A large, abstract graphic on the left side of the page, consisting of overlapping, semi-transparent, curved shapes in shades of gray, creating a sense of depth and movement.

Sun StorEdge™
**5210/5310 NAS HOME
DIRECTORY STORAGE
CONSOLIDATION**

White Paper
November 2005

Table of Contents

Introduction	2
The Distributed Desktop Workstation and Server Architecture.	3
Network Attached Storage Consolidation	4
Benefits to NAS based Home Directory Storage Consolidation.	4
Home Directory Data Connectivity	4
Logical Volume Management	4
General File Sharing	4
Data Protection	5
Data Integrity	6
Information Technology Management	6
Backup/Restore and Disaster Recovery	7
Return on Investment and Reduced Total Cost of Ownership	7
Conclusion	9
Additional Information	9

Introduction

This paper provides an approach to implementing network attached storage (NAS) consolidation with respect to Home Directories. Information Technology (IT) environments require home directories so that users can store their data in a location that is ideal for easy and fast access. During the first wave of distributed computing, it was standard practice to assign home directories to each individual desktop and/or server on their associated direct access storage devices (DAS). The result of this decision was that protecting and preserving data availability became an often insurmountable problem due to the difficulty and expense associated with implementing a successful backup, restore, and disaster recovery plan due to the distributed nature of such a storage implementation.

This document provides a framework for implementing and consolidating home directories on a Sun StorEdge™ 5000 NAS Appliance. The Sun StorEdge 5000 product family are a simple, fast, reliable and affordable implementation of a NAS appliance(s).

This document's intended audience is system and/or storage administrators, and requires a basic working knowledge of UNIX®, NFS, Microsoft Windows, CIFS/SMB, RAID and the Sun StorEdge 5000 NAS Appliance.

The Distributed Desktop Workstation and Server Architecture

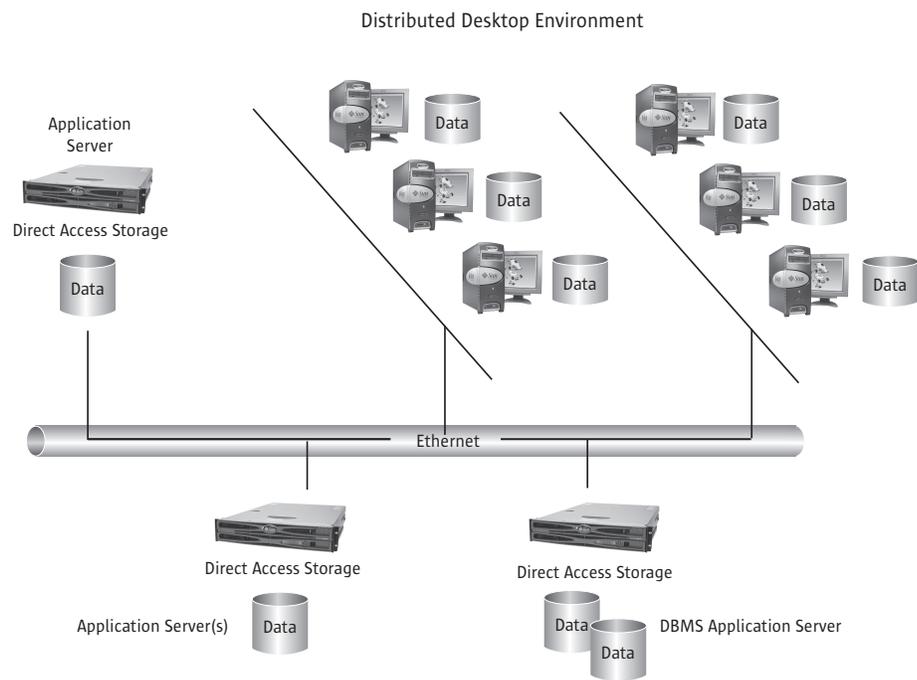
Desktop workstations typically do not have data protection capabilities like RAID or tape backup which are present in typical storage environments. This is because it is cost prohibitive to implement these types of data protection capabilities in a distributed workstation and server environment. Labor, backup/restore tooling and software licensing and other hidden system administration costs tend to increase exponentially when implementing these types of capabilities in a distributed environment.

Nevertheless, preservation of data when a particular desktop client fails or loses data still requires a nightly backup of all pertinent data on every workstation in the computing and workstation infrastructure. Typically, in an organization's infrastructure, 50-60% of all relevant mission related data is stored on the desktop with the balance of mission critical data being stored on a typical centralized server computing storage infrastructure such as direct access storage (DAS) or a storage area network (SAN).

Finally, during the architectural life cycle of a typical desktop device or workstation, its ability to provide storage capacity is usually exceeded. When this occurs a replacement desktop device or workstation is usually required in order to provide more storage capacity thus reducing the effective life of the workstation unnecessarily.

The combination of all of these downside attributes increases the total cost of ownership (TCO) of a distributed storage architecture such that an alternative storage architecture is necessary.

Figure 1: Distributed Computing and Storage Architecture prior to NAS Storage Consolidation



Network Attached Storage Consolidation

The alternative is to implement an architecture where all desktop and server components can utilize a common consolidated pool of storage. This is accomplished easily and simply on the Sun StorEdge 5000 NAS Appliance utilizing the Home Directory feature ('autohome'). This feature leverages the existing capabilities of the Microsoft Windows desktop environment often called 'Home Directory(s)' as well as the typical feature utilized in Solaris™ and other UNIX environments called 'Automount'. Deploying the Home Directory capability facilitates the consolidation of all distributed storage from the desktop and server infrastructure to a centralized data content repository that resides on the Sun StorEdge 5000 NAS Appliance.

Benefits to NAS based Home Directory Storage Consolidation

Home Directory Data Connectivity

After data and content are moved to the Sun StorEdge 5000 NAS Appliance, Microsoft Windows based 'shares' are allocated that relate each desktop user to their corresponding data and content. In addition, Solaris and/or UNIX based NFS Mounts are allocated utilizing the Automount feature to connect these users and servers to their corresponding data and content. This facilitates automatic connectivity of a user to their corresponding data and content when they login to the computing domain. The WEB based GUI on the Sun StorEdge 5000 NAS Appliance provides for easy allocation and administration of these Home Directories. This is convenient and allows for substantial benefit to the organization's computing and storage infrastructure and administration.

Logical Volume Management

Included with the Sun StorEdge 5000 NAS Appliance is a logical volume manager which is employed to create logical volumes with file systems and subdirectories that are, as mentioned earlier, exported intelligently to associated desktop and server elements in the infrastructure.

The Home Directory feature of the Sun StorEdge 5000 NAS Appliance was provided utilizing the easy to use Java based graphical user interface (GUI) to allocate a logical volume with its corresponding filesystem and home directory account subdirectories. This filesystem essentially contains a root level directory which in turn provides a subdirectory that is the corresponding Home Directory for all desktop and server elements in the storage architecture.

General File Sharing

The Sun StorEdge 5000 NAS Appliance can also be used to provide a general file sharing capability. Typical distributed desktop workstation and server architectures often require a centralized storage medium that can be used to share data and content among users, projects and certain mission critical applications. Again, the logical volume manager included with the Sun StorEdge NAS Operating System (NAS OS) can be used to allocate a volume for file sharing and general purpose storage. This filesystem then in turn can be suballocated to each group of users, projects or applications and can be secured with the industry standard security tooling included with the NAS Filer. The NAS OS includes support for Active Directory, Kerberos as well as typical UNIX based security paradigms and allows for sharing of data among different OS architectures.

Data Protection

As mentioned earlier, typical DAS environments which are implemented in a distributed desktop environment have little or no data protection capabilities. Also, system administration costs associated with providing these capabilities is too expensive and labor intensive to practically accomplish.

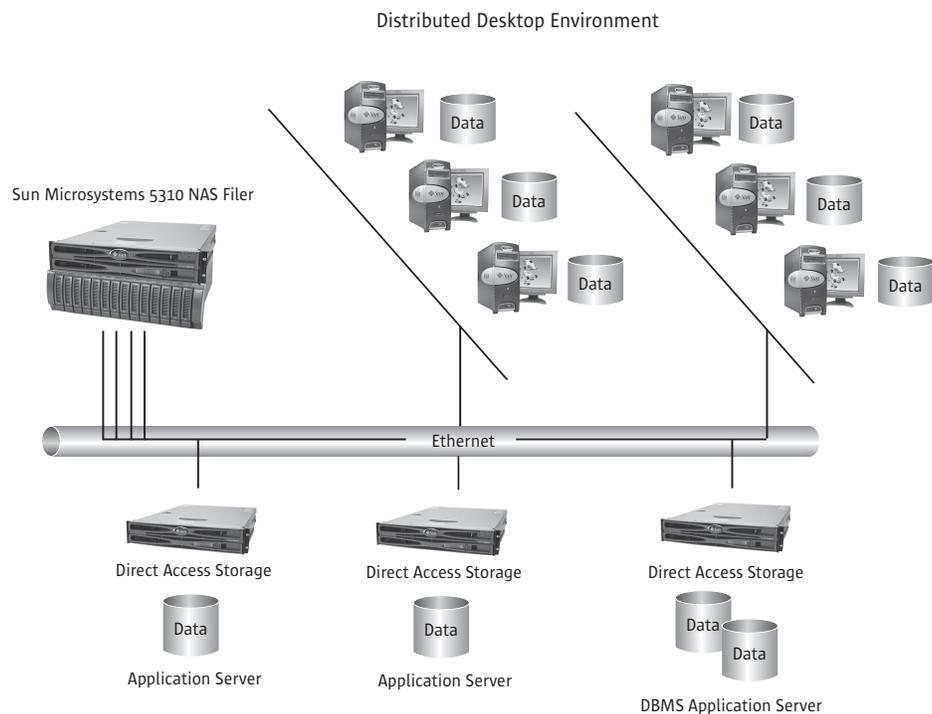
When NAS based Home Directory Storage Consolidation is accomplished, all data and content is protected on the Sun StorEdge 5000 NAS Appliance. The NAS Appliance is equipped with Hardware RAID which is typically deployed as RAID-5 or Striped Distributed Parity. This RAID architecture provides two(2) copies of the data, one which is the primary copy and the 2nd of which is a logical copy which can be reproduced after a disk failure utilizing the remaining disks in the RAID set as well as the distributed parity.

In addition, the Sun StorEdge 5000 NAS Appliance has what is called a Checkpoint feature. This feature provides a logical point in time copy of data and content when implemented by the system or storage administrator. It can be accomplished manually or on a scheduled basis. Currently, the Sun StorEdge 5000 NAS Appliance provides up to 64 Checkpoints per logical volume and as well as an unlimited number of checkpoints per NAS Filer and is only constrained by the available amount off storage. Uniquely, the NAS Filer does not require storage to be reserved for Checkpoints but rather uses storage from the filesystem from which it is providing a logical point in time copy.

Note – The terms ‘snapshot’ and ‘checkpoint’ are used synonymously.

Thus, the system or storage administrator will be relieved of concerns regarding a shortage of reserved allocated snapshot capacity.

Figure 2: Distributed Computing and Storage Architecture after NAS Storage Consolidation



When data or content is accidentally lost, erased or corrupted by an application, either the system/storage administrator or the user depending on organizational policy can restore the data or content instantaneously.

These capabilities provide for a level of data protection not typically found in a distributed desktop workstation environment.

Data Integrity

Along with data protection the Sun StorEdge 5000 NAS Appliance provides for exceptional data integrity capabilities by implementing a high performance journaling filesystem (JFS).

A unique feature of the Sun StorEdge NAS OS JFS is that it is implemented in an internal architecture where it is protected from data loss. Typical JFS implementations store journal data and meta-data on RAM, NVRAM or DAS all of which do not provide for power loss and subsequent possible loss of data due to data corruption.

The NAS JFS deployment on the Sun StorEdge 5000 NAS Appliance utilizes exceptional data protection measures in order to provide for data integrity and preserve performance. The filesystem journal is allocated on a RAID-5 protected filesystem, with its associated mechanisms for data protection as well as RAID Controllers (RC) that include one(1) Gbyte of mirrored cache each. If elements of the Input-Output(I-O) infrastructure fail, data integrity is preserved via RAID based parity and mirrored cache. These elements facilitate successful failover and recovery when the Sun StorEdge 5000 NAS Appliance is deployed in either a High Availability or No Single Point of Failure configuration thus preserving data integrity. Thus to preserve data integrity, the Sun StorEdge 5000 NAS Appliance provide for data integrity, data protection as well as data availability all of which are required for mission critical applications.

Information Technology Management

Hidden costs associated with maintaining a mission critical computing application and storage infrastructure can be challenging in a distributed desktop workstation environment. Hidden costs can include excessive system and storage administration labor costs as well as costs associated with evolving mission critical application requirements.

System and Storage Administration labor necessary to support a distributed desktop and server architecture can be arduously time consuming and very expensive. A large portion of administrative labor is associated with providing data protection, data availability, and quick recoverability. Consolidating all production data onto a single Sun StorEdge™ 5210 or 5310 NAS Appliance allows system administration to be accomplished quickly and easily with dramatically less administrative overhead.

Mission critical applications often increase in capability, which in turn increases computational and storage requirements. The desktop and server infrastructure must be able to quickly and seamlessly adapt to changing dynamic requirements over time. Sun StorEdge 5000 NAS Appliances enable cost savings in the face of exponential storage growth by facilitating 'on the fly' storage growth by providing a feature that expands a production logical volume without incurring an outage and by the very nature of storage consolidation in that it separates data from the mission critical application platforms. This allows the desktop and server infrastructure to evolve without placing mission critical data and content at risk during application and infrastructure upgrades.

Backup/Restore and Disaster Recovery

The daily, weekly and monthly backup procedures are a chore that challenges about every data center today. The exponential growth of data and the associated storage infrastructure that has accommodated this growth requires special consideration to facilitate the daily task of protecting data.

As mentioned earlier, Checkpoints provide a logical backup of data and content by providing a logical point in time copy of a logical volume and its associated filesystem. It is important to note that this is not a 2nd physical copy of the data and thus should not be considered a permanent backup of mission critical data and content. Backup and Restore of data though, is facilitated using the Checkpoint feature of the Sun StorEdge 5000 NAS Appliance. This is because in order to provide a coherent backup of data, the filesystem must be static at the time of the backup. Checkpoints provide this critical element to the backup procedure. Typically, the backup tooling requires this and the Sun StorEdge 5000 NAS Appliance provides this feature when working in concert with commercially available industry standard backup/restore tools such as Veritas NetBackup and Backup Exec, Legato Networker and Bakbone Netvault. The Sun StorEdge 5000 NAS Appliance includes an NDMP Backup/Restore agent at no charge and is imbedded in the NAS OS of the Filer. This agent is certified to function with the aforementioned commercially available backup/restore tooling.

Disaster Recovery has become extremely visible and important in today's mission critical data center. The Sun StorEdge 5000 NAS Appliance includes a capability that provides replication as well as near real-time logical volume/filesystem mirroring. This means that the Sun StorEdge 5000 NAS Appliance can replicate and mirror a production filesystem to another Sun StorEdge 5000 NAS Appliance. One-to-One, Many-to-One and Bi-Directional Replication and Mirroring are supported. Easy volume mirror allocation, promotion and recovery are provided via the Web based GUI. This capability provides a relatively inexpensive way to protect data in the event of a disaster that affects data center availability by providing an up to the transaction level mirror of the NAS based filesystem.

Note – Mirror coherency and performance are highly dependent on network connectivity.

Return on Investment and Reduced Total Cost of Ownership

Typically, ROI and TCO are very difficult to measure. Statements regarding expected ROI/TCO results before an organization makes an investment often are not realized because the difficulty in measurement of the metrics associated with them.

When implementing NAS based storage consolidation real ROI and TCO can be realized. As mentioned earlier, expensive desktop computing is often retired early only because the DAS associated with the asset is not capable of supporting the requirements of the application over time. The DAS on each of the desktops and servers is retained, but only for purposes of warehousing the operating system and any other localized required software tooling. The implementation of Home Directories and Generalized File Sharing remove the application storage requirement from the desktop, thus extending the life of the desktop workstation and postponing the costs associated with its replacement.

In addition, often costly application servers serve double duty on a daily basis not only providing generalized file sharing storage but also supporting production applications. By removing this storage and computing overhead from the production application server, the purchase of additional computing and storage resources is eliminated and the effective life of the server has been extended.

Server based DAS or SAN resources are often very expensive to manage. Placing Home Directory and general file server data on this resource is counterproductive and costly over time. Often, mission critical applications require working storage to execute on a daily basis and usually cannot find this critical resource thus wasted computing cycles and time. Since Home Directory data is moved to the Sun StorEdge™ 5000 NAS Appliance, critical application storage resources are now available to facilitate daily mission critical application execution. This postpones the purchase of additional storage for the consolidated server infrastructure providing real savings to the organization.

Conclusion

This document provided information regarding implementing a NAS based storage consolidation paradigm utilizing the Sun StorEdge 5000 NAS Appliance. In addition, it discussed the benefits to deploying this type of storage architecture. Please feel free to contact Sun Microsystems, Inc. to inquire about the Sun StorEdge 5000 NAS Appliance and how it can help you consolidate storage, ease administration and reduce your TCO.

Additional Information:

For more information, refer to the following URLs:

<http://www.sun.com/storage/5000>

<http://docs.sun.com/app/docs/prod/storedge#hic>

Notes

© 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054 USA

All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California.

Sun, Sun Microsystems, the Sun logo, StorEdge and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a). DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS HELD TO BE LEGALLY INVALID.