security-assessment.com

White Paper

Access over Ethernet: Insecurities in AoE

Prepared by: Carl Purvis
Independent Security Researcher
&
Morgan Marquis-Boire
Security Consultant
Security-Assessment.com

Date: August 24, 2006

Abstract

This paper investigates the insecurities present in the ATA over Ethernet (AoE) protocol. It is clear that this protocol has been designed with ease and simplicity as its primary focus, and as such there are security concerns which require understanding before AoE is used with critical information.

What is AoE?

ATA over Ethernet (AoE) is an open standards based protocol that allows direct network access to disk drives by client hosts. It has been developed by Coraid™ ("The Linux Storage People") as a SAN technology and it has been adopted for use by many Universities and US Government agencies[1]. Coraid provides a hardware AoE cluster implementation called EtherDrive™. The Coraid website has downloadable case studies from NASA and the University of Alaska. The claim is that "AoE delivers a simple, high performance, low cost alternative to iSCSI and FibreChannel for networked block storage by eliminating the processing overhead of TCP/IP." Support for AoE is native in the linux kernel as of version 2.6.11.

AoE is a stateless protocol which consists of request messages sent to the AoE server and reply messages returned to the client host. Some messages contain ATA commands, and any data associated with the transaction. Other messages relate to the Config/Query feature of the protocol, to set and query a small amount of out of band data. The formats of these messages are simple and have two forms: ATA messages, and Config/Query messages. Both share a common header format that facilitates network delivery. AoE utilizes the standard Ethernet MAC header for IEEE 802.3 Ethernet frames. AoE has a registered Ethernet type of 0x88A2.[2]

[1] http://www.coraid.com/
[2] http://www.coraid.com/documents/AoEDescription.pdf

AoE Header Format:[3]

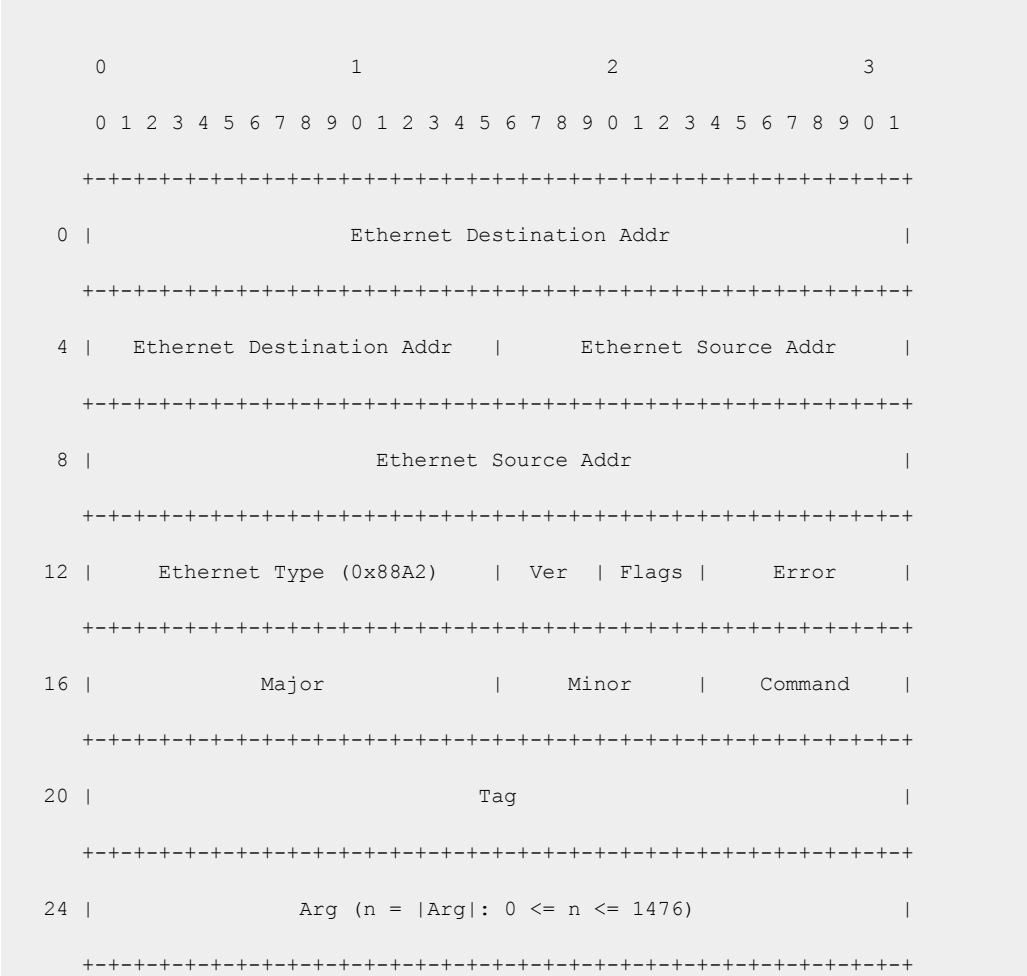

The AoE software provided at http://aoetools.sourceforge.net/ comes in two main parts:

Aoetools – this is a suite of client tools for using AoE exported disk.
Vblade – This is the AoE server software

Kernel module support is available for Linux, Solaris, Windows and soon OS X.

[3] http://www.coraid.com/pdfs/documentation/AoEr8.txt

Why Attack AoE?

While there is increasing industry adoption of this technology, there is no publicly available security research about AoE. This means that there is a general lack of customer awareness regarding the security implications of AoE deployment in the enterprise. There are no vendor recommendations with regard to proper segregation of security domains. If an AoE server is connected to multiple security domains the extent of compromise could be widespread. The compromise of a single account or operating system is trivial compared to the possible damage done if the disk store of an entire organization is compromised. In a commercial shared SAN environment where several enterprises share the same disk, an attack on AoE infrastructure could lead to multiple hosts, or even multiple organizations being compromised.

Security Claims

The AoE protocol appears to have been designed primarily with simplicity in mind. This is, as a rule, a good thing. Unfortunately, there appears to have been little concern given to security. The fact that the protocol is (in theory) not routable is given as a method of ensuring data security;

"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." [4]

The security of AoE relies on the security of the infrastructure surrounding it. The protocol itself lacks strong authentication. This is, in fact, a feature of the protocol. It is specifically stated that "storage is not captive to one server, therefore if a server failure occurs its storage can be mounted by a backup server".[5]

Coraid's hardware AoE solution, Etherdrive™ implements basic security via MAC address filtering;

"MAC address filtering allows the storage appliance to restrict access to disks based upon host MAC addresses. With MAC address filtering, only hosts with allowed MAC addresses can access specified logical devices within the SR420 Storage appliance."[6] Future support for VLANs is planned.

Additionally, Etherdrive™ appliances offer restricted disk access via a "configuration string" which is written to the storage device. Only requests which contain this string will be honored by the device.[7]

[4] http://www.linuxjournal.com/article/8149
[5] http://www.coraid.com/simple_connections.html
[6] http://www.coraid.com/documents/EtherDriveSR420.pdf
[7] http://www.coraid.com/technology.html

Attacks

Due to the stateless nature of the protocol and lack of authentication the following attacks on AoE are possible:

- Replay attacks
- Unauthenticated disk access
- AoE Proxying
- Malicious Server
- DOS attacks

For the purposes of the following:

"Client" describes the host mounting the disk exported to the network via AoE.
"Server" describes the host exporting disk via "vblade" software.

For the attacks described below it is useful to imagine that they occur in the following post-intrusion scenario. The below diagram shows a B2B (business to business) connection.

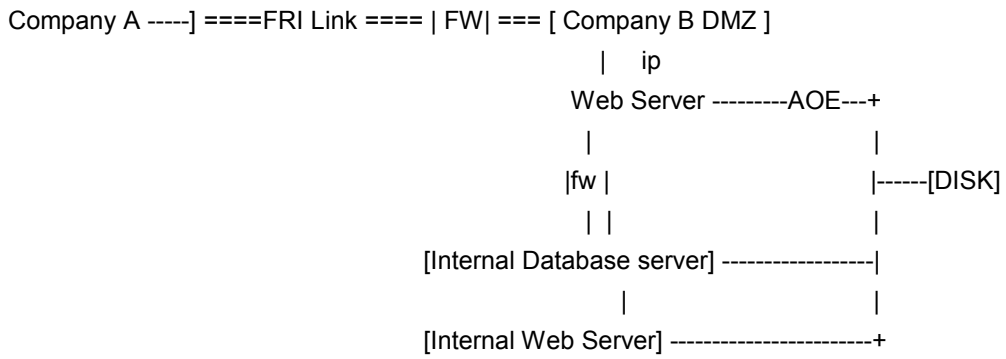

Company A gains access to Company B DMZ via an attack on Company B webserver. This webserver is attached to the same AoE switch fabric as the Company B internal database, allowing for an attacker to access disk on both internal hosts. From this point, several attacks are possible.

Replay Attacks:

This is the most straight forward attack that can be performed on an AoE server. A packet sniffer (such as tcpdump[8]) and a tool capable of replaying network conversations (such as tcpreplay[9]) can be used to facilitate this. AoE traffic is simply captured to a pcap file and then replayed over the network

[8] http://www.tcpdump.org
[9] http://tcpreplay.synfin.net/trac/

tcpdump ether proto 0x88a2 –s 1532 –w file
tcpdump ether host <mac address sending requests> -s 1532 –r file -w file2
tcpreplay –i <interface> --pps=10 file2

This type of attack can be useful for rewriting deleted data to the server, reverting log files to prior states etc.

Unauthenticated Disk Access:

It is possible to read directly from the raw disk which is exported to the Ethernet network. In order to read from the exported drive as raw disk we use the forensic tools for raw disk access available from the Sleuthkit[10]:

fls –lr /dev/etherd/e0.0

This gives us the inode name of the file we wish to read.

icat –v /dev/etherd/e0.0 <inode no of file to be read>

It is also possible to write to the disk. It is, however, important to remember that we are writing to the disk directly, without the knowledge of the client. The impact of this is that the file system will not be updated in the same manner as it would be if it were written to while mounted. In this case, it is recommended to over write only fixed-length strings with data of the same length (password hashes are a good example of something that is easy to over-write in this scenario). File system corruption may occur if more ambitious writes are attempted. Client kernel disk cache is unaware of writes that occur in this manner. It may take some time for such changes to become apparent. If the file is cached in the client's kernel disk cache and the client updates the file, then this will overwrite the illegitimate changes that have been made.

A write maybe attempted as below:

dd if=/dev/etherd/e0.0 bs=1 skip=<disk offset> count=<file size> | dd of=dodgyfile
dd if=dodgyfile of=/dev/etherd/e0.0 bs=1 seek=<diskoffset> count=<file size>

(ie. assume "dodgyfile" is /etc/shadow and we are replacing a password hash)

[10] http://www.sleuthkit.org/

AoE Proxying:

Making the AoE protocol routable is possible for a multi-homed client. The client need simply run up an instance of vblade (AoE server software) and export the network available disk out of another interface. It is also possible to do this via tunneling interfaces (watch for MTU mis-matches).

vblade 0 0 <network device> /dev/etherd/e0.0

(where /dev/etherd/e0.0 is the disk that has been exported to the client)

Malicious Server:

This attack is similar to the AoE Proxying attack. A man-in-the-middle attack is possible by flooding the client with config reponses which match the same shelf and blade number of the real disk, but instead, have the MAC address of the malicious server. The malicious server then proxies the read/write requests of the client back to the real server. Handcrafting of AoE config/discover packets is required.

DoS Attacks:

Attacks of this nature are trivial given the earlier demonstration of unauthorized disk access. It is simple to either overwrite the file system (removal of all data) or swap (crash system).

Overwrite filesystem:
cat /dev/zero > /dev/etherd/e0.0

Overwrite swap partition (system crash but no data damaged)

fdisk /dev/etherd/e0.0 (find swap partition)
cat /dev/zero > /dev/etherd/e0.0p<partition number of swap>

The nature of the Denial of Service attack that can be performed is constrained by what disk we have available to us with which to perform our attack.

Mitigations

Coraid's hardware AoE product, EtherDrive supports MAC filtering.[11] If MAC filtering is also enabled correctly on the switch infrastructure this provides a certain level of security. In this case however, it is possible that various attacks on the switch (such as cam table flooding), could be possible to bypass this security feature. These attacks however, are outside of the scope of this whitepaper. If MAC filtering is not enabled on the switch layer, then client MAC theft is possible. This is an active and invasive attack which will result in lack of client service. After successfully performing this attack, it is subsequently possible to utilize the "Malicious Server" technique described earlier as the client will have to reconnect to the server.

The EtherDrive disk restriction mechanism via "configuration string" described earlier can be easily bypassed with packet forgery. The packet containing the configuration string can be sniffed and replayed, or once the configuration string is captured, it can be embedded in a forged packet. It may also be possible to either guess or brute-force the "configuration string" used for authentication in order to gain unauthorized access to the disk.

Securing the AoE infrastructure to ensure separation between clients in different security domains will alleviate the problems described herein. If both the server and the switch support 802.1q VLAN trunking then the following process will provide an AoE infrastructure which is resistant to the attacks described in this whitepaper:
Configure an AoE server with multiple physical interfaces and export one logical array per interface per client. Configure VLAN trunking on both the server and the switch. Each AoE connected client will be in a separate VLAN.
While these steps may provide adequate protection for the data on your SAN, management overhead is increased and the inherent insecurity of the AoE protocol remains. The security lies with the infrastructure which itself needs to be configured correctly in order to be properly resistant to attack.

[11] A patch to provide MAC filtering for the linux userland server was published on the AoE mailing list by Fran Firman. http://aoetools.sourceforge.net

Final Summary

It is important to realize that the main risk with AoE is the possibility of deployment in a shared SAN environment. Use of AoE across multiple security domains is highly discouraged.

ATA over Ethernet is not yet a mature and secure technology. It should be utilized with care, especially if the data in question is viewed as sensitive.
The security concerns raised in this paper should be part of any discussion on the deployment of this technology into a production enterprise environment.

References

ATA over Ethernet Tools
http://aoetools.sourceforge.net

Tcpreplay - Trac
http://tcpreplay.synfin.net/trac

TCPdump
http://www.tcpdump.org

The Sleuth Kit & Autopsy: Digital Investigation Tools for Linux and other Unixes
http://www.sleuthkit.org

Coraid:: The Linux Storage People
http://www.coraid.com

Linux Journal
http://www.linuxjournal.com

Security-Assessment.com
www.security-assessment.com

About Security-Assessment.com

Security-Assessment.com is Australasia's leading team of Information Security consultants specialising in providing high quality Information Security services to clients throughout the Asia Pacific region. Our clients include some of the largest globally recognised companies in areas such as finance, telecommunications, broadcasting, legal and government. Our aim is to provide the very best independent advice and a high level of technical expertise while creating long and lasting professional relationships with our clients.

Security-Assessment.com is committed to security research and development, and its team continues to identify and responsibly publish vulnerabilities in public and private software vendor's products. Members of the Security-

Assessment.com R&D team are globally recognised through their release of whitepapers and presentations related to new security research.

Security-Assessment.com is an Endorsed Commonwealth Government of Australia supplier and sits on the Australian Government Attorney-General's Department Critical Infrastructure Project panel. We are certified by both Visa and MasterCard under their Payment Card Industry Data Security Standard Programs.

Copyright Information