


Presentation by ANML  
May 2003



pervasive **technology** labs

AT INDIANA UNIVERSITY

Automatic DDoS Mitigation

(Where's the Red Button?)

# Automatic DDoS Mitigation

---

- Motivation  
(Why it would be nice)
- Technical and Political Issues  
(Why we don't have it)
- Potential Solutions  
(What might work)



# Motivation

---

- The tools we've examined involve *analysis* and *reporting* but not *automatic control*
- These tools allow engineers to detect and locate an ongoing DDoS
- They do not offer any automatic protection from the effects of a DDoS attack
- We have no “fire and forget” system



# Why Not Automatic?

---


- Some of the tools we've looked at even output filtering code for major routers
- Why not apply these filters automatically?
- Wouldn't this provide excellent response time?



# Why Network Engineers Say “No”

---

- Engineers are generally very conservative within their field
- *Proven* carries a lot more weight than *promising*
- Their mission is to keep the network running, and they don't like to experiment with it



# Once Bitten, Twice Shy

---

- A large proportion of network downtime is related to bugs in routing software
- Because of the complexity of modern routers and economic factors, engineers are often given immature code
- Result: network engineers do not trust the work of software engineers



# Is This Reasonable?

---

- Estimates of Code Quality (in bugs per thousand lines of code):
  - Standard SW: 25 bugs / kLOC
  - Good SW: 2 bugs / kLOC
  - Critical SW: 0.1 bugs / kLOC  
(think Space Shuttle)



# Is This Reasonable?

---

- A modern router OS may have 5,000,000 lines of code
- Estimated number of bugs:
  - Standard SW: 125,000 bugs
  - Good SW: 10,000 bugs
  - Critical SW: 500 bugs
- Engineers' fear is not just superstition



# Issues to Bear in Mind

---

- Any automatic system must prove itself before being accepted
- Resistance will be strong if there are large demands on engineer time
- False positives undermine system credibility



# Administrative Domains



---

- Remember that the Internet is a collection of many interconnected networks
- For routing, the Internet is broken up into *autonomous systems* (ASes)
  - Each AS has responsibility for some portion of the Internet address space
  - This includes routing configuration, security, and preventative techniques



# Hierarchy in Routing

- The AS system is not hierarchical.
- While there are “Tier 1” and “Tier 2” ISPs, traffic between host A and host B generally does not travel like this:



# Why Doesn't It Look Like That?

---

- There are lateral connections between ASes at each “level” of this path
- This is especially true of Tier 1 ISPs
- Traffic is likely to pass through direct competitors when going from A to B



# Example Routing Paths

---

1. Bowman Capital
2. Verio
3. SprintLink
4. GHS, Inc.
5. RUNNET
6. Trident
7. RUSNET

(from the Route Views Project at U. Oregon)



# Example Routing Paths

---

1. Yuma Proving Ground
2. AT&T
3. DREN
4. UUNet
5. DISAUN

(from the Route Views Project at U. Oregon)



# Example Routing Paths

---

1. Hewlett-Packard
2. ERX-SingNet
3. Singapore Telecom
4. Genuity
5. SprintLink
6. Sprint
7. Sprint-CA

(from the Route Views Project at U. Oregon)



# Features of Routing Paths

---

- Between two edge networks, there are often multiple ASes of various tiers
- These peering connections are critical to the Internet and will continue to be around, even with IPv6
- The forward and reverse paths may transit different sets of ASes



# What Does This Mean for DDoS?

---

- DDoS mitigation requires at least the involvement of the AS closest to the target
- If traffic is spoofed, finding the source of the attack requires the involvement of every AS in the path
- For distributed attacks, there may be multiple AS paths for the attack traffic



# So What is “Involvement”?

---

“Involving” the next AS in the chain means gathering sufficient proof; calling the NOC for that AS (assuming you can find accurate contact information); hoping the person on the phone speaks English and has authority to act; explaining the problem; and hoping that they believe you, that they have the expertise to do something, and that they’ll actually do it.



# Consequences

---

- For locating the attack source, this must happen between every pair of ASes
- This “human engineering” time can easily become the largest component in incident response time
- Also remember that DDoS attacks rarely consider or care about international boundaries



# Another Problem: Spoofed IPs

---

- Anti-spoofing filters are now common
- Edge networks should not allow packets onto the Internet with addresses not in their local address space (*egress* filtering)
- This is not difficult to accomplish
- But we *still* see randomly spoofed IP addresses in current attacks



# Why Does That Matter?

---

- Tracing the source of the attack is much more difficult with fully spoofed addresses
- Each AS can really just hope to find the next AS in the chain, not the whole path
- In practice, a complete trace is unlikely to happen



# Ingress Filtering

---

- Use routing information to prevent spoofed address from *entering* your network
- This is possible for transit ASes
- This saves you from having to trust that everyone is doing proper egress filtering



# Does This Fix Things?

---

- The problem is a lot smaller
- Attackers can still spoof within the local address space
- Identifying the source AS is fairly straightforward
- Within that AS, large switched campus networks can make finding the exact source difficult



# What do ISPs do about DDoS?

---

- The standard practice is to combine a “blackhole route” and a DNS name change:
  - Customer changes address of server under attack.
  - ISP updates DNS with the new address.
  - ISP discards all traffic to the old address.



# Problems with this Approach

---

- This doesn't remove all effects of the attack; some systems cache DNS lookups
- If attackers begin to target hostnames rather than addresses, the technique won't work
- The DNS entries for the server may be an integral part of load-sharing



# Automating DDoS Defenses

---

- Suppose that ISPs can trust all routes received from peers
  - (This is a big “suppose”)
- Suppose also that a QoS infrastructure exists to give preference to BGP traffic
  - (This is less of a “suppose”)



# Automating DDoS Defenses

---

- We could use BGP to propagate blackhole routes across AS boundaries
- Pushing the blackhole route toward the traffic source reduces network impact
- Single-homed customers could either run minimal BGP or still contact their ISP directly



# Additional Benefits

---

- When routers send traffic to a blackhole, this can be reported to a management system
- The management system can use this data as an indication of a ingress point for attack traffic
- This allows better tracing of the origin of a DDoS attack



# Why Can't This Be Done Now?

---

- If ISPs blindly accept blackhole routes from peers, the routing infrastructure has a DDoS vulnerability
- An attacker could introduce a blackhole route for an address over which they have no authority
- The network would also be vulnerable to accidental misconfiguration



# What Needs to Happen?

---

1. All Tier-1 ISPs must adopt an infrastructure for per-prefix routing policy.
2. Standards organizations (NANOG/IETF) must agree on the BGP extensions for propagating blackhole routes.
3. A shared infrastructure must be developed for backtracing analysis.



# Per-Prefix Routing Policy

---

- This involves the use of a Routing Registry.
- RRs have existed in the past, but a number of Tier-1 ISP refuse to use them.
  - UUNet and Sprint are examples
- Automatic blackhole propagation is too dangerous without RR infrastructure
- Attitudes must change for this to work



# Standards Organizations

---

- Solution must work across international boundaries
- BGP extensions introduced by a single vendor are unlikely to be accepted
- Network engineers will resist any system not adopted by a standards organization (IETF in particular)



# Shared Infrastructure

---

- Accurate backtracing requires data from ASes around the Internet
- Many ASes are direct competitors
- Some ASes cross international borders
- Some pairs of ASes are in mutually unfriendly nations
- Only a neutral party can be given this data



# Making This Happen

---

- The greatest challenges are personal and political, not technical
- Network engineers are very hesitant to use automated systems
  - Loss of autonomy
  - Experience with buggy beta-test code
  - Lack of faith in software engineering



# Making This Happen

---

- Any successful system must satisfy these fears
  - Cannot threaten network infrastructure
  - Cannot compromise information felt to be proprietary
  - Every false positive or mistaken blackhole is itself a DoS attack



# Open Issues

---

- When and how do blackhole routes go away?
  - The same infrastructure that propagates them needs to be able to recall them as well
  - At the end of the AS path, we can't tell if the attack is continuing
  - We must rely on either time or the judgment of the AS closer to the attacker



# Open Issues

---

- What if attackers focus on names rather than addresses?
  - This makes the blackhole routing defense ineffective
  - There is *no technical barrier* to this sort of attack appearing on the Internet tomorrow
  - If it can happen, it will happen

