

EXPEDITING ACCESS TO V6 SERVICES:

GETTING WEB CONTENT AVAILABLE OVER IPV6 QUICKLY AND AT LOW COST

Tim LeMaster

lemaster@juniper.net



IPV6 REALITY CHECK: THE IPV4 LONG TAIL

Post IPv4 allocation completion:

- Many hosts & applications in customer residential networks (eg Win 95/98/2000/XP, Playstations, consumer electronic devices) are IPv4-only.
- Many applications, software & servers in enterprise network are IPv4-only
 - They will not function well in an IPv6-only environment.
 - Some of those cannot or will not upgrade to IPv6.
- Content servers (web, email,...) are hosted on the Internet by many different parties. It will take time to upgrade those to IPv6.


Current measurement:

0.4% of Alexa top 1-million web sites are available via IPv6

Source: <http://ipv6monitor.comcast.net>

COMCAST IPV6 MONITOR /1

IPv6 Reachability Among Top 1M



JUNIPER VISION: A VOICE OF REASON

Religious debate:

- Ipv6 will never come [Focus only on CGN]
- Ipv6 is here now [We don't need CGN]



Juniper Vision:

- Solution provider
- CGN will be deployed
- Final goal is still Ipv6




What is the right transition technology?


- There is no right answer. Each customer is unique - therefore we provide a toolbox and help each Service Provider to make their right choice.

TRADITIONAL ARCHITECTURES

Cannot offer IPv6 services until everything is IPv6 ready.




IPv6 servers
Dedicated IPv6 load balancers




IPv6 servers
Shared IPv4/IPv6 load balancers

PREVIOUS ATTEMPTS AT ENABLING IPV6 SERVICE DELIVERY BASED ON IPV4 SERVERS

Decouple IPv6 deployment in the network from IPv6 deployment on the servers




IPv4-only servers
IPv4/IPv6 load balancers
with NAT64




IPv4-only servers
IPv4-only load balancer
NAT64 as a front-end to load balancer

Still requires deployment of IPv6 into the datacenter...

IN-PATH (ISP-BASED) TRANSLATOR ARCHITECTURE



CLOUD TRANSLATOR ARCHITECTURE



PROBLEM STATEMENT: GETTING CONTENT AVAILABLE OVER IPV6 QUICKLY

How to get example.com web site available over IPv6 quickly and at the lowest possible cost?

- Get everything dual-stack (Network, Load-balancer, Servers...)
- Get the network dual-stack and leave the servers IPv4
(Easier, as the engineering teams dealing with servers are often not the same as the ones dealing with the network)
- Don't touch anything and let some else handle the problem...

An IPv6->IPv4 translator in the cloud can do this translation for you.

IPV6 WEB SERVICE IN 3 STEPS

Step 1: Add a DNS AAAA record for www.example.com that points to a translator-in-the-cloud IPv6 address.


Step 2: Configure the translator to point back to the IPv4 address of the real web server www.example.com.

Step 3: Provide IPv6 analytics

EXAMPLE OF DEPLOYMENT

Deployed in a week in Juniper environment.

The NAT64 device can be anywhere. In our case in another datacenter with only a normal internet connection.



<http://www.juniper.net> is back to IPv4 only

<http://ipv6.juniper.net> remains IPv6 reachable

ADVANTAGES OF A ROUTER BASED SOLUTION

Can be combined with other managed services

- Stateful firewall/DPI
- Quality of Services
- VPNs (L3vpn, Ipsec, ...)
- Tiered Services

Network Design

- Possibility to share the NAT64 device between multiple end customers
- Avoid the need to had any IPv6 capacity on the end customer site

Easy integration into existing Juniper Routers (M, MX, T series)

- No backhauling to a separate device

SCALING PERFORMANCE IN 11.2

Per card (MS-DPC) performance – on average 19Gbps throughput

Metrics	NAPT44(4) PBA ¹	NAT64
Throughput	19Gbps	18Gbps
Total Flows	17M	15M
Peak Flow ² Ramp-up Rate	1.2M Flows/sec	540K Flows/sec
Public Port Pool	4B ports	4B ports
Ramp-up time (4M Flows)	4sec	8sec

¹Port Block Allocation (PBA): When PBA is configured, ports for a host are allocated in blocks. Subsequent port allocations for the same host come from the previously allocated block.

²Flow = Uni-directional flow through the Router

SCALING

A) IPv6 traffic can be scaled using classic techniques (over IPv6, of course).

- Multiple AAAA records for multiple translator
- IPv6 load balancer in front of translator
- DNS/CDN/Caching techniques (eg Akamai & others)

B) Traffic between the translator (or set of translators) to the IPv4 servers can be engineered with proper QoS.

IP FAMILY TRANSITION SERVICES ON MS-PIC/MS-DPC

IPv6 Features

- IPv6 NAT and IPv6 Stateful Firewall
- NAT-PT Supported (ICMP ALG)
- NAT-PT DNS ALG (10.4)
- NAT66 supported
- NAT64 (10.4)

NAT44

- Support CGN requirement
(draft-ietf-behave-lsn-requirements-00)

IPv6 Softwire

- DS-Lite (10.4)
- 4over6 (10.4)
- 6rd/6to4/6to4-pmt (1H2011)



6 MS-DPC supported by Single MX Chassis

8 MS-DPC per Chassis in 2H2011



NAT FEATURE SET

Support for a large type of NAT (NAT44, NAPT44, NAT66, NAT-PT, NAT64, NAPT66, Twice-NAT)

Standard NAT Features

- ❖ TCP/UDP/ICMP configurable timeouts and TCP Keep-Alives
- ❖ Large number of Application Level Gateways (Bootp, RPC, rsh, FTP, H323, ICMP, IIOP, SMB, Netshow, Realaudio, RTSP, Snmp, Sqlnet, TFTP, Traceroute, Winframe, DNS, SIP, PPTP)
- ❖ NAT MIB
- ❖ Port Limit per private IP

draft-ietf-behave-lsn-requirements

- ❖ EIM/EIF
- ❖ Air pinning
- ❖ Address Pooling paired

Logging Improvement

- ❖ Port bucket allocation (11.2)

Load-Balancing across Service Cards

- ❖ 1 + 1 Warm Standby
- ❖ 1 + N Warm Standby
- ❖ Active/Active Stateless load balancing

O&M commands

- ❖ alarms to monitor NAT pool, mapping, session state, etc
- ❖ monitor total sessions, sessions/sec, sessions lifetime, etc

Tight Routing integration


- ❖ VRF/6PE/6VPE support
- ❖ CGN Bypass

Service Chaining (IDS/IDP, Stateful Firewall, ...)

SUMMARY OF CGN TECHNOLOGIES

Method	Pros	Cons *
NAT44(4)	<ul style="list-style-type: none"> • Extends the life of IPv4 • No change to CPE/least impact to existing infrastructure 	<ul style="list-style-type: none"> • Two layers of NAT breaks some applications • Session state requirements • Address overlap
NAT64	<ul style="list-style-type: none"> • Extends the life of IPv4 • Transparent to end-users/cost savings 	<ul style="list-style-type: none"> • Requires ALGs • Depends on external DNS64 translation
6rd	<ul style="list-style-type: none"> • Allows for rapid deployment of IPv6 without significant infrastructure changes 	<ul style="list-style-type: none"> • CPE must be updated/replaced to add support for 6rd • Proximity to 6rd relay/AFTR affects geo-location
DS-LITE	<ul style="list-style-type: none"> • Requires a IPv6 access network • Only one layer of NAT 	<ul style="list-style-type: none"> • Proximity to 6rd relay/AFTR affects geo-location • DHCP servers will like require enhancement to support DHCP option • CPE must be upgraded/replaced to support DS-LITE

* Logging requirements extensive due to CALEA (i.e. not longer 1:1 mapping of public IP address to Subscriber)



everywhere

NOTES ABOUT JUNIPER & WORLD IPV6 DAY

Juniper was only using AKAMAI to insert the IPv6 AAAA records, not for traffic caching or acceleration

- No ssl support. Akamai only supported for v6 day

IPv6 traffic was directed to a Juniper M10i acting as “translator-in-the-cloud”

Juniper experienced a **global DNS outage** around 6am EDT. This was not related to the world IPv6 day activities.

Juniper IT also deployed native IPv6 on the Sunnyvale campus, on select wireline jacks and on the secure wireless network.


AKAMAI VIEW OF HTTP://WWW.JUNIPER.NET

DNS Traffic


IPv6 AAAA

IPv4 A

% DNS Requests per Traffic Target



DNS Requests per Second, per Traffic Target



ANALYSIS OF DNS REQUEST ORIGIN

Total number of unique DNS resolver: 49715

Number of unique DNS resolver asking for A: 48895


Number of unique DNS resolver asking for AAAA/ANY: 14467 **(30%)**

Number of unique DNS resolvers asking A data over IPv6 transport: 0

Number of unique DNS resolvers asking AAAA/ANY data over IPv6 transport: 144

RIPE-NCC VIEW

Juniper lost DNS



JUNIPER STATISTICS


Translator:

- **Outages: 0**
 - Peak Total TCP Flows Active : 4772
 - Peak Total UDP Flows Active : 35168
 - A lot of traceroute traffic...
 - Total packet translated: **13.5 million**

Web server: (stats for traffic coming from translator)

- **Outages: 0**
 - Total Hits 298,906
 - Visitor Hits 268,661
 - Spider Hits 30,245
 - Page Views 62,264
 - **Total Bandwidth 5 GB**


BANDWIDTH UTILIZATION ON TRANSLATOR




Input [bits/s]	Avg: 956.9 k	Min: 413.7 k	Max: 4.1 M	Cnt: 286
	SD: 318.0 k	P99: 2.1 M	P95: 1.5 M	P90: 1.3 M
Output [bits/s]	Avg: 961.1 k	Min: 415.1 k	Max: 4.1 M	Cnt: 286
	SD: 319.4 k	P99: 2.1 M	P95: 1.5 M	P90: 1.3 M


SUNNYVALE CAMPUS IPV6 UTILIZATION

2 Day View



8 Day View





Tim LeMaster

lemaster@juniper.net

