Network Working Group   S. Jiang
Internet Draft   D. Guo
Intended status: Informational   Huawei Technologies Co., Ltd
Expires: March 22, 2010   B. Carpenter
                                                University of Auckland
                                                    September 24, 2009

An Incremental Carrier-Grade NAT (CGN) for IPv6 Transition
draft-jiang-v6ops-incremental-cgn-03.txt

Status of this Memo

Copyright Notice

Abstract

Global IPv6 deployment was slower than originally expected in the last ten years. As IPv4 address exhaustion gets closer, the IPv4/IPv6 transition issues become more critical and complicated. Host-based transition mechanisms are not able to meet the requirements while most end users are not sufficiently expert to configure or maintain these transition mechanisms. Carrier Grade NAT with integrated transition mechanisms can simplify the operation of end users during the IPv4/IPv6 migration or coexistence period. This document proposes an incremental Carrier-Grade NAT (CGN) approach for IPv6 transition. It can provide IPv6 access services for IPv6-enabled end hosts and IPv4 access services for IPv4 end hosts while remaining most of legacy IPv4 ISP networks unchanged. It is suitable for the initial stage of IPv4/IPv6 migration. Unlike CGN alone, it also supports and encourages transition towards dual-stack or IPv6-only ISP networks.

Table of Contents

1. Introduction.....3
2. An Incremental CGN Approach.....4
   2.1. Incremental CGN Approach Overview.....4
   2.2. Choice of tunnelling technology.....5
   2.3. Behaviour of Dual-stack Home Gateway.....5
   2.4. Behaviour of Dual-stack Carrier-Grade NAT.....6
   2.5. Impact for existing end hosts and remaining networks....6
   2.6. Discussion.....6
3. Migration towards IPv6 Core Network.....7
   3.1. Legacy communication in Phase 2.....8
4. Security Considerations.....8
5. IANA Considerations.....8
6. Acknowledgements.....9
7. Change Log.....9
8. References.....9
   8.1. Normative References.....9
   8.2. Informative References.....9
Author's Addresses.....11

1. Introduction

Up to now, global IPv6 deployment does not happen as was expected 10 years ago. The progress was much slower than originally expected. Network providers were hesitant to take the first move while IPv4 was and is still working well. However, IPv4 address exhaustion is now confirmed to happen soon. The dynamically-updated IPv4 Address Report [IPUSAGE] has analyzed this issue. It predicts early 2011 for IANA unallocated address pool exhaustion and middle 2012 for RIR unallocated address pool exhaustion. Based on this fact, the Internet industry appears to have reached consensus that global IPv6 deployment is inevitable and has to be done quite quickly.

IPv4/IPv6 transition issues therefore become more critical and complicated for the soon-coming global IPv6 deployment. Host-based transition mechanisms alone are not able to meet the requirements in all cases. Therefore, network supporting functions and/or new transition mechanisms with simple user-side operation are needed.

Carried Grade NAT (CGN) alone creates operational problems, but does nothing to help IPv4/IPv6 transition. In fact it allows ISPs to delay the transition, and therefore causes double transition costs (once to add CGN, and again to support IPv6).

Carrier-Grade NAT that integrates multiple transition mechanisms can simplify the operation of end user services during the IPv4/IPv6 migration or coexistence period. CGNs are deployed on the network side and managed/maintained by professionals. On the user side, new CPE devices may be needed too. They may be provided by network providers, depending on the specific business model. Dual-stack lite [DSLite] is a CGN-based solution that supports transition, but it requires the ISP to upgrade its network to IPv6 immediately. Many ISPs hesitate to do this as the first step. Theoretically, DS-Lite can be used with double encapsulation (IPv4-in-IPv6-in-IPv4) but this seems even less likely to be accepted by an ISP and is not discussed further.

This document proposes an incremental CGN approach for IPv6 transition. The approach is similar to DSLite, but the other way around. Technically, it mainly combines v4-v4 NAT with v6-over-v4 tunnelling functions along with some minor adjustment. It can provide IPv6 access services for IPv6-enabled end hosts and IPv4 access services for IPv4 end hosts, while leaving most of legacy IPv4 ISP networks unchanged. The deployment of this technology does not affect legacy IPv4 hosts with global IPv4 addresses at all. It is suitable

for the initial stage of IPv4/IPv6 migration. It also supports transition towards dual-stack or IPv6-only ISP networks.

2. An Incremental CGN Approach

Most ISP networks are still IPv4. Network providers are starting to provide IPv6 access services for end users. However, at the initial stage of IPv4/IPv6 migration, IPv4 connectivity and traffic would be the majority for most ISP networks. ISPs would like to minimize the changes on their IPv4 networks. Switching the whole ISP network into IPv6-only would be considered as a radical strategy. Switching the whole ISP network to dual stack is less radical, but introduces operational costs and complications. Although some ISPs have successfully deployed dual stack routers, others prefer not to do this as their first step in IPv6. However, they currently face two urgent pressures - to compensate for an immediate shortage of IPv4 addresses by deploying some method of address sharing, and to prepare actively for the deployment of IPv6 address space and services. The approach described in this draft addresses both of these pressures by proceeding in two phases.

2.1. Incremental CGN Approach Overview

The incremental CGN approach we propose is illustrated as the following figure.

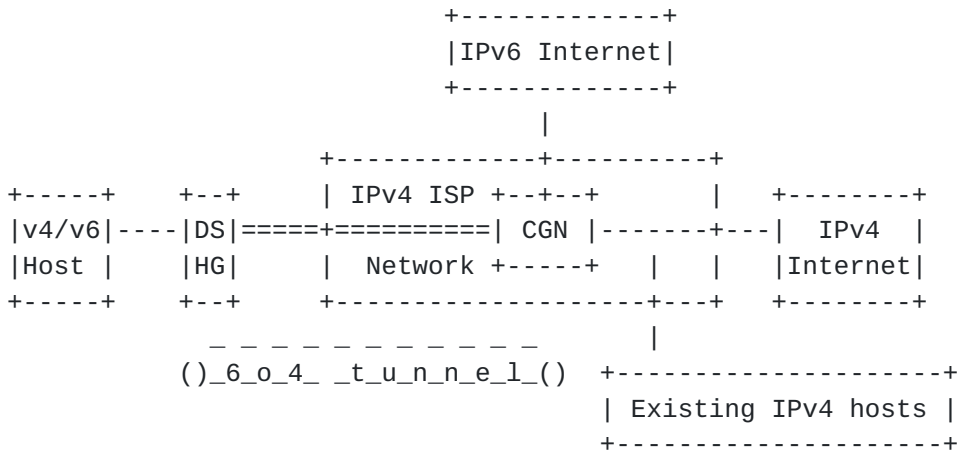

Figure 1: Phase 1 of incremental CGN approach with IPv4 ISP network

DS HG = Dual-Stack Home Gateway (CPE).

The above figure shows only Phase 1, in which the ISP has not significantly changed its IPv4 network. This approach enables IPv4 hosts to access the IPv4 Internet and IPv6 hosts to access the IPv6

Internet. A dual stack host can be treated as an IPv4 host when it uses IPv4 access service and as an IPv6 host when it uses IPv6 access service. In order to enable IPv4 hosts to access IPv6 Internet and IPv6 hosts to access IPv4 Internet, NAT-PT [RFC2766, RFC4966] (or its replacement) can be integrated with CGN. The integration of such mechanisms is out of scope for this document

Two new types of devices need to be deployed in this approach: a dual-stack home gateway, which may follow the requirements of [6CPE], and dual-stack Carrier-Grade NAT. The dual-stack home gateway integrates IPv4 forwarding and v6-over-v4 tunnelling functions. It may integrate v4-v4 NAT function, too. The dual-stack CGN integrates v6-over-v4 tunnelling and carrier-grade v4-v4 NAT functions.

2.2. Choice of tunnelling technology

In principle, this model will work with any form of tunnel between the DS HG and the dual-stack CGN. However, tunnels that require individual configuration are clearly undesirable because of their operational cost. Configured tunnels based directly on [RFC4213] are therefore not suitable. A tunnel broker according to [RFC3053] would also have high operational costs.

Modified 6RD [6RD] technology appears suitable to support v6-over-v4 tunnelling with low operational cost. Modified GRE [RFC2784] with additional auto-configuration mechanism is also suitable to support v6-over-v4 tunnelling. Other tunnelling mechanisms such as 6over4 [RFC2529], 6to4 [RFC3056], the Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) [RFC5214] or Virtual Enterprise Traversal (VET) [VET] are also considered. If the ISP has an entirely MPLS infrastructure between the CPE and the dual-stack CGN, it would also be possible to consider a 6PE [RFC4798] tunnel directly over MPLS. This would, however, only be suitable for an advanced CPE that is unlikely to be found as a home gateway, and is not further discussed here.

2.3. Behaviour of Dual-stack Home Gateway

When a dual-stack home gateway receives a data packet from an end host, it firstly checks whether the packet is IPv4 or IPv6. For IPv4 data, the HG can directly forward it if there is no v4-v4 NAT running on the HG. Or the HG translates packet source address from a HG-scope private IPv4 address into a CGN-scope private IPv4 address. The HG records the v4-v4 address mapping information for inbound packets, just like normal NAT does.

For IPv6 data, the HG needs to encapsulate the data into an IPv4 tunnel, which has the dual-stack CGN as the other end. Then the HG sends the new IPv4 packet towards CGN.

The HG records the mapping information between the tunnel and the source IPv6 address for inbound packets if HG uplinks to more than one CGN. Detailed considerations for the use of multiple CGNs by one HG are for further study.

2.4. Behaviour of Dual-stack Carrier-Grade NAT

When a dual-stack CGN receives a data packet from a dual-stack home gateway, it firstly checks whether the packet is a normal IPv4 packet or a v6-over-v4 tunnel packet. For a normal IPv4 packet, the CGN translates packet source address from a CGN-scope private IPv4 address into a public IPv4 address, and then send it to IPv4 Internet. The CGN records the v4-v4 address mapping information for inbound packets, just like normal NAT does. For a v6-over-v4 tunnel packet, the CGN needs to decapsulate it into the original IPv6 packet and then send it to IPv6 Internet. The CGN records the mapping information between the tunnel and the source IPv6 address for inbound packets.

Depending on the deployed location of the CGN, it may use v6-over-v4 tunnels to connect to the IPv6 Internet.

2.5. Impact for existing end hosts and remaining networks

This approach does not affect the remaining networks at all. Legacy IPv4 ISP networks and their IPv4 devices remain in use. The existing IPv4 hosts, shown as the right box in Figure 1, either having global IPv4 addresses or behind v4-v4 NAT can connect to IPv4 Internet as it is now. Of course, these hosts, if they are upgraded to become dual-stack hosts, can access IPv6 Internet through IPv4 ISP network by using IPv6-over-IPv4 tunnel technologies.

2.6. Discussion

For IPv4 traffic, this approach inherits all the problems of CGN (e.g., scaling, and the difficulty of supporting well-known ports for inbound traffic). Application layer problems created by double NAT are for further study.

If a different technology than v4-v4 NAT is chosen for IPv4 address sharing, for example [APLUSP], the present approach could be suitably modified, for example replacing the v4-v4 NAT function by the A+P gateway function.

However, for IPv6 traffic, a user behind the DS HG will see normal IPv6 service. We therefore observe that an IPv6 tunnel MTU of at least 1500 bytes would ensure that the mechanism does not cause excessive fragmentation of IPv6 traffic nor excessive IPv6 path MTU discovery interactions.

However, for IPv6 traffic, a user behind the DS HG will see normal IPv6 service. This, and the absence of NAT problems for IPv6, will create an incentive for users and application service providers to prefer IPv6.

ICMP filtering [RFC4890] function may be included as part of CGN functions.

3. Migration towards IPv6 Core Network

If the core network transits to IPv6, this approach can easily be transited into Phase 2, in which the ISP network is either dual-stack or IPv6-only.

For dual-stack ISP networks, dual-stack home gateways can simply switch off the v6-over-v4 function and forward both IPv6 and IPv4 traffic directly while the dual-stack CGN only keeps its v4-v4 NAT function. However, this is considered an unlikely choice, since we expect ISPs to choose the approach described here because they want to avoid dual-stack deployment completely.

For IPv6-only ISP networks, the dual-stack lite solution [DSLite], which also needs dual-stack home gateway and CGN devices, can be adopted for Phase 2. The best business model for this approach is that CPE has integrated the functions for both Phase 1 and 2, and can automatically detect the change. For example, the DS HG can use the appearance of IPv6 Route Advertisement messages or DHCPv6 messages as a signal that Phase 2 has started. Then when ISPs decide to switch from Phase 1 to Phase 2, it may be that only a configuration change or a minor software update is needed on the CGNs. The DS HG will then switch automatically to DSLite mode. The only impact on the home user will be to receive a different IPv6 prefix.

It will not be necessary for all customers of a given ISP to switch from Phase 1 to Phase 2 simultaneously; in fact it will be operationally better to switch small groups of customers (e.g. all those connected to a single point of presence). This is a matter of planning and scheduling.

3.1. Legacy communication in Phase 2

We do not expect to see IPv6-only public services as long as there is an IPv4-only customer base in the world, for obvious commercial reasons. However, especially in Phase 2, IPv4/IPv6 intercommunication may become issues. [DSLInter] describes a proposal to enhance DS-lite solution with an additional feature to ease interconnection between IPv4 and IPv6 realms. Furthermore, home users may encounter the problem of reaching legacy IPv4-only public services from IPv6-only clients. This problem could already exist in Phase 1, but will become more serious as time goes on. Each ISP can provide its IPv6-only customers with a network-layer translation service to satisfy this need. Such a service is not fully defined at this time, so we refer to it non-specifically as "NAT64". Current work in the IETF is focussed on one particular proposal [NAT64].

The NAT64 service can be provided as a common service located at the border between the ISP and the IPv4 Internet, beyond the dual stack CGN from the customer's viewpoint. It may be integrated into CGN devices too. The question has been asked why it is better to do this than to distribute the NAT64 function by locating it in (or near) the home gateway, so that relevant translation state resides only in the HG. While this might be suitable in Phase 1, when the ISP still provides full IPv4 connectivity, it would force all translated traffic into DSLite tunnels in Phase 2. This seems undesirable.

4. Security Considerations

Security issues associated with NAT have been documented in [RFC2663] and [RFC2993].

Further security analysis will be needed to understand double NAT security issues and tunnel security issues. However, since the tunnel proposed here exists entirely within a single ISP network, between the CPE and the CGN, the threat model is relatively simple. [RFC4891] describes how to protect tunnels using IPSec, but it is not clear whether this would be an important requirement. An ISP could deem its infrastructure to have sufficient security without additional protection of the tunnels.

The dual-stack home gateway will need to provide basic security for IPv6 [6CPESec]. Other aspects are described in [RFC4864].

5. IANA Considerations

This draft does not request any IANA action.

6. Acknowledgements

Useful comments were made by Fred Baker, Dan Wing, Fred Templin, Seiichi Kawamura, Remi Despres, Janos Mohacsi, Mohamed Boucadair, Shin Miyakawa and other members of the IETF V6OPS working group.

7. Change Log [RFC Editor please remove]

draft-jiang-incremental-cgn-00, original version, 2009-02-27

draft-jiang-v6ops-incremental-cgn-00, revised after comments at IETF74, 2009-04-23

draft-jiang-v6ops-incremental-cgn-01, revised after comments at v6ops mailing list, 2009-06-30

draft-jiang-v6ops-incremental-cgn-02, remove normative parts (to be documented in other WGs), 2009-07-06

draft-jiang-v6ops-incremental-cgn-03, revised after comments at v6ops mailing list, 2009-09-24

8. References

8.1. Normative References

[RFC2529] B. Carpenter, and C. Jung, "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels", RFC2529, March 1999.

[RFC2784] D. Farinacci, T. Li, S. Hanks, D. Meyer and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000.

8.2. Informative References

[RFC2663] P. Srisuresh and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, August 1999.

[RFC2766] G. Tsirtsis and P. Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)", RFC 2766, February 2000.

[RFC2993] T. Hain, "Architectural Implications of NAT", RFC 2993, November 2000.

[RFC3053] A. Durand, P. Fasano, I. Guardini and D. Lento, "IPv6 Tunnel Broker", RFC 3053, January 2001.

[RFC3056] B. Carpenter and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.

[RFC4213] E. Nordmark and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.

[RFC4798] J. De Clercq, D. Ooms, S. Prevost and F. Le Faucheur, "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)", RFC 4798, February 2007.

[RFC4864] G. Van de Velde, T. Hain, R. Droms, B. Carpenter and E. Klein, "Local Network Protection for IPv6", RFC 4864, May 2007.

[RFC4890] E. Davies and J. Mohacsi, "Recommendations for Filtering ICMPv6 Messages in Firewalls", RFC 4890, May 2007.

[RFC4891] R. Graveman, "Using IPsec to Secure IPv6-in-IPv4 Tunnels", RFC4891, May 2007.

[RFC4966] C. Aoun and E. Davies, "Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status", RFC 4966, July 2007.

[RFC5214] F. Templin, T. Gleeson and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, March 2008.

[DSLite] A. Durand, R. Droms, B. Haberman and J. Woodyatt, "Dual-stack lite broadband deployments post IPv4 exhaustion", draft-durand-softwire-dual-stack-lite-01, work in progress.

[IPUSAGE] G. Huston, IPv4 Address Report, March 2009, http://www.potaroo.net/tools/ipv4/index.html.

[6RD] R. Despres, "IPv6 Rapid Deployment on IPv4 infrastructures (6rd)", draft-despres-6rd, work in progress.

[6CPE] H. Singh, "IPv6 CPE Router Recommendations", draft-wbeebee-ipv6-cpe-router, work in progress.

[6CPESec] J. Woodyatt, "Recommended Simple Security Capabilities in Customer Premises Equipment for Providing Residential IPv6 Internet Service", draft-ietf-v6ops-cpe-simple-security, work in progress.

[APLUSP] R. Bush, O. Maennel, J. Zorz, S. Bellovin and L. Cittadini, "The A+P Approach to the IPv4 Address Shortage", draft-ymbk-aplusp, work in progress.

[VET] F. Templin, "Virtual Enterprise Traversal (VET)", draft-templin-autoconf-dhcp, work in progress.

[DSLInter] M. Boucadair, et al, "Stateless IPv4-IPv6 Interconnection in the Context of DS-lite Deployment", draft-boucadair-dslite-interco-v4v6, work in progress.

[NAT64] M. Bagnulo, P. Matthews and I. van Beijnum, "NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", draft-bagnulo-behave-nat64, work in progress.

Author's Addresses

Sheng Jiang
Huawei Technologies Co., Ltd
KuiKe Building, No.9 Xinxi Rd.,
Shang-Di Information Industry Base, Hai-Dian District, Beijing 100085
P.R. China
Phone: 86-10-82836774
Email: shengjiang@huawei.com

Dayong Guo
Huawei Technologies Co., Ltd
KuiKe Building, No.9 Xinxi Rd.,
Shang-Di Information Industry Base, Hai-Dian District, Beijing 100085
P.R. China
Phone: 86-10-82836284
Email: guoseu@huawei.com

Brian Carpenter
Department of Computer Science
University of Auckland
PB 92019
Auckland, 1142
New Zealand
Email: brian.e.carpenter@gmail.com