None.
Internet-Draft
Expires: April 14, 2004

G. Jones, Editor
The MITRE Corporation
October 15, 2003

Operational Security Requirements for IP Network Infrastructure: Advanced Requirements draft-jones-opsec-info-00

Status of this Memo

Copyright Notice

Abstract

This document defines a list of operational security requirements for the infrastructure of large IP networks (such as routers and switches). The goals of this document are to serve as a collection of ideas for security features that would improve operational security and to assist consumers of network equipment in communicating their security requirements to vendors. The requirements in this document are NOT considered to be best current practice (BCP). Comments to: "opsec-comment@ops.ietf.org".

Table of Contents

1. Introduction

1.1 Goals

The goals of this document are to serve as a collection of ideas for security features that would improve operational security and to assist consumers of network equipment in communicating their security requirements to vendors.

1.2 Scope

The primary scope of these requirements is intended to cover the infrastructure of large IP networks (e.g. routers and switches).

General purpose hosts (including infrastructure hosts such as name/time/log/AA servers, etc.), unmanaged, or customer managed devices (e.g. firewalls, Intrusion Detection System, dedicated VPN devices, etc.) are explicitly out of scope.

Confidentiality and integrity of customer data are outside the scope.

While, the examples given are written with IPv4 in mind, most of the requirements are general enough to apply to IPv6.

1.3 Definition of a Secure Network

For the purposes of this document, a secure network is one in which:

- the network keeps passing legitimate customer traffic (availability)
- traffic goes where it's supposed to go (availability)
- the network elements remain manageable (availability)
- only authorized users can manage network elements (authorization)
- there is record of all security related events (accountability)
- the network operator has the necessary tools to detect and respond to illegitimate traffic

The following assumptions are made:

- Devices are physically secure.
- The management infrastructure (AAA/DNS/log server, SNMP management stations, etc.) is secure.

1.4 Intended Audience

There are two intended audiences: the end user (consumer) who selects, purchases, and operates IP network equipment, and the vendors who create them.

1.5 Format

The individual requirements are listed in one of the three sections listed below.

- Section 2 lists functional requirements.
- Section 3 lists documentation requirements.
- Section 4 lists assurance requirements.

Within these areas, requirements are grouped in major functional areas (e.g., logging, authentication, filtering, etc.)

Each requirement has the following subsections:

- The Requirement (What)
- The Justification (Why)
- Examples (How)
- Warnings (if applicable)

The requirement describes a policy to be supported by the device. The justification tells why and in what context the requirement is important. The examples section is intended to give examples of implementations that may meet the requirement. Examples cite technology and standards current at the time of this writing. It is expected that the choice of implementations to meet the requirements will change over time. The warnings list operational concerns, deviation from standards, caveats, etc.

Security requirements will vary across different device types and different organizations, depending on policy and other factors. A desired feature in one environment may be a requirement in another. Classifications must be made according to local need.

1.6 Intended Use

It is anticipated that this document will be used in the following manners:

Documenting Useful, non-BCP Features This document is a collection of security features that would be useful in improving operational security. The features listed herein are not considered to best current practice (BCP) at this time. It is anticipated that the features listed here may, over time, become widely implemented and thus be candidates for migration to a BCP document.

Security Capability Checklist The requirements in this document may be used as a checklist when evaluating networked products.

Communicating Requirements This document may be referenced, to clearly communicate security requirements.

Basis For Testing and Certification This document may form the basis for testing and certification of security features of networked products.

1.7 Definitions

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].

Unless otherwise indicated, "IP" refers to IPv4

2. Functional Requirements

The requirements in this section are intended to list testable, functional requirements that are needed to operate devices securely.

2.1 Device Management Requirements

2.1.1 Restrict Management to Local Interfaces

Requirement. The device MUST have the ability to restrict management traffic to sources within one hop of the device in cases where management done over IP.

Justification. Restricting management traffic to devices attached locally reduces the risk of unauthorized configuration of the device from across the Internet.

This requirement applies primarily to SOHO equipment, where out-of-band management may not be feasible, and additional security for management traffic is most effectively applied by restricting it to local only.

Examples. This requirement MAY be satisfied by reducing the TTL on return TCP management traffic to 1, or by filtering all traffic to the management service not sourced from local subnets. See [I-D.gill-gtsh]

Warnings. None.

2.2 In-Band Management Requirements

This section lists security requirements for devices that are managed In-band. "In-band management" is defined as any management done over the same channels and interfaces used for user/customer data. In-band management has the advantage of lower cost (no extra interfaces or lines), but has significant security disadvantages:

- saturation of customer lines or interfaces can make the device unmanageable
- since public interfaces/channels are used, it is possible for attackers to directly address and reach the device and to attempt management functions
- in-band management traffic on public interfaces may be intercepted
- Since the same networking code and interfaces are shared for

management and customer data, it is not possible to isolate management functions from failures in other areas (for example, a "magic packet" or buffer overrun that causes the data forwarding portions of a router to crash will also likely make it impossible to manage...this would not necessarily be the case if the management and data forwarding elements were completely separated)

2.2.1 Key Management Must Be Scalable

Requirement. The number of keys and passwords that must be managed to support other requirements in this document MUST scale well. Specifically, The number of keys and passwords managed MUST increase, at most, linearly as the number of devices and users.

Justification. In large networks, or in networks with a large number of users, the key/password space could quickly grow to unmanageable size.

Examples. The use of AAA protocols such as RADIUS or the use of Kerberos greatly increases the manageability of keys and passwords however, someone still needs to configure the databases and periodically ensure that the databases have not become compromised. The use of a Public Key Infrastructure (PKI), which utilizes digital certificates to automate the secure distribution of passwords and keys, should be a consideration for networks with a large set of keys/passwords to manage.

Warnings. None.

2.3 Out-of-Band (OoB) Management Requirements

See Section 2.2 for a discussion of the advantages and disadvantages of In-band vs. Out-of-Band management.

2.3.1 Enforce Separation of Data and Management Planes

Requirement. The device MUST support separation of data and management plane. It MUST support complete physical and logical separation of management and non-management traffic.

Justification. Separation of management and data plane enables the application of separate and appropriate controls to each channel, and reduces the possibility that a vulnerability in one area/ environment (data forwarding) could have an adverse impact on another area (control/management). For example, imagine that a "killer packet" or buffer overrun is discovered that allows

arbitrary users of a public network to crash the data forwarding elements of a router. If data forwarding and management elements are separated, it is likely that the management elements will continue to function, allowing the network operator to evaluate and respond to the problem. If they are not separated (e.g., they both use the same interfaces and share an operating system and IP stack), then it is likely that the entire device will crash or become unmanageable.

Examples. One way to satisfy this requirement would be to do all of the following

* Implement management and forwarding planes using separate Operating Systems and IP stacks.

* Do not allow forwarding between management and data planes.

* Disable (or do not implement) all management functions (e.g., telnet, FTP, TFTP, SSH, SNMP, HTTP, etc.) on the data plane.

Warnings. None.

2.4 User Interface Requirements

2.4.1 Display All Configuration Settings

Requirement. The device MUST provide a mechanism to display a complete listing of all possible configuration settings and their current values. This MUST include values for any "hidden" commands. It MUST be possible to display all values, even those that are disabled, "off," or set to default values.

Justification. It is not possible to perform thorough audits without a complete listing of all possible configuration settings and their current values.

Examples. Sometimes default settings change between releases, for example an older release of software may enable directed broadcasts by default while the newer one disables it. If the device only displays non-default settings, then the customer/ auditor must keep a list of software versions and default settings in order to insure that device configuration complies with local policy (e.g. "directed broadcasts must be disabled"). The task of auditing for policy compliance is made much simpler if there is a way to display *all* settings, default or otherwise.

Warnings. It has been stated that it may be unreasonable to expect vendors to expose all settings, as this would lead to confusion due to customers changing settings that did not apply to their situation, and could drive up support costs.

2.5 IP Stack Requirements

2.5.1 Ability to Disable Processing of Packets Utilizing IP Options

Requirement. The device MUST provide a means to disable processing of all packets utilizing IP Options. This option MUST be available on a per-interface basis. It MUST be possible to individually configure which options are processed. Source routing SHOULD be disabled by default.

Justification. Options can be used to alter normal traffic flows and thus circumvent network-based access control mechanisms (such as firewalls). They can also be used to provide information (such as routes taken) that could be useful to an attacker mapping a network.

Examples. None.

Warnings. RFC791 says "The Options provide for control functions needed or useful in some situations but unnecessary for the most common communications... [options] must be implemented by all IP modules (host and gateways). What is optional is their transmission in any particular datagram, not their implementation"

2.5.2 Support Denial-Of-Service (DoS) Tracking

Requirement. The device MUST include native "spoofed" packet tracking. This feature:

* MUST be able to capture data to a tracking table that shows how many packets match a configurable layer 3/4 header pattern or list of patterns from each previous hop router.

* MUST display the interface on which a matching packet arrived.

* MUST display the layer-2 header information.

* MUST implement "unknown source" as an optional part of the header pattern where "unknown" is the set of all addresses that are unreachable by the router (i.e., not in the forwarding table).

* MUST be able to display the tracking table showing the pattern that is being tracked and how many matches were received from each previous hop.

This feature MUST be implemented with minimal impact to system performance.

Justification. This applies in situations where DoS attacks, possibly utilizing spoofed source addresses, must be tracked across one or more routers. Without the capability to track DoS packets, it is possible that an attacker could adversely impact the availability of resources (hosts, routers, network links, etc.) leaving network administrators little to no capability to track and stop the attack. Layer 2 header information is particularly useful for identifying spoofed sources coming in over an Ethernet interface at a peering point and you want to track the source back to a particular ISP so you can ask them to trace the source.

Examples.

These features must allow the customer to quickly and easily ask the router which packets matching a given profile came into the router, from where, and how many from each source.

Warnings. None.

2.5.3 Traffic Monitoring

Requirement. The device MUST provide a means to monitor selected traffic through the system. It MUST provide the ability to select specific traffic patterns for monitoring based on arbitrary IP header patterns and layer 4 (TCP and UDP) header patterns. This includes: source and destination IP address, IP header flags, layer 4 source and destination ports (TCP, UDP), ICMP type and code fields, and other IP protocol types (e.g., 50 - ESP, 47 - GRE, etc.). It MUST provide the ability to monitor the full contents of the packets. This feature MUST be implemented with minimal impact on system performance. In addition, the device MUST provide a means to remotely capture the data being monitored.

Justification. This requirement applies in contexts where traffic headers and content must be monitored. This enables characterization of malicious (and non-malicious) traffic, which may be essential to enable effective response and maintain normal operations.

Examples.

The addition of any traffic monitoring facility must be implemented with minimal impact on system performance.

Remote capture of header data could be implemented by sending it via syslog or SNMP. For the full packet capture, the device may send this information over the network for small data streams, or provide a "port mirroring" capability for large data streams where the data would be duplicated out a second configurable port.

Warnings. Monitoring data can add significant network traffic, processor, and memory use.

2.5.4 Traffic Sampling

NOTE: there is a proposed IETF working group active in this area. See the mailing list archives at https://ops.ietf.org/lists/psamp/. It is possible this section may just reference the product of that working group.

Requirement. The device MUST provide a means to sample traffic through the system and summarize data from the layer 3 and 4 headers.

It MUST be possible to dump the cache at specified intervals to a collection host. It MUST be possible to specify device behavior when the cache is full. Options SHOULD include: dumping the cache to the specified collection host(s), clearing the cache, overwriting the cache, and disabling further sampling. The cache SHOULD be implemented as a circular buffer such that older entries are overwritten first. The device SHOULD provide options to manually dump or clear the cache.

The device SHOULD provide a means of summarizing sampled data. The following IP layer header information SHOULD be summarized appropriately: type of service (or DS field), total length, protocol, source, and destination. The following TCP/UDP header information SHOULD be summarized appropriately: source port, destination port, UDP packet length, TCP header length, and TCP flag bits.

The device MUST provide the ability to select the traffic-sampling rate. For instance, there MUST be a way to sample every nth packet, where n is a number determined by an authorized user and entered into the system configuration file. This feature must be implemented with minimal impact on system performance.

Justification. This requirement enables accurate characterization of data transiting the device. This supports identification of and response to malicious traffic.

Examples. This requirement MAY be satisfied by allowing the user to specify that 1 in every N packets should be sampled.

Warnings. Traffic sampling can add significant network traffic, processor, and memory use.

2.5.5 Ability To Remove In-Band Visibility

Requirement. The device MUST provide a mechanism to allow it to become a "black box" as seen from public interfaces. Specifically this means:

* The device SHOULD not accept any packets beyond those required to support routing information transfer.

* The device SHOULD NOT generate any packets beyond those required to support routing information transfer. This includes ICMP error messages.

While the default configuration of the device SHOULD be fully RFC compliant (including the sending of ICMP messages), it MUST be possible to alter the default configuration such that the device is "stealthed" (i.e., does not send ICMP messages or otherwise respond directly to packets directed to it on non-management interfaces).

Justification. This applies to devices comprising the core network infrastructure. This enforces out of band only access, and ensures that risk to the core infrastructure from end users is minimized.

Examples. Some specific capabilities important to stealthing include:

* Ability to filter/deny/ignore pings (ICMP echo requests)

* Ability to filter on individual protocol header bits

* Ability to control the generation of ICMP messages, including port unreachable and timeouts

It MUST be possible to configure each of these settings individually.

Warnings. Although some STEALTHING MECHANISMS MAY BE IN VIOLATION OF SOME RFCs, they are desirable/necessary in certain circumstances for security and operational reasons.

2.6 Basic Filtering Capabilities

2.6.1 Ability to Filter Without Performance Degradation

Requirement. The device MUST provide a means to filter packets without performance degradation. The device MUST be able to filter on ALL interfaces (up to the maximum number possible) simultaneously and with multiple filters per interface (e.g., inbound and outbound).

Justification. This is important because it enables the implementation of filtering wherever and whenever needed. To the extent that filtering causes degradation, it may not be possible to apply filters that implement the appropriate policies.

Examples. Another way of stating the requirement is that filter performance should not be the limiting factor in device throughput. If a device is capable of forwarding, say, 30Mb/sec without filtering, then it should be able to forward the same amount with filtering in place. This requirement most likely implies a hardware-based solution (ASIC).

Warnings. Without hardware based filtering, it may be possible for the implementation of filters to degrade the performance of the device or to cause it to cease functioning.

2.7 Packet Filtering Criteria

2.7.1 Ability to Filter on Layer 2 MAC Addresses

Requirement. Filters in layer 2 devices MUST be able to filter based on Media Access Control (MAC) addresses.

Justification. This provides a level of control that may be needed to enforce policy and respond to malicious activity.

Examples. Policy may require, for example, that personal systems not be allowed to connect to the internal desktop network. Restricting the MAC addresses on a port is one way of enforcing this.

Warnings. None.

2.8 Event Logging Requirements

2.8.1 Ability to Log All Security Related Events

Requirement. The logging facility MUST be capable of logging any event that affects system security.

Justification. Having the device log all events that might impact system security promotes accountability and enables audit-ability.

Examples.

The list of items that must be logged includes, but is not limited to, the following events:

* Filter matches."
* Authentication failures (e.g., bad login attempts)
* Authentication successes (e.g., user logins)
* Authorization changes (e.g., User privilege level changes)
* Configuration changes (e.g., command accounting)
* Device status changes (interface up/down, etc.)

Warnings. None.

2.8.2 Ability to Select Reliable Delivery

Requirement. It MUST be possible to select reliable, sequenced delivery of log messages. .

Justification. Reliable delivery is important to the extent that log data is depended upon to make operational decisions and forensic analysis. Without reliable delivery, log data becomes a collection of hints.

Examples. One example of reliable syslog delivery is defined in [RFC3195]. Syslog-ng provides another example, although the protocol has not been standardized.

Warnings. None.

2.8.3 Ability to Classify Events

Requirement. The device SHOULD provide a mechanism for assigning classifications to all messages. At a minimum, it MUST provide the ability to assign a chosen classification to all security related messages, and different classification(s) to all other messages.

Justification. This is important because it allows messages of certain types to be sent to different servers for processing. This is important in environments with large numbers of devices, large numbers of log messages, and/or where responsibilities for certain classes of messages are divided.

Examples. This requirement MAY be satisfied by providing a mechanism to assign specific syslog facility codes to specific messages or groups of messages. For example, all security events could be assigned to one facility code, all network routing issues to another, and all physical (power, line card) to another.

Warnings. None.

2.8.4 Logs Do Not Contain DNS Names by Default

Requirement. By default, log messages MUST NOT contain DNS names resolved at the time the message was generated. The device MAY provide a facility to incorporate translated DNS names in addition to the IP address.

Justification. This is important because IP to DNS mappings change over time and mappings done at one point in time may not be valid later. Also, the use of the resources (memory, processor, time, bandwidth) required to do the translation could result in *no* data being sent/logged, and, in the extreme case could lead to degraded performance and/or resource exhaustion.

Examples. None.

Warnings. DNS name translation can impose significant performance delays.

2.9 Authentication, Authorization, and Accounting (AAA) Requirements

2.9.1 Enforce Selection of Strong Local Static Authentication Tokens (Passwords)

Requirement. Strength checks for static passwords fall into three types:

1. computational checks against the password itself (length, character set, upper/lower case)

2. comparison checks against static data sets (dictionary tests)

3. comparison checks against dynamic data sets (history checks, username tests)

The device MUST support at least computational checks with the following minimum requirements: The password MUST be at least [6] characters long and MUST contain at least [3] of the following elements

* At least [1] Lower case alphabetic character

* At least [1] Upper case alphabetic character

* At least [1] Numeric character

* At least [1] Special character

The device MAY enforce the selection of "strong" local passwords through comparison checks against dynamic and/or static data sets.

Justification. Trivial passwords are easily guessed, increasing the likelihood of unauthorized access.

Examples. An initial configuration dialog may require the user to set a password to control initial access. If the user enters a password that is not strong (e.g. "123") then the configuration dialog should inform the user that the chosen password is weak and provide another opportunity to select a strong password.

Warnings.

2.9.2 Support Device-to-Device Authentication

Requirement. The device MUST support device-to-device authentication for all non-interactive management protocols.

Justification. This is required to allow automated management functions to operate with a reasonable level assurance that updates and sharing of management information is occurring only with authorized devices.

Examples. Examples of protocols that implement device to device authentication are: SNMP (community strings), NTP and BGP (shared keys).

Warnings. None.

2.10 Layer 2 Requirements

2.10.1 Filtering MPLS LSRs

Requirement. The device MUST provide a method to filter packets based on layer 3 and 4 criteria on Label Switch Routers (LSRs) regardless of whether they are encapsulated using Multi Protocol Label Switching (MPLS). The MPLS encapsulated packets MUST NOT be allowed to bypass IP filters. Logging facilities MUST provide sufficient information so that the previous hop for a logged packet can be determined. Packets tagged with MPLS labels MUST be treated as IP packets when crossing an interface on which a filter is applied. Encapsulation/decapsulation MAY take place before or after the filter as long as it does not cause the filters to be ignored. When logging the input interface information for hits on outgoing filter list rules, any MPLS label that was present when the packet was received MUST be logged with the input interface. This functionality is equivalent to the requirement that all layer 2 source information must be logged when the input interface is logged. Also, the addition of any filtering and logging MUST be implemented with no significant performance degradation to the normal system operations.

Justification. This is important because it may be necessary to filter traffic encapsulated in a LSP. This applies primarily to backbone and large core networks.

Examples. None.

Warnings. None.

2.10.2 VLAN Isolation

Requirement. The device MUST NOT allow VLAN Hopping. This applies to the insertion of falsified VLAN IDs or 802.1Q (or equivalent) tags into frames in an attempt to hop from one VLAN to another while traversing the switch. Many VLAN implementations allow hopping if the native VLAN (usually VLAN 1) is set up as the trunk port. If this is the case then the default configuration on the switch MUST NOT allow the trunk port to be set as the native VLAN. Also the switch MUST NOT broadcast ARP requests across VLANs.

Justification. This requirement is intended to ensure that layer 2 traffic remains isolated to designated VLANs. It applies in situations where data on different VLAN segments have different sensitivity classification.

Examples. None.

Warnings. None.

2.10.3 Layer 2 Denial-of-Service

Requirement. It MUST NOT be possible for users connected to a switch port to perform an action which results in denial of service to other users connected to the switch. Examples of denial of service would include:

* Causing the switch to crash

* Causing long delays (e.g., by forcing spanning tree recalculations)

* Redirecting/stealing traffic

Justification. This requirement is needed to ensure the confidentiality and availability of data transmitted via the switch.

Examples. None.

Warnings. None.

3. Documentation Requirements

The requirements in this section are intended to list information that will assist operators in evaluating and securely operating a device.

3.1 Provide a List of All Protocols Implemented

Requirement. The vendor SHOULD provide a concise list all protocols implemented by the device.

Justification. This facilitates thorough and appropriately targeted testing.

Examples. The documentation should contain a concise list in the system/release documentation describing the protocols implemented (link,network,transport,management, routing, etc.)

Warnings. None.

3.2 Provide Documentation for All Protocols Implemented

Requirement. The vendor SHOULD provide references to publicly available specifications for all protocols implemented.

Justification. Security thorough obscurity is bad policy. Closed, undocumented protocols that have not undergone through public review may contain undiscovered (by the vendor) vulnerabilities that can easily be exploited. Open, documented protocols facilitate thorough and appropriately targeted testing.

Examples. None.

Warnings. It is acknowledged that there may be valid business or other non-technical reasons for not releasing documentation for protocols. This requirement should be evaluated on a case-by-case basis.

3.3 Catalog of Log Messages Available

Requirement. The vendor SHOULD specify a catalog of all messages that a device can emit. This SHOULD be included with every release of software for the device. The contents of variable portions of each message (IP address, hostname, timestamp, etc.) SHOULD be documented.

Justification. A complete catalog of all possible messages permits the customer to automate response to possible events.

Examples. If the device sends syslog messages, then the documentation should contain a list of all possible syslog messages.

Warnings. None.

4. Assurance Requirements

The requirements in this section are intended to

- identify behaviors and information that will increase confidence that the device will meet the security functional requirements.

- Provide information that will assist evaluation

4.1 Ability to Withstand Well-Known Attacks and Exploits

Requirement. The vendor MUST provide software updates or configuration advice "in a timely fashion" to mitigate the effect of "well know vulnerabilities" in the device itself and "well known exploits" directed to the device. These updates or configuration changes MUST NOT result in a reduced feature set - except in cases where removing a feature entirely is the ONLY way to stop the exploit. Updates SHOULD NOT introduce new features. Vendors MUST NOT require customers to pay a fee or purchase support (or other) contracts in order to obtain exploit fixes. These requirements only apply to devices that are supported that the time the exploit or vulnerability becomes "well known".

For the purpose of this document, well-known vulnerabilities and exploits are defined as those that have been published by the following:

* Computer Emergency Response Team Coordination Center [CERT/CC] Advisories

* Common Vulnerabilities and Exposures [CVE] entries

* Standard Nessus [Nessus] Plugins

* Vendor security bulletins for the device in question.

* The [PROTOS] test suite

While "in a timely fashion" is open to interpretation, one measurable, customer-centric metric is "before the vulnerability is exploited in my device causing loss of confidentiality, integrity or availability".

Justification. Product vulnerabilities and tools to exploit vulnerabilities are all constantly evolving. A configuration that is secure one day may be insecure the next due to the discovery of a new vulnerability or the release of a new exploit script.

Devices that are vulnerable to known exploits may be easily compromised or disabled. This can affect confidentiality, availability, and data integrity.

Examples. Take for example the SNMP vulnerabilities described in [CERT.2002-03]. These vulnerabilities were discovered and a toolkit for exploiting them was publicly released. What this requirement is saying is that known vulnerabilities such as this should be fixed.

It is up to the customer/operator to verify to their satisfaction that the system is "bug free" and free of known exploits. Some possible methods of doing this include

* Taking the vendors word

* Testing for themselves

* Relying on 3rd party testing/certification

Warnings. It is acknowledged that the number of known vulnerabilities is constantly expanding and that it is not possible to prove that any system is completely bug and vulnerability free. Any test or "certification" of a device to show compliance with this requirement will be an approximation at a point in time. The most that can be shown is that a given list of exploits failed.

4.2 Vendor Responsiveness

Requirement. The vendor MUST be responsive to current and future security requirements as specified by the customer. When new security exploits are discovered, either by the customer or the public, the vendor MUST provide patches or workarounds in a timely fashion to mitigate the threat from any existing vulnerability in the system. The vendor MUST ensure that it remains actively aware of security threats.

Justification. This is important because new vulnerabilities are regularly discovered. Slow vendor response to vulnerabilities increase the level of risk/window of opportunity for exploit. This requirement applies to ALL devices.

Examples. This is a non-technical requirement. The implementation involves process, customer support, engineering, etc.

Warnings. This "requirement" has a large element of subjectivity. When evaluating vendor responsiveness, objective data (such as mean time to releasing patches for new exploits) should be evaluated.

5. Security Considerations

Security is the subject matter of this entire memo. It might be more appropriate to list operational considerations. Operational issues are mentioned as needed in the examples and warnings sections of each requirement.

References

[CERT.2002-03] CERT/CC, "Multiple Vulnerabilities in Many Implementations of the Simple Network Management Protocol (SNMP)", 2002, <http://www.cert.org/advisories/CA-2002-03.html>.

[CERT/CC] CERT/CC, "CERT/CC Advisories", 2003, <http://www.cert.org/advisories/>.

[CVE] The MITRE Corporation, "MITRE Common Vulnerabilities and Exposures", 2003, <http://www.cve.mitre.org>.

[I-D.gill-gtsh] Gill, V., Heasley, J. and D. Meyer, "The Generalized TTL Security Mechanism (GTSM)", draft-gill-gtsh-04 (work in progress), October 2003.

[Nessus] Deraison, R., "Nessus Security Scanner", 2003, <http://www.nessus.org>.

[PROTOS] University of Oulu, "PROTOS Test Suites", 2003, <http://www.ee.oulu.fi/research/ouspg/protos/>.

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

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

[RFC3195] New, D. and M. Rose, "Reliable Delivery for syslog", RFC 3195, November 2001.

Author's Address

George M. Jones, Editor
The MITRE Corporation
7525 Colshire Dr., WEST
McLean, VA 22102
U.S.A.

Phone: +1 703 488 9740
EMail: gmjones@mitre.org
URI: http://www.port111.com/opsec/

Appendix A. Acknowledgments

This document grew out of an internal security requirements document used by UUNET for testing devices that were being proposed for connection to the backbone.

The editor gratefully acknowledges the contributions of:

- Greg Sayadian, author of a predecessor of this document.
- Eric Brandwine, a major source of ideas/critiques.
- The MITRE Corporation for supporting continued development of this document. NOTE: The editor's affiliation with The MITRE Corporation is provided for identification purposes only, and is not intended to convey or imply MITRE's concurrence with, or support for, the positions, opinions or viewpoints expressed by the editor.
- UUNET's entire network security team (past and present): Jared Allison, Eric Brandwine, Clarissa Cook, Dave Garn, Tae Kim, Kent King, Neil Kirr, Mark Krause, Michael Lamoureux, Maureen Lee, Todd MacDermid, Chris Morrow, Alan Pitts, Greg Sayadian, Bruce Snow, Robert Stone, Anne Williams, Pete White.
- Others who have provided significant feedback at various stages of the life of this document are: Ran Atkinson, Fred Baker, Steve Bellovin, Michael H. Behringer, Matt Bishop, Scott Blake, Randy Bush, Steven Christey, Sean Donelan, Robert Elmore, Barry Greene, Dan Hollis, Merike Kaeo, John Kristoff, Chris Liljenstolpe, James W. Laferriere, Alan Paller, Rob Pickering, Gregg Schudel, Rodney Thayer, David Walters, Anthony Williams, Neal Ziring
- Madge B. Harrison, technical writing review.
- This listing is intended to acknowledge contributions, not to imply that the individual or organizations approve the content of this document.
- Apologies to those who commented on/contributed to the document and were not listed...contact the editor to be credited in future versions

Version: $Id: draft-jones-opsec-01.cpp,v 1.2 2003/08/13 17:45:25 george Exp $

Intellectual Property Statement

Full Copyright Statement

HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgment

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