

## Introduction to IP Storage: Overview, Methods and Advantages of the iSCSI (Internet SCSI) Protocol



---

© Copyright 1998-2008 American Megatrends, Inc.

All rights reserved.

American Megatrends, Inc.

5555 Oakbrook Parkway, Building 200

Norcross, GA 30093

## 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](http://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.

## EXECUTIVE SUMMARY

*iSCSI, the acronym for Internet SCSI (Small Computer System Interface), is an Internet Protocol (IP) -based storage networking standard for linking data storage facilities developed by the Internet Engineering Task Force (IETF). Designed to carry SCSI commands over IP networks, iSCSI is used to facilitate data transfer over networks and manage storage needs over long distances. The iSCSI protocol is one of the key technologies expected to continue to push the rapid development of the Storage Area Network (SAN) market, by increasing the capabilities and performance of storage data transmission, while contributing to the declining cost and expense of storage technologies. Previously, SAN implementations required Fibre Channel (FC) technology and its costly infrastructure of training, support, and proprietary components. Due to the growing cost of FC deployments and the significant management commitment required to keep them afloat, IP-based SAN technology has continued to gain momentum in the industry, making enterprise-level storage accessible to small and medium sized business (SMB) users for the first time. Thanks to the ubiquity of IP networks, iSCSI can be used to transmit data over local area networks (LAN), wide area networks (WAN), or the Internet, and can enable location-independent data storage, backup, and retrieval. IP-SAN solutions are indeed becoming the preferred fast storage access solution for both enterprise and SMB users, thanks to their innovative features, high performance, and reasonable cost.*

## Network Storage Overview

To many individuals both inside and outside the IT world, the term “network storage” can often seem to carry an ambiguous meaning. Is it merely a means of keeping data in a widely accessible network? Is it the archiving of sensitive data in a secure and fault-tolerant system in order to comply with increased regulatory requirements, and the need for oversight? Is it the ability to keep the servers handling mission critical financial or order processing applications performing with “five nines” (99.999%) reliability? Is it the power to perform a quick, complete recovery of compromised data in the event of a catastrophe? As this paper will show, through the proper combination of reliable hardware and robust, elegantly designed software, network storage is all of that and more; moreover, these diverse needs can now be met by a solution contained in a single, affordable appliance. Read on to find out how IP storage solutions leveraging iSCSI technology can do all this and more.

## The Storage Market Today

The IP storage market segment has grown at a significant pace in recent years, for a number of reasons. Low cost and increased optimization of data delivery speed and capability has provided the opportunity for network storage solutions with enterprise-class features to be deployed in a staggering array of different configurations, for organizations large and small. The requirements that organizations are facing, some of them outlined above, have served to accelerate the demand for reliable, expandable storage and file sharing.

Let’s work our way up the food chain, and take a look at the major types of storage solutions that exist for companies today, pointing out the pros and cons of each, particularly as they compare to IP-SAN Storage.

## DAS

Direct attached storage, or DAS, is the simplest and most common type of storage in use, giving it first place at the low end of the storage list in terms of performance, cost, and capacity. Most everyone should be familiar with this type of storage, as it is simply a storage device, whether internal hard drive, or USB or FireWire external hard drive, ZIP drive, etc., attached directly to a workstation, desktop, or laptop unit. Rather inexpensive and easy to operate, these lack complexity, and for the purposes of this discussion, the most important characteristic of network storage, that is, the networked function. Direct attached storage is by its very definition only accessible to the user to whose computer the storage device is attached, meaning that multiple users cannot share the data or applications that reside on the device. In this sense they are “islands of data” and despite their low cost, their value is somewhat limited, beyond increasing the storage capacity of a workstation or laptop machine, as their data repository is isolated.

## JBOD

Shorthand for “Just a Bunch of Disks”, or JBOD, this approach gives an increase in performance by increasing the size of the storage capacity, but does not offer much else. A JBOD configuration still needs additional hardware in order to make it useful, namely to be managed by a server storage head in order to give any kind of file sharing or application management capability. Without a storage head, a JBOD is otherwise pretty much just a glorified DAS.

## NAS

Taking the first steps into the area of managed storage, Network Attached Storage (NAS) appliances provide file-based storage over an IP Network. These devices allow multiple users the ability to access the same files across a network, without interrupting the application server. The benefit of NAS devices is their capability to provide networked, shareable, and fault-tolerant storage in an attractive price range.

Some of the advantages NAS appliances offer include freedom from the limitations of the number of disks a server can hold, the use of well-established file server protocols such as NFS and CIFS, and the ability to make use of existing Ethernet infra-

structure. NAS appliances continue to gain momentum by combining increasingly sophisticated storage server technologies with pre-existing network functionality. SATA and SAS drives are the most common form of storage being used in NAS appliances; combined with dual or even quad Gigabit Ethernet (GbE) ports, NAS solutions have evolved into an efficient and most economical solution to network sharing.

The drawbacks of NAS appliances, however, leave many users looking for more. The main hindrance is that, as file servers, they have no capability for block-level (I/O) data transfer. That means that they are limited to file serving, and not capable of any sort of application hosting. Additionally, NAS devices are somewhat limited in their backup and data protection capabilities in comparison to more sophisticated storage area network (SAN) appliances. Because of this, organizations with greater demands for performance and, most importantly, the ability to connect geographically disparate storage devices to transfer block-level data are compelled to move to the next level, which is a SAN appliance using either Fibre Channel or IP (Ethernet) as the communication level.

## **Fibre Channel SAN**

Now we come to the primary competitor to IP Storage, namely SAN solutions using Fibre Channel. Fibre Channel is a high-speed communication protocol for the transfer of block-level data in storage networks at speeds of up to 4 gigabits per second (Gbps). SAN environments using Fibre Channel are in general specialized, expensive solutions, typically targeted at enterprise-wide implementations in large corporations. For many years, Fibre Channel SANs have enjoyed a large market share in comparison to IP-SAN solutions, due mainly to their high performance (data transfer speed) and reliability, and could be considered in many respects to be the main competitor to iSCSI. However, although they do indeed provide high performance, Fibre Channel SANs are burdened with several drawbacks that often lead to difficult and costly implementations, including:

- Lack of standardization in IP storage implementations (iFC or FCIP question not resolved)
- Cost Commitment
- Requires Specialized, Often Costly Equipment
- Administrators must be Trained to Operate FC Appliances and Switches
- Performance Difference not Readily Apparent as iSCSI Matures

### *Lack of Standardization in IP Implementations*

First, although Fibre Channel solutions have been available for a considerable period of time, the industry has yet to agree upon one open standard when it comes to how to leverage Fibre Channel connectivity over geographically distant SANs networked by TCP/IP connections (in MAN or WAN environments, for example). Though FC solutions offer excellent performance in LAN environments, this drawback has served to limit their usefulness to a degree in these situations. The first solution for extending FC over long distances, FCIP (Fibre Channel over IP, also known as Fibre Channel-tunneling or storage tunneling) can only be used in conjunction with Fibre Channel fabrics. FCIP translates Fibre Channel control codes and data into IP packets for transmission between geographically distant Fibre Channel sites by “tunneling” via IP traffic to create a fabric between two sites. One drawback here is that typically, FCIP is vendor-specific, and intercommunication between FCIP devices from different manufacturers is rare or non-existent. iFCP is intended as an improvement to FCIP by allowing it to communicate with both FC and iSCSI SAN devices, in order to enable data sharing over distance for devices using disparate protocols. As of this writing it is still an emerging standard as different players in the industry seek to leverage their flavor of iFCP as the one to be chosen as the standard.

### *Cost, Equipment, Personnel Commitments*

Also, Fibre Channel networks also require specialized personnel, cables, and network switches to operate. These networks also have physical limitations, in that devices cannot be distributed across network connections exceeding 10 km in length. In particular, this limitation generates problems in implementing replication and disaster recovery strategies to another storage site. Because Fibre Channel solutions are deployed across specialized dedicated networks, these networks must be managed by trained professionals, and require a great amount of planning to operate successfully.

To reduce cost, some companies have reduced costs by using multi-stage switches within their networks. By saving this cost, the Fibre Channel network does not have dedicated bandwidth, and the switches actually become a performance bottleneck reducing the network’s efficiency.

### *Performance Difference Not Readily Apparent as iSCSI Matures*

Thanks in part to the cost and complexity of Fibre Channel, iSCSI has been able to gain ground and increase its number of users. Figure 1 illustrates the increase in market share year-on-year since 200x that IP storage has gained on its costly main competitor, Fibre Channel. As you can see, IP storage continues to gain momentum as more and more organizations realize the advantages it has over the specialized solution of Fibre Channel and its orientation to the top tier of the largest organizations who are willing to absorb the cost and time investment it requires. As a senior analyst from one of the premier industry study groups observed in March of 2007, “IP-based networking, with iSCSI as a fundamental technology, has the potential to raise the level of adoption

to reach the masses.”<sup>1</sup> Several industry magazines report that in 2006, for the first time in their surveys, the majority of their respondents who are thinking about installing networked storage are planning to install iSCSI-based solutions rather than Fibre Channel. Though the numbers vary slightly depending on to whom you are speaking, the idea to keep in mind is that Fibre Channel still has about an 80% market share, but the scales are slowly tipping towards iSCSI as more and more SMBs realize that they need to implement networked storage solutions.

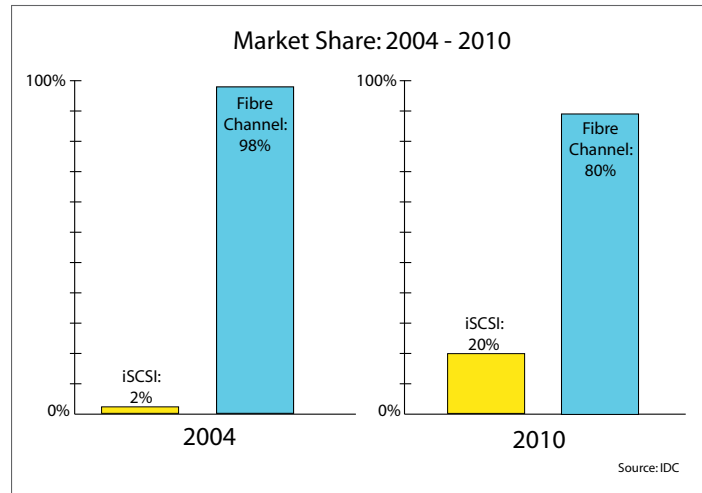


Figure 1: iSCSI Market Share Growth, 2004 - 2010

As Figure 2 below illustrates, an iSCSI-based IP storage area network (IP-SAN) utilizes a company’s existing standard Gigabit Ethernet to create a SAN or NAS environment (or both) with ease, at lower than expected cost when compared to a similar Fibre Channel solution. Initial cost is significantly lower, and operating cost over several years, typically estimated at approximately 30% of hardware cost annually, remains consistently lower as well.

	iSCSI	Fibre Channel
Initial Hardware Cost (IHC):	\$50,000	\$500,000
Annual Total Operating Cost (TOC): [IHC x 30%]	\$15,000	\$150,000
Number of Years of Service:	5 years	5 years
Lifetime TOC:	\$75,000	\$750,000

Figure 2: Sample TOC Comparison, IP SAN vs. Fibre Channel

Figure 3 shows in more detail exactly how iSCSI storage solutions take advantage of an existing Gigabit network to provide the storage companies need. The need to invest in expensive new cabling, training, personnel, or products disappears, because iSCSI technology can provide virtually all of the same features, at a fraction of the cost and time required for deployment.

1 Tony Asaro, "The State of iSCSI 2007," Enterprise Strategy Group, February 2007.




Figure 3: Sample Architecture of an SMB IP-SAN Deployment

### ***The Case for Deploying iSCSI-based IP Storage***

iSCSI-based IP storage is a fast growing alternative to the traditional Fibre Channel SAN environment. Fibre Channel issues are frequently seen as the reason widespread acceptance of SANs has yet to be realized. As noted earlier, with the emergence of IP-SAN, taking advantage of common network hardware and technologies will provide a solution that is less expensive, less complicated, and provide the features needed in a SAN environment. Utilizing iSCSI allows for fewer interoperability issues, and enables users to connect to SANs around the world. Disaster Recovery and Replication can be performed from any location, to provide the security and stability needed in high availability, high reliability data storage.

As described above, iSCSI is one of two main approaches to storage data transmission over IP networks. We have seen the limitations of the other methods, Fibre Channel over IP (FCIP), and Internet Fibre Channel Protocol (iFCP). As pointed out before, but bears repeating, in contrast to Fibre Channel solutions, iSCSI can be deployed over existing Ethernet networks.

### **How iSCSI Works**

Having outlined the other technologies currently competing for market space with iSCSI, let's take a closer look at iSCSI to gain a better understanding of the cost and performance advantages it offers. Leaving out the overtly technical details, the basic iSCSI workflow progresses in the following manner:

When an end user or application sends a request, the operating system generates the appropriate SCSI commands and data request, which then go through encapsulation and, if necessary, encryption procedures. A packet header is added before the resulting IP packets are transmitted over an Ethernet connection. When a packet is received, it is decrypted (if it was encryption was applied before transmission) and disassembled, to separate the SCSI commands and request. The SCSI commands are next sent on to the SCSI controller, and from there to the SCSI storage device. Because iSCSI is bi-directional, the protocol can also be used to return data in response to the original request.

### **Transaction Overview**

The following figure describes the transactions and data flow to and from an iSCSI storage device. This also shows how an initiator issues an iSCSI command, and how the target responds to the command.

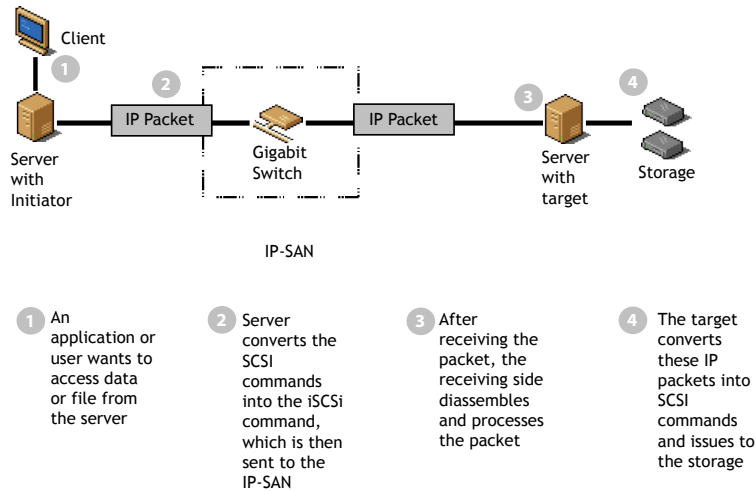


Figure 4: iSCSI Transaction Architecture

To move a bit more deeply into the details of how iSCSI provides data over IP, it is useful to understand the concept of initiators and targets and how they interact in an iSCSI exchange.

### iSCSI Initiator

In the iSCSI transaction model, the initiator is the party requesting that data to be read from or written to storage. The important point to remember is that, as the name suggests, the initiator always initiates the transactions. In read operations, the initiator sends a READ command to a peer, who acts as a target; the target then returns the requested data to the initiator. In a write operation, initiator sends a WRITE command followed by the data packets to the target.

### iSCSI Target

Conversely, a target always processes the commands from the initiator, and cannot independently initiate a transaction. Upon receiving a command from the initiator, a target will execute the command, and then send the requested data, response, etc.

### Protocol Model

iSCSI provides data using TCP/IP over both reliable and unreliable networks. The iSCSI layer in both the initiator and the target respectively encapsulates or reassembles the SCSI data respectively. The transaction uses the familiar OSI model in order to resolve all potential quality of service issues. The following figure outlines the SCSI command block (or CDB transaction) between the initiator and the target in an iSCSI exchange.

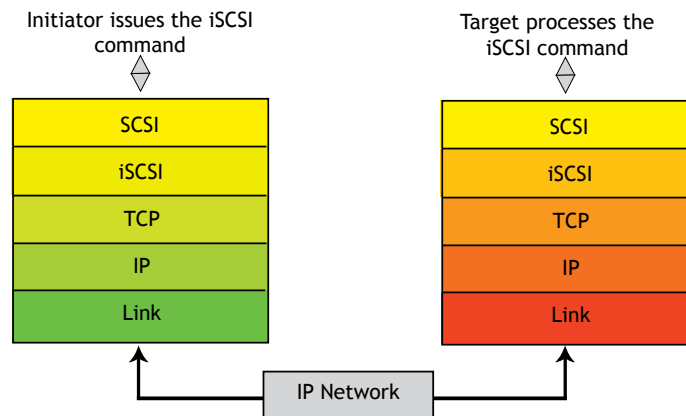


Figure 5: iSCSI Protocol Model

On the initiator side, the host server finds the remote iSCSI target as a SCSI device and issues the SCSI commands to access the storage. These commands are known as SCSI CDB. The iSCSI layer in the OSI model of the server encapsulates the CDB inside an iSCSI packet. This packet is then carried by the TCP, IP and the Link layer and sent to the other side. Similarly the iSCSI layer at the target side reassembles iSCSI packet into the original CDB and issues that through the SCSI layer.

### IP Storage and iSCSI

Having looked in some detail at how iSCSI operates, let us now turn our focus to how IP Storage takes advantage of the simplicity

and reliability of the iSCSI Protocol. As discussed above, iSCSI is a protocol that efficiently makes use of TCP/IP connectivity to enable data transfer between a server and a storage network. It is precisely this capability, combined with the fact that it meets the storage needs of small and medium sized businesses at a price within their means, and provides them with enterprise-level data storage networking capabilities, that makes iSCSI so effective. Having established an understanding of the iSCSI technology that underpins IP-SAN solutions, we can take a closer look at IP Storage appliances, see how they solve critical data management needs, and resolve some of the common storage situations that are made simple with iSCSI.

IP Storage uses the Internet Protocol (IP) over a Gigabit Ethernet link to connect to secondary or remote servers for backup purposes. In most corporate environments today, Ethernet cables are pre-wired in the work area, allowing for an IP Storage Appliance to be plugged directly into the existing corporate network with no costly preparation, conversion, or extra hardware involved. Rather, IP Storage takes advantage of common off-the-shelf hardware and network technologies that make IP-SANs less complicated to deploy than Fibre Channel solutions. The hardware components are less expensive, and because the underlying technology is widely known and used by practically all IT Administrators, there are few interoperability issues. Likewise, training costs are much lower than those involved with implementing a highly complicated, costly Fibre Channel network environment. Furthermore, the ubiquity of TCP/IP networks makes it possible to extend or connect SANs worldwide to multiple remote offices for corporate users. An IP storage network makes it extremely easy and cost effective to deploy an IP-SAN Storage Solution, thanks to iSCSI.

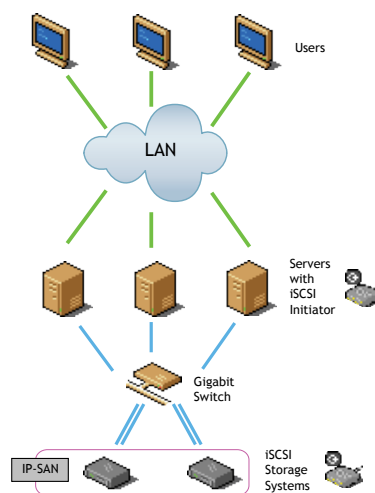


Figure 6: A Sample IP Storage Network Deployment

### iSCSI in Action

Now that we have a good understanding of how IP storage networks can take advantage of iSCSI, it is time to match this with a manufacturer that has the experience and resources to leverage this relationship to optimize IP-SAN devices. The StorTrends® line of IP-SAN storage appliances by American Megatrends makes excellent use of the benefits of iSCSI to provide one of the most powerful and flexible storage appliances on the market, yet at a fraction of the cost. The purpose of the StorTrends line of IP-SAN devices is to provide enterprise-level features to SMB users at a price level they can afford.

### Storage Needs Made Simple with StorTrends

With a single StorTrends IP-SAN storage appliance, an administrator can address the complete storage requirements of today's increasingly complex IT architectures. StorTrends appliances incorporate a wide range of advanced storage technologies that can be easily configured, yet feature enterprise-class functionality normally found only in high-end storage systems.

Talk to any storage professional today, from the SMB to enterprise level organization, and most likely she or he will point out the following issues as the solutions or situations that are deemed today's most urgent required features for any storage appliance or storage software stack:

- Storage Consolidation
- Back-up and Restore Options
- Continuous Data Protection
- Disaster Recovery
- High Availability

### Summary and Conclusion

What makes American Megatrends' line of StorTrends® storage appliances the right choice for IP Storage?

---

AMI relies on standard, proven technologies like Linux to reduce cost and time-to-market (TTM). The defining feature of all StorTrends solutions is the iTX storage software stack found in each StorTrends IP-SAN appliance. Created by the engineering team that developed the award-winning MegaRAID RAID controller, found in virtually every server today, StorTrends' distinction has been the careful and meticulous design of the storage technologies. As mentioned previously, StorTrends iTX runs on Commercial Off the Shelf hardware, thereby guaranteeing the user the latest in hardware technology at the best possible price.

AMI StorTrends storage appliances come in 1U or 3U form factors. The 1U StorTrends 1100i and 1300i can be configured with 250, 500GB or 1TB SAS or SATA disk drives, for storage capacity of 1 to 4 TB, while the 3U StorTrends 3200i can be configured for storage of 4 to 16 TB. Moreover, it is scalable with support for additional JBODs, for total capacity of up to 256TB. These devices can act as NAS and IP-SAN solutions within the same appliance, reducing the number of data servers required in a storage network. With fewer servers to manage, AMI StorTrends appliances help in consolidating storage networks to reduce space and management costs.

#### *Find Value with StorTrends, by Lowering Capital Expenditures & Operational Costs (TCO)*

AMI StorTrends solutions help organizations lower the cost and effort required to support storage resources by consolidating resources, simplifying management and reducing unused spindles. Some of the many advantages of storage consolidation include:

- Higher utilization of available storage
- Use of standard hardware and networking technologies, simplifying maintenance
- Management of more with less, by reducing the time it takes for existing staff to manage your data
- Allowing for heterogeneous connectivity between servers & storage
- Standardization of storage, eliminating multiple points of support
- Elimination of system downtime by seamlessly & non-disruptively adding or expanding storage to production servers

#### *Reduce Complexity*

StorTrends use of iSCSI technology allows existing network personnel manage the new storage environment, foregoing the need for extensive additional training or new storage administrators as with typical FC SAN environments. AMI also deploys Commercial Off-the-Shelf (COTS) hardware, to offer the assurance that the latest developments in hardware technology are in the box. AMI operates on the premise that custom designed hardware has long development cycles, and that many of the features in custom designed hardware can often be provided by creative implementation in software. For example, AMI's ManageTrends® software allows monitoring and control of the StorTrends iTX product, to simplify configuration and set-up, and make quick work of normally complex, time consuming tasks.

#### **Why AMI?**

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 reduce costs and complexity of your storage environment.

#### **How can I contact American Megatrends to find out more about StorTrends?**

For more information on AMI StorTrends solutions, visit [www.StorTrends.com](http://www.StorTrends.com), or contact a StorTrends® Sales Representative by email ([sales@ami.com](mailto:sales@ami.com)) or by calling (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.

© 2008 American Megatrends, Inc.

All Rights Reserved



**American Megatrends Inc.**  
**5555 Oakbrook Parkway, BUilding**  
**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**