Access over Ethernet:

Insecurities in AoE

Presented By Morgan Marquis-Boire

Ruxcon2006

security-assessment.com

Hi

- I'm Morgan Marquis-Boire.
- I work as a Security Consultant for Security-Assessment.com
- I'm generally fond of big gear, UNIX, cryptography, privacy (will rant, just add beer).
- Some of you may remember me from Ruxcon 2k5…

Intro

- SAN Basics
- New SAN Tech
- ISCSI/FCIP/iFCP
- What's AoE?
- Attacks on AoE
- Mitigations
- Q & A

SAN vs. NAS

- It’s probably easiest to explain SAN technologies (for those hitherto unfamiliar with them) by contrasting them with another popular disk storage access method (NAS).
- Both generally involve lots of disk (which is cool). Who doesn’t <3 big storage??

Mmmm... Disk.

What is NAS?

- **A Quick primer**
  - Network Attached Storage
  - NAS uses well-understood, universal file based protocols –SMB(CIFS)/NFS.
  - Advantages over local storage – Data Available to many machines. Data present on a host dedicated to file sharing. Probably using a backup solution if you're sane. High Availability of Data.
  - Can be proprietary NAS array (NetApp 'filer'): an embedded appliance
    or
    simply a server running the appropriate software (Linux box + RAID + NFS + SAMBA).

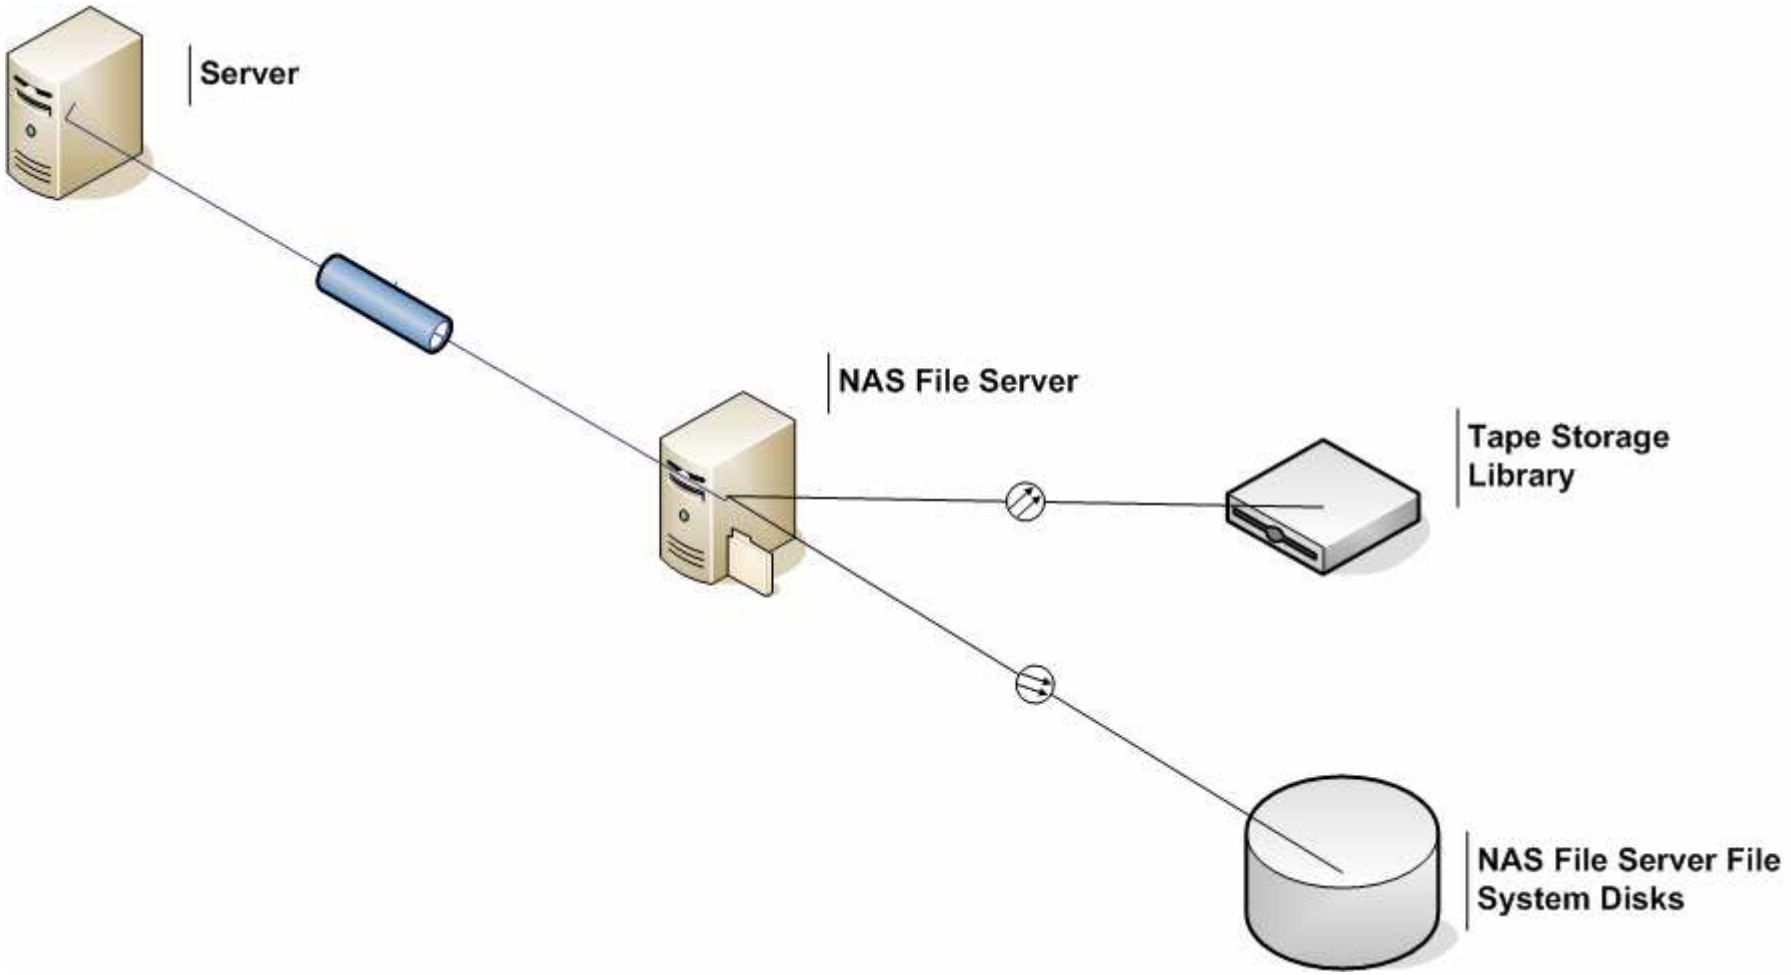
Server
NAS File Server
Tape Storage Library
NAS File Server File System Disks

What is SAN?

- **A Quick primer**
  - Storage Area Network. Conventionally, it's a big load of disk tied together with fibre.
  - A SAN is distinguished from being merely network available storage by the low-level disk access used. (Not just a big share you can copy your warez to)
  - Data traffic on a SAN is very similar to that used for internal disk drives (disk in your local PC), i.e. ATA, SCSI etc.
  - A data request is issued for specific blocks from specific drives.

What is SAN?

- **A Quick primer cont.**
  - Usually built on an separate network infrastructure specifically designed to handle storage communications.
  - Faster and more reliable access than higher level protocols
  - Most commonly SCSI over Fibre Channel. (Traditionally SCSI has been perceived as more reliable).
  - New alternatives are iSCSI and AoE
    - iSCSI – SCSI command set over TCP/IP, most typically Ethernet.
    - AoE – ATA command set inside raw Ethernet frames

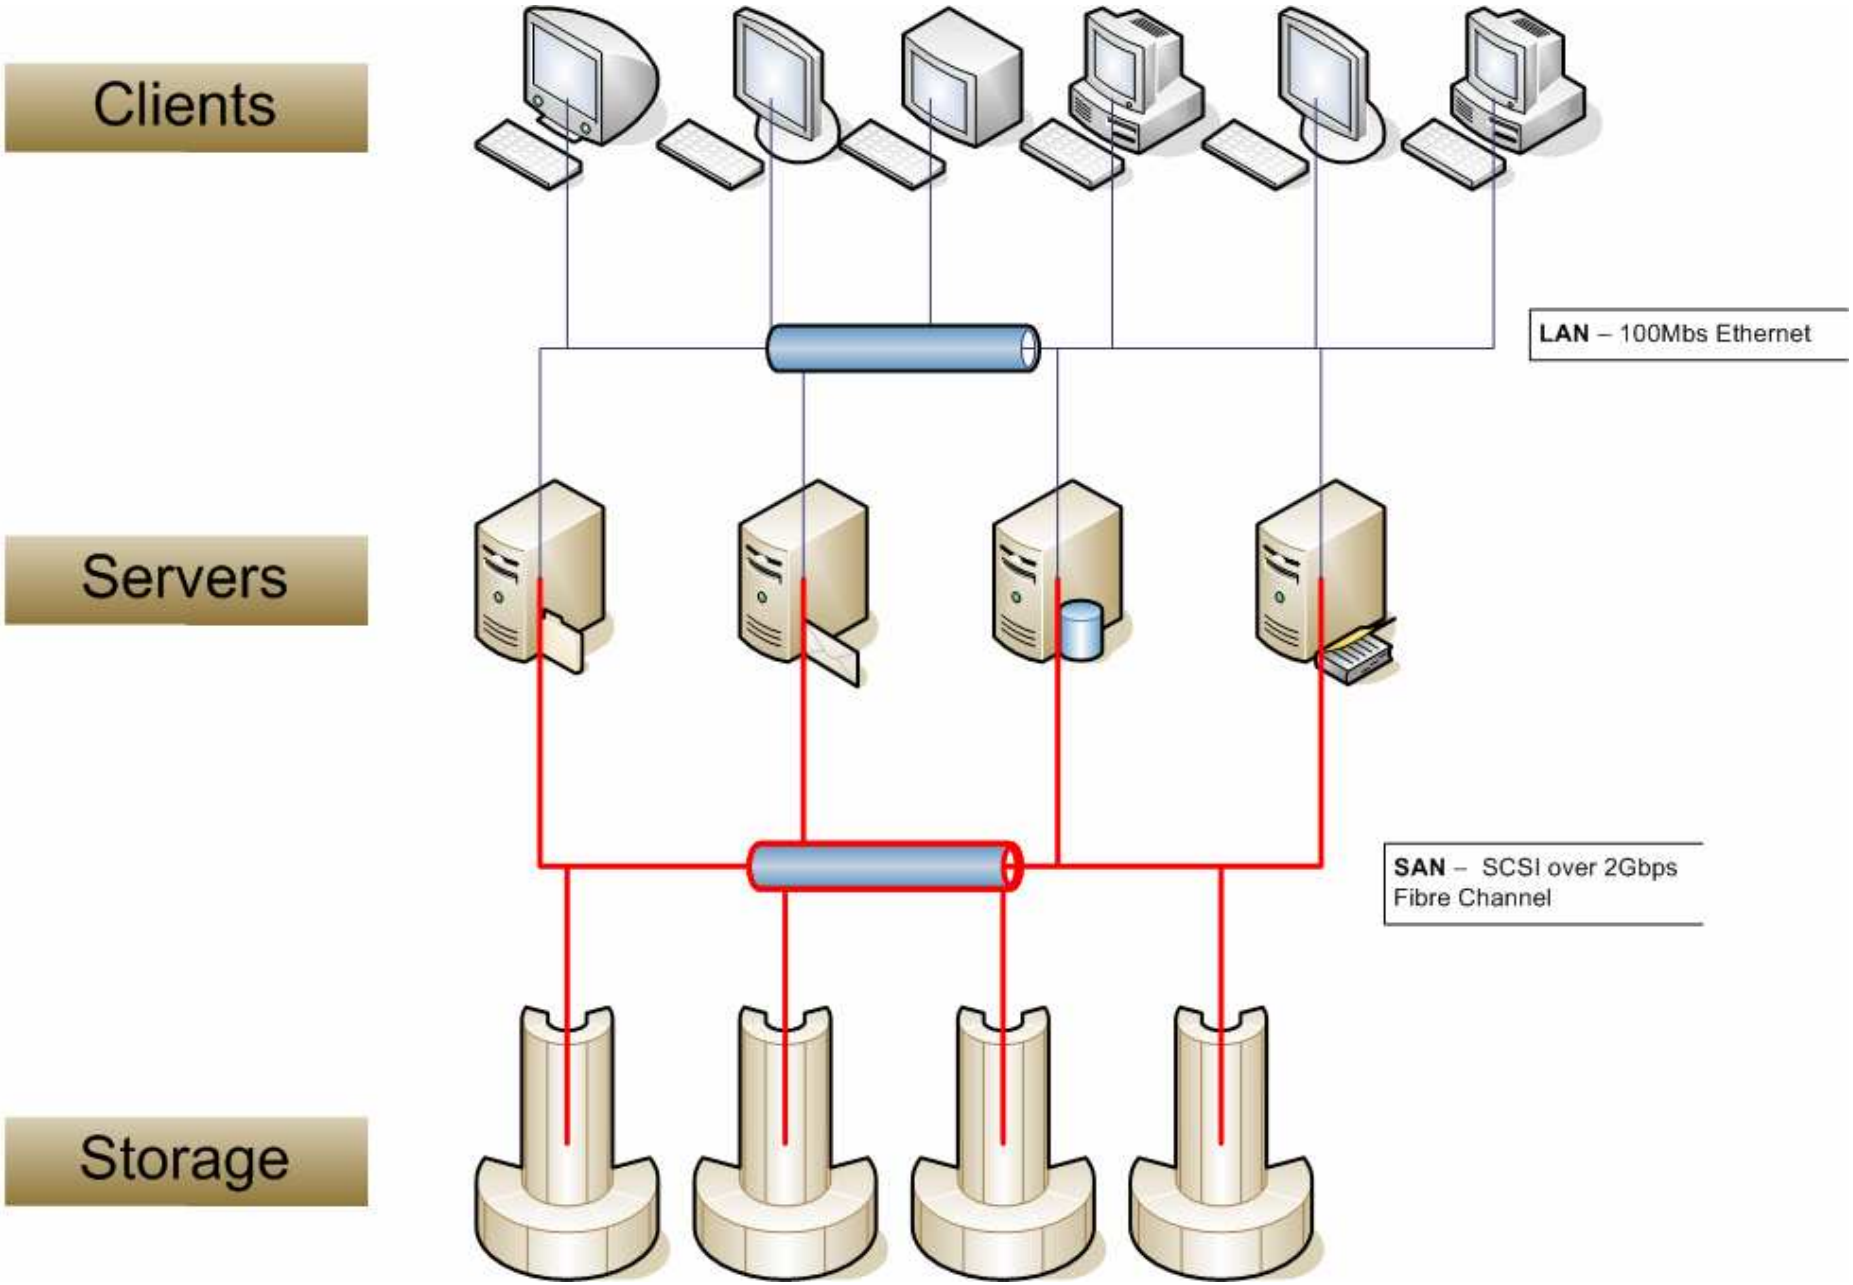
Clients
Servers
Storage
LAN – 100Mbs Ethernet
SAN – SCSI over 2Gbps
Fibre Channel

Times are a chagin’

- **New SAN protocols emerging over the last few years…**

SAN over Ethernet

- **Why all this new SAN tech?**

  – Ethernet was initially deemed to be too slow.

  – Gigabit Ethernet is now more or less ubiquitous. It came with my laptop and my workstation as standard

  – 10 Gigabit Ethernet (10GbE) is the most recent (2006) and fastest of the Ethernet standards. 10GBase-T provides 10Gigabit/second Connections over twisted pair cable using RJ-45 connectors. Category 7 cable is required for full speed at 100 meters. (IEEE 802.3an)

  – WAN network is much faster, you can access storage at useful speeds from geographically separate locations.

FCIP

- Fibre Channel over IP (FCIP, FC/IP, Fibre Channel tunneling, or storage tunneling)
- Enables the transmission of Fibre Channel information over IP networks between SANs.
- Fibre Channel is tunneled.
- FCIP can only be used in conjunction with fibre channel technology (contrast with iSCSI and AoE – existing Ethernet)

iFCP

- Internet Fibre Channel Protocol
- Supports Fibre Channel Layer 4 over IP
- Gateway-to-Gateway Protocol where IP switching and routing can replace Fibre Channel fabric
- The lower-layer of Fibre Channel transport is replaced with IP

Even Newer…

- Still fibre channel….

- We're interested in Ethernet due the ubiquity of existing Ethernet infrastructure

- So, let's move on to Ethernet SAN protocols…

iSCSI (circa 2002)

- Uses IP for it's data transfer.
- Access SCSI devices over Ethernet via an iSCSI initiator to provide block level I/O. (That means you can talk low-level to SCSI disk over an Ethernet network)
- Acceptance of iSCSI has accelerated with the proliferation of Gigabit Ethernet (now we're almost in the same world as fibre)
- Support for almost everything AIX/Cisco/Linux/Windows etc.

Be hip with the kids - iSCSI Lingo

- **iSCSI Initiator - An iSCSI client**
- **iSCSI Target – An iSCSI storage device**
- **iSNS - iSCSI Name Service (groups targets and initators)**
- **Domain Sets – Logical segregation of Targets and Initators**
- **iQN – Iniator Node Name (Unique identifier similar to MAC addresses)**
- **LUNs – Logical Unit Number. Address for disk device. (Today not usually disk, but partition of a RAID set).**

iSCSI (circa 2002)

- **It's worth noting that iSCSI gains significant speed benefits from utilisation of dedicated hardware**
- **iSCSI host bus adapter (HBA)**
  - A network interface controller that incorpates a TCP Offload Engine with onboard iSCSI processing.

I is for Insecurity

- **At BlackHat2k5 Himanshu Dwivedi wrote a paper about iSCSI**

- **iSCSI authentication is disabled by default**
- **iSCSI authentication uses CHAP**
- **iSCSI is a clear text protocol**
- **iQN Values are trusted and are sniffable, spoofable and can be brute forced.**

Fake server attacks, Man-in-the-middle, all the good stuff…

- **Wonder how that will work against newer SAN protocols??**

What is AoE?

- **Definition**
  ATA over Ethernet (AoE) is an open standards based protocol that allows direct network access to disk drives by client hosts

- **Native in the Linux Kernel as of 2.6.11**

- "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."

- **Layer 2 Protocol which encapsulates ATA (the command set used by most commodity disk) in Ethernet Frames**
  - An Ethernet request which has in it, give me block 'foo' from disk 'bar' on shelf 'baz' on blade 'bim'

What is AoE?

- **Protocol**
  AoE is a stateless protocol which consists of request messages sent to the AoE server and reply messages returned to the client host

- Messages have two formats:
  - ATA Messages
  - Config/Query Messages

- AoE utilises the standard Ethernet MAC header for IEEE 802.3 Ethernet frames and has a registered Ethernet type of 0x88A2

AoE Packet Header

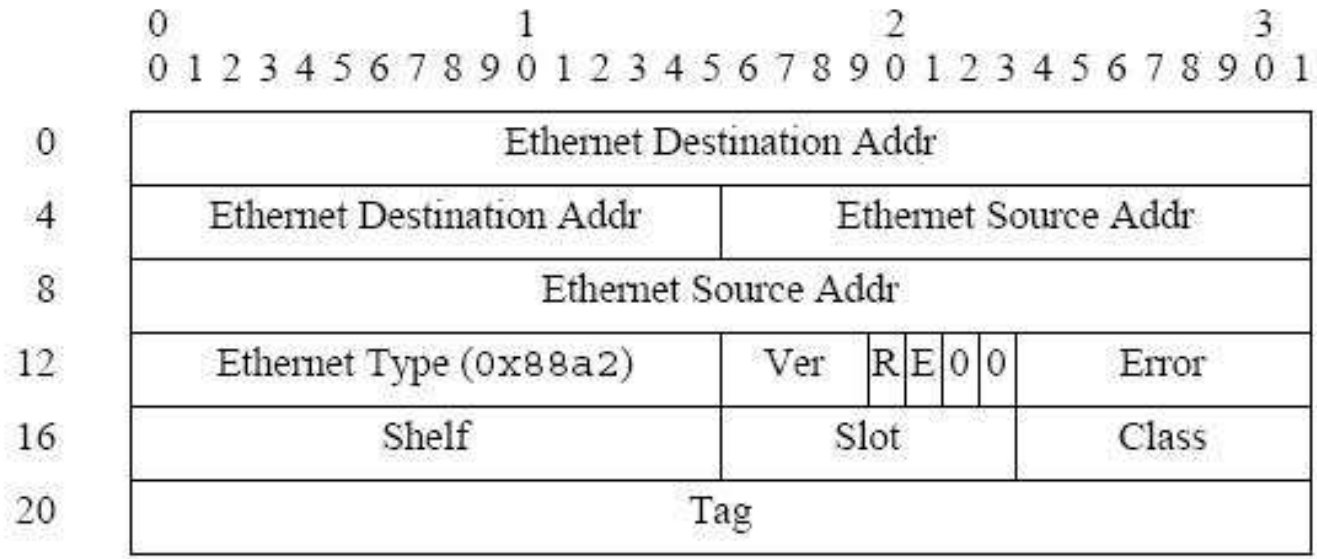

AoE Tools

- **Userland Tools:**
  - Aoetools – suite of client tools for using AoE exported disk
  - Vblade – AoE server software

  Available from:

  http://aoetools.sourceforge.net

  Kernel module support for Linux, Solaris, Windows, and OS X (client needed in order to mount the disk, disk is exported by userland tools, probably faster if done in kernel)

- **Hardware AoE Sold by Coraid**
  - EtherDrive – Hardware AoE solution. Clustered blades. Looks cool. Don't have an AoE blade cluster. Wish I did.

AoE Device Discovery

- **AoE Device Discovery**
  - AoE device discovery is done in two ways. When the AoE device initializes it will broadcast a message so that potential users will know it exists. Secondly, the client may send out a message to probe for devices.
- **Any device on the AoE switch fabric can see all disks that are exported.**
  - When an aoe disk is exported, this information is broadcast and all aoe-capabale clients on the network will see it
- **This is why it's a REALLY dumb idea to deploy AoE across security domains!**

Why AoE is Cool

- **Heavily Reduces Cost**
  This is pretty much a no brainer. ATA vs SCSI, Ethernet vs Fibre switch fabric. The costs are a fraction.

- **Open Standards**
  While most SANs work with proprietary technology, AoE is an open standards solution. Source code to server and userland tools are available to you. (GPL)

- **Easy to Implement**
  AoE is really simple. REALLY SIMPLE.

- **My First Home SAN**
  You can implement this at home with a bunch of commodity disks and a Linux box.

My Home SAN

AoE Setup

- **Let’s see how easy it is to setup a server with a video!!!**

AoE Setup

- **Let’s see how easy it is use as a client with another video!!**

Why Attack AoE?

- **Increasing Industry Adoption**
  The Coraid website has downloadable case studies from NASA and the University of Alaska. Recent Slashdot exposure.

- **Lack of Security Documentation**
  Lack of customer awareness regarding security implications of AoE deployment. No vendor recommendations with regard to proper segregation of security domains

- **Widespread Effects of Compromise**
  The compromise of a single account or operating system is trivial compared with the possible damage done if the disk store of an entire organisation is compromised.

Security Claims

- **Non-routable**
  This protocol is layer-2 only.
  - "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."

- **MAC Filtering**
  EtherDrive Only. Although there is a patch for the free linux tools

- **Configuration String**
  EtherDrive Only.

- **Designed for simplicity and ease of use, not security.**

- **The following attacks made possible due to the stateless nature of the AoE protocol and the lack of authentication**

Attacks

- **For the purposes of the attacks:**
  "Client" describes the host mounting the disk exported to the network via AoE

  "Server" describes the host exporting AoE disk via "vblade" software.

  It is perhaps easiest to imagine the attacks occurring in the following scenario…

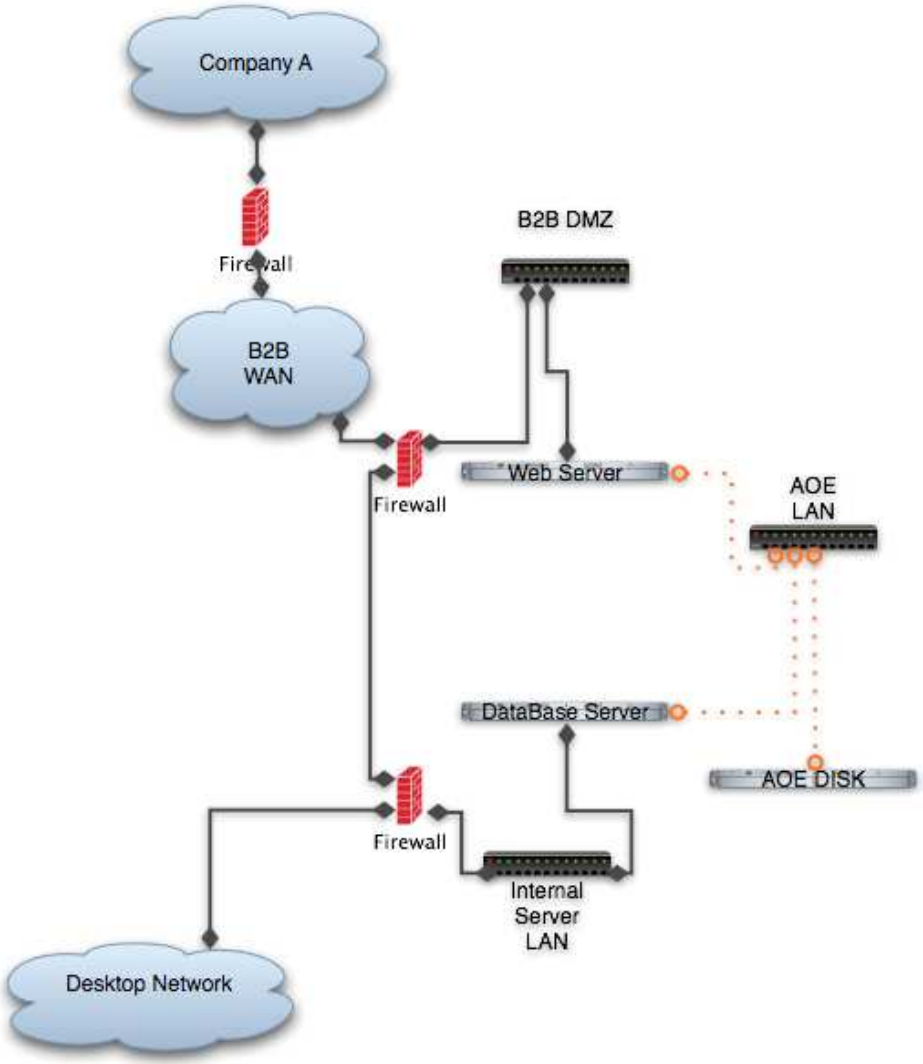

Attack Scenario

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

Attacks

- Replay Attacks
- Unauthenticated Disk Access
- AoE Proxying
- Malicious Server
- DOS Attacks

Closer look at the AoE LAN

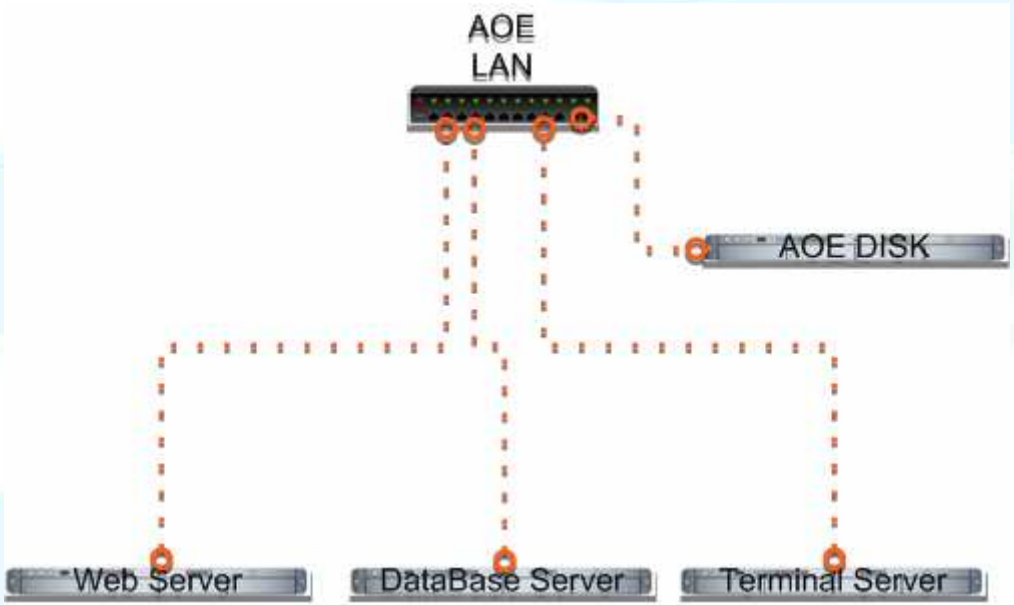

- **Remember, we're the compromised web server ☺**

Replay Attacks

- **This is the most straight forward attack that can be performed on an AoE server**
- **Use a packet sniffer (wireshark, tcpdump) and tool for replaying network conversation (tcpreplay)**
- **Capture the traffic to a file then replay over the network**

```
# 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
```

Useful for rewriting deleted data to the server (reverting logs or re-adding user accounts)

Unauthenticated Disk Access

- **It is possible to read directly from the raw disk which is exported to the network**

- **We want to read the disk without out mounting it. Fortunately there are already existing tools to do this**

- **Enter the Sleuthkit.**
  This is a set of forensic tools provided by Brian Carrier which allows the access of disk at a low-level without mounting the disk or using existing file system drivers.

  # fls –lr /dev/etherd/e0.0

  This provides us with the inode of the file we wish to read

Unauthenticated Disk Access

- # icat –v /dev/etherd/e0.0 <inode no of file to be read>
- Now that we can do read, we can now attempt a disk write.
- This is a more difficult proposition than unauthenticated read. The file system will not be updated in the same manner as it would be if written to while mounted.
- It is probably a good idea to limit write attempts to fixed-length strings (ie password hashes). File system corruption may occur if more ambitious writes are attempted.

Unauthenticated Disk Access

- **Malicious Write attempt**

Pull file off disk:

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

Modify file

Write back to disk

- # dd if=dodgyfile of=/dev/etherd/e0.0 bs=1 seek=<diskoffset> count=<file size>

(assume that "dodgyfile" is /etc/shadow and we are replacing password hash)

Unauthenticated Disk Access

- Let’s see another video!!!

AoE Proxying

- **Making the AoE protocol routable is possible for a multi-homed client**
- **The client need simply run up an instance of vblade and export the network available disk out of another interface.**
- # vblade 0 0 <network device> /dev/etherd/e0.0

  (where /dev/etherd/e0.0 is the network available disk that shows up on the client)

  You could do this via a tunneling interface and export the disk via L2TP!!

Malicious Server

- **Similar to the AoE Proxying attack. Different aim.**
- **Man-in-the-middle is possible by flooding the client (not you, another client) with config responses which match the same shelf and blade number as 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 server.**
- **You'll have to forge the config/discover packets.**
  (But that's really easy, you can just sniff them and change the MAC address)

DOS Attacks

- **This is pretty trivial given that we've established that we have disk access.**
- **What sort of DOS attack we can do depends on what we can write to.**
- **We can always mindlessly overwrite the network available disk:**
  # cat /dev/zero > /dev/etherd/e0.0
- **It's pretty sweet if we can write to swap:**
  # cat /dev/zero > /dev/etherd/e0.0p<partition number of swap>

Mitigations - EtherDrive

- **MAC Filtering**
  Coraid's hardware AoE product, EtherDrive supports MAC filtering.
- This only affords some level of protection if the switch infrastructure also supports MAC filtering otherwise client MAC theft is possible.

- **Configuration String**
  EtherDrive allows for the use of a configuration string. Packets which do not have this string will be ignored by the server.
- This can be bypassed by packet forgery. If it is not possible to sniff the original string then guessing or brute-forcing the string may be possible

Mitigations – AoE General

- **Securing the Infrastructure**
  If both server and switch support 802.1q VLAN trunking then you can implement securely.

- 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 connected AoE client will have it's own VLAN.

- The protocol is without security. The security lies with the infrastructure and it's correct configuration.

Remember

- **Protocol doesn't have any security. Don't deploy across security domains!!!**

Final Word

- New SAN technologies emerging with the ubiquity of high speed networks
- Protocols which remain un-examined are likely to have flaws
- Anyone out there got an new fancy FCAL array that they wanna lend me?

References

- **Access over Ethernet Whitepaper**
  http://www.security-assessment.com/files/whitepapers/Insecurities_in_AoE.pdf

- **iSCSI Security – I is for Insecurity**
  www.blackhat.com/presentations/bh-usa-05/bh-us-05-Dwivedi-update.pdf

Thanks

- Carl Purvis
- Security-Assessment.com
- Tmasky (people just love giving this guy props)

Questions ?

http://www.security-assessment.com

morgan@security-assessment.com