IJRSET Volume 2, Issue 4 Pages: 25-31

ISSN 2394-739X
International Journal for Research in Science Engineering and Technology

Investigations on Network-Attached Storage Devices

[1] A. Jasmine Antony Raj
[1] Assistant Professor, PG & Research Dept of Computer Science
[1] Hindusthan College of Arts & Science, Coimbatore

Abstract:-

The need in the business for network storage mediums bring about the interest for storage limit. Our reality is completely reliant on developing business for networked storage is result for storage limit in our undeniably web subordinate world and its tight work market. Storage area networks (SAN) and network attached storage (NAS) are two demonstrated ways to deal with networking storage. actually, incorporating a document framework in a storage subsystem differentiaates nas, which has one, from san, which doesn't. practically speaking, then again, it is frequently nas' nearby relationship with Ethernet network hadware and san with Fiber Channel network equipment that has a more prominent impact on a client's buying choices. This article is about how rising innovation may obscure the network-driven refinement in the middle of NAS and SAN. For instance, the diminishing specialization of SAN conventions guarantees SAN-like gadgets on Ethernet network equipment. On the other hand, the expanding specialization of NAS frameworks may implant a great part of the document framework into storage gadgets. For clients, it is progressively beneficial to explore networked storage center and developing advances.

Keywords: - network-attached Storage, SAN

1. INTRODUCTION

Network-attached storage (NAS) is a type of dedicated file storage device that provides local-area network local area network (LAN) nodes with file-based shared storage through a standard Ethernet connection. Network-attached storage uses a standard Ethernet connection to provide network nodes with file-based shared storage services. NAS devices, which typically do not have a keyboard or display, are configured and managed with a Web-based utility program. Each NAS resides on the LAN as an independent network node and has its own IP address. In the home, NASes are often used for storing and serving multimedia files and for automated backups. Many smart homes rely on NAS to provide centralized storage for smart TVs, security systems and other Internet of Things (IoT) components in the home. In the enterprise, a NAS array can be used as a backup target for archiving and disaster recovery. If a NAS device has a server mode, it can also function as an email, multimedia, database or print server for a small business. Some higher-end NAS products can hold enough disks to support RAID, a storage technology that turns multiple hard disks into one logical unit in order to provide better performance times, high availability and redundancy. One of the main advantages of a NAS storage solution is its simplicity of setup. Many NAS products ship with pre-installed disks that allow non-IT personnel, guided by setup wizards, to quickly add storage to a network often in less than 15 minutes. Since the NAS adds storage by simply attaching to the network rather than to a server, there is no server downtime.

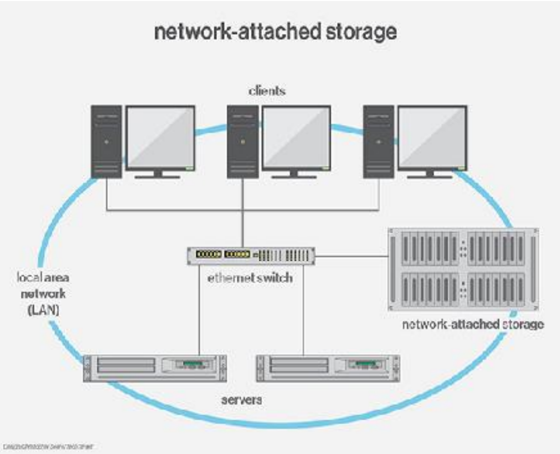

Figure: 1 network-attached storage

Often, NAS products use Linux as the operating system, eliminating the licensing costs associated with Windows or Unix-based servers. The Linux kernel, stored in flash memory on the device, rarely has to be patched and is less vulnerable to viruses. This eliminates the down-time associated with patch installation and anti-virus updates for servers using general purpose operating systems. In addition, NAS products support both CIFS, as well as NFS, file systems and are completely compatible with Windows, Mac, and Linux and UNIX clients. No specialized IT experience is required to ensure cross platform compatibility. Advances in both NAS and disk technology have made NAS storage an attractive, cost-effective alternative to the more expensive Storage Area Networks (SANs) that are favored by larger enterprises. Currently, the inexpensive SATA drives used in many NAS products have capacities of up to 2 Terabytes, so 4 drive/8 Terabyte configurations are fairly standard. Yet for the simplicity of configuration, NAS products targeted at small to medium enterprises possess the features to seamlessly blend into an existing environment as they provide high data availability. Some of these features are

Active Directory support: Many NAS products support Active Directory. This allows the NAS to be easily designated as a storage device if the business is using Active Directory.

Multiple gigabit Ethernet ports: NAS products designed for the enterprise have at least one gigabit Ethernet port, and many have two. The gigabit ports ensure that the network connection will not be a performance bottleneck. A second port can be used to provide additional network capacity into and out of the NAS, or as a failover if the primary port should fail.

RAID support: RAID technology is commonly used to provide fault tolerance in a disk system by storing information across multiple disks. Should a disk fail, depending on the RAID level, the data will not be lost. Here is a brief listing of the common RAID levels:

o **RAID 0**: RAID 0 improves performance by "striping" data across two or more drives. Performance improves, but if any of the drives fail, all data is lost

o **RAID 1**: Data is "mirrored" by writing copies of the same file to two or more drives, essentially creating a real time backup.

o **RAID 5:** Data is striped across three or more disks, along with parity information needed for recovery. If a single drive fails, the missing data can be recreated from the data and parity information on the remaining good drives. If more than one drive fails in a RAID 5 configuration, however, the data will be lost.

o RAID 6: Data is striped across multiple four or more drives and uses a dual parity information system so that data will not be lost even if two drives fail. o RAID 10: RAID 10 combines RAID 0, which stripes the data across drives, and RAID 1, which mirrors the drive contents. RAID 10 requires a minimum of four drives to provide a storage system that can withstand multiple failures.

Hot Swappable Hard Drives: If a drive fails in a NAS that has "hot swappable" capabilities, the failed drive can be removed and replaced with a new drive without powering down the system. RAID systems restore their fault tolerance by rebuilding the array as a background task after a defective drive has been replaced. In any RAID configuration except RAID 0, clients will retain access to their data even as the process takes place.

Data Replication: Some NAS products support "share-level replication" with other compatible NAS products. With share level replication, folders from one NAS device can be automatically and immediately backed up to another NAS device on the network, providing yet another layer of data redundancy to ensure that no data is lost.

ISCSI: In the past, NAS products provided only file level access, and were used primarily for serving client requests for files such as spreadsheets and word processed documents. With the installation of software called an iSCSI initiator, a server can use a designated iSCSI target as though it were directly attached storage. Today, many NAS products support iSCSI, an IP-based standard protocol used for linking storage devices. iSCSI-capable NAS products allow for some of all of the storage capacity to be designated as an iSCSI target. The server, rather than using a SCSI host bus adapter, uses its Ethernet connection to communicate with the iSCSI storage target. Thus, and iSCSI enabled NAS can provide block level access to a server in a different physical location. Since iSCSI operates over an existing network infrastructure, no specialized additional cabling is required such as would be found in a Fiber Channel SAN. Current NAS implementation of iSCSI also allows multiple servers to share iSCSI NAS-based storage.

Requirements for Emerging Systems Future network storage systems must provide features to meet all existing requirements, including resource consolidation, rapid deployment, central management, convenient backup, high availability, and data sharing, as well as the following emerging requirements.

Geographic separation of system components.

Because online commerce is increasingly global and competitive, remote data must be continuously available, and all data must have remote copies updated frequently to protect against regional disasters. Moreover, with the Internet infrastructure's bandwidth growing at upward of 300% per year, employing the Internet in an organization's own internal network is increasingly affordable and effective.

Increasing risk of unauthorized access to storage.

Security is an increasingly critical storage property as online commerce becomes important and as electronic crime (hacking) becomes increasingly sophisticated and common. Moreover, storage-system interconnects, including most SANs today, were originally designed as extensions of the internal buses of their hosts; their security provisions are limited or nonexistent. Extended buses may be easily secured from hostile traffic through physical means, but they are also limited in geographic distribution and the number of attached devices they can support. Linking extended buses to other networks, now possible with storage interconnects like Fibre Channel, greatly weakens these physical controls. Mechanisms for restricting access to network-storage servers are even more important on the Internet than in standalone Fiber Channel networks. Need for performance to scale with capacity.

Performance, measured either as accesses per second or megabytes per second, needs to scale with storage capacity to accommodate the increasing power and number of client machines, as well as the increasing size of datasets (for applications manipulating such data as sensor traces, transaction records, still images, and video).

In general, the increasing scale and scope of the use of storage systems drives these emerging requirements. mechanisms most likely to be stressed beyond their design goals by the increasing scale and scope of how these systems are used.

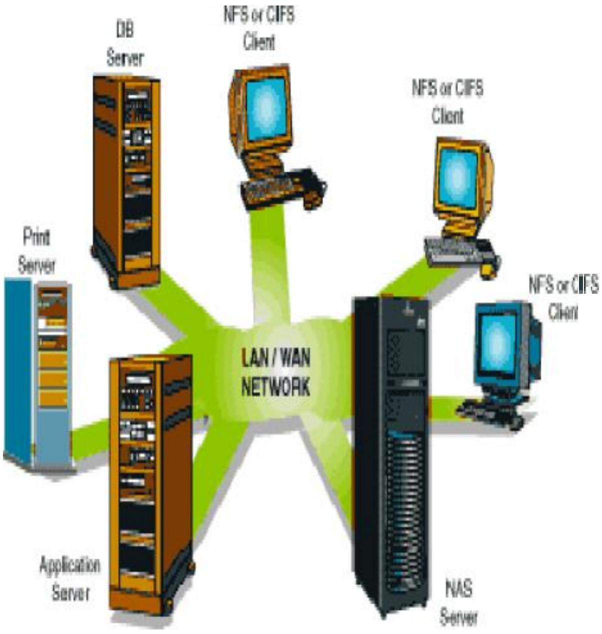

Figure 1.2 LAN/WAN Network

Advantages

The benefit of a NAS over a SAN or DAS is that multiple clients can share a single volume, whereas SAN or DAS volumes can be mounted by only a single client at a time. NAS devices allow administrators to implement simple, low cost load-balancing and fault-tolerant systems.

Disadvantages

The downside to a NAS is that not all applications will support it because they're expecting a block-level storage device, and most clustering solutions are designed to run on a SAN. Besides the backup solution is more expensive than the storage system. And even, any constrictions in the local area network will slow down the storage access time.

2. DIFFERENCE BETWEEN SAN AND NAS, ARCHITECTURE

A storage area network (SAN) is storage connected in a fabric (usually through a switch) so that there can be easy access to storage from many different servers. From the server application and operating system standpoint, there is no visible difference in the access of data for storage in a SAN or storage that is directly connected. A SAN supports block access to data just like directly attached storage.

Network-attached storage (NAS) is really remote file serving. Rather than using the software on your own file system, the file access is redirected using a remote protocol such as CIFS or NFS to another device (which is operating as a server of some type with its own file system) to do the file I/O on your behalf. This enables file sharing and centralization of management for data.

So from a system standpoint, the difference between SAN and NAS is that SAN is for block I/O and NAS is for file I/O. One additional thing to remember when comparing SAN vs. NAS is that NAS does eventually turn the file I/O request into a block access for the storage devices attached to it.

3. NAS QUALITY

The quality of the NAS is a major factor for ongoing success in the enterprise. When choosing a NAS, consider that up-front price may not be as important as the capability of the device. Lower-end NAS devices may have a much higher total cost of ownership when considering downtime, future upgrades, or lackluster performance.

If using a NAS, Dell recommends enterprise-grade network attached storage for best performance. The higher-end NAS hardware that you use, the less likely you are to encounter NAS hardware issues (providing reasonable network load and environmental factors). Consider a device with features such as redundant Gigabit Ethernet connections or 10Gbit Ethernet connections; consider whether you need access to the Fibre Channel storage-area network (SAN).

Consider if the NAS appliance allows you to upgrade capacity. Perform research before purchasing a NAS if possible; considering searching the internet using a phrase such as Guide to Network Storage.

For a NAS device to be supportable, the data saved to the repository must remain in the exact state in which the Core stored it. For this reason, for AppAssure 5, Dell does not support NAS devices that have their own built-in deduplication features if those features are enabled.

Sufficient input/output (I/O) transfer speed will yield the best results for backing up to the repository.

Dell recommends hard drives of at least 7200 RPM with good access speeds. For transfer speeds, Dell recommends transfer speeds of at least 30 megabytes per second, with a minimum of at least 10 MB per second. If transfer speeds appear to be below 10 MB/second, the issues are most likely to be (a) a result of insufficient hardware, (b) hardware that would be sufficient but is being too heavily tasked (multi-purposed, or poor with multiple operations), or (c) a network that is saturated and is acting as a bottleneck for the transfer. The administrators should be aware that NAS devices are susceptible to the same environmental stresses as other systems on the network. Factors that affect network performance include number of concurrent users, network load, number of operations, frequency of backups, and other issues familiar to network administrators.

You may consider optimizing your retention policy.

4. INTERMITTENT FAILURES OF NAS

If you experience NAS issues are intermittent, check that any mount failures may be occurring at a time when there is particularly high I/O activity. For example, there could be multiple processes occurring such as agent backups during VM Export, or nightly job occurring while also performing rollups, etc. Intermittent failures that are caused by too much I/O traffic are difficult to diagnose and may appear as unrelated issues such as replication failures or Exchange log truncations, but that are actually symptoms of an overtaxed NAS with too many I/O operations being attempted simultaneously. Changing the order of some I/O operations or otherwise reducing I/O is appropriate in these cases. Refer to troubleshooting steps. After trying other steps, consider reducing the rate of transfer speed to allow the NAS to catch up, as described below.

CONCLUSION

Network centralization and network decentralization are about rolling out improvements. At the point when rolling out such pervasive improvements, specialists prescribe that you consider these controlling standards: realize what particular issue you are settling; comprehend your inspiration for rolling out the improvement; concentrate as much as bodes well throughout today; perceive that as in revealing any new administration, it obliges cautious arranging; and, most essential, listen to the clients. The innovation is still moderately youthful and there is additionally an absence of acknowledged mechanical norms.

Information trustworthiness is a vital issue. Today, numerous present usage depend on the physical wiring just for security, however other security plans are being created. For instance, a few sellers are adding to a consistent parceling methodology called zoning.

REFERENCE

1. Benner, A. Fibre Channel: Gigabit Communications and I/O for Computer Networks. McGraw Hill, New York, 1996.
2. Callaghan, B. NFS Illustrated. Addison Wesley Publishing Co., Reading, Mass., 2000.
3. Gibson, G., et al. A cost-effective, high-bandwidth storage architecture. In Proceedings of the ACM 8th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS) (San Jose, Calif., Oct). ACM Press, New York, 1998, 92–103; see also www.pdl.cs.cmu.edu.
4. Hartman, J. and Ousterhout, J. The Zebra striped network file system. In Proceedings of ACM Symposium on Operating Systems Principles (SOSP) (Ashville, N.C., Dec.). ACM Press, New York, 1993, 29–43.
5. Hitz, D., Lau, J., and Malcolm, M. File systems design for an NFS file server appliance. In USENIX Winter 1994 Technical Conference Proceedings (San Francisco, Jan. 1994).
6. Kronenberg, N., et al. VAXclusters: A closely coupled distributed system. ACM Transact. Comput. Syst. (TOCS) 4, 2 (May 1986), 130–146.
7. Lee, E. and Thekkath, C. Petal: Distributed virtual disks. In Proceedings of the ACM 7th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS) (Cambridge, Mass., Oct). ACM Press, New York, 1996, 84–92.
8. McKusick, M., et al. A fast file system for Unix. ACM Transact. Comput. Syst. (TOCS) 2, 3 (Aug. 1984).
9. Network Appliance, Inc. DAFS: Direct Access File System Protocol, Version 0.53 (July 31, 2000); see www.dafscollaborative.org.
10. Sachs, M., Leff, A., and Sevigny, D. LAN and I/O convergence: A survey of the issues. IEEE Comput. 27, 12 (Dec. 1994), 24–32
11. Satran, J., et al. iSCSI (Internet SCSI), IETF draft-satran-isci-01.txt (Jul.10, 2000); see www.ece.cmu.edu/~ips.
12. Van Meter, R., Finn, G., and Hotz, S. VISA: Netstation's virtual Internet SCSI adapter. In Proceedings of the ACM 8th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS) (San Jose, Calif., Oct.). ACM Press, New York, 1998, 71–80; see also www.isi.edu/netstation/.