Insiders View: Network Security Devices

Dennis Cox
CTO @ BreakingPoint Systems

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

Is it Hardware or Software?

- 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!

Hardware Security Devices

Not only does God play dice, but... he sometimes throws them where they cannot be seen - *Stephen Hawking*

Our Virtual Device

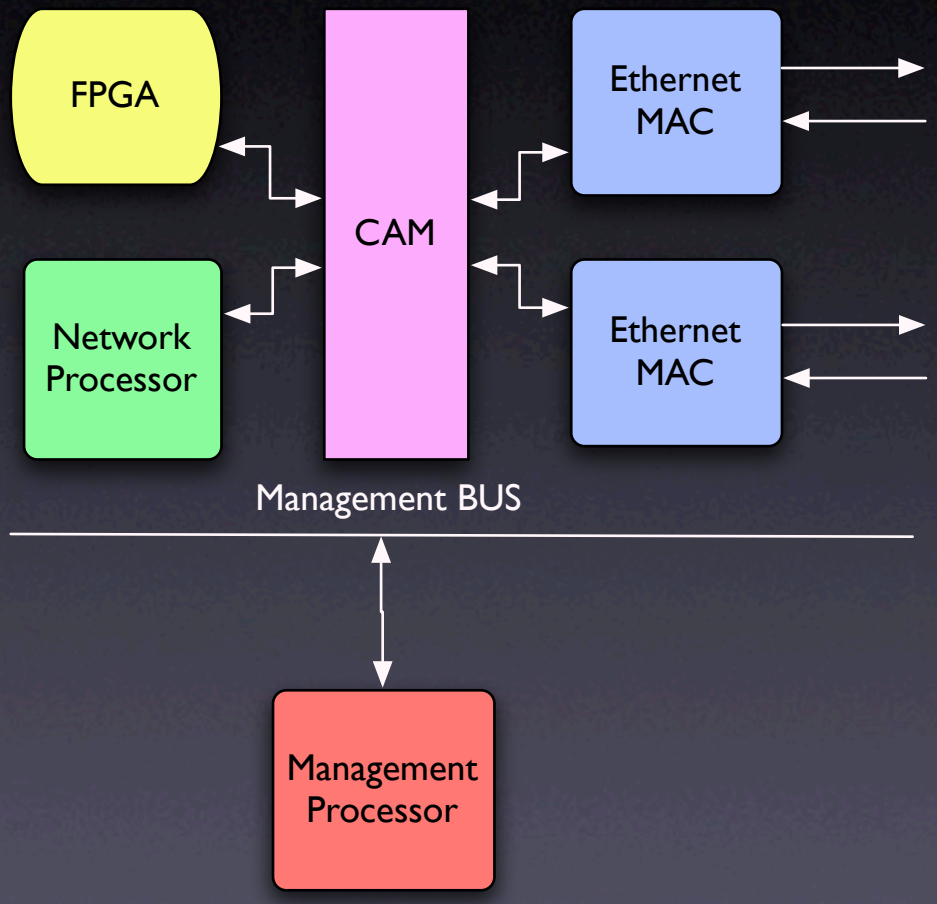

Ethernet MAC

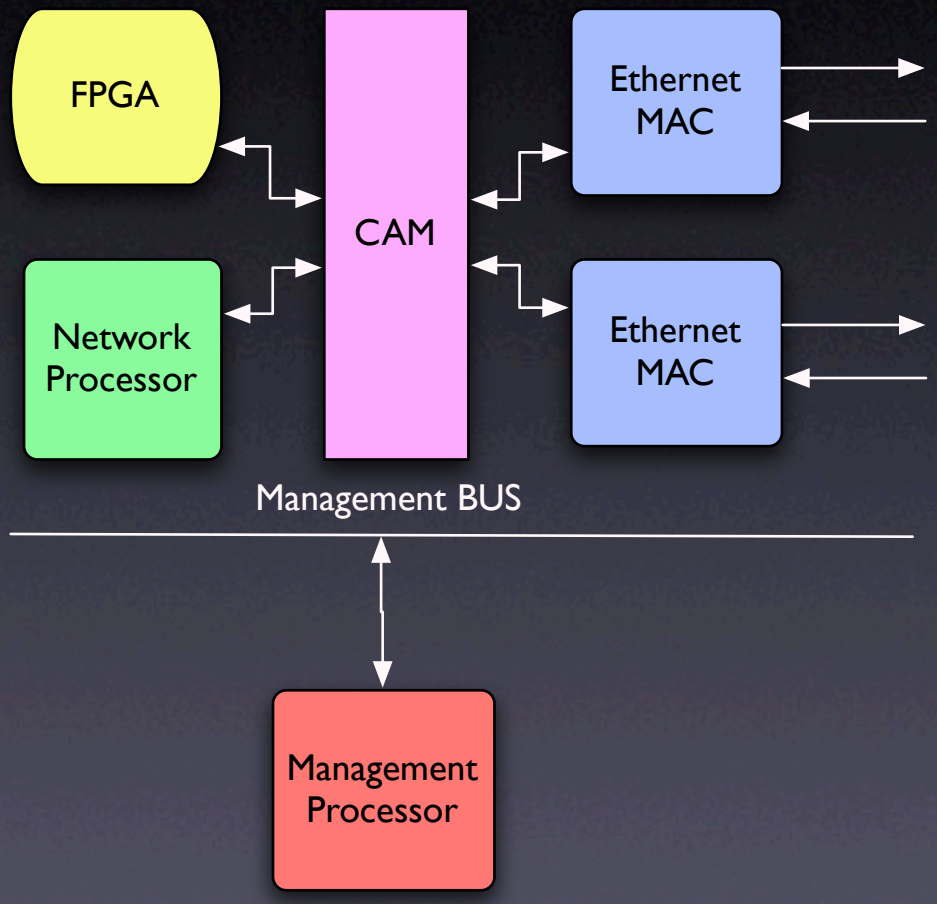

- Who is the vendor?
- What are the specs?
- What revision is the chip? (A0 is sweet, sweet love)
- ESIC will get you true love
- Everybody uses the same driver - audit the driver code

Ethernet Frames

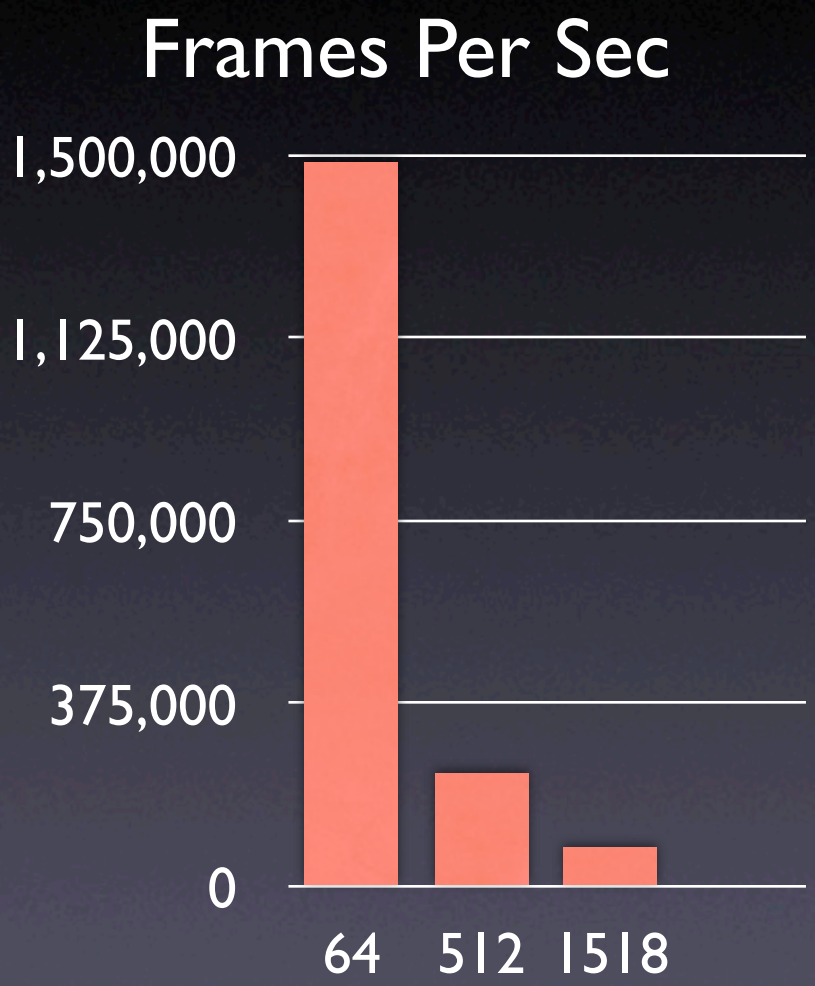

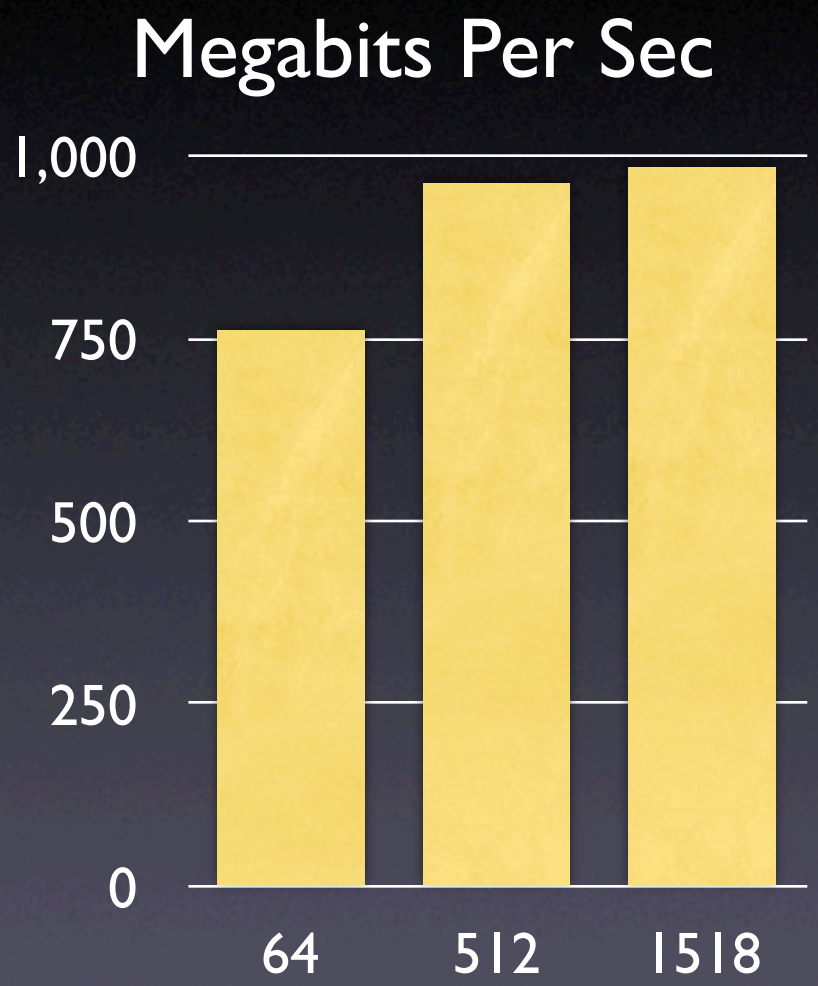

Content Addressable Memory

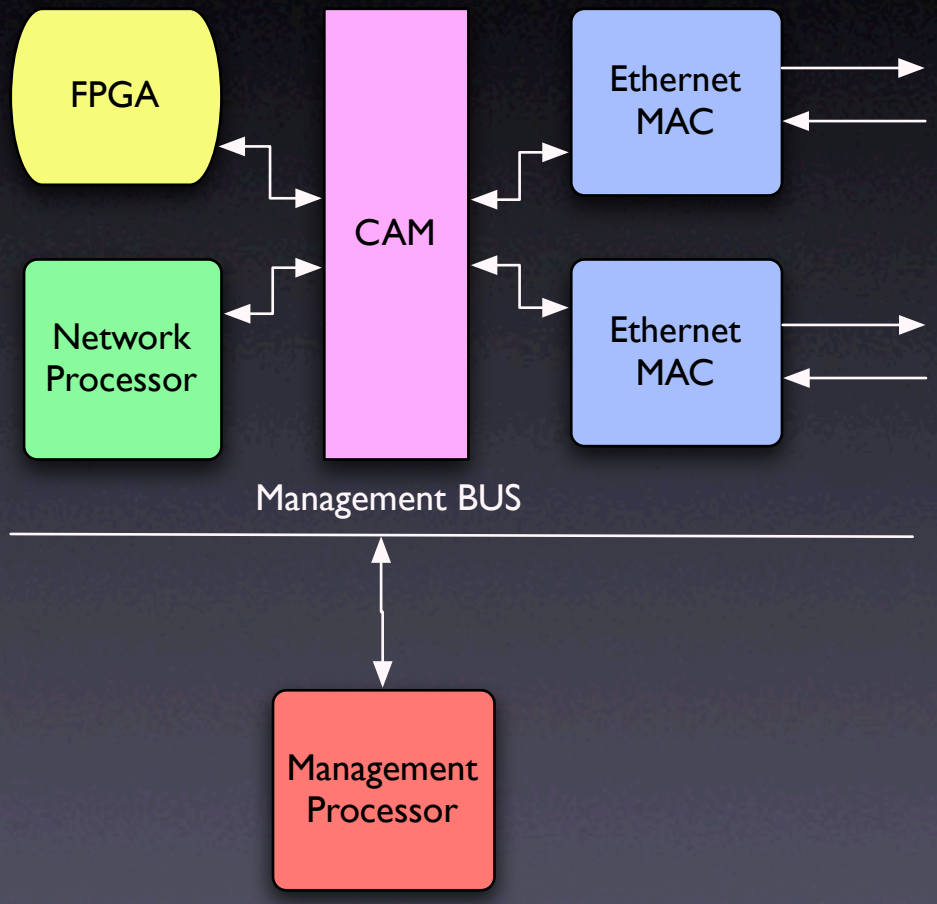

- Same Questions
- Semi Programmable
- Super Fast, Little Flexibility
- Cisco Switches are CAM Based - accessible via SNMP
- Overflow the CAM

Field Programmable Gate Array

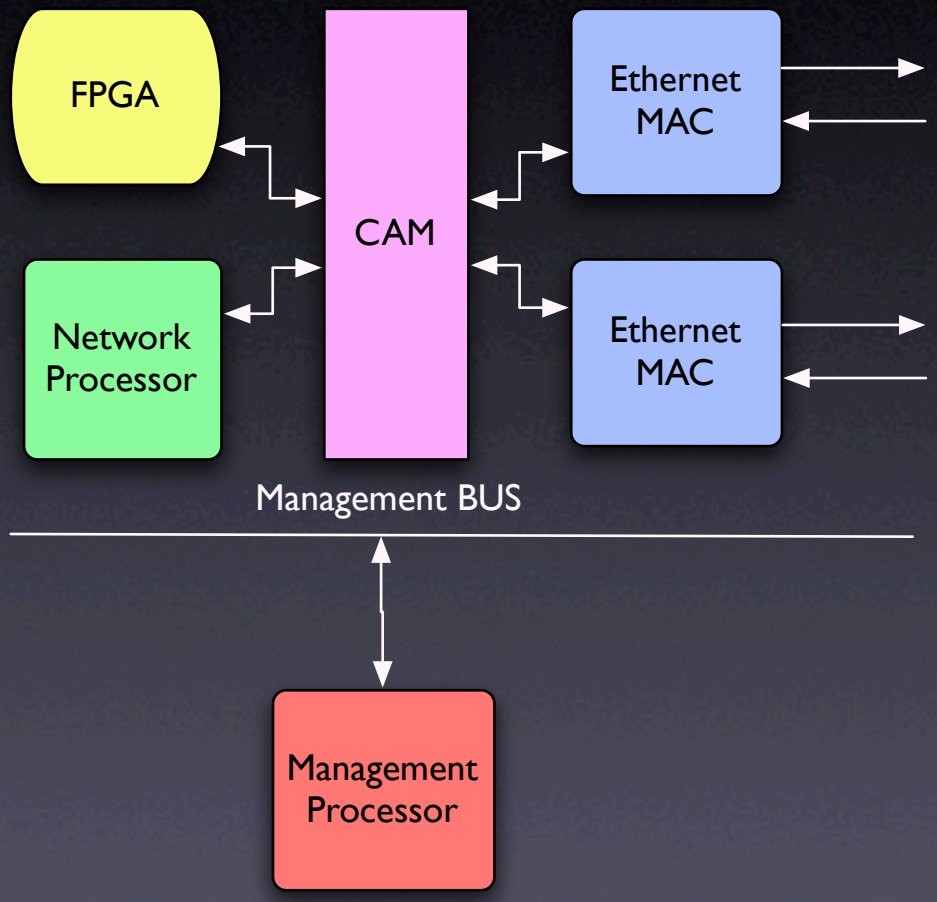

- Questions don't apply
- Very Programmable
- It's a Processor (custom)
- Some Security Guy -> Some Software Engineer -> Some Requirements Documents -> Some Design Engineer
- Attack State Machine and Parsing Engine
- Abnormal QA cycle

Network Processors

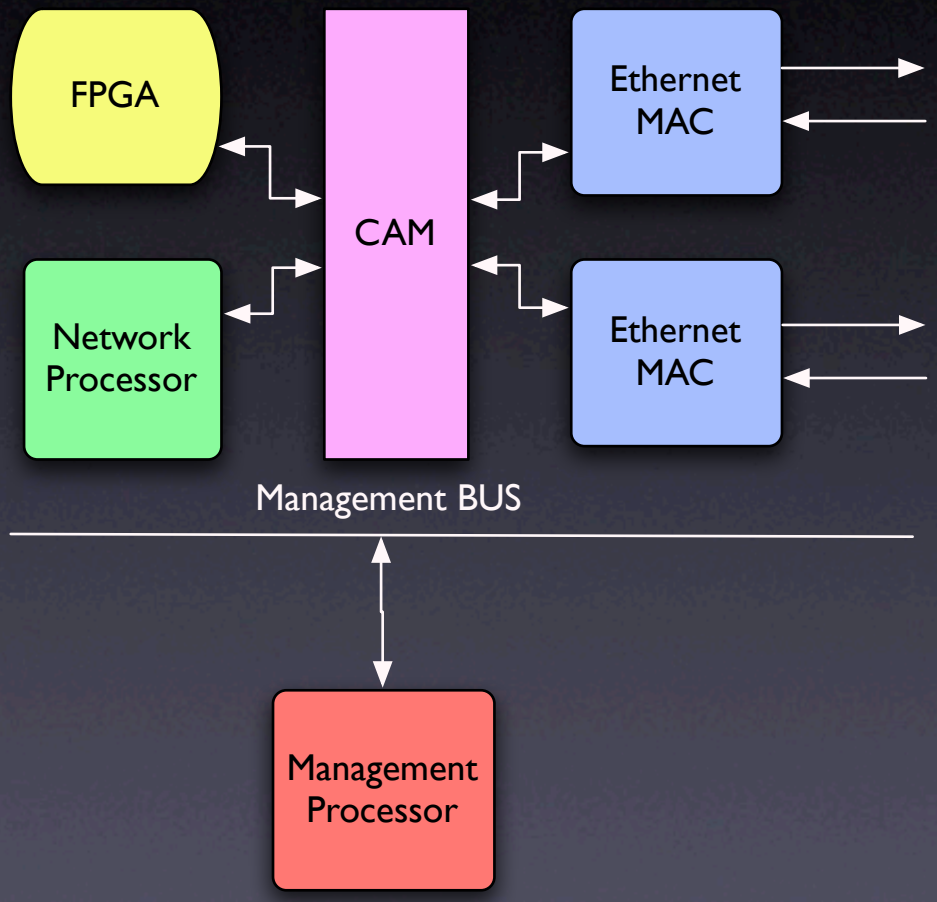

- Questions don't apply
- Programmability is based on the Vendor
- It's a fix field pattern parser
- State, State and more State
- Much stronger on bugs
- Really bad on memory
- Use it's abuse of memory to your advantage

Management Processor

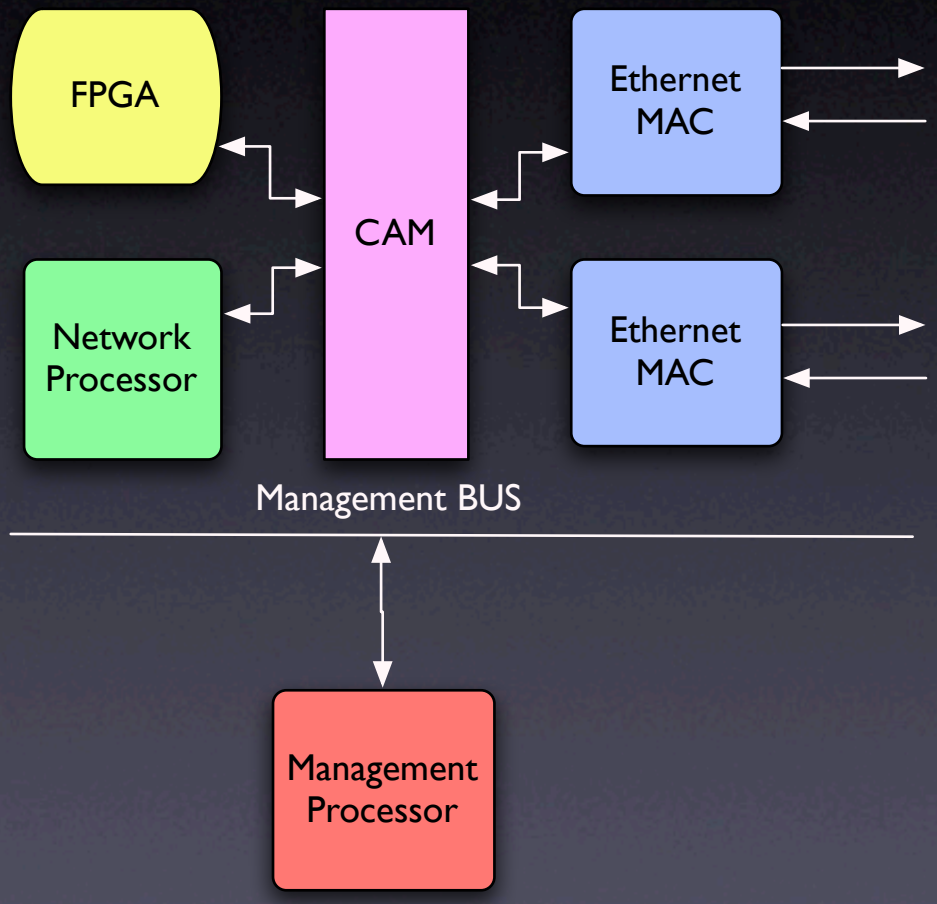

- Just your average, ordinary chip
- If you cause the management interface to be busy, do packets slow down?
- Really bad on memory
- Use it's abuse of memory to your advantage

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

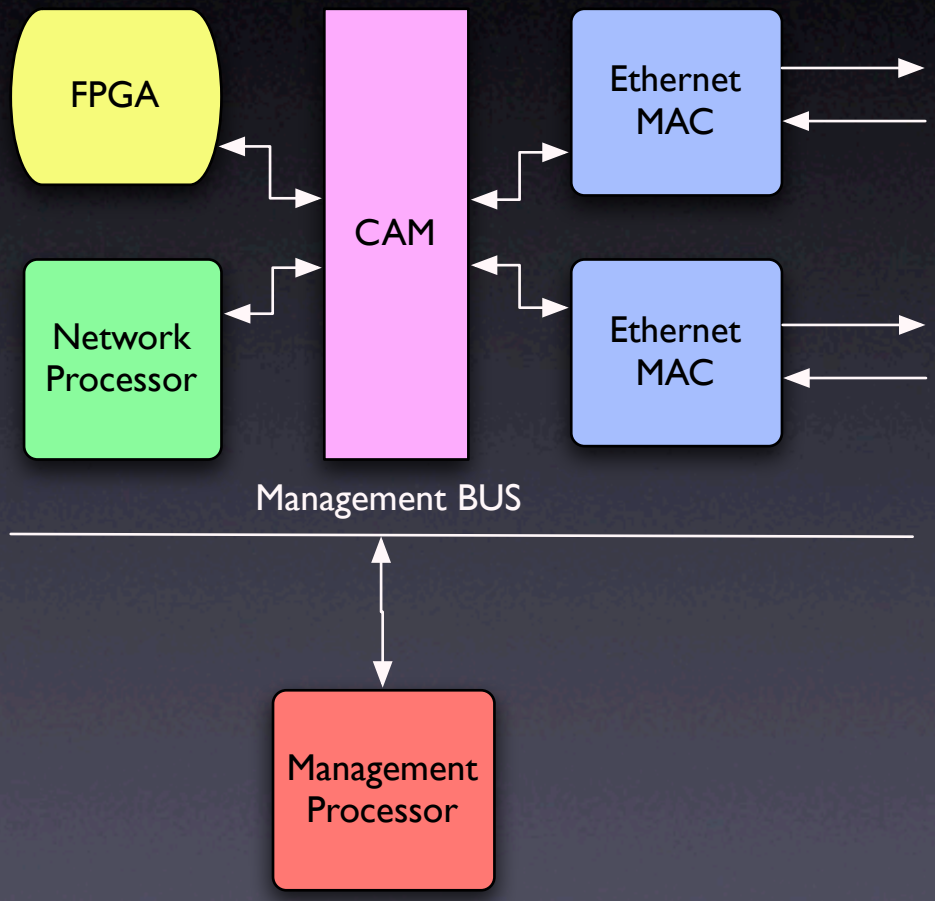

- Multiple BUSes sometimes
- If they are interconnected doesn't matter still weakest link the chain
- Some buses can't handle interleaved packets
- Could you force interleaving of packets?
- Buses use wimpy identifiers - can you modify that identifier?
- A bus has two elements: Max Performance, Max #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

A man's got to know his limitations.
Dirty Harry

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 requires 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
- The above can only handle half-duplex, let alone full-duplex
- 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

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 due to slow access of the buffers
- 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...

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!

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 ...

Example - ISS

- 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

- They use a PCI Bus on that Dell Platform
  - Bus limited to 528 Mbits/s full duplex (472 due to overhead)
- 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

- 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

- (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

Juniper Inspection

- Never saw one before up close
- Got it on eBay IDP-50 (new!)
- 1U PC [Pentinum 4 2.8 Ghz] [ATI RAGE]
- Linux Kernel 2.4.31
- Using Intel e1000 cards [w/ Silicom Bypass]


Juniper
IDP 50

Juniper Inspection.2

- IDP 10, 50, 100, 200, 500, 600C/600F, 1000, 1100C, 1100CF ... same box?
- e1000 cards set to 4096 descriptors
- Went from 3.1 to latest release (drivers changed multiple times)
- They are secure enough in their manhood - root


CD-ROM CABLE

Juniper Inspection.3

- Requires management server - loaded it on the box itself
- scio and sctop are your friends
  - /usr/idp/device/...
  - scio - IO control (set/get all sorts of cmds)
  - sctop - Monitor everything
- [/usr/idp/device/bin] - attach and watch
- Box was too loud - what can I do about that?

Virtualize

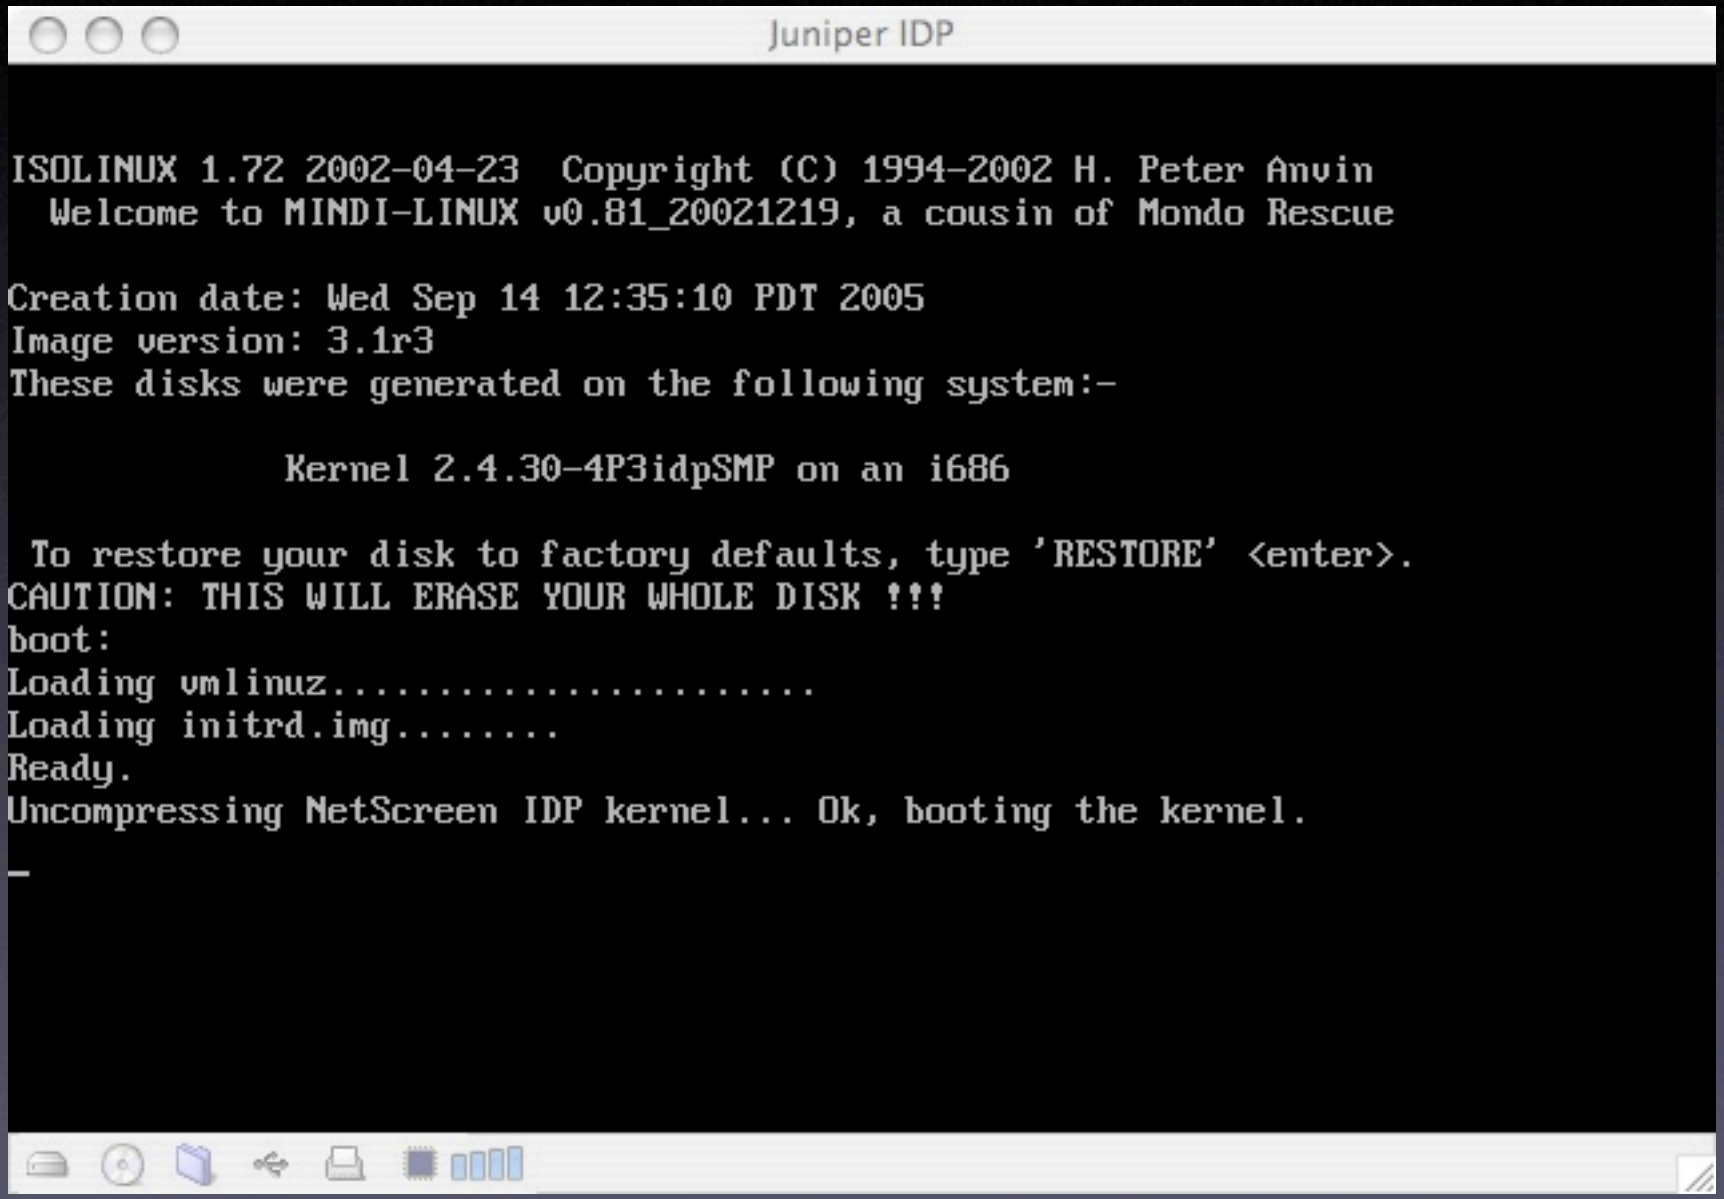
Juniper IDP
ISOLINUX 1.72 2002-04-23 Copyright (C) 1994-2002 H. Peter Anvin
  Welcome to MINDI-LINUX v0.81_20021219, a cousin of Mondo Rescue
Creation date: Wed Sep 14 12:35:10 PDT 2005
Image version: 3.1r3
These disks were generated on the following system:-
Kernel 2.4.30-4P3idpSMP on an i686
To restore your disk to factory defaults, type 'RESTORE' <enter>.
CAUTION: THIS WILL ERASE YOUR WHOLE DISK !!!
boot:
Loading vmlinuz.....
Loading initrd.img.....
Ready.
Uncompressing NetScreen IDP kernel... Ok, booting the kernel.

False Positives

- Mythical to me
- Two out of the box
  - IP: Microsoft IGMPv3 DOS (uh?)
  - SSH: PuTTY SSH2 MSG_DEBUG Overflow (dropped!!!)

Juniper Sig Dive

- The signature constructs [context + regex]
- IGMP
  - Packet with IP Options
- SSH
  - Client to Server using SSH looking for \(SSH.2 \0.PUTTY\].*
  - Then Server to Client in a packet looking for ([^ \00]..|.[^\00].|..[\040-\0377])..\x04\x.[\0200-\0377].*

Juniper Thoughts

- Now I'm curious - if it's all regex...
- RECURSION - let's see how to handles it
- [dig dig dig dig dig dig dig]
- hmmm - what's this?

./scio counter get flow

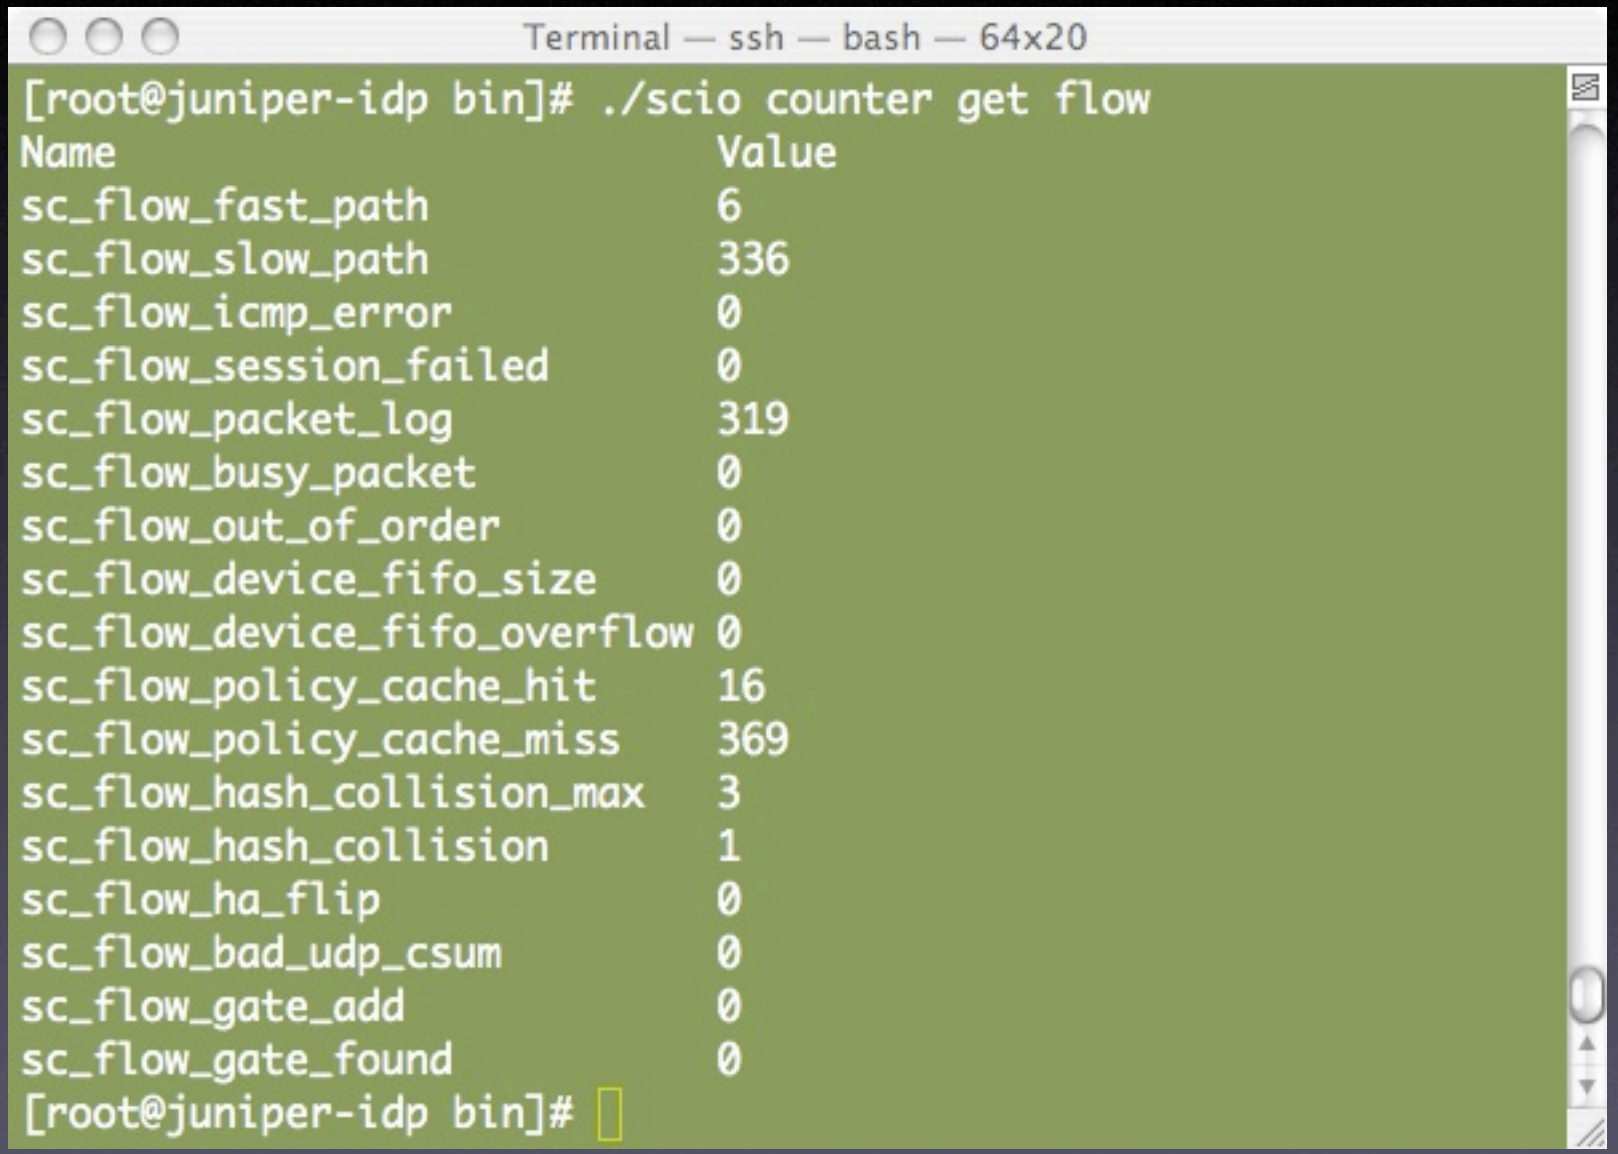

./scio const list

```
[root@juniper-idp bin]# ./scio const list
  sc_debug_features              = 0x10        [ 0...ffffffff ]
  sc_debug_qmodules              = 0x0         [ 0...ffffffff ]
  sc_debug_services              = 0x0         [ 0...ffffffff ]
  sc_debug_services2             = 0x0         [ 0...ffffffff ]
  sc_debug_level                 = 0x1         [ 0...3 ]
  sc_debug_detail                = 0x0         [ 0...1 ]
  sc_malloc_debug                = 0x0         [ 0...1 ]
  sc_malloc_debug_size           = 0x200       [ 0...fc17 ]
  sc_log_cache_size              = 0x3200      [ 1...ffff ]
  sc_log_chunk_size              = 0x4000      [ 400...4000 ]
  sc_log_chunk_timeout           = 0x186a0     [ 1...f4240 ]
  sc_pktlog_cache_size           = 0x100000    [ 400...ffffffff ]
  sc_pktlog_chunk_size           = 0x1f82e     [ 400...ffffffff ]
  sc_pktlog_chunk_timeout        = 0x186a0     [ 1...f4240 ]
  sc_sam_cache_size              = 0x80        [ 1...ffff ]
  sc_flow_hash_table_size        = 0x186a0     [ 400...f4240 ]
  sc_memory_limit_percent        = 0x3c        [ a...5a ]
  sc_tsig_hash_table_size        = 0x10000     [ 100...100000 ]
  sc_policy_lookup_cache         = 0x1         [ 0...1 ]
  sc_enable_packet_pool          = 0x1         [ 0...1 ]
  sc_enable_all_qmodules         = 0x1         [ 0...1 ]
  sc_enable_ha_lb                = 0x0         [ 0...1 ]
  sc_ha_lb_sniff                 = 0x0         [ 0...1 ]
  sc_mgt_svr_ui_port             = 0x1c23      [ 1...fc17 ]
  sc_ha_heartbeat_port           = 0x1581      [ 1...fc17 ]
  sc_enable_bypass_unit          = 0x0         [ 0...1 ]
  sc_enable_layer2_bypass        = 0x0         [ 0...1 ]
  sc_enable_udp_csum             = 0x0         [ 0...1 ]
  sc_dump_szblocks               = 0x0         [ 0...ffffffff ]
  sc_dump_szblocks_times         = 0x0         [ 0...ffffffff ]
  sc_log_enable_thresholding     = 0x1         [ 0...1 ]
  sc_log_threshold_use_dst       = 0x0         [ 0...1 ]
  sc_log_first_n                 = 0x1         [ 1...80 ]
  sc_log_threshold_count         = 0x4000      [ 100...10000 ]
  sc_log_threshold_timeout       = 0xa         [ 1...3c ]
  sc_dfa_run_merged              = 0x1         [ 0...1 ]
  sc_pcre_recursion_limit        = 0x7         [ 1...20 ]
  sc_ids_process_ignore_s2c      = 0x0         [ 0...1 ]
  sc_log_implicit_pkt_drop       = 0x0         [ 0...1 ]
  sc_reass_ha_sync               = 0x0         [ 0...1 ]
```

How did it handle strikes?

- Backdoors [0 out of 4]
- Network Worms [3 out of 6]
- Exploits [21 out of 155]
- Recon [5 out of 78]
- Hostile [33 out of 37]
- Denial of Service [1 out of 20]

Score

21%

Example - Juniper

- Juniper Filter
  - HTTP (".*/cvsweb\.cgi/.*;.*")
- Running on a 1.5 GHZ G4 using PCRE v6.4
- Standard run (after initial) (100 bytes)
  - Match: 66 usecs || 15,151 PPS
  - Miss: 4 usecs || 250,000 PPS

Example - Juniper 2

- Increase Data to 1500 bytes
  - Match: 179 usecs || 5,586 pps
  - Miss: 191 usecs || 5,235 pps
- Multiple Packets (15k)
  - Miss: 1452 usecs* || 688 pps

Build your own 200/600

- Buy one Super Microboard
- Install two XEON 2.8 CPU's
- Install 2 Gigabytes of memory
- Install Silicom Ethernet cards (e1000)
- ./scio const -s s0:reass set sc_tcp_max_flow_mem_kb 0x4000 [insert]
- ./scio const -s s0:reass set sc_tcp_max_packet_mem_kb 0x100000 [insert]

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

Questions?

Dennis Cox
dcox@bpointsys.com