



# Prefix Filtering, Black Holes, and Protecting Your Business

David Barak [thegameiam@yahoo.com](mailto:thegameiam@yahoo.com)

Barry Raveendran Greene [bgreene@senki.org](mailto:bgreene@senki.org)

Mark Prior [mprior@juniper.net](mailto:mprior@juniper.net)



# Free Use

- This slide deck can be used by any operator to help empower their teams, teach their staff, or work with their customers.
- It is part of the next generation of **NANOG Security Curriculum** .... providing tools that can improve the quality of the Internet.



# Learning Objectives




- Business Impact to Poor Prefix Management - Hijack & Processor Attacks/Collateral Damage
- Know your Network – Whys to understand and manage the prefixes you use in your network.
- Principles of Prefix Filtering to protect the Business
- Approaches to Prefix Filtering
- What should your Business Practice & how provide your service?



# Outline

- Why are we talking about this topic again?
- Threat & Problem
- Know Your Network
- Have Your Own Prefix Database (IRR)
- Principles of Prefix Filtering
- Enables Capabilities

Why are we talking about this topic again?





# This is yet another Tutorial .....

- [Tutorial: Routing Policy Specification Language/IRR](#), by Cengiz Alaettinoglu, ISI, and Gerald Winters, Merit. NANOG 17, Oct. 1999.
- [BGP Configuration From the IRR](#) (Tutorial), by Cengiz Alaettinoglu, ISI. NANOG 19, June 2000.
- [Routing Policy Implementation Guide](#) (Tutorial), by Dan Golding, Sockeye. NANOG 24, February 2002.
- [Introduction to RPSL](#) (Tutorial), by Ambrose Magee, Ericsson. NANOG 25, June 2002.
- [IRR tutorial](#) [Nurani Nimpuno](#) & [Champika Wijayatunga](#) APNIC15 2003
- [IRR](#) - Practical use of RPSL and IRR tools [Andy Linton](#) APNIC15 2003



## Of a long list of sessions ....

- [Routing Policy System Status](#), by Curtis Villamizar, ANS. NANOG 13, June 1998.
- [Research Forum: Nemecis: A Tool to Analyze the IRR Registries](#), by Georgos Siganos, UC Riverside. NANOG 30, February 2004.
- [RPSLng Status Update](#), by Larry Blunk, Merit. NANOG 32, October 2004.



# With very details materials ...

- RIPE IRR 1 Day Training Program:
  - <http://www.ripe.net/training/rr/index.html>



# Pakistan and YouTube

## Pakistan Blocks YouTube Video Access

February 24, 2008 12:17 PM PST

YouTube blames Pakistan outage

Posted by Greg Sandoval

Updated, 9:40 p.m. to add YouTube's explanation outage.

YouTube suffered a two-hour long, system-wide outage the company said was triggered by a network bas

SADAQAT JAN | February 24, 2008 09:04 AM EST | **AP**

Read More: [Pakistan](#), [Pakistan Blocks Youtube](#), [Pakistan Elections](#), [Pakistan Youtube](#), [Pervez Musharraf](#), [Youtube](#), [Youtube Pakistan Anti-Islamic Movies](#), [Breaking Politics News](#)



Email ▶  
Print ▶  
Comments ▶




ISLAMABAD, Pakistan — Pakistan's government has banned access to the video-sharing Web site YouTube because of anti-Islamic movies that users have posted on the site, an official said Sunday.

The Pakistan Telecommunication Authority told the country's 70 Internet service providers Friday that the popular Web site would be blocked until further notice.

The authority did not specify what the offensive material was, but a PTA official said the ban concerned a movie trailer for an upcoming film by

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



# Our Approach

- Drive Simple Recommendations
- Collect the wisdom of the materials already presented.
- Make sure everyone is clear about the consequences of in action.
- Internet Draft which could express “BCP.”

# Business Impact to Poor Prefix Policies

1



# Malicious Route Injection

## *Perceive Threat*

- Bad Routing Information does leak out. This has been from mistakes, failures, bugs, and intentional.
- Intruders are beginning to understand that privileged access to a router means route tables can be altered
- CERT/CC is aware of a small number of incidents involving malicious use of routing information.
- Perceived Threat is that this will be a growth area for attackers.

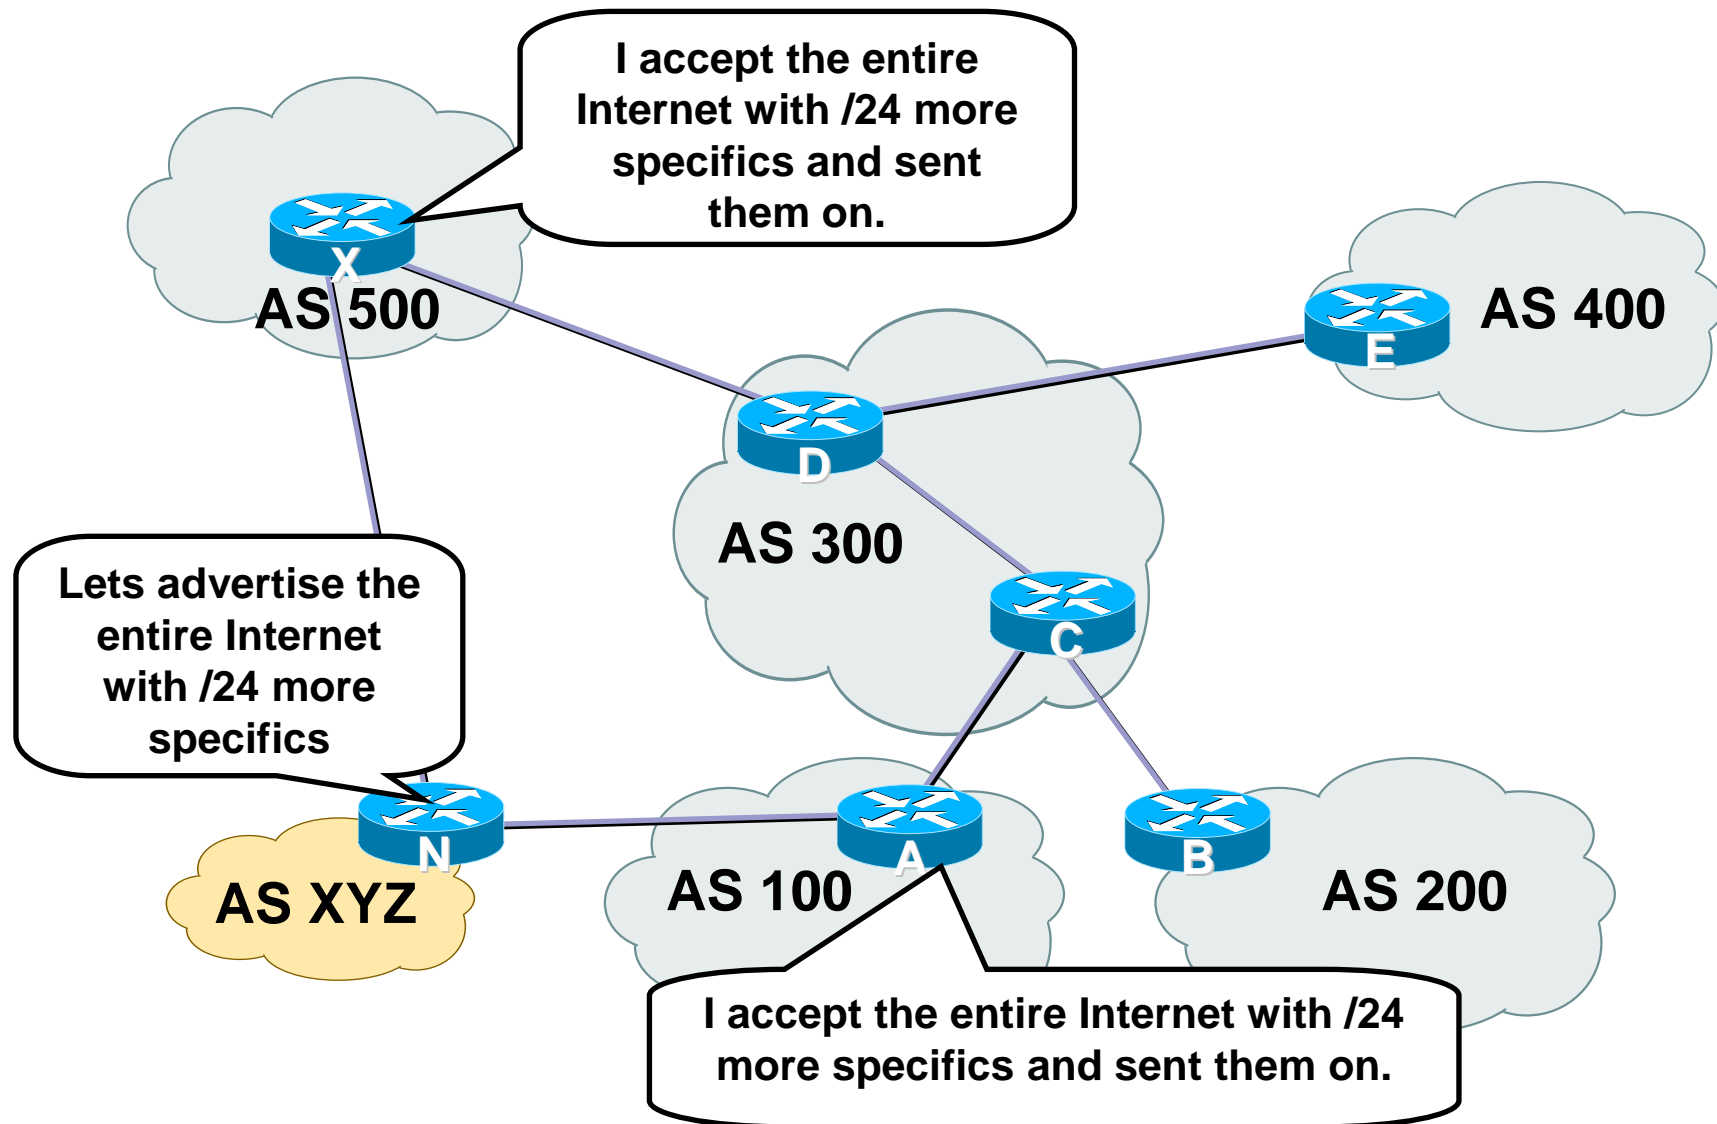


# Malicious Route Injection

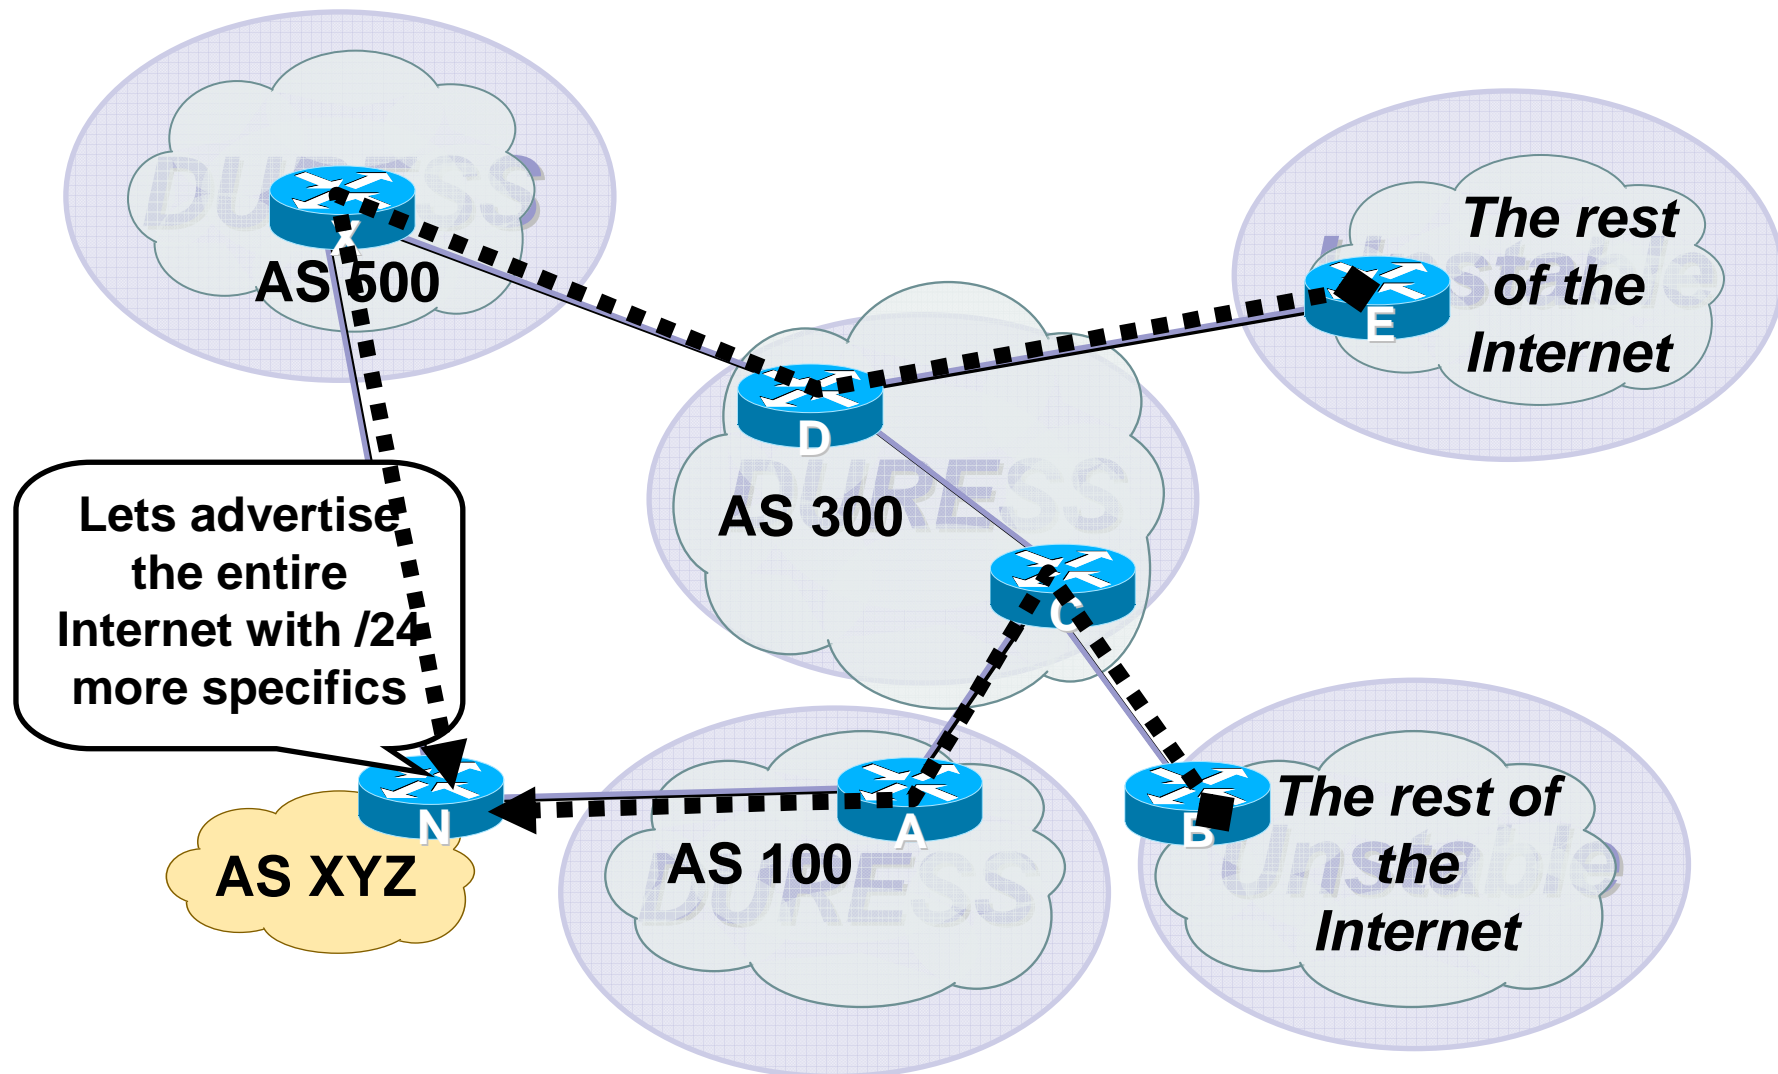
## *Reality – an Example*

- AS 7007 incident used as an attack.
- Multihomed CPE router is violated and used to “de-aggregate” large blocks of the Internet.
- Evidence collected by several CERTs that hundreds of CPEs are violated.

# Garbage in – Garbage Out: What is it?




# Garbage in – Garbage Out: Results



# Garbage in – Garbage Out: Impact


- Garbage in – Garbage out does happen on the Net
- AS 7007 Incident (1997) was the most visible case of this problem.
- Key damage are to those ISPs who pass on the garbage.
- Disruption, Duress, and Instability has been an Internet wide effect of Garbage in – Garbage out.





# Garbage in – Garbage Out: What to do?

- Take care of your own Network.
  - Filter your customers
  - Filter your advertisements
- Net Police Filtering
  - Mitigate the impact when it happens
- Prefix Filtering and Max Prefix Limits





# Malicious Route Injection

## *Attack Methods*

- Good News – Risk is mainly to BGP speaking Routers.
- Bad News – Multihomed BGP Speaking customers are increasing!
- Really Bad News – Many of these routers have no passwords!
- Local layer 3 configuration alteration on compromised router
- Intra-AS propagation of bad routing information
- Inter-AS propagation of bad routing information

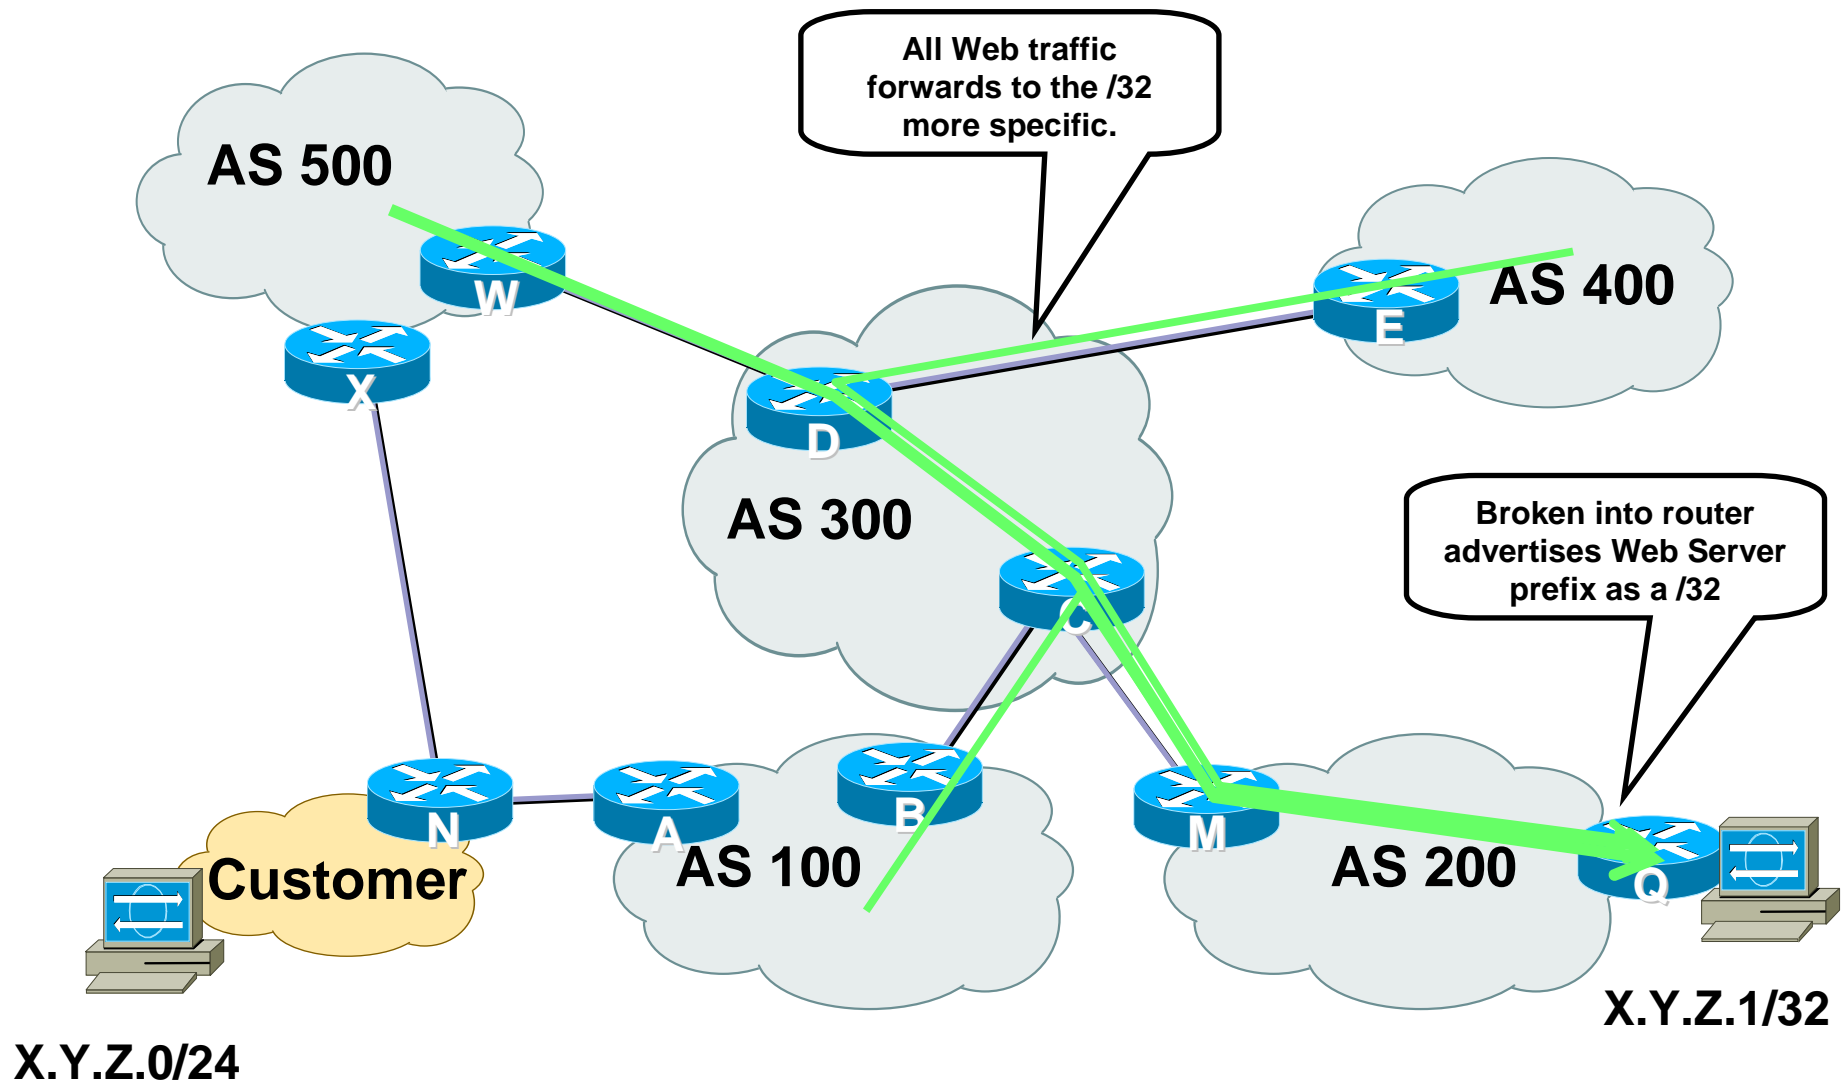


# Malicious Route Injection

## *Impact*


- Denial-Of-Service to Customer(s), ISP(s), and the Internet.
- Traffic Redirection / Interception
- Prefix Hijacking
- AS Hijacking

# What is a prefix hijack?



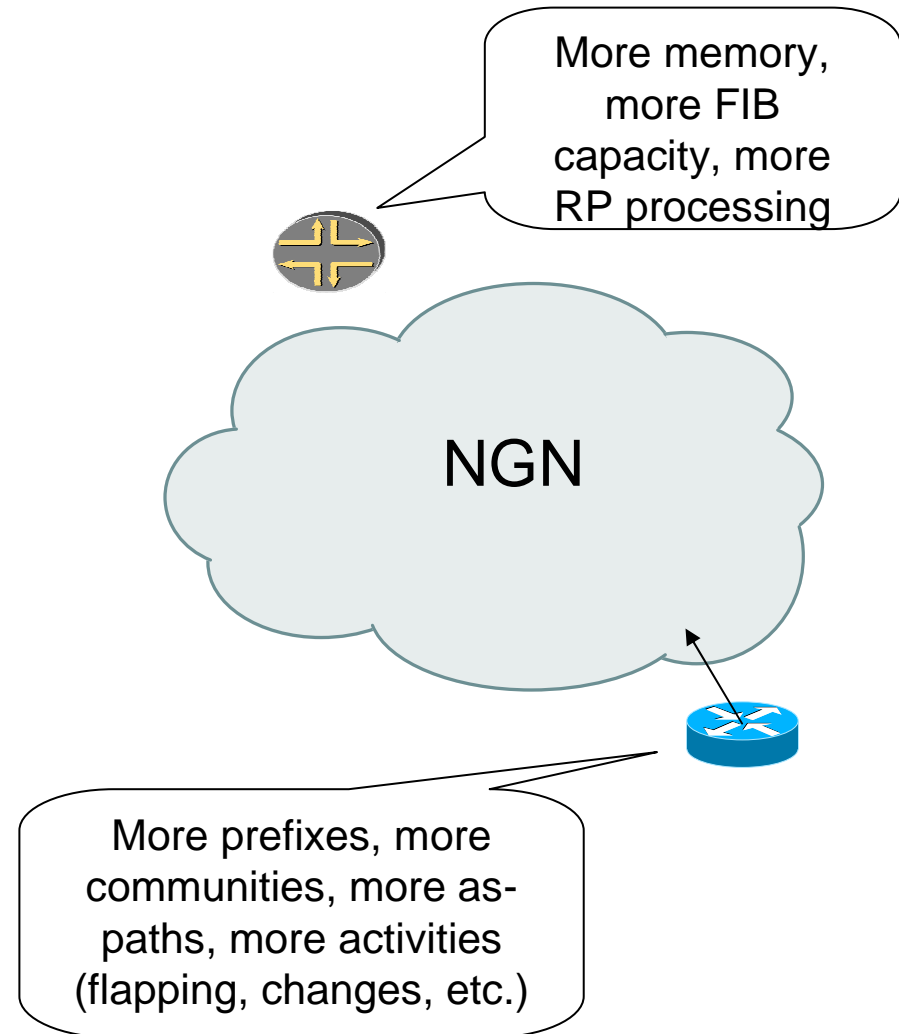
# What could be worse?

- The Miscreant Economy Trades violated “BGP Speaking” routers. Get 20 in different parts of the Internet.
- Take each, pick your targets, and start disaggregating.



# Why?

- Today's (and tomorrow's) NGN will be different from the past
- A business on one side of the planet will force you into OPEX and CAPEX expenditure!



Know Your Network

2



## If you run a network ....

- You should know which routes you've been allocated to route.
- You should know which routes your customer has contracted you to route.
- This information should be kept in a central place to be used by the company.



Have Your Own  
Prefix Database  
(IRR)

3



# Why have multiple route DBs?

- If you have to keep everything in one routing DB, why not use one that all parts of the organization can tap into, other SPs tap into, and used to maintain the routing policies of your network?

- ☐ [www.ird.net](http://www.ird.net)

- ☐ <http://www.ripe.net/db/cvs-bugzilla.html>

- ☐ <http://www.isc.org/index.pl>



## Keeping it simple ...

- Run your own whois daemon and keep control
  - IRRD
  - RIPE whois
- RPSL can be used to describe peering relationships in extreme detail
- But we don't want to scare people (too much :-)
- Just need one object type to start ...



# Route object

route: key]	[mandatory]	[single]	[primary/look-up
descr:	[mandatory]	[multiple]	[ ]
origin: key]	[mandatory]	[single]	[primary/inverse
holes:	[optional]	[multiple]	[ ]
country:	[optional]	[single]	[ ]
member-of:	[optional]	[multiple]	[ ]
inject:	[optional]	[multiple]	[ ]
aggr-mtd:	[optional]	[single]	[ ]
aggr-bndry:	[optional]	[single]	[ ]
export-comps:	[optional]	[single]	[ ]
components:	[optional]	[single]	[ ]
remarks:	[optional]	[multiple]	[ ]
notify:	[optional]	[multiple]	[inverse key]
mnt-lower:	[optional]	[multiple]	[inverse key]
mnt-routes:	[optional]	[multiple]	[inverse key]
mnt-by:	[mandatory]	[multiple]	[inverse key]
changed:	[mandatory]	[multiple]	[ ]
source:	[mandatory]	[single]	[ ]



# Example route object

```
route:          192.189.54.0/24
descr:          connect.com.au Pty Ltd
                (0)
origin:         AS2764
notify:         routing@connect.com.au
mnt-by:         MAINT-AS2764
changed:        nobody@connect.com.au
                19970923
source:         RADB
```



# Software to play with route objects

- Some people roll their own but let's be simple
- IRRToolSet (looked after by ISC)
  - Supports talking to IRRd & RIPE whois servers
  - RtConfig is normally used to write code fragments



# Policy fragment (Junos)

```
policy-statement as6939-ipv4-import {  
    term as6939 {  
        from policy rs-as6939;  
        then {  
            local-preference 90;  
            community add peers;  
            next policy;  
        }  
    }  
    term reject {  
        then reject;  
    }  
}
```



## Policy fragment (Junos) [2]

```
policy-statement rs-as6939 {  
    term prefixes {  
        from {  
        }  
        then accept;  
    }  
    then reject;  
}
```





## RtConfig fragment (Junos)

```
policy-statement rs-as6939 {  
    replace:  
    term prefixes {  
        from {  
@RtConfig printPrefixRanges "\t\troute-filter %p/%l  
    upto /24;\n" filter AS6939 AND NOT { 0.0.0.0/0^25-  
    32 } AND NOT fltr-martian  
        }  
        then accept;  
    }  
}
```

# Principles of Prefix Filtering to Protect the Business

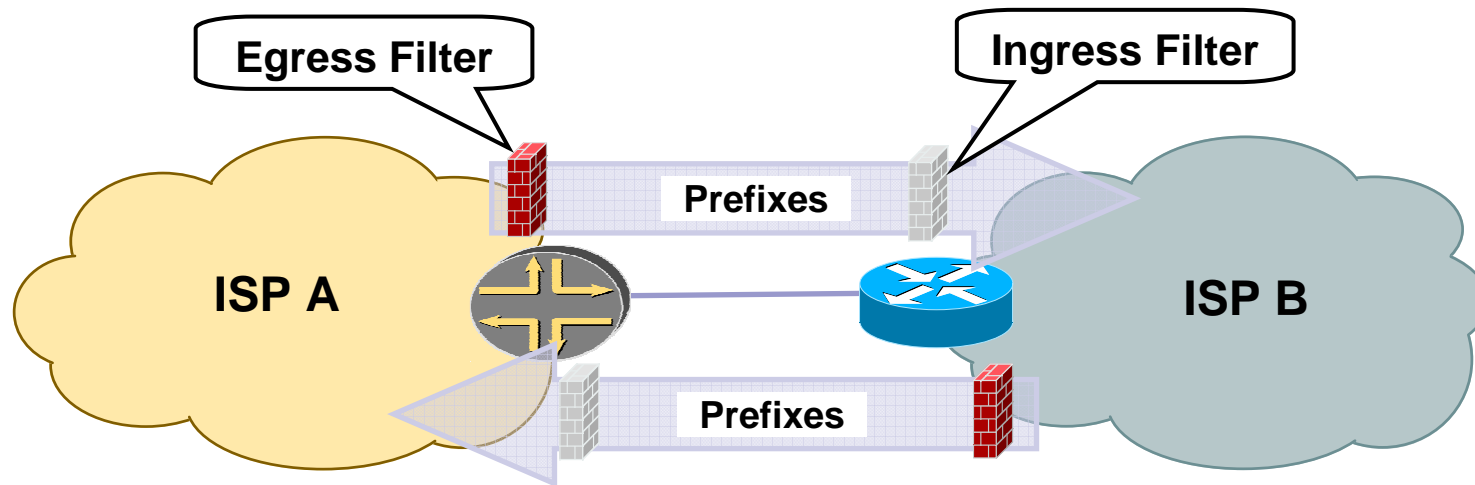
4



# BGP Peering Fundamentals

- BGP Peering assumes that something could go wrong with the policy filters between the neighboring routers.
- Filters are all created to mutually reinforce each other. If one policy filter fails, the policy filter on the neighboring router will take over – providing redundancy to the policy filters.
- This mutually reinforcement concept used BGP peering filters are created are also called guarded trust, mutual suspicion, or Murphy Filtering.

# Guarded Trust



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



# Malicious Route Injection

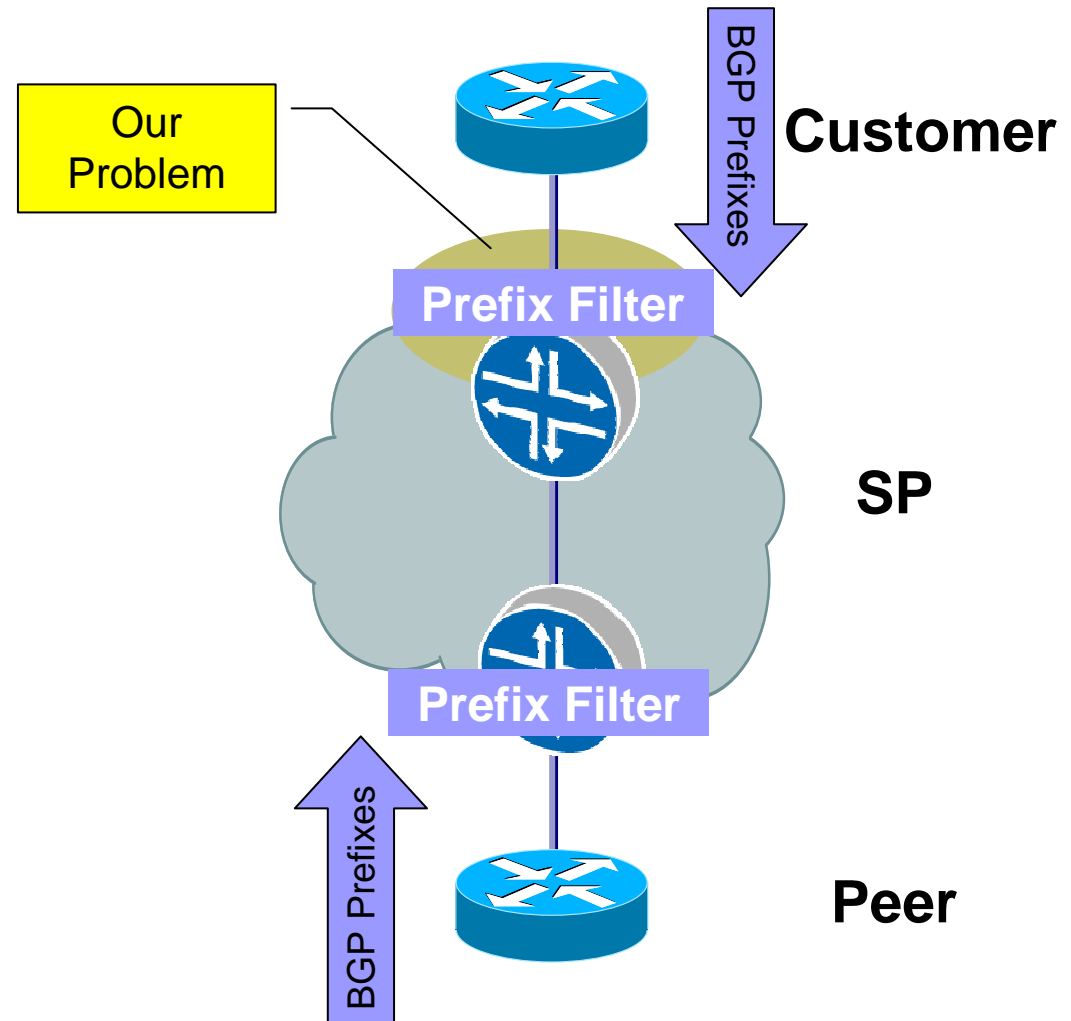
## *What can SPs Do?*

- Know your network – What to filter, where to filter.
- Customer Ingress Prefix Filtering!
- SPs should only accept customer prefixes which have been assigned or allocated to their downstream customers.
- For example
  - Downstream customer has 220.50.0.0/20 block.
  - Customer should only announce this to peers.
  - Upstream peers should only accept this prefix.

# Prefix Filters: In

## Apply Prefix Filters to All eBGP Neighbors

- From Customers
- From Peers & Upstreams





# Malicious Route Injection

## *What can ISPs Do?*

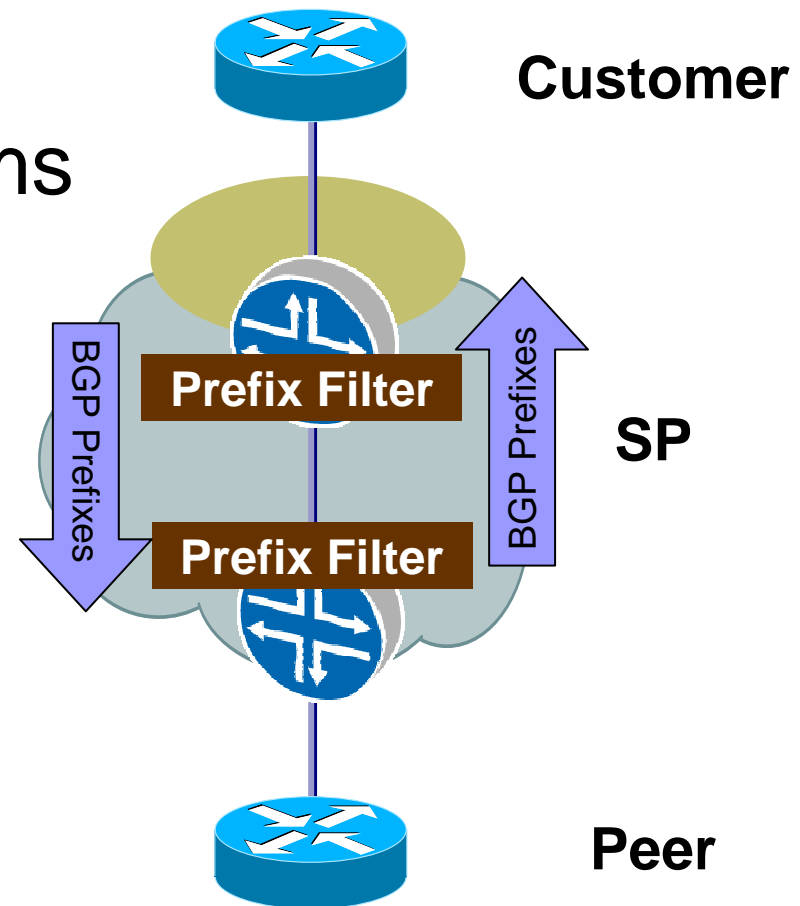
### ■ Containment Filters!

- Design your network with the principles of of survivability.
- Murphy's Law of Networking implies that the customer ingress prefix filter will fail.
- Remember 70% to 80% of ISP problems are maintenance injected trouble (MIT).
- Place Egress Prefix Filters on the Network to contain prefix leaks.

# Prefix Filters: Out

**Apply Prefix Filters to All eBGP Neighbors**

- To Customers
- To Peers & Upstreams





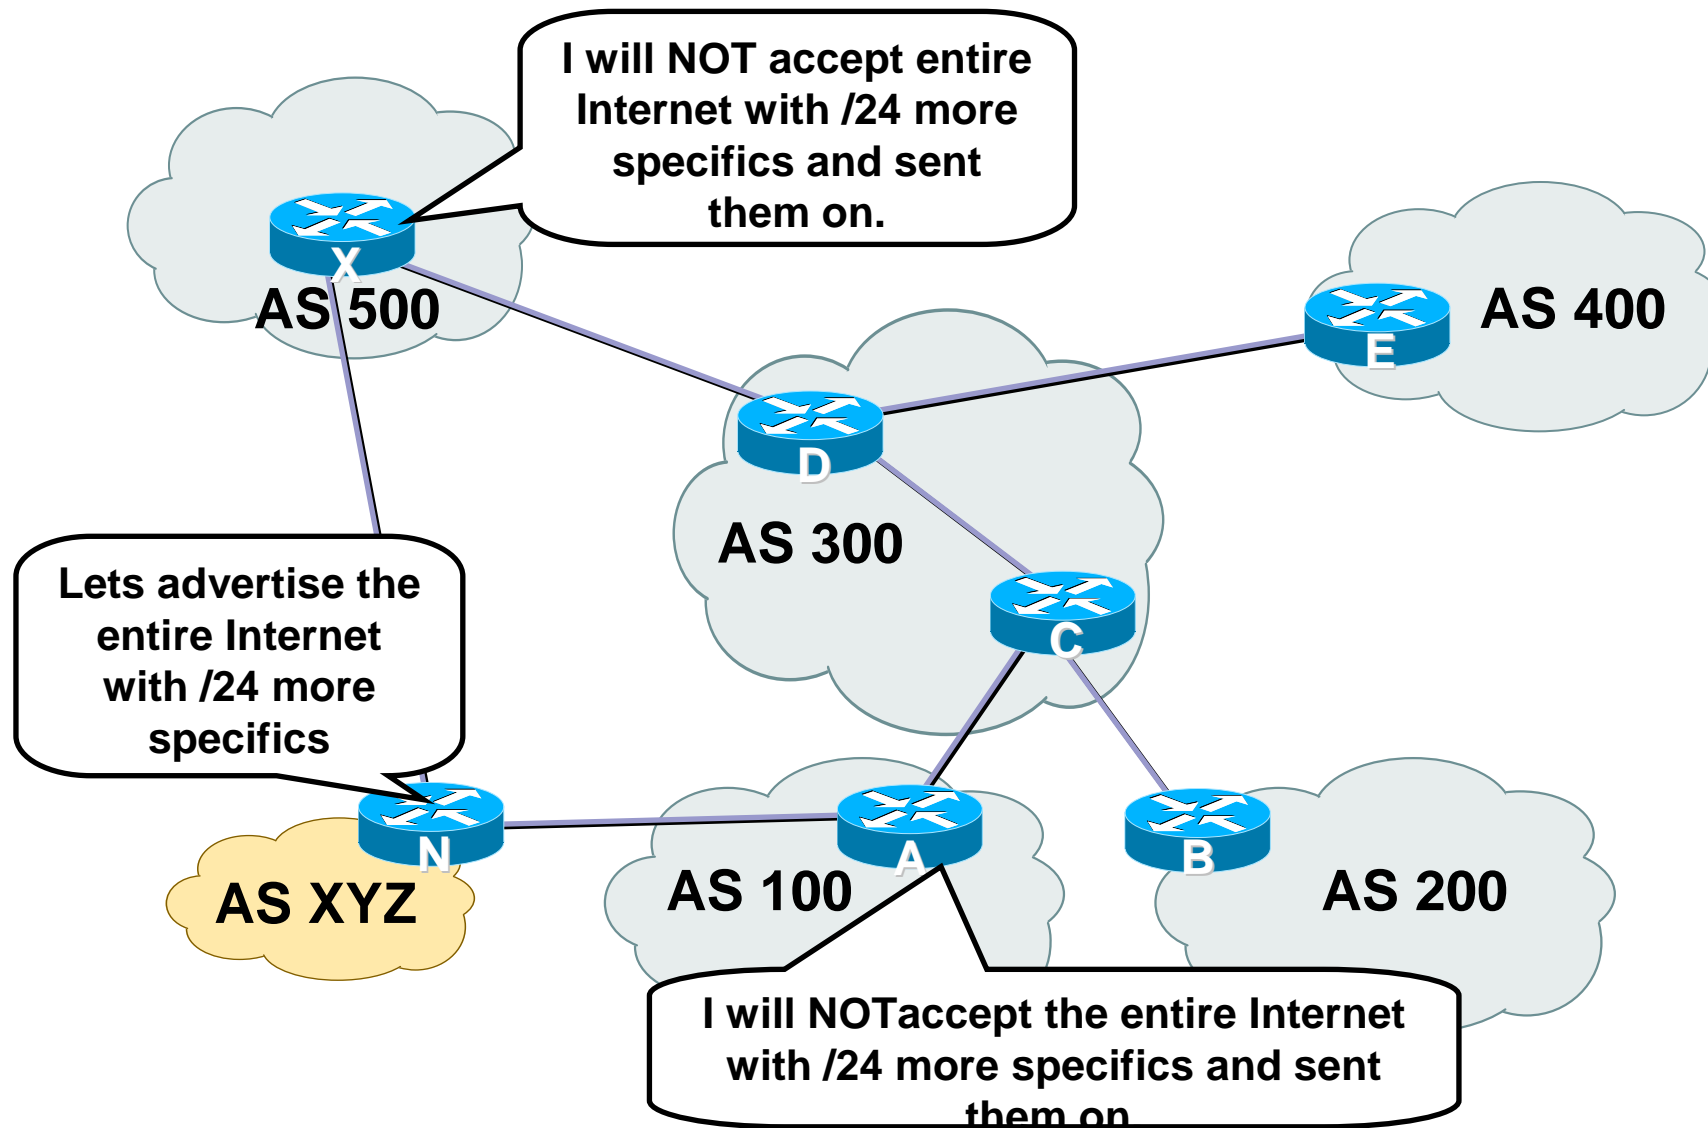


## *What can ISPs Do?*

# Containment Egress Prefix Filters

- What about all my multihomed customers with prefixes from other ISPs?
- Add them to the customer ingress prefix filter.
  - You should know what you will accept.
- Add them to the master egress prefix-filter.
  - You should know what you're advertising to everyone else.
  - *Bigness* is not an excuse.

# Containment Filters





# Malicious Route Injection

## *What can ISPs Do?*

- Customer Ingress Prefix Filtering
- Prefix filtering between intra-AS trust zones
- Route table monitoring to detect alteration of critical route paths
- SPAMers are using route-hijacking.



# Summary

- Understand the risk
  - Take infrastructure protection into account in network design
- Want to deploy voice? Want to deploy video? Want to deploy xyz?
  - All services deployment depend on an available infrastructure
- Understand the techniques/features and apply them appropriately
  - Edge filters: iACLs
  - Control plane traffic filtering: rACL
  - Next-phase of control plane filtering (including policing): CoPP
- Each feature has pros/cons
  - Ultimately, mix and match as needed: remember defense in depth



# Summary

- Review your current protection schemes
  - Identify gaps and areas of exposure
  - Develop a plan for protection
- Next steps:
  1. Begin to classify network traffic
  2. Use classification data and platform mix to determine appropriate protection schemes
- Start planning your deployments!
  - Can be difficult but certainly worthwhile!
  - Many customers have widespread deployments and have seen the benefits

# What to Prefix Filter?



# Documenting Special Use Addresses (DUSA)

- There are routes that should NOT be routed on the Internet
  - RFC 1918 and “Martian” networks (DUSA)
  - 127.0.0.0/8 and multicast blocks (DUSA)
- **RFC 3330 Special-Use IPv4 Addresses**  
<http://tools.ietf.org/html/rfc3330>
- BGP should have filters applied so that these routes are not advertised to or propagated through the Internet



# Documenting Special Use Addresses (DUSA)

## ■ Quick review

- 0.0.0.0/8 and 0.0.0.0/32—Default and broadcast
- 14.0.0.0/8 - This block is set aside for assignments to the international system of Public Data Networks
- 24.0.0.0/8 - This block was allocated in early 1996 for use in provisioning IP service over cable television systems.
- 39.0.0.0/8 - This block was used in the "Class A Subnet Experiment" that commenced in May 1995, as documented in [\[RFC1797\]](#).
- 128.0.0.0/16 - This block, corresponding to the numerically lowest of the former Class B addresses, was initially and is still reserved by the IANA.
- 127.0.0.0/8—Host loopback
- 192.0.2.0/24—TEST-NET for documentation
- 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16—RFC 1918 private addresses
- 169.254.0.0/16—End node auto-config for DHCP





# Summary Table – RFC 3330

Address Block	Present Use	Reference
0.0.0.0/8	"This" Network	[RFC1700, page 4]
10.0.0.0/8	Private-Use Networks	[ <a href="#">RFC1918</a> ]
14.0.0.0/8	Public-Data Networks	[RFC1700, page 181]
24.0.0.0/8	Cable Television Networks	--
39.0.0.0/8	Reserved but subject to allocation	[ <a href="#">RFC1797</a> ]
127.0.0.0/8	Loopback	[RFC1700, page 5]
128.0.0.0/16	Reserved but subject to allocation	--
169.254.0.0/16	Link Local	--
172.16.0.0/12	Private-Use Networks	[ <a href="#">RFC1918</a> ]
191.255.0.0/16	Reserved but subject to allocation	--
192.0.0.0/24	Reserved but subject to allocation	--
192.0.2.0/24	Test-Net	
192.88.99.0/24	6to4 Relay Anycast	[ <a href="#">RFC3068</a> ]
192.168.0.0/16	Private-Use Networks	[ <a href="#">RFC1918</a> ]
198.18.0.0/15	Network Interconnect Device Benchmark Testing	[ <a href="#">RFC2544</a> ]
223.255.255.0/24	Reserved but subject to allocation	--
224.0.0.0/4	Multicast	[ <a href="#">RFC3171</a> ]
240.0.0.0/4	Reserved for Future Use	[RFC1700, page 4]



# Example Prefix List (Cisco)

```
ip prefix-list rfc1918-dsua deny 0.0.0.0/8 le 32
ip prefix-list rfc1918-dsua deny 10.0.0.0/8 le 32
ip prefix-list rfc1918-dsua deny 127.0.0.0/8 le 32
ip prefix-list rfc1918-dsua deny 169.254.0.0/16 le 32
ip prefix-list rfc1918-dsua deny 172.16.0.0/12 le 32
ip prefix-list rfc1918-dsua deny 192.0.2.0/24 le 32
ip prefix-list rfc1918-dsua deny 192.168.0.0/16 le 32
ip prefix-list rfc1918-dsua deny 224.0.0.0/3 le 32
ip prefix-list rfc1918-dsua deny 0.0.0.0/0 ge 25
ip prefix-list rfc1918-dsua permit 0.0.0.0/0 le 32
```



# Bogons

- IANA has reserved several blocks of IPv4 that have yet to be allocated to a RIR:
  - <http://www.iana.org/assignments/ipv4-address-space>
- These blocks of IPv4 addresses should never be advertised into the global Internet Route Table.
- Filters should be applied on the AS border for all inbound and outbound advertisements.



# Ingress Prefix Filter Template

- “It is hard to build the list.” --- “OK, we’ll build the community a template. Next excuse.”
- Bogon List by CYMRU Bogon Team
  - <http://www.cymru.com/Bogons/>
  - Starting point for putting together the Bogon Filtering.
  - Supplies up to date templates for Cisco and Juniper



# Ingress Prefix Filter Template

- Cisco Template by Barry Greene

- <ftp://ftp-eng.cisco.com/cons/isp/security/Ingress-Prefix-Filter-Templates/>

- Juniper Template by Steven Gill

- <http://www.qorbit.net/documents.html>

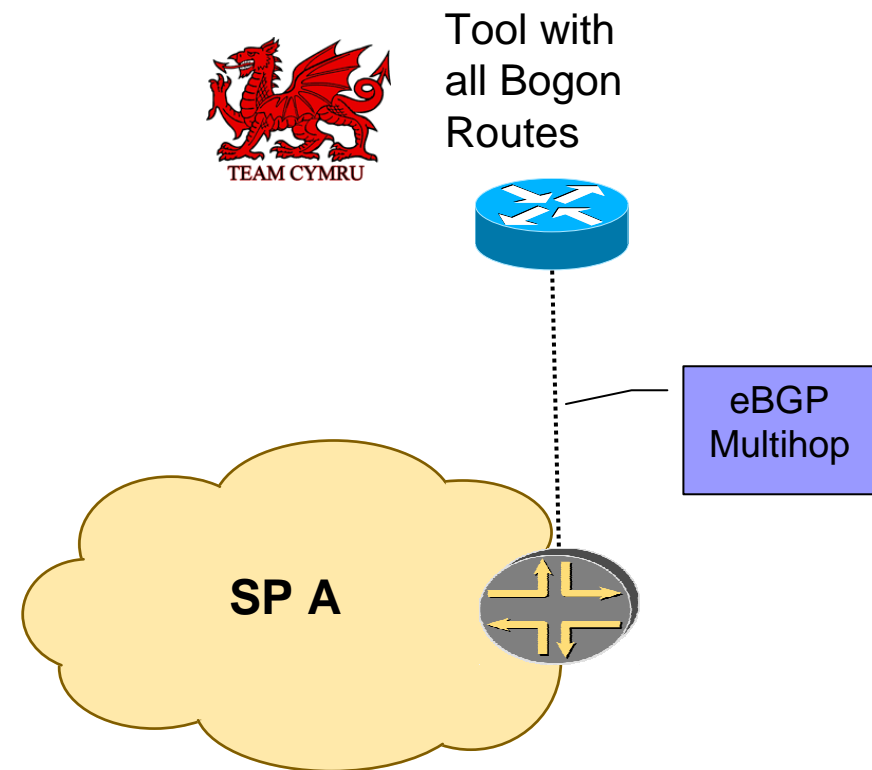


# Three Modes

- Build your own Bogon filter.
- Pull down one of the Bogon templates to get started.
- Sign up to the Bogon Router Server

# BOGON Project eBGP Feed

- The BOGON Project provides a eBGP Multihop feed which black holes Bogons.
- <http://www.team-cymru.org/Services/Bogons/routeserver.html>
- Uses a special community which allows you to match and set to null0



# *Net Police* Route Filtering





# “Net Police” Route Filtering

- *Net Police* route filtering describes ingress peering filtering that only allows the *minimum practical allocation* from a RIR.
  - So if APNIC’s minimum practical allocation is a /20, then the Net Police filter will only allow a /8 to a /20/. Any prefix larger than a /20 (i.e. a /21) will get dropped by the filter.
- Net Police Filtering has two effects:
  - Reduces the number of prefixes in an ISP’s RIB.
  - Protects the ISP from *Garbage in Garbage out* problems/incidents on the Net.



# “Net Police” Route Filtering

- Two Techniques:

- ☐ Permit only prefixes on the RIR's *minimum practical allocations*.
- ☐ Permit prefixes allocated by the RIRs with a lower boundary set by the ISP (i.e. /24 vs a /20).



# Net Police Filter Technique #1

- Permit Only Allocated IPv4 Blocks
- Need to check with each of the RIR's for details on which networks they are allocating from and what the specific *minimum practical allocation* for each block.
  - RIRs are announcing changes to the Internet Operations Aliases.
- ARIN - [http://www.arin.net/reference/ip\\_blocks.html](http://www.arin.net/reference/ip_blocks.html)
- RIPE - <https://www.ripe.net/ripe/docs/ripe-ncc-managed-address-space.html>
- APNIC - <http://www.apnic.net/db/min-alloc.html>
- LACNIC - <http://lacnic.net/en/registro/index.html>
- AFRINIC <http://www.afrinic.net/index.htm>



# Technique #1 Net Police Prefix List

(check for update)

```
!! APNIC
ip prefix-list FILTER permit 61.0.0.0/8 ge 9 le 20
ip prefix-list FILTER permit 202.0.0.0/7 ge 9 le 20
ip prefix-list FILTER permit 210.0.0.0/7 ge 9 le 20
ip prefix-list FILTER permit 218.0.0.0/7 ge 9 le 20
!! ARIN
ip prefix-list FILTER permit 63.0.0.0/8 ge 9 le 20
ip prefix-list FILTER permit 64.0.0.0/7 ge 9 le 20
ip prefix-list FILTER permit 66.0.0.0/8 ge 9 le 20
ip prefix-list FILTER permit 199.0.0.0/8 ge 9 le 20
ip prefix-list FILTER permit 200.0.0.0/8 ge 9 le 20
ip prefix-list FILTER permit 204.0.0.0/6 ge 9 le 20
ip prefix-list FILTER permit 208.0.0.0/7 ge 9 le 20
ip prefix-list FILTER permit 216.0.0.0/8 ge 9 le 20
!! RIPE NCC
ip prefix-list FILTER permit 62.0.0.0/8 ge 9 le 20
ip prefix-list FILTER permit 80.0.0.0/7 ge 9 le 20
ip prefix-list FILTER permit 193.0.0.0/8 ge 9 le 20
ip prefix-list FILTER permit 194.0.0.0/7 ge 9 le 20
ip prefix-list FILTER permit 212.0.0.0/7 ge 9 le 20
```



## Net Police Prefix List Deployment Issues

- Objective – protect the network from ISPs who won't and don't aggregate
- Impacts *more specific* style multihoming
- Impacts regions where domestic backbone is unavailable or costs \$\$\$ compared with international bandwidth
- Maintenance Overhead – requires updating when RIRs start allocating from new address blocks
- Understand the Consequences!



## Technique #2 Net Police Prefix List Alternative

- Permit Only Allocated IPv4 Blocks
- Move the minimal allocation prefix to a /24
- Most Operators agree that blocks longer than /24 should not be seen on the Net.
- This minimizes some of the operational impact to customer multihoming.



# Technique #2 Net Police Prefix List Alternative (check for update)

```
!! APNIC
ip prefix-list FILTER permit 61.0.0.0/8 ge 9 le 24
ip prefix-list FILTER permit 202.0.0.0/7 ge 9 le 24
ip prefix-list FILTER permit 210.0.0.0/7 ge 9 le 24
ip prefix-list FILTER permit 218.0.0.0/7 ge 9 le 24
!! ARIN
ip prefix-list FILTER permit 63.0.0.0/8 ge 9 le 24
ip prefix-list FILTER permit 64.0.0.0/7 ge 9 le 24
ip prefix-list FILTER permit 66.0.0.0/8 ge 9 le 24
ip prefix-list FILTER permit 199.0.0.0/8 ge 9 le 24
ip prefix-list FILTER permit 200.0.0.0/8 ge 9 le 24
ip prefix-list FILTER permit 204.0.0.0/6 ge 9 le 24
ip prefix-list FILTER permit 208.0.0.0/7 ge 9 le 24
ip prefix-list FILTER permit 216.0.0.0/8 ge 9 le 24
!! RIPE NCC
ip prefix-list FILTER permit 62.0.0.0/8 ge 9 le 24
ip prefix-list FILTER permit 80.0.0.0/7 ge 9 le 24
ip prefix-list FILTER permit 193.0.0.0/8 ge 9 le 24
ip prefix-list FILTER permit 194.0.0.0/7 ge 9 le 24
ip prefix-list FILTER permit 212.0.0.0/7 ge 9 le 24
```



## Bottom Line

- Net Police filtering effectively protects networks from garbage in garbage out problems on the Net.
  - ISPs using Net Police filters did not have the AS 7007 or 129/8 incidents effect their network.
- While Net Police filters are controversial, their use as a security tool has been proven.





# Looking for examples?

- Cisco Template by Barry Greene

- <ftp://ftp-eng.cisco.com/cons/isp/security/Ingress-Prefix-Filter-Templates/>


- Juniper Template by Steven Gill

- <http://www.qorbit.net/documents.html>


*Saying it Again*  
Prefix Filters on  
Customers

# Prefix Filters on Customers

- Prefix filter all routes from your customers!



# BGP with Customer Infers Multihoming





# Receiving Customer Prefixes


- ISPs should only accept prefixes which have been assigned or allocated to their downstream peer/customer.
- For example
  - Downstream has 220.50.0.0/20 block
  - Should only announce this to peers
  - Peers should only accept this from them
  - Explicitly permit prefixes from other ISPs (i.e. multihomed to two or more ISPS).



# Receiving Customer Prefixes


- Configuration example on upstream:

```
router bgp 100
  neighbor 222.222.10.1 remote-as 101
  neighbor 222.222.10.1 prefix-list customer in
  !
ip prefix-list customer permit 220.50.0.0/2
ip prefix-list customer deny 0.0.0.0/0 le 32
```



## Excuses – Why providers are not prefix filtering customers.

- “Some of my customers are multihomed, so they want to advertise more specifics.”
- “These are down stream ISPs, so their advertisements will change.”
- Filtering customer is just too hard!




# What if you do not filter your customer?

- Not filtering your customers puts your network at risk to:
  - Bogon Prefix Insertion (sucks down backscatter)
  - Un-Authorized Route Insertion (sucks down traffic)
  - Re-advertise other ISP's routes (customer's T1 becomes the peering link).




# What if you do not filter your customer?



# *Saying it Again* Prefixes to Peers

# Prefixes to Peers

- Prefix filter all routes to your peers!





# Prefixes to Peers

- What do you send to the Internet?
  - Your prefixes.
  - More specific customers prefixes (customers who are multihoming)
- What do you not send to the Internet?
  - DUSA Prefixes – assume junk will leak into your iBGP.
  - Bogons – assume garbage will leak into your iBGP.
  - Lower Prefix Boundary – Unless absolutely necessary, Do not allow anything in the /25 - /32 range.



# Egress Filter to ISP Peers - Issues

- The egress filter list can grow to be very large:
  - More specifics for customers.
  - Specific blocks from other ISPs




# Policy Questions

- Will you allow customers to announce IP prefixes from other ISPs?
- Will the customer be required to tell you these prefixes?
- Will you advertise these prefixes back to the ISP?
- Will you advertise these prefixes to the entire Internet?
- Will you transit communities from your customers?

# What if the filter is too big?

- Egress filter to peers could get large – too large for the router to handle.
- One approach – use BGP Community Filtering.




*Saying it Again*  
Ingress Prefix  
Filtering from Peers



# Prefixes from Peers

- Prefix filter all routes from your peers!






# Ingress Routes from Peers or Upstream

- Ingress Routes from Peers and/or the Upstream ISP are the nets of the Internet.
- Ideally, the peering policy should be specific so that exact filters can be put in place.
  - Dynamic nature of the peering makes it hard to maintain specific route filters.



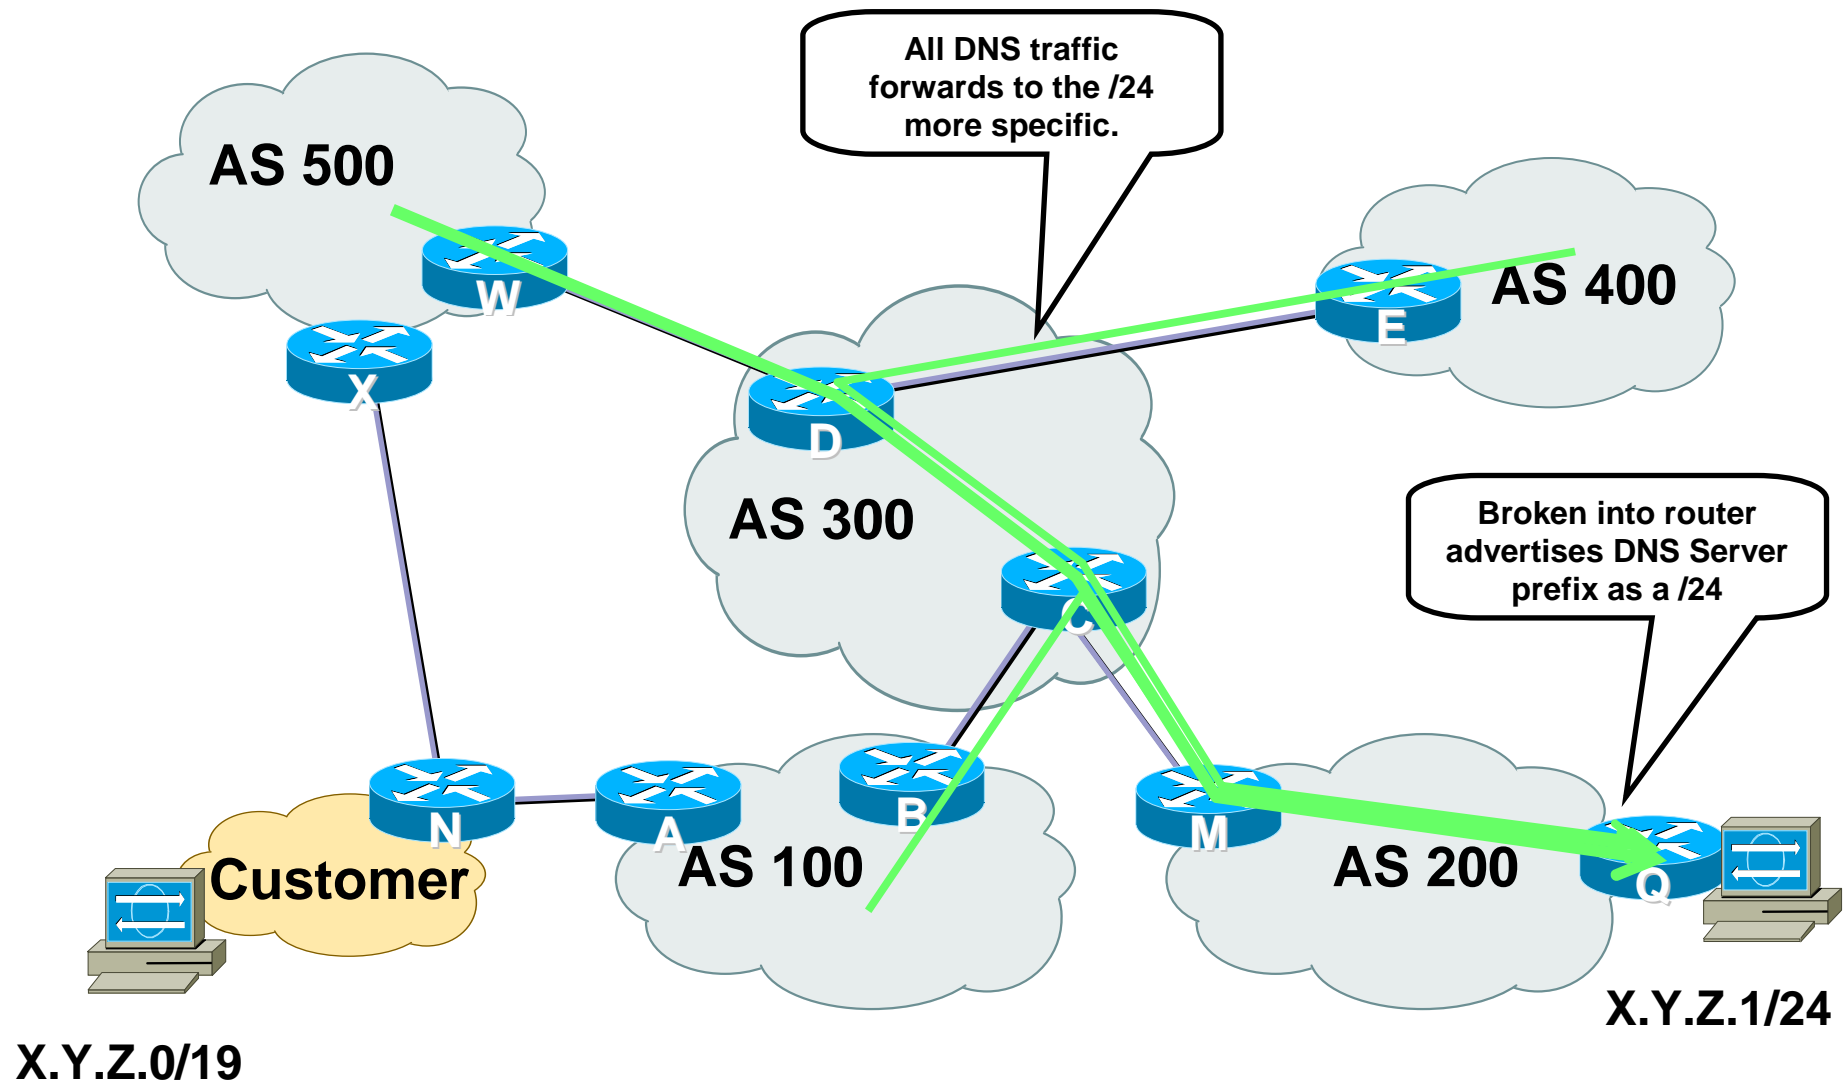
## Receiving Prefixes from Upstream & Peers (ideal case)

- ☐ Don't accept RFC1918 etc prefixes
- ☐ Don't accept your own prefix
- ☐ Don't accept default (unless you need it)
- ☐ Don't accept prefixes longer than /24
- ☐ Don't accept prefixes on IXPs your  
whom you have membership
- ☐ Consider *Net Police* Filtering
- ☐ Consider denying MED or communities



# How do you work a Hijack?

# What is a prefix hijack?





# What do you do?

- How do you detect a hijack is happening?
- Who is doing the hijacking?
- How are they doing the hijacking?
- Who do you call?
- How do you influence change?
- How do you find people at the hijacking ASN to help?
- What if it is an ASN which will not help?



# Preparation is Critical

- You need the phone numbers, E-mails, Chat ID, etc of all your upstreams, peers, and other peering business partners.
- While groups are trying to build tools, the #1 technique of stopping a hijack is to get on the phone and work with the SP community.




# Is De-aggregation an Option?

- You are a good CIDR Citizen advertising a /19.
- The Hijacker injects a /24.
- You inject a /24.
- They inject a /25.
- Could you inject a /25?
- Who has the better AS Path?



# What do you do?



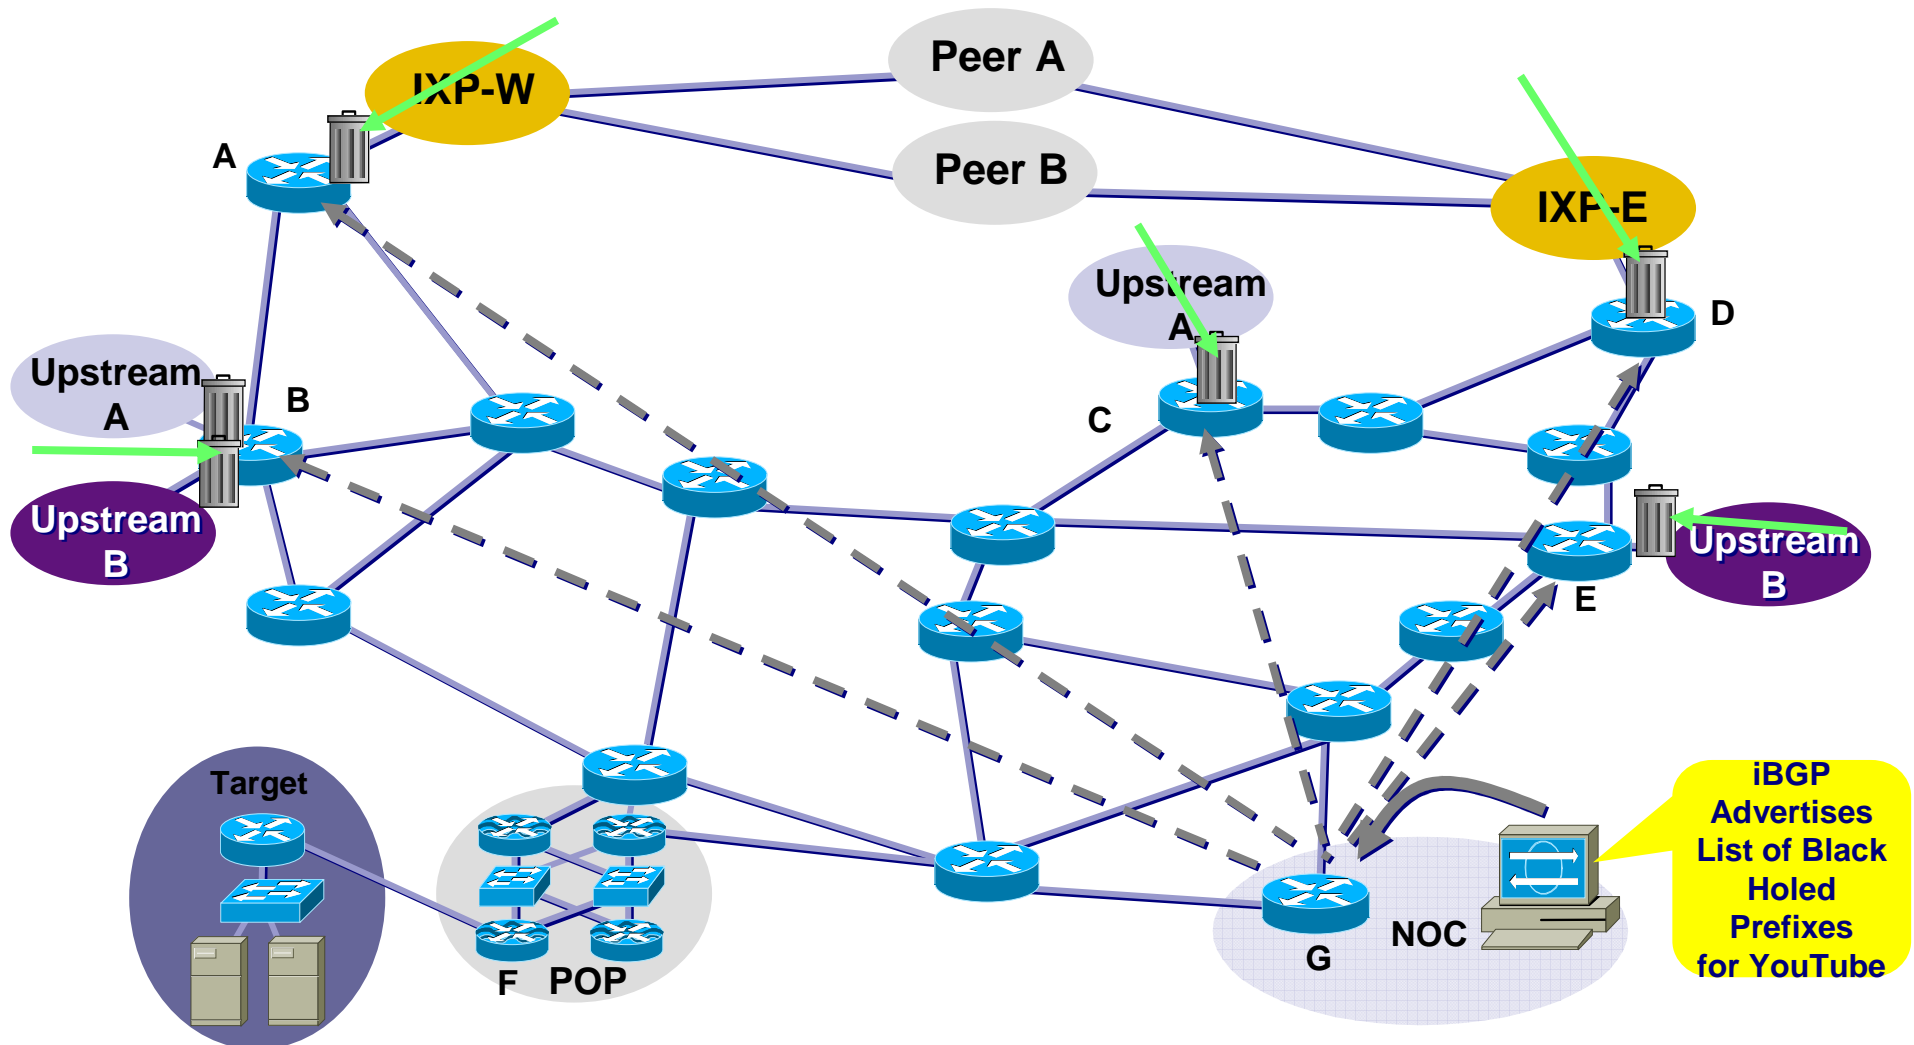
# What Happened in Pakistan?




# SITREP

- Government is upset, they don't want anyone in Pakistan seeing Videos
- Order comes from to the government – block YouTube
- How would an SP do it quickly?

# RTBH - YouTube





# The problem – no Egress Filter

- BGP Community Triggered
- There wasn't a “no-export” community on the Black Hole
- So the Black Hole prefix (YouTube) gets advertised to the rest of the world.
- Any egress filter of not sending out your own community would have caught this.
- An egress prefix-filter would have caught this.