OPSEC Working Group
Internet-Draft
Expires: April 21, 2005

G. Jones
The MITRE Corporation
R. Callon
Juniper Networks
M. Kaeo
Double Shot Security
October 21, 2004

Framework for Operational Security Capabilities for IP Network Infrastructure
draft-jones-opsec-framework-01

Status of this Memo

Copyright Notice

Abstract

This document outlines work to be done and documents to be produced by the Operational Security Capabilities (OPSEC) Working Group. The goal of the working group is to codify knowledge gained through

operational experience about feature sets that are needed to securely deploy and operate managed network elements providing transit services at the data link and IP layers. The intent is to provide clear, concise documentation of capabilities necessary for operating networks securely, to assist network operators in communicating their requirements to vendors, and to provide vendors with input that is useful for building more secure devices. The working group will produce a list of capabilities appropriate for large Internet Service Provider (ISP) and Enterprise Networks. This work is intended to refine [RFC3871].

Table of Contents

1. Introduction

1.1 Goals

The goal of the Operational Security Working Group is to codify knowledge gained through operational experience about feature sets that are needed to securely deploy and operate managed network elements providing transit services at the data link and IP layers.

It is anticipated that the codification of this knowledge will be an aid to vendors in producing more securable network elements, and an aid to operators in increasing security by deploying and configuring more secure network elements.

This framework document provides an overview of the work to be done by the working group, and describes the documents to be produced in this effort.

1.2 Motivation

Network operators need the appropriate feature sets and tools on their infrastructure devices to ensure that they can effectively deploy and manage their networks securely while maintaining the ability to provide reliable service to their customers. Vendors need guidelines on which security features and functionality are critical for operators to be able to reach that goal.

1.3 Threat Model

1.3.1 Threats Addressed, Threats Not Addressed

This section describes the general classes of threats that this work intends to address. Specific threats and attacks will be discussed in the documents which are referred to in this framework. Each of those documents will enumerate the capabilities which are required to mitigate the risk of these specific threats.

The intent is to address real-world threats to and attacks on network infrastructure devices which have severely impacted network

operations or have immediate potential to do so. The intent is NOT to build a complete theoretical threat model or list every possible attack.

The threats will be limited to those that affect the management of network infrastructure and its ability to transit traffic. Threats to the confidentiality and integrity of transit traffic will not be addressed.

1.3.2 Active, Passive and Combined Attacks

[RFC3552] describes a general Internet threat model which readers of this document should be familiar with. It defines a threat model to describes the capabilities that an attacker is assumed to be able to deploy against a resource. [RFC3552] classifies attacks into two main categories: passive attacks and active attacks. Passive attacks are ones where an attacker simply reads information off the network and obtains confidential and/or private information which can be used to compromise network systems. Active attacks are ones where the attacker writes data to the network and can include replay attacks, message insertion, message deletion, message modification and man-in-the-middle attacks. Often, these passive and active attacks are combined. For example, routing information is diverted via a man-in-the-middle attack to force confidential information to transit a network path on which the attacker is able to perform eavesdropping.

1.3.3 Categories of Threats

The following sections provide a model that can be used to further categorize attacks on infrastructure devices and/or the operating behavior of these devices, and also gives some examples of attacks which fall into each classification.

It is common to categorize threats based on the effects or damage caused by associated attacks. For example, threats generally fall under one of the three categories as defined in [RFC2196]:

- Unauthorized access to resources and/or information
- Unintended and/or unauthorized disclosure of information
- Denial of service

There are a number of attacks, any one of which, if exploited, can lead to any of the above mentioned threats. As one example, if an intruder has taken control of a router (for example by guessing the password) then he could potentially obtain unauthorized access to resources, could gain unauthorized disclosure of information, and could also deny service to legitimate users. This method of categorizing threats based on the result of the threat therefore results in categories which are orthogonal to the cause of the effect, and thus orthogonal to the device capabilities which are

needed.

Categorization of attacks based on the capabilities required to mount the attack will allow the analysis and description of the attacks to be more closely aligned with the product capabilities required to defeat or mitigate the attack.

1.3.4 Threat Sources

The sources of threats in an operational network take many forms. Some sources can be intentional, such as a malicious intruder actively gaining access to an unauthorized resource or causing a denial of service attack. Other sources can be unintentional but still render the network unusable, such as software bugs or configuration mistakes. Many of the unintentional threat sources can be difficult to recognize or prevent. However wherever possible, capabilities and functionality will be defined which minimize the extent of the damage done under these circumstances.

Threats can originate from outside or inside and can be due to vulnerabilities in a device or weaknesses in operational processes. Inside threats pertain to an authorized participant in the operation of the network performing unauthorized actions. Outside threats pertain to any unauthorized network devices or person causing havoc with normal network operations.

On Path network devices are able to read, modify, or remove any datagram transmitted along a given path. Off-path hosts can transmit arbitrary datagrams that appear to come from any hosts but cannot necessarily receive datagrams intended for other hosts.

1.4 Attacks

This section specifies attack categories based on the capabilities required to mount the attack and provides more granular detail of many of the identifiable and recognized threats to which network infrastructure devices are susceptible.

1.4.1 Passive attacks

Passive attacks are ones where an attacker simply reads information off the network and obtains confidential and/or private information which can be used to compromise network systems.

1.4.2 Eavesdropping/Sniffing

The most common form of passive attack is eavesdropping, where the

attacker is able to read the data which is being transmitted from the sender to the receiver. In any operational network, the entire data path and every device involved in the data path must be considered for this type of attack. Any information which could be used to potentially gain unauthorized access to a device or is private must be protected. This includes passwords, configuration files and log files. It is common to think only of protecting the data path and to make sure that data is not diverted along a different path which may

be easier to eavesdrop on, such as a wireless network. In many instances it would be wise to consider cryptographically protecting data confidentiality wherever sensitive information is involved.

1.4.3 Off-line Cryptographic Attacks

These attacks typically capture some data which has been cryptographically protected and then use varying means to try and recover the original data. Poor password protection protocols can easily be reverse engineered and poorly chosen passwords can also be easily deciphered. As described in [RFC3552], a number of popular password-based challenge response protocols are vulnerable to a dictionary attack. The attacker captures a challenge-response pair and then proceeds to try entries from a list of common words (such as a dictionary file) until he finds a password that produces the right response.

1.4.4 Active Attacks

Active attacks are ones where the attacker writes data to the network. Generally, any part of a data packet can be forged. When the source IP address is forged, the attack is generally referred to as a spoofing attack. These attacks can be mitigated by filtering traffic based on IP addresses to only allow legitimate traffic to/from a network.

Not all active attacks require forged addresses and most systems are susceptible to a number of common attack patterns which are described in the next sections. Note that any type of active attack can be used for Denial of Service if the traffic is sent at such a rate that it exceeds a networks link capacity or exhausts device resources.

1.4.5 Replay Attacks

A replay attack is a combination of a passive and an active attack. In this type of attack, the attacker records some number of messages off of the wire and then plays them back to the original recipient. Note that the attacker does not need to be able to understand the messages. He merely needs to capture and re-transmit them.

1.4.6 Message Insertion

In a message insertion attack, the attacker forges one or more messages and injects them into the network. Often these messages will have a forged source address in order to disguise the identity of the attacker.

Message insertion attacks can be used to exploit known

vulnerabilities in protocol software. Routers and switches implement protocols which in some cases make use of software which is well known and widely deployed. Malicious attackers therefore may be familiar with the protocol software and be able to exploit known vulnerabilities.

1.4.7 Message Modification

In a message modification attack, the attacker removes a message from the wire, modifies it, and then resends it. The contents of the message may be modified and/or the intended recipient. [need example specific to network operations where this would be harmful]

1.4.8 Message Deletion

In a message deletion attack, the attacker simply removes a message from the wire. [need example specific to network operations where this is harmful]

1.4.9 Man-In-The-Middle

A Man-In-The-Middle attack combines the above techniques in a special form: The attacker subverts the communication stream in order to pose as the sender to receiver and the receiver to the sender. This differs fundamentally from the above forms of attack because it attacks the identity of the communicating parties, rather than the data stream itself. Consequently, many techniques which provide integrity of the communications stream are insufficient to protect against man-in-the-middle attacks.

Man-in-the-middle attacks are possible whenever peer entity authentication is not used. For example, it is trivial to mount man-in-the-middle attacks on local networks via ARP spoofing where the attacker forges an ARP with the victim's IP address and his own MAC address to gain access to a network. The attacker can then do further damage by sending forged messages. Imagine if the victim^Ôs IP address was that of a tftp server. The attacker could potentially download invalid system images or configuration files to a network device and subsequently compromise that network device.

[Ed. - Need to review existing capabilities. Do the threats and attack types listed above cover them all ? Are there capabilities that imply threats and attack classes not listed above]

1.5 Scope

The working group will produce a list of capabilities appropriate for:

- Internet Service Provider (ISP) Networks
- Enterprise Networks

The following are explicitly out of scope:

- general purpose hosts that do not transit traffic including infrastructure hosts such as name/time/log/AAA servers, etc.,
- unmanaged devices,
- customer managed devices (e.g. firewalls, Intrusion Detection System, dedicated VPN devices, etc.),
- SOHO (Small Office, Home Office) devices (e.g. personal firewalls, Wireless Access Points, Cable Modems, etc.),
- confidentiality of customer data,
- integrity of customer data,
- physical security.

These limitations have been made to keep the amount of work and size of documents manageable. While the capabilities listed here may apply to systems outside the scope, no capabilities have been added to account for their unique needs.

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

1.6 Intended Audience

There are two intended audiences: the network operator who selects, purchases, and operates IP network equipment, and the vendors who create these devices.

1.7 Format and Definition of Capabilities

A separate document will be created for specific categories of capabilities. Each individual capability will have the following element:

Capability (what)
: The capability describes a policy to be supported by the device.

For example, "The device MUST support secure channels that allow in-band access to all management and configuration functions."

Capabilities should not refer to specific technologies. It is expected that desired capability will change little over time.

Justification (why)
The justification tells why and in what context the capability is important.

For example, "Secure channels are important because they insure confidentiality, and integrity. This is important in contexts where management is performed in-band over networks with potentially hostile users."

The justification is intended to give operators information needed to determine the applicability of a capability their local environment.

Examples (how)
Examples are intended to give examples of technology and standards current at the time of writing that implement the capability. Examples of configuration and usage may also be given.

For example, "SSH provides access to management and configuration functions via secure channels. One way to provide this capability would be to enable SSH for in-band management and to disable all insecure in-band management mechanisms (e.g. telnet, SNMPv1, etc.)"

It is expected that the choice of implementations to provide the capability will change over time. See [RFC3631] for a list of some current mechanisms.

Warnings (if applicable)
The warnings list operational concerns, deviation from standards, caveats, etc.

For example, "If SSH is chosen as the mechanism to provide secure channels for remote management and configuration, then there are a number of issues which must be considered including key distribution and known vulnerabilities in various protocol versions."

1.8 Applicability

These capabilities are intended to give guidance on how best to protect communications infrastructure. Service Providers, Network Operators, and Equipment Suppliers are encouraged to study these capabilities, and prioritize the extent and manner in which they may implement and/or deploy equipment supporting these capabilities.

Decisions of whether or not to support a specific capabilities are intended to be left with the responsible organization (e.g., Service

Provider, Network Operator, or Equipment Supplier). Due to the continuously evolving nature of security threats to networks, and due to significant variations in the specific security threats and requirements in different network environments, it is not appropriate to mandate implementation of these capabilities through legislation or regulation, nor would any mandate be consistent with their intent.

1.9 Intended Use

It is anticipated that the capabilities in these documents will be used for the following purposes:

- o as a checklist when evaluating networked products,
- o to create profiles of different subsets of the capabilities which describe the needs of different devices, organizations, and operating environments,
- o to assist operators in clearly communicating their security requirements,
- o as high level guidance for the creation of detailed test plans.
- o as guidance for vendors to make appropriate decisions for engineering feature roadmaps.

1.10 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].

NOTE: The following definitions are take from RFC3871. Unless otherwise stated, the working group documents will use these terms as defined below.

Bogon.

A "Bogon" (plural: "bogons") is a packet with an IP source address in an address block not yet allocated by IANA or the Regional Internet Registries (ARIN, RIPE, APNIC...) as well as all addresses reserved for private or special use by RFCs. See [RFC3330] and [RFC1918].

CLI.

Several capabilities refer to a Command Line Interface (CLI). While this refers at present to a classic text oriented command interface, it is not intended to preclude other mechanisms which may provide all the capabilities that reference "CLI".

Console.

Several capabilities refer to a "Console". The model for this is the classic RS232 serial port which has, for the past 30 or more years, provided a simple, stable, reliable, well-understood and nearly ubiquitous management interface to network devices. Again, these capabilities are intended primarily to codify the benefits

provided by that venerable interface, not to preclude other mechanisms that provide the same capabilities.

Filter.

In this document, a "filter" is defined as a group of one or more rules where each rule specifies one or more match criteria.

In-Band management.

"In-Band management" is defined as any management done over the same channels and interfaces used for user/customer data. Examples would include using SSH for management via customer or Internet facing network interfaces.

High Resolution Time.

"High resolution time" is defined in this document as "time having a resolution greater than one second" (e.g. milliseconds).

IP.

Unless otherwise indicated, "IP" refers to IPv4.

Management.

This document uses a broad definition of the term "management". In this document, "management" refers to any authorized interaction with the device intended to change its operational state or configuration. Data/Forwarding plane functions (e.g. the transit of customer traffic) are not considered management. Control plane functions such as routing, signaling and link management protocols and management plane functions such as remote access, configuration and authentication are considered to be management.

Martian.

Per [RFC1208] "Martian: Humorous term applied to packets that turn up unexpectedly on the wrong network because of bogus routing entries. Also used as a name for a packet which has an altogether bogus (non-registered or ill-formed) Internet address." For the purposes of this document Martians are defined as "packets having a source address that, by application of the current forwarding tables, would not have its return traffic routed back to the sender." "Spoofed packets" are a common source of martians. Note that in some cases, the traffic may be asymmetric, and a simple forwarding table check might produce false positives. See [RFC3704]

Out-of-Band (OoB) management.

"Out-of-Band management" is defined as any management done over channels and interfaces that are separate from those used for user/customer data. Examples would include a serial console interface or a network interface connected to a dedicated management network that is not used to carry customer traffic.

Open Review.

"Open review" refers to processes designed to generate public discussion and review of technical solutions such as data communications protocols and cryptographic algorithms with the goals of improving and building confidence in the final solutions. For the purposes of this document "open review" is defined by [RFC2026]. All standards track documents are considered to have been through an open review process.

It should be noted that organizations may have local requirements that define what they view as acceptable "open review". For example, they may be required to adhere to certain national or international standards. Such modifications of the definition of the term "open review", while important, are considered local issues that should be discussed between the organization and the vendor.

It should also be noted that section 7 of [RFC2026] permits standards track documents to incorporate other "external standards and specifications".

Service.

A number of capabilities refer to "services". For the purposes of this document a "service" is defined as "any process or protocol running in the control or management planes to which non-transit packets may be delivered". Examples might include an SSH server, a BGP process or an NTP server. It would also include the transport, network and link layer protocols since, for example, a TCP packet addressed to a port on which no service is listening will be "delivered" to the IP stack, and possibly result in an ICMP message being sent back.

Secure Channel.

A "secure channel" is a mechanism that ensures end-to-end integrity and confidentiality of communications. Examples include TLS [RFC2246] and IPsec [RFC2401]. Connecting a terminal to a console port using physically secure, shielded cable would provide confidentiality but possibly not integrity.

Single-Homed Network.

A "single-homed network" is defined as one for which

* There is only one upstream connection
* Routing is symmetric.

See [RFC3704] for a discussion of related issues and mechanisms for multi-homed networks.

Spoofed Packet.

A "spoofed packet" is defined as a packet that has a source address that does not correspond to any address assigned to the system which sent the packet. Spoofed packets are often "bogons" or "martians".

Secure Network

For the purposes of these documents, a secure network is one in which:

* The network keeps passing legitimate customer traffic (availability).
* Traffic goes where it is supposed to go, and only where it is supposed to go (availability, confidentiality).
* The network elements remain manageable (availability).
* Only authorized users can manage network elements (authorization).
* There is a record of all security related events (accountability).
* The network operator has the necessary tools to detect and respond to illegitimate traffic.

2. Documents

[Ed. This list of documents is likely to be consolidated/reduced]

The following is a list of documents to be produced by OPSEC working group. Each document is intended to cover an area important to secure operation of large network infrastructure.

2.1 Standards Survey (info)

Overview

This document provides an overview of other efforts in developing standards, guidelines, best practices, or other information intended to facilitate improvement in network security. Any effort which is known, such as the ANSI T1.276, the NRIC V "Best Practices", ITU-T M.3016 and X.805, the T1S1 effort on securing signalling will be included. The intent is to provide a clear understanding of which efforts are complementary and/or contradictory such that any efforts of future cross-certification of standards may be facilitated.

Security Considerations

Many contradictory security requirements from varying standards bodies would seriously impact operator or vendor understanding of which features and functionalities are the most effective to deploy and operate secure networks. This documented survey will help to ensure that there is a consistent set of product requirements to follow.

2.2 In-Band management capabilities (BCP)

Overview

Although there are known security issues with in-band management, there are many situations where in-band management makes sense, is used, and/or is the only option. The features recommended in this document will provide for a more secure means of using any in-band management functionality.

Security Considerations

Although in-band management has the advantage of lower cost (no extra interfaces or lines), it has significant security disadvantages:

* Saturation of customer lines or interfaces can make the device unmanageable unless out-of-band management resources have been reserved
* 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, however this would typically require a significant compromise in the routing system.

* Public interfaces used for in-band management may become unavailable due to bugs (e.g. buffer overflows being exploited) while out-of-band interfaces (such as a serial console device) remain available

The capabilities from this document are meant to provide the means of securing in-band management traffic.

2.3 Out-of-Band management capabilities (BCP)

Overview

This document will describe capabilities related to out of band management of networked devices.

Security Considerations

Out-of-band management often provides a more secure means of managing networked devices. To ensure that all devices have the appropriate support, this document will list capabilities as to what functionality is needed to effectively use out-of-band management.

2.4 Filtering capabilities (BCP)

Overview

This document will describe capabilities related to stateless filtering for network elements providing transit service at link and transport layers.

Security Considerations

Filtering is an important security functionality to permit or deny forwarding of traffic, or to specify special treatment of packets, depending on layer 2 or layer 3 header and forwarding information. It provides a basic means of implementing policies, such as policies that specify which traffic is allowed and which is not, and policies which specify special treatment such as setting CoS, rate limiting, or packet copying. It also provides a basic tool for responding to malicious traffic.

2.5 Event Logging Capabilities document (BCP)

Overview

[Ed. The basic questions here are "what gets logged", "how does it get logged", "what are the security issues". There is work in progress (syslog) for the last two that can be cited. The "what gets logged" question needs work]

This document will describe the recommended features when logging network device traffic and anomalies. The goal is to make it possible to correlate logging information from varying systems and making sure that logged information is useful and effective.

Security Considerations

Logging data provides a means for detecting malicious behavior.

The logged information can also be used as evidence in legal prosecution cases against illegal network access and device compromises. Ineffective logging practices due to inconsistent

functionality in many devices make it hard to get effective data. This document will help provide consistent logging functionality for more effective auditing. It will also point to privacy or legal considerations when logging/monitoring user activity.

2.6 Configuration and Management Interface Capabilities (BCP)

Overview

This document lists the security capabilities necessary for interfaces which allow for configuring and managing the network device. In most cases, this currently involves some sort of command line interface (CLI) and configuration files. It may be possible to provide the capabilities with other mechanisms, for instance SNMP or a script-able HTML interface that provides full access to management and configuration functions. In the future, there may be others (e.g. XML based configuration).

Security Considerations

The interfaces used to manage and configure network elements need to be effectively secured to avoid a malicious user from being able to logically gain illegal access. In the past, many security vulnerabilities have been discovered, especially with SNMP and HTTP access to devices. These recommendations will help the user and vendor community mitigate any known risks in this area.

2.7 AAA capabilities document (BCP)

Overview

This document will list the capabilities needed for centralized authentication, authorization and accounting functionality.

Security Considerations

Keeping track of who has access to network devices is critical to any secure infrastructure. Mechanisms to provide authorized access upon successful authentication and also keeping track of what was done can provide important information in case of a device compromise.

2.8 Documentation and Assurance capabilities document (BCP)

Overview

These capabilities will list information which should be documented that will assist operators in evaluating and securely operating a device.

Security Considerations

Devices many times have default behavior which can cause a severe security vulnerability. Knowing which services are enabled by default or which commands impact other default behavior is essential knowledge that is necessary to effectively mitigate security risks.

2.9 Miscellaneous capabilities document (BCP)

Overview

This document will describe capabilities which do not fit into any of the other documents, and which are brief enough that they don't justify their own document, but which are important enough that they should be documented.

2.10 Large ISP Operational Security Capabilities Profile (info)

Overview

This document will provide a profile specifying which of the capabilities outlined in the set of documents described above are most applicable to large Internet Service Providers offering transit service.

2.11 Large Enterprise Operational Security Capabilities Profile (info)

Overview

This document will provide a profile specifying which of the capabilities outlined in the set of documents described above are most applicable to large Enterprise networks.

2.12 OPSEC Deliberation Summary document (info)

Overview

This document will provide a summary of discussions that have taken place within the OPsec working group. The intent is to document ideas that were "left on the cutting room floor" in order to provide a possible starting point for future work.

3. Security Considerations

Security is the entire focus of this document.

4 Normative References

[RFC1208] Jacobsen, O. and D. Lynch, "Glossary of networking terms", RFC 1208, March 1991.

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G. and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

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

[RFC2196] Fraser, B., "Site Security Handbook", RFC 2196, September 1997.

[RFC2246] Dierks, T. and C. Allen, "The TLS Protocol Version 1.0", RFC 2246, January 1999.

[RFC2401] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.

[RFC3013] Killalea, T., "Recommended Internet Service Provider Security Services and Procedures", BCP 46, RFC 3013, November 2000.

[RFC3330] IANA, "Special-Use IPv4 Addresses", RFC 3330, September 2002.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July

2003.

[RFC3631] Bellovin, S., Schiller, J. and C. Kaufman, "Security Mechanisms for the Internet", RFC 3631, December 2003.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.

[RFC3871] Jones, G., "Operational Security Requirements for Large Internet Service Provider (ISP) IP Network Infrastructure", RFC 3871, September 2004.

Authors' Addresses

George M. Jones
The MITRE Corporation
7515 Colshire Drive, M/S WEST
McLean, Virginia 22102-7508
U.S.A.

Phone: +1 703 488 9740
EMail: gmjones@mitre.org

Ross Callon
Juniper Networks
10 Technology Park Drive
Westford, MA 01886
U.S.A.

Phone: +1 978 692 6724
EMail: rcallon@juniper.net

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

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

Appendix A. Acknowledgments

The authors gratefully acknowledge the contributions of:

- Acknowledgments to be determined.
- The MITRE Corporation for supporting development of this document. NOTE: The author'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.
- 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.

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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