

STORAGE SWITZERLAND BRIEFING REPORT

CORAID ETHERDRIVE SRX-SERIES STORAGE SYSTEMS

ETHERNET SAN ARRIVES IN THE ENTERPRISE



George Crump, Senior Analyst

There's an old adage when selecting a storage system that says: "you can have simplicity, reliability or performance - pick two". Traditional NAS and iSCSI storage solutions, both running at the TCP/IP layer, have been trying to ignore this 'wisdom' and provide storage managers with all three. Reality is that these technologies often fall short on performance and their initial simplicity becomes complex as the IP infrastructure scales. A new category of Ethernet SAN solutions may enable customers to truly pick all three.

[Coraid](#) is the company behind an emerging Ethernet SAN offering that uses a lightweight storage networking protocol, ATA over Ethernet (AoE), utilizing raw Ethernet frames for transport. As discussed in our earlier article "[Storage Evolution: FC SAN, IP SAN, Ethernet SAN](#)", there's a big difference between Layer 3 and Layer 2 Ethernet when it comes to performance and complexity. NAS, iSCSI, and FCoE storage networking approaches typically add layers of expensive optimization features to improve performance. In contrast, Ethernet SAN achieves impressive speed and simplicity by eliminating layers of complexity in the network and leveraging the native performance and availability attributes of Ethernet. As enterprises increasingly move to virtualized and private cloud architectures, Ethernet SAN may become a logical first choice.

Coraid grew organically in the Linux market, and attracted a large customer base of over 1100 customers largely by word of mouth. Its AoE-based EtherDrive storage systems were initially focused on the mid-range Linux and research computing markets which benefited from cost effective, high performance, scalable systems. This year, the company added venture financing and an experienced management team with a focus on enterprise and cloud markets. Coraid followed up by launching EtherDrive SRX Series storage arrays with 10 Gb Ethernet connectivity, support for SATA/SAS/SSD drives, and throughput performance of 500-900 MB/sec per shelf, with pricing under \$1000 per terabyte. In addition to a 5-10x price-performance win, EtherDrive storage appliances offer unprecedented simplicity for the storage industry: configuration that can be completed in as little as sixty seconds.

The SRX series comes in three models, the 2800, 3200 and 3500. The 2800 can handle 16 3.5" drives in a 3U shelf, the 3200, 24 3.5" drives in a 4U shelf and the 3500, 24 2.5" drives in a space-efficient 2U shelf. SAS, SATA or SSD can be mixed into the same shelf providing impressive flexibility in a single system. RAID protection is provided on each shelf and can include RAID 0, 1, 5, 6 or 10 configurations depending on the type or protection required. Each shelf can support a variety of Ethernet port configurations up to a maximum of 4x10GbE ports or 6x1GbE ports.

Coraid's initial offering lacked a few key storage management features, but this wasn't a barrier in the Linux, research and virtualization markets where users were able to leverage file system and hypervisor-based features. This makes perfect sense, with an increasing number of storage features like snapshots, thin provisioning, and de-duplication moving up the stack. Even capabilities like replication can often be better handled through virtual machine aware products like Vizioncore's vReplicator, as evidenced by the number of environments that have storage systems with replication yet still purchase VM-aware third party software. In fact, many customers are growing increasingly frustrated with paying for these features twice – once at the virtualization layer, and again in the array. For these customers, the ideal solution is a high-performance SAN that installs easily, without a dozen expensive software line items tacked on.

Coraid is moving aggressively to close the gap on missing enterprise array features in 2010, and has already launched enhanced support for Windows, Solaris, Linux, and VMware environments. Adding advanced replication and storage management capabilities at the array level will extend the reach of Ethernet SAN into a broader set of non-virtualized application environments including database and email.

In addition to developing the industry's first scale-out Ethernet SAN, Coraid has been working on other innovations. For example, instead of offering a costly program for 8-hour and 4-hour product replacement, Coraid has developed a "zero-hour" support solution with hot spare units. Coraid EtherDrive lays down data in a way that disks can be pulled from a broken unit and inserted directly into a hot-spare chassis in any order, getting customers back in business in a matter of minutes. This feature also provides flexibility as needs change. For example, if the initial investment is an all disk system with a mix of SAS and SATA, later another system can be added with a few SSDs installed. The SAS drives could then be moved to the new system with the SSDs, creating a performance-centric shelf, and then the original shelf could be filled out with additional high capacity SATA drives. No-single-point-of-failure configurations are also available using Coraid's ESX mirroring appliances.

Coraid drivers have been native in the Linux kernel since 2005, and the company provides dedicated Ethernet HBAs for servers running Windows, Solaris, Linux/Xen, and VMware. Coraid provides HBAs in two configurations; either a dual 10GbE NIC or a dual 1GbE NIC. Multiple HBAs can be installed in a server for linear scaling; Layer 2 Ethernet automatically provides multi-pathing, with no need for additional software. The port flooding technology available on AoE allows for both ports to be used seamlessly as one combined port for 20GbE or 2GbE of available storage I/O bandwidth depending on the card. Again, because of AoE's lack of dependence on IP there is less overhead to limit performance or chew up server CPU cycles, and near line-rate speeds are common. The interconnection between the servers and the storage is a standard Layer 2 Ethernet switch running jumbo frames. That switch can be part of the existing network core, but many customers use one or more inexpensive top-of-rack switches. This makes 10GbE connectivity available for less than the price of today's 4Gb fibre switches, and simplifies network management.

Once connected, LUNs are created on each system based on the desired RAID configuration and drive type. For example, a high performance LUN can be created using two mirrored SSDs, a medium performance LUN can be created with 5 x 15K SAS drives running RAID 5 for protection and then the remaining drive slots populated with SATA drives running RAID 6 for a capacity LUN. Each LUN is referred to by its shelf number and LUN number.

Storage Switzerland's Take

Storage can never be too fast, too reliable or too simple. The Coraid SRX line may be the perfect balance of those capabilities, especially for users that are leveraging storage management features at the file system and hypervisor layers. As Coraid delivers additional features at the array level, their platform becomes a real contender across the rest of the enterprise. For storage managers that are looking to lower costs, improve performance and simplify storage operations, Ethernet SAN may be uniquely positioned to deliver on all three.

About Storage Switzerland

Storage Switzerland is an analyst firm focused on the virtualization and storage marketplaces. For more information please visit our web site: <http://www.storage-switzerland.com>.