Opsec Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: October 8, 2007

J. Gill
Verizon Business
D. Lewis
P. Quinn
Cisco Systems Inc.
P. Schoenmaker
NTT America
April 6, 2007

Service Provider Infrastructure Security
draft-ietf-opsec-infrastructure-security-01

Status of this Memo

Copyright Notice

Abstract

This RFC describes best current practices for implementing Service Provider network infrastructure protection for network elements. This RFC complements and extends RFC 2267 and RFC 3704. RFC 2267 provides guidelines for filtering traffic on the ingress to service

provider networks. RFC 3704 expands the recommendations described in RFC 2267 to address operational filtering guidelines for single and multi-homed environments. The focus of those RFCs is on filtering packets on ingress to a network, regardless of destination, if those packets have a spoofed source address, or if the source address fall within "reserved" address space. Deployment of RFCs 2267 and 3704 has limited the effects of denial of service attacks by dropping ingress packets with spoofed source addresses, which in turn offers other benefits by ensuring that packets coming into a network originate from validly allocated and consistent sources. This document focuses solely on traffic destined to elements of the the network infrastructure itself. This document presents techniques that, together with network edge ingress filtering and RFC 2267 and RFC 3704, provides a defense in depth approach for infrastructure protection. This document does not present recommendations for protocol validation (i.e. "sanity checking") nor does it address guidelines for general security configuration.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Table of Contents

1. Introduction

This RFC describes best current practices for implementing Service Provider network infrastructure protection. This document both refines and extends the filtering best practices outlined in RFC 2267 [RFC2267] and RFC 3704 [RFC3704] and focuses only on traffic destined to the network infrastructure itself to protect the service provider network from denial of service and other attacks. This document presents techniques that, together with network edge ingress filtering and RFC 2267 [RFC2267] and RFC 3704 [RFC3704], provides a defense in depth approach for infrastructure protection.

Attacks targeting the network infrastructure can take many forms, including bandwidth saturation to crafted packets destined to a router. These attacks might use spoofed or non spoofed source addresses. Regardless of the nature of the attack, the network infrastructure must be protected from both accidental floods and intentional attacks. Additionally, this protection will assist in preventing the network elements from being used as reflectors in attacks against others.

The techniques outlined in this document and described in section 2 below, describe best practices for infrastructure protection: edge policy (filtering and precedence), per device traffic policy enforcement for packets destined to a device and, limiting of address and routing visibility to reduce exposure to limit core network -- that is provider (P) and provider edge (PE) infrastructure -- attacks. This document is targeted at network operators seeking to limit their exposure to risks associated with denial of service

targeting the infrastructure. These techniques are designed to be used in addition to specific protocol or application security features implemented in network devices.

Infrastructure protection is a complex topic. While the best practices outlines in this document do not provide perfect protection, they can significantly improve the protection of the network infrastructure.

2. Overview of Infrastructure Protection Techniques

This section provides an overview of recommended techniques that may be used to protect network infrastructure. The areas described below are not exhaustive; other mechanisms can be used to provide additional protection. The techniques discussed in this document have been widely deployed and have proven operational security benefits in large networks.

2.1. Edge Remarking

Edge Remarking, detailed in section 4, ensures that ingress IP precedence or DSCP values match expected values within the context of security. This provides another layer of defense, particularly for traffic permitted through any of the Edge Infrastructure Access Control Lists. This document focuses only on using Edge Remarking best practices to enforce security policies.

2.2. Device and Element Protection

Each network infrastructure device should enforce local rules for traffic destined to itself. These rules can take the form of filters (permit/deny) or rate limiting rules that allow ingress traffic at specified rates. These should complement any existing Edge Infrastructure Access Control Lists and are described in more detail in section 5. The deployment of these local device protection rules complements the edge techniques by protecting the device from traffic that: i) was permitted but violates device policy, ii) could not be filtered at the edge, iii) entered the network on an interface that did not have ingress filtering enabled.

2.3. Infrastructure Hiding

Hiding the infrastructure of the network provides an elegant mechanism for protecting the network infrastructure. If the destination of an attack is to an infrastructure address that is unreachable, attacks become far more difficult. Infrastructure hiding can be achieved in several ways: i) MPLS techniques ii) IGP configuration iii) Route advertisement control. Section 6 addresses these infrastructure hiding techniques in detail.

3. Edge Infrastructure Access Control Lists

Edge Infrastructure Access Control Lists (EIACLs) are a specific implementation of the more general Ingress Access List. As opposed to generic ingress filtering which denies data (sometimes referred to as user) plane traffic, edge infrastructure access control lists do not attempt to deny transit traffic; rather, this form of access control limits traffic destined to infrastructure equipment while permitting -- if needed, explicitly -- traffic through the network.

3.1. Constructing the Access List

Edge Infrastructure Access Control Lists permit only required traffic destined to the network infrastructure, while allowing data plane traffic to flow through unfiltered. The basic premise of EIACLs is

that only a relatively limited subset of traffic sourced from outside an AS should be allowed to transit towards a core router and that by explicitly permitting only that known and understood traffic the core devices are not subjected to unnecessary traffic that might result in a denial of service. Since edge infrastructure access control lists protect only the infrastructure, the development of the list differs somewhat from "traditional" access filter lists:

1. Review addressing scheme, and identify address block(s) that represent core devices.

2. Determine what traffic must be destined to the core devices from outside the AS.

3. Create a filter that allows the required traffic, denies all

traffic destined to the core address block and then finally, permits all other traffic.

4. Optionally, prior to implementing the deny action for all traffic destined to the core address block, a log action may be used and the results of the deny actions evaluated.

As with other ingress filtering techniques, EIACLs are applied on ingress interfaces, as close to the edge as possible. Comprehensive coverage (i.e. on as many interfaces as possible) yields the most protection.

3.2. Other Traffic

In addition to the explicitly permitted traffic, EIACLs can be combined with other common edge filters such as: 1. Source spoof prevention (as per RFC 3704 [RFC3704]) by denying internal AS addresses as external sources. 2. Filtering of reserved addresses (e.g. RFC 1918 [RFC1918] addresses) since traffic should not be sourced from reserved address. 3. Other unneeded or unnecessary traffic. Filtering this traffic can be part of the list explicitly or implicitly; explicit filters often provide log-able information that can be of use during a security audit.

3.3. Edge Infrastructure Conclusion

Edge Infrastructure Access Control Lists provide a very effective first line of defense. EIACLs are not perfect and cannot protect the network against every attack. Furthermore, to be manageable, EIACLs must be able to clearly and simply identify infrastructure address space. To be effective, the EIACLs should be deployed as widely as possible at the edge of the network on devices that support the required filtering performance characteristics.

The potential impact on the device's performance must be taken into consideration when deploying EIACLs.

4. Edge Rewrite/Remarking

Typically edge packet rewrite/remarking deals primarily with traffic passing through a device. However, IP Precedence/DSCP values are

used in prioritizing traffic sent to the devices as well. RFC 1812 [RFC1812] section 5.3 defines the use of IP Precedence in IPv4 packets for routing, management and control traffic. In addition, the RFC [RFC1812] recommends that devices use a mechanism for providing preferential forwarding for packets marked as routing, management or control traffic using IP Preference bits 6 or 7 (110 or 111 in binary). RFC 2474 [RFC2474] defines DSCP and the compatibility of IP Preference bits when using DSCP.

All packets received by customer- and peer-facing Provider Edge (PE) router interfaces with IP Preference values of 6 or 7 or DSCP bits of 11xxxx, as specified in RFC2474 [RFC2474] Differentiated Services Field Definition, should have the IP Preference bits rewritten. Routing traffic received from customer- and peer-facing interfaces can safely have the IP Preference bits rewritten because only a limited number of protocols are transmitted beyond the first PE router. The bits may be rewritten to any value other than IP Preference values 6 or 7, or any DSCP value other than 11xxxx. The new value can be based on the network operators IP Preference or DSCP policy. If no policy exists the bits should be rewritten to 0. In cases where control, management, and routing traffic enters the provider network via the customer and peer facing interfaces policy should be employed to ensure proper prioritization of critical traffic. EIACLs may be be used facilitate the proper classification of traffic. To offer fully transparent service, a provider may not wish to modify the IP precedence on transit traffic through the network. If a provider has alternate means of applying different prioritization to router management and control traffic and transit traffic then rewriting IP precedence bits is not required.

4.1. Edge Rewrite/Remarking Discussion

By default router vendors do not differentiate an interface on a PE router connected to a P router from an interface connected to a CE router. As a result any packet with the proper IP Preference or DSCP bits set may receive the same preferential forwarding behavior as legitimate routing, management, and control traffic. A malicious attack may be able to take advantage of the vulnerability to increase the effectiveness of the attack or to attack the routing, management, and/or control traffic directly. The forwarding prioritization given

to routing, management, and control traffic by default leaves devices

vulnerable to indirect attacks to the infrastructure. By rewriting the IP Precedence at the PE protection is provided for both traffic through the network along with traffic that is to the network that is not blocked by other methods discussed in this document. This document assumes that all customer and peer facing interfaces cannot be trusted for inter- domain diff-serv. In cases where a trust relationship exists for inter-domain diff-serv, diff-serv bits 1xxxxxxx do not have to be rewritten.

4.2. Edge Rewriting/Remarking Performance Considerations

The potential impact on the device's performance must be taken into consideration when rewriting/remarking IP Precedence/DSCP bits. Devices may require additional resources to rewrite/remark packets compared to merely forwarding them.

5. Device/Element Protection

Even with the widest possible deployment of the techniques described above in section 3, Infrastructure Edge Access Control, the individual devices of the network must implement access control mechanisms. In addition to the cases incomplete or imperfect deployment of edge infrastructure controls, threats may come from from trusted sources within the perimeter of the network.

5.1. Service Specific Access Control

Many vendor's implementation of service specific controls are not made with overall system availability as a primary concern. With this in mind, it is recommended that these controls be used in conjunction with any aggregate mechanisms to control device access as well as techniques like EIACLs and Core Hiding. There are many practical examples available of vendor specific service security mechanisms, the references section provides links to several of them. These should guide the operator in securing the services that they enable.

5.1.1. Common Services

While each service implemented by network equipment manufacturers differs in its available security features there are some common services and security features for those services that have been widely deployed. The most important first step for the operator is to disable any unneeded/unused services. This reduces the device's profile. If the device is not listening to a port, it is much more difficult to attack via that port. Second, the operator should

utilize the services access control mechanisms to limit the access to the devices service to only required sources. Examples of per service security are using virtual terminal access control lists, or SNMP Community access control lists.

5.2. Aggregate Device Access Control

Aggregating the security policy -- as opposed to defining it on a per service basis -- allows for a simplified view of the access policies traffic going to the device. Many vendors leverage this simplified view to allow for the policy to be implemented in hardware, protecting the device's control plane. A key requirement of these mechanisms is that it must not impact transit data plane traffic.

5.2.1. IP Fragments

Traffic destined to a router is not typically fragmented. Use of mechanisms to deny fragments to the device are recommended.

5.2.2. Performance Considerations

Care should be taken to understand a vendors implementation of aggregate device access control and to make sure that device operation is not impaired during DoS attacks against the device.

5.2.3. Access Control Implementation Guide

Implementing a complex set of access controls for all traffic going to and from a router is non trivial. The following is a recommended set of steps that has been used successfully by many carriers.

1. Develop list of required protocols.

2. Develop source address requirements: Determine destination interface on router Does the protocol access a single interface? Does the protocol access many interfaces? Does the protocol access a virtual or physical interfaces?

3. Prior to implementing with a deny, it is recommended to test the behavior with the action of "log" and observe the results

4. Deployment should be an iterative process: Start with relatively open lists then tighten as needed

5.3. Device Access Authorization and Accounting

Operators should use per command authorization and accounting wherever possible. Aside from their utility in mitigating other security threats, they provide an invaluable tool in the post event forensics.

6. Infrastructure Hiding

Hiding the infrastructure of the network provides one layer of protection to the devices that make up the network core. By hiding those devices (making them unreachable) successful execution of denial of service attacks becomes far more difficult. Before implementing measures to make network infrastructure unreachable from outside consider carefully that these actions can create limitations that staff, customers, and applications may not expect and weigh these results against the additional security afforded by a hiding the network core. The following sections present different options for accomplishing infrastructure hiding. The operator should carefully consider the merits of each approach based on their network's architecture.

6.1. Use Less IP

One way to reduce exposure of network infrastructure is to use unnumbered links wherever possible. This is particularly useful in the simple case of a customer with single link used as their default path to the Internet. Not only can such a configuration reduce the exposure of the equipment on both ends of the link to malicious attack, the overall effort required to manage a link can be reduced considerably with a simplified configuration and without the additional overhead and expense of managing the IP addresses.

6.2. MPLS Techniques

While it may not be feasible to hide the entire infrastructure of large networks, it is certainly possible to reduce exposure of critical core infrastructure by hiding the existence and complexity

of that infrastructure using an MPLS mesh where TTL is not decremented as packets pass through it as described in RFC 3443 [RFC3443]. In this manner the number, addresses, and even existence of intermediary devices can be hidden from traffic as it passes through the core. As pointed out by RFC 4381 [RFC4381], not only can this method hide the structure, it simplifies access restrictions to core devices. e.g. Core equipment addresses are inaccessible from the "Public Internet" VPN. The fact that this technique is transparent from a layer-3 viewpoint recommends it to providers of

public services. Because basic external troubleshooting and presents to all external views a simplified network structure with few potential target addresses exposed it offers a better balance of security against effort than most techniques for public networks.

6.3. IGP Configuration

Using a non-IP control plane for the core routing protocol can substantially reduce the number of IP addresses that comprise (and therefore, expose) the core. This simplifies the task of maintaining edge ACLs or route announcement filters. IS-IS is an elegant and mature protocol that some operators have found suitable for this task.

6.4. Route Advertisement Control

6.4.1. Route Announcement Filtering

Inasmuch as it is unavoidable that some network elements must be configured with IP addresses, it may be possible to assign these address out of netblocks for which the routing advertisement can be filtered out, thereby limiting possible sources of traffic to core netblocks down to customers for whom you provide a default route, or direct peers who would make the effort to create a static route for your core netblock into your AS. By assigining address for network infrastructure out of a limited number of address blocks which are well known to internal network administrators, the operator can greatly simplify EIACL configuration. This can also minimise the frequency with which ACLs need to be updated based on changes in the network. This can also have performance implications, especially for equipment where the length of ACLs is limited. By keeping ACLs short they may be deployable on a wider range of existing equipment.

Further, it may be possible in those situations where customer point-to-point links must be numbered, to address such links out of another range of addresses for which announcements could be similarly filtered. While this has implications for a customer's ability to remote-monitor their circuit, this can often be overcome with application of an address from the customer's routed space to the CPE loopback.

6.4.2. Address Core Out of RFC 1918 Space

In addition to filtering the visibility of core addresses to the wider Internet, it may be possible to use private RFC 1918 [RFC1918] netblocks for numbering infrastructure when IP addresses are required (eg, loopbacks). This added level of obscurity takes prevention of wide distribution of your infrastructure address space one step further. Many networks filter out packets with RFC 1918 [RFC1918]

address at ingress/ egress points as a matter of course. In this circumstance, tools such as traceroute can work for operations and support staff but not from outside networks. Care should be taken to limit reverse-resolution of descriptive DNS names to queries from internal/support groups. The fact that this technique can break simple troubleshooting tools such as traceroute may frustrate customers who expect to be able to perform basic troubleshooting tasks on their own. Campus or corporate networks may find some advantage to this configuration, on balance. Use caution when employing this technique, particularly in public internet service provider environments.

6.5. Further obfuscation

The strategy of changing services to run on ports different from the default and well-known ones will not protect you from a determined attacker. It can, however, provide some level of protection from many attack tools, worms, auto-rooters, etc. Should they find access to the infrastructure equipment in some way. Again, this does nothing to restrict access, nor to make network devices more difficult to reach. As with the other methods, a careful consideration of how much effort and management each strategy requires must be weighed against the protection that it provides and the necessity of that protection in light of all measures taken to protect a network.

7. IPv6

IPv6 Networks contain the same infrastructure security risks as IPv4. All techniques described in this document for IPv4 should be directly applicable to IPv6 networks. Limitations exist where devices do not have feature parity between IPv4 and IPv6. Different techniques maybe required where IPv4 and IPv6 networks deviate in implementation. Multi-vendor networks can create greater difficulties when each vendor does not have feature parity with each other. Hardware differences in devices that support both IPv4 and IPv6 must also be taken into consideration. Because IPv6 uses a longer address space the scaling, and performance characteristics of ACLs maybe lower for IPv6 vs IPv4. The fields or number of fields that an ACL can match on may also differ. The fact that all PE devices do not support all the recommended IPv6 security features should not preclude the implementation of the recommendations in this document on the devices that do support the security features. With the number of Network Operators deploying IPv6 growing, along with the continued availability of IPv6 Tunnel services, connecting to the IPv6 Internet is less difficult. Dual stack IPv6 networks run on Networks with speeds equal to IPv4. Neither the edge nor the core

limit potential IPv6 attacks. Despite the increased deployment of IPv6 it still does not have the same level of operational experience as IPv4.

7.1. Use IPv6 Edge Infrastructure Access Control Lists

IPv6 Infrastructure security policy will be similar to the IPv4 policy for EIACLs, Edge remarking and Device and Element protection. Construction of the IPv6 EIACL should use the same process as the IPv4 EIACL. The construction of the EIACL can be made less difficult with IPv6 because of the sparse address assignment capability given the larger total address space. IPv6 DSCP bits should be rewritten in the same manner that IPv4 DSCP bits. Differences between DSCP rewriting of IPv4 and IPv6 will minimal except in cases where the device capabilities differ between IPv4 and IPv6. Device and Element protection should be created using the same methods described in this document for IPv4. The policy may differ for IPv6 from IPv4 in cases where services are exclusively IPv4 or exclusively IPv6. Services

not used with IPv6 should be disabled.

7.2. IPv6 Infrastructure Hiding

Network operators may deploy IPv4 differently from IPv6 in their network. Providers may use native forwarding for IPv6 while using MPLS for IPv4, other combinations. IPv6 infrastructure hiding should have parity with IPv4 infrastructure hiding even if the technique used is different. Implementation of IPv6 route advertisement control for infrastructure hiding is difficult when using global address space. Registeries assign fewer large blocks of IPv6 space compared to IPv4. Providers cannot control the announcement of infrastructure global IPv6 blocks for infrastructure hiding without deaggregating their IPv6 announcements.

8. IP Multicast

IP Multicast behaves differently from IP unicast therefore must be secured in a different manner. Some of the protocols used with multicast rely on IP unicast to transport the routing, and control information. Unicast based protocols should be secured using the technique described in much of this document. Multicast security is better addressed in a multicast specific security document.

9. Security Considerations

This entire document is concerned with security.

10. Acknowledgements

Don Smith provided invaluable comments and suggestions. Pat Cain, Ross Callon, Vince Fuller, Barry Greene, George Jones, David Meyer, Peka Savola reviewed this document and provided feedback.

11. References

11.1. Normative References

[RFC1812] Baker, F., Ed., "Requirements for IP Version 4 Routers", June 1995.

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., Lear, E., "Address Allocation for Private Internets", February 1996.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2474] Nichols, K., Blake, S., Baker, F., Black, D., "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", December 1998.

[RFC3443] Agarwal, P., Akyol, B., "Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks", January 2003.

[RFC3704] Baker, F., Savola, P., "Ingress Filtering for Multihomed Networks", March 2004.

[RFC4381] Behringer, M., "Analysis of the Security of BGP/MPLS IP Virtual Private Networks (VPNs)", February 2006.

11.2. Informative References

[AKIN] "Hardening Cisco Routers", T. Akin, O'Reilly Media, 2002, ISBN 0596001665.

[CYMRU-J] "JUNOS Secure Template", S. Gill, Team Cymru, March 2005, http://www.cymru.com/gillsr/documents/junos-template.pdf

[CYMRU-C] "Secure IOS Template", R. Thomas, Team Cymru, March 2007, http://www.cymru.com/Documents/secure-ios-template.html

[GREENE] "Cisco ISP Essentials", B. Greene and P. Smith, Cisco Press, 2002, ISBN 1587050412.

[NANOG-M] "Implications of Securing Backbone Router Infrastructure", R. McDowell, NANOG 31, May 2004. http://www.nanog.org/mtg-0405/mcdowell.html

[RFC2334] Narten, T., and H. Alverstand, "Guidelines for Writing an IANA Considerations Section in RFCs", October 1998.

[RFC2827] Ferguson, P., Senie, D., "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", May 2000.

[RFC3669] Bradner, S., "Intellectual Property Rights in IETF Technology", February 2004.

[RFC3871] Jones, G., Ed., "Operational Security Requirements for Large Internet Service Provider (ISP) IP Network Infrastructure", September 2004.

[RFC3978] Bradner, S., Ed., "IETF Rights in Contributions", February 2004.

[RFC4778] Kaeo, M., "Current Operational Security Practices in Internet Service Provider Environments", January 2007.

Authors' Addresses

James Gill
Verizon Business
22001 Louden County Parkway
Ashburn, VA 20147
US

Phone: +1-703-886-3834
Email: james.gill@verizonbusiness.com
URI: www.verizonbusiness.com

Darrel Lewis
Cisco Systems Inc.
170 West Tasman Dr.
San Jose, CA 95134
US

Phone: +1-408-853-3653
Email: darlewis@cisco.com

URI: www.cisco.com

Paul Quinn
Cisco Systems Inc.
170 West Tasman Drive
San Jose, CA 95134
US

Phone: +1-408-527-3560
Email: paulq@cisco.com
URI: www.cisco.com

Peter Schoenmaker
NTT America
101 Park Ave., FL 41
New York, NY 10178
US

Phone: +1-202-808-2298
Fax:
Email: pds@ntt.net
URI:

Full Copyright Statement

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).