Database Configuration: SAN or NAS

Discussion of Fibre Channel SAN and NAS

Version: 1.0

Greg Schulz

DBSANNAS.DOC

Preface

This publication is for general guidance and informational purposes only. The information is current as of the date of publication. MTI, however, has no obligation to correct or update the information herein and makes no representations or warranties concerning its accuracy. References to MTI products herein do not imply that only MTI's products may be used. This document does not represent a commitment that MTI or any other entity will make generally available any specific products at a future date.

Printing Instructions

If you experience only part of a page printing or error messages on the printer's operator panel, such as "Low Memory," "Memory Overflow," "Print Overrun," etc., click "Properties." Next to your printer name in the print dialog box, click the "Graphics" tab and select the "Raster graphics" option instead of "Vector graphics." Double-click here for more help.

MTI Corporate Headquarters

4905 E. La Palma Avenue, Anaheim, CA 92807
(714) 970-0300 or (800) 999-9MTI
Fax: (714) 693-2256
http://www.mti.com

Table of Contents

1 INTRODUCTION

There is a lot of discussion about where best to store and how to access database information. This discussion is centered on Storage Area Networks (SANs) and Network Attached Storage (NAS) storage networking technologies. As a leading provider of advanced storage and information management solutions, MTI has prepared this document to help you understand the issues and facts in order to make informed decisions. This document is laid out with an executive overview and summary of issues on the first two pages. A detailed discussion, with examples, follows, along with a Meta group flash for best practices for databases.

1.1 Executive Summary

NAS technology has a place in the IT organization for general file sharing applications, using industry standard protocols such as NFS over traditional networks. NAS is not well suited nor should it be used for database applications. The reasons shown below pertain to performance, data integrity and management. General-purpose NAS or NFS filers, and specific products, do not provide basic data integrity safeguards to support databases.

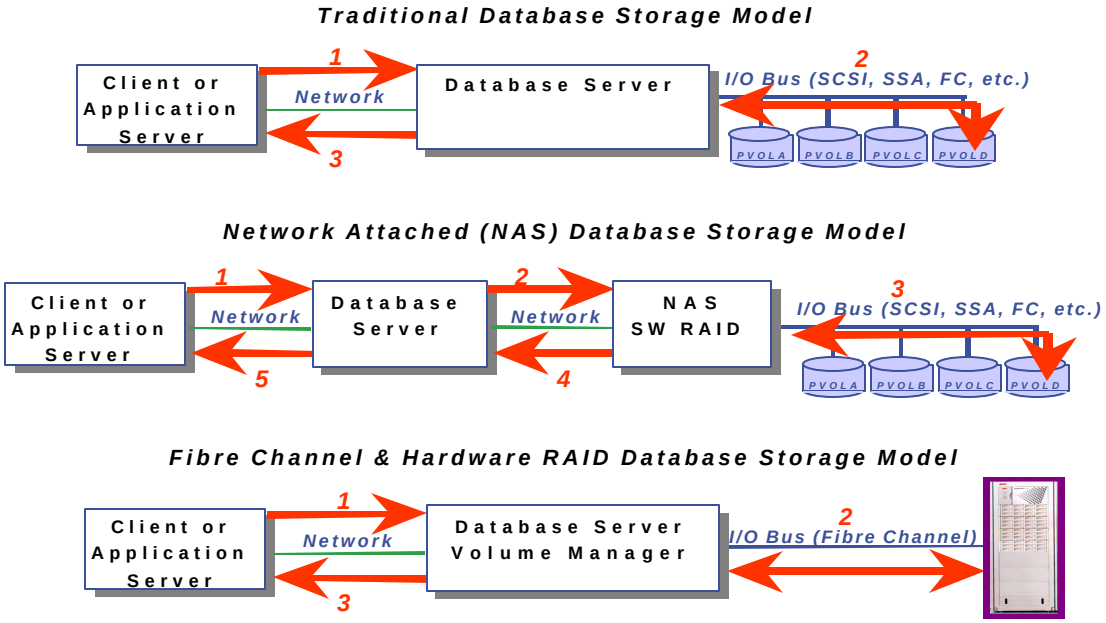

Figure-1 (Various database storage models)

While vendor-specific changes or enhancements can be found on some solutions, they do not provide full end-to-end data integrity for database systems or fit in all database applications. For example, NAS-based solutions do not support shared database access from two or more database servers of the same-shared database, such as an Oracle Parallel server.

The traditional database storage model (Figure-1 Top) consists of disk drives physically attached to a database server system using SCSI, SSA, or other storage interface. This is a

time-proven approach to data storage however it is not without pitfalls. These include lack of scalability, distance restrictions, performance, and limited multiple host storage sharing.

An improvement on the traditional storage model is to use Network Attached Storage (NAS) (Figure-1 Middle) where physical disk drives are attached to a server, allowing the storage to be made accessible over a network using such protocols as NFS and TCP/IP. This approach improves distance, sharing, ease of management and, in some cases, performance. This approach however adds additional workload to already overloaded and crowded networks, and requires extra steps to process an I/O request, which results in lower performance.

The solution is a Fibre Channel-based Storage Area Network (SAN) (Figure-1 Bottom) to off load I/Os from busy networks that uses interfaces optimized for storage functions. Fibre Channel and Hardware RAID provide mirrored cache, end-end data protection, multiple hardware RAID levels, dynamic storage allocation and movement, simplified backups and support for all databases on all platforms, including clustered database systems.

1.2 What do the industry analysts have to say?

META Group recently recommended against running databases on NAS*

"…we view NAS as a "storage appliance" for file serving, and organizations should not, in general, run DBMS (e.g., Oracle, SQL Server, UDB) on NAS solutions."

Sean Derrington, program director with META Group

Please refer to the attached Meta Group fax found at the end of this document.

1.3 The MTI integrated database solution

Integrated storage solutions such as those that MTI DataServices and MTI Vivant provide to enable database systems and applications to share storage and data. MTI Vivant reduces technology risk by using fully tested and integrated SAN components. MTI Vivant and DataServices architecture help to improve your overall productivity with performance, simplified management, end-to-end data integrity, and investment protection.

For additional information on SAN, NAS, enterprise storage, and MTI Vivant and DataServices refer to the MTI Web Site (www.mti.com) and these specific documents:

- SAN/NAS Complementary Technology – MTI Web Site and Enterprise Systems Journal
- Enterprise Storage Networks, NAS and SAN – MTI Web Site
- http://www.mti.com/documents/reports/whitepapers/wp2000/wp20002.pdf

* For complete META Fax article, please see end of this document

2 UNDERSTANDING THE TECHNICAL ISSUES

The following section details the specific technical issues involved with each of the three described database storage access models (traditional direct attached, network attached and Fibre Channel SAN).

2.1 Traditional database storage access model

In figure-2, you can see how database servers access storage in the traditional model involving three basic steps—client requests data from the database server over the network, database server does disk I/Os and the result is sent to the user or client over the network.

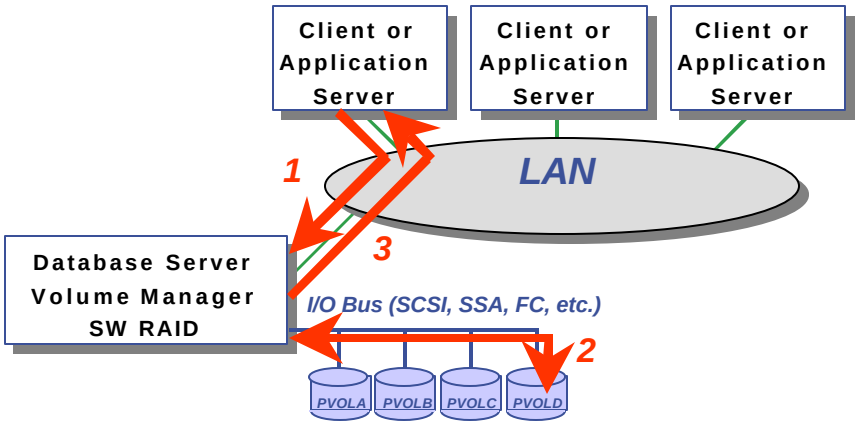

Figure-2 (Traditional storage and Database Server model)

Anatomy of a traditional database request:

Step	Action or activity
1	Client makes data request over network to database server
2	Database server performs I/O to disks over I/O bus (SCSI, IDE, SSA, etc.)
3	Database server sends results back to client via network

Note that database-specific activities. including transaction integrity, buffering and other functions, are present in all access methods discussed in this document.

Pros (+) and Cons (-) of this access method:

+ This is a time tested and proven access method regardless of the storage device, database, application, I/O bus interface and host computer platform.
- Physical storage can not be shared with other systems or applications without being served by the host system it is attached to.
- Storage must be located physically close to host system due to I/O bus restrictions on distance and performance.
- The scalability of number and capacities of devices is limited, depending on I/O bus and storage technology being used.

- If a host volume manager is not being used, there will be a limited or no dynamic load balancing, expansion or RAID.
- There will be limited dynamic or snapshot backup capability.

2.2 Network Attached Storage (NAS) database access model

To solve the problem of shared storage and distance, and to help justify NAS products, specific NAS products and databases have been modified to support database operations in an NFS or CIFS NAS environment. Compared to the traditional model, the NAS model (figure-3), has five basic steps to access the data. In this example, the database server must make a request over the LAN to a NAS server that, in turn, must issue the disk I/O. This is an extra, and unnecessary, step that adds to overhead and response time for a database environment. While this approach is good for general file sharing, for applications such as a database that just requires a single host accesses and maintains data integrity, this is not a good fit. Analysts, including the Meta group, have issued reports and quotes recommending that NAS should not be used for database applications.

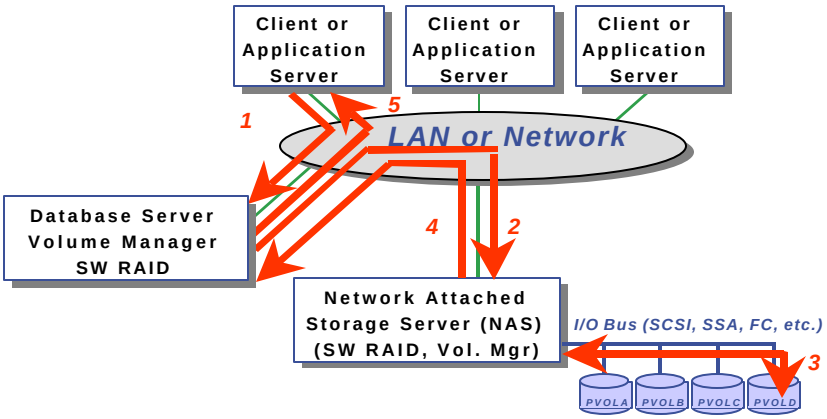

Figure-3 (Network Attach Database Server model)

The Network Attached Storage (NAS) model attempts to solve these issues by using general-purpose networks and common protocols such as NFS and TCP/IP. The NAS model, however, introduces new issues, including more overhead, extra traffic on the network, more steps to process and I/O, and lack of data integrity for database systems without special custom fixes.

Anatomy of a NAS Database request:

Step	Action or activity
1	Client makes data request over network to database server.
2	Database server performs I/O to Network Attached Server (NAS) over network.
3	NAS performs I/O to disks over I/O bus (SCSI, IDE, SSA, Fibre Channel).
4	NAS returns I/O to Database server over regular LAN or Network.
5	Database server sends results back to client over network.

Pros (+) and Cons (-) of this access method:

+ Storage sharing among two or more host systems is possible.
+ Storage can be physically located further away from host systems.
+ Storage can be brought online and reconfigured dynamically.
+ Backups can be simplified.
- Extra workload is placed on LAN network.
- Networks may need upgrading to support extra I/O activity.
- Extra steps are involved in performing I/O operations.
- There is a lack of mirrored cache for end-end data protection.
- There is no direct or server attach storage capability exists.
- There is no support for RAW volume partitions for fast I/O.
- There is no support for parallel or clustered databases like Oracle Parallel Server.
- I/O sizes and RAID configurations are limited or fixed.
- Extra burden is placed on NAS server to perform I/O, RAID and disk rebuilds.

2.3 Fibre Channel SAN, Server Attached Storage, Block Access

Fibre Channel-based Hardware RAID is the perfect fit for supporting database systems with built-in end-to-end data protection such as those from available from MTI. With built-in mirrored read and write cache, data integrity is maintained along with complete power protection.

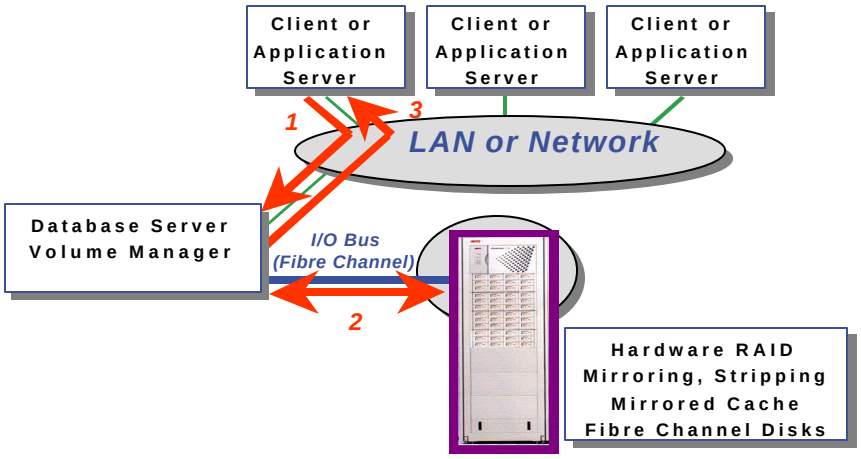

Figure-4 (Server Attached Storage/Block or Fibre Channel Attached Database Model)

In figure-4, Fibre Channel Hardware-based RAID solution such as those that MTI Vivant provides, feature end-to-end data protection. In this model, you have the low latency and high performance characteristics of the traditional proven model. However, you also have the flexibility of a networking solution with distance, sharing and multiple-host connectivity. This model also gives a level of data sharing not possible with NAS products today in the database world by supporting shared applications such as Oracle Parallel Servers--something that NAS solutions from NetApp can not do!

Anatomy of a Server Attached Storage (SAS) Block Database request:

Step	Action or activity
1	Client makes data request over network to database server.
2	Database server performs I/O over Fibre Channel I/O bus to hardware RAID.
3	Database server sends results back to client over network.

Pros (+) and Cons (-) of the Server Attached/Block access method:

- \- File or data sharing requires global file system/NAS like NFS or CIFs.
- \+ Storage sharing among two or more host systems is possible.
- \+ Storage can be physically located further away from host systems.
- \+ Storage can be brought on-line and reconfigured dynamically.
- \+ Backups can be simplified.
- \+ It is possible to expand and add more positives.
- \+ Workload is off-loaded from LAN networks.
- \+ You can postpone or delay updates and improve performance of LANs.
- \+ There are no extra steps involved in performing I/O to reduce latency.
- \+ There is fully redundant hardware RAID with mirrored cache.
- \+ Integration with volume managers and other storage utilities is possible.
- \+ There is support for RAW volumes for Oracle and other applications.
- \+ Dynamic storage allocation and expansion can be done.
- \+ There is support for parallel or clustered databases like Oracle Parallel Server.
- \+ Variable I/O size, RAID levels and other performance enhancement tools are avail able.
- \+ Hardware RAID offloads server from RAID operations and disk drive rebuilds.

While Fibre Channel-based Hardware RAID provides a simplified approach to database server storage, the Fibre Channel SAN capability, combined with Oracle Parallel Server or other software, enables databases to be clustered and shared, something that under the covers NAS can not do!

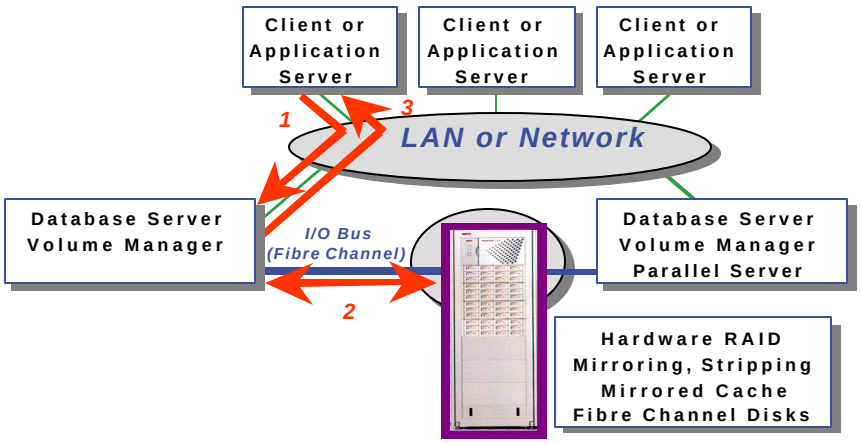

Figure-5 (Redundant or Parallel Access Database Storage Model)

Building on the example shown in Figure-4, while NAS provides storage and data sharing, database applications do not support shared access to their underlying data files without using parallel or sophisticated extensions such as Oracle Parallel server. These features are not available with NFS or NAS storage, as they require RAW or direct physical access via a traditional I/O bus for data integrity. The access steps are the same as in Figure-4, with the addition that a second or more database server(s) can be configured to share the workload and provide additional failover.

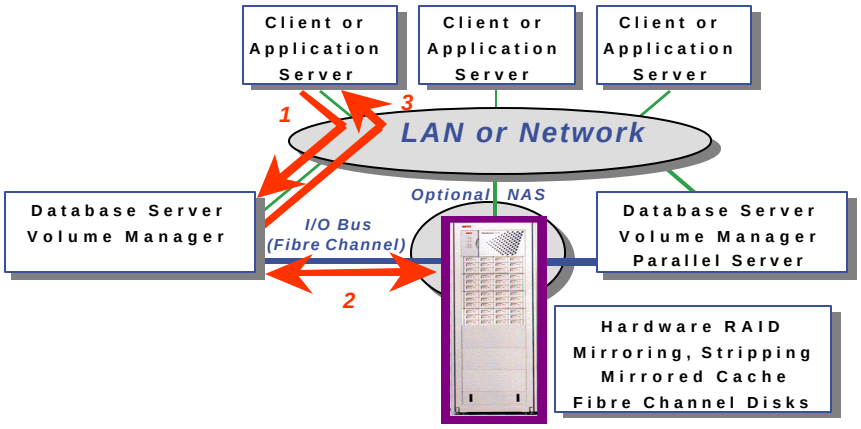

Figure-6 (Advanced integrated SAN and NAS model)

Combing the best of SAN and NAS access capabilities and functionality enables an IT organizations to configure and utilize technology in best-fit mode.

2.4 When to use NAS vs. Fibre Channel SAN

NAS is ideal for file sharing and Web hosting when multiple systems need to read the same files. SAN is better suited for databases or applications that do not require file sharing. In a nutshell, a SAN looks and performs like (or better than) internal or direct-attached storage. NAS, on the other hand, adds significant overhead. Will SAN and NAS merge or converge? Yes. In fact, vendors such as MTI have been delivering integrated NAS and SAN solutions for some time now. With MTI DataServices Architecture and Vivant family of products, storage can be tailored, tuned and managed to meet your applications' needs.

Today, SANs are well suited for storage sharing and building infrastructures for server and storage consolidation, as well as to address performance- and capacity-intensive applications. NAS should be used to solve data sharing issues and sharing of storage for small systems. SAN should be used for storage sharing and high performance, low latency applications, such as databases. A common mistake is to try and use SAN or NAS for everything, rather than match the proper technology or "access method" to the specific business and application issue.

Application or environment	Access method
Storage or server consolidation	Use SAN

Data sharing, staging, and movement between various host systems	Use NAS with underlying SAN
Data access by Unix, Linux, NT, and others	Use NAS with underlying SAN or SAN
Data sharing including Internet Web content for Web server farms	Use NAS with underlying SAN
Performance sensitive with low latency including database and OLTP	Use SAN
Large I/Os or data transfer applications	Use SAN
LAN-free or Serverless backup	Use SAN

Table-1, Comparison of when to use SAN and NAS or a combination of the two

In the above examples, SAN indicates that SAN storage sits “behind” the NAS server or appliance as a storage backbone. This differs from hosts systems that directly access SAN storage as block-level applications.

Whether SAN or NAS is better should not be the decision criteria for solving business issues; rather, one should consider how SAN and NAS could be combined to provide solutions. Take a look at your applications to see how SAN and NAS can collectively address various business issues. SAN and NAS implemented in a complimentary solution can help to reduce costs and simplify storage and data management. MTI provides a completely integrated SAN and NAS solution with the DataServices architecture and Vivant family of storage solutions. For more information on SAN, NAS and enterprise storage solutions, refer to the MTI Web site www.mti.com/documents for the various white papers on these and other related technologies.

META FAX

Summary of META Group's Weekly Research Meeting 3 July 2000
208 Harbor Drive • Stamford, CT 06912-0061 • (203) 973-6700 • Fax (203) 359-8066 • metagroup.com

NAS: Nearly Acceptable Storage

With significant hype around network-attached storage (NAS), users must understand its applicability and complementary nature to attach and storage-area network (SAN) solutions. Indeed, while channel-attached (e.g., Ultra SCSI, Fibre Channel) storage will be necessary for all users, we believe NAS (IP) storage is acceptable for high-performance file serving, file sharing, low-end/workgroup file serving, caching appliances, etc. Moreover, we view NAS, as a "storage appliance" for file serving, and organizations should not, in general, run DBMS (e.g., Oracle, SQL Server, and UDB) on NAS solutions. Organizations must weigh the impact/differences of storage administration (e.g., backup/recovery, scalability, manageability) of SAN and NAS solutions when selecting between the two.

Bottom Line: IT organizations should examine NAS solutions for file services, and channel-attached (SAN) solutions for block services, understanding the limitations/benefits of each architecture.