

IPv6 Operations Working Group (v6ops)
Internet-Draft
Intended status: Informational
Expires: February 8, 2015

F. Gont
SI6 Networks / UTN-FRH
J. Linkova
Google
T. Chown
University of Southampton
W. Liu
Huawei Technologies
August 7, 2014

IPv6 Extension Headers in the Real World
draft-gont-v6ops-ipv6-ehs-in-real-world-00

Abstract

IPv6 Extension Headers allow for the extension of the IPv6 protocol, and provide support for some core functionality such as IPv6 fragmentation. However, IPv6 Extension Headers are deemed to present a challenge to IPv6 implementations and networks, and are known to be intentionally filtered in some existing IPv6 deployments. This summarizes the issues associated with IPv6 extension headers, and presents real-world data regarding the extent to which packets with IPv6 extension headers are filtered in the public Internet, and where in the network such filtering occurs. Additionally, it provides some guidance to operators in troubleshooting IPv6 blackholes resulting from the use of IPv6 extension headers. Finally, this document provides some advice to protocol designers, and discusses areas where further work might be needed.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on February 8, 2015.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Previous Work on IPv6 Extension Headers	3
3.	Operational Implications	4
3.1.	Performance Issues	4
3.2.	Security Implications	4
4.	Support of IPv6 Extension Headers in the Public Internet . .	5
5.	Implications of Widespread IPv6 Extension Header Filtering .	7
5.1.	Advice to Protocol Designers	8
5.2.	A possible attack vector	8
5.3.	Possible Future Work	8
6.	Troubleshooting Packet Drops due to IPv6 Extension Headers .	9
7.	IANA Considerations	9
8.	Security Considerations	9
9.	Acknowledgements	9
10.	References	9
10.1.	Normative References	9
10.2.	Informative References	10
Appendix A.	Measurements Caveats	12
A.1.	Isolating the Dropping Node	12
A.2.	Obtaining the Responsible Organization for the Packet Drops	13
	Authors' Addresses	14

[1.](#) Introduction

IPv6 Extension Headers (EHs) allow for the extension of the IPv6 protocol, and provide support for core functionality such as IPv6 fragmentation. However, IPv6 Extension Headers have been deemed to present a challenge to IPv6 implementations and networks, and have been assumed/known to be intentionally filtered in some existing IPv6 deployments.

Discussions over the operational implications of IPv6 extension headers and their usability in the public Internet come up over and over again at both in IETF circles and other venues, and not infrequently some key aspects involving IPv6 extension headers are overlooked.

This document tries raise awareness about the operational implications of IPv6 Extension Headers, and their usability in the public Internet. Additionally, it provides some guidance in troubleshooting IPv6 blackholes resulting from the filtering of packets that employ IPv6 extension headers. Finally, it aims to raise awareness about the operational reality of IPv6 extension headers to protocol designers, and trigger discussion within the IETF community regarding areas where future work might be required.

[Section 2](#) of this document summarizes the work that has been done in the area of IPv6 extension headers. [Section 3](#) discusses the operational implications of IPv6 Extension Headers. [Section 4](#) presents real-world data regarding the extent to which IPv6 Extension Headers are usable in the public Internet. [Section 5](#) provides advise to protocol designers regarding the use of IPv6 extension headers, and aims to raise awareness about the possible interoperability implications on existing protocols. Finally, [Section 6](#) provides some guidance in troubleshooting of problems that may arise as a result of filtering packets that employ IPv6 Extension Headers.

2. Previous Work on IPv6 Extension Headers

Some of the implications of IPv6 Extension Headers have been discussed in IETF circles. For example, [\[I-D.taylor-v6ops-fragdrop\]](#) discusses a rationale for which operators filter IPv6 fragments. [\[I-D.wkumari-long-headers\]](#) discusses possible issues arising from "long" IPv6 header chains. [\[RFC7045\]](#) clarifies how intermediate nodes should deal with IPv6 extension headers. [\[RFC7112\]](#) discusses the issues arising in a specific case where the IPv6 header chain is fragmented into two or more fragments (and formally forbids such case). [\[RFC6980\]](#) analyzes the security implications of employing IPv6 fragmentation with Neighbor Discovery for IPv6, and formally recommends against such usage. Finally, [\[RFC7123\]](#) discusses how some popular RA-Guard implementations are subject to evasion by means of IPv6 extension headers.

While packets employing IPv6 Extension Headers have been "known" to be dropped in some IPv6 deployments, there was not much concrete data on the topic. Some preliminary measurements have been presented in [\[PMTUD-Blackholes\]](#), [\[Gont-IEPG88\]](#) and [\[Gont-Chown-IEPG89\]](#), whereas [\[Linkova-Gont-IEPG90\]](#) presents more comprehensive results on which [Section 4](#) of this document is based.

3. Operational Implications

3.1. Performance Issues

Many IPv6 router implementations suffer from a negative performance impact when IPv6 Extension Headers are employed.

In the most trivial case, a packet that includes a Hop-by-Hop Options header will typically go through the slow forwarding path, and be processed by the router's CPU. Another case is that in which a router that has been configured to enforce an ACL based on upper-layer information (e.g., upper layer protocol or TCP Destination Port). In such case, the router will need to process the entire IPv6 header chain in order to find the required information, and this may cause the packet to be processed in the slow path [[Cisco-EH-Cons](#)].

Processing a large amounts of traffic in the slow patch may cause the router to be unable to handle the same traffic loads when compared to normal packets, and may result in Denial of Service (DoS) scenarios.

We note that, for obvious reasons, the aforementioned performance issues may also affect other devices such as firewalls, Network Intrusion Detection Systems (NIDS), etc. [[Zack-FW-Benchmark](#)].

3.2. Security Implications

The security implications of IPv6 Extension Headers generally fall into one or more of these categories:

- o Evasion of security controls
- o DoS due to processing requirements
- o DoS due to implementation errors
- o Extension Header-specific issues

Different from IPv4, where the upper-layer protocol can be found after the variable-length IPv4 header, the structure of IPv6 packets is both more flexible and complex. Namely, finding the upper-layer information may imply processing the (daisy-chain like) entire IPv6 header chain. This has been often overlooked, and a number of security devices have been found to be trivially evasible by inserting one or more IPv6 Extension Headers between the main IPv6 header and the upper layer protocol. [[RFC7113](#)] describes this issue for the RA-Guard case, but same same techniques can be employed for circumventing e.g. some IPv6 firewalls. Additionally, inconsistencies in how some packets may be processed may result in

evasion of security controls [[I-D.kampanakis-6man-ipv6-eh-parsing](#)] [[Atlasis2014](#)].

As noted in [Section 3.1](#), packets that employ IPv6 Extension Headers may have a negative performance impact on the handling devices. Unless appropriate mitigations are put in place (e.g., packet filtering and/or rate-limiting), an attacker could simply send a large amount of IPv6 traffic employing IPv6 Extension Headers with the purpose of performing a Denial of Service (DoS) attack.

IPv6 implementations, as virtually every piece of software, tend to mature over time. While the IPv6 protocol itself (and many implementations) have been around for a long time already, bugs in IPv6 Extension Header processing have been recently found in a number of implementations. Because there is currently almost no reliance on IPv6 Extension headers, the corresponding code paths are rarely exercised, and there is the potential that bugs still remain to be discovered in some implementations.

Besides the general implications of IPv6 Extension Headers, each Extension Header tends to its own specific implications. One particular case is that of the Fragment Header, which is employed to provide the fragmentation function in IPv6. While many of the security implications of the fragmentation/reassembly mechanism are known from the IPv4 world, many of the related issues have crept into IPv6 implementations. They range from Denial of Service attacks to information leakage (see e.g. [[I-D.ietf-6man-predictable-fragment-id](#)], [[Bonica-NANOG58](#)], [[Atlasis2012](#)]).

4. Support of IPv6 Extension Headers in the Public Internet

This section summarizes the results obtained when measuring the support of IPv6 Extension Headers on the path towards different types of public IPv6 servers. Two sources were employed for the list of public IPv6 servers: the "World IPv6 Launch Day" site (<http://www.worldipv6launch.org/>) and Alexa's top 1M web sites (<http://www.alexa.com>). For each list of domain names, the following datasets were obtained:

- o Web servers (AAAA records of the aforementioned list)
- o Mail servers (MX -> AAAA of such list)
- o Name servers (NS -> AAAA of such list)

Duplicate and unreachable addresses were eliminated from each of those lists prior to obtaining the results below.

For each of the aforementioned address sets, three different types of probes were sent:

- o IPv6 packets with a Destination Options header of 8 bytes
- o IPv6 packets resulting in two IPv6 fragments of 512 bytes each (approximately)
- o IPv6 packets with a Hop-by-Hop Options header of 8 bytes

In the case of packets with Destination Options Header and Hop-by-Hop Options header, the desired EH size was achieved by means of PadN options [[RFC2460](#)]. The upper-layer protocol of the probe packets was, in all cases, TCP [[RFC0793](#)] segments with the Destination Port set to the service port [[IANA-PORT-NUMBERS](#)] of the corresponding dataset. For example, the probe packets for all the measurements involving web servers were TCP segments with the destination port set to 80.

Besides obtaining the packet drop rate when employing the aforementioned IPv6 extension headers, we tried to identify whether the Autonomous System (AS) dropping the packets was the same as the Autonomous System of the destination/target address. This is of particular interest since it essentially reveals whether the packet drops are under the control of the intended destination of the packets. Packets dropped by the destination AS are less of a concern, since they can be assumed to be the result of an explicit policy of the organization to which the packets are destined (who can make its own decision regarding what kind of traffic is "acceptable"). On the other hand, packets dropped by transit ASes are more of a concern, since they affect the deployability and usability of IPv6 extension headers (including IPv6 fragmentation) regardless of the intent of the communicating end-points. Thus, when packet drops do occur, the "best-case scenario" is that in which the packets are dropped by the destination AS, whereas the "worst-case scenario" is that in which the packets are dropped by a transit AS. Since there is some ambiguity when identifying the autonomous system to which a specific router belongs (see [Appendix A.2](#), our measurements result in a percentage *range*: the lowest percentage value represents the "best case scenario" (where, when in doubt, we assume the packet drops occur in the same AS as the destination AS), and the highest percentage value represents the "worst case scenario" (where, when in doubt, we assume the packet drops occur at different AS than the destination AS).

Dataset	D08	HBH8	FH512
Web	11.88% (17.60%-20.80%)	40.70% (31.43%-40.00%)	30.51% (5.08%-6.78%)
Mailservers	17.07% (6.35%-26.98%)	48.86% (40.50%-65.42%)	39.17% (2.91%-12.73%)
Namerservers	15.37% (14.29%-33.46%)	43.25% (42.49%-72.07%)	38.55% (3.90%-13.96%)

Table 1: WIPv6LD dataset: Packet drop rate for different destination types

Dataset	D08	HBH8	FH512
Web	10.91% (46.52%-53.23%)	39.03% (36.90%-46.35%)	28.26% (53.64%-61.43%)
Mailservers	11.54% (2.41%-21.08%)	45.45% (41.27%-61.13%)	35.68% (3.15%-10.92%)
Namerserver s	21.33% (10.27%-56.80%)	54.12% (50.64%-81.00%)	55.23% (5.66%-32.23%)

Table 2: Alexa's top 1M sites dataset: Packet drop rate for different destination types

There are a number of observations to be made based on the results presented above. Firstly, while it has been generally assumed that it is IPv6 fragments that are dropped by operators, our results indicate that it is IPv6 extension headers in general that are dropped. Secondly, our results indicate that a significant percentage of such packet drops occur in transit Autonomous Systems; that is, the packet drops are not under the control of the same organization as the final destination.

5. Implications of Widespread IPv6 Extension Header Filtering

The results presented in [Section 4](#) indicate that at least for part of the public Internet, communication employing IPv6 extension headers is unreliable. The following subsections discuss specific implications arising from this conclusion.

5.1. Advice to Protocol Designers

New protocols that are to operate in the public Internet should consider the effect of widespread filtering of IPv6 extension headers in the public Internet. If IPv6 extension headers are at all employed, a fall-back mechanism that does not rely on IPv6 extension headers should be considered.

5.2. A possible attack vector

The widespread filtering of IPv6 packets employing IPv6 Extension Headers could, in some scenarios, be exploited for malicious purposes: if packets employing IPv6 EHs are known to be filtered on the path from one system (say, "A") to another (say, "B"), an attacker could cause packets sent from A to B to be dropped by sending a forged an ICMPv6 Packet Too Big [[RFC4443](#)] error message to A (with a Next-Hop MTU smaller than 1280), such that subsequent packets from A to B include a fragment header (i.e., they result in atomic fragments [[RFC6946](#)]).

A problem with this attack scenario is that a node cannot simply "filter/drop all incoming ICMPv6 Packet Too Big error messages", or else it might not be able to properly reduce the assumed path MTU when communicating with other IPv6 nodes.

Possible mitigations for this issue include:

- o Filtering incoming ICMPv6 Packet Too Big (PTB) error messages that advertise a Next-Hop MTU smaller than 1280 bytes.
- o Artificially reducing the MTU to 1280 bytes and filter incoming ICMPv6 PTB error messages

Filtering only those ICMPv6 PTB messages that advertise a Next-Hop MTU smaller than 1280 would prevent the generation of IPv6 atomic fragments without breaking Path-MTU Discovery. However, such filtering would require deep packet inspection, and such functionality (if at all desirable) might not be available.

5.3. Possible Future Work

The impact of widespread filtering of IPv6 EHs on existing protocols should be considered. In particular, the effect of widespread filtering of IPv6 fragments on the Domain Name System (DNS) [[RFC1034](#)] should be evaluated (particularly when it is expected that reliance on IPv6 transport will increase over time).

6. Troubleshooting Packet Drops due to IPv6 Extension Headers

Isolating IPv6 blackholes essentially involves performing IPv6 traceroute for a destination system with and without IPv6 extension headers. The (EH-free) traceroute would provide the full working path towards a destination, while the EH-enabled traceroute would provide the address of the last-responding node for EH-enabled packets (say, "M"). In principle, one could isolate the dropping node by looking-up "M" in the EH-free traceroute, with the dropping node being "M+1" (see [Appendix A.1](#) for caveats).

At the time of this writing, most traceroute implementations do not support IPv6 extension headers. However, the path6 tool [[path6](#)] and RIPE Atlas [[RIPE-Atlas](#)] provide such support. Additionally, the blackhole6 tool [[blackhole6](#)] automates the troubleshooting process and can readily provide information such as: dropping node's IPv6 address, dropping node's Autonomous System, etc.

7. IANA Considerations

There are no IANA registries within this document. The RFC-Editor can remove this section before publication of this document as an RFC.

8. Security Considerations

The security implications of IPv6 extension headers are discussed in [Section 3.2](#). This document does not introduce any new security issues.

9. Acknowledgements

Fernando Gont would like to thank Jan Zorz and Go6 Lab <<http://go6lab.si/>> for providing access to systems and networks that were employed to produce some of the measurement results presented in this document. Additionally, he would like to thank SixXS <<https://www.sixxs.net>> for providing IPv6 connectivity.

10. References

10.1. Normative References

- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), September 1981.
- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](#), November 1987.

- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", [RFC 4443](#), March 2006.
- [RFC6946] Gont, F., "Processing of IPv6 "Atomic" Fragments", [RFC 6946](#), May 2013.
- [RFC6980] Gont, F., "Security Implications of IPv6 Fragmentation with IPv6 Neighbor Discovery", [RFC 6980](#), August 2013.
- [RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", [RFC 7045](#), December 2013.
- [RFC7112] Gont, F., Manral, V., and R. Bonica, "Implications of Oversized IPv6 Header Chains", [RFC 7112](#), January 2014.
- [RFC7113] Gont, F., "Implementation Advice for IPv6 Router Advertisement Guard (RA-Guard)", [RFC 7113](#), February 2014.
- [RFC7123] Gont, F. and W. Liu, "Security Implications of IPv6 on IPv4 Networks", [RFC 7123](#), February 2014.

10.2. Informative References

- [Atlasis2012]
Atlasis, A., "Attacking IPv6 Implementation Using Fragmentation", BlackHat Europe 2012. Amsterdam, Netherlands. March 14-16, 2012, <https://media.blackhat.com/bh-eu-12/Atlasis/bh-eu-12-Atlasis-Attacking_IPv6-Slides.pdf>.
- [Atlasis2014]
Atlasis, A., "A Novel Way of Abusing IPv6 Extension Headers to Evade IPv6 Security Devices", May 2014, <<http://www.insinuator.net/2014/05/a-novel-way-of-abusing-ipv6-extension-headers-to-evade-ipv6-security-devices/>>.
- [Bonica-NANOG58]
Bonica, R., "IPv6 Extension Headers in the Real World v2.0", NANOG 58. New Orleans, Louisiana, USA. June 3-5, 2013, <<https://www.nanog.org/sites/default/files/mon.general.fragmentation.bonica.pdf>>.

[Cisco-EH-Cons]

Cisco, , "IPv6 Extension Headers Review and Considerations", October 2006,
<http://www.cisco.com/en/US/technologies/tk648/tk872/technologies_white_paper0900aecd8054d37d.pdf>.

[Gont-Chown-IEPG89]

Gont, F. and T. Chown, "A Small Update on the Use of IPv6 Extension Headers", IEPG 89. London, UK. March 2, 2014,
<<http://www.iepg.org/2014-03-02-ietf89/fgont-iepg-ietf89-eh-update.pdf>>.

[Gont-IEPG88]

Gont, F., "Fragmentation and Extension header Support in the IPv6 Internet", IEPG 88. Vancouver, BC, Canada. November 13, 2013, <<http://www.iepg.org/2013-11-ietf88/fgont-iepg-ietf88-ipv6-frag-and-eh.pdf>>.

[I-D.ietf-6man-predictable-fragment-id]

Gont, F., "Security Implications of Predictable Fragment Identification Values", [draft-ietf-6man-predictable-fragment-id-01](#) (work in progress), April 2014.

[I-D.kampanakis-6man-ipv6-eh-parsing]

Kampanakis, P., "Implementation Guidelines for parsing IPv6 Extension Headers", [draft-kampanakis-6man-ipv6-eh-parsing-01](#) (work in progress), August 2014.

[I-D.taylor-v6ops-fragdrop]

Jaeggli, J., Colitti, L., Kumari, W., Vyncke, E., Kaeo, M., and T. Taylor, "Why Operators Filter Fragments and What It Implies", [draft-taylor-v6ops-fragdrop-02](#) (work in progress), December 2013.

[I-D.wkumari-long-headers]

Kumari, W., Jaeggli, J., and R. Bonica, "Operational Issues Associated With Long IPv6 Header Chains", [draft-wkumari-long-headers-02](#) (work in progress), October 2013.

[IANA-PORT-NUMBERS]

IANA, "Service Name and Transport Protocol Port Number Registry", <<http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.txt>>.

[Linkova-Gont-IEPG90]

Linkova, J. and F. Gont, "IPv6 Extension Headers in the Real World v2.0", IEPG 90. Toronto, ON, Canada. July 20, 2014, <<http://www.iepg.org/2014-07-20-ietf90/iepg-ietf90-ipv6-ehs-in-the-real-world-v2.0.pdf>>.

[PMTUD-Blackholes]

De Boer, M. and J. Bosma, "Discovering Path MTU black holes on the Internet using RIPE Atlas", July 2012, <<http://www.nlnetlabs.nl/downloads/publications/pmtu-black-holes-msc-thesis.pdf>>.

[RIPE-Atlas]

RIPE, , "RIPE Atlas", <<https://atlas.ripe.net/>>.

[Zack-FW-Benchmark]

Zack, E., "Firewall Security Assessment and Benchmarking IPv6 Firewall Load Tests", IPv6 Hackers Meeting #1, Berlin, Germany. June 30, 2013, <<http://www.ipv6hackers.org/meetings/ipv6-hackers-1/zack-ipv6hackers1-firewall-security-assessment-and-benchmarking.pdf>>.

[blackhole6]

blackhole6, , "blackhole6 tool manual page", <<http://www.si6networks.com/tools/ipv6toolkit>>, 2014.

[path6]

path6, , "path6 tool manual page", <<http://www.si6networks.com/tools/ipv6toolkit>>, 2014.

Appendix A. Measurements Caveats

A number of issues have needed some consideration when producing the results presented in [Section 4](#). These same issues should be considered when troubleshooting connectivity problems resulting from the use of IPv6 Extension headers.

A.1. Isolating the Dropping Node

Let us assume that we find that IPv6 EHs are dropping on their way to the destination system 2001:db8:d::1, and the output of running traceroute towards such destination is:

1. 2001:db8:1:1000::1
2. 2001:db8:2:2000::4
3. 2001:db8:2:4000::1
4. 2001:db8:3:4000::1
5. 2001:db8:3:1000::1
6. 2001:db8:4:4000::1
7. 2001:db8:4:1000::1
8. 2001:db8:5:5000::1
9. 2001:db8:5:6000::1
10. 2001:db8:d::1

and the output of EH-enabled traceroute to the same destination is:

1. 2001:db8:1:1000::1
2. 2001:db8:2:2000::4
3. 2001:db8:2:4000::1
4. 2001:db8:3:4000::1
5. 2001:db8:3:1000::1
6. 2001:db8:4:4000::1

For the sake of brevity, let us refer to the last-responding node in the EH-enabled traceroute ("2001:db8:4:4000::1" in this case) as "M". Assuming both packets in both traceroutes employ the same path, we'll refer to "the node following the last responding node in the EH-enabled traceroute" ("2001:db8:4:1000::1" in our case), as "M+1", etc.

Based on traceroute information above, which node is the one actually dropping the EH-enabled packets will depend on whether the dropping node filters packets on ingress or the egress. If the former, the dropping node will be M+1. If the latter, the dropping node will be "M".

Throughout this document (and our measurements), we assume that nodes perform ingress-filtering. Thus, in our example above the last responding node to the EH-enabled traceroute ("M") is "2001:db8:4:4000::1", and therefore we assume the "node" dropping node to be "2001:db8:4:1000::1" ("M+1").

Additionally, we note that when isolating the dropping node we assume that both the EH-enabled and the EH-free traceroutes result in the same paths. However, this might not be the case.

A.2. Obtaining the Responsible Organization for the Packet Drops

In order to identify the organization operating the dropping node, one would be tempted to lookup the ASN corresponding to the dropping node. However, assuming that M and M+1 are two peering routers, any

of these two organizations could be providing the address space employed for such peering. Or, in the case of an Internet eXchange Point (IXP), the address space could correspond to the IXP AS, rather than to any of the participating ASes. Thus, the organization operating the dropping node (M+1) could be the AS for M+1, but it might as well be the AS for M+2. Only when the ASN for M+1 is the same as the ASN for M+2 we have certainty about who the responsible organization for the packet drops is (see slides 21-23 of [[Linkova-Gont-IEPG90](#)]).

In the measurement results presented in [Section 4](#), the aforementioned ambiguity results in "percentage ranges" (rather than a specific ratio). This same ambiguity should be considered when troubleshooting and reporting IPv6 packet drops.

Finally, we note that a specific organization might be operating more than one Autonomous System. However, our measurements assume that different Autonomous System Numbers imply different organizations.

Authors' Addresses

Fernando Gont
SI6 Networks / UTN-FRH
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

Phone: +54 11 4650 8472
Email: fgont@si6networks.com
URI: <http://www.si6networks.com>

J. Linkova
Google
1600 Amphitheatre Parkway
Mountain View, CA 94043
USA

Email: furry@google.com

Tim Chown
University of Southampton
Highfield
Southampton , Hampshire S017 1BJ
United Kingdom

Email: tjc@ecs.soton.ac.uk

Will (Shucheng) Liu
Huawei Technologies
Bantian, Longgang District
Shenzhen 518129
P.R. China

Email: liushucheng@huawei.com