hp storage

july 2002

SAN
business blueprint

NAS on SAN – universal networked storage

executive summary

IT managers are facing ever-increasing quantities of data that need to be stored and managed. Networked storage in the form of Storage Area Networks (SANs) and Network Attached Storage (NAS) is widely accepted as the solution to these issues. Unfortunately, competing vendors have positioned NAS and SAN as either/or technologies, when in reality they are complimentary.

Hp is a strong proponent of merging NAS and SAN - we're calling this Universal Networked Storage. As a first proof-point of this vision, this Solution Blueprint explains how to add the NAS8000 to any of hp's existing blueprinted SAN configurations.

Introduction

This document is one of a series of Solution Blueprints from HP. Storage Solution Blueprinting is HP's concept of defining a configuration for a specific storage problem, and providing all the information necessary to implement. A blueprint represents a fully tested and supportable configuration, which is orderable as a set of individual components from HP's standard price list. In some instances HP has tested third party products required to complete the described solution. Recognizing that one size does not fit all, guidance on flexibility and scalability is given. Standard product support is provided for each HP component in a blueprint configuration, although custom implementation and support services are also available. Third-party products are supported by the respective vendor.

The blueprints are split into two levels – business and technical. The business blueprint is intended as a first-level guide that defines a storage problem, demonstrates what type of storage solution can solve that problem, and recommends an HP configuration. The technical blueprint provides the next step by giving detailed schematics and part lists to enable implementation of the recommended configuration.

Time to solution

Hp's solution blueprints enable users to build a more competitive storage infrastructure, in a quicker time, and with less risk. Blueprints reduce time-to-solution, and uncertainty, by offering pre-tested configurations to specific storage problems. Unlike traditional "white papers", blueprints actually give the detail on what and how to implement, and unlike fixed solution bundles, blueprints offer scalability/flexibility to meet specific needs.

Depth of experience

Storage solutions can be complex, and comprise more than simply a collection of components. Channel partners and end-users need the backing of an experienced vendor. The new hp is the leading supplier of SAN attached disk arrays[1] with over 55% of shipments in 2001.

Fewer vendor relationships

Ultimately any IT infrastructure solution is only as good as its component part, and frequently end-users need to source from multiple vendors to get all components. Hp uniquely offers the most complete solution, ranging from a storage

[1] Gartner Dataquest; 2001 RAID-Based Disk Storage Worldwide Market Share & Forecast. Combined unit market share figures for pre-merge hp and pre-merge Compaq

portfolio offering choice and flexibility that is second to none, to servers, networking, and a complete life cycle of services. In the long term, managing fewer vendor relationships instead of many leads to lower costs for the end-user. And by choosing hp, you can feel confident you'll get best-of-breed components with roadmaps offering investment protection.

The Networked Storage dilemma/the SAN & NAS dilemma

Storage management faces the pressures of exponential growth in capacity coupled with declining, or at best stagnant, IT budgets. Combine this challenge with the fact that businesses are increasingly reliant on the data in their on-line (e.g. order processing) and back-office systems (e.g. ERP/CRM), and we see IT managers fighting to provide reliable access to ever increasing quantities of critical data with shrinking resources. IT Managers are saying:

"I don't know what my storage environment looks like."
"I need to assign storage automatically, quickly and efficiently without interruptions."
"I need to know ahead of time if I am going to run out of space."
"I don't know if I'm efficiently using the storage I already have."
"I want to identify where my bottlenecks are in order to remove them."

Networked storage is widely accepted as the solution to these problems. By networked storage we mean providing multiple hosts (servers) with access to a shared pool of disk and/or tape storage resources. This is in contrast to traditional Direct Attached Storage (DAS) approaches where each server has its own dedicated storage.

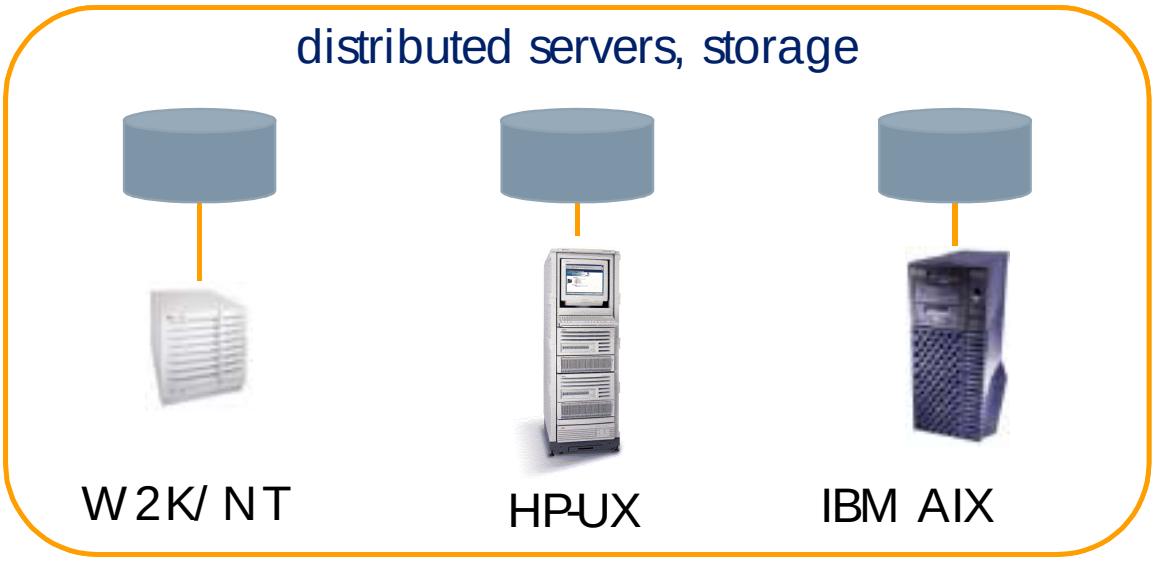

DAS is characterized by isolated islands of storage, with each server having its own disk and tape storage. This one-to-one, server centric model simply does not meet the demands of capacity growth, declining IT budgets, and need for higher levels of data availability. The weaknesses of DAS are:

1) If you add storage to one server, it does not benefit any other server in your environment
2) If a server fails, no other servers can get to the data associated with that server
3) And, it never seems to fail that excess storage on one server is desperately needed by another server.
4) Since each operating system requires expert administrators, management across platforms is complicated and labor intensive.
5) No overall/global view of the storage pool – for status monitoring and capacity planning
6) Finally, if storage is added, removed, or in some way changed on a server, it often has to be taken down denying business applications access to data.

The solution to many of these problems is to establish a storage architecture that allows management of storage as a centralized pool rather than isolated island of storage, a network centric approach instead of a server centric approach.

Storage Area Networks (SANs), and Network Attached Storage (NAS) are the two available methodologies for networking storage. There will be a brief overview of each.

NAS

A NAS configuration separates the storage from the server and makes the storage resource shareable among many servers. There are two components to a NAS configuration – the NAS controller (often called the "NAS head") and the disk storage component. These two components can either be deployed as a single NAS appliance, or deployed separately.

NAS is shareable as it connects directly to the standard Ethernet LAN, and any server on that LAN (provided it has permissions from the NAS controller) can access the NAS storage.

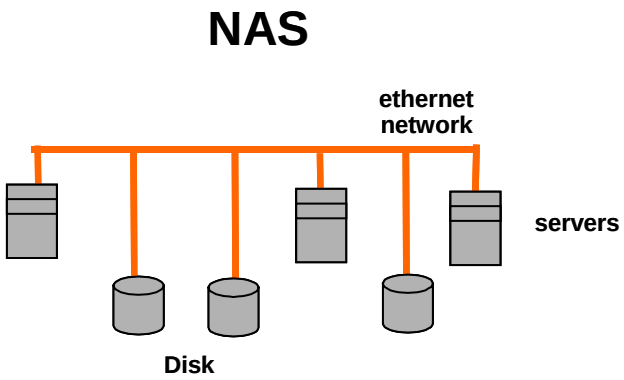

The key benefit of sharing the disk resource among many servers is an increase in overall storage efficiency. In a standard DAS environment, IT managers will add "white-space" (excess storage capacity to allow for future data growth) to each DAS installation. As a result each device has excess capacity so the overall capacity utilization (e.g. the proportion of total disk space that is actually used) is relatively low. In fact research shows[2] that this can be around 40%.

However, when storage becomes a networked resource, any server can access any storage device; it then becomes possible to pool these areas of white space and actually run with less spare capacity. Particularly as capacity additions are easier to make in a networked pool. Hence capacity utilization in a networked environment has been shown[3] to be at 80% or above.

The other strong benefit of NAS comes from the way in which it handles data – storage is contained and served at the file level. As explained by Goldman Sachs[4], file level storage has significant access and sharing benefits.

The key benefit of file access is the intelligence that is associated with file-based data – intelligence that provides benefits such as file sharing and ease of management. With file access systems, storage becomes accessible to any client via a common file access protocol such as Network File System (NFS) or Common Internet File System (CIFS). Because of the intelligence associated with file-based data, the amount of data used to transfer and process files can be less, sometimes significantly so, than what is used in block-based access. In addition, NAS devices, through the access methods of their file systems, are able to allow direct and easy access to the same set of data by users with different operating systems. This cross-platform access to shared data is an extremely powerful capability.

---

[2] Merrill Lynch/McKinsey The Storage Report June 01
[3] Merrill Lynch/McKinsey The Storage Report June 01
[4] Goldman Sachs Global Equity Research – Technology: Storage Networking – Nov 2000

Finally, NAS devices tend to offer faster access to shared files than a standard server – this is because NAS devices have operating systems solely dedicated to file serving and carry no other overhead. And, their inherent simplicity brings easy installation and ongoing management.

NAS is the optimal solution for applications with heterogeneous clients needing to access and share the same files e.g. CAD/CAM or generic departmental/workgroup file sharing.

In summary, the key benefits of NAS are:

- Improved storage efficiency compared to DAS
- Optimized solution for sharing files among multi OS servers
- Easy to deploy
- Simpler management than DAS as one administrator could manage storage associated with multiple servers.

Despite its obvious benefits, NAS does have limitations. Primarily its file-level storage serving is not appropriate for enterprise database applications, which require large volumes of data to be served. And, in general, NAS devices do not offer advanced backup or high availability features that mission critical applications require. Gartner[5] warns:

Users must avoid deploying NAS to support mission-critical or high-volume transaction database applications, especially online transaction processing, which requires immediate response to each request and results in high I/O intensity.

SAN

A SAN provides shared storage by creating a network of storage devices separate from the standard Ethernet LAN, and letting servers access that shared storage. SANs are usually built on a switched fibre channel network and data is stored and served at the block level. This means large volumes of data can be accessed, moved, and manipulated, very quickly.

SAN

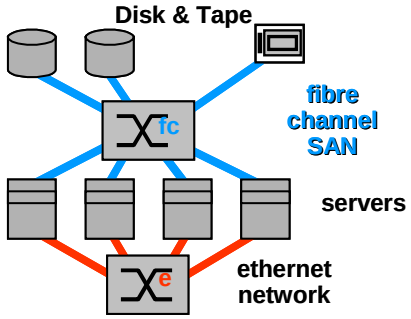

Like NAS, SANs allow storage resources to be pooled across many servers. One difference being that it's possible to have pooled tape as well as disk – NAS tends to be disk only.

SANs then go a step further by enabling pools of storage to run over campus-wide distances – the fibre channel protocol allows up to 10km between devices. Hence, SAN can provide a true enterprise-wide storage network.

The key difference between NAS and SAN is in the way the data is stored and served – NAS operates at the file level, and SAN runs at the block level. Goldman Sachs[6] provide a clear description of block level data access:

[5] Gartner Using SAN & NAS Technologies to consolidate storage (Mar 22, 2002)
[6] Goldman Sachs Global Equity Research – Technology: Storage Networking – Nov 2000

Block-based access deals instead with managing volumes – or blocks – of disk space with less importance assigned to identifying individual files on the disk. In its most basic application, block-based access provides high-speed access to large quantities of data. Block-based access is optimally used when the objective is to consolidate storage and data and then to duplicate it, back it up, or otherwise manage it en masse.

Hence SANs provide fast access to large quantities of data, and are thus ideally suited to database applications, such as order processing or ERP. And, when using a switched fibre-channel network, it's actually simpler, compared to a DAS environment, to build in redundant paths for high data availability.

Differences between NAS and SAN

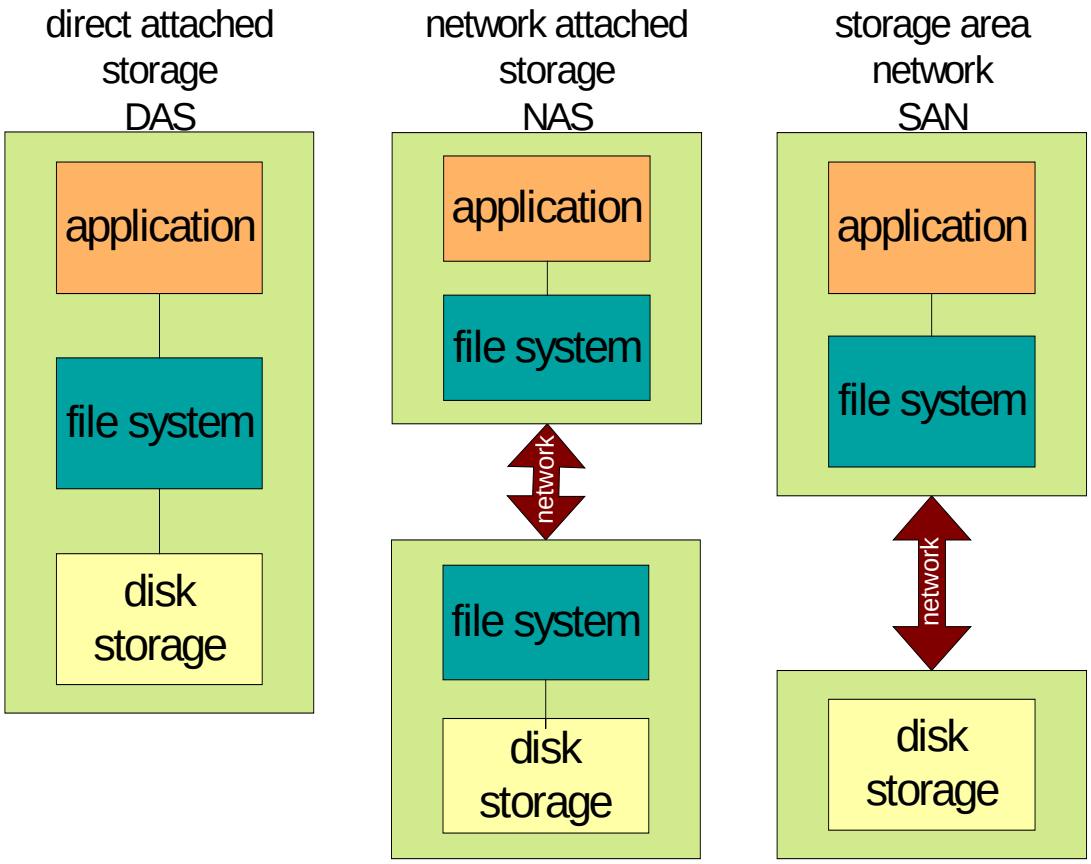

Thus the key difference comes down to file or block level transfers. Goldman Sachs[7] summarise the differences between file and block level storage management as:

File access	Block access
Allows for data sharing across platforms	Provides rapid access to large amounts of data
Makes storage addressable using common protocol (NFS or CIFS)	Enables high-speed direct access to data volumes
Uses less data to transfer and process information	Allows for centralized management of storage
Offers ease of installation and manageability	

[7] Goldman Sachs Global Equity Research – Technology: Storage Networking – Nov 2000

So, the conclusion is that NAS and SAN both offer distinct benefits, and enterprises need both. And, in fact, merging NAS and SAN together can deliver enhanced benefits to the organization. See the following section.

Combining NAS & SAN

It can actually be straightforward to add NAS to an existing SAN installation. As outlined early in this document, a NAS configuration comprises two key components – the NAS controller (or “NAS head”) and the disk system. In the early days of NAS development the two components were delivered as a combined appliance, now more sophisticated NAS solutions are available, which offer the NAS controller as a separate unit. This enables users to connect the NAS head to existing storage pools – either direct attached or on a SAN.

The diagram below shows an outline view of NAS controllers attached to a SAN. The basic premise being, the NAS heads can utilize the SAN storage pool and be managed by the central SAN management software.

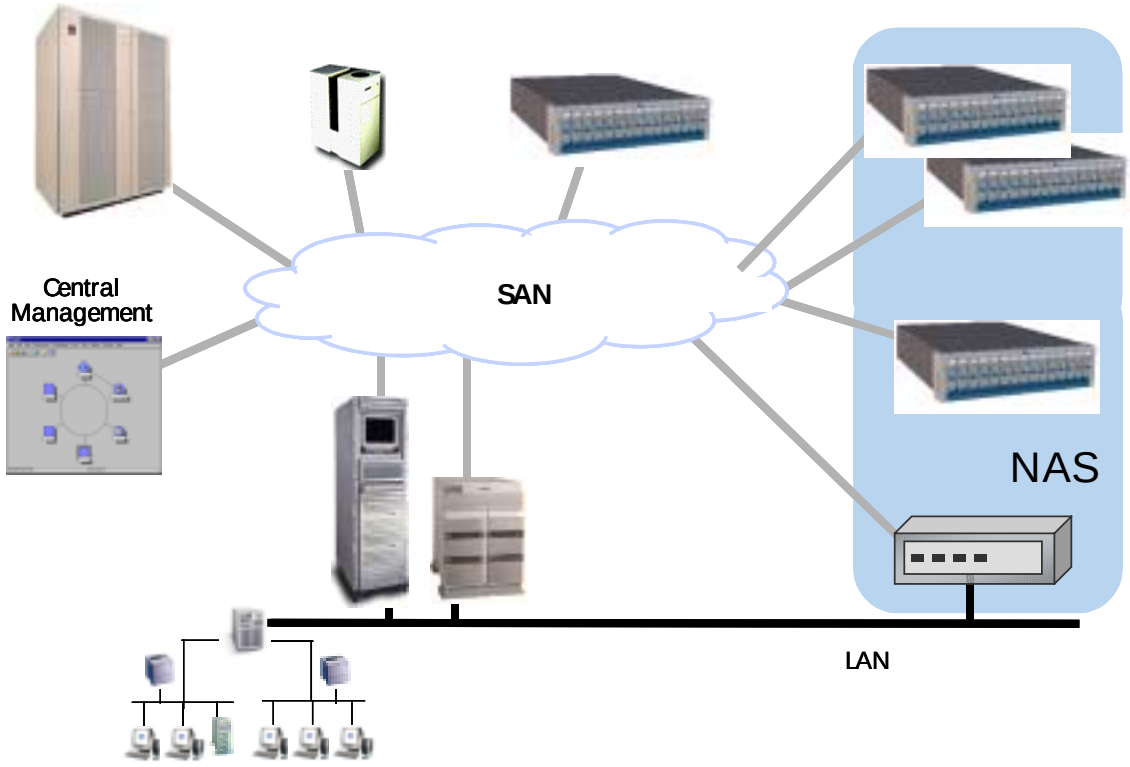

Benefits of combined NAS and SAN

- Fully centralized management
- More usage of SAN storage
- Higher availability of the NAS
- Pooled backup for the NAS

Fully centralized management

Connecting the NAS controller to the SAN brings yet another part of the enterprise's storage resource pool into the centralized monitoring and management of the SAN software. This improves IT's overall control of storage resources and helps existing staff to manage more capacity.

Greater utilization of SAN storage

By de-coupling the NAS controller from the NAS storage, there is no longer a need to have separate and dedicated NAS storage arrays. Instead, capacity on the shared SAN pool can be used. So, for example some LUNs on a SAN disk array could be run as file-level NAS storage, and others as block-level storage. This brings the benefit of greater utilization of the SAN pool and the ability to make more rational decisions around capacity upgrades e.g. only adding new capacity to one pool rather than separate SAN and NAS pools.

Centralized backup for the NAS

As the SAN disk arrays are now being used for the NAS primary storage, it follows that the SAN's centralized backup resources can also be used by the NAS. This actually presents a significant benefit. In traditional NAS installations (e.g. a NAS appliance comprising NAS controller and storage array) backup is often a straightforward direct-attached tape drive. The other option is to backup to a centralized backup server and library over the LAN – but this can heavily burden the network. So being able to use the SAN's shared backup resources and backup the NAS data on the separate SAN network is a significant improvement.

Higher levels of availability for the NAS

Finally, connecting the NAS to the SAN makes it easier to provide high levels of data availability for the NAS users. By de-coupling the NAS controller from the NAS storage, it becomes possible to build in redundant links and fail-over units more easily than in a traditional NAS appliance set-up. If fail-over is required with a standard NAS appliance, then the only option is to buy a second complete NAS unit. With a NAS-on-SAN configuration redundancy can be built in at the NAS controller or storage level as required – see two following diagrams.

Without redundant paths or fail-over

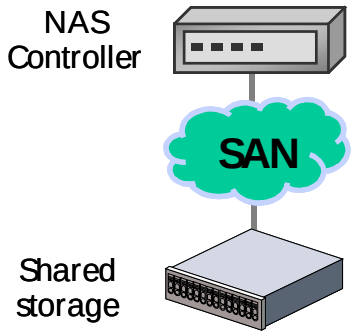

With redundancy and fail-over

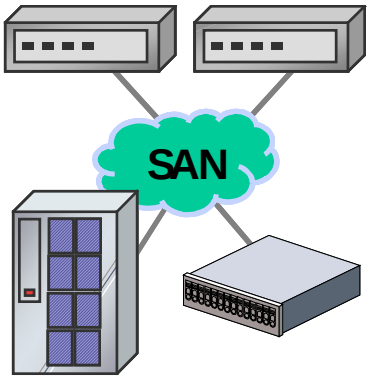

the solutions approach

HP is offering users the ability to add a NAS capability to any of the existing SAN blueprint configurations. Hp's latest NAS solution – the NAS8000 is available in SAN-attach configuration, and is compatible with the OpenView Storage Area Manager suite.

In the accompanying technical blueprint readers will find details on how to attach a NAS8000 to existing SAN blueprint configurations.

additional information

SAN solutions

Additional HP SAN Solution business blueprints

- Storage Efficiency
- High Availability
- Starter SAN
- Open Storage Management

For technical-level information on implementing the SAN configurations outlined in our business-level blueprints, take a look at our SAN Solution technical blueprints at www.hp.com/products1/storage/san/index.html

- Entry-level SAN
- Entry-level HA SAN
- Enterprise SAN
- Enterprise HA SAN
- Open Storage Management

To get answers on further SAN implementation questions contact your hp sales representative who will be able to consult our regularly updated SAN interoperability matrices and provide guidance on additional OS, fabric topology and 3rd party/legacy device interoperability.

hp SAN components

To get further information on the individual components in an HP SAN, go to www.hp.com/go/storage

hp services

HP offers a complete lifecycle of SAN support and consultancy services, go to www.hp.com/hps/storage for full details. Specific SAN services, which may be of interest to customers looking into our Storage Efficiency Blueprints are:

performance & capacity planning services

We will analyze the performance and capacity usage of your storage environment including all major system components. You will receive a detailed Performance and Capacity report with recommendations how to tune your performance and optimize your capacity usage.

data migration services

We offer a stress-free data migration from mission critical HP-UX, Windows NT/2000, SUN legacy and EMC storage systems to the HPStorageWorks SAN platform based on an end-to-end management of the entire data migration process.

storage/SAN integration services

HP also offers Storage/SAN integration services, which provide a trouble-free and quick on-site installation of your SAN solution. For full details see www.hp.com/hps/gds.

07/02