

APPLICATION NOTE

March 2006

Using the StorageTek™ 5000 Family of NAS Appliances for consolidation and data disaster recovery

ABSTRACT

The StorageTek™ 5000 Family of NAS Appliances allows your organization to implement an effective storage consolidation strategy and establish reliable data disaster recovery capabilities. This paper discusses how the 5000 NAS appliances running StorageTek File Replicator software can support near real-time mirroring of mission-critical data to a disaster recovery site.

- 1 Executive summary 2**
- 2 Benefits 2**
- 3 Data disaster recovery challenges. 3**
- 4 Data disaster recovery strategy. 3**
 - 4.1 Storage consolidation 3
 - 4.2 Replication and mirroring 4
 - 4.3 Mirror types 5
 - 4.3.1 One-to-one mirroring 5
 - 4.3.2 Hub mirroring 5
 - 4.3.3 Bidirectional mirroring 6
- 5 Recovering data availability 6**
- 6 Conclusion 6**

1 Executive summary

In the past, protecting your business from catastrophic data loss was an expensive and complex undertaking. The Sun StorageTek™ 5000 Family of NAS Appliances and add-on StorageTek File Replicator software create a reliable, affordable, and robust storage infrastructure that not only simplifies access, management, and replication of mission-critical data, but in the event of a catastrophic system failure, enables your organization to recover a near real-time mirror of your vital information.

When using 5000 NAS appliances and File Replicator software, the two key steps to successfully implementing a disaster recovery solution are:

- **Consolidating storage** — In this step, data is moved from network servers and individual workstations to a centralized NAS device, such as a 5000 NAS appliance. This allows your organization to streamline data management practices and enables heterogeneous systems to access data based on policies established by the system administrator.
- **Mirroring data** — Once data consolidation is complete, 5000 NAS appliances and File Replicator software copy and mirror all data from selected volumes to a remote disaster recovery (DR) site that is equipped with a mirror Sun NAS appliance. Using IP-based technology, the primary Sun NAS device sends the selected data to the mirrored DR site in near real time to help provide complete business continuity in the event of a catastrophic event.

This paper describes, in detail, the requirements necessary to design and deploy a DR solution using 5000 NAS appliances and File Replicator software. It also provides several examples of mirroring options based on a range of business needs.

This document's intended audience is system and storage administrators with a working knowledge of UNIX®, network bandwidth, NFS, CIFS/SMB, and the 5000 NAS appliance.

2 Benefits

The 5000 NAS appliances allow your organization to move away from storing mission-critical data on disparate servers and individual workstations. When using NAS devices, such as the 5000 NAS appliances, you can consolidate and store production data on a centralized storage device, making it easier to access, manage, replicate, and if necessary, restore vital information.

Once the storage consolidation is complete, the 5000 NAS appliances support the replication of file systems and near real-time volume mirroring using File Replicator software. File Replicator duplicates selected volumes to an alternate disaster recovery site. In the event of system failure, administrators can provide file-based services instantaneously by promoting a NAS mirror-based logical volume to production status using a simple-to-use Web interface. And, depending on the disaster recovery plan in place, administrators can promote the mirror-based logical volume without reconfiguring network servers or interrupting the end user's workflow.

In addition, production applications can access legacy data from point-in-time snapshots (checkpoints) that have been archived based upon an administrator's predetermined policies before, during, and after a disaster recovery event.

The 5000 NAS appliances running File Replicator software comprise an affordable data disaster recovery (DDR) solution to help your organization manage, replicate, and restore business-critical information, such as collaborative, Home Directory, and compliance-related data.

The key elements of data disaster recovery are the preservation of the computing infrastructure and the ability to provide up-to-date versions of mission-critical data.

3 Data disaster recovery challenges

Disaster recovery is an implementation that is often cost-prohibitive and reserved only for Fortune 1000 data centers. Many organizations expend large amounts of budget so that their computing and storage infrastructures will be operational and available after a flood, fire, earthquake, or other catastrophic event.

At the highest level and scope, the two essential elements of disaster recovery are:

- Preserving and duplicating the computing infrastructure
- Protecting and providing up-to-date versions of mission-critical, if not all, stored data and content

The basic concept of DR is that your organization can provide essential, minimal IT services until production capability has been completely restored either at the primary or backup site. Typically, with comprehensive disaster recovery planning, this minimalist approach can work for production applications. But, more often than not, complete and up-to-the-transaction-level data recovery is required for a business to survive a catastrophic event.

4 Data disaster recovery strategy

The easiest approach to DDR is to duplicate a minimal computing infrastructure at an alternate site as well as provide up-to-date data that is ready for use by production applications in case of a disaster scenario. First, however, steps must be taken to minimize the cost of a disaster recovery implementation.

The approach to DR for the computational infrastructure is far different from the storage infrastructure. It is much easier to duplicate processing capability than to replicate up-to-date — and usable — mission-critical data. As such, your organization must separate these infrastructure elements in a survivable DR implementation. Undertaking a storage consolidation effort must be the first step in this process; when production data is separated from the computing infrastructure, less dependency exists between the two elements. This separation allows the data to become, essentially, interoperable, allowing the computing DR infrastructure the flexibility it needs to minimize DR costs. It also simplifies the design required to provide a DR application-sensitive infrastructure.

Storage consolidation using a purpose-built storage infrastructure is the first step in implementing an effective disaster recovery strategy.

Once the data has been moved to a purpose-built storage infrastructure, duplicating the computing ability is a relatively straightforward process.

Mission-critical data can be handled separately with a variety of toolsets available. The 5000 NAS appliances provide storage consolidation tools as well as reliable DDR capabilities that can be deployed in a DR implementation.

4.1 Storage consolidation

The 5000 NAS appliances have many purpose-built capabilities that can be used to easily facilitate a storage consolidation effort. The built-in capacity to provide a network storage-based home directory (such as `autohome`) for Microsoft Windows, Solaris™ Operating System, and UNIX environments is a simple and effective way of implementing a storage consolidation effort. Market analysts have estimated that as much as 60 percent of mission critical-related data found in any computing infrastructure is stored in a distributed paradigm: on desktop, laptop, or other workstation computing devices.

Moving home directory data to a centralized NAS device is easy to implement in most modern operating systems and is a critical step in protecting your end users' vital data.

Moving all home directory data to a centralized NAS device such as a 5000 NAS appliance is essential to protecting this data. With a storage consolidation architecture in place, the File Replicator software can provide not only replication capabilities, which duplicate complete file systems in place, but also provide an up-to-the-transaction-level mirror of these same file systems. There are server-based replication products available, but rarely do they provide up-to-date data coherency without a large budget and bandwidth expenditure.

4.2 Replication and mirroring

One objective of a storage consolidation effort is to provide a storage capability that allows for replicating mission-critical data to an alternate DR site. File Replicator software is an add-on to the 5000 NAS appliance that provides two capabilities. First, it replicates all data stored on selected storage volumes to a DR site using IP-based technology, and second, it provides for an up-to-date, near real-time mirroring capability. During the replication phase, a mirror journal is allocated based on administrator input, and an exact duplicate of the logical volume is created on the mirror Sun NAS appliance. During the engineering phase of a disaster recovery project, the IT engineering organization should take great care so that enough bandwidth is provisioned to support replication as well as the follow-on mirroring process. Important data points to consider are logical volume size, amount of data on the logical volume, and end-to-end network infrastructure as well as the line bandwidth shown in Table 1 below.

Table 1. Expected Leased Line Bandwidth Table.

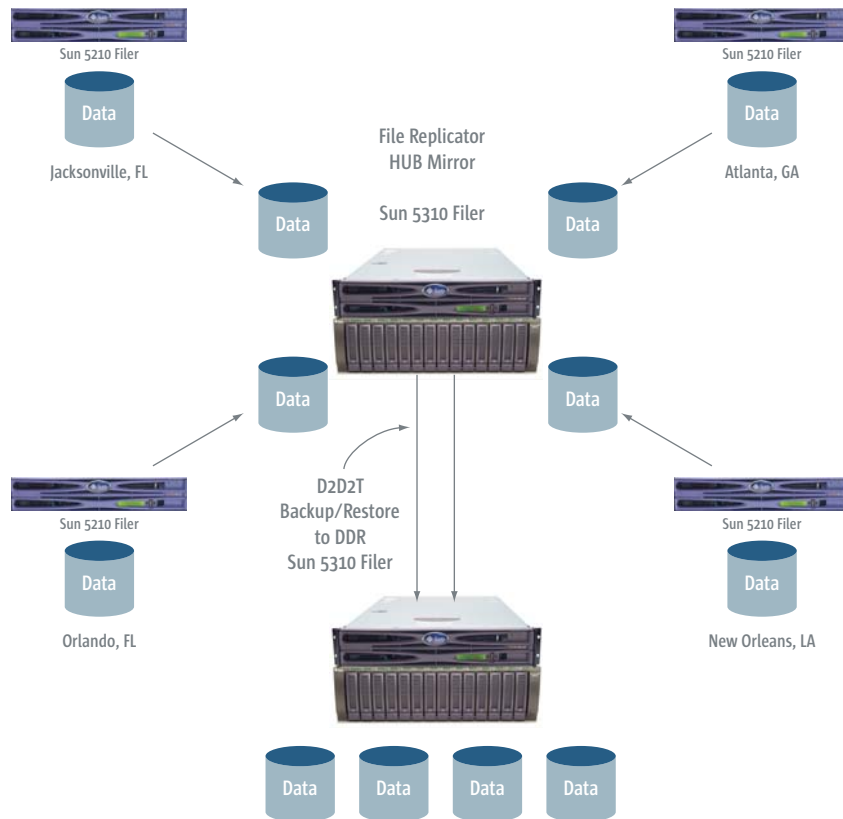
Typical leased line	Expected bandwidth
T1	1-5 MB/sec
T3	4-7 MB/sec
OC3	19 MB/sec

Once the replication of the logical volume is complete, the 5000 NAS appliance, using File Replicator, extends the journal component of the file system such that all input/output (I/O) write transactions are also extended asynchronously in real time to the 5000 NAS appliance that is acting as the NAS Mirror Filer.

The 5000 NAS appliances replicate all data stored on selected storage volumes to a DR site using IP-based technology to provide near real-time mirroring capability.

Before implementing this solution, confirm that the write I/O workload taking place on the primary NAS Filer can also be supported by the bandwidth that is provisioned between the primary NAS Filer and the Mirror NAS Filer. Once the I/O workload is properly configured and a DDR capability deployed, the organization will benefit from the storage consolidation phase in which all Home Directory, eligible mission-critical, and collaborative data can be replicated and kept up-to-date at the DR site by utilizing the File Replicator feature on the 5000 NAS appliances.

Figure 1. File Replicator Hub Mirror.



Note — For a mirrored file system to remain coherent, LAN, campus WAN, or leased line bandwidth must be able to support the overall update rate that is occurring on the primary NAS Filer.

4.3 Mirror Types

Three types of mirroring are of interest in DR efforts, as described below.

4.3.1 One-to-one mirroring

In its simplest form, a DR implementation can provide for one-to-one replication and mirroring capability. Thus, a 5000 NAS appliance can replicate and mirror data to another 5000 NAS appliance. Note that multiple one-to-one mirror instances of diverse logical volumes can take place between 5000 NAS appliances.

4.3.2 Hub mirroring

Often organizations use a centralized repository for backup/restore data as well as disaster recovery data. In these scenarios, as illustrated in Figure 1, many edge sites use 5000 NAS Filers at remote locations and, with File Replicator, duplicate and keep near real-time mirrors of file systems at a central site. This affords the flexibility of the central site functioning in a recovery scenario if one or more remote 5000 NAS Filers experience disaster scenarios, as well as providing up-to-date, multiple, versioned mirrors of mission-critical data on demand. Thus, production applications can reference not only up-to-date data, but also legacy data from previous point-in-time snapshots (checkpoints) that have been mirrored over time. This is the most powerful and comprehensive File Replicator type of deployment.

4.3.3 Bidirectional mirroring

Often DR sites are not only used in a passive sense, awaiting a catastrophic event. They are also used as additional computing and storage capacity for day-to-day operations. As such, bidirectional mirroring is a method commonly deployed that protects data at both sites while providing near real-time capability.

In the event of a system failure, an administrator can promote a mirrored NAS Filer to production status with the click of a mouse.

5 Recovering data availability

When and if a DR event occurs, the 5000 NAS appliance designated as the Mirror Filer can easily be promoted to production status. By a simple mouse-click in the Web-based NAS OS interface, the Mirror NAS Filer makes all mirrored data available for real-time use without any adjustment by users or servers that are connected to the data in the infrastructure.

6 Conclusion

A catastrophic loss of data can threaten the future of your organization. The StorageTek 5000 Family of NAS Appliances and File Replicator software work in tandem to consolidate your mission-critical data and mirror, in near real time, all data from selected volumes to a remote disaster recovery site. Implementing 5000 NAS appliances and File Replicator software helps keep your mission-critical data available and up-to-date in the event of a major system failure.

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN **Web** sun.com

© 2006 Sun Microsystems, Inc. All rights reserved. Sun, Sun Microsystems, the Sun logo, StorageTek, the StorageTek logo, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

NT 0003 A 03/06

