v6ops Working Group	E. Vyncke
Internet-Draft	G. Van de Velde
Intended status: Informational	Cisco Systems
Expires: September 9, 2009	March 8, 2009

TOC

IPv6 Deployment and Statistics at a Conference

Status of this Memo

Copyright Notice

Abstract

During the Cisco [Cisco] (, “Cisco Systems (http://www.cisco.com),” .) European networkers Conference 2009 that ran from 26th to 29th January in Barcelona native IPv6 was added to the traditional IPv4 infrastructure. During this conference the 3500 attendees had dual stack access to both IPv4 and IPv6 simultaneously. The goal of this IPv6 deployment project was to gather usage statistics in a situation where the end-user just wants to access his/her enterprise VPN or simply get onto the Internet. The collected statistics are not only useful per se but this document presents easy ways to measure the quality of the IPv6 connectivity offered on such events. In essence the users were not conducting IPv6 technology tests, but were just using Internet services. The statistics collected give some pieces of information on the size and impact of IPv6 onto the normal userbase and will also derive the importance of IPv6 onto the infrastructiure and end-user operating systems and firewall technologies. The experiment ran in collaboration with Google [Google] (, “Google (http://www.google.com/),” .) and Tata-Communications [Tata] (, “Tata Communications (http://www.tatacommunications.com/),” .).

Table of Contents

TOC

1. Introduction

Traditionally there is at this conference in Europe about 3500 attendees. The nature of the attendees is technical with a heavy focus on networkes and 90% will have a laptop but will not use them all the time (lack of time and lack of electrical power). The main elements the attendees expect from the network infrastructure is that its stable and easy accessible. The attendees use the network Infrastructure to access their enterprise network via a VPN and Internet content (hotels, restaurants, email, etc..).
During the conference, about 10% of the users were made aware at the IPv6 session slots that IPv6 connectivity was available in native dual stack, however, the vast majority did not make special efforts to run IPv6 or even enable IPv6 on their end-user PC's.
The experiment was run in three steps: (1) Networkers infrastructure was enabled for dual stack, (2) The local DNS server was white-listed by Google as content provider to give AAAA records for Google content and (3) the Router Advertisements were set to ask the users to use DHCPv6 for address configuration (M-flag) instead of IPv6 Stateless Address Autoconfiguration (SLAAC). Each of the steps was run with a lenght of 24 hours.
This paper gives a summary insight in the statistics and the topology used for the IPv6 connectivity during each step. It also reports that everything kept working as expected and that the end-users were not aware they were using IPv6 as a foundation communication protocol in addition to IPv4: it was completely transparent for them.

---

2. Conference Topology Summary

TOC

The Network infrastructure at Cisco European Networkers conference is deployed in a four-storey building. For each floor has a dedicated routed IPv4 subnet available with a dedicated number of RFC 1918 IPv4 addresses. These addresses are translated at the edge of the network from private into a public address by a Network Address Translator. The wireless Access Points do also have smart services installed including end-system security (captive portal) and inter-AP roaming capability.
For the IPv6 infrastructure, it was selected to create a single Layer-2 domain for the full Networkers conference spanning all 4 floors. In contrast, the number of stations per floor is not an issue with an IPv6 /64 subnet. This made the IPv6 deployment more simpler, eventhough it was needed to add some Ethernet protocol type filtering in place at the layer 2 of the OSI layer to separate IPv4 from IPv6: on the same topology (4 wireless LAN), there was both a 4 IPv4 subnets and a single IPv6 subnet.
At the edge of network, an IPv6 router provided IPv6 connectivity through an IPv6-in-IPv4 tunnel to Tata Communication POP in Paris.

Note: initially we were planning for a real native dual-stack connectivity over the local loop, but it appeared a too costly option when only a very small incremental budget is available (like almost 0 USD incremental cost). It was because of the kind sponsorship of Tata Communications that IPv6 was made available at zero incremental cost involved towards the conference leadership.
To gather statistics from the deployment Netflow v9 and SNMP were used as well as regular shell access to network equipments. For the management, an old laptop was used running a Linux distribution. The statistics and traffic data from the event can be found http://www.cisconetworkers6.com/. A graphical representation of the topology can be found at http://www.cisconetworkers6.com/network/.
The same Linux laptop run a DNS server and was offering the statistics over HTTP access. All HTTP accesses were logged including the User-Agent header in order to collect statistics about the browser and the operating system. The HTTP 64 aware [HTTP-64AWARE] (Vyncke, E., "IPv6 Connectivity Check and Redirection by HTTP Servers," 2008.) technique was also used to force stations to bypass the address selection policy and use IPv6.
The Following applications were used as supporting infrastructure: (1) MRTG (SNMP poll) [MRTG] (Oetiker, Tobi., "The Multi Router Traffic Grapher (http://oss.oetiker.ch/mrtg/)," .), (2) NFsen (supporting Netflow v9 with IPv6 support) [NFSen] (, "NfSen - Netflow Sensor (http://nfsen.sourceforge.net/)," .), (3) NDPMON (to monitor ND activities) [NDPMON] (, "NDPMon - IPv6 Neighbor Discovery Protocol Monitor (http://ndpmon.sourceforge.net/)," .), (4) RAMOND (to monitor RA activities) [RAMOND] (Morse, James., "RAMOND (http://ramond.sourceforge.net/), University of Southampton," .), DHCPv6 Server (Cisco IOS DHCPv6 Server).

---

3. Testing Steps and Procedure

TOC

*Monday 26th of January 9:00: the MRTG & Netflow collector are connected to the Network;

*Tuesday 27th in the late afternoon, Google applies the DNS-trick and starts serving A & AAAA to all laptops using the local DNS server (only announced over DHCPv6)

*Wednesday 28th in the morning, the local DHCPv4 also serves a Google-trick-enabled DNS servers to all DHCPv4 clients, i.e., all laptops actually received the A and AAAA for Google

*Wednesday 28th 16:00, Router Advertisements include the M-flag for a few minutes;

*Thursday 29th 9:0, RA includes the M-flag all day;

*Thursday 29th 14:00, RA prefix is advertised with a calendar lifetime to 16:15;

*Thursday 29th 16:15, no more IPv6 traffic (thanks to RA prefix expiration), the router & collector are removed from the site;

---

4. Statistics

TOC

4.1. Round-Trip Time

TOC

The RTT was measured with ICMP echo requests & replies to www.google.com and www.6net.org over IPv6. In general the time to google remained constant from the conference (65 msec), while the RTT to 6net was lower 50% until the AAAA to Google was enabled (potentially due to higher IPv6 traffic load?).
The RTT to Google was also measured over IPv4 (average 64 msec, peak 218 msec) and IPv6 (average 65 msec, peak 72 msec). The use of a 6-in-4 tunnel was not really impacting the latency to access Google if we assume that ICMP is a good measurement methodology.

4.2. Stats for IPv6 Traffic to the Internet

TOC

What can be seen in the graphs at http://www.cisconetworkers6.com/mrtg/tunnel.html for the 6-in-4 tunnel is that there was a continuous growth in IPv6 traffic. There was many less people on Monday than on the rest of the week due to the conference organization and agenda.
Monday the inbound traffic (Internet to the conference) was around 100kbps, Tuesday 400kbps, Wednesday 1.5Mbps and Thursday also around 1.5 Mbps.
Outbound (conference to the Internet) can be seen for the 6-in-4 tunnel; there was also ongoing growth in IPv6 traffic. Monday the traffic was low and around 10kbps, Tuesday 50kbps, Wednesday 100kbps and Thursday also around 100kbps.

TOC

4.3. IPv6 DHCP Clients

Details to be found at http://www.cisconetworkers6.com/mrtg/dhcpv6.html. From day 1, the Cisco IOS router was configured as a DHCPv6 IA server. However, only from Wednesday onwards, the M-bit was set in the Router Advertisements. This setting had a very interresting result because it made the number of DHCPv6 assigned addresses grow from 4 to a total 151 systems using DHCPv6 IPv6 address allocation. This is a clear indication that there were only 4 laptops with a IPv6 stack always trying to use statefull DHCPv6 while the vast majority of the laptops were only using RA for SLAAC. It can be assumed that the majority of the laptops were running Microsoft Windows XP or Vista.

4.4. IPv6 Neighbours

TOC

Details regarding the amount of IPv6 neighbors of the router can be found at http://www.cisconetworkers6.com/mrtg/neighbours.html. The neighbors were split in two categories based on the IPv6 address: link-local or global addresses.
A couple of interesting observations:

- the number of IPv6 neighbors grew on every single day during the conference, with a high peak on Wednesday 28th January.

- many systems had IPv6 enabled (555 stations) as seen by the Link-local neighborships, however at maximum 358 were actually using IPv6 to access the Internet.

- the number of link-local neighbors was similar on all days (except on Monday because there were less attendees). But, the number of global address neighbors changed dramatically on Wednesday morning when the AAAA for Google was served: it doubled from 180 to 358.

4.5. DNS Requests

TOC

The traffic of DNS requests was also measured: 6 DNS requests/sec over IPv4 and 2 DNS requets/sec over IPv6. This is probably linked to the OS used on the laptops where the majority was probably Windows XP which only use IPv4 for DNS access.

4.6. Web Server Access

TOC

In order to collect operating system information of the attendees, a challenge was announced in order to attract users on a dual-stack web server. Based on the number of attendees and the number of IPv6 neighbors, it is clear that the number of visitors (about 100) was only a small parts of the local IPv6 hosts (about 555) or even attendees (about 3,500).

- *IPv6-enabled visitors: 20 *nix, 20 Microsoft Windows XP, 16 Microsoft Windows CE (smart-phones), 16 Apple Mac OS/X, 13 Microsoft Windows Vista, 6 Symbian (smart-phones)

- *IPv4-only visitors: 70 Microsoft Windows XP, 13 Apple Mac OS/X, 9 Windows Vista and 25 *nic

---

4.7. Netflow Information

TOC

Based on Netflow information, some data-points were collected over the 4 days:

- *The IPv6 protocols: 84% for TCP, 10% for ICMPv6, the rest (7%) for UDP. There were also a very small amount of IPv6 datagrams with 59 (No Next Header for IPv6): 73 packets on a total of more than 8.7 millions.

- *The layer-4 ports: 31% for NNTP, 30% for HTTP, 11% for SSH, 5% for HTTPS, 3% for DNS. The rest was not clearly identified. This may not be statiscally significant, but, the HTTP traffic went from 22% to 42% as soon as the AAAA for Google was served.

- *The Internet hosts: 31% for 2001:888::119 (newszilla6.xs4all.nl), 7% for 2001:4860:0:1001::68 (Google), 5% for 2001:4860:0:1001::53 (Google), 5% for a site in the .NO domain, 5% for a site in the .SE domain, 2% for a site in a .CZ domain (as the addresses seem to relate to a non-public site, the authors preferred to keep those addresses non-public as well).

4.8. Security

TOC

RAMOND and NDPMON detected not a single attack against the Neighbor Discovery Protocol.

5. Areas for Improvements

TOC

A couple of areas for improvement have been identified after the experiment:

- *Compare the Netflow information for IPv4 and IPv6.

- *Collect information about all the operating systems (both IPv4 and IPv6), this could be done by sniffing the HTTP traffic and collecting the User-Agent. This measurement will probably require some legal advice in some countries...

6. IANA Considerations

TOC

There are no extra IANA consideration for this document.

7. Security Considerations

TOC

There are no extra Security consideration for this document.

8. Acknowledgements

TOC

Many thanks go to Tata Communications and Yves Poppe for the sponsorship of the IPv6 Connectivity. We would also like to thank Erik Kline and Lorenzo Colitti from Google to have supported the deployment to enable AAAA DNS records for the Networkers Conference. All of this experimenting would not of been possible without the help from Cisco Networkers NOC team under leadership of Andy Phillips.

9. References

TOC

9.1. Normative References

TOC

9.2. Informative References

TOC

[Tata]	“Tata Communications (http://www.tatacommunications.com/).”
[Google]	“Google (http://www.google.com/).”
[Cisco]	“Cisco Systems (http://www.cisco.com).”
[MRTG]	Oetiker, Tobi., “The Multi Router Traffic Grapher (http://oss.oetiker.ch/mrtg/).”
[NFSen]	“NfSen - Netflow Sensor (http://nfsen.sourceforge.net/).”
[NDPMON]	“NDPMon - IPv6 Neighbor Discovery Protocol Monitor (http://ndpmon.sourceforge.net/).”
[RAMOND]	Morse, James., “RAMOND (http://ramond.sourceforge.net/), University of Southampton.”
[HTTP-64AWARE]	Vyncke, E., “IPv6 Connectivity Check and Redirection by HTTP Servers,” 2008.

Authors' Addresses

TOC

	Eric Vyncke
	Cisco Systems
	De Kleetlaan 6a
	Diegem 1831
	Belgium
Phone:	+32 2 778 4677
Email:	evyncke@cisco.com
	Gunter Van de Velde
	Cisco Systems
	De Kleetlaan 6a
	Diegem 1831
	Belgium
Phone:	+32 2704 5473
Email:	gunter@cisco.com