Experience in fighting DDoS attacks

Nicolas FISCHBACH [nico@colt.net]
Senior Manager - Network Engineering/Security
SwiNOG-9

Agenda

- One year later…
- Some real examples
- Routers and network hardening
- Arbor Peakflow DoS
- Botnet detection
- Filtering attacks
- Cisco/Riverhead Guard w/ MPLS-based diversion

One year later

- Attacks
  - DDoS extortion (traditional model applied to the Internet)
  - Underground becoming commercial
  - Trojans (ab)used for SPAM, C/C checking, etc.
  - “Phising” (fake forms, spoofed URLs, client bugs, etc): e-commerce, e-banks and software editors
  - “Safe” types not safe anymore...
    - ZIP
    - PDF
    - BMP/PNG/JPEG/etc.
      … what do you tell your users ?
  - Is XP2 (XP SP2) going to change anything ?

One year later

- Worms and wormability

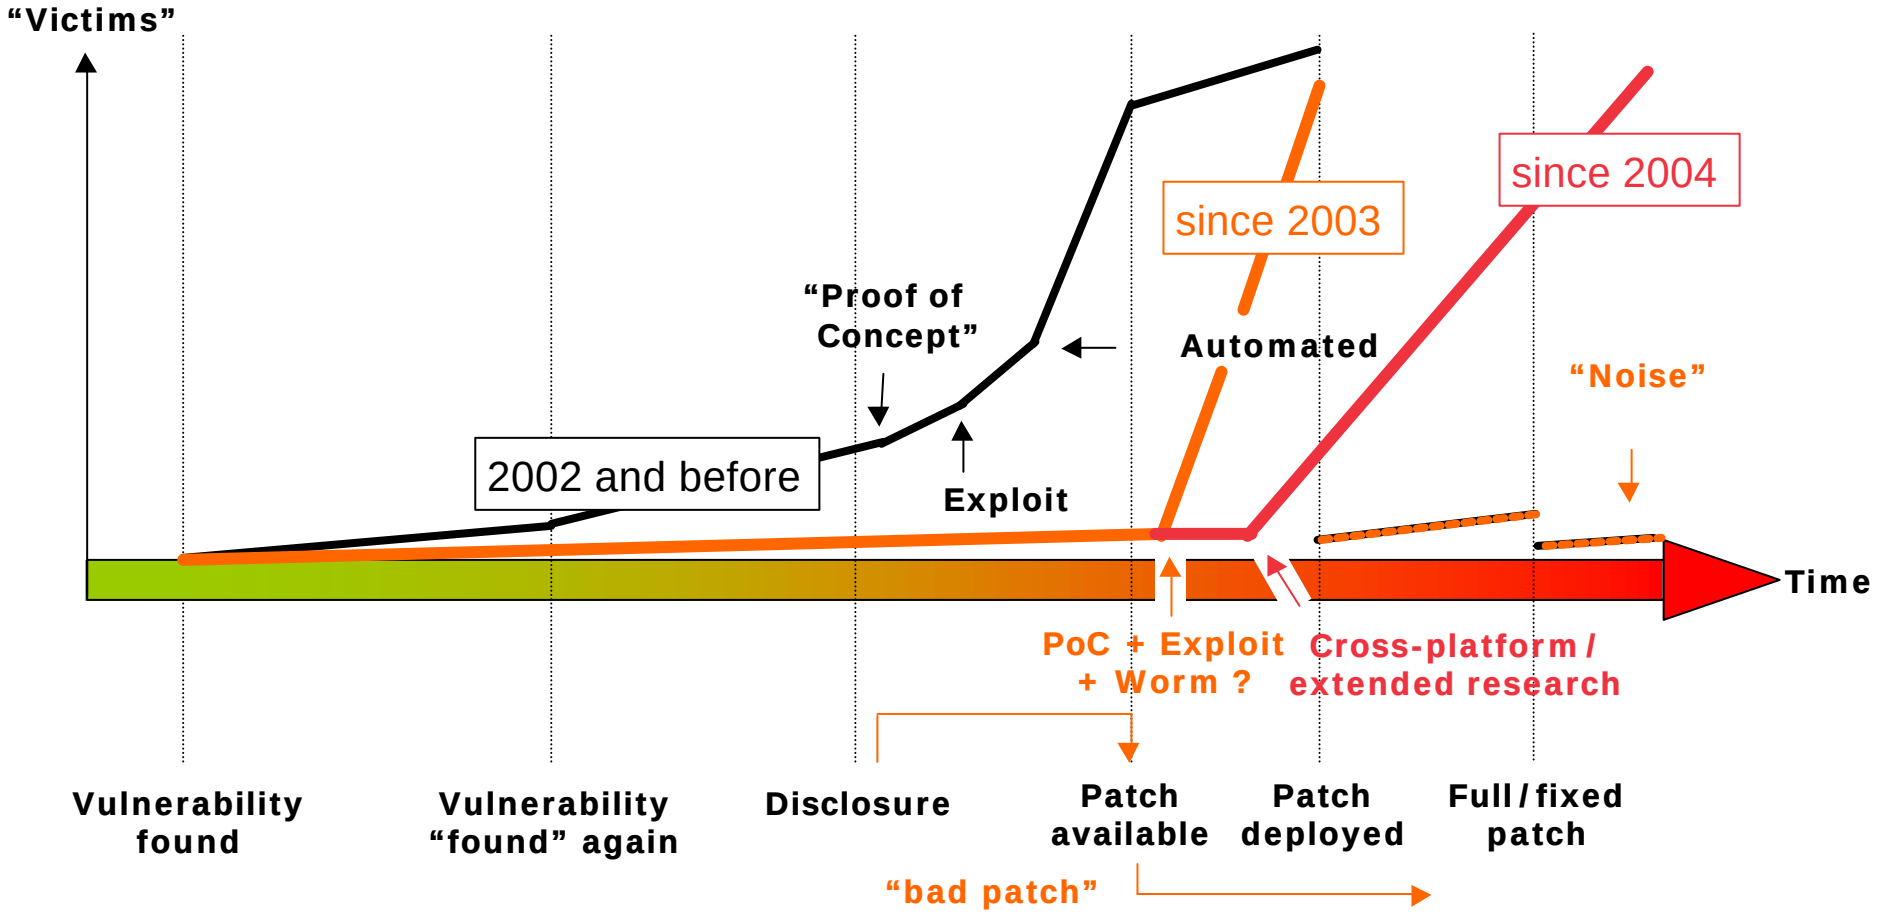

Some data

- Last ~4 months (May-Sep'04)
- 21962 anomalies detected
  - 5302 High ("displayed")
  - 15513 Medium
  - 1147 Low
- Per day:
  - ~40 anomalies make it to the "Security" screen in the NMC/NOC
  - Some are duplicates for the same attack, some are false positives, etc.
  - Overall 20 "real attacks" a day…
  - A third of them are probably "business affecting" from the customer's perspective

Some examples

- Aug/13: Eat this...

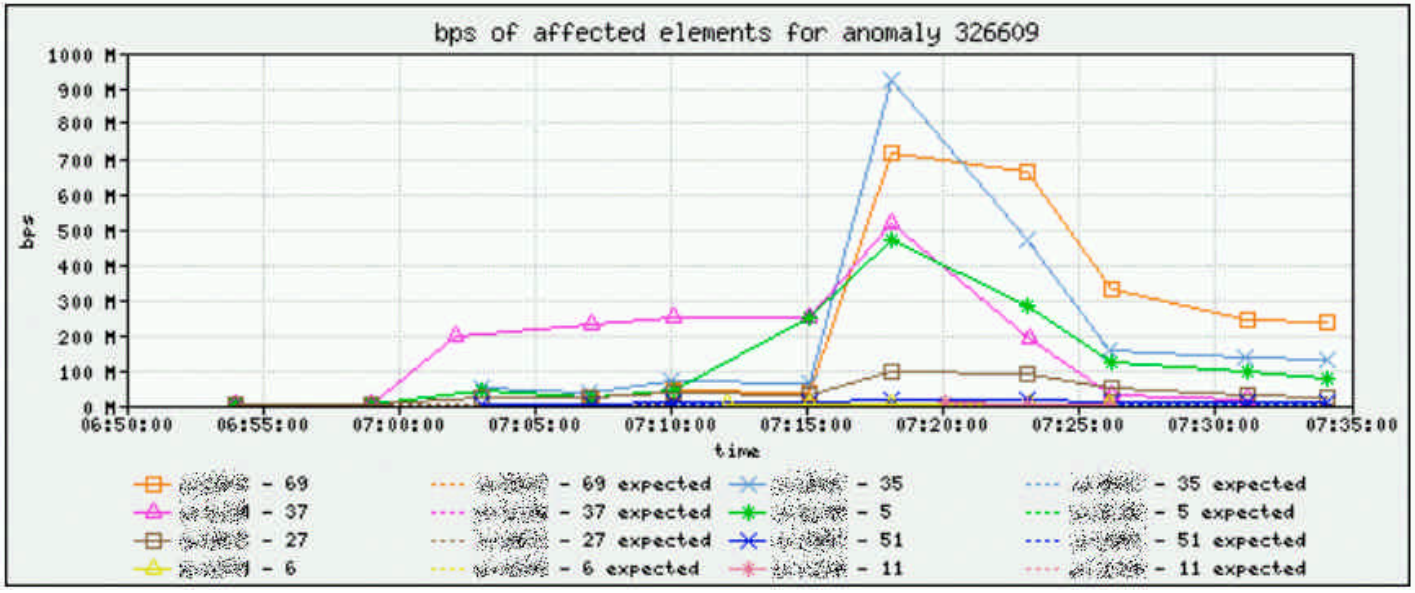

- Nearly 3Gb/s of traffic
- Mainly from America/AP
- 0/UDP (to filter or not to filter “by default” ?)

Some examples

- Sep/28: Eat this...
- 200kPPS of TCP SYNs
- Really distributed…
- 80/TCP

- On Sep/6:
- Nearly 1MPPS

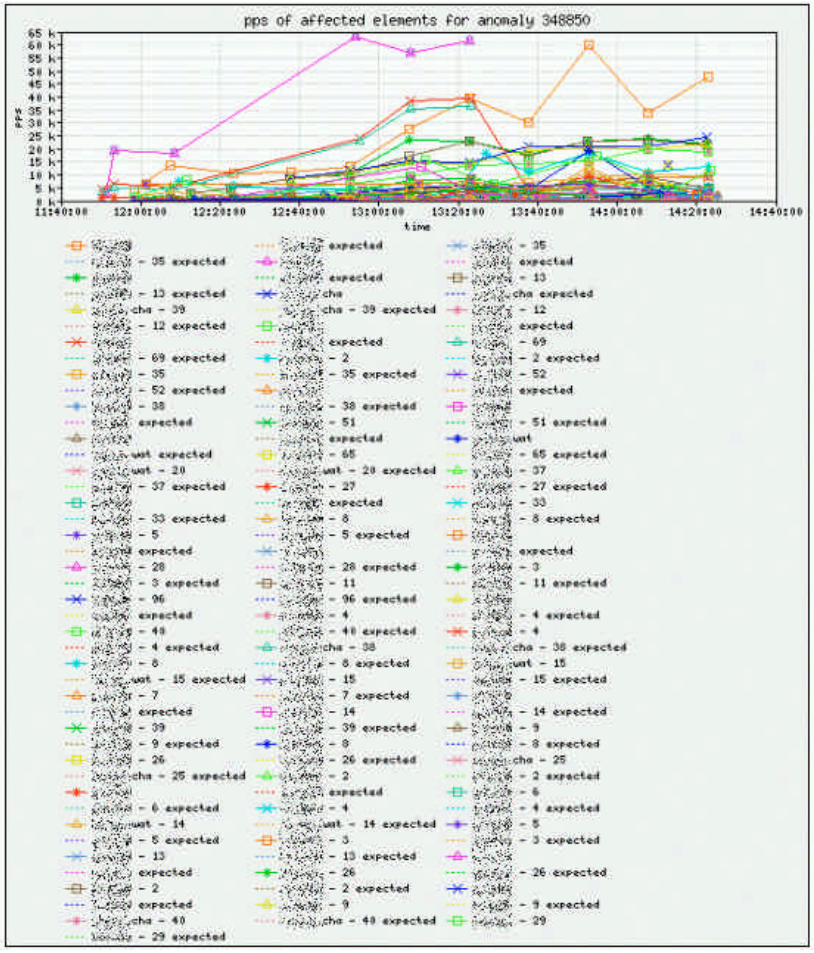

Some examples

- Jun/18: A mix of the two

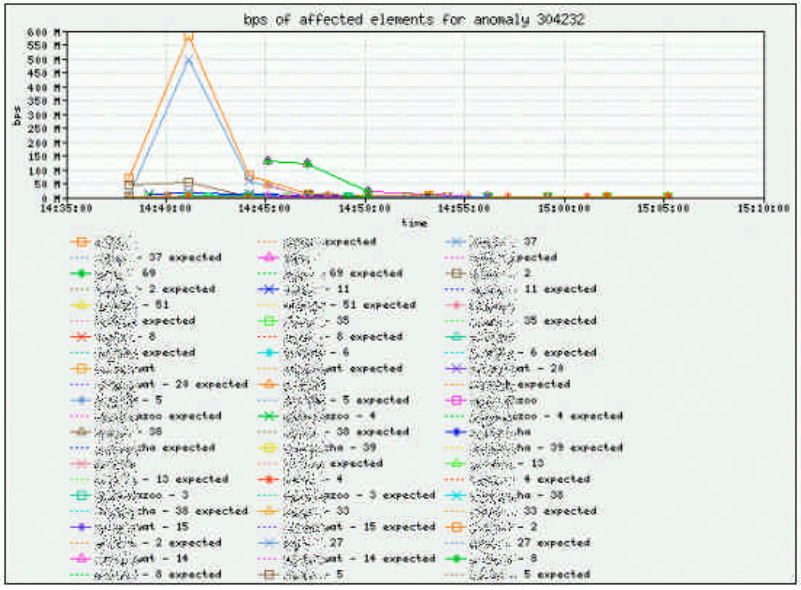

- ~1.3Gb/s of traffic
- Distributed (coming mainly from America/AP and over LINX/UK)
- 11677/UDP

Some examples

- Jul/10: Let's try over different peering points

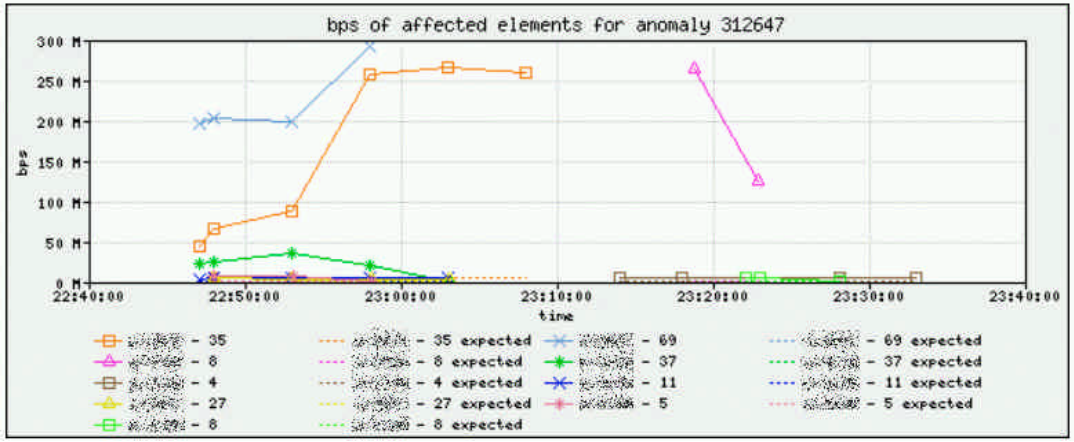

Some examples

- Jul/6: Nothing is perfect, a false positive

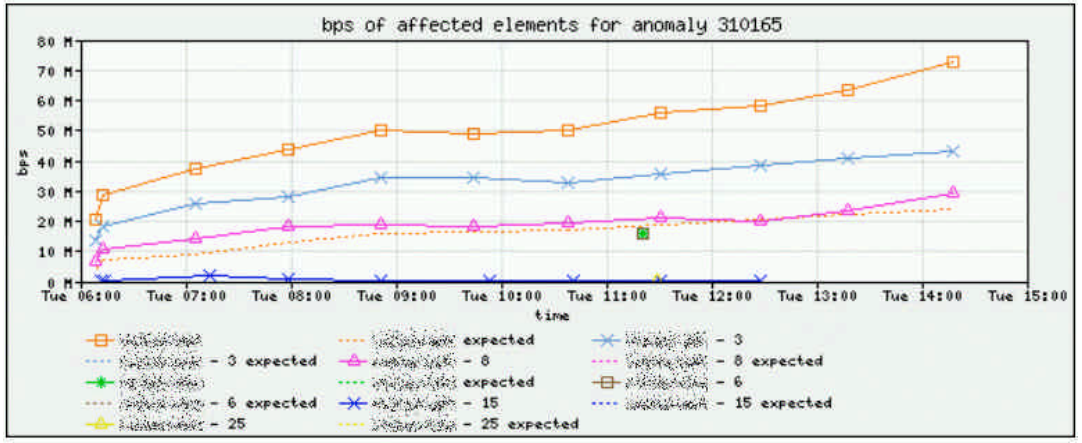

- 80/TCP

Cisco and network hardening

- Router hardening
  - rACLs
  - GSR:
    - limiting tofab and frfab queues proved effective
    - 3GE cards can’t do Netflow+QoS+ACLs...
  - 75xx:
    - Buffer carving calculation not easy to understand
    - Like the GSR, requires queue limit to prevent backpressure and memory consumption on the VIP
    - fair-queue isn’t enough: WRED for =< E1
  - 76xx:
    - Difficult to understand because of different HW modules and different SW/configuration depending on HW: LAN, OSM, Flex-Wan

Cisco and network hardening

- Network hardening
  - Router level
    - iACLs
    - tACLs (edge, access)
    - uRPF
    - QoS
  - See http://www.securite.org/presentations/secip/
  - Systems
    - Sinkholes
    - Backscatter
    - Honeypots
    - BGP UPDATEs accounting

Arbor Peakflow DoS

- Netflow-based detection solution
  - Only solution that “fits” for large SPs
  - Lots of (security) vendors adding Netflow support in their products
- Pros/cons/experience
  - Netflow vs payload
  - Distributed vs inline/data center
  - As complex to configure and fine tune as an NIDS !
  - NMS integration (SNMP trap “high alarms” only)
  - Scalability
    - From Edge to Access (number of sources)
    - Netflow re-use (billing and traffic engineering)
  - Sharing data: DoS your DDoS detection system ?

Arbor Peakflow DoS

- Netflow
  - GSR: sampled Netflow (1/100)
  - Issues with
    - 4xSTM-4: 12.0.23(S)
    - 3xGE: can't fit all features (ACLs, QoS and Netflow)
    - DPT: mix of sampled and unsampled (TAC)
  - 72xx/75xx: random sampled Netflow in 12.0.26(S)
  - MPLS-aware Netflow and tools ?

- Detecting outgoing attacks
  - ISE Output sampled Netflow: 12.0.24(S)
  - Edge/Peering to Core interface (MPLS PHP - pure IP)

Botnets detection

- Zombies C&C

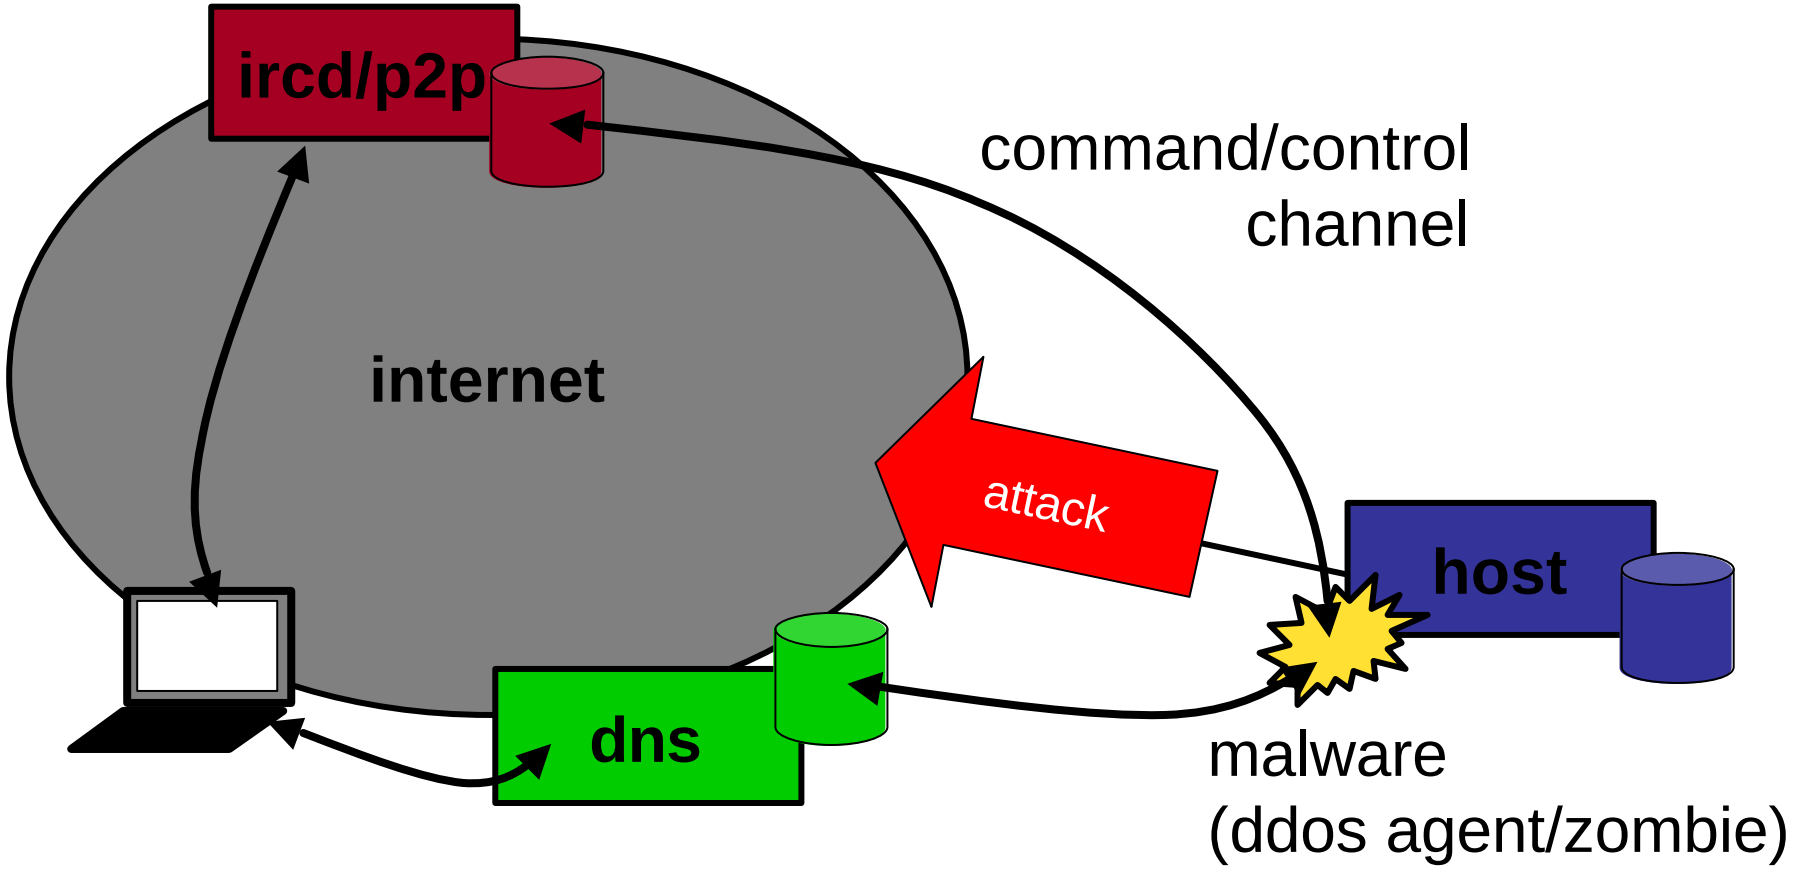

Filtering an attack

- How to filter ?
  - Old school: blackhole the destination and call your upstreams' NOC
  - Mess up your network: deploy ACLs
  - Divert and clean...
    - Configure the "victim"
    - Reroute
    - Clean
    - Re-inject
  - … and use RTBH (remote triggered blackhole) if some of your key links look like they won't survive the attack
    - in specific regions/with specific peerings
    - working with the peer/upstream (if they don't support RTBH)

An example of “better” filtering

- Sep/6: Filtering a 80/TCP SYN flood…

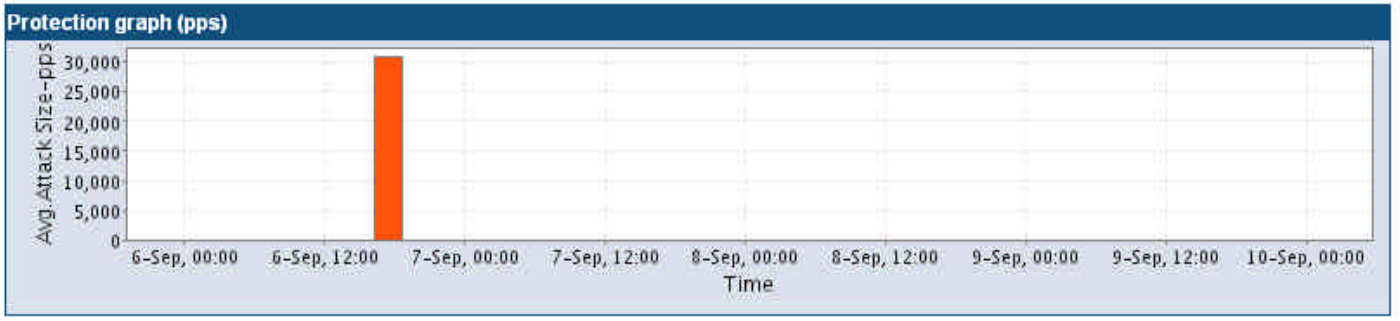

Total Attack Statistics

Attacks Mitigated	Attacks Duration	Max. Traffic Rate	Total Rx	Total Blocked	Legitimate vs Malicious Traffic
54	46:39:29	91,476.92	265,995,609	255,654,832	3.89% / 96.11%

- … while keeping the site up&running :-)

Cisco/Riverhead Guard

- Traffic filtering
  - No baseline needed
  - Baseline helps to build a filter template (“ACLs”-like)
  - No permanent rerouting through the Guard needed
  - Regional deployment, linked via GE to the Core and BGP sessions to Edge routers
- Pros/cons/experience
  - No learning period
  - Identify the attack and adapt the filters (high/low thresholds with drop/no drop policy)
  - As complex to configure/fine tune than a NIDS (on a per attack basis) !
  - Reporting and customer “experience”
  - Performance 250kPPS -> 1M+ PPS

Cisco/Riverhead Guard

- Traffic diversion:
  - Traffic diversion using MPLS LSPs (the Guard doesn't have to speak MPLS)
    - LSP from Edge to DCR (Directly Connected Router)
    - DCR doesn't perform a routing lookup and "sends" the IP packet to the Guard (penultimate hop)
  - Issue(s) with the 76xx (PFC2)
    - Traffic enters via the OSM-4GE-WAN-GBIC
    - LC pops the label but needs additional layer 2 information (PFC2 is used for this)
    - LSP dies – routing lookup
    - Result: traffic not send to the Guard but back into the network (to the customer)
    - Fix: Sup720 (PFC3) or use a 7206VXR+NPE-G1 ;-)

Cisco/Riverhead Guard

- Traffic diversion:
  - Do you have:
    - Enough inter-city spare capacity (STM-1+/GE) ?
    - No engine 0/1 cards on the divert path ?
    - No software based systems on the path ?
    - A DDoS capacity planning team ?
  - Mitigation:
    - Strong need for regional DDoS “treatment” centers
      - IP anycast-like engineering
      - HA solution on key transit/peerings
    - Engine 3 migration programme

Thank you