

Network Working Group
Request for Comments: 4778
Category: Informational

M. Kaeo
Double Shot Security, Inc.
January 2007

Current Operational Security Practices in
Internet Service Provider Environments

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This document is a survey of the current practices used in today's large ISP operational networks to secure layer 2 and layer 3 infrastructure devices. The information listed here is the result of information gathered from people directly responsible for defining and implementing secure infrastructures in Internet Service Provider environments.

Table of Contents

- 1. Introduction 2
 - 1.1. Scope 2
 - 1.2. Threat Model 3
 - 1.3. Attack Sources 4
 - 1.4. Operational Security Impact from Threats 5
 - 1.5. Document Layout 7
- 2. Protected Operational Functions 8
 - 2.1. Device Physical Access 8
 - 2.2. Device Management - In-Band and Out-of-Band (OOB) 10
 - 2.3. Data Path 16
 - 2.4. Routing Control Plane 18
 - 2.5. Software Upgrades and Configuration Integrity/Validation 22
 - 2.6. Logging Considerations 26
 - 2.7. Filtering Considerations 29
 - 2.8. Denial-of-Service Tracking/Tracing 30
- 3. Security Considerations 32
- 4. Acknowledgments 32
- 5. References 33
 - 5.1. Normative References 33
 - 5.2. Informational References 33
- Appendix A. Protocol Specific Attacks 34
 - A.1. Layer 2 Attacks 34
 - A.2. IPv4 Protocol-Based Attacks 34
 - A.3. IPv6 Attacks 36

1. Introduction

Security practices are well understood by the network operators who have, for many years, gone through the growing pains of securing their network infrastructures. However, there does not exist a written document that enumerates these security practices. Network attacks are continually increasing and although it is not necessarily the role of an ISP to act as the Internet police, each ISP has to ensure that certain security practices are followed to ensure that their network is operationally available for their customers. This document is the result of a survey conducted to find out what current security practices are being deployed to secure network infrastructures.

1.1. Scope

The scope for this survey is restricted to security practices that mitigate exposure to risks with the potential to adversely impact network availability and reliability. Securing the actual data traffic is outside the scope of the conducted survey. This document

focuses solely on documenting currently deployed security mechanisms for layer 2 and layer 3 network infrastructure devices. Although primarily focused on IPv4, many of the same practices can (and should) apply to IPv6 networks. Both IPv4 and IPv6 network infrastructures are taken into account in this survey.

1.2. Threat Model

A threat is a potential for a security violation, which exists when there is a circumstance, capability, action, or event that could breach security and cause harm [RFC2828]. Every operational network is subject to a multitude of threat actions, or attacks, i.e., an assault on system security that derives from an intelligent act that is a deliberate attempt to evade security services, and violate the security policy of a system [RFC2828]. Many of the threats to a network infrastructure occur from an instantiation (or combination) of the following:

Reconnaissance: An attack whereby information is gathered to ascertain the network topology or specific device information, which can be further used to exploit known vulnerabilities

Man-In-The-Middle: An attack where a malicious user impersonates either the sender or recipient of a communication stream while inserting, modifying, or dropping certain traffic. This type of attack also covers phishing and session hijacks.

Protocol Vulnerability Exploitation: An attack that takes advantage of known protocol vulnerabilities due to design or implementation flaws to cause inappropriate behavior.

Message Insertion: This can be a valid message (it could be a reply attack, which is a scenario where a message is captured and resent at a later time). A message can also be inserted with any of the fields in the message being spoofed, such as IP addresses, port numbers, header fields, or even packet content. Flooding is also part of this threat instantiation.

Message Diversion/Deletion: An attack where legitimate messages are removed before they can reach the desired recipient, or are re-directed to a network segment that is normally not part of the data path.

Message Modification: This is a subset of a message insertion attack where a previous message has been captured and modified before being retransmitted. The message can be captured using a man-in-the-middle attack or message diversion.

Note that sometimes denial-of-service attacks are listed as separate categories. A denial-of-service is a consequence of an attack and can be the result of too much traffic (i.e., flooding), exploiting protocol exploitation, or inserting/deleting/diverting/modifying messages.

1.3. Attack Sources

These attacks can be sourced in a variety of ways:

Active vs Passive Attacks

An active attack involves writing data to the network. It is common practice in active attacks to disguise one's address and conceal the identity of the traffic sender. A passive attack involves only reading information off the network. This is possible if the attacker has control of a host in the communications path between two victim machines, or has compromised the routing infrastructure to specifically arrange that traffic pass through a compromised machine. There are also situations where mirrored traffic (often used for debugging, performance monitoring, or accounting purposes) is diverted to a compromised machine, which would not necessarily subvert any existing topology, and could be harder to detect. In general, the goal of a passive attack is to obtain information that the sender and receiver would prefer to remain private [RFC3552].

On-path vs Off-path Attacks

In order for a datagram to be transmitted from one host to another, it generally must traverse some set of intermediate links and routers. Such routers are naturally able to read, modify, or remove any datagram transmitted along that path. This makes it much easier to mount a wide variety of attacks if you are on-path. Off-path hosts can transmit arbitrary datagrams that appear to come from any host but cannot necessarily receive datagrams intended for other hosts. Thus, if an attack depends on being able to receive data, off-path hosts must first subvert the topology in order to place themselves on-path. This is by no means impossible, but is not necessarily trivial [RFC3552]. A more subtle attack is one where the traffic-mirroring capability of a device is hijacked and the traffic is diverted to a compromised host since the network topology may not need to be subverted.

Insider vs Outsider Attacks

An "insider attack" is initiated from inside a given security perimeter by an entity that is authorized to access system resources, but uses them in a way not approved by those who granted the authorization. An "outside attack" is initiated from outside the perimeter by an unauthorized or illegitimate user of the system.

Deliberate Attacks vs Unintentional Events

A deliberate attack is where a miscreant intentionally performs an assault on system security. However, there are also instances where unintentional events cause the same harm, yet are performed without malicious intent. Configuration errors and software bugs can be as devastating to network availability as any deliberate attack on the network infrastructure.

The attack source can be a combination of any of the above, all of which need to be considered when trying to ascertain the impact any attack can have on the availability and reliability of the network. It is nearly impossible to stop insider attacks or unintentional events. However, if appropriate monitoring mechanisms are in place, these attacks can also be detected and mitigated as with any other attack source. The amount of effort it takes to identify and trace an attack is, of course, dependent on the resourcefulness of the attacker. Any of the specific attacks discussed further in this document will elaborate on malicious behavior, which are sourced by an "outsider" and are deliberate attacks. Some further elaboration will be given to the feasibility of passive vs active and on-path vs off-path attacks to show the motivation behind deploying certain security features.

1.4. Operational Security Impact from Threats

The main concern for any of the potential attack scenarios is the impact and harm it can cause to the network infrastructure. The threat consequences are the security violations that results from a threat action, i.e., an attack. These are typically classified as follows:

(Unauthorized) Disclosure

A circumstance or event whereby an entity gains access to data for which the entity is not authorized.

Deception

A circumstance or event that may result in an authorized entity receiving false data and believing it to be true.

Disruption

A circumstance or event that interrupts or prevents the correct operation of system services and functions. A broad variety of attacks, collectively called denial of service attacks, threaten the availability of systems and bandwidth to legitimate users. Many such attacks are designed to consume machine resources, making it difficult or impossible to serve legitimate users. Other attacks cause the target machine to crash, completely denying service to users.

Usurpation

A circumstance or event that results in control of system services or functions by an unauthorized entity. Most network infrastructure systems are only intended to be completely accessible to certain authorized individuals. Should an unauthorized person gain access to critical layer 2/layer 3 infrastructure devices or services, they could cause great harm to the reliability and availability of the network.

A complete description of threat actions that can cause these threat consequences can be found in [RFC2828]. Typically, a number of different network attacks are used in combination to cause one or more of the above-mentioned threat consequences. An example would be a malicious user who has the capability to eavesdrop on traffic. First, he may listen in on traffic for a while, doing reconnaissance work and ascertaining which IP addresses belong to specific devices, such as routers. Were this miscreant to obtain information, such as a router password sent in cleartext, he can then proceed to compromise the actual router. From there, the miscreant can launch various active attacks, such as sending bogus routing updates to redirect traffic or capture additional traffic to compromise other network devices. While this document enumerates which countermeasures ISPs are deploying today, a useful generic analysis of actual backbone infrastructure attacks and the appropriate countermeasures can be found in [RTGWG].

1.5. Document Layout

This document is a survey of current operational practices that mitigate the risk of being susceptible to any threat actions. As such, the main focus is on the currently deployed security practices used to detect and/or mitigate attacks. The top-level categories in this document are based on operational functions for ISPs and generally relate to what is to be protected. This is followed by a description of which attacks are possible and the security practices currently deployed. This will provide the necessary security services to help mitigate these attacks. These security services are classified as follows:

- o User Authentication
- o User Authorization
- o Data Origin Authentication
- o Access Control
- o Data Integrity
- o Data Confidentiality
- o Auditing/Logging
- o DoS Mitigation

In many instances, a specific protocol currently deployed will offer a combination of these services. For example, Authentication, Authorization, and Accounting (AAA) can offer user authentication, user authorization, and audit/logging services, while the Secure SHell (SSH) Protocol can provide data origin authentication, data integrity, and data confidentiality. The services offered are more important than the actual protocol used. Note that access control will refer basically to logical access control, i.e., filtering. Each section ends with an additional considerations section that explains why specific protocols may or may not be used, and also gives some information regarding capabilities, which are not possible today due to bugs or lack of usability.

2. Protected Operational Functions

2.1. Device Physical Access

Device physical access pertains to protecting the physical location and access of the layer 2 or layer 3 network infrastructure device. Physical security is a large field of study/practice in and of itself, arguably the largest, oldest, and most well-understood area of security. Although it is important to have contingency plans for natural disasters, such as earthquakes and floods, which can cause damage to networking devices, this is out of the scope of this document. Here, we concern ourselves with protecting access to the physical location and how a device can be further protected from unauthorized access if the physical location has been compromised, i.e., protecting the console access. This is aimed largely at stopping an intruder with physical access from gaining operational control of the device(s). Note that nothing will stop an attacker with physical access from effecting a denial-of-service attack, which can be easily accomplished by powering off the device or just unplugging some cables.

2.1.1. Threats/Attacks

If any intruder gets physical access to a layer 2 or layer 3 device, the entire network infrastructure can be under the control of the intruder. At a minimum, the intruder can take the compromised device out of service, causing network disruption, the extent of which depends on the network topology. A worse scenario is where the intruder uses this device to crack the console password, gaining complete control of the device (perhaps without anyone detecting such a compromise, or to attach another network device onto a port and siphon off data with which the intruder can ascertain the network topology) and the entire network.

The threat of gaining physical access can be realized in a variety of ways, even if critical devices are under high security. Cases still occur where attackers have impersonated maintenance workers to gain physical access to critical devices that have caused major outages and privacy compromises. Insider attacks from authorized personnel also pose a real threat and must be adequately recognized and addressed.

2.1.2. Security Practices

For physical device security, equipment is kept in highly restrictive environments. Only authorized users with card-key badges have access to any of the physical locations that contain critical network infrastructure devices. These card-key systems keep track of who

accessed which location and at what time. Most cardkey systems have a fail-back "master key" in case the card system is down. This "master key" usually has limited access and its use is also carefully logged (which should only happen if the card-key system is NOT online/functional).

All console access is always password protected and the login time is set to time out after a specified amount of inactivity - typically between 3-10 minutes. The type of privileges that you obtain from a console login varies between separate vendor devices. In some cases you get initial basic access and need to perform a second authentication step to get more privileged access (i.e., enable or root). In other vendors, you get the more privileged access when you log into the console as root, without requiring a second authentication step.

How ISPs manage these logins vary greatly, although many of the larger ISPs employ some sort of AAA mechanism to help automate privilege-level authorization and utilize the automation to bypass the need for a second authentication step. Also, many ISPs define separate classes of users to have different privileges while logged onto the console. Typically, all console access is provided via an out-of-band (OOB) management infrastructure, which is discussed in Section 2.2 of this document.

2.1.3. Security Services

The following security services are offered through the use of the practices described in the previous section:

- o User Authentication - All individuals who have access to the physical facility are authenticated. Console access is authenticated.
- o User Authorization - An authenticated individual has implicit authorization to perform commands on the device. In some cases, multiple authentication is required to differentiate between basic and more privileged access.
- o Data Origin Authentication - Not applicable.
- o Access Control - Not applicable.
- o Data Integrity - Not applicable.
- o Data Confidentiality - Not applicable.

- o Auditing/Logging - All access to the physical locations of the infrastructure equipment is logged via electronic card-key systems. All console access is logged (refer to Section 2.2 of this document for more details).
- o DoS Mitigation - Not applicable.

2.1.4. Additional Considerations

Physical security is relevant to operational security practices as described in this document, mostly from a console-access perspective. Most ISPs provide console access via an OOB management infrastructure, which is discussed in Section 2.2 of this document.

The physical and logical authentication and logging systems should be run independently of each other and should reside in different physical locations. These systems need to be secured to ensure that they themselves will not be compromised, which could give the intruder valuable authentication and logging information.

Social engineering plays a big role in many physical access compromises. Most ISPs have set up training classes and awareness programs to educate company personnel to deny physical access to people who are not properly authenticated or authorized to have physical access to critical infrastructure devices.

2.2. Device Management - In-Band and Out-of-Band (OOB)

In-band management is generally considered to be device access, where the control traffic takes the same data path as the data that traverses the network. Out-of-band management is generally considered to be device access, where the control traffic takes a separate path as the data that traverses the network. In many environments, device management for layer 2 and layer 3 infrastructure devices is deployed as part of an out-of-band management infrastructure, although there are some instances where it is deployed in-band as well. Note that while many of the security concerns and practices are the same for OOB management and in-band management, most ISPs prefer an OOB management system, since access to the devices that make up this management network are more vigilantly protected and considered to be less susceptible to malicious activity.

Console access is always architected via an OOB network. Presently, the mechanisms used for either in-band management or OOB are via virtual terminal access (i.e., Telnet or SSH), Simple Network Management Protocol (SNMP), or HTTP. In all large ISPs that were interviewed, HTTP management was never used and was explicitly

disabled. Note that file transfer protocols (TFTP, FTP, and SCP) will be covered in Section 2.5 of this document.

2.2.1. Threats/Attacks

For device management, passive attacks are possible if someone has the capability to intercept data between the management device and the managed device. The threat is possible if a single infrastructure device is somehow compromised and can act as a network sniffer, or if it is possible to insert a new device that acts as a network sniffer.

Active attacks are possible for both on-path and off-path scenarios. For on-path active attacks, the situation is the same as for a passive attack, where either a device has to already be compromised or a device can be inserted into the path. For off-path active attacks, where a topology subversion is required to reroute traffic and essentially bring the attacker on-path, the attack is generally limited to message insertion or modification.

2.2.1.1. Confidentiality Violations

Confidentiality violations can occur when a miscreant intercepts any management data that has been sent in cleartext or with weak encryption. This includes interception of usernames and passwords with which an intruder can obtain unauthorized access to network devices. It can also include other information, such as logging or configuration information, if an administrator is remotely viewing local logfiles or configuration information.

2.2.1.2. Offline Cryptographic Attacks

If username/password information was encrypted but the cryptographic mechanism used made it easy to capture data and break the encryption key, the device management traffic could be compromised. The traffic would need to be captured either by eavesdropping on the network or by being able to divert traffic to a malicious user.

2.2.1.3. Replay Attacks

For a replay attack to be successful, the management traffic would need to first be captured either on-path or diverted to an attacker to later be replayed to the intended recipient.

2.2.1.4. Message Insertion/Deletion/Modification

Data can be manipulated by someone in control of intermediary hosts. Forging data is also possible with IP spoofing, where a remote host sends out packets that appear to come from another, trusted host.

2.2.1.5. Man-In-The-Middle

A man-in-the-middle attack attacks the identity of a communicating peer rather than the data stream itself. The attacker intercepts traffic that is sent from a management system to the networking infrastructure device and traffic that is sent from the network infrastructure device to the management system.

2.2.2. Security Practices

OOB management is done via a terminal server at each location. SSH access is used to get to the terminal server from where sessions to the devices are initiated. Dial-in access is deployed as a backup if the network is not available. However, it is common to use dial-back, encrypting modems, and/or one-time-password (OTP) modems to avoid the security weaknesses of plain dial-in access.

All in-band management and OOB management access to layer 2 and layer 3 devices is authenticated. The user authentication and authorization is typically controlled by an AAA server (i.e., Remote Authentication Dial-in User Service (RADIUS) and/or Terminal Access Controller Access-Control System (TACACS+)). Credentials used to determine the identity of the user vary from static username/password to one-time username/password schemes such as Secure-ID. Static username/passwords are expired after a specified period of time, usually 30 days. Every authenticated entity via AAA is an individual user for greater granularity of control. Note that often the AAA server used for OOB management authentication is a separate physical device from the AAA server used for in-band management user authentication. In some deployments, the AAA servers used for device management authentication/authorization/accounting are on separate networks to provide a demarcation for any other authentication functions.

For backup purposes, there is often a single local database entry for authentication that is known to a very limited set of key personnel. It is usually the highest privilege-level username/password combination, which in most cases is the same across all devices. This local device password is routinely regenerated once every 2-3 months, and is also regenerated immediately after an employee who had access to that password leaves the company or is no longer authorized to have knowledge of that password.

Each individual user in the AAA database is configured with specific authorization capability. Specific commands are either individually denied or permitted, depending on the capability of the device to be accessed. Multiple privilege levels are deployed. Most individuals are authorized with basic authorization to perform a minimal set of commands, while a subset of individuals are authorized to perform more privileged commands. Securing the AAA server is imperative and access to the AAA server itself is strictly controlled. When an individual leaves the company, his/her AAA account is immediately deleted and the TACACS/RADIUS shared secret is reset for all devices.

Some management functions are performed using command line interface (CLI) scripting. In these scenarios, a dedicated user is used for the identity in scripts that perform CLI scripting. Once authenticated, these scripts control which commands are legitimate, depending on authorization rights of the authenticated individual.

SSH is always used for virtual terminal access to provide for an encrypted communication channel. There are exceptions due to equipment limitations which are described in the additional considerations section.

If SNMP is used for management, it is for read queries only and restricted to specific hosts. If possible, the view is also restricted to only send the information that the management station needs, rather than expose the entire configuration file with the read-only SNMP community. The community strings are carefully chosen to be difficult to crack and there are procedures in place to change these community strings between 30-90 days. If systems support two SNMP community strings, the old string is replaced by first configuring a second, newer community string and then migrating over from the currently used string to the newer one. Most large ISPs have multiple SNMP systems accessing their routers so it takes more than one maintenance period to get all the strings fixed in all the right systems. SNMP RW is not used and is disabled by configuration.

Access control is strictly enforced for infrastructure devices by using stringent filtering rules. A limited set of IP addresses are allowed to initiate connections to the infrastructure devices and are specific to the services to which they are to be limited (i.e., SSH and SNMP).

All device management access is audited and any violations trigger alarms that initiate automated email, pager, and/or telephone notifications. AAA servers keep track of the authenticated entity as well as all the commands that were carried out on a specific device. Additionally, the device itself logs any access control violations (i.e., if an SSH request comes in from an IP address that is not

explicitly permitted, that event is logged so that the offending IP address can be tracked down and investigations made as to why it was trying to access a particular infrastructure device)

2.2.3. Security Services

The security services offered for device OOB management are nearly identical to those of device in-band management. Due to the critical nature of controlling and limiting device access, many ISPs feel that physically separating the management traffic from the normal customer data traffic will provide an added level of risk mitigation and limit the potential attack vectors. The following security services are offered through the use of the practices described in the previous section:

- o User Authentication - All individuals are authenticated via AAA services.
- o User Authorization - All individuals are authorized via AAA services to perform specific operations once successfully authenticated.
- o Data Origin Authentication - Management traffic is strictly filtered to allow only specific IP addresses to have access to the infrastructure devices. This does not alleviate risk the from spoofed traffic, although when combined with edge filtering using BCP38 [RFC2827] and BCP84 [RFC3704] guidelines (discussed in Section 2.5), then the risk of spoofing is mitigated, barring a compromised internal system. Also, using SSH for device access ensures that no one can spoof the traffic during the SSH session.
- o Access Control - Management traffic is filtered to allow only specific IP addresses to have access to the infrastructure devices.
- o Data Integrity - Using SSH provides data integrity and ensures that no one has altered the management data in transit.
- o Data Confidentiality - Using SSH provides data confidentiality.
- o Auditing/Logging - Using AAA provides an audit trail for who accessed which device and which operations were performed.
- o DoS Mitigation - Using packet filters to allow only specific IP addresses to have access to the infrastructure devices. This limits but does not prevent spoofed DoS attacks directed at an infrastructure device. However, the risk is lowered by using a separate physical network for management purposes.

2.2.4. Additional Considerations

Password selection for any device management protocol used is critical to ensure that the passwords are hard to guess or break using a brute-force attack.

IP security (IPsec) is considered too difficult to deploy, and the common protocol to provide for confidential management access is SSH. There are exceptions for using SSH due to equipment limitations since SSH may not be supported on legacy equipment. In some cases, changing the host name of a device requires an SSH rekey event since the key is based on some combination of host name, Message Authentication Code (MAC) address, and time. Also, in the case where the SSH key is stored on a route processor card, a re-keying of SSH would be required whenever the route processor card needs to be swapped. Some providers feel that this operational impact exceeds the security necessary and instead use Telnet from trusted inside hosts (called 'jumphosts' or 'bastion hosts') to manage those devices. An individual would first SSH to the jumphost and then Telnet from the jumphost to the actual infrastructure device, fully understanding that any passwords will be sent in the clear between the jumphost and the device to which it is connecting. All authentication and authorization is still carried out using AAA servers.

In instances where Telnet access is used, the logs on the AAA servers are more verbose and more attention is paid to them to detect any abnormal behavior. The jumphosts themselves are carefully controlled machines and usually have limited access. Note that Telnet is NEVER allowed to an infrastructure device except from specific jumphosts; i.e., packet filters are used at the console server and/or infrastructure device to ensure that Telnet is only allowed from specific IP addresses.

With thousands of devices to manage, some ISPs have created automated mechanisms to authenticate to devices. As an example, Kerberos has been used to automate the authentication process for devices that have support for Kerberos. An individual would first log in to a Kerberized UNIX server using SSH and generate a Kerberos 'ticket'. This 'ticket' is generally set to have a lifespan of 10 hours and is used to automatically authenticate the individual to the infrastructure devices.

In instances where SNMP is used, some legacy devices only support SNMPv1, which then requires the provider to mandate its use across all infrastructure devices for operational simplicity. SNMPv2 is primarily deployed since it is easier to set up than v3.

2.3. Data Path

This section refers to how traffic is handled that traverses the network infrastructure device. The primary goal of ISPs is to forward customer traffic. However, due to the large amount of malicious traffic that can cause DoS attacks and render the network unavailable, specific measures are sometimes deployed to ensure the availability to forward legitimate customer traffic.

2.3.1. Threats/Attacks

Any data traffic can potentially be attack traffic and the challenge is to detect and potentially stop forwarding any of the malicious traffic. The deliberately sourced attack traffic can consist of packets with spoofed source and/or destination addresses or any other malformed packet that mangle any portion of a header field to cause protocol-related security issues (such as resetting connections, causing unwelcome ICMP redirects, creating unwelcome IP options, or packet fragmentations).

2.3.2. Security Practices

Filtering and rate limiting are the primary mechanism to provide risk mitigation of malicious traffic rendering the ISP services unavailable. However, filtering and rate limiting of data path traffic is deployed in a variety of ways, depending on how automated the process is and what the capabilities and performance limitations of the existing deployed hardware are.

The ISPs that do not have performance issues with their equipment follow BCP38 [RFC2827] and BCP84 [RFC3704] guidelines for ingress filtering. BCP38 recommends filtering ingress packets with obviously spoofed and/or 'reserved' source addresses to limit the effects of denial-of-service attacks, while BCP84 extends the recommendation for multi-homed environments. Filters are also used to help alleviate issues between service providers. Without any filtering, an inter-exchange peer could steal transit just by using static routes, and essentially redirect data traffic. Therefore, some ISPs have implemented ingress/egress filters that block unexpected source and destination addresses not defined in the above-mentioned documents. Null routes and black-hole triggered routing [RFC3882] are used to deter any detected malicious traffic streams. These two techniques are described in more detail in Section 2.8 below.

Most ISPs consider layer 4 filtering useful, but it is only implemented if performance limitations allow for it. Since it poses a large administrative overhead and ISPs are very much opposed to acting as the Internet firewall, Layer 4 filtering is typically

implemented as a last option. Netflow is used for tracking traffic flows, but there is some concern whether sampling is good enough to detect malicious behavior.

Unicast Reverse Path Forwarding (RPF) is not consistently implemented. Some ISPs are in the process of doing so, while other ISPs think that the perceived benefit of knowing that spoofed traffic comes from legitimate addresses are not worth the operational complexity. Some providers have a policy of implementing uRPF at link speeds of Digital Signal 3 (DS3) and below, which was due to the fact that all hardware in the network supported uRPF for DS3 speeds and below. At higher-speed links, the uRPF support was inconsistent and it was easier for operational people to implement a consistent solution.

2.3.3. Security Services

- o User Authentication - Not applicable.
- o User Authorization - Not applicable.
- o Data Origin Authentication - When IP address filtering per BCP38, BCP84, and uRPF are deployed at network edges it can ensure that any spoofed traffic comes from at least a legitimate IP address and can be tracked.
- o Access Control - IP address filtering and layer 4 filtering is used to deny forbidden protocols and limit traffic destined for infrastructure device itself. Filters are also used to block unexpected source/destination addresses.
- o Data Integrity - Not applicable.
- o Data Confidentiality - Not applicable.
- o Auditing/Logging - Filtering exceptions are logged for potential attack traffic.
- o DoS Mitigation - Black-hole triggered filtering and rate-limiting are used to limit the risk of DoS attacks.

2.3.4. Additional Considerations

For layer 2 devices, MAC address filtering and authentication is not used in large-scale deployments. This is due to the problems it can cause when troubleshooting networking issues. Port security becomes unmanageable at a large scale where thousands of switches are deployed.

Rate limiting is used by some ISPs, although other ISPs believe it is not really useful, since attackers are not well-behaved and it doesn't provide any operational benefit over the complexity. Some ISPs feel that rate limiting can also make an attacker's job easier by requiring the attacker to send less traffic to starve legitimate traffic that is part of a rate limiting scheme. Rate limiting may be improved by developing flow-based rate-limiting capabilities with filtering hooks. This would improve the performance as well as the granularity over current capabilities.

Lack of consistency regarding the ability to filter, especially with respect to performance issues, cause some ISPs not to implement BCP38 and BCP84 guidelines for ingress filtering. One such example is at edge boxes, where up to 1000 Tls connecting into a router with an OC-12 (Optical Carrier) uplink. Some deployed devices experience a large performance impact with filtering, which is unacceptable for passing customer traffic through, though ingress filtering (uRPF) might be applicable at the devices that are connecting these aggregation routers. Where performance is not an issue, the ISPs make a tradeoff between management versus risk.

2.4. Routing Control Plane

The routing control plane deals with all the traffic that is part of establishing and maintaining routing protocol information.

2.4.1. Threats/Attacks

Attacks on the routing control plane can be from both passive or active sources. Passive attacks are possible if someone has the capability to intercept data between the communicating routing peers. This can be accomplished if a single routing peer is somehow compromised and can act as a network sniffer, or if it is possible to insert a new device that acts as a network sniffer.

Active attacks are possible for both on-path and off-path scenarios. For on-path active attacks, the situation is the same as for a passive attack, where either a device has to already be compromised or a device can be inserted into the path. This may lead to an attacker impersonating a legitimate routing peer and exchanging routing information. Unintentional active attacks are more common due to configuration errors, which cause legitimate routing peers to feed invalid routing information to other neighboring peers.

For off-path active attacks, the attacks are generally limited to message insertion or modification, which can divert traffic to illegitimate destinations, causing traffic to never reach its intended destination.

2.4.1.1. Confidentiality Violations

Confidentiality violations can occur when a miscreant intercepts any of the routing update traffic. This is becoming more of a concern because many ISPs are classifying addressing schemes and network topologies as private and proprietary information. It is also a concern because the routing protocol packets contain information that may show ways in which routing sessions could be spoofed or hijacked. This in turn could lead into a man-in-the-middle attack, where the miscreants can insert themselves into the traffic path or divert the traffic path and violate the confidentiality of user data.

2.4.1.2. Offline Cryptographic Attacks

If any cryptographic mechanism was used to provide for data integrity and confidentiality, an offline cryptographic attack could potentially compromise the data. The traffic would need to be captured either by eavesdropping on the network or by being able to divert traffic to a malicious user. Note that by using cryptographically protected routing information, the latter would require the cryptographic key to already be compromised anyway, so this attack is only feasible if a device was able to eavesdrop and capture the cryptographically protected routing information.

2.4.1.3. Replay Attacks

For a replay attack to be successful, the routing control plane traffic would need to first be captured either on-path or diverted to an attacker to later be replayed to the intended recipient. Additionally, since many of these protocols include replay protection mechanisms, these would also need to be subverted, if applicable.

2.4.1.4. Message Insertion/Deletion/Modification

Routing control plane traffic can be manipulated by someone in control of intermediate hosts. In addition, traffic can be injected by forging IP addresses, where a remote router sends out packets that appear to come from another, trusted router. If enough traffic is injected to be processed by limited memory routers, it can cause a DoS attack.

2.4.1.5. Man-In-The-Middle

A man-in-the-middle attack attacks the identity of a communicating peer rather than the data stream itself. The attacker intercepts traffic that is sent from one routing peer to the other and communicates on behalf of one of the peers. This can lead to a diversion of the user traffic to either an unauthorized receiving

party or cause legitimate traffic to never reach its intended destination.

2.4.2. Security Practices

Securing the routing control plane takes many features, which are generally deployed as a system. Message Digest 5 (MD5) authentication is used by some ISPs to validate the sending peer and to ensure that the data in transit has not been altered. Some ISPs only deploy MD5 authentication at the customers' request. Additional sanity checks to ensure with reasonable certainty that the received routing update was originated by a valid routing peer include route filters and the Generalized TTL Security Mechanism (GTSM) feature [RFC3682] (sometimes also referred to as the TTL-Hack). The GTSM feature is used for protocols such as the Border Gateway Protocol (BGP), and makes use of a packet's Time To Live (TTL) field (IPv4) or Hop Limit (IPv6) to protect communicating peers. If GTSM is used, it is typically deployed only in limited scenarios between internal BGP peers due to lack of consistent support between vendor products and operating system versions.

Packet filters are used to limit which systems can appear as a valid peer, while route filters are used to limit which routes are believed to be from a valid peer. In the case of BGP routing, a variety of policies are deployed to limit the propagation of invalid routing information. These include: incoming and outgoing prefix filters for BGP customers, incoming and outgoing prefix filters for peers and upstream neighbors, incoming AS-PATH filter for BGP customers, outgoing AS-PATH filter towards peers and upstream neighbors, route dampening and rejecting selected attributes and communities. Consistency between these policies varies greatly and there is a definite distinction whether the other end is an end-site vs an internal peer vs another big ISP or customer. Mostly ISPs do prefix-filter their end-site customers, but due to the operational constraints of maintaining large prefix filter lists, many ISPs are starting to depend on BGP AS-PATH filters to/from their peers and upstream neighbors.

In cases where prefix lists are not used, operators often define a maximum prefix limit per peer to prevent misconfiguration (e.g., unintentional de-aggregation or neighbor routing policy mis-configuration) or overload attacks. ISPs need to coordinate with each other what the expected prefix exchange is, and increase this number by some sane amount. It is important for ISPs to pad the max-prefix number enough to allow for valid swings in routing announcements, preventing an unintentional shut down of the BGP session. Individual implementation amongst ISPs are unique, and depending on equipment supplier(s), different implementation options

are available. Most equipment vendors offer implementation options ranging from just logging excessive prefixes being received, to automatically shutting down the session. If the option of reestablishing a session after some pre-configured idle timeout has been reached is available, it should be understood that automatically reestablishing the session may potentially introduce instability continuously into the overall routing table if a policy mis-configuration on the adjacent neighbor is causing the condition. If a serious mis-configuration on a peering neighbor has occurred, then automatically shutting down the session and leaving it shut down until being manually cleared, is sometimes best and allows for operator intervention to correct as needed.

Some large ISPs require that routes be registered in an Internet Routing Registry (IRR), which can then be part of the Routing Assets Database (RADb) - a public registry of routing information for networks in the Internet that can be used to generate filter lists. Some ISPs, especially in Europe, require registered routes before agreeing to become an eBGP peer with someone.

Many ISPs also do not propagate interface IP addresses to further reduce attack vectors on routers and connected customers.

2.4.3. Security Services

- o User Authentication - Not applicable.
- o User Authorization - Not applicable.
- o Data Origin Authentication - By using MD5 authentication and/or the TTL-hack, a routing peer can be reasonably certain that traffic originated from a valid peer.
- o Access Control - Route filters, AS-PATH filters, and prefix limits are used to control access to specific parts of the network.
- o Data Integrity - By using MD5 authentication, a peer can be reasonably certain that the data has not been modified in transit, but there is no mechanism to prove the validity of the routing information itself.
- o Data Confidentiality - Not implemented.
- o Auditing / Logging - Filter exceptions are logged.

- o DoS Mitigation - Many DoS attacks are mitigated using a combination of techniques including: MD5 authentication, the GTSM feature, filtering routing advertisements to bogons, and filtering routing advertisements to one's own network.

2.4.4. Additional Considerations

So far the primary concern to secure the routing control plane has been to validate the sending peer and to ensure that the data in transit has not been altered. Although MD5 routing protocol extensions have been implemented, which can provide both services, they are not consistently deployed amongst ISPs. Two major deployment concerns have been implementation issues, where both software bugs and the lack of graceful re-keying options have caused significant network down times. Also, some ISPs express concern that deploying MD5 authentication will itself be a worse DoS attack victim and prefer to use a combination of other risk mitigation mechanisms such as GTSM (for BGP) and route filters. An issue with GTSM is that it is not supported on all devices across different vendors' products.

IPsec is not deployed since the operational management aspects of ensuring interoperability and reliable configurations is too complex and time consuming to be operationally viable. There is also limited concern to the confidentiality of the routing information. The integrity and validity of the updates are of much greater concern.

There is concern for manual or automated actions, which introduce new routes and can affect the entire routing domain.

2.5. Software Upgrades and Configuration Integrity/Validation

Software upgrades and configuration changes are usually performed as part of either in-band or OOB management functions. However, there are additional considerations to be taken into account, which are enumerated in this section.

2.5.1. Threats/Attacks

Attacks performed on system software and configurations can be both from passive or active sources. Passive attacks are possible if someone has the capability to intercept data between the network infrastructure device and the system which is downloading or uploading the software or configuration information. This can be accomplished if a single infrastructure device is somehow compromised and can act as a network sniffer, or if it is possible to insert a new device that acts as a network sniffer.

Active attacks are possible for both on-path and off-path scenarios. For on-path active attacks, the situation is the same as for a passive attack, where either a device has to already be compromised or a device can be inserted into the path. For off-path active attacks, the attacks are generally limited to message insertion or modification where the attacker may wish to load illegal software or configuration files to an infrastructure device.

Note that similar issues are relevant when software updates are downloaded from a vendor site to an ISPs network management system that is responsible for software updates and/or configuration information.

2.5.1.1. Confidentiality Violations

Confidentiality violations can occur when a miscreant intercepts any of the software image or configuration information. The software image may give an indication of exploits which the device is vulnerable to while the configuration information can inadvertently lead attackers to identify critical infrastructure IP addresses and passwords.

2.5.1.2. Offline Cryptographic Attacks

If any cryptographic mechanism was used to provide for data integrity and confidentiality, an offline cryptographic attack could potentially compromise the data. The traffic would need to be captured either by eavesdropping on the communication path or by being able to divert traffic to a malicious user.

2.5.1.3. Replay Attacks

For a replay attack to be successful, the software image or configuration file would need to first be captured either on-path or diverted to an attacker to later be replayed to the intended recipient. Additionally, since many protocols do have replay protection capabilities, these would have to be subverted as well in applicable situations.

2.5.1.4. Message Insertion/Deletion/Modification

Software images and configuration files can be manipulated by someone in control of intermediate hosts. By forging an IP address and impersonating a valid host which can download software images or configuration files, invalid files can be downloaded to an infrastructure device. This can also be the case from trusted vendors who may unbeknownst to them have compromised trusted hosts. An invalid software image or configuration file can cause a device to

hang and become inoperable. Spoofed configuration files can be hard to detect, especially when the only added command is to allow a miscreant access to that device by entering a filter allowing a specific host access and configuring a local username/password database entry for authentication to that device.

2.5.1.5. Man-In-The-Middle

A man-in-the-middle attack attacks the identity of a communicating peer rather than the data stream itself. The attacker intercepts traffic that is sent between the infrastructure device and the host used to upload/download the system image or configuration file. He/she can then act on behalf of one or both of these systems.

If an attacker obtained a copy of the software image being deployed, he could potentially exploit a known vulnerability and gain access to the system. From a captured configuration file, he could obtain confidential network topology information, or even more damaging information, if any of the passwords in the configuration file were not encrypted.

2.5.2. Security Practices

Images and configurations are stored on specific hosts that have limited access. All access and activity relating to these hosts are authenticated and logged via AAA services. When uploaded/downloading any system software or configuration files, either TFTP, FTP, or SCP can be used. Where possible, SCP is used to secure the data transfer and FTP is generally never used. All SCP access is username/password authenticated but since this requires an interactive shell, most ISPs will use shared key authentication to avoid the interactive shell. While TFTP access does not have any security measures, it is still widely used, especially in OOB management scenarios. Some ISPs implement IP-based restriction on the TFTP server, while some custom written TFTP servers will support MAC-based authentication. The MAC-based authentication is more common when using TFTP to bootstrap routers remotely.

In most environments, scripts are used for maintaining the images and configurations of a large number of routers. To ensure the integrity of the configurations, every hour the configuration files are polled and compared to the previously polled version to find discrepancies. In at least one environment these, tools are Kerberized to take advantage of automated authentication (not confidentiality). 'Rancid' is one popular publicly available tool for detecting configuration and system changes.

Filters are used to limit access to uploading/downloading configuration files and system images to specific IP addresses and protocols.

The software images perform Cyclic Redundancy Checks (CRC) and the system binaries use the MD5 algorithm to validate integrity. Many ISPs expressed interest in having software image integrity validation based on the MD5 algorithm for enhanced security.

In all configuration files, most passwords are stored in an encrypted format. Note that the encryption techniques used in varying products can vary and that some weaker encryption schemes may be subject to off-line dictionary attacks. This includes passwords for user authentication, MD5-authentication shared secrets, AAA server shared secrets, NTP shared secrets, etc. For older software that may not support this functionality, configuration files may contain some passwords in readable format. Most ISPs mitigate any risk of password compromise by either storing these configuration files without the password lines or by requiring authenticated and authorized access to the configuration files that are stored on protected OOB management devices.

Automated security validation is performed on infrastructure devices using Network Mapping (Nmap) and Nessus to ensure valid configuration against many of the well-known attacks.

2.5.3. Security Services

- o User Authentication - All users are authenticated before being able to download/upload any system images or configuration files.
- o User Authorization - All authenticated users are granted specific privileges to download or upload system images and/or configuration files.
- o Data Origin Authentication - Filters are used to limit access to uploading/downloading configuration files and system images to specific IP addresses.
- o Access Control - Filters are used to limit access to uploading/downloading configuration files and system images to specific IP addresses and protocols.
- o Data Integrity - All systems use either a CRC-check or MD5 authentication to ensure data integrity. Also, tools such as rancid are used to automatically detect configuration changes.

- o Data Confidentiality - If the SCP protocol is used then there is confidentiality of the downloaded/uploaded configuration files and system images.
- o Auditing/Logging - All access and activity relating to downloading/uploading system images and configuration files are logged via AAA services and filter exception rules.
- o DoS Mitigation - A combination of filtering and CRC-check/MD5-based integrity checks are used to mitigate the risks of DoS attacks. If the software updates and configuration changes are performed via an OOB management system, this is also added protection.

2.5.4. Additional Considerations

Where the MD5 algorithm is not used to perform data-integrity checking of software images and configuration files, ISPs have expressed an interest in having this functionality. IPsec is considered too cumbersome and operationally difficult to use for data integrity and confidentiality.

2.6. Logging Considerations

Although logging is part of all the previous sections, it is important enough to be covered as a separate item. The main issues revolve around what gets logged, how long are logs kept, and what mechanisms are used to secure the logged information while it is in transit and while it is stored.

2.6.1. Threats/Attacks

Attacks on the logged data can be both from passive or active sources. Passive attacks are possible if someone has the capability to intercept data between the recipient logging server and the device from which the logged data originated. This can be accomplished if a single infrastructure device is somehow compromised and can act as a network sniffer, or if it is possible to insert a new device that acts as a network sniffer.

Active attacks are possible for both on-path and off-path scenarios. For on-path active attacks, the situation is the same as for a passive attack, where either a device has to already be compromised, or a device can be inserted into the path. For off-path active attacks, the attacks are generally limited to message insertion or modification that can alter the logged data to keep any compromise from being detected, or to destroy any evidence that could be used for criminal prosecution.

2.6.1.1. Confidentiality Violations

Confidentiality violations can occur when a miscreant intercepts any of the logging data that is in transit on the network. This could lead to privacy violations if some of the logged data has not been sanitized to disallow any data that could be a violation of privacy to be included in the logged data.

2.6.1.2. Offline Cryptographic Attacks

If any cryptographic mechanism was used to provide for data integrity and confidentiality, an offline cryptographic attack could potentially compromise the data. The traffic would need to be captured either by eavesdropping on the network or by being able to divert traffic to a malicious user.

2.6.1.3. Replay Attacks

For a replay attack to be successful, the logging data would need to first be captured either on-path or diverted to an attacker and later replayed to the recipient.

2.6.1.4. Message Insertion/Deletion/Modification

Logging data could be injected, deleted, or modified by someone in control of intermediate hosts. Logging data can also be injected by forging packets from either legitimate or illegitimate IP addresses.

2.6.1.5. Man-In-The-Middle

A man-in-the-middle attack attacks the identity of a communicating peer rather than the data stream itself. The attacker intercepts traffic that is sent between the infrastructure device and the logging server or traffic sent between the logging server and the database that is used to archive the logged data. Any unauthorized access to logging information could lead to the knowledge of private and proprietary network topology information, which could be used to compromise portions of the network. An additional concern is having access to logging information, which could be deleted or modified so as to cover any traces of a security breach.

2.6.2. Security Practices

When it comes to filtering, logging is mostly performed on an exception auditing basis (i.e., traffic that is NOT allowed is logged). This is to assure that the logging servers are not overwhelmed with data, which would render most logs unusable. Typically the data logged will contain the source and destination IP

addresses and layer 4 port numbers as well as a timestamp. The syslog protocol is used to transfer the logged data between the infrastructure device to the syslog server. Many ISPs use the OOB management network to transfer syslog data since there is virtually no security performed between the syslog server and the device. All ISPs have multiple syslog servers - some ISPs choose to use separate syslog servers for varying infrastructure devices (i.e., one syslog server for backbone routers, one syslog server for customer edge routers, etc.)

The timestamp is derived from NTP, which is generally configured as a flat hierarchy at stratum1 and stratum2 to have less configuration and less maintenance. Consistency of configuration and redundancy is the primary goal. Each router is configured with several stratum1 server sources, which are chosen to ensure that proper NTP time is available, even in the event of varying network outages.

In addition to logging filtering exceptions, the following is typically logged: routing protocol state changes, all device access (regardless of authentication success or failure), all commands issued to a device, all configuration changes, and all router events (boot-up/flaps).

The main function of any of these log messages is to see what the device is doing as well as to try and ascertain what certain malicious attackers are trying to do. Since syslog is an unreliable protocol, when routers boot or lose adjacencies, not all messages will get delivered to the remote syslog server. Some vendors may implement syslog buffering (e.g., buffer the messages until you have a route to the syslog destination), but this is not standard. Therefore, operators often have to look at local syslog information on a device (which typically has very little memory allocated to it) to make up for the fact that the server-based syslog files can be incomplete. Some ISPs also put in passive devices to see routing updates and withdrawals and do not rely solely on the device for log files. This provides a backup mechanism to see what is going on in the network in the event that a device may 'forget' to do syslog if the CPU is busy.

The logs from the various syslog server devices are generally transferred into databases at a set interval that can be anywhere from every 10 minutes to every hour. One ISP uses Rsync to push the data into a database, and then the information is sorted manually by someone SSH'ing to that database.

2.6.3. Security Services

- o User Authentication - Not applicable.
- o User Authorization - Not applicable.
- o Data Origin Authentication - Not implemented.
- o Access Control - Filtering on logging host and server IP address to ensure that syslog information only goes to specific syslog hosts.
- o Data Integrity - Not implemented.
- o Data Confidentiality - Not implemented.
- o Auditing/Logging - This entire section deals with logging.
- o DoS Mitigation - An OOB management system is used and sometimes different syslog servers are used for logging information from varying equipment. Exception logging tries to keep information to a minimum.

2.6.4. Additional Considerations

There is no security with syslog and ISPs are fully cognizant of this. IPsec is considered too operationally expensive and cumbersome to deploy. Syslog-ng and stunnel are being looked at for providing better authenticated and integrity-protected solutions. Mechanisms to prevent unauthorized personnel from tampering with logs is constrained to auditing who has access to the logging servers and files.

ISPs expressed requirements for more than just UDP syslog. Additionally, they would like more granular and flexible facilities and priorities, i.e., specific logs to specific servers. Also, a common format for reporting standard events so that modifying parsers after each upgrade of a vendor device or software is not necessary.

2.7. Filtering Considerations

Although filtering has been covered under many of the previous sections, this section will provide some more insights to the filtering considerations that are currently being taken into account. Filtering is now being categorized into three specific areas: data plane, management plane, and routing control plane.

2.7.1. Data Plane Filtering

Data plane filters control the traffic that traverses through a device and affects transit traffic. Most ISPs deploy these kinds of filters at customer facing edge devices to mitigate spoofing attacks using BCP38 and BCP84 guidelines.

2.7.2. Management Plane Filtering

Management filters control the traffic to and from a device. All of the protocols that are used for device management fall under this category and include: SSH, Telnet, SNMP, NTP, HTTP, DNS, TFTP, FTP, SCP, and Syslog. This type of traffic is often filtered per interface and is based on any combination of protocol, source and destination IP address, and source and destination port number. Some devices support functionality to apply management filters to the device rather than to the specific interfaces (e.g., receive ACL or loopback interface ACL), which is gaining wider acceptance. Note that logging the filtering rules can today place a burden on many systems and more granularity is often required to more specifically log the required exceptions.

Any services that are not specifically used are turned off.

IPv6 networks require the use of specific ICMP messages for proper protocol operation. Therefore, ICMP cannot be completely filtered to and from a device. Instead, granular ICMPv6 filtering is always deployed to allow for specific ICMPv6 types to be sourced or destined to a network device. A good guideline for IPv6 filtering is in the Recommendations for Filtering ICMPv6 Messages in Firewalls [ICMPv6].

2.7.3. Routing Control Plane Filtering

Routing filters are used to control the flow of routing information. In IPv6 networks, some providers are liberal in accepting /48s due to the still unresolved multihoming issues, while others filter at allocation boundaries, which are typically at /32. Any announcement received that is longer than a /48 for IPv6 routing and a /24 for IPv4 routing is filtered out of eBGP. Note that this is for non-customer traffic. Most ISPs will accept any agreed upon prefix length from its customer(s).

2.8. Denial-of-Service Tracking/Tracing

Denial-of-Service attacks are an ever-increasing problem and require vast amounts of resources to combat effectively. Some large ISPs do not concern themselves with attack streams that are less than 1G in bandwidth - this is on the larger pipes where 1G is essentially less

than 5% of an offered load. This is largely due to the large amounts of DoS traffic, which continually requires investigation and mitigation. At last count, the number of hosts making up large distributed DoS botnets exceeded 1 million hosts.

New techniques are continually evolving to automate the process of detecting DoS sources and mitigating any adverse effects as quickly as possible. At this time, ISPs are using a variety of mitigation techniques including: sinkhole routing, black hole triggered routing, uRPF, rate limiting, and specific control plane traffic enhancements. Each of these techniques will be detailed below.

2.8.1. Sinkhole Routing

Sinkhole routing refers to injecting a more specific route for any known attack traffic, which will ensure that the malicious traffic is redirected to a valid device or specific system where it can be analyzed.

2.8.2. Black Hole Triggered Routing

Black hole triggered routing (also referred to as Remote Triggered Black Hole Filtering) is a technique where the BGP routing protocol is used to propagate routes which in turn redirects attack traffic to the null interface where it is effectively dropped. This technique is often used in large routing infrastructures since BGP can propagate the information in a fast, effective manner, as opposed to using any packet-based filtering techniques on hundreds or thousands of routers (refer to the following NANOG presentation for a more complete description <http://www.nanog.org/mtg-0402/pdf/morrow.pdf>).

Note that this black-holing technique may actually fulfill the goal of the attacker if the goal was to instigate black-holing traffic that appeared to come from a certain site. On the other hand, this black hole technique can decrease the collateral damage caused by an overly large attack aimed at something other than critical services.

2.8.3. Unicast Reverse Path Forwarding

Unicast Reverse Path Forwarding (uRPF) is a mechanism for validating whether or not an incoming packet has a legitimate source address. It has two modes: strict mode and loose mode. In strict mode, uRPF checks whether the incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix. If the incoming packet fails the unicast RPF check, the packet is not accepted on the

incoming interface. Loose mode uRPF is not as specific and the incoming packet is accepted if there is any route in the routing table for the source address.

While BCP84 [RFC3704] and a study on uRPF experiences [BCP84-URPF] detail how asymmetry, i.e., multiple routes to the source of a packet, does not preclude applying feasible paths strict uRPF, it is generally not used on interfaces that are likely to have routing asymmetry. Usually for the larger ISPs, uRPF is placed at the customer edge of a network.

2.8.4. Rate Limiting

Rate limiting refers to allocating a specific amount of bandwidth or packets per second to specific traffic types. This technique is widely used to mitigate well-known protocol attacks such as the TCP-SYN attack, where a large number of resources get allocated for spoofed TCP traffic. Although this technique does not stop an attack, it can sometimes lessen the damage and impact on a specific service. However, it can also make the impact of a DoS attack much worse if the rate limiting is impacting (i.e., discarding) more legitimate traffic.

2.8.5. Specific Control Plane Traffic Enhancements

Some ISPs are starting to use capabilities that are available from some vendors to simplify the filtering and rate limiting of control traffic. Control traffic here refers to the routing control plane and management plane traffic that requires CPU cycles. A DoS attack against any control plane traffic can therefore be much more damaging to a critical device than other types of traffic. No consistent deployment of this capability was found at the time of this writing.

3. Security Considerations

This entire document deals with current security practices in large ISP environments. It lists specific practices used in today's environments and as such, does not in itself pose any security risk.

4. Acknowledgments

The editor gratefully acknowledges the contributions of: George Jones, who has been instrumental in providing guidance and direction for this document, and the insightful comments from Ross Callon, Ron Bonica, Ryan McDowell, Gaurab Upadhaya, Warren Kumari, Pekka Savola, Fernando Gont, Chris Morrow, Ted Seely, Donald Smith, and the numerous ISP operators who supplied the information that is depicted in this document.

5. References

5.1. Normative References

- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.
- [RFC2828] Shirey, R., "Internet Security Glossary", RFC 2828, May 2000.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.
- [RFC3682] Gill, V., Heasley, J., and D. Meyer, "The Generalized TTL Security Mechanism (GTSM)", RFC 3682, February 2004.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.
- [RFC3882] Turk, D., "Configuring BGP to Block Denial-of-Service Attacks", RFC 3882, September 2004.

5.2. Informational References

- [BCP84-URPF] Savola, P., "Experiences from Using Unicast RPF", Work in Progress, November 2006.
- [ICMPv6] Davies, E. and J. Mohacsi, "Recommendations for Filtering ICMPv6 Messages in Firewalls", Work in Progress, July 2006.
- [RTGWG] Savola, P., "Backbone Infrastructure Attacks and Protections", Work in Progress, July 2006.

Appendix A. Protocol Specific Attacks

This section will list many of the traditional protocol-based attacks that have been observed over the years to cause malformed packets and/or exploit protocol deficiencies. Note that they all exploit vulnerabilities in the actual protocol itself and often, additional authentication and auditing mechanisms are now used to detect and mitigate the impact of these attacks. The list is not exhaustive, but is a fraction of the representation of what types of attacks are possible for varying protocols.

A.1. Layer 2 Attacks

- o ARP Flooding

A.2. IPv4 Protocol-Based Attacks

- o IP Addresses, either source or destination, can be spoofed which in turn can circumvent established filtering rules.
- o IP Source Route Option can allow attackers to establish stealth TCP connections.
- o IP Record Route Option can disclose information about the topology of the network.
- o IP header that is too long or too short can cause DoS attacks to devices.
- o IP Timestamp Option can leak information that can be used to discern network behavior.
- o Fragmentation attacks which can vary widely - more detailed information can be found at <http://www-src.lip6.fr/homepages/Fabrice.Legond-Aubry/www.ouah.org/fragma.html>.
- o IP ToS field (or the Differentiated Services (DSCP) field) can be used to reroute or reclassify traffic based on specified precedence.
- o IP checksum field has been used for scanning purposes, for example when some firewalls did not check the checksum and allowed an attacker to differentiate when the response came from an end-system, and when from a firewall.
- o IP TTL field can be used to bypass certain network-based intrusion detection systems and to map network behavior.

A.2.1. Higher Layer Protocol Attacks

The following lists additional attacks, but does not explicitly numerate them in detail. It is for informational purposes only.

- o IGMP oversized packet
- o ICMP Source Quench
- o ICMP Mask Request
- o ICMP Large Packet (> 1472)
- o ICMP Oversized packet (>65536)
- o ICMP Flood
- o ICMP Broadcast w/ Spoofed Source (Smurf Attack)
- o ICMP Error Packet Flood
- o ICMP Spoofed Unreachable
- o TCP Packet without Flag
- o TCP Oversized Packet
- o TCP FIN bit with no ACK bit
- o TCP Packet with URG/OOB flag (Nuke Attack)
- o SYN Fragments
- o SYN Flood
- o SYN with IP Spoofing (Land Attack)
- o SYN and FIN bits set
- o TCP port scan attack
- o UDP spoofed broadcast echo (Fraggle Attack)
- o UDP attack on diagnostic ports (Pepsi Attack)

A.3. IPv6 Attacks

Any of the above-mentioned IPv4 attacks could be used in IPv6 networks with the exception of any fragmentation and broadcast traffic, which operate differently in IPv6. Note that all of these attacks are based on either spoofing or misusing any part of the protocol field(s).

Today, IPv6-enabled hosts are starting to be used to create IPv6 tunnels, which can effectively hide botnet and other malicious traffic if firewalls and network flow collection tools are not capable of detecting this traffic. The security measures used for protecting IPv6 infrastructures should be the same as in IPv4 networks, but with additional considerations for IPv6 network operations, which may be different from IPv4.

Author's Address

Merike Kaeo
Double Shot Security, Inc.
3518 Fremont Avenue North #363
Seattle, WA 98103
U.S.A.

Phone: +1 310 866 0165
EMail: merike@doubleshotsecurity.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.