

Comparing Fibre Channel, Serial Attached SCSI (SAS) and Serial ATA (SATA)

by

Allen Hin Wing Lam
Bachelor of Electrical Engineering
Carleton University 1996

PROJECT SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF ENGINEERING

In the
School of Engineering Science

© Allen Hin Wing Lam 2009

SIMON FRASER UNIVERSITY

Fall 2009

All rights reserved. However, in accordance with the *Copyright Act of Canada*, this work may be reproduced, without authorization, under the conditions for *Fair Dealing*. Therefore, limited reproduction of this work for the purposes of private study, research, criticism, review and news reporting is likely to be in accordance with the law, particularly if cited appropriately.

Approval

Name: Allen Hin Wing Lam
Degree: Master of Engineering
Title of Project: Comparing Fibre Channel, Serial Attached SCSI (SAS) and Serial ATA (SATA)

Examining Committee:

Chair:

Dr. Daniel Lee
Chair of Committee
Associate Professor, School of Engineering Science
Simon Fraser University

Dr. Stephen Hardy
Senior Supervisor
Professor, School of Engineering Science
Simon Fraser University

Jim Younger
Manager, Product Engineering
PMC- Sierra, Inc.

Date of Defence/Approval

Dec 2, 09

Declaration of Partial Copyright Licence

The author, whose copyright is declared on the title page of this work, has granted to Simon Fraser University the right to lend this thesis, project or extended essay to users of the Simon Fraser University Library, and to make partial or single copies only for such users or in response to a request from the library of any other university, or other educational institution, on its own behalf or for one of its users.

The author has further granted permission to Simon Fraser University to keep or make a digital copy for use in its circulating collection (currently available to the public at the "Institutional Repository" link of the SFU Library website <www.lib.sfu.ca> at: <<http://ir.lib.sfu.ca/handle/1892/112>>) and, without changing the content, to translate the thesis/project or extended essays, if technically possible, to any medium or format for the purpose of preservation of the digital work.

The author has further agreed that permission for multiple copying of this work for scholarly purposes may be granted by either the author or the Dean of Graduate Studies.

It is understood that copying or publication of this work for financial gain shall not be allowed without the author's written permission.

Permission for public performance, or limited permission for private scholarly use, of any multimedia materials forming part of this work, may have been granted by the author. This information may be found on the separately catalogued multimedia material and in the signed Partial Copyright Licence.

While licensing SFU to permit the above uses, the author retains copyright in the thesis, project or extended essays, including the right to change the work for subsequent purposes, including editing and publishing the work in whole or in part, and licensing other parties, as the author may desire.

The original Partial Copyright Licence attesting to these terms, and signed by this author, may be found in the original bound copy of this work, retained in the Simon Fraser University Archive.

Simon Fraser University Library
Burnaby, BC, Canada

Abstract

Serial Attached SCSI (SAS), Serial ATA (SATA) and Fibre Channel (FC) are the main storage interconnect technologies at present. FC is the main interconnect in Storage Area Network (SAN) with a well-known high performance and strong foundation of SCSI. SATA is the most cost-effective and high capacity in the hard drive market. On the other hand, SAS is just picking-up the momentum in both storage and hard drive areas and positions itself as a potential player. All three of these technologies have their own strengths and weaknesses and also their own targeted markets. The following paper will show the analysis of their abilities and the comparison of their strengths and weaknesses in different areas and categories. The main objective is to allow readers to have a better understanding of these storage interfaces and their potentials in their future markets by comparing their performances in different realistic scenarios.

Acknowledgements

I would like to take this opportunity to thank a few important individuals that have taken a great part in assisting me in my M.Eng Project.

First, Professor Dr. Stephen Hardy who has guided me throughout my whole project. He has taken a great time and effort in advising me for the past few months. I would like to sincerely thank him for all the supports and appreciate him in attending my defence even during his sick leave.

Second, I would like to thank both Professor Dr. Daniel Lee and my work supervisor Jim Younger in taking their precious time in attending my defence and reviewing my final paper.

Lastly, I would like to thank my wife Amanda who has mentally supporting me and has tolerated a busy husband during her pregnancy with our first child, Alivia Lam. I would really love to send this special moment to my beloved wife and daughter.

Table of Contents

Approval	ii
Abstract	iii
Acknowledgements	iv
Table of Contents	v
List of Figures	vii
List of Tables	viii
Introduction	1
Section I: Technology Background	2
1.1 Fibre Channel	2
1.1.1 Background	2
1.1.2 Fibre Channel Standards	3
1.1.3 Road Map for Fibre Channel	4
1.1.4 Applications and Characteristics	5
1.2 Serial Attached SCSI (SAS)	8
1.2.1 Background	8
1.2.2 SAS Standards	10
1.2.3 Application and Characteristic	11
1.2.4 Road Map for SAS	13
1.3. SERIAL ATA (SATA)	14
1.3.1 Background	14
1.3.2 SATA Standards	16
1.3.3 Applications and Characteristics	17
1.3.4 Road Map for SATA	18
Section II Comparison of Fiber Channel, Serial Attached SCSI and Serial ATA	19
2.1 Speed and Bandwidth Comparison:	19
2.2 Transmission Distance	20
2.3 Media Limitation	21
2.4 Compatibility and Flexibility	26
2.5 Connectivity and Extensibility	29
2.6 Cost	33
2.7 Performance	35
2.8 Reliability	36
2.9 Lab Test Result	37
2.10 Performance summary	41
Section III Application and Future Potential	42
3.1 Disk Drive Application	42
3.1.1 SATA hard disk	44
3.1.2 FC and SAS drives	45
3.2 Network Interconnect Application	46
3.3 Future Potential	48
3.3.1 Future Potential for Hard Drive Storage	48
3.3.2 Future Potential Network Connection	50
3.3.3 Green power data center	55
3.3.4 Disaster Backup	56

Section IV Conclusion 59
Section V Reference 61

List of Figures

Figure 1: Basic Network Structure [17].....	5
Figure 2: FC connection vs SCSI connection	6
Figure 3: Parallel SCSI connection vs Serial SCSI connection	9
Figure 4: SAS configuration with RAID controllers and expanders connection [14].....	12
Figure 5: SAS Roadmap	13
Figure 6: Typical SATA connection [15]	16
Figure 7: Typical SATA application in desktop PC	17
Figure 8: SATA in SAN application.....	17
Figure 9 - Random jitter for SFP module from different manufacturers	24
Figure 10: Deterministic jitter for SFP module from different manufacturers.....	24
Figure 11: Total jitter for SFP module from different manufacturers	25
Figure 12: Compatibility between SAS and SATA connectors [12].....	27
Figure 13: Compatibility on SAS and SATA on SAS system bankplane [14].....	27
Figure 14: SAS and SATA drives are compatible in SAS network setup [12]	28
Figure 15: SAS system connection work with SAS and SATA drive.....	29
Figure 16: SAS network connection with edge expanders and fanout expanders[14]	30
Figure 17: SATA connection between host and targets [15].....	31
Figure 18: FC ring loop connection [18]	32
Figure 19: FC ring loop connection with hub, Port Bypass Circuit (PBC)[18].....	32
Figure 20: FC switch fabric loop connection [18].....	32
Figure 21: SAS Tx measured jitter at 6Gb/s	37
Figure 22: FC Tx measured jitter at 8Gb/s	38
Figure 23: SAS measured Rx jitter tolerance at 6Gb/s data rate	39
Figure 24: Rx jitter tolerance at 6Gb/s data rate	40
Figure 25: Disk drive market distribution.....	43
Figure 26: FC SAN network with SAS/SATA disk interface [14].....	47
Figure 27: Hard drive market distribution in the past 5 years	49
Figure 28: Disk Drive Characteristics for different technologies	50
Figure 29: I/O consolidation to reduce network interfaces [19].....	52
Figure 30: Network connection with and without IO Consolidation.....	52
Figure 31: Deployment of Converged Fabric integrated with existing FC [20].....	53
Figure 32: FCoE layers mapping	54
Figure 33: Format of FCoE frame encapsulation [21].....	54

List of Tables

Table 1: Fibre Channel Roadmap	4
Table 2: Speed comparison on FC, SAS and SATA	19
Table 3: Transmission distance comparison on FC, SAS and SATA.....	20
Table 4: Transmission distance and speed over different types of cables	22
Table 5: Compatibility on FC, SAS and SATA.....	26
Table 6: Connectivity comparison on FC, SAS and SATA.....	29
Table 7: Cost comparison on FC, SAS and SATA	33
Table 8: Cost for hardware comparison between SAS and Fibre Channel.....	34
Table 9a: Hard disk price comparison	34
Table 10: Performance comparison between FC, SAS and SATA	35
Table 11: Reliability comparison on FC, SAS and SATA	36
Table 12: Summary of performance on FC, SAS and SATA.....	41
Table 13: SATA and SAS/FC disk drive comparison	44
Table 14: Cost of downtime due to natural disaster	58

Introduction

Section I will give a brief background introduction about FC, SAS and SATA in terms of their history, protocol standards and characteristics, applications and future roadmap.

Section II will highlight the main comparisons among FC, SAS and SATA in terms of different categories such as bandwidth, transmission distance, compatibility, extensibility, cost, performance and reliability. In the end of this comparison section, the lab data and jitter measurement performed from one of the 8G FC controllers and 6G SAS controllers chip by PMC-Sierra, the semiconductor chipmaker company, will be analyzed thoroughly.

Following the comparison of characteristics of the three interfaces, Section III will then discuss their performances in each target-markets and applications. In addition, the analysis of the reasons why each dominates their own current market will be revealed. This section will also focus on two different market areas, the hard drive market and storage network market. In conclusion of this section, a specific technology will be recommended in domination of the current market and the other will be seen as faded-out.

In addition to the present market potentials, Section III will also suggest the future market potential for the next generation of these technologies. With FC long transmission distance advantage and the FC protocol integration with Ethernet network, there will be a definite potential in the storage area with higher bandwidth and enhanced transfer distance capability.

The conclusions will be presented in Section IV with the summary of recommendations and potentials for future applications and IT cost savings.

Section I: Technology Background

1. 1 Fibre Channel

1.1.1 Background

Fibre Channel is an interconnect technology for high-performance computer peripherals and networks. It is also used with storage media such as disk drive for important data storage and allows network communication over copper or optical cable. In a network interconnect environment, Fibre Channel transfers data by separating the delivery data from its content to define a mechanism for the transmission of SCSI, TCP/IP and other types of data between two devices. This mechanism allows the data to be transported between two devices without being manipulated or translated between formats. Therefore it is very flexible in transporting different data types.

Today's data communication includes a wide range of applications including network data storage, transaction processing, data imaging or backup, real time network access and server disk drive data transfer. Fibre Channel (FC) can provide a safe, reliable, high performance and high speed data transmission solution in today's data transfer technology. Simultaneously multi-access of data and long data transmission distance are both the main features of Fibre Channel compared to other technologies. FC applications range from small production systems on FC loop to very large systems linking thousands of users, servers, and storage devices into a switched FC fabric network. Its high speed data transfer rate is suitable for server to storage and server to server networking application. Together with FC's dual channels capability, FC is the main interconnect technology for Storage Area Network (SAN) which provides an expandable, high-speed network of storage to IT professionals for multi-terabytes storage and retrieval.

Fibre Channel standards were developed by the American National Standards Institute (ANSI) to overcome the shortcomings of the current SCSI infrastructure, and are used to provide high-speed connections between servers and storage devices.

1.1.2 Fibre Channel Standards

In general, there are five layers to the Fibre Channel standard. Each layer is responsible for a certain set of functions or capabilities. The layers are numbered FC-0 to FC-4 from bottom to top. The following is a brief explanation of the standards and their functions.

- FC-0 - Physical Layer: Defines cables, connectors and the signals that control the data.
- FC-1 - Transmission Protocol Layer: Responsible for procedures such as error detection, maintenance of links, data synchronization, data encoding and decoding
- FC-2 - Framing and Signalling Protocol Layer: Responsible for segmentation and reassembly of data packets that are sent and received by the device. Sequencing and flow controlling are also performed at this layer
- FC-3 - Common Services Layer : Provides services such as multi-casting, striping, encryption or RAID
- FC-4 - Upper Layer Protocol Mapping Layer: Provides the communication point between upper layer protocols (such as SCSI) and the lower FC layers. The FC-4 layer makes it possible for more than SCSI data to travel over a Fibre Channel link

The physical layer consists of copper and fibre-optic cables that carry Fibre Channel signals between transceiver pairs. Interconnect devices, such as hubs and switches define the route for Fibre Channel frames at Giga bit rates. Translation devices, such as Host Bus Adapters (HBA), routers, gateways, and bridges are the bridging between Fibre Channel protocols and upper layer protocols such as SCSI, Ethernet, ATM, and SONET. By conforming to the layer format, products and applications that perform at one layer can be automatically compatible with products and applications that reside at another layer.

1.1.3 Road Map for Fibre Channel

FC Speed	Throughput (Mbps)	Line Rate (Gbaud)	T11 Spec Technically Completed (Year)	Market Availability (Year)
1G	200	1.0625	1996	1997
2G	400	2.125	2000	2001
4G	800	4.25	2003	2005
8G	1600	8.5	2006	2008
16G	3200	14.025	2009	2011
32G	6400	28.5	2012	Market Demand
64G	12800	57	2016	Market Demand
128G	26500	114	2020	Market Demand

Table 1: Fibre Channel Roadmap

The above table [20] reflects the road map for Fibre Channel development in the next couple of generations. 8G FC disk drives and controllers are already available in current market and have been deployed in some of the system networks and applications. The next generation of FC development will have a data rate of 16G bit/sec and will be expected to be in the market in 2011. Fibre Channel specification guarantees at least two generations of forward and backward compatibility, future-proofing storage and as well as providing the best backward and forward compatibility of any data transport.

1.1.4 Applications and Characteristics

Storage Area Networking (SAN) is a term used to describe one of the most popular uses of Fibre Channel. Storage area networking involves enormous terabyte of data transfer and storage networks that are based on Fibre Channel technology known as Fibre Channel SANs.

As shown in Fig 1 [17], a basic networking structure where Fibre Channel SAN covers the connection between storage and servers. Multiple servers are connected in the SAN network where data are transferred and then distributed to a large number of users over Internet Protocol (IP) network such as Ethernet. A typical user would not have a direct connection to the SAN but would access the data stored in the SAN via a server on the IP Ethernet network.




Figure 1: Basic Network Structure [17]

The increase in accessibility and manageability of data offered by the Fibre Channel architecture are the benefits of FC SAN. Fibre Channel switch fabric allows hundreds of storage devices and servers to be connected and also provides highly accessible and available structure for multiple concurrent data transactions at the same time. Compared to the traditional SCSI that was limited to one controller access to device, FC fabric

switch structure has an advantage in accessing the storage device with multiple paths and connections without being affected by failure of other devices. The independence between server and storage enables optimization of storage devices and increases performance over longer distances. Therefore FC switch fabric architecture improves the management of terabytes of data.

Directly attached SCSI:



Fibre Channel Storage:




Figure 2: FC connection vs SCSI connection

Remote Mirroring

Disk mirroring is the process of data written or duplicated to two or more storage devices as backup when it is saved. When one disk fails, an identical copy still exists and is ready for access instantaneously. This process has been with SCSI technology for many years. However, with SCSI's distance limitations, the mirrored data usually remains in the same

room as the primary copy. Fibre Channel has changed the mirroring by its distance advantage with data mirrored and stored in backup devices in a remote area hundreds of kilometres apart [17]. Nowadays, whole storage subsystems and Redundant Arrays of Independent Disks (RAIDs) can be mirrored at remote sites over multi-kilometres to prevent data loss due to localized disasters.

Storage Backup

Nowadays, disk drive backup is a more common and efficient way than tape backup to prevent data loss from human error and viruses. Although the cost for FC disk drive is much higher, companies are still willing to budget for a much more reliable and high performance source for critical data backup. Old technique of server backup over LAN requires at least two dedicated servers, two SCSI buses, most of the LAN's bandwidth and a significant network downtime. This results in great reduction in efficiency and data availability. However, FC SAN technology will avoid these problems due to the independent servers and storage architecture. FC SAN backup will become LAN-free backup when no backup traffic travels on the LAN and in addition to no downtime in the system network.

1.2 Serial Attached SCSI (SAS)


1.2.1 Background

Serial Attached SCSI (SAS) is a technology that carries SCSI forward to new generation that is deployed for high speed data transfer between devices such as hard drives, CD-ROMs, and tapes. SAS uses SCSI protocol, which has been popular and powerful for 20+ years of reliable technology with many enhancements added over the years. Due to the SCSI solid background, SAS benefits storage management, reduction in the risk of storage technology change and increase system interoperability, flexibility and scalability.

Serial Attached SCSI is a serialized operation of SCSI with lots of enhancement features compared to the conventional parallel operation. SAS not only eliminated the use of a parallel bus but also required much fewer physical connections. In serial data transfer, data move linearly in a single path or a pair of cable. Compared to parallel mode, where multiple data streams are packed in the bus and throughput efficiency would be affected as a result of the clock skewing issue. It is because the data streams that start out together on the transmitter side may not arrive simultaneously on receiver side due to variable gating, buffer delays and varying signal path lengths. This will highly degrade the transfer efficiency, accuracy and reachable distance especially at high data transfer speed. Since there are no transmit clock involved in serial transfer operation, the clock skew and asynchronies arrival data issue will then be eliminated and data integrity is preserved.

The SAS standard and protocol was developed and promoted by T10 committee of the International Committee for Information Technology Standards (INCITS) and SCSI Trade Association (SCSITA) respectively. When SAS products entered to the market in 2004, the enterprise storage industry had already been focusing on this new standard and technology due to SAS point-to-point technology with expander architecture. The combination offered high performance, reliability and compatibility in the enterprise storage market. Companies providing the SAS interface promotion are Compaq, IBM, LSI Logic, Maxtor and Seagate.

Parallel SCSI Bus



Point to Point SAS Connections




Figure 3: Parallel SCSI connection vs Serial SCSI connection

1.2.2 SAS Standards

Serial Attached SCSI (SAS) draws attention in both disk drive storage market and storage network markets. Its architecture can fit into most storage environments and other computing environments such as servers and workstations. A Serial Attached SCSI system mainly includes the following components [5]:

Initiator: Initiator initiates task management and device service requests. Generally, the initiator is part of the host computer and devices are the target. The initiator is available as part of the motherboard or host bus adaptor (HBA).

Target: Requests for processing are sent to the target. The target contains logical units and target ports which processes the device service and task management requests. An example of a target device would be the hard disk.

Service Delivery Subsystem: The transfer of data between an initiator and a target takes place in a service delivery subsystem which connected initiator and target with cables.

Edge Expanders: An edge expander allows for communication with up to 128 SAS initiators. Without a fanout expander, a maximum of two edge expanders are allowed in a delivery subsystem. To solve this bottleneck issue, fanout expander will then be required.

Fanout expander: A fanout expander can connect up to 128 sets of edge expanders, allowing up to $128 \times 128 = 16,384$ SAS devices to be addressed. The subtractive routing port of each edge expanders will then be connected to the phys of fanout expander. A fanout expander can only forward subtractive routing requests to the connected edge expanders.

SAS Topology with Expanders

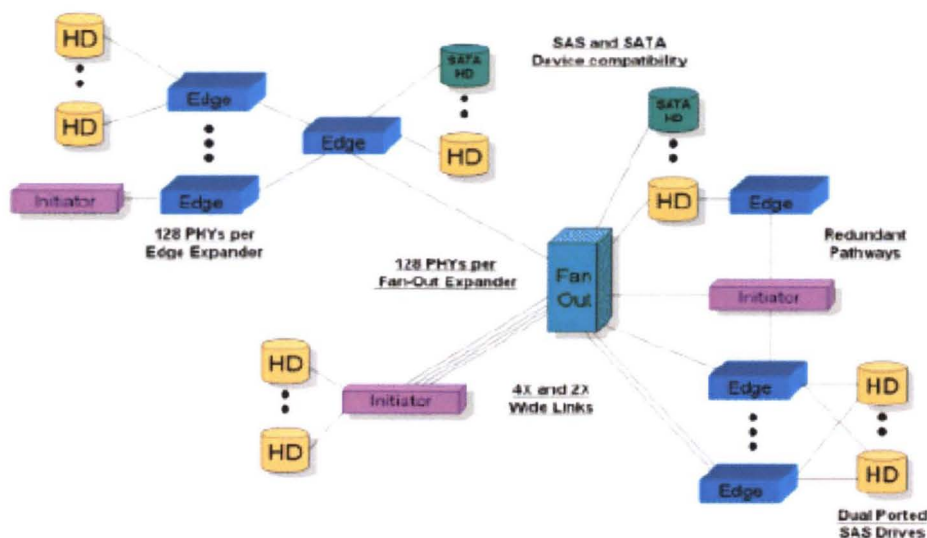


Figure 5: SAS Topology with Expanders [5]

1.2.3 Application and Characteristic

Although SAS has been widely used in the hard drive market for simple or complex RAID arrays setup, SAS has a lot more to offer and may also be applied in storage network applications. SAS also allows operating complex storage topologies as network storage as well as individual hard drives or storage boxes.

When SAS is connected as a storage network or in a SAN environment, the setup will consist multiple disks drives and expanders. The data transactions will then be controlled by the SAS controller. For example with the RAID controllers, any SAS devices in storage arrays can be setup, extended and reconfigured depending on the bandwidth requirements. SAS wide port capabilities allow multiple high-speed physical links to be combined into a single faster high-speed port to enhance the bandwidth of those physical links to a controller. Data will then be transferred to the external network such as Fibre Channel to other networks such as LAN or SAN.

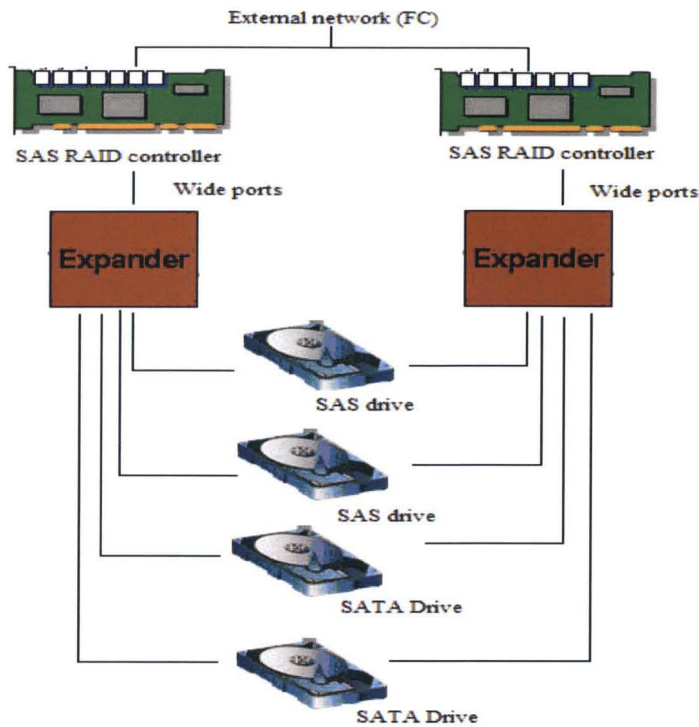


Figure 4: SAS configuration with RAID controllers and expanders connection [14]

In terms of disk drive application, with the SATA Tunnelling Protocol (STP), SAS controller supports the low-cost SATA drives for data backup as well as the high performance SAS drives for critical data storage. The SAS backplane offers the IT manager and the end user customer the choice of drivers. In other words, storage configuration can be effectively done for both low end backup and high end storage in specific application.

1.2.4 Road Map for SAS

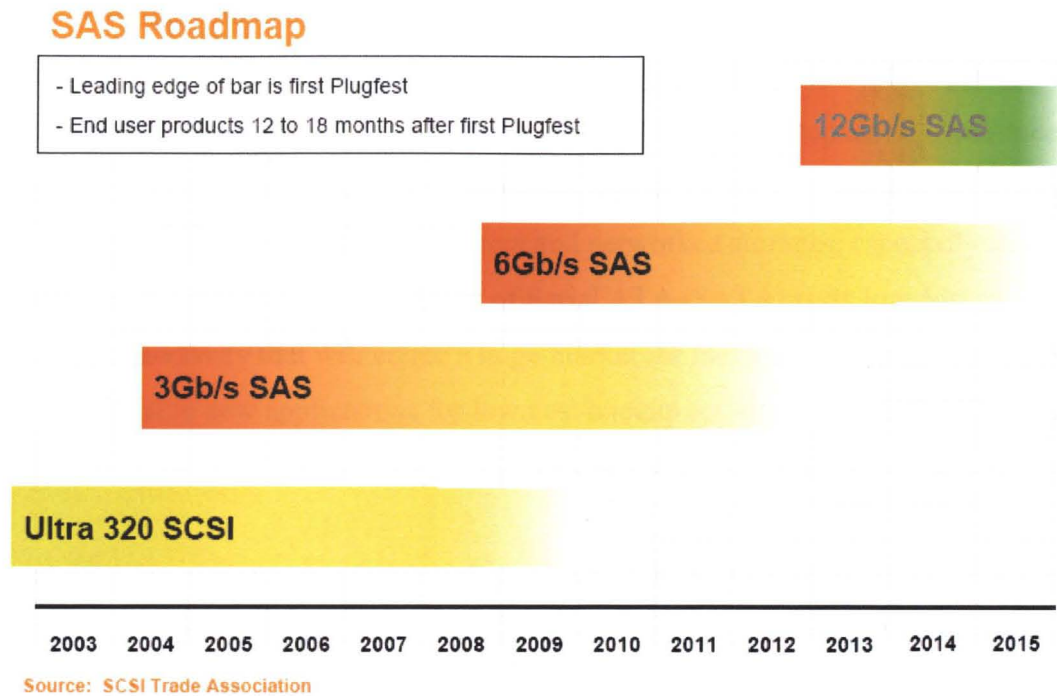


Figure 5: SAS Roadmap

As shown in the above chart [24], 6 G SAS had been already deployed in 2009 whereas the market for SAS application had just began ramping rapidly for the 6G SAS market. The SCSI Trade Association (STA) roadmap calls for 12Gb/sec SAS in late 2012 while the 12G SAS design work will be commence shortly in 2013.

1.3. SERIAL ATA (SATA)

1.3.1 Background

Serial ATA (SATA) is a storage interface technology developed to replace parallel ATA. While parallel ATA has been traditionally implemented as desktop storage interface technology, Serial ATA technology introduces the capabilities not only for desktop usage but also provides an alternative for servers and networked storages, especially when price and cost are the key factors. The vision of Serial ATA (SATA) is its low cost and scalable connectivity that will create a huge market for inexpensive storage solutions as well as to enable new applications for low cost backup solutions such as RAID protected data.

Serial ATA is an evolution of the parallel ATA interface that was developed for hard disk to connect to desktop PCs, servers, and enterprise systems. The transition from a parallel bus interconnect to a serial interconnect was largely driven by the demand for increased data transfer rates. With a 16 bit parallel data bus, issues such as signal crosstalk and skew become significant as data transfer rates increased. Compared to Serial ATA (SATA) serial operation, SATA has a solid position in the desktop industry back in 2001 with a data rate of 1.5Gb/s. With the strong and firm background of ATA technology which firmly grounded in the internal desktop storage industry, serial ATA is positioned to continue development into the desktop, entry-level server and entry-level network storage systems. As a serial connection, serial ATA provides storage vendors and end-users the benefit of point-to-point signalling that provides full bandwidth to device-to-device connection. Hot plug capability, point-to-point connection (vs master/slave in parallel), smaller connectors, simpler cabling, longer reachable distance and cyclical redundancy checking (CRC) are features that are enhanced in the serial version of data transfer than in the parallel mode. The cost advantage and the improve performance of Serial ATA present it as competitive alternative to other technologies.

Today, SATA drive can be found in many entry Network Attached Storage (NAS) products and entry servers. Moreover, there are numerous applications such as backup and near-line storage that SATA drives are providing in inexpensive solutions for video surveillance, near-line storage, high-bandwidth back up and storage backup.

SATA 1.0 spec released in 2001 was developed by SATA working group which was then disbanded as per membership agreement. SATA II specification released in 2002 was developed by SATA II working group. The SATA-IO organization now leads future releases.

1.3.2 SATA Standards

SATA devices include initiators (SATA controllers), port multipliers, and targets (SATA drives) as shown in the figure below [15]. Port multipliers connect initiators to targets in a SATA domain.




Figure 6: Typical SATA connection [15]

Initiators --- SATA initiator is a controller that can be embedded into the motherboard or a host bus adapter (HBA) plugged into a PCI expansion slot.

Port multipliers --- SATA port multipliers require host controllers that are port-multiplier-aware, such as SATA 1.5 Gb/s (with extensions) and SATA 3.0-Gb/s. Therefore, port multipliers are not compatible with original SATA 1.5 Gb/s controllers.

Targets --- In a SATA domain, targets are limited to SATA hard drives, each with a single link port.

Physical size for SATA cables and connectors are significantly smaller than the parallel ATA. SATA connectors use only 25% of the PCB space required by parallel ATA connectors. That will not only provide spacing benefits for more drives in the system, but will also be beneficial for cooling solution in the system.

1.3.3 Applications and Characteristics

In Desktop Application

In a desktop application as shown in figure below [8], the host is the controller card installed in the PC. The SATA drives are the targets connecting the SATA connectors and cables.




Figure 7: Typical SATA application in desktop PC

In Network Application

Storage Area Network (SAN)

SATA drives are also used in enterprise servers and storage arrays, where multiple drives used in each system to increase overall storage capacity. The Serial ATA protocol is used as an interconnection between the host and the drive in the system, and as well as the communication to and from the SATA drives through SATA standard protocol. An application using SATA drives in a storage area network is shown in the figure below [8]




Figure 8: SATA in SAN application

1.3.4 Road Map for SATA

The SATA specification was released in three parts: SATA 1.5 Gb/s, SATA 1.5 Gb/s with extensions, and SATA 3.0 Gb/s. 6.0 Gb/ SATA is the latest standard and fastest spec in the current market. Higher transfer rate for SATA standard is possible but will depend on market demand. For SATA, the electrical interface specifications are defined as follows:

Gen1i: 1.5 Gbps interface for internal host to device applications

Gen1m: 1.5 Gbps interface for short backplane and external desktop applications

Gen 2i: 3.0 Gbps interface for internal host to device applications

Gen2m: 3.0 Gbps interface for short backplane and external desktop applications

Gen2x: 3.0 Gbps Extended length interface for long backplane and system to system applications

Section II Comparison of Fiber Channel, Serial Attached SCSI and Serial ATA

All three interfaces have their own characteristics for their dedicated applicable market area. In order to have an apple to apple comparison, this section of the report is going to analyze and compare their strengths and weaknesses in different categories. Based on direct comparison, it is easier to understand the characteristics for FC, SAS and SATA. Comparing areas include data transfer rate, transmission distance, compatibility, extensibility, cost and performance to determine and recommend the corresponding best-fit application.

2.1 Speed and Bandwidth Comparison:

Speed	FC	SAS	SATA
1st Gen	2Gb/s	1.5Gb/s	1.5Gb/s
2nd Gen	4Gb/s	3Gb/s	3Gb/s
3rd Gen	8Gb/s	6Gb/s	6Gb/s
Next Gen	16Gb/s	12Gb/s	-

Table 2: Speed comparison on FC, SAS and SATA

In term of speed and bandwidth, FC appears to have a slight advantage over SAS and SATA due to higher speed and transfer rate over different generations and standards. The highest data rate standard and specification that is available in the current market for FC and SAS/SATA are 8Gb/s and 6G/bs respectively. In terms of standard road map, FC 16Gb/s data rate will begin the design work in 2010/2011. By just comparing bandwidth and speed, FC is much more successful in this category. However, although the speed and bandwidth for FC are faster than SAS and SATA, the real benefit of FC in storage area network is its impressive distance capability, not the speed. This advantage allows FC to be widely used in storage area networks. More detail comparison on the reachable distance will be shown in the next section.

2.2 Transmission Distance

Media	FC	SAS	SATA
Copper cable	30 meters depends on cable quality	8 meters	1 meter
Optic Cable	Few 100 meter to 10's km depends on optic	N/A	N/A

Table 3: Transmission distance comparison on FC, SAS and SATA

Shown above, SATA systems are designed to support connectivity with 1m cables to a single drive with a host port. SAS was designed to support connections with up to 8m cables, and will also support the use of expanders which allows for connection to multiple SAS drives to a single host port. However, compared to the distance that FC can support with optic media, FC has a huge transmission range advantage, which defeats SAS and SATA in this category.

Fibre Channel is most often associated with fibre optic media for SAN network application, while traditional copper media with twisted pair are used for smaller network implementations such as hard disk to servers data transfer. This is due to Fibre Channel over copper introducing the similar media issues as the other two technologies in short transmission distances.

While copper based media may be suitable for some small network environments, fibre optic cable is the better alternative for larger storage network implementations using Fibre Channel. FC with optic media can support connections up to 10's km apart (single mode fiber) while SAS/SATA with copper cable only support 10's m distance range. Therefore, it would reasonable to run FC SAN area regardless whether a good network connection is successful across various rooms, buildings, towns or even cities.

Recent reports [9] also indicated that FC can extend up to 600km if protocol goes over Gigabit Ethernet network. Throughput capacity of a Fibre Channel system falls very quickly over link lengths greater than 100 kilometers due to the inherent time delay in fibre link length and insufficient buffering in Fibre Channel fabric switches. FC's advantage over Ethernet is merely the cost-effectiveness and its widely availability of Gigabit Ethernet networks that can then be used to interconnect SAN applications. It allows FC data frame to transmit in Gigabit Ethernet network which reduces costs as well as increases the range of options to have data transfer over national rather than regional distances. A more detailed FC over Ethernet application will be described later this report.

Summary of the Trail Test Setup and Results [9]:

“Vtesse and HDS worked with Nishan Systems on the synchronous SAN. The HDS TrueCopy 9900 disk storage system was used to synchronously update an Oracle database over a Vtesse Networks Gigabit Ethernet transmission system, using single-mode fiber supplied by Corning. A Nishan IPS 3300 multi-protocol switch was used to convert native Fibre Channel from the HDS equipment into Gigabit Ethernet and thereby remove the usual link length restrictions of the Fibre Channel format. Timings for a representative Oracle database operation involving 82.5GB of data were measured both for local operation with zero link length and over a range of link distances between 50 to 600 kilometers. Throughput capacity of 42MB/second in local mode and within 83% of that at the maximum link length of 600 kilometers, with timings for the operation of the data replication application increasing by less than 33% over this distance, the companies said.”

2.3 Media Limitation

Optic media can expand the data transmission range in FC protocol in SAN. However, optic media is not constantly a worry-free option as Fibre Channel cabling infrastructures and the distance limitations will occur by continued increases in FC speeds. This is one of the issues that can be easily but inappropriately overlooked in a data center design.

Cables quality and optic modules will directly affect the length of transmission especially at a high data rate such as 8G or even 16G.

As demonstrated by the numbers below [10], a normal short wave SFP (Small Form-factor Pluggable) will drive 500 meters on 50um (diameter) of cable at 1 Gbps, and only 300m at 2 Gbps. When 4 Gbps and faster data rates are available, transmission distances will drop even further due to the short wave length signalling that gets blurred and limits the transmission length as speed increases. If running 62.5um fibre in the cable plant, serious considerations must be made when moving into 8Gb/s. Upgrade from 62.5um to 50um fibre to accommodate the higher data transfer speed may be required. This problem will continually worsen for the deployment of 16 Gb/s beginning in 2011 and scale up to 128 Gb/s in 2020. Therefore, serious considerations must be considered before deployment in cable plant structure and infrastructure design.

Multi Mode 62.5um	
Data speed	Distance
8Gb/s	2m - 21m
4Gb/s	2m - 50m
2Gb/s	2m - 90m
1Gb/s	2m - 300m

Multi Mode 50um	
Data speed	Distance
8Gb/s	2m - 150m
4Gb/s	2m - 175m
2Gb/s	2m - 300m
1Gb/s	2m - 500m

Single Mode 9um	
Data speed	Distance
8Gb/s	2m - 1.4km
4Gb/s	2m - 2km

Table 4: Transmission distance and speed over different types of cables

Fibre optic cable is categorized by its diameter, which is measured in microns, and by its 'mode'. There are two modes of cable. Single-mode fibre with smaller diameter provides higher bandwidth and higher transmission rate, while multimode is able to carry more signal modes by bouncing the signal against the sides of the glass core of the cable. Single mode transfers data up to 50 times more distance than multi mode. However, the cost for single mode is much more than the multiple mode cable and is currently unsupported by most HBA and storage manufacturers due to the default standard of multi-mode in the current market. Therefore the wide usage of single mode cable as the future FC generation media is not an option.

Another concern of optical transmission is the quality of SFP (Small Form-factor Pluggable) Transceiver that could add another parameter in degradation in performance. The quality of SFP varies widely. Different manufacturer will have different performance and frequency response even at the same data rate. If the short wavelength SFP is not a reliable source to handle the required distance at specific speeds, then a longer wave length SFP or optical repeater will be most likely necessary. However, longer wavelength SFP will draw more power consumption and repeater will add cost in the data transmission cost. In a result, they may be unlikely the best solution for FC infrastructure. Below are the diagrams for SFP measured jitter (UI)¹ vs different SFP from different manufacturers at 8Gb/s FC setting. Tests were done on one of the PMC-Sierra 8G product and results showed that some of the SFP failed the FC 8G jitter specification. Therefore SFP module quality will be another serious consideration when set up the FC connection planning especially in high frequency.

¹ UI or Unit Interval is the unit used for jitter which quantifies the jitter in terms of a fraction of the ideal period of the clock (Jitter/Clock period).

Random Jitter vs. SFP




Figure 9 - Random jitter for SFP module from different manufacturers

Deterministic Jitter vs. SFP




Figure 10: Deterministic jitter for SFP module from different manufacturers

Total Jitter vs. SFP




Figure 11: Total jitter for SFP module from different manufacturers

Although SFP manufacturer Gennum Corporation just announced the release of world's first 16G FC SFP [11] which confirms product has been available although module performance will depend on further testing.

Optic media has an advantage in transmission distance length. However, at the same time, it intensifies the limitation with increasing bandwidth and data rate generation that needs to be considered in FC network infrastructure planning.

2.4 Compatibility and Flexibility

	FC	SAS	SATA
Device		SAS and SATA	
Compatibility	Fibre Channel	devices	SATA only
Backward compatibility	8G, 4G and 2G	6G, 3G and 1.5G	6G, 3G and 1.5G

Table 5: Compatibility on FC, SAS and SATA

From the above table, all three protocols showed backward compatibility to their previous 2 or more generations. However, SAS controller has an additional capability to work with either SAS or SATA drives. This promotes SAS to become the most successful in this category as FC and SATA controller can solely operate with their own dedicated devices.

SAS is a protocol that can be used to interconnect between disk drives and host controllers. With SAS's STP (SATA Tunneling Protocol), SAS controllers will still be able to connect to SATA drives that are attached to an expander. Therefore, SAS was designed to be backwards compatible with SATA systems and supports both SCSI, SAS and SATA. However, SAS devices cannot be operated on SATA controller or backplane due to the Serial SCSI Protocol (SSP) being used in SAS. This procedure cannot be reversed.




Figure 12: Compatibility between SAS and SATA connectors [12]

The above diagram demonstrates the SATA device being compatible with SAS controller connector but not in a vice versa environment [12].




Figure 13: Compatibility on SAS and SATA on SAS system bankplane [14]

Both SAS and SATA can be plugged into the SAS controller with system backplane [14]. With the capability to work with both SAS and SATA drive, SAS infrastructure widely increases the compatibility and flexibility in storage system setup. It is solely the user's

preference to either use the reliable and high performance SAS storage or the cost-effective and high capacity SATA storage, or even both.

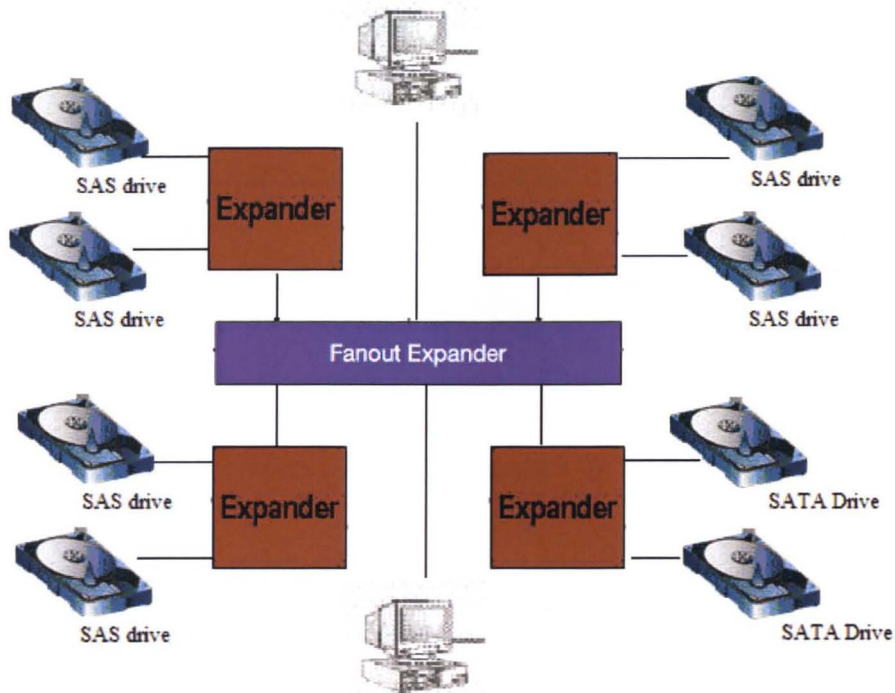


Figure 14: SAS and SATA drives are compatible in SAS network setup [12]

Together with SAS edge expander and fanout expander, both SAS and SATA drives can be used in SAS storage system [12]. This backward compatibility gives the user an option to choose higher end SAS drives or lower end SATA drives if SAS infrastructure environment is being deployed.




Figure 15: SAS system connection work with SAS and SATA drive

The diagram shown above [14] illustrates wide range of applications on SAS drive, whereas FC and SATA controllers or HBA are both limited to work with their own drives

2.5 Connectivity and Extensibility

	FC	SAS	SATA
Max devices supported	127 devices with loop (16777216 with switch)	1 device (16384 devices with expander)	1 device (15 devices with multiplier)

Table 6: Connectivity comparison on FC, SAS and SATA

In addition to high performance and high availability, the storage systems must also be manageable and expandable. Manageability allows setting the appropriate configuration for storage and enabling hardware-related problems to be fixed easily. Alternatively, extensibility allows expandable growth as requirements for bandwidth change. These factors are important for external disk for cluster storage as the external disk arrays are optimized to address these challenges.

In the table shown above, SAS and SATA follows the point-to-point architecture design that allows only one device to be connected to the controller directly. While in connection with expanders and fanout expander, SAS can support up to 16K devices while SATA can support up to 15 devices if multiplier is applied.

A SAS Domain consists of a set of SAS devices that would communicate with one another by means of a service delivery subsystem. Each edge expander can connect up to 128 devices or initiators, while each fanout expander can connect up to 128 edge expanders. As a result, a total of 16,256 devices can be hooked up within a SAS domain.

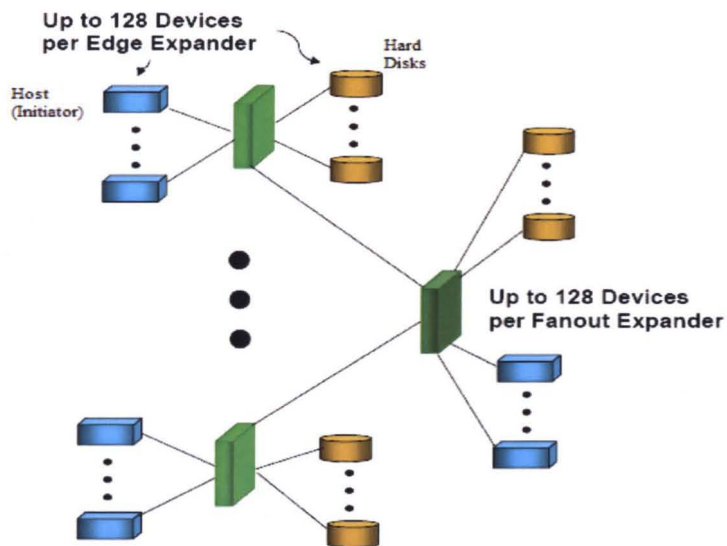


Figure 16: SAS network connection with edge expanders and fanout expanders[14]




Figure 17: SATA connection between host and targets [15]

A SATA system consists of one Initiator host or SATA Controller connected with up to 15 disk drive or target through a SATA port multiplier. This limited expandability has left SATA behind from the other two technologies which also prevents SATA from becoming a main storage network interconnect.

In a FC environment, there are typically 3 different types of connections that are permitted in a FC network structure design. This includes the point-to-point connection, arbitrated loop and switch fabric. The point-to-point is the simplest topology with limited connectivity which is very similar to the other two protocols with direct connection between two devices. Arbitrated loop allows maximum of 127 devices connected in a loop or ring network at the same time. However, adding or removing a device from the loop will cause all activities on the ring to be interrupted and stalled. Any failure of a single device could even cause a disconnection in the ring. Therefore, FC arbitrated loop with centralized hubs are usually added to construct a Port Bypass Circuit (PBC). The advantage of this setup is to bypass the failure port and allows the loop to remain intact. Switch fabric is third FC connection that allows 2^{24} (16 Millions) devices to be connected to FC switches with simultaneous access across multiple ports. Failure ports or devices will be isolated in precaution of affecting the operation of the whole switch network.

Hence, FC is a much preferred option compared to the other two technologies when dealing with huge network of thousands or millions of devices connected in SAN [18].




Figure 18: FC ring loop connection [18]




Figure 19: FC ring loop connection with hub, Port Bypass Circuit (PBC)[18]




Figure 20: FC switch fabric loop connection [18]

2.6 Cost

	FC	SAS	SATA
Cost	High	Medium	Low

Table 7: Cost comparison on FC, SAS and SATA

In both hard disk and network markets, setup cost for FC has always been the most expensive option, whereas SATA is ranked as the lowest cost amongst all three protocols. Compared to SAS network setup, the cost to network a server to the SAN for FC would normally cost three to six times higher than prices in other technologies that are dependant on hardware components qualities. The difference would normally come from the high end FC optical requirements such as HBA server interface, switch ports, and optical cable.

In network interface, FC requires a higher performance HBA as the server interfaces while SAS requires less expensive HBA. Switches are another item contributed to the price difference between FCs, which normally requires the high end optical switch where other technologies do not. In general, the cost of FC switches is normally three to five times more expensive than the SAS switches. Lastly, cabling also highly contributes to the price difference with SAS usually only around 20% the price per meter of optical fibre. As a result, a network of 100 servers with FC would definitely cause in a significant difference in setup cost. Below is a comparison of the network setup cost between SAS and FC [22].

Hardware Component	SAS	Fibre Channel
Server Interface (HBA)	\$200	\$1,500
Switch Ports (advanced/enterprise)	\$3,600	\$12,000
Cabling (SAS) (1Meter)	\$16	-
Cabling (fiber optics) (2Meters)	-	\$96
Total FC : SCSI Cost	\$3,816	\$13,596
Ratio	-	3.5:1

Table 8: Cost for hardware comparison between SAS and Fibre Channel

The above concern also applies to the hard disk market. Although FC hard drives have the highest performance and reliability level, it is as well the most expensive disk compared to SAS and SATA hard disks. SATA always provides the least expensive solution of storage and therefore, SATA drives have dominated the high capacity consumer PC and desktop market. Table 8a showed the price per GB comparison for FC, SAS and SATA hard disk. [23]

Hard Drive	FC	SAS	SATA
Price/GB	\$2.5 - \$5.5	\$1.1 - \$1.7	\$0.24 - \$0.4

Table 9a: Hard disk price comparison

2.7 Performance

	FC	SAS	SATA
RPM	15K	10-15K	5.4-7.2K
Dual Ports	Yes	Yes	No
Full Duplex	Yes	Yes	No
Overhead	Very low	High	High
8b/10B coding	Yes	Yes	Yes
32-bit CRC	Yes	Yes	Yes
Hot pluggable	Yes	Yes	Yes

Table 10: Performance comparison between FC, SAS and SATA

In terms of hard drive performance, SATA drives only operate at 7,200 rpm, which is not even match compared to the performance of FC and SAS drives that runs at 15K and 10K rpm respectively. In addition, FC and SAS's full duplex architecture allows simultaneous bi-directional data and command transfers, which effectively doubling throughput than SATA drives.

FC and SAS dual porting capabilities and their ability to support I/O requests from more than one controller at a given time also enables the design of dynamic load balancing systems. SATA drives only supports single port connections and are only able to connect dual ports if a dual ports switch is connected. However, this only applies to scenarios of one port at a time but not both simultaneously.

Fibre Channel protocol is specifically designed for highly efficient operations using hardware for protocol offload engines (POE's) to reduce overhead for efficiency improvements. Therefore, FC has very little transmission overhead, low-latency switching, and minimal interruptions to the data flow. FC effectively uses almost 98% of the available bandwidth with minimal CPU utilization, while Ethernet with SAS is

perhaps only 30% efficient with more processor overhead. This significantly-reduced overhead supports longer strings and more instructions per host port.

All FC, SAS and SATA use 8b10b encoding which then converts 8-bit data byte to 10-bit transmission characters. The main benefit of 8b/10b encoding is that the clocking signal will be encoded or embedded into the data stream to eliminate the skew problem and proper alignment between data and clock. Another benefit of 8b/10b encoding is to balance the number of 1's and 0's in a row. The special coding would allow special transmission control characters (e.g. K28.5, K28.3) and allows better error detection.

All the hard disk drives are hot pluggable, As a result, they can be inserted or removed without harming the data or the system while the entire system is still powered on.

2.8 Reliability

	FC	SAS	SATA
MTBF	1.4M	1.2M	600K

Table 11: Reliability comparison on FC, SAS and SATA

SAS and FC drives are designed and built for the stress of enterprise use that runs on a 24/7 basis. They both have warranties of 5 years and Mean-Time-Between-Failure (MTBF) ratings exceed 1.4 to 1.2 million hours while SATA only offers MTBF with approximately of 600K hours. MTBF is a common term usually used to indicate the reliability, lifespan of a disk drive and how often a failure will occur.

By comparing, SATA is evidently not as reliable as SAS and FC. With such a high demand and volume for ATA disk drives in the consumers market, it would be impossible to burn in 100's of millions of SATA disks prior to shipment. Therefore, it is still critical for SATA drive to focus on technology improvements to increase reliability and remain competitive. Due to low MTBF rate in SATA, they would not be

recommended in non-fault tolerant applications with mission-critical applications, or extreme environments with excessive vibration.

2.9 Lab Test Result

PMC-Sierra, a broadband and storage chip provider, released their 6G SAS/SATA controllers and 8G FC controllers in 2008. Complete characterization tests were done on both interfaces and results were compared to their corresponding industrial specification. Measurements were done on alpha-point that are closest to high speed link interface of the product. Below are some of the measured results on 6G SAS/SATA as well as 8G FC on both Tx jitter and Rx jitter tolerance.

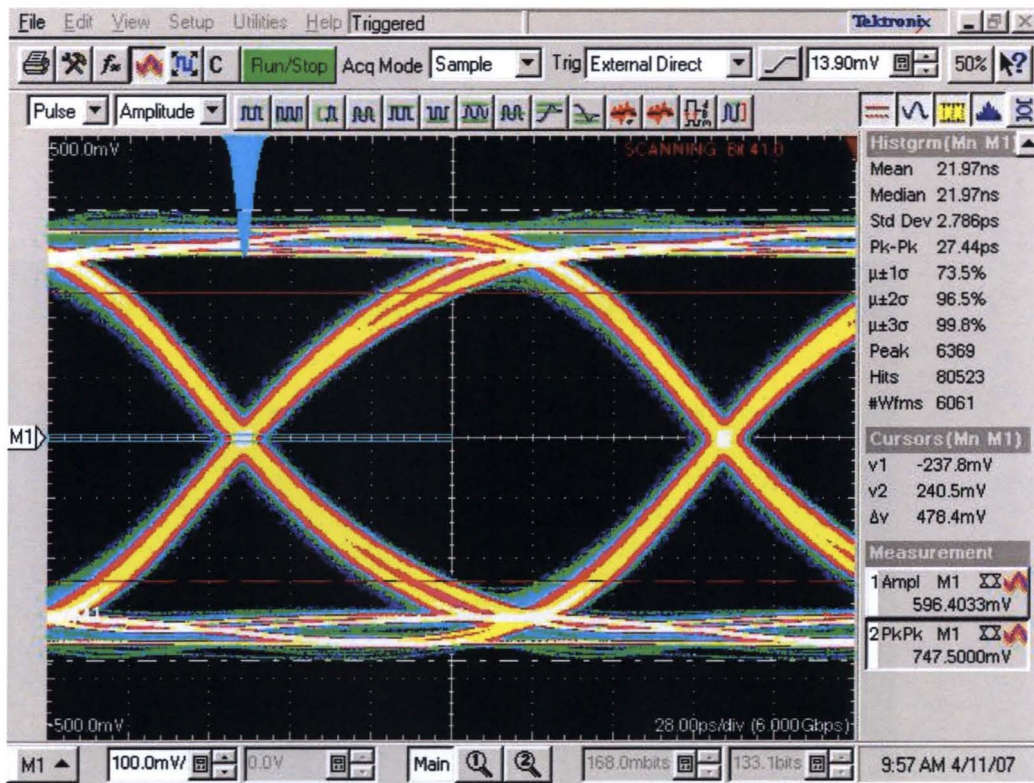


Figure 21: SAS Tx measured jitter at 6Gb/s

Above is the Tx eye diagram for 6G SAS pk-pk differential signal pairs. Measured Tx jitter is around 27.4ps or $0.165UI^2$.

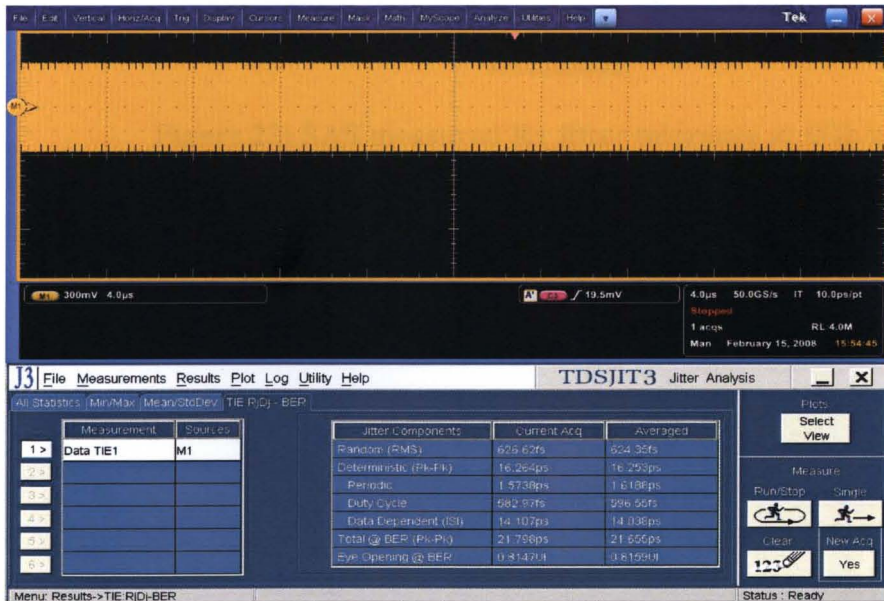
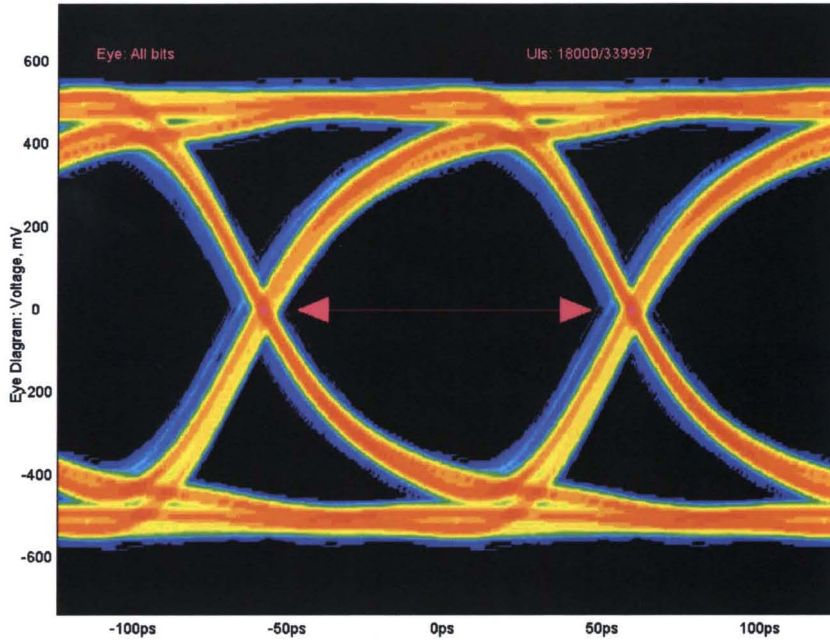


Figure 22: FC Tx measured jitter at 8Gb/s

² UI or Unit Interval is the unit used for jitter which quantifies the jitter in terms of a fraction of the ideal period of the clock (Jitter/Clock period).

Above is the measured Tx jitter and eye diagram for 8G FC. Measured Tx jitter is around 21.79ps or 0.186UI

Rx Jitter tolerance test were performed on 6G SAS with the worst frequency respond at around 10.5M with tolerance of 0.2UI.

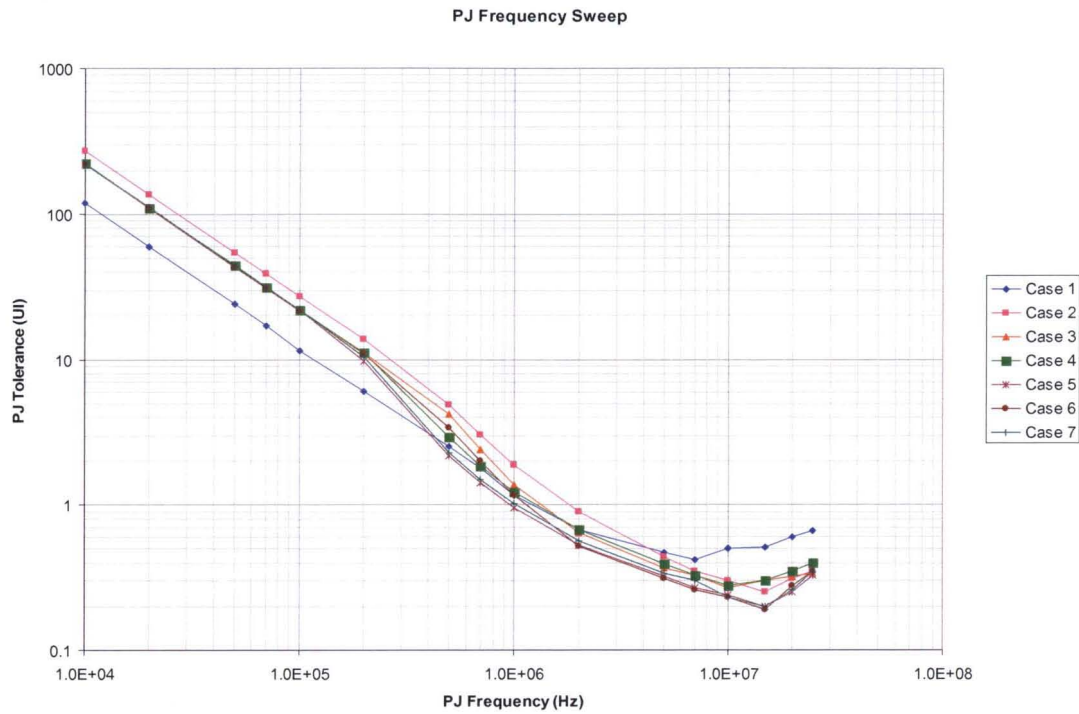


Figure 23: SAS measured Rx jitter tolerance at 6Gb/s data rate

The following diagram is the Rx jitter tolerance on 8G FC and the worst frequency response happened at around 10.5M and with worst tolerance of ~0.3UI.

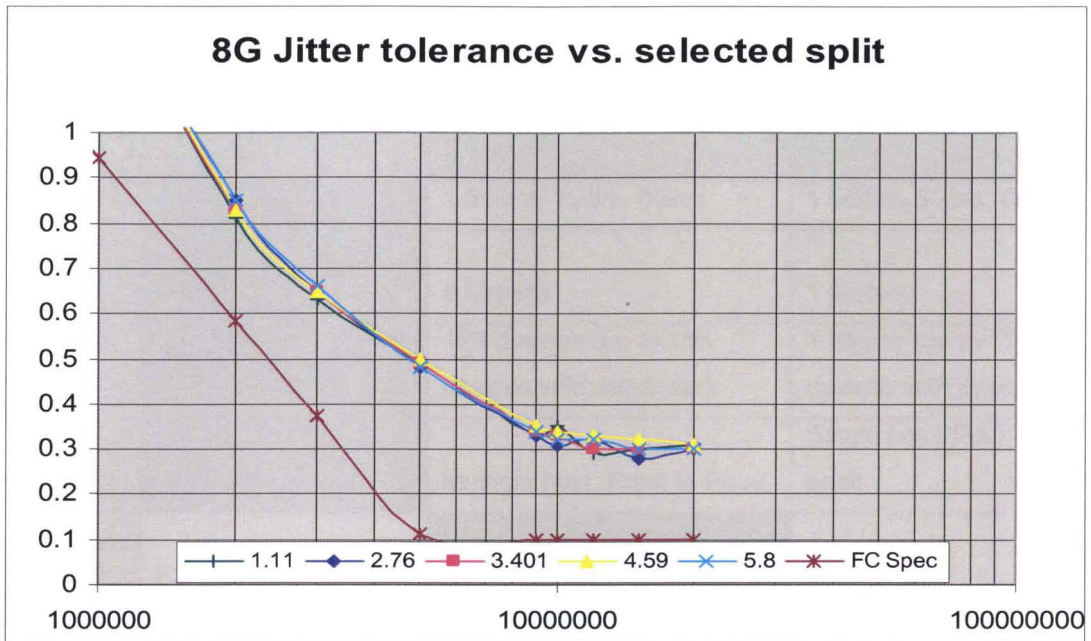


Figure 24: Rx jitter tolerance at 6Gb/s data rate

In general overall results, the Tx jitter and Rx tolerance data showed both are comparable and no significant advantage over one another between PMC 6G SAS and 8G FC controllers.

2.10 Performance summary

	FC	SAS	SATA
Bandwidth	2Gb/s, 4Gb/s, 8G/s	1.5Gb/s, 3Gb/s, 6Gb/s	1.5Gb/s, 3Gb/s, 6Gb/s
Cable Length	10's Meter - 100's Km	8 Meters	1 Meter
Scalability	up to 16M devices with fabric switch)	126 devices (up to 16K devices with expander)	1 device (up to 15 devices with multipler
Connectivity	Multiple host, Loop or switch fabric	Multiple host, Point to Point	Single host, Point to point
Device compatibility	FC only	SAS and SATA	SATA only
Performance	Dual port	Dual ports	Single Port
	Full duplex	Full duplex	Half duplex
	Hot swappable	Hot swappable	Hot swappable
	8b/10b encode	8b/10b encode	8b/10b encode
Reliability	1.4M MTBF	1.2M MTBF	600K MTBF
Cost	High	Medium	Low

Table 12: Summary of performance on FC, SAS and SATA

The above table is a summary comparison in the categories we have discussed in the previous sections. The highlighted cells are the winners in each of the categories which showed both the strength and advantage on the selected technologies.

Section III Application and Future Potential

The purpose of this paper is to analyse and compare the different areas and applications of FC, SAS and SATA, and to also recommend the potential players in the future or next generation of data storage and data transfer technology. Given that all FC, SAS and SATA have their strengths and weaknesses in different areas and applications, it would be most reasonable to compare in the area of Hard Drive Storage and Enterprise Network Storage application

3.1 Disk Drive Application

In today's economy, cost-effective, performance-oriented, scalable and reliable storage solutions are fundamental to success. Today's IT managers as well as storage vendors need to keep their competitive edge by continuing to reduce cost of their operation and solutions. However, how to keep the balance between performance and cost and to choose the best technology and setup in the specific application is the critical factor.

In terms of performance in the disk drive market, FC drive has always been the top pick even if not all storage in market requires such a high performance drive especially in the high demand and volume customer markets. A "good enough" SATA drive with reasonable pricing would probably be a more important factor to end-users in the market. For those who are looking for high performance but with a budget constraint, SAS would probably be the alternative choice.

FC Disks: Used for online, mission critical applications in large enterprises where performance is a major priority coupled with high availability and reliability.

SAS Disks: SAS is used as a high performance and more affordable alternative to FC. Used for online, mission critical applications in mid-size enterprises network where performance is still a priority coupled with high availability and reliability. SAS can also be mixed-and-matched with SATA drives, allowing a flexible in storage budget control for IT manager.

SATA Disks: Not suitable for mission critical applications where performance is a priority. SATA has higher capacity compared to FC and SAS and is used as a near-line storage. In server market, SATA drives should apply in application with low-workload entry level servers and multi-drive storage configurations, such as JBOD or RAID. Due to its low cost high volume advantage, it is a good choice for PC and desktop market as well as disk backup or reference data, low-cost storage and video surveillance applications.




Figure 25: Disk drive market distribution

From the diagram shown above [16], the momentums of SAS and SATA drives have been much faster and stronger than the FC drive in the past 4 years. The trend for SAS and SATA disk drives usage have been continuously growing rapidly while FC drive has been increasing slightly or even remaining flat.

The table below summarizes the main differences between Desktop SATA Drive and Enterprise SAS or FC drive [12].

Drive Comparison Table	Desktop SATA	Enterprise SAS or FC
Performance (Access to Data)		
Latency + Seek Time	13msec @7200 PRM	5.7 msec @ 15K RPM
Command Queuing and Reordering	LBA based	LBA and RPS based
Rotational Vibration	5 to 12 rad/sec/sec	21 rad/sec/sec
Typical I/O per sec/drive (no RV)	77	319

Typical I/O per sec/drive (10 rad/sec)	35	319
Typical I/O per sec/drive (20 rad/sec)	< 7	310
Duplex Operation	Half	Full
Customization		
Unique Code and Hardware	Limited	Extensive
Variable Sector Sizes	No	Yes
Mode Page Parameter Control	No	Yes
Inquiry Data	No	Yes
Diagnostic Pages	No	Yes
Capacity Controls	No	Yes
Indicators		
Activity LED	No	Yes
Fault LED	No	Yes
Reliability		
MTBF	600K Hrs	1.2M Hrs
Duty Cycle	8x5	24x7
Interactive Error Management	No	Yes
Internal Data Integrity Checks	No	IOEDC
Dual Port	No	Yes

Table 13: SATA and SAS/FC disk drive comparison

It is quite easy to differentiate or compare between the high capacity SATA disk and the high performance SAS/FC disk. Below is a quick summary of the SATA disk:

3.1.1 SATA hard disk

The most obvious benefit of using SATA hard disk is its minimal initial hardware costs. In addition, it is also the most sensitive and critical factor in customer or end-user market. These are the main reasons why SATA drives have a good position in the PC or desktop market that has been always “good enough” to handle those non-critical mission and control disk backup. In enterprise storage, SATA starts to draw more attention due to the RAID application. In RAID application, the RAID redundancy will take over drive

failures and compensate for the lower reliability of SATA drives while the hot spares disk will kick in automatically to replace the failed units. Therefore, SATA drive has been gaining momentum in the enterprise storage with RAID application.

However, in terms of performance compared to SAS/FC disk, SATA is still left behind not only due to the low RPM rate that reduces efficiency, but also due to the half duplex data transfer which only allows single sequential. In addition, SATA's low reliability with high MTBF rate, lack of error recovery management on disk and lack of internal data integrity check also contributes to being behind in the performance comparison to SAS/FC disk. All these factors will increase the system down time possibility and as well as the possibility of data lost after disk failure. The more critical the application, the more significant the downtime and the greater the impact on total cost and value. Therefore, for critical mission and high performance required applications, FC and SAS drives would still be the better alternatives.

3.1.2 FC and SAS drives

The comparison between SAS and FC drives will definitely be much tougher. Both drives offer high performance and high reliability which are suitable for critical mission and important data backup.

FC technology is mature and has a reputation of being best in class for performance and reliability. Furthermore, the FC interface provides a slightly higher data transfer rate advantage. In SAN enterprise storage, FC drive and network connection is strong and moderately dominates due to its superior performance and high reliability. However, SAS drive has moved closer into this storage market recently with similar high performance and reliability to meet the essential requirements for storage systems as well as cost savings for users and solution providers due to SAS high performance/cost ratio. This is where SAS comes into the game for enterprise storage. The SAS 2.0 interface boasts a 6Gb/sec data transfer rate, which is close to 8Gb/sec FC drives that are available today.

3.2 Network Interconnect Application

As discussed, SAS drive has led FC drive in storage due to its comparable high performance, high reliability, scalability and its cost-effectiveness. However, the scenario is much different in storage network connection.

The fundamental problem of how to connect terabytes of information to hundreds of servers reliably has been solved by SAN. Direct attached storage such as SAS or SATA networks operates well for a few servers with a few hundred gigabytes of storage. However, to manage and back up terabytes of storage data would be a nightmare. With exponential storage growths and 24/7 operations, Fibre Channel SANs would still be the most reliable solution for large storage networks. Fibre Channel connection would be most suitable for extremely large systems linking thousands of users, servers, and storage systems into a switch Fibre Channel network. Over 90% of all SAN (Storage Area Network) installations throughout the world are FC networks.

Here, it should be emphasized that FC storage infrastructure are not limited to solely working with fibre channel disks. Certainly FC disk will have a direct attach connection if FC network interface is used. In fact, SCSI, SAS or SATA disk drives for storage could be used in FC network provided that a protocol bridging is applied. This is due to the fact that FC separates the delivery of data from the content to define a mechanism for different protocols and commands transmission. Through bridging, data in storage devices with protocol other than FC can be transmitted in FC network structure for storage connectivity. Interconnecting at the network (switch) level would then be handled by FC, whereas a converter in the controller (bridge) would handle the SATA/SAS translation to the drive. Therefore, FC can do a better job in network interface while SAS and SATA are both well-positioned and most suitable to be recognized as a network storage server and end-user disk storage.




Figure 26: FC SAN network with SAS/SATA disk interface [14]

3.3 Future Potential

3.3.1 Future Potential for Hard Drive Storage

Primary storage systems that stored business-critical information-data with the highest value and importance still prefer Fibre Channel disk drives even at a higher price. This data requires continuous availability and typically has high-performance requirements. Business-critical data will continue to be stored on Fibre Channel-based disk arrays.

However, with the high cost in FC disk, increasing storage capacity and shrinking IT budgets, more and more organizations have started turning towards alternative technologies in SAN storage. This is where SAS and SATA technologies come into the picture in enterprise storage.

Fibre Channel as a disk drive interface is expected to level off in volume owing to the rise of both SAS (at the high end) and SATA (at customer end) beginning in 2007. Only a few disk drive manufacturers are committed to producing Fibre Channel drives. Almost every drive vendor is offering both SAS and SATA to increase its competitiveness. Although the 8G FC drives are already available in the market with a higher transfer rate than 6G SAS/SATA drives, hard disk manufacturer such as Fujitsu, Hitachi and Seagate predict the demand for 8G FC drives market is shrinking compared to SAS drives. Storage system and server vendors are using 6G SAS as the replacement for 4G FC and do not predict a move to 8G FC drive. By end of 2009, SAS cut into FC market share will provide faster connectivity at a lower cost than FC drives. SAS drive will possibly replace direct-attached FC and replace FC in external RAID systems.

At one of the SCSI Trade Association (STA) event earlier this year, Dell, LSI and Seagate have already demonstrated 6Gb interoperability with Dell servers, LSI SAS RAID-on-chip and expander components, and Seagate 6Gb SAS and SATA hard drives.

The expectation of 6G SAS server and drive will be extensively shipped to users in late 2009 [13].

As shown in the diagram below [3], although the demand for FC drive picked up earlier than SAS and SATA, FC's demand had remained flat for the last couple years and is projected to continue to be flat. Alternatively, SAS and SATA drives have been growing rapidly since 2006 and continues to show a trend of continuous growth in the near future.

Enterprise Class HDD Interfaces

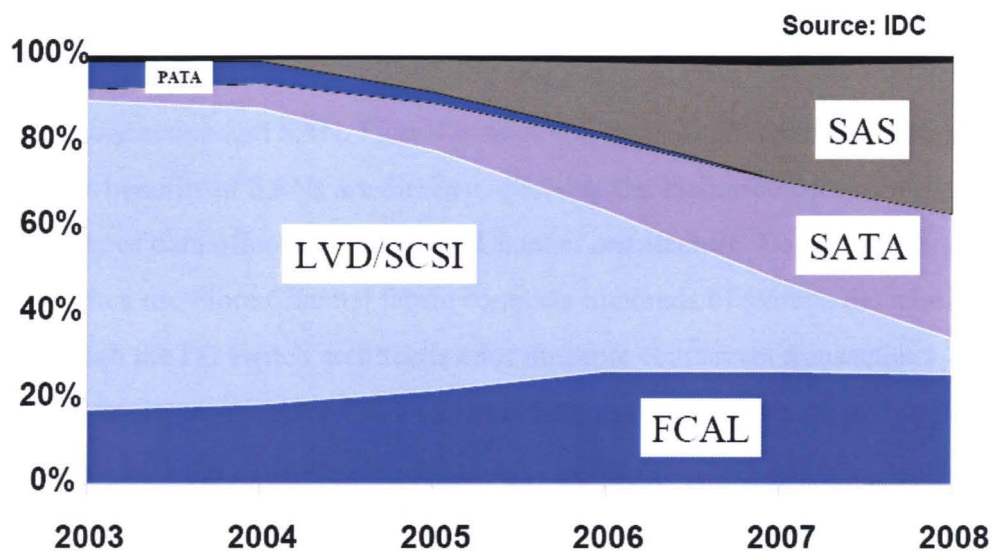


Figure 27: Hard drive market distribution in the past 5 years

In short, the SAS architecture enables customers to incorporate with either SAS or SATA hard drives and is thus suitable for either high-capacity SATA or high-performance SAS environments - or both at the same time. In addition, SAS drive is able to achieve an optimal balance between capacity, cost, reliability, and performance. Hence, SAS will be the leading future player in the hard drive storage market [3].

SATA		Fibre Channel		SAS	
High Performance	Reliability	High Performance	Reliability	High Performance	Reliability
Scalability	Low Cost	Scalability	Low Cost	Scalability	Low Cost

Figure 28: Disk Drive Characteristics for different technologies

3.3.2 Future Potential Network Connection

In network connection and SAN, Fibre Channel will most likely continue to take the lead position. The benefits of SANs are directly related to the increased accessibility and manageability of data offered by the Fibre Channel architecture. Data becomes more accessible when the Fibre Channel fabric connects hundreds of storage devices and servers through the FC switch architecture for multiple concurrent transactions. The strength and future potential for Fibre Channel is to overcome distance limitations when Fibre Channel links span hundreds of kilometres or are sent over a WAN. An additional reason of having interconnect with FC is its well developed and matured infrastructure in terms of both hardware and software. FC's superior switching capabilities provided by FC protocol will also overcome the other comparators, SAS and SATA.

FC over Ethernet

A FC future potential is to combine the FC architecture together with Ethernet network to become a Fibre Channel over Ethernet network that allows FC packets to run in 10 Giga Ethernet network. In today's datacenters, companies normally maintain both Ethernet for TCP/IP networks for local area networking (LAN) and Fibre Channel (FC) for storage area networks (SANs), each dedicated to specific purposes. Ethernet networks are generally implemented when end-users need to transfer relatively small amounts of information over both local and global distances devices. Storage area networks are implemented for access to block I/O for applications such as mail servers, file servers and

large databases. FC over Ethernet is able to combine the two individual networks onto a common Ethernet infrastructure, hence cost saving.

One of the benefits for FCoE is I/O consolidation which allows user to use multi-function network/storage adapters CNA (Converged Network Adapters) to replace single-function network and storage-specific cards, thereby reducing the number of server slots and switch ports. At the server level, when using a FC driver, the CNA functionally represents a traditional FC HBA or initiators just as if that same server were connected on a SAN. Alternatively, when the NIC driver is being used, the CNA will also functionally represent a traditional networking device just like as if it is connected to LAN. The enhanced server CNS (Converged Network Servers) contains fewer server slots and switch ports, thus will reduce the number of power consumed for I/O and reduce cooling overhead required, which will then result in cost and power saving (Fig.31) [19]. In addition, FCoE allows Internet Protocol (IP) and Fibre Channel network traffic to be carried over existing drivers, NICs, and switches that simplifies network topology by reducing cabling cost and complexity, reducing I/O adapter cards, switches and HBA, but remaining the same security and traffic management offered by FC. As shown in Fig. 32 [21], the diagram on the left shows a current data centre with servers containing separate Ethernet NIC and FC HBA for separate Ethernet and FC switches, whereas the diagram on the right shows an implementation of CNA connected to FCoE switches or CNS. Both diagrams demonstrate how CNA reduces the number of adaptors, cables and switches in FCoE network.




Figure 29: I/O consolidation to reduce network interfaces [19]

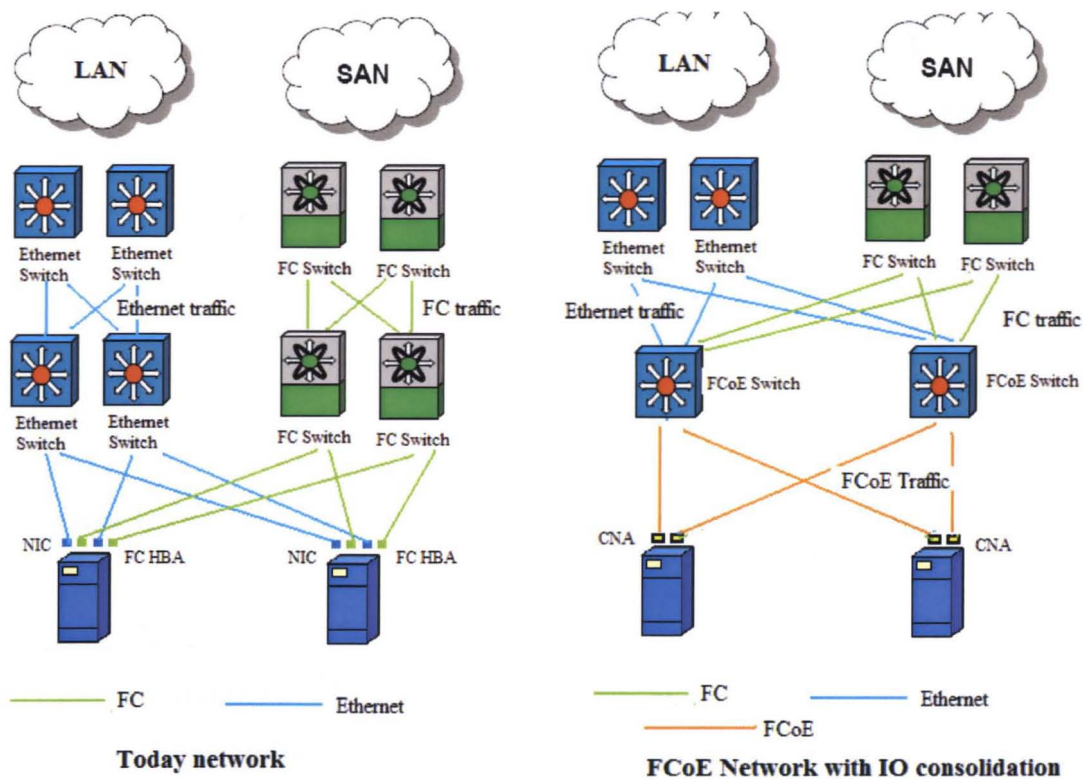


Figure 30: Network connection with and without IO Consolidation

Fibre Channel over Ethernet extends rather than replaces Fibre Channel, allowing companies to integrate their Ethernet and Fibre Channel networks. FCoE's compatibility will leverage the economics of Ethernet networks while preserving the infrastructure of

the existing Fibre Channel storage management framework. The same management tools that customers use to manage and maintain their SANs today will be used in an FCoE environment. Therefore, FC investment would be protected. The diagram below showed the deployment of converged FCoE fabric integrated with existing FC SAN.

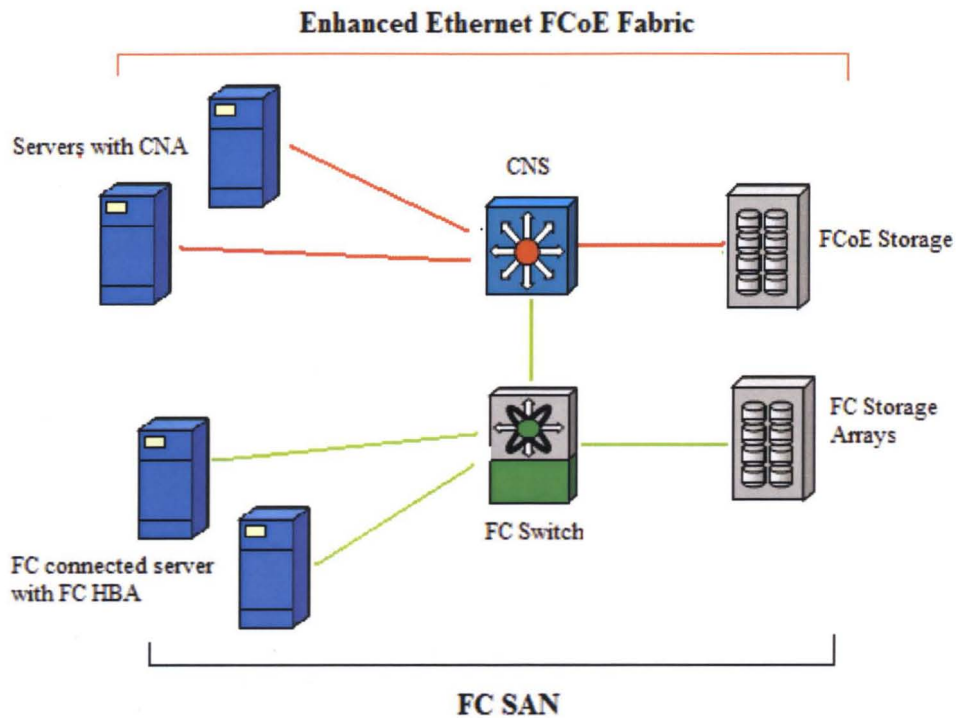


Figure 31: Deployment of Converged Fabric integrated with existing FC [20]

In addition, FCoE can run in higher bandwidth 10G Ethernet than FC alone and enhances efficiency data transfer in datacentres. Furthermore, FCoE positions Fibre Channel as the storage networking protocol of choice and extends the reach of Fibre Channel throughout the data center to all servers. Through the vast Ethernet infrastructure, FC package and data can reach much longer distance than what it can handle alone, and could cover much wider area for high speed transfer application since Ethernet is basically everywhere and inexpensive. Below is a brief description of how FCoE functions and how the obstacles are resolved [19].

• **Fibre Channel Frame Encapsulation**

The encapsulation of the Fibre Channel frame occurs through the mapping of FC onto Ethernet. Fibre Channel and traditional networks have stacks of layers where each layer in the stack represents a set of functionality. The Fibre Channel stack consists of five layers, FC-0 through FC-4 whereas Ethernet protocol consists of seven layers OSI stack that defines the physical and data link layers. FCoE replaces the FC0 and FC1 layers of FC stack with Ethernet and provides the capability to carry the FC-2 Framing layer over the Ethernet layer. This allows Ethernet to transmit the upper Fibre Channel layers FC-3 (Service Layer) and FC-4 (Protocol Mapping) over the IEEE 802.3 Ethernet layers. The Ethernet MAC address layer and physical layer will be maintained in the FCoE stack layer structure. See Fig. 34 which demonstrated the layer mapping for FCoE.

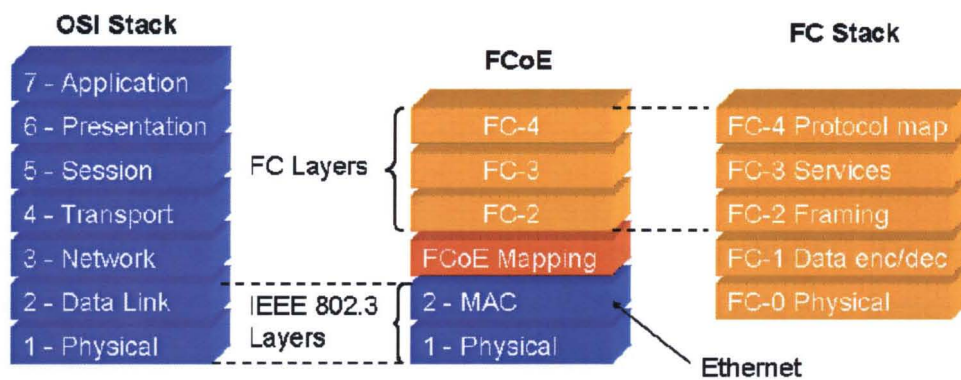


Figure 32: FCoE layers mapping

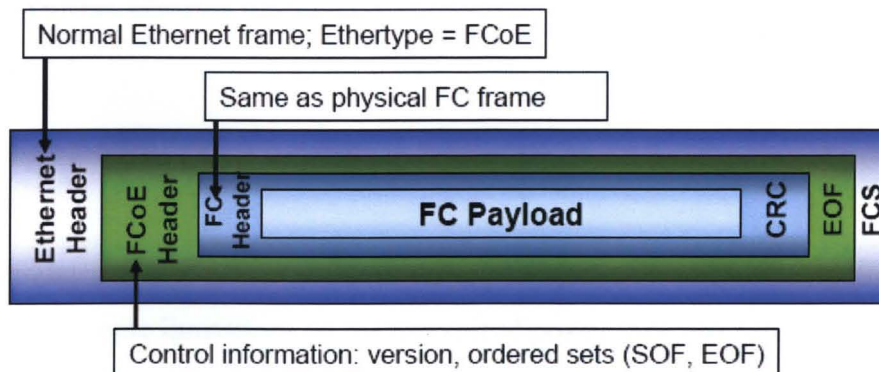


Figure 33: Format of FCoE frame encapsulation [21]

- **Lossless Fabric in Enhanced Ethernet for FCoE**

One of the challenges of passing Fibre Channel frames over Ethernet is that FC provides a lossless transport while Ethernet lacks this feature. Ethernet technology will drop the Ethernet frames under network congestion, which causes a “lossy” network that must be resolved in Enhanced Ethernet for FCoE. Traditional Fibre Channel manages congestion through the credit based flow control that guarantees no loss of a frame under normal conditions. A device cannot send any additional frames until the Receiver side indicates that it is OK to do so. Typical Ethernet without an adequate flow control mechanism will drop packets when traffic congestion occurs. However, Ethernet does have a flow control PAUSE mechanism that can be used to prevent packet loss by asking the Transmitter to stop sending more frames until the receiver buffers are cleared. With the enhanced IEEE 802.1Qbb Priority Flow Control (PFC) in FCoE, a new PAUSE function will pause the lower priority traffic during heavy congestion period while allowing higher priority and latency sensitive tasks to continue, hence no loss in congestion.

- **MAC-Addresses in Ethernet Fabric**

Traditional Fibre Channel fabric switches maintain forwarding tables FC_IDs. FC switches use these forwarding tables to select the best link available for a frame so that the frame reaches its destination port. Fibre Channel links in switch fabric are typically point-to-point and do not need an address at the link layer. An Ethernet network is different as it does not form a point to point connection similar to how FC functions. This requires FCoE to rely on Ethernet MAC (Media Access Control) addresses to direct a frame to its correct Ethernet destination.

3.3.3 Green power data center

The world today is rather concerned over energy usage for the reason being that many businesses have some limitations on the amount of power they can access due to legal restraints, physical access or spacing limitation. Another concern is the budget limitation due to high power or electricity cost in supporting large server networks and cooling

systems. Therefore, companies are starting to seek ways in energy saving to cut cost in IT but still sustain the normal operation of the daily network system. One of the solutions that can save companies lots of electricity cost by shifting computing power across their data centers to a location with cheaper energy prices. By setting up a company's data centers in an area where natural green energy like solar, wind or hydro are unlimited with low cost, companies will then save costs. This would also be an excellent action to reduce global greenhouse effect by reducing carbon release in cities where coal may still be the main energy source. Economy wise, companies can save budgets due to the fact that energy generated by natural energy source could provide cheaper electricity for datacentre usage. Setting up datacentre in cheaper energy location is an advantageous plan, however there is still the question of how to access the data when remote datacentre sites could be hundreds or thousands of km away from the IT centre. A solution would be to apply Fibre Channel long distance transfer with its high data transfer rate. FC connections with optical media and boosters can link the remote datacentre and IT office together effortlessly. With FC combining Ethernet network as FC over Gigabit Ethernet, data can travel universally through the Ethernet network with high data transfer rate, high performance, reliability and flexibility.

In addition, the setup of a datacentre in a natural disaster free location can also assist company IT centre or datacentre to be secured from disasters such as earthquake and hurricane. Therefore, companies will then be worry-free from their data backup and recovery problem. This scenario shows another excellent FC application with the benefit of a long distance high transfer rate.

3.3.4 Disaster Backup

All efforts of having the best technology, reliable hard disk and strive of perfection of network or circuit designs are geared towards for the safety of the critical data and to avoid any possible failures. However, inevitable disasters like hurricanes, floods, fires, earthquakes and extended power outages could definitely threaten the life of any company. As a disaster occurs, recovery time will become critical and IT operations must

deliver support in the shortest recovery time possible. These backup and data recovery requirements drive the need for high-performance solutions as any hours or even any minutes of system down time could cause a huge loss in business. Therefore, the ability to store, retrieve data, existence of a remote datacentre backup and ability to system recovery from failure are critical to businesses and IT nowadays. Fibre Channel is definitely an efficient solution to provide high-performance networks that enable businesses to prepare and execute an effective recovery and survival plan from disasters.

Investment in Gigabit Fibre Channel infrastructure prior to the striking of a natural disaster, can dramatically improve the efficiency of system backup and recovery. With its long distance transfer and superior transfer data rate, Fibre Channel solutions enable real-time backups and the ability to move massive amounts of data from the remote backup datacentre within critical backup windows rapidly after the occurrence of natural disaster. The higher the accelerate recovery, the lower the amount of lost productivity and profitability. Fibre Channel is the definitive option for this long distance backup and had been ultimately field-tested during the 9/11/01 disaster and subsequent natural disasters.

According to a report from Contingency Planning Resources (White Plain, NY), in the last ten years, an estimated \$80B has been lost due to disasters affecting computing services. The following chart shows typical costs of downtime where the hourly cost of downtime is surprisingly high [1]. As shown, Fibre Channel investment can provide a quick return on your investment if you can greatly reduce your downtime in disaster.

Business	Industry	Hourly Downtime Costs
Brokerage Operations	Finance	\$6,450,000
Credit card sales authorizations	Finance	\$2,600,000
Pay-per-view	Media	\$150,000
Home shopping (TV)	Retail	\$113,000
Catalogue sales	Retail	\$90,000
Airline reservations	Transportation	\$90,000
Tele-ticket sales	Media	\$69,000

Package shipping	Transportation	\$28,000
ATM fees	Finance	\$14,500

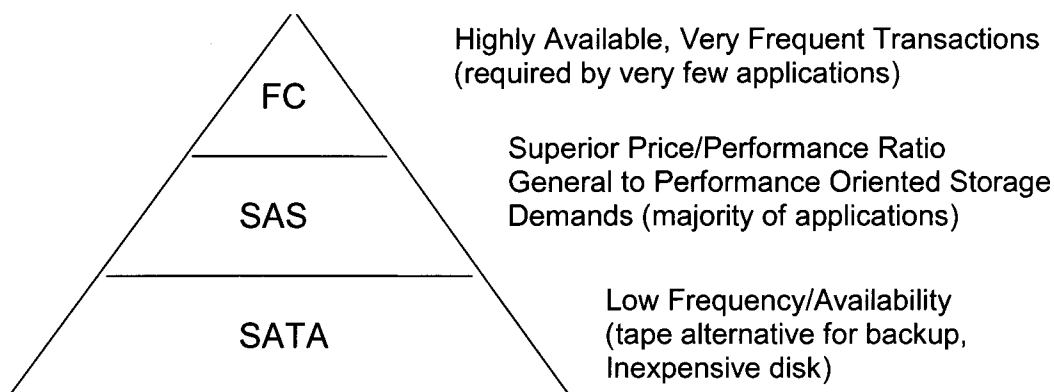
Table 14: Cost of downtime due to natural disaster

FC provides a high bandwidth and data transfer rate between recovery sites and data center as far as hundreds km away to restore data rapidly. This is particularly attractive to companies and organizations that are critically dependent on their computer systems. If the primary system should fail, such as a in-house fire, the remote site would be immediately ready to take over processing in order to eliminate the system downtime and reduce loss.

Section IV Conclusion

According to all analysis comparisons from different areas and categories on SAS, SATA and FC, it is concluded that different technologies have their own advantages and market share in their own industry.

In the hard drive market, in general, as a rule of thumb, it is best to choose SATA hard drives for their low cost and high capacity storage where cost is a much more important factor than performance, such as PC and desktop application. SAS provides an attractive price/performance and will gain most of the market for external storage as it is much more economical than Fibre Channel and much sturdier than SATA. SAS's high performance and high reliability will lead SAS drives to become the hottest storage media in the servers market. Alternatively, FC drives will provide superior performance and reliability that is most suitable to be installed in applications only where performance requirements exceed the price-to-performance attributes SAS offers. However, the demand for FC drives are expected to be flatted or reduced mainly due to its high-cost.



In network interconnect or network connectivity to SAN, FC is currently taking the lead position due to FC superior switching capability and the fact that applications required high data throughput, long transmission distances, low maintenance, and high reliability.

Especially with the advantage of the long distance transmission through optic media, FC SAN connection changed the nature of data transfer from range within building to range in national. Backup data with FC connection can be re-archived from remote sites located in thousands km away from the main data centre. Thus, remote sites are safe and located in energy-efficient areas where it is natural disaster proof or rich in nature energy sources. When FC operates together with Ethernet (FCoE), data can reach even longer distance by penetrating the Ethernet network. Combining the storage network and Ethernet together will further simplify both network interfaces into a single network as well as reducing I/O and cabling management. FCoE's compatibility allows companies to integrate the existing FC infrastructure in FCoE setup without in huge infrastructure modification expense. The advantage for FCoE will not only further enhance the bandwidth and speed for data transfer by going through 10 Gigabit Ethernet, it will also extend the reach of transmission distance from storage network to Ethernet network. FCoE is expected to be the new potential network structure in the next coming years.

Section V Reference

- [1] Fibre Channel Solutions, “Business Continuity When Disaster Strikes”, P.2 by Fibre Channel Association <http://www.scribd.com/doc/17469021/Fibre-Channel-Solutions>
- [2] “Serial ATA- Next Generation Storage Interface” White paper from HITACHI, by Frank Chu.
[http://www.hitachigst.com/tech/techlib.nsf/techdocs/88B8092A094253CD86256D4E005544BD/\\$file/sata_interface_white_paper_091605.pdf](http://www.hitachigst.com/tech/techlib.nsf/techdocs/88B8092A094253CD86256D4E005544BD/$file/sata_interface_white_paper_091605.pdf)
- [3] “Serial Attached SCSI – Better Performer, Scalability, and Reliability for better Storage Solutions”, by SUPERMICRO
http://www.supermicro.com/downloadables/pdf/Supermicro_SAS_LinuxWorld.pdf
- [4] “iSCSI vs Fibre Channel: A cost Comparison”, Processor Editorial Article
<http://www.processor.com/editorial/article.asp?article=articles/p3014/31p14/31p14.asp>
- [5] “Serial Attached SCSI”, <http://www.bestpricecomputers.co.uk/glossary/serial-attached-scsi.htm>”
- [6] “Comparing SAS, SATA and Fibre Channel”,
<http://www.networkworld.com/newsletters/stor/2004/0906stor2.html>”
- [7] “SAS vs FC” white paper by R/Evolution
http://www.dothill.com/assets/pdfs/dothillwp_FCvsSAS_25Sep06_letter.pdf
- [8] “Introduction to Serial Attached SCSI (SAS) and Serial SATA (SATA)”, white paper by MINDSPEED
<http://www.mindspeed.com/web/download/download.jsp?docId=28618>
- [9] “Synchronous SAN Sets Fibre Channel Distance Record”, March 28, 2003 by Paul Shread <http://www.enterprisestorageforum.com/industrynews/article.php/2171801>
- [10] “Can Fibre Channel (FC) Go the Distance?” By Tim Anderson on May 30, 2008
<http://www.dciginc.com/2008/05/can-fibrechannel-go-the-distance.html>
- [11] “Gennum Unveils IC Solution for 16G Fiber Channel SFP+” by Jai C.S March 25, 2009-10-29 <http://it.tmcnet.com/topics/it/articles/52963-gennum-unveils-ic-solution-16g-fiber-channel-sfp.htm>
- [12] “Serial Interfaces in the Enterprise Environment” White paper by Willis Whittington, ESG Interface Planning, Seagate. Dec 2002, Number: TP-306

http://www.lsi.com/DistributionSystem/AssetDocument/files/docs/marketing_docs/storage_stand_prod/Technology/tp306_serial_interfaces.pdf

[13] “LSI, Dell and Seagate First to Demonstrate End-To-End 6Gb/s SAS Server Interoperability”, by Milpitas, Calif., May 6, 2008.

http://www.lsi.com/news/product_news/2008/2008_05_06.html

[14] “SAS and SATA Storage Technologies”, SNIA Education, by Harry Mason, LSI Logic, Marty Czekalski, Maxtor and Ahmad Zamer, Intel.

<http://www.snwusa.com/images/SASSATA.pdf>

[15] “Serial ATA Technology”, technology brief, 2nd edition by HP

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00301688/c00301688.pdf>

[16] “Comparing Fibre Channel and Serial ATA” by Ram Gopalan

<http://www.dentistryiq.com/index/display/article-display/226135/articles/infostor/volume-9/issue-4/features/comparing-fibre-channel-and-serial-ata.html>

[17] FCIA (Fibre Channel Industry Association)

<http://www.fibrechannel.org/overview/fcbasics/topologies>

[18] “Fibre Channel Fundamentals” by Tom Weimer

http://www.unylogix.com/data_storage/raid_san/PDFs/White_Paper_Fibre_Channel_Fundamentals.pdf

[19] “Fibre Channel over Ethernet in the Data Center: An Introduction” by FCIA,

http://download.intel.com/technology/comms/unified_networking/white_paper_FCI AFC_oE_RHC_Unified_Networking.pdf

[20] “Fibre Channel Solutions Guide” by FCIA,

www.fibrechannel.org/documents/doc_download/1-fcia-solution-guide

[21] “Fibre Channel over Ethernet (FCoE)” by EMC, Mark Lippitt, Erik Smith, Erik

Paine, 2009 <http://www.emc.com/collateral/hardware/technical-documentation/h6290-fibre-channel-over-ethernet-techbook.pdf>

[22] Price comes from website: www.pcsuperstore.com

[23] Price comes from website: www.nextag.com

[24] STA SCSI Trade Association

http://www.scsita.org/aboutscsi/sas/SAS_roadmap.html