                 Operational Security Current Practices
                 draft-ietf-opsec-current-practices-00

Status of this Memo

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

   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 Threat Model

   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.

1.2 Operational Security Impact from Threats

   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]. 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. 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 gateways. Such gateways 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 hosts 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]

Insider or outsider attacks
   An "insider attack" is one which 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 one 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 malice in mind. 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 what 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 be as easily detected and mitigated as with any
other attack source. Any of the specific attacks discussed further
in this document will elaborate on attacks 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.

The threat consequences are the security violations which 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 to do some
reconnaissance work and ascertain which IP addresses belonged 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.

1.3 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 which will provide the necessary security services
to help mitigate these attacks. These security services are
classified as:

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, AAA can offer user
authentication, user authorization and audit / logging services while
SSH can provide data origin authentication, data integrity and data
confidentiality. The services offered are more important than the
actual protocol used. Each section ends with an additional
considerations section which 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 ease
of use.

1.4 Definitions
RFC 2119 Keywords

> 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 [RFC2119].
> The use of the RFC 2119 keywords is an attempt, by the editor, to
> assign the correct requirement levels ("MUST", "SHOULD",
> "MAY"...). It must be noted that different organizations,
> operational environments, policies and legal environments will
> generate different requirement levels.

2. Protected Operational Functions

2.1 Device Physical Access

Device physical access pertains to protecting the physical location
of the layer 2 or layer 3 network infrastructure device. 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-scope for 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.

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 can use this device to crack the console password and have
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 take control of 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. There still
occur cases 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 dealt with.

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.

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. Individual users are authentication to get
basic access. For privileged (i.e. enable) access, a second
authentication step needs to be completed. Typically all console
access is provided via an out-of-band (OOB) management infrastructure

which is discussed in the section on OOB management.

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. Console access
      is usually granted via at least two privilege levels:
      authorization for performing a basic set of commands vs
      authorization for performing all commands.
   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 the OOB
      management section 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 the OOB management section of this document.

2.2 Device In-Band Management

   In-band management is generally considered to be device access where
   the control traffic takes the same data path as the data which
   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. Presently, the
   mechanisms used for in-band management are via virtual terminal
   access (i.e. Telnet or SSH), SNMP, or HTTP. In all large ISP
   environments, HTTP management is never used and is explicitly
   disabled. Note that file transfer protocols (TFTP, FTP, SCP) will
   be covered in the 'Software Upgrades and Configuration
   Integrity/Validation' section.

2.2.1 Threats / Attacks

   For in-band 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 which 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 attack is generally limited to message insertion or
   modification.

2.2.1.1 Confidentiality Violations

   Confidentiality violations can occur when a miscreant intercepts
   confidential data that has been sent in cleartext. 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, in-band 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 which 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 an in-band management system to the
networking infrastructure device and traffic that is sent from the
network infrastructure device to the in-band management system.

[2.2.2](#) Security Practices

All in-band management access to layer 2 and layer 3 devices is
authenticated. The user authentication and authorization is
typically controlled by a AAA server (i.e. RADIUS and/or TACACS+).
Credentials used to determine the identity of the user vary from
static username/password to one-time username/password scheme 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. In some
deployments, The AAA servers used for in-band management
authentication/authorization/accounting are on separate out-of-band
networks to provide a demarcation for any other authentication
functions.

For backup purposes, there is often a single local database entry for
authentication which 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.

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 in-band management, it is for read queries only
and restricted to specific hosts. 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 strings, a second new string is set and then
migrate over from the 1st to the 2nd. Most large ISPs have multiple
SNMP systems accessing their routers so it takes more then one
maintenance period to get all the strings fixed in all the right
systems. SNMP RW is not used and 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 which they are to limited to (i.e. SSH and
SNMP).

All in-band device management access is audited. The AAA server
keeps 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 which 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 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 from spoofed
      traffic. Using SSH for device access ensures that noone can spoof
      the traffic during the SSH session.
   o Access Control - In-band 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 noone 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 - Filtering to allow only specific IP addresses to
      have access to the infrastructure devices. This does not defend
      against spoofed traffic which may be used to source a DoS attack.

      Often OOB management is used to lower that risk.

2.2.4 Additional Considerations

   IPsec is considered too difficult to implement and the common
   protocol to provide for confidential in-band management access is
   SSH. There are exceptions for using SSH due to equipment limitations
   since SSH may not be supported on legacy equipment. 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') to manage those devices.
   An individual would first SSH to the jumphost and then Telnet from
   the jumphost to the actual infrastructure device. 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. Note that Telent is NEVER allowed to an
   infrastructure device except from specific jumphosts whose IP
   addresses are filtered at the infrastructure device.

   With thousands of devices to manage, some ISPs have created automated
   mechanisms to authenticate to devices. Kerberos is used to automate
   the authentication process. 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 Device Out-of-Band Management

Out-of-band management is generally considered to be device access
where the control traffic takes a separate path as the data which
traverses the network. Console access is always architected via an
OOB network. SNMP management is also usually carried out via that
same OOB network infrastructure.

2.3.1 Threats / Attacks

For OOB 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 which 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 attack is generally limited to message insertion or
modification.

2.3.1.1 Confidentiality Violations

Confidentiality violations can occur when a miscreant intercepts any
of the OOB management data that has been sent in cleartext. 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.

2.3.1.2 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 OOB 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.

2.3.1.3 Replay Attacks

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

2.3.1.4 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 which appear to come from another, trusted host.

2.3.1.5 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 an OOB management system to the networking
   infrastructure device and traffic that is sent from the network

   infrastructure device to the OOB management system.

2.3.2 Security Practices

   OOB 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.

   All OOB management access to layer 2 and layer 3 devices is
   authenticated. The user authentication and authorization is

typically controlled by a AAA server (i.e. RADIUS and/or TACACS+).
Credentials used to determine the identity of the user vary from
static username/password to one-time username/password scheme 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.

For backup purposes, there is often a single local database entry for
authentication which 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.

Some OOB 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 OOB management, it is for read queries only and
restricted to specific hosts. 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 strings, a second new string is set and then migrate
over from the 1st to the 2nd. Most large ISPs have multiple SNMP
systems accessing their routers so it takes more then one maintenance
period to get all the strings fixed in all the right systems. SNMP
RW is not used and 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 which they are to limited to (i.e. SSH and
SNMP).

All OOB device management access is audited. The AAA server keeps
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 which 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)

2.3.3 Security Services

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 from spoofed
   traffic. Using SSH for device access ensures that noone can spoof
   the traffic during the SSH session.
o Access Control – In-band 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 noone 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 – Filtering to allow only specific IP addresses to
   have access to the infrastructure devices. This does not defend
   against spoofed traffic which may be used to source a DoS attack.

   The risk is lowered by using a separate network for management
   purposes.

2.3.4 Additional Considerations

   IPsec is considered too difficult to implement and the common
   protocol to provide for confidential OOB management access is SSH.
   There are exceptions for using SSH due to equipment limitations
   since SSH may not be supported on legacy equipment. 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') to manage those device. An
   individual would first SSH to the jumphost and then Telnet from the
   jumphost to the terminal server before logging in to the device
   console. 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. Note that Telent is NEVER
   allowed to a console server or infrastructure device except from
   specific jumphosts whose IP addresses are filtered at the console
   server and/or infrastructure device.

   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.

2.4 Data Path

   This section refers to how traffic is handled which traverses the
   network infrastructure device. The primary goal of ISPs is to
   forward customer traffic. However, due to the large amount of attack
   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.

2.4.1 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.

2.4.2 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 the reliability of existing deployed hardware.

The ISPs which do not have performance issues with their equipment
follow BCP38 guidelines. Null routes and black-hole filtering are
used to deter any detected malicious traffic streams. Most ISPs
consider layer 4 filtering useful but it is only implemented if there
is no performance limitations on the devices. Netflow is used for
tracking traffic flows but there is some concern whether sampling is
good enough to detect malicious behavior.

Unicast RPF is not consistently implemented. Some ISPs are in
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 DS3 and below.

2.4.3 Security Services

   o User Authentication - Not applicable
   o User Authorization - Not applicable
   o Data Origin Authentication - When IP address filtering per BCP38
      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.
   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.

2.4.4 Additional Considerations

   For layer 2 devices, MAC address filtering and authentication is not
   used. This is due to the problems it can cause when troubleshooting

networking issues. Port security becomes unmanageable at a large
scale where 1000s 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. Rate limiting
can be improved by (need info)

Performance issues cause some ISPs to not implement BCP38 guidelines
for ingress filtering. One such example is at edge boxes where you
have up to 1000 T1's connecting into a router with an OC-12 uplink.
Some ISP's experience a large performance impact with filtering which
is unacceptable for passing customer traffic through. Where
performance is not an issue, the ISPs make a tradeoff between
management versus risk.

2.5 Routing Control Plane

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

2.5.1 Threats / Attacks

Attacks on the routing control plane can be both from 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 which 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 which can divert traffic to illegitimate destinations
and cause traffic to never reach its intended destination.

2.5.2 Confidentiality Violations

Confidentiality violations can occur when a miscreant intercepts any

of the routing update traffic. [is this an issue?]

2.5.3 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 eavesdrop and
capture the cryptographically protected routing information.

2.5.4 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.

2.5.5 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
which 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.

2.5.6 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.

2.5.7 Security Practices

Securing the routing control plane takes many features which are
generally deployed as a system. 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 do MD-5 authentication
at customer's request. Many ISPs also deploy sanity checks to ensure
with some certainty that the received routing update has been
authorized to be sent by the sending party. In the case of BGP
routing, this is accomplished through the use of routing registries
and prefix limits. Additionally, route filters and the ttl-hack
(politically correct name? BTSH?) ensure with reasonable probability
that the routing update came from a valid peer.

[editor's note: should more info be included on secure BGP policy?
Rejecting advertisements for your own backbone, advertisements to
bogons, route damping, rejecting selected attributes and communities,
etc. (Will need more specific input from provider deployment)]

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

2.5.9 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 MD-5 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
(ttl-hack and route filters).

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.

[2.6](#) 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.

[2.6.1](#) 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 which 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.

2.6.2 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.

2.6.3 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.

2.6.4 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.

2.6.5 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. 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.

2.6.6 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 and
acts on behalf of one or both of these systems.

[2.6.7](#) Security Practices

Images and configurations are stored on specific hosts which 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 TFTP and SCP access is
username/password authenticated and in most environments scripts are
used for maintaining 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).

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

The software images perform CRC-checks but many ISPs expressed
interest in having software image integrity validation based on the
MD5 algorithm for enhanced security. The system binaries use the MD5
algorithm to validate integrity.

In all configuration files, the passwords are stored in an obfuscated
format. This includes passwords for user authentication, MD5 shared
secrets, AAA server shared secrets, NTP shared secrets, etc. For
older software which may not support this functionality,
configuration files are stored with passwords in readable format.
[is this true? are configuration files then protected? if passwords
in readable format, is the thought that an OOB management network
with SCP will be enough protection?]

Automated security validation is performed on infrastructure devices
using nmap and nessus to ensure valid configuration against many of

the well-known attacks.

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 - TBD
   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.
   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 - TBD

Additional Considerations

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

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.

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
   the logged data originated from. 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 which
   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 which can alter the logged data to keep any compromise
from being detected or to destroy any evidence which could be used
for criminal prosecution.

2.7.1.1 Confidentiality Violations

Confidentiality violations can occur when a miscreant intercepts any
of the logging data which 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.

2.7.1.2 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.

2.7.1.3 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. [is reply handled by syslog protocol?]

2.7.1.4 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.

2.7.1.5 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 which is used to archive the logged data.

2.7.2 Security Practices

Logging is mostly performed on an exception auditing basis when it

comes to filtering (i.e. traffic which is NOT allowed is logged).
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. Each router is configured with one stratum1
peer both locally and remotely.

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. Some ISPs put in passive
devices to see routing updates and withdrawals and 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 which 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.

2.7.3 Security Services

   o User Authentication - Not applicable
   o User Authorization - Not applicable
   o Data Origin Authentication - TBD
   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 - TBD
   o DoS Mitigation - TBD

2.7.4 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.

   ISPs expressed requirements for more than just UDP syslog. They
   would also like more granular and flexible facilities and priorities,
   i.e. specific logs to specific servers.

2.8 Filtering Considerations

   [Editor note: Already covered but could be a sum up. This section to
   be added if enough proponents for it]

2.8.1 Threats / Attacks

   To be added.

2.8.2 Security Mechanisms

   To be added.

2.8.3 Security Services

To be added.

    o TBD
    o TBD

2.8.4 Additional Considerations

    TBD.

2.9 Denial of Service Tracking / Tracing

    [special section for describing security techniques to avoid well
    known attacks? Includes sink hole routing, black-hole filtering,
    uRPF, rate limiting]

2.9.1 Threats / Attacks

    To be added.

2.9.2 Security Mechanisms

    To be added.

2.9.3 Security Services

    To be added.

    o TBD
    o TBD

2.9.4 Additional Considerations

    TBD

3. Security Considerations

   This entire document deals with current security practices in large
   ISP environments. As a synopsis, a table is shown below which
   summarizes the operational functions which are to be protected and
   the security services which currently deployed security practices
   offer: [Table to be added]

4. Normative References

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

   [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.

Author's Address

   Merike Kaeo
   Double Shot Security, Inc.
   520 Washington Blvd. #363
   Marina Del Rey, CA 90292
   U.S.A.

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

Appendix A. Acknowledgments

   The editor gratefully acknowledges the contributions of:

   o George Jones, who has been instrumental in providing guidance and
      direction for this document.
   o To be named

o To be named

This listing is intended to acknowledge contributions, not to imply
that the individual or organizations approve the content of this
document.

Intellectual Property Statement