Opsec Working Group
Internet Draft
Expires: January

Z. Ye
M. Fuyou
Huawei Technologies
R. Callon
Juniper Networks
June d, yyyy

Routing Control Plane Security Capabilities
draft-zhao-opsec-routing-capabilities-01.txt

Status of this Memo

Copyright Notice

Abstract

The document lists the security capabilities needed for the routing control plane of an IP infrastructure to support the practices defined in Operational Security Current Practices [OSCP]. In particular this includes capabilities for route filtering and for authentication of routing protocol packets. GMJ> isn't authentication of routing protocols the subject of entire working

groups (RPSEC ?). Are we duplicating ? Should we cite (to show awareness ?)

Conventions used in this document

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 [RFC-2119]

Table of Contents

1. Introduction

This document is defined in the context of [FRAMEWORK] and [OSCP].

The Framework for Operational Security Capabilities [FRAMEWORK] outlines the effort of the IETF OPSEC working group. This includes producing a series of drafts to codify knowledge gained through operational experience about capabilities that are needed to securely deploy and operate managed network elements providing transit services at the data link and IP layers. GMJ> "and network layers" ?

This document lists the security capabilities needed for the routing control plane of IP infrastructure to support the practices defined in Operational Security Current Practices [OSCP]. In particular this includes capabilities for route filtering and for authentication of routing protocol packets.

Note that this document lists capabilities that can reasonably be expected to be currently deployed in the context of existing standards. Extensions to existing protocol standards and development of new protocol standards are outside of the scope of this effort. The preferred capabilities needed for securing the routing infrastructure may evolve over time.

There will be other capabilities which are needed to fully secure a router infrastructure. For example, network management of devices must be secured in order to prevent unauthorized access to or denial of service to the device [NMASC]. The reader should refer to [FRAMEWORK] for a more complete list of documents describing operational capabilities for network and link layer devices supporting IP Network Infrastructure.

Operational Security Current Practices [OSCP] defines the goals, motivation, scope, definitions, intended audience, threat model, potential attacks and give justifications for each of the practices.

1.1. Threat model

The capabilities listed in this document are intended to aid in preventing or mitigating the threats outlined in [FRAMEWORK] and [OSCP].

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 [OSCP] will assist operators in identifying their security capability requirements and communicating them clearly to vendors.

The capabilities described in this document follow the format outlined in section 1.7 of [FRAMEWORK].

1.3. Packet Filtering versus Route Filtering

It is useful to make a distinction between Packet Filtering versus Route Filtering.

The term "packet filter" is used to refer to filters that routers apply to network layer packets that they are forwarding. In general packet filters are based on contents of the network (IP) and transport (TCP,UDP) layers, and are mostly stateless, in the sense that whether or not a filter applies to a particular packet is a function of that packet (including the contents of IP and transport layer headers, size of packet, incoming interface, and similar characteristics), but does not depend upon the contents of other packets which might be part of the same stream (and thus which may also be forwarded by the same router). One commonly minor exception to the "stateless" nature of packet filters is that packets that fit a particular filter may be counted and/or rate limited (the act of counting therefore represents a very simple "state" associated with the filter).

Because of the simplicity and stateless nature of packet filters, they can typically be implemented with very high performance. It is not unusual for them to be implemented on line cards and to perform at or near full line rate. For this reason they are very useful to counter very high bandwidth attacks, such as large DDoS attacks.

Packet filtering capabilities are outside of the scope of this document. A detailed description of packet filtering capabilities can be found in "Filtering Capabilities for IP Network Infrastructure" [FILTER].

The Term "route filter" is used to refer to filters that routers apply to the content of routing protocol packets that they are either sending or receiving. Typically these therefore occur at the application layer (although which route filters are applied to a particular packet may be a function of network layer information, such as what interface the packet is received on, or the source address for the packet -- indicating the system that transmitted the packet).

Route filters are typically implemented in some sort of processor. GMJ> How else are they implemented ? ASICs ? This seems to me to be a redundant statement. In many cases the total bandwidth which can be received by the processor is considerably less than the sum of the rate that packets may be received on all interfaces to a router. Therefore in general route filters cannot handle the same

bandwidth as packet filters. Route filters are however very useful in that they can be applied to the contents of routing packets. GMJ> Implication (maybe state it) "This is manageable because the volume of routing packets is much lower than the total volume of incoming packets"

1.4. RFC 2119 Keywords

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

The use of the RFC 2119 keywords is an attempt, by the 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.

2. Route Filtering Capabilities

2.1. Route Filtering of External Routing Protocols

GMJ> formatting nit. Can you do something like:

Capability. The device MUST…

I can help you convert to xml2rfc if you like. It handles the formatting automatically.

Capability.
The device MUST provide a means to filter routing updates for all protocols used to exchange external routing information. Generally this includes BGP [RFC4271], as well as static routes.

Supported Practices.
See [RFC3013] and section 3.2 of [RFC2196] and section 2.5 of [OSCP].
GMJ> Formatting nit. I would prefer [draft-ietf-opsec-xxx-##]

Current Implementations.
Typically BGP implementations allow operators to apply a variety of filters to restrict which incoming updates are accepted from BGP peers, as well as to limit which updates are sent out to BGP peers.

GMJ> Can you give hard examples, maybe from Cisco or Juniper ?

In general packet filters may be used in the following ways:

GMJ> The following are phrased as requirements (MUST, etc), not descriptions of implementations. These may be the things you want, but you need to phrase them differently. Also, if these are different capabilities, then they should be split out into separate sections.

- ——Routers MUST allow operators to configure route filters which restrict which routes are accepted from other peer routers. Route filters MUST be capable of being individually configured on a per-neighbor basis.

For example:

Capability:

Router allows operators to configure route filters which restrict…

Supported practices

Currently operators use this capability to…

Current Implementations

Several vendors currently support the application of ACL style filters to received routes. Examples include Cisco's FOO and Junipers BAR…

Considerations

Be advised that if you apply filters incorrectly the you may stop routing traffic to customers…there may be heavy performance penalty, etc., etc.

- Routers MUST allow operators to configure route filters which restrict which routes are sent to other peer routers. Route filters MUST be capable of being individually configured on a per-neighbor basis.

- Routers SHOULD allow operators to configure whether Outbound Route Filters [ORF] are accepted from other peer routers. This SHOULD be configurable on a per-neighbor basis.

- Routers SHOULD allow operators to configure which (if any) Outbound Route Filters [ORF] are sent to other peer routers. This SHOULD be configurable on a per-neighbor basis.

In general route filters determine whether a route is accepted from or sent to a neighboring router. Filters MAY be based upon any combination of route attributes, such as:

- ~~Accept or reject s~~Specific route prefixes

This may include a list of specific prefixes to be accepted or rejected. This may alternately include a list of prefixes, such that more specific (longer) prefixes which are included in the ~~shorter (~~more inclusive (shorter)~~)~~ prefix are accepted, rejected, or summarized into the shorter prefix.

- Maximum length of route prefix

- Maximum number of routes to be accepted from a particular peer router

If too many routes are sent, then the router may reset the BGP session, or may reject excess routes. In either case the failure event should be logged.

- Restrictions on the AS_PATH

Restrictions on the contents of the AS PATH are frequently used: for example if you get a prefix from AS X, then you might want to make sure that X is in the AS PATH.

- Restrictions on BGP Community and Extended Community

Route redistribution is used to exchange routing information between different protocols. Although route redistribution bridges between different route domains and improves the flexibility of routing system, it may lead to looping or black hole as well.

- Routers SHOULD provide method to limit the scope of route redistribution between different route protocols. Unfiltered redistribution SHOULD be forbidden.

Considerations.
Operators may wish to ignore advertisements for routes to specially used addresses, such as private addresses, reserved addresses and multicast addresses, etc. The up-to-date allocation of IPv4 address space can be found in [IANA].

2.2. Route Filtering Within an IGP Area

GMJ> Are these capabilities ? Looks more like current practice. Either get Merike to add it to the capabilities draft and cite it or add the supported practices under your own capability. From the framework:

Supported Practices (why)

The Supported Practice section cites practices described in [I-D.ietf-opsec-current-practices] that are supported by this capability. The need to support the cited practices provides the justification for the feature.

In a few cases, practices not listed in [I-D.ietf-opsec-current-practices] may be listed at the end of the capability document and cited as justification for a capability. This may be necessary if a practice becomes common after [I-D.ietf-opsec-current-practices]

is finished or if there is widespread consensus that the practice would improve security but it is not, for whatever reason, in widespread deployment.

This section describes route filtering as it may be applied to OSPF[RFC 2328] and IS-IS [RFC1195] when used as the interior routing protocol (Internal Gateway Protocol or "IGP") used within a routing domain. Route filtering with RIP [RFC2453] is TBD.

2.2.1. Route Filtering Within an IGP Area

A critical design principle of OSPF and IS-IS is that each router within an area has the same view of the topology, thereby allowing consistent routes to be computed by all routers within the area. For this reason, all properly authentication (if applicable) routing topology advertisements (Link State Advertisements or LSAs in OSPF, or Link State Packets or LSPs in IS-IS) are flooded unchanged throughout the area. Route filtering within an OSPF or IS-IS area is therefore not appropriate.

2.2.2. Route Filtering Between IGP Areas

Capability

GMJ> Again, capabilities should be of the form:

Capability:
The device is able to…

This looks like a practice.

It is normal when passing routes into the backbone area (area 0.0.0.0 in OSPF, or the level 2 backbone in IS-IS) for routes to be summarized, in the sense that multiple more specific (longer) address prefixes that are reachable in an area may be summarized into a smaller number of less specific (shorter) address prefixes. This provides important scaling improvements, but is generally not

primarily intended to aid in security and is therefore outside of the scope of this document.

Gmj> now you're back to requirements/cabilities.

Routers MAY implement the capability to allow the network operator the option of configuring route filters which restrict which routes (ie, address prefixes) are advertised into areas from outside of the area (ie, from other OSPF or IS-IS areas).

Supported Practices
TBD.

Current Implementations
TBD.

Considerations
If filters are used which restrict the passing of routes between IGP areas, then this may result in some addresses being unreachable from some other areas within the same routing domain.

2.3. Ability to Filter Routing Update by TTL

Capability GMJ> Good!
The device should provide a means to filter route packets based on the value of the TTL field in the IPv4 header or the Hop-Limit field in the IPv6 header.

Gmj> this is a supported practice.

For example, in many cases routing protocol packets should only be arriving from immediate neighboring routers. In these cases, packets SHOULD be dropped if the TTL is not equal to 255. In these cases filtering on TTL prevents any system which is not immediately physically adjacent to a router from sending that router spoofed routing packets.

Note that "Filtering Capabilities for IP Network Infrastructure" [FILTER] specifies:

> Capability.
>
> The filtering mechanism supports filtering based on the value(s) of any portion of the protocol headers for IP, ICMP, UDP and TCP.

The ability to filter based on TTL is therefore a packet filtering capability which is already implicitly covered by the capabilities listed in [FILTER]. Since this capability is particularly important for routing protocols, we felt that it is worth mentioning here.

Supported Practices.
See [OSCP] section 2.5.7 Good ! You could move your additional practices here or have Merike add them.

Current Implementations.
GMJ: this is more of a how. How about "several vendors provide mechanisms for specifying TTL decrements…the interfaces provided are command line…
When a router forwards a packet, it will decrement the TTL value (Hop-Limit for IPv6) of the packet by one. Thus, TTL spoofing is considered nearly impossible. Furthermore, the vast majority of routing peers are adjacent. This capability is therefore quite useful, and is widely implemented in routers.

Considerations.
There will be situations in which the distance to the neighboring router is more than one hop away. This for example is common for I-BGP.

3. Route Flap Dampening

GMJ> If you have definitions, let's add them to the framework.

"Route flap" means that a route's state changes from up to down or down to up. In some cases a route may come up and go down multiple times in a short period of time (for example due to an unstable link somewhere in the global Internet). When repeated route flapping occurs, the route process has to insert or delete an item and the advertised update. If large amounts of routes continue to go up and down multiple times in a short time period this may result in significant load on CPUs and could result in DoS (whether intentional or not).

Capabilities.
GMJ> again, using the MUST/MAY/SHOULD are appropriate for requirements. Here we are talking about capabilities. Rephrase as: "The devide provides the ability to…" (do the same throughout).

The device MUST provide ability to dampen route flap.

Route Flap dampening MUST be configurable. For example, some operators may want to change the timers, and others may want to turn it off altogether.

Supported Practices.

The function to dampen route flap may enhance the stability of routing system and minimize the influence of flapping. It is useful to counter against some DoS attacks.

~~Route flapping dampening is the primary mechanism to mitigate the influence caused by flapping.~~

Current Implementations.

In BGP, route flapping dampening is the primary mechanism to mitigate the influence caused by flapping. Most of current implements support this capability.

~~The function to dampen route flap may enhance the stability of routing system and minimize the influence of flapping. It is useful to counter against some DoS attacks.~~

Consideration.

None

4. Authentication of Routing Protocols

As mentioned in [RFC4272], the authentication mechanism specified in [TCPMD5] can counter several types of attacks on BGP, such as message insertion, modification, deletion, man-in-the-middle, and some types of DOS attack. Even though an assailant can guess TCP sequence numbers of a BGP session, he will fail to launch the attack mentioned above. Most other routing protocols adopt similar authentication mechanism.

Capabilities.

- MUST provide a mechanism through which operators can manually configure a sequence of keys on peer systems

- MUST provide a mechanism through which peer systems can transition from one key to another based upon system time

- MUST provide a mechanism through which peer systems can transition from one key to another without resetting the neighboring session

- MUST support authentication alogorithms that are stronger than MD5 (e.g., CMAC-AES-128-96, HMAC-SHA-1-96).

- SHOULD support automatic generation and encrypted distribution of key material.

Supported Practices.
See [OSCP] section 2.5.7.~~Route authentication ([PRACTICES] section 2.5.7)~~

Current Implementations.
~~Because of its simpleness and high efficiency,~~ [TCPMD5] is deployed widely in BGP. Other routing protocols, such as OSPF, adopt similar technology.

In most of current implements, neither the authentication mechanism nor key can be negotiated. An operator has to configure it manually.

Consideration.
~~Some protocols, such as~~ OSPF~~,~~ supports plain text authentication which is not able to counter attacks above. Most OSPF implementations also support MD5 authentication. In this section the authentication mechanism refers to the technology using cryptographic hash functions.

In order to counter key-guessing attack, ~~As mentioned above, we have to share and rekey the secret manually. Thus the most important consideration about authentication is secret key management.~~

~~W~~when manual key management is used, ~~a operator can follow the advices below:~~a device SHOULD support a proper length of a key

~~1. A key SHOULD have a proper length in order to counter key guessing attack~~.

~~2. It is not recommended to select a common word to be a long time key, because it is vulnerable to a dictionary attack.~~

~~3. A key SHOULD only be shared between peers in case of exposure.~~

~~4. A key SHOULD have a proper lifetime and rekeying operations SHOULD be performed periodically.~~

5. Security Considerations

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 [OSCP] defines the threat model and practices, and lists justifications for each practice.

6. Acknowledgements

The authors gratefully acknowledge the contributions of:

o tbd, xxx, yyy, ...

o We would like to thank Ron Bonica and Pat Cain for their helpful comments and suggestions.

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

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

[RFC1812] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.

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

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

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

[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.

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

[RFC3309] Stone, J., Stewart, R., and D. Otis, "Stream Control Transmission Protocol (SCTP) Checksum Change", RFC 3309, September 2002.

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

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

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

[RFC1195] R. Callon, “Use of OSI IS-IS for Routing in TCP/IP and Dual Environments”, RFC1195, December 1990.

[RFC2328] J. Moy, “OSPF Version 2”, RFC2328, April 1998.

[RFC2453] G. Malkin, “RIP Version 2”, RFC2453, November 1998.

[RFC3682] V. Gill, J. Heasley, D. Meyer, “The Generalized TTL Security Mechanism (GTSM)”, RFC3682, February 2004.

[TCPMD5] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC 2385, August 1998.

7.2. Informative References

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

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

[FILTER] Morrow, C., "Filtering Capabilities for IP Network Infrastructure", draft-ietf-opsec-filter-caps-00 (work in progress), October 2005.

[NMASC] Bonica, R. and S. Ahmed, "Network Management Access Security Capabilities", draft-bonica-opsec-nmasc-00 (work in progress), October 2005.

[IANA] IANA, “INTERNET PROTOCOL V4 ADDRESS” SPACE,http://www.iana.org/assignments/ipv4-address-space

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

GMJ> No citation to 3871 in body . You should only list references that have citations. Again, xml2rfc (xml.resource.org) does a very nice job of this. Let me know if you want me to hack this into xml2rfc format.

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

[ORF] Enke Chen, Yakov Rekhter, “Cooperative Route Filtering Capability for BGP-4”, draft-ietf-idr-route-filter-13.txt, (work in progress), March 2006.

[RFC4272] S. Murphy., “BGP Security Vulnerabilities Analysis”, RFC 4272, January 2006

Author's Addresses

Zhao Ye
Huawei Technologies
No.3, Xinxi Road, Shangdi Information Industry Base
Haidian District, Beijing City 100085
Email: yezhao@huawei.com

Miao Fuyou
Huawei Technologies
No.3, Xinxi Road, Shangdi Information Industry Base
Haidian District, Beijing City 100085
Email: miaofy@huawei.com

Ross W. Callon
Juniper Networks
10 Technology Park Drive
Westford, MA 01886
USA
Email: rcallon@juniper.net

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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