



WHITE PAPER

INFRASTRUCTURE SECURITY ON THE SPRINT END-TO-END CISCO NETWORK

INTRODUCTION

Large organizations are using the Internet and networking technology to generate revenue, serve customers, improve communication, and support a range of business activities. Many companies are also consolidating their numerous, disparate networks to simplify capital expenditures and network management and to control costs. Some of these organizations are out-tasking network functions to service providers, such as Sprint.

Sprint and its business customers share a common concern about network security, as daily threats to networks and connected devices use up valuable company resources. Some malicious distributed denial-of-service (DDoS) attacks today target the network infrastructure as well as customer assets. In fact, the 2004 annual survey by the Computer Security Institute and Federal Bureau of Investigation places DDoS attacks among the most costly security events affecting businesses, second only to viruses. The annual revenue losses reported in this survey averaged \$1.4 million, with a high of \$60 million¹.

Within the context of increasing network threats, the growing reliance on the Internet and IP network core creates new challenges for Sprint—both to protect its own core network infrastructure and to provide mechanisms and value-added services that will help protect the customer's IP traffic.

The purpose of this technical white paper is to highlight, in detail, some of the techniques that Sprint uses to control packet flows and protect its network core against security risks associated with such threats as DDoS and other types of attacks. In addition, it provides an overview of some of the security solutions available to mitigate network threats against high-profile Internet-based businesses. Because the Sprint network is built end-to-end with Cisco Systems® equipment, and thus many of Sprint's services possess the Cisco Powered Network designation, many of the techniques in this paper are based on Cisco technology and innovation.

AUDIENCE

This case study should be considered essential reading for technical leaders at service providers and large enterprise businesses, to obtain a better understanding of the security best practices required in today's networks.

Some of the techniques demonstrated in this case study focus on the service provider-centric Interior Gateway Protocol (IGP) and Intermediate System-to-Intermediate System (IS-IS) Protocol. It must be noted that these techniques are equally applicable to other popular routing protocols such as Open Shortest Path First (OSPF) and Enhanced Interior Gateway Routing Protocol (EIGRP).

¹ CSI 2004 survey available from <http://www.gocsi.com>.

OVERVIEW OF INFRASTRUCTURE SECURITY FEATURES: CISCO IOS SOFTWARE-BASED PLATFORMS

Sprint uses multiple mechanisms to protect its core networks and its IT infrastructure from attacks to maintain a very high level of network availability and performance. This allows Sprint to offer industry-leading service-level agreements (SLAs) to its customers.

Sprint is able to select from a wide range of protection mechanisms available in Cisco routers and switches to secure its infrastructure. Some of the protection mechanisms available include the following:

- Receive access control lists (rACLs)
- Control plane policing (CPP)
- Unicast Reverse Path Forwarding (URPF) checks
- Authentication, authorization, and accounting (AAA) features
- IS-IS advertisement features
- IP options handling features
- Protocol-specific ACLs

A short description of each of these features follows.

RECEIVE ACCESS CONTROL LIST (rACL)

The rACL feature provides a level of protection against the resource exhaustion of the high-importance router CPU.

Network traffic received by a router can be divided into two broad categories:

- Transit traffic—Traffic that is destined “through” the router—that is, traffic that passes through the router via the forwarding path to some other destination.
- Receive-path traffic—Traffic that is destined “to” the router itself—that is, traffic that must be sent via the *receive path* to the router CPU for processing or further analysis.

In normal operations, the vast majority of traffic simply flows through a router en route to other destinations. However, certain types of data must be handled by the router CPU, most notably routing protocols, remote router access, and network management traffic (Secure Shell [SSH] Protocol, Simple Network Management Protocol [SNMP], etc.). In addition to the aforementioned traffic, other Layer 3 packets might require the processing flexibility of the CPU. These include packets with certain IP options set, and certain forms of Internet Control Message Protocol (ICMP) packets, for example.

The rACL feature was added to service provider core routers, such as the Cisco® 12000 Series Router and Cisco route switch processor (RSP) based platforms, to provide the capability to filter traffic destined “to” the router without affecting the majority of the network traffic that simply transits the router.

The Cisco 12000 Series, for example, has several data paths, each servicing different forms of traffic. Transit traffic is forwarded from the ingress line card (LC) to the fabric and then to the egress LC for next-hop delivery. In addition to the transit traffic data path, a Cisco 12000 Series has two other paths for traffic requiring local processing: the LC-to-LC CPU path and the LC-to-LC CPU-to-fabric-to-GRP path.




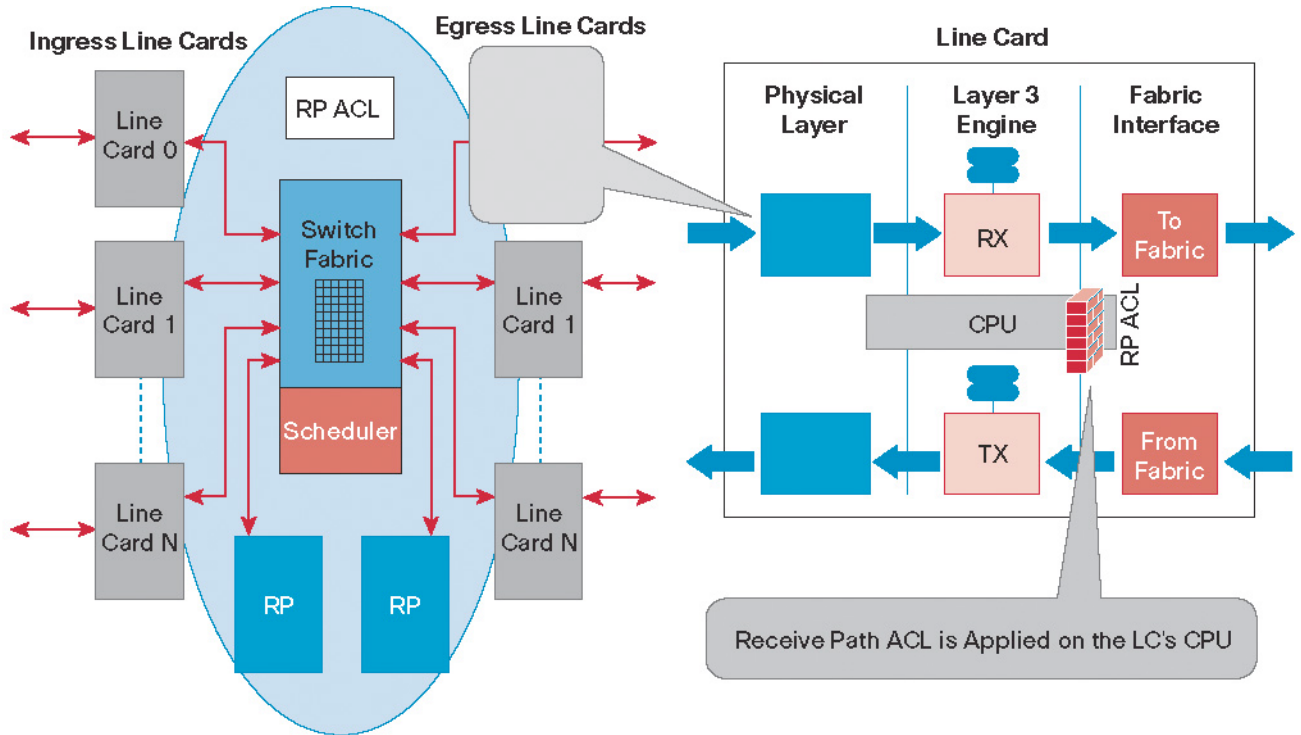
Table 1 shows the paths for several commonly used features and protocols.

Table 1. Common Paths for Various Traffic Types on Cisco 12000 Series Routers

Traffic Type	Data Path
Normal (transit) traffic	LC-to-fabric-to-LC
Routing protocols/SSH/SNMP	LC-to-LC CPU-to-fabric-to-GRP
ICMP echo (ping) logging	LC-to-LC CPU

As shown in Figure 1, rACL execution occurs on each LC before the packet is transmitted to the GRP. Further, rACL execution does not affect normal transit path traffic and hence does not affect the performance of forwarded traffic.

Figure 1
Receive ACL Functionality on the Cisco 12000 Series Platform



When deploying rACLs, the existing control and management plane traffic requirements must be well understood. Initial deployment should follow a very conservative approach using iterative rACL configurations to help identify and eventually filter traffic. The following is a brief list of typical traffic types that should be considered when deploying rACLs:

- Routing protocol and management traffic sourced from legitimate addresses must be permitted.
- Other required traffic, such as Network Time Protocol (NTP), Internet Group Management Protocol (IGMP), ICMP, etc., should be permitted as required. When possible, address blocks should be narrowed to the minimum required range.
- Explicit deny statements should be included as the last entries in the rACL to prevent unauthorized traffic from reaching the GRP.

Benefits of rACL

When properly implemented, rACLs can provide the following security benefits:

- rACLs protect the CPU from having to expend resources on the processing of undesired traffic coming in from any router interface using one simple configuration line.
- A Standard, Extended, or Turbo ACL is created on the GRP, and this ACL is then downloaded to all the line cards.
- The rACL execution occurs on the line cards for all received adjacency packets before they are queued to be sent to the GRP.
- Deploying rACLs has helped defend against various security advisories in all U.S. service provider network infrastructures [NANOG 31 October 2004—U.S. service provider feedback].

For more information about rACLs, see: http://www.cisco.com/en/US/tech/tk648/tk361/technologies_white_paper09186a00801a0a5e.shtml.

CONTROL PLANE POLICING

The control plane policing (CPP) feature is a follow-up to, and significant improvement upon, the rACL feature. Whereas rACLs merely allowed the configuration of basic “permit” and “deny” filters for traffic destined to the router CPU, the CPP feature extends this by allowing users to configure a quality of service (QoS) filter that can also “rate-limit” these receive-path traffic flows. Thus CPP further protects the control plane of Cisco IOS® Software-based routers and switches against many attacks, including reconnaissance and denial-of-service (DoS) attacks. In this manner, the control plane (CP) can maintain packet forwarding and protocol state despite an attack or heavy load on the router or switch.

The route processor is critical to network operation, so any service disruption of the control plane traffic can lead to network outages that affect business operations. A DoS attack targeting the route processor, either of an inadvertent or malicious nature, typically involves high rates of traffic that result in excessive CPU utilization on the route processor itself. Such an attack can be devastating to network stability and availability and may include the following symptoms:

- High route processor CPU utilization (near 100 percent)
- Loss of line protocol keepalives and routing protocol updates, leading to route flaps and major network transitions
- Interactive sessions via the command-line interface (CLI) are slow or completely unresponsive because of high CPU utilization
- Route processor resources such as memory and buffers are unavailable for legitimate IP data packets
- Packet queue backup, leading to indiscriminate drops (or drops due to lack of buffer resources) of other incoming packets

Certain types of traffic often must be permitted to reach the route processor, but unlimited access to the router processor is not desirable. ICMP echo requests are one recognizable example. With rACLs, the granularity of control was limited to simple permit and deny statements. CPP addresses this need to protect the control plane by bringing the modular QoS CLI (MQC) capability to the problem. In this way, CPP provides filtering and rate-limiting capabilities for control plane packets.

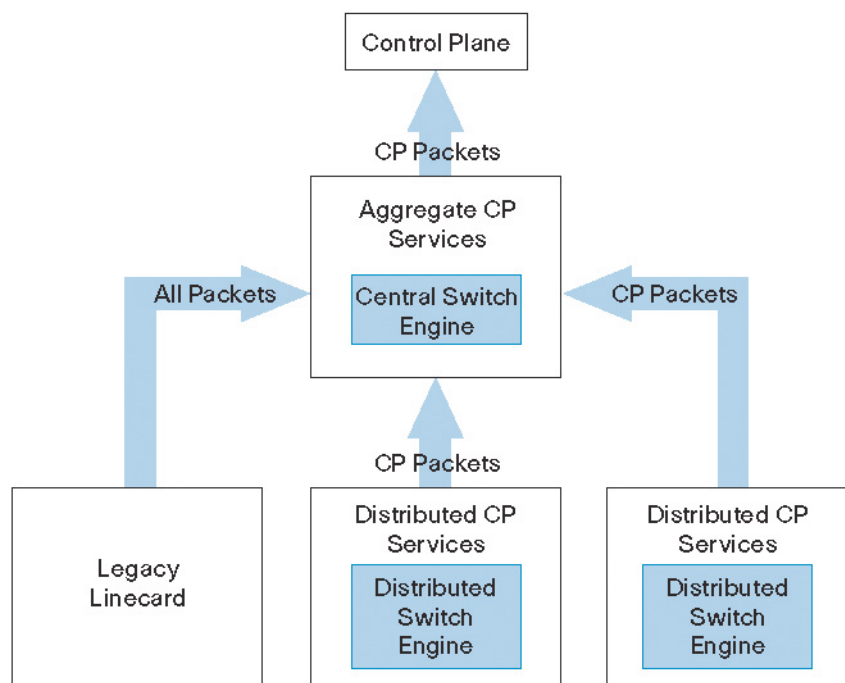
CPP treats the control plane as a separate entity with its own ingress (input) and egress (output) ports, much like ports on a router or switch. Policies are developed that include a set of rules defining which packets may reach the route processor and at what rate. These policies are then associated with the ingress and egress port of the control plane.

Input CPP services are executed after both router input port services and a routing decision on the input path have been made. As shown in Figure 2, CPP can be applied:

- On an aggregate (platform) level by the central switch engine and applied to all control plane packets received from all line cards on the router
- On a distributed level by the distributed switch engine of a line card and applied to all control plane packets received from all interfaces on that one line card

Figure 2

Input Control Plane Services: Aggregate and Distributed Services



Infrastructure attacks are becoming increasingly common, highlighting the need for infrastructure protection. Control plane policing provides a hardware-independent mechanism for defining and implementing router protection schemes of varying sophistication. In its most basic form, CPP can be used to permit or deny traffic destined to the router's processor in a manner similar to that provided by rACLs. As operational experience grows, the complexity of the CPP policy may be increased to provide more granular and tighter control by configuring rate-limits for certain traffic types that must reach the route processor. The rate-limiting features provide extremely flexible policy implementation.

Benefits of CPP

CPP deployment provides several important security benefits:

- Protection against DoS attacks targeted toward the network infrastructure by traffic flows and protocols that must be permitted, but where rate limiting offers substantial protection
- Ease of deployment—CPP leverages the existing MQC infrastructure, which allows customers to preserve the existing interface configurations and add global commands to address security goals
- Consistent implementation strategy across all Cisco hardware
- Increased reliability, security, and availability of the network

For more information about CPP, see: www.cisco.com/univercd/cc/td/doc/product/software/ios122s/122snwft/release/122s18/gtrtlmt.htm.

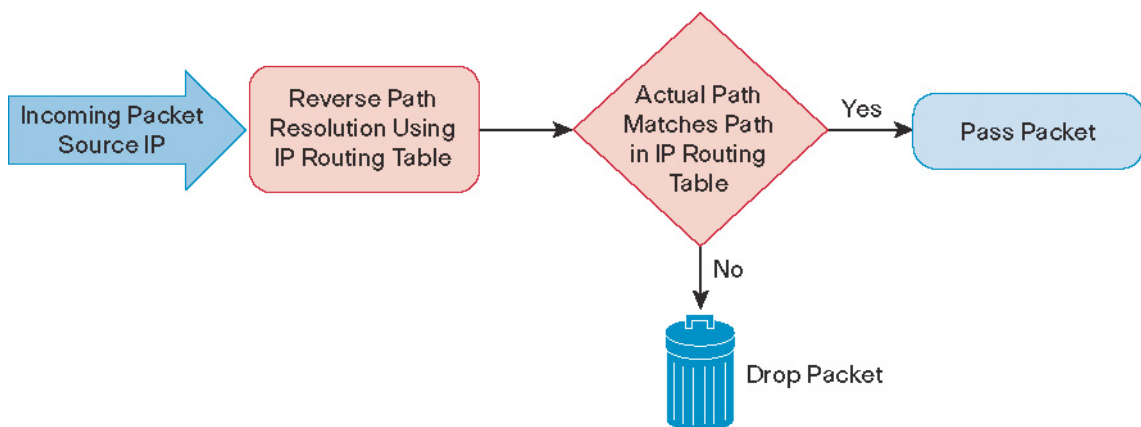
UNICAST REVERSE PATH FORWARDING (URPF)

URPF is a feature originally created to implement the requirements of BCP 38/RFC 2827. (*Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing*, by P. Ferguson, D. Senie.) As such, URPF was originally designed for the customer–Internet service provider (ISP) network edge. The objective was to develop a feature that could be easily automated in the customer provisioning system, scale as new address blocks were allocated to the customer, and work with MTRIE-based Cisco Express Forwarding switching path.

Several common types of DoS attacks take advantage of forged or rapidly changing source IP addresses, allowing attackers to thwart efforts by ISPs to filter these attacks using ACLs. By taking advantage of the information stored in the forwarding information base (FIB) that is created by the Cisco Express Forwarding switching process, URPF can determine whether IP packets are spoofed or malformed by matching the IP source address and ingress interface against the FIB entry that reaches “back” to this source (a “reverse lookup”). If URPF does not find a reverse path for the packet, it drops the packet. Once enabled on an interface, URPF checks all packets arriving on that interface. This process is illustrated in Figure 3.

Figure 3

URPF Drops Packets with Spoofed IP Source Addresses



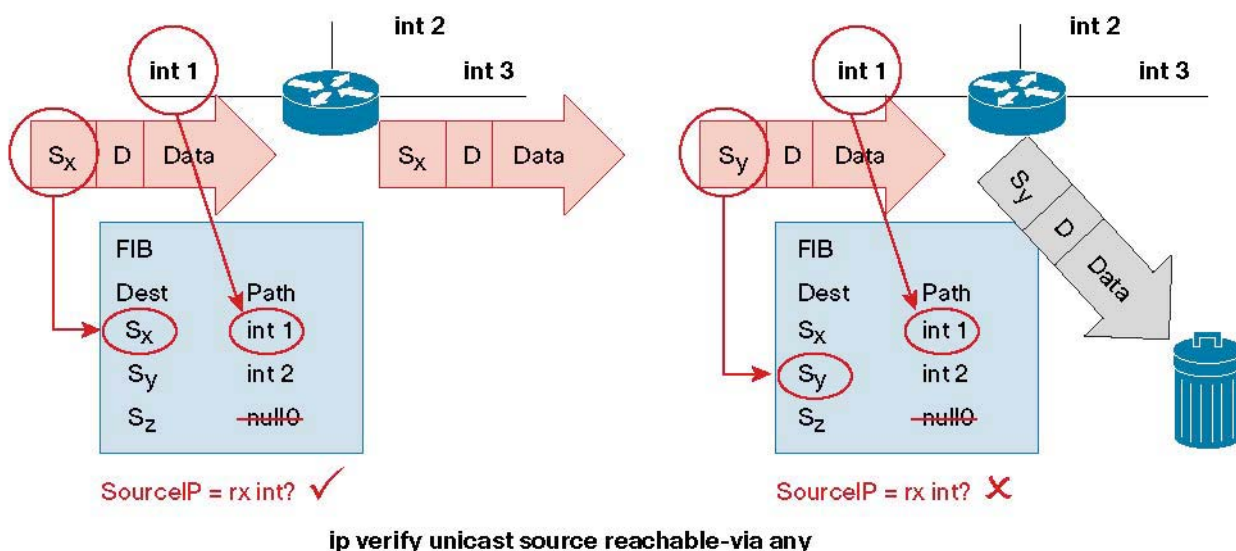
The particular path validation criteria used to determine path consistency depends on the particular URPF mode enabled on an interface. Two modes of URPF are currently available:

- URPF strict mode
- URPF loose mode

URPF Strict Mode

URPF strict mode requires that the source IP address of an incoming packet has a FIB return path via the exact same interface as that on which the packet arrived. If the FIB return path does not exist, or refers to a different interface than the one on which the packet arrived, the packet is dropped. This functionality is illustrated in Figure 4.

Figure 4
URPF Strict Mode Drops Packets with Spoofed IP Source Addresses That Fail the Interface Test



URPF strict mode can be used only in deployments where the FIB entries match the traffic paths.

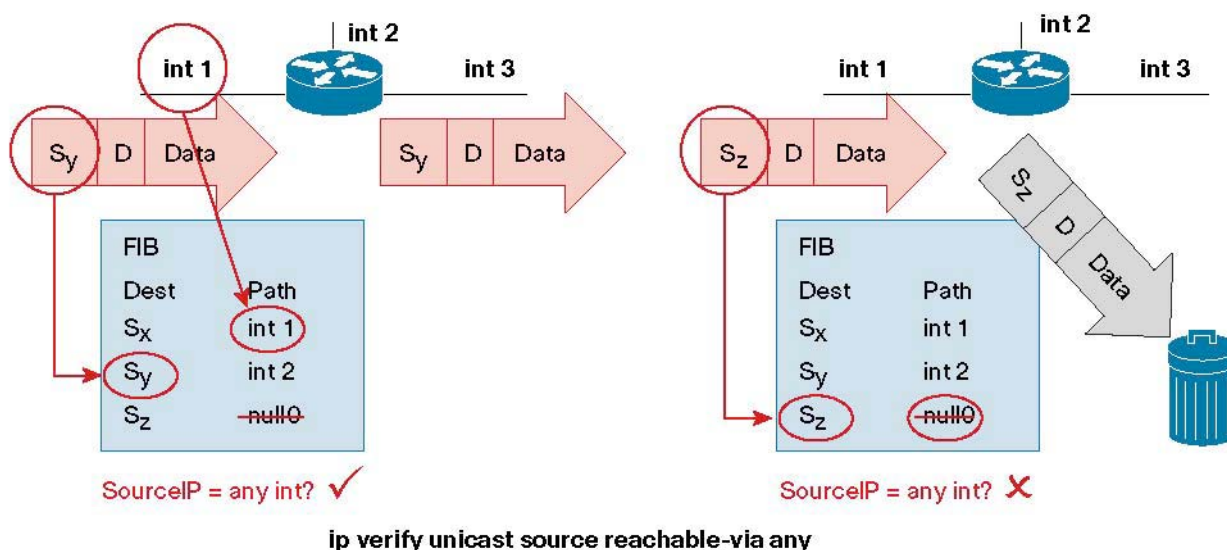
Note: If the Cisco Express Forwarding FIB paths do not or cannot be forced to match the traffic paths, URPF strict mode must not be enabled, and URPF loose mode should be considered as an alternative.

URPF Loose Mode

URPF loose mode requires that a valid FIB path entry via any interface exists, which excludes Null0, for the source IP address of an incoming packet. That is, in URPF loose mode, strict interface adherence is no longer enforced, but a valid FIB entry simply must exist. If a valid FIB entry does not exist, or if the FIB entry refers to the Null0 interface, the packet is dropped. This functionality is illustrated in Figure 5.

Figure 5

URPF Loose Mode Drops Packets with IP Source Addresses That Do Not Have FIB Return Path Entries or That Refer to Null0



Loose mode URPF, by virtue of its design, does not offer the same degree of protection from source IP address spoofing as URPF strict mode. However, for scenarios where URPF strict mode is not feasible, such as on a router multihomed to multiple autonomous systems where asymmetric traffic flows are common, URPF loose mode does enable filtering of certain undesirable traffic. For example, source IP addresses that do not exist in the routing table, such as RFC 1918 and unallocated addresses, as well as those not advertised by a Border Gateway Protocol (BGP) peer, will be dropped.

In addition, a return path may be purposely made invalid by associating it with the Null0 interface to cause the URPF path validation check to fail and those packets with that source IP address to be dropped. This is the technique employed by the “source-based” remote triggered black hole (RTBH) filtering techniques, which are used to drop attack traffic from a source IP address identified as the source of an attack.

Benefits of URPF

URPF offers many security benefits, including:

- Operational efficiency—When URPF strict mode can be applied, for example, on customer interfaces, the ability to implement BCP 38/ RFC 2827 spoofed IP source address filtering is trivial. The IP routing table and FIB provide the entire context from which decisions are made. No operational maintenance is necessary, such as would be required using ACLs. Network addresses changes are automatically taken into account.
- Minimal performance impacts—In most applications, the implementation of URPF introduces minimal performance impacts to the data forwarding process.
- Enablement of other security features—URPF is an enabling feature for source-based RTBH filtering, one of the essential network security features deployed by all service providers as a means of mitigating certain attacks.

For more information about URPF, see:

Unicast Reverse Path Forwarding Enhancements—Cisco IOS Software Release 12.1T Documentation

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121t/121t2/rpf_plus.htm

11.1CC Unicast RPF Documentation

http://www.cisco.com/univercd/cc/td/doc/product/software/ios111/cc111/uni_rpf.htm

AUTHENTICATION, AUTHORIZATION, AND ACCOUNTING (AAA)

Access control is the way in which companies control who is allowed access to the network server and what services they are allowed to use once they have access. Authentication, authorization, and accounting (AAA) network security services provide the primary framework through which access control is set up on the router or access server.

AAA is an architectural framework for configuring a set of three independent network security functions in a consistent manner. AAA provides a modular way of performing the following services:

- **Authentication**—Provides the method of identifying users, including login and password dialog, challenge and response, messaging support, and, depending on the security protocol you select, encryption. Authentication is the way a user is identified before being allowed access to the network and network services.
- **Authorization**—Provides the method for remote access control, including one-time authorization or authorization for each service, per-user account list and profile, user group support, and support of IP, Internetwork Packet Exchange (IPX), AppleTalk Remote Access (ARA) protocol, and Telnet. AAA authorization works by assembling a set of attributes that describe what the user is authorized to perform. These attributes are compared to the information contained in a database for a given user and the result is returned to the AAA server to determine the user's actual capabilities and restrictions.
- **Accounting**—Provides the method for collecting and sending security server information used for billing, auditing, and reporting, such as user identities, start and stop times, executed commands (such as Point-to-Point Protocol [PPP]), number of packets, and number of bytes. Accounting enables you to track the services users are accessing as well as the amount of network resources they are consuming.

In most circumstances, AAA uses protocols such as RADIUS, TACACS+, or Kerberos to administer its security functions.

Benefits of Using AAA

AAA provides the following benefits:

- **Increased flexibility and control of access granted to network devices**—Granular controls are available to provide access to only those devices, and within those devices, only those features required by that user.
- **Accountability**—Each action taken by a user can be logged and time-stamped, providing a useful record for troubleshooting or, if required, documenting chains of events.
- **Scalability**—AAA configurations simplify the overall access control process, resulting in a more robust, scalable solution.
- **Standardized authentication methods**—Industry-standard AAA protocols, such as RADIUS, TACACS+, and Kerberos, permit interoperability and user consistency in configuration.

For more information about AAA, see: http://www.cisco.com/univercd/cc/td/doc/product/access/acs_soft/csacs4nt/acs32/user02/z.pdf.

IS-IS ADVERTISE PASSIVE ONLY

As part of engineering efforts to improve routing protocol convergence times, additional configuration parameters were added to the Intermediate System-to-Intermediate System (IS-IS) protocol to allow for the removal of IP addresses of specified links from the link state database. A smaller link state database results in more efficient Dykstra algorithm processing, and hence, faster convergence times.

Previously, to accomplish a similar result, interfaces needed to be configured as unnumbered. However, for network management reasons, *unnumbered* interfaces are more difficult to administer.

IS-IS now provides a mechanism to exclude connected IP prefixes from LSP advertisements. This feature can be implemented in one of two ways—either as an interface-level command or as a global IS-IS routing process-level command. These two methods are described in the following sections.

Small-Scale Method to Reduce IS-IS Convergence Time

An IS-IS interface can be explicitly configured not to advertise its IP network to neighbors by using the **no isis advertise-prefix** interface command. This method is feasible for a small network, but it does not scale well. When dozens or hundreds of routers exist in the network, with possibly tens of thousands of physical interfaces involved, it becomes difficult to add this command to each interface on every router.

Large-Scale Method to Reduce IS-IS Convergence Time

An easier way to reduce IS-IS convergence is to configure the IS-IS instance on a router to advertise only passive interfaces. This is done by using the **advertise-passive-only** command within the IS-IS router process. This command relies on the fact that when enabling IS-IS on a loopback interface, it is usually configured as passive (to prevent sending unnecessary hello packets out through it because there is no chance of finding a neighbor behind it). Thus, to advertise only the loopback IP address in ISIS, and if it has already been configured as passive, configuring the advertise-passive-only command per IS-IS instance prevents the overpopulation of the routing tables.

Benefit of Excluding IP Prefixes in IS-IS Network

In addition to the fast convergence benefits provided by implementing either of the IS-IS prefix exclusion methods described above, additional security benefits are achieved. For example, the point-to-point links between core routers became virtually unreachable for packets entering the network at the edge. That is, packets arriving from a customer-edge router can reach only interfaces directly connected to its provider-edge router and the next directly connected provider router, but no deeper into the core network.

Note: Implementing this technique may result in network troubleshooting operational changes depending on the methods and procedures being used. Individual link endpoints will no longer be accessible from more than one hop away. Traceroute will still work, but pinging interfaces and using Telnet to connect interfaces will not work.

For more information about IS-IS **advertise-passive-only**, see:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a00800ad395.html.

IP OPTIONS IGNORE

IPv4 options are part of the Internet. Some IP options are used in day-to-day operations, while others are defined for a new protocol or service and never get widely adopted. The core character of IP options is that they add length to the IP header of the packet, making the IP header of the packet variable in length. In traditional CPU-oriented software forwarding, these variable-length headers are manageable, but in today's ASICs-oriented forwarding paradigm, variable-length headers are a challenge. Most (if not all) forwarding ASICs punt packets with IP options to either the forwarding ASICs supporting CPU (i.e., the LC CPU) or to the route processor. This punt opens a set of security vulnerabilities. Specific attack flows can be designed with IP options—forcing the ASICs to punt—resulting in an overload of the ASICs supporting CPU or the router's route processor.

To close this attack vector, a generic (global configuration) **[no] ip options [drop | ignore]** capability is available. This single global command applies to all interfaces on the router—providing a quick and easy way to turn off IP options processing. When the “drop” action is specified, all received IP packets that have IP options configured will be dropped. When the “ignore” action is specified, all received IP packets that have IP options configured will be forwarded, but without the normal punt and RP handling of the packet—that is, the ASIC will still forward these IP packets based on the first 20 bytes of the IP header.

Note that some protocols, such as Resource Reservation Protocol (RSVP), which is used by Multiprotocol Label Switching-Traffic Engineering (MPLS-TE), for example, Internet Group Management Protocol Version 2 (IGMPV2), and others, use IP options packets. These protocols may not function in drop or ignore mode if this feature is configured.

Benefits of Using IP Options [drop | ignore]

The **ip options [drop | ignore]** feature provides the following benefits:

- Ability to drop spurious packets containing IP options, which may be sent by attackers to cause network disruptions
- Ability to identify the source IP addresses sending packets with IP options set

For more information about **ip options [drop | ignore]**, see:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a00801d4a94.html

http://www.cisco.com/en/US/products/sw/iosswrel/ps5207/products_feature_guide09186a00801d4a7d.html

CASE STUDY: SPRINT'S INFRASTRUCTURE SECURITY AND DESIGN CONSIDERATIONS FOR THE INTERNET CORE

Before considering how to architect a highly secure IP backbone, one must consider the possible threats to the infrastructure. After the threats are understood and prioritized, security mechanisms can be built in as part of the network design, not as an afterthought or add-on. Sprint, along with many other service providers, has traditionally deployed best practice features such as AAA, VTY access-classes, SSH, and MD5 checksums on routing protocol updates, SNMP community access lists and good community strings, and SNMP views. While these features do provide protection against some attacks, they are not sufficient in today's world. First, they can fail to protect against some vulnerabilities because they are enforced on the router CPU. Second, under some attacks, if not protected against, the router CPU can become overwhelmed. Hence to that point, there is a need for a mechanism to protect the router CPU.

Sprint recognized that architectural changes and new security mechanisms were required to stop unauthorized packets before they reached critical components. There are essentially three ways to architect a highly secure infrastructure:

- Virtual Routing and Forwarding (VRF) MPLS can be implemented, and all traffic—including the “untrusted” Internet traffic—can be contained within separate VRFs. In this case, the entire network core becomes opaque to the outside world. The security afforded by this approach is relatively high since traffic is controlled throughout. The upfront network engineering effort is relatively high, but operational expense is relatively low for supporting long-term operations. The provider-edge devices are not secured by this method and warrant another solution.
- ACLs can be deployed at the entire network perimeter to control traffic flow for traffic destined to everything inside the network. In effect, traffic destined “to” the infrastructure routers is denied. This configuration is difficult to deploy, and even more difficult to maintain due to the significant number of ACLs that must be maintained on a large number of interfaces. It also requires that every edge-facing interface have the capability to support ACLs at line rate.
- Packet-routing strategies coupled with control plane filters can be deployed to control traffic flows destined to everything inside the network. Specifically, control plane packet filters are used to thwart rogue packets and IGP routing enhancements are used to make it more difficult to reach control plane IP addresses. This technique can also help mitigate the effects of massive denial-of-service attacks. The security afforded by this approach is relatively good, since traffic flow is controlled reasonably well. The upfront network engineering effort is relatively small for this approach, and the operating expense is also small for long-term support.

For Sprint, the choice was easy. These first two options required capital expenditures to implement, and the capital wasn’t available. Sprint decided to begin working toward the third option and identified two external types of attacks that could potentially target its IP router infrastructure. The first attack involves exploiting software vulnerabilities or design flaws with a very small amount of traffic. For example, the SNMP PROTOS vulnerability involved a single malformed SNMP packet that could crash a router, or a specially crafted packet with IP options set, which causes the router CPU to spend extensive processing cycles. The second type of attack is the more traditional volume-based DoS attack. This attack simply overwhelms the router CPU or circuits with a high volume of traffic.

To protect against the low-volume packet attacks, Sprint needed a way to drop packets before they reached the CPU. The answer was rACLs and its successor, control plane policing. rACLs were deployed to control which packets were permitted to reach the CPU based on protocol, source address, destination address, and UDP/TCP port. rACLs were very successful at closing off unused ports and protocols and restricting access to authorized sources.

Attackers can still circumvent rACLs, but they would be required to spoof a packet from an authorized source IP address to the router. In order to further reduce the risk of spoofing, Sprint has deployed URPF) on many of its customer interfaces to drop spoofed packets. There are some hardware limitations with this technology that prohibit its use on all line cards. Therefore, URPF is not universally deployed, but even a limited deployment makes it more difficult to spoof IP packets.

IP options presented a possible vulnerability that the rACL could not block as well, since these packets are not necessarily not destined to the router itself, but rather are punted to the RP due to the necessity to handle the IP options field in CPU. To counter this threat, Sprint is currently deploying the **ip options ignore** mechanism, which will ignore the IP options flag and pass it through the forwarding path without processing the IP option. This is similar to the **no ip source-route** that many ISPs use to prevent the router from obeying IP source routing options.

Sprint next focused on the volume-based DoS attacks and the spoofed packet attacks that can still sneak through rACLs and URPF. rACLs can protect the routers from DoS attacks, but only to a certain point. Sprint quickly realized that it must find a way to drop the unauthorized packets before they reached its routers.

Sprint's first concern was the TCP session for eBGP between a Sprint customer and Sprint. The customer IP addresses can be obtained with a simple traceroute operation. To help prevent these eBGP sessions from being attacked, Sprint has removed the specific route to the /30 prefix on the point-to-point networks between Sprint and its customers. Routers that are not directly connected use the aggregate route to Null0 instead. This makes it impossible to attack an eBGP session between Sprint and a Sprint customer unless the attacker is connecting to Sprint through the router being attacked. While this is not 100 percent effective, it greatly reduces the risk. Sprint is also investigating deploying the BGP TTL Security Hack (BTSH) mechanism (IETF draft draft-gill-gtsh-04.txt) to further secure eBGP sessions from the corner case where the attacker enters the Sprint network on the router that is directly connected to the /30.

Building on this concept, Sprint next focused on securing its router-to-router links in the same way that the Sprint-to-customer links were secured. Sprint is currently deploying the IS-IS **advertise-passive-only** global and **no isis advertise-prefix** interface configurations, which will achieve the same effect on router-to-router links in the core of the network. While these IS-IS features were originally designed to decrease IS-IS convergence times, they also play a valuable role in securing the network. Once complete, the /30 and /31 networks connecting Sprint's routers will no longer be reachable from most of the Internet.

RECEIVE ACCESS CONTROL LISTS

Sprint deployed rACLs in a slow, phased approach. The timeline is as follows:

1. Approximately four weeks of lab testing to examine performance impacts across hardware under various different levels of attack.
2. Approximately two days of designing the rACL. This first pass at the receive access list should not attempt to be all-encompassing but fairly general to provide the easiest first migration. Start by covering just the protocols and ports used, rather than source and destination IP addresses. The rACL had a **permit ip any any log** statement at the end to log all packets not explicitly permitted.
3. Approximately one month of deploying the rACL on various routers at different points in the network and examining the syslogs to find any missed protocols. This version of the rACL was applied for only several hours at a time so an engineer could constantly monitor it in case the router generated a huge number of syslog messages. The rACL was updated to reflect any missed protocols and ports. Additionally, the rACL was updated to temporarily explicitly deny known bad traffic to reduce the syslog volume. This process was constantly repeated until the rACL could be left on the routers for hours without generating any syslog messages. Some common missed protocols include Hot Standby Routing Protocol (HSRP), Cisco Protection Gateway Protocol (PGP), Multicast Source Discovery Protocol (MSDP), traceroute, and "Telnet" from the router to TCP/80 and TCP/25 so the NOC can troubleshoot by manually talking to HTTP or SMTP servers.
4. At this point, the **permit ip any any log** statement was replaced with a deny statement, and the explicit denies that were installed to drop known bogus traffic to minimize the syslog messages were removed. It then took about three months to deploy the finished rACL point of presence (POP) by POP across the network.
5. The next two years were spent slowly refining the rACL. This included limiting ports and protocols to source and destination IPs.

Below is a very early version of Sprint's rACL.

```
!  
!-----Receive ACL-----  
!  
access-list 111 deny ip any any fragments  
!  
!-----SSH/TELNET-----  
!  
!Firewall  
access-list 111 permit tcp host <firewall ip> any eq 22  
!Router loopbacks
```

```

access-list 111 permit tcp <router loopbacks> any eq 22
!
!-----Outbound Telnet/SSH-----
!
access-list 111 permit tcp any range 22 telnet any established
!
!-----SNMP-----
!
access-list 111 permit udp host <SNMP Server> any eq snmp
!
!-----DNS-----
access-list 111 permit udp host <DNS server> eq domain any
!
!-----TACACS-----
!
access-list 111 permit tcp host <TACACS+ server> any established
!
!-----NTP-----
!
access-list 111 permit udp host <NTP server> any eq ntp
access-list 111 permit udp <router loopbacks> any eq ntp
!
!-----APS-----
!
access-list 111 permit udp <router interfaces> any eq 1972
!
!-----BGP-----
!
access-list 111 permit tcp any gt 1024 any eq bgp
access-list 111 permit tcp any eq bgp any gt 1024 established
!
!-----MSDP-----
!
access-list 111 permit tcp any any eq 639
access-list 111 permit tcp any eq 639 any gt 1024 established
!
!-----FTP/TFTP-----
!
access-list 111 permit tcp host <Config server> eq ftp any
access-list 111 permit tcp host <Config server> any established
access-list 111 permit udp host <Config server> eq tftp any
!
!-----IGMP-----
!
access-list 111 permit igmp any any
!
!-----PIM-----
!
access-list 111 permit pim any any
!

```

```

!-----ICMP-----
!
access-list 111 permit icmp any any echo-reply
access-list 111 permit icmp any any ttl-exceeded
access-list 111 permit icmp any any unreachable
access-list 111 permit icmp any any echo
!
!-----TRACEROUTE-----
!
access-list 111 permit udp any gt 10000 any gt 10000
!
!-----HSRP-----
!
access-list 111 permit udp any host 224.0.0.2 eq 1985
!
!-----L2TPv3-----
!
access-list 111 permit 115 <router loopbacks> any
!
!-----GRE-----
!
access-list 111 permit gre <router loopbacks> any
!
!-----Classify Drops-----
!
access-list 111 deny tcp any any
access-list 111 deny udp any any
access-list 111 deny icmp any any
access-list 111 deny ip any any
!

```

As Sprint slowly began to refine their rACL, some valuable lessons were learned.

- An organized IP addressing strategy is essential. Sprint had a dozen loopback ranges deployed on its routers. This would require one line of rACL per SNMP server host to limit SNMP. If three SNMP servers existed, this would result in 36 ACL lines. To limit SSH between router loopbacks, it then requires N^2 ACL lines. In Sprint's case, it would have taken 144 ACL lines to control access with SSH alone. This would have to be duplicated for all the other protocols limited to router loopbacks. The rACL could easily grow to over 1000 lines. To remedy the problem, Sprint consolidated its loopbacks into three blocks (ARIN, RIPE, and APNIC). This required only 9 ACL lines, instead of 144. The result is an approximately 100-line ACL rather than one in the thousands.
- Change control posed a problem because incorrectly altering the rACL could cause a major outage. For example, if the network operations center (NOC) engineer accidentally misspells a command in the BGP line of the rACL, a new rACL could be created with BGP missing, which would cause all BGP sessions to be blocked if applied. Sprint developed several "Expect" scripts to help ensure that during the rACL update, the router was always in a safe state. Sprint also has a router in the production network that is used as a preproduction deployment test site for new rACLs. This is used in the deployment process immediately prior to deploying the rACL on production routers.

It remains imperative not to disrupt the production network even when a thoroughly tested feature is being deployed. Hence, if any disruption is experienced during the deployment of rACL, a rollback strategy was practiced. The rollback strategy was not to investigate the failure of rACL deployment but just to reverse the changes on an individual router where the failure was experienced. An investigation was carried out

later and the router experiencing rACL deployment, was manually changed to overcome any failure. It is important to note that all the improvements learned from these failures were incorporated for the next round of the automated rACL deployment.

- A decision must be made regarding security versus maintenance of the rACL. The main factor influencing this decision is eBGP peers. In order to have an adaptive rACL incorporating changes such as new eBGP peer addition, rACL should be able to detect and act on the router configuration change and add the two lines per eBGP peer.

Two main issues arise when the rACL is altered to explicitly permit each peer:

- The rACL will be different across various routers, which can make consistency checking slightly more difficult.
- For routers with hundreds of eBGP sessions, there could be a lot of changes required to the rACL. The risk of constantly changing the rACL must be weighed against a more general and weaker rACL.

Sprint's NOC was initially not prepared to deal with the questions from customers who sent unexpected packets to Sprint routers. Here are some examples of the complaints:

- "Sprint's routers are broken because I cannot ping your router with a 1501/4471 byte (fragmented) packet."

Because Sprint blocked fragments, they had to explain to their customers that their routers still work even though customers could not ping the routers with fragmented packets and generate responses. Many people have a hard time understanding the difference between sending a packet to a router and sending it through a router. The Sprint NOC has had to educate customers about these differences, which puts an extra burden on the NOC.

- "I cannot NTP peer with Sprint's routers."

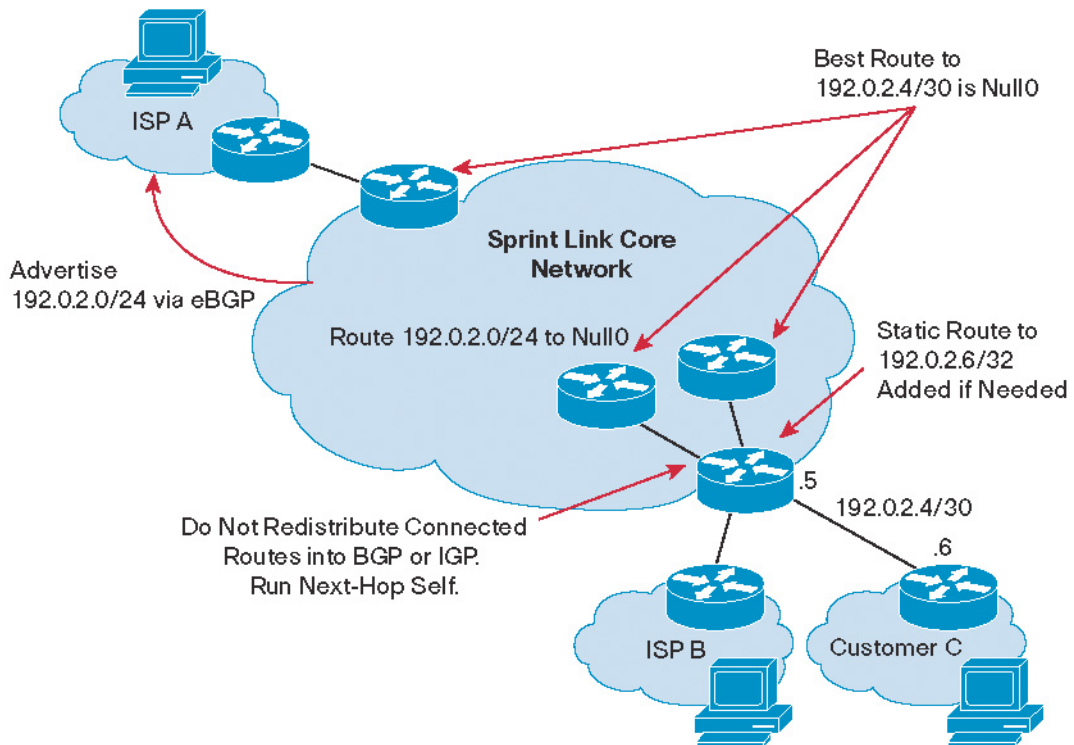
Sprint never supported NTP to its customers, but because there was no way to prevent it before, many customers used the routers to provide an NTP clock source. Even though a service is not supported, it is very difficult to explain to a customer that they need to reconfigure their device. Beware of unsupported services that customers use and rely on.

HIDING THE INTERNET CORE

Customer/Provider Point-to-Point Hiding

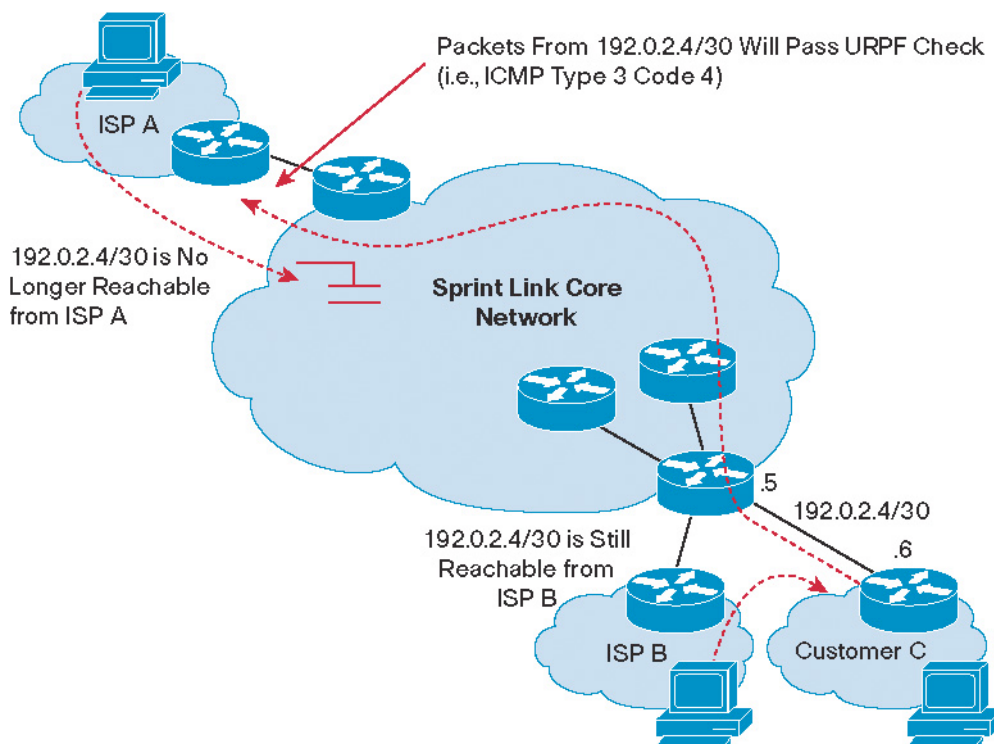
Sprint's first target of trying to hide the core was the /30 point-to-point link between the Sprint aggregation router and the customer premises equipment (CPE). The goal of most attacks is to deny service, and this is usually accomplished by directly attacking the service. However, when the service is adequately protected, it is often easier to attack the router supporting the service. Most of these attacks target the IP addresses of the routers obtained from a traceroute to the service. Sprint realized that most often these IP addresses are used only to communicate between the two directly connected routers (for example, eBGP sessions). With this assumption, the point-to-point networks do not have to be globally reachable, and therefore, do not have to be announced to the rest of the network.

Figure 6
Hiding the Core to Increase Infrastructure Security



Sprint's aggregate IP allocations (192.0.2.0/24 in this example) are null-routed on every router in the network. The aggregation routers run BGP next-hop-self to ensure that all eBGP learned routes have the next hop of the aggregation router's loopback instead of the eBGP neighbor IP address. This is required to ensure that the BGP next hop is always reachable. The aggregation router then does not need to redistribute its connected routes into BGP or IS-IS. At this point, ISP A can no longer reach the point-to-point network for customer C, as the only route its aggregation router has to 192.0.2.4/30 is the static null route to 199.0.2.0/24. If reachability is required to 192.0.2.6/32, Sprint installs a static route to 192.0.2.6/32 on its aggregation router to the customer interface. This static route is then redistributed into BGP. This is required when the CPE router runs Network Address Translation (NAT), IP Security (IPSec) tunnels, VoIP, etc. to the 192.0.2.6 point-to-point address.

Figure 7
Deploying URPF in Loose Mode with an Aggregate Announcement



Another important concern was breaking path maximum transmission unit (MTU) discovery (PMTUD). Because Sprint still announces the aggregate block (192.0.2.0/24 in this example), packets from 192.0.2.6 will not fail a URPF loose-mode check at ISP A since Sprint advertises 192.0.2.0/24 to ISP A.

In the example in Figure 7, although ISP A cannot reach customer C's 192.0.2.4/30 address, ISP B can because it is directly connected to the same Sprint aggregation router. This method is not 100 percent effective, but it greatly reduces the risk because it requires the attacker to be located on a very specific part of the Internet.

Some important observations about this conversion include:

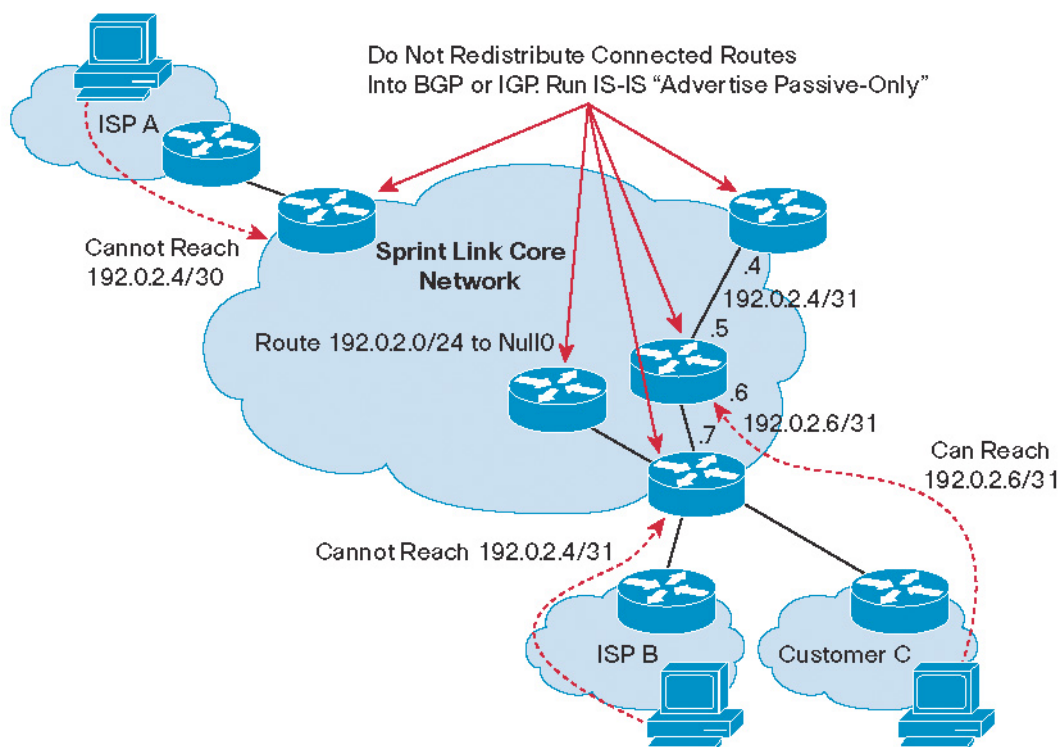
- Twenty-five percent of Sprint customers required the use of their point-to-point address. These customers operate applications or features such as IPSec, generic routing encapsulation (GRE), or voice over IP (VoIP) and utilize the point-to-point address.
- This change took more than a year to implement. This was due to the amount of time to properly inform customers of the change and its potential effects and to allow them to opt out by requesting a /32 static route. This change was also rolled out very slowly to about a dozen routers a week. This action was taken to minimize calls from customers who did not realize they would be affected until their service was down.

Provider/Provider Point-to-Point Hiding

Sprint also realized that to provide Frame Relay, ATM, and VPN services over its IP backbone, the backbone must be made more secure. As it happened, while working on an unrelated IS-IS fast convergence project, Sprint recognized that an intrinsic security benefit could also be gained as a by-product. One of the mechanisms for IS-IS fast convergence is the **advertise-passive-only** command. Sprint is still in the process of deploying the **advertise-passive-only** feature.

Figure 8

Enabling IGP Passive Only Advertisement to Further Secure the Infrastructure



With IS-IS **advertise-passive-only** enabled, the point-to-point links between Sprint's routers became virtually unreachable. In Figure 8, ISP A cannot reach the 192.0.2.4/31 or 192.0.2.6/31 point-to-point links. ISP B and Customer C cannot reach the 192.0.2.4/31 link, but can reach the 192.0.2.6/31 link because they connect to the Sprint aggregation router that the link is directly connected to.

Implementing this technique may result in network troubleshooting operational changes depending on the methods and procedures being used. Individual link endpoints will no longer be accessible from more than one hop away. Traceroute will still work, but pinging and using Telnet to interfaces will not work.

When a customer uses a traceroute operation through the Sprint network, they will see all the IP addresses of the Sprint point-to-point links in the path, but when the customer tries to ping or traceroute to those IPs themselves, they will not be able to reach them. Furthermore, many network management systems rely on ping to determine if a router is operational or not. Many of these systems will report Sprint routers down and this is expected to cause concern as Sprint completes the deployment of the **advertise-passive-only** command.

**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

All contents are Copyright © 1992–2005 Cisco Systems, Inc. All rights reserved. Cisco, Cisco IOS, Cisco Systems, and the Cisco Systems logo are registered trademarks or trademarks of Cisco Systems, Inc. and/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. (0502R)
HE/LW8573 05/05