

**APPLICATION NOTE**

March 2006

## Using the StorageTek™ 5000 Family of NAS Appliances in Web serving environments

**ABSTRACT**

This paper describes how administrators can use the StorageTek™ 5000 Family of NAS Appliances in Web serving environments. Specifically, it outlines how to avoid the duplication of processing and storage resources that often result when IT organizations are forced to quickly implement expanded Web site services.

- 1.0 Executive summary ..... 2
- 2.0 Web site deployment ..... 2
- 3.0 NAS storage consolidation ..... 3
  - 3.1 Simplified storage ..... 4
  - 3.2 Data protection and availability ..... 4
  - 3.3 Cost savings ..... 4
- 4.0 Conclusion ..... 5

### 1.0 Executive summary

This paper describes the advantages of implementing a network-attached storage (NAS) consolidation architecture in a Web serving environment. IT organizations, when initially implementing a Web serving environment, often deploy multiple Web servers as quickly as possible due to time-to-market pressures, requiring duplicated processing and storage resources. This causes excessive system, application, and storage administration overhead as well as significant deployment risk.

These penalties against total cost of ownership (TCO) and return on investment (ROI) can be easily avoided by implementing a storage consolidation architecture utilizing the Sun StorageTek™ 5000 Family of NAS Appliances, which provide a simple, fast, reliable, and affordable implementation of NAS.


This document is intended for system, Web application, and/or storage administrators; it requires a basic working knowledge of UNIX®, NFS, Microsoft Windows, CIFS/SMB, RAID, Web application deployment, and the 5000 NAS appliance.

### 2.0 Web site deployment

Historically, Web serving environments have ranged from small “closet” data centers to large enterprise data centers. Web serving environments have many things in common. Typically, they start out on purpose-built servers configured to provide key elements of the Web site for the site’s development. Often, each Web server is initially deployed with direct-attached storage (DAS) or sometimes on a more expensive storage area network (SAN).

As usage of the Web site grows, administrators add Web servers to the computing infrastructure along with more attached storage, duplicating data from the original Web servers. Network firewall appliances protect the Web site as well as distribute traffic to the multiple Web servers with duplicated content.

Figure 1. Web serving and islands of data. Working under market pressure, Web administrators may have to add duplicate servers to handle increased Web traffic. NAS can eliminate that duplication.



This approach to addressing exponential growth in Web transactions and storage presents problems in two different areas:

First, “islands of data,” as illustrated in Figure 1, are introduced into the Web site. With duplicative servers and storage, each time the Web site is updated or changed, all content must be copied exactly to preserve consistency among the servers as well as the appearance of the Web site to the public. This increases the probability of error during Web site cut-over because administrators must duplicate content as well as any programmatic materials employed in the Web site. Any errors during distribution of the updated Web site introduce risk of Web site downtime.

Second, from an architectural perspective, the TCO of the Web service begins to increase because the infrastructure is being duplicated to handle the volume of traffic required to service Web site customers.

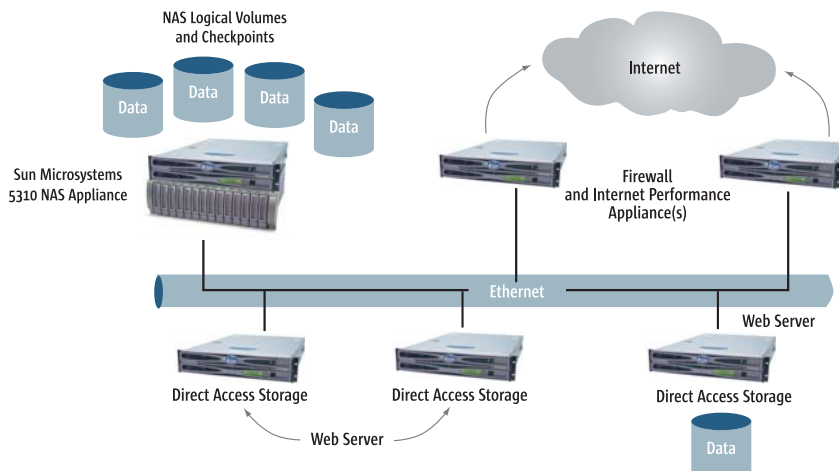
Adding more Web servers may not even address the performance requirements of the Web site itself. The actual effectiveness of the Web servers depends not only on the Web servers’ processing performance, but also on their ability to perform disk I/O. In some cases, Web servers may have computational headroom available, but because of the overall bandwidth limitations they are not able to deliver the amount of DAS/SAN storage required by the Web site. As a result, some Web servers may be taken out of service prematurely even though they are not fully utilized, thereby increasing overall TCO. The inability to fully use the capacity of the Web server can also be caused by operating system constraints. The OS may limit I/O capability architecturally through the lack of HBA slots and/or LUN size, which in turn affects the performance in the Web site.

### 3.0 NAS storage consolidation

The alternative is to implement Web servers (as demand requires) while consolidating content. The components of the solution, shown in Figure 2, include:

- One or more 5000 NAS appliance, depending on service levels required to maintain Web site performance and availability.
- One or more Sun Microsystems Web servers to provide the Web site with superior performance and reduced TCO.

Figure 2. Web serving and storage consolidation.



### 3.1 Simplified storage

Centrally locating data offers several advantages: Content and storage management are simplified by eliminating redundant content. Additionally, data can be organized using the built-in logical volume manager (LVM) on logical volumes and made available to each Web server as it is added to accommodate growth.

This logical approach allows data to be shared between the multiple Web servers, greatly simplifying the content management task.

### 3.2 Data protection and availability

The 5000 NAS appliances offer a checkpoint feature that can be used to create logical copies of the entire Web site. This allows Web site updates and changes to be made without interfering with Web service operations.

The 5000 NAS appliances also provide centralized configuration management of all Web site elements. The built-in LVM in the Sun StorageTek NAS OS allows the Web site administrator to allocate additional volumes for Web site test and development purposes. After testing is complete, the administrator can use the simple Web graphical user interface (GUI), which is based on Java™ technology, to rename these volumes to replace the current production Web site volumes. If it is necessary to return the Web site to its former state, this same procedure can be performed in reverse to recover the Web site almost instantaneously.

Administrators can easily back up the Web site using the 5000 NAS appliance's GUI. The checkpoint feature provides a logical backup of the Web site, enabling instantaneous restore in case of accidental data erasure or erroneous update. Physical backup of the Web site can then take place using these checkpoints as a point-in-time copy.

Backup tools typically require some sort of system "snapshot," a need met by checkpoints. The 5000 NAS appliance provides backup/restore capabilities in concert with commercially available industry-standard tools. The 5000 NAS appliance embeds an NDMP backup/restore agent in the NAS OS at no charge. This agent is certified to function with these commercially available backup/restore tools.

### 3.3 Cost savings

Web server consolidation on NAS decreases TCO through the simplified administration and elimination of redundant content while also allowing for non-intrusive Web server upgrades.

At the outset, eliminating the need for duplicative storage for each of the Web servers decreases TCO. NAS storage and data content are shared among all of the production Web servers, as well as any development and test Web servers used in the ongoing evolution of the organization's Web site. Labor costs are significantly decreased during ongoing Web site development because a duplicate storage infrastructure does not have to be maintained. Development and testing can take place on logical volumes allocated and prepared for production usage on the 5000 NAS appliance.

As development and testing continue, administrators can use the checkpoint feature to "versionize" the Web site and establish milestones prior to cut-over to production. This increases the efficiency and accuracy of Web site testing and shakedown prior to pushing newly developed content into production. Finally, this approach to Web site deployment significantly decreases the probability of error during production cut-over, thereby eliminating risk and costs associated with Web site downtime.

*The 5000 NAS appliance checkpoint feature allows administrators to quickly make logical copies of entire sites for backup, testing, or update purposes.*

*The 5000 NAS appliance works in conjunction with industry-standard backup and recovery tools.*

#### 4.0 Conclusion

Implementing a NAS-based storage consolidation architecture utilizing the 5000 NAS appliance simplifies IT operations and maintenance while reducing TCO. By not having to maintain multiple copies of the same content to accommodate growing Web site traffic, management overhead is greatly reduced. Centralized resource management means better utilization of existing resources and less administrative overhead trying to align data across distributed servers.

Other benefits of storage consolidation in a Web serving environment range from enabling enhanced development and testing to reducing the risk of Web site downtime. Consolidating on the StorageTek 5000 Family of NAS Appliances means that whatever Web services are being delivered, the infrastructure cost will not outweigh the business benefits. The scalability and flexibility of the Sun NAS architecture means data can grow without disruptive additions of servers or software. Businesses can therefore focus on delivering business value rather than worrying about bandwidth and capacity.

**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 Java 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 0004 A 03/06