

Internet Engineering Task Force
Internet-Draft
Intended status: BCP
Expires: June 2, 2012

S. Perreault, Ed.
Viagenie
I. Yamagata
S. Miyakawa
NTT Communications
A. Nakagawa
Japan Internet Exchange (JPIX)
H. Ashida
IS Consulting G.K.
November 30, 2011

**Common requirements for Carrier Grade NATs (CGNs)
draft-ietf-behave-lsn-requirements-05**

Abstract

This document defines common requirements for Carrier-Grade NAT (CGN).

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [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 June 2, 2012.

Copyright Notice

Copyright (c) 2011 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 BSD License.

Table of Contents

1.	Introduction	3
2.	Terminology	3
3.	Requirements for CGNs	4
4.	Logging	10
5.	Bulk Port Allocation	10
6.	Deployment Considerations	12
7.	IANA Considerations	12
8.	Security Considerations	12
9.	Acknowledgements	12
10.	References	13
10.1.	Normative References	13
10.2.	Informative Reference	13
Appendix A.	Change Log (to be removed by RFC Editor prior to publication)	14
A.1.	Changed in -05	14
A.2.	Changed in -04	15
A.3.	Changed in -03	15
A.4.	Changed in -02	16
A.5.	Changed in -01	17
Authors'	Addresses	17

1. Introduction

With the shortage of IPv4 addresses, it is expected that more ISPs may want to provide a service where a public IPv4 address would be shared by many subscribers. Each subscriber is assigned a private address, and a NAT situated in the ISP's network translates between private and public addresses. When a second IPv4 NAT is located at the customer edge, this results in two layers of NAT.

This is not to be considered a solution to the shortage of IPv4 addresses. It is a service that can conceivably be offered alongside others, such as IPv6 services or regular, un-NATed IPv4 service. Some ISPs started offering such a service long before there was a shortage of IPv4 addresses, showing that there are driving forces other than the shortage of IPv4 addresses.

This document describes behavioral requirements that are to be expected of those multi-subscriber NATs. Meeting this set of requirements will greatly increase the likelihood that subscribers' applications will function properly.

Readers should be aware of potential issues that may arise when sharing a public address between many subscribers. See [\[RFC6269\]](#) for details.

This document builds upon previous works describing requirements for generic NATs [\[RFC4787\]](#)[\[RFC5382\]](#)[\[RFC5508\]](#). These documents, and their updates if any, still apply in this context. What follows are additional requirements, to be satisfied on top of previous ones.

2. Terminology

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 [\[RFC2119\]](#).

Readers are expected to be familiar with [\[RFC4787\]](#) and the terms defined there. The following additional term is used in this document:

Carrier-Grade NAT (CGN): A NAT-based [\[RFC2663\]](#) functional element operated by an administrative entity (e.g., operator) to share the same address among several subscribers. A CGN is managed by the administrative entity, not the subscribers.

Note that the term "carrier-grade" has nothing to do with the quality of the NAT; that is left to discretion of implementers.

Rather, it is to be understood as a topological qualifier: the NAT is placed in an ISP's network and translates the traffic of potentially many subscribers. Subscribers have limited or no control over the CGN, whereas they typically have full control over a NAT placed on their premises.

Figure 1 summarizes a common network topology in which a CGN operates.




Figure 1: CGN network topology

Another possible topology is one for hotspots, where there is no customer premise or customer-premises equipment (CPE), but where a CGN serves a bunch of customers who don't trust each other and hence fairness is an issue. One important difference with the previous topology is the absence of a second layer of NAT. This, however, has no impact on CGN requirements since they are driven by fairness and robustness in the service provided to customers, which applies in both cases.

3. Requirements for CGNs

What follows is a list of requirements for CGNs. They are in addition to those found in other documents such as [RFC4787], [RFC5382], and [RFC5508].

REQ-1: A CGN MUST support at least the following transport protocols: TCP (MUST support [\[RFC5382\]](#)), UDP (MUST support [\[RFC4787\]](#)), and ICMP (MUST support [\[RFC5508\]](#)). Support for additional transport protocols is OPTIONAL.

Justification: These protocols are the ones that NATs traditionally support. The IETF has documented the best current practices for them.

REQ-2: A CGN MUST have a default "IP address pooling" behavior of "Paired" (as defined in [\[RFC4787\] section 4.1](#)). A CGN MAY provide a mechanism for administrators to change this behavior on an application protocol basis.

- * When multiple overlapping internal IP address ranges share the same external IP address pool (e.g., DS-Lite [\[RFC6333\]](#)), the "IP address pooling" behavior applies to mappings between external IP addresses and internal subscribers rather than between external and internal IP addresses.

Justification: This stronger form of REQ-2 from [\[RFC4787\]](#) is justified by the stronger need for not breaking applications that depend on the external address remaining constant.

Note that this requirement applies regardless of the transport protocol. In other words, a CGN must use the same external IP address mapping for all sessions associated with the same internal IP address, be they TCP, UDP, ICMP, something else, or a mix of different protocols.

The justification for allowing other behaviors is to allow the administrator to save external addresses and ports for application protocols that are known to work fine with other behaviors in practice. However, the default behavior MUST be "Paired".

REQ-3: The CGN function SHOULD NOT have any limitations on the size nor the contiguity of the external address pool. In particular, the CGN function SHOULD be configurable with contiguous or non-contiguous external IPv4 address ranges.

Justification: Given the increasing rarity of IPv4 addresses, it is becoming harder for an operator to provide large contiguous address pools to CGNs. Additionally, operational flexibility may require non-contiguous address pools for reasons such as differentiated services, routing management, etc.

REQ-4: A CGN SHOULD support limiting the number of external ports (or, equivalently, "identifiers" for ICMP) that are assigned per subscriber.

- A. Limits SHOULD be configurable by the CGN administrator.
- B. Limits MAY be configurable independently per transport protocol.
- C. Additionally, it is RECOMMENDED that the CGN include administrator-adjustable thresholds to prevent a single subscriber from consuming excessive CPU resources from the CGN (e.g., rate limit the subscriber's creation of new mappings).

Justification: A CGN can be considered a network resource that is shared by competing subscribers. Limiting the number of external ports assigned to each subscriber mitigates the DoS attack that a subscriber could launch against other subscribers through the CGN in order to get a larger share of the resource. It ensures fairness among subscribers. Limiting the rate of allocation mitigates a similar attack where the CPU is the resource being targeted instead of port numbers.

REQ-5: A CGN SHOULD support limiting the amount of state memory allocated per mapping and per subscriber. This may include limiting the number of sessions, the number of filters, etc., depending on the NAT implementation.

- A. Limits SHOULD be configurable by the CGN administrator.
- B. Additionally, it SHOULD be possible to limit the rate at which memory-consuming state elements are allocated.

Justification: A NAT needs to keep track of TCP sessions associated to each mapping. This state consumes resources for which, in the case of a CGN, subscribers may compete. It is necessary to ensure that each subscriber has access to a fair share of the CGN's resources. Limiting TCP sessions per subscriber and per time unit is an effective mitigation against inter-subscriber DoS attacks. Limiting the rate of allocation is intended to prevent against CPU resource exhaustion.

REQ-6: It SHOULD be possible to administratively turn off translation for specific destination addresses and/or ports.

Figure 2 illustrates this use-case.

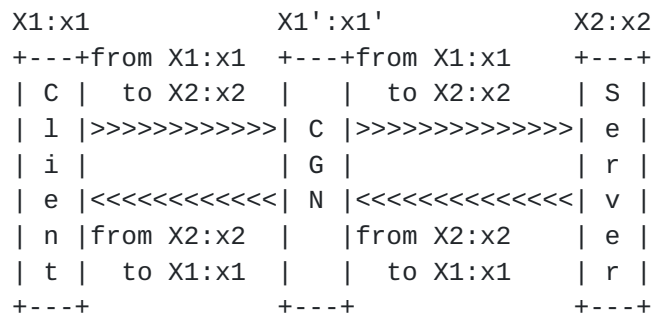


Figure 2: CGN pass-through

REQ-7: It is RECOMMENDED that a CGN have an "Endpoint-Independent Filtering" behavior (as defined in [\[RFC4787\] section 5](#)). If it is known that "Address-Independent Filtering" does not cause the application-layer protocol to break (how to determine this is out of scope for this document), then it MAY be used instead.

This is based on the observation that some games and peer-to-peer applications require EIF for the NAT traversal to work. In the context of a CGN it is important to minimize application breakage.

REQ-8: When a CGN loses state (due to a crash, reboot, failover to a cold standby, etc.), it MUST NOT reuse the same external address+port pairs for new dynamic mappings for at least 120 seconds, except for any of the following cases:

- A. If the CGN tracks TCP sessions (e.g., with a state machine, as in [\[RFC6146\] section 3.5.2.2](#)), TCP ports MAY be reused immediately.
- B. If external ports are statically assigned to internal addresses (e.g., address X with port range 1000-1999 is assigned to subscriber A, 2000-2999 to subscriber B, etc.), and the assignment remains constant across state

loss, then ports MAY be reused immediately.

Justification: This is necessary in order to prevent collisions between old and new mappings and sessions. It ensures that all established sessions are broken instead of redirected to a different peer.

The exceptions are for cases where reusing a port immediately does not create a possibility that packets would be redirected to the wrong peer.

The 120 seconds value corresponds to the Maximum Segment Lifetime (MSL) from [[RFC0793](#)].

One way that this requirement could be satisfied would be to have two distinct address pools: one dormant and one active. When rebooting, the CGN would swap the dormant pool with the active pool. Another way would be simply to wait 120 seconds before resuming NAT activity.

REQ-9: Once an external port is deallocated, it SHOULD NOT be reallocated to a new mapping until at least 120 seconds have passed. The length of time and the maximum number of ports in this state SHOULD be configurable by the CGN administrator. The following exceptions apply:

- A. If the CGN tracks TCP sessions (e.g., with a state machine, as in [[RFC6146](#) section 3.5.2.2]), TCP ports MAY be reused immediately.
- B. If the allocated external ports used address-dependent or address-and-port-dependent filtering before state loss, they MAY be reused immediately.

Justification: This is to prevent users from receiving unwanted traffic. It also helps prevent against clock skew when mappings are logged.

The exceptions are for cases where reusing a port immediately does not create a possibility that packets would be redirected to the wrong peer.

The 120 seconds value corresponds to the Maximum Segment Lifetime (MSL) from [[RFC0793](#)].

REQ-10: A CGN SHOULD include a Port Control Protocol server [[I-D.ietf-pcp-base](#)].

Justification: Allowing subscribers to manipulate the NAT state table with PCP greatly increases the likelihood that applications will function properly.

REQ-11: A CGN SHOULD support [[RFC4008](#)].

Justification: It is anticipated that CGNs will be primarily deployed in ISP networks where the need for management is critical.

Note also that there are efforts within the IETF toward creating a MIB tailored for CGNs (e.g., [[I-D.perreault-opsawg-natmib-bis](#)]).

REQ-12: When a CGN is unable to create a mapping due to resource constraints or administrative restrictions (i.e., quotas):

- A. it MUST drop the original packet;
- B. it SHOULD send an ICMP Destination Unreachable message with code 3 (Port Unreachable) to the session initiator;
- C. it SHOULD send a notification (e.g., SNMP trap) towards a management system (if configured to do so);
- D. and it SHOULD NOT delete existing mappings in order to "make room" for the new one. (This only applies to normal CGN behavior, not to manual operator intervention.)

Justification: This is a slightly different form of REQ-8 from [[RFC5508](#)]. Code 3 is preferred to code 13 because it is listed as a "soft error" in [[RFC5461](#)], which is important because we don't want TCP stacks to abort the connection attempt in this case. Sending an ICMP error may be rate-limited for security reasons, which is why requirement B is a SHOULD, not a MUST.

Applications generally handle connection establishment failure better than established connection failure. This is why dropping the packet initiating the new connection is preferred over deleting existing mappings. See also the rationale in [[RFC5508](#) [section 6](#)].

4. Logging

It may be necessary for CGN administrators to be able to identify a subscriber based on external IPv4 address, port, and timestamp in order to deal with abuse and lawful intercept requests. When multiple subscribers share a single external address, the source address and port that are visible at the destination host have been translated from the ones originated by the subscriber.

In order to be able to do this, the CGN would need to log the following information for each mapping created:

- o subscriber identifier (e.g., internal source address or tunnel endpoint identifier)
- o external source address
- o external source port
- o timestamp

By "subscriber identifier" we mean information that uniquely identifies a subscriber. For example, in a traditional NAT scenario, the internal source address would be sufficient. In the case of DS-Lite, many subscribers share the same internal address and the subscriber identifier is the tunnel endpoint identifier (i.e., the B4's IPv6 address).

A disadvantage of logging mappings is that CGNs under heavy usage may produce large amounts of logs, which may require large storage volume.

REQ-13: A CGN SHOULD NOT log destination addresses or ports.

Justification: Destination logging at the CGN creates privacy issues. Furthermore, readers should be aware of logging recommendations for Internet-facing servers [[RFC6302](#)]. With compliant servers, the destination address and port do not need to be logged by the CGN. This can help reduce the amount of logging.

5. Bulk Port Allocation

So far we have assumed that a CGN allocates one external port for every outgoing connection. In this section, the impacts of allocating multiple external ports at a time are discussed.

There is a range of things a CGN can do:

Traditional: For every outgoing connection, allocate one external port.

Scattered port set: For an outgoing connection, create a set of several non-consecutive external ports. Subsequent outgoing connections will use ports from the set. When the set is exhausted, a new connection causes a new set to be created. A set is smaller or equal to the user's maximum port limit.

Consecutive port set: Same as the scattered port set, but the ports allocated to a set are consecutive.

Note that this list is not exhaustive. There is a continuum of behavior that a CGN may choose to implement. For example, a CGN could use scattered port sets of consecutive port sets.

The impacts of bulk port allocation are as follows.

Port Utilization: The mechanisms at the top of the list are very efficient in their port utilization. In that sense, they have good scaling properties (nothing is wasted). The mechanisms at the bottom of the list will waste ports. The number of wasted ports is proportional to size of the "bin".

Logging: Traditional allocation creates a lot of log entries as compared to allocation by port sets which creates much fewer entries. Scattered and consecutive port sets generate the same number of log entries. In the case of consecutive port sets, entries can be expressed very compactly by indicating a range (e.g., "12000-12009"). Some scattered port set allocation schemes can also generate small log entries containing the parameters and algorithm used for the port set generation (see, e.g., [[I-D.boucadair-pppext-portrange-option](#)]).

With large set sizes, the logging frequency for scattered and consecutive port sets can approach that of DHCP servers.

Logging destination addresses and ports can only be done on a per-session basis. This means that destination logging for a CGN implementing bulk port allocation would create one log entry per session containing the destination address and port. Other information could still be logged in one entry per port set.

Security: Traditional and scattered port sets provide very good security in that ports numbers are not easily guessed. Easily guessed port numbers put subscribers at risk of the attacks described in [[RFC6056](#)]. Consecutive port sets provides poor security to subscribers, especially if the set size is small.

6. Deployment Considerations

Several issues are encountered when CGNs are used [[RFC6269](#)]. There is current work in the IETF toward alleviating some of these issues. For example, see [[I-D.boucadair-intarea-nat-reveal-analysis](#)].

The address sharing ratio is the ratio between the number of external addresses and the number of internal addresses that a CGN is configured to handle. See [[RFC6269](#)] [section 26.2](#) for guidance on picking an appropriate ratio.

7. IANA Considerations

There are no IANA considerations.

8. Security Considerations

If a malicious subscriber can spoof another subscriber's CPE, it may cause a DoS to that subscriber by creating mappings up to the allowed limit. Preventing this can be accomplished with ingress filtering, as described in [[RFC2827](#)].

Endpoint-Independent Filtering has security considerations which are discussed in [[RFC4787](#)].

NATs sometimes perform fragment reassembly. CGNs would do so at presumably high data rates. Therefore, the reader should be familiar with the potential security issues described in [[RFC4963](#)].

9. Acknowledgements

Thanks for the input and review by Arifumi Matsumoto, Benson Schliesser, Dai Kuwabara, Dan Wing, Dave Thaler, Francis Dupont, Joe Touch, Lars Eggert, Kousuke Shishikura, Mohamed Boucadair, Nejc Skoberne, Reinaldo Penno, Senthil Sivakumar, Takanori Mizuguchi, Takeshi Tomochika, Tina Tsou, Tomohiro Fujisaki, Tomohiro Nishitani, Tomoya Yoshida, and Yasuhiro Shirasaki. Dan Wing also contributed much of [section 5](#).

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4008] Rohit, R., Srisuresh, P., Raghunarayan, R., Pai, N., and C. Wang, "Definitions of Managed Objects for Network Address Translators (NAT)", [RFC 4008](#), March 2005.
- [RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", [BCP 127](#), [RFC 4787](#), January 2007.
- [RFC5382] Guha, S., Biswas, K., Ford, B., Sivakumar, S., and P. Srisuresh, "NAT Behavioral Requirements for TCP", [BCP 142](#), [RFC 5382](#), October 2008.
- [RFC5508] Srisuresh, P., Ford, B., Sivakumar, S., and S. Guha, "NAT Behavioral Requirements for ICMP", [BCP 148](#), [RFC 5508](#), April 2009.
- [I-D.ietf-pcp-base]
Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", [draft-ietf-pcp-base-16](#) (work in progress), October 2011.

10.2. Informative Reference

- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), September 1981.
- [RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", [RFC 2663](#), August 1999.
- [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.
- [RFC4963] Heffner, J., Mathis, M., and B. Chandler, "IPv4 Reassembly Errors at High Data Rates", [RFC 4963](#), July 2007.
- [RFC5461] Gont, F., "TCP's Reaction to Soft Errors", [RFC 5461](#), February 2009.
- [RFC6056] Larsen, M. and F. Gont, "Recommendations for Transport-Protocol Port Randomization", [BCP 156](#), [RFC 6056](#), January 2011.

- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", [RFC 6146](#), April 2011.
- [RFC6269] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", [RFC 6269](#), June 2011.
- [RFC6302] Durand, A., Gashinsky, I., Lee, D., and S. Sheppard, "Logging Recommendations for Internet-Facing Servers", [BCP 162](#), [RFC 6302](#), June 2011.
- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", [RFC 6333](#), August 2011.
- [I-D.boucadair-intarea-nat-reveal-analysis]
Boucadair, M., Touch, J., Levis, P., and R. Penno, "Analysis of Solution Candidates to Reveal a Host Identifier in Shared Address Deployments", [draft-boucadair-intarea-nat-reveal-analysis-04](#) (work in progress), September 2011.
- [I-D.boucadair-pppext-portrange-option]
Boucadair, M., Levis, P., Bajko, G., Savolainen, T., and T. Tsou, "Huawei Port Range Configuration Options for PPP IPCP", [draft-boucadair-pppext-portrange-option-09](#) (work in progress), September 2011.
- [I-D.perreault-opsawg-natmib-bis]
Perreault, S. and T. ZOU, "Definitions of Managed Objects for Network Address Translators (NAT)", [draft-perreault-opsawg-natmib-bis-00](#) (work in progress), September 2011.

Appendix A. Change Log (to be removed by RFC Editor prior to publication)

A.1. Changed in -05

- o Removed DSCP requirement since it applies to non-CG NATs as well.
- o Removed instances of "NAT444".
- o Filtering has no effect on the requirement for a hold down pool. Removed REQ-8-B.

- o Statically assigned port ranges do not need to go in the hold down pool. Added a new REQ-8-B.
- o Fixed various nits. More precise text in some places.

A.2. Changed in -04

- o Fixed nits, spelling, updated references.
- o CGNs SHOULD NOT log destinations.
- o Allow address-dependent filtering when it does not cause the application protocol to break.
- o Refer to [RFC4787](#) security considerations on EIF.
- o Clarify REQ-12 point D (it does not apply to operator intervention).
- o Changed "CGNs SHOULD limit ..." to "SHOULD support limiting" to make it clear that the operator is in control.
- o Added reference to [RFC 4963](#).
- o Added requirement for non-contiguous external address pools.

A.3. Changed in -03

- o Added exceptions for which it is not necessary to wait 120 seconds before reusing a port.
- o Renamed "random port set" to "scattered port set", which is more accurate.
- o Log "subscriber identifier" instead of internal address+port to allow for overlapping internal address ranges (DS-Lite).
- o Adjusted logging text and added reference to I-D.boucadair-pppext-portrange-option.
- o Adjusted destination logging text for bulk port allocation schemes.
- o Removed requirement for I-D.ietf-intarea-ipv4-id-update.
- o Made PCP support a SHOULD-level requirement.

- o Lowered the level of requirement for not dropping existing mappings in order to "make room" to SHOULD level, and added rationale.

A.4. Changed in -02

- o CGNs MUST support at least TCP, UDP, and ICMP.
- o Add requirement from I-D.ietf-intarea-ipv4-id-update.
- o Add informative reference to [[RFC6269](#)].
- o Add requirement (SHOULD level) for a port forwarding protocol.
- o Allow any pooling behavior on a per-application protocol basis.
- o Adjust wording for external port allocation rate limiting.
- o Add requirement for [RFC4008](#) support (SHOULD level).
- o Adjust wording for swapping address pools when rebooting.
- o Add DSCP requirement (stolen from [draft-jennings-behave-nat6](#)).
- o Add informative reference to [draft-boucadair-intarea-nat-reveal-analysis](#).
- o Add requirement for hold-down pool.
- o Change definition of CGN.
- o Avoid usage of "device" loaded word throughout the document.
- o Add requirement about resource exhaustion.
- o Change title.
- o Describe additional CGN topology where there is no NAT444.
- o Better justification for "Paired" pool behavior.
- o Make it clear that rate limiting allocation is for preserving CPU resources
- o Generalize the requirement for limiting the number of TCP sessions per mapping so that it applies to all memory-consuming state elements.

- o Change CPE to subscriber where it applies throughout the text.
- o Better terminology for bulk port allocation mechanisms.
- o Explain how external address pairing works with DS-Lite.

A.5. Changed in -01

- o Terminology: LSN is now CGN.
- o Imported all requirements from RFCs 4787, 5382, and 5508. This allowed us to eliminate some duplication.
- o Added references to [draft-ietf-intarea-server-logging-recommendations](#) and [draft-ford-shared-addressing-issues](#).
- o Incorporated a requirement from [draft-xu-behave-stateful-nat-standby-06](#).

Authors' Addresses

Simon Perreault (editor)
Viagenie
2875 boul. Laurier, suite D2-630
Quebec, QC G1V 2M2
Canada

Phone: +1 418 656 9254
Email: simon.perreault@viagenie.ca
URI: <http://www.viagenie.ca>

Ikuhei Yamagata
NTT Communications Corporation
Gran Park Tower 17F, 3-4-1 Shibaura, Minato-ku
Tokyo 108-8118
Japan

Phone: +81 50 3812 4704
Email: ikuhei@nttv6.jp

Shin Miyakawa
NTT Communications Corporation
Gran Park Tower 17F, 3-4-1 Shibaura, Minato-ku
Tokyo 108-8118
Japan

Phone: +81 50 3812 4695
Email: miyakawa@nttv6.jp

Akira Nakagawa
Japan Internet Exchange Co., Ltd. (JPIX)
Otemachi Building 21F, 1-8-1 Otemachi, Chiyoda-ku
Tokyo 100-0004
Japan

Phone: +81 90 9242 2717
Email: a-nakagawa@jpix.ad.jp

Hiroyuki Ashida
IS Consulting G.K.
12-17 Odenma-cho Nihonbashi Chuo-ku
Tokyo 103-0011
Japan

Email: assie@hir.jp

