R. Callon
Juniper Networks
G. Jones
The Mitre Corporation

February 2, 2006

Miscellaneous Capabilities for IP Network Infrastructure
draft-ietf-opsec-misc-cap-00.txt

Status of this Memo

This Internet-Draft will expire on August 6, 2006.

Copyright Notice

Abstract

The Framework for Operational Security Capabilities [11] outlines the proposed effort of the IETF OPSEC working group. This includes producing a series of drafts 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. Current plans include separate capabilities documents for Packet Filtering; Event Logging;

In-Band and Out-of-Band Management; Configuration and Management Interfaces; AAA; and Documentation and Assurance. This document describes some additional miscellaneous capabilities which do not fit into any of these specific catagories, and whose descriptions are brief enough that it does not seem appropriate to create a separate document for each.

Operational Security Current Practices [12] lists current operator practices related to securing networks. This document lists miscellaneous capabilities needed to support those practices.

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 may be cited. Special considerations are discussed as appropriate listing operational and resource constraints, limitations of current implementations, tradeoffs, etc.

Table of Contents

1. Introduction

This document is defined in the context of [11] and [12].

The Framework for Operational Security Capabilities [11] outlines the proposed effort of the IETF OPSEC working group. This includes producing a series of drafts 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. Current plans include separate capabilities documents for Packet Filtering; Event Logging; In-Band and Out-of-Band Management; Configuration and Management Interfaces; AAA; and Documentation and Assurance. This document describes some additional miscellaneous capabilities which do not fit into any of these specific catagories, and whose descriptions are brief enough that it does not seem appropriate to create a separate document for each.

Operational Security Current Practices [12] 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 [14]

EDITORS NOTE: This is an early draft. Additional work will be needed to further refine the listed practices, to respond to comments, and to further align the supported practices with the practices listed in [12]. Editor's notes listed in this document are intended to be removed prior to final publication.

1.1. Threat model

The capabities listed in this document are intended to aid in preventing or mitigating the threats outlined in [11] and [12].

1.2. Capabilities versus Requirements

Capabilities may or may not be requirements. That is a local determination that must be made by each operator with reference to the policies that they must support. It is hoped that this document, together with [12] will assist operators in identifying their security capability requirements and communicating them clearly to vendors.

1.3. 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 [CurPrc] 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. See rfc3631 [13]. 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, tradeoffs, etc.

1.4. Terms Used in this Document

The following terms are used in this document. These definitions are taken from rfc3871 [14].

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 [9] and rfc1918 [3].

Martian

Per rfc1208 [1] "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 [10].

Service

A number of requirements 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.

Single-Homed Network.

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

* There is only one upstream connection

* Routing is symmetric.

See rfc3704 [10] for a discussion of related issues and mechanisms for multihomed 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".

1.5. 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 [4].

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

NOTE: This document defines capabilities. This document does not define requirements, and there is no requirement that any particular capability be implemented or deployed. The use of the terms MUST, SHOULD, and so on are in the context of each capability in the sense that if you conform to any particular capability then you MUST or SHOULD do what is specified for that capability, but there is no requirement that you actually do conform to any particular capability.

EDITOR'S NOTE: An earlier contribution towards this document (draft-callon-misc-caps-00.txt) included a section on route filtering. This section has been removed since it seems likely that there may be a new document proposed giving significant more text which includes this topic. It is intended and expected that some opsec document will contain a description of route filtering capabilities. details are tbd.

2. IP Stack Capabilities

EDITOR'S NOTE: This is taken from section 2.5 of RFC3871.

2.1. Ability to Identify All Listening Services

Capability.

The vendor MUST:

* Provide a means to display all services that are listening for network traffic directed at the device from any external source.

* Display the addresses to which each service is bound.

* Display the addresses assigned to each interface.

* Display any and all port(s) on which the service is listing.

* Include both open standard and vendor proprietary services.

Supported Practices.

This information is necessary to enable a thorough assessment of the security risks associated with the operation of the device (e.g., "does this protocol allow complete management of the device without also requiring authentication, authorization, or accounting?"). The information also assists in determining what steps should be taken to mitigate risk (e.g., "should I turn this service off?")

Current Implementations.

tbd.

Considerations.

If the device is listening for SNMP traffic from any source directed to the IP addresses of any of its local interfaces, then

this requirement could be met by the provision of a command which displays that fact.

2.2. Ability to Disable Any and All Services

Capability.

The device MUST provide a means to turn off any "services" (see section 1.4.1).

Supported Practices.

The ability to disable services for which there is no operational need will allow administrators to reduce the overall risk posed to the device.

As an example of how this is used, many service providers restrict which network management protocols may be used to access the device (see section 2.2 of [12]).

Current Implementations.

tbd.

Considerations.

Processes that listen on TCP and UDP ports would be prime examples of services that it must be possible to disable.

2.3. Ability to Control Service Bindings for Listening Services

Capability.

The device MUST provide a means for the user to specify the bindings used for all listening services. It MUST support binding to any address or net-block associated with any interface local to the device. This must include addresses bound to physical or non-physical (e.g., loopback) interfaces.

Supported Practices.

It is a common practice among operators to configure "loopback" pseudo-interfaces to use as the source and destination of management traffic. These are preferred to physical interfaces because they provide a stable, routable address. Services bound to the addresses of physical interface addresses might become unreachable if the associated hardware goes down, is removed, etc.

This requirement makes it possible to restrict access to management services using routing. Management services may be bound only to the addresses of loopback interfaces. The loopback interfaces may be addressed out of net-blocks that are only routed between the managed devices and the authorized management networks/hosts. This has the effect of making it impossible for anyone to connect to (or attempt to DoS) management services from anywhere but the authorized management networks/hosts.

It also greatly reduces the need for complex filters. It reduces the number of ports listening, and thus the number of potential avenues of attack. It ensures that only traffic arriving from legitimate addresses and/or on designated interfaces can access services on the device.

Current Implementations.

tbd.

Considerations.

If the device listens for inbound SSH connections, this requirement means that it should be possible to specify that the device will only listen to connections destined to specific addresses (e.g., the address of the loopback interface) or received on certain interfaces (e.g., an Ethernet interface designated as the "management" interface). It should be possible in this example to configure the device such that the SSH is NOT listening to every address configured on the device. Similar effects may be achieved with the use of global filters, sometimes called "receive" or "loopback" ACLs, that filter traffic destined for the device itself on all interfaces.

2.4. Ability to Control Service Source Addresses

Capability.

The device MUST provide a means that allows the user to specify the source addresses used for all outbound connections or transmissions originating from the device. It SHOULD be possible to specify source addresses independently for each type of outbound connection or transmission. Source addresses MUST be limited to addresses that are assigned to interfaces (including loopbacks) local to the device.

Supported Practices.

This allows remote devices receiving connections or transmissions to use source filtering as one means of authentication. For example, if SNMP traps were configured to use a known loopback address as their source, the SNMP workstation receiving the traps (or a firewall in front of it) could be configured to receive SNMP packets only from that address.

Current Implementations.

tbd.

Considerations.

The operator may allocate a distinct block of addresses from which all loopbacks are numbered. NTP and syslog can be configured to use those loopback addresses as source, while SNMP and BGP may be configured to use specific physical interface addresses. This would facilitate filtering based on source address as one way of rejecting unauthorized attempts to connect to peers/servers.

Care should be taken to assure that the addresses chosen are routable between the sending and receiving devices, (e.g., setting SSH to use a loopback address of 10.1.1.1 which is not routed between a router and all intended destinations could cause problems).

Note that some protocols, such as SCTP [8], can use more than one IP address as the endpoint of a single connection.

Also note that rfc3631 [13] lists address-based authentication as an "insecurity mechanism". Address based authentication should be replaced or augmented by other mechanisms wherever possible.

2.5. Support Automatic Anti-Spoofing for Single-Homed Networks

Capability.

The device MUST provide a means to designate particular interfaces as servicing "single-homed networks" (see Section 1.4.1) and MUST provide an option to automatically drop "spoofed packets" (Section 1.4.1) received on such interfaces where application of the current forwarding table would not route return traffic back through the same interface. This option MUST work in the presence of dynamic routing and dynamically assigned addresses.

Supported Practices.

See section 3 of rfc1918 [3], sections 5.3.7 and 5.3.8 of rfc1812 [2], and rfc2827 [6].

Current Implementations.

This requirement could be satisfied in several ways. It could be satisfied by the provision of a single command that automatically generates and applies filters to an interface that implements anti-spoofing. It could be satisfied by the provision of a command that causes the return path for packets received to be checked against the current forwarding tables and dropped if they would not be forwarded back through the interface on which they were received.

Considerations.

See rfc3704 [10].

This requirement only holds for single-homed networks. Note that a simple forwarding table check is not sufficient in the more complex scenarios of multi-homed or multi-attached networks, i.e., where the traffic may be asymmetric. In these cases, a more extensive check such as Feasible Path RPF could be very useful.

2.6. Support Automatic Discarding of Bogons and Martians

Capability.

The device MUST provide a means to automatically drop all "bogons" (Section 1.4.1) and "martians" (Section 1.4.1). This option MUST work in the presence of dynamic routing and dynamically assigned addresses.

Supported Practices.

These sorts of packets have little (no?) legitimate use and are used primarily to allow individuals and organization to avoid identification (and thus accountability) and appear to be most often used for DoS attacks, email abuse, hacking, etc. In addition, transiting these packets needlessly consumes resources and may lead to capacity and performance problems for customers.

See section 3 of rfc1918 [3], sections 5.3.7 and 5.3.8 of rfc1812 [2], and rfc2827 [6].

Current Implementations.

This requirement could be satisfied by the provision of a command that causes the return path for packets received to be checked against the current forwarding tables and dropped if no viable return path exists. This assumes that steps are taken to assure that no bogon entries are present in the forwarding tables.

Considerations.

See rfc3704 [10].

This requirement only holds for single-homed networks. Note that a simple forwarding table check is not sufficient in the more complex scenarios of multi-homed or multi-attached networks, i.e., where the traffic may be asymmetric. In these cases, a more extensive check such as Feasible Path RPF could be very useful.

2.7. Support Counters for Dropped Packets

Capability.

The device MUST provide accurate, per-interface counts of spoofed packets dropped in accordance with Section 3.5 and Section 3.6.

Supported Practices.

Counters can help in identifying the source of spoofed traffic.

Current Implementations.

Generally the hardware that is required to drop packets includes specific support for counters. Details vary greatly based on the wide variety of data plane hardware implementations.

Considerations.

An edge router may have several single-homed customers attached. When an attack using spoofed packets is detected, a quick check of counters may be able to identify which customer is attempting to send spoofed traffic.

3. Performance and Prioritization

EDITOR'S NOTE: This section is taken from section 2.15 and a slightly expanded section 2.2.5 of RFC3871.

3.1. Security Features Should Have Minimal Performance Impact

Capability.

Security features specified by the requirements in this document and related OPSEC documents SHOULD be implemented with minimal impact on performance. Other sections of this document or other OPSEC capabilities documents may specify different performance requirements (e.g., "MUST"s).

Supported Practices.

Security features which significantly impact performance may leave the operator with no mechanism for enforcing appropriate policy.

Current Implementations.

Here again how this is implemented depend upon the details of the hardware. In some cases this may require using faster processors than would otherwise be needed, using operating systems that allow resources to be guaranteed to particular processes, or using parallel hardware. There is a very wide range of possible implementations that are possible in order to ensure that security features can be turned on with minimal or no performance impact.

Considerations.

If the application of filters is known to have the potential to significantly reduce throughput for non-filtered traffic, there will be a tendency, or in some cases a policy, not to use filters.

Assume, for example, that a new worm is released that scans random IP addresses looking for services listening on TCP port 1433. An operator might want to investigate to see if any of the hosts on their networks were infected and trying to spread the worm. One way to do this would be to put up non-blocking filters counting and logging the number of outbound connection 1433, and then to block the requests that are determined to be from infected hosts. If any of these capabilities (filtering, counting, logging) have the potential to impose severe performance penalties, then this otherwise rational course of action might not be possible.

Requirements for which performance is a particular concern include: filtering, rate-limiting, counters, logging and anti-spoofing.

3.2. Prioritization of Management Functions

Capability.

Management functions SHOULD be processed at higher priority than non-management traffic. This SHOULD include ingress, egress, internal transmission, and processing. This SHOULD include at least protocols used for configuration, monitoring, configuration backup, logging, time synchronization, and authentication.

Supported Practices.

Certain attacks (and normal operation) can cause resource saturation such as link congestion, memory exhaustion or CPU overload. In these cases it is important that management functions be prioritized to ensure that operators have the tools needed to recover from the attack.

Current Implementations.

Here again how this is implemented depend upon the details of the hardware. There is a very wide range of possible implementations that are possible in order to give priority to management functions. This requirement can potentially implement any of processing, memory, choice of operating system or other software architecture issues, as well as internal and external data transmission.

Considerations.

Imagine a service provider with 1,000,000 DSL subscribers, most of whom have no firewall protection. Imagine that a large portion of these subscribers machines were infected with a new worm that enabled them to be used in coordinated fashion as part of large denial of service attack that involved flooding. It is entirely possible that without prioritization such an attack would cause processor saturation or other internal resource saturation on routers causing the routers to become unmanageable. A DoS attack against hosts could therefore become a DoS attack against the network.

Prioritization is not a panacea. Control packets may not make it across a saturated link. This requirement simply says that the device should prioritize management functions within its scope of control (e.g., ingress, egress, internal transit, processing). To the extent that this is done across an entire network, the overall effect will be to ensure that the network remains manageable.

3.3. Prioritization of Routing Functions

Capability.

Routing functions SHOULD be processed at higher priority than user data traffic. This SHOULD include ingress, egress, internal transmission, and processing. This SHOULD include all packets necessary for routing protocol operation, and specifically MUST include priority processing of routing HELLO packets for BGP, IS-IS, and OSPF.

Supported Practices.

Certain attacks (and normal operation) can cause resource saturation such as link congestion, memory exhaustion or CPU overload. In these cases it is important that routing functions be prioritized to ensure that the network continues to operate (for example, that routes can be computed in order to allow management traffic to be delivered). For many routing protocols the loss of HELLO packets can cause the protocol to drop adjacencies and/or to send out additional routing packets, potentially adding to whatever congestion may be causing the problem.

Current Implementations.

Here again how this is implemented depend upon the details of the hardware. There is a very wide range of possible implementations that are possible in order to give priority to routing functions. This requirement can potentially implement any of processing, memory, choice of operating system or other software architecture issues, as well as internal and external data transmission.

Considerations.

If routing HELLO packets are not prioritized, then it is possible during DoS attacks or during severe network congestion for routing protocols to drop HELLO packets, causing routing adjacencies to be lost. This in turn can cause overall failure of a network. A DoS attack against hosts can therefore become a DoS attack against the network.

Prioritization within routers is not a panacea. Routing update packets may not make it across a saturated link (thus for example it may also be desirable to prioritize routing packets for transmission across link layer devices such as Ethernet switches). This requirement simply says that the device should prioritize routing functions within its scope of control (e.g., ingress,

egress, internal transit, processing). To the extent that this is done across an entire network, the overall effect will be to ensure that the network continues to operate.

3.4. Resources used by IP Multicast

Capability.

Routers SHOULD provide some mechanism(s) to allow the control plane resources used by IP multicast, including processing and memory, to be limited to some level which is less than 100% of the total available processing and memory. The maximum limit of resources used by multicast MAY be configurable. Routers SHOULD also provide a mechanism(s) to allow the amount of link bandwidth consumed by IP multicast on any particular link to be limited to some level which is less than 100% of total available bandwidth on that link.

Supported Practices.

IP multicast has characteristics which may potentially impact the availability of IP networks. In particular, IP multicast requires that routers perform control plane processing and maintain state in response to data plane traffic. Also, the use of multicast implies that a single packet input into the network can result in a large number of packets being delivered throughout the network. Also, it is possible in some situations for a multicast traffic to *both* enter a loop, and also be delivered to some destinations (implying that many copies of the same packet could be delivered).

Current Implementations.

tbd.

Considerations.

If the amount of resources used by multicast are not limited, then it is possible during an attack for multicast to consume potentially as much as 100% of available memory, processing, or bandwidth resources, thereby causing network problems.

4. Security Features Must Not Cause Operational Problems

EDITOR'S NOTE: This is taken from section 2.14 of RFC3871.

Capability.

The use of security features specified by the requirements in this document SHOULD NOT cause severe operational problems.

Supported Practices.

Security features which cause operational problems are not useful and may leave the operator with no mechanism for enforcing appropriate policy.

Current Implementations.

Again this capability potentially impacts many aspects of the implementation.

Considerations.

Some examples of severe operational problems include:

* The device crashes.

* The device becomes unmanageable.

* Data is lost.

* Use of the security feature consumes excessive resources (CPU, memory, bandwidth).

Determination of compliance with this requirement involves a level of judgement. What is "severe"? Certainly crashing is severe, but what about a %5 loss in throughput when logging is enabled? It should also be noted that there may be unavoidable physical limitations such as the total capacity of a link.

5. Security Considerations

General

Security is the subject matter of this entire document. This document lists device capabilities intended to improve the ability of the network to withstand security threats. Operational Security Current Practices [12] defines the threat model and practices, and lists justifications for each practice.

6. Acknowledgements

The authors gratefully acknowledge the contributions of:

- o xxx, yyy, ...

- o The MITRE Corporation for supporting development of this document. NOTE: An 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 authors.

- o We note that there are many people from multiple network operators who have contributed to the OPSEC effort, but who wish to remain anonymous. We would like to thank them for their considerable help.

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

- o Apologies to those who commented on/contributed to the document and were not listed.

7. References

7.1. Normative References

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

[2]   Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.

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

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

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

[6]   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.

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

[8] Stone, J., Stewart, R., and D. Otis, "Stream Control Transmission Protocol (SCTP) Checksum Change", RFC 3309, September 2002.

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

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

7.2. Informational References

[11] Jones, G., "Framework for Operational Security Capabilities for IP Network Infrastructure", draft-ietf-opsec-framework-01 (work in progress), October 2005.

[12] Kaeo, M., "Operational Security Current Practices", draft-ietf-opsec-current-practices-02 (work in progress), October 2005.

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

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

Authors' Addresses

Ross W. Callon
Juniper Networks
10 Technology Park Drive
Westford, MA 01886
USA

Email: rcallon@juniper.net

George Jones
The Mitre Corporation
7515 Colshire Drive
McLean, Virginia, VA 22102-7508
USA

Email: gmjones@mitre.org

Intellectual Property Statement

Disclaimer of Validity

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.

Copyright Statement

Acknowledgment

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