

StorTrends® Vs NetApp® implementation for SQL services in a Virtualized Server infrastructure

© Copyright 1998-2007 American Megatrends, Inc.

All rights reserved.

American Megatrends, Inc.

6145-F Northbelt Parkway

Norcross, GA 30071

TRADEMARK AND COPYRIGHT ACKNOWLEDGMENTS

This publication contains proprietary information that is protected by copyright. No part of this publication can be reproduced, transcribed, stored in a retrieval system, translated into any language or computer language, or transmitted in any form whatsoever without the prior written consent of the publisher, American Megatrends, Inc. Trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. American Megatrends, Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

FOR ADDITIONAL INFORMATION

Call American Megatrends at 1-800-246-8600 for additional information. You can also visit us online at ami.com.

LIMITATIONS OF LIABILITY

In no event shall American Megatrends be held liable for any loss, expenses, or damages of any kind whatsoever, whether direct, indirect, incidental, or consequential, arising from the design or use of this product or the support materials provided with the product.

LIMITED WARRANTY

No warranties are made, either express or implied, with regard to the contents of this work, its merchantability, or fitness for a particular use. American Megatrends assumes no responsibility for errors and omissions or for the uses made of the material contained herein or reader decisions based on such use.

DISCLAIMER: Although efforts have been made to assure the accuracy of the information contained here, American Megatrends expressly disclaims liability for any error in this information, and for damages, whether direct, indirect, special, exemplary, consequential or otherwise, that may result from such error, including but not limited to the loss of profits resulting from the use or misuse of the information contained herein (even if American Megatrends has been advised of the possibility of such damages). Any questions or comments regarding this document or its contents should be addressed to American Megatrends at the address shown on the back cover of this document.

American Megatrends provides this publication “as is” without warranty of any kind, either expressed or implied, including, but not limited to, the implied warranties of merchantability or fitness for a specific purpose. Some states do not allow disclaimer of express or implied warranties or the limitation or exclusion of liability for indirect, special, exemplary, incidental or consequential damages in certain transactions; therefore, this statement may not apply to you. Also, you may have other rights that vary from jurisdiction to jurisdiction. This publication could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. American Megatrends may make improvements and/or revisions in the product(s) and/or the program(s) described in this publication at any time.

Introduction

Both StorTrends® and NetApp® storage servers come with industry leading features making them very strong contenders as external virtualized storage servers. Each has its own innovations and offer almost parallel and compelling feature sets. However, StorTrends® brings in certain distinctions in its storage offering which is not seen in other vendor's solutions. This is a brief document that touches upon some of these innovations that StorTrends® brings in. Because of these innovations, StorTrends® has positioned itself as an unrivaled solution. Here we discuss the merits of StorTrends® over NetApp® in a virtualized infrastructure (VMware®) implementing SQL services.

Storage section alignments

When running Microsoft® SQL server® under a virtual infrastructure like VMware® ESX server®, one has to pay a lot of attention for creating the SAN LUNs and database files. The two top criteria in this regard are partition alignment and selection of file-system block and cluster sizes. When you use virtualized SAN Storage in such an environment, a lot of layers for serving storage capacity come into the picture. Each layer has its own groupings or boundaries which have to be aligned and matched in order to ensure best performance. For example, SAN servers have RAID which work on storage entities called RAID Stripes. Typically the Stripe size is 64KB. On top of this, the Virtualized LUNs may employ certain mapping granularity. For StorTrends this is called a 'chunk' which is also 64KB. VMware® VMFS file system® has a similar allocation unit that is called a 'block'. On top of this resides the Windows® file system (NTFS®) which works with entities called 'clusters'. The next item we need to consider is the SQL I/O page size. Microsoft® SQL servers use 8KB I/O units for reading and writing the database files. In order to extract the best performance, it is very crucial to align these various allocation entities so that they are aligned. Failing this, extra disk I/Os will be generated causing a very perceptible performance hit. The following figures pictorially depict the benefits of having aligned storage boundaries across the various storage layers.

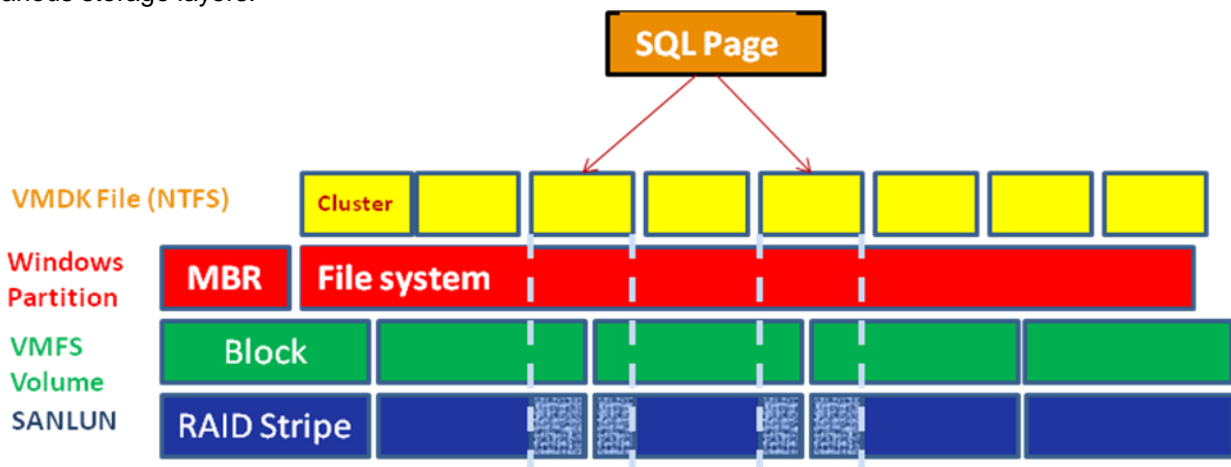


Figure: Effect of misaligned file system and small NTFS® cluster

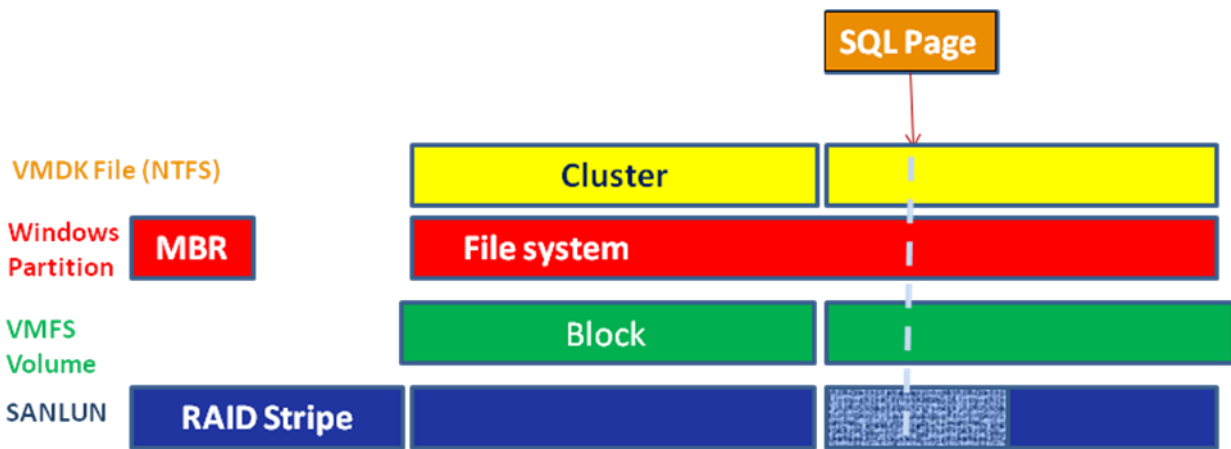


Figure: Merits of aligned file system and large NTFS® cluster

Both Microsoft® and VMware® are very particular about this issue and have independently published their recommendations of ensuring such alignments. In a study published by VMware®, it has been claimed that a properly aligned storage subsystem may offer up to a max of 62% performance advantage. StorTrends assures 64KB storage granularity and, therefore, fits in very well inside a storage hierarchy serving SQL® database to offer the best performance.

NetApp® uses a mapping file system called WAFL® (Write Anywhere File Layout), which employs a block size of 4KB. It is easy to visualize that with usage, an 8KB SQL page I/O may get split into two WAFL® blocks which are not physically contiguous on the RAID disks. This therefore will create extra I/Os, causing a performance hit.

RAID performance

NetApp® offers RAID 4 and RAID – DP for disk fault tolerance. RAID 4 covers single disk failure and utilizes a dedicated parity disk. RAID – DP protects against 2 disk failures and uses diagonal parity configuration with dedicated disks. It is a well known fact that dedicated disks become a bottleneck for random writes. StorTrends offers RAID 5 with striped parity to guard against single disk failure. It also offers RAID 6 with P parity and Q syndrome striped across the disks. This definitely gives a very big performance advantage during random writes.

However, NetApp® uses WAFL® on top of RAID to camouflage some of the performance concerns. It collates the random I/Os into a sequential pattern on the disk layout. This therefore hides the random I/O penalty for RAID 4 or RAID – DP, which would otherwise be very prominent in the absence of WAFL®, but overtime, it may not be possible to place the random I/Os on sequential space due to fragmentation of free blocks for WAFL®. So, with usage, you may gradually see the performance penalties in this architecture.

The other disadvantage for RAID-DP is during a degraded state where two disks have failed. A good understanding of RAID-DP architecture will reveal that this will generate more I/Os during read operations. This, therefore, will hurt the performance in such a degraded environment. Even during normal I/O writes, RAID-DP utilizes more memory bandwidth as opposed to a RAID 6 implementation. Though RAID 6, on the contrary, uses a more compute-intensive algorithm, the RAID controller in StorTrends uses H/W assist to ensure top-notch performance.

Important

The highest performing RAID level for random I/Os is RAID 10. Microsoft also strongly recommends the use of RAID 10 in SQL databases for performance reasons. Though RAID 10 requires more disks than RAID 5 (or RAID 4), it reigns supreme for random write performance. StorTrends supports RAID 10 configuration, whereas NetApp does not support RAID 10 configuration.

Thin Provisioned Storage

Microsoft® SQL 2005 server® and beyond, now support 'instant initialization' for database files. This feature, apart from having various time and performance advantages makes Storage thin-provisioning viable and attractive. Both StorTrends® and NetApp® offer industry leading thin-provision architecture and, therefore, are very suitable for external virtualized storage servers in such deployments. But what sets StorTrends® apart is its rich SRM offering where StorTrends® keeps track of the capacity utilization and offers a capacity planner tool called Outreach™.

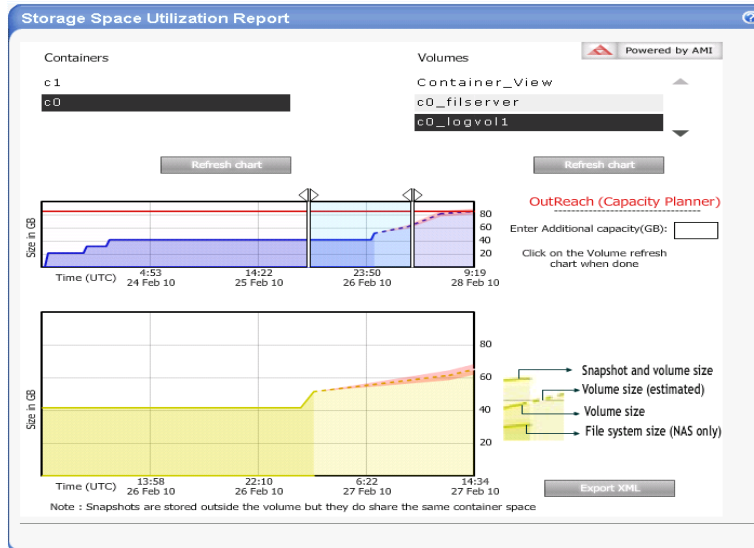


Figure: Storage Space Utilization – StorTrends® Outreach™

It is proven beyond doubt that a capacity planner is essential if you implement thin-provisioning. Most of the Storage servers lack this feature and recommend the use of external 3rd party tools. The built-in capacity planning feature clearly tips the scales in favor of StorTrends® when it comes to thin-provisioning.

WAN Acceleration

Disaster Recovery (DR) is yet another very important consideration for any datacenter. For long distance replications, WAN over Ethernet is widely used. By default the software architecture employed in this context is TCP/IP (note that iSCSI natively uses TCP/IP also). Unfortunately TCP/IP performs very poorly in networks that have high latency and packet drops. WAN, unfortunately, because of the distances involved, is characterized by high latency transport and also incurs a certain percentage of packet drops.

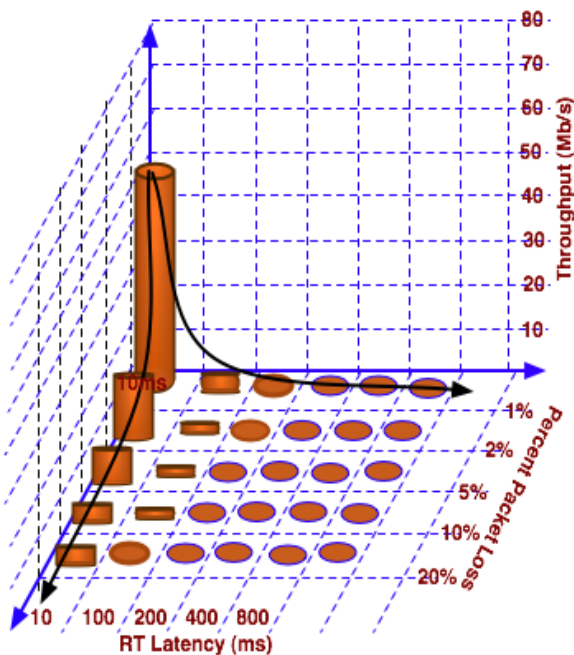



Figure: Throughput with and without WAN Acceleration

To counter this, other vendors use external WAN optimization appliances (Steel head®, Silver Peak® and the likes) at every transport endpoint. This not only adds to the cost (CAPEX and OPEX), but also brings in added challenges for management. StorTrends® natively supports WAN acceleration and de-duplication for WAN replication. These are some stellar features which are not commonly seen in Storage stacks including NetApp®.

Summary

The demands on a datacenter have become more stringent than ever before. With the technology landscape changing to adapt itself to the diverse needs of the day, the IT Industry is witnessing certain game changing strategies. There is a big move towards consolidation and virtualization in every segment of a datacenter. For these reasons, virtualized Application Servers are achieving the state of ubiquity and

this is matched on equal footing with the mass scale deployment of virtualized Storage Servers. The success of a datacenter that employs virtual application servers on shared virtual storage servers largely depends on the strength of the individual servers and how well they mate with and complement each other. StorTrends® outdoes its competition when it comes to the deployment of SQL servers in such a virtualized infrastructure. The NetApp® Storage Server has many features and bells and whistles to tout it as a strong contender, yet it is lacking in some very vital areas. Most noteworthy of these, in the present context, are storage alignment boundaries, RAID support, capacity planning tools and WAN optimization. These features are very carefully thought of and implemented in StorTrends® to give it a clear edge over NetApp®.

Glossary

That glossary makes use of the explanations given within Wikipedia.org. Please see for further information and literature the complete article posted at wikipedia.org including the remarks of the respective author

VMDK File

The **VMDK (Virtual Machine Disk)** file format is a type of virtual appliance developed for VMware products.

MBR

A **master boot record (MBR)**, or **partition sector**, is the 512-byte boot sector that is the first sector ("LBA Sector 0") of a partitioned data storage device such as a hard disk. (The boot sector of a non-partitioned device is a Volume Boot Record. These are usually different, although it is possible to create a record that acts as both; it is called a multi boot record.) The MBR may be used for one or more of the following:

- Holding a disk's primary partition table.
- Bootstrapping operating systems, after the computer's BIOS passes execution to machine code instructions contained within the MBR.
- Uniquely identifying individual disk media, with a 32-bit *disk signature*; even though it may never be used by the machine the disk is running on.

Due to the broad popularity of IBM PC-compatible computers, this type of MBR is widely used, to the extent of being supported by and incorporated into other computer types including newer cross-platform standards for bootstrapping and partitioning.

RAID

RAID, an acronym for **Redundant Array of Inexpensive Disks** or **Redundant Array of Independent Disks**, is a technology that allows high levels of storage reliability from low-cost and less reliable PC-class disk-drive components, via the technique of arranging the devices into arrays for redundancy. Marketers representing industry RAID manufacturers later reinvented the term to describe a *Redundant array of independent disks* as a means of dissociating a "low cost" expectation from RAID technology.^[2]

RAID is now used as an umbrella term for computer data storage schemes that can divide and replicate data among multiple hard disk drives. The different schemes/architectures are named by the word RAID followed by a number, as in RAID 0, RAID 1, etc. RAID's various designs involve two key design goals: increase data reliability and/or increase input/output performance. When multiple physical disks are set up to use RAID technology, they are said to be *in* a RAID array^[3]. This array distributes data across multiple disks, but the array is seen by the computer user and operating system as one single disk. RAID can be set up to serve several different purposes.

SANLUN

As a further detail which differentiates the ESX from other VMware virtualization products: ESX supports the VMware proprietary cluster file system VMFS. VMFS enables multiple hosts to access the same SAN LUNs simultaneously, while file-level locking provides simple protection to file-system integrity.

WAN

A **wide area network (WAN)** is a computer network that covers a broad area (i.e., any network whose communications links cross metropolitan, regional, or national boundaries ^[1]). This is in contrast with personal area networks (PANs), local area networks (LANs), campus area networks (CANs), or metropolitan area networks (MANs) which are usually limited to a room, building, campus or specific metropolitan area (e.g., a city) respectively. WANs are used to connect LANs and other types of networks together, so that users and computers in one location can communicate with users and computers in other locations. Many WANs are built for one particular organization and are private. Others, built by Internet service providers, provide connections from an organization's LAN to the Internet. WANs are often built using leased lines. At each end of the leased line, a router connects to the LAN on one side and a hub within the WAN on the other. Leased lines can be very expensive. Instead of using leased lines, WANs can also be built using less costly circuit switching or packet switching methods. Network protocols including TCP/IP deliver transport and addressing functions. Protocols including Packet over SONET/SDH, MPLS, ATM and Frame relay are often used by service providers to deliver the links that are used in WANs. X.25 was an important early WAN protocol, and is often considered to be the "grandfather" of Frame Relay as many of the underlying protocols and functions of X.25 are still in use today (with upgrades) by Frame Relay.

Academic research into wide area networks can be broken down into three areas: Mathematical models, network emulation and network simulation.

Performance improvements are sometimes delivered via WAFS or WAN optimization.

TCP/IP

The **Internet Protocol Suite** (commonly known as **TCP/IP**) is the set of communications protocols used for the Internet and other similar networks. It is named from two of the most important protocols in it: the Transmission Control Protocol (TCP) and the Internet Protocol (IP), which were the first two networking protocols defined in this standard. Today's IP networking represents a synthesis of several developments that began to evolve in the 1960s and 1970s, namely the Internet and LANs (Local Area Networks), which emerged in the mid- to late-1980s, together with the advent of the World Wide Web in the early 1990s.

CAPEX

Capital expenditures (CAPEX or capex) are expenditures creating future benefits. A capital expenditure is incurred when a business spends money either to buy fixed assets or to add to the value of an existing fixed asset with a useful life that extends beyond the taxable year. Capex are used by a company to acquire or upgrade physical assets such as equipment, property, or industrial buildings. In accounting, a capital expenditure is added to an asset account ("capitalized"), thus increasing the asset's basis (the cost or value of an asset as

adjusted for tax purposes). Capex is commonly found on the Cash Flow Statement as "Investment in Plant Property and Equipment" or something similar in the Investing subsection.

OPEX

An **operating expense, operating expenditure, operational expense, operational expenditure** or **OPEX** is an ongoing cost for running a product, business, or system. Its counterpart, a capital expenditure (CAPEX), is the cost of developing or providing non-consumable parts for the product or system. For example, the purchase of a photocopier is the CAPEX, and the annual paper, toner, power and maintenance cost is the OPEX. For larger systems like businesses, OPEX may also include the cost of workers and facility expenses such as rent and utilities

Why AMI?

Since 1996, AMI has been a leader in the data storage industry, with the inception of the MegaRAID® RAID Controller. AMI was the largest third party RAID Controller manufacturer in 1997, and by 2001 had reached complete market share. At that time, AMI sold its MegaRAID® Division to LSI Logic in 2001. This asset sale allowed AMI to develop the StorTrends® IPStorage product family. The entirety of AMI's many years of expertise in data storage has been rolled into the StorTrends® product. AMI clearly understands the industry's needs and requirements for data storage and has used this understanding to develop a complete line of "best in class" IP-Storage products. AMI looks forward to entertaining any questions regarding the suitability of the StorTrends® IP-Storage Array for the Microsoft® Exchange environment.

AMI offers a wide array of disaster recovery and high availability solutions for your business needs. We provide services that range from storage needs analysis to the design and implementation of a custom disaster recovery solution. We can help your business plan for when things are at their worst while reducing the cost and complexity of your storage environment. For more information on AMI StorTrends solutions, visit www.StorTrends.com, email to sales@ami.com, or call (800) U.Buy.AMI.