

None.

Internet-Draft

Intended status: Informational

Expires: March 5, 2007

C. Morrow

UUNET Technologies

G. Jones

The MITRE Corporation

V. Manral

IP Infusion

September 1, 2006

**Filtering and Rate Limiting Capabilities for IP Network Infrastructure
draft-ietf-opsec-filter-caps-04**

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on March 5, 2007.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

[I-D.ietf-opsec-current-practices] lists operator practices related to securing networks. This document lists filtering and rate limiting capabilities needed to support those practices. Capabilities are limited to filtering and rate limiting packets as they enter or leave the device. Route filters and service specific filters (e.g. SNMP, telnet) are not addressed.

Capabilities are defined without reference to specific technologies. This is done to leave room for deployment of new technologies that implement the capability. Each capability cites the practices it supports. Current implementations that support the capability are cited. Special considerations are discussed as appropriate listing operational and resource constraints, limitations of current implementations, trade-offs, etc.

Table of Contents

1.	Introduction	4
1.1.	Threat Model	4
1.2.	Format	5
2.	Packet Selection for Management and Data Plane Controls	6
3.	Packet Selection Criteria	7
3.1.	Select Traffic on All Interfaces	7
3.2.	Select Traffic To the Device	8
3.3.	Select Transit Traffic	8
3.4.	Select Inbound and/or Outbound	9
3.5.	Select by Protocols	10
3.6.	Select by Addresses	11
3.7.	Select by Protocol Header Fields	12
4.	Actions	14
4.1.	Specify Filter Actions	14
4.2.	Specify Rate Limits	14
4.3.	Specify Log Actions	15
4.4.	Specify Log Granularity	16
4.5.	Ability to Display Filter Counters	17
5.	Counters	18
5.1.	Filter Counters Displayed Per Application	18
5.2.	Ability to Reset Filter Counters	18
5.3.	Filter Hits are Counted	19
5.4.	Filter Counters are Accurate	20
6.	Minimal Performance Degradation	21
7.	Additional Operational Practices	23
7.1.	Profile Current Traffic	23
7.2.	Block Malicious Packets	23
7.3.	Limit Sources of Management	23
7.4.	Respond to Incidents Based on Accurate Data	23
7.5.	Implement Filters Where Necessary	24
8.	Security Considerations	25
9.	Non-normative References	26
Appendix A.	Acknowledgments	27
	Authors' Addresses	28
	Intellectual Property and Copyright Statements	29

1. Introduction

This document is defined in the context of

[[I-D.ietf-opsec-current-practices](#)].

[[I-D.ietf-opsec-current-practices](#)] defines the goals, motivation, scope, definitions, intended audience, threat model, potential attacks and give justifications for each of the practices. Many of the capabilities listed here refine or add to capabilities listed in [[RFC3871](#)].

Also see [[I-D.lewis-infrastructure-security](#)] for a useful description of techniques for protecting infrastructure devices, including the use of filtering.

1.1. Threat Model

Threats in today's networked environment range from simple packet floods with overwhelming bandwidth toward a leaf network to subtle attacks aimed at subverting known vulnerabilities in existing applications. The attacked network or host might not be an end user, it may be the networking device or links inside the provider core.

Networks must have the ability to place mitigation in order to limit these threats. These mitigation steps could include routing updates, traffic filters, and routing filters. It is possible that the mitigation steps might have to affect transit traffic as well as traffic destined to the device on which the mitigation steps are activated.

The scope of the threat includes simply denying services to an individual customer on one side of the scale to exploiting a newly discovered protocol vulnerability which affects the entire provider core. The obvious risk to the business requires mitigation capabilities which can span this range of threats.

Threat: An indication of impending danger or harm to the network or its parts. This could be formed from the projected loss of revenue to the business. Additionally, it could be formed from the increased cost to the business caused by the event. (more interfaces, more bandwidth, more personnel to support the increased size or complexity)

Risk: The possibility of suffering harm or loss of network services due to a threat.

Attack: To set upon with violent force the network or its parts. Typically this is a form of flood of packets to or through a network. This could also be a much smaller stream of packets created with the

intent of exploiting a vulnerability in the infrastructure of the network.

Asset: Either a customer, network device or network link. Any of these could be assets from a business perspective.

These terms are more completely defined in [[RFC2828](#)] we have added some scope specific information only.

Also see [[I-D.savola-rtgwg-backbone-attacks](#)] for a list of attacks on backbone devices and counter measures.

1.2. Format

Each capability has the following subsections:

- o Capability (what)
- o Supported Practices (why)
- o Current Implementations (how)
- o Considerations (caveats, resource issues, protocol issues, etc.)

The Capability section describes a feature to be supported by the device. The Supported Practice section cites practices described in [[I-D.ietf-opsec-current-practices](#)] that are supported by this capability. The Current Implementation section is intended to give examples of implementations of the capability, citing technology and standards current at the time of writing. It is expected that the choice of features to implement the capabilities will change over time. The Considerations section lists operational and resource constraints, limitations of current implementations, trade-offs, etc.

2. Packet Selection for Management and Data Plane Controls

In this document [Section 3](#) describes a number of criteria for performing packet selection. It is assumed in this document that

- o all of these criteria can be used to select packets for both filtering and rate limiting packets,
- o management plane controls can be implemented by applying these criteria to filter/rate limit traffic destined for the device itself,
- o data plane controls can be implemented by applying these criteria to filter/rate limit traffic destined through the device
- o multiple packet selection criteria can be used to select a single set of packets for filtering action

3. Packet Selection Criteria

This section lists packet selection criteria that can be applied to both filtering and rate limiting.

3.1. Select Traffic on All Interfaces

Capability.

The device provides a means to filter IP packets on any interface implementing IP.

Supported Practices.

- * Security Practices for Device Management ([[I-D.ietf-opsec-current-practices](#)], Section 2.2.2)
- * Security Practices for Data Path ([[I-D.ietf-opsec-current-practices](#)], [Section 2.3.2](#))
- * Security Practices for Software Upgrades and Configuration Integrity/Validation ([[I-D.ietf-opsec-current-practices](#)], Section 2.5.2)
- * Data Plane Filtering ([[I-D.ietf-opsec-current-practices](#)], Section 2.7.1)
- * Management Plane Filtering ([[I-D.ietf-opsec-current-practices](#)], Section 2.7.2)
- * Profile Current Traffic ([Section 7.1](#))
- * Block Malicious Packets ([Section 7.2](#))

Current Implementations.

Many devices currently implement access control lists or filters that allow filtering based on protocol and/or source/destination address and or source/destination port and allow these filters to be applied to interfaces.

Considerations.

None.

3.2. Select Traffic To the Device

Capability.

It is possible to apply the filtering mechanism to traffic that is addressed directly to the device via any of its interfaces - including loopback interfaces.

Supported Practices.

- * Security Practices for Device Management ([[I-D.ietf-opsec-current-practices](#)], Section 2.2.2)
- * Security Practices for Software Upgrades and Configuration Integrity/Validation ([[I-D.ietf-opsec-current-practices](#)], Section 2.5.2)
- * Management Plane Filtering ([[I-D.ietf-opsec-current-practices](#)], Section 2.7.2)

Current Implementations.

Many devices currently implement access control lists or filters that allow filtering based on protocol and/or source/destination address and or source/destination port and allow these filters to be applied to services offered by the device.

Examples of this might include filters that permit only BGP from peers and SNMP and SSH from an authorized management segment and directed to the device itself, while dropping all other traffic addressed to the device.

Considerations.

None.

3.3. Select Transit Traffic

Capability.

It is possible to apply the filtering mechanism to traffic that will transit the device via any of its interfaces.

Supported Practices.

- * Security Practices for Data Path ([[I-D.ietf-opsec-current-practices](#)], Section 2.3.2)

- * Data Plane Filtering ([[I-D.ietf-opsec-current-practices](#)], Section 2.7.1)

Current Implementations.

Many devices currently implement access control lists or filters that allow filtering based on protocol and/or source/destination address and or source/destination port and allow these filters to be applied to the interfaces on the device in order to protect assets attached to the network.

Examples of this may include filtering all traffic save SMTP (tcp/25) destined to a mail server. A common use of this today would also be denying all traffic to a destination which has been determined to be hostile.

Considerations.

This allows the operator to apply filters that protect the networks and assets surrounding the device from attacks and unauthorized access.

[3.4.](#) Select Inbound and/or Outbound

Capability.

It is possible to filter both incoming and outgoing traffic on any interface.

Supported Practices.

- * Security Practices for Device Management ([[I-D.ietf-opsec-current-practices](#)], Section 2.2.2)
- * Security Practices for Data Path ([[I-D.ietf-opsec-current-practices](#)], Section 2.3.2)
- * Security Practices for Software Upgrades and Configuration Integrity/Validation ([[I-D.ietf-opsec-current-practices](#)], Section 2.5.2)
- * Data Plane Filtering ([[I-D.ietf-opsec-current-practices](#)], Section 2.7.1)
- * Management Plane Filtering ([[I-D.ietf-opsec-current-practices](#)], Section 2.7.2)

Current Implementations.

It might be desirable on a border router, for example, to apply an egress filter outbound on the interface that connects a site to its external ISP to drop outbound traffic that does not have a valid internal source address. Inbound, it might be desirable to apply a filter that blocks all traffic from a site that is known to forward or originate large amounts of junk mail.

Considerations.

This allows flexibility in applying filters at the place that makes the most sense. It allows invalid or malicious traffic to be dropped as close to the source as possible with the least impact on other traffic transiting the interface(s) in question.

3.5. Select by Protocols

Capability.

The device provides a means to filter traffic based on the value of the protocol field in the IP header.

Supported Practices.

- * Security Practices for Device Management
([[I-D.ietf-opsec-current-practices](#)], Section 2.2.2)
- * Security Practices for Data Path
([[I-D.ietf-opsec-current-practices](#)], Section 2.3.2)
- * Security Practices for Software Upgrades and Configuration Integrity/Validation
([[I-D.ietf-opsec-current-practices](#)], [Section 2.5.2](#))
- * Data Plane Filtering ([[I-D.ietf-opsec-current-practices](#)], Section 2.7.1)
- * Management Plane Filtering ([[I-D.ietf-opsec-current-practices](#)], Section 2.7.2)

Current Implementations.

Some denial of service attacks are based on the ability to flood the victim with ICMP traffic. One quick way (admittedly with some negative side effects) to mitigate the effects of such attacks is to drop all ICMP traffic headed toward the victim.

Considerations.

Being able to filter on protocol is necessary to allow implementation of policy, secure operations and for support of incident response. Filtering all traffic to a destination host is not often possible, business requirements will dictate that critical traffic be permitted if at all possible.

3.6. Select by Addresses

Capability.

The device is able to control the flow of traffic based on source and/or destination IP address or blocks of addresses such as Classless Inter-Domain Routing (CIDR) blocks.

Supported Practices.

- * Security Practices for Device Management
([[I-D.ietf-opsec-current-practices](#)], Section 2.2.2)
- * Security Practices for Data Path
([[I-D.ietf-opsec-current-practices](#)], Section 2.3.2)
- * Security Practices for Software Upgrades and Configuration Integrity/Validation ([[I-D.ietf-opsec-current-practices](#)], Section 2.5.2)
- * Data Plane Filtering ([[I-D.ietf-opsec-current-practices](#)], Section 2.7.1)
- * Management Plane Filtering ([[I-D.ietf-opsec-current-practices](#)], Section 2.7.2)

Current Implementations.

One example of the use of address based filtering is to implement ingress filtering per [[RFC2827](#)]

Considerations.

The capability to filter on addresses and address blocks is a fundamental tool for establishing boundaries between different networks.

3.7. Select by Protocol Header Fields

Capability.

The filtering mechanism supports filtering based on the value(s) of any portion of the protocol headers for IP, ICMP, UDP and TCP by specifying fields by name (e.g., "protocol = ICMP") rather than bit- offset/length/numeric value (e.g., 72:8 = 1).

It supports arbitrary header-based filtering (possibly using bit- offset/length/value) of all other protocols.

Supported Practices.

- * Security Practices for Device Management
([\[I-D.ietf-opsec-current-practices\]](#), Section 2.2.2)
- * Security Practices for Data Path
([\[I-D.ietf-opsec-current-practices\]](#), Section 2.3.2)
- * Security Practices for Software Upgrades and Configuration Integrity/Validation ([\[I-D.ietf-opsec-current-practices\]](#), Section 2.5.2)
- * Data Plane Filtering ([\[I-D.ietf-opsec-current-practices\]](#), Section 2.7.1)
- * Management Plane Filtering ([\[I-D.ietf-opsec-current-practices\]](#), Section 2.7.2)

Current Implementations.

This capability implies that it is possible to filter based on TCP or UDP port numbers, TCP flags such as SYN, ACK and RST bits, and ICMP type and code fields. One common example is to reject "inbound" TCP connection attempts (TCP, SYN bit set+ACK bit clear or SYN bit set+ACK,FIN and RST bits clear). Another common example is the ability to control what services are allowed in/out of a network. It may be desirable to only allow inbound connections on port 80 (HTTP) and 443 (HTTPS) to a network hosting web servers.

Supporting arbitrary offset/length/value filtering allows filtering of unknown (possibly new) protocols, e.g. filtering RTP even when the device itself does not support RTP.

Considerations.

Being able to filter on portions of the header is necessary to allow implementation of policy, secure operations, and support incident response.

4. Actions

4.1. Specify Filter Actions

Capability.

The device provides a mechanism to allow the specification of the action to be taken when a filter rule matches. Actions include "permit" (allow the traffic), "reject" (drop with appropriate notification to sender), and "drop" (drop with no notification to sender).

Supported Practices.

- * Data Origin Authentication ([[I-D.ietf-opsec-current-practices](#)], Section 2.3.3)

Current Implementations.

Assume that your management devices for deployed networking devices live on several subnets, use several protocols, and are controlled by several different parts of your organization. There might exist a reason to have disparate policies for access to the devices from these parts of the organization.

Actions such as "permit", "reject", and "drop" are essential in defining the security policy for the services offered by the network devices.

Considerations.

While silently dropping traffic without sending notification may be the correct action in security terms, consideration should be given to operational implications. See [[RFC3360](#)] for consideration of potential problems caused by sending inappropriate TCP Resets.

Also note that it might be possible for an attacker to effect a denial of service attack by causing too many rejection notifications to be sent (e.g. syslog messages). For this reason it might be desirable to rate-limit notifications.

4.2. Specify Rate Limits

Capability.

The device provides a mechanism to allow the specification of the action to be taken when a rate limiting filter rule matches. The actions include "transmit" (permit the traffic because it's below the specified limit), "limit" (limit traffic because it exceeds the specified limit). Limits should be applicable by both bits per second and packets per timeframe (possible timeframes might include second, minute, hour). Limits should be able to be placed in both inbound and outbound directions.

Supported Practices.

- * Denial of Service Tracking/Tracing with Rate Limiting ([[I-D.ietf-opsec-current-practices](#)], Section 2.8.4)

Current Implementations.

Assume that your management devices for deployed networking devices live on several subnets, use several protocols, and are controlled by several different parts of your organization. There might exist a reason to have disparate policies for access to the devices from these parts of the organization with respect to priority access to these services. Rate Limits may be used to enforce these prioritizations.

Considerations.

This capability allows a filter to be used to rate limit a portion of traffic through or to a device. It maybe desirable to limit SNMP (UDP/161) traffic to a device, but not deny it completely. Similarly, one might want to implement ICMP filters toward an external network instead of discarding all ICMP traffic.

While silently dropping traffic without sending notification may be the correct action in security terms, consideration should be given to operational implications. See [[RFC3360](#)] for consideration of potential problems caused by sending inappropriate TCP Resets.

[4.3.](#) Specify Log Actions

Capability.

It is possible to log all filter actions. The logging capability is able to capture at least the following data:

- * permit/reject/drop status
- * source and destination IP address
- * source and destination ports (if applicable to the protocol)
- * which network element received or was sending the packet (interface, MAC address or other layer 2 information that identifies the previous hop source of the packet).

Supported Practices.

- * Logging Security Practices([\[I-D.ietf-opsec-current-practices\]](#), Section 2.6.2)

Current Implementations.

Actions such as "permit", "reject", "drop" are essential in defining the security policy for the services offered by the network devices. Auditing the frequency, sources and destinations of these attempts is essential for tracking ongoing issues today.

Considerations.

Logging can be burdensome to the network device, at no time should logging cause performance degradation to the device or services offered on the device.

Also note logging itself can be rate limited so as to not cause performance degradation of the device or the network(in case of syslog or other similar network logging mechanism).

4.4. Specify Log Granularity

Capability.

It is possible to enable/disable logging on a per rule basis.

Supported Practices.

- * Logging Security Practices([\[I-D.ietf-opsec-current-practices\]](#), Section 2.6.2)

Current Implementations.

If a filter is defined that has several rules, and one of the rules denies telnet (tcp/23) connections, then it should be possible to specify that only matches on the rule that denies

telnet should generate a log message.

Considerations.

The ability to tune the granularity of logging allows the operator to log the information that is desired and only the information that is desired. Without this capability, it is possible that extra data (or none at all) would be logged, making it more difficult to find relevant information.

4.5. Ability to Display Filter Counters

Capability.

The device provides a mechanism to display filter counters.

Supported Practices.

- * Profile Current Traffic ([Section 7.1](#))
- * Respond to Incidents Based on Accurate Data ([Section 7.4](#))

Current Implementations.

Assume there is a router with four interfaces. One is an up-link to an ISP providing routes to the Internet. The other three connect to separate internal networks. Assume that a host on one of the internal networks has been compromised by a hacker and is sending traffic with bogus source addresses. In such a situation, it might be desirable to apply ingress filters to each of the internal interfaces. Once the filters are in place, the counters can be examined to determine the source (inbound interface) of the bogus packets.

Considerations.

None.

5. Counters

5.1. Filter Counters Displayed Per Application

Capability.

If it is possible for a filter to be applied more than once at the same time, then the device provides a mechanism to display filter counters per filter application.

Supported Practices.

- * Profile Current Traffic ([Section 7.1](#))
- * Respond to Incidents Based on Accurate Data ([Section 7.4](#))

Current Implementations.

One way to implement this capability would be to have the counter display mechanism show the interface (or other entity) to which the filter has been applied, along with the name (or other designator) for the filter. For example if a filter named "desktop_outbound" applied two different interfaces, say, "ethernet0" and "ethernet1", the display should indicate something like "matches of filter 'desktop_outbound' on ethernet0 ..." and "matches of filter 'desktop_outbound' on ethernet1 ..."

Considerations.

It may make sense to apply the same filter definition simultaneously more than one time (to different interfaces, etc.). If so, it would be much more useful to know which instance of a filter is matching than to know that some instance was matching somewhere.

5.2. Ability to Reset Filter Counters

Capability.

It is possible to reset counters to zero on a per filter basis.

Supported Practices.

- * Profile Current Traffic ([Section 7.1](#))
- * Respond to Incidents Based on Accurate Data ([Section 7.4](#))

Current Implementations.

For the purposes of this capability it would be acceptable for the system to maintain two counters: an "absolute counter", C[now], and a "reset" counter, C[reset]. The absolute counter would maintain counts that increase monotonically until they wrap or overflow the counter. The reset counter would receive a copy of the current value of the absolute counter when the reset function was issued for that counter. Functions that display or retrieve the counter could then display the delta (C[now] - C[reset]).

Considerations.

Assume that filter counters are being used to detect internal hosts that are infected with a new worm. Once it is believed that all infected hosts have been cleaned up and the worm removed, the next step would be to verify that. One way of doing so would be to reset the filter counters to zero and see if traffic indicative of the worm has ceased.

5.3. Filter Hits are Counted

Capability.

The device supplies a facility for counting all filter matches.

Supported Practices.

- * Profile Current Traffic ([Section 7.1](#))
- * Respond to Incidents Based on Accurate Data ([Section 7.4](#))

Current Implementations.

Assume, for example, that a ISP network implements anti-spoofing egress filters (see [[RFC2827](#)]) on interfaces of its edge routers that support single-homed stub networks. Counters could enable the ISP to detect cases where large numbers of spoofed packets are being sent. This may indicate that the customer is performing potentially malicious actions (possibly in violation of the ISPs Acceptable Use Policy), or that system(s) on the customers network have been "owned" by hackers and are being (mis)used to launch attacks.

Considerations.

None.

5.4. Filter Counters are Accurate

Capability.

Filter counters are accurate. They reflect the actual number of matching packets since the last counter reset. Filter counters are be capable of holding up to $2^{32} - 1$ values without overflowing and should be capable of holding up to $2^{64} - 1$ values.

Supported Practices.

- * Respond to Incidents Based on Accurate Data ([Section 7.4](#))

Current Implementations.

If N packets matching a filter are sent to/through a device, then the counter should show N matches.

Considerations.

None.

6. Minimal Performance Degradation

Capability.

The device provides a means to filter packets without significant performance degradation. This specifically applies to stateless packet filtering operating on layer 3 (IP) and layer 4 (TCP or UDP) headers, as well as normal packet forwarding information such as incoming and outgoing interfaces.

The device is able to apply stateless packet filters on ALL interfaces (up to the total number of interfaces attached to the device) simultaneously and with multiple filters per interface (e.g., inbound and outbound).

Supported Practices.

- * Implement Filters Where Necessary ([Section 7.5](#))

Current Implementations.

Another way of stating the capability is that filter performance should not be the limiting factor in device throughput. If a device is capable of forwarding 30Mb/sec without filtering, then it should be able to forward the same amount with filtering in place.

Considerations.

The definition of "significant" is subjective. At one end of the spectrum it might mean "the application of filters may cause the box to crash". At the other end would be a throughput loss of less than one percent with tens of thousands of filters applied. The level of performance degradation that is acceptable will have to be determined by the operator.

Repeatable test data showing filter performance impact would be very useful in evaluating this capability. Tests should include such information as packet size, packet rate, number of interfaces tested (source/destination), types of interfaces, routing table size, routing protocols in use, frequency of routing updates, etc. This capability does not address stateful filtering, filtering above layer 4 headers or other more advanced types of filtering that may be important in certain operational environments. Finally, if key infrastructure devices crash or experience severe performance degradation when filtering under heavy load, or even have the reputation of doing so, it is likely that security personnel will be forbidden, by policy, from using filtering in

ways that would otherwise be appropriate for fear that it might cause unnecessary service disruption.

7. Additional Operational Practices

This section describes practices not covered in [[I-D.ietf-opsec-current-practices](#)]. They are included here to provide justification for capabilities that reference them.

7.1. Profile Current Traffic

This capability allows a network operator to monitor traffic across an active interface in the network at a minimal level. This helps to determine probable cause for interface or network problems.

The ability to separate and distinguish traffic at a layer-3 or layer-4 level allows the operator to characterize beyond simple interface counters the traffic in question. This is critical because often the operator has no tools available for protocol analysis aside from interface filters.

7.2. Block Malicious Packets

Blocking or limiting traffic deemed to be malicious is a key component of application of any security policy's implementation. Clearly it is critical to be able to implement a security policy on a network.

Malicious packets could potentially be defined by any part of the layer-3 or layer-4 headers of the IP packet. The ability to classify or select traffic based on these criteria and take some action based on that classification is critical to operations of a network.

7.3. Limit Sources of Management

Management of a network should be limited to only trusted hosts. This implies that the network elements will be able to limit access to management functions to these trusted hosts.

Currently operators will limit access to the management functions on a network device to only the hosts that are trusted to perform that function. This allows separation of critical functions and protection of those functions on the network devices.

7.4. Respond to Incidents Based on Accurate Data

Accurate counting of filter rule matches is important because it shows the frequency of attempts to violate policy. Inaccurate data can not be relied on as the basis for action. Under-reported data can conceal the magnitude of a problem. This enables resources to be focused on areas of greatest need.

7.5. Implement Filters Where Necessary

This enables the implementation of filters on whichever services are necessary. To the extent that filtering causes degradation, it may not be possible to apply filters that implement the appropriate policies.

8. Security Considerations

General

Security is the subject matter of this entire memo. The capabilities listed cite practices in [\[I-D.ietf-opsec-current-practices\]](#) that they are intended to support. [\[I-D.ietf-opsec-current-practices\]](#) defines the threat model, practices and lists justifications for each practice.

9. Non-normative References

- [I-D.ietf-opsec-current-practices]
Kaeo, M., "Operational Security Current Practices",
[draft-ietf-opsec-current-practices-04](#) (work in progress),
June 2006.
- [I-D.lewis-infrastructure-security]
Lewis, D., "Service Provider Infrastructure Security",
[draft-lewis-infrastructure-security-00](#) (work in progress),
June 2006.
- [I-D.savola-rtgwg-backbone-attacks]
Savola, P., "Backbone Infrastructure Attacks and
Protections", [draft-savola-rtgwg-backbone-attacks-01](#) (work
in progress), June 2006.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering:
Defeating Denial of Service Attacks which employ IP Source
Address Spoofing", [BCP 38](#), [RFC 2827](#), May 2000.
- [RFC2828] Shirey, R., "Internet Security Glossary", [RFC 2828](#),
May 2000.
- [RFC3360] Floyd, S., "Inappropriate TCP Resets Considered Harmful",
[BCP 60](#), [RFC 3360](#), August 2002.
- [RFC3871] Jones, G., "Operational Security Requirements for Large
Internet Service Provider (ISP) IP Network
Infrastructure", [RFC 3871](#), September 2004.

[Appendix A](#). Acknowledgments

The authors gratefully acknowledge the contributions of:

- o Merike Kaeo for help aligning these capabilities with supported practices
- o The MITRE Corporation for supporting 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.

Authors' Addresses

Christopher L. Morrow
UUNET Technologies
21830 UUNet Way
Ashburn, Virginia 21047
U.S.A.

Phone: +1 703 886 3823
Email: chris@uu.net

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

Vishwas Manral
IP Infusion
Ground Floor, 5th Cross Road, Off 8th Main Road
Bangalore, 52
India

Phone: +91-80-4113-1268
Email: vishwas@ipinfusion.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

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

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

Intellectual Property

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

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

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

Acknowledgment

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

