



Internet Architecture WG:
**DoS-resistant Internet
Subgroup Report**

Mark Handley
University College London



DoS-Resistant Internet Working Group

Initial meeting held January 27th in London:

The objective of the initial meeting is to share experience and concerns, explore what the Working Group can usefully do, and (we hope) kick off that work. The emphasis is on understanding the real issues and looking at near and long term approaches.

Chaired by Jon Crowcroft, Cambridge University/CII

Who was there?

- Chatham House Rule:
 - I can share the information presented and discussed, but I can't attribute it to anyone or tell you who was there.

What sort of people were there?

- About 50 people:
 - Ten ISPs: several large, several medium size (somewhat UK centric), major internet exchange point, cellular operator.
 - Several "victims": online gambling, major bank.
 - Major network equipment vendors.
 - Major OS vendors (desktop and mobile).
 - Several vendors in anti-DoS space.
 - Telecoms Regulator.
 - Police.
 - Intelligence Community.
 - Academics in networking and public policy.
 - CII members.
- 16 presentations representing almost all these communities



Outline

- Summary of the Workshop
 - The Nature of the Problem
 - Current Defense Techniques
 - Future Architectures

- Next Steps for the WG



Outline

- Summary of the Workshop
 - The Nature of the Problem
 - Current Defense Techniques
 - Future Architectures

- Next Steps for the WG

Denial of Service



The Register » [Security](#) » [Network Security](#) »

US credit card firm fights DDoS attack

By [John Leyden](#)

Published Thursday 23rd September 2004 11:13 GMT

US credit card processing firm [Authorize.Net](#) is fighting a sustained distributed denial of service (DDoS) attack that has left it struggling to stay online.

In a statement to users posted yesterday, Authorize.Net said it "continues to experience intermittent distributed denial of service (DDoS) attacks. Our system engineers have successfully minimised the impact of each attack and have quickly restored services to affected merchants. Industry experts are onsite and working with Authorize.Net to expedite a resolution. Please be aware that the stability and reliability of the Authorize.Net platform remains our top priority; and we are doing everything we can to restore and maintain secure transaction processing despite these unforeseen attacks."

Denial of Service

- The Internet does a great job of transmitting packets to a destination.
 - Even if the destination doesn't want those packets.
 - Overload servers or network links to prevent the victim doing useful work.

- Distributed Denial of Service becoming commonplace.
 - Automated scanning results in armies of compromised zombie hosts being available for coordinated attacks.

ISP's view of the problem.

ISP1 (very large ISP)

- 6-7 ongoing DoS attacks at any time.
- Peak bandwidth seen in UK: 3Gb/s
- Peak bandwidth known to be seen in US: 5Gb/s (flatlined 2 OC48 links)

ISP2 (large ISP)

- >22000 anomalies in May-Sept 2004
- 5000 high rate
- 20 real attacks per day - perhaps 1/3 seriously affect customers.

ISP's view of the problem

ISP 3: (large international ISP)

- Sees attacks from 300 to 10000+ simultaneous hosts.
- Sophisticated full spectrum attacks:
 - SYN flood
 - TCP connection flood
 - URL flood
 - UDP flood
 - ICMP flood
 - DNS attacks
 - Malformed packets
- It's not getting any better.

ISP's view of the problem

Major security vendor:

- Lack of data encourages speculation, confusion and hyperbole....
- But trends are worrying:
 - DoS attacks greater than 10Gbps aggregate.
 - Of 1127 customer-impacting DDoS attacks seen in 2004 on a large network, only 4 employed source address spoofing.
 - 80K+ node botnet largest seen this year.
 - DoS attack vectors are changing (eg application level, Ack with simulated sequence numbers)

ISP's view of the problem

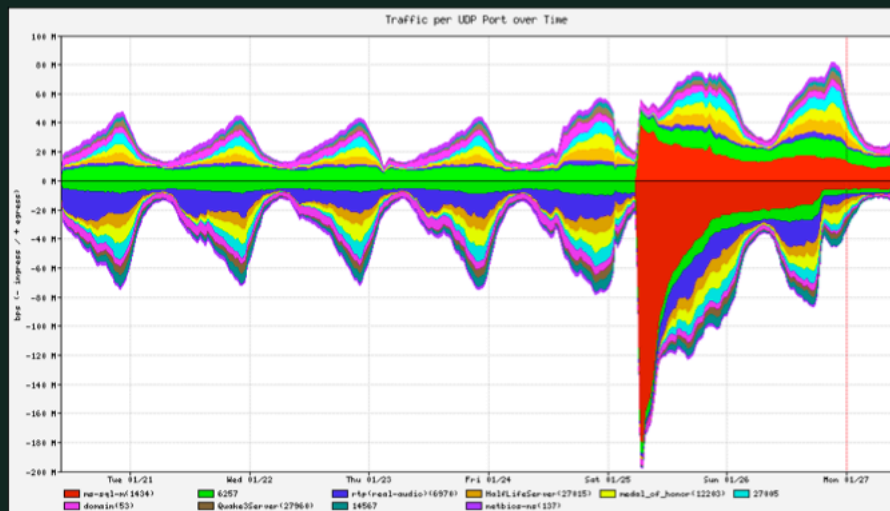
- ISP 4 (medium size national ISP)
 - **Problem, what problem?**
 - This ISP has no high-profile DoS targets.
 - Mostly home users.
 - Their backbone and peerings are over-provisioned.

 - DoS mostly only noticed when another ISP complains one of their customers is being DoSed.
 - Dealt with on a case-by-case basis.
 - Not worth them investing in a detection infrastructure.

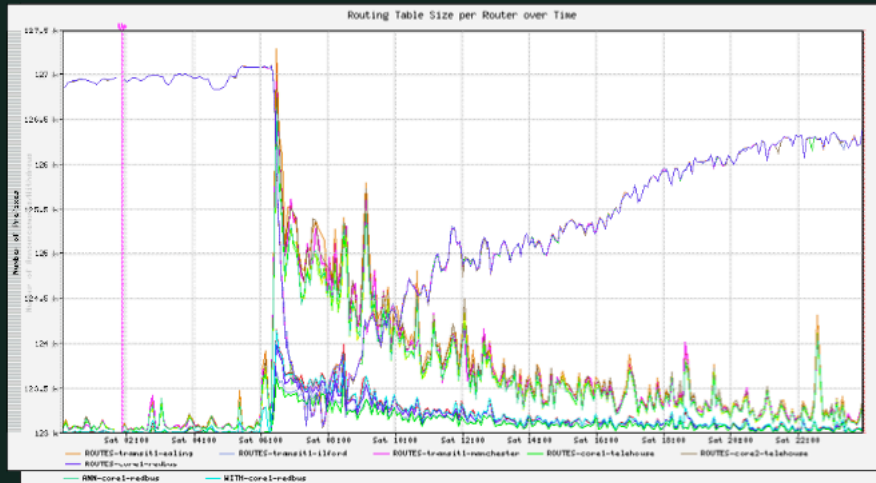
The nature of the attacks

- Pretty broad range.
 - Wide range of attacks on end-hosts (CPU, memory exhaustion)
 - Attacks on edge routers (bandwidth exhaustion, forwarding power, CPU cycles)
 - Very little source address spoofing.
- Range of possible attacks is much broader, but the simple attacks mostly work well enough.

SLAMMER – THE UDP TRAFFIC PICTURE



SLAMMER – THE BGP PICTURE



Motivation of the attackers *today*

- MEECES (Max Kilger, Honeynet)
 - Money
 - Ego
 - Entertainment
 - Cause
 - Entrance into Social Groups
 - Status



Profile of attackers today

- Asia-Pacific and South America are main sources.
 - Not just Eastern Europe and Russia anymore.
 - Mostly poor countries, where a few hundred/thousand dollars is a year's salary.
 - Usually good education, but in a country with high unemployment.
- Groups communicate mostly in-band (Internet).
 - But most ISPs don't have the resources to analyze TBs/day of IRC logs in many languages.
- Many groups are well organized and highly skilled.
 - Mostly not for fun on free time anymore.



Potential Perpetrators

- "Traditional" hackers
- Script kiddies
- Spammers
- Organized crime
- Terror Groups
- Hostile States

Significance of New Classes of Perpetrator

Additional skills and resources

- Better planning and testing
- Better planning and software engineering

Capability to Combine Attacks

- To assist the electronic attack
 - Eg. infiltration, corruption of insiders
- To amplify the electronic attack
 - Simultaneous physical attack

Different target selection

Bots and Botnets

Bot

- application that performs some action on behalf of a remote controller
- installed on a victim machine (zombie)
- modular (plug in your own functionality/exploit/payload)

Botnets

- Linkage of "Owned" machines into centrally controlled armies
- literally roBOT NETworks

Control channel

- Method for communicating with an army

Herder

- Owns control channel, commands botnet army

Botnets

- Mass acquisition tools used for initial compromise.
 - Losing a botnet isn't a tragedy - can quickly re-compromise new hosts.
- Variety of communication channels used to control botnets, but IRC and P2P protocols are most common.
- After compromise, protect host to prevent multiple zombies/agents on the same host.

Botnet Spammer Rental Rates

>20-30k always online SOCKs4, url is de-duped and updated every
>10 minutes. 900/weekly, Samples will be sent on request
>Monthly payments arranged at discount prices

3.6 cents per bot week

>\$350.00/weekly - \$1,000/monthly (USD)
>Type of service: Exclusive (One slot only)
>Always Online: 5,000 - 6,000
>Updated every: 10 minutes

6 cents per bot week

>\$220.00/weekly - \$800.00/monthly (USD)
>Type of service: Shared (4 slots)
>Always Online: 9,000 - 10,000
>Updated every: 5 minutes

2.5 cents per bot week

What are the effects?

Application-Level Attacks:

- Use expected behaviour of protocols to cause victim to spend resources.
- Difficult to filter - looks like real transactions or requests.
- Load prevents victim from processing real requests.

Attack	Resource Threshold	Requests/bot	Bots needed to exhaust
static http GET	60,000/sec	93 requests/sec at 250 bytes/request	645
dynamic http GET	3,000/sec	93 requests/sec at 250 bytes/request	40
SSL handshake	600/sec	10 requests/sec	60

What are the effects?

Flooding Attacks:

- SYN flood: attacker sends TCP connect requests faster than victim can process them.
- Victim responds then waits for confirmation.
- Victim's connection table fills up, new connections ignored

Attack	Resource Threshold	Requests/bot	Bots needed to exhaust
SYN flood	18,000/sec	450 SYNs/sec	40
SYN flood, tuned server	200,000/sec	450 SYNs/sec	440
SYN flood, dedicated hardware	1,000,000/sec	450 SYNs/sec	2,200

What are the effects?

Bandwidth Attacks:

- Attacker fills the pipe to the victim with high volume of traffic.
- Downlink to victim: must be filtered upstream, and tailored to the specific attack.
- Uplink from victim: small requests causing large responses.

Attack	Resource Threshold	Requests/bot	Bots needed to exhaust
Downlink T1 flood	1.54Mb/s	186Kb/s	8
Downlink T3 flood	43Mb/s	186Kb/s	231
Uplink T1 flood	1.54Mb/s	450Kb/s	3.4
Uplink T3 flood	43Mb/s	450Kb/s	3.95

Outline

- Summary of the Workshop
 - The Nature of the Problem
 - Current Defense Techniques
 - Future Architectures

- Next Steps for the WG

First, Secure the Core Network

- 1. Don't let packets into the core**
 - No way to attack core routers, except through routing. → Still "open": routing protocol
- 2. Secure the routing protocol**
 - Neighbor authentication, maximum routes, dampening, ... → Only attack vector: Transit traffic
- 3. Design for transit traffic**
 - QoS to give VPN priority over Internet → Now only insider attacks possible
 - Choose correct router for bandwidth
- 4. Operate Securely** → Avoid insider attacks

Incident Response Methodology

- 1. Preparation:** Best Practices / Planning
 - 2. Detection:** Something is wrong
 - 3. Classification:** What is wrong?
 - 4. Traceback:** Find ingress path
 - 5. Reaction:** Counter measures
 - ACLs upstream
 - Re-direction
 - spoofed packet trace back
 - 6. Post Mortem** Review
- Time Critical

Step 1: Detection

- Customer Call
 - SNMP: Line/CPU overload
 - Netflow: Counting Flows
 - ACLs with Logging
 - Backscatter
 - Cisco Detector
 - Arbor's Peakflow DoS
 - Similar products
- } “manual”
- } “automatic”

Step 2: Traceback

Non-spoofed: Technically trivial

- Internet Routing Registry: RIPE, ARIN, APNIC, LACNIC
- But: Potentially tracing 1000's of sources...

Spoofed:

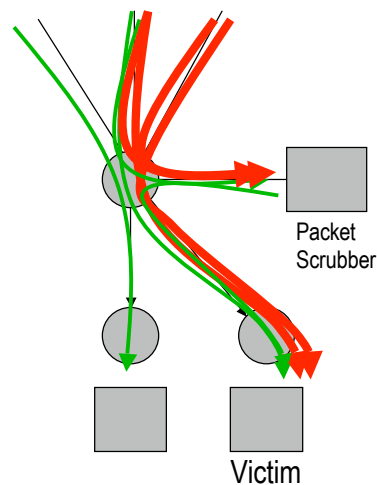
- IP Source Tracker: router by router
- Netflow:
 - Trivial if mechanisms are installed
 - Used by tools (commercial and free)
- ACLs:
 - Has performance impact on some routers.
 - Mostly manual: router-by-router
- Backscatter technique
 - One step, fast. Only for spoofed sources.

Step 3: DoS Containment

- ACLs:
 - Manual, performance impact
- uRPF:
 - Stops non-existing sources
 - Automated with BGP for specific shunning
- Blackholing / Sinkholing / Redirection on Ingress
 - Triggered by BGP
- CAR:
 - Limit attack flow, performance impact

Packet Scrubbers: Distinguish “good” from “bad” traffic

1. Detect
2. Activate: Auto/Manual
3. Divert only victim’s traffic
4. Filter only DoS traffic



DDoS Challenges


- Operationalization of detection, mitigation, and forensics
 - DDoS part of doing business today
 - Automated processes
 - How much rope to give the customer
- Very Large Distributed DoS
 - Inter-provider fingerprinting remains ad-hoc.
- Sophistication of attacks driving arms race with intelligent detection and intelligent hardware filtering.
- VoIP bringing increased pressure on provider detection and filtering capabilities.

Beyond DDoS


- Understanding large, topologically diverse, heterogeneous networks is hard.
 - “What the heck is going on now?”
- Configuration change management continues to represent more outage hours than DDoS.
- BGP and DNS security fragile
- And of course, worms.

Identification/Correlation


BGP Flaps



Packet Size



CPU



Outline

- Summary of the Workshop
 - The Nature of the Problem
 - Current Defense Techniques
 - Future Architectures

- Next Steps for the WG

Architectural Ossification

- The net is already hard to change in the core.
- IP Options virtually useless for extension.
 - Slow-path processed in fast hardware routers.
- NATs make it hard to deploy many new applications.
- Firewalls make it make to deploy anything new.
 - But the alternative seems to be worse.

- Now consider the effect of DoS mitigation solutions....

The Big Challenges

- How can we mitigate DoS attacks and other security threats without sacrificing the future?
 - How to enable application innovation?
 - How to provide robust network services in the face of attack?

Extrapolation of current trends does not bode well....

Future: “The Intelligent Network”

- Network “understands” clients, controls “bad” traffic.
 - Network-based application recognition
 - Deep packet inspection
 - Network admission control
 - Packet scrubbing
- Possibility to define and control policy
 - Security policy
 - QoS policy
- Control and enforcement end-to-end
 - “Pervasive” security

Architectural Solutions

- Can we change the Internet architecture in such a way that the low-level easy attacks become hard/impossible?
 - Tilt the balance of power towards the victim?
 - Reduce the threat from worms?
 - Prevent bandwidth flooding?
 - Prevent spoofing?
 - Provide safe automated pushback mechanisms?
- Preserve the general-purpose nature of the Internet to allow future innovation.

Architectural Solutions

Steps towards a DoS-resistant Internet architecture.

Mark Handley, Adam Greenhalgh, UCL and CII

<http://www.cs.ucl.ac.uk/staff/M.Handley/papers/>

Downstream Knowledge Upstream: Re-Feedback

Bob Briscoe, BT Research and CII


<http://nrg.cs.ucl.ac.uk/mjh/tmp/refeedback.pdf>

Steps towards a DoS-resistant Internet architecture.

General strategy: To solve a problem at the IP level, you need IP-based solutions. *Addressing is the main handle.*

- Separate address space into “client” and “server” addresses. Can only initiate a connection from a client to a server.
- Client addresses are not globally unique, but built up along the path.
 - Similar to the assymetry introduced by NAT, but makes it an explicit part of the architecture.
- Provide ways to enable client-to-client communication only when both clients simultaneously consent.

Path-based Addressing



Claimed advantages...

- No rapidly spreading worms.
- No source address spoofing.
- No reflection attacks.
- Clients completely protected from direct attack.
- Servers protected from attack by servers (and clients are much harder to compromise)
- Simple pushback mechanisms against known malicious clients.
- No per-flow state, except when actively solicited by servers.
- Puzzles make all but the largest DDoS attacks unsustainable.
- Large DDoS attacks cannot use unidirectional traffic.
- The remaining attacks mostly look like a flash crowd.

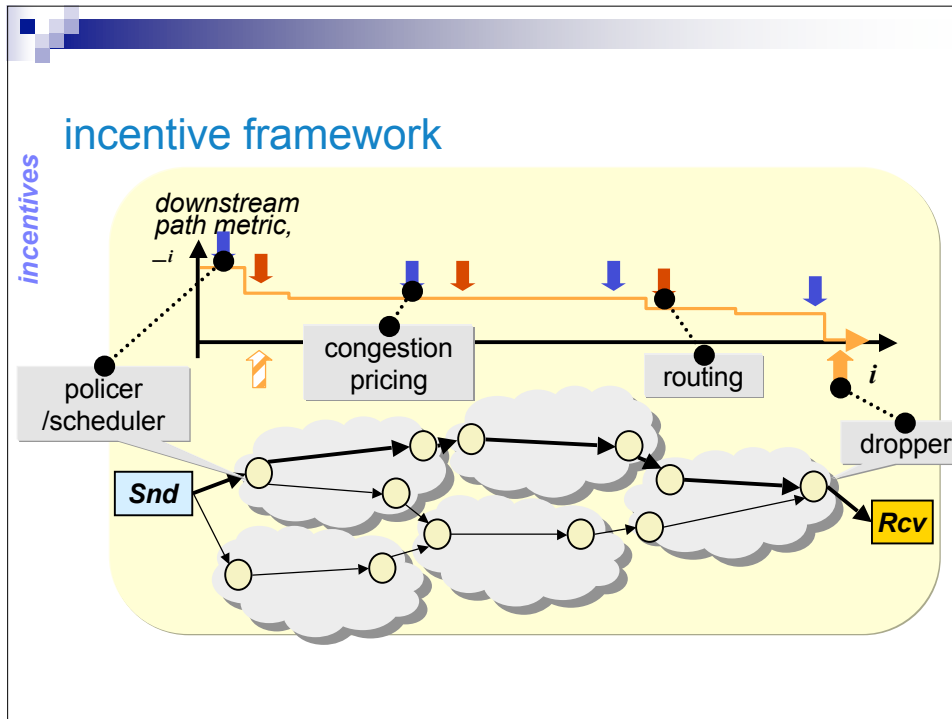
Summary: Preventing DoS by addressing

- Simple architectural changes can make a big difference to the DoS threat space.
- Making asymmetry an integral part of the architecture seems key.
 - “Client” vs “Server” split is a big win.
- Symmetric applications supported through simultaneously setup.
 - More complicated, but not disasterously so.
 - Peer-to-peer may be just too risky though, as it permits fast spreading worms.

Re-feedback (Briscoe)

General strategy:

- Tackle flooding attacks as part of a larger incentive framework.
- Routers provide explicit information about congestion levels by decrementing a congestion field in packets.
 - Feedback explicit information about downstream congestion to the data sender.
 - Data sender-reinserts this feedback information into the packets.
 - Goal is for the sender to set the field correctly so the remaining value is zero at the receiver.
- Policy at ingress and egress to provide incentive for sender to send at the correct rate for the network congestion level.



Summary: re-feedback

- Long-term approach to re-architecting the congestion-control framework for the Internet.
 - Nice alignment of incentive with mechanism
 - Pushes the costs for misbehaviour towards the origin network of the malicious traffic, but provides the mechanism to throttle it.

- Incremental deployment should be possible, but not trivial.



Summary: Future Architectures

- There are architectural changes that could be made that would significantly change the balance of power in the DoS wars.
- Architectural change is *really* hard.
 - It's no-one's responsibility.
 - It requires broad consensus on where to go.
 - It requires strong incentive to get there.
- *Incremental* architectural change is happening anyway.
 - It looks like it takes us to a really bad place.
 - Avoiding stifling future innovation requires an *active injection of energy* into alternatives.



Outline

- Summary of the Workshop
 - The Nature of the Problem
 - Current Defense Techniques
 - Future Architectures
- Next Steps for the WG

Feedback

"Good session"

"Good mix of parties"

"Perhaps more time for discussion... "

"What next ?"

Have volunteers for co-chairs.

Next Steps

- No doubt about the level of interest in this space.
- What can CII do to make a difference?

- Three timescales:
 - *Short*: information sharing
 - *Medium*: work with people on the front line to avoid too much collateral damage from "solutions"
 - *Long*: push for architectural change.



Immediate Next Steps

- Maintain the energy level from the Jan meeting.
 - Set up a framework for anonymized information sharing.

- Next meeting is likely to be February 25th in London.
 - Smaller - probably 20 participants.
 - Invitation-only.
 - Chatham House Rule
 - Please contact Chris Hall about participation
 - chris.hall@thecii.org