



WHITE PAPER

REMOTELY TRIGGERED BLACK HOLE FILTERING— DESTINATION BASED AND SOURCE BASED

Remotely triggered black hole (RTBH) filtering is a technique that provides the ability to drop undesirable traffic before it enters a protected network. This document describes RTBH filtering and its merits, operational gains, applications, and deployment considerations and provides sample router configurations. This document describes the mitigation of distributed-denial-of-service (DDoS) attacks within a single Interior Gateway Protocol (IGP) domain. This document does not describe how to use RTBH filtering for mitigating the attack across multiple providers.

This document is intended for network design architects, support engineers, and marketing professionals who are responsible for planning, designing, implementing, and operating networks.

OVERVIEW

This section describes RTBH filtering and how it is used for both destination-based and source-based filtering. This section includes the following topics:

- Benefits of Remotely Triggered Black Hole Filtering
- Remotely Triggered Black Hole Filtering Within the Service Provider Security Framework
- Destination-Based
- Source-Based

Benefits of Remotely Triggered Black Hole Filtering

Black holes, from a network security perspective, are placed in the network where traffic is forwarded and dropped. Once an attack has been detected, black holing can be used to drop all attack traffic at the edge of an Internet service provider (ISP) network, based on either destination or source IP addresses. RTBH filtering is a technique that uses routing protocol updates to manipulate route tables at the network edge or anywhere else in the network to specifically drop undesirable traffic before it enters the service provider network.

RTBH filtering provides a method for quickly dropping undesirable traffic at the edge of the network, based on either source addresses or destination addresses by forwarding it to a null0 interface. Null0 is a pseudointerface that is always up and can never forward or receive traffic. Forwarding packets to null0 is a common way to filter packets to a specific destination.

RTBH filtering is not a specific Cisco IOS® Software feature, but rather a technique that incorporates a set of well-coordinated configurations across multiple routers. RTBH filtering is one of the many techniques in the security toolkit that can be used together to enhance network security in the following ways:

- Effectively mitigate DDoS and worm attacks
- Quarantine all traffic destined for the target under attack
- Enforce blacklist filtering

A typical deployment scenario for RTBH filtering would require running internal Border Gateway Protocol (iBGP) at the access and aggregation points and configuring a separate device in the network operations center (NOC) to act as a trigger. The triggering device sends iBGP updates to the edge, that cause undesirable traffic to be forwarded to a null0 interface and dropped.


Destination-Based Remotely Triggered Black Hole Filtering

With a denial-of-service (DoS) attack, in addition to service degradation of the target, there is possible collateral damage such as bandwidth consumption, processor utilization, and potential service loss elsewhere in the network. One method to mitigate the damaging effects of such an attack is to black hole (drop) traffic destined to the IP address or addresses being attacked and to filter the infected host traffic at the edge of the network closest to the source of the attack.

The challenge is to find a way to quickly drop the offending traffic at the network edge, document and track the black holed destination addresses, and promptly return these addresses to service once the threat disappears. Destination-based IP black hole filtering with remote triggering allows a network-wide destination-based black hole to be propagated by adding a simple static route to the triggering device (trigger).

The trigger sends a routing update for the static route using iBGP to the other edge routers configured for black hole filtering. This routing update sets the next hop IP address to another preconfigured static route pointing to the null interface. This process is illustrated in Figure 1.

Figure 1. Destination-Based Black Hole Filtering with Remote Triggering



The three steps in destination-based black hole filtering are summarized below.

Step 1. The setup (preparation)

A trigger is a special device that is installed at the NOC exclusively for the purpose of triggering a black hole. The trigger must have an iBGP peering relationship with all the edge routers, or, if using route reflectors, it must have an iBGP relationship with the route reflectors in every cluster. The trigger is also configured to redistribute static routes to its iBGP peers. It sends the static route by means of an iBGP routing update.

The Provider Edges (PEs) must have a static route for an unused IP address space. For example, 192.0.2.1/32 is set to Null0. The IP address 192.0.2.1 is reserved for use in test networks and is not used as a deployed IP address.

Step 2. The trigger

An administrator adds a static route to the trigger, which redistributes the route by sending a BGP update to all its iBGP peers, setting the next hop to the target destination address under attack as 192.0.2.1 in the current example.

The PEs receive their iBGP update and set their next hop to the target to the unused IP address space 192.0.2.1. The route to this address is set to null0 in the PE, using a static routing entry in the router configuration. The next hop entry in the forwarding information base (FIB) for the destination IP (target) is now updated to null0.

All traffic to the target will now be forwarded to Null0 at the edge and dropped.

Step 3. The withdrawal

Once the trigger is in place, all traffic to the target destination is dropped at the PEs. When the threat no longer exists, the administrator must manually remove the static route from the trigger, which sends a BGP route withdrawal to its iBGP peers. This prompts the edge routers to remove the existing route for the target that is pointed to 192.0.2.1 and to install a new route based on the IGP routing information base (RIB).

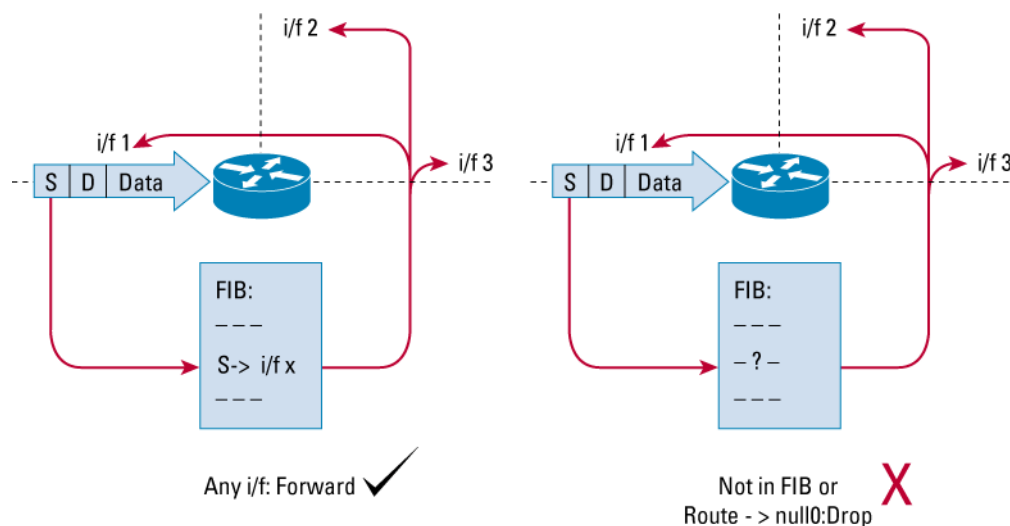
Source-Based Remotely Triggered Black Hole Filtering

Source-based black holes provide the ability to drop traffic at the network edge based on a specific source address or range of source addresses. With destination-based black holing, all traffic to a specific destination is dropped once the black hole has been activated, regardless of where it is coming from. Obviously, this could include legitimate traffic destined for the target.

If the source address (or range of addresses) of the attack can be identified (spoofed or not), it would be better to drop all traffic at the edge based on the source address, regardless of the destination address. This would permit legitimate traffic from other sources to reach the target. Implementation of source-based black hole filtering depends on Unicast Reverse Path Forwarding (URPF), most often loose mode URPF.

Loose URPF checks the packet and forwards it if there is a route entry for the source IP of the incoming packet in the router FIB. If the router does not have an FIB entry for the source IP address, or if the entry points to Null0, the Reverse Path Forwarding (RPF) check fails, and the packet is dropped, as shown in Figure 2.

Figure 2. Loose URPF with Source-Based Black Hole Filtering



Step 2. The trigger

An administrator adds a static route to the trigger. The trigger redistributes the route by sending a BGP update to all its iBGP peers, which sets the next hop to the source IP of the attacker (as 192.0.2.1 in the current example).

Each PE receives an iBGP update and sets its next hop to the source IP to the unused IP address space 192.0.2.1. The next hop to this address is set to Null0 using a static routing entry in the router configuration. The next hop entry in the FIB for the source IP address is now updated to Null0.

All traffic from the source IP will fail the loose URPF check at the PEs and as a consequence will be dropped.

Step 3. The withdrawal

Once the trigger is in place, all traffic from the source IP addresses will be dropped at the PEs. When the threat no longer exists, the administrator must manually remove the static route from the triggering device, which sends a BGP route withdrawal to its iBGP peers. This prompts the edge routers to remove the existing route for the source IP that points to 192.0.2.1 and to install a new route in the FIB based on the IGP RIB. If this new route is successful, loose URPF checks will pass, and traffic from the blocked source will be forwarded normally.

CONFIGURATION GUIDELINES

This section provides general configuration guidelines for the components of a remotely triggered black hole filtering solution. It includes the following topics:

- Configuring the Trigger
- Configuring the Provider Edge Routers

Configuring the Trigger

A trigger sends iBGP updates to its peers, which set the next hop for a network or host to a predetermined unused IP address. The trigger can be any device that runs BGP and does not necessarily have to be a router. If an Arbor Networks Peakflow SP device is used for anomaly detection, it can also be used as a trigger. A UNIX workstation running BGP can also be used as a triggering device.

A typical deployment would have multiple regionalized triggering devices for redundancy, with each device protected by dual power supplies. A key deployment consideration is to make sure that the trigger advertisement is contained within the ISP domain and does not impact other routing policies within the ISP itself. The following are three recommended ways to contain the iBGP routing updates used to trigger black holes at the network edge:

- Using BGP communities, which are attributes that are applied to prefixes and used to group and filter routes. Specifically, the no-export community is used within an autonomous system (AS) to make sure that the edge routers (PEs) do not export or readvertise the prefix used by the trigger advertisement to their external BGP (eBGP) peers.

Note: When using confederations, use the local-as community instead of the no-export community. Functionally, this has the same impact as restricting the advertisements to the local autonomous systems.

- As an added safety precaution, an extra community should be set for the prefix used in the trigger advertisement. A community filter must then be added to all eBGP routers to ensure that a prefix with this community is not advertised to its peers. This helps contain the prefix within the AS even if the no-export community is accidentally deleted for the prefix.
- An egress prefix filter should be used with eBGP peers that will filter all prefixes less than a given length, such as /24. The trigger advertisement is typically a prefix from /25 to /32, so it will not be advertised to external peers.

Once the trigger advertisement is sent to the iBGP peers, you need to make sure that the peers prefer this route to one that they already have for the destination network. One way is to set a local-preference value greater than the default value of 100 and set the origin to igp. The preferred method

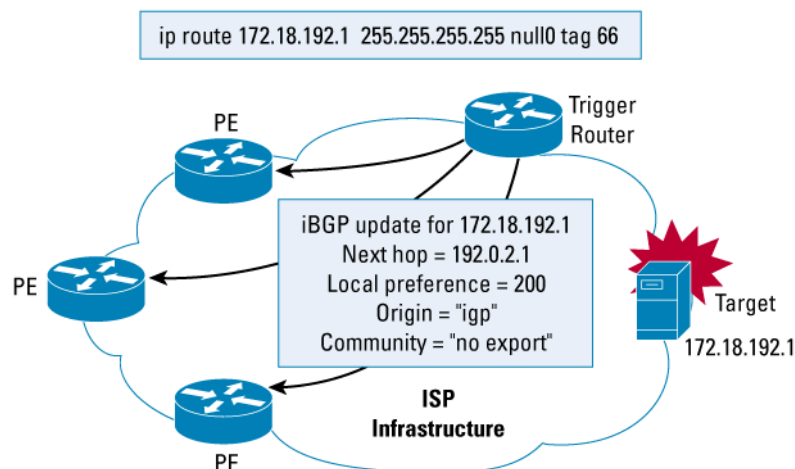
is to set a path with a higher value for local-preference and a lower value of origin (IGP is lower than Exterior Gateway Protocol [EGP], which is lower than incomplete).

The following is a sample configuration of a router that is used as a triggering device:

```
interface loopback0
ip address 192.168.1.1 255.255.255.255
interface Null0
no ip unreachable
router bgp 64555
no synchronization
no bgp client-to-client reflection
bgp log-neighbor-changes
redistribute static route-map black-hole-trigger
neighbor black-hole peer-group
neighbor black-hole remote-as 64555
neighbor black-hole update-source loopback0
neighbor black-hole route-reflector-client
neighbor x.x.x.x peer-group black-hole
route-map black-hole-trigger permit 10
match tag 66
set ip next-hop 192.0.2.1
set local-preference 200
set origin igp
set community no-export
route-map black-hole-trigger deny 20
```

In the configuration shown in Figure 4, the trigger consists of adding a static route with a tag value of 66.

Figure 4. Triggering Destination-Based Remotely Triggered Black Hole Filtering



The route map black-hole-trigger is applied prior to redistributing static routes into BGP. The following occurs in the route map before redistributing the route to iBGP peers:

- Match on a tag value of 66.
- Set the next hop to 192.0.2.1.
- The local-preference is set to 200.
- The origin is set to igp.
- The community is set to no-export.

A sample trigger is shown below:

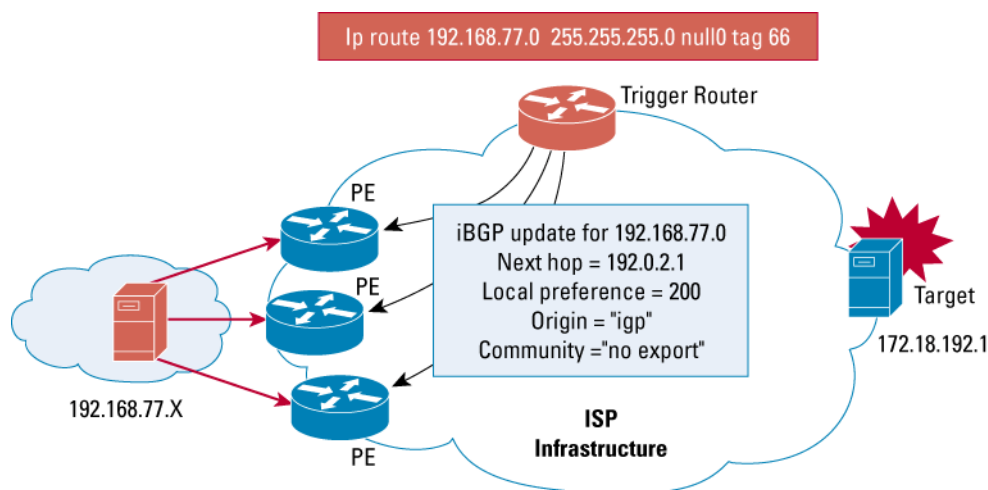
```
ip route 172.18.192.1 255.255.255.255 Null0 tag 66
```

This line causes the trigger to send an iBGP route update to all its iBGP peers for the host 172.18.192.1, setting its next hop to 192.0.2.1. This type of trigger can be used to drop all traffic at the edges of the network closest to the target under attack, which in this example is 172.18.192.1.

On the other hand if the attacker is using source addresses in a predictable range, such as within the network 192.168.77.0, the following trigger could be used to black hole traffic from these sources, as shown in Figure 5.

```
ip route 192.168.77.0 255.255.255.0 null0 tag 66
```

Figure 5. Triggering Source-Based Remotely Triggered Black Hole Filtering



Configuring the Provider Edge Routers

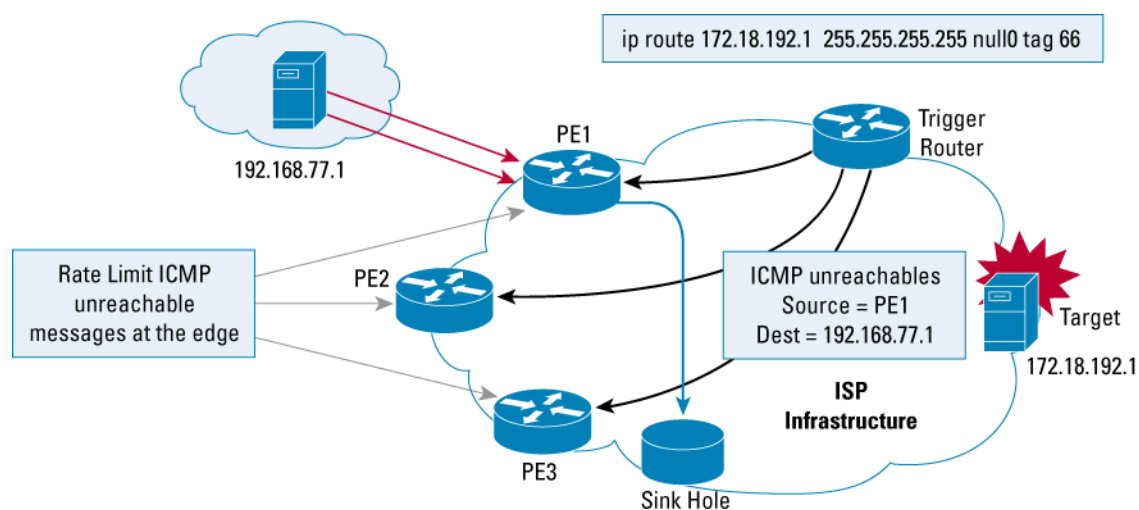
The edge routers essentially drop suspicious traffic by forwarding it to a Null0 interface. Null0 is an invalid interface in Cisco® Express Forwarding tables, so all traffic forwarded to a null interface will be dropped by Cisco Express Forwarding and does not require process switching. Hence, using Null0 as a way to filter traffic adds minimal overhead to the edge routers. A static route is configured at the edge, and the next hop is set to Null0.

The edge router receives an iBGP route update, which sets the next hop of the destination under attack to this static route. All future traffic to the destination is dropped at the edge because the next hop to the destination under attack is equated to Null0. The static route that is used should be for an address in a network that will never be used. In our examples, we use an address in network 192.0.2.0.

Typically, when an IP datagram is dropped, an Internet Control Message Protocol (ICMP) unreachable message is sent back to the source giving the reason why the packet could not be delivered to its final destination. In most cases, when traffic is deliberately dropped by being forwarded to a null interface, you do not want to overburden the router by making it send this unreachable message to the source address. Also, these messages would create additional traffic on the network and inform the source that the packets are being dropped. So, it is recommended that when a Null0 interface is created at the edges, the ICMP unreachable message is disabled for this interface.

Sometimes ICMP unreachable messages are useful for tracking the entry point of the attack. The source address of the attack traffic is hijacked by a sinkhole device within the ISP domain, which collects ICMP unreachable messages sent from the router that drops the traffic. The entry point of the attack can be determined by the source IP address in the unreachable messages, which identifies the edge router that sent the message (see Figure 6).

Figure 6. Identifying the Entry Point of an Attack Using Internet Control Message Protocol Unreachable Messages



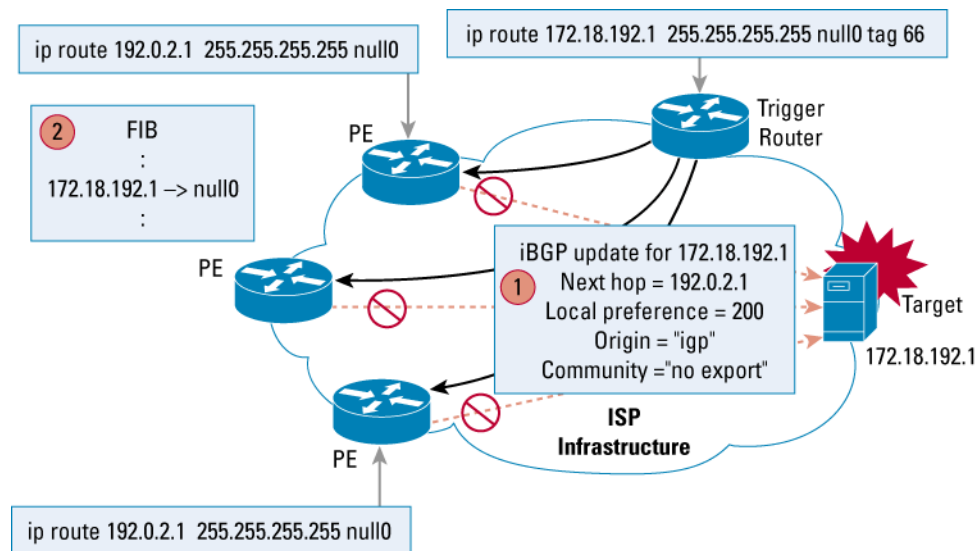
Analyzing the sinkhole traffic in this example, the entry point of the attack can be determined as PE-1 by looking at the source of the ICMP unreachable messages.

If ICMP unreachable messages are not disabled, it is strongly recommended that they be rate limited so they will not flood the network. One way to rate-limit ICMP unreachable messages is by using the **ip icmp rate-limit unreachable n** command. In this command, n specifies the number of milliseconds between two consecutive ICMP unreachable messages. A sample configuration for an edge router is shown below.

```
interface loopback0
ip address x.x.x.x 255.255.255.255
interface null0
no ip unreachables
router bgp 64555
  no synchronization
  bgp log-neighbor-changes
  neighbor black-hole peer-group
  neighbor black-hole remote-as 65535
  neighbor black-hole update-source loopback0
  neighbor a.a.a.a peer-group black-hole
  no auto-summary
ip route 192.0.2.1 255.255.255.255 null 0
```


When an edge router receives the triggered advertisement, it becomes the preferred route to the destination because it has a higher local preference and a lower origin of igp. The next hop to the destination network is set to 192.0.2.1. However, because the next hop destination equates to Null0, the next hop route to the destination is set to Null0, as shown in Figure 7.

Figure 7. Provider Edge Router Configuration



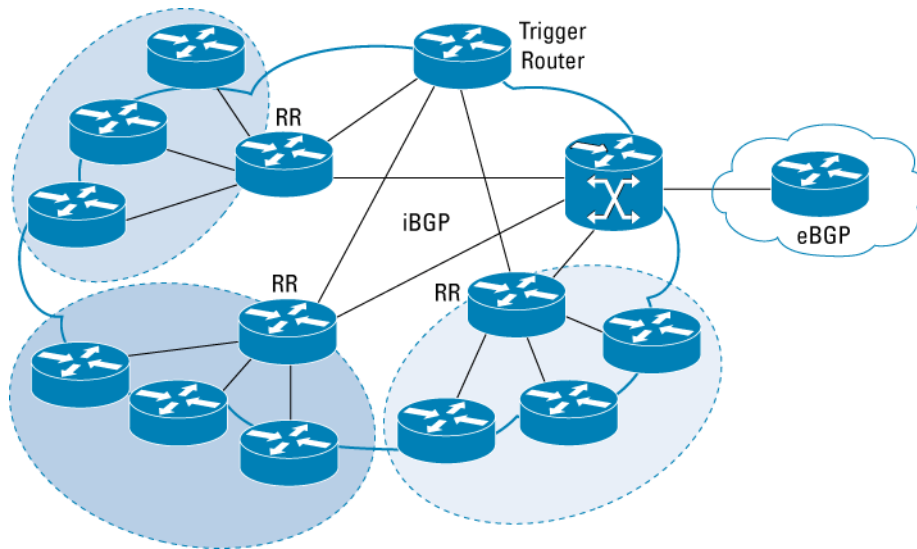
DEPLOYMENT CONSIDERATIONS

RTBH filtering is typically deployed by configuring iBGP on the provider edge and a new triggering device at the NOC. The trigger maintains an iBGP peering relationship with all the edge routers or with all the route reflectors in the ISP environment (Figure 8).

Manual entry of the static route is done on the trigger as well as setting all the necessary characteristics such as the next hop, local preference, and BGP communities.

There are two methods for deploying RTBH filtering. The first, and easier, way is called the next hop method, in which the next hop attribute is set on the trigger and is sent in the route update to its iBGP peers at the edge. The second method is to use BGP communities. In the latter method, the trigger sets the BGP community for a route and sends it to the edge routers using iBGP. The edge routers use a route map to match this community and set attributes locally, such as next hop and other routing metrics.

Figure 8. Typical Remotely Triggered Black Hole Filtering Deployment

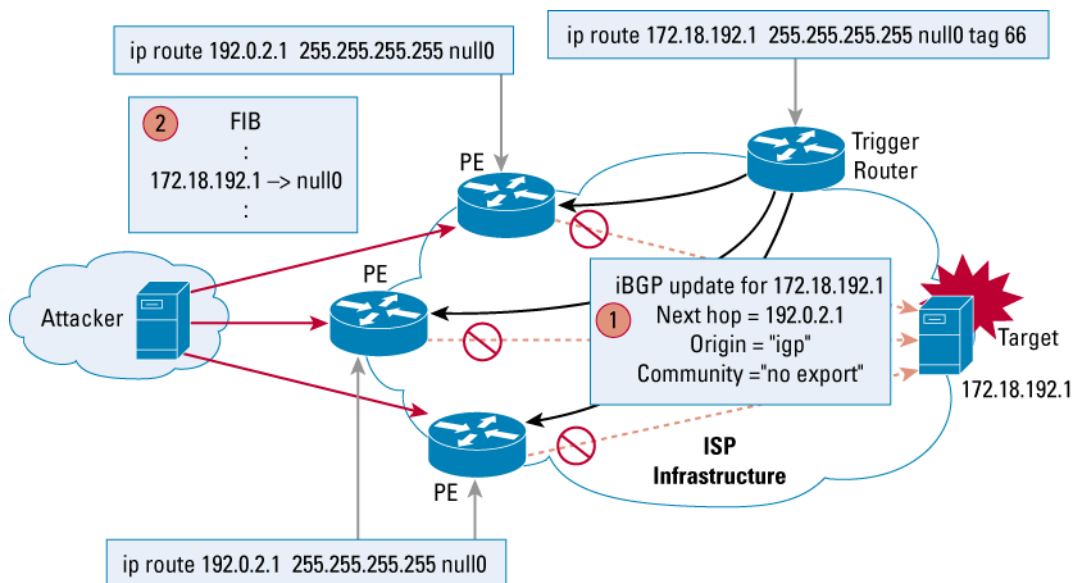


Next Hop Method

In the next hop method the trigger device sets the next hop route for the destination IP address that is to be black holed and then uses iBGP route update to propagate this route to all the edge routers (iBGP peers). The iBGP peers then set their next hop to the destination based on this update from the trigger. On every edge router, there is a static route for this next hop set to null0. Upon receiving a route update for the destination IP address, all edge routers set their next hops accordingly. The static route for the next hop effectively forwards all traffic for the black holed destination IP address to null0.

This process is illustrated in Figure 9.

Figure 9. Deploying Remotely Triggered Black Hole Filtering—Next Hop Method

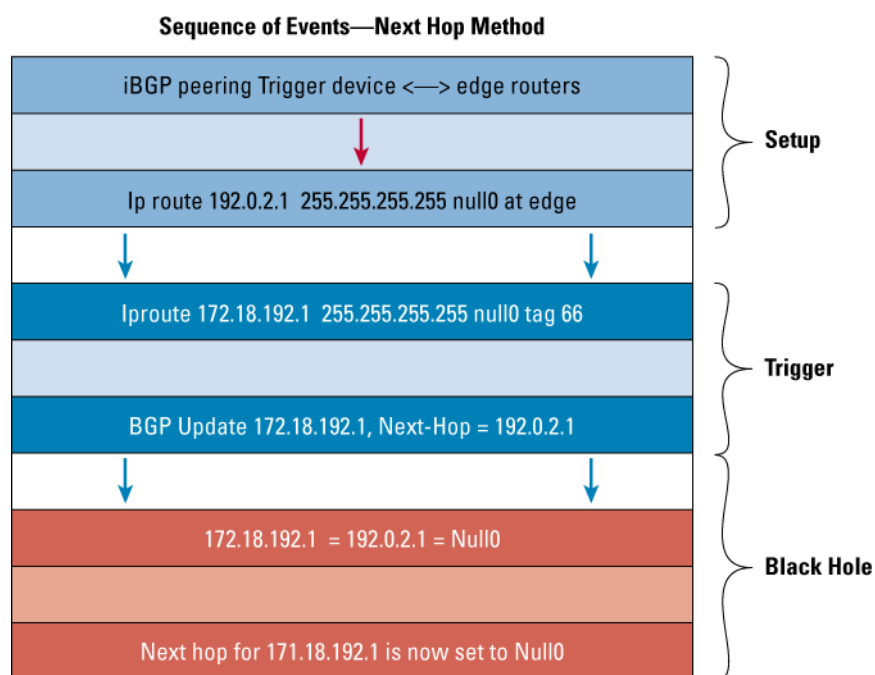


The sequence of events for next hop RTBH filtering is as follows:

1. An attack is targeted at 172.18.192.1.
2. A static route to the target IP address, 172.18.192.1, is added to the triggering device with the tag 66.
3. A route map in the trigger matches the tag 66 and sets the next hop to 192.0.2.1, origin to IGP, and community to no-export.
4. The trigger sends the route as an iBGP route update to its peers.
5. The edge routers that are peers receive the update and accordingly set the next hop to the target, 172.18.192.1, as 192.0.2.1.
6. Because the edge routers have static routes of 192.0.2.1 set to null0, the final FIB entry for the target IP address (172.18.192.1) is set to null0.
7. All future traffic to 172.18.192.1 will be forwarded to null0 and dropped.

This sequence of events is summarized in Figure 10.

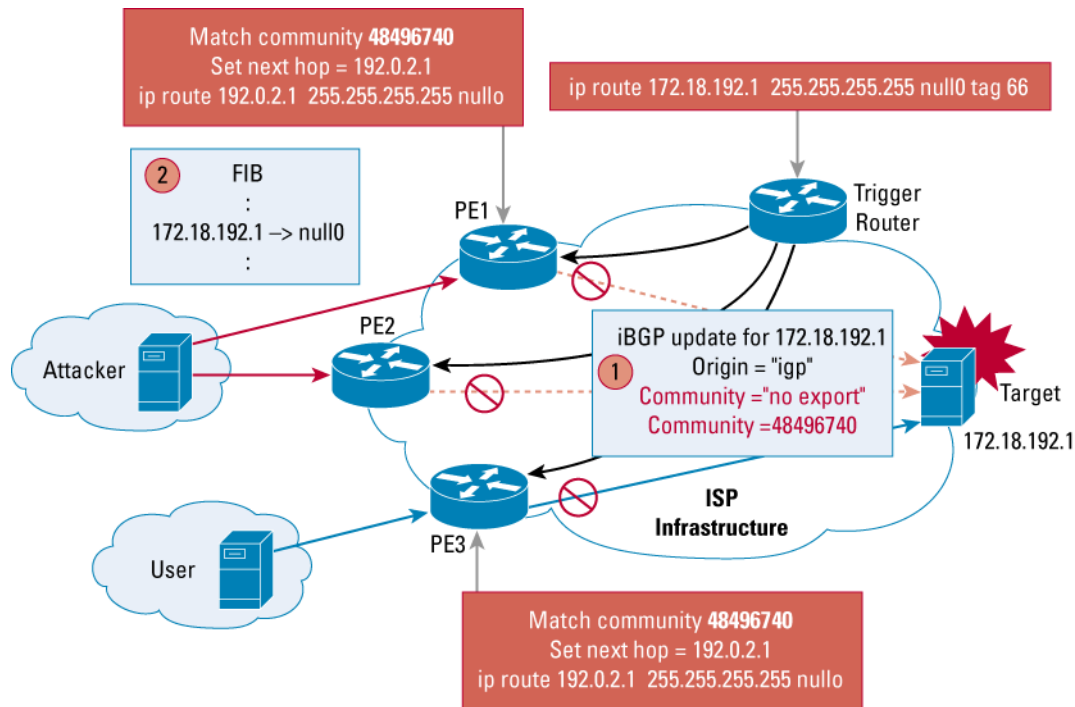
Figure 10. Sequence of Events—Next Hop Method



The Communities-Based Approach

Community-based triggering allows for better control over the drop process. In this technique, the trigger router is used to set different BGP community values and then send these values to its iBGP peers (BGP send-community) in its route update. The edge routers have the flexibility to act or not act on updates based on the community values. So, the decision to act or change route attributes such as next hop is based on community values, and the decision-making process is pushed out to the edge of the network, making it a highly flexible solution that can be used to selectively drop traffic, as shown in Figure 11 Deploying Remotely Triggered Black Hole Filtering—Communities Method.

Figure 11. Deploying Remotely Triggered Black Hole Filtering—Communities Method



Changing the community values effectively controls the points in the network where the undesirable traffic is dropped.

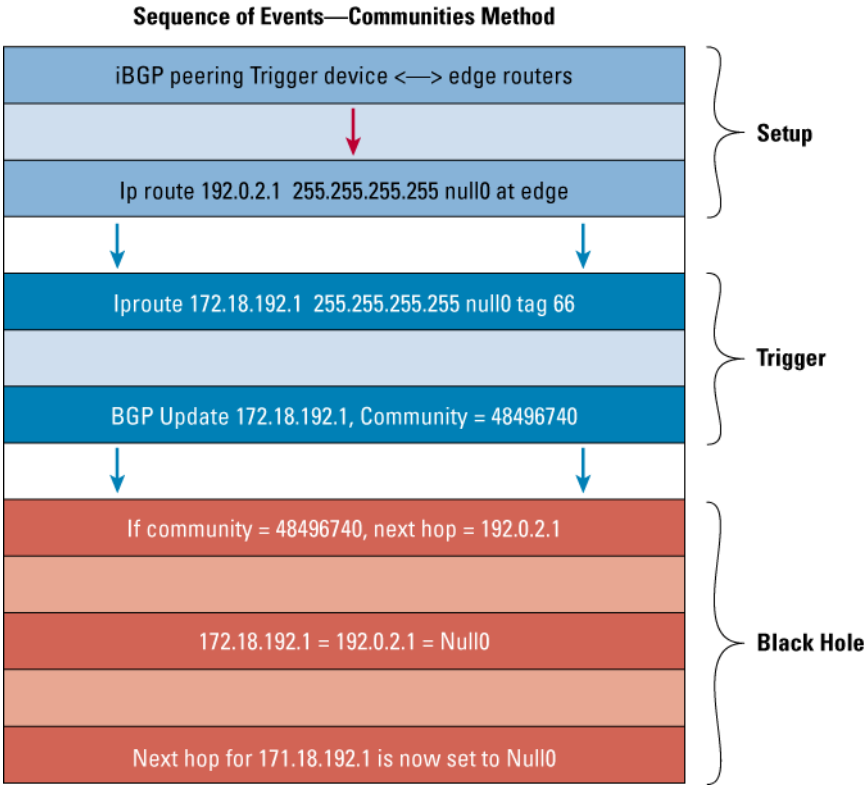
The sequence of events for using BGP community values for RTBH filtering is as follows:

1. The ingress points of the attack toward the target are determined to be PE-1 and PE-2.
2. A static route to the target IP address, 172.18.192.1, is added to the trigger with the tag 66.
3. A route map on the trigger matches the tag 66 and sets the BGP community to no-export and 48496740. A different value of the tag can be used to set a different value of community attribute for the route.
4. The updated route along with the community values are sent to all the iBGP peers. All iBGP peers apply a route map to incoming route updates.
5. The route maps on PE-1 and PE-2 match on the community value 48496740 and set the next hop of the route to 192.0.2.1. Since there is no match in the route map applied at PE-3, the incoming route update is ignored on PE-3.
6. Because the edge routers have static routes of 192.0.2.1 set to null0, the final FIB entry for the target IP address, 172.18.192.1, is set to null0.
7. All traffic destined to the target IP address, 172.18.192.1, is dropped at PE-1 and PE-2. Traffic is forwarded normally by PE-3.

If the community value is set to 48496840, then PE-3 will drop traffic. However, in this case, PE-1 and PE-2 will continue to forward traffic.

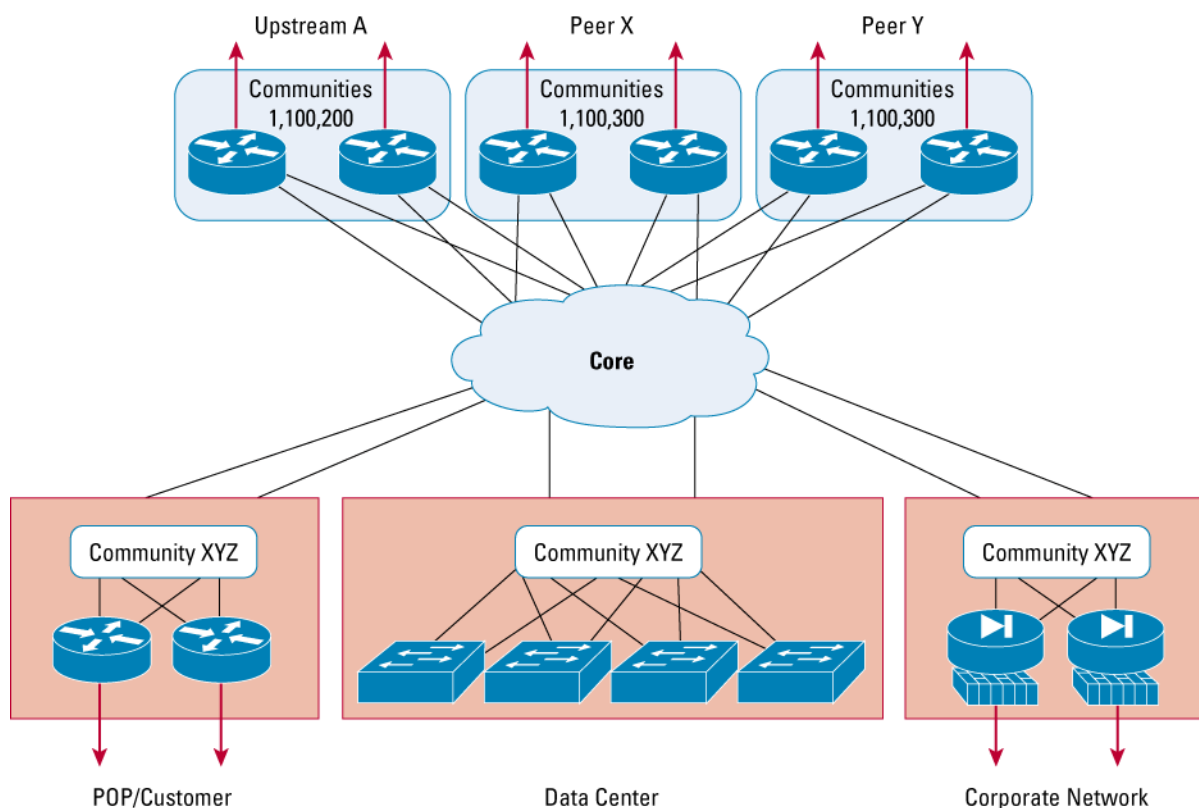
This sequence of events is illustrated in Figure 12.

Figure 12. Sequence of Events—Communities Method



A sample deployment of RTBH filtering using communities is shown in Figure 13. This example uses community 100 to drop all traffic coming from peering and upstream ISPs, a community of 300 to drop traffic coming from peering ISPs, a community value of 200 to drop traffic from upstream ISPs, and a community value of xyz if the source of the attack is from a customer or a point of presence (POP) site. A special community value of 1 will drop all traffic coming into the core.

Figure 13. Deploying the Communities Method for Remotely Triggered Black Hole Filtering



Operational Gains and Considerations

RTBH filtering is designed to drop undesirable traffic at the network policy edge. However, if the source of the attack can be traced easily, consider other techniques to quickly take out the suspect machine.

When used, this technique should be universally deployed to all the edge routers, or an attacker might find a path to the target by finding a hole in the SP filtering shield.

When a particular destination is globally black holed using destination-based RTBH filtering, all traffic to the target host is dropped at the edge. Although this limits the effect of the attack on the target and collateral damage to the SP network, it also prevents any legitimate traffic to the target host.

This type of blanket technique has serious implications for high-value targets such as core servers or voice gateways that are under contractual obligations of high availability. However, only the service under attack is not available, and other services are not affected. So, it provides for partial service recovery and an opportunity for the provider to address the threat without any further collateral damage to the network and other services.

BGP community-based triggering allows for more controlled drops. As an example, all routers for remote triggering purposes could be classified as "peering routers" or "customer edge routers." Three communities could then be used to trigger "all routers," "peering routers," or "customer routers." This type of deployment gives you better control as to where to drop the undesirable traffic. For example, if the source of the DoS attack is a peer, then only traffic from outside the AS is dropped at all the peering routers. Partial service is still maintained for all customers from within the AS until the threat has been mitigated.

Source-based RTBH filtering should be employed only if the source IP of the attacker can be identified and predicted within a specific address range. This is quite rare, as most attack traffic uses spoofed source IP addresses that constantly change.

RTBH filtering is a technique that is used to drop undesirable traffic, and it is absolutely critical that the technique, once deployed, not be exploited intentionally or unintentionally for dropping legitimate traffic. To protect from exploitation, deploying the secure BGP techniques is highly recommended. Secure BGP implementation techniques include features such as neighbor authentication, prefix filtering, and time to live (TTL) security hack. For further information, refer to the BGP documentation available on www.cisco.com.

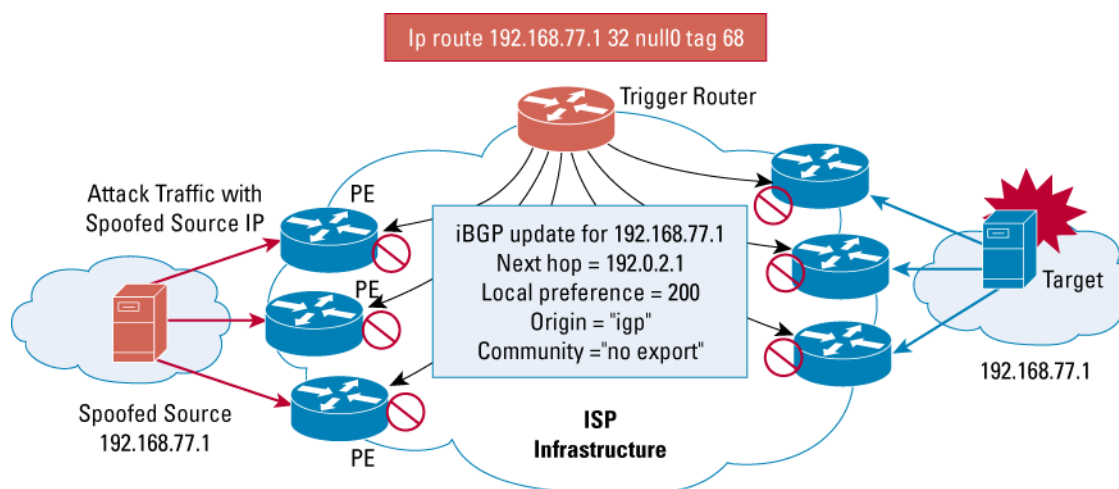
Prefix filters must be used at both the trigger routers and edge routers to make sure that essential services such as DNS are not black holed even by mistake.

There is no built-in application layer awareness with RTBH filtering. Traffic is filtered and dropped only based on source and destination IP addresses and not based on transport-specific or application-specific ports. So, it is not possible to drop only application-specific traffic, such as Telnet or HTTP. If a Web server is the intended target and the attack traffic is flooding the HTTP port on the server, all traffic to this server will be dropped at the network edge. If this server is also an anonymous FTP server, this service will also be affected because RTBH filtering lacks any application layer awareness.

When using source-based RTBH filtering, all traffic to and from the black holed source IP will be dropped. It is very important to double-check the source IP being black holed when doing source-based black holes to make sure that the address being black holed is the actual address of the attacker and is not being spoofed.

As shown in Figure 14, an attacker could target a legitimate IP address by spoofing it as the source of an attack and counting on the ISP to black hole the source using sourced-based RTBH filtering. For example, in Figure 15, an attacker targets a destination within the ISP with a spoofed IP address of 192.168.77.1, which is the real target. The ISP responds by black holing the source address, 192.168.77.1. This causes all traffic from the source, 192.168.77.1, to be dropped at the edges. However, in this scenario, it also causes all traffic to the destination, 192.168.77.1, to be dropped. This accomplishes the original intent of the attack, which was targeting the legitimate service at 192.168.77.1.


Figure 14. Exploiting Source-Based Remotely Triggered Black Hole Filtering



Drop Placement

Once a decision has been made to deploy this technique, various points in the ISP network need to be considered for dropping traffic. The red dots in Figure 15 indicate possible drop locations.

Figure 15. Remotely Triggered Black Hole Filtering Drop Placement



The following list summarizes some of the locations at which traffic might need to be dropped:

- Worm traffic from infected hosts in customer networks needs to be dropped at the distribution layer.
- Incoming attacks from outside the AS need to be dropped at the peering edges.
- Data centers need to be contained or protected from all undesirable traffic.

REMOTELY TRIGGERED BLACK HOLE FILTERING APPLICATIONS

This section describes the different applications for RTBH filtering. It includes the following topics:

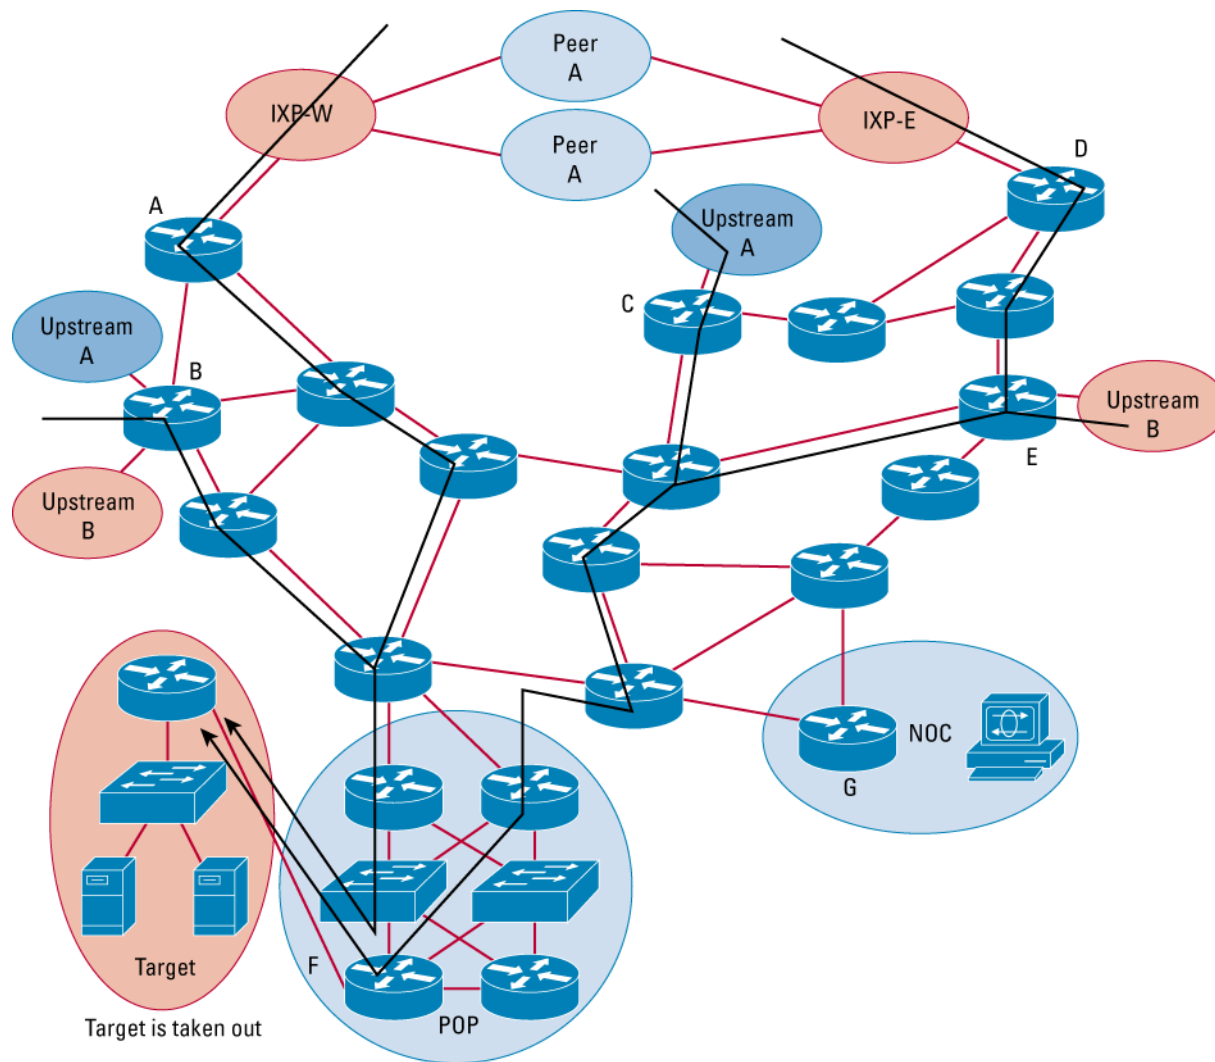
- Distributed Denial of Service and Worm Mitigation
- Quarantine or Redirection of Attack Traffic
- Blacklist Filters

Distributed Denial of Service and Worm Mitigation

Once a DDoS attack or a worm attack has been detected, RTBH filtering can be used to selectively drop undesirable traffic destined to the attack destination, as shown in the following diagrams.

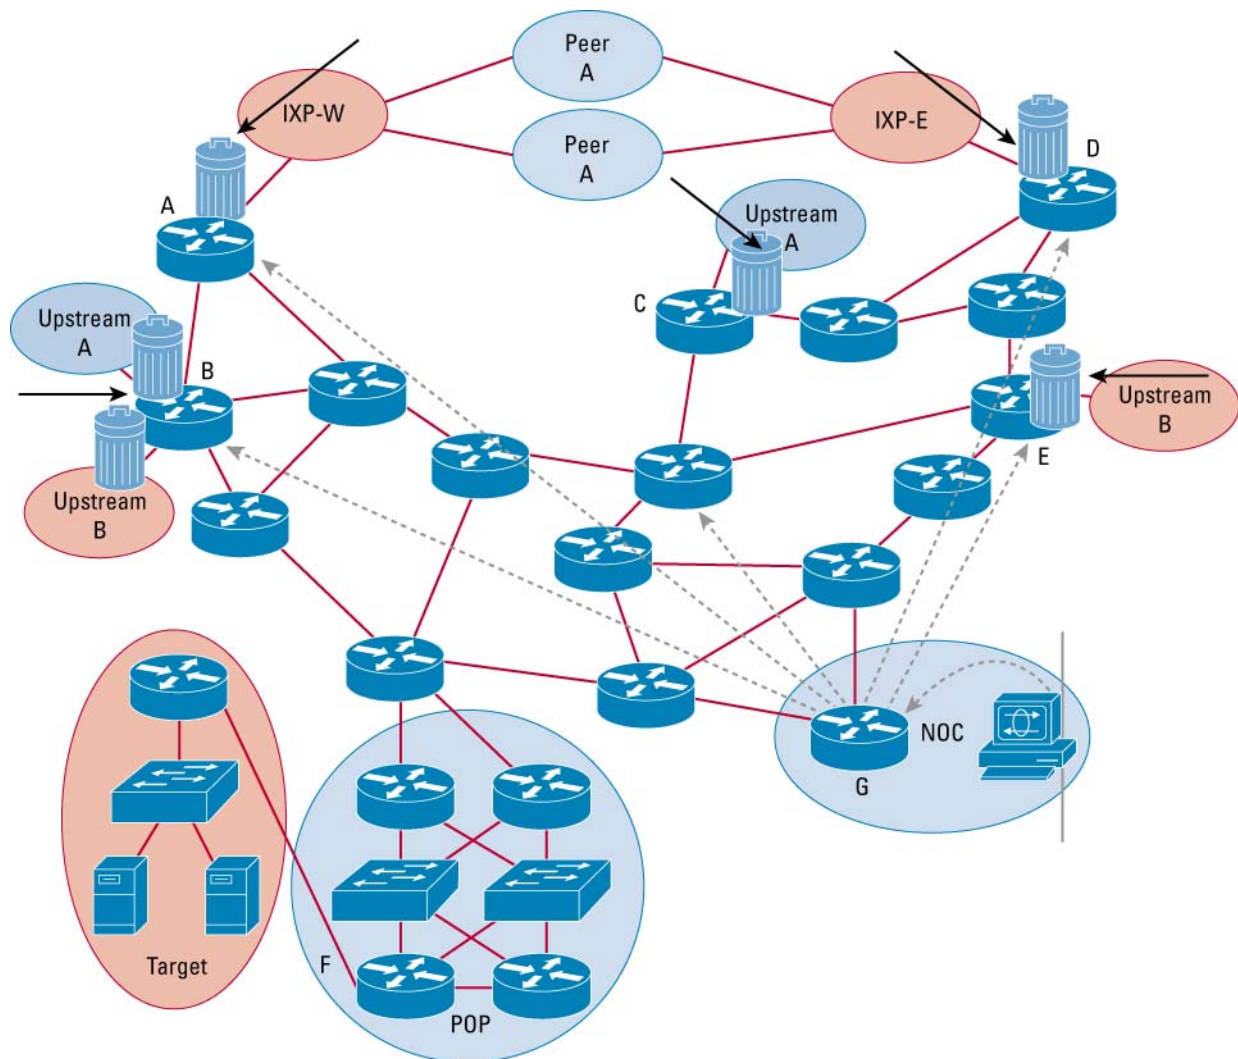
In Figure 16, the target is a victim of a DDoS attack. Destination-based RTBH filtering is employed in this example to drop all traffic to the target, thus preventing collateral damage to the ISP infrastructure and giving the customer and the service provider time to respond to the attack and to restore service.

Figure 16. Remotely Triggered Black Hole Filtering with a Distributed-Denial-of-Service Attack



In Figure 17, after the black hole has been remotely triggered, the traffic is being dropped at the edge of the network.

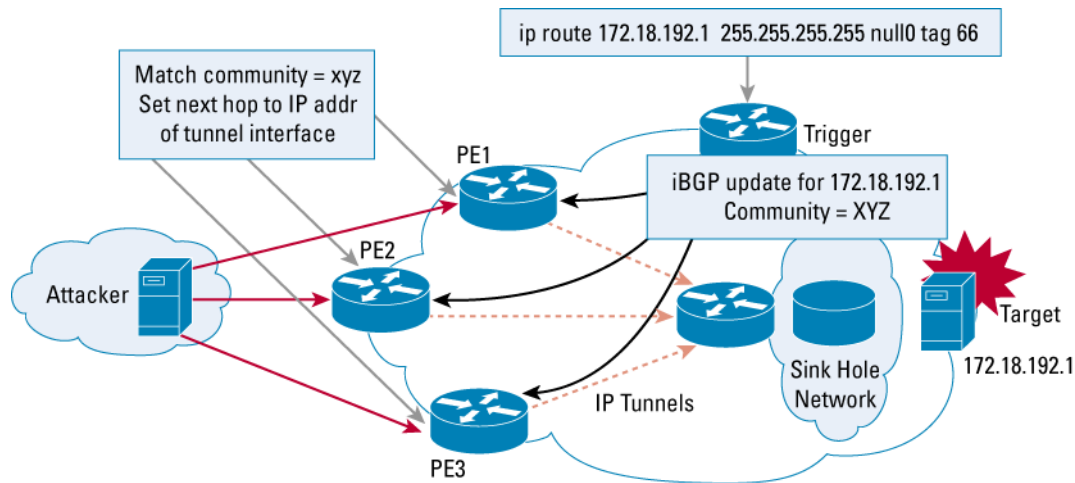
Figure 17. Traffic Being Dropped to Black Hole Destination



Quarantine or Redirection of Attack Traffic

Dropping traffic at the edges has its benefits, but to further understand the attack, it is useful to analyze some of the undesirable traffic. This is done by redirecting attack traffic to a sinkhole for further analysis, as shown in Figure 18. A BGP trigger is used to set the next hop to the target as the tunnel interface IP address. The other end of the tunnel is terminated on a router in the sinkhole.

Figure 18. Attack Traffic Quarantine



The following are some of the key points to consider when redirecting attack traffic:

- The edge routers have an IP tunnel to a router in the sinkhole network.
- The bandwidth across the tunnel interface is policed at a predetermined rate to limit the attack traffic toward the sinkhole.
- A static route for the target is manually entered at the triggering device with a tag. A route map at the triggering device matches the tag and sets the community for the route.
- Edge routers receive the updated route for the target and set the next hop for the destination to the IP address of the tunnel interface based on the community value.
- All attack traffic toward the destination is now forwarded to the tunnel toward the sinkhole and is quarantined. The traffic across the tunnel is limited by the policy enforced on the tunnel.

Instead of redirecting attack traffic to a sinkhole, it could be redirected to a packet scrubber such as the Cisco Guard DDoS Mitigation Appliance, which analyses traffic and reinjects legitimate traffic back into the network. Cisco Guard drops undesirable traffic directed to the attack target by recognizing well-known characteristics of attack traffic.

Blacklist Filters

The ability to drop traffic based on destination addresses can also be applied to enforce blacklists at the customer edge. A blacklist is a set of destinations to which an organization does not want to allow access. These destination IP addresses can be black holed at the customer edge, which prevents access from all users within the organization. If a destination is known to be a source of worms or is infected, it could be temporarily blacklisted to protect ISP customers and users. However, in many cases, the source address could be spoofed, in which case blacklisting will drop legitimate traffic to a valid destination. Therefore, potential blacklist destinations must be carefully checked prior to being blacklisted.

SAMPLE CONFIGURATIONS

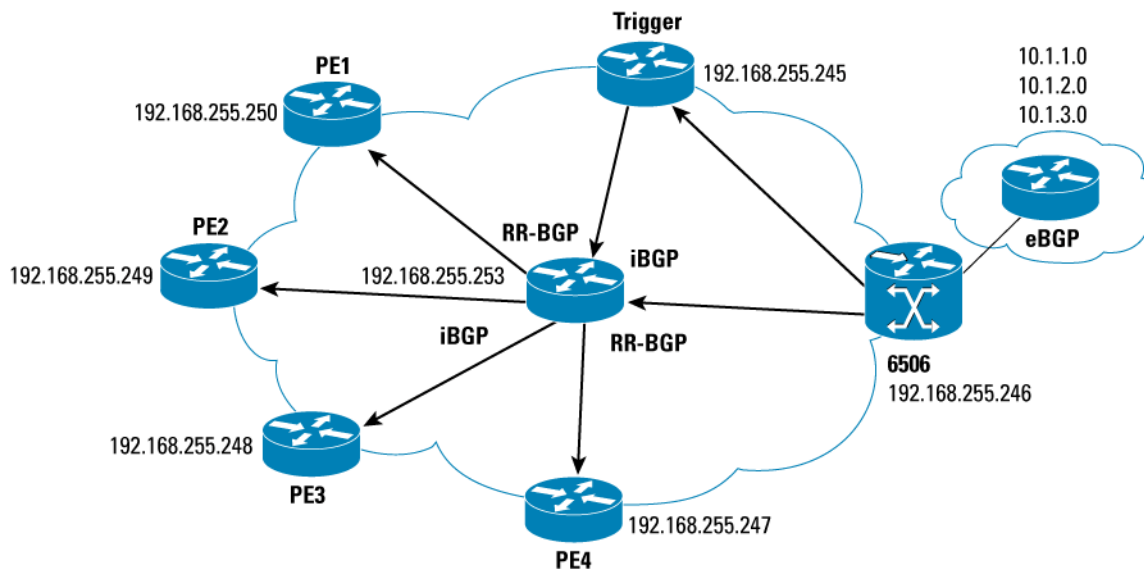
This section provides sample configurations for different RTBH filtering implementations. It includes the following topics:

- Global Black Hole
- Regionalized Black Holes
- Alternate Configuration for Regionalized Black Holes

Global Black Holes

This section includes working and tested configurations used for setting up RTBH filtering (Figure 19). Once the trigger is applied by installing a static route on the route reflector, ALL traffic to the target destination is dropped at ALL the edges of the network.

Figure 19. Remotely Triggered Black Hole Filtering Implementation—Global Black Hole



Configuration Procedure

The following is the general configuration procedure for implementing a global black hole. Set up the trigger router exclusively to trigger the application of black holes.

Step 1. Create a route map called “black-hole-trigger” that matches specific route tags and sets route characteristics.

A route with tag 66 is matched, and the next hop is set to 192.0.2.1, origin set to igp, community set to no-export, and local preference set to 200.

Step 2. The trigger router has an iBGP peer relationship with all edge routers.

If route reflectors are used for scalability reasons, then the trigger router has iBGP peer relationships with all route reflectors in the network.

Step 3. Static routes are redistributed into BGP after applying the route map described in step 1.

Trigger Router

This section provides a sample configuration listing for the trigger router. In this configuration, peer groups are used because it is the preferred way to configure a large number of BGP peers with similar characteristics. You must set the **send-community** for all these peers so they receive the no-export community and respect it by not advertising this redistributed route to any of their external peers.

Also, make sure to set **no auto-summary** so that specific host routes can be black holed. Otherwise BGP will automatically summarize the route based on class boundaries. Static routes are then redistributed into BGP after applying the black hole-trigger route map.

```
trigger#sh run
Building configuration...
version 12.1
!
hostname trigger
!
ip subnet-zero
!
interface Loopback0
 ip address 192.168.255.245 255.255.255.255
!
interface Null0
 no ip unreachable
!
interface Ethernet0/0
 ip address 192.168.4.3 255.255.255.0
 half-duplex
!
router ospf 100
 log-adjacency-changes
 redistribute connected subnets
 network 192.168.4.0 0.0.0.255 area 0

router bgp 740
 no synchronization
 bgp log-neighbor-changes
 redistribute static route-map black-hole-trigger
 neighbor black-hole peer-group
 neighbor black-hole remote-as 740
 neighbor black-hole update-source Loopback0
 neighbor black-hole send-community
 neighbor 192.168.255.246 remote-as 740
 neighbor 192.168.255.246 update-source Loopback0
 neighbor 192.168.255.253 peer-group black-hole
 no auto-summary
```

```
!  
ip route 192.0.2.1 255.255.255.255 Null0  
  
route-map black-hole-trigger permit 10  
  match tag 66  
  set ip next-hop 192.0.2.1  
  set local-preference 200  
  set origin igp  
  set community no-export  
!  
route-map black-hole-trigger deny 25  
!  
no scheduler allocate  
end
```

The last part of this configuration creates a route map to match the route tag 66 and sets route characteristics. A higher value of local preference is desired for choosing a route, so it is set to 200, which is greater than the default value of 100. Also, to make sure that other static routes are not affected by this route map, a deny statement is placed at the end.

Edge Router (Route Reflector Client)

The following is a sample configuration for the edge router. This configuration creates a static route to 192.0.2.1 pointed to Null0; 192.0.2.1 is the next hop set by the trigger router to the black holed destination.

```
PE-2#sh run  
Building configuration...  
hostname PE-2  
!  
interface Loopback0  
  ip address 192.168.255.249 255.255.255.255  
!  
!Create a Null0 interface and turn off ICMP unreachable generation.  
interface Null0  
  no ip unreachable  
!  
!  
interface Ethernet2/0  
  ip address 192.168.2.2 255.255.255.0  
  half-duplex  
  no cdp enable  
!  
router ospf 100  
  log-adjacency-changes  
  redistribute connected subnets  
  network 192.168.2.0 0.0.0.255 area 0
```

```
!  
router bgp 740  
  no synchronization  
  bgp log-neighbor-changes  
  neighbor 192.168.255.253 remote-as 740  
  neighbor 192.168.255.253 update-source Loopback0  
  no auto-summary  
  
ip classless  
ip route 192.0.2.1 255.255.255.255 Null0  
!
```

Verifying Configuration

To verify that the configuration works, try to black hole a test destination. Verification requires adding the static route at the trigger router and then observing that the route update is received at the edge and installed with the appropriate next hop in the routing table.

It is equally important to verify that the target can be put back into service by deleting the static route at the trigger router. The steps involved with the verification are summarized below.

Step 1. A static route is added on the route reflector to trigger the black hole at the edge.

```
trigger(config)#ip route 192.168.1.100 255.255.255.255 null0 tag 66
```

Step 2. On the edge router, you can observe the update being received and the route to the target destination being inserted into the IP routing table with the next hop set to 192.0.2.1. The following shows the sample output:

```
PE-2#debug ip routing  
IP routing debugging is on  
PE-2#debug bgp up in  
BGP updates debugging is on (inbound)  
*Mar  1 22:26:27.750: BGP(0): 192.168.255.253 rcvd UPDATE w/ attr: nexthop 192.0.2.1, origin i,  
localpref 200, metric 0, community no-export  
*Mar  1 22:26:27.754: BGP(0): 192.168.255.253 rcvd 192.168.1.100/32  
*Mar  1 22:26:27.754: BGP(0): Revise route installing 1 of 1 route for 192.168.1.100/32 -> 192.0.2.1  
to main IP table  
*Mar  1 22:26:27.754: RT: network 192.168.1.0 is now variably masked  
*Mar  1 22:26:27.754: RT: add 192.168.1.100/32 via 192.0.2.1, bgp metric [200/0]  
*Mar  1 22:26:27.754: RT: NET-RED 192.168.1.100/32  
*Mar  1 22:26:27.754: RT: NET-RED queued, Queue size 1  
*Mar  1 22:26:31.606: RT: NET-RED 0.0.0.0/0  
*Mar  1 22:26:31.606: RT: NET-RED queued, Queue size 1  
The route can be seen in the IP routing table pointing to the next hop of 192.0.2.1  
PE-2#sh ip route
```

Codes: C-connected, S-static, R-RIP, M-mobile, B-BGP
D-EIGRP, EX-EIGRP external, O-OSPF, IA-OSPF inter area
N1-OSPF NSSA external type 1, N2-OSPF NSSA external type 2
E1-OSPF external type 1, E2-OSPF external type 2
i-IS-IS, L1-IS-IS level-1, L2-IS-IS level-2, ia-IS-IS inter area
*-candidate default, U-per-user static route, o-ODR
P-periodic downloaded static route

Gateway of last resort is 192.168.2.1 to network 0.0.0.0

```
O    192.168.4.0/24 [110/12] via 192.168.2.1, 21:34:37, Ethernet2/0
O    192.168.5.0/24 [110/12] via 192.168.2.1, 21:34:37, Ethernet2/0
    10.0.0.0/24 is subnetted, 3 subnets
B      10.1.3.0 [200/0] via 192.168.255.246, 17:12:19
B      10.1.2.0 [200/0] via 192.168.255.246, 17:12:19
B      10.1.1.0 [200/0] via 192.168.255.246, 17:12:19
    192.168.255.0/32 is subnetted, 8 subnets
O E2   192.168.255.247 [110/20] via 192.168.2.1, 21:34:37, Ethernet2/0
O      192.168.255.246 [110/13] via 192.168.2.1, 21:34:38, Ethernet2/0
O E2   192.168.255.253 [110/20] via 192.168.2.1, 21:34:38, Ethernet2/0
O      192.168.255.252 [110/12] via 192.168.2.1, 21:34:38, Ethernet2/0
O      192.168.255.251 [110/11] via 192.168.2.1, 21:34:38, Ethernet2/0
O E2   192.168.255.250 [110/20] via 192.168.2.3, 21:34:38, Ethernet2/0
C      192.168.255.249 is directly connected, Loopback0
O E2   192.168.255.248 [110/20] via 192.168.2.1, 21:34:41, Ethernet2/0
    192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
B      192.168.1.100/32 [200/0] via 192.0.2.1, 00:00:22
O      192.168.1.0/24 [110/11] via 192.168.2.1, 21:34:41, Ethernet2/0
    192.0.2.0/32 is subnetted, 1 subnets
S      192.0.2.1 is directly connected, Null0
C      192.168.2.0/24 is directly connected, Ethernet2/0
O      192.168.100.0/24 [110/13] via 192.168.2.1, 21:34:41, Ethernet2/0
O      192.168.3.0/24 [110/11] via 192.168.2.1, 21:34:41, Ethernet2/0
O*E2   0.0.0.0/0 [110/1] via 192.168.2.1, 21:34:41, Ethernet2/0
B      192.168.0.0/16 [200/0] via 192.168.255.246, 17:12:24
```


Step 3. The route has a community set to no-export to make sure it is not readvertised by the edge router to any external peers it may have.

```
PE-2#sh ip bgp community no-export
```

```
BGP table version is 63, local router ID is 192.168.255.249
```

```
Status codes: s-suppressed, d-damped, h-history, * valid, > best, i-internal,  
               r RIB-failure
```

```
Origin codes: i-IGP, e-EGP, ?-incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i192.168.1.100/32	192.0.2.1	0	200	0	i

Step 4. The target does not need to be black holed any longer, so the static route is removed from the trigger router.

```
trigger# no ip route 192.168.1.100 255.255.255.255 null0 tag 66
```

Step 5. The iBGP update withdrawing the route can be observed at the edge, and the route is removed from the IP routing table.

```
PE-2#debug ip routing
```

```
IP routing debugging is on
```

```
PE-2#debug ip bgp up in
```

```
BGP updates debugging is on (inbound)
```

```
PE-2#
```

```
*Mar 1 22:31:19.338: BGP(0): 192.168.255.253 rcv UPDATE about 192.168.1.100/32 -- withdrawn
```

```
*Mar 1 22:31:19.338: BGP(0): no valid path for 192.168.1.100/32
```

```
*Mar 1 22:31:19.338: BGP(0): nettable_walker 192.168.1.100/32 no best path
```

```
*Mar 1 22:31:19.338: RT: del 192.168.1.100/32 via 192.0.2.1, bgp metric [200/0]
```


```
*Mar 1 22:31:19.338: RT: delete subnet route to 192.168.1.100/32
```

Regionalized Black Holes

With a simple black hole, all traffic to the target destination is dropped at the edge of the network. The service provider may want to be able to control the application of the black holes and limit them to a predefined region of edge routers if they can determine that the attack is coming in from a specific region. This would allow uninterrupted traffic from other regions of the network to the target host or network while mitigating the attack. This section provides sample configurations for such a deployment.

As shown in Figure 20, the set of four edge routers is divided into two regions. To black hole a target in region 1, the tag required to be attached to the static route is 66. If an edge router in region 2 is suspected to be forwarding attack traffic, the tag to be attached to the static route is 76. If the attack is coming in from all the edge routers or cannot be regionalized, the static route with an attached tag of 86 will drop all traffic destined to the target from all four edge routers.

Figure 20. Remotely Triggered Black Hole Filtering Implementation—Regionalized Black Holes



Based on the tag for each of the static routes, different communities are set using a route map on the trigger router. The route updates are sent to all iBGP peers from the route reflector (RR-BGP). Each iBGP peer applies a route map on all incoming advertisements and either accepts or rejects the route after matching the community string. To avoid the possibility of externally generated routes having the same community, only routes that are originated in the local AS are accepted. The community is reset to no-advertise to make sure that this route is never readvertised using BGP.

Configuration Procedure

The following is the general configuration procedure for implementing RTBH filtering with regionalized black holes.

Step 1. Create a route map called “black-hole-trigger” that matches the three tags 66, 76, and 86.

This route map sets the next hop to 192.0.2.1, the origin to igp, and the local preference to 200. If the route tag matches 66, the community is set to 740:100 and no-export; if the route tag matches 76, the community is set to 740:200 and no-export; and if the route tag matches 86, the community is set to 740:300 and no-export. No other route is modified in any way.

Step 2. Create a peer group called “black-hole” with the update source set to loopback0.

Communities are sent to neighbors associated with this peer group. This is the peer group used to peer with all route reflectors or all iBGP peers if route reflectors are not used.

Step 3. Create a statement on the trigger to redistribute static routes into BGP after the route map “black-hole-trigger,” created in step 1, is applied to it.

Step 4. On the edge routers, create a route map called “black-hole-create.”

All routes are denied that match the community assigned to another region. Therefore, on a router in region 1, all routes are denied that match the community 740:200. All routes are permitted that match either of the communities 740:100 and 740:300 and originating from the same AS, which is 740 in this example. The next hop route is set to 192.0.2.1, and the community is reset to no-advertise. No other route is permitted or altered in any way.

Step 5. On the edge routers, apply the route map created in step 4 to all incoming routes.

Trigger Router

The following is a configuration listing for the trigger router when implementing regionalized black holes. In this listing, a peer group called “black-hole” is created with the update source set to loopback0. Communities are sent to neighbors associated with this peer group. This is the peer group used to peer with all route reflectors or all iBGP peers if route reflectors are not used. Static routes are redistributed into BGP using the route map “black-hole-trigger.”

```
trigger#sh run
hostname trigger
!
interface Loopback0
 ip address 192.168.255.245 255.255.255.255
!
interface Null0
 no ip unreachable
!
interface Ethernet0/0
 ip address 192.168.4.3 255.255.255.0
!
router ospf 100
 log-adjacency-changes
 redistribute connected subnets
 network 192.168.4.0 0.0.0.255 area 0

router bgp 740
 no synchronization
 bgp log-neighbor-changes
 redistribute static route-map black-hole-trigger
 neighbor black-hole peer-group
 neighbor black-hole remote-as 740
 neighbor black-hole update-source Loopback0
 neighbor black-hole send-community
 neighbor 192.168.255.246 remote-as 740
 neighbor 192.168.255.246 update-source Loopback0
 neighbor 192.168.255.253 peer-group black-hole
 no auto-summary
```

```
!  
ip route 192.0.2.1 255.255.255.255 Null0
```

```
route-map black-hole-trigger permit 10  
  match tag 66  
  set ip next-hop 192.0.2.1  
  set local-preference 200  
  set origin igp  
  set community 48496740 no-export
```

```
!  
route-map black-hole-trigger permit 15  
  match tag 76  
  set ip next-hop 192.0.2.1  
  set local-preference 200  
  set origin igp  
  set community 48496840 no-export
```

```
!  
route-map black-hole-trigger permit 20  
  match tag 86  
  set ip next-hop 192.0.2.1  
  set local-preference 200  
  set origin igp  
  set community 48496940 no-export
```

```
!  
route-map black-hole-trigger deny 25
```

The last part of this configuration creates a route map called “Black-hole-trigger” that matches the three tags 66, 76, and 86. It sets the next hop to 192.0.2.1, the origin to “igp,” and the local preference to 200. If the route tag matches 66, the community is set to 740:100 and no-export; if the route tag matches 76, the community is set to 740:200 and no-export; and if the route tag matches 86, the community is set to 740:300 and no-export. No other route is modified in any way.

Edge Router Configuration for PE-1

The following is a sample configuration listing for the edge router (PE-1) used in this example. This listing applies the route map “black-hole-create” to all incoming routes from the route reflector.

```
PE-1#sh run  
Building configuration...  
  
Current configuration : 1818 bytes  
!  
version 12.2  
no parser cache  
no service single-slot-reload-enable
```

```
service timestamps debug uptime
service timestamps log uptime
!
hostname PE-1
!
enable password XXXX
ip subnet-zero
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
interface Loopback0
 ip address 192.168.255.250 255.255.255.255
!
interface Null0
 no ip unreachable
!
interface FastEthernet2/0
 ip address 192.168.2.3 255.255.255.0
 duplex auto
 speed auto
!
router ospf 100
 log-adjacency-changes
 redistribute connected subnets
 network 192.168.2.0 0.0.0.255 area 0

router bgp 740
 no synchronization
 bgp log-neighbor-changes
 neighbor 192.168.255.253 remote-as 740
 neighbor 192.168.255.253 update-source Loopback0
 neighbor 192.168.255.253 route-map black-hole-create in
!
ip classless
ip route 192.0.2.1 255.255.255.255 Null0
ip http server
ip community-list 1 permit 740:100
ip community-list 1 permit 740:300
ip community-list 2 permit 740:200
ip as-path access-list 1 permit ^$
```

```
route-map black-hole-create deny 5
  match community 2
!
route-map black-hole-create permit 10
  match community 1
  match as-path 1
  set ip next-hop 192.0.2.1
  set community no-advertise
!
route-map black-hole-create permit 15
!
!
dial-peer cor custom
!
line con 0
line aux 0
line vty 0 4
  password XXXX
  login
end
```

The last part of this configuration creates a route map called “black-hole-create.” This denies all routes that match the community 740:200. All routes are permitted that match either of the communities 740:100 and 740:300 and that originate from the same AS, which is 740 in this example. The next hop route is set to 192.0.2.1, and the community is reset to no-advertise. All other routes are permitted and are not altered in any way.

Configuration for PE-3

The following is a sample configuration listing for PE-3 used in this example.

```
PE-3#sh run
hostname PE-3
!
interface Loopback0
  ip address 192.168.255.248 255.255.255.255
!
interface FastEthernet0/0
  ip address 192.168.3.2 255.255.255.0
  duplex auto
  speed auto
!
router ospf 100
  log-adjacency-changes
  redistribute connected subnets
  network 192.168.3.0 0.0.0.255 area 0
```

```
!  
router bgp 740  
  no synchronization  
  bgp log-neighbor-changes  
  neighbor 192.168.255.253 remote-as 740  
  neighbor 192.168.255.253 update-source Loopback0  
  neighbor 192.168.255.253 route-map black-hole-create in  
!  
ip classless  
ip http server  
ip community-list 1 permit 740:200  
ip community-list 1 permit 740:300  
ip community-list 2 permit 740:100  
ip as-path access-list 1 permit ^$  
  
route-map black-hole-create deny 5  
  match community 2  
!  
route-map black-hole-create permit 10  
  match as-path 1  
  match community 1  
  set ip next-hop 192.0.2.1  
  set community no-advertise  
!  
route-map black-hole-create permit 15  
!
```

The last part of this configuration creates a route map called “black-hole-create.” This route map denies all routes that match the community 740:100. All routes are permitted that match either of the communities 740:200 and 740:300 and that originate from the same AS, which is 740 in this example. The next hop route is set to 192.0.2.1, and the community is reset to no-advertise. All other routes are permitted and are not altered in any way.

Verifying Configuration

With regionalized black holing, it is important to make sure that only the PEs in the suspect region black hole the destination and not all edge routers. As before, we would need to test route insertion and withdrawal. These verification steps are shown below.

Step 1. A static route is added on the route reflector to trigger the black hole at the edge trigger.

```
(config)#ip route 192.168.1.100 255.255.255.255 null0 tag 76
```

Step 2. On the edge routers (PE-3 and PE-4) in region 2, you can observe the update being received and the route to the target destination being inserted into the IP routing table, with the next hop set to 192.0.2.1.

```
PE-3#debug ip routing
IP routing debugging is on
PE-3#debug ip bgp up in
BGP updates debugging is on (inbound)
PE-3#
1w3d: BGP(0): 192.168.255.253 rcvd UPDATE w/ attr: nexthop 192.0.2.1, origin i,localpref 200, metric
0, originator 192.168.255.245, clusterlist 192.168.255.253, community no-export
1w3d: BGP(0): 192.168.255.253 rcvd 192.168.1.100/32
1w3d: BGP(0): Revise route installing 192.168.1.100/32 -> 192.0.2.1 to main IP table
1w3d: RT: network 192.168.1.0 is now variably masked
1w3d: RT: add 192.168.1.100/32 via 192.0.2.1, bgp metric [200/0]
```

Step 3. On PE-1 and PE-2, make sure that the update is denied and not added to the IP routing table.

```
PE-1#debug ip routing
IP routing debugging is on
PE-1#debug ip bgp up in
BGP updates debugging is on (inbound)
PE-1#
1w4d: BGP(0): 192.168.255.253 rcvd UPDATE w/ attr: nexthop 192.0.2.1, origin i,localpref 200, metric
0, originator 192.168.255.245, clusterlist 192.168.255.253, community no-export
1w4d: BGP(0): 192.168.255.253 rcvd 192.168.1.100/32 -- DENIED due to: route-map;
```

Step 4. The black hole can be removed by removing the static route at the trigger router as follows.

On PE-3 the route to the specific host (192.168.1.100) is deleted from its routing table. On PE-1, you see the update coming in, but since the route was never installed, nothing else is done.

```
trigger(config)#no ip route 192.168.1.100 255.255.255.255 null0 tag 76
```

```
PE-3#
1w3d: BGP(0): 192.168.255.253 rcv UPDATE about 192.168.1.100/32 -- withdrawn
1w3d: BGP(0): no valid path for 192.168.1.100/32
1w3d: BGP(0): nettable_walker 192.168.1.100/32 no best path
```



```
1w3d: RT: del 192.168.1.100/32 via 192.0.2.1, bgp metric [200/0]
```

```
1w3d: RT: delete subnet route to 192.168.1.100/32
```

```
PE-1#
```

```
1w4d: BGP(0): 192.168.255.253 rcv UPDATE about 192.168.1.100/32 -- withdrawn
```

Alternate Configuration for Regionalized Black Holes

If the provider does not have default routes at the network edges, it is possible to regionalize the application of black hole triggers by simply using a different next hop on the trigger router for different tags. The edge router will black hole traffic to this destination, depending on if it has a route to the next hop or not.

If a BGP speaker receives a route with the next hop set, but does not have a route to the next hop, it will not install the route because it fails the next hop test. The following example illustrates how this is done.

Step 1. On the trigger router, set the next hop address as follows:

```
Match tag 66, set next hop 192.0.2.1, set community no-export, set origin igp
```

```
Match tag 76, set next hop 192.0.2.2, set community no-export, set origin igp
```

```
Match tag 86, set next hop 192.0.2.3, set community no-export, set origin igp
```

Step 2. On the edge routers PE-1 and PE-2 (region 1), set static routes to 192.0.2.1 and 192.0.2.2 to null0.

Step 3. On the edge routers PE-2 and PE-3 (region 2), set static routes to 192.0.2.1 and 192.0.2.3 to null0.

When a static route is added at the trigger for a host route that is to be black holed with a tag of 66, the following occurs:

- All routers in both regions have a static route to the next hop set to null0.
- All routers in both regions will install the route in the routing table following the BGP update.
- Traffic to the target host will be black holed.

However, when a static route is added with a tag of 76, all routers will receive the routing update from the trigger, but only routers in region 1 have a static route to the next hop and will install the route in their routing table. Routers in region 2 (PE-3 and PE-4) will not install the route because there is no route to the BGP next hop (next hop check fails), and they continue forwarding traffic to the host. However, as mentioned earlier, this is only possible if the routers in region 2 do not have a default route. If there is a default route, the route will be installed, because the BGP next hop check passes, and all traffic to the host will be dropped.

The Trigger Router

The following is the sample listing for the trigger router in the alternate configuration for regionalized black holes.

```
trigger#sh run
Building configuration...
hostname trigger
!
interface Loopback0
 ip address 192.168.255.245 255.255.255.255
```

```
!  
interface Null0  
  no ip unreachable  
!  
interface Ethernet0/0  
  ip address 192.168.4.3 255.255.255.0  
  half-duplex  
!  
router ospf 100  
  log-adjacency-changes  
  redistribute connected subnets  
  network 192.168.4.0 0.0.0.255 area 0  
!  
router bgp 740  
  no synchronization  
  bgp log-neighbor-changes  
  redistribute static route-map black-hole-trigger  
  neighbor black-hole peer-group  
  neighbor black-hole remote-as 740  
  neighbor black-hole update-source Loopback0  
  neighbor black-hole send-community  
  neighbor 192.168.255.246 remote-as 740  
  neighbor 192.168.255.246 update-source Loopback0  
  neighbor 192.168.255.253 peer-group black-hole  
  no auto-summary  
!  
ip kerberos source-interface any  
ip classless  
ip route 192.0.2.1 255.255.255.255 Null0  
ip route 192.0.2.2 255.255.255.255 Null0  
ip route 192.0.2.3 255.255.255.255 Null0  
ip route 192.168.1.100 255.255.255.255 Null0 tag 86  
no ip http server  
!  
route-map black-hole-trigger permit 10  
  match tag 66  
  set ip next-hop 192.0.2.1  
  set local-preference 200  
  set origin igp  
  set community no-export  
!  
route-map black-hole-trigger permit 15  
  match tag 76
```

```
set ip next-hop 192.0.2.2
set local-preference 200
set origin igp
set community no-export
!
route-map black-hole-trigger permit 20
  match tag 86
  set ip next-hop 192.0.2.3
  set local-preference 200
  set origin igp
  set community no-export
!
route-map black-hole-trigger deny 25
```

The Edge Router

The following is the sample listing for the edge router in the alternate configuration for regionalized black holes.

```
PE-2#sh run
Building configuration...

hostname PE-2
!
interface Loopback0
  ip address 192.168.255.249 255.255.255.255
!
interface Null0
  no ip unreachable
!
interface Ethernet2/0
  ip address 192.168.2.2 255.255.255.0
  half-duplex
  no cdp enable
router ospf 100
  log-adjacency-changes
  redistribute connected subnets
  network 192.168.2.0 0.0.0.255 area 0
!
router bgp 740
  no synchronization
  bgp log-neighbor-changes
  neighbor 192.168.255.253 remote-as 740
  neighbor 192.168.255.253 update-source Loopback0
  no auto-summary
```

```
!  
ip classless  
ip route 192.0.2.1 255.255.255.255 Null0  
ip route 192.0.2.2 255.255.255.255 Null0  
ip http server  
!
```

Verifying Configuration

The following procedure summarizes how to verify the alternate configuration for regionalized black holes.

Step 1. Set the trigger by adding a static route on the trigger router.

```
trigger(config)#ip route 192.168.1.100 255.255.255.255 null0 tag 76
```

Step 2. On the edge router, you can see the route being installed:

```
PE-2#  
*Mar  8 03:13:48.853: BGP(0): 192.168.255.253 rcvd UPDATE w/ attr: nexthop 192.0.2.2, origin i,  
localpref 200, metric 0, originator 192.168.255.245, clusterlist 192.168.255.253, community no-  
export  
*Mar  8 03:13:48.853: BGP(0): 192.168.255.253 rcvd 192.168.1.100/32  
*Mar  8 03:13:48.853: BGP(0): Revise route installing 1 of 1 route for 192.168.1.100/32 -> 192.0.2.2  
to main IP table  
*Mar  8 03:13:48.853: RT: network 192.168.1.0 is now variably masked  
*Mar  8 03:13:48.853: RT: add 192.168.1.100/32 via 192.0.2.2, bgp metric [200/0]
```

Step 3. Now you add a route with tag 86.

This should only cause the host to be black holed in region 2, with routers PE-3 and PE-4.

```
ip route 192.168.1.100 255.255.255.255 null0 tag 86
```

Step 4. You can verify that the route does not get added to the forwarding table in PE-2 because there is no valid path to the next hop.

```
PE-2#  
*Mar  8 03:16:31.513: BGP(0): 192.168.255.253 rcvd UPDATE w/ attr: nexthop 192.0.2.3, origin i,  
localpref 200, metric 0, originator 192.168.255.245, clusterlist 192.168.255.253, community no-  
export  
*Mar  8 03:16:31.513: BGP(0): 192.168.255.253 rcvd 192.168.1.100/32  
*Mar  8 03:16:31.513: BGP(0): no valid path for 192.168.1.100/32
```

ADDITIONAL REFERENCES

For further information about the topics discussed in this paper and related issues, refer to the documentation available at the following Websites:

<http://grc.com/dos/drds.htm>

http://www.cert.org/tech_tips/denial_of_service.html

<http://www.networkmagazine.com/article/NMG20000829S0003>



Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel
Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan
Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2005 Cisco Systems, Inc. All rights reserved. Cisco, Cisco Systems, the Cisco Systems logo, Cisco IOS, and IOS are registered trademarks of Cisco Systems, Inc. or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0501R)

204171.d_ETMG_AE_2.05

Printed in the USA