



Welcome to the new IP reality

Best practices that failed for YouTube

Łukasz Bromirski
Channel Systems Engineer, CCIE #15929
lbromirski@cisco.com



Agenda

- Intro
- Incident analysis
- Best practices 101
- Few thoughts about state of best practices
- Q&A



What did happen on 24 feb 2008?



Let's take a look what really did happen

YouTube case

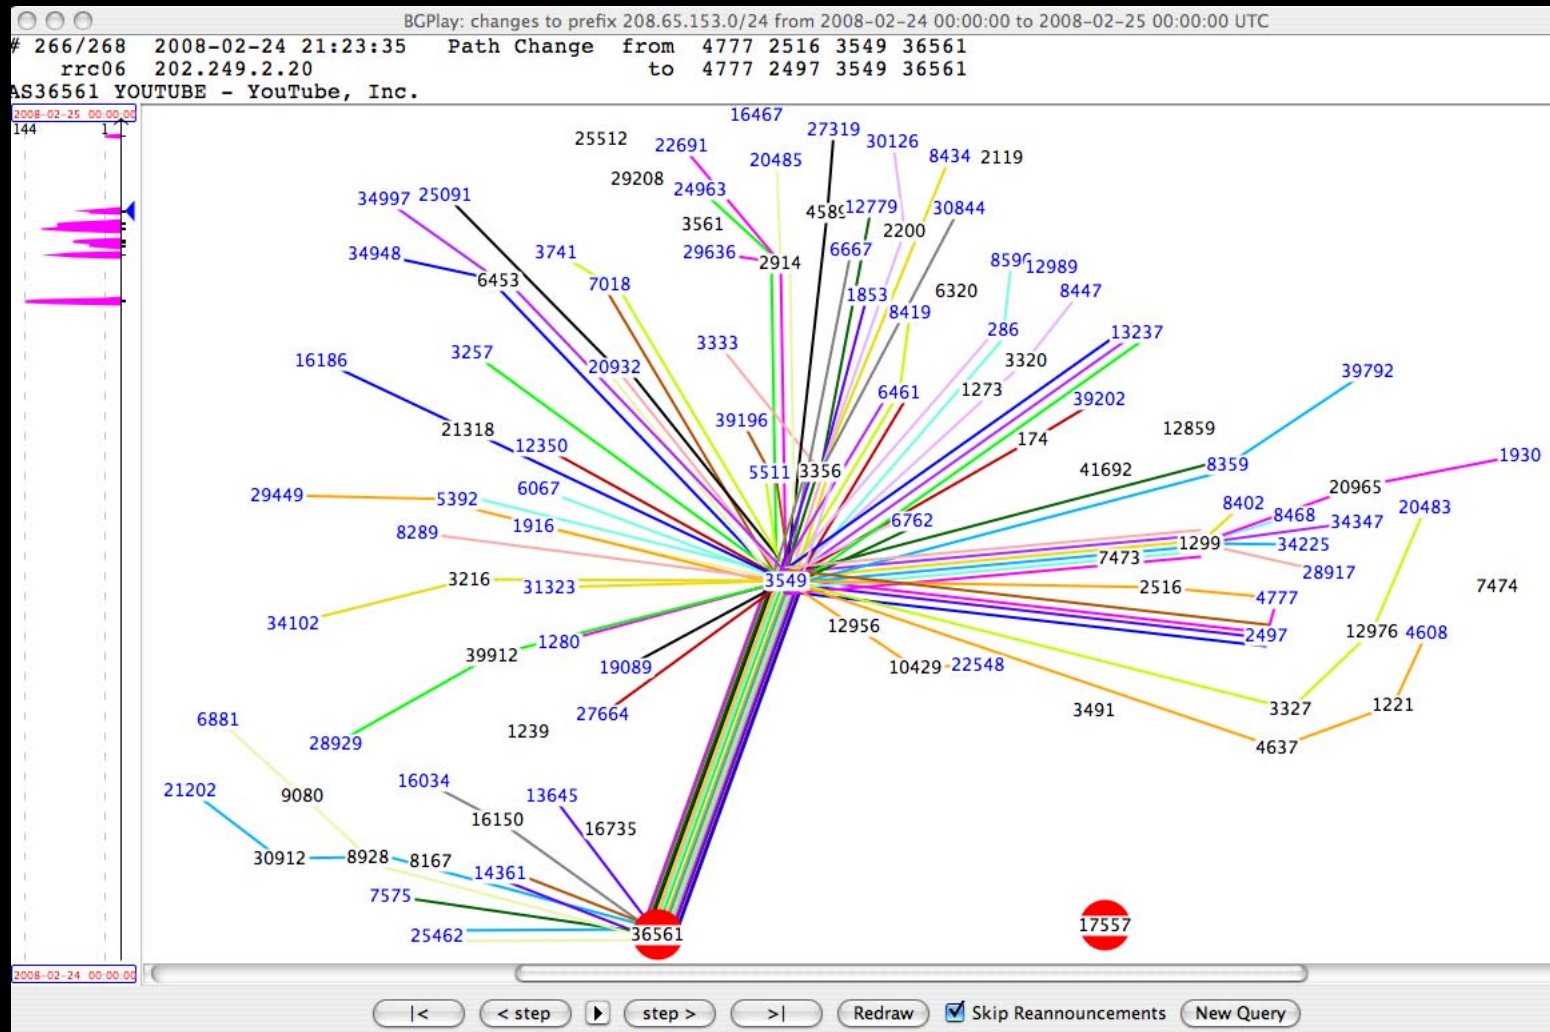
- On Sunday, 24 feb 2008, Pakistan Telecom (AS17557) started an **unauthorised announcement** of the prefix 208.65.153.0/24.
- One of PT providers, PCCW Global (AS3491) forwarded this announcement to the rest of the Internet, which resulted in the **hijacking of YouTube traffic** on a global scale

BGP blackholing technique that went off control?

Was it PT or PCCW fault?


<http://www.ripe.net/news/study-youtube-hijacking.html>

YouTube – it should look this way




Screenshots from RIPE's BGPlay

YouTube – it looked like this



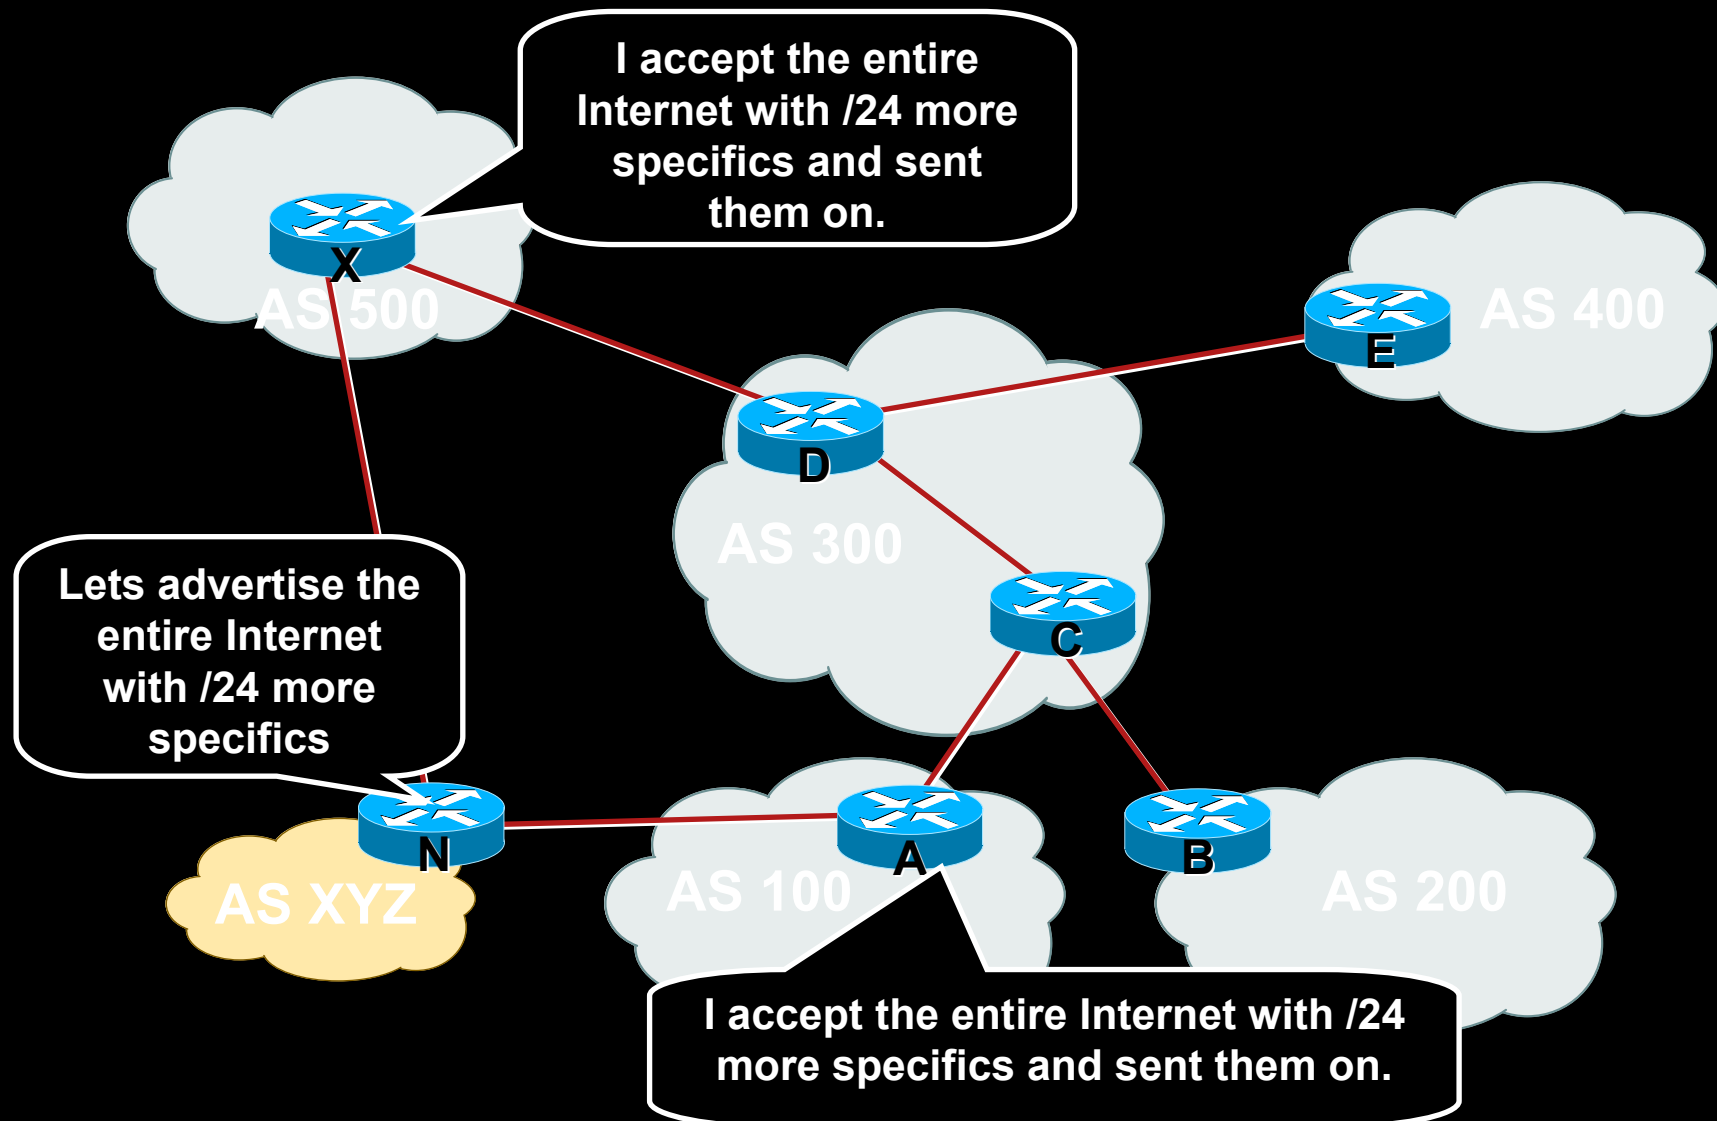
Screenshots from RIPE's BGPlay

Guarded Trust




- ISP A trust ISP B to send X prefixes from the Global Internet Route Table.
- ISP B Creates a egress filter to insure only X prefixes are sent to ISP A.
- ISP A creates a mirror image ingress filter to insure ISP B only sends X prefixes.
- ISP A's ingress filter reinforces ISP B's egress filter.

Garbage in – Garbage Out: What is it?



Garbage in – Garbage Out: Results



What went wrong?

- BGP is by design decentralized to scale to hundreds of thousands of prefixes...

...so any centralized tools are „artificial by design”

Synchronizing the efforts of hundreds of NOC engineers needs capable infrastructure (which is in place NSP-SEC)


- As always, people were the **weakest link** – from the operational and security perspective

How many of the PT and PCCW engineers previously attended „design” and „best practices” sessions just like You right now?

„OK, but here in Poland we're good, man”

- Not exactly, err, man
- Why I see such advertisements from my favourite SP?
Actually, from most of the SPs...


```
router# show ip bgp
* 192.168.1.0/24      x.x.x.x      0 xxxxx i
* 192.168.2.0/24      x.x.x.x      0 xxxxx i
* 10.10.1.0/24        x.x.x.x      0 xxxxx i
* 10.10.2.0/24        x.x.x.x      0 xxxxx i
* 10.40.0.0/16        x.x.x.x      0 xxxxx i
* 172.16.0.0/24       x.x.x.x      0 xxxxx i
* 172.16.9.0/24       x.x.x.x      0 xxxxx i
[...]
```



Let's do it by the book

Service Provider operational best practices 101

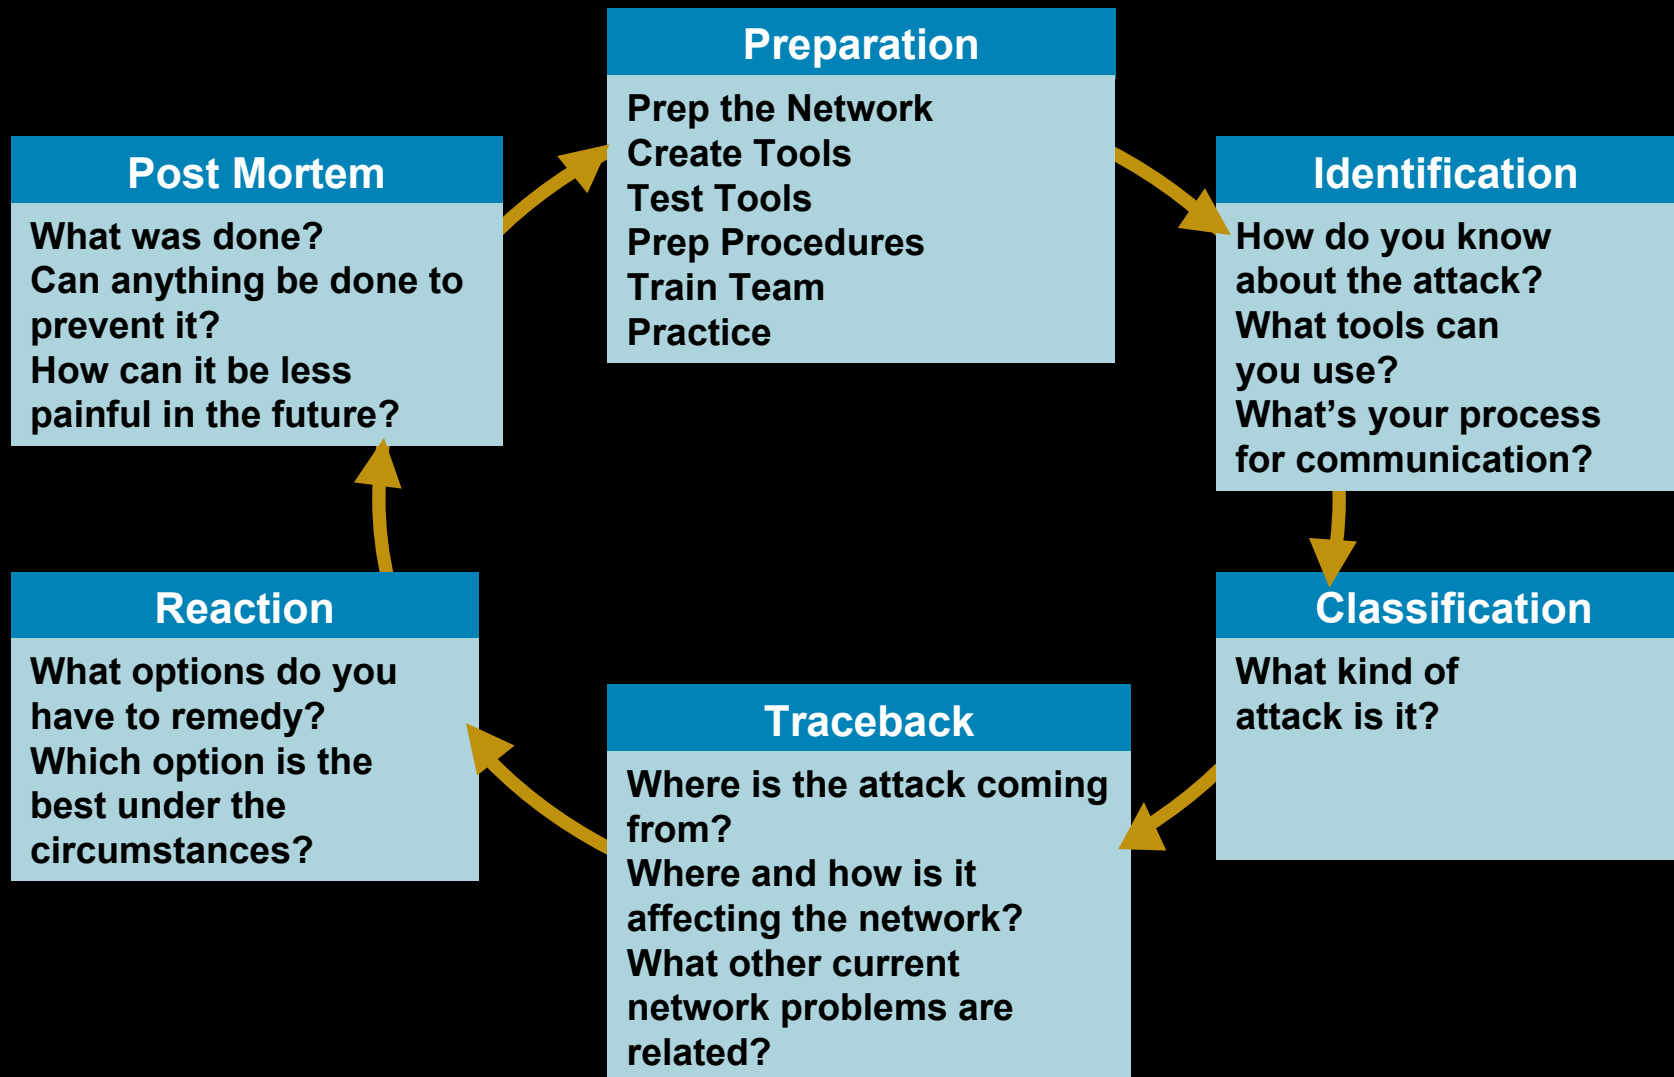
It is all about the packet...




Once a packet gets into the Internet, **some device, somewhere** has to do one of two things:

- **Deliver** the packet
- **Drop** the packet

Procedures in place...




First things first... the „old world” of SP



- Core routers individually secured
- Every router accessible from outside

First things first...the „new world” of SP



- Core routers individually secured PLUS
- Infrastructure protection
- Routers generally NOT accessible from outside

Then comes a loooooong list...

- Filter out the junk from network traffic on the edges

(L3) Packets with IP source and IP destination belonging to my own, reserved or not yet allocated address space, or clearly routed in via wrong interface (uRPF check)

(L3) Encrypt and authenticate your routing sessions

(L4) BGP prefixes announcements having typical errors:

- my own AS in AS_PATH (by default)
- looped AS in AS_PATH (by default)
- my own address space and address space of known „golden” networks (like DNS root zone servers, etc)
- tools exist to automatically build filter expressions based on actual RIR databases for transit providers

(L4) [...]

RFC3704/BCP84 Ingress Packet Filtering

- Packets should be sourced from valid, allocated address space, consistent with the topology and space allocation

Our goal here is to bind the problem and reduce the requirements for implementing security

- No BCP 84 means that:

Devices can (wittingly or unwittingly) send traffic with spoofed and/or randomly changing source addresses out to the network

Complicates trace back immensely


Sending bogus traffic is **not** free!

Attacks can be much more devious with spoofing

I could spoof You...

State of IP Spoofing

[\[news\]](#) [\[results\]](#) [\[methodology\]](#) [\[download\]](#) [\[feedback\]](#) [\[MailList\]](#) [\[FAQ\]](#)



* Spooftable and unspooftable counts represent *actual* client reports while indicated estimates are *extrapolated* from the number of globally routable netblocks, addresses and ASes respectively. Individual clients are counted singly regardless of the number of tests performed.

<http://spoofer.csail.mit.edu/summary.php>

BCP 84 Packet Filtering Principles

- Filter as close to the edge as possible
- Filter as precisely as possible
- Filter both source and destination where possible
- Can be implemented in various ways

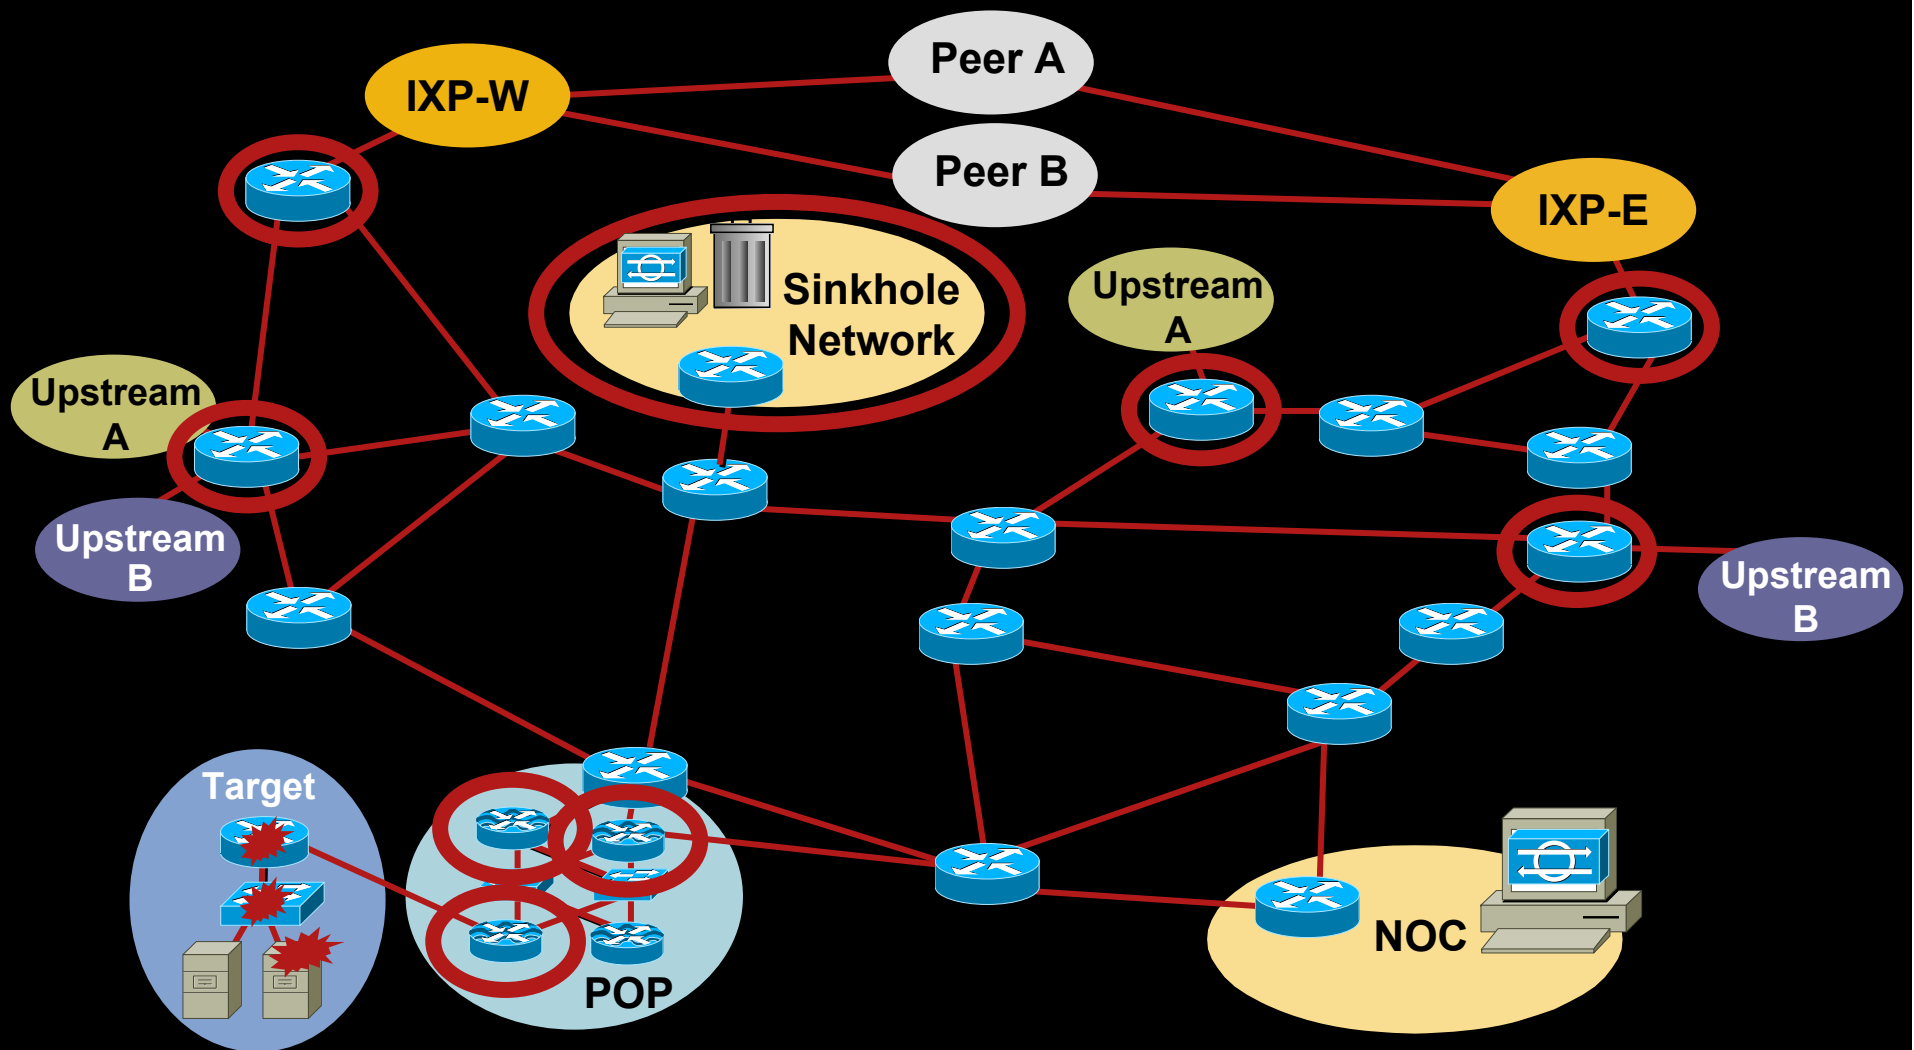
Infrastructure ACLs

unicast Reverse Path Forwarding

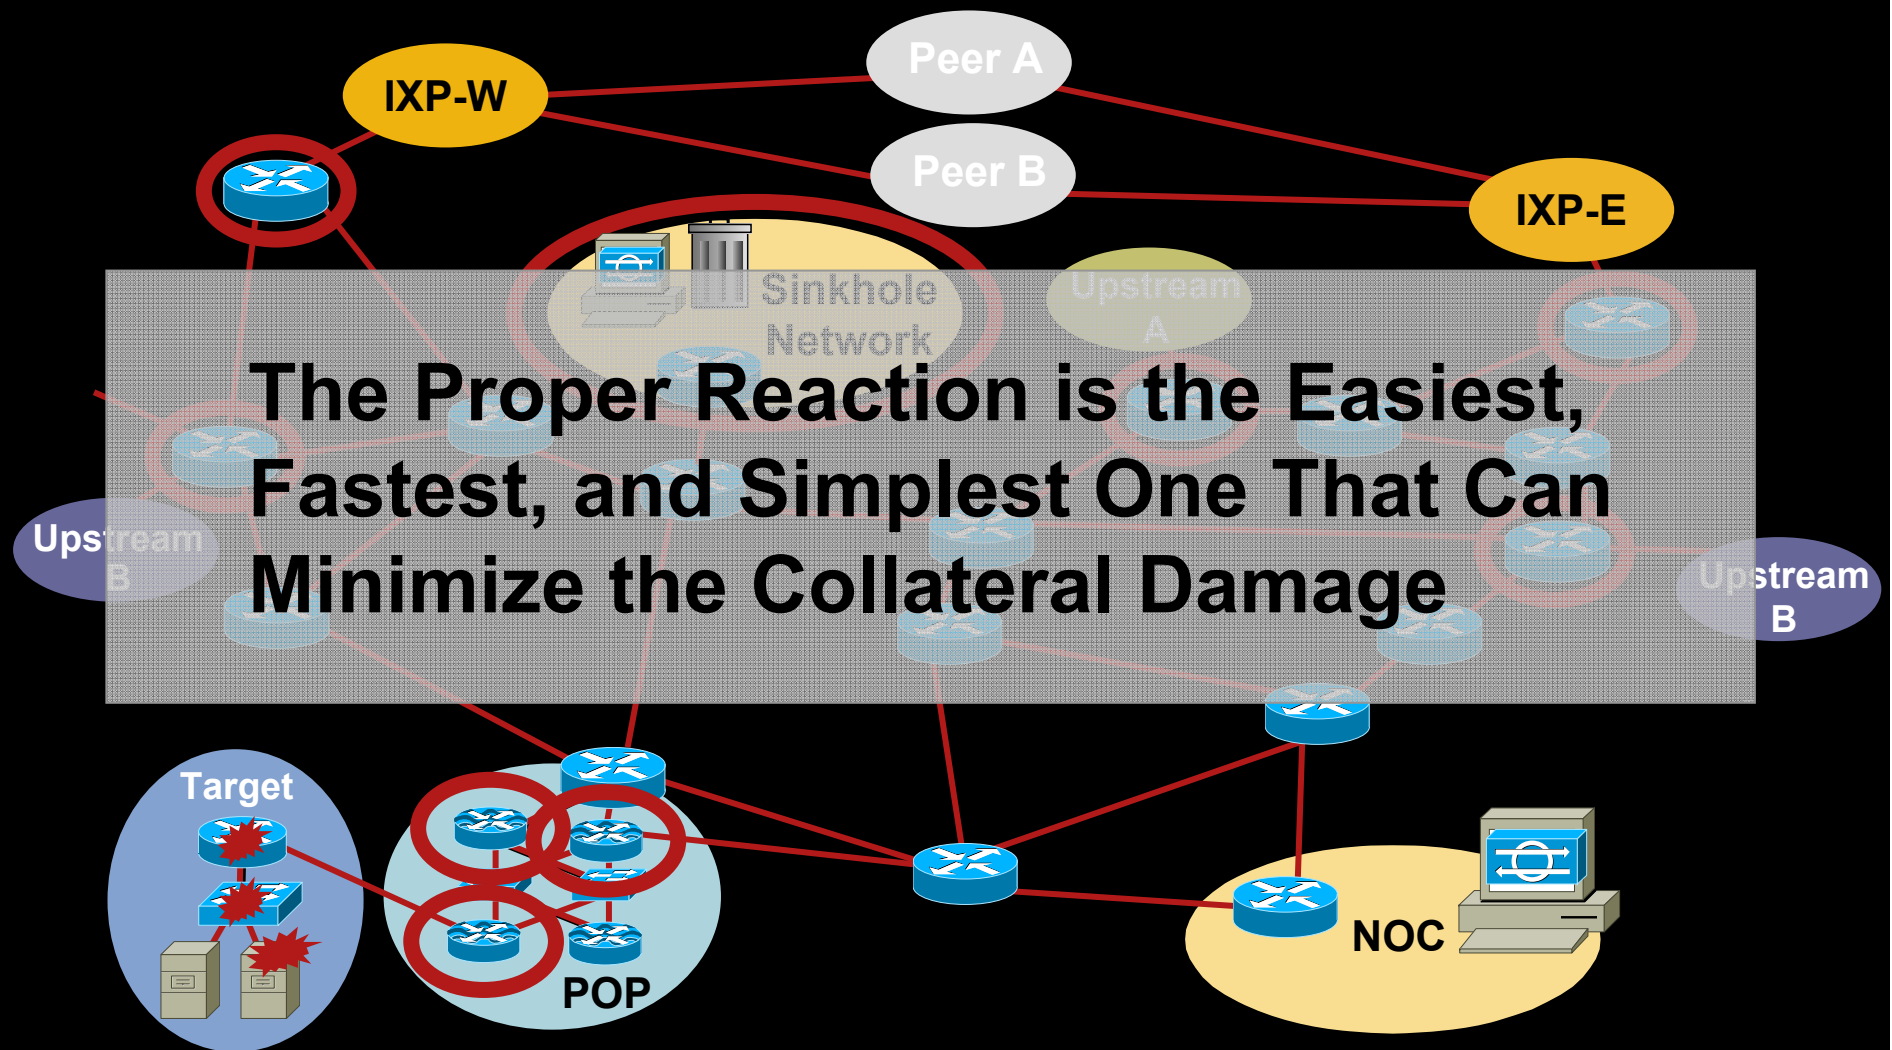
Cable source verify DHCP

IP source guard/DHCP snooping

Where to React?



Where to React?





Let's do it by the book

Service Provider security best practices 101

Take it to the next level

- NOT skipping the basic security measures
 - HARDENING the infrastructure
 - MAKING SURE clients can't reach Your network (why for?)
- Fighting the botnets (for DDoS, for e-crime, etc)
 - Blackholing (PL ☺, Cymru)
 - Anycast techniques
 - DNS blackholing, Borys Łacki & team
- Filtering the encrypted P2P traffic

Anycast is great...

- Two main uses:


- Distributed monitoring of traffic directed to greyspace – already assigned but not advertised anywhere as source of content/services

- high-performance workstation sniffing all the traffic, and then analyzing the logs and binary dumps

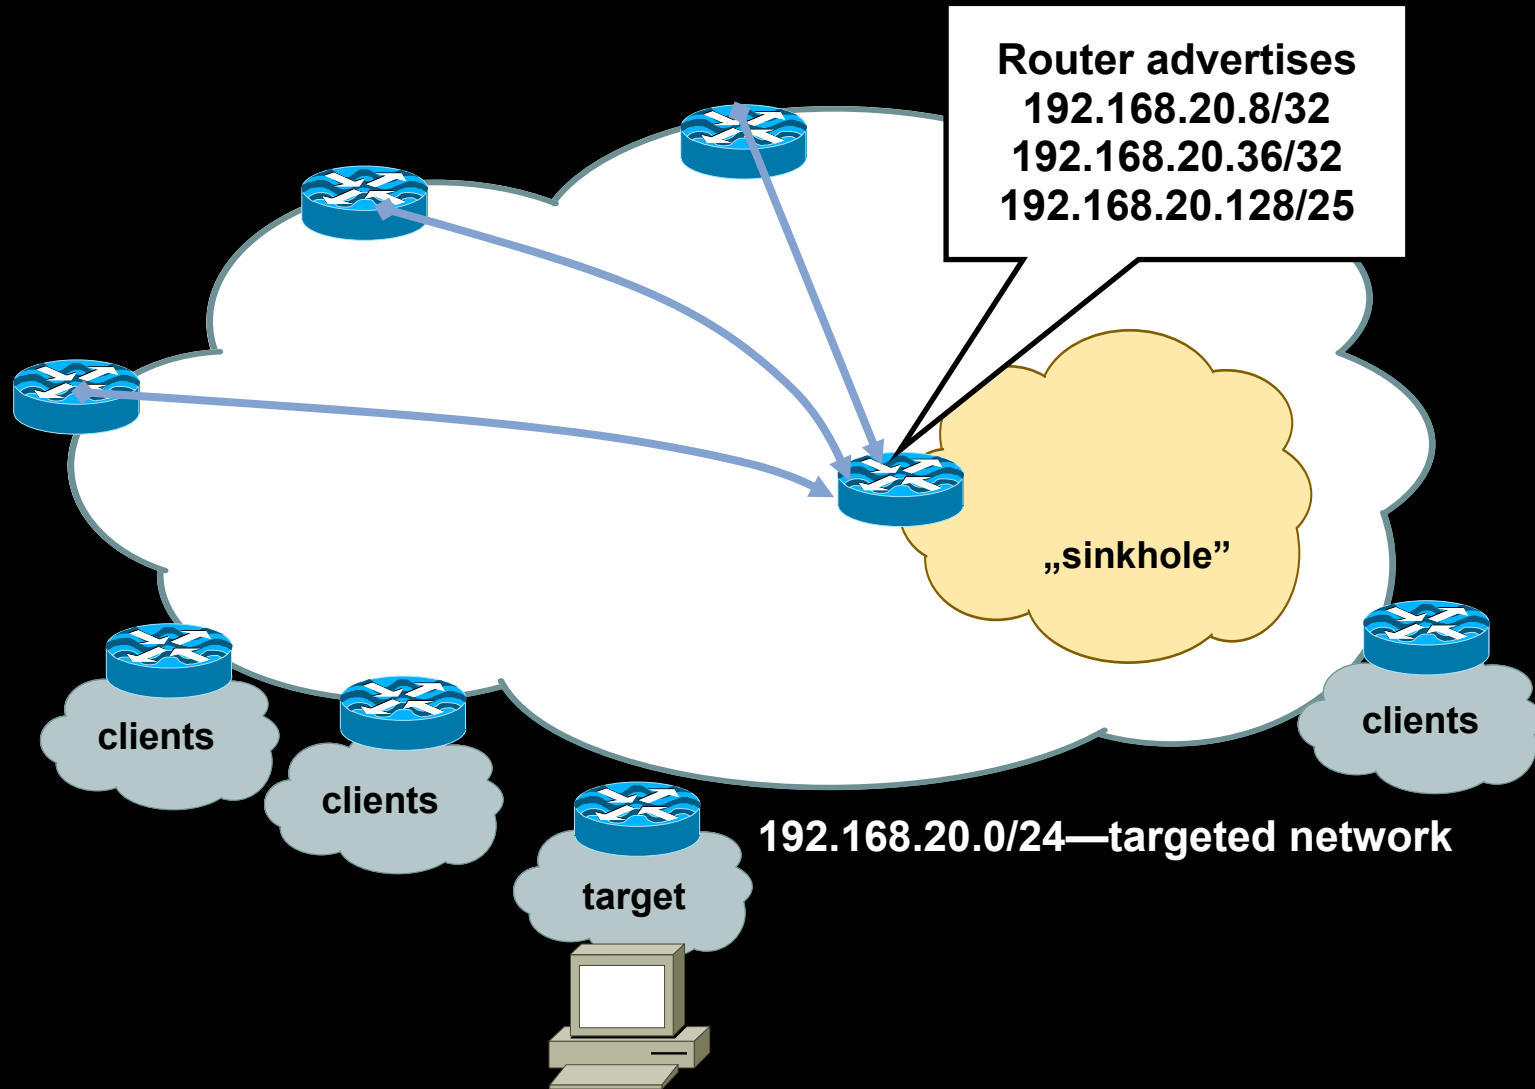
- Distribute the high traffic load to your services, to closest input point from perspective of your own network

- great for DNS for example, where one IP address (let's say – 192.168.10.5/32), can be assigned to a number of physical machines and then advertised simultaneously by routers in different places of the network – traffic by itself will stick to closest reachable host

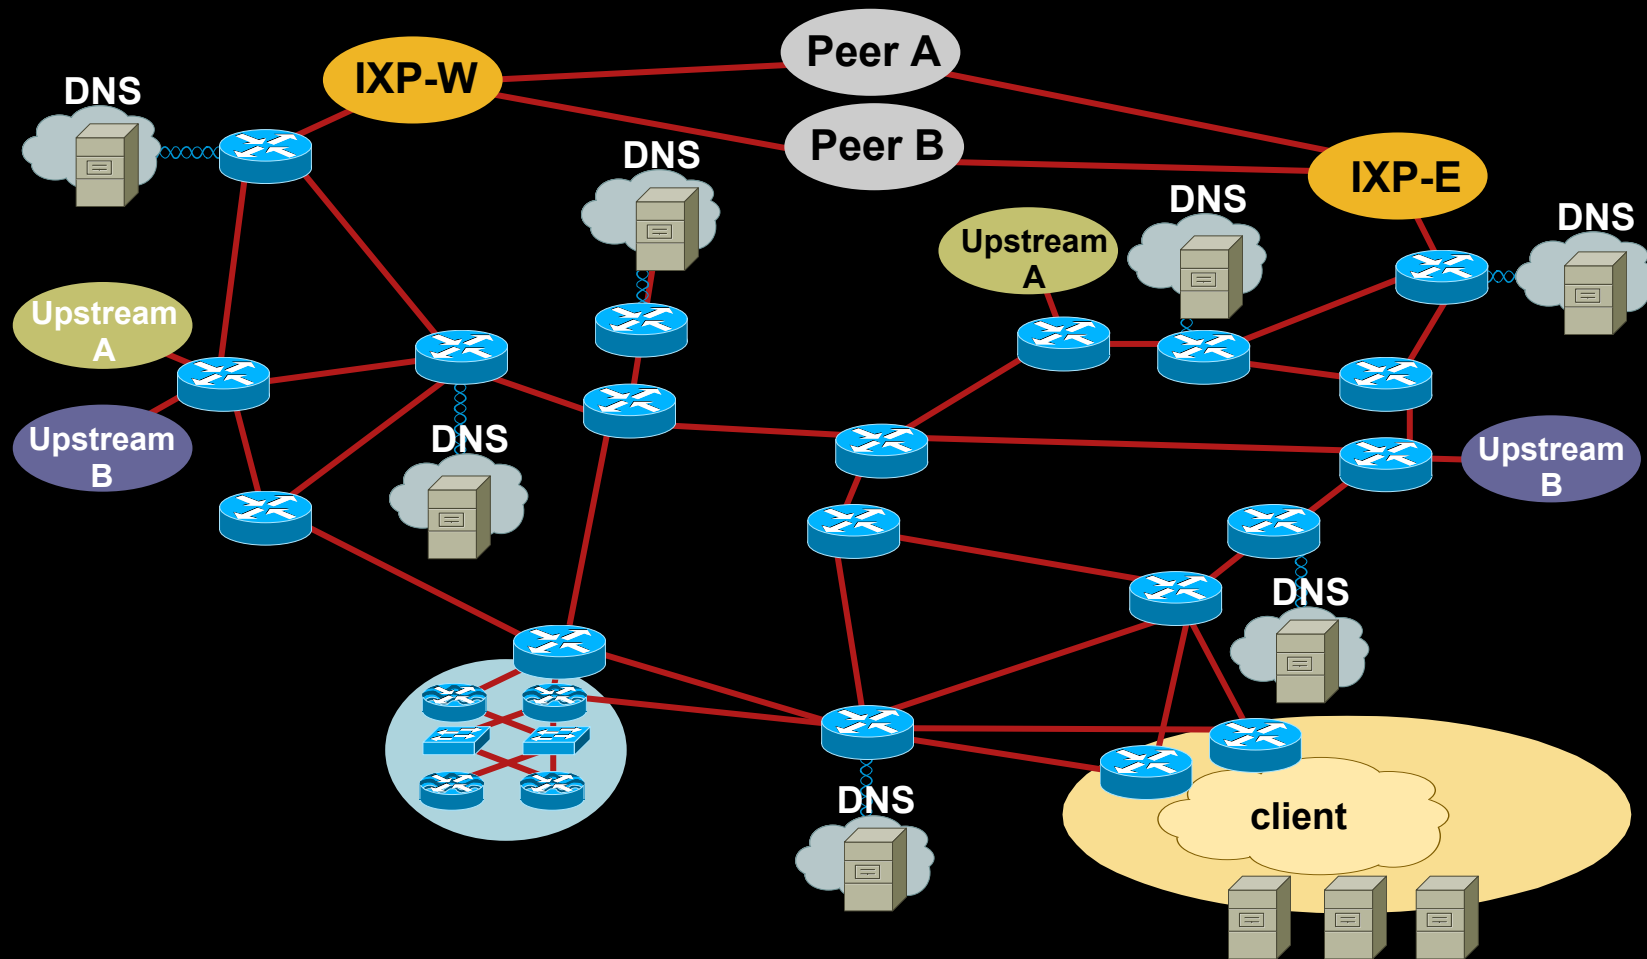
Blackholing/sinkholing – why for?




Anycast sinkhole – monitoring, analyzing



Anycast sinkhole – scaling



DDoS anyone? 1k hosts for 50\$?




Статистика

Total:	35001
Online:	2354

Новые за последние 2 часа:	244
Новые за последние 24 часа:	5238

Статистика по СТ	
Страна	Доступно за последние сутки/2 часа
AU	167/7
DE	43/0
GB	72/1
IT	293/1
NZ	7/0
ES	297/8
US	29183/1460
BG	5/0
DK	94/0
FR	41/3



Цены

Country	Price for 1k	
AU	300\$	Order now
DE	220\$	Order now
GB	210\$	Order now
IT	200\$	Order now
NZ	200\$	Order now
ES	200\$	Order now
US	110\$	Order now
BG	100\$	Order now
DK	100\$	Order now
FR	100\$	Order now
PT	100\$	Order now
NL	100\$	Order now
CA	80\$	Order now
JP	80\$	Order now
SE	70\$	Order now
BR	60\$	Order now
TR	60\$	Order now
NO	50\$	Order now
---	---	---

6/0
100/2
52/5

p2p encryption

- More and more traffic is encrypted and has distributed nature instead of simple, easy to understand and to sniff, transactional traffic (i.e. HTTP, FTP, SMTP)

Comparison of anonymous networks

Network	Purpose	Nodes	TCP/UDP	Encryption	Bootstrapping	Connected IPs	Known IPs	Supernodes	Websites	Content Cache	Webpage
ANts P2P	anon-p2p	50-150	TCP	Point-to-Point & End-to-End (D-H + AES)	Webcache, IRC, Buddies	4-5	about 10	yes	yes	no	antisp2p.sf.net
Entropy	anon-p2p		TCP		seednode file	~ 20 in, ~ 20 out	~ 40	not needed	yes	yes	entropy.stop1984.com
Freenet 0.5	anon-net anon-p2p	~1500	TCP	yes	seednode file	~ 200	many	not needed	yes	yes	freenetproject.org
Freenet 0.7	anon-net anon-p2p	~400	UDP	yes	IRC, Buddies			not needed	yes	yes	freenetproject.org
GNUnet	anon-p2p	150+	TCP/UDP	Point-to-Point & End-to-End (RSA + AES)	http or manual	10-80	350+	no needed	no	yes	gnunet.org
I2P	anon-net anon-p2p	700+	UDP	Point-to-Point & End-to-End (D-H + AES)	Net DB	~ 50	~ 800+	not needed	yes	no	i2p.de
MUTE	anon-p2p	300+	TCP	Point-to-Point (RSA + AES)	Webcache	1-10	~ 500+	not needed	no	no	mute-net.sourceforge.net
RShare	anon-p2p		TCP	Point-to-Point (RSA + Rijndael)	WebCaches	10+	200+	not needed	no	no	rshare.de
SUMI	anon-p2p		UDP								sumi.berlios.de
Tor	anon-net anon-p2p		TCP								tor.eff.org

p2p encryption

- SPs can be required to comply with monitoring the traffic due to international copyright laws or signed contracts
 - „Lawful intercept” is the keyword
- Not a really problem for transit traffic unless there's no infrastructure to actually sniff the traffic
- But if it happens on your edge – why not take control of everything You can't identify?

Stateless QoS mapping – in most of the cases useless

Stateful QoS classification/mapping to traffic classes

p2p encryption – possible solution

- Firewalls [CBR03], packet filters, intrusion detection systems, and the like often **have difficulty distinguishing between packets that have malicious intent and those that are merely unusual**. The problem is that making such determinations is hard. To solve this problem, we define a security flag, known as the "evil" bit, in the IPv4 [RFC791] header. **Friendly packets have this bit set to 0; those that are used for an attack will have the bit set to 1.**

RFC 3514 - The Security Flag in the IPv4 Header
<http://www.ietf.org/rfc/rfc3514.txt>



How to learn, to think outside the box...

Let me recommend

- ISP Essentials:

<ftp://ftp-eng.cisco.com/cons/isp/essentials/>

- ISP Security Essentials (NANOG):

<http://www.nanog.org/ispsecurity.html>









- Philip Smith presentations

<ftp://ftp-eng.cisco.com/pfs/seminars/>

- Packet Clearing House

<http://www.pch.net>

And the printed books...



Questions?

