


Debunking the Myths: iSCSI vs. Fibre Channel for Microsoft® Exchange Server



StorTrends® **3100i**

© 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.

Original Release: 05/30/2007

Introduction

Observations and predictions from the experts following the data storage market project that the iSCSI install base will grow beyond that of Fibre Channel by 2008. This phenomenon is no accident, as iSCSI has gained its enormous popularity and endorsements on the strength of its solid performance and reliability in the field. In spite of this, some in the industry still continue to claim that Fibre Channel (FC) is “faster” than iSCSI, because of the fact that it has higher native bandwidth (2GB) in comparison to Gigabit (GigE) iSCSI.

This “speed question” seems to be the major lingering controversial factor, and may mislead some to think that FC is therefore twice as fast as iSCSI. Because this stigma still continues to linger to a certain extent, some organizations or decision-makers might think twice about an iSCSI implementation, over concerns that critical applications like Microsoft® Exchange will not be supported properly or could suffer a performance impact if hosted on an iSCSI appliance. The truth of the matter is, however, that the perceived speed difference will be more apparent in bandwidth-hungry applications such as video servers, than it will be for transactional I/O-intensive applications such as Microsoft® Exchange servers. In fact, testing conducted by well-known industry analyst firm ESG demonstrated that only a meager improvement in the range of 5 to 15% increased performance is possible with FC over iSCSI, even for high bandwidth applications (which Exchange is not).¹ This paper will debunk some of the “myths” that surround the supposed unsuitability of iSCSI for Microsoft® Exchange. In particular, it will demonstrate how careful stack design and configuration like that found in the StorTrends® iTX software stack from American Megatrends (AMI) provides a solution that is far superior to Fibre Channel, when performance benchmarks are combined with outlying concerns such as solution cost and implementation.

The Voice of the Industry Speaks in Favor of iSCSI

The ESG study cited above also pointed out another important consideration, namely that for high bandwidth applications where that small performance gain is most apparent, the key factor is disk performance, as opposed to transport performance. ESG actually went a step further and unequivocally documented how a carefully designed iSCSI stack implementation can outperform a FC setup.

A separate survey by ESG in 2007 entitled “The State of iSCSI 2007” showed the most telling results about the feelings of both early-adopters and more recent users towards their adoption of iSCSI, and their confidence in using it in critical applications. First, ESG’s data showed that over 50% of implementations of iSCSI were for Tier I application primary storage, and 36% had implemented iSCSI for Tier I application data.²

iSCSI Implementation Breakdown	
Tier I Application Primary Storage	50%
Tier I Application Primary Data	36%

Table 1: iSCSI Percentage of Tier I Storage / Application Data Implementations, 2007

Moreover, between 43% to 50% of the surveyed iSCSI implementations were earmarked for e-mail administration, and we can be sure that the overwhelming majority of these implementations are running Microsoft® Exchange, due to its strong portion of market share.³ This rate of adoption of iSCSI for critical application administration is remarkable when looking at the short period of time that iSCSI has been in the market in comparison to Fibre Channel.

¹ Lee Dorrier, “Comparison of Storage Protocol and Disk Drive Performance with Microsoft® Exchange Server 2003 Workloads,” Network Appliance, Inc. Publication TR-3514, November 2006.

² Tony Asaro, “ESG Report: The State of iSCSI,” Enterprise Strategy Group, February, 2007, p.2.

³ Ibid, p.3.

iSCSI Email Implementation	
Early iSCSI Adopters	50%
New iSCSI Adopters	43%

Table 2: iSCSI Percentage of Email Implementations, 2007

Additionally, the survey results showed that close to 90% of iSCSI adopters described themselves as either “satisfied” or “very satisfied” with both iSCSI application performance and network performance; clearly, the “voice of the industry” can’t be wrong on this point.⁴ Going beyond survey results describing users’ satisfaction with iSCSI, however, a closer look at the performance issues will reveal the same fact, that the “myths” of FC performance superiority over iSCSI are no more than that, when examined under the clear light of hard data.

Packet Transfer Disparity

Besides misgivings over speed disparities, a smaller, yet equally misguided complaint about iSCSI is that unlike Fibre Channel (which was conceived to handle block I/Os), iSCSI works over Ethernet, which was devised for small packets in the range of 1.5 KB. The general line of thinking is that since more packets mean more interrupts and, by extension, more CPU overhead, iSCSI would seem to be at a disadvantage here. Continuing along this line of thought, the block I/O size for Microsoft® Exchange is 4KB, and stands at 8KB for other transactional servers (such as SQL), which would seem to present a problem for iSCSI because of the higher CPU overhead it would incur.

The storage experts at American Megatrends however, have anticipated this potential problem and have created the ideal response. Their StorTrends® iTX stack offers support for *jumbo* frame sizes of up to 9KB in size (beyond the block size used by Exchange), in order to adequately address this issue. This feature allows for enhanced I/O performance and CPU utilization, to the extent that by merely increasing frame size support alone, the performance differences observed between iSCSI and FC will disappear.

Block I/O Size Comparison	
iSCSI	1.5 KB
Microsoft Exchange	4 KB
SQL	8 KB
StorTrends iTX	9 KB

Table 2: Block I/O Size Handling Comparison by Application / Protocol

CPU Utilization

The StorTrends iTX software stack is a very efficient, performance-tuned stack that uses fine-grained locking and multi-CPU (including Hyper-threading) support to effectively squeeze out every bit of processing power available. In fact, the iTX stack has been tested to reach the I/O throughput threshold that bottlenecks with disk performance with more than 50% of CPU processing power in reserve.

Conversely, if the CPU becomes an issue on the server side, then an iSCSI Host Bus Adapter (HBA) like those manufactured by QLogic, Adaptec, and other vendors can be used to increase performance. However, again recent findings show that for

transactional servers like those running Exchange, this is not the case. The fact that transactional servers do not generally experience CPU bottlenecks is the same reason why server virtualization (using VMWare) has become so popular for these applications.

iSCSI and Replication

For remote replication sites separated by short distances (in the low tens of miles) connected by a GigE link, synchronous replication is an attractive proposition. Here, once again some administrators may feel a reservation that iSCSI may become a bottleneck. After all, each and every write has to be replicated over iSCSI, and even if the percentage of writes is relatively small, one could imagine that the latency during those writes could become a point of concern. In reality, synchronous replication is no problem, since the StorTrends iTX replication stack guarantees low latency and fast resynchronization. If doubts still persist about the suitability of synchronous replication, then AMI can offer asynchronous replication as a feasible alternative. StorTrends features one of the most intelligent methods of asynchronous replication available in the market, namely what AMI calls “snapshot-assisted” or “snap-assisted” replication. With snap-assisted replication only the snapshots or delta changes in the primary device are replicated on the secondary; this powerful mechanism provides the benefits of full replication at a fraction of the network bandwidth and secondary storage capacity required by synchronous replication or other, typical methods of asynchronous replication.

Distinguishing Features of the iSCSI Implementation by AMI

American Megatrends has developed the StorTrends IP-Storage software stack based on the iSCSI protocol to offer superior performance at a cost-effective pricing level. StorTrends was designed to provide users with enterprise-class features at a fraction of the cost of an all-in-one IP-Storage environment. Some of the key differences between the iSCSI implementation from AMI and other iSCSI providers are:

- The StorTrends iTX stack from AMI uses iSCSI implementation at the kernel-level, as opposed to application-level implementations found in many contemporary stacks.
- AMI has specifically designed the StorTrends iTX stack to use *Zero-Memory-Copy* architecture wherever possible.
- AMI has developed an Advanced Caching Module to assist in automatically tuning disk spindles to yield the best possible I/O performance.

Most importantly, the StorTrends iTX stack offers a Multipath I/O Device Specific Module (MPIO DSM), and also supports multiple connections to the same target from an initiator, meaning that multiple GigE links can communicate with the same target. With this capability, bandwidth disparity from FC to iSCSI in the transport speed is a non-issue.

Microsoft® Exchange and StorTrends iSCSI: Truth in Numbers

Clearly, faultless execution of e-mail services is one of the most important applications for organizations today. The management and backup of Microsoft® Exchange Data is a high priority for IT managers and administrators. Therefore storage array setup, configuration and performance are key components of a robust and effective MS Exchange® Server deployment. This section of the document outlines the key indicators of how iSCSI storage is comparable or better than Fibre Channel for Microsoft® Exchange Storage.

Configuration is Critical

To highlight just how important configuration is to an iSCSI-based implementation of Microsoft Exchange®, the following excerpt from the ESG performance study mentioned at the beginning of this document describes a test where a storage appliance from Network Appliance, with either a single 2GB Fibre Channel link or a single 1GBE link, was tested for throughput with merely average results:

“The [NetApp] FAS3050 storage system contained two storage controllers running in an active/active configuration. Each Jetstress server had a single connection (2Gb/sec Fibre Channel or 1Gb/sec iSCSI) to a unique storage controller on the FAS3050. Each Fibre Channel connection supported approximately 190MB/

sec of bandwidth and each iSCSI connection over Gigabit Ethernet supported approximately 90MB/sec of bandwidth.”⁵

In contrast to this limited connection capability, however, the AMI StorTrends iTX stack supports multiple NICs. This means that if 2 GBE NICs are used, there will be a substantial performance enhancement for sequential I/Os (log writes). Furthermore, multiple iSCSI connections will also generate higher performance due to load balancing across the links, and concurrent transport utilization.

The same Protocol Comparison mentioned above showed in performance tests with an Exchange 2003 workload running on Fibre Channel and iSCSI protocols that:

- Both Fibre Channel and iSCSI protocols successfully handled the Exchange workload at all user counts within Microsoft’s latency recommendations with no errors.
- At the lowest user count tested (up to 4,000 mailboxes), the results showed similar performance between the two protocols.
- At the highest user count tested, the Fibre Channel configuration showed lower latencies, resulting in better Jetstress IOPS and throughput results than the iSCSI configuration, although the performance difference was less than 10%.
- The iSCSI configurations consistently showed higher host processor utilization primarily due to the additional processing requirements of the Microsoft iSCSI Software Initiator compared to the Fibre Channel configurations.⁶

Again, it is important to stress here that with AMI StorTrends iTX, thanks to the capability to utilize two GBE links and jumbo frames the difference of 10% noted above will be substantially less. Using Jumbo Frames in iSCSI reduces CPU utilization by substantially lowering interrupts. CPU utilization can be increased even further by utilizing an iSCSI HBAs instead of the Microsoft® software iSCSI initiator.

As the test results indicate, the difference between the FC and iSCSI protocols is extremely minimal, and configuration optimizations can negate the remaining difference. Therefore from a performance standpoint, there should be no concerns regarding the safe and successful implementation of Microsoft® Exchange environments on IP-Storage Arrays utilizing the iSCSI protocol. However, beyond the test results that indicate near-parity on performance, what about “incidental” issues like cost?

Cost Advantage of iSCSI

Clearly, setting performance considerations aside, cost is a major factor in favor of iSCSI. Put simply, iSCSI enjoys a massive cost advantage over Fibre Channel not only for the initial hardware cost, but also for installation, management, maintenance and training (TCO). Major drawbacks of Fibre Channel with respect to cost are the fact that implementations can be very costly and time consuming, and also that dedicated personnel are required to implement and maintain the FC infrastructure. In comparison, with iSCSI the network administrator or system administrator can manage the complete IP-Storage solution due to the familiarity of most IT personnel with Ethernet technology.

Conclusion

To summarize, FC may enjoy some slight performance advantages over iSCSI, but in the case of Microsoft® Exchange implementations, these advantages are not so clearly apparent.

First of all, it is clear that the performance differences between iSCSI and FC can be minimized through careful implementation and configuration of an iSCSI-based software stack like StorTrends iTX. With the chunk size of I/Os utilized by Exchange, and the fact that StorTrends iTX can handle jumbo frame sizes of up to 9KB, latency is not an issue. The ability of StorTrends iTX to support multiple NICs means that the roughly 10% edge in bandwidth that FC once enjoyed has also disappeared. Finally, the cost differential of iSCSI solutions like StorTrends iTX is also a major advantage. One could pay up to \$100K for a FC solution, with the added cost of integration, training, maintenance, and specialized staff, but this choice makes little sense when an iSCSI solution costs significantly less, and as this document

⁵ Lee Dorrier, “Comparison of Storage Protocol and Disk Drive Performance with Microsoft® Exchange Server 2003 Workloads,” Network Appliance, Inc. Publication TR-3514, November 2006.

⁶ Ibid.

shows, offers more than acceptable performance for the task.

Testing and survey results by one of the leading industry analyst firms only serve to underscore the reality of the remarks presented here. It is clear that for critical applications such as Microsoft® Exchange, well-implemented iSCSI solutions present comparable performance to Fibre Channel, and when the cost and maintenance advantages are added to the discussion, the choice of iSCSI is easier than ever.

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® IP-Storage 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.

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.

© 2007 American Megatrends, Inc.

All Rights Reserved



American Megatrends Inc.
6145-F Northbelt Parkway
Norcross GA 30071 | t: 770.246.8600
Sales & Product Information
sales@ami.com | t: 800.828.9264
Technical Support
support@ami.com | t: 770.246.8645
www.ami.com