SCTE CABLE-TEC EXPO® '13
OCTOBER 21-24 / ATLANTA, GA

SCTE Society of Cable Telecommunications Engineers

Quantifying the User Experience Behind Carrier Grade NAT

A Technical Paper prepared for the Society of Cable Telecommunications Engineers
By

Bill Coward
Principal Network Architect
Cox Communications
6305-B Peachtree Dunwoody Road
Atlanta, Georgia 30328
404.269.6643
bill.coward@cox.com

Overview

Internet providers are facing the inevitable depletion of IPv4 space, and they must implement and support IPv6 transition technologies soon, as IPv4 will continue to be with us for many years to come. One of the most common transition technologies examined is Carrier Grade NAT (CGN) aka NAT444. Much has been discussed and documented about CGN but how will an ISP truly comprehend its behavior in their production environment. After extensive lab testing and live market trials we at Cox have partnered with Spirent to deliver concrete analytics to demonstrate what a user will experience after CGN is introduced in the network. This will allow our leadership to prescribe the appropriate policies and procedures in line with values regarding the application of CGN or other transition technologies.

Contents

Carrier Grade NAT, and in particular NAT444 is a relatively new technique in which service providers (SPs) are prepared to deploy in the hopes of extending the support of IPv4 long after public IPv4 resources have been exhausted. As with any application of a new technology, a SP must understand the business and technical repercussions of said presentation. Consequently, before a business strategy and policy can be articulated a clear technical comprehension must be obtained. This paper will attempt to address many of the greater technical challenges of CGN.

The Basics

NAT444 is an IPv4 life extension technique that allows multiple end users to ‘share’ a single public IPv4 address. This is made possible by the SP introducing a CGN appliance in the users’ traffic path. Figure 1 below depicts a central CGN appliance intercepting and translating traffic sourced from both users private 100.64.0.1 and 10.64.0.2 addresses and translating both user’s traffic to 174.24.96.1 before forwarding it to the Internet. However simple translation of the traffic is not enough to achieve our goal of sharing/multiplexing public IPv4 space. Notice that each user is allocated a unique block of TCP source ports, user 1 is allocated 2000 thru 2999 and user 2 is allocated 3000:3999, this TCP port block allocation allows the CGN to map each TCP conversation and thereby allowing multiple users to share a single public IPv4 address.

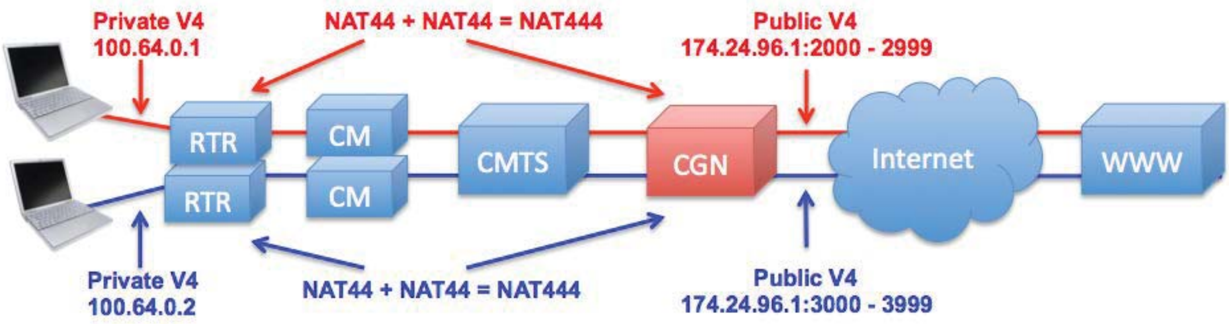

Figure 1 Basic NAT444 Concept

CGN Variables

There has been great effort in documenting and understanding this new technology particularly in the MSO community, Internet Engineering Task Force (IETF) draft “Accessing the impact of NAT444 on Network Applications” is one of the best examples of these efforts. However, several technical decisions may lead to varying results compared to that documented in the IETF draft. These technical decisions include but not limited to: CGN vendor of choice, traffic separation scheme, CGN NAT configuration and deployment architecture. Given some of these decisions have less bearing on how a particular network application will operate with CGN and others have a greater bearing on the overall CGN performance and therefore user experience.

CGN Vendor

This might be one of the most important decisions to sort, as the right partner in CGN as well as other IPv6 transition technologies is paramount to a comprehensive end-to-end IPv6 program. I would suggest that your CGN vendor must be flexible as CGN and other IPv6 transition technologies are relatively new in the Service Provider network. New protocols and processes are constantly being developed and perfected such as port control protocol (PCP) as well as application layer gateways (ALGs) that can increase the stability and reliability of CGN and your vendor of choice must possess the agility to adapt to the IPv6 transition landscape.

Traffic Separation

This technique insures that outbound and subsequent inbound IPv4 traffic from users designated to be NAT'd flows thru your CGN appliance, as not all user source traffic will be NAT'd. Traditionally this is accomplished in one of two ways, virtual routing and forwarding (VRF) or policy based routing (PBR). Both of these traffic separation techniques will meet your isolation goals but each technique has long-term management and dual-stack considerations. In that VRF separation can be considerable more technically challenging to troubleshoot furthermore delivering dual-stacked IPv6 via VRF can compound the operational challenges.

CGN Configuration

Specific CGN configurations can greatly effect network application behavior, too numerous to address in this document we will highlight the most prominent. Port allocation, dynamic of fixed, we experienced varying degrees of results; particularly when testing gaming use cases, utilizing each configuration but found fixed NAT for stability, performance and logging size. The ratio of private IP address users that will share a single public IP address may also affect the overall CGN experience, as network host expect access to all 65.535 tcp ports and not a subset allocated by CGN and may affect the application functionality and behavior. We refer to this ratio as the 'compression' ratio and it is suggested to start on the conservative side and allow ample ports per application, host and household, something in the order of 20:1. Initial deployment towards the conservative approach allows for adjustments in the future as IPv6 adoption increases in a dual-stacked strategy. Hair-pinning filtering configuration and options supported by your vendor can also have an impact on users behind CGN particularly in gaming scenarios.

Deployment Architecture

Basically we are considering a centralized deployment with the CGN appliance closer to the destination versus a distributed model with the CGN appliance closer to the source and or variations of the two models. Again not much affect on how network applications will behave behind CGN but another technical element that determines the overall user experience particular when dealing with geo-location services.

Setting User Expectations

Tradition dictates lab examination followed by user trials when deploying a new service or application into the cable network. However CGN is a bit different in that the user trial feedback will be subjective to their experience and environment. User trials have their place but with the expanse of network applications and home environments how do we formulate a comprehensive CGN business policy based on subjective user information. Consequently, we refocused our efforts to create a hybrid location of part lab and part production. This hybrid environment brings the test gear to a live production system to help us measure, simulate and document what a user can expect behind a CGN in an objective approach. Armed with this data we can then articulate the impact of CGN to the business where leadership can construct a well-informed policy.

Measuring the Experience Gap

Now with our testing topology established, depicted in figure 2 below, we can now record, compare and contrast users' traffic traversing the CGN and compare against a non-CGN user or baseline traffic. This topology also allows us to measure any internal round-trip network latency introduced by the CGN in our network and then use the latency constant with any external round trip times.

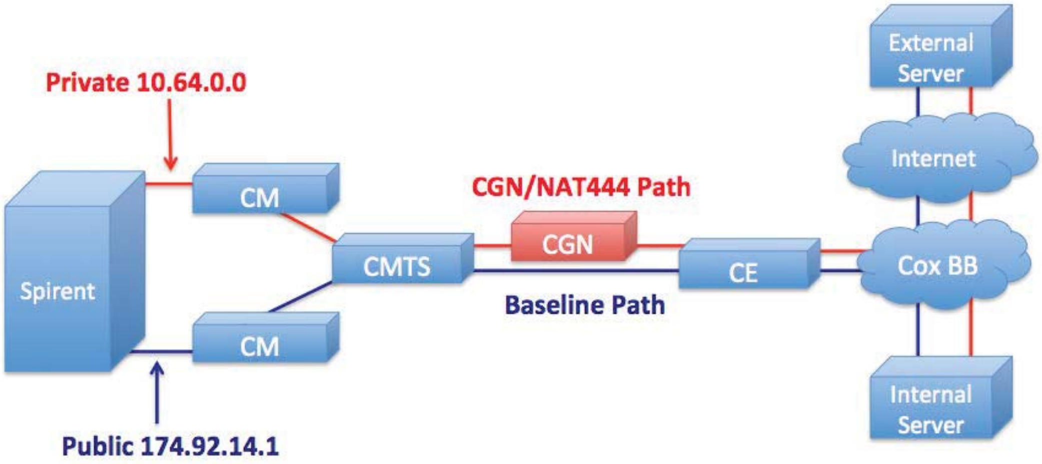

Figure 2 NAT444 Test Environments

With hundreds of networking applications enabled on the Internet we worked with our "Network Intelligence" group, the deep packet inspection (DPI) fellows, to prioritize our applications to be tested according to user demand. The top 20 networking applications are listed in Figure 3 below.

NetFlix
YouTube
HTTP
SSL
Bit Torrent
Flash
Video
RTMP
HTTPD
Amazon IV

Facebook
Skype
HLS
RTMPE
Instagram
MPEG
DASH
PlasyStation
Xbox Live
Pandora

Figure 2 NAT444 Test Environments

CGN Summation

After many hours of testing, we can report the all of the top 20 Internet network applications that transverse our network work with CGN with insignificant latency and delay. However we continue to examine the balance of applications and of those that work we expect similar results as networking applications operations with CGN are binary they work or don't. Users behind CGN will not have support for applications that require inbound originated traffic on a specific TCP port - i.e. web hosting, most server applications and some multimedia applications such as Slingbox. This list of unsupported applications is decreasing as CGN vendors, developers and SPs are implementing CGN aware solutions such as ALG, proxy servers and session traversal utilities for NAT (STUN)/ interactive connectivity establishment (ICE) like functionality. Additionally, deploying a shared IPv4 address via CGN with a native IPv6 globally unique address (GUA) in a dual-stack mode will further compensate for the 'brokenness' of CGN and thereby closing the experience gap further.

End Game

Now that we can measure and document the full effects of CGN, this does not eliminate the need and value offered by an extended user trial period. However, we are able to formulate business policies that will in part resemble the following…

Deploy Post v4 Exhaustion
Dual-Stacked with Native IPv6
New Residential Customers Only
Opt-Out Option

Abbreviations and Acronyms

ALG	Application Layer Gateway
CGN	Carrier Grade NAT
DPI	Deep Packet Inspection
ICE	Interactive Connectivity Establishment
GUA	Globally Unique Address
IETF	Internet Engineering Task Force
LSN	Large Scale NAT
MSO	Multi-System Operator
NAT	Network Address Translation
NAT44	Network Address Translation, address translation from an IPv4 address to another IPv4 address, usually private to public.
NAT444	Network Address Translation, address translation from an IPv4 address to an IPv4 address and then another IPv4 address, with the last NAT being within the SP network as the case of CGN.
PBR	Policy Based Routing
PCP	Port Control Protocol
SP	Service Provider
STUN	Session Traversal Utilities for NAT
TCP	Transport Control Protocol
VRF	Virtual Routing and Forwarding