Internet Engineering Task Force
Internet-Draft
Expires: January 3, 2009

T. Nishitani
S. Miyakawa
NTT Communications
July 2, 2008

Carrier Grade Network Address Translator (NAT) Behavioral Requirements for Unicast UDP, TCP and ICMP
draft-nishitani-cgn-00

Status of this Memo

Copyright Notice

Abstract

This document defines basic terminology for describing different types of carrier-grade Network Address Translation (NAT) behavior when handling Unicast UDP, TCP and ICMP. Developing carrier-grade NATs that meet this set of requirements increase transparency of data between carrier networks.

Table of Contents

1. Introduction

Global IPv4 address from the IANA pool will run out in a few years, thus carriers need to shift from IPv4 services to IPv6 ones. However, IPv6 deployment seems to take a long time.

NAT [RFC3022] is a key technology to utilize IPv4 global address effectively in current practice. ISP may have to place NAT devices between end-users and the public Internet to suppress global IPv4 address consumption.

In this document, we call carrier's NAT device Carrier Grade NAT (CGN). This document describes behavioral requirements of CGN for unicast UDP, TCP and ICMP. [RFC4787], [I-D.ietf-behave-tcp] and [I-D.ietf-behave-nat-icmp] describes requirements of unicast UDP, TCP and ICMP for NAT which is placed on network edge and is intended for high transparency of NAT. CGNs also need interoperability and high transparency among carriers to make end-users be able to use various services like Peer-to-Peer(P2P) applications and Instant Messenger. [RFC5128] is nominated for an NAT traversal condition in P2P.

The main target of this document is 4-4-4 model which uses IPv4 address both internal and external side of CGN. [I-D.durand-v6ops-natv4v6v4] describes 4-6-4 model, and CGN may apply to 4-6-4 model.

Interaction of this requirements and security of Customer Premises Equipment(CPE) is out of scope because CPE should defend itself.

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 term are used in this document:

Carrier-grade NAT(CGN): NAT devices placed between CPE and public Internet by a carrier. CGN converts CPE IP Address, CPE Port, and CPE Identifier into CGN external IP Address, CGN external Port and CGN external Identifier in communication between CPE and GGN external.

CGN external realm: The realm where IPv4 global addresses are assigned

CGN internal realm: The realm placed between CGN and CPEs

CGN external IP address: The IP address on CGN in CGN external realm corresponding to CPE IP address

CGN external port: The port on CGN in CGE external realm corresponding to CPE port

CGN external identifier: The identifier of ICMP on CGN in CGN external realm corresponding to CPE identifier

Customer Premises Equipment(CPE): The terminal which is placed in CGN internal realm and may establish TCP sessions to CGN external realm

CPE IP address: The IP address on CPE in CGN internal realm

CPE port: The port on CPE in CGN internal realm

CPE identifier: CPE's identifier of ICMP in CGN internal realm

CPE 3-tuple: The tuple of TCP/UDP, CPE IP address, and CPE Port Carrier Service Server (CSS) The server a carrier supplies various services for CPE

Carrier Service Server (CSS): The server a carrier supplies various services for CPE

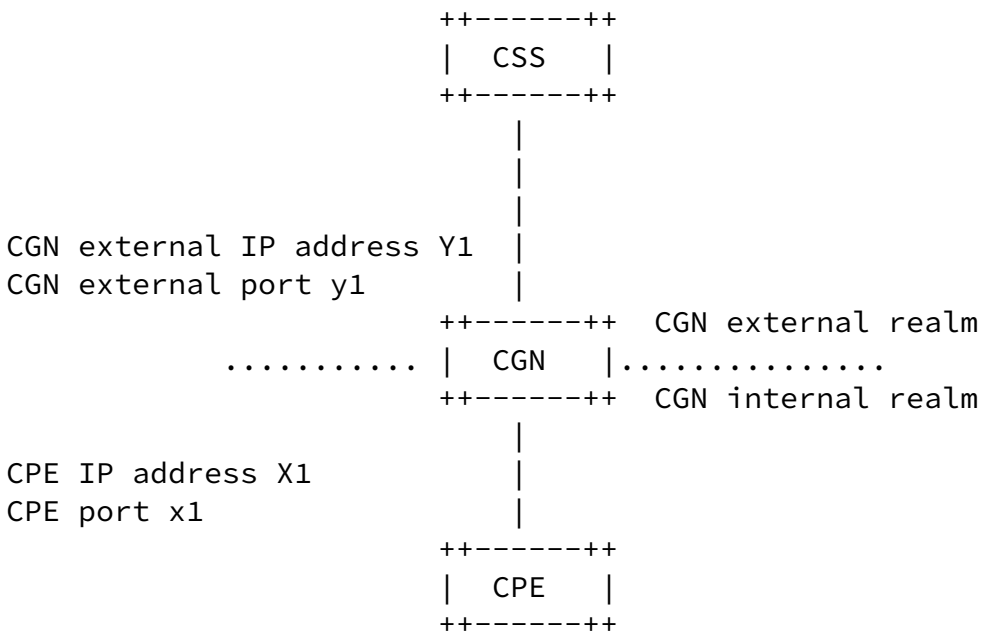

CGN network

3. The policy of assignment of CGN external IP address, port and identifier

A CGN has a pool of CGN external IP addresses, ports and identifiers. CPEs share CGN external IP addresses. Each CGN occupies combination of CGN external IP address and CGN external port exclusively. For a fair use of limited resources, CGN has a limitation for the number of the CGN external ports per CPE. CGNs need to keep high transparency to continue existing services after a carrier introduces CGN. Requirement of high transparency for CGN leads to high scalability of CGN. High transparency means CGN basically keeps communications among CPEs except effect of limitations of the number of CGN external ports and TCP sessions.

A CPE MAY apply UDP hole punching or TCP hole punching for interactive services among CPEs like Voice over IP and P2P. CGN SHOLUD NOT interfere in services using UDP hole punching or TCP hole punching.

REQ-1: A CGN MUST allocate one external IP address to each CPE.

a) CGN external IP address of the UDP, TCP and ICMP MUST be same.

Justification: If a CGN allocates multiple CGN external IP addresses to each CPE, some applications might not work.

REQ-2: A CGN MUST allocate CGN external ports corresponding to CPE ports of UDP.

a) A CGN MUST NOT overload CGN external port while a NAT UDP mapping timer does not expire.

b) A CGN MAY overload CGN external port after a NAT UDP mapping timer expires.

c) A CGN SHOULD limit the number of the CGN external ports of UDP per CPE.

d) The number of the CGN external ports of UDP per CPE which CGN can allocate SHOULD be configurable for the administrator of CGN.

e) A CGN SHOULD NOT allocate well-known ports as CGN external ports.

Justification: CPEs can communicate to CPE external realm fairly by

limiting the number of CGN external ports per CPE.

REQ-3: A CGN MUST allocate CGN external ports corresponding to CPE ports of TCP.

a) A CGN MUST NOT overload CGN external port while the port is allocated for one or more TCP sessions originated by another CPE.

b) A CGN MAY reuse CGN external port while the port is allocated for no session originated by any CPE.

c) A CGN SHOULD limit the number of the CGN external ports of TCP per CPE.

d) The number of the CGN external ports of TCP per CPE SHOULD be

an administratively configurable option.

e) A CGN SHOULD limit the number of the new sessions of TCP per time unit and per CPE.

f) A CGN SHOULD NOT allocate well-known ports as CGN external ports.

Justification: CPEs can communicate to CPE external realm fairly by limiting the number of CGN external ports per CPE. In addition, TCP CGN external port MAY have TCP sessions, and therefore the TCP session timer is necessary for every 5-Tuple. CGN can have not only the limitations of the number of CGN external ports but also TCP sessions per CPE. Thus a CGN can prevent denial of service attacks with the tons of TCP open and close by malicious CPEs.

REQ-4: A CGN MUST allocate CGN external identifiers corresponding to CPE identifiers.

a) A CGN MUST NOT overload CGN external identifier before an ICMP Query session timer expires.

b) A CGN MAY overload CGN external identifier after an ICMP Query session timer expires.

c) A CGN SHOULD limit the number of the CGN external identifier allocated per CPE.

d) The number of the CGN external identifiers per CPE which CGN can allocate SHOULD be an administratively configurable option.

Justification: CPEs can communicate to CPE external realm fairly by limiting the number of CGN external identifiers every CPE.

When a CGN limits the number of CGN external ports and TCP sessions, CPE may not use TCP services during using web and P2P services. For example, some services using Ajax demand few dozens of TCP sessions. P2P software like BitTorrent demands also TCP sessions more than few dozens. Some CPEs MAY use E-mail services like POP3 and SMTP even though CPE uses the services which demand many TCP sessions at the same time. Therefore it is important to reserve CGN external ports for such administratively configured services.

REQ-5: Reserving CGN external ports per CPE for the always-available services are RECOMENDED.

a) The destination port which is used for reservation of CGN external ports SHOULD be administratively configurable.

Justification: To reserve the CGN external ports for specific services, CPE can avoid the effect of the limitation of CGN external ports by CGN.

In addition, it MAY not be necessary to set a limit to the number of CGN external ports for the communications between CPEs and CSS. The reason is because CGN should pass-through the communications between CPEs and CSS.

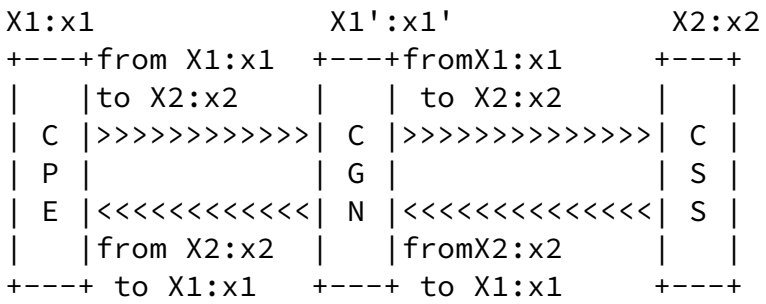

pass-through

REQ-6: A CGN SHOULD pass-through the communication between CPEs and CSS.

Justification: Using pass-through, CGN does not have to assign CGN external IP address, ports, and identifiers and limit to the number of ports and TCP sessions for the services that a carrier manages.

4. Unicast UDP Requirements

[RFC4787] describes requirements of the Unicast UDP of a NAT, and the behavior of "Endpoint-Independent Filtering "is RECOMMNEDED, and a NAT MUST have an "Endpoint-Independent Mapping" behavior to ensure transparency of CGN.

To have "Endpoint-Independent Filtering" and "Endpoint-Independent Mapping" behaviors for CGNs, CGNs help to establish UDP Hole Punching among CPEs. In other words, the possibility of the establishment of UDP Hole Punching among CPEs which have CGN is equal to the possibility among CPEs which don's t have CGN. If CGNs have an "Address-Dependent Mapping" or "Address and Port-Dependent Mapping" behavior, the possibility that establishment of UDP Hole Punching is less than when CGNs have an "Endpoint-Independent Mapping" behavior. And if CGNs have an "Address and Port-Dependent Filtering" behavior, the possibility that establishment of UDP Hole Punching is less than when CGNs have an "Endpoint-Independent Filtering" or "Address Dependent Filtering" behavior. Because a CSS is placed external CGN realm, the source IP address and port of the communication from CPE to CSS is CGN external IP address and port. It is RECOMMENDED to use STUN[I-D.ietf-behave-rfc3489bis] if CPEs check the CGN external IP address and port for CPE.

A carrier MAY introduce TURN [I-D.ietf-behave-turn] to support communications among CPEs. If CGN supports "Hairpinning", CGN can hairpin the communications between CPEs in the same CGN. Therefore the requirements of Hairpinning for CGN MAY reduce requirements for the performance of TURN servers. When CPEs decide the course of UDP between CPEs, CPE MAY use [I-D.ietf-mmusic-ice] .

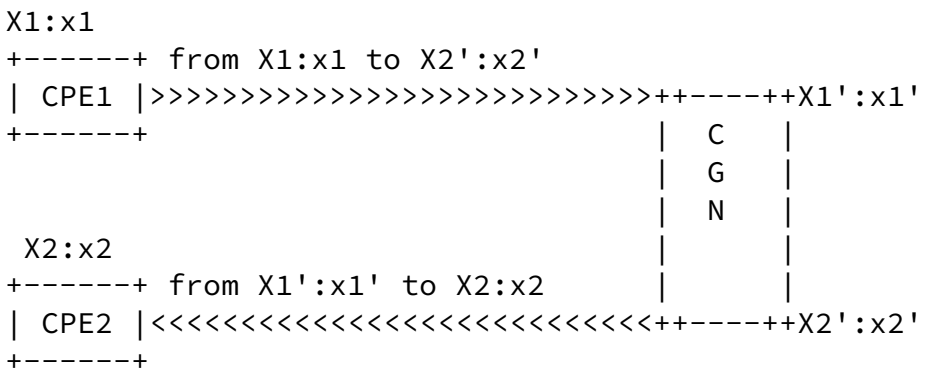

Haripinning

REQ-7: A CGN SHOULD comply with [RFC4787] for unicast UDP.

Justification: CGN SHOULD have to keep high transparency for unicast UDP communications. And CPE MAY use P2P and interactive services between CPEs after a carrier introduces CGN.

5. TCP Requirements

[I-D.ietf-behave-tcp] describes requirements of TCP of a NAT, and the behavior of "Endpoint-Independent Filtering" is RECOMMNEDED, and a NAT MUST have an "Endpoint-Independent Mapping" behavior to ensure transparency of CGN

To have "Endpoint-Independent Filtering" and "Endpoint-Independent Mapping" behaviors for CGNs, CGNs help to establish TCP Hole Punching among CPEs. In other words, the possibility of the establishment of TCP Hole Punching among CPEs which have CGN is equal to the possibility among CPEs which don's t have CGN. If CGNs have an "Address-Dependent Mapping" or "Address and Port-Dependent Mapping" behavior, the possibility that establishment of TCP Hole Punching is less than when CGNs have an "Endpoint-Independent Mapping" behavior. And if CGNs have an "Address and Port-Dependent Filtering" behavior, the possibility that establishment of TCP Hole Punching is less than when CGNs have an "Endpoint-Independent Filtering" or "Address Dependent Filtering" behavior. Because a CSS is placed external CGN realm, the source of IP address and port of the communication from CPE to CSS is CGN external IP address and port. It is RECOMMENDED to use STUN[I-D.ietf-behave-rfc3489bis] if CPEs want to check the CGN external IP address and port for CPE.

A carrier MAY introduce TURN [I-D.ietf-behave-turn] to support communications among CPEs. If CGN supports "Hairpinning", CGN can hairpin the communications between CPEs in the same CGN. Therefore the requirements of Hairpinning for CGN MAY reduce requirements for the performance of TURN servers. When CPEs decide the course of TCP between CPEs, CPE MAY use [I-D.ietf-mmusic-ice] .

REQ-8: A CGN SHOULD comply with [I-D.ietf-behave-tcp] for TCP.

Justification: CGN SHOULD have to keep high transparency for TCP communications. And CPE MAY use P2P and interactive services between CPEs after a carrier introduces CGN.

6. ICMP Requirements

[I-D.ietf-behave-nat-icmp] describes requirements of ICMP of a NAT.

And there MAY be a case that CPE cannot establish communication from CPEs to CGN external realm because CGN limits the number of CGN

external ports, identifiers and TCP sessions per CPE. It is useful if CPE can distinguish an error to occur by the limitation of the CGN external ports, identifiers and TCP sessions from other errors.

REQ-9: A CGN SHOULD comply with [I-D.ietf-behave-nat-icmp] for ICMP.

a) When a CGN can't establish new session of TCP/UDP by limiting of TCP/UDP ports per user, the CGN sends an ICMP destination unreachable message, with code of 13 (Communication administratively prohibited) to the sender.

Justification: CGN SHOULD have to keep high transparency for ICMP. And CPE MAY use P2P and interactive services between CPEs after a carrier introduces CGN. And it is necessary to be able to distinguish an error to occur by the limitation of the CGN external ports and TCP sessions from a network error.

7. Summary of Requirements

REQ-1: A CGN MUST allocate one external IP address to each CPE.

a) CGN external IP address of the UDP, TCP and ICMP MUST be same.

REQ-2: A CGN MUST allocate CGN external ports corresponding to CPE ports of UDP.

a) A CGN MUST NOT overload CGN external port while a NAT UDP mapping timer does not expire.

b) A CGN MAY overload CGN external port after a NAT UDP mapping timer expires.

c) A CGN SHOULD limit the number of the CGN external ports of UDP per CPE.

d) The number of the CGN external ports of UDP per CPE which CGN can allocate SHOULD be configurable for the administrator of CGN.

e) A CGN SHOULD NOT allocate well-known ports as CGN external ports.

REQ-3: A CGN MUST allocate CGN external ports corresponding to CPE ports of TCP.

a) A CGN MUST NOT overload CGN external port while the port is allocated for one or more TCP sessions originated by another CPE.

b) A CGN MAY reuse CGN external port while the port is allocated for no session originated by any CPE.

c) A CGN SHOULD limit the number of the CGN external ports of TCP per CPE.

d) The number of the CGN external ports of TCP per CPE SHOULD be an administratively configurable option.

e) A CGN SHOULD limit the number of the new sessions of TCP per time unit and per CPE.

f) A CGN SHOULD NOT allocate well-known ports as CGN external ports.

REQ-4: A CGN MUST allocate CGN external identifiers corresponding to CPE identifiers.

a) A CGN MUST NOT overload CGN external identifier before an ICMP Query session timer expires.

b) A CGN MAY overload CGN external identifier after an ICMP Query session timer expires.

c) A CGN SHOULD limit the number of the CGN external identifier allocated per CPE.

d) The number of the CGN external identifiers per CPE which CGN can allocate SHOULD be an administratively configurable option.

REQ-5: Reserving CGN external ports per CPE for the always-available services are RECOMENDED.

a) The destination port which is used for reservation of CGN external ports SHOULD be administratively configurable.

REQ-6: A CGN SHOULD pass-through the communication between CPEs and CSS.

REQ-7: A CGN SHOULD comply with [RFC4787] for unicast UDP.

REQ-8: A CGN SHOULD comply with [I-D.ietf-behave-tcp] for TCP.

REQ-9: A CGN SHOULD comply with [I-D.ietf-behave-nat-icmp] for ICMP.

a) When a CGN can't establish new session of TCP/UDP by limiting of TCP/UDP ports per user, the CGN sends an ICMP destination unreachable message, with code of 13 (Communication

administratively prohibited) to the sender.

8. IANA Considerations

There are no IANA considerations.

9. Security Considerations

If malicious CPE can camouflage CPE 3-Tuple, the malicious CPE MAY prevent a normal CPE from sending data to external realm. Therefore, a carrier SHOULD make p olicies to prevent a spoofing of CPE 3-tuple.

10. Acknowledgements

Thanks for the input and review by Yasuhiro Shirasaki, Takeshi Tomochika, Kousuke Shishikura, Dai Kuwabara, Tomoya Yoshida, Takanori Mizuguchi.

11. References

11.1. Normative References

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

[RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, January 2001.

[RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, January 2007.

[RFC5128] Srisuresh, P., Ford, B., and D. Kegel, "State of Peer-to-Peer (P2P) Communication across Network Address Translators (NATs)", RFC 5128, March 2008.

[I-D.ietf-behave-tcp] Guha, S., "NAT Behavioral Requirements for TCP", draft-ietf-behave-tcp-07 (work in progress), April 2007.

[I-D.ietf-behave-nat-icmp] Srisuresh, P., Ford, B., Sivakumar, S., and S. Guha, "NAT Behavioral Requirements for ICMP protocol",

draft-ietf-behave-nat-icmp-08 (work in progress), June 2008.

[I-D.ietf-behave-rfc3489bis] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for (NAT) (STUN)", draft-ietf-behave-rfc3489bis-15 (work in progress), February 2008.

[I-D.ietf-mmusic-ice] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", draft-ietf-mmusic-ice-19 (work in progress), October 2007.

[I-D.ietf-behave-turn] Rosenberg, J., Mahy, R., and P. Matthews, "Traversal Using Relays around NAT (TURN): Relay Extensions to Session

Traversal Utilities for NAT (STUN)",
draft-ietf-behave-turn-08 (work in progress), June 2008.

11.2. Informative Reference

[I-D.shirasaki-isp-shared-addr]
Miyakawa, S., Nakagawa, A., Yamaguchi, J., and H. Ashida,
"ISP Shared Address after IPv4 Address Exhaustion",
draft-shirasaki-isp-shared-addr-00 (work in progress),
June 2008.

[I-D.durand-v6ops-natv4v6v4]
Durand, A., "Distributed NAT for broadband deployments
post IPv4 exhaustion", draft-durand-v6ops-natv4v6v4-01
(work in progress), February 2008.

Authors' Addresses

Tomohiro Nishitani
NTT Communications Corporation
Tokyo Opera City Tower 21F, 3-20-2 Nishi-Shinjuku, Shinjuku-ku
Tokyo 163-1421
Japan

Phone: +81 3 6800 3214
Email: tomohiro.nishitani@ntt.com

Shin Miyakawa
NTT Communications Corporation
Tokyo Opera City Tower 21F, 3-20-2 Nishi-Shinjuku, Shinjuku-ku
Tokyo 163-1421
Japan

Phone: +81 3 6800 3262
Email: miyakawa@nttv6.jp

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).