



# Large Scale Denial of Service Mitigation

**Darrel Lewis** ([darlewis@cisco.com](mailto:darlewis@cisco.com)),

**Technical Leader, Cisco Systems**

**Paul Quinn** ([paquinn@cisco.com](mailto:paquinn@cisco.com)),

**Technical Leader, Cisco Systems**

# Agenda

Cisco.com

- **Introduction**
- **Prerequisites to Mitigation**
- **Mitigation**
- **Large Scale Discussion**
- **Future Developments**



# Introduction

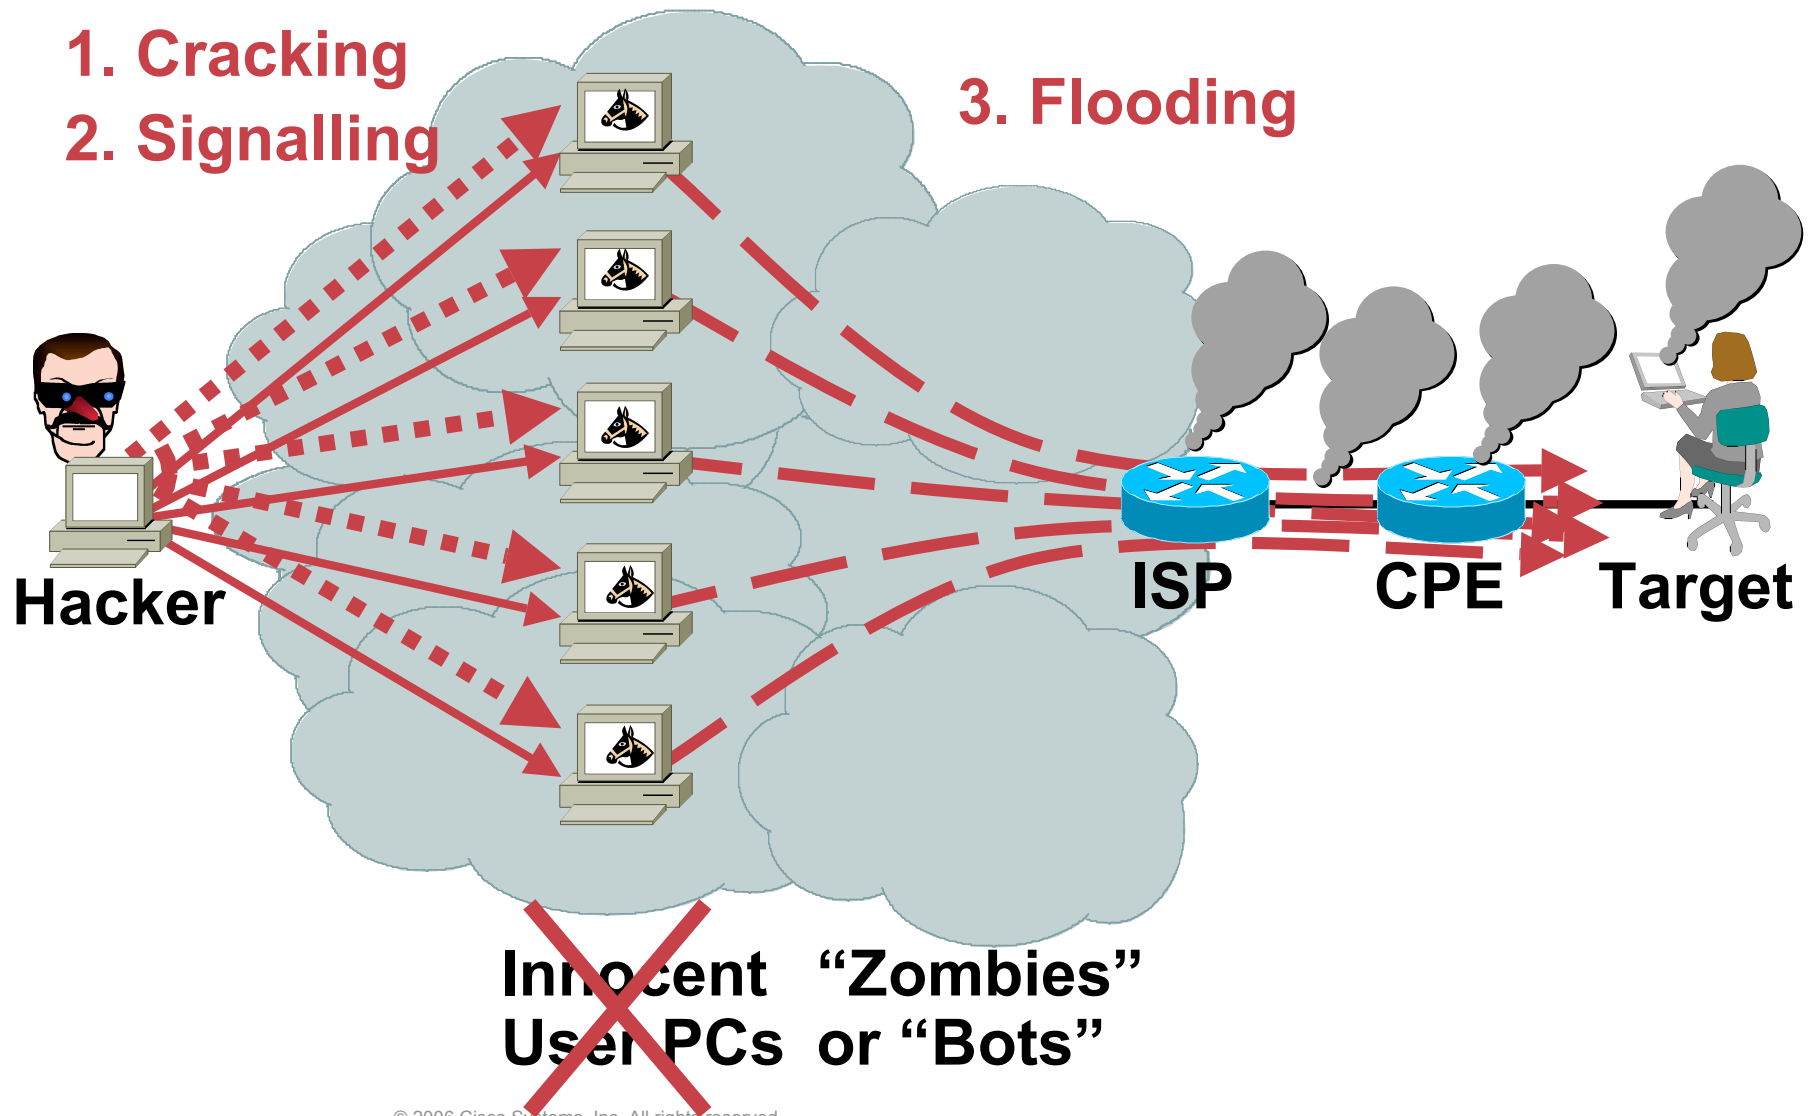
# Denial of Service Attacks

Cisco.com

- **We understand intrusions (patch, patch, patch ;-))**
- **What about DoS? Do “the right things” and still suffer**
- **The vast majority of modern DoS attacks are distributed**  
**DDoS IS DoS**
- **DoS is often driven by financial motivation**  
**DoS for hire :-(**  
**Economically-driven miscreant community**
- **DoS cannot be ignored; your business depends on effective handling of attacks**


# DoS: The Procedure

Cisco.com



# An SP View: Denial of Service

Cisco.com



# Denial of Service Trends

Cisco.com

- **Multi-path**  
**Truly distributed**  
**Inbound and Outbound**
- **Multi-vector**  
**SYN AND UDP AND...**
- **Financial incentive**  
**SPAM, DoS-for-hire**  
**Large, thriving business**  
**Forces us to reassess the risk profile**

# Infrastructure Security

Cisco.com

- **All of the techniques talked about today also assume that the infrastructure is available to route and forward packets!**
- **The infrastructure can either be targeted, or be impacted indirectly**



# Infrastructure Attacks

Cisco.com

- **Infrastructure attacks are increasing in both volume and sophistication**

**Sites with Cisco documents and presentations on routing protocols (and I don't mean Cisco.com)**

**Marked increase in presentations about routers, routing and Cisco IOS vulnerabilities at conferences like Blackhat, Defcon and Hivercon**

**Router attack tools and training are being published**

- **Why mount high-traffic DDOS attacks when you can take out your target's gateway routers?**
- **Hijacked routers are valuable in the spam world, which has a profit driver**
- **Router compromise (0wn3d) due to weak password**

# From Bad to Worms

- **Worms have emerged as the new security reality**
- **Old worms never die!**
  - Millions of UPnP and Slammer packets still captured daily
- **Most worms are intended to compromise hosts**
- **Worm propagation is dependant on network availability**
- **Worms and DoS are closely related**
  - Secondary worm effects can lead to denial of service
  - Worms enable DoS by compromising hosts → BOTnets
- **Perimeters are crumbling under the worm onslaught (VPN/mobile workers, partners, etc.)**
- **How do you Assess your backbone's risk?**
  - Per flow forwarding
  - Excess capacity
  - SLAs

# Worms and the Infrastructure

Cisco.com

- **Worms typically infect end-stations**
- **To date, worms have not targeted infrastructure BUT secondary effects have wreaked havoc**
  - Increased traffic**
  - Random scanning for destination**
  - Destination address is multicast**
  - TTL and other header variances**
- **At the core SP level, the aggregate effects of a worm can be substantial**
- **Worm severity is escalating and evolving**

# What we are and are not Covering

Cisco.com

- **In scope:**

**Large-scale attacks: large number of customer impacted, multiple sites**

**Collateral damage to the network**

- **Out of scope:**

**'Small scale' deployments**

**Content evaluation techniques**

**DDoS needs to cover cases where the content is valid!**

**This isn't going to be an IPS signature discussion**



# Prerequisites to Mitigation

# Six Phases to Large-Scale Security Incident Response

Cisco.com

- 1. Preparation**
- 2. Identification**
- 3. Classification**
- 4. Traceback**
- 5. Reaction**
- 6. Post mortem**

# Preparation

Cisco.com

## Preparation—Develop and Deploy a Solid Security Foundation

- Includes technical and non-technical components
- Encompasses best practices
- The hardest, yet most important phase
- Without adequate preparation, you are destined to fail
- The midst of a large attack is not the time to be implementing foundational best practices and processes

# Preparation

- **Know the enemy**
  - Understand what drives the miscreants**
  - Understand their techniques**
- **Create the security team and plan**
  - Who handles security during an event? Is it the security folks? The networking folks?**
- **Harden the devices**
- **Prepare the tools**
  - Network telemetry**
  - Reaction tools**
  - Understand performance characteristics**



# You Are Under Attack: It's Usually Too Late

Cisco.com

```
TCP
Local Address      Remote Address      State
-----
*.*                *.*                IDLE
*.sunrpc           *.*                LISTEN
*.ftp              *.*                LISTEN
*.telnet           *.*                LISTEN
*.finger           *.*                LISTEN
target.telnet      10.10.10.11.41508   SYN_RCVD
target.telnet      10.10.10.12.41508   SYN_RCVD
target.telnet      10.10.10.13.41508   SYN_RCVD
target.telnet      10.10.10.14.41508   SYN_RCVD
target.telnet      10.10.10.10.41508   SYN_RCVD
target.telnet      10.10.10.15.41508   SYN_RCVD
target.telnet      10.10.10.16.41508   SYN_RCVD
target.telnet      10.10.10.17.41508   SYN_RCVD
target.telnet      10.10.10.18.41508   SYN_RCVD
target.telnet      10.10.10.19.41508   SYN_RCVD
target.telnet      10.10.10.20.41508   SYN_RCVD
*.*                *.*                IDLE
```

**Output From**  
**netstat -an**  
**on Target**  
**Host**

**Once the Connection Queue Is Full of Waiting-to-Be-Completed Connections, **All** SYN+RCVDs Get FIFOed out!**

# Ways to Detect

Cisco.com

- **Customer call**
  - “The Internet is down”
- **Unexplained changes in network baseline**
  - SNMP: line/CPU overload, drops**
  - Bandwidth**
  - NetFlow**
- **ACLs with logging**
- **Backscatter**
- **Packet capture**
- **Network IDS**
- **Anomaly detection**

# Network Baselines

Cisco.com

- **NMS baselines**
- **Unexplained changes in link utilization**  
Worms can generate a lot of traffic, sudden changes in link utilization can indicate a worm
- **Unexplained changes in CPU utilization**  
Worm scans can effect routers/switches resulting in increased CPU both process and interrupt switched
- **Unexplained syslog entries**
- **These are examples**  
Changes don't always indicate a security event!  
**Need to know what's normal in order to identify abnormal behavior**

# Source vs. Destination Detection

Cisco.com

- **Destination detection focuses on the victim**

**This is often easy!**

- **How do we distinguish between good and bad sources?**

**Furthermore, what about a given source's behavior when it is both good and bad?**

**Flash crowds**

- **Comprehensive detection deployments need to take src and dst detection into account**

# Classification

- **Classification—understand the details and scope of the attack**

Identification is not sufficient; once an attack is identified, details matter

Guides subsequent actions

- **Identification and classification are often simultaneous**

# Classification

- **Qualify and quantify the attack without jeopardizing services availability (e.g., crashing a router):**
  - What type of attack has been identified?**
  - What's the effect of the attack on the victim(s)?**
  - What next steps are required (if any)?**
- **At the very least:**
  - Source and destination address**
  - Protocol information**
  - Port information**

# Reaction

## Reaction—Do Something to Counter the Attack

- **Should you mitigate the attack?**  
Where? How?
- **No reaction is a valid form of reaction in certain circumstances**
- **Reaction often entails more than just throwing an ACL onto a router**

# Traceback

- **Operational challenge**
  - Know your peers!**
  - Use iNOC phone network to communicate with them directly**
- **Traceback—what are the sources of the attack?**
  - How to trace to network ingress points**
  - Understand your topology!**
- **Traceback to network perimeter**
  - NetFlow**
  - Backscatter**
  - Packet accounting**
- **Retain attack data**
  - Use to correlate interdomain traceback**
  - Required for prosecution**
  - Clarify billing and other disputes**
  - Post mortem analysis**



# Reacting to Attacks

Cisco.com

- **Many varying reaction mechanisms**
- **No one tool or technique is applicable in all circumstances**

**Think “toolkit”**

**Automate where possible**

**Don't forget about the operational costs!**

- **Choose your techniques wisely**

# Post Mortem

Cisco.com

## Post Mortem—Analyze the Event

- **The step everyone forgets!**
- **What worked? What didn't? How can we improve?**
- **What can be done to build build defense against repeat occurrences**
- **Was the DOS attack you just handled the real threat? Or was it a smoke screen for something else that just happened?**
- **What can you do to make it faster, easier, less painful in the future?**
- **Metrics are important!**
  - Resources, headcount, etc.**



# SP Community Efforts

# Inter-Provider Cooperation

Cisco.com

- **Increased inter-provider cooperation is required when handling attacks**
  - Competitors working together**
  - Traceback requires inter-AS work**
- **Ad-hoc buddy network didn't work**
  - iNOC IP phone network**
  - NSP-SEC mailing list**

# NSP-SEC

Cisco.com

- **NSP-SEC – Closed Security Operations Alias for engineers actively working with NSPs/ISPs to mitigate security incidents.**
- **Multiple Layers of sanity checking the applicability and trust levels of individuals**
- **Not meant to be perfect – just better than what we had before**
- <http://puck.nether.net/mailman/listinfo/nsp-security>

# NSP-SEC: Daily DDOS Mitigation Work

Cisco.com

**I've been working an attack against XXX.YY.236.66/32 and XXX.YY.236.69/32. We're seeing traffic come from <ISP-A>, <ISP-B>, <IXP-East/West> and others.**

**Attack is hitting both IP's on tcp 53 and sourced with x.y.0.0.**

**I've got it filtered so it's not a big problem, but if anyone is around I'd appreciate it if you could filter/trace on your network. I'll be up for a while :/**

# NSP-SEC-DISCUSS

Cisco.com

- **NSP-SEC** is where the mitigation takes place. You do not learn anything, you are already expected to know.
- **NSP-SEC-DISCUSS** is the place to learn, consult, work on new mitigation techniques, and lurk (if you want to).

<http://puck.nether.net/mailman/listinfo/nsp-security-discuss>

# Technical Efforts

Cisco.com

- **Security forums and topics at large events**
  - NANOG**
  - IETF**
- **Sharing the wealth**
  - Service providers are sharing their experience with others**
  - UU published traceback technique**
  - AOL talked about ISIS migration**



# What Do ISPs Need to Do?

Cisco.com

- **Implement Best Common Practices (BCPs)**
  - ISP infrastructure security**
  - ISP network security**
  - ISP services security**
- **Work with operations groups, standards organizations, and vendors on new solutions**
- **NSP-SEC mailing list is a great example of inter-provider communication and mitigation**

# iNOC DBA – Why ?

Cisco.com

- **ISPs need to coordinate for attacks. They need to talk.**
- **It is not easy to reach the right contact. The engineer you are trying to reach will not likely pick up the phone.**
- **Solution: Dedicated NOC Hotline System**

**INOC-DBA: *Inter-NOC Dial-by-ASN***

# What is the problem?

- **ISPs needed to talk to each other in the middle of the attack.**
- **Top Engineers inside ISPs often do not pick up the phone and/or screen calls so they can get work done. If the line is an outside line, they do not pick up.**
- **Potential solution – create a dedicated NOC Hotline system. When the *NOC Hotline* rings, you know it is one of the NOC Engineer's peers.**

# iNOC DBA Hotline

Cisco.com

- The iNOC Hotline is used to get directly to peers.
- Numbering system based on the Internet:  
ASnumber:phone  
109:100 is Barry's house.
- SIP Based VoIP system, managed by Packet Clearing House ([www.pch.net](http://www.pch.net)), and sponsored by Cisco.
- [www.pch.net/inoc-dba](http://www.pch.net/inoc-dba)



# How to Participate

Cisco.com

- **With your own phones:**

**PCH needs your MAC address, contact info, ASNs, and extension number.**

- **With PCH phones:**

**PCH need your contact and shipping address, ASNs, and extension number.**

# How is iNOC being used today?

Cisco.com

- **Used during attacks like Slammer**
- **Coordination during large DoS attacks targeting multi-homed customers**
- **Many DNS Root Servers are using the iNOC Hotline for their phone communication**
- **General engineering consultation – SP engineers working on inter-SP issues**

# More Information

Cisco.com

- **General information:**  
<http://www.pch.net/inoc-dba/>
- **Mailing-list archive:**  
<http://www.pch.net/resources/discussion/inoc-dba/archive/>
- **Who's participating:**  
<http://www.pch.net/inoc-dba/directory/>

Exchanges	Carriers		Associations	
LINX	SD-NAP	UUnet	AT&T	ARIN
PAIX	LAIIX	Sprint	SBC	APNIC
Equinix	NSP-IXP2		C&W	AOL/T-WRIPE/NCC
AMS-IX	NOTA	Genuity	RCN	ICANN
MAEs	OIX	Verio/NTT	TDS	ISC



# Mitigation



# Capacity as a Solution

Cisco.com

- **For many types of attacks, a common solution is to add more capacity**
- **Not every problem gets solved this way**
  - Think about collateral damage
- **Challenge is to solve all the problems in the most economically feasible way**



# Data Plane Mitigation

# RFC 2827/BCP 38 Ingress *Packet* Filtering

Cisco.com

**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**

# BCP 38: Consequences of No Action

Cisco.com

## No BCP 38 Means That:

- **Devices can (wittingly or unwittingly) send traffic with spoofed and/or randomly changing source addresses out to the network**
- **Complicates traceback immensely**
- **Sending bogus traffic is NOT free!**

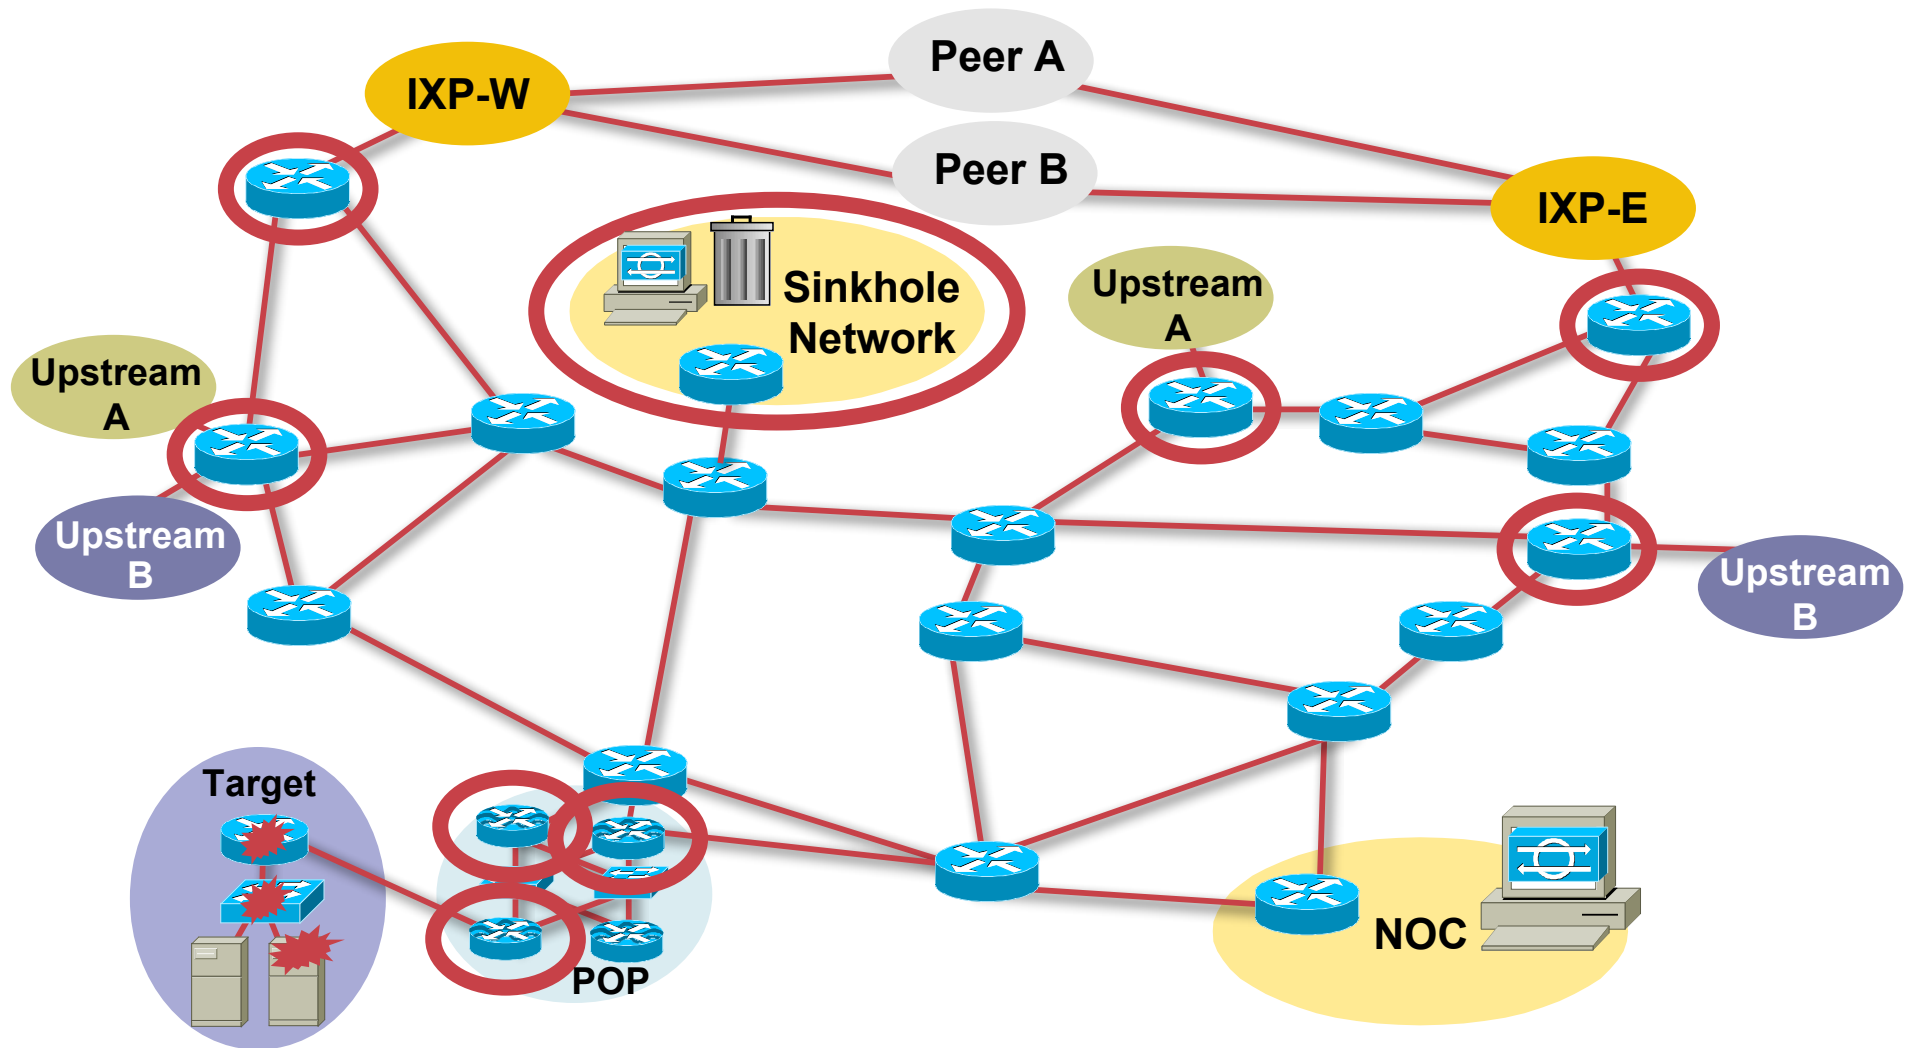
# BCP 38 *Packet* Filtering Principles

Cisco.com

- **Filter as close to the edge as possible**
- **Filter as precisely as possible**
- **Filter both source and destination where possible**


# Where to React?

Cisco.com



# Reacting to an Attack with CAR

Cisco.com



- Layer-3 input and output rate limits—specifically **input rate limits**
- Security filters use the input rate limit to drop *Packets* before they are forwarded through the network
- Aggregate and granular limits
  - Port, MAC address, IP address, application, precedence, QOS\_ID
- Excess burst policies

# Reacting to an Attack with ACLs

Cisco.com

- **Traditional method for stopping attacks**
- **Scaling issues encountered:**
  - Operational difficulties**
  - Changes on the fly**
  - Multiple ACLs per interface**
  - Performance concerns**



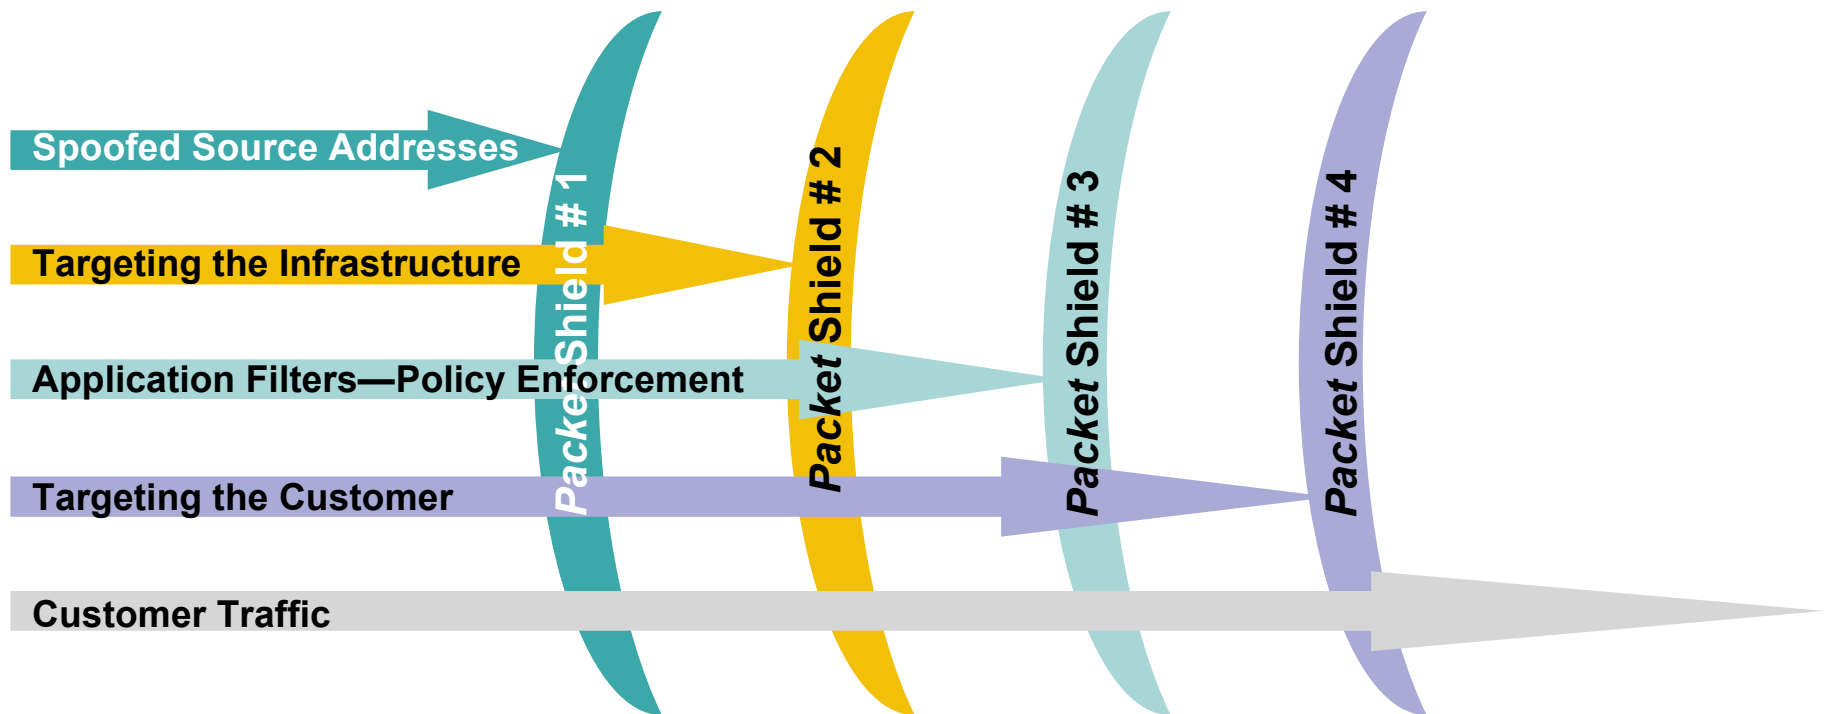
# ACLs: Deployment Considerations

Cisco.com

- How does the ACL load into the router? Does it interrupt *Packet* flow?
- How many ACEs can be supported in hardware? In software?
- How does ACL depth impact performance?
- How do multiple concurrent features effect performance?

# Packet Filtering Viewed Horizontally

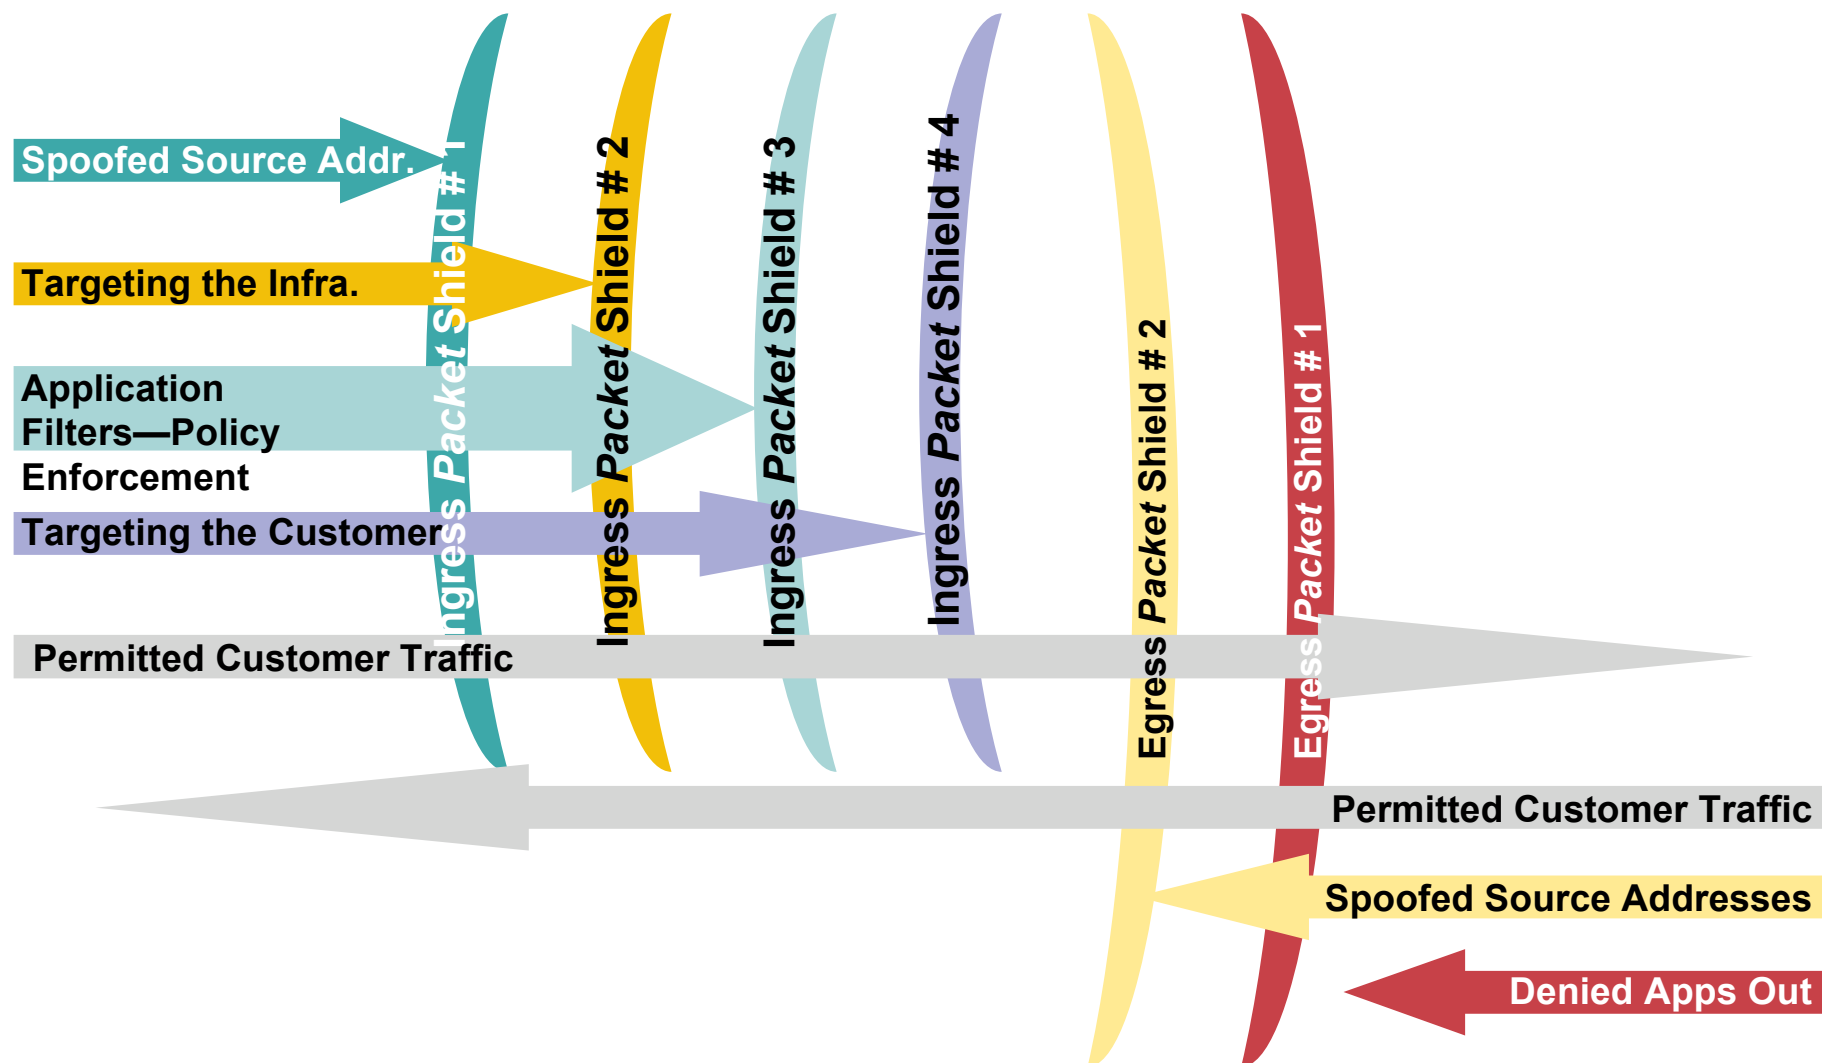
Cisco.com



# Packet Filtering

## Remember to Filter the Return Path

Cisco.com



# ACL Construction

- **Most common problems:**
  - Poorly constructed ACLs**
  - Ordering matters**
- **Scaling and maintainability issues with ACLs are commonplace**
- **Make your ACLs as modular and simple as possible**

# ACL Categories: Hybrid Philosophy

Cisco.com

## Hybrid Permit/Deny

- **Anti-spoofing**
- **Anti-bogon (source)**
- **Infrastructure**
- **Explicit deny specific L3**
- **Explicit deny specific L4**
- **Incident reaction**
- **Explicit permit L3 (good traffic)**
- **Explicit permit L4 (good traffic)**
- **Explicit deny everything else (auditing)**

# ACL Summary

- **ACLs are widely deployed as a primary containment tool**
- **Prerequisites: identification and classification—need to know what to filter**
- **Apply as specific an ACL as possible**
- **ACLs are good for static attacks, not as effective for rapidly changing attack profiles**
- **Understand ACL performance limitations before an attack occurs**
- **Operational efficiencies are important—scripted**

# The Pros and Cons of ACLs

- **ACLs' key strengths:**
  - Detailed *Packet* filtering (ports, protocols, ranges, fragments, etc.)
  - Relatively static filtering environment
  - Clear filtering policy
- **ACLs can have issues when faced with:**
  - Dynamic attack profiles (different sources, different entry points, etc.)
  - Frequent changes
  - Quick, simultaneous deployment on a multitude of devices
  - Operationally hard to remove
- **Because of these weaknesses, another tool was developed—using the control plane to signal the action**

# ACL Mitigation: SQL Slammer

Cisco.com

- **Slammer Worm**
  - Rapidly spreading worm in Jan 2003**
  - Targeted vulnerability in MS SQL and MSDE**
  - UDP packets destined to port 1434**
- **Infected not only hosts but networks as well**
  - Random scanning**
  - Surge in traffic**
- **ACL here widely used by SP to stop worm traffic**
  - Easy, static ACL: deny UDP/1434**
  - SP deployed an ACL at their edges**
- **Philosophical debate: can / should SPs act as the world's firewall?**



# Blaster: ACLs Were Not Effective

- **Blaster Worm**

Multiple propagation methods, open port: TCP/135, 4444, etc.

Once in place, worm launched a SYN flood on a specific date

- **With a high number of infected hosts, a significant DoS attacks was launched**

Effects not only SYN flood target but has collateral effect as well

- **ACL could help limit the propagation but were not effective at stopping the DoS**

Cannot distinguish between “good” and “bad” SYN


Large number of hosts, often sending valid and invalid traffic, made per host blocking ineffective



# Reacting with the Control Plane

# Routers Drop Data, Often!

Cisco.com



- An AS collects all the garbage (backscatter, scans, etc.) destined for 10.1/19, 10.1.96/19, and 10.1.128/17 addresses
- Routers that source those aggregates drop the data to unreachable parts of the networks, and are required to process data, send ICMP unreachables, etc.

# Black Hole Filtering

Cisco.com

- **Black Hole Filtering** or **Black Hole Routing** forwards a *Packet* to a router's **bit bucket**

Also known as “route to Null0”

- Works only on destination addresses, since it is really part of the forwarding logic
- Forwarding ASICs are designed to work with routes to Null0—dropping the packet with minimal to no performance impact
- Used for years as a means to “black hole” unwanted *Packets*

# Remotely Triggered Black Hole Filtering

Cisco.com

- **We will use BGP to trigger a network-wide response to an attack**
- **A simple static route and BGP will enable a network-wide destination address black hole as fast as iBGP can update the network**
- **This provides a tool that can be used to respond to security-related events and forms a foundation for other remotely triggered uses**
- **Often referred to as RTBH**

# Remotely Triggered Black Hole (RTBH)

Cisco.com

- **Configure all edge routers with static route to Null0 (must use “reserved” network)**

**ip route 192.0.2.1 255.255.255.255 Null0**

- **Configure trigger router**

**Part of iBGP mesh**

**Dedicated router recommended**

- **Activate black hole**

**Redistribute host route for victim into BGP with next-hop set to 192.0.2.1**

**Route is propagated using BGP to all BGP speakers and installed on routers with 192.0.2.1 route**

**All traffic to victim now sent to Null0**

# Step 1: Prepare All the Routers with Trigger

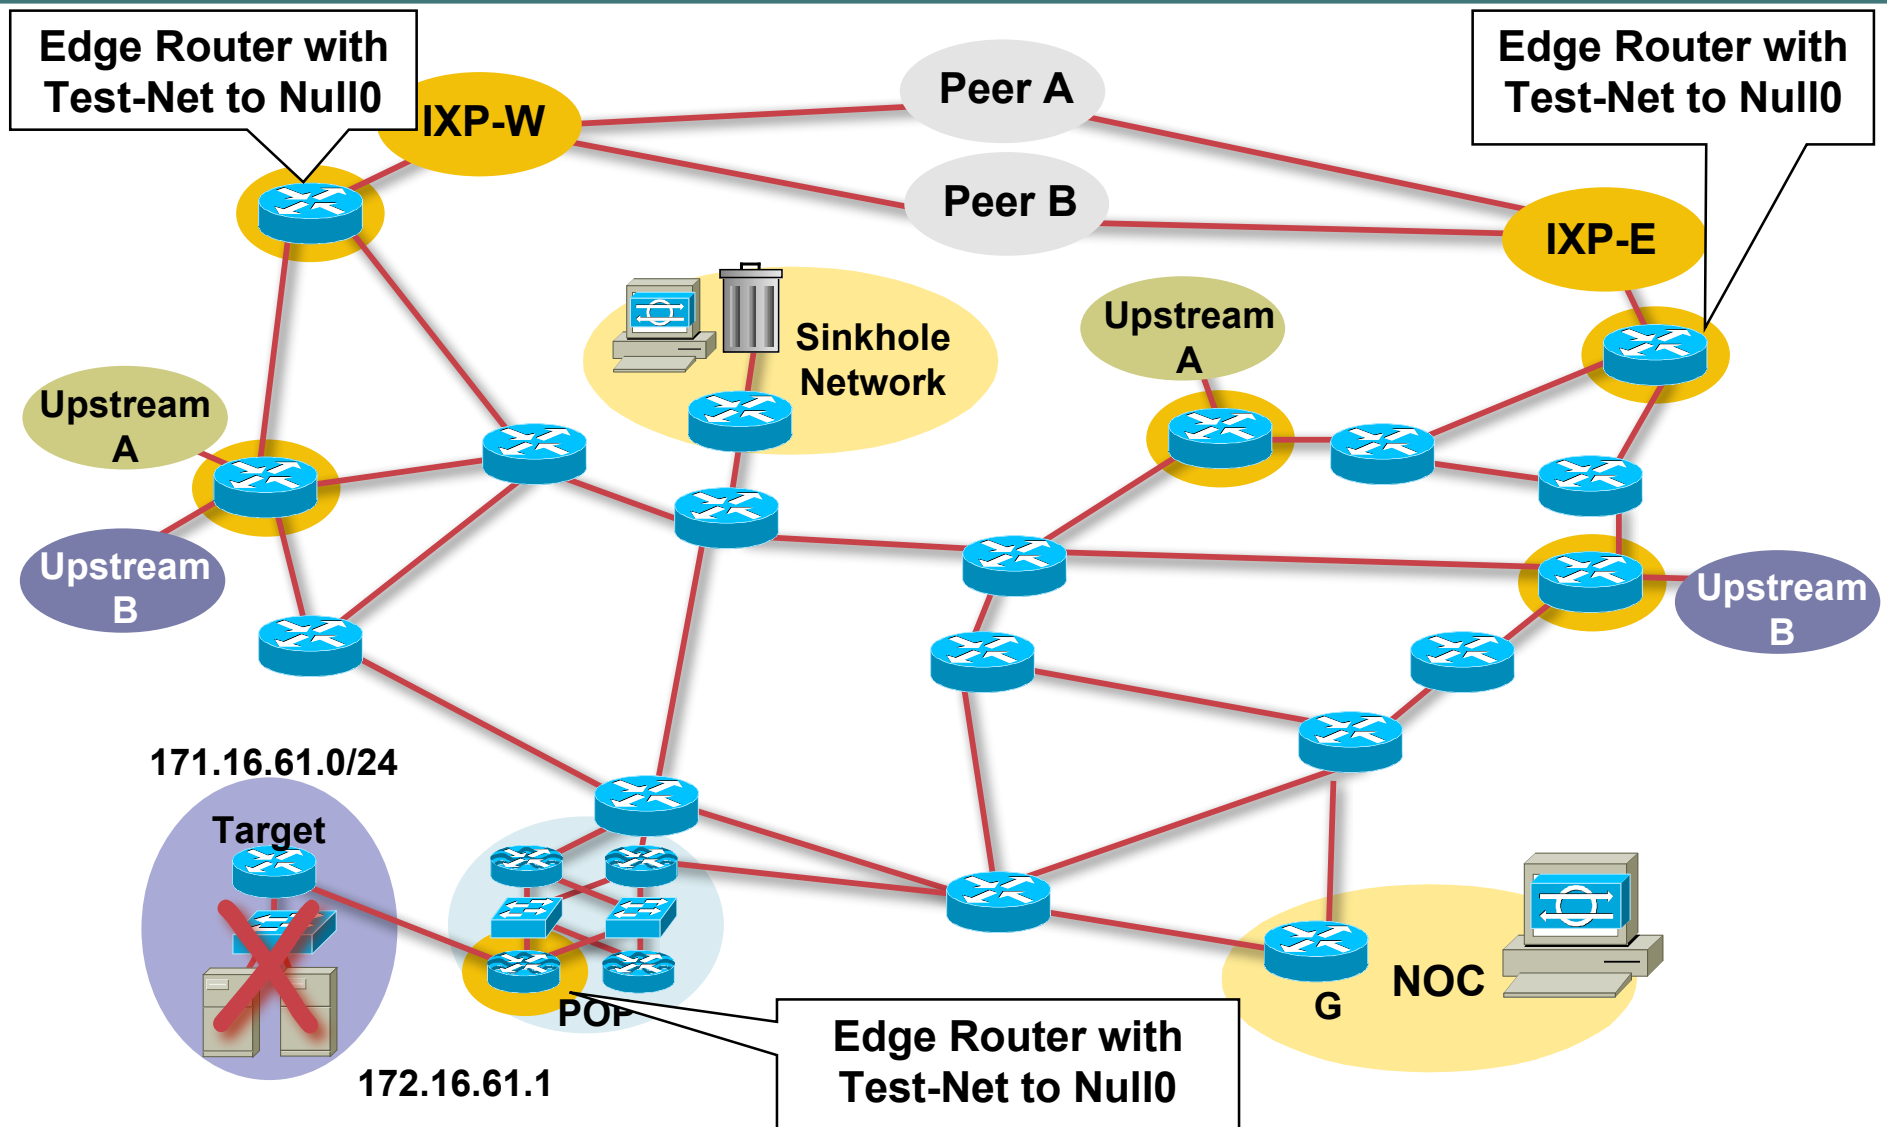
Cisco.com

- **Select a small block that will not be used for anything other than black hole filtering; Test-Net (192.0.2.0/24) is optimal since it should not be in use**
- **Put a static route with Test-Net—192.0.2.0/24 to Null0 on every edge router on the network**

```
ip route 192.0.2.1 255.255.255.255 Null0
```

# Step 1: Prepare All the Routers with Trigger

Cisco.com





## Step 2: Prepare the Trigger Router

Cisco.com

**The Trigger Router Is the Device That Will Inject the iBGP Announcement into the ISP's Network**

- **Should be part of the iBGP mesh—but does not have to accept routes**
- **Can be a separate router (recommended)**
- **Can be a production router**
- **Can be a workstation with Zebra/Quagga (interface with Perl scripts and other tools)**

# Trigger Router's Config

Cisco.com

**Redistribute  
Static with a  
Route-Map**

```
router bgp 65535
·
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 200
set community no-export
set origin igp
!
Route-map static-to-bgp permit 20
```

**Match  
Static  
Route Tag**

**Set Next-  
Hop to the  
Trigger**

**Set Local-Pref**

## Step 3: Activate the Black Hole

- **Add a static route to the destination to be blackholed; the static is added with the “tag 66” to keep it separate from other statics on the router**

```
ip route 172.16.61.1 255.255.255.255 Null0 Tag 66
```

- **BGP advertisement goes out to all BGP-speaking routers**
- **Routers receive BGP update and “glue” it to the existing static route; due to recursion, the next-hop is now Null0**

## Step 3: Activate the Black Hole

Cisco.com

**BGP Sent—172.16.61.1 Next-Hop = 192.0.2.1**

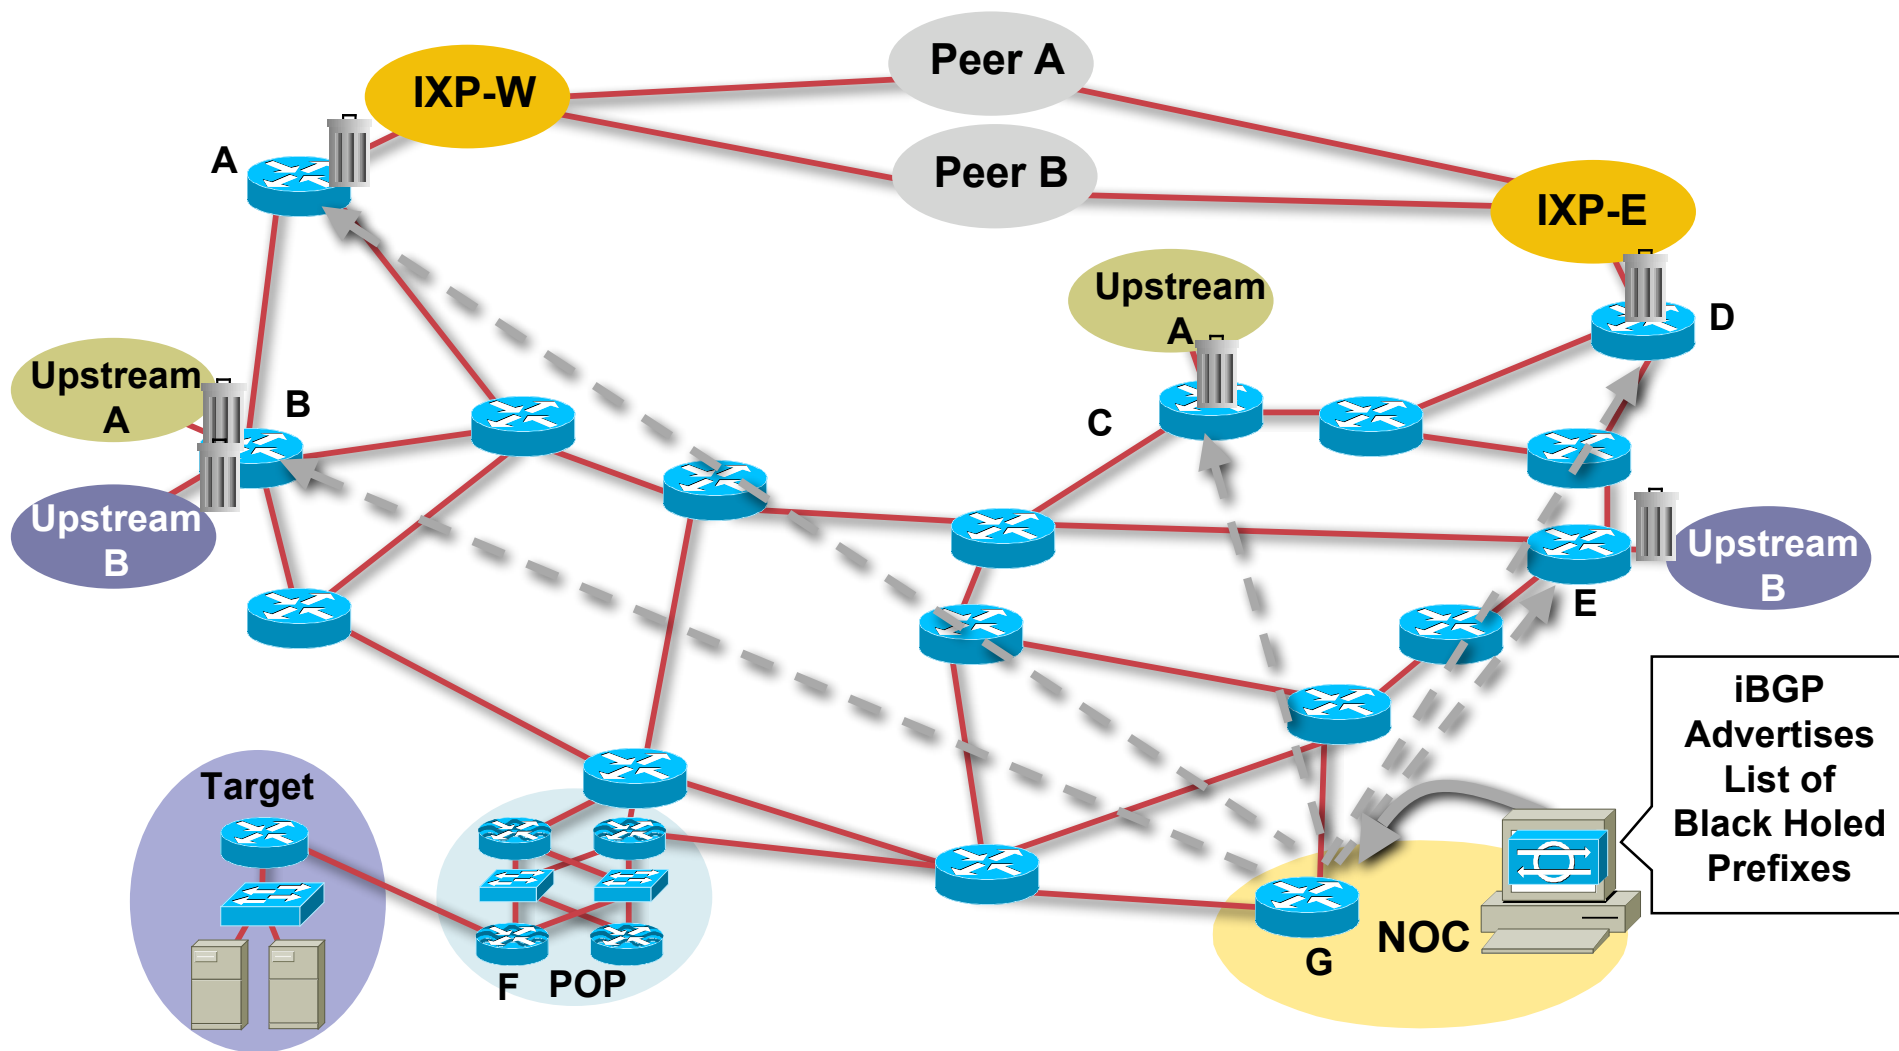
**Static Route in Edge Router—192.0.2.1 = Null0**

**172.16.61.1 = 192.0.2.1 = Null0**

**Next-Hop of 172.16.61.1  
Is Now Equal to Null0**

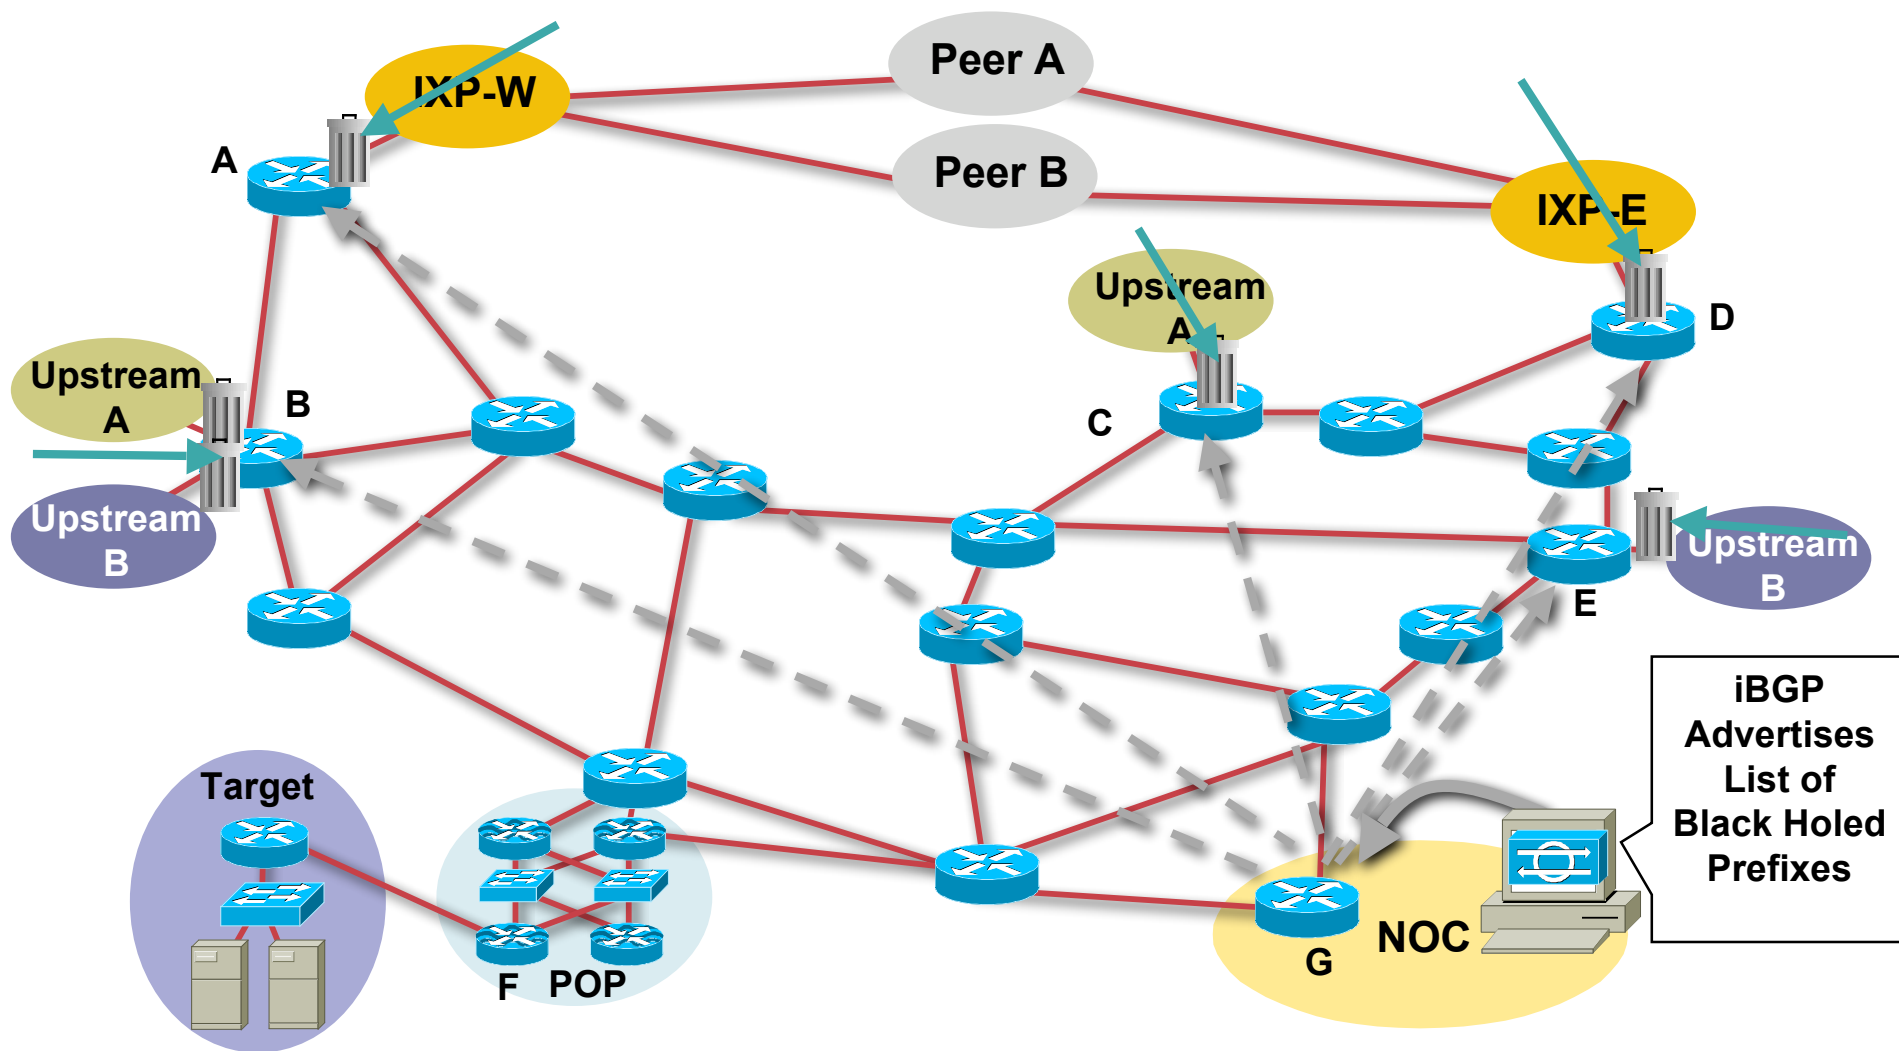
# Step 3: Activate the Black Hole

Cisco.com



# Customer is DOSed (After) *Packet Drops* Pushed to the Edge

Cisco.com



# Customer Based triggering

Cisco.com

- **Allow your customers to trigger the drops directly**

**No call to SP --> they decide when to drop traffic**

**Customer injects route that triggers null0 route**

- **Can be offered as a service**

**Who detects?**

- **Proper configuration required to ensure that the customer can only trigger drops for appropriate addresses**

**You don't want Customer A dropping Customer B's traffic!**

# RTBH: Triggered Source Drops

Cisco.com


- **Dropping on destination is very important**  
Dropping on source is often what we really need
- **Reacting using source address provides some interesting options:**
  - Stop the attack without taking the destination offline
  - Filter command and control servers
  - Filter (contain) infected end stations
- **Must be rapid and scalable**
  - Leverage pervasive BGP again



# Strict uRPF Check (Unicast Reverse Path Forwarding)

Cisco.com

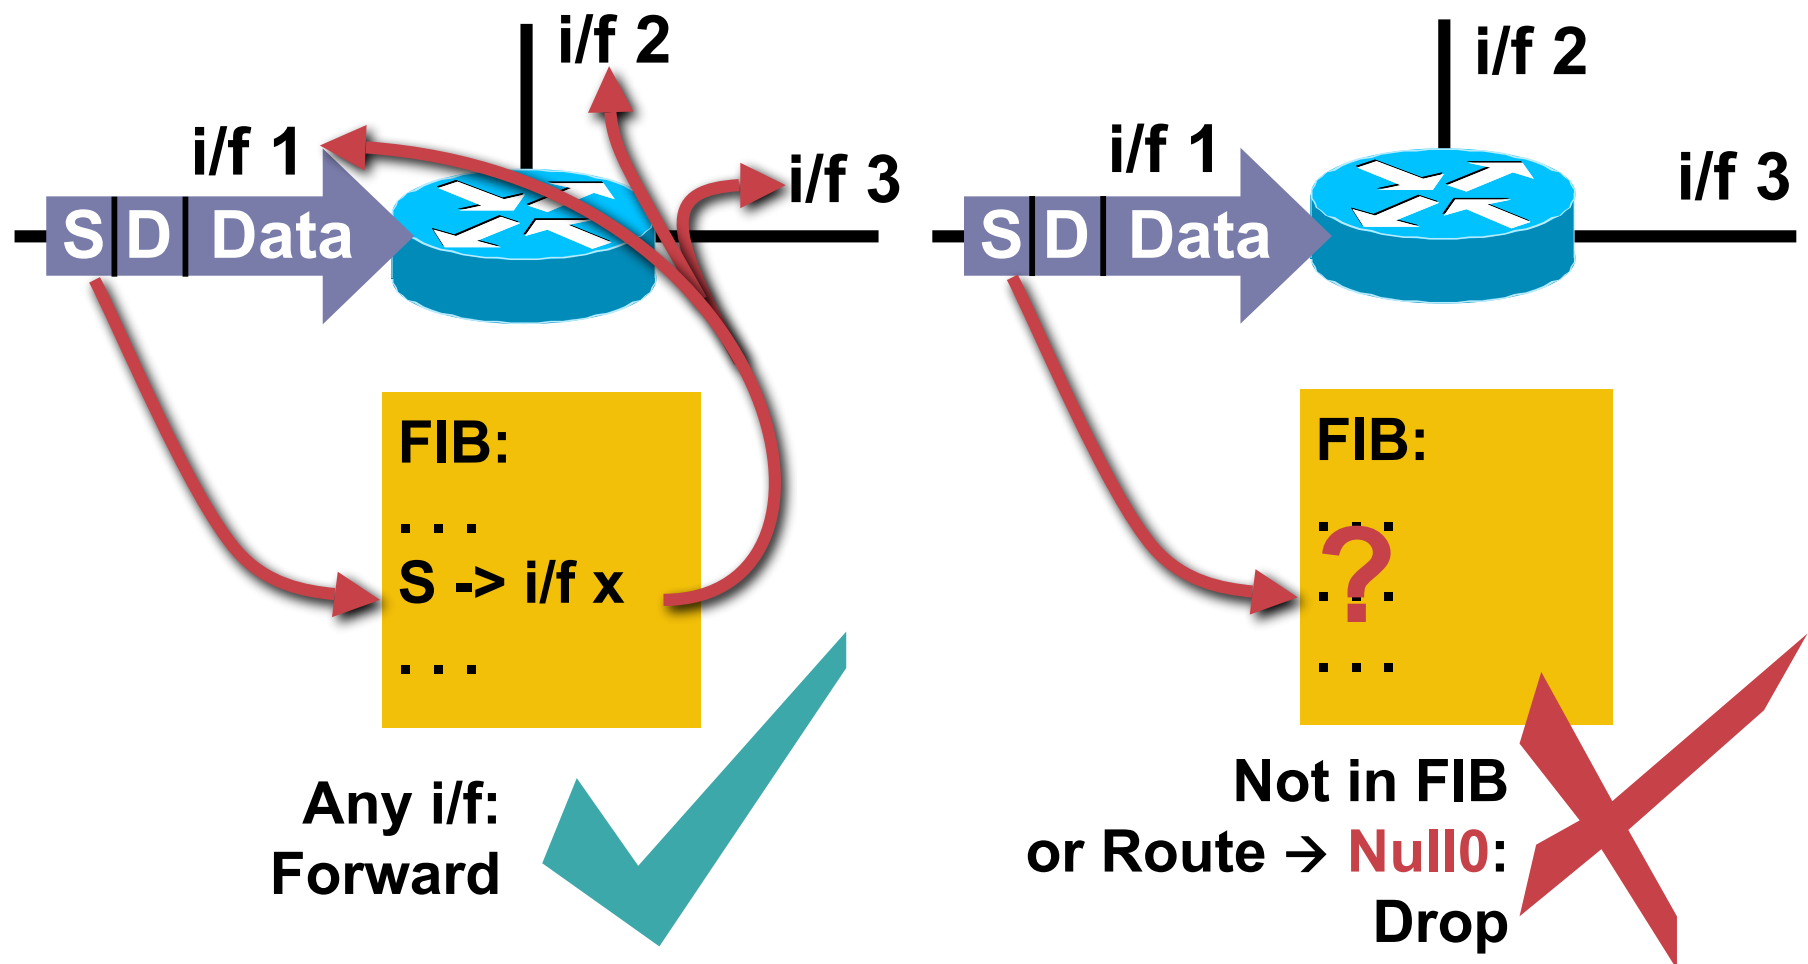
`router(config-if)# ip verify unicast reverse-path`  
or: `ip verify unicast source reachable-via rx`



# Loose uRPF Check

Cisco.com

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



# Source-Based Remotely Triggered Black Hole Filtering

Cisco.com

**Uses the same architecture as destination-Based Filtering and unicast RPF**

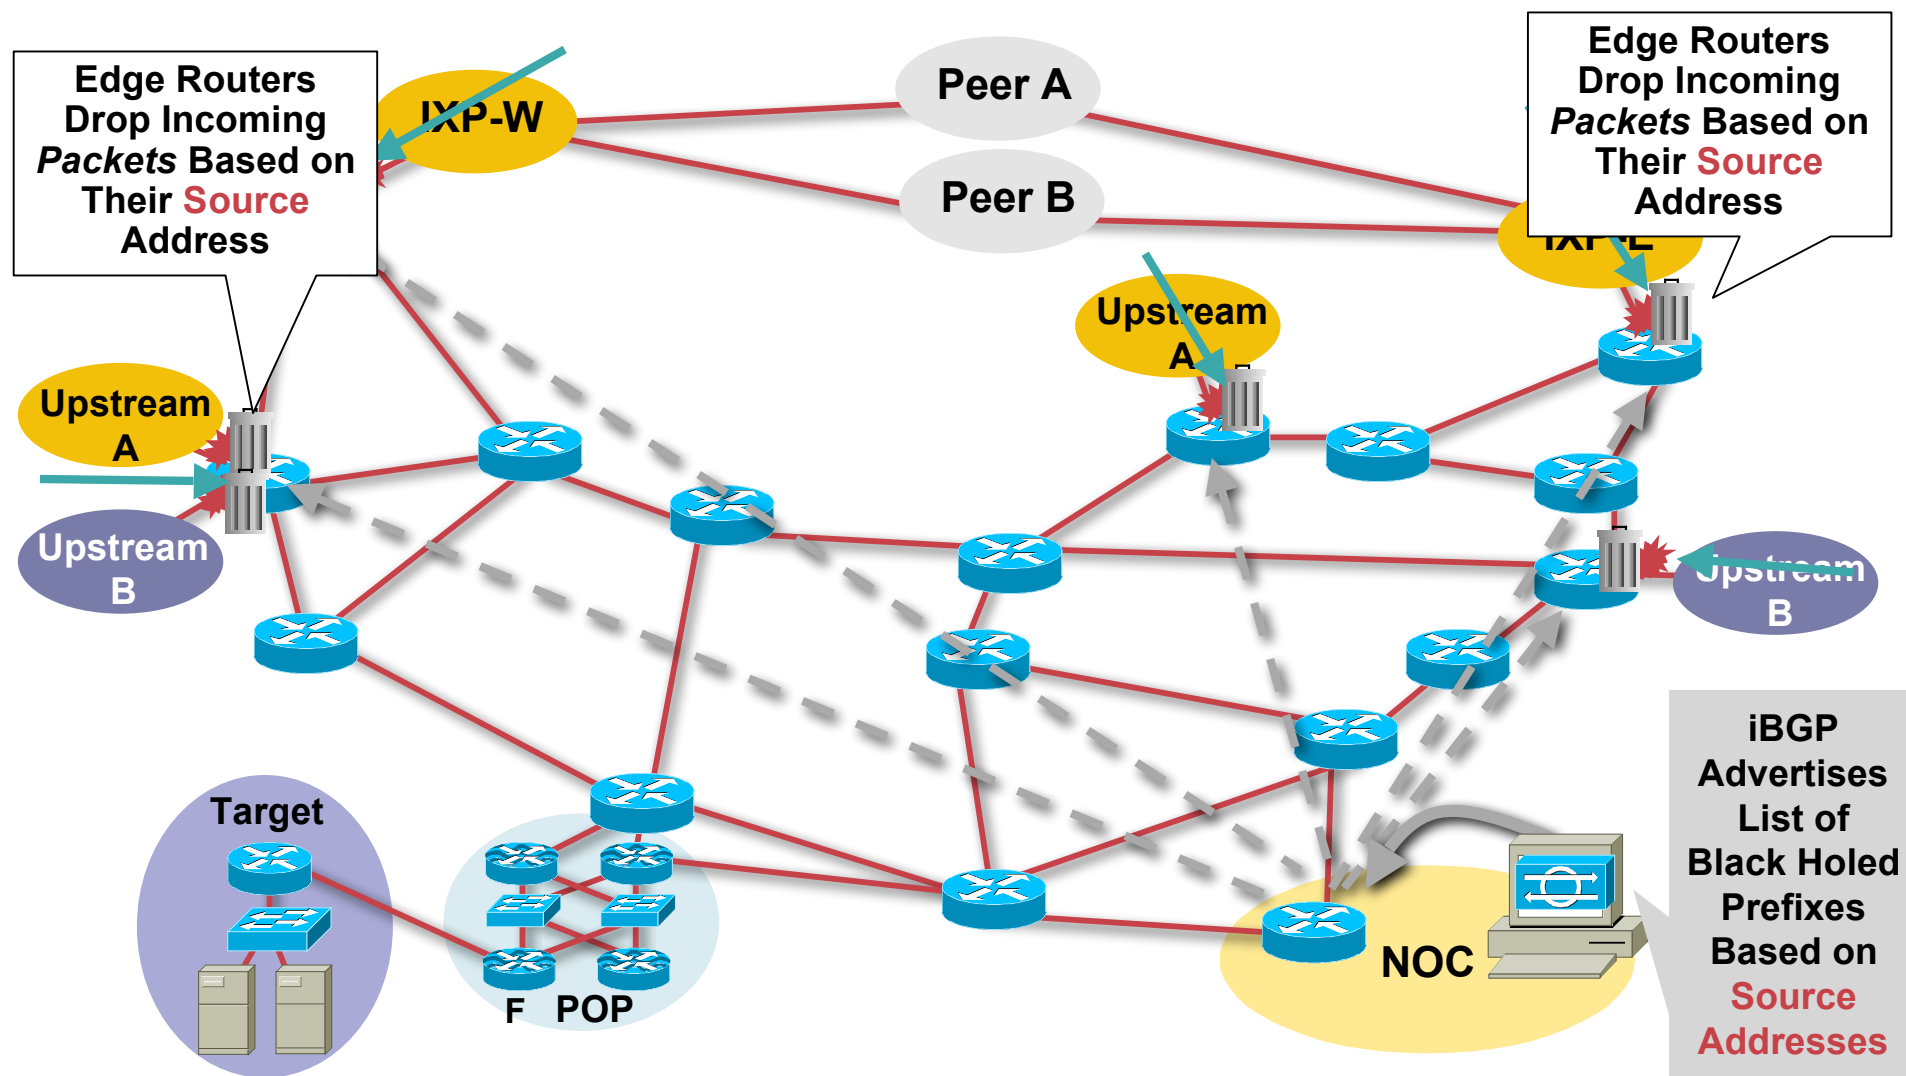
- **Edge routers must have static in place**
- **They also require unicast RPF**
- **BGP trigger sets next-hop—in this case the “victim” is the source we want to drop**

# Source-Based Remotely Triggered Black Hole Filtering

- What do we have?
  - Black Hole Filtering**—if the **destination** address equals Null0, we drop the *Packet*
  - Remotely Triggered**—trigger a prefix to equal Null0 on routers across the network at iBGP speeds
  - uRPF Loose Check**—if the **source** address equals Null0, we drop the *Packet*
- Put them together and we have a tool to trigger a drop for any packet coming into the network whose source or destination equals Null0!

# Customer Is DOSed (After) *Packet Drops* Pushed to the Edge

Cisco.com



# Source Dropping Caution

Cisco.com

- **Caution: you will drop all *packets* with that source**
- **Remember spoofing**

# Community-Based Trigger

Cisco.com

- **BGP community-based triggering allows for more fine-tuned control over where you drop the packets**
- **Three parts to the trigger:**
  - Static routes to Null0 on all the routers**
  - Trigger router sets the community**
  - Reaction router (on the edge) matches community and sets the next-hop to the static route to Null0**

# Why Community-Based Triggering?

Cisco.com

## Allows for More Control on the Attack Reaction

- **Trigger community #1 can be for all routers in the network**
- **Trigger community #2 can be for all peering routers; no customer routers—allows for customers to talk to the DOSed customer within your AS**
- **Trigger community #3 can be for all customers; used 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**



# (Source-Based) RTBH

- **Advantages:**

- No ACL update**

- No change to the router's configuration**

- Drops happen in the forwarding path**

- Frequent changes when attacks are dynamic  
(for multiple attacks on multiple customers)**

- **Limitations:**

- Source detection and enumeration**

- Attack termination detection (reporting)**

- Resource utilization: finite resources**

- Effects all traffic, on all triggered interfaces, regardless of actual intent**

# Attack Example: Mitigation W/ Control Plane

Cisco.com

- **Botnet Controllers are commonly blocked with RTBH – disrupting the communication channel between attacker and his bots**
- **Works better the nearer you are to the bot controller**
- **This prevents the attacker from modifying the attack and lets the target of the attack defend itself – by say changing their IP address in DNS.**
- **Attacks may use multiple controllers, or switch controllers quickly – RTBH lets an operator react just as quickly**

# Attack Example: Mydoom - RTBH was problematic

Cisco.com

- **Mydoom.C**
  - Spread as a Virus Attached to email**
  - DDoS payload attacked on specific date**
  - Opened 14 tcp connections/sec**
  - Made one GET request ( GET/) on each connection**
- **Size and Scope**
  - Over 1 million concurrent attackers**
  - 5 Gigabytes + of traffic peak**
  - Attack lasted for several months**
- **Why RTBH was problematic**
  - Would block sources completely – not allow patching**
  - Proxies would be blocked... including all hosts behind it**
  - Telling who was a good vs bad source just on opening tcp connections is problematic with most telemetry systems**



# Dedicated Devices

# Network IDS/IPS Overview

Cisco.com

- **Content inspection technology**
- **Originally designed to prevent 'intrusion' type attacks**
- **Content filtering can directly apply to DDoS attacks**
  - Some DoS attacks use perfectly legitimate content!**
  - Performance problematic – not optimized for DDoS**
- **IPS assume that you are inline and see both directions...**

# Network IDS/IPS

- **False positives: system mistakenly reports certain benign activity as malicious; also called false alarms**
- **False negatives: system does not detect and report actual malicious activity**
- **For many, false positives are the bane of IDS technology**
- **Additionally, you require a signature in order to stop the attack**

# Firewalls

## Modern Stateful Firewall: The Security Keystone

Cisco.com

### What It Is

- Sometimes called a hybrid
- Combines features of other firewall approaches such as:
  - Access control lists
  - Application-specific proxies/inspections
  - Stateful inspection
- Plus features of other devices:
  - Web (HTTP) cache
  - Specialized servers
    - SSH, SOCKS, NTP
  - Most include VPN, some include IDS

### Pros and Cons

- Pro: Pretty high performance in modern implementations
- Pro: Application layer gateway services provide application security while resolving the NAT issue
- Con: Does not provide complete session termination, as would a full proxy
- Con: Actively tracks the state of incoming connections—a DoS issue

# Load Balancers

- **Stateful by design**
  - One to many
  - Distributes capacity
- **Some have DDoS features being added**
  - Syn Cookies
  - Basic Behavior Recognition
- **Require inline/symmetry to work**
- **Since load balancers are stateful, problems can occur when an attack uses a protocol without inherent state (e.g. UDP)**
  - Not so good with DNS!



# Given Everything Said, What Remains?

Cisco.com

- **Raise the bar! Stop ONLY bad traffic**
- **In asymmetric environments, especially across peers, packet spoofing is still problematic**
- **Detection of exactly “who” is attacking is problematic**
- **Doing all this in the core requires specialized hardware, which has scaling and availability problems**

# Packet Scrubbing Technologies

Cisco.com

**Shunting the packets**

**Scrubbing the packets**

# Formal Requirements for a Core Security Device

Cisco.com

- **Need to avoid state**

Constant state tracking leaves us vulnerable to DDoS attacks

- **Doesn't rely on signatures**

If I get an attack with no signature, I cannot block it

Possibly can use signature-like filters, however, after the fact

- **Doesn't have to be in-line when it isn't needed**

- **Scales easily**

- **Doesn't require traffic symmetry**

The Internet is a very asymmetrical place!

# Define Packet Scrubbing

- **Packet scrubbers have the following core characteristics**

- **Classification**
- **Source verification**
- **Measurement**
- **Recognition**
- **Enforcement**

**A reference architecture for packet scrubbing is presented later**

# Packet Scrubbers: Details

Cisco.com

- **Scale by using traffic shunting and anycast**
- **Packet scrubbing**
  - 1) **Validate incoming traffic to make sure it comes from the source IPs that are in the SRC IP field of the *Packet***
  - 2) **Evaluate these validated sources against a baseline and then recommend either further processing or dropping for sources that misbehave**
- **Availability is the key security metric!**
  - Pad thresholds to reduce likelihood of false positives**
  - Better to stop 90% of the attack than 110%**
  - Implement content filtering only when necessary**

# Packet Scrubbing Via Shunts

Cisco.com

- **Advantages:**
  - Not on critical path during normal operation**
  - Anomaly-based detection with base lining**
  - Optimized for high-performance blocking**
  - Is resistant to state limitations of most other devices**
- **Limitations:**
  - Not designed to stop *single-packet* attacks**
  - Inherent is an assumption of a “destination” being protected**
  - Resource utilization: finite resources in the scrubber complex**
  - Requires up-front network engineering to implement**

# Traffic Shunts

- **Intercept and shunt traffic to the mitigation device—the “scrubber”**
- **Return good traffic back to the customer**
- **Need to avoid forwarding loops—means some sort of tunneling**


# Shunts in the Data Center

- **All devices on the same subnet**
  - Either scrubber-driven or configured in router**
  - May use remotely triggered shunt trick**
    - All traffic in core to target goes to the scrubber**
- **Optionally, you can use VLANs to avoid loops**
  - Bypassing the “modified” router is trivial with VLANs and .1Q trunking**



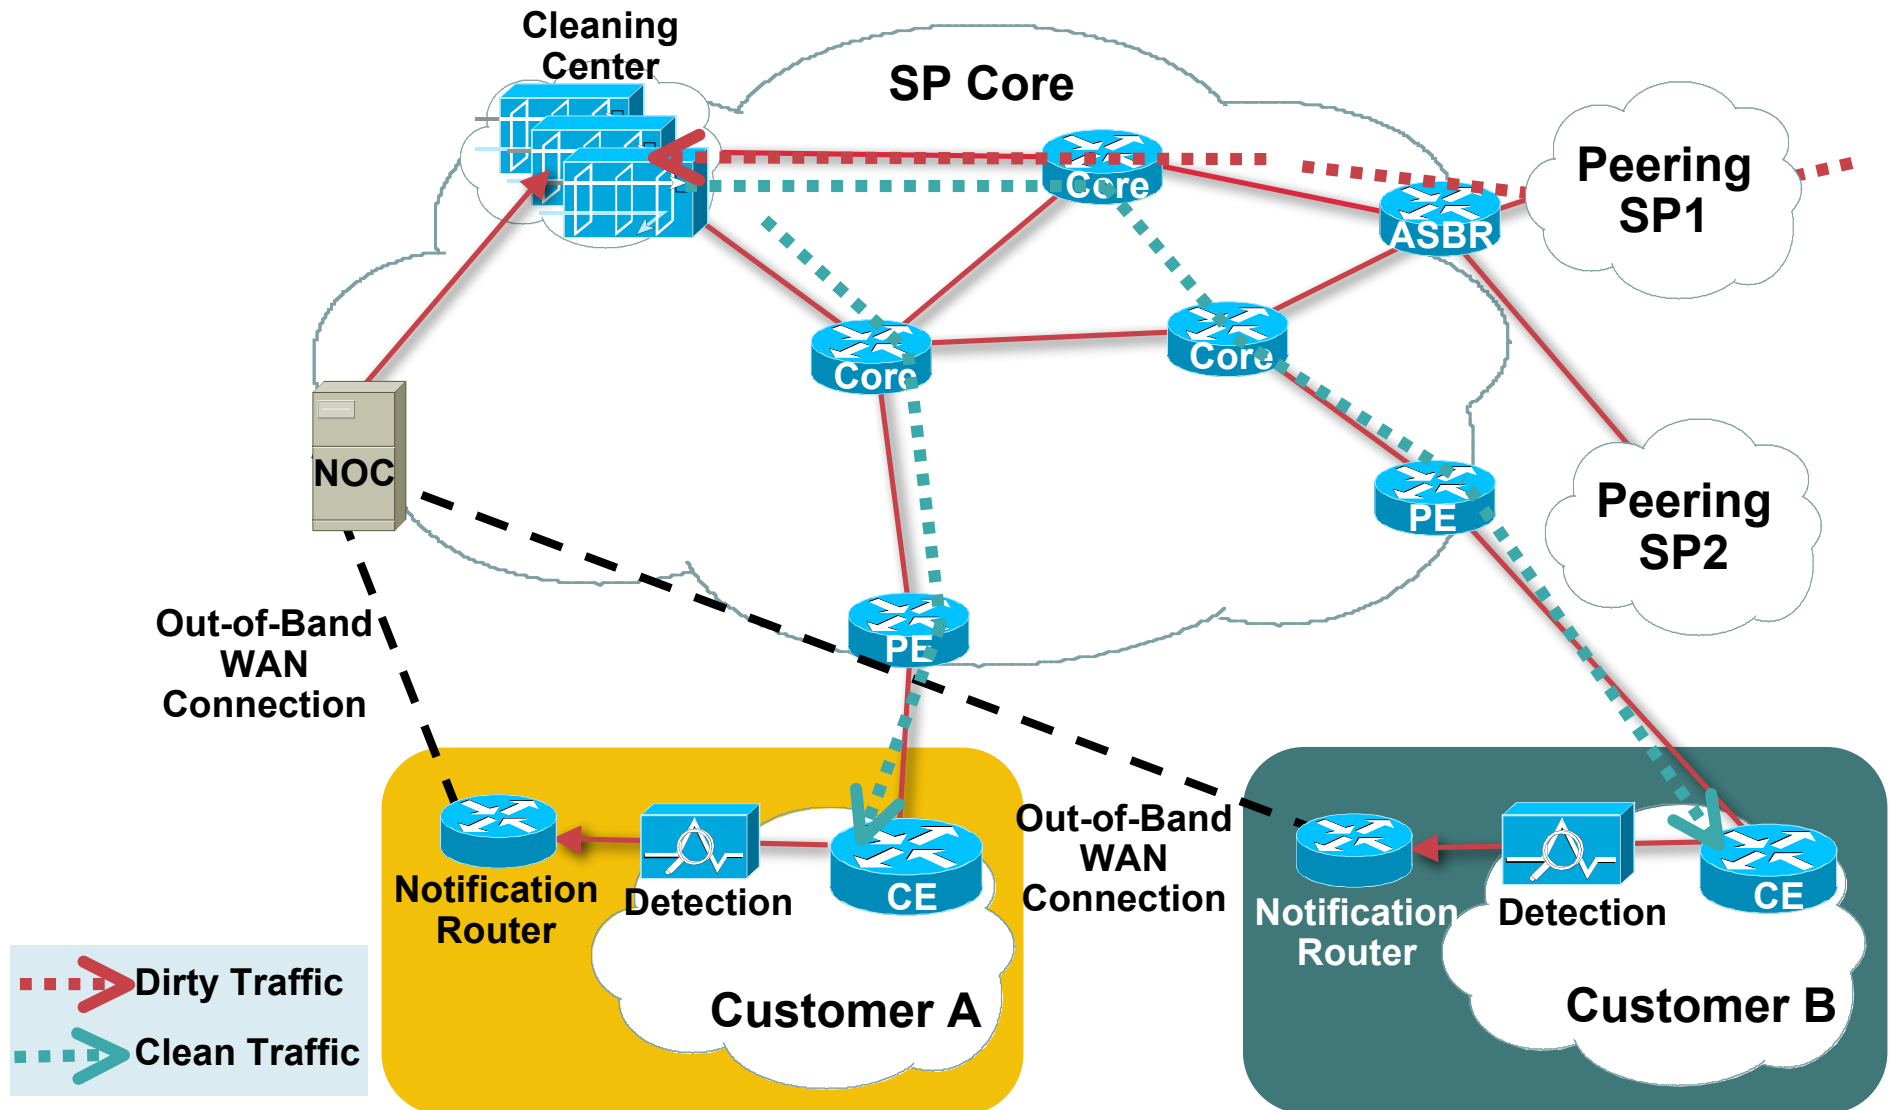
# Hosting/SP Data Center

Cisco.com



# Packet Scrubbing in the Core: The Cleaning Center

Cisco.com



# Shunts in the Backbone: Scrubbing Centers

Cisco.com

- **Question is: how many?**
  - Most national providers have decided to start with two**
  - Geographic redundancy**
  - Adequate incoming bandwidth in key locations**
  - Limit the backhaul of traffic across expensive links**
- **Once you decide on where, then the hard part starts**
  - Getting traffic to and from the center(s)**

# Getting Traffic to and from the Scrubbing Center

Cisco.com

- **GRE**
  - Requires GRE start**
  - Requires GRE endpoint**
- **MPLS VPN shunt**
  - Requires a single MPLS VPN for injection, for all customers**
  - Easy config, no performance hits**
  - Works if peering PEs also connect potential victims**
- **MPLS proxy egress LSP**
  - Requires MPLS, but not VPN.**
  - Very simple, easy config**
  - But: Does not work if peering PEs connect potential victim**



## **Option 1: IP Core and GRE injection**

# IP Core and GRE Injection

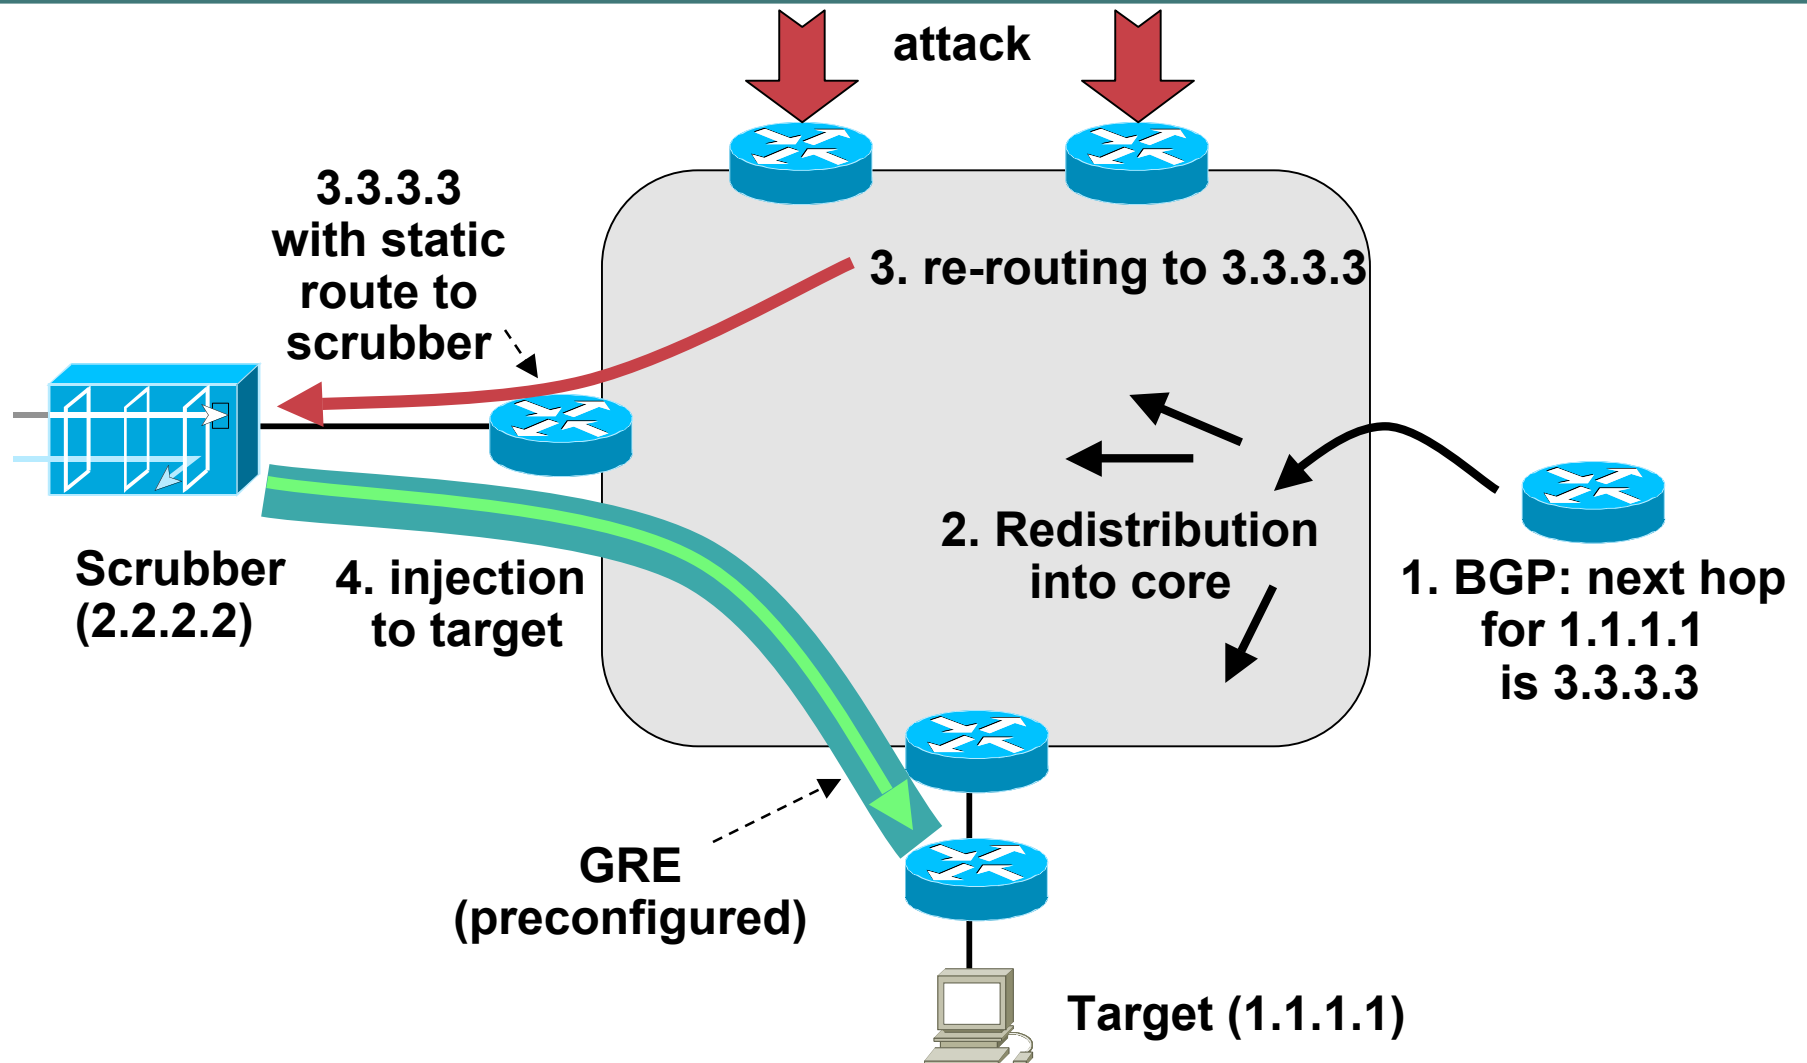
Cisco.com

- **Core routes “target” IP to the scrubber**  
Either scrubber driven or configured in router  
May use remote triggered shunt trick  
→ All traffic in core to target goes to the scrubber
- **Injection into GRE tunnel**  
Bypassing the “modified” core routing  
GRE starts on scrubber, terminates on CPE, which has  
“clean” routing to target

# IP Core and GRE Injection


## Triggered by a router or other device

Cisco.com



# mGRE Injection: Simplified Hub Config!

Cisco.com





# IP Core + GRE

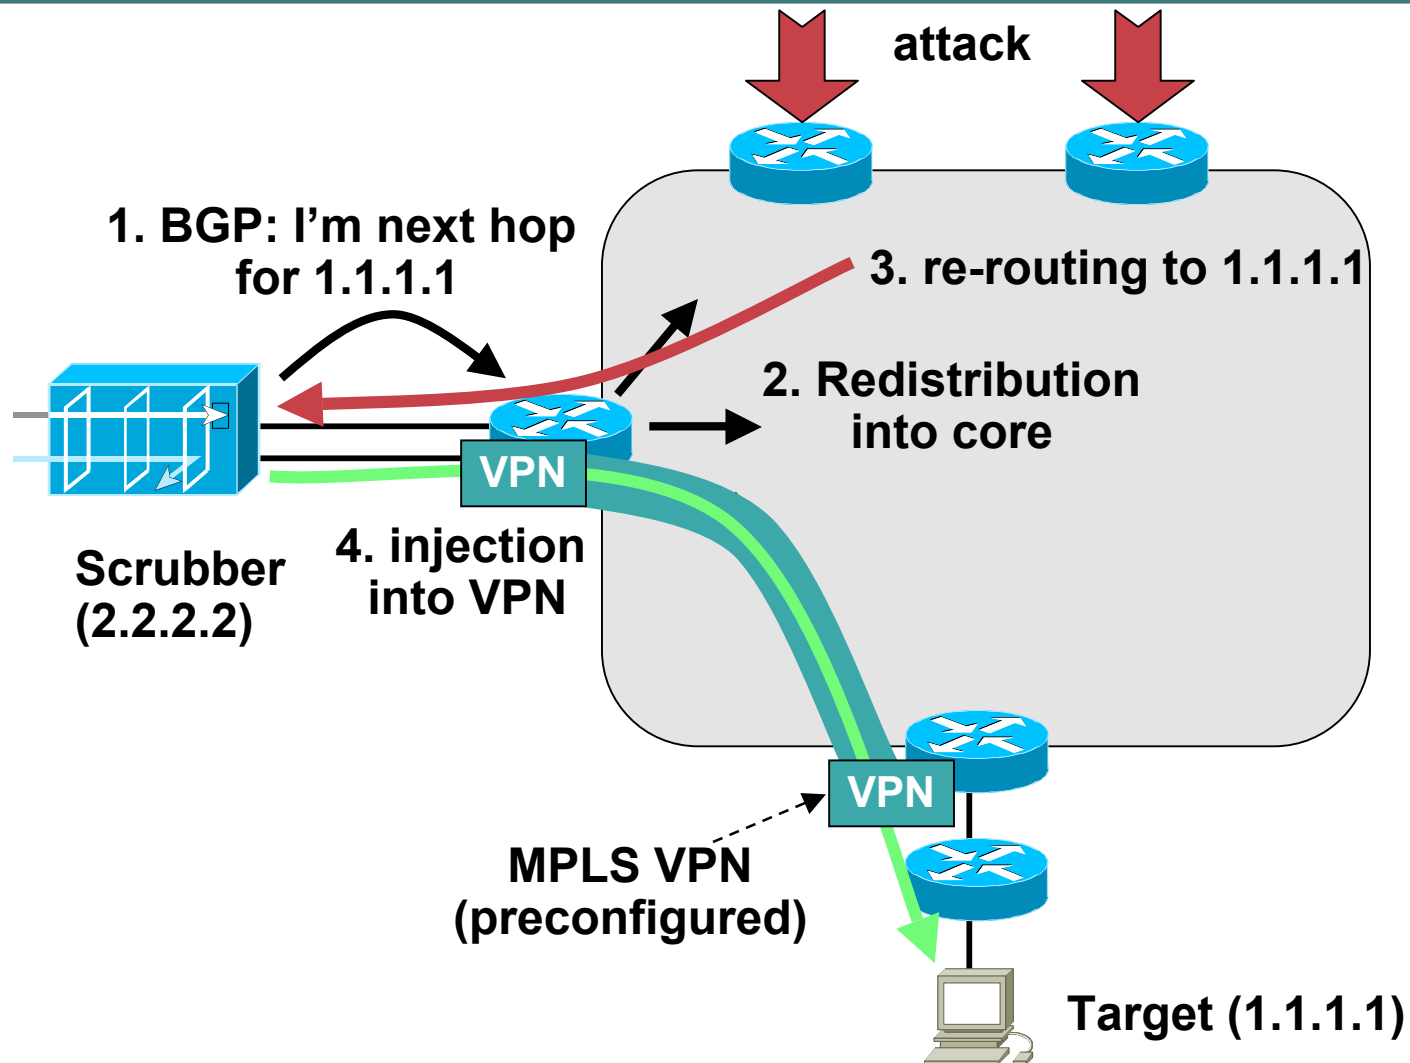
- **Easy to deploy:**
  - Core remains untouched**
  - GRE tunnel invisible to core**
- **GRE might have performance impact**
- **GRE endpoint required:**
  - PE?**
  - CE?**



## **Option 2: MPLS VPN Shunt**

# MPLS VPN Shunt

Cisco.com



# MPLS VPN Shunt

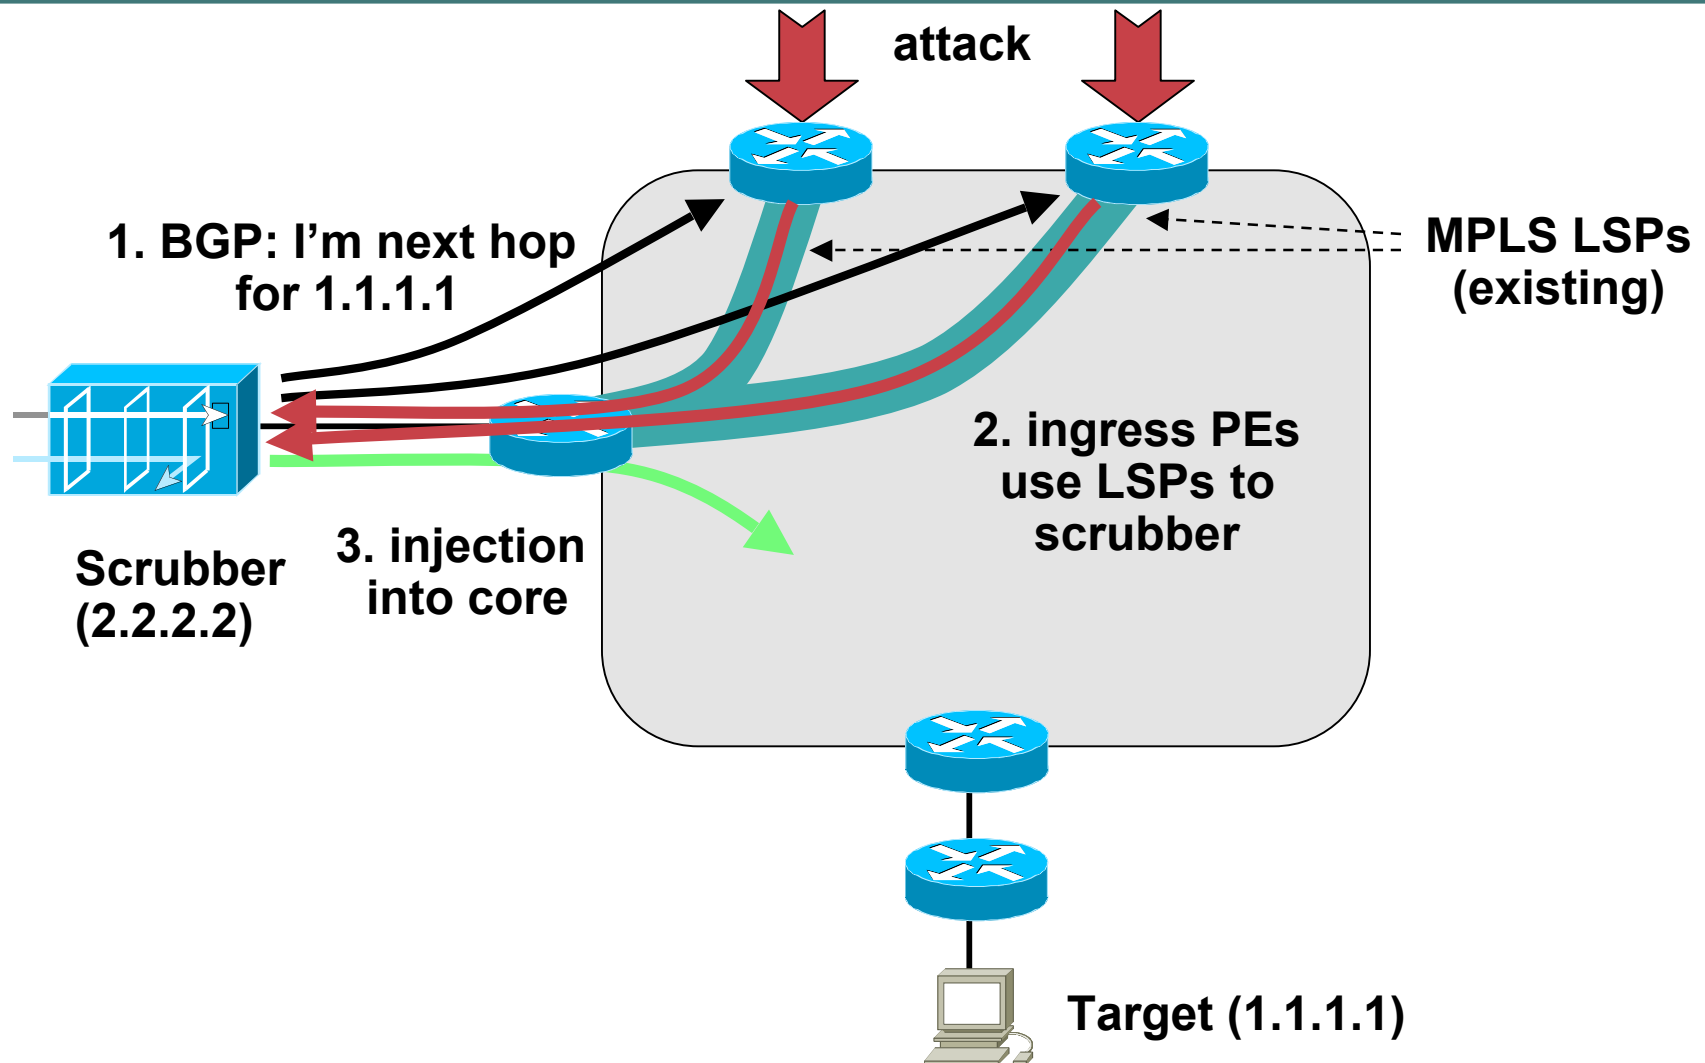
- **Easy to deploy:**
  - Core remains untouched, injection VPN pre-configured**
  - Need one VPN only**
  - VPN invisible to core**
- **No performance impact**
- **No need to touch CPE**
- **But: MPLS VPN required on core**



## **Option 3: MPLS Proxy Egress LSP Shunt**

# MPLS LSP Shunt

Cisco.com



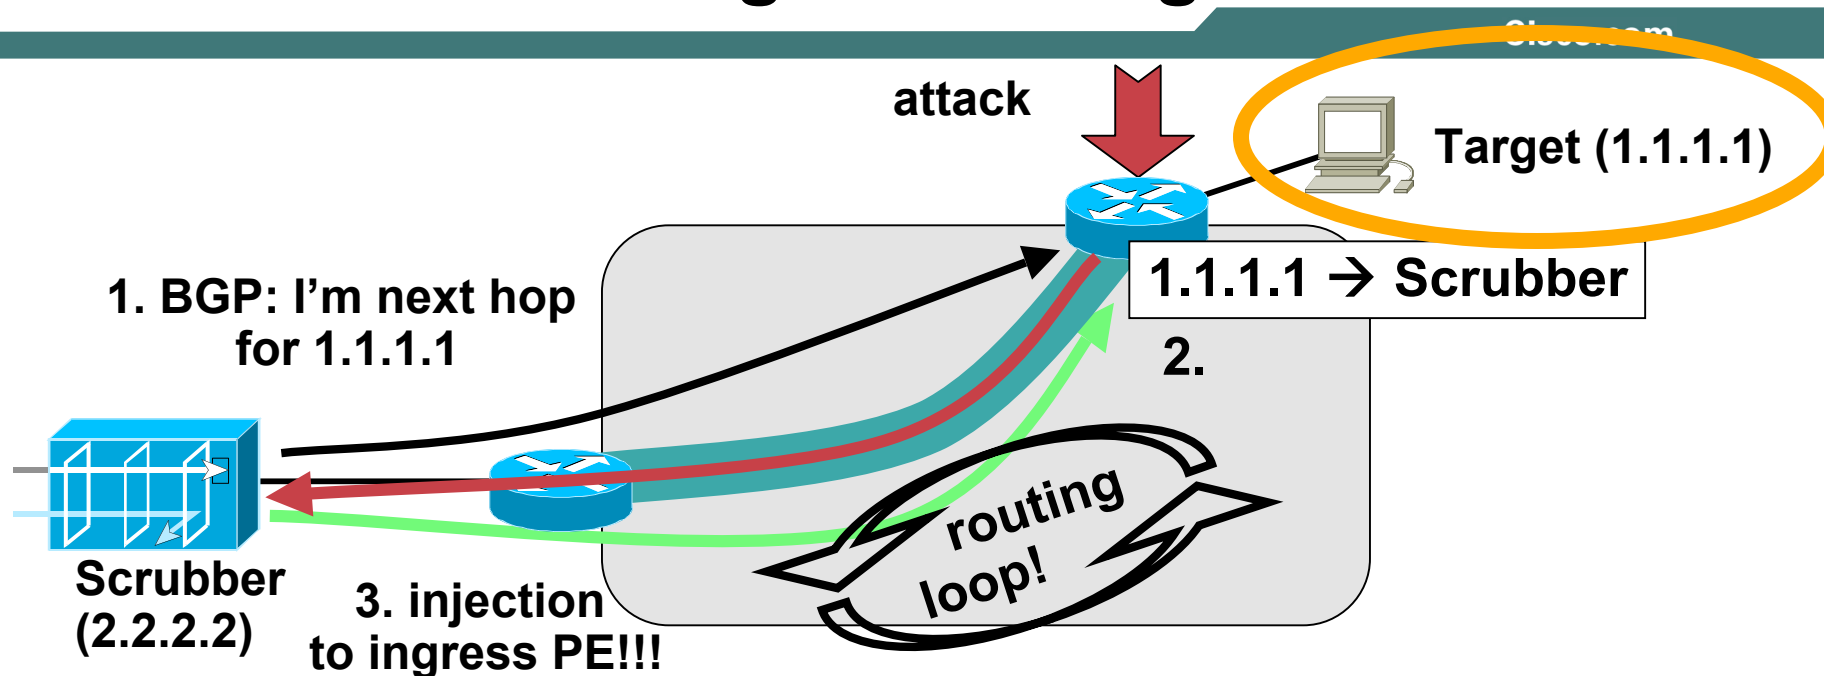
# MPLS Proxy Egress LSP Shunt

Cisco.com

- **Preparation: scrubber loopback is in IGP → LSPs exist to the scrubber**
  - Filtering PE acts as “proxy”, popping the label, sending IP only
- **Scrubber peers with ingress PEs (and only with ingress PEs!)**
  - Scrubber is iBGP “next hop” for victim (on ingress PEs).
    - Network sends traffic to the victim to the iBGP next hop
    - Follows the LSPs to the scrubber (pre-established)
- **Filtering PE does not route this traffic, but label switches**
  - Routing table unchanged, still pointing to victim
  - Scrubber sends cleaned traffic to the filtering PE
  - Traffic goes onto LSP to egress PE

# MPLS LSP Shunt:

## Cannot have “Targets” on Ingress PE!



- Ingress PE: Routing to “target” now → Scrubber  
Cannot reach directly connected Target any more!
- Careful! Routing loop with several Gig of traffic!!!!
- Workaround: PE based communities (a bit risky...)



# MPLS LSP Shunt

- **Using existing LSPs from ingress PEs to scrubbing center**  
Only config: static route on filtering PE to scrubber loopback
- **Scrubber iBGP peers with all ingress PEs**
- **Other PEs (including egress PE) maintain normal routing to target**
- **No performance impact**
- **No need to touch CPE**
- **But: ingress PE must not connect the victim!**  
Otherwise: Routing loop. Workarounds:  
Use MPLS VPN method! or communities per PE
- **MPLS required, but not MPLS VPN**

# Shunting and DDoS: Summary

Cisco.com

- **Many options for deployment:**
  - One or several scrubbing centers**
    - For several centers: Static or anycast**
    - Trigger by scrubber, Arbor, other devices (BGP speakers)**
    - Either re-direction or injection must bypass normal routing table**
    - Options: GRE, MPLS VPN, MPLS LSPs, other tunnels, ...**
- **Highly scalable**
  - Several scrubbing centers**
  - Several scrubbers in each**



# A Reference Architecture for Mitigation

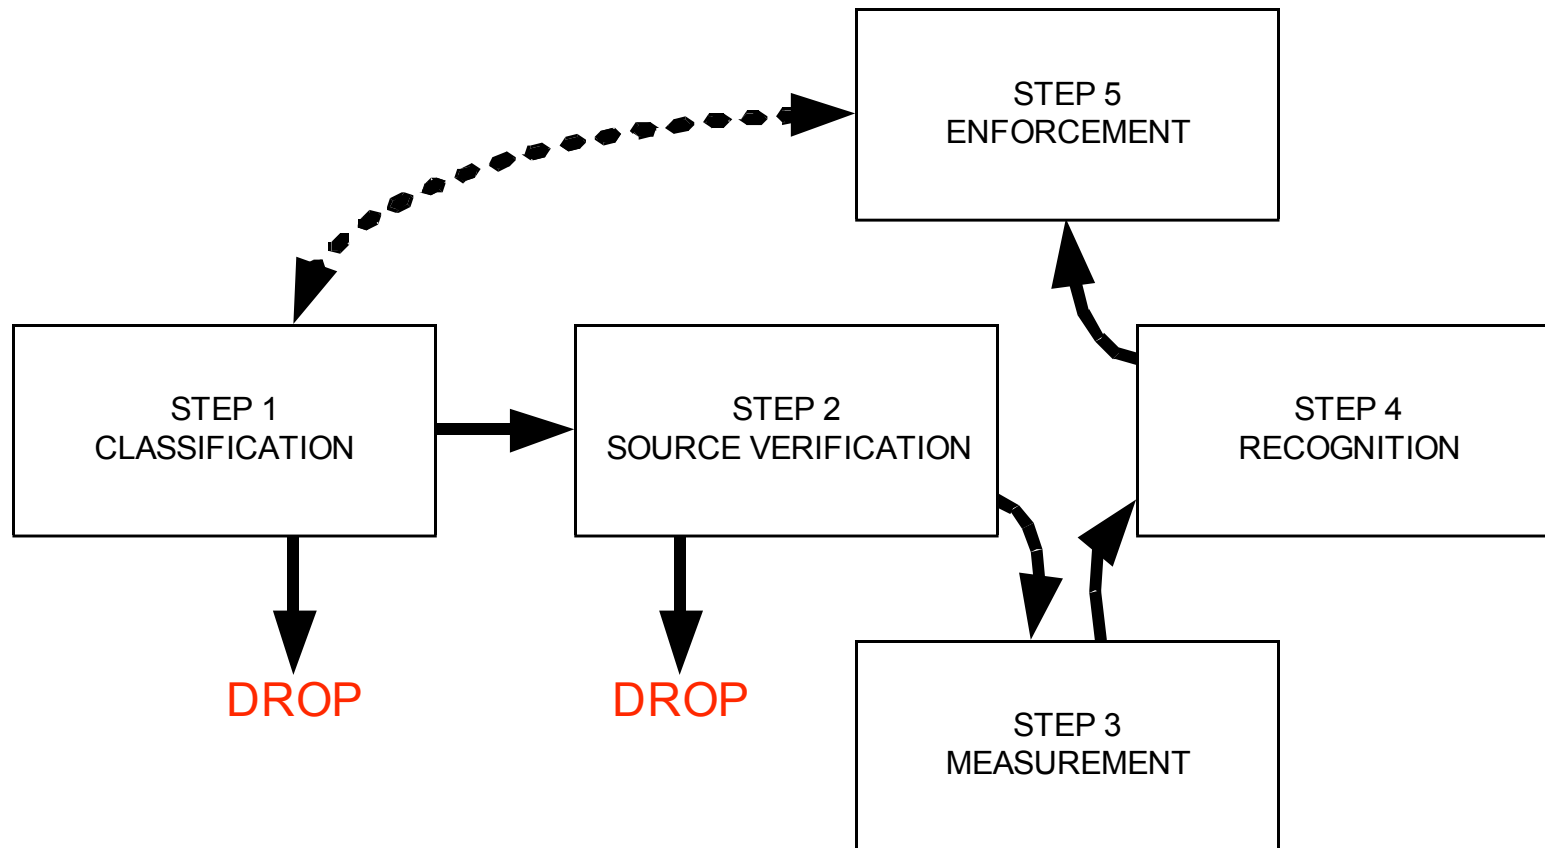
# Mitigation Reference Architecture

Cisco.com

- **Alternative to a “Product Pitch”**
- **Goal is to define a standard for DDoS technologies to be judged**
- **This is a implementation guide based on the requirements outlined in the previous section**
- **A five step architecture that addresses all the functional components in modern DDoS mitigation technologies**

# Mitigation Reference Architecture

Cisco.com



# Mitigation Reference Architecture

Cisco.com

- **Step 1: Classification**

**Classifies traffic for future inspection**

**Implements 'inline' enforcement:**

- **blocking malicious sources**
- **blocking malicious content**

- **Step 2: Source Verification**

**Defeats spoofed attacks**

**Must operate at 'line rate' or will be vulnerable to DoS**

# Mitigation Reference Architecture

Cisco.com

- **Step 3: Measurement**  
Counts the traffic
  - Can be on the resource itself
  - Can be on an offline device like a NetFlow collector
- **Step 4: Recognition**  
Uses input from measurement to determine who is an attacker  
Suggests to Enforcement who to block
- **Step 5: Enforcement**  
Informs classifier what to block  
Gets information back from classifier when attack is over



# Future Developments



# A Look at How We Got Here

Cisco.com

- **End-to-end principle**

**Network devices should not replicate end system functions**

**Corollary: Every IP address on the Internet is supposed to be able to communicate to every other IP address**

**Controlling this behavior is the key challenge**

- **Attacks are getting smarter based on economically motivated factors**

- **Attack frequency is determined by two factors:**

**Background malicious activity**

**Financial opportunity**

# The Right and Wrong Way to Control Traffic

Cisco.com

- Today there is a natural, accepted mechanism for blocking traffic—filtering ports!
- This leads to a situation where applications avoid filtering by using common ports of other applications (e.g. Port 80 for P2P)
- Now the security professional decides to invest in deep *Packet* inspection technologies
- This will lead to a situation where applications start using encryption to avoid detection
- Imagine everything encrypted over 443!

This isn't as far away as you might think.

# Operational Security and Encryption

Cisco.com

- **Two ways today to mitigate attacks**

**Packet content evaluation:** evolving to content filtering

**Flow-behavior analysis:** and anomaly recognition

- **Encryption prevents effective use of either mechanism**

**How can you evaluate the content—classify it, police it, etc.**

**All encrypted traffic tends to look alike, behaviorally**

**Since decryption usually requires CPU, it in itself usually makes a system more vulnerable**

# An Alternative Model to Port Blocking

Cisco.com

- **Classify, and mark *packets***
- **Once marked, treat traffic in three to four broad buckets:**
  - Control**
  - High quality**
  - Best effort**
  - Bulk traffic**
- **DiffServ as a strategic security tool!**
- **Classification is the key!**



# Summary

# Summary

- **Implementing DDoS mitigation in large, complex environments is non trivial**
- **A mitigation system requires both security and networking resources – can't be done with one alone!**
- **A variety of tools are available- these must be integrated into your Operational environment**



# Thank You!

**Questions?**

# References

Cisco.com

- **ISP Essentials**

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

- **Team Cymru**

<http://www.cymru.com/>

- **Nanog Security Curriculum**

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

- **Cisco DDos Protection Solution**

[http://www.cisco.com/cdc\\_content\\_elements/networking\\_solutions/service\\_provider/ddos\\_protection\\_sol/](http://www.cisco.com/cdc_content_elements/networking_solutions/service_provider/ddos_protection_sol/)

- **DDos Links**

<http://www.honeypots.net/incidents/ddos-mitigation>

- **ICMP Backscatter Traceback**

<http://www.secsup.org/Tracking/>

- **Customer Triggered Blackholes**

<http://www.secsup.org/CustomerBlackHole/>



