

OPSEC Working Group
Internet-Draft
Intended status: Best Current
Practice
Expires: October 7, 2007

Y. Zhao
F. Miao
Huawei Technologies
R. Callon
Juniper Networks
April 5, 2007

Routing Control Plane Security Capabilities
draft-ietf-opsec-routing-capabilities-02.txt

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/lid-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 October 7, 2007.

Copyright Notice

Copyright (C) The IETF Trust (2007).

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. In particular this includes capabilities for route filtering, for authentication of routing protocol packets, and for ensuring resource availability for control functions.

Internet-Draft

routing-capabilities

April 2007

Table of Contents

1.	Introduction	3
1.1.	Threat model	3
1.2.	Format and Definition of Capabilities	3
1.3.	Packet Filtering versus Route Filtering	3
2.	Route Filtering Capabilities	4
2.1.	General Route Filtering Capabilities	4
2.1.1.	Ability to Filter Inbound or Outbound Routes	4
2.1.2.	Ability to Filter Routes by Prefix	5
2.2.	Route Filtering of Exterior Gateway Protocol	6
2.2.1.	Ability to Filter Routes by Route Attributes	6
2.2.2.	Ability to Filter Routing Update by TTL	7
2.2.3.	Ability to Limit the Number of Routes from a Peer	8
2.2.4.	Ability to Limit the Length of Prefixes	9
2.2.5.	Ability to Cooperate in Outbound Route Filtering	9
2.3.	Route Filtering of Interior Gateway Protocols	10
2.3.1.	Route Filtering Within an IGP Area	10
2.3.2.	Route Filtering Between IGP Areas	10
2.4.	Route Filtering during Redistribution	11
3.	Route Authentication Capabilities	11
3.1.	Ability to configure an authentication mechanism	11
3.2.	Ability to support authentication key chains	12
4.	Ability to Damp Route Flap	12
5.	Resource Availability for Router Control Functions	13
5.1.	Ensure Resources for Management Functions	13
5.2.	Ensure Resources for Routing Functions	14
5.3.	Limit Resources used by IP Multicast	15
6.	Security Considerations	16
7.	IANA Considerations	16
8.	Acknowledgements	16
9.	References	17
9.1.	Normative References	17
9.2.	Informative References	17
	Authors' Addresses	18
	Intellectual Property and Copyright Statements	20

1. Introduction

This document is defined in the context of [[I-D.ietf-opsec-framework](#)] and [[RFC4778](#)].

This document lists the security capabilities needed for the routing control plane of IP infrastructure to support the practices defined in [[I-D.ietf-opsec-framework](#)]. 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. [[RFC4778](#)] 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 [[I-D.ietf-opsec-framework](#)].

1.2. Format and Definition of Capabilities

Each individual capability will be defined using the four elements, "Capability", "Supported Practices", "Current Implementations", and "Considerations", as explained in section 1.7 of [[I-D.ietf-opsec-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 the filter that a router applies to network layer packets passing through or destined to it. 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

Zhao, et al.

Expires October 7, 2007

[Page 3]

Internet-Draft

routing-capabilities

April 2007

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 common minor exception to the "stateless" nature of packet filters is that packets that match 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 [[I-D.ietf-opsec-filter-caps](#)].

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. 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.

[2.](#) Route Filtering Capabilities

[2.1.](#) General Route Filtering Capabilities

[2.1.1.](#) Ability to Filter Inbound or Outbound Routes

Capability.

The device provides the ability to filter which routes may be received with [[RFC4271](#)], as well as the ability to filter which routes are announced into each routing protocol.

Zhao, et al.

Expires October 7, 2007

[Page 4]

Internet-Draft

routing-capabilities

April 2007

Supported Practices.

See [section 2.4.2 of \[RFC4778\]](#).

It is a beneficial practice to configure routing filters in both directions, which will counter potential misconfiguration in either peer. Also, incoming route filters will prevent a deliberate attacker from injecting invalid routing information.

Current Implementations.

The unicast routing protocols used with IP can be classified into path vector routing protocols (such as BGP), distance vector protocols (such as [[RFC2453](#)]) and link state protocols (such as [[RFC2328](#)] and [[RFC1195](#)]). Because of differences in the protocols, route filters will affect them in different ways.

Take BGP for example, an implementation may check a received route against inbound filters to determine whether to install it into the overall route table or not. Also, it will restrict the routes which will go out to neighbors against outbound route filters.

However, as to link state protocols, such as OSPF, a router maintains a topology database and exchanges link state information with neighbors. Since route filters do not influence the link state database, route filters will only affect which routes are advertised into the routing protocol. That is to say, only inbound route filters are effective on link state protocols.

Most of the routing protocols support methods to configure route filters which permit or deny learning or advertising of specific routes.

Considerations.

None.

[2.1.2.](#) Ability to Filter Routes by Prefix

Capability.

The device supports filtering routes based on prefix.

Supported Practices.

See [section 2.4.2 of \[RFC4778\]](#).

Current Implementations.

Zhao, et al.

Expires October 7, 2007

[Page 5]

Internet-Draft

routing-capabilities

April 2007

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

Considerations.

Operators may wish to ignore advertisements for routes to specific 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 of Exterior Gateway Protocol

An exterior gateway protocol is used to exchange external routing information between different autonomous systems. Since BGP is the most widely used protocol, this section mainly depicts special routing filter capabilities of BGP.

2.2.1. Ability to Filter Routes by Route Attributes

Capability.

The device supports filtering routing updates by route attributes.

Supported Practices.

See [[RFC3013](#)] , [section 3.2](#) of[RFC2196] and [section 2.4.2 of \[RFC4778\]](#).

Current Implementations.

In comparison with other routing protocols, BGP defines various path attributes to describe characteristics of routes. Besides filtering by specific prefixes, BGP also provides capability to use some path attributes to precisely filter routes to determine whether a route is accepted from or sent to a neighboring router.

These filters may be based upon any combination of route attributes, such as:

- * Restrictions on the Content of AS_PATH. Restrictions on the contents of the AS PATH are frequently used: for example, the received AS_PATH may be checked to ensure the sending AS is actually contained in the received AS_PATH.

- * Restrictions on Communities. Implementations could filter received routes based on the set of communities [[RFC1997](#)] or extended communities [[RFC4360](#)].

Considerations.

None.

[2.2.2.](#) Ability to Filter Routing Update by TTL

Capability.

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.

Note that [[I-D.ietf-opsec-filter-caps](#)] 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 Filtering Capabilities. Since this capability is particularly important for BGP, we felt that it is worth mentioning here.

Supported Practices.

See [[RFC4778](#)] [section 2.4.2](#) and [[RFC3682](#)].

Current Implementations.

Most current BGP implementations support this capability to protect BGP sessions.

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.

[2.2.3.](#) Ability to Limit the Number of Routes from a Peer

Capability.

The device should provide a means to configure the maximum number of routes (prefixes) to accept from a peer.

Supported Practices.

Both routing policy misconfiguration and a deliberate attack from a peer may cause too many routes to be sent to a peer which may exhaust the memory resources of the router, introduce routing instability into the overall routing table, or both. Therefore, operators may want to restrict the amount of routes received from a particular peer router through a maximum prefix limitation approach.

Current Implementations.

Most BGP implementations support this capability. If too many routes are sent, then the router may reset the BGP session or may reject excess routes. In either case the device may log the failure event (at a minimum), or shut down the BGP session.

Considerations.

Operators must be cognizant of the need to allow for valid swings in routing announcements between themselves, and as such should always set the max-prefix limit to some agreed upon number plus a sane amount for overhead to allow for these necessary announcement swings. Individual implementations amongst ISPs are unique, and depending on equipment supplier(s) different implementation options are available. Most equipment vendors offer implementation options ranging from just logging excessive prefixes being received to automatically shutting down the session. If the option of reestablishing a session after some pre-configured idle timeout has been reached is available, it should be understood that automatically reestablishing the session may continuously introduce instability into the overall routing table if a policy misconfiguration on the adjacent neighbor is causing the condition. If a serious misconfiguration on a peering neighbor has occurred then automatically shutting down the session and leaving it shut down until being manually cleared is perhaps best and allows for operator intervention to correct as needed.

[2.2.4.](#) Ability to Limit the Length of Prefixes

Capability.

The device has the capability to allow filtering of route updates by prefix length.

Supported Practices.

Some large ISPs declare in their peer BGP policies that they will not accept the announcements whose prefix length is longer than a specific threshold.

Current Implementations.

Most BGP implementations support this capability.

Considerations.

None.

[2.2.5.](#) Ability to Cooperate in Outbound Route Filtering

Capability.

A device provides the capability to allow operators to configure whether Outbound Route Filtering/ORF defined in [[I-D.ietf-idr-route-filter](#)] are accepted from or sent to other peer routers.

Supported Practices

"Outbound Route Filtering" defines a BGP mechanism to reduce the number of BGP updates between BGP peers. It will conserve the resource in both sides of peers, since the BGP speaker will not generate updates that will be filtered and the neighbor router will not process them as well. A router with limited resource may need this feature to prevent overfull routes from peers.

Current Implementations.

ORF may be based on prefix, path attributes. Currently, most implementations support prefix-based ORF.

Considerations.

None.

[2.3.](#) Route Filtering of Interior Gateway Protocols

This section describes route filtering as it may be applied to OSPF and IS-IS when used as the interior gateway protocol (Internal Gateway Protocol or IGP) used within a routing domain.

[2.3.1.](#) 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 authenticated (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.

[2.3.2.](#) Route Filtering Between IGP Areas

Capability.

The device provides the capability to allow the network operator to configure route filters which restrict which routes (i.e., address prefixes) are advertised into areas from outside of the area (i.e., from other OSPF or IS-IS areas).

Supported Practices.

See [section 2.4.2 of \[RFC4778\]](#).

Current Implementations.

Some OSPF/IS-IS implementations support this capability.

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.

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.

[2.4.](#) Route Filtering during Redistribution

Capability.

The device provides a means to filter routes when distributing them between routing protocols or between routing protocol processes running in the single device.

Supported Practices.

Route redistribution bridges between different route domains and improves the flexibility of routing system. This allows for the transmission of reachable destinations learned in one protocol through another protocol. However, without careful consideration it may lead to looping or black holes as well.

Filters are always needed when routes redistributing between IGP and EGP. For example, it is infeasible to inject all Internet routes from EGP to IGPs, since IGPs are not able to deal with such a large number of routes.

Current Implementations.

Most implementations allow applying a filter based on a prefix list to control redistribution.

Considerations

None.

[3.](#) Route Authentication Capabilities

[3.1.](#) Ability to configure an authentication mechanism

Capabilities.

The device has one or more methods to allow an authentication mechanism to be configured for the routing protocol.

Supported Practices.

See [\[RFC4778\] section 2.4.2.](#)

Current Implementations.

[RFC2385] is deployed widely in BGP. Other routing protocols, such as OSPF, adopt similar technology.

Zhao, et al.

Expires October 7, 2007

[Page 11]

Internet-Draft

routing-capabilities

April 2007

Consideration.

In most of current implementations, neither the authentication mechanism nor key can be negotiated. An operator has to configure it manually, which will affect scalability.

As of this writing, MD5 is the only cryptographic hash function used in route authentication. However, recent research revealed weakness of MD5, which means stronger algorithms are necessary.

[3.2.](#) Ability to support authentication key chains

Capabilities.

The device provides a key chain mechanism to update authentication keys of routing protocols.

Supported Practices.

Using a fixed authentication key is vulnerable to a compromise. A key chain is a series of keys which will be used in configured time intervals. A device can transit keys based on system time and configured key chain. In this way, it reduces possibility of leakage of an authentication key.

Current Implementations.

This mechanism is implemented in most routing protocols. Different vendors provide this feature in different routing protocols, such as RIP, OSPF and BGP.

Consideration.

Since the rollover of the key is based on system time on different routers, it requires clock synchronization across the routers.

4. Ability to Damp Route Flap

Capability.

The device provides the capability to damp route flaps.

Supported Practices.

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

Current Implementations.

BGP RFD (Route Flap Damping) of [[RFC2439](#)] defines the primary mechanism in BGP to mitigate the influence caused by flapping. Most of current BGP implementations support this capability.

Other routing protocols may be vulnerable to route flaps as well. Some vendors introduce SPF (shortest path first) algorithm timers in OSPF to control parameters, such as the amount of minimal time between consecutive SPF computations, which may be used to mitigate excessive resource exhaustion caused by link flaps.

Consideration.

[MAO] described a flaw of current BGP RFD standard [RFC2439](#), which shows that route flap damping could suppress relatively stable routes and affect routing convergence.

Since, at the time of this writing, no vendors are known to have fixed this problem, [\[RIPE378\]](#) proposes that, with the current implementations of BGP flap damping, the application of flap damping in ISP networks is not recommended.

[5.](#) Resource Availability for Router Control Functions

[5.1.](#) Ensure Resources for Management Functions

Capability.

This capability specifies that device implementations ensure that at least a certain minimum sufficient level of resources are available for management functions. This may include such resources as memory, processor cycle, data structure and/or bandwidth at ingress to the device, on egress from the device, for internal transmission and processing. This may include at least protocols used for configuration, monitoring, configuration backup, logging, time synchronization, authentication and authorization.

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 resources be available for management functions in order to ensure that operators have the tools needed to recover from the attack.

Zhao, et al.

Expires October 7, 2007

[Page 13]

Internet-Draft

routing-capabilities

April 2007

Current Implementations.

How this is implemented depend upon the details of the device. There are a variety of ways that this may be ensured such as prioritizing management functions in comparison with other functions performed by the device, providing separate queues for management traffic, use of operating systems or other methods that partition resources between processes or functions, and so on.

Consideration.

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 such an attack could in some cases 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.

Guarantee of resources within an individual device is not a panacea. Control packets may not make it across a saturated link. This requirement simply says that the device should ensure resources for 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.

[5.2.](#) Ensure Resources for Routing Functions

Capability.

This capability specifies that a device implementation ensures at least a certain minimum sufficient level of resources are available for routing protocol functions. This may include such resources as memory, processor cycle, data structure and bandwidth at ingress to the device, on egress from the device, for internal transmission, and processing. This may include at least protocols used for routing protocol operation, including resources used for 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 resources be available for the operation of routing protocols in order 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 destabilizing the routing protocol and/or adding to whatever congestion may be causing the problem.

Current Implementations.

How this is implemented depend upon the details of the device. There are a variety of ways that this may be ensured such as prioritizing routing functions in comparison with other functions performed by the device, providing separate queues for routing traffic, use of operating systems or other methods that partition resources between processes or functions, and so on.

Consideration.

For example, 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.

Guaranteeing resources within routers is not a panacea. Routing 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.

[5.3.](#) Limit Resources used by IP Multicast

Capability.

This capability specifies that some mechanism(s) is provided 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. In some cases the maximum limit of resources used by multicast may be configurable. Routers may 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.

None.

Consideration.

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.

[6.](#) 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 defines the threat model and practices, and lists justifications for each practice.

[7.](#) IANA Considerations

This document has no actions for IANA.

[8.](#) Acknowledgements

The authors gratefully acknowledge the contributions of Ron Bonica, Ted Seely, Pat Cain, George Jones, and Russ White etc for their contributed texts, useful comments and suggestions.

9. References

Zhao, et al.

Expires October 7, 2007

[Page 16]

Internet-Draft

routing-capabilities

April 2007

9.1. Normative References

- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", [RFC 1195](#), December 1990.
- [RFC1997] Chandrasekeran, R., Traina, P., and T. Li, "BGP Communities Attribute", [RFC 1997](#), August 1996.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), April 1998.
- [RFC2385] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", [RFC 2385](#), August 1998.
- [RFC2439] Villamizar, C., Chandra, R., and R. Govindan, "BGP Route Flap Damping", [RFC 2439](#), November 1998.
- [RFC2453] Malkin, G., "RIP Version 2", STD 56, [RFC 2453](#), November 1998.
- [RFC3013] Killalea, T., "Recommended Internet Service Provider Security Services and Procedures", [BCP 46](#), [RFC 3013](#), November 2000.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", [RFC 4360](#), February 2006.

9.2. Informative References

- [RFC2196] Fraser, B., "Site Security Handbook", [RFC 2196](#), September 1997.
- [RFC3682] Gill, V., Heasley, J., and D. Meyer, "The Generalized TTL Security Mechanism (GTSM)", [RFC 3682](#), February 2004.
- [RFC4778] Kaeo, M., "Operational Security Current Practices in Internet Service Provider Environments", [RFC 4778](#), January 2007.

[I-D.ietf-opsec-framework]

Jones, G., "Framework for Operational Security Capabilities for IP Network Infrastructure", [draft-ietf-opsec-framework-05](#) (work in progress), April 2007.

[I-D.ietf-opsec-filter-caps]

Zhao, et al.

Expires October 7, 2007

[Page 17]

Internet-Draft

routing-capabilities

April 2007

Morrow, C., "Filtering and Rate Limiting Capabilities for IP Network Infrastructure", [draft-ietf-opsec-filter-caps-06](#) (work in progress), April 2007.

[I-D.ietf-idr-route-filter]

Chen, E. and Y. Rekhter, "Outbound Route Filtering Capability for BGP-4", [draft-ietf-idr-route-filter-16](#) (work in progress), September 2006.

[IANA]

IANA, "INTERNET PROTOCOL V4 ADDRESS SPACE", <http://www.iana.org/assignments/ipv4-address-space> , 2007.

[MAO]

Mao, Z., Govindan, R., Varghese, G., and R. Katz, "Route Flap Damping Exacerbates Internet Routing Convergence", Sigcomm , 2002.

[RIPE378]

RIPE, "Recommendations on Route-flap Damping", RIPE , May 2006.

Authors' Addresses

Zhao Ye
Huawei Technologies
No. 3, Xixi Rd
Shangdi Information Industry Base
Haidian District, Beijing 100085
P. R. China

Email: ye.zhao_ietf@hotmail.com

Miao Fuyou
Huawei Technologies
No. 3, Xinxu Rd
Shangdi Information Industry Base
Haidian District, Beijing 100085
P. R. China

Phone: +86 10 8288 2008
Email: miaofy@huawei.com

Zhao, et al. Expires October 7, 2007 [Page 18]

Internet-Draft routing-capabilities April 2007

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

Email: rcallon@juniper.net

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

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, THE IETF TRUST 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).