

SP Infrastructure Security Survey & Attack Classification

Danny McPherson danny@arbor.net & Ray Hunt ray.hunt@canterbury.ac.nz

Goals

- Given time constraints, focus will be given to providing details of a few popular techniques, rather than providing overly terse information on many techniques – full slide deck provides considerably more detail
- Nothing new or especially exciting here, just information on how some techniques service providers are using to protect their customers and their own infrastructure

Agenda


- 3 Discrete Planes
- DDOS Traceback Techniques
- DDOS Mitigation Techniques
- Infrastructure Security Survey
- IMS Data - If time permits

Internet Address Spaces

- Bogon:
 - Regional Internet Registries
 - RIPE NCC, APNIC, ARIN, LACNIC, AFNIC?
 - RFC 1918/Reserved
 - Unallocated – IANA or an RIR
- Dark Address Space – Allocated and advertised but unused/not sub-allocated
- Active Address Space – In Use

Three Discrete Planes

- Management Plane
 - SNMP, Telnet, Out of Band Access, Etc..
- Control Plane
 - Routing & Signaling Protocols; BGP, OSPF/IS-IS, LDP, Etc..
- Data Plane
 - Packet forwarding functions



Management Plane

Management Plane

- Device Access & Management Functions
- Protocols include:
 - Telnet
 - SSH
 - SNMP
- Also consider console & OOBA, etc..

Control Plane

Control Plane

- Inter-domain routing in the Internet: BGP
- Interior Routing: IS-IS, OSPF, EIGRP, RIP
- MPLS: LDP & RSVP-TE
- Multicast: PIM SSM, MSDP, MP-BGP

Control Plane

- TCP employed for transport of BGP/LDP
 - Makes session vulnerable to many attack vectors (e.g., SYN, RST, etc..)
 - Protection?
 - MD5 TCP Signature Option
 - IPSEC
 - Infrastructure ACLs (iACLs)
 - GTSH
 - IGP support MD5 for many functions
 - Neighbor discovery & adjacency establishment
 - LSA/LSP/Update authentication
 - Etc..
 - Control Plane Policing
 - filter/limit who/what/how much can gain access to a router or switch control plane/route processor

Route Hijacking

- What is it?
 - Announcing Internet address space that belongs to someone else – without their permission
 - Typically via BGP
 - Result of misconfiguration or malicious intent, more often the latter
- Why do it?
 - Anonymous IP space for spamming
 - Launching non-spoofed (e.g., Application Layer) attacks from source addresses within the space
 - Sharing materials anonymously
 - Breaking connectivity to rightful owners of address space (i.e., Denial of Service)

Route Hijacking

- Why is it possible?
 - Routing on the Internet always prefer “longest match” (most specific route) for a given destination
 - No central authoritative source for who owns what addresses, and who provides transit services for address space owners, etc..
 - As such, very little inter-domain prefix filtering, mostly limited to customer/subscriber routing sessions (as opposed to ‘peer’ sessions), if employed at all!

Route Hijacking


- What to do about it?
 - Prefix filtering
 - Need accurate central repository for route ownership data
 - Internet Routing Registries (e.g., RADB)?
 - Regional Internet Registries (e.g., RIPE, ARIN, APNIC)?
 - Secure the routing system – hrmmm..?
 - SBGP- Secure BGP
 - soBGP- Secure Origin BGP
 - IETF:
 - SIDR WG – Secure Inter-Domain Routing IETF WG
 - RPSEC WG – Routing Protocol Security Requirements WG

Route Hijacking

- NANOG 36: Short-lived Prefix Hijacking on the Internet:
 - <http://www.nanog.org/mtg-0602/pdf/booth.pdf>
- ***“Result: between 26 and 95 successful prefix hijackings occurred in December of 2005”***
- Note: prefix hijackings do not include events which appear to be the result of misconfiguration


Slammer Data Plane Impact - A European SPs View

- Some DDOS/worms easier to detect than others...



Apricot 2006 - Perth, Australia

Slammer Control Plane Impact – THE BGP PICTURE



Apricot 2006 - Perth, Australia

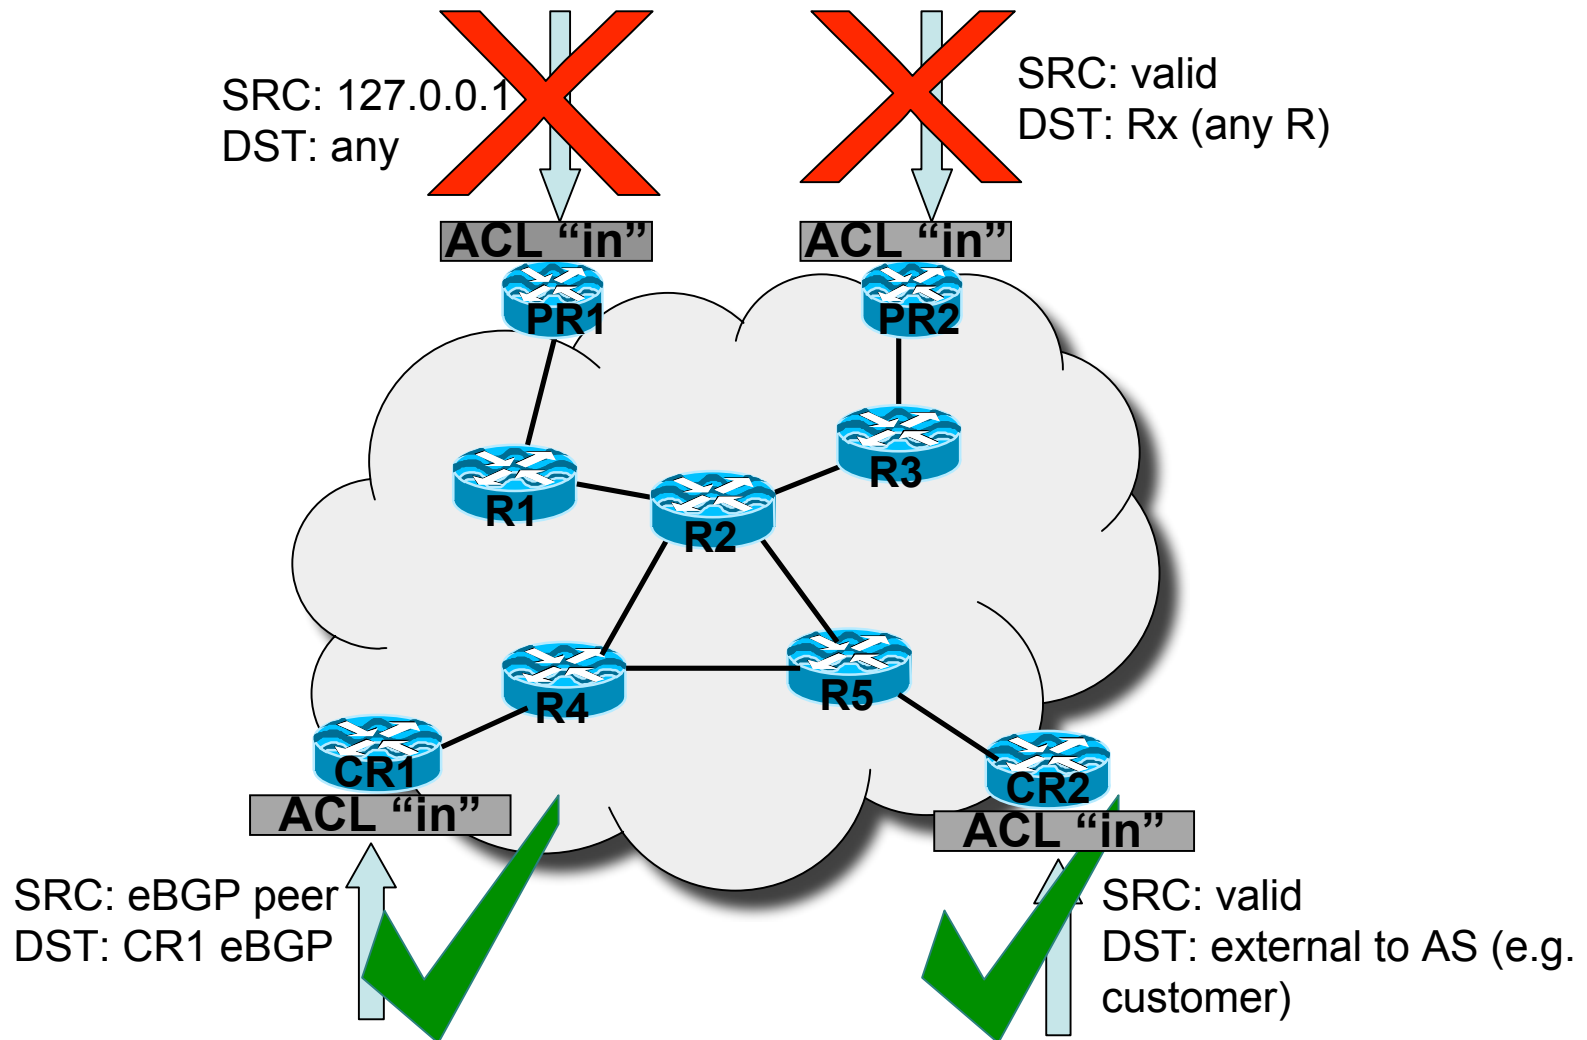
16

Data Plane

Infrastructure ACLs (iACLs)

- Simple concept: instigate policies on the network perimeter that do not allow traffic to enter my network if it is destined for addresses allocated to network infrastructure devices (e.g., routers, switches, etc..)
- Exceptions may be required in order to permit legitimate traffic such as ICMP Echo Requests, etc.. (although you may desire to rate-limit this traffic)
- Never allow packets with source addresses of your own address space to enter your network (could be used for control plane attacks, etc..)

Infrastructure ACLs in Action



Infrastructure ACL Example (Cisco)

—! Deny our internal space as a source of external packets

```
-access-list 101 deny ip our_CIDR_block any
```

—! Deny src addresses of 0.0.0.0 and 127/8

```
-access-list 101 deny ip host 0.0.0.0 any
```

```
-access-list 101 deny ip 127.0.0.0 0.255.255.255 any
```

—! Deny RFC1918 space from entering AS

```
-access-list 101 deny ip 10.0.0.0 0.255.255.255 any
```

```
-access-list 101 deny ip 172.16.0.0 0.0.15.255 any
```


```
-access-list 101 deny ip 192.168.0.0 0.0.255.255 any
```

TTL Security Hack

- Formerly known as BTSH (BGP TTL Security Hack), then GTSH (Generalized TTL Security Hack), and finally, GTSM (Generalized TTL Security Mechanism)
- Defined in RFC 3682
- Can be performed in hardware data path (in forwarding ASICs)
- Initially applied to BGP, but can be employed for any IP-based protocols
- Exploits routers native TTL decrement behavior

TTL Security Hack

- Protect peers from multi-hop attacks
- Routers are configured to transmit packets with TTL of 255 and reject received packets with TTL of < 254
- Removes possibility of injected packets affecting session
- Applied on external BGP peering sessions where iACLs could not be applied

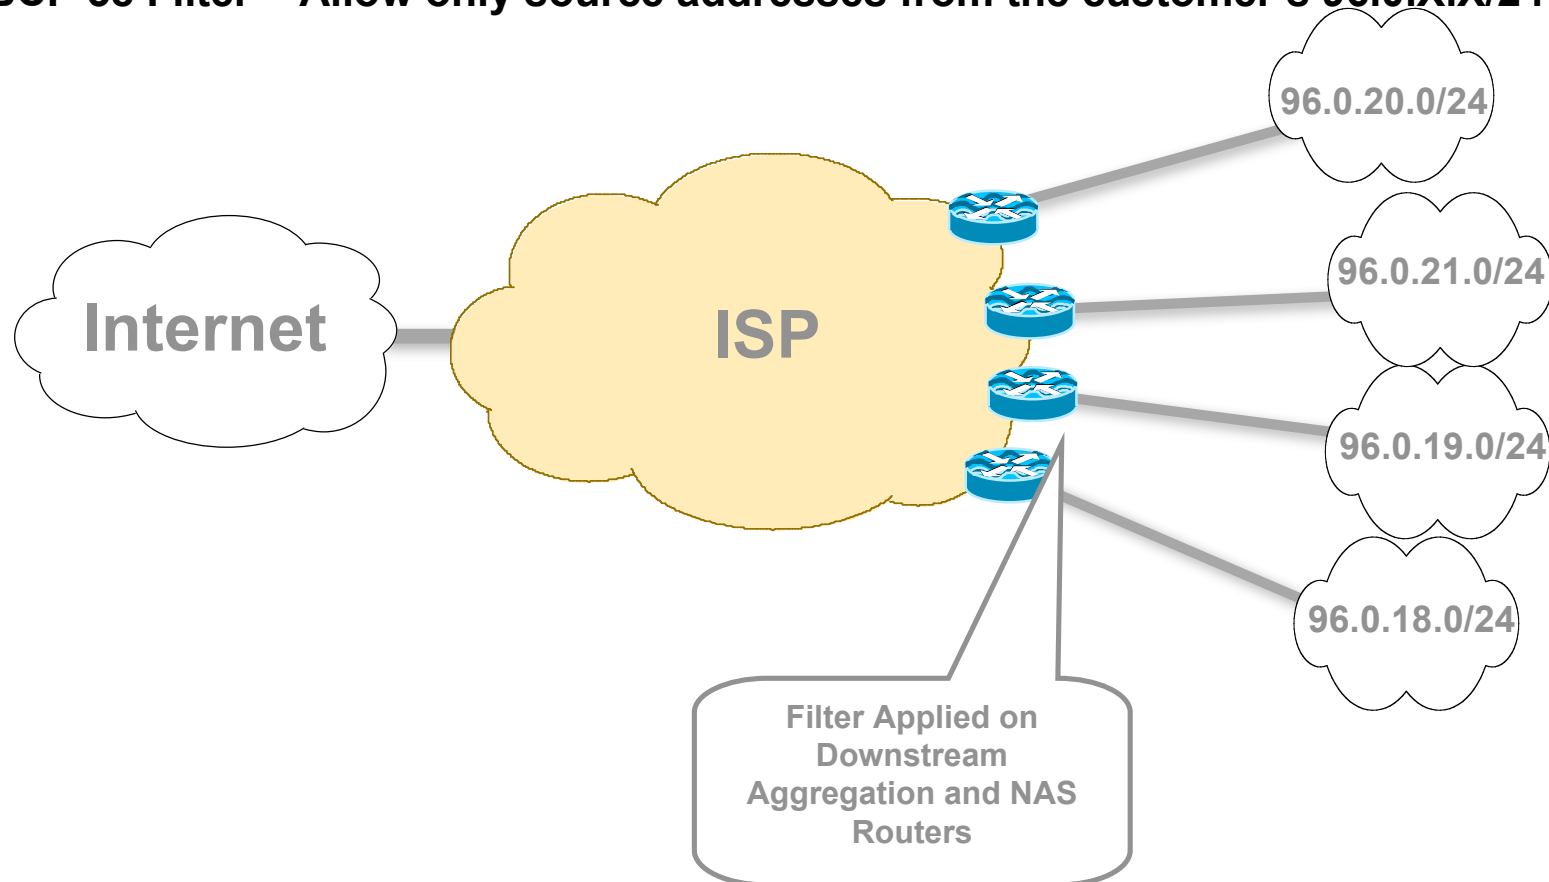


Ingress Filtering

- RFC 3704/BCP 84 updates RFC 2827/BCP 38 - mitigate address spoofing and packets destined to bogon space
- Employ packet filtering mechanisms such that subscribers/customers are only allowed to source packets from addresses which they've been allocated – apply filters as close to the edge as possible, filter as precisely as possible
- Extremely difficult to maintain filters for customers with large numbers of routes
- Rarely applied to “peers” on the Internet, per ACL generation is extremely difficult and hardware would be required to support hundreds of thousands of filters
- Removes plausibility of spoofing – makes tracing attacks/malicious activity back to actual source much simpler

Ingress Packet Filtering

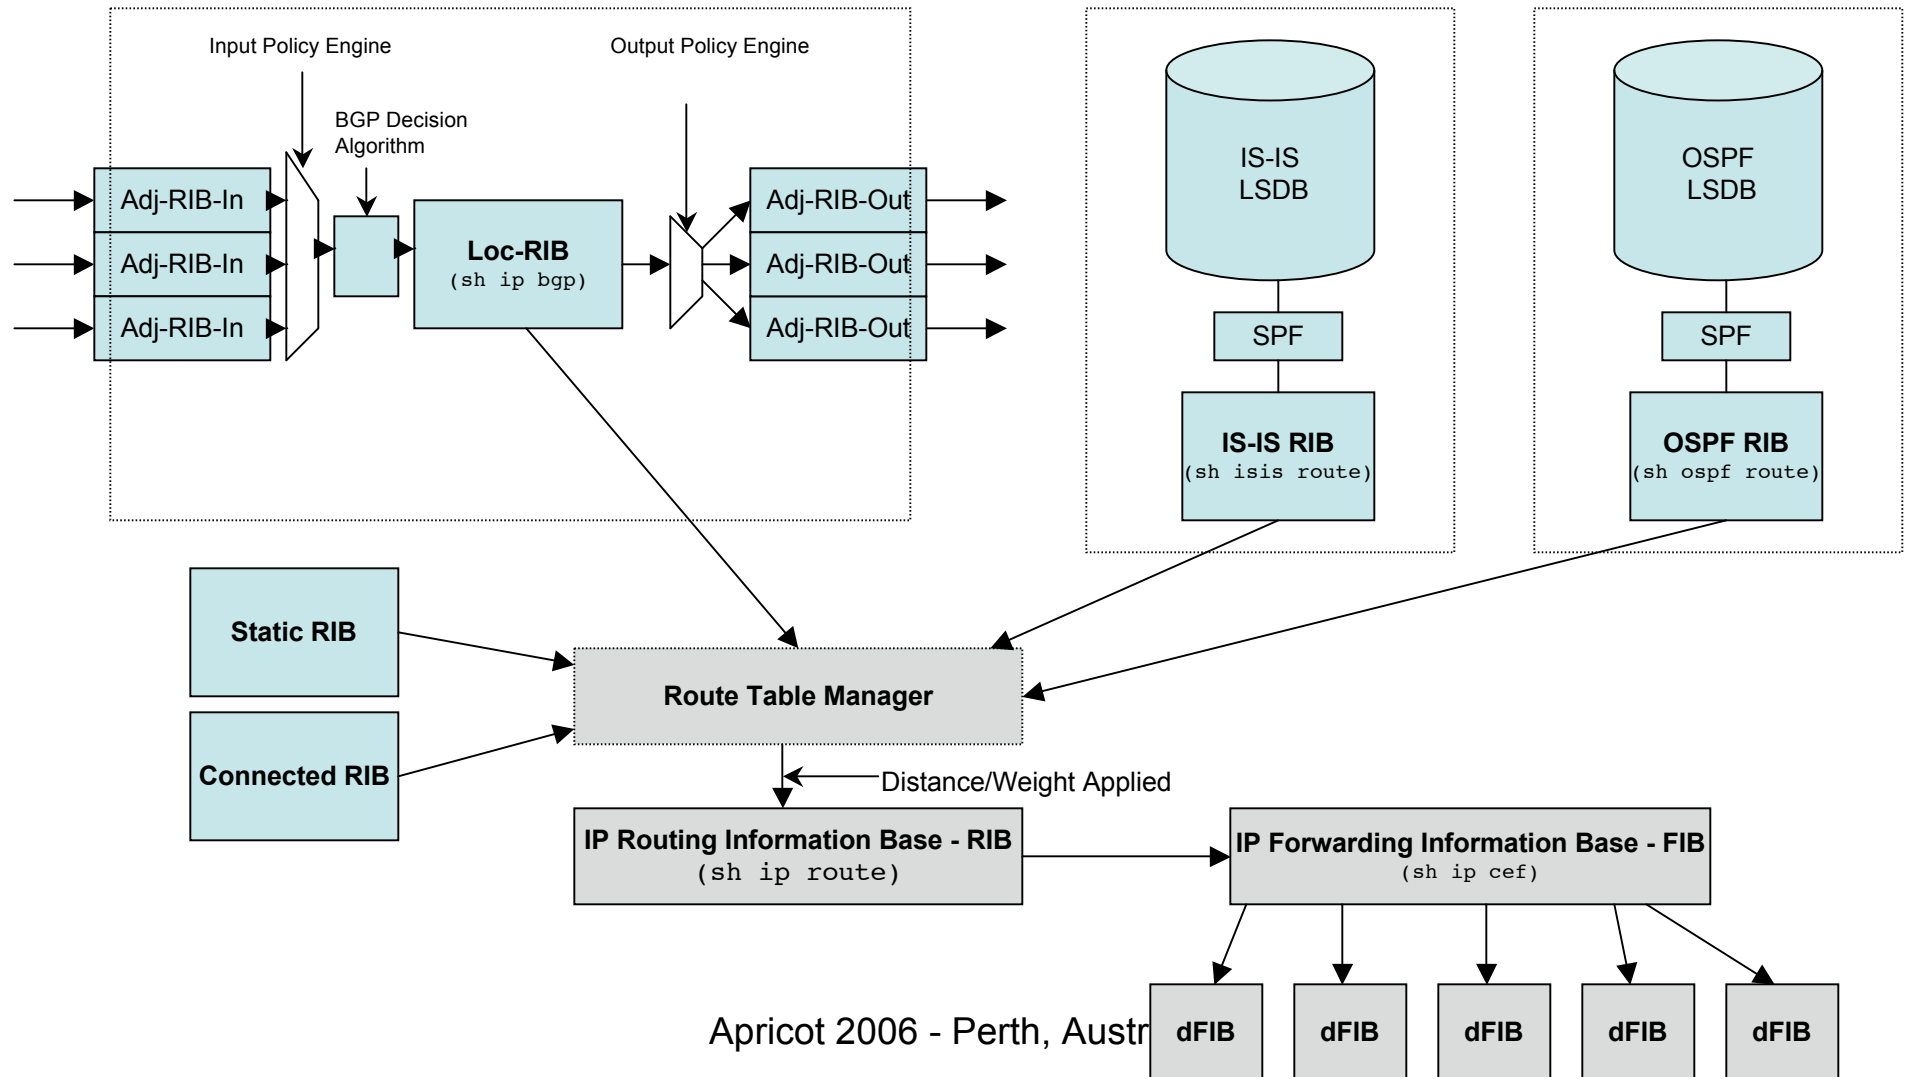
ISP's Customer Allocation Block: 96.0.0.0/19
BCP 38 Filter = Allow only source addresses from the customer's 96.0.X.X/24



What's in a FIB?

- FIB == Forwarding Information Base (i.e., forwarding table)
- Correspondingly, RIB == Routing Information Base (i.e., Routing Table)

Conceptual Router Architecture (RIBs & FIBS)

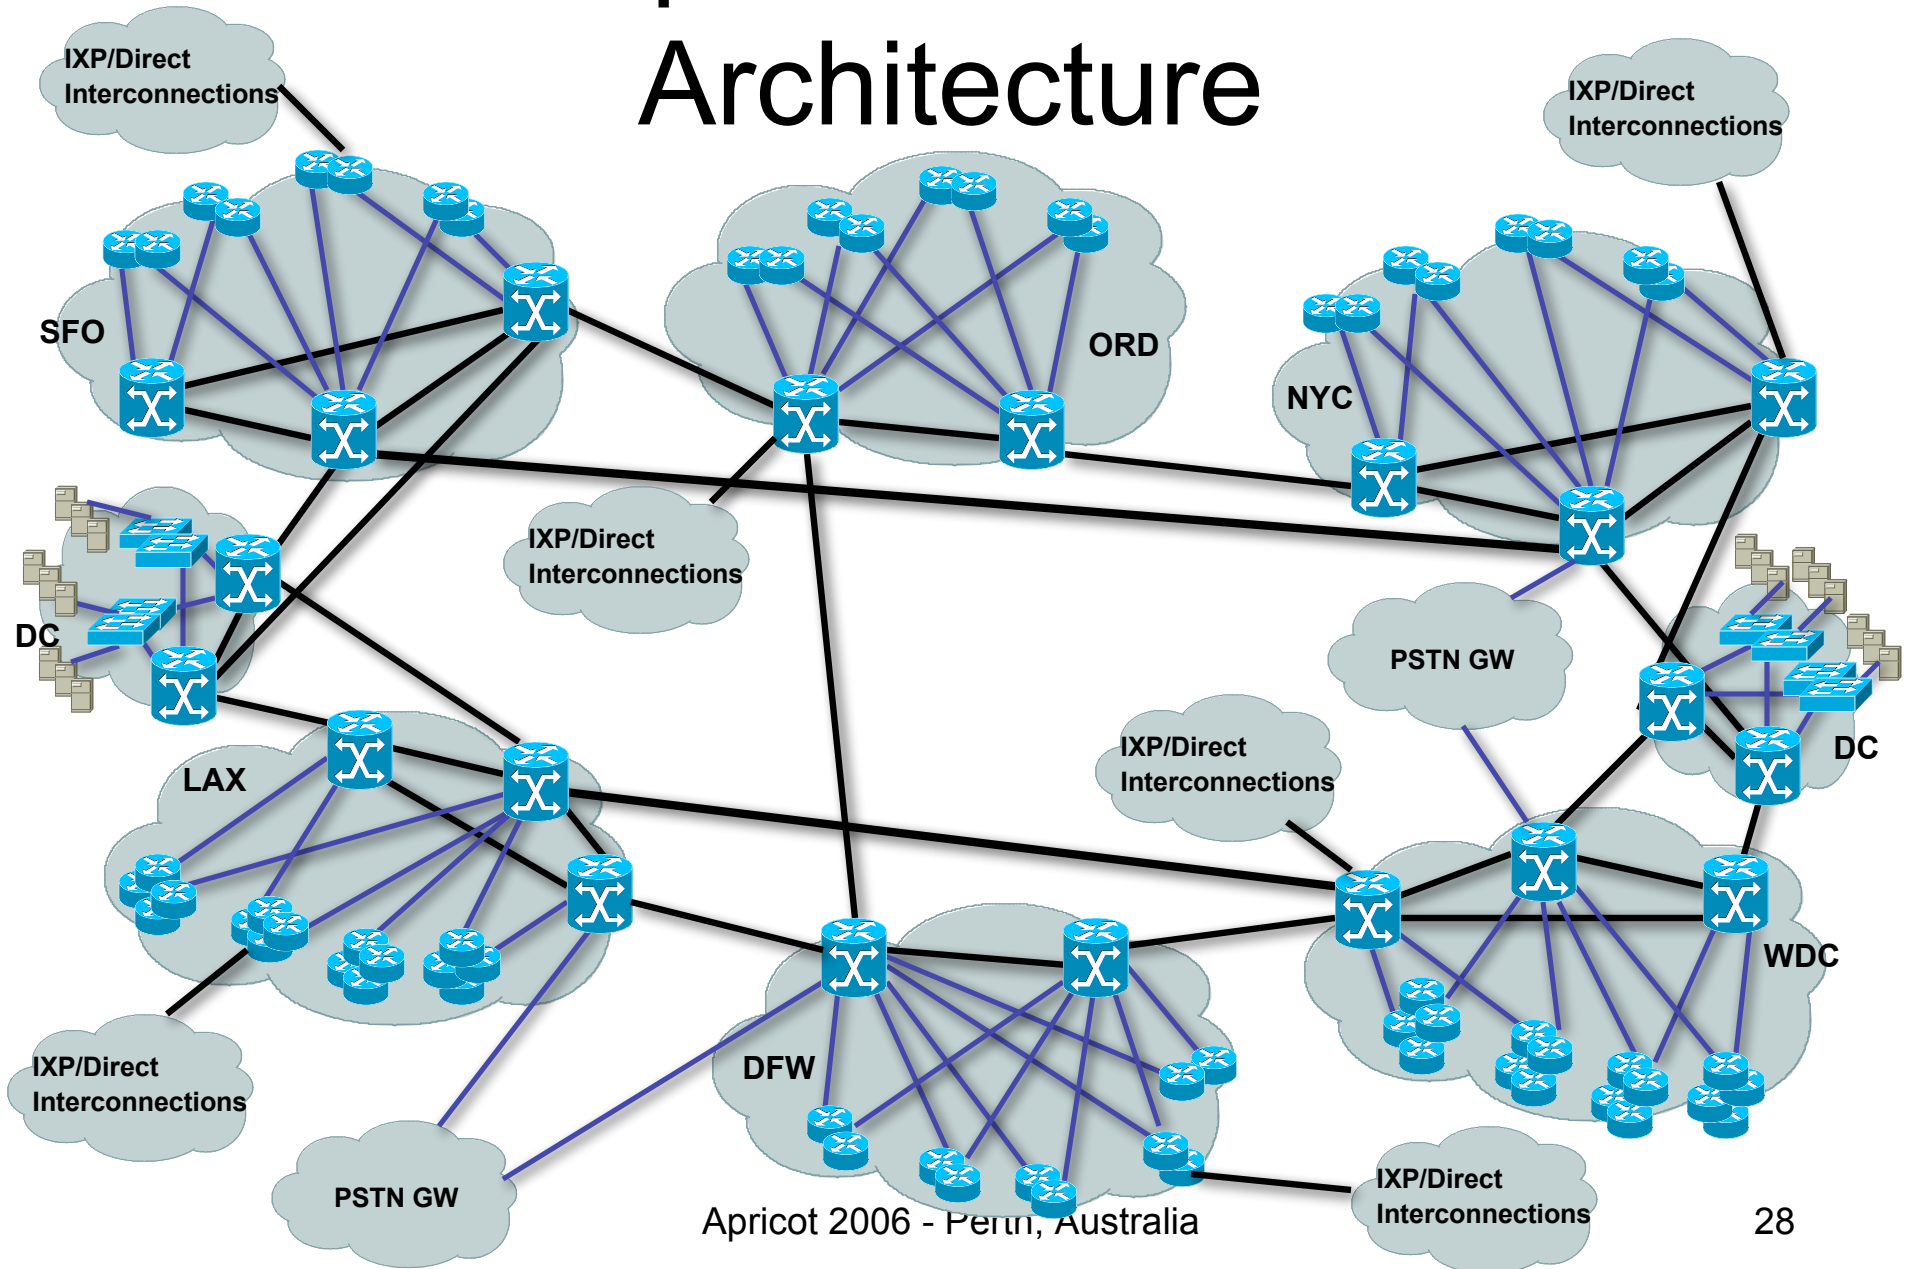


Apricot 2006 - Perth, Austr

RIB/FIB Generation

- Shortest Path First (SPF) algorithm ran on Link State Database (LSDB) to determine next hop node to reach each destination for link state protocols (e.g., IS-IS or OSPF)
- BGP only [typically] installs a single best path to any given destination, even if multiple paths are presented via Adj-RIBs-In. BGP [typically] only advertises a single best path for each reachable destination prefix.
- RTM applies local weights that result in routes from different sources having varying degrees of preference (e.g., connected -> static -> IS-IS -> BGP). Only a single route is typically installed in RIB – even if multiple paths exist!
- RIB contains route origination information that is not necessary in FIB (e.g., route came from IS-IS, has weight of n , etc..)
- FIB is in essence a subset of RIB, but contains next hop forwarding information (e.g., next hop Link Layer address, such as Ethernet MAC address). FIB is akin to CEF table in Cisco-Speak..
- FIB-based forwarding can be performed locally, or FIB can be distributed to linecards to perform distributed forwarding functions

Sample ISP Network Architecture

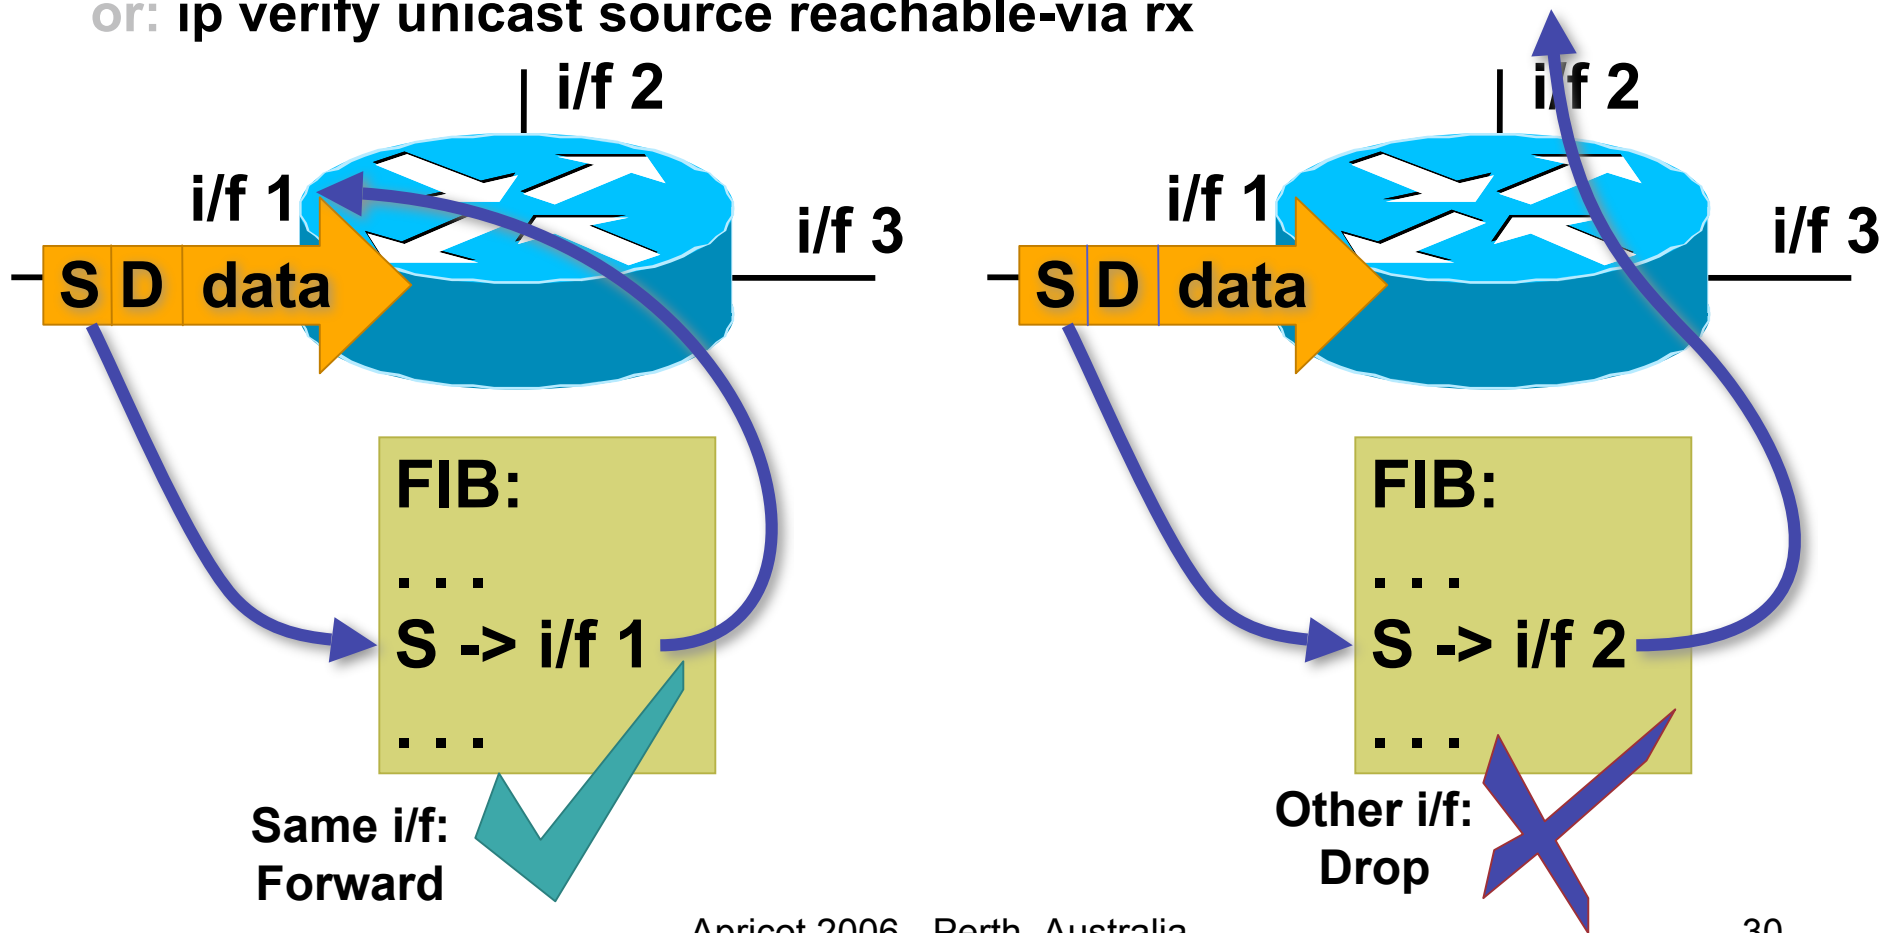


uRPF

- Traditional Unicast Reverse Path Forwarding (RPF)
 - RPF functions akin to that of multicast RPF-based forwarding; forward packet only if received on preferred interface from which source address is considered reachable
 - Works fine for unicast IF multiple paths don't exist

Strict uRPF Check

`router(config-if)# ip verify unicast reverse-path`
or: `ip verify unicast source reachable-via rx`



Effects of Asymmetric Routing

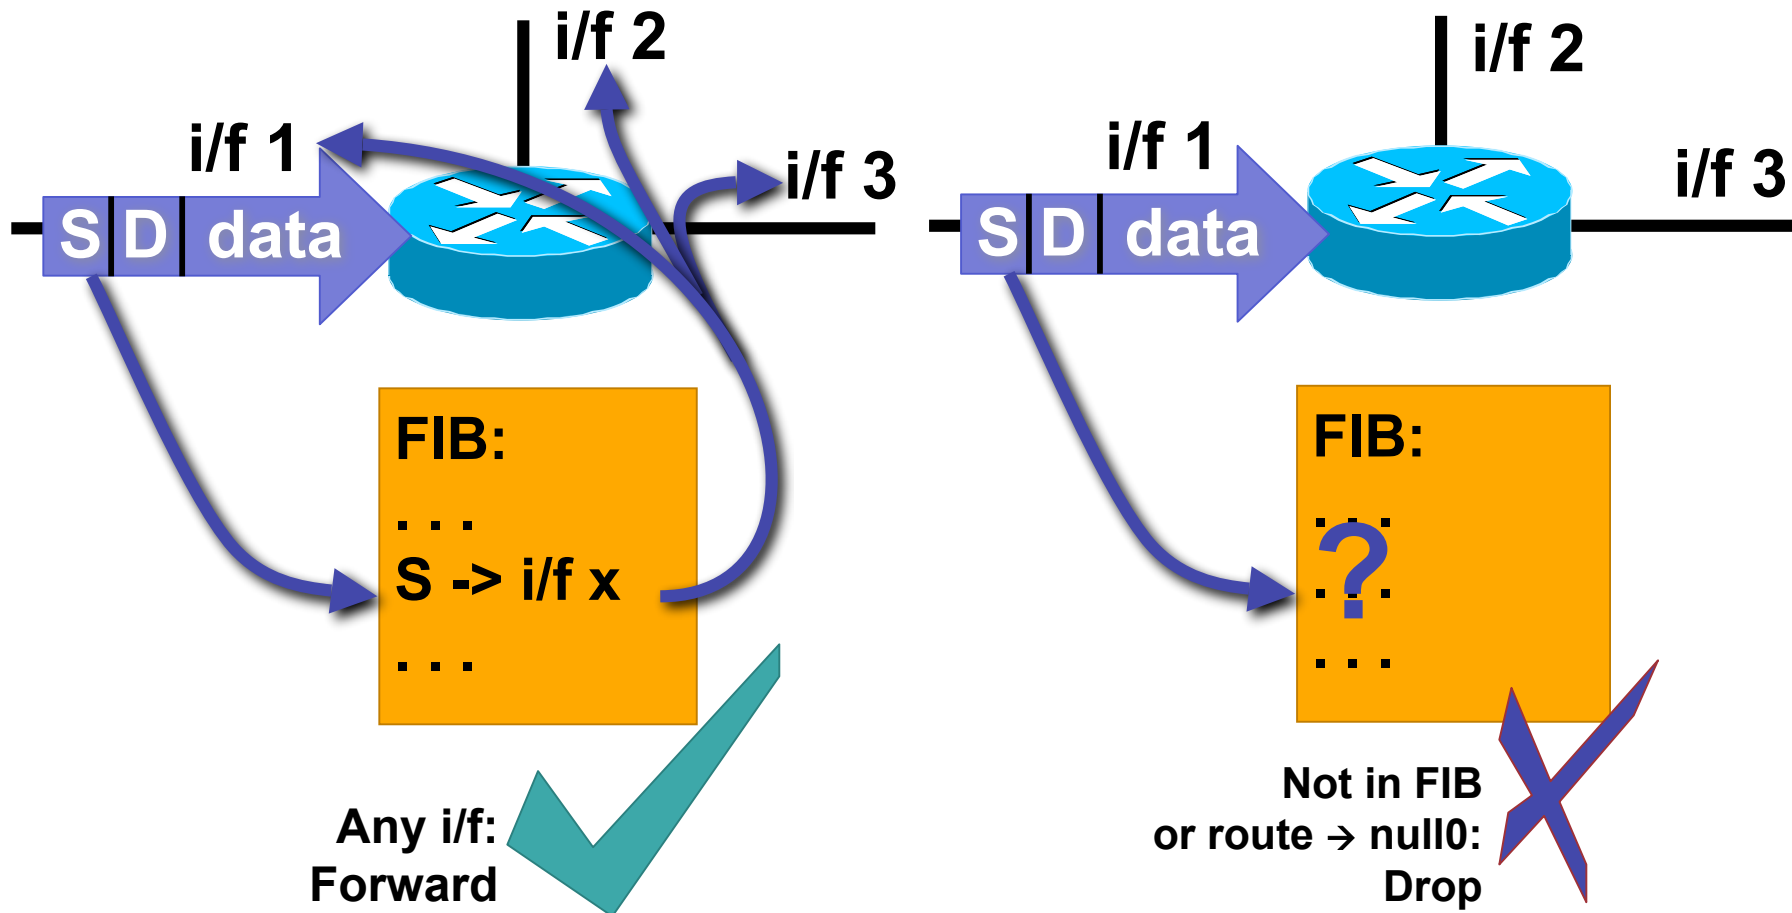
- Traditional uRPF becomes problematic if asymmetric routing is possible
- Packets received on interfaces that aren't the preferred interface associated with reaching the prefix listed in the source address field of the packet – the packet will be discarded
- Dense interconnection models and multi-homing on the Internet therefore make “strict mode” uRPF problematic

“Loose mode” uRPF

- Remember those different types of Internet Address Spaces...?
- Let's at least nuke packets sourced from bogon address spaces – i.e., If NO FIB entry exists for the address prefix from which the source of the packet is defined, discard the packet
- If ANY FIB entry exists, regardless of the ingress interface, forward the packet – perhaps encouraging spoofing of addresses that are routed on the Internet?

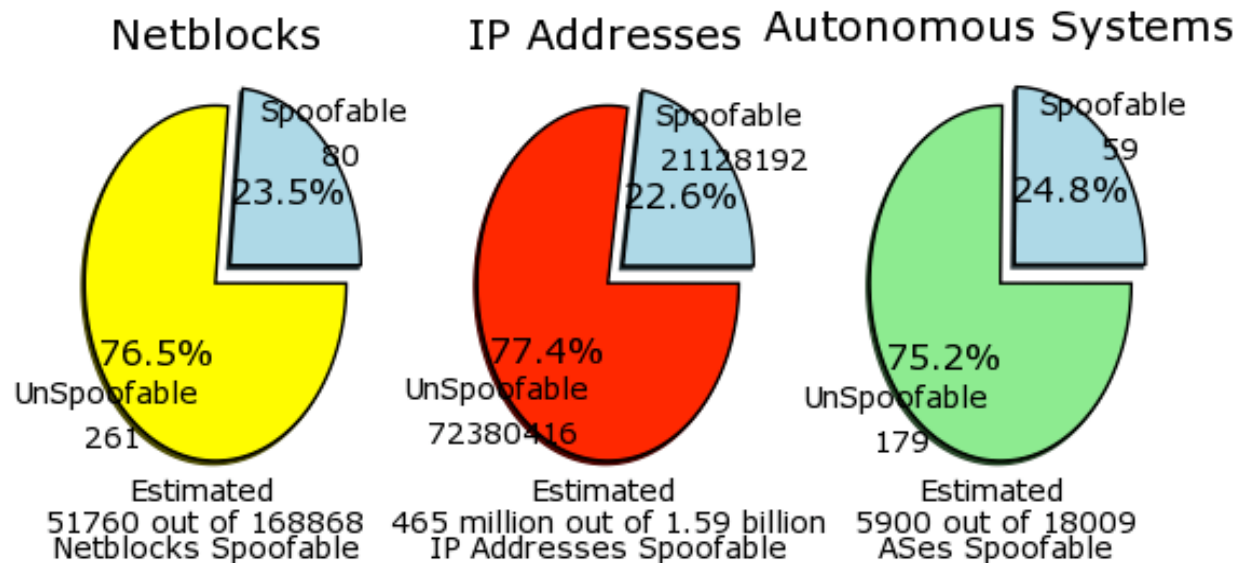
Loose uRPF Check

`router(config-if)# ip verify unicast source reachable-via any`




MIT ANA Spoofer Project

- <http://momo.lcs.mit.edu/spoofer>
- ~23% of observed netblocks corresponding to ~24% of observed ASes allow spoofing



DDOS Traceback

Attack Detection Capabilities



- Most operators had some commercial tools in place, though not covering the entire network perimeter
- Most provided employed multiple mechanisms for attack detection
- ISPs in wholesale/transit mostly rely on NOC trouble tickets (i.e., customer calls)

Traceback

Traceback to ingress network perimeter

Manual

- Packet filters (ACLs)
- IP accounting
- Disable interfaces

Backscatter

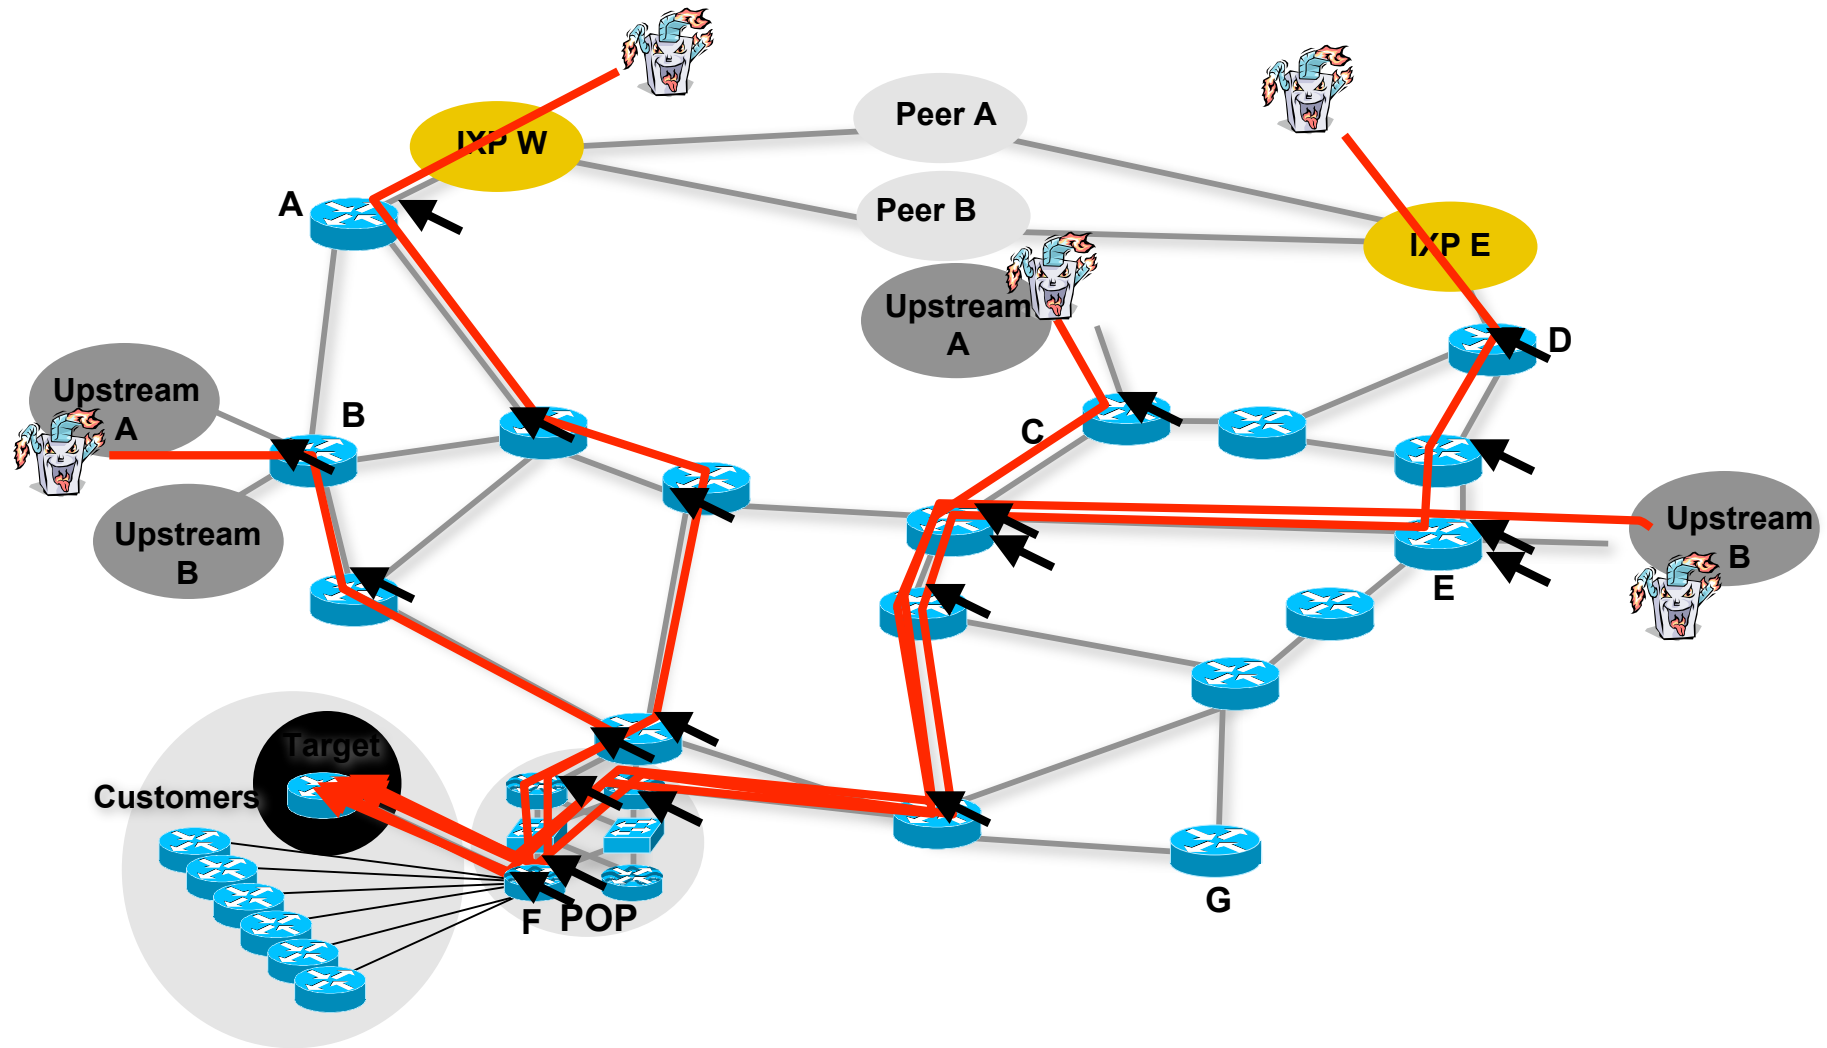
Packet/CEF (Cisco Express Forwarding) Accounting

NetFlow/JFlow/sFlow/IPFIX

Traceback: Manual

- Steps
 - Began with classification ACLs and counters at network egress to customer
 - Filtered attack traffic as it was destined for customer premise
 - Manually traced back through the network, hop-by-hop, interface by interface (automated with ACL scripting tools; i.e., dostracker.pl)
 - ACLs applied at network ingress to drop traffic destined for victim IPs
- Limitations
 - Error-prone
 - May impact service availability
 - Tedious & Very time consuming; especially for well-distributed attacks
 - Fully characterizing and accounting for full impact of attack is still unlikely

Traceback: Manual



Traceback: Manual

- Classification ACL (cACLs) applied to customer interface:

```
access-list 101 permit icmp any any echo
access-list 101 permit icmp any any echo-reply
access-list 101 permit udp any any eq echo
access-list 101 permit udp any eq echo any
access-list 101 permit tcp any any established
access-list 101 permit tcp any any range 0 65535
access-list 101 permit ip any any
```

```
interface serial 10/1/1
ip access-group 101 out
```

```
router# sh ip access-list 101
```

```
Extended IP access list 101
```

```
permit icmp any any echo (2 matches)
permit icmp any any echo-reply (2171374 matches)
permit udp any any eq echo
permit udp any eq echo any
permit tcp any any established (150 matches)
permit tcp any any (15 matches)
permit ip any any (45 matches)
```

- Once attack type is classified, Traceback ACL (tACLs) applied to egress then subsequent upstream interfaces back towards network ingress

```
access-list 170 permit icmp any any echo-reply log-input
access-list 170 permit ip any any
```

```
interface serial 10/1/1
ip access-group 170 out
```

```
router# sh log
```

```
%SEC-6-IPACCESSLOGDP: list 170 permit icmp 1.1.1.1 (Serial0/1/1
*HDLC*) -> 192.168.1.1 (0/0), 1 packet
%SEC-6-IPACCESSLOGDP: list 170 permit icmp 2.2.2.2 (Serial0/1/1
*HDLC*) -> 192.168.1.1 (0/0), 1 packet
%SEC-6-IPACCESSLOGDP: list 170 permit icmp 3.3.3.3 (Serial0/1/1
*HDLC*) -> 192.168.1.1 (0/0), 1 packet
%SEC-6-IPACCESSLOGDP: list 170 permit icmp 4.4.4.4 (Serial0/1/1
*HDLC*) -> 192.168.1.1 (0/0), 1 packet
%SEC-6-IPACCESSLOGDP: list 170 permit icmp 5.5.5.5 (Serial0/1/1
*HDLC*) -> 198.168.1.1 (0/0), 1 packet
```


Traceback: Flow-based

- Trace attack by matching fingerprint/signature at each interface via passive monitoring:
 - Flow data (e.g., NetFlow, cflowd, sFlow, IPFIX)
 - Span Data
 - PSAMP (Packet Sampling, IETF PSAMP WG)
- Number of open source and commercial products evolving in market
- Non-intrusive, widely supported

Flow-based Detection

- Monitor flows (i.e., Network and Transport Layer transactions) on the network and build baselines for what normal behavior looks like:
 - Per interface
 - Per prefix
 - Per Transport Layer protocol & ports
 - Build time-based buckets (e.g., 5 minutes, 30 minutes, 1 hours, 12 hours, day of week, day of month, day of year)

Detect Anomalous Events: SQL “Slammer” Worm



Apricot 2006 - Perth, Australia

Flow-based Detection (cont)

- Once baselines are built anomalous activity can be detected
 - Pure **rate-based** (pps or bps) anomalies may be legitimate or malicious
 - Many **misuse** attacks can be immediately recognized, even **without** baselines (e.g., TCP SYN or RST floods)
 - **Signatures** can also be defined to identify “interesting” transactional data (e.g., proto udp and port 1434 and 404 octets(376 payload) == slammer!)
 - Temporal compound signatures can be defined to detect with higher precision


Flow-based Commercial Tools...

Anomaly 150228	Get Report: PDF XML
----------------	---

ID	Importance	Duration	Start Time	Direction	Type	Resource
150228	High 130.0% of 2 Kpps	17 mins	03:34, Aug 16	Incoming	Bandwidth (Profiled)	Microsoft 207.46.0.0/16 windowsupdate.com

Traffic Characterization

Sources	204.38.130.0/24
	204.38.130.192/26
	1024 - 1791
Destination	207.46.248.234/32
	80 (http)
Protocols	tcp (6)
TCP Flags	S (0x02)

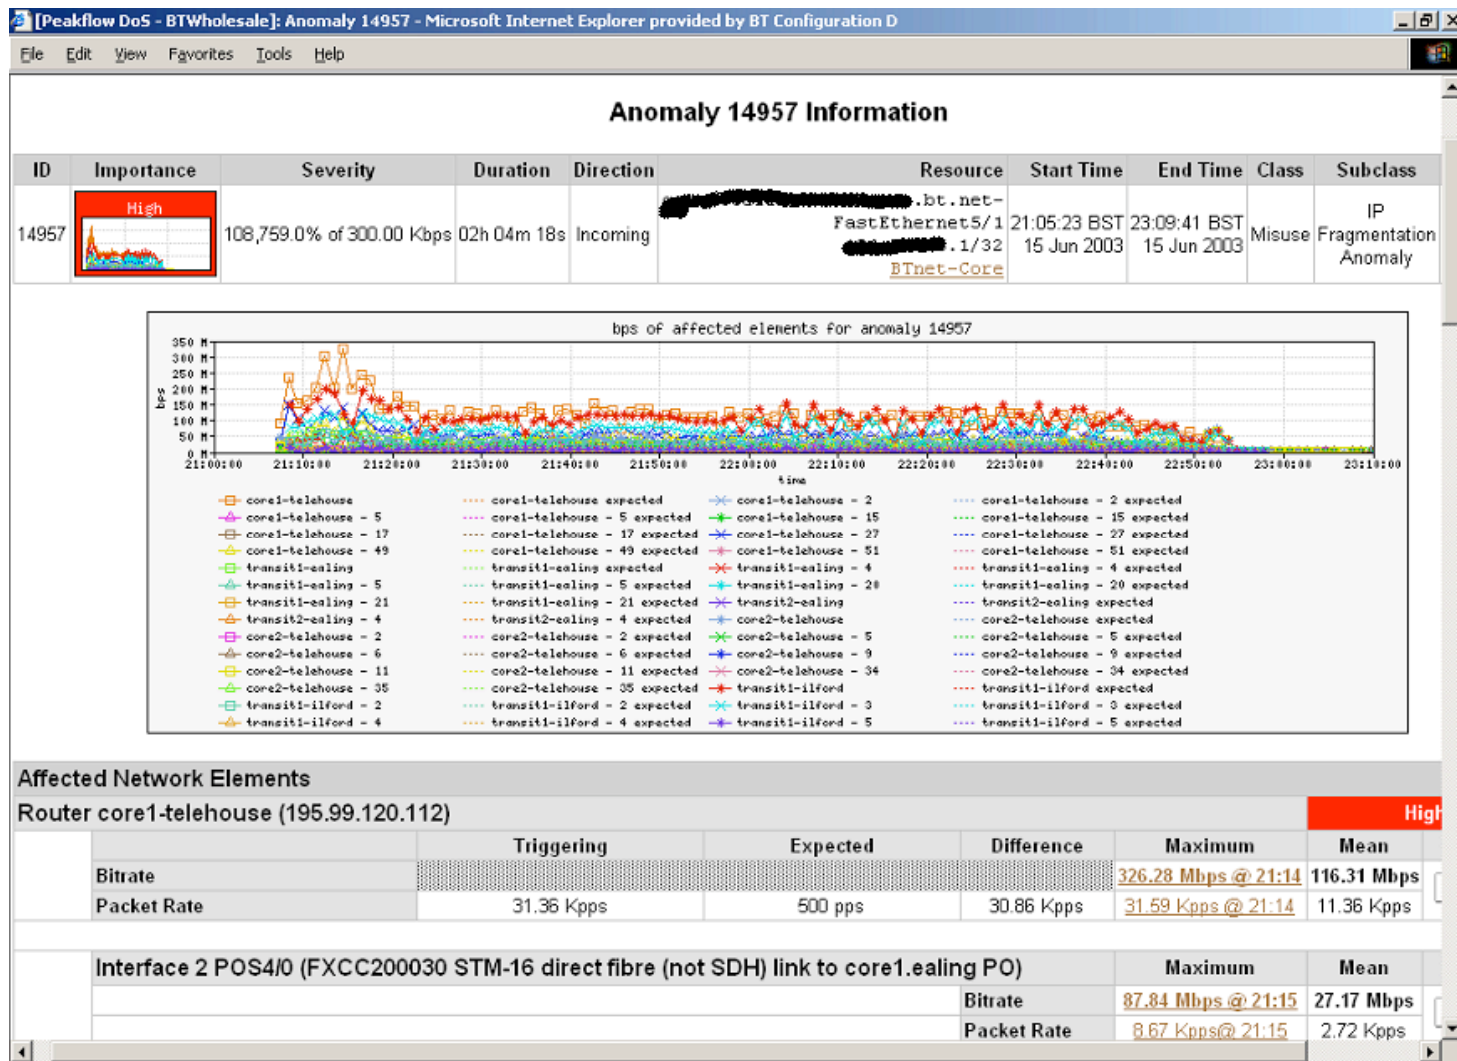


Affected Network Elements		Expected	Observed bps		Observed pps		
	Importance	pps	Max	Mean	Max	Mean	
Router nl-chi3 198.110.131.125	High						
Interface 67 at-1/1/0.14 <i>pvc to WMU</i>		26	832 K	563.1 K	2.6 K	1.7 K	Details

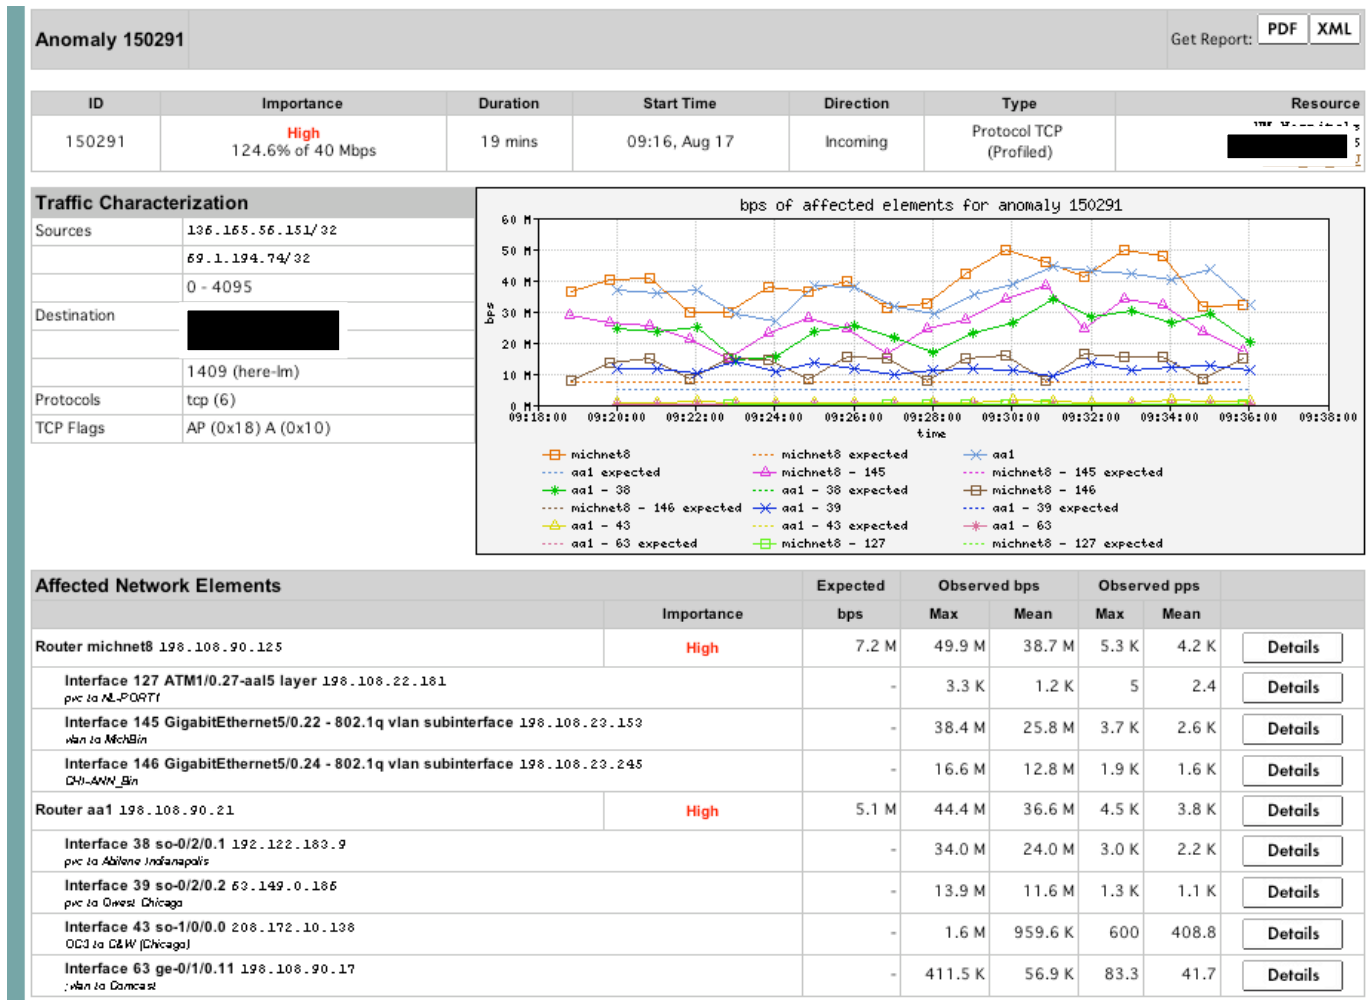
Anomaly Comments

Commercial Detection

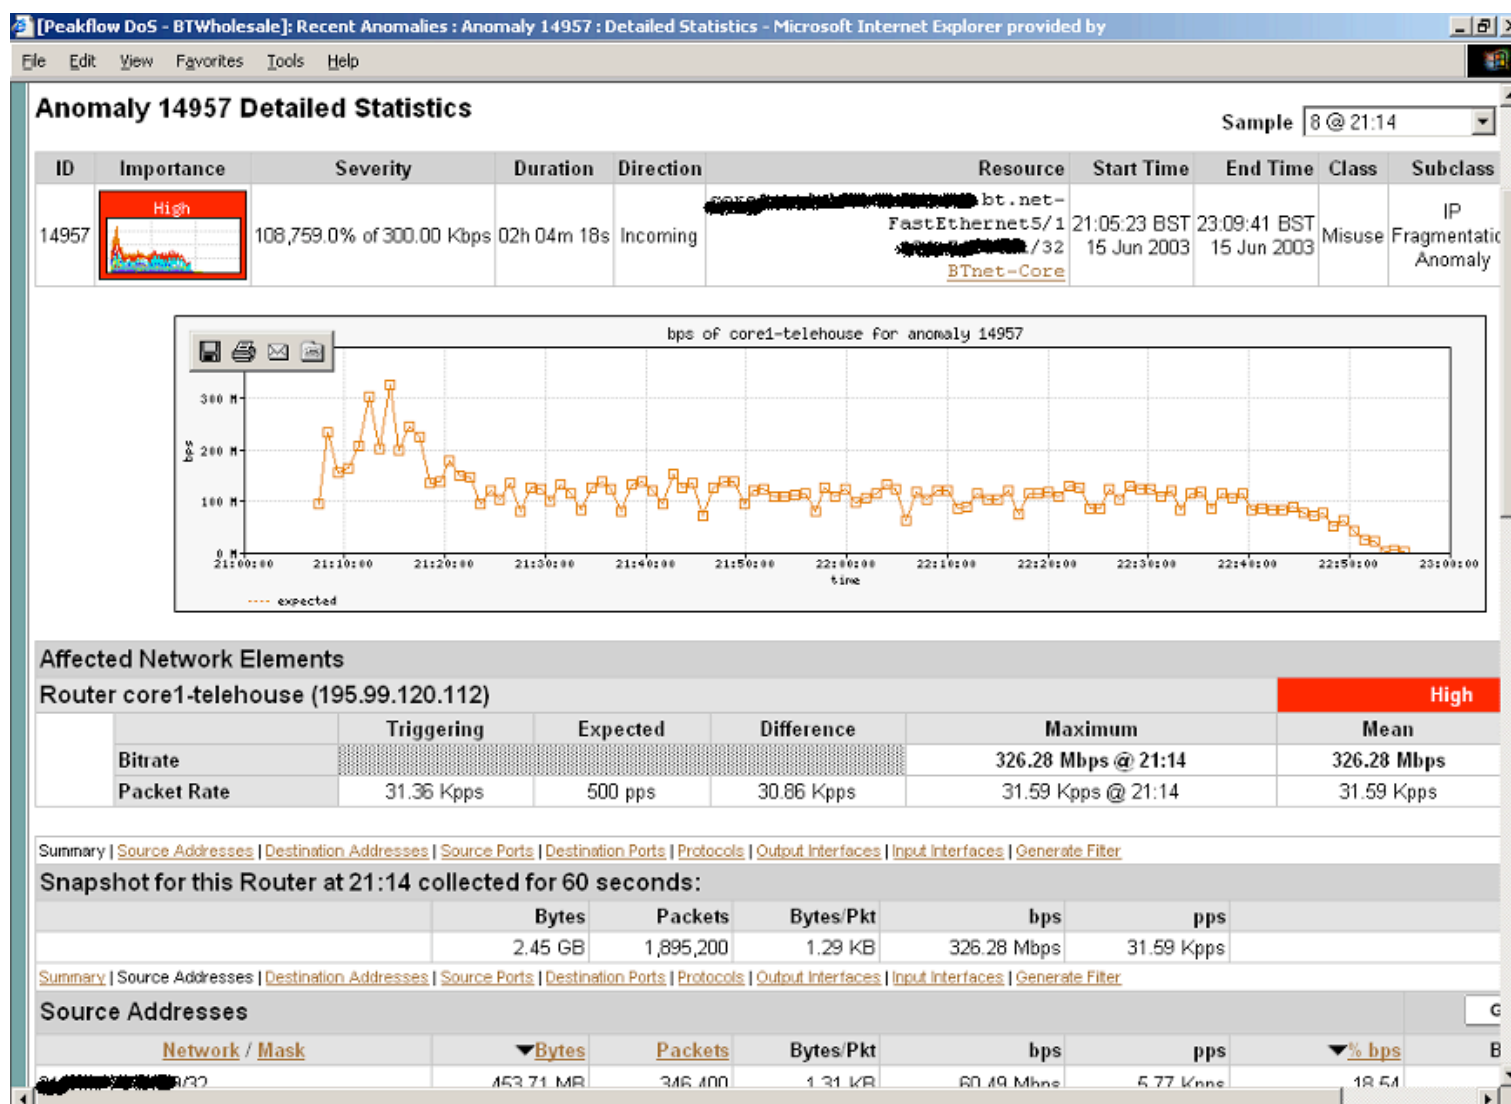
A Large Scale DOS attack



Traceback: Commercial








Commercial Traceback: More Detail



Apricot 2006 - Perth, Australia


DDOS Mitigation Techniques

Potential Mitigation Options

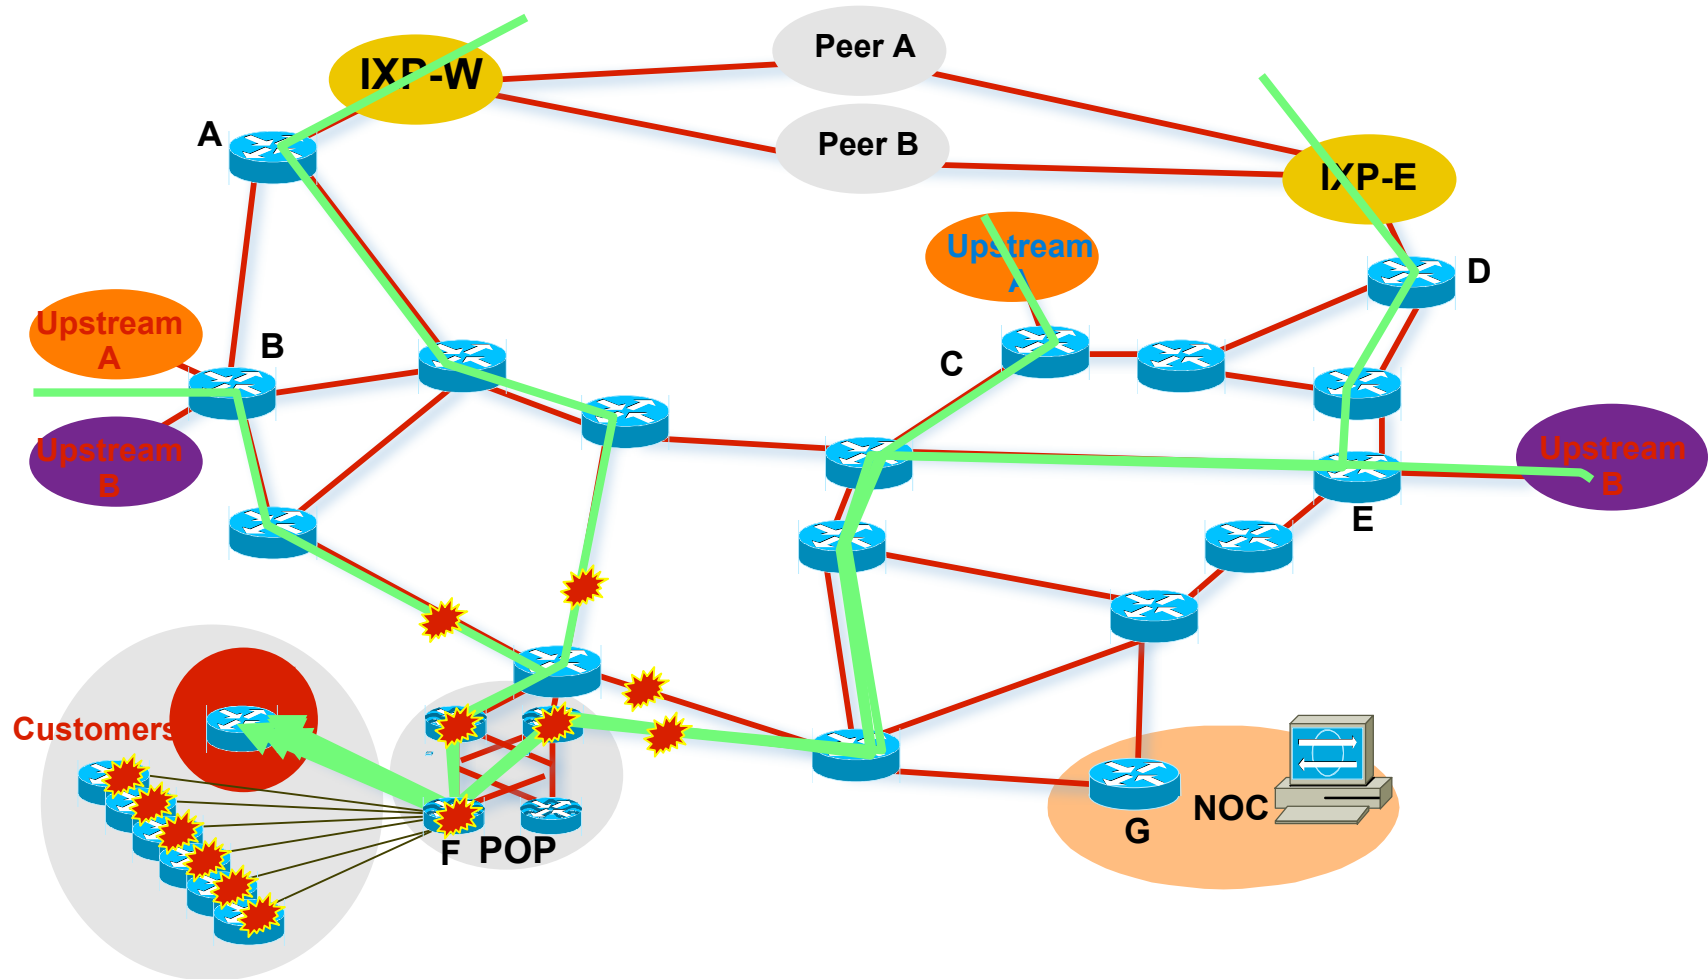
- Do Nothing
- Actively respond:
 -  Packet filters (e.g., ACLs) or rate-limit (e.g., CAR)
 -  BGP remote-triggered drop
 - Blackhole (dst == Null 0/discard interface)
 - uRPF loose check (src == Null 0/discard interface)
 - Customer-performed
 - FLOW_SPEC
 -  Intelligent filtering (e.g. divert to CloudShield, Cisco Guard)
 -  Peer/upstream filtering
 -  CPE filtering firewall, IDS or similar

Most Common Mitigation Approaches

- UMich/Arbor Survey of 40+ tier1/tier2 ISPs
- Most common approach is to BGP null-route destination
- BGP destination more scalable than ACLs and most common mitigation approach



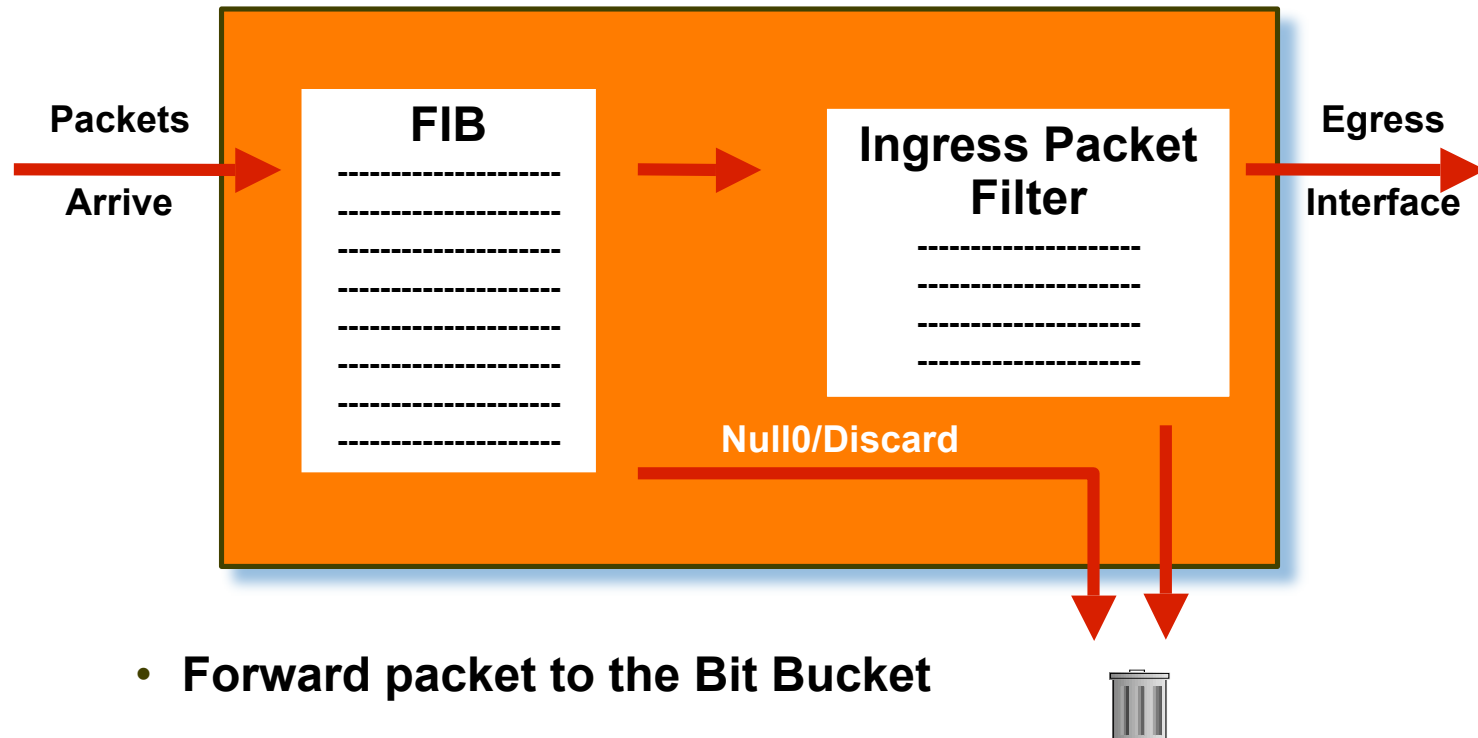
Mitigation Issues: Avoid Collateral Damage



Blackhole Routing

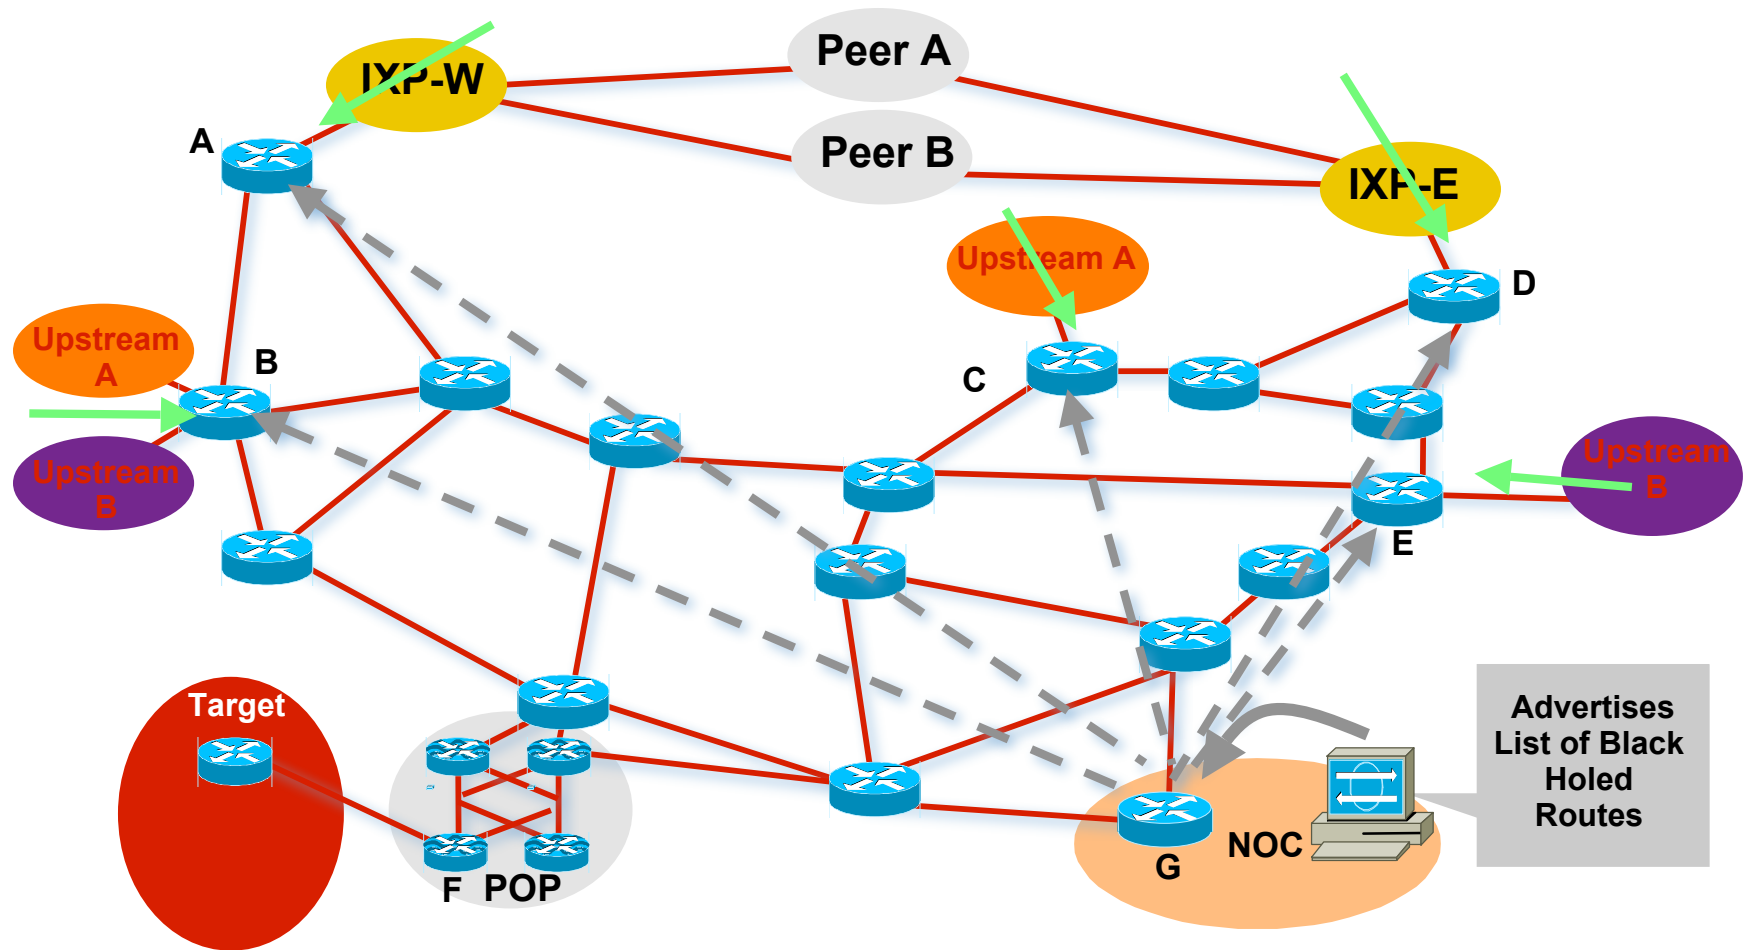
- *Blackhole Routing* or *Blackhole Filtering* results in packets being forwarded to a router's *bit bucket*, also known as:
 - Null interface
 - Discard Interface
- Initially worked only based on IP destination address, per it's exploit of a router's forwarding logic (can work based on source as well w/uRPF)
- Typically results in desired packets being dropped with minimal or no performance impact
- At any given time, tier-1 providers average 500 active BGP null routes

Exploits Forwarding Logic



- Forward packet to the Bit Bucket
- Saves on CPU and ACL processing

Blackhole Routing



Blackhole Trigger

- Select a small unused block (e.g. TestNet 192.0.2.0/24)
- Configure static route with TestNet to Null 0 on every router
- Prepare BGP speaking router to act as trigger device (next slide)

Blackhole Trigger Configuration

Redistribute
Static with a
route-map

Match
Route Tag

```
router bgp 65501
!
redistribute static route-map static-to-bgp
!
route-map static-to-bgp permit 10
match tag 66
set ip next-hop 192.0.2.1
set local-preference 50
set community no-export
set origin igp
!
route-map static-to-bgp permit 20
!
ip route 192.0.2.1 255.255.255.255 null 0
```

Set BGP
NEXT_HOP to
the Trigger

Set LOCAL_PREF

Blackhole: Community Based Trigger

- BGP Community-based triggering allows for more granular control over where you drop the packets.
- Examples of flexibility
 - Community #1 can be for all routers in the network.
 - Community #2 can be for all peering routers. No customer routers – Preserves customer-customer connectivity if the victim is within your AS.
 - Community #3 can be for all customers (e.g., to push a inter-AS traceback to the edge of your network).
 - Trigger Communities per ISP Peer can be used to only black hole on one ISP Peer's connection. Allows for the DOSed customer to have partial service.
- Three parts to the trigger:
 - Static routes to Null 0 on all the routers.
 - Trigger router sets the community and advertises the BGP update.
 - Reaction Routers (on the edge) matches community and sets the next-hop to the static route which maps to Null0.

Customer Initiated Mitigation

- Several providers accept more-specifics of customer routes with destination-based BGP blackholing community attached
- No source-based blackholing
- Only accept more-specifics of customer prefixes

BGP FLOW_SPEC

- Use BGP to specify explicit Network & Transport Layer filters
- Basic idea:
 - Use BGP to distribute more specific information about flows beyond destination and/or source address
 - A flow specification is an n-tuple consisting of several matching criteria that can be applied to IP packet data.
 - May or May not include reachability information (e.g., NEXT_HOP).
 - Well-known or AS-specific COMMUNITIES can be used to encode/trigger a pre-defined set of actions (e.g., blackhole, PBR, rate-limit, divert, etc..)
 - Application is identified by a specific (AFI, SAFI) pair and corresponds to a distinct set of RIBs.
 - BGP itself treats the NLRI as an opaque key to an entry in its database.

Some Good Resources

- <http://www.securite.org/presentations/ddos/COLT-SwiNOG9-ExpDDoS-NF-v1.ppt>
- <http://www.securite.org/presentations/secip/SwiNOG7-iSecurityDDoS-v101b.ppt>
- <ftp://ftp-eng.cisco.com/cons/isp/security>
- <http://arbor.net>
- <http://www.nanog.org>


Infrastructure Security Survey

Background


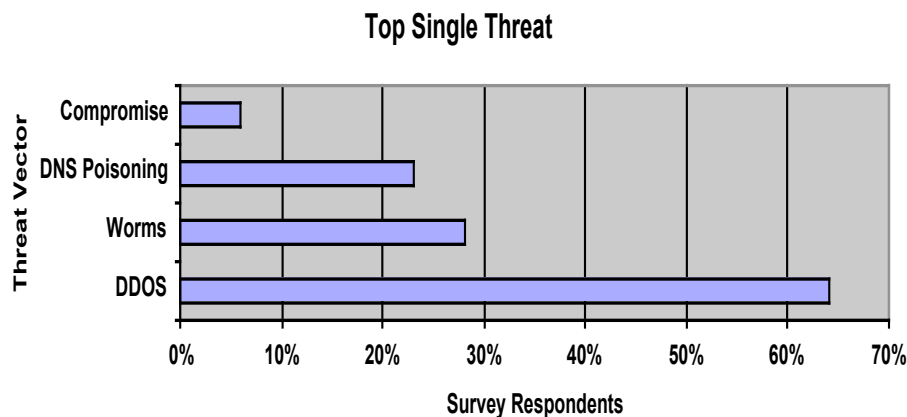
- Earlier this year a survey was conducted among network security operators
- The survey was targeted at obtaining an understanding of some of the operational security aspects occurring in large Internet networks today
- 36 network operators responded to the survey - some responses were, hmmm.. less than trivial to parse
- The survey was composed of 32 multiple choice and free response questions
- The findings of this survey are reflects in the following slides

Survey Respondents

Respondent Distribution




Primary Threat Concerns



- DDOS was top concern, with worms coming in second
- Implicit DOS impacts of worm more concerning than worm payload itself
- BGP vulnerabilities weren't listed as anyone's top concern

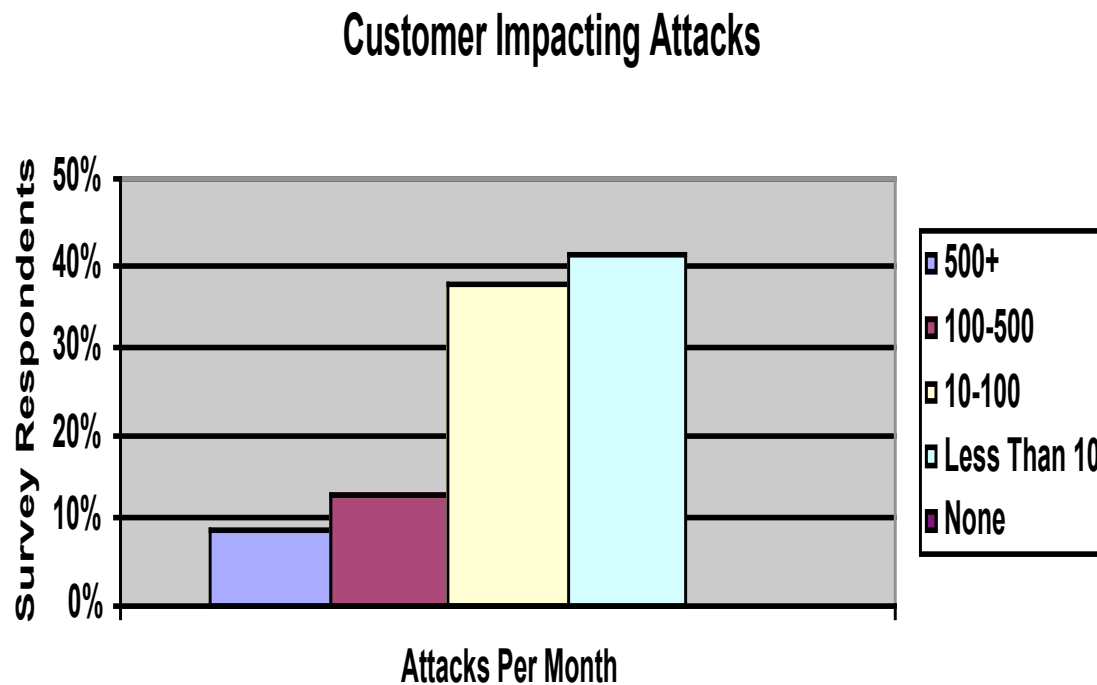
Attack Vectors

Primary Observed Attack Vectors



- While TCP SYN and UDP flooding “brute-force” attacks were most commonly observed actionable attacks, more sophisticated attacks such as multi-modal and Application Layer attacks were reported as well


Customer Impacting Attacks




- An average of 40 actionable customer impact attacks per month were reported

Infrastructure Impacting Attacks

- Infrastructure impacting attacks were far less common, on the order of 1-2 per month on average
- These attacks were both directly at the infrastructure, as well as a result of collateral damage from customer attacks




Largest Attacks Observed



- Attacks greater than 10 Gbps sustained bandwidth were reported
- Not a large differential in largest attack ever v. largest in past six months - perhaps indicative of worsening problem

Attacks Reported to Law Enforcement

- Of actionable attacks, only ~1.5% are reported to law enforcement agencies
- Some of the reasoning provided:
 - Jurisdictional issue
 - Online gambling technically illegal in US
 - IRC users unloved
 - Customer profiles - they don't want attacks recorded
 - Lack of evidence and forensics data
 - Large amount of uncertainty from legal department



Botnet Observations

- No noticeable trends in sizing of botnets from respondents - although attacks are appearing to be better organized
- Few reported any tools track botnets
- One provider indicated that the botnets appeared smaller, but much better organized. This provider described large pools of “reinforcements” that joined the attack as the provider initiated different mitigation efforts. Another provider described armies comprised of “divisions” of smaller groups, noting: “the little bastards appear to be learning actual military tactics.”

DDoS Overview: What is Under Attack?

- Most frequently attacked sites include:
 - IRC servers
 - Gambling, especially offshore
 - Porn sites
- Additional survey reports included:
 - Residential users
 - Web hosting
 - The Chinese
 - RIAA related sites



Security Teams

- Quite a variation in size and reporting structure for security teams across respondent organizations
- Some tier-1s had dedicated infrastructure security teams of as many as 9 full-time employees, others had only 2-4, many of whom were also responsible for backbone engineering functions
- Residential broadband and dial-up providers seemed to have the largest security-related organizations
- Across all respondents, approximately 50% of the security teams were part of network engineering, 25% were part of operations, and 25% were an independent entity
- Some respondents privately complained that the design/architecture teams have no responsibility for the edge and beyond

Social Engineering

- Large European provider had internal tiger team successfully phish security/authentication information from NOC
- Social Engineering will always be a factor

Conclusions....

- DDOS is still the primary concern for network security operations
- Brute-force attacks most popular and clearly effective
- Detection and mitigation mechanisms need to improve and be deployed ubiquitously
- Until miscreants are prosecuted it's unlikely things will get better
- Tools and staffing are a major factor in operator response capabilities

About the Survey

- Plan to conduct bi-annually
- Thanks to all those that responded or reviewed the results
- Hope to get more details and pose less ambiguous questions in future revisions
- Full survey report can be found here:
 - http://www.arbor.net/sp_security_report.php

Thanks!


Internet Motion Sensor

Backscatter Analysis

For more information on the Internet Motion Sensor:

<http://ims.eecs.umich.edu> ims@umich.edu

IMS Overview




- Much of this non-productive traffic is observed by unused addresses

Internet Motion Sensor

- The Internet Motion Sensor monitors almost 100 darknets globally:
 - Deployed at Tier 1 ISPs, Large Enterprise, Broadband, Academic, National & Regional ISPs
- **17,096,192** IPs monitored
- **1.15%** of routed IPv4 space
- **31** /8 blocks with an IMS sensor
- **21%** of *all* routable /8 blocks have at least one sensor


About that Backscatter?

- One method of quantifying spoofing is to analyze backscatter data:




How much backscatter?

- About 3,000,000 packets/hour on a /8 darknet!




What kinds of spoofing?

- Dominated by spoofed TCP:



What kinds of spoofing?

- Dominated by spoofed SYNs:



Top 10 Ports Targeted by Spoofed TCP SYNs:

TCP Port	Service	Packets
80	HTTP (HyperText Transfer Protocol)	38805062
7000	W32.Gaobot, Spyboter, W32.Mydoom, W32.Mytob	2342659
6904	-	919211
300	-	828651
100	-	757745
25	SMTP (Simple Mail Transfer Protocol)	563894
6000	X11 - X-Windows	480937
3389	Microsoft Terminal Server (RDP)	391371
22	SSH	161991
7777	cbt/Omicron-HTTP-Server	155682 ⁸⁵

Data Plane Filtering Issues

- Capabilities of linecard, router or switch impact where and what can be filtered
 - Number of ACLs severely constrained (e.g. at most 1K and usually in the 100s)
 - ACLs may impact forwarding performance (element specific as possible)
 - Flexibility of filter language
 - Usually IP 5 tuple
 - E.g., Juniper supports packet length
- Related issues:
 - Sequence of filters may impact performance (higher hit counts earlier in path)
 - Configuration management (humans prone to error (e.g., employ *tool* or *rancid*)
 - Impact of installing ACLs (e.g., application forwarding hit, recompilation to take effect, etc..)
 - Many ACLs do not filter fragments
 - Avoid collateral damage