




The Linux Storage People



**SR1520 and SR420 Storage Appliances
&
PB1020 Storage Blades**



EtherDrive® Storage

Product Description & Applications

Updated 4 March 2006

Product Description


The AoE Storage Protocol

Coraid storage products use a simple protocol called ATA-over-Ethernet (AoE). Coraid developed AoE as an “open” protocol based upon standards like Ethernet and the ATA disk access protocol. You can find the complete protocol definition in an 8 page specification document at <http://www.coraid.com/support.htm>. This simple yet powerfully efficient protocol is designed to reliably transport ATA disk commands over Ethernet frames. AoE is a command/response protocol with robust features needed for storage traffic. AoE can work with any Operating System by using an AoE driver. Drivers are available for: Linux, FreeBSD, Solaris, Apple OS X, and Windows.

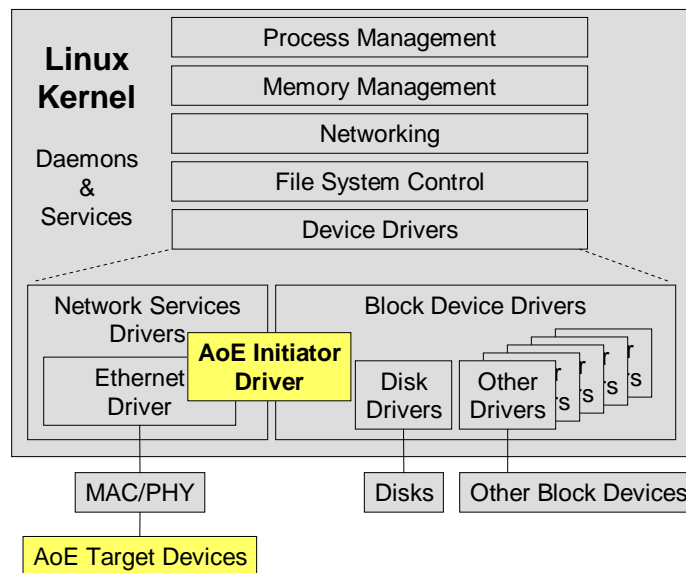
AoE uses Ethernet frames (AoE is registered Ethernet type 0x88A2) and does not require TCP/IP or iSCSI protocol layers. This means AoE is not a routable protocol, and therefore provides excellent security for the storage network. And because AoE is simple, it requires less host processing load, and can work with standard Ethernet networks and NIC interfaces. This can save a considerable amount of money in large networked storage systems.

ATA-over-Ethernet (AoE)

- AoE sits directly on top of Ethernet
- AoE is not “routed” therefore no IP or iSCSI layers are required
- AoE is simple, secure and requires less processing



AoE storage devices are block storage, just like a locally attached disk. However since the disk is really out on the LAN you can add an unlimited number of disks and share them with multiple hosts. The AoE initiator driver is shown in the drawing below, bridging the block storage device drivers with network drivers in a Linux kernel. AoE is now included in the main Linux kernel 2.6.11 and higher, so Linux devices are now natively able to use AoE devices without additional drivers.




File System (EXT3, JFS, XFS, GFS, etc.)	
Logical Volume Manager (LVM, LVM2)	Optional
RAID Software (RAID tools, MDADM)	Optional
Block Devices	
AoE Devices	

AoE target devices are block devices and can be treated like any other block storage devices even though they are attached to the LAN. AoE devices can be put into RAID sets, and managed by volume manager software layers, and mounted with any disk filesystem. AoE devices behave like locally attached block storage devices.

AoE can operate at wire speed with or without jumbo Ethernet frames. Utilizing interrupt coalescing techniques in the host network interface allows the host to provide high throughput without excess processing load. This eliminates the need for networking equipment to support jumbo frames. AoE can also work with jumbo Ethernet frames (8K frames) and reduce the packet per second load even further on the host systems. Coraid's storage appliances (not blades) support both standard and jumbo Ethernet frames

Coraid's product line of AoE target devices are called EtherDrive® Storage Appliances and EtherDrive® Storage Blades.

A Simple System Configuration



- **Simply connect storage and server with Ethernet**
- **Server and storage communicate using AoE protocol**
- **Server discovers the new EtherDrive block storage devices in its /dev directory**

AoE target devices appear in the /dev directory on Linux systems. They are physically located on the network by a unique shelf number and slot number. If disks are accessed individually, the shelf and slot number locate the disk. If the disks are assembled into a RAID set inside the EtherDrive storage appliance, then a "logical" shelf and slot number is used. For example disks located in shelf #1 slots 0 through 6 could be assembled into a RAID5 set and viewed by the host as /dev/etherd/1.0 where the 0 device has the capacity and performance of the RAID5 set. More information about setting RAID can be found in the software manual for the appliance.

An AoE target device (a volume) can be shared by multiple servers attached to the network. The volume can be partitioned and each server can have its own partition or the volume can be shared using a cluster aware filesystem like GFS.

Coraid Product Categories

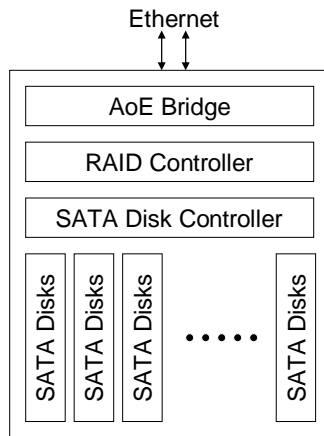
Coraid currently manufactures and sells three categories of EtherDrive Storage products; (1) Storage Appliances (which have internal RAID), (2) Storage Blades and (3) RAID Appliances.

EtherDrive® Storage Appliances

EtherDrive® Storage Appliances are AoE target devices. The appliance provides a common RAID controller and Ethernet interface for disks installed in the chassis. There are 2 variants available; SR1520

is a 3U 15 disk appliance and SR420 is a 1U 4 disk appliance. The disks within the unit can be accessed individually as JBOD (just a bunch of disks) but generally they are assembled into RAID sets and accessed as one or more “logical” AoE devices. A simple command line user interface is used to create one or more RAID sets. RAID sets can be expanded by concatenation. Global spare disks can be assigned, to protect any RAID set within the appliance.

EtherDrive® Storage Appliances



EtherDrive® Storage Appliances include:

- Dual GigE Interface
- AoE Protocol Processing
- Internal RAID Controller
- SATA Interface (4 or 15 Disk)

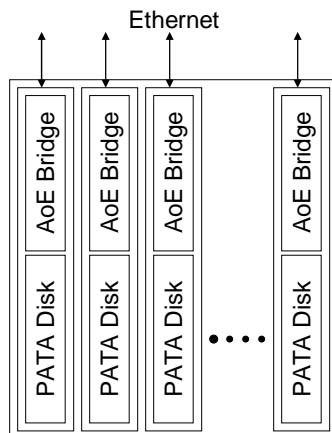
EtherDrive® Storage Appliances have multiple Ethernet ports but do not require an IP address. Messages indicating disk failure within a RAID set are reported to the hosts “syslog” using UDP datagrams from a fixed IP address of 205.185.197.30. These messages convey the shelf number and detail status message. A syslog demon on the host can then be used to process the messages into email or other system alerts.

EtherDrive® Storage Appliances allow the user to restrict access to any logical AoE device based upon server MAC address filtering. Port assigned VLANs is another effective means of limiting server access to networked AoE devices.

EtherDrive® Storage Blades

EtherDrive® Storage Blades are AoE target devices. Disks are mounted to each blade. Multiple blades are installed in a shelf with each blade having its own Ethernet connection. Therefore each blade operates independently from the other blades within the shelf.

EtherDrive® Storage Blades



EtherDrive® Storage Blade include:

- Ethernet Interface/Blade
- AoE Protocol Processing/Blade
- Each Blade is an independent AoE device

EtherDrive® Storage Blades can be assembled into RAID sets using Coraid's RA20 RAID Appliance or by using software RAID on the host. Each blade is uniquely identified by its shelf and slot address as well as its Ethernet MAC address.

AoE RAID Appliances

AoE RAID Appliances are AoE target devices for hosts and AoE initiators for AoE storage devices. The EtherDrive® RA20 RAID Appliances create RAID sets from multiple EtherDrive® Storage Appliances or from EtherDrive® Storage Blades.

Coraid's EtherDrive® RA20 is placed between the hosts and the AoE storage. Using a simple command line interface, and unlimited number of AoE storage devices can be put into RAID sets and presented to hosts as "logical" AoE devices of arbitrary size. Each EtherDrive® RA20 can export (virtualize) up to 63 separate AoE target devices, thus creating very large single volumes of storage. Each logical AoE device can be expanded by RAID concatenation.

EtherDrive® RA20 RAID Appliances have multiple Ethernet ports but do not require an IP address. Messages indicating disk failure within a RAID set are reported to the hosts "syslog" using UDP datagrams from a fixed IP address of 205.185.197.30. These messages convey the shelf number and detail status message. A syslog demon on the host can then be used to process the messages into email or other system alerts.

RAID appliances allow the user to restrict access to any logical AoE device based upon server MAC address filtering. Port assigned VLANs is another effective means of limiting server access to networked AoE devices.

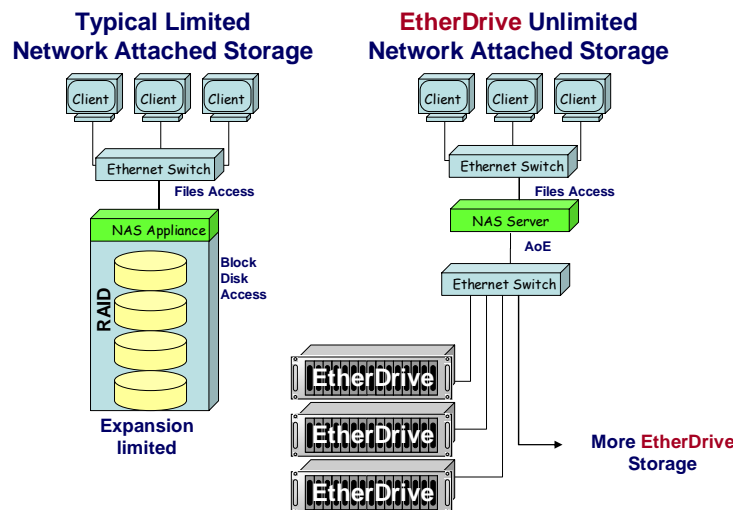
Applications

EtherDrive® Storage is perfect companion storage technology for blade server clusters, and underneath NAS filers, Virtual Tape Library servers, as well as Database, Email and Web Application servers.

NAS

In a typical Network Attached Server (NAS - as shown below) the storage is directly attached to the NAS and constrained by the size of the NAS chassis. As the NAS storage is increased, the user has no choice but to purchase another NAS system when the chassis capacity is reached.

In the typical NAS a failed server or RAID controller results in captured storage that can not be readily accessed by any other servers on the network. To recover the data the user has no choice but to repair the failed server or controller.

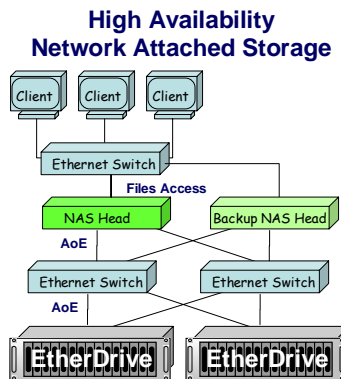


Using EtherDrive Storage, under a NAS head/gateway, can yield many benefits. An EtherDrive storage array under a NAS head/gateway has unlimited expandability. To add storage to the NAS filer, all that is needed is EtherDrive Storage and an Ethernet connection.

The CORAID Linux NAS Gateway product CLN20 uses this method to provide a low cost NAS system with unlimited expansion capacity.

High Availability NAS

The figure below illustrates how EtherDrive Storage can enable a simple configuration to achieve a high availability NAS configuration.




Using a backup server that is constantly monitoring the health of the primary server, the backup server can quickly mount the primary servers storage if a failure is detected.

The CORAID Linux NAS Gateway product CLN20 can be configured as a redundant NAS with automatic fail over protection for high availability applications.

Shared Storage Pool NAS

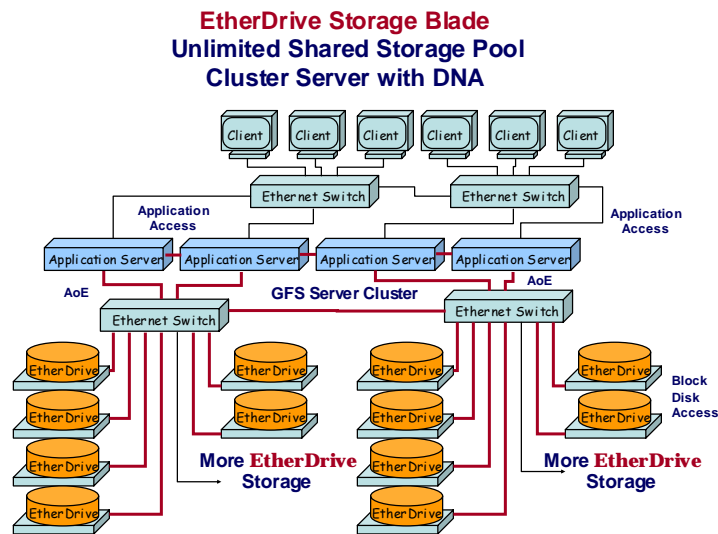
Another variant of NAS is shown here, where a pool of storage is freely shared by multiple NAS heads/gateways. Each server having its own set of EtherDrive Storage Blades on a common Ethernet network.



This eliminates wasted or unused disks that would be captured inside a typical NAS chassis.

Shared Storage Pool GFS

Shown below is an example application for the Global File System (GFS). EtherDrive Storage is uniquely suited for this application and can provide the networked storage underneath a cluster of servers using GFS to have simultaneous read/write access to the AoE storage devices.



Summary

EtherDrive Storage provides a uniquely flexible storage alternative for any server application. Using familiar Ethernet networking technology networked storage can now be afforded by any application.