Protecting Your Infrastructures: Current Deployment Practices & Latest Trends

Merike Kaeo
Double Shot Security
merike@doubleshotsecurity.com

Agenda

- Network Threat Model
  - Threat Actions (i.e. Attacks)
  - Threat Consequence
- Latest Security Risk Mitigation Techniques
  - Proactive Security Measures
    - Filtering Practices
    - Routing Security
  - Mitigating DDoS Risk / Impact
    - Sinkholes
    - Remotely Triggered Blackhole Routing
    - Backscatter Traceback
    - Packet Scrubbing

Network Threat Model

Consider Attack Sources

- Passive vs Active
  - Writing and/or reading data on the network
- On-Path vs Off-Path
  - How easy is it to subvert network topology?
- Insider or Outsider
  - What is definition of perimeter?
- Deliberate Attack vs Unintentional Event
  - Configuration errors and software bugs are as harmful as a deliberate malicious network attack

Passive vs Active Attacks

- Passive Attacks
  - Eavesdropping
  - Offline cryptographic attacks
- Active Attacks
  - Replay
  - Man-In-The-Middle
  - Message Insertion
  - Spoofing (device or user)
  - Denial of Service
  - Protocol specific attacks

Threat Consequences

- Unauthorized Disclosure
  - circumstance or event whereby entity gains access to data for which it is not authorized
- Deception
  - circumstance or event that may result in an authorized entity receiving false data and believing it to be true
- Disruption
  - circumstance or event that interrupts or prevents the correct operation of system services and functions
- Usurpation
  - circumstance or event that results in control of system services or functions by an unauthorized entity

How Can Router Infrastructure Threats Be Realized ?

- Protocol error
  - Routing protocol itself
  - TCP issues for BGP
- Software bugs
  - Is it a bug or feature ?
- Active attack
  - Target control/management plane
  - Target data plane
  - More probable than you think !
- Configuration mistakes
  - Most common form of problem

Routing Threat Consequence

- Traffic is sent along invalid path
- Traffic is dropped

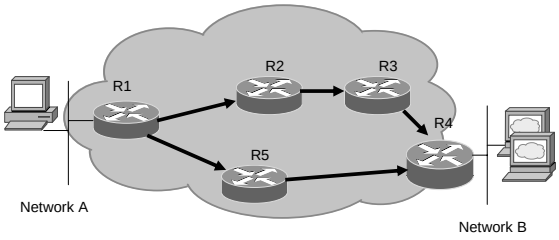

What Can We Do To Protect The Routing Infrastructure ?

- Understand the Problem (Risk Analysis)
- Establish an Effective Routing Infrastructure Security Policy
  - physical security
  - logical security
  - control/management plane
  - routing plane
  - data plane
- Have Procedures In Place For Incident Response
  - procedures for assessing software vulnerability risk
  - auditing configuration modifications

How Mitigate Most Threats?

- Secure end-system hosts
- Limit access to network
- Authenticate (device vs user)
- Based on claimed identity, allow (authorize) access to specific resources
- Audit network traffic
- Use confidentiality if needed

Agenda

- Network Threat Model
  - Threat Actions (i.e. Attacks)
  - Threat Consequence
- Latest Security Risk Mitigation Techniques
  - Proactive Security Measures
    - Filtering Practices
    - Routing Security
  - Mitigating DDoS Risk / Impact
    - Sinkholes
    - Remotely triggered Blackhole Routing
    - Backscatter Traceback
    - Packet Scrubbing

Proactive Security Measures

Know Your Traffic

- Data Plane
  - Traffic going through the router
- Management Plane
  - Traffic used to monitor and log information
  - Traffic used to manage device
- Control Plane
  - Traffic specific to routing protocols

Filter Classification

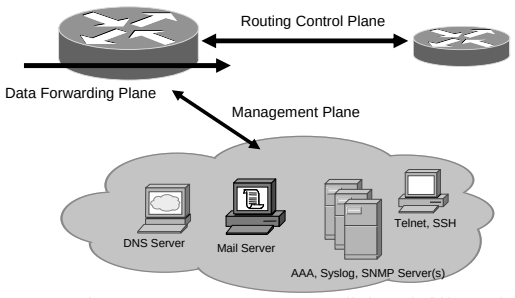

Filtering Deployment Considerations

- How does the filter load into the router? Does it interrupt packet flow?
- How many filters can be supported in hardware? In software?
- How does filter depth impact performance?
- How do multiple concurrent features affect performance?

Data Plane (Packet) Filters

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

Management Plane Filters

- Define Explicit Access To/From Management Stations
  - SNMP, Syslog, TFTP, NTP, AAA Protocols, DNS, SMTP, SSH, Telnet, etc.
- Authenticate Access
- Think of Using Out-of-Band Management Network

Control Plane (Routing) Filters

- Filter traffic destined TO your core routers
- Develop list of required protocols that are sourced from outside your AS and access core routers
  - Example: eBGP peering, GRE, IPSec, etc.
  - Use classification filters as required
- Identify core address block(s)
  - This is the protected address space
  - Summarization is critical for simpler and shorter filter lists

BGP Prefix Filtering

- All BGP Prefixes coming into your network and leaving your network need to be filtered to enforce a policy.
- The problem is most ISPs are not:
  - Filtering Comprehensively
  - Filtering their customer's prefixes
  - Filtering prefixes going out of their network.

Example: No Prefix Filtering

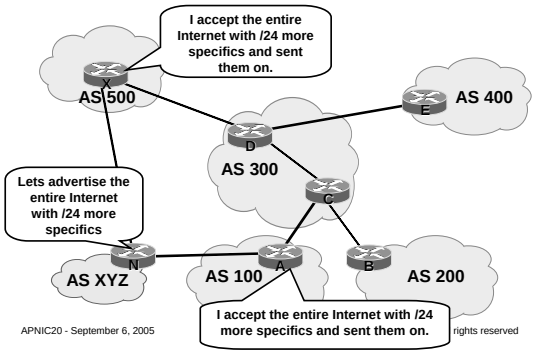

Impact of No Prefix Filtering

- AS 7007 Incident (1997) was very visible case of problem.
- Key damage are to those ISPs who pass on the garbage.
- Disruption, Duress, and Instability has been an Internet wide effect.

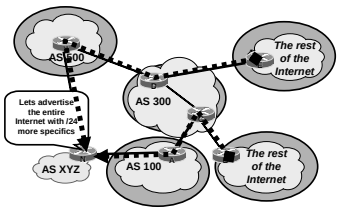

Where to Prefix Filter ?

- Customer's Ingress/Egress
- ISP Ingress on Customer (may Egress to Customer)
- ISP Egress to Peer and Ingress from Peer
- Peer Ingress from ISP and Egress to ISP

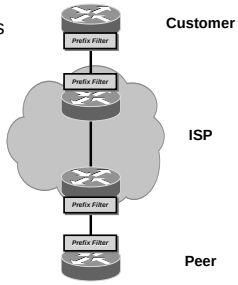

Receiving Customer Prefixes

- Configuration example on upstream:

```
router bgp 100
  neighbor 123.123.6.1 remote-as 101
  neighbor 123.123.6.1 prefix-list customer in
  !
ip prefix-list customer permit 121.60.0.0/2
ip prefix-list customer deny 0.0.0.0/0 le 32
```

Peering With Other ISPs

- Similar to eBGP customer aggregation except inbound prefix filtering is rarely used (lack of global registry)
- Use maximum-prefix and prefix sanity checking instead
- Still use per-neighbor passwords!

Example of ISP-Peers (peer group)

neighbor nap peer-group
neighbor nap description for peer ISPs
neighbor nap remove-private-AS
neighbor nap version 4
neighbor nap prefix-list sanity-check in
neighbor nap prefix-list cidr-block out
neighbor nap route-map nap-out out
neighbor nap maximum prefix 30000

Example of ISP Peers (route-map)

route-map nap-out permit 10
match community 1 ; customers only
set metric-type internal ; MED = IGP metric
set ip next-hop peer-address ; our own

Peer Groups for NAPs Sanity Check Prefix List

```
# FIRST - FILTER OUT YOUR IGP ADDRESS SPACE!!
# deny the default route
ip prefix-list sanity-check seq 5 deny 0.0.0.0/32
# deny anything beginning with 0
ip prefix-list sanity-check seq 10 deny 0.0.0.0/8 le 32
# deny masks > 20 for all class A nets (1-127)
ip prefix-list sanity-check seq 15 deny 0.0.0.0/1 ge 20
# deny 10/8 per RFC1918
ip prefix-list sanity-check seq 20 deny 10.0.0.0/8 le 32
```

Peer Groups for NAPs Sanity Check Prefix List

```
# reserved by IANA - loopback address
ip prefix-list sanity-check seq 25 deny 127.0.0.0/8 le 32
# deny masks >= 17 for all class B nets (129-191)
ip prefix-list sanity-check seq 30 deny 128.0.0.0/2 ge 17
# deny net 128.0 - reserved by IANA
ip prefix-list sanity-check seq 35 deny 128.0.0.0/16 le 32
# deny 172.16 as RFC1918
ip prefix-list sanity-check seq 40 deny 172.16.0.0/12 le 32
```


Peer Groups for NAPs Sanity Check Prefix List

```
# class C 192.0.20.0 reserved by IANA
ip prefix-list sanity-check seq 45 deny 192.0.2.0/24 le 32
# class C 192.0.0.0 reserved by IANA
ip prefix-list sanity-check seq 50 deny 192.0.0.0/24 le 32
# deny 192.168/16 per RFC1918
ip prefix-list sanity-check seq 55 deny 192.168.0.0/16 le 32
# deny 191.255.0.0 - IANA reserved
ip prefix-list sanity-check seq 60 deny 191.255.0.0/16 le 32
```


Peer Groups for NAPs Sanity Check Prefix List

```
# deny masks > 25 for class C (192-222)
ip prefix-list sanity-check seq 65 deny 192.0.0.0/3 ge 25
# deny anything in net 223 - IANA reserved
ip prefix-list sanity-check seq 70 deny 223.255.255.0/24 le 32
# deny class D/Experimental
ip prefix-list sanity-check seq 75 deny 224.0.0.0/3 le 32
```

Prefix Filter Bogons and RIR Blocks

- Templates available from the Bogon Project:
  - http://www.cymru.com/Bogons/index.html
- 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

Other BGP Security Techniques

- BGP Community Filtering
- MD5 Keys on the eBGP and iBGP Peers
- Max Prefix Limits
- Prefer Customer Routes over Peer Routes (RFC 1998)
- BGP Dampening with RIPE-299
- BTSH (i.e. TTL Hack)

Mitigating DDoS Risk / Impact

Sinkhole Routers/Networks

- Sinkholes are the network equivalent of a ***honey pot***, also commonly referred to as a ***tar pit***, sometimes referred to as a ***blackhole***.
  - Router or workstation built to *divert traffic* and assist in analyzing attacks and determine the source.
  - Used to redirect attacks away from the customer – working the attack on a router built to withstand the attack.
  - Used to monitor *attack noise, scans, data from mis-configuration* and other activity (via the advertisement of default or unused IP space)
  - Traffic is typically diverted via BGP route advertisements and policies.

Before Sinkhole Activated

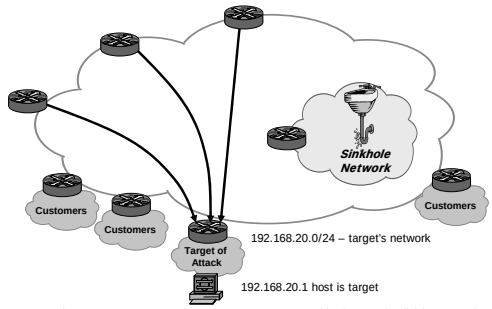

Sinkhole Routers/Networks

• Attack is pulled away from customer/aggregation router.

• Can now apply classification ACLs, Packet Capture, Etc…

• Objective is to minimize the risk to the network while investigating the attack incident.

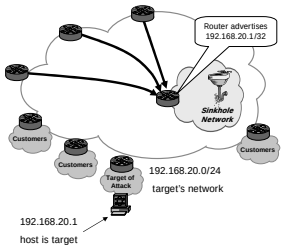

Sinkhole Routers/Networks

- Advertising "default" from the Sinkhole will pull down all sorts of *innocuous* traffic:
  - Customer Traffic when circuits flap
  - Network Scans to unallocated address space
  - Code Red/NIMDA/Worms
  - Backscatter
- Can place tracking tools in the Sinkhole network to monitor the noise.

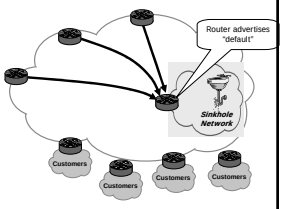

Scaling Sinkhole Networks

- Multiple Sinkholes can be deployed within a network
- Combination of IGP with BGP Trigger
- Regional deployment
  - Major PoPs
- Functional deployment
  - Peering points
  - Data Centers
- Note: Reporting more complicated, need aggregation and correlation mechanism

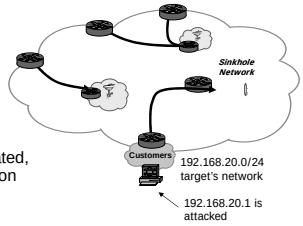

The Basic Sinkhole

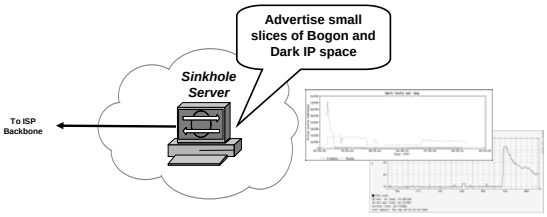

- Sinks Holes do not have to be complicated.
- Some large providers started their Sinkhole with a spare workstation with free unix, Zebra, and TCPdump.
- Some GNU or MRTG graphing and you have a decent sinkhole.

Expanding The SinkHole

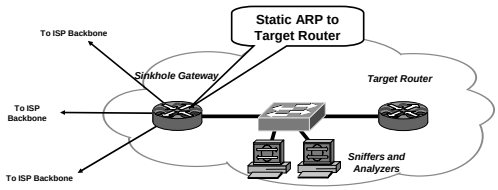

- Expand the Sinkhole with a dedicated router into a variety of tools.
- Pull the DOS/DDOS attack to the sinkhole and forwards the attack to the target router.
- Static ARP to the target router keeps the Sinkhole Operational – Target Router can crash from the attack and the static ARP will keep the gateway forwarding traffic to the Ethernet switch.

Typical Aggregate Sources

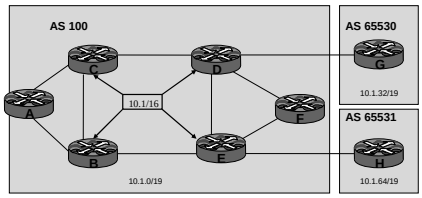

- 10.1/16 allocated to AS 100
- 10.1.0/19 used for infrastructure
- 10.1.32/19 AS 65530
- 10.1.64/19 AS 65531
- 10.1/16 (10.1.96-10.1.255.255) implicitly nailed to null interface on core routers (C,B,D&E)

Routers Collect Garbage

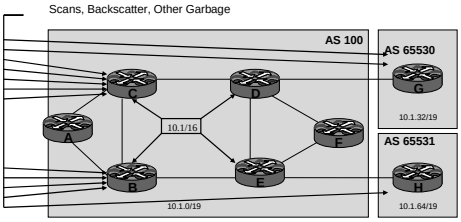

- Routers collect all the garbage (backscatter, scans, etc..) destined for 10.1/19, 10.1.96/19 & 10.1.128/17 addresses
- Routers are required to process data, send ICMP unreachables, etc..

Sinkhole is Useful Here

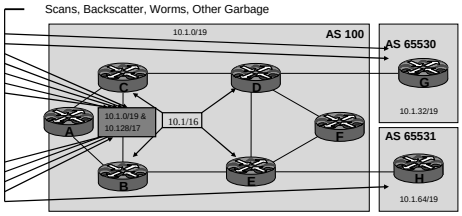

- Divert garbage to sinkhole, if not for further analysis, at least to off-load data processing from routers
- Traffic forwarded to sinkhole for analysis, removes processing overhead from routers
- Provide collection point for further analysis

Sinkholes - Advertising Dark IP Space

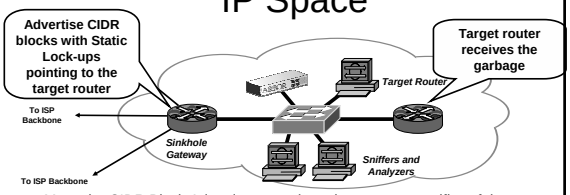

- Move the CIDR Block Advertisements (or at least more-specifics of those advertisements) to Sinkholes
- Does not impact BGP routing – route origination can happen anywhere in the iBGP mesh (careful about MEDs and aggregates)
- Control where you drop the packet
- Turns networks inherent behaviors into a security tool!

Sinkhole Monitoring

- Scans on Dark IP (allocated & announced but unassigned address space).
  - Who is scoping out the network – pre-attack planning.
- Scans on Bogons (unallocated).
  - Worms, infected machines, and Bot creation
- Backscatter from Attacks
  - Who is getting attacked
- Backscatter from Garbage traffic (RFC-1918 leaks)
  - Which customers have misconfiguration or "leaking" networks.

Sinkhole Considerations

- Do not allow advertisements to leak:
  - BGP no-export, no-advertise, additive communities
  - Explicit egress prefix policies (community, prefix, etc.)
- Do not allow traffic to escape the sinkhole:
  - Backscatter from a sinkhole defeats the function of a sinkhole (egress ACL on the sinkhole router)
- Advanced sinkhole designs
  - True honeypot potential →protect resources in the sinkhole
  - Don't become part of the attack
  - Filter/rate limit outgoing connections

Monitoring Scan Rates

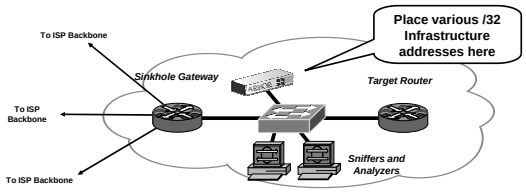

- Select /32 (or larger) address from different block of your address space. Advertise them out the Sinkhole
- Assign them to a workstation built to monitor and log scans. (Arbor Network's *Dark IP* Peakflow module is one turn key commercial tool that can monitor scan rates via data collected from the network.)

Monitoring Backscatter

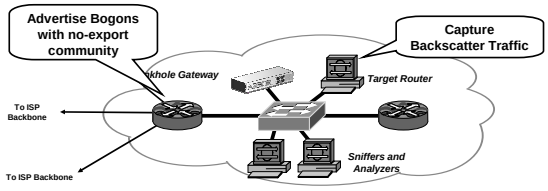

- Advertise bogon blocks with NO_EXPORT community and an explicit safety community (plus prefix-based egress filtering on the edge)
- Static/set the BGP NEXT_HOP for the bogon to a backscatter collector workstation (as simple as TCPdump).
- Pulls in backscatter for that range – allows monitoring.

Monitoring Spoof Ranges

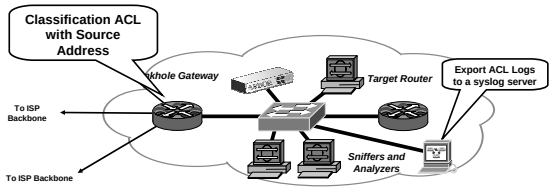

- Attackers use ranges of valid (allocated blocks) and invalid (bogon, martian, and RFC1918 blocks) spoofed IP addresses.
- Extremely helpful to know the spoof ranges.
- Set up a classification filter on source addresses.

Monitoring Spoof Ranges

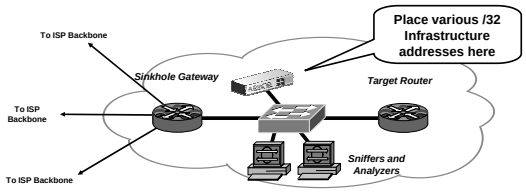

- Select /32 address from different block of your address space. Advertise them out the Sinkhole
- Assign them to a workstation built to monitor and log scans.
- Home grown and commercial tools available to monitor scan rates

Sinkhole Impact

- In the past, this issue of pulling down garbage packets has not been a big deal.
- GigBots and Turbo Worms change everything
- Even ASIC-based forwarding platforms get impacted from the *overhead.*

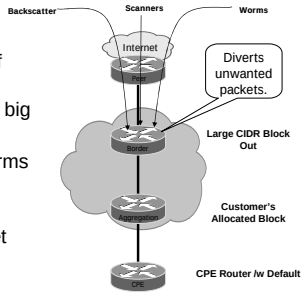

Protecting Backbone Point-to-Point Addresses

Do you really need to reach the Backbone router's Point to Point Address from any router other than a directly connected neighbor?

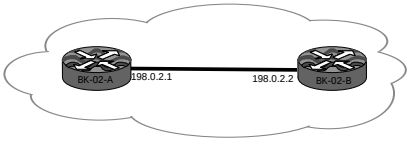

Protecting Backbone Point-to-Point Addresses

- What could break?
  - Network protocols are either loopback (BGP, NTP, etc.) or adjacent (OSPF, IS-IS, EIGRP).
  - NOC can Ping the Loopback (alhough some tools such as HP OV may have issues).
  - Traceroutes reply with the correct address in the reply. Reachability of the source is not required.

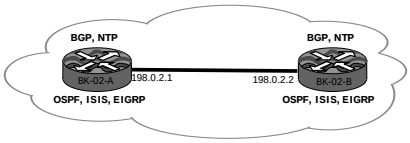

Protecting Backbone Point-to-Point Addresses

- What have people done in the past:
  - ACLs – Long term ACL management problems.
  - RFC 1918 – Works against the theme of the RFC – Traceroute still replies with RFC 1918 source address.
  - Does not protect against a reflection attack.

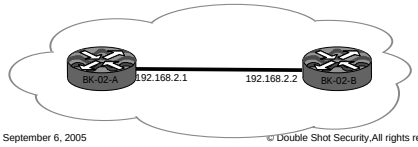

Protecting Backbone Point-to-Point Addresses

- Move the Point to Point Address blocks to IGP based Sinkholes.
  - All packets to these addresses will be pulled into the Sinkhole.
  - People who could find targets with traceroute cannot now hit the router with an attack based on that intelligence.
  - Protects against internal and reflection based attacks.

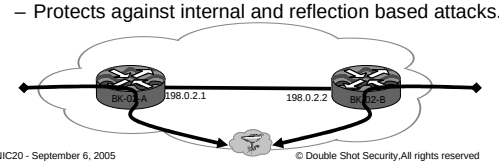

Just Another Hurdle

- Will not work with the routers on the border.
  - By default, C (Connected) prefixes override all BGP injected prefixes from the Sinkhole (you want this to happen).
  - Basic security principle – increment layers of security – there is never a perfect solution – just additional hurdles – the more hurdles the better.

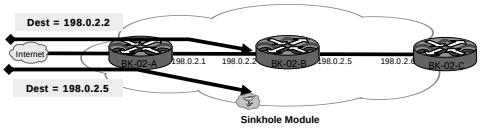

If Using ISP Edge Filters ?

- Anti-Spoof and Anti-Infrastructure ACLs are encouraged on the edge. But ….
- Need to be everywhere to achieved desired effect – including the customer edge (this is beyond the BCP 38 requirements).

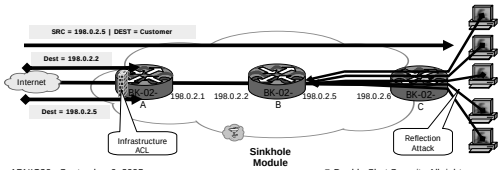

If Using ISP Edge Filters ?

- Anti-Spoof and Anti-Infrastructure ACLs can be combined with Sink Holing the Infrastructure Blocks.
- Remember – it's a layered approach

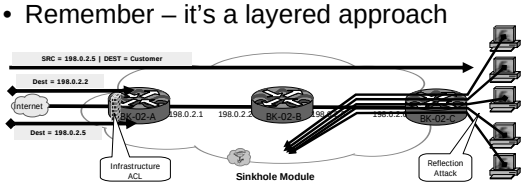

Sinkhole Routing Components

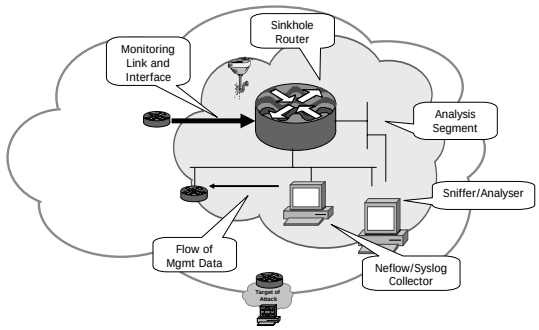

Sample TEST-NET Allocation

Address Block	Purpose
192.0.2.1/32	All iBGP routers for "Drop to NULL0"
192.0.2.2/32	All Peering Edge routers drop
192.0.2.3/32	All Customer Edge routers drop
192.0.2.4/30	Monitor Link addresses NOTE: provision these addresses in all Sinkholes
192.0.2.254	ANYCAST Sinkhole Address
192.0.2.8 -> balance	Sinkhole Diversion Addresses

Sinkhole Routing

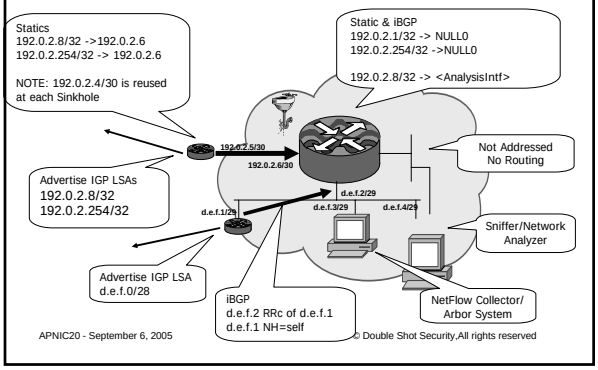

Sinkhole Routing Considerations

- No Default static route in Sinkhole.
  - –Sinkhole must not loop traffic back out Management Interface.
  - –Telnet access via router servicing the Sinkhole's Management Segment.

What Is A Blackhole?

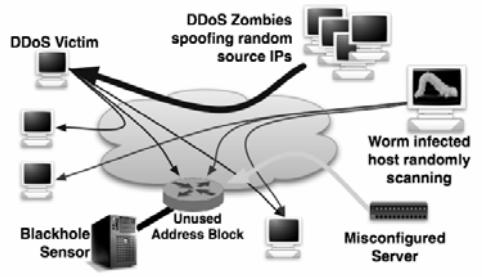

- A Blackhole sensor monitors an unused globally advertised address block that contains ***no active hosts***
- Traffic is the result of ***DDoS backscatter***, ***worm propagation***, ***scanning…***

BGP Blackhole Routing

- Commonly referred to as *BGP Real-Time (or Remotely-Triggered) Blackhole Routing (RTBH),* or *Blackhole Filtering*; results in packets being forwarded to a:
  - Null Interface
  - Discard Interface
- Several Techniques:
  - Destination-based BGP Blackhole Routing
  - Source-based BGP Blackhole Routing (coupling uRPF)
  - Customer-triggered
- Exploits router's forwarding logic - typically results in desired packets being dropped with minimal or no performance impact
- Enables BGP Backscatter Traceback Technique

Blackhole Filtering

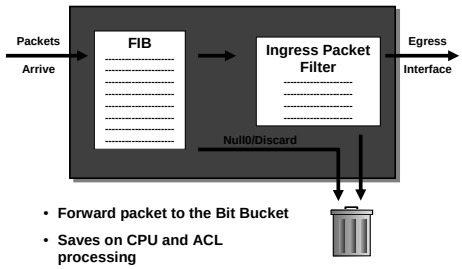

- **Forward packet to the Bit Bucket**
- **Saves on CPU and ACL processing**

RTBH Basics

- Use BGP routing protocol to trigger network wide response to an attack flow.
- Simple static route and BGP allows ISP to trigger network wide black holes as fast as iBGP can update the network.
- Unicast RPF allows for the black hole to include any packet whose source or destination address matches the prefix.
  - Effective against spoofed and valid source addresses.

RTBH - Steps

- 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 speaker and installed on routers with 192.0.2.1 route
  - –All traffic to victim now sent to Null0

Preparing Routers with Filtering Trigger

- 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 Null 0 on every edge router on the network

```
ip route 192.0.2.1 255.255.255.255 Null0
```


Preparing Routers with Filtering Trigger

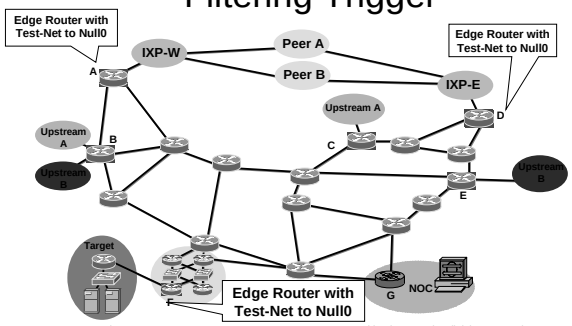

Preparing Trigger Router

- 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 Configuration

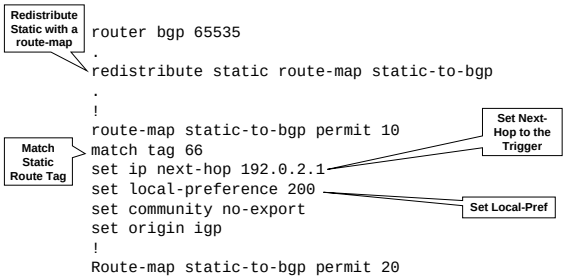

Activate Black Hole

- 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 received BGP update, and "attach" it to the existing static route; the next-hop is now Null0

Activating the Black Hole

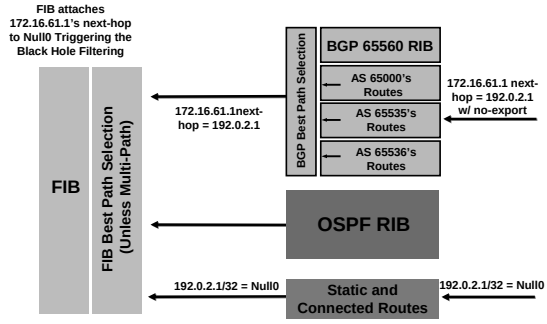

Activating the Black Hole

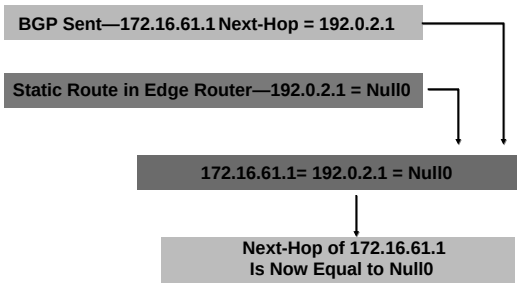

DDoS Before RTBH

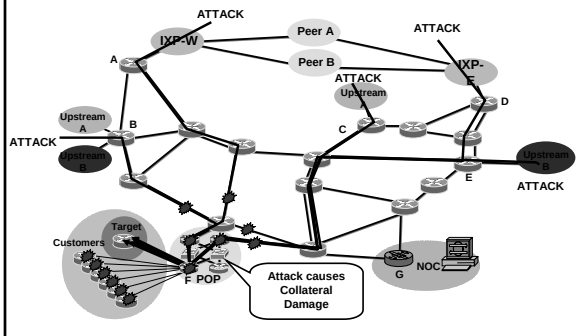

DDoS After RTBH
Packet Drops Pushed to Edge

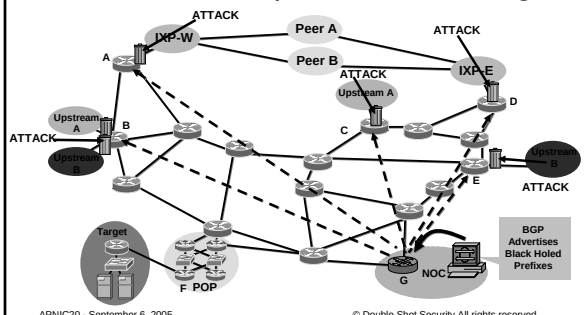

Triggered Source Drops

- Dropping on destination is very important
  - Dropping on source if 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

Source Based RTBHF

- What do we have?
  - *Black Hole Filtering* **– If the *destination* address equals Null 0 we drop the packet.**
  - *Remote Triggered* – **Trigger a prefix to equal Null 0 on routers across the Network at iBGP speeds.**
  - uRPF Loose Check – **If the *source* address equals Null 0, we drop the packet.**
- Put them together and we have a tool to trigger drop for any packet coming into the network whose source or destination equals Null 0!

uRPF Loose Mode

- Originally created to scale BCP 38 ingress filtering on the ISP
  - Customer Edge of an ISP's network.
- Loose Check Mode added to provide ISPs with means to trigger a network wide, source based black hole filter activated at BGP update speeds.
- uRPF Loose Check will passively drop any packet whose source address is not in the router's FIB.
- Effective way to drop Bogon addresses.

uRPF - Loose Mode

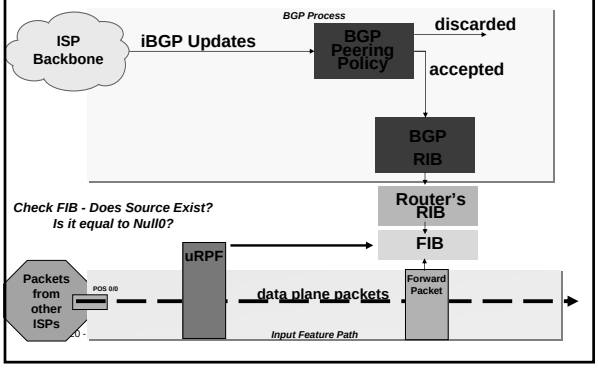

uRPF Example

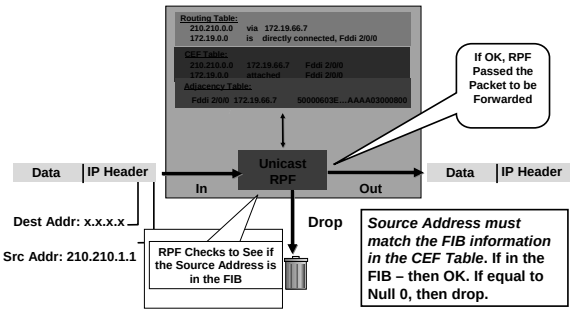

Community Based Trigger

- BGP community-based triggering allow for more fined 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 routers (on the edge) matches community and sets the next-hop to the static route to Null0

Community Based Trigger Examples

- 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 filter 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
  - Resource utilization: finite resources
  - Effects all traffic, on all triggered interfaces, regardless of actual intent

Filtering vs Remotely Triggered Drops

- Filtering key strengths:
  - –Detailed packet filtering (ports, protocols, ranges, fragments, etc.)
  - –Relatively static filtering environment
  - –Clear filtering policy
- Filters 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
- Combining filters with uRPF remote-triggered drops allows for filters to handle the strict static policies while uRPF remote-triggered black hole handles the dynamic source-based drops

BGP Flow Specification

- Defined in:
  - http://www.ietf.org/internet-drafts/draft-marques-idr-flow-spec-02.txt
- Specifies procedures for the distribution of flow specification rules via BGP
- Defines an application for the purpose of packet filtering in order to mitigate (distributed) denial of service attacks
- Defines procedure to encode flow specification rules as BGP NLRI which can be used in any way the implementer desires

What Is A Flow Specification?

- 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

Use of Flow Specification

- Primarily/Initially: DDOS/Worm Mitigation
- Continue evolution from:
  - Destination-based blackhole routing
  - uRPF/source-based BGP blackhole routing
- To:
  - Much more precise/granular mechanism that contains all the benefits of it's predecessors
- At least one implementation complete, another (more?) on the way

Traditional Traceback

- Hop-by-hop
- Error-prone
- May impact service availability
- Tedious
- Very time consuming
- Fully characterizing and accounting for full impact of attack is still unlikely.

Backscatter Traceback

- Pioneered by Chris Morrow and Brian Gemberling @ UUNET as a means of finding the entry point of a spoofed DOS/DDOS.
  `http://www.secsup.org/Tracking/`
- Combines the Sink Hole router, Backscatter Effects of Spoofed DOS/DDOS attacks, and remote triggered Black Hole Filtering to create a traceback system that provides a result within ~10 minutes.

Backscatter Traceback

- What is backscatter?

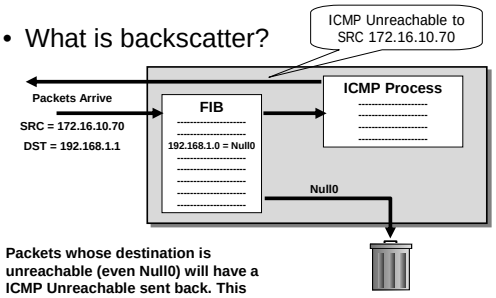

Packets whose destination is unreachable (even Null0) will have a ICMP Unreachable sent back. This "unreachable noise" is backscatter.

Backscatter Traceback Preparation

1. Sink Hole Router/Network connected to the network and ready to classify the traffic. Like before, BGP Route Reflector Client, device to analyze logs, etc.
   - Can use one router to do both the route advertisement and logging OR break them into two separation routers – one for route advertisement and the other to accept/log traffic
   - Can be used for other Sink Hole functions while not using the traceback technique.
   - Sink Hole Router can be a iBGP Route Reflector into the network.

Backscatter Traceback Preparation

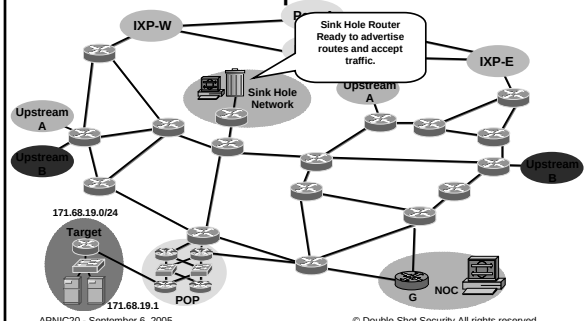

Backscatter Traceback Activation

```
!
router bgp 31337
!
! set the static redistribution to include a route-map so we can filter
! the routes somewhat... or at least manipulate them
! redistribute static route-map static-to-bgp
!
! add a stanza to the route-map to set our special next hop
!
 route-map static-to-bgp permit 5
 match tag 666
 set ip next-hop 172.20.20.1
 set local-preference 50
 set origin igp
```


Backscatter Traceback Activation

```
# Setup the bgp protocol to export our special policy, like redistributing
# NOTE: "XXX" # is the IBGP bgp group... we don't want to send this to customers
#
set protocols bgp group XXX export BlackHoleRoutes
#
# Now, setup the policy option for BlackHoleRoutes, like a route-map if static route
# with right tag, set local-pref low, internal, no-export  can't leak these or Tony Bates
# will have a fit, and set the nexthop to the magical next-hop.
#
set policy-statement BlackHoleRoutes term match-tag666 from protocol static tag 666
set policy-statement BlackHoleRoutes term match-tag666 then local-preference 50
set policy-statement BlackHoleRoutes term match-tag666 then origin igp
set policy-statement BlackHoleRoutes term match-tag666 then community add no-export
set policy-statement BlackHoleRoutes term match-tag666 then nexthop 172.20.20.1
set policy-statement BlackHoleRoutes term match-tag666 then accept
```


Backscatter Traceback Preparation

2. All edge devices (routers, NAS, IXP Routers, etc) with a static route to Null0. The Test-Net is a safe address to use (192.0.2.0/24) since no one is using it.

   Cisco: ip route 172.20.20.1 255.255.255.255 Null0
   Juniper: set routing-options static route 172.20.20.1/32 reject install

   – Routers also need to have ICMP Unreachables working. If you have ICMP Unreachables turned off (i.e. *no ip unreachables* on a Cisco), then make sure they are on.
   – If ICMP Unreachable Overloads are a concern, use a ICMP Unreachable Rate Limit (i.e. *ip icmp rate-limit unreachable* command on a Cisco).

Backscatter Traceback Preparation

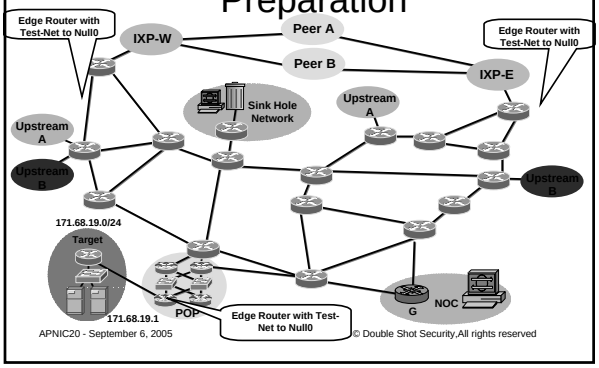

Backscatter Traceback Preparation

3. Sink Hole Router advertising a large block of un-allocated address space with the BGP no-export community and BGP Egress route filters to keep the block inside. (Ex: 96.0.0.0/3)
   - Check with IANA for unallocated blocks: **`www.iana.org/assignments/ipv4-address-space`**
   - BGP Egress filter should keep this advertisement inside your network.
   - Use BGP ***no-export*** community to insure it stays inside your network.

Backscatter Traceback Preparation

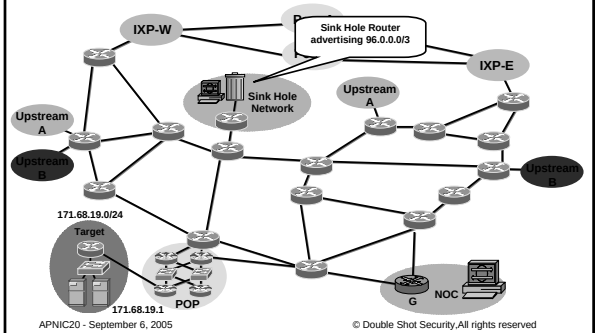

Backscatter Traceback Activation

- Activation happens when an attack has been identified.
- Basic Classification should be done to see if the backscatter traceback will work:
  - May need to adjust the advertised block.
  - Statistically, most attacks have been spoofed using the entire Internet block.

Backscatter Traceback Activation

Sink Hole Router Advertises the /32 under attack into iBGP with static route with the “666” tag:

```
ip route victimip 255.255.255.255 Null0 tag 666
```

or

```
set routing-options static route victimip/32 discard tag 666
```

The static triggers the routers to advertise the customer’s prefix

Backscatter Traceback Activation

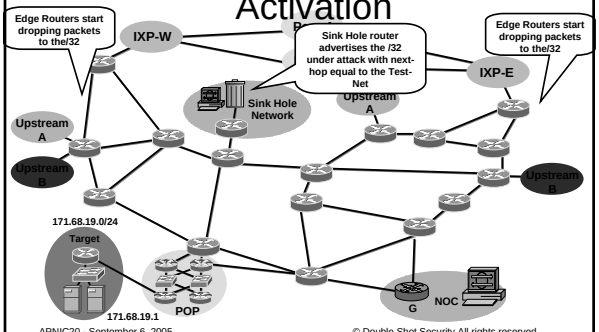

Backscatter Traceback Activation

Black Hole Filtering is triggered by BGP through out the network. Packets to the target get dropped. ICMP Unreachable Backscatter starts heading for 96.0.0.0/3.

- Access list is used on the router to find which routers are dropping packets.
- access-list 101 permit icmp any any unreachables log
  access-list 101 permit ip any any

Backscatter Traceback Activation

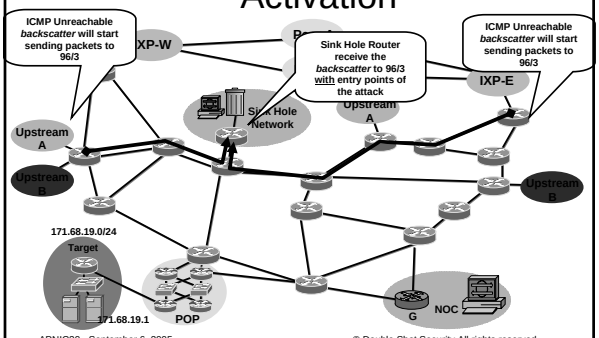

Backscatter Traceback Log

SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.47.251.104 (3/1), 1 packet
SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.70.92.28 (3/1), 1 packet
SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.222.127.7 (3/1), 1 packet
SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.96.223.54 (3/1), 1 packet
SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.14.21.8 (3/1), 1 packet
SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.105.33.126 (3/1), 1 packet
SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.77.198.85 (3/1), 1 packet
SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.50.106.45 (3/1), 1 packet

Packet Scrubbing

- Build more granularity and selectivity into the discard process
  - Not all traffic from a given source is “bad”
  - Validate sources
  - Per source detection and enforcement
- Can use the same BGP mechanism to redirect traffic to scrubbing devices
- Activate redirection:
  - Redistribute host route for victim into BGP with next-hop set to scrubbing devices
  - Route is propagated using BGP to all BGP speaker and traffic redirected
- When attack is over, BGP route can be removed to return to normal operation

Packet Scrubbing

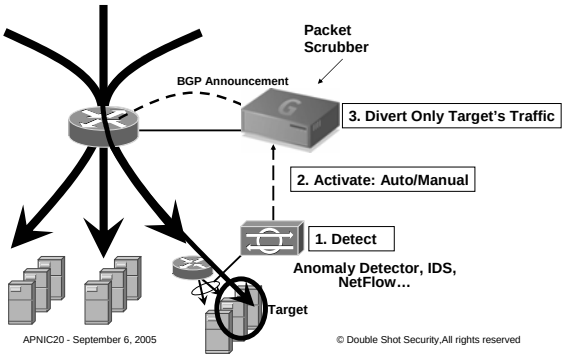

Packet Scrubbing

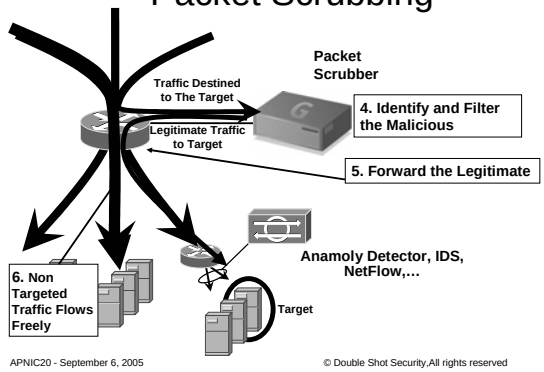

DDoS Mitigation Summary

- Many varying reaction mechanisms
  - Sinkhole
  - Remotely Triggered Blackhole Routing
  - Backscatter Traceback
  - Packet Scrubbing
- No one tool or technique is applicable in all circumstances
  - Use combination of tools
  - Automate where possible
- Choose your techniques wisely

QUESTIONS ?