Insiders View: Network Security Devices

Dennis Cox
CTO @ BreakingPoint Systems

CanSecWest/Core06
Vancouver, April 2006

Who am I?

- Chief Technology Officer - BreakingPoint Systems
- Director of Engineering - TippingPoint
- Engineering - Cisco Systems
- Operated an ISP

Today's Talk

- Fact vs Fiction of today's security devices
- How to approach testing the validity of claims
- Some simple math
- Example cases

Approach

- What type of box is it?
  - Look at the mechanical design?
  - Who's runs the Hardware Team?
  - What silicon is it using?
- How big is the company?
  - Sub Contractor?
  - Check for posts!

Finding the kill spot

- Something's cost more than others
- What costs the Box the most?
- Latency is the easiest way...
- The secret is the ...

Our Virtual Device

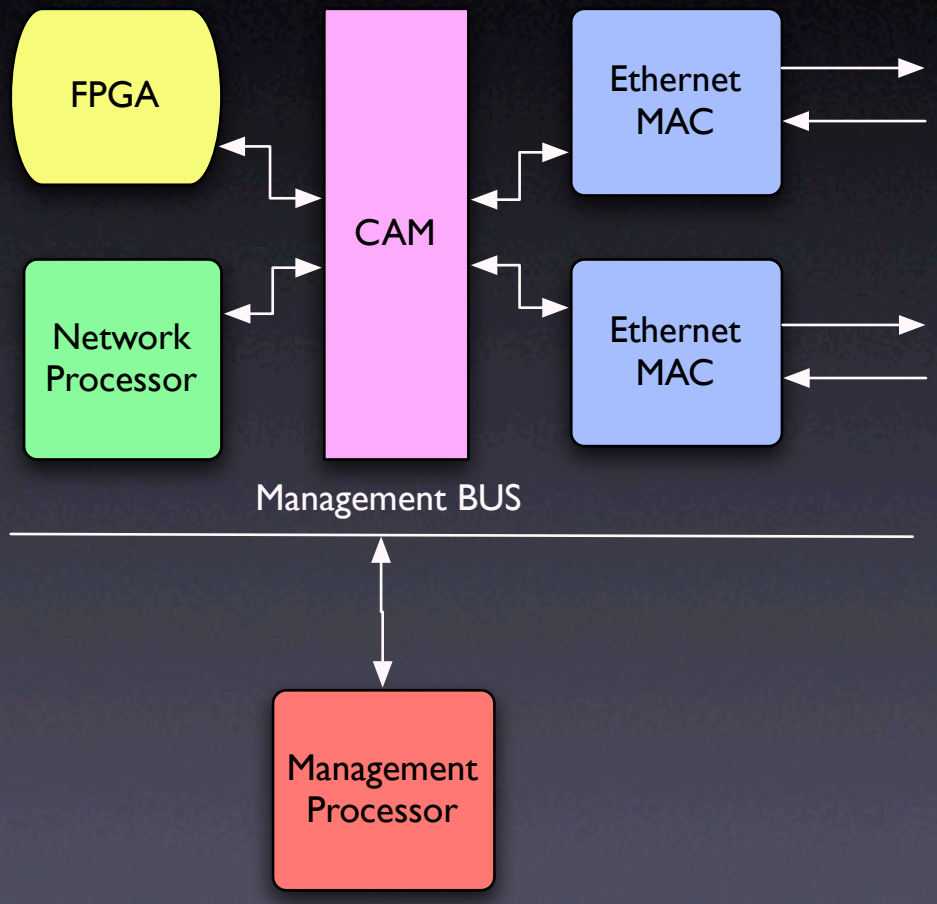

Ethernet MAC

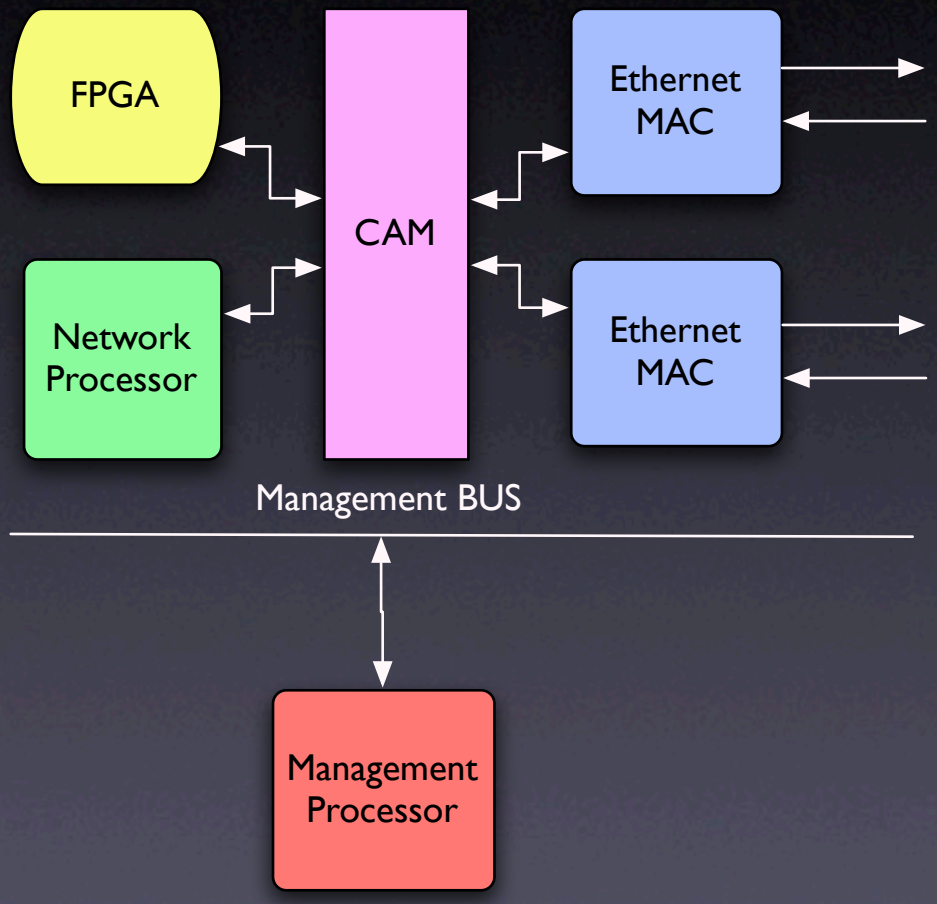

- Who is the vendor?
- What are the specs?
- What revision is the chip? (A0 is sweet, sweet love)
- ETHSIC will get you true love
- Everybody uses the same driver - audit the driver code

Ethernet Frames

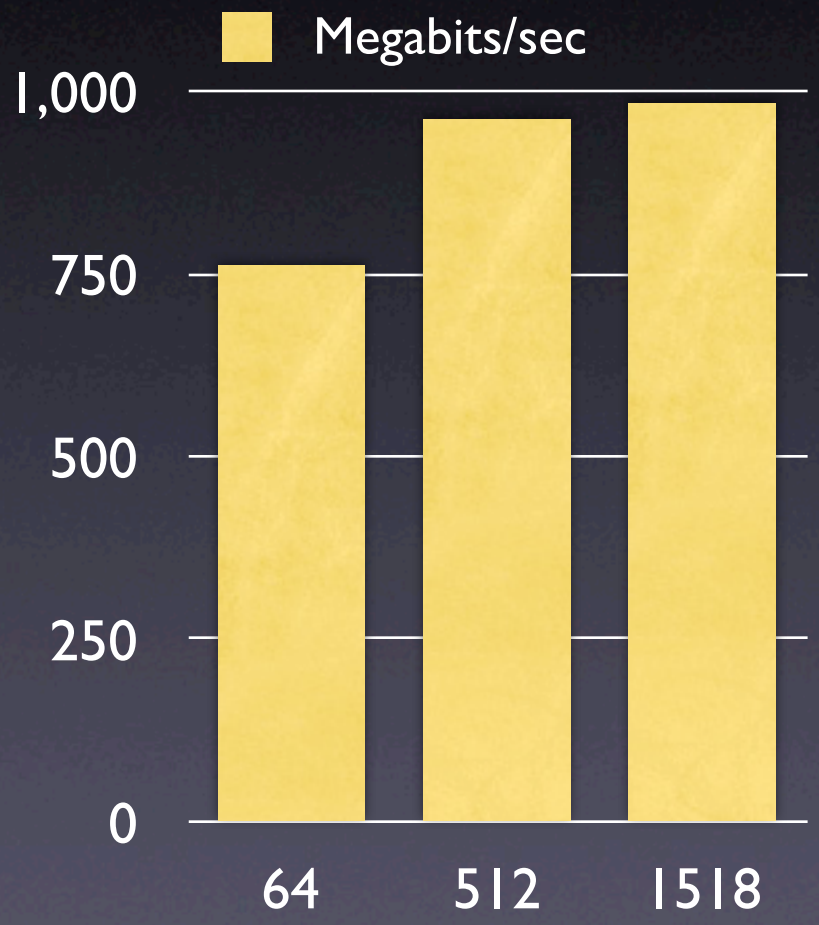

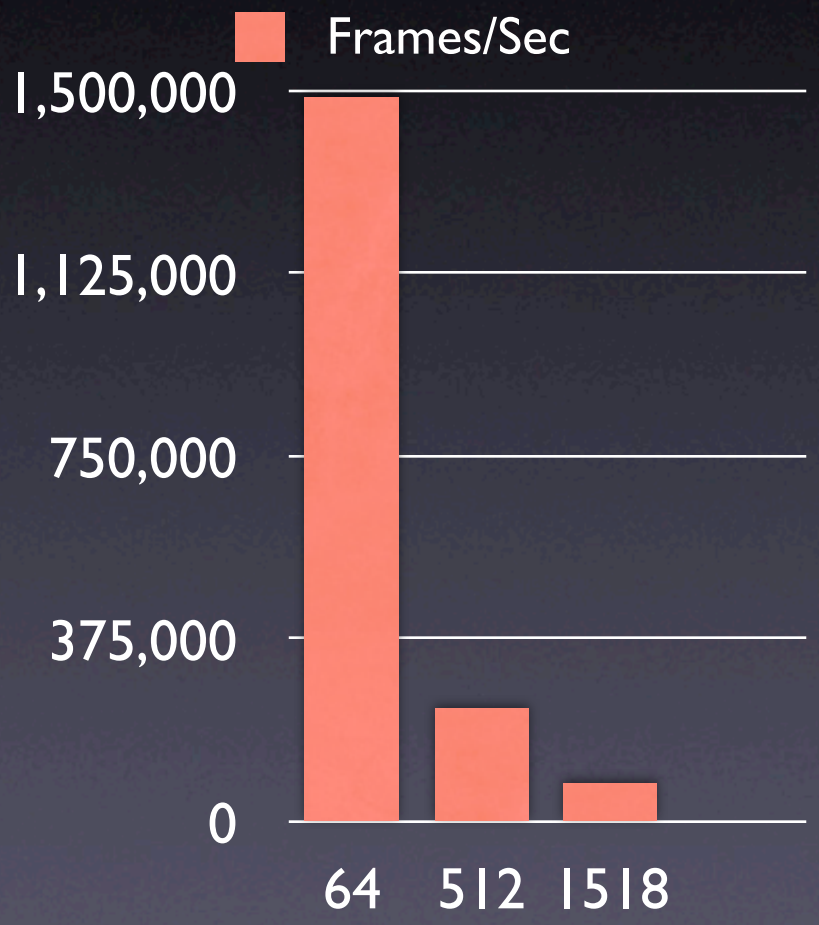

Connection Math

- 70 percent of traffic is TCP (location matters)
- Average TCP packet size ~ 512 bytes
  - (99% < 70 bytes and > 1400)
- 1 Gigabit at 512 bytes equals 244k connections
  - (1,000,000,000 / 8) / 512 = 244k
  - TCP setup is under 3 packets under 70 bytes (generally) which means...
  - Gigabit Ethernet wires can have 1.4 million connections per second happening at any moment in time

* The stats change per about every 9 -12 months. These stats are from November 2004. Source: More sites that I can list (Cable Companies, Telcos, Major Universities and Corporations)

Software Interrupt Stats

- A super high end Ethernet Card
  - (Intel Pro/1000 Server)
- Receive 680,000 pps
- Transmit 840,000 pps
- Full Duplex is still 1.45 million away
- Conclusion: Hardware Systems don't suffer this fate (depending on the hardware system)

Software Performance

- If your using a “Dude it’s a Dell”...
- Your at 761M divided by 2 roughly
- ... 380 Megabits per second

Content Addressable Memory

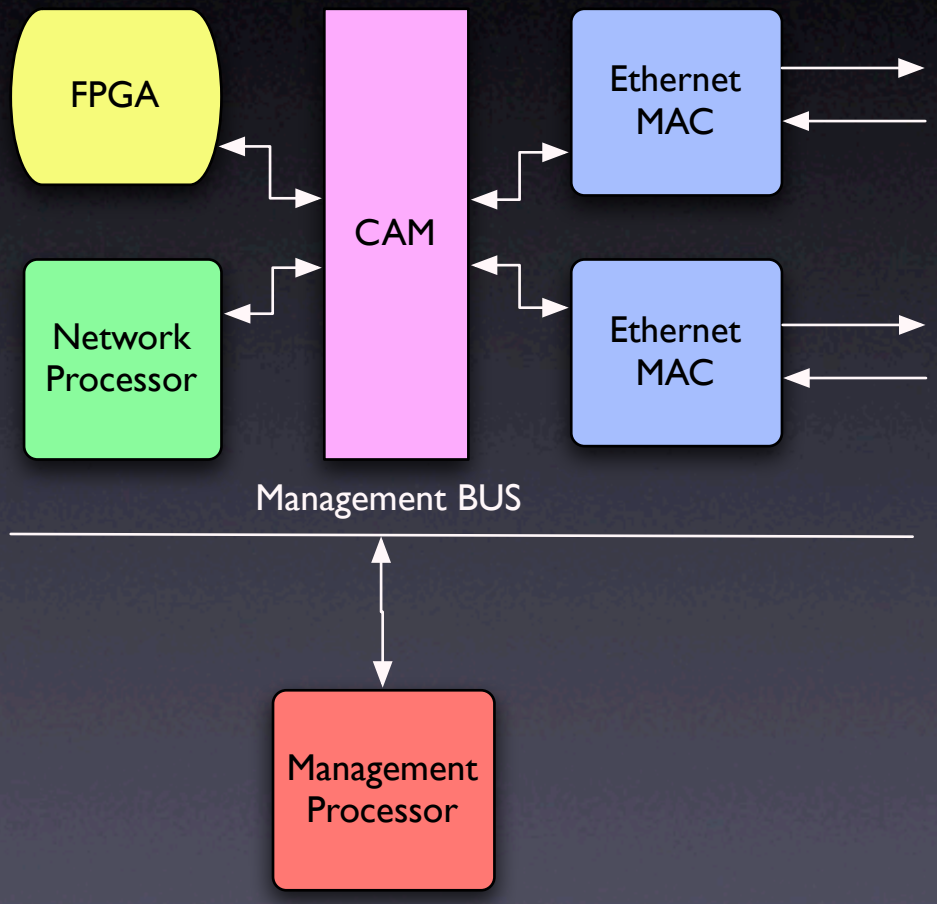

- Same Questions
- Semi Programmable
- Super Fast, Little Flexibility
- Cisco Switches are CAM Based - accessible via SNMP
- Overflow the CAM

Field Programmable Gate Array

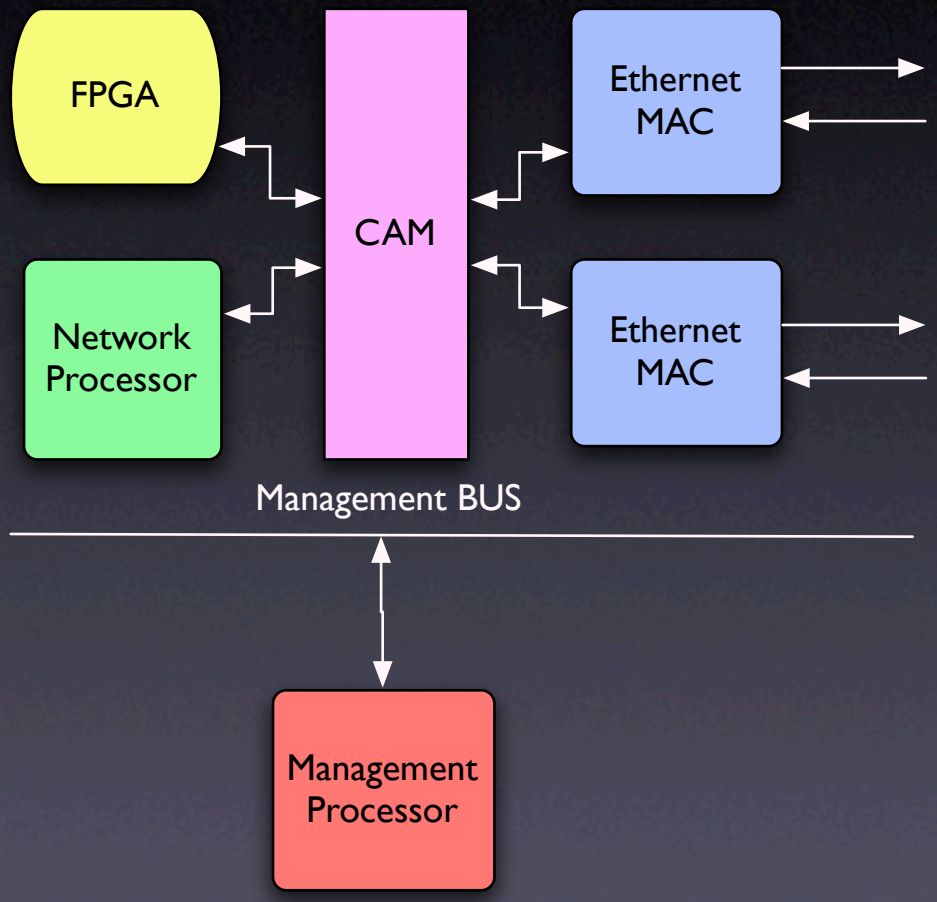

- Questions don't apply
- Very Programmable
- It's a Processor (custom)
- State, State and more State
- Some Security Guy -> Some Software Engineer -> Some Requirements Documents -> Some Design Engineer
- Attack State Machines

Network Processors

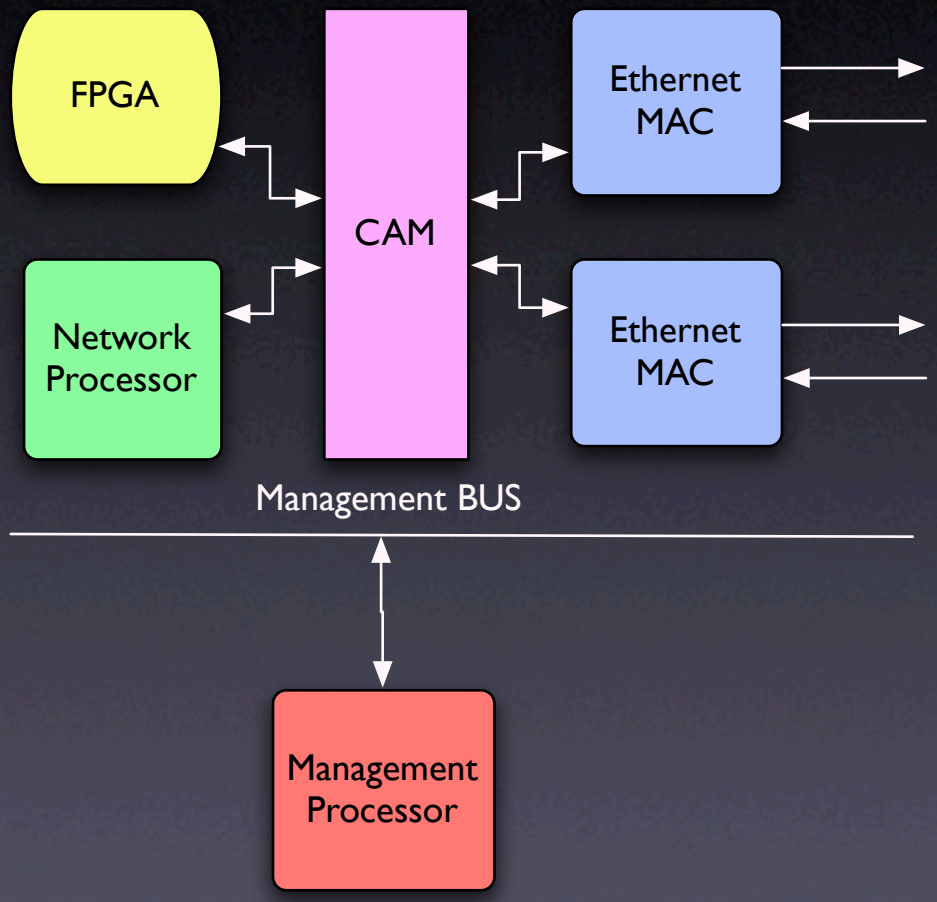

- Questions don't apply
- Programmability is based on the Vendor
- It's a fix field pattern parser
- State, State and more State
- Much stronger on bugs
- Really bad on memory
- Use it's abuse of memory to your advantage

Management Processor

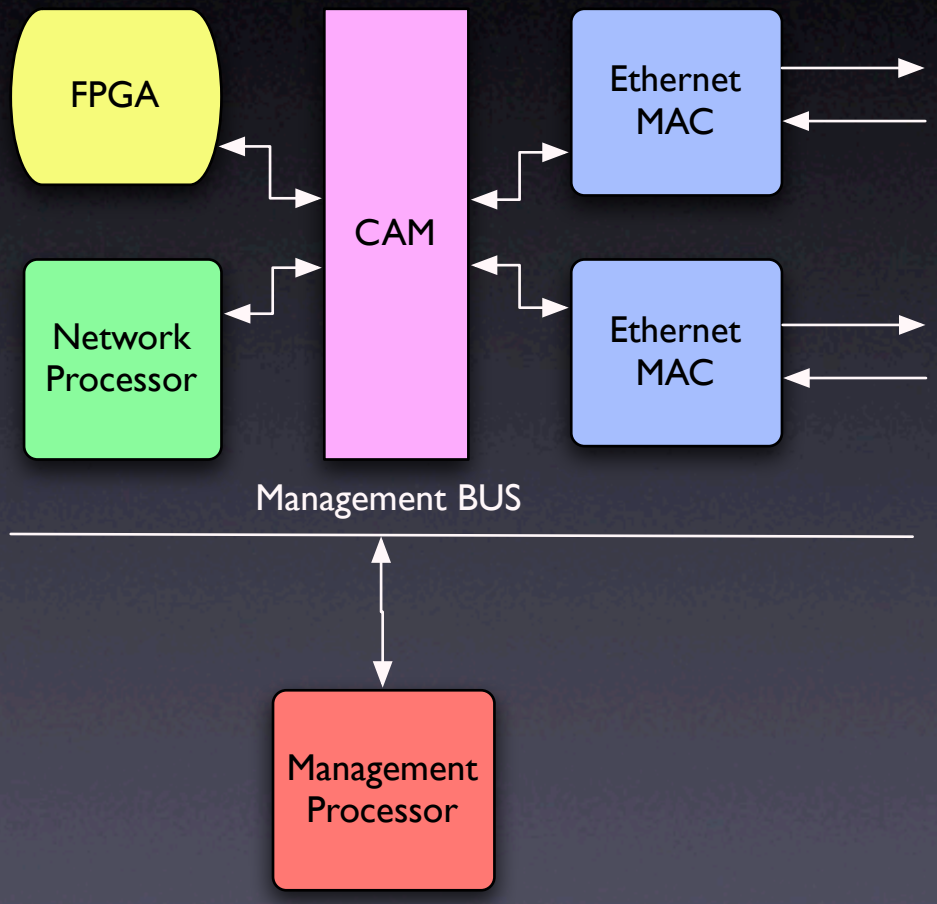

- Just your average, ordinary chip
- If you cause the management interface to be busy, do packets slow down?
- Really bad on memory
- Use it's abuse of memory to your advantage

Exception Processing

- Exception processing or “SlowPath”
  - Most complex devices have one
  - The more complex the request, the better chance it goes there
  - If you can get to the Management Processor via Exception you can root the box or denial of service the box
- Tip: If a device supports encryption, exception handling is constant. You can DDoS with a few Kbytes of traffic.

BUS

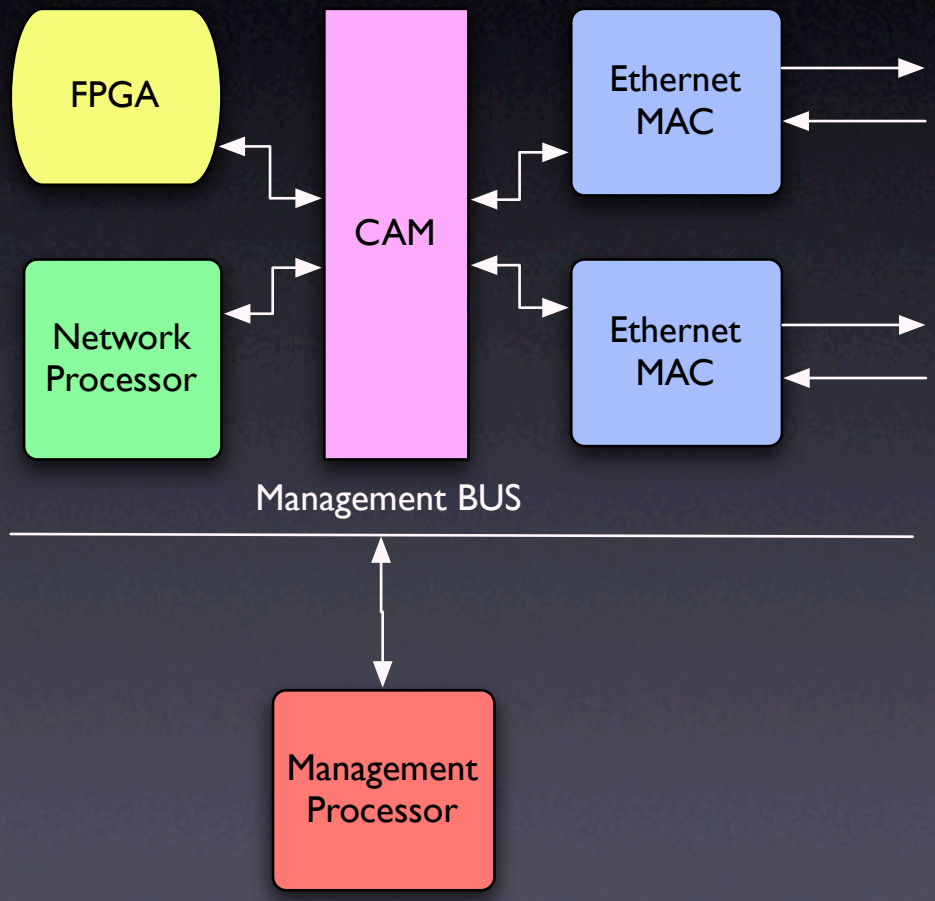

- Multiple BUSes sometimes
- If they are interconnected doesn't matter still weakest link the chain
- Some buses can't handle interleaved packets
- Could you force interleaving of packets?
- Buses use wimpy identifiers - can you modify that identifier?
- A bus has two elements: Max Performance, Max # of Frames
  - Max Frame Size + Max Frames = Max Performance

Bus Math

Bus and Frequency	Peak 32 bit Transfer Rate	Peak 64 bit Transfer Rate	Reality
33-MHz PCI	133 MB/sec	266 MB/sec	972 Mb/s
66-MHz PCI	266 MB/sec	532 MB/sec	N/A
100-MHz PCI-X	N/A	800 MB/sec	2 Gb/s
133-MHz PCI-X	N/A	1 GB/sec	N/A
AGP8X	2.1 GB/sec	N/A	

* Parts of the data are from Dell and Intel's website

Software Boxes

- We already know - limited by BUS
- We already know - limited by Interrupts
- What else do we need to know?

Software Optimizations

- Buffers are the key
- Having too many buffers causes latency
- Buffers are generally not malloc'd
  - Too Slow
- Buffers are set to max packet size
  - If the device supports jumbo frames that's 9k size...
- Too many buffers means slow to access buffers

Buffers Continued

- Fragmentation and TCP Reassembly take up buffers (64k IP + ???? TCP)
  - Generally an additional pool of memory
- Attacks over time based on # of buffers - or worse yet they drop when buffers are full!
- Regular Expressions or Protocol Decoders
  - They take up buffers!

Example - ISS Q1

- First Questions:
  - What type of box is it?
  - Look at the mechanical design?
  - Who's runs the Hardware Team?
- Answers:
  - G1000 has Two Gigabit Ethernet Ports *
  - Repackaged "Dell" Server with a logo on it
  - Nobody runs hardware - they don't have a team **

* Information can be found at http://documents.iss.net/literature/proventia/ProventiaGSeries_Datasheet.pdf

Example - ISS A1

- They use a PCI Bus on that Dell Platform
  - Bus limited to 972 Mbits/s full duplex
- Using Software - so Interrupts come into play
  - 368 Mbits/s full duplex (64 byte packets)
- Using Two Ethernet Controllers
  - Double the Interrupt fun! 184 Mbits/s
- Requires at least double buffering
  - Ethernet 1 to PC to Ethernet 2
- A Dell Server costs $3k (US) max
  - ISS charges $36k (US) for the product

Example - ISS Q2

- Second Questions:
  - What is the rated max concurrent sessions?
  - How does it handle buffers?
- Answers:
  - Rated 1,000,000 Concurrent Sessions
  - TCP Reassembly and Flow Reassembly supported
  - Jumbo Frames Supported

Example ISS A2

- (Flow Reassembly + TCP Reassembly + Max Packet Size) * Max Sessions
- (64k + 9k + 9k) = 82k * 1,000,000
- 82,000,000,000 = 82 Gigabytes of memory
  - Max addressable memory - 4 Gigabytes
- 1,000,000 sessions concurrent can be overflowed on a single Ethernet Wire

ISS - Knowing that

- It most likely can't hit 1 Gigabit per second since it would get killed on small packets
- It can't handle 1 Million connections
  - Can't address that much memory
  - Too many buffer copies
  - No memory for anything else!
  - Even if they could they need to handle more (1.48M)
- Homework: Narrow done which area of memory is the smallest - send partial attack thru that area of memory - fill it up then send the rest of the attack

Example - Netscreen

- Netscreen Filter
  - HTTP (".*/cvsweb\.cgi/.*;.*")
- Running on a 1.5 GHZ G4 using PCRE v6.4
- Standard run (after initial) (100 bytes)
  - Match: 66 usecs
  - Miss: 4 usecs

Example - Netscreen 2

- Increase Data to 1500 bytes
  - Match: 179 usecs
  - Miss: 191 usecs
- Repeating Partial Match (15k)
  - 1.452 seconds*

Example - TopLayer

- “Leader of Intrusion Prevention”
- 4.4 Gbs raw firewall throughput
- 2.0 Gbs rated firewall throughput
- 50k new sessions per second
- 50k sessions tear-down per second
- 1 million Concurrent Sessions
- 1.5 million SYN Flood DOS Protection Rate

* Reference TopLayer Website

Math, Math, Math

- 50,000 is the max session setup
  - 50,000 Connections * 64 Bytes
- Can only achieve 3.2 Mbits per second of new traffic (being conservative)
- Real world testing shows that a TopLayer box can handle 2.5 Mbits of traffic before being DDoS itself
- Math proved it out! Now checkout a Netscreen box!

Device Discovery

- Most inline devices modify packets
- Some change TTL's
- Others reorder TCP Packets
- Did you know some devices even set unique values in packets that come there way?
  - Can you figure out what device does what?
  - Example: TopLayer sets TTL to 255 and TCP Options are changed to MSS=1460

Remember!

- Somewhere on every device the box trusts the packet in some way
- Find that location and you'll get your exploit
- ISS, Netscreen and Toplayer are just examples - no offense to those poor bastards
- Every box has it's Breaking Point

Thank You

dcox@bpointsys.com