

LinuxDevices.com printer-friendly story display of:

<http://www.linuxdevices.com/articles/AT3656850450.html> . . .

Ethernet-based Network Storage -- Reality or Pipe Dream?

by Patrick B. T. Khoo, Wilson Y. H. Wang, and H. N. Yeo (Nov. 11, 2002)

FAQ . . .

What is HyperSCSI?

- **What exactly is HyperSCSI anyway?**

transmission of SCSI commands and data across a network. To put this in "ordinary" terms, it can allow one to connect to and use SCSI and SCSI-based devices (like IDE, USB, Fibre Channel) over a network as if it was directly attached locally.

- **Why was HyperSCSI created?**

transport protocol to run over Ethernet as opposed to Fibre Channel for network storage. The primary technology being developed at the time was called iSCSI, and encapsulated SCSI on TCP/IP. A major problem was the TCP/IP overheads that was slowing the adoption of Ethernet as a viable alternative to Fibre Channel. As a result, we decided to explore putting SCSI on top of raw Ethernet instead, thereby bypassing the TCP/IP stack. What started as a research project has now become HyperSCSI.

- **What can HyperSCSI be used for?**

across a wide variety of networking systems. You can use it to access storage devices like hard disks, CDROMs, tape drives and so on across a network. It is a low level protocol, which means it even supports advanced functions like CD-writing over a network. For more examples, please refer to our documentation.

- **How and who is developing HyperSCSI? --** HyperSCSI is being developed by the Modular Connected Storage Architecture (MCSA) group under the Network Storage Technology Division of the Data Storage Institute [of Singapore]. Of course, as an open source project, we welcome any and all developers to contribute to this technology.



**HyperSCSI demonstration environment
for Ethernet-SAN and wireless storage**

Release of HyperSCSI as open source

On Aug. 28, 2002, the Modular Connected Storage Architecture (MCSA), group within Singapore's national Data Storage Institute (a research institute under the national Agency for Science, Research, and Technology) released the HyperSCSI protocol and reference implementation code into open source under the GNU General Public License (GPL). The full announcement of HyperSCSI GPL release is available [here](#). HyperSCSI's protocol specifications, reference code, benchmarking results, and other information are available at the [HyperSCSI](#) website.

HyperSCSI whitepaper

Excerpts from the beginning and conclusion of the newly released HyperSCSI whitepaper follow. You can obtain the full whitepaper [here](#) (PDF download).

A HyperSCSI whitepaper:

Ethernet-based Network Storage -- Reality or Pipe Dream?

by Patrick B. T. Khoo, Wilson Y. H. Wang, and H. N. Yeo

Abstract

The basic problem in network storage today is how to implement solutions that are cost effective and efficient but without using a lot of new components and equipment. Implementing Storage Area Networks today generally implies using Fibre Channel, however, a lot of work is being done to try to use Ethernet instead for network storage. But these new Ethernet solutions still seem a little short.

In this paper, we would like to first introduce some of the key features of the new open source network storage protocol, HyperSCSI, and how it is different from existing solutions. This will be followed by benchmark and test results to show how HyperSCSI is capable of using an existing Ethernet-based network infrastructure, common-off-the-shelf hardware and well-established storage technologies and turning that into a high-performance and reliable network storage solution. Since specialized hardware and custom software are not required for HyperSCSI, we believe this is a step in the right direction to building cost effective and efficient network storage solutions.

Finally, we want to draw just one conclusion in our paper, that the existing network infrastructure and technologies can be successfully exploited to meet the requirements of network storage. The lesson learned in the search for this answer is that one must be ready to look for innovative new methods and perhaps, use unconventional thinking to meet the requirements of network storage. The resultant HyperSCSI protocol is proof of this.

Introduction

The concept of a network for storage has not been a new one. Before the terms Storage Area Networks (SAN) and Network Attached Storage (NAS) became commonplace, mainframes, and later on, simple file servers have been doing network storage for years. Of course, it was not quite in the form that we recognize today, but it was certainly a kind of network-ed storage. In fact, the very concept of transferring data over a wire, which is fundamental to network storage, is older than networking itself.

Fibre Channel (FC) is today a key technology for the deployment of the modern SAN. The explosive growth of the Internet, applications like CRM and ERP and so on, has been supported in part by the ability of storage to scale. However, its adoption has not been as stellar as many had hoped. There was

even a time when people talked about running IP over FC, but not anymore. The deployment of SANs has not grown as fast as the analysts had predicted. Of course, part of the reason for this is due to socio-economic events and downturns, but certainly not the whole reason.

With new advancements like Fast and Gigabit Ethernet (and 10GE around the corner), as well as new high performance wire-speed Layer 3 switching and so on, it is a good time to analyze if Ethernet-based network storage can do the job. However, it is not as simple, nor as easy as that. A lot of work needs to be done and the industry developed various technologies like iSCSI, mFCP, iFCP, FCIP and iSNS to fill this gap. The two leading core protocols are of course Internet SCSI (iSCSI) and Fibre Channel over IP (FCIP). While they may seem to compete, they actually meet different needs. FCIP is needed for bridging FC-SANs over an IP-based network, while iSCSI is more for storage over an IP network. iSCSI is probably the leading contender that allows the deployment of pure Ethernet-SANs without the use of FC. However, its performance has not been adequate. To solve this problem, development work has shifted towards hardware acceleration like TCP/IP Offload Engines (TOEs) and iSCSI HBAs. These new developments are needed to push iSCSI to comparable FC speeds. But the reality is, with all these add-ons, what will be the difference between a TOE-NIC and an iSCSI HBA from a FC HBA then? TOE and iSCSI HBA manufacturers point to the cost savings you will get by using Ethernet switches instead of FC-based switches, and the fact that TOEs and iSCSI HBAs will get cheaper as adoption increases. But does that sound a little like vapor-ware? As a result, adoption is still rather slow. Perhaps a new solution is required.

Designing a network storage protocol is not as straightforward as it might seem. We must consider that the characteristics of data storage are different from the conventional data traffic. Furthermore, the quality of the network has improved greatly compared to many years ago. As such, we believe that a non-conventional approach to solving this problem is required. In this paper, we present precisely such a solution, a new open source network storage protocol, which we have named HyperSCSI.

HyperSCSI is designed for the transmission of SCSI commands and data across a network in a simple and efficient way. The current implementation runs over an Ethernet network, uses existing common-off-the-shelf hardware and components and does not require any additional customized software or expensive hardware accelerators. This ultimately, will reduce the cost of an overall network storage implementation. We believe that using existing hardware and technologies does not compromise on performance and reliability. In this paper, we will present key features of HyperSCSI and how it is different from existing solutions as well as various test results to demonstrate its capabilities.

Selected Key Features of HyperSCSI

HyperSCSI is a new open source network storage protocol designed for the transmission of SCSI commands and data across a network. To put this in ordinary terms, it can allow one to connect to and use SCSI and SCSI-based devices (like IDE, USB, Fibre Channel) over a network as if it was directly attached locally. This section focuses on a few key features of the HyperSCSI protocol, and how they differ from existing solutions . . .

(. . . snip . . .)

Discussion and Conclusion

In the introduction of this paper, we looked back and understood that network storage is in fact not really a new idea. Neither are SCSI, Ethernet and TCP/IP. But why reminisce? Looking back in history shows us that in fact, the primary drivers for growth in both economic and technical terms has been very simply -- Can I do more for less? If we understand this very important point, then all other

predictions on trends and evaluations on technologies can be put into its proper focus.

Also in the introduction, we covered how users were not completely satisfied with Fibre Channel and iSCSI-type solutions. Whether it was complexity or cost, the adoption has just not been as strong as forecasted. Do users really get more for less? The naysayers therefore, point to the possibility that Ethernet can't really do storage.

However, in the rest of this paper, we tried to illustrate a new solution called HyperSCSI, which attempts to truly give users more for less. It is not over-engineered, is simple in concept, does not require special hardware or customized software and manages to provide a reasonable level of performance.

Can people really get more for less? We believe the answer is an unqualified yes. However, as we outlined in the beginning of this paper, doing so requires us to change our mindsets.

For example, if TCP/IP is really the bottleneck for storage over Ethernet, then why use it? Do we really need it? Of course, for certain wide-area connectivity applications, like disaster recovery and so on, a solution like iSCSI is quite a good idea. Some long distance tests were successfully conducted, demonstrating iSCSI storage between Israel and California. But in fact, most SAN implementations are not for use in wide-area long-distance applications. If it were, FC would have already died out by now since FC is itself also a local-area networking technology that needs something like FCIP to bridge the router divide. Eliminating TCP/IP would eliminate the need for hardware accelerators while still achieving high performance Ethernet-SANs. A nice side effect of eliminating TCP/IP is that the disk array providing Ethernet-SAN can't be hacked from the Internet.

Furthermore, HyperSCSI gives rise to entirely new applications and markets. For example, HyperSCSI runs quite well on wireless LAN, thus allowing the development of a wireless HDD or CDRW for your laptop instead of through USB or Firewire. Or how about watching a movie on a webpad from the patio by directly accessing the DVD player in your living room wirelessly? With such new developments, it's no wonder that there is a renewed sense of optimism for the network storage industry.

No, we believe that Ethernet storage is not coming it's already here. And yes, you will get more for less.

Obtain the full HyperSCSI whitepaper (173KB PDF download) here:

[Ethernet-based Network Storage -- Reality or Pipe Dream?](#)

[by Patrick B. T. Khoo, Wilson Y. H. Wang, and H. N. Yeo](#)

Related stories:

- **[DSI releases new network storage protocol into Open Source](#)**

... For further information

Except where otherwise specified, the contents of this site are copyright © 1999-2001 DeviceForge LLC. All rights reserved. DeviceForge, LinuxDevices, and LinuxDevices.com are trademarks of DeviceForge LLC. The LinuxDevices.com logo is a service mark of DeviceForge LLC. All other trademarks are the property of their respective owners. Site coding by Webgurus.com