Thanks to John Mathews for the suggestion

To SAN or to NAS?

Author: Yogi Schulz

We've all heard the terms Storage Area Network (SAN) and Network Attached Storage (NAS). But what are the advantages and shortcomings of the two concepts? When should you deploy one over the other? How are these technologies evolving?

Background

Once upon a time, servers relied exclusively on directly attached storage (DAS). This approach worked well enough so long as the data volume was small and the need to share data among applications running on different servers was rare.

However, as data volumes grew and the need to integrate data from multiple applications grew, operating systems gained the ability to make storage volumes connected to one server visible to another server. This approach became unwieldy to manage as the number of devices continued to grow.

Backups also became problematic. Buying a dedicated tape drive for each server became increasingly expensive and resulted too much running around to mount and retrieve tape cartridges. Backing up the storage of one server through the local area network to a tape drive connected to another server is excruciatingly slow and creates server availability problems.

Origin of SAN and NAS

Both SAN and NAS technologies were developed to help organizations manage growing volumes of data without compromising performance and to contain operational complexity. SAN tended to be adopted by larger organizations with investments in mainframes and UNIX servers. NAS was often acquired by organizations with investments in Intel servers and with data stored using the Network File System (NFS). NAS also provided a way to easily share data between UNIX and Windows workstations.

SAN and NAS Technology

In support of describing the advantages and shortcomings of the two technologies, we need to define their technical characteristics. The data transmission technology of NAS is TCP/IP over Ethernet while SAN uses Fibre Channel. NAS uses a file approach to data transmission while SAN uses a block approach. The protocol of NAS is NFS while SAN uses encapsulated SCSI.

SAN and NAS Benefits

SAN advantages include enhanced data management, high performance through the block approach to data transmission and scalability in smaller increments.NAS advantages include easy file serving, simplicity of management and scalability in quite large increments.

Both technologies offer a high-speed approach to backup that does not require local area network resources.

As to shortcomings, NAS owners are concerned about the limits in scalability, availability and functionality of their NAS solutions. Some SAN owners are concerned about implementation complexity and the lack of interoperability among the components of different SAN suppliers.

SAN and NAS Developments

As a result of mergers and data center consolidation, many organizations operate both NAS and SAN in their computing infrastructures. Having both technologies, while offering all the advantages, also tends to create operational complexity.

To address the shortcomings and to maximize the strengths of both technologies, some storage vendors have begun to offer a hybrid solution. Such a storage configuration uses a SAN on the back end with NAS on the front end. Most industry observers suggest that the ultimate goal of seamless interoperability and data sharing can and will be reached.

Conclusions

If your organization is just starting to think about implementing NAS or SAN, a hybrid solution should receive serious consideration. If your requirement is truly just shared file access, then a NAS will be most suitable. If shared access to databases, perhaps with failover, is your requirement, then a SAN is the better choice.

SHORTCOMINGS:

Throughput is limited by the network the device is connected to, so it will be slower than if connected to the computer through most other means.

The benefits of NAS

By *Ryan* – January 12, 2011**Posted in:** Product News

I love all of the technology I have in the house and will always look for the next gadget to own, but I spend a lot of time with friends who aren't particularly interested in motherboards, BIOS settings or the reasons to upgrade to solid state drives (here's my blog post on this: SSHD). Therefore, I find myself trying to explain the vast amount of technology I have lying around and, more importantly, why I'm replacing a perfectly good graphics card, hard drive or [insert any type of technology, here...] with a newer technology. The most recent occurrence of this has been the inspiration for this blog post, namely, trying to explain why I have chosen to use NAS (Network Attached Storage) over a simple and perfectly functional external hard drive. So, let's start by taking a look at the benefits of NAS:

One of the main benefits of NAS over an external HD is that it is a standalone device so it does not have to be directly attached to a computer, rather it is attached to the network and is therefore easier and quicker to access from multiple

places or computers directly connected to the network. If you have multiple computers in the house (laptop, desktop or workstation, home media system in the lounge, etc) or a few flatmates that also want access, then NAS is the fastest way for all of these devices to access the data.

Most NAS drives give you the option to add more hard drives to the configuration. This has two benefits. Firstly, you can easily upgrade your storage without having to transfer everything across to a new and larger drive. You simply slot in a second drive and go from there. Most consumer NAS systems have space for up to two hard drives, so the ability to add a second can be invaluable. The second benefit comes from when you have this second drive in place, namely, setting the two drives up into a RAID configuration. I won't get too technical here (there are scores of great sites and blogs that tell you exactly how to do this), but by having the drives set up in this configuration they can "mirror" the data stored on one drive to allow you to access each file far quicker (by downloading half the data from one drive and half from the other, at the same time) and will be protected from drive failures or data loss. If your external hard drive fries itself you're in a lot of trouble, but if one of your NAS drives fries, the other has the data stored so you can recover your library.

If you pride yourself on your huge movie or music library, are storing large amounts of data for your home office or would like to have faster data access and a more secure place to store your digital stuff, then take a serious look at Network Attached Storage drives. We have a few great options at the moment, such as the Seagate 2TB GoFlex Home NAS (£119.98 in our sale), the very impressive LG Electronics N2R1 2x1TB NAS with DVD-RW (for £174.99), or the cost effective, entry level D-Link ShareCentre Pulse 2-bay NAS (for only £69.99).

Network-attached storage

From Wikipedia, the free encyclopedia

Not to be confused with Storage area network.

Network-attached storage	
Connects to	Local Area Network via one of: • Ethernet • Wi-Fi Hard Drives via one of: • SATA • SAS • USB
Common manufacturers	EMC Corporation NetApp HP Buffalo Technology Netgear Sun Microsystems

	LaCie
Introduced	1980s

Network-attached storage (**NAS**) is file-level computer data storage connected to a computer network providing data access to a heterogeneous group of clients. NAS not only operates as a file server, but is specialized for this task either by its hardware, software, or configuration of those elements. NAS is often manufactured as a computer appliance – a specialized computer built from the ground up for storing and serving files – rather than simply a general purpose computer being used for the role.[nb 1]

As of 2010 NAS devices are gaining popularity, as a convenient method of sharing files among multiple computers.[1] Potential benefits of network-attached storage, compared to file servers, include faster data access, easier administration, and simple configuration.[2]

NAS systems are networked appliances which contain one or more hard drives, often arranged into logical, redundant storage containers or RAID. Network-attached storage removes the responsibility of file serving from other servers on the network. They typically provide access to files using network file sharing protocols such as NFS, SMB/CIFS, or AFP.

Contents

- 1 Description
  - 1.1 NAS vs. DAS
  - 1.2 NAS vs. SAN
- 2 History
- 3 Implementation
- 4 Uses
- 5 Examples
  - 5.1 Open source server implementations
  - 5.2 List of network protocols used to serve NAS
- 6 Clustered NAS
- 7 See also
- 8 Notes
- 9 References

Description

A Netgear NAS

A NAS unit is a computer connected to a network that provides only file-based data storage services to other devices on the network. Although it may technically be possible to run other software on a NAS unit, it is not designed to be a general purpose server. For example, NAS units usually do not have a keyboard or display, and are controlled and configured over the network, often using a browser.[3]

A full-featured operating system is not needed on a NAS device, so often a stripped-down operating system is used. For example, FreeNAS, an open source NAS solution designed for commodity PC hardware, is implemented as a stripped-down version of FreeBSD.

NAS systems contain one or more hard disks, often arranged into logical, redundant storage containers or RAID.

NAS uses file-based protocols such as NFS (popular on UNIX systems), SMB/CIFS (Server Message Block/Common Internet File System) (used with MS Windows systems), AFP (used with Apple Macintosh computers), or NCP (used with OES and Novell NetWare). NAS units rarely limit clients to a single protocol.

NAS vs. DAS

The key difference between direct-attached storage (DAS) and NAS is that DAS is simply an extension to an existing server and is not necessarily networked. NAS is designed as an easy and self-contained solution for sharing files over the network.

Both DAS and NAS can potentially increase availability of data by using RAID or clustering.

When both are served over the network, NAS could have better performance than DAS, because the NAS device can be tuned precisely for file serving which is less likely to happen on a server responsible for other processing. Both NAS and DAS can have various amount of cache memory, which greatly affects performance. When comparing use of NAS with use of local (non-networked) DAS, the performance of NAS depends mainly on the speed of and congestion on the network.

NAS is generally not as customizable in terms of hardware (CPU, memory, storage components) or software (extensions, plug-ins, additional protocols) as a general-purpose server supplied with DAS.

NAS vs. SAN

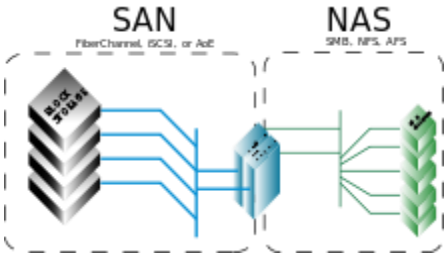

Visual differentiation of NAS vs. SAN use in network architecture

NAS provides both storage and a file system. This is often contrasted with SAN (Storage Area Network), which provides only block-based storage and leaves file system concerns on the "client" side. SAN protocols include Fibre Channel, iSCSI, ATA over Ethernet (AoE) and HyperSCSI.

One way to loosely conceptualize the difference between a NAS and a SAN is that NAS appears to the client OS (operating system) as a file server (the client can map network drives to shares on that server) whereas a disk available through a SAN still appears to the client OS as a disk, visible in disk and volume management utilities (along with client's local disks), and available to be formatted with a file system and mounted.

Despite their differences, SAN and NAS are not mutually exclusive, and may be combined as a SAN-NAS hybrid, offering both file-level protocols (NAS) and block-level protocols (SAN) from the same system. An example of this is Openfiler, a free software product running on Linux-based systems. A shared disk file system can also be run on top of a SAN to provide filesystem services.

History

In the early 1980s, the "Newcastle Connection" by Brian Randell and his colleagues at Newcastle University demonstrated and developed remote file access across a set of UNIX machines.[4][5] Novell's NetWare server operating system and NCP protocol was released in 1983. Following the Newcastle Connection, Sun Microsystems' 1984 release of NFS allowed network servers to share their storage space with networked clients. 3Com and Microsoft would develop the LAN Manager software and protocol to further this new market. 3Com's 3Server and 3+Share software was the first purpose-built server (including proprietary hardware, software, and multiple disks) for open systems servers.

Inspired by the success of file servers from Novell, IBM, and Sun, several firms developed dedicated file servers. While 3Com was among the first firms to build a dedicated NAS for desktop operating systems, Auspex Systems was one of the first to develop a dedicated NFS server for use in the UNIX market. A group of Auspex engineers split away in the early 1990s to create the integrated NetApp filer, which supported both the Windows CIFS and the UNIX NFS protocols, and had superior scalability and ease of deployment. This started the market for proprietary NAS devices now led by NetApp and EMC Celerra.

Starting in the early 2000s, a series of startups emerged offering alternative solutions to single filer solutions in the form of clustered NAS – Spinnaker Networks (acquired by NetApp in February 2004), Exanet (acquired by Dell in February 2010), Gluster (acquired by RedHat in 2011), ONStor (acquired by LSI in 2009), IBRIX (acquired by HP), Isilon, (Comanter), (acquired by EMC), PolyServe (acquired by HP in 2007), and Panasas, to name a few.

In 2009, NAS vendors (notably CTERA Networks[1][6] and NETGEAR) began to introduce online backup solutions integrated in their NAS appliances, for online disaster recovery.[7][8]

Implementation

The way manufacturers make NAS devices can be classified into three types:

1. Computer based NAS—Using a computer (Server level or a personal computer), installs FTP/SMB/AFP... software server. The power consumption of this NAS type is the largest, but its functions are the most powerful. Some large NAS manufacturers like Synology, QNAP and ASUStor make these types of devices. Max FTP throughput speed varies by computer CPU and amount of RAM.
2. Embedded system based NAS—Using an ARM, MIPS... embedded system and RTOS to run a NAS server. The power consumption of this NAS type is fair, and functions in the NAS can fit most end user requirements. Marvell, Oxford, and Storlink make chipsets for this type of NAS. Max FTP throughput varies from 20 MB/s to 120 MB/s.
3. ASIC based NAS—Provisioning NAS through the use of a single ASIC chip, using hardware to implement TCP/IP and file system. There is no OS in the chip, as all the performance-related operations are done by hardware acceleration circuits. The power consumption of this type of NAS is low, as functions are limited to only support SMB and FTP. Layerwalker is the only chipset manufacturer for this type of NAS. Max FTP throughput is 40 MB/s.

Uses

NAS is useful for more than just general centralized storage provided to client computers in environments with large amounts of data. NAS can enable simpler and lower cost systems such as load-balancing and fault-tolerant email and web server systems by providing storage services. The

potential emerging market for NAS is the consumer market where there is a large amount of multi-media data. Such consumer market appliances are now commonly available. Unlike their rackmounted counterparts, they are generally packaged in smaller form factors. The price of NAS appliances has plummeted in recent years, offering flexible network-based storage to the home consumer market for little more than the cost of a regular USB or FireWire external hard disk. Many of these home consumer devices are built around ARM, PowerPC or MIPS processors running an embedded Linux operating system.

Examples Open source server implementations

Open source NAS-oriented distributions of Linux and FreeBSD are available, including FreeNAS, NAS4Free, CryptoNAS, NASLite, Gluster, Openfiler, OpenMediaVault, EasyNAS and the Debian-based TurnKey File Server.[9] These are designed to be easy to set up on commodity PC hardware, and are typically configured using a web browser.

They can run from a virtual machine, Live CD, bootable USB flash drive (Live USB), or from one of the mounted hard drives. They run Samba (an SMB daemon), NFS daemon, and FTP daemons which are freely available for those operating systems.

NexentaStor, built on the Nexenta Core Platform, is similar in that it is built on open source foundations; however, NexentaStor requires more memory than consumer-oriented open source NAS solutions and also contains most of the features of enterprise class NAS solutions, such as snapshots, management utilities, tiering services, mirroring, and end-to-end checksumming due, in part, to the use of ZFS.

List of network protocols used to serve NAS

- AFS
- AFP
- CIFS
- FTP
- HTTP
- NFS
- rsync
- SFTP
- UPnP

Clustered NAS

Main article: Clustered file system

A **clustered NAS** is a NAS that is using a distributed file system running simultaneously on multiple servers. The key difference between a clustered and traditional NAS is the ability to distribute[*citation needed*] (e.g. stripe) data and metadata across the cluster nodes or storage devices. Clustered NAS, like a traditional one, still provides unified access to the files from any of the cluster nodes, unrelated to the actual location of the data.