

Routing Area WG
Internet-Draft
Intended status: Informational
Expires: July 20, 2007

P. Savola
CSC/FUNET
January 16, 2007

Backbone Infrastructure Attacks and Protections
draft-savola-rtgwg-backbone-attacks-03.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/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 July 20, 2007.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

A number of countermeasures for attacks against service provider backbone network infrastructure have been specified or proposed, each of them usually targeting a subset of the problem space. There has never been a more generic analysis of the actual problems, and which countermeasures are even necessary (and where). This document tries to provide that higher-level view.

Table of Contents

- [1. Introduction](#) [3](#)
- [1.1. Abbreviations](#) [3](#)
 - [1.2. Assumptions](#) [4](#)
 - [1.3. Threat Model](#) [4](#)
- [2. Attack Vectors](#) [5](#)
- [2.1. Lower-layer Attacks](#) [5](#)
 - [2.2. Generic DoS on the Router](#) [5](#)
 - [2.3. Generic DoS on a Link](#) [6](#)
 - [2.4. Cryptographic Exhaustion Attacks](#) [6](#)
 - [2.5. Unauthorized Neighbor or Routing Attacks](#) [6](#)
 - [2.6. TCP RST Attacks](#) [7](#)
 - [2.7. ICMP Attacks](#) [7](#)
- [3. Typical Countermeasures](#) [7](#)
- [3.1. Filtering Addresses in Packets](#) [7](#)
 - [3.2. Filtering Addresses in Routing Updates](#) [8](#)
 - [3.3. GTSM](#) [8](#)
 - [3.4. TCP-MD5 and Other Custom Authentication](#) [9](#)
 - [3.5. IPsec and IKE](#) [9](#)
- [4. Protocol Analysis](#) [9](#)
- [4.1. OSPF](#) [10](#)
 - [4.2. IS-IS](#) [10](#)
 - [4.3. BFD](#) [10](#)
 - [4.4. BGP](#) [11](#)
 - [4.5. Multicast Protocols \(PIM, MSDP\)](#) [11](#)
- [5. Summary](#) [12](#)
- [6. IANA Considerations](#) [12](#)
- [7. Acknowledgements](#) [12](#)
- [8. Security Considerations](#) [13](#)
- [9. References](#) [13](#)
- [9.1. Normative References](#) [13](#)
 - [9.2. Informative References](#) [14](#)
- [Author's Address](#) [15](#)
- [Intellectual Property and Copyright Statements](#) [16](#)

1. Introduction

A number of countermeasures for attacks against service provider backbone network infrastructure have been specified or proposed, each of them usually targeting a subset of the problem space. There has never been a more generic analysis of the actual problems, and which countermeasures are even necessary (and where). This document tries to provide that higher-level view.

The scope of this document are backbone infrastructures and the critical protocols that are required to function for legitimate traffic to be correctly forwarded through the network. As such, other important services or applications required by infrastructure elements such as RADIUS, NTP, remote access, syslog, SNMP, and DNS are out of scope. All such components should be adequately protected through appropriate measures, the most important of which are proper address and route filtering and restricting authorized access.

Additionally, the network might run additional routing protocols that are not described in this memo, such as (G)MPLS, RSVP-TE or LDP.

1.1. Abbreviations

We exclude the common abbreviations such as TCP, ICMP and DNS.

BGP	Border Gateway Protocol
bfd	Bidirectional Forwarding Detection
DoS	Denial of Service
DSCP	DiffServ Code Point
GTSM	Generalized TTL Security Mechanism
IGP	Interior Gateway Protocol
IKE	Internet Key Exchange
IRR	Internet Routing Registry
IS-IS	Integrated System - Integrated System (routing protocol)
LDP	Label Distribution Protocol
(G)MPLS	(Generalized) Multi-Protocol Label Switching
MSDP	Multicast Source Discovery Protocol
NTP	Network Time Protocol
OSPF	Open Shortest Path First
PIM	Protocol Independent Multicast
RADIUS	Remote Authentication Dial-In User Service
RSVP-TE	Resource Reservation Protocol - Traffic Engineering
SNMP	Simple Network Management Protocol

1.2. Assumptions

This document assumes that the service provider is doing at least some form of address filtering at its border devices, i.e., by ensuring that only the infrastructure nodes can use infrastructure source IP addresses to talk to the other nodes in the infrastructure. So, for example, if a router sees an IP packet coming from a source address assigned to another router in the backbone, it can be sure the packet has been originated inside the backbone (assuming the physical security of nodes in the backbone have not been subverted).

NOTE: many SP networks do not fulfill this assumption, often due to (1) legacy equipment which is not capable of line-rate filtering, and/or (2) very large network with hundreds or even thousands of devices is considered just too big to guard at the borders (and sometimes can't be broken down to several smaller ones). Analysis of this document does not and will not intend to cover these networks as the problem space is substantially different and other approaches are warranted. For example, [[I-D.zinin-rtg-dos](#)] suggested an alternative and provides good analysis; cryptographic protection of all the control traffic may be an option if "all bets are off".

This requirement can be satisfied by applying ingress filtering at all the ISP borders [[RFC2827](#)][[RFC3704](#)] for example, using feasible path strict uRPF towards customers and ingress access lists towards peers and upstreams. However, just filtering the infrastructure IP addresses used as source addresses from the outside is also sufficient. Some may even implement this by blocking access to the infrastructure destination addresses at the border, but this document doesn't describe this approach as that has a number of other issues.

Current operational practices are described in [[I-D.ietf-opsec-current-practices](#)]. Various filtering capabilities have been discussed at more length in [[I-D.ietf-opsec-filter-caps](#)].

1.3. Threat Model

In the context of this document, threats are assumed to come from external sources, either from customers or other networks. The typical attacks are either meant to cause some form of denial of service or simply cause collateral damage, such as:

- o DoS attacks directed at infrastructure (e.g., TCP RSTs, ICMP attacks),
- o Collateral damage from DoS and other attacks directed at someone else but causing harm to infrastructure or service (e.g., too much traffic exceeds forwarding or control processor capacity), or

- o Hijacking attacks (e.g., unauthorized routing advertisement, access control attempt with a spoofed address).

Other possible attack vectors but which are considered out of scope include:

- o ISP's systems being compromised through unauthorized access, system vulnerability, etc.,
- o Inside attacks (e.g., compromised personnel),
- o Lower-layer attacks as described in [Section 2.1](#).

While not perfect solutions, these can all be mitigated to some degree by controls and automatic configuration audits. As such the first order priority problems typically come from external sources.

[2.](#) Attack Vectors

This Section describes the most obvious attack vectors. Many of these are also described in [[RFC4732](#)].

[2.1.](#) Lower-layer Attacks

If an attacker has access to a (physical) link, it can obviously cause downtime for the link. In many cases the downtime is not a critical threat, as it can be quickly noticed, traffic rerouted, and the problem fixed. Some ISPs are more concerned about other forms of attacks: insertion of eavesdropping or man-in-the-middle devices. Fortunately, installing such would require downtime, and insertion could be noticeable, e.g., as an unscheduled issue gets fixed on its own.

However, a lower-layer attack is not specific to routing protocols. An attacker could just violate integrity or confidentiality of regular packets, instead of tampering with routing. As such, if a lower-layer attack is deemed a concern, full protection for all the traffic should be provided and therefore this threat is not addressed in this document.

[2.2.](#) Generic DoS on the Router

A typical attack is to overload a router using various techniques, e.g., by sending traffic exceeding the router's forwarding capacity, sending special transit packets that go through a "slow-path" processing (such functions may also come with problems of their own [[BLOCKED](#)]), or by sending some packets directed at the router itself

(e.g., to exceed the input queue for CPU processing).

Many of these techniques can be mitigated using implementation-specific rate-limiting mechanisms, so they are not addressed further in this memo. However, protocol designers should be advised to avoid any designs that require noticing and processing any special packets from the transit traffic (e.g., messages marked with router alert option).

2.3. Generic DoS on a Link

Overloading the capacity of a link is often more difficult to prevent than a router DoS. Traffic is typically not automatically rerouted and even if it was, doing so could make the issue worse unless there is ample spare capacity.

Mitigation methods include monitoring the usage status of links, prioritizing or deprioritizing certain kinds of traffic using DSCPs, or devising some form of rate-limiters.

2.4. Cryptographic Exhaustion Attacks

A special form of DoS are attacks which target a protocol that uses cryptographic mechanisms, for example TCP-MD5 or IPsec. The attacker sends valid protocol messages with cryptographic signatures or other properties to the router, which is forced to perform cryptographic validation of the message. If the cryptographic operations are computationally expensive, the attack might succeed easier than with other generic DoS mechanisms. Cryptographic protocols employing primitives such as stateless cookies, puzzles or return routability are typically more resistant to this kind of attacks.

Some implementation-specific mitigation techniques (rate-limiting etc.) have been deployed. Protocol design should take these attacks into account.

2.5. Unauthorized Neighbor or Routing Attacks

Unauthorized nodes can obtain a routing protocol adjacency on links where an IGP has been enabled by misconfiguration, or where authentication is not used. This may result in many different kinds of attacks, for example traffic redirection [[RFC4593](#)].

At least in theory, while it may not be possible to establish an adjacency from outside the link, it may be possible to inject packets as if the adjacency had been established (e.g., OSPF in Section 4.1.2 of [[I-D.ietf-rpsec-ospf-vuln](#)]).

Protocols such as BGP and MSDP that process routing information from untrusted, external sources may also be attacked, for example by an unauthorized advertisement of a prefix.

Special care needs to be made to ensure that unauthorized neighbors are prevented (e.g., by regular configuration audits and OSPF protocol filtering at borders). On the other hand, routing attack threats from valid neighbors can be slightly mitigated via appropriate route filtering.

2.6. TCP RST Attacks

TCP sessions can be closed by attackers that can send a TCP RST packet with guessed spoofed endpoint identifiers and a sufficiently close sequence number. The attacks and defenses have been described at length in [[I-D.ietf-tcpm-tcp-antispoof](#)]. One particular approach is modifying the TCP state machine [[I-D.ietf-tcpm-tcpsecure](#)].

2.7. ICMP Attacks

A slightly newer attack is employing ICMP by sending an ICMP type that indicates a hard error condition. ICMP errors must be propagated to the upper layer, and most applications heed the errors as they should by closing a connection or session. ICMP attacks and defenses against TCP have been extensively described in [[I-D.ietf-tcpm-icmp-attacks](#)]. Most TCP stacks have since then been fixed [[CVE](#)].

It is also possible to execute ICMP attacks against other protocols such as UDP or IPsec, but the impact and whether/how these protocols demultiplex received errors have not been extensively studied. IPsec is protected by ICMP attacks through a number of assumptions (e.g., that only ICMP errors from the end-point are accepted) or manual configuration.

3. Typical Countermeasures

This Section describes some of the most common countermeasures applied today. This just introduces the techniques; the afforded protection is analyzed in [Section 4](#) in the context of each protocol.

3.1. Filtering Addresses in Packets

As described in the first section, this document assumes that the internal infrastructure is secure from spoofed messages that purport to come from inside the infrastructure. More fine-grained, router-specific filters are sometimes deployed as well.

It is possible to hide the infrastructure by using private or non-advertised addressing, but this has numerous drawbacks such as breaking address filtering and traceroute, not protecting from the ISP's customers that use a default route, etc. so this document doesn't recommend doing so.

In addition, it may also make sense to ensure that egress packets have the ISP's own source addresses and/or that ingress packets arrive with either multicast/broadcast or ISP's own destination addresses. These ensure that in case your own filtering fails, no bad traffic leaks out and prevent certain classes of abuse from peers (e.g., stealing transit by static routing).

3.2. Filtering Addresses in Routing Updates

Similar principles as used in address filtering can be used to mitigate routing attacks. Specifically, reject any equal or more specific incoming routing advertisements to the ISP's address space unless explicitly authorized. Further, monitor the filtered prefixes and use public services (such as RIPE's MyASN [[MYASN](#)]) to monitor the correctness of advertisements globally.

As with address filtering, such routing advertisements might still be processed by other networks, but at least these steps prevent hijacking inside the ISP's own network and allow monitoring of most unauthorized attempts.

It may also make sense to filter out in a similar fashion the advertisements or more specifics of IX peering blocks where the ISP connects to. These could be advertised by an attacker to mess up forwarding next-hops.

In addition, especially in regions where the operational practice is to keep Internet Routing Registry (IRR) in sync, it may be possible to restrict the prefixes accepted from a peer or a customer to an automatically generated list. In any case, many operators define a maximum prefix limit per peer (which typically resets the session if exceeded) to prevent misconfiguration (e.g., unintentional deaggregation) or overload attacks.

3.3. GTSM

GTSM [[I-D.ietf-rtgwg-rfc3682bis](#)] is a technique where the sender of a packet sets the TTL/Hop Count to 255 and the receiver verifies it's still 255 (or some other preconfigured value). GTSM can be used to protect from off-link attacks (especially spoofing). This applies when GTSM-enabled control traffic is inside a single link: any packets coming from outside the link can summarily be discarded as

they have a TTL/Hop Count smaller than 255.

The open issue at the moment is how GTSM handles TCP RSTs. I.e., should it require that RSTs for a GTSM-enabled session should be sent with TTL=255 and verified to come with TTL=255 (or a configured value)? Some implementations already send out all packets with TTL=255, but receipt verification is not performed. Is there a sensible transition plan or need to make a change if any? Note that this has only limited impact on GTSM's security as other TCP RST mitigation techniques still apply.

NOTE IN DRAFT: the following paragraph should be removed in a future revision, to be placed to the GTSMbis draft.

We suggest that the GTSM spec is amended so that TCP RSTs relating to a GTSM-enabled protocol port MUST be sent with TTL=255. The recipient's behaviour SHOULD be configurable, and it is RECOMMENDED that the default be to discard messages where TTL is not 255 (or 255-TrustRadius).

3.4. TCP-MD5 and Other Custom Authentication

At least BGP, MSDP, and LDP are able to use the TCP-MD5 signature option to verify the authenticity of control packets. TCP-MD5 uses manually configured static keys, and changing them must be a coordinated event to prevent session reset. Due to the operational cost of re-keying, the solution is sub-optimal in cases where (rather paranoid) security procedures require (e.g., after an employee leaves the organization) that the keys must be easily and often changeable.

Using TCP-MD5 and other similar authentication mechanisms (e.g., for IGP or BFD) also opens an attack vector for cryptographic exhaustion attacks unless implementations have appropriate mechanisms to throttle or otherwise manage heavy cryptographic operations.

3.5. IPsec and IKE

IPsec and IKE have been proposed as a more comprehensive countermeasure, but these protocols also require a lot of heavyweight protocol machinery, lots of configuration, and cryptographic processing. Vendors have also expressed difficulty in applying IPsec to control traffic protection.

4. Protocol Analysis

This Section briefly discusses the protocol-specific attack properties below.

ICMP attacks apply to all the IP protocols at least to some degree. There is no reasonable way to appropriately protect from these attacks by operative methods such as filtering: the vendors should implement countermeasures described in [[I-D.ietf-tcpm-icmp-attacks](#)] to mitigate these attacks.

[4.1.](#) OSPF

OSPF attacks have already been analyzed [[I-D.ietf-rpsec-ospf-vuln](#)]. In this context the most important of them are preventing (1) misconfiguration and unauthorized neighbors, and (2) off-path directed attacks as described in Section 4.1.2 of [[I-D.ietf-rpsec-ospf-vuln](#)].

The former requires configuration change procedures and regular audits of OSPF configuration, and disabling OSPF adjacencies on customer-facing links, or adding authentication when there are multiple routers. The latter requires using OSPF authentication, dropping all OSPF traffic at all the borders, or moving to another, less vulnerable protocol (e.g., IS-IS).

OSPF is also used to some degree with provider-provisioned VPNs by the customers. In such scenarios, strict route filtering needs to be applied to ensure only the valid prefixes are accepted.

[4.2.](#) IS-IS

Routing IP with IS-IS has gained popularity in the backbone networks lately. As IS-IS does not use IP as its control protocol, external attackers cannot attack IS-IS in the same way as they can attack OSPF. Hence it is sufficient to prevent misconfiguration and unauthorized neighbors, using similar countermeasures as with OSPF: configuration change procedures and regular configuration audits and disabling IS-IS adjacencies on customer-facing links, or adding authentication when there are multiple routers.

[4.3.](#) BFD

Bidirectional Forwarding Detection (BFD) detects faults in the forwarding path between two endpoints. As a generic mechanism, it can be applied to a number of protocols (e.g., OSPF, IS-IS, BGP, MPLS, or static routes).

When BFD is in use for a single-hop scenario, it uses GTSM to protect from off-link attackers. Authentication can also be used for example on untrusted links.

[4.4.](#) BGP

Internal BGP sessions run between loopback addresses. There is no need to run TCP-MD5 for outsider protection as address filtering will avoid TCP RST attacks.

External BGP sessions may run multi-hop between loopback addresses or single-hop between interface addresses. The latter case is much more common and easier to protect and applying GTSM provides first-order resistance to off-link attackers.

In any case, assuming address filtering, the session can only be reset by the peer, or by attacks from the direction of the peer's network (e.g., through lack of peer's border filtering). One can therefore question the necessity of further protection as the peer can only shoot itself in the foot by killing the BGP session or allowing the BGP session be killed through negligence.

There is one exception to the above: if the customer is multihomed through multiple ISPs and the addresses used for the peering session are from the customer's address block. In such scenarios, using each ISP's respective addresses for the peering link might be the simplest approach.

If the link is not trusted (e.g., in some large Ethernet-based Internet Exchange points), it may also be desirable to ensure that peers are not able to reset others' sessions, so a mechanism like TCP-MD5 may be appropriate. One should note that the security requirements are not necessarily very high as the attacker should already be easily traceable on a single link, and thus re-keying may not be worth the trouble.

As BGP processes data heard from external sources, the routing data can be modified in numerous ways, e.g., to create arbitrarily complex advertisements using path attributes to crash naive BGP implementations. These and many other BGP attacks are described in [\[RFC4272\]](#). Techniques described in [Section 3.2](#) can mitigate the attack vectors to some degree, but a more comprehensive solution to securing routing data is needed.

[4.5.](#) Multicast Protocols (PIM, MSDP)

Multicast routing is typically achieved by PIM-SM [\[I-D.ietf-mboned-routingarch\]](#). MSDP is used for IPv4 source discovery. Multicast routing protocol threats have been analyzed separately in [\[RFC4609\]](#) (backbone perspective) and [\[I-D.ietf-pim-lasthop-threats\]](#) (last-hop perspective).

In summary, most of the multicast threats pertain to overloading control processors via too much state. Implementation-specific rate-limiters can help in mitigating the risk. If resetting MSDP sessions is a concern, TCP-MD5 option similar to BGP can be used. Address filtering can be applied in particular in PIM Unicast-Register message decapsulation; other messages use multicast and already employ reverse path forwarding checks.

5. Summary

IGPs require a great deal of care to ensure that they are not enabled on links where they shouldn't be. Preventing external OSPF attacks also requires OSPF authentication everywhere or filtering OSPF packets at the edges.

ICMP attacks are able to cause a great deal of harm to almost all the protocols, including IPsec, and there is little to do to mitigate the risk except to implement enhanced ICMP payload verification/processing techniques. More study of the impact on connectionless protocols and IPsec should be conducted.

With border address filtering in place, internal sessions are reasonably safe. With additional GTSM protection, external private interconnection links are also reasonably safe, as the session can only be reset by the neighbor or due to lack of filtering, someone through the neighbor's network. TCP-MD5 protection is most appropriate for Internet Exchange points with multiple neighbors or multihop eBGP sessions, but it's worth remembering that the security requirements for the solution are not very high as the attackers have very strict topological restrictions.

IPsec and IKE are obviously an option for heavy-weight protection, but impractical (yet) due to configuration complexity and processing overhead. Simplifications in configuration, implementation, and cryptographic hardware offloading might help the situation for the cases where the use of heavier protection (e.g., possibly Internet Exchange points) could be warranted.

6. IANA Considerations

This memo makes no request to IANA.

7. Acknowledgements

George Jones suggested improvements to the initial version of this

draft. Further feedback was received from Sean P. Turner, Seo Boon NG, Warren Kumari, Hank Nussbacher, Jonathan Trostle, Iljitsch van Beijnum, and Barry Greene.

8. Security Considerations

This document does not define a protocol but rather describes and analyzes the security properties and countermeasures in existing service provider backbone network infrastructures.

The most important issues that should be noted are its security assumptions:

- o We require at least certain degree of address filtering at borders, or else all bets are off. This assumption is notably NOT satisfied by a number of networks.
- o The main concern is an external attack (from customers or some other network); lower-layer attacks are not considered a particular concern for routing protocols.
- o Generic DoS attacks against routers can be mitigated using implementation-specific measures.

There are a number of actions for network operators in order to protect the network (e.g., filtering OSPF packets at the edges or auditing IGP configurations). There are also lessons to be learned for protocol designers (e.g., OSPF external attacks, ICMP attacks against non-TCP, use of GTSM). Many of the issues listed also depend on vendors to implement effective, vendor-specific rate-limiting techniques.

9. References

9.1. Normative References

- [I-D.ietf-opsec-current-practices]
Kao, M., "Operational Security Current Practices",
[draft-ietf-opsec-current-practices-07](#) (work in progress),
August 2006.
- [I-D.ietf-ospf-vuln]
Jones, E. and O. Moigne, "OSPF Security Vulnerabilities
Analysis", [draft-ietf-ospf-vuln-02](#) (work in
progress), June 2006.

- [I-D.ietf-rtgwg-rfc3682bis]
Gill, V., "The Generalized TTL Security Mechanism (GTSM)",
[draft-ietf-rtgwg-rfc3682bis-08](#) (work in progress),
December 2006.
- [I-D.ietf-tcpm-icmp-attacks]
Gont, F., "ICMP attacks against TCP",
[draft-ietf-tcpm-icmp-attacks-01](#) (work in progress),
October 2006.
- [I-D.ietf-tcpm-tcp-antispoof]
Touch, J., "Defending TCP Against Spoofing Attacks",
[draft-ietf-tcpm-tcp-antispoof-05](#) (work in progress),
October 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.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed
Networks", [BCP 84](#), [RFC 3704](#), March 2004.
- [RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis",
[RFC 4272](#), January 2006.
- [RFC4593] Barbir, A., Murphy, S., and Y. Yang, "Generic Threats to
Routing Protocols", [RFC 4593](#), October 2006.
- [RFC4609] Savola, P., Lehtonen, R., and D. Meyer, "Protocol
Independent Multicast - Sparse Mode (PIM-SM) Multicast
Routing Security Issues and Enhancements", [RFC 4609](#),
October 2006.

9.2. Informative References

- [BLOCKED] Cisco Systems, "Cisco Security Advisory: Cisco IOS
Interface Blocked by IPv4 Packets", 2004, <[http://
www.cisco.com/warp/public/707/
cisco-sa-20030717-blocked.shtml](http://www.cisco.com/warp/public/707/cisco-sa-20030717-blocked.shtml)>.
- [CVE] CVE-2004-0790, "Multiple TCP/IP and ICMP implementations
allow remote attackers to cause a denial of service (reset
TCP connections) via spoofed ICMP error messages, aka the
"blind connection-reset attack."", 2004, <[http://
cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0790](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0790)>.
- [I-D.ietf-mboned-routingarch]
Savola, P., "Overview of the Internet Multicast Routing

Architecture", [draft-ietf-mboned-routingarch-07](#) (work in progress), November 2006.

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

Morrow, C., "Filtering and Rate Limiting Capabilities for IP Network Infrastructure", [draft-ietf-opsec-filter-caps-04](#) (work in progress), September 2006.

[I-D.ietf-pim-lasthop-threats]

Savola, P. and J. Lingard, "Last-hop Threats to Protocol Independent Multicast (PIM)", [draft-ietf-pim-lasthop-threats-00](#) (work in progress), October 2006.

[I-D.ietf-tcpm-tcpsecure]

Ramaiah, A., "Improving TCP's Robustness to Blind In-Window Attacks", [draft-ietf-tcpm-tcpsecure-06](#) (work in progress), November 2006.

[I-D.zinin-rtg-dos]

Zinin, A., "Protecting Internet Routing Infrastructure from Outsider DoS Attacks", [draft-zinin-rtg-dos-02](#) (work in progress), May 2005.

[MYASN]

RIPE NCC, "MyASn System", <http://www.ris.ripe.net/myasn.html>.

[RFC4732]

Handley, M., Rescorla, E., and IAB, "Internet Denial-of-Service Considerations", [RFC 4732](#), December 2006.

Author's Address

Pekka Savola
CSC/FUNET
Espoo
Finland

Email: psavola@funet.fi

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).

