

ULA Experience in JANOG34

IETF 90th TORONTO v6ops

Yasuhiro Ohara

2014/07/21

Standpoint

- This is nothing to do with NTT.
- Delegated by Shishio Tsuchiya, Cisco.
- Will not discuss whether the network design is GOOD or BAD.
- We just wanted to check
 - ULA with translators (both stateless/stateful)
 - Simple users (just issue outgoing connections. No servers)
- Identify the problems when using the network design
 - Irrespective of whether the design is GOOD or BAD.

Summary


- Adding ULA to GUA+IPv4(RFC1918) seemed harmless
- Coexistence of SLAAC/ULA, DHCPv6/ULA, Stateless prefix NAT66, and stateful DNS64/NAT64 in a segment is possible.
- Tested mainly by the Mac laptop.
- Not tested:
 - Mixture of nodes having different address types
 - Incoming connections (for servers)
 - SIP
 - And others.

About JANOG

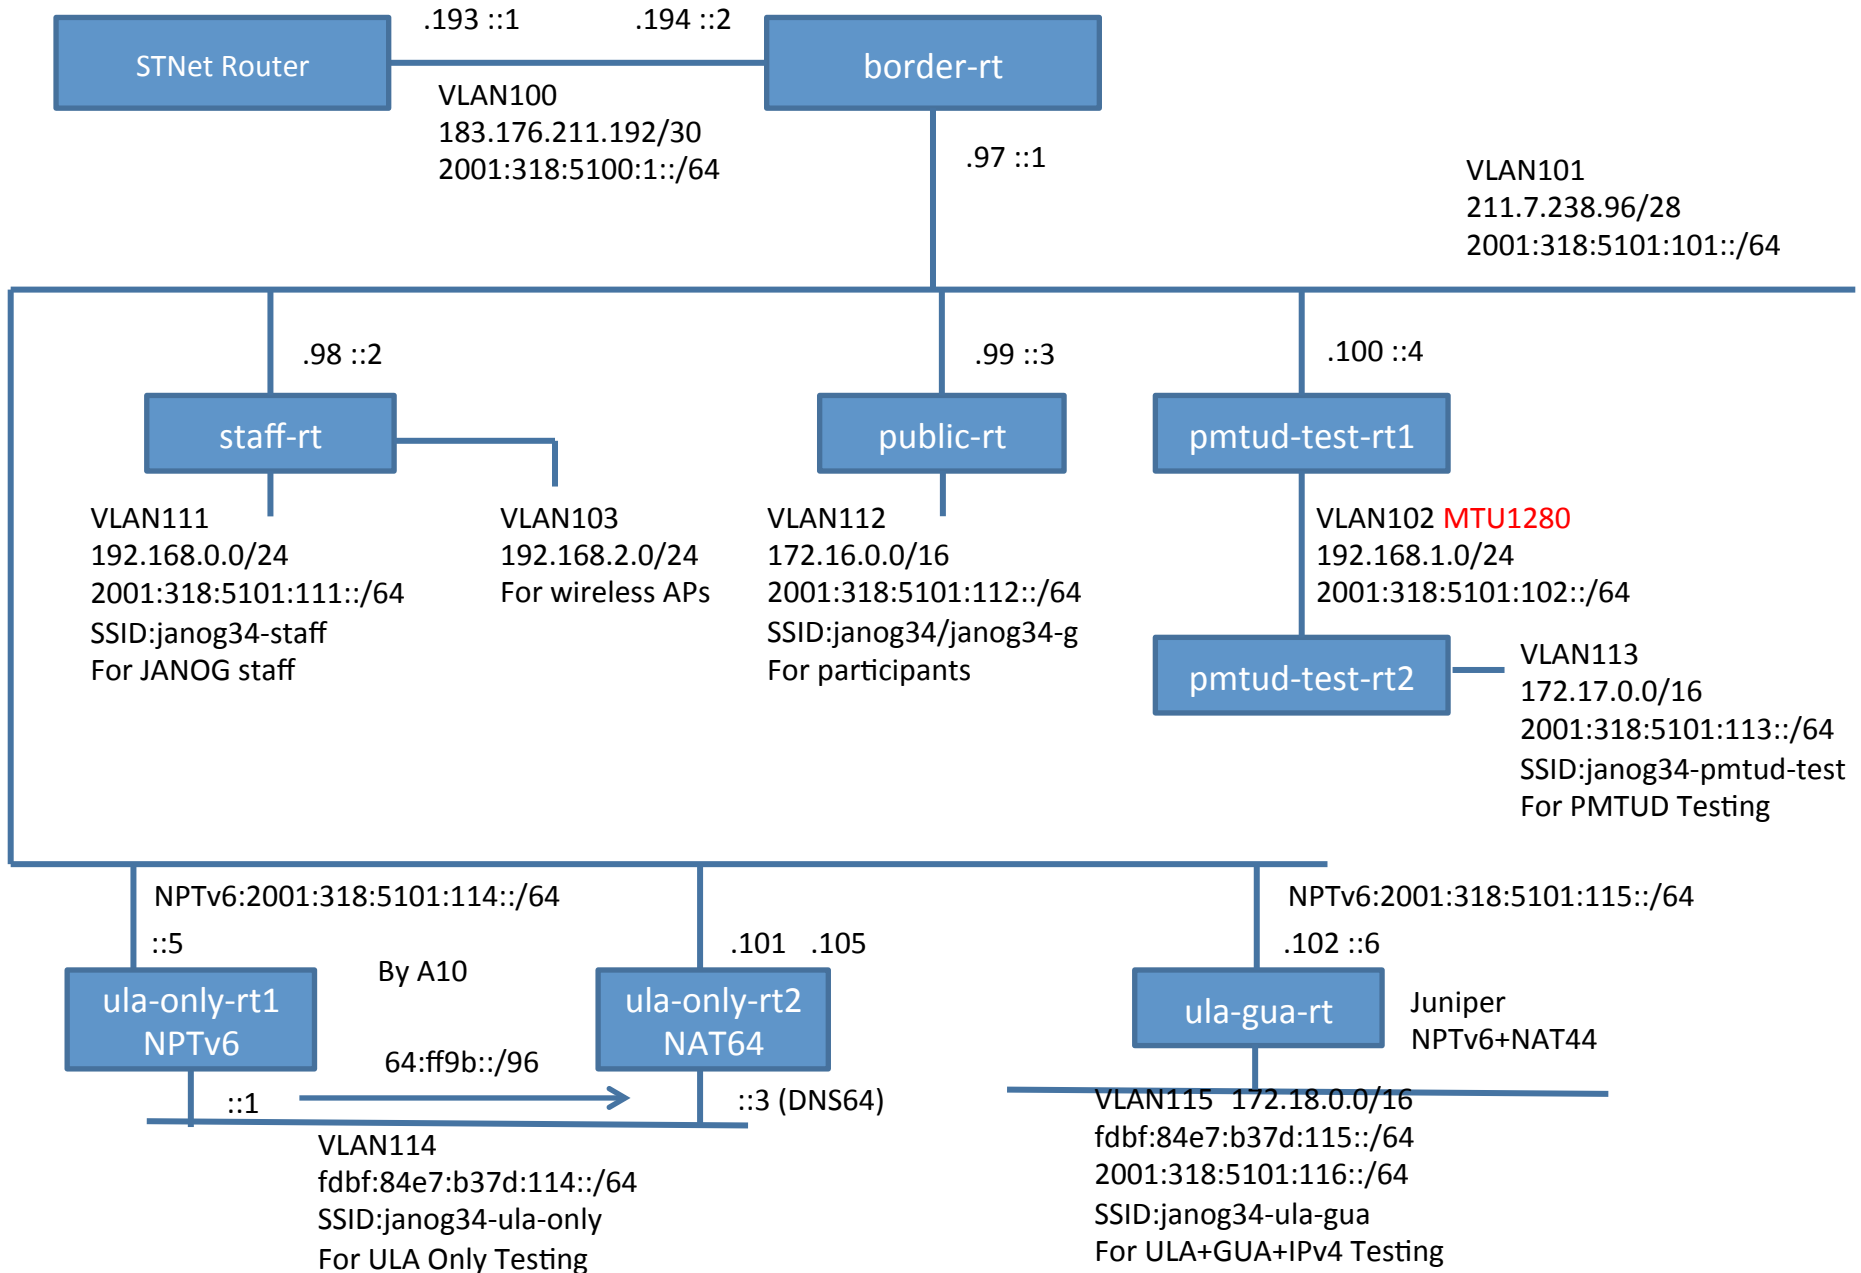
- Japan Network Operators Group (JANOG)
- Meetings are held three times a year.
- 7/17-18 Takamatsu, Kagawa-Pref, Japan.
- 547 participants. Streaming 217(max).

- Served by only one PC
 - Fujitsu RX200 S7
 - VMWare with several VMs.
 - Juniper FireFly, A10 vThunder

Peak traffic is approx. (140+40) Mbps




JANOG34 Logical Network Topology




The ULA+GUA+IPv4 segment

- janog34-ula-gua
- ULA/GUA/Private-IPv4 addresses are distributed by SLAAC and DHCP.
- DHCPv6 is used to inform the DNS server.
- Even ULA addresses are translated to GUA by stateless prefix NAT.
- Apps:
 - OK: Mail.app, iCal, ssh, Chrome (Yahoo, Google, Amazon, YouTube, Facebook, Rakuten, Wikipedia, Twitter, Goo, Nicovideo), Skype
- Maximum number of SSID associated: 13
- No communication happened using ULA (only appeared in the contents of ND and mDNS)



NAT Sessions: peak 300



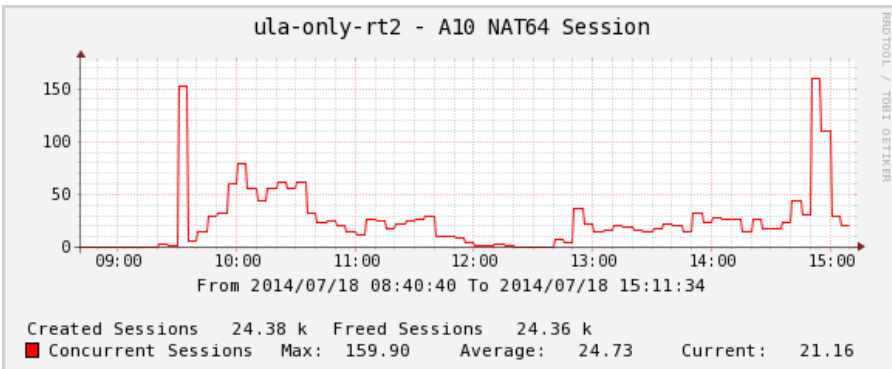
Traffic: peak (1 + 3) Mbps

The ULA+GUA+IPv4 segment


```
en0: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500
  ether 78:31:c1:c4:f6:4c
  inet6 fe80::7a31:c1ff:fec4:f64c%en0 prefixlen 64 scopeid 0x4   IPv6 Linklocal
  inet 172.18.1.16 netmask 0xffff0000 broadcast 172.18.255.255 IPv4 Private(RFC1918)
  inet6 2001:318:5101:116:7a31:c1ff:fec4:f64c prefixlen 64 autoconf SLAAC/GUA
  inet6 2001:318:5101:116:c916:1500:60fc:60d1 prefixlen 64 autoconf temporary SLAAC/GUA
  privacy
  inet6 fdbf:84e7:b37d:115:7a31:c1ff:fec4:f64c prefixlen 64 autoconf SLAAC/ULA
  inet6 fdbf:84e7:b37d:115:a8eb:7b88:af10:a490 prefixlen 64 autoconf temporary SLAAC/
  ULA privacy
  nd6 options=1<PERFORMNUD>
  media: autoselect
  status: active
```

The ULA only segment

- SSID janog34-ula-only
- Only ULA addresses are distributed by both SLAAC and DHCPv6.
- No IPv4 addresses.
- Apps:
 - OK: Mail.app, iCal, ssh, Chrome (Yahoo, Google, Amazon, YouTube, Facebook, Rakuten, Wikipedia, Twitter, Goo, Nicovideo)
 - NG: Skype
- Maximum number of SSID associated: 8.



NAT Sessions: peak 150

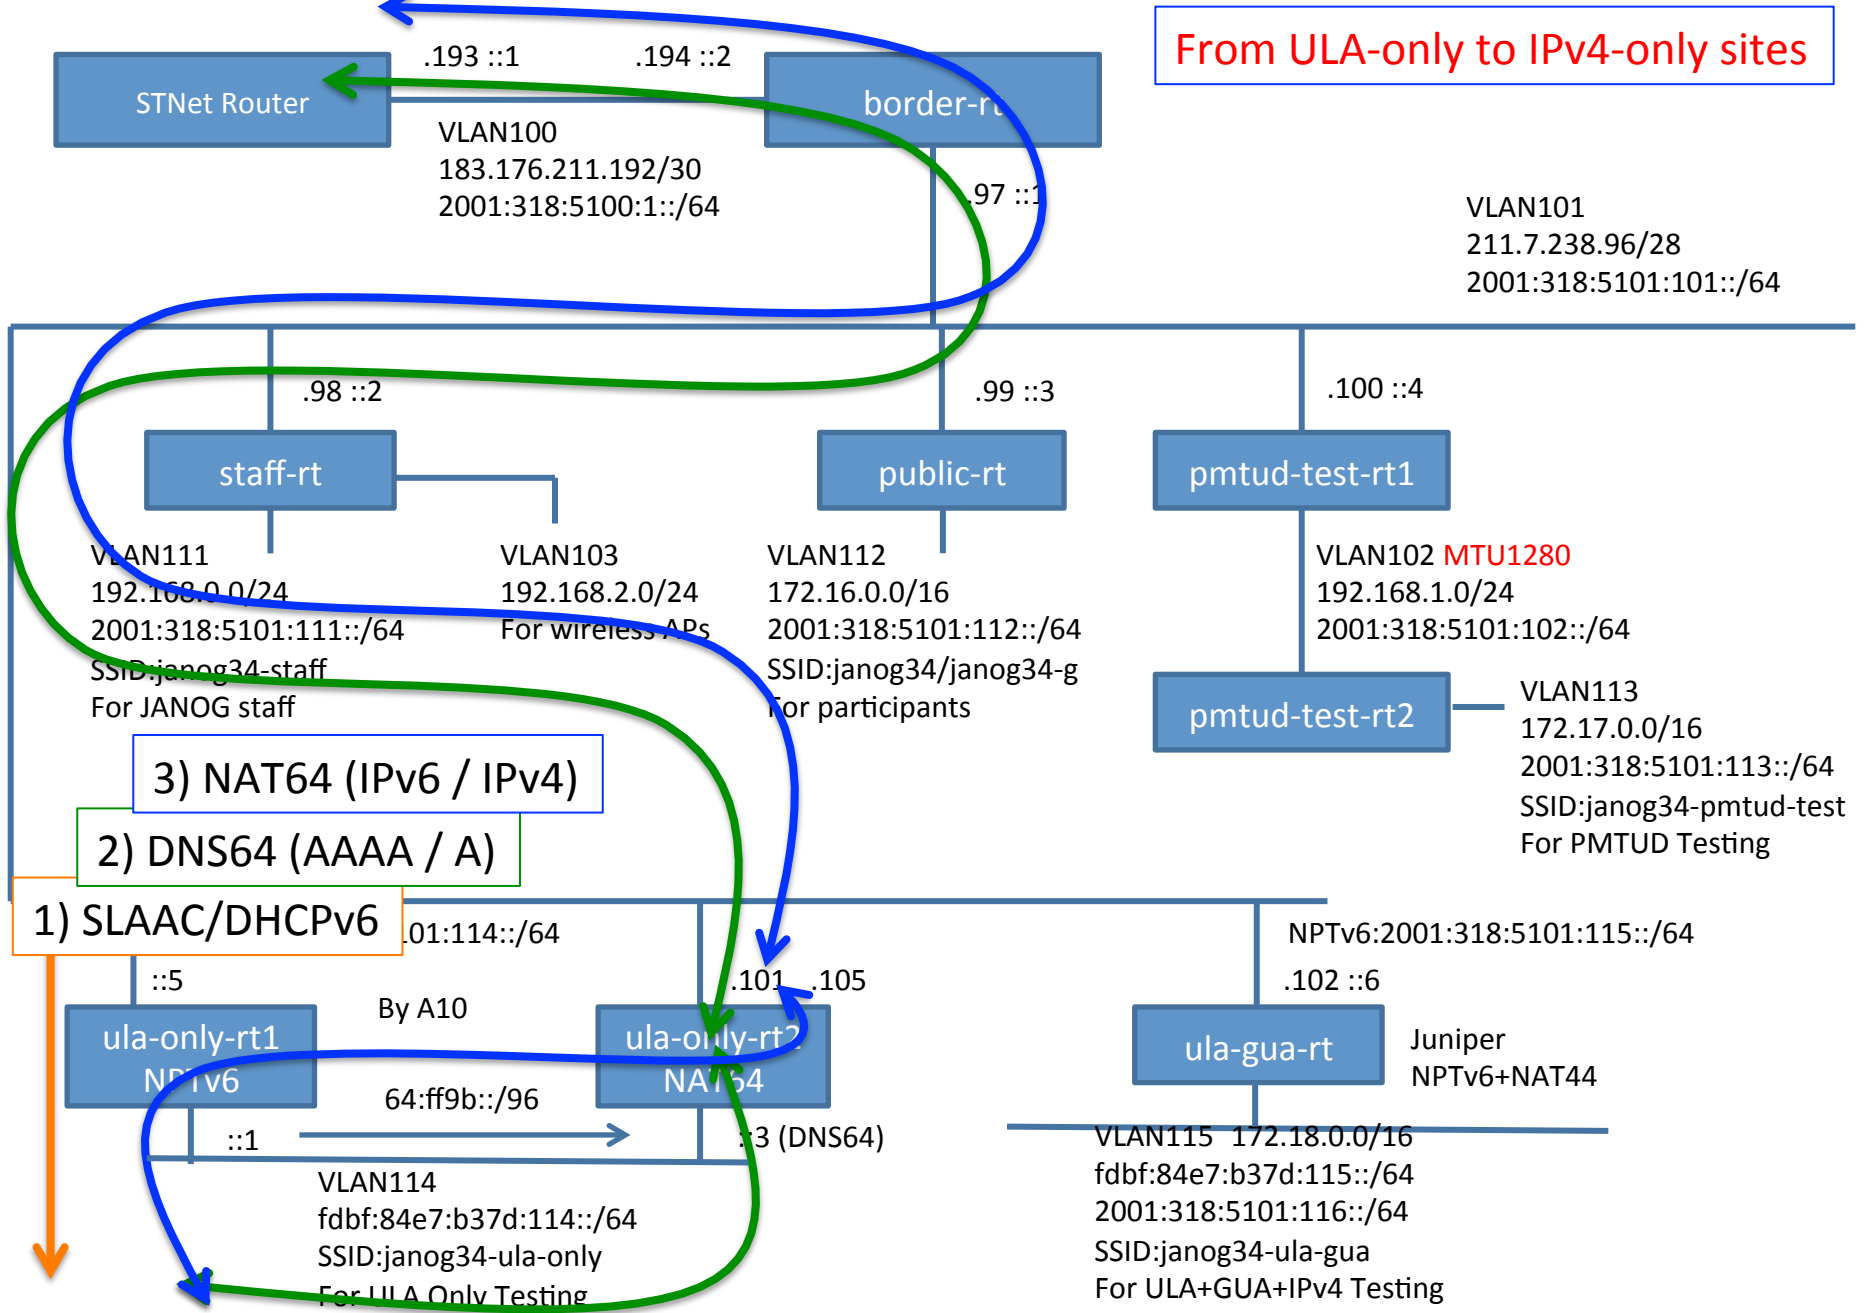


Traffic: peak (2 + 2)Mbps

The ULA only segment

```
en0: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST>  
mtu 1500  
    ether 78:31:c1:c4:f6:4c  
    inet6 fe80::7a31:c1ff:fec4:f64c%en0 prefixlen 64 scopeid 0x4 IPv6  
LinkLocal  
    inet6 fdbf:84e7:b37d:114:7a31:c1ff:fec4:f64c prefixlen 64 autoconf  
SLAAC/ULA  
    inet6 fdbf:84e7:b37d:114:b95f:55e:8939:dfbd prefixlen 64 autoconf  
temporary SLAAC/ULA privacy  
    inet6 fdbf:84e7:b37d:114:a10:a10:a10:f15f prefixlen 64 DHCPv6/ULA  
    inet 169.254.126.93 netmask 0xffff0000 broadcast 169.254.255.255 IPv4  
LinkLocal (RFC3927)  
    nd6 options=1<PERFORMNUD>  
    media: autoselect  
    status: active
```

JANOG34 Logical Network Topology



DNS64@ULA-only

```
% dig www.yahoo.co.jp AAAA
```

```
;; ANSWER SECTION:
```

```
www.yahoo.co.jp.      893    IN     CNAME  www.g.yahoo.co.jp.  
www.g.yahoo.co.jp.   60     IN     AAAA   64:ff9b::7697:e7e7
```

64:ff9b::/96 was used as DNS64/NAT64 prefix.

The IPv4 address of IPv4-only site is encoded in the 64:ff9b::/96 space.

The route to this prefix is redirected to NAT64 GW in the same segment.

The DHCPv6/ULA was used as the source for (and only for) ND and DNS ?

Questionable. Why doesn't it use the SLAAC/ULA ?

20:31:15.879134 IP6 fdbf:84e7:b37d:114:c9c4:6553:c915:f183 > ff02::1:ff00:3: ICMP6, neighbor solicitation, who has fdbf:84e7:b37d:114::3, length 32

20:31:16.745847 IP6 fdbf:84e7:b37d:114:7a31:c1ff:fec4:f64c > ff02::1:ff00:3: ICMP6, neighbor solicitation, who has fdbf:84e7:b37d:114::3, length 32

20:31:17.757020 IP6 fdbf:84e7:b37d:114:a10:a10:a10:f15f > ff02::1:ff00:3: ICMP6, neighbor solicitation, who has fdbf:84e7:b37d:114::3, length 32

20:31:19.163810 IP6 fdbf:84e7:b37d:114:a10:a10:a10:f15f > ff02::1:ff00:3: ICMP6, neighbor solicitation, who has fdbf:84e7:b37d:114::3, length 32

20:31:19.779419 IP6 fdbf:84e7:b37d:114:a10:a10:a10:f15f > ff02::1:ff00:3: ICMP6, neighbor solicitation, who has fdbf:84e7:b37d:114::3, length 32

:
:

20:31:35.875976 IP6 fdbf:84e7:b37d:114:a10:a10:a10:f15f > ff02::1:ff00:3: ICMP6, neighbor solicitation, who has fdbf:84e7:b37d:114::3, length 32

(I think I failed to capture the ND advertisement here due to my tcpdump filtering config)

20:31:36.650332 IP6 fdbf:84e7:b37d:114:a10:a10:a10:f15f.59169 > fdbf:84e7:b37d:114::3.53: 46769+ AAAA? talk.google.com. (33)

20:31:36.653917 IP6 fdbf:84e7:b37d:114::3.53 > fdbf:84e7:b37d:114:a10:a10:a10:f15f.59169: 46769 2/0/0 CNAME talk.l.google.com., AAAA 2404:6800:4008:c01::7d (82)

Time is wrong due to timezone issue

The DHCPv6/ULA was used as the source for (and only for) ND and DNS ?

Questionable. Why doesn't it use the SLAAC/ULA ?

```
20:31:37.841252 IP6 fdbf:84e7:b37d:114:a10:a10:a10:f15f.23712 > fdbf:84e7:b37d:114::3.53: 28011+
A? mtalk.google.com. (34)
20:31:37.841322 IP6 fdbf:84e7:b37d:114:a10:a10:a10:f15f.4213 > fdbf:84e7:b37d:114::3.53: 52684+
AAAA? mtalk.google.com. (34)
20:31:37.841505 IP6 fdbf:84e7:b37d:114:c9c4:6553:c915:f183.64838 > 2404:6800:4008:c01::7d.5222:
Flags [F.], seq 1, ack 1, win 16384, length 0
20:31:37.842024 IP6 fdbf:84e7:b37d:114:a10:a10:a10:f15f.46727 > fdbf:84e7:b37d:114::3.53: 6954+ A?
www.google.com. (32)
20:31:37.842185 IP6 fdbf:84e7:b37d:114:a10:a10:a10:f15f.29999 > fdbf:84e7:b37d:114::3.53: 61308+
AAAA? www.google.com. (32)
20:31:37.842277 IP6 fdbf:84e7:b37d:114:c9c4:6553:c915:f183.64839 > 2404:6800:4008:c01::7d.5222:
Flags [F.], seq 909, ack 567, win 16384, length 0
20:31:37.843552 IP6 fdbf:84e7:b37d:114::3.53 > fdbf:84e7:b37d:114:a10:a10:a10:f15f.23712: 28011
2/0/0 CNAME mobile-gtalk.l.google.com., A 74.125.23.188 (79)
20:31:37.843554 IP6 fdbf:84e7:b37d:114::3.53 > fdbf:84e7:b37d:114:a10:a10:a10:f15f.4213: 52684
2/0/0 CNAME mobile-gtalk.l.google.com., AAAA 2404:6800:4008:c01::bc (91)
20:31:37.843555 IP6 fdbf:84e7:b37d:114::3.53 > fdbf:84e7:b37d:114:a10:a10:a10:f15f.46727: 6954
5/0/0 A 173.194.117.209, A 173.194.117.211, A 173.194.117.208, A 173.194.117.212, A
173.194.117.210 (112)
20:31:37.843555 IP6 fdbf:84e7:b37d:114::3.53 > fdbf:84e7:b37d:114:a10:a10:a10:f15f.29999: 61308
1/0/0 AAAA 2404:6800:4004:80c::1010 (60)
```

Time is wrong due to timezone issue

Other discussions

- Androids did not bring the NICs up without an IPv4 address assigned.
- IPv4 only Apps: Skype didn't work in the ULA-only (without IPv4). Dropbox wouldn't neither.

Summary

- Adding ULA to GUA+IPv4(RFC1918) seemed harmless
- Coexistence of SLAAC/ULA, DHCPv6/ULA, Stateless prefix NAT66, and stateful DNS64/NAT64 in a segment is possible.
- Tested mainly by the Mac laptop.
- Not tested:
 - Mixture of nodes having different address types
 - Incoming connections (for servers)
 - SIP
 - And others.