. After that, you need only update the new and changed blocks incrementally over the network. By replicating only a subset of the entire filer data set, SnapMirror technology can significantly reduce network bandwidth requirements. In addition, SnapMirror technology sets checkpoints during data transfers. If the system goes down, the transfer restarts from the most recent checkpoint. SnapMirror technology can also perform intelligent resynchronization, which virtually eliminates the need for full transfers when recovering from a broken mirror or loss of synchronization. If data on the mirrored copy was modified during application testing, it can be quickly resynchronized with the production data by copying the new and changed data blocks from the production system to the mirrored copy.

Configuration flexibility

IBM TotalStorage N series systems with SnapMirror technology can be deployed into any networking infrastructure with enough bandwidth to handle the data transfers. Support for multiple transports paths allows for greater use of existing equipment and better availability since you can failover between paths.

SnapMirror technology supports data protection by allowing you to choose the right level of synchronicity (sync, semi-sync, and async) between source and target copies. For instance, the sync option supports the replication of data at the remote site such that the remote data can be up-to-date and ready for use after a failure. This can facilitate disaster recovery efforts and reduce system downtime. Semi-sync allows you to determine the number of I/O operations or how long the replicated site can be out of sync with the source, based on their site needs. Or with async, schedule transfers whenever you want every minute, hour, or day. You can establish the frequency that works best for each site. The schedule can be easily modified and changes can be made effective immediately. You can also choose different filer configurations for the source and mirrored systems. The source system can even be a clustered filer with 6TB of storage mirroring 2TB of mission-critical data to a different filer model.

In addition, cascade and multihop mirroring let an IBM TotalStorage N series system with SnapMirror technology target volume serve as a source to other targets, with each mirror pair running on its own schedule to meet site-specific requirements. Cascade mirroring can help you replicate data over a distance—for instance from New York to Paris, Rome, and London. You can duplicate the New York data to London, and then use lower-cost links to replicate the data from London to Paris and Rome.

Filers are SnapMirror software-ready and need no added software installation. Built-in SNMP support facilitates integration with an SNMP framework.

Using SnapMirror, data can be efficiently and cost-effectively replicated to remote sites for disaster recovery or data distribution.

For more information

Contact your IBM representative or IBM Business Partner or visit:

ibm.com/totalstorage/nas

IBM Systems and Technology Group
5600 Cottle Road
San Jose, CA 95193
U.S.A.

Produced in the United States
July 2005

This document could include technical inaccuracies or typographical errors. IBM may make changes, improvements or alterations to the products, programs and services described in this document, including termination of such products, programs and services, at any time and without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. The information contained in this document is current as of the initial date of publication only and is subject to change without notice. IBM shall have no responsibility to update such information.

IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein. Performance data for IBM and non-IBM products and services contained in this document was derived under specific operating and environmental conditions. The actual results obtained by any party implementing such products or services will depend on a large number of factors specific to such party's operating environment and may vary significantly. IBM makes no representation that these results can be expected or obtained in any implementation of any such products or services.

MB, GB and TB equal 1,000,000, 1,000,000,000 and 1,000,000,000,000 bytes, respectively, where referring to storage capacity. Actual storage capacity will vary based upon many factors and may be less than stated. Some numbers given for storage capacities give capacity in native mode followed by capacity using data compression technology.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS-IS" WITHOUT ANY WARRANTY, EITHER EXPRESSED OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided.

References in this document to IBM products, programs or services do not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM program or product in this document is not intended to state or imply that only that program may be used. Any functionally equivalent program or product that does not infringe IBM's intellectual property rights may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

TSD01269-USEN-00