Lack of Classification Ability Considered Harmful

Vijay Gill

vijaygill9@aol.com

NANOG 27, Phoenix, AZ

February 2003

Vendors Please Pay Attention

- Security

Security

- Routers are optimized for traffic **through** the hardware
  - Not traffic **for** the hardware
- Designing a cost efficient router implies:
  - Cross-sectional bandwidth capacity dominates budget
  - No cost-effective way to engineer a router that can absorb and usefully process data at the rate it can arrive

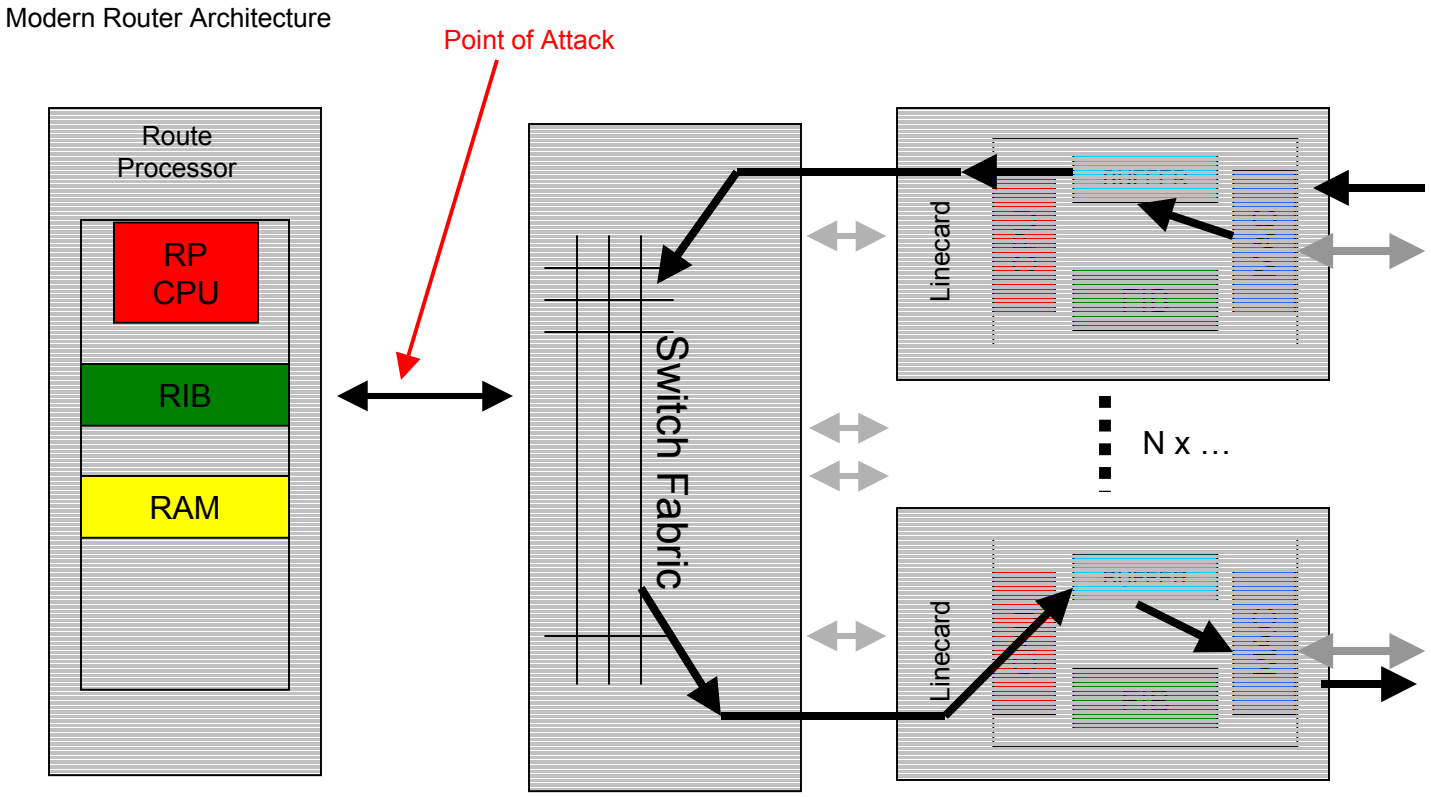

Hardware – Queuing of Control Plane Traffic

- This one should be easy to get but surprisingly few can do it
- Simple, unambiguous parsing
  - Filter on stuff that is for the router
    - What I deem interesting goes onto the high priority queue
    - Everything else goes onto the low priority queue
- Simple discriminator function/ACL etc.
- Rate-limit on low priority queues
- Apply discriminator on linecard/forwarding engines BEFORE it hits the brain
- Why?

Outside Context Problem

- Attackers are seizing this weak link as a point of attack
  - DoS attacks targeted at infrastructure are increasing
  - Hackers will evolve – Have seen port 179 attacks already (and MSDP can't be far behind)
- Problem
  - Need some way to disambiguate between invalid and valid control traffic (e.g. BGP updates)
  - Rate-limiting on control traffic is not sufficient
    - Enough false data will swamp legitimate data
      - Connection flaps/resets
    - Need to focus on BGP (MSDP)– other traffic is not control, thus will not cause control plane issues

Security

- IGP traffic can be safely blocked
- MD5 on neighbors will not prevent the Router CPU from being inundated with packets that must be processed
- Solution
  - Short term - Dynamic Filtering on the line cards
  - Long term – outboard processing of SHA1/HMAC-MD5
  - This is very long term indeed – not necessarily solving a known problem today (replay or wire sniffing)
  - Vendors have to implement priority queuing for control traffic from line cards to control plane

Dynamic Filtering

- Filtering on the 4-tuple
  - Use the BGP 4-tuple to dynamically build a filter that is executed on the line card or packet forwarding engine
  - Packets destined for the router are matched against the filter
    - If the packet matches the filter
      - Place into the high priority queue
    - Else
      - Place into the low priority queue

Analysis

- On average, will need to try 32000 times to find correct 4-tuple
  - Attacker resources will need to be on average 32000 times greater to adversely affect a router
  - Cost of attacking infrastructure has risen
  - Cost to defender minor
    - Each configured BGP session already has all the state needed above to populate the filter
    - Can use the same solution to protect against MSDP spoofing
- Implementation (sort of)
  - In JunOS (apply-path)

Thoughts

- Stability is most important
  - Only place the high priority queue filter for a neighbor once the session is established
    - Before session is established, place neighbor packets in low priority queue
  - We'll take time for a session to come up over knocking existing sessions down

Thoughts

- Future Goals
  - Use BGP over SSL/TLS (will prevent replay attacks)
    - Can use the filter list along with SSL/TLS to reduce number of valid packets making it to the RP CPU to a comfortable number
- Vendor Feedback
  - Please ensure that your TCP/IP stack chooses randomly when picking a source port (currently most do not)

Analysis

- Any valid BGP packet arriving on any line card will have the right 4-tuple, and should be placed into the high priority queue
- Most spoofed DoS BGP packets will not match the filter and will be placed into the low priority queue
- Route Processor CPU services the high priority queue first
  - Mitigates packet flooding

The BGP TTL Security Hack (BTSH)

- BGP TTL Hack
  - Uses TTL as input into the discriminator
  - http://ietfreport.isoc.org/ids/draft-gill-btsh-01.txt
  - Set TTL to 255
    - Most BGP sessions are between direct neighbors
      - Only allow BGP packets with TTL in 254-255 range
      - Reduces attack diameter dramatically