



Coraid EtherDrive — The Shortest Path to Low-Cost Enterprise SAN Storage

Analyst: Michael Fisch

Management Summary

One of the first things you learn in geometry is that the shortest distance between two points is a straight line. It is relatively easy to find that line on a blank sheet of paper with a ruler and pencil in hand. However, finding something approximating a straight line between two points in a city like Boston with curvy, convoluted streets and dizzying traffic is much more difficult, unless you know the city well. This is why visitors to Boston are prone to getting lost or off-track while driving to their destination. However, if you can find the shortest path, you arrive as fast as possible, saving time, hassle, and gasoline.

In a similar way, you could say the Redwood City, CA, company Coraid Inc. (pronounced “Co-raid”) set out to simplify SAN storage by finding the “shortest path” between a server’s CPU and a storage device connected over a network. Conceptually, a straight line is simple, but finding it between two points in a complex datacenter can be difficult. It requires understanding the entire system and its interrelated components, down to the minutia, and then stepping back and thoughtfully answering the question: Which pathways and parts are nonessential? After stripping away the unnecessary complexity – the detours, doodads, and diversions – what is left is the simplest possible solution. It is clean, elegant, efficient, and, speaking in practical terms, faster and less costly.

We dare say Coraid has succeeded remarkably. It developed a lightweight protocol for storage networking called *ATA-over-Ethernet (AoE)* that runs over standard Ethernet switches. AoE is now published as an open standard. Then Coraid developed its *EtherDrive* line of scalable, high-availability SAN storage solutions based on the AoE protocol. The result is networked storage with solid performance at a price point 5-to-8 times less than Fibre Channel SAN storage. (Yes, that is nearly an order of magnitude less.) Read on for details.

ATA-over-Ethernet

ATA-over-Ethernet or AoE is a lightweight, open protocol for disk sharing over a network. It was developed by a team of engineers at Coraid (led by CTO Brantley Coile) as an efficient, cost-effective alternative to Fibre Channel and iSCSI block storage protocols. The constituent technologies are ATA, the standard way a computer CPU communicates with a disk drive, and Ethernet, the standard networking technology for local area networks that is available in 10 Gbit/s speeds and moving to 40 and 100 Gbit/s. Both technologies are mature and mainstream and benefit from massive economies of scale, facilitating a low price point for systems based on ATA-over-Ethernet.

AoE encapsulates ATA commands and sends them over Layer 2 Ethernet. Packets are addressed to devices using their MAC addresses. It does not use the high-level TCP and IP protocols with their associated high overhead, so it is streamlined and fast. AoE is also non-routable, which confers a degree of security because the data packets lack an IP address and cannot travel beyond the local SAN.

IN THIS ISSUE

➤ ATA-over-Ethernet	1
➤ EtherDrive SAN Storage	2
➤ Conclusion	2

A SAN based on AoE is easy to implement. In this case, it requires an EtherDrive storage system, a standard Ethernet switch, and a standard Ethernet NIC and software driver (i.e., an AoE initiator) on the host server. Software drivers are OS-specific and currently available for VMware ESX and ESXi, Linux, Windows, Mac OS X, FreeBSD, and Solaris. The software driver automatically discovers AoE storage LUNs attached to a server. Since Ethernet is a datagram protocol as opposed to a connection-based protocol (like Fibre Channel), the driver sends datagrams to all available ports on the server and targeted storage to maximize throughput, minimize latency, and provide fault tolerance. Additional software for port bonding and multi-path I/O are unnecessary. The best performance is achieved if the Ethernet switches and NICs support jumbo frames.

iSCSI is the closest mainstream competitive alternative to AoE. While both are based on Ethernet, iSCSI uses the TCP and IP protocols and requires more processing and latency. In the analogy of straight line between two points, AoE is the shortest route while iSCSI takes a couple extra loops along the way.

EtherDrive SAN Storage

Coraid developed its EtherDrive line of SAN storage solutions based on the AoE protocol. The EtherDrive storage appliances are densely scalable, offer good throughput performance, and are relatively easy to install and manage. Most impressive of all is a price tag that starts under \$600 per raw terabyte – 5-to-8 times less than competitive Fibre Channel storage systems (some of which are even more costly).

The *EtherDrive SRX-Series* is the highest performing with throughput speeds up to 1,800 MB/s. These appliances scale to 36 drives in either 3.5" or 2.5" form factors in a 4U chassis – which is up to 72 TB of raw capacity per system when using 2 TB SATA drives. Multiple units can be deployed to scale capacity well beyond the petabyte range. They support SATA drives for capacity storage, SAS drives for performance, and SSD drives (i.e., flash memory) for ultra-fast performance. RAID 0, 1, 5, 6, and 10 are available for data resilience. Connectivity options include multiple GigE or 10GigE ports.

The EtherDrive SR-Series is the low-cost option. These appliances support 3.5" SATA drives and scale to 24 drives in a 4U chassis, 16 drives in 3U, or 8 drives in 2U, depending on the model.

RAID configurations and two GigE ports are supported.

Coraid provides AoE host drivers for Linux, Windows, VMware ESX, XEN, OpenSolaris, and Solaris. Drivers for other operations systems are available from third parties. EtherDrive supports MAC address filtering for restricting logical volume access to host servers with specific MAC addresses. Similar to LUN masking in Fibre Channel SANs, it prevents the possibility of a host server overwriting another's data by mistake.

For high availability configurations, the *EtherDrive EMX Mirror Appliance* creates synchronous mirrors between EtherDrive systems.

For centralized management, the *EtherDrive SAN Manager* appliance provides a GUI client for discovering, monitoring, and managing all EtherDrive storage on a SAN. A *RESTful* control interface is available for HTTP scripting in cloud applications.

Founded in 2004, Coraid has accumulated a long list of notable EtherDrive customers, including Ford, Dunkin Donuts, MIT, NASA, U.S. Navy, and Lockheed Martin.

Conclusion

EtherDrive and the AoE protocol are disruptive technologies – disruptive in a good way because they facilitate a new low-cost point for SAN storage. If simplicity is the new sophistication, then AoE will be the next storage chic because it charts the most direct, efficient and inexpensive route between servers and storage over a network. Think of it as iSCSI, but without the *SCSI* or the *i*, because they found a way to network storage without the overhead of SCSI and TCP/IP. After all, who needs the complexity?

Coraid EtherDrive is suitable for a range of applications from backup and archiving to e-mail, file servers, and server virtualization. If your enterprise has a shared storage requirement where price per unit capacity is a critical factor, but you cannot sacrifice availability and throughput performance, then you should consider Coraid.



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Michael Fisch is a Senior Contributing Analyst for The Clipper Group. He brings over 14 years of experience in the computer industry working in marketing, sales, and engineering, the last nine of which he has been an analyst with Clipper. Before Clipper, Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems (since acquired by Sybase). Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

- ***Reach Michael Fisch via e-mail at mike.fisch@clipper.com at 781-235-0085 Ext. 211. (Please dial "211" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, and "*clipper.com*" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "*Navigating Information Technology Horizons*", and "*teraproductivity*" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.