Dealing with reality

Alain Durand

April 22nd, 2009

comcast

IPv4 reality check: completion of allocation is real

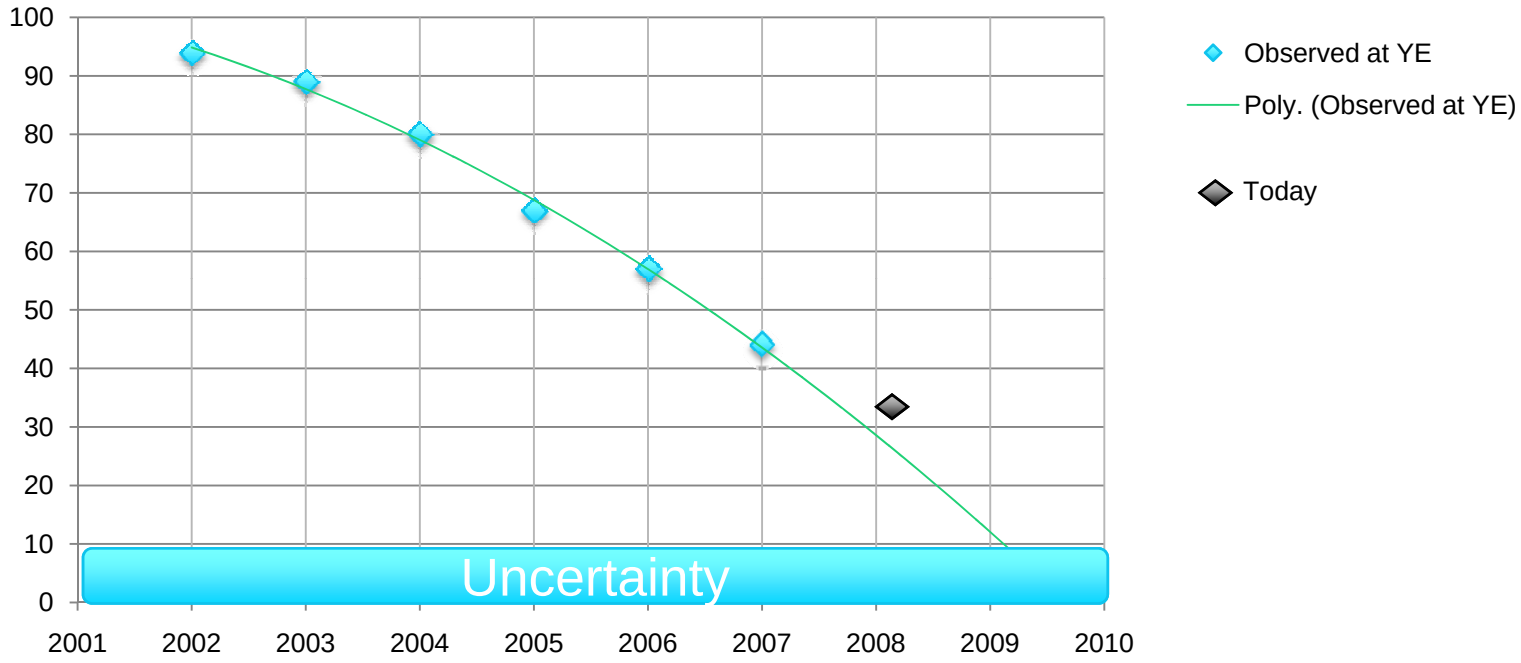

After completion:

Existing IPv4 addresses will not stop working.
Current networks will still operate.

comcast

IPv6 reality check: the IPv4 long tail

- Post IPv4 allocation completion:
  - Many hosts in the home (eg Win 95/98/2000/XP, Playstations, consumer electronic devices) are IPv4-only.
    - They will not function in an IPv6-only environment.
    - Few of those hosts can and will upgrade to IPv6.
  - Content servers (web, email,…) hosted on the Internet by many different parties will take time to upgrade to support IPv6.

comcast

Dealing with both realities: a two prong approach

① **Embrace IPv6**

- Move as many devices/services to IPv6 as possible to lower dependency on IPv4 addresses

② **Build an IPv6 transition bridge for the IPv4 long tail**

- Goal:
  - Provide IPv4 service without providing a dedicated IPv4 address
- Technology:
  - Leverage IPv6 access infrastructure
  - Provide only IPv6 addresses to endpoint
  - Share IPv4 addresses in the access networks
  - DS-lite: IPv4/IPv6 tunnel + provider NAT

comcast

DS-lite update

Draft-ietf-softwire-dual-stack-lite-00.txt

comcast

IPv4 port distribution

- Measurements:
  - Average #ports/customer < 10 (per transport protocol)
  - Peak #ports/customer > 100? > 1000? > 5000?

- Do not dimension for peaks, but for average!
  - No cookie cutter approach
  - Large dynamic pool of ports shared by many customers

- Customers want to choose their own applications
  - CGN MUST not interfere with applications, eg avoid ALGs,…
  - Need to support incoming connections
  - Small static pool of reserved ports under the control of customers

comcast

Port forwarding & A+P extensions

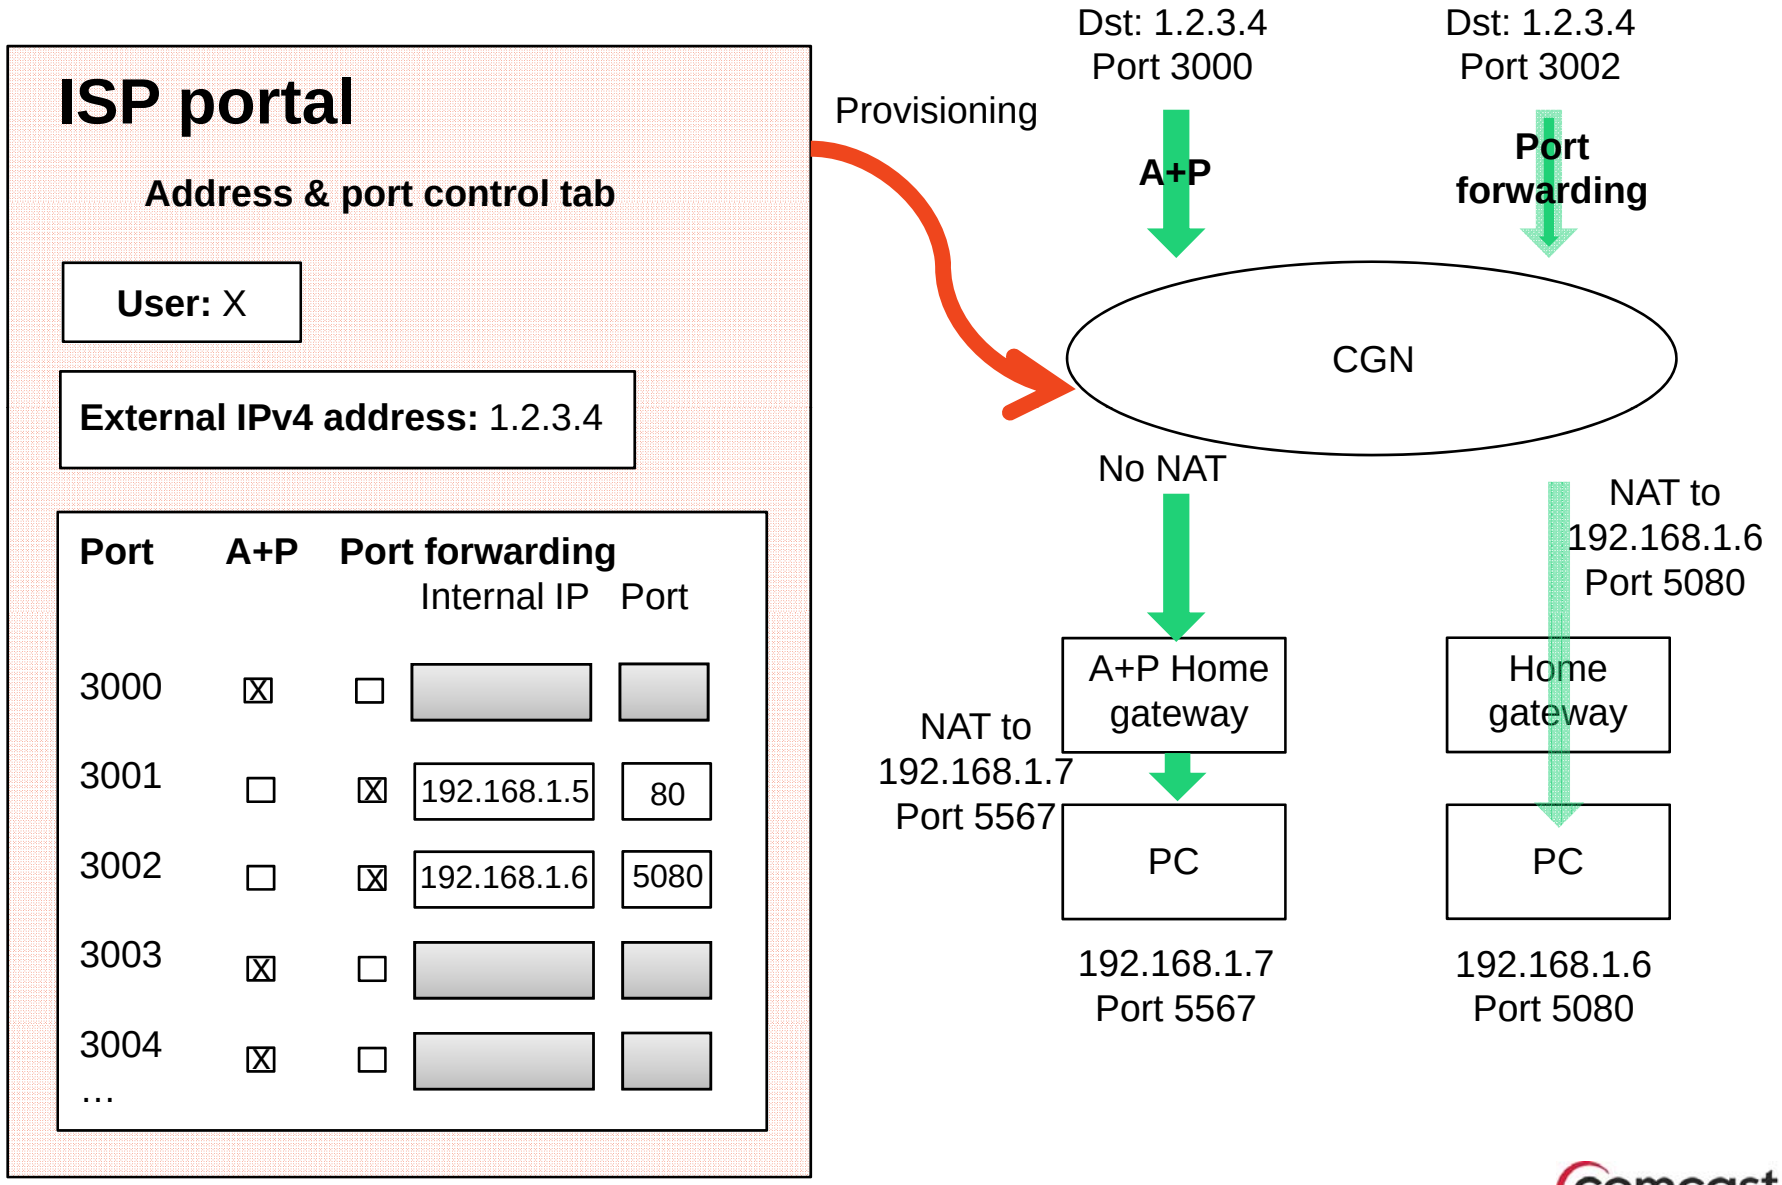

comcast

UPnP

- Typical UPnP application will:
  - Decide to run on port X
  - Ask IGD to forward port X traffic
  - If IGD declines, try again with X+1
    - After 10 or so attempts, abort

- This will NOT work with any IPv4 address sharing mechanism (NAT444, DS-lite, NAT64, IVI, A+P,…)

- NAT-PMP has a better semantic: IGD can redirect the application to use an alternate available port

- UPnP forum is reported to be addressing this issue

Security issues relative to CGN

- Port number information is necessary for full identification
  - Need to log port numbers on the receiving side
  - Need to log NAT bindings on CGN
- CGN needs to enforce per customer limits either on new connection rate or maximum number of sessions
- User authentication on service provider CGN may not be necessary, users get authenticated at the IPv6 access layer. A simple ACL on the CGN to limit access to the service provider customers seems to be sufficient. 3rd party CGNs may have different requirements.
- HGW & CGN need to enforce that customer IPv4 addresses inside of IPv6 tunnel are indeed RFC1918 addresses

Other security issues

- The Internet community needs to deal with Web sites that put IPv4 address in penalty box after a number of unsuccessful login attempts.

- More generally, the community need to revisit notion that an IPv4 address uniquely identifies a customer.

Scaling issue

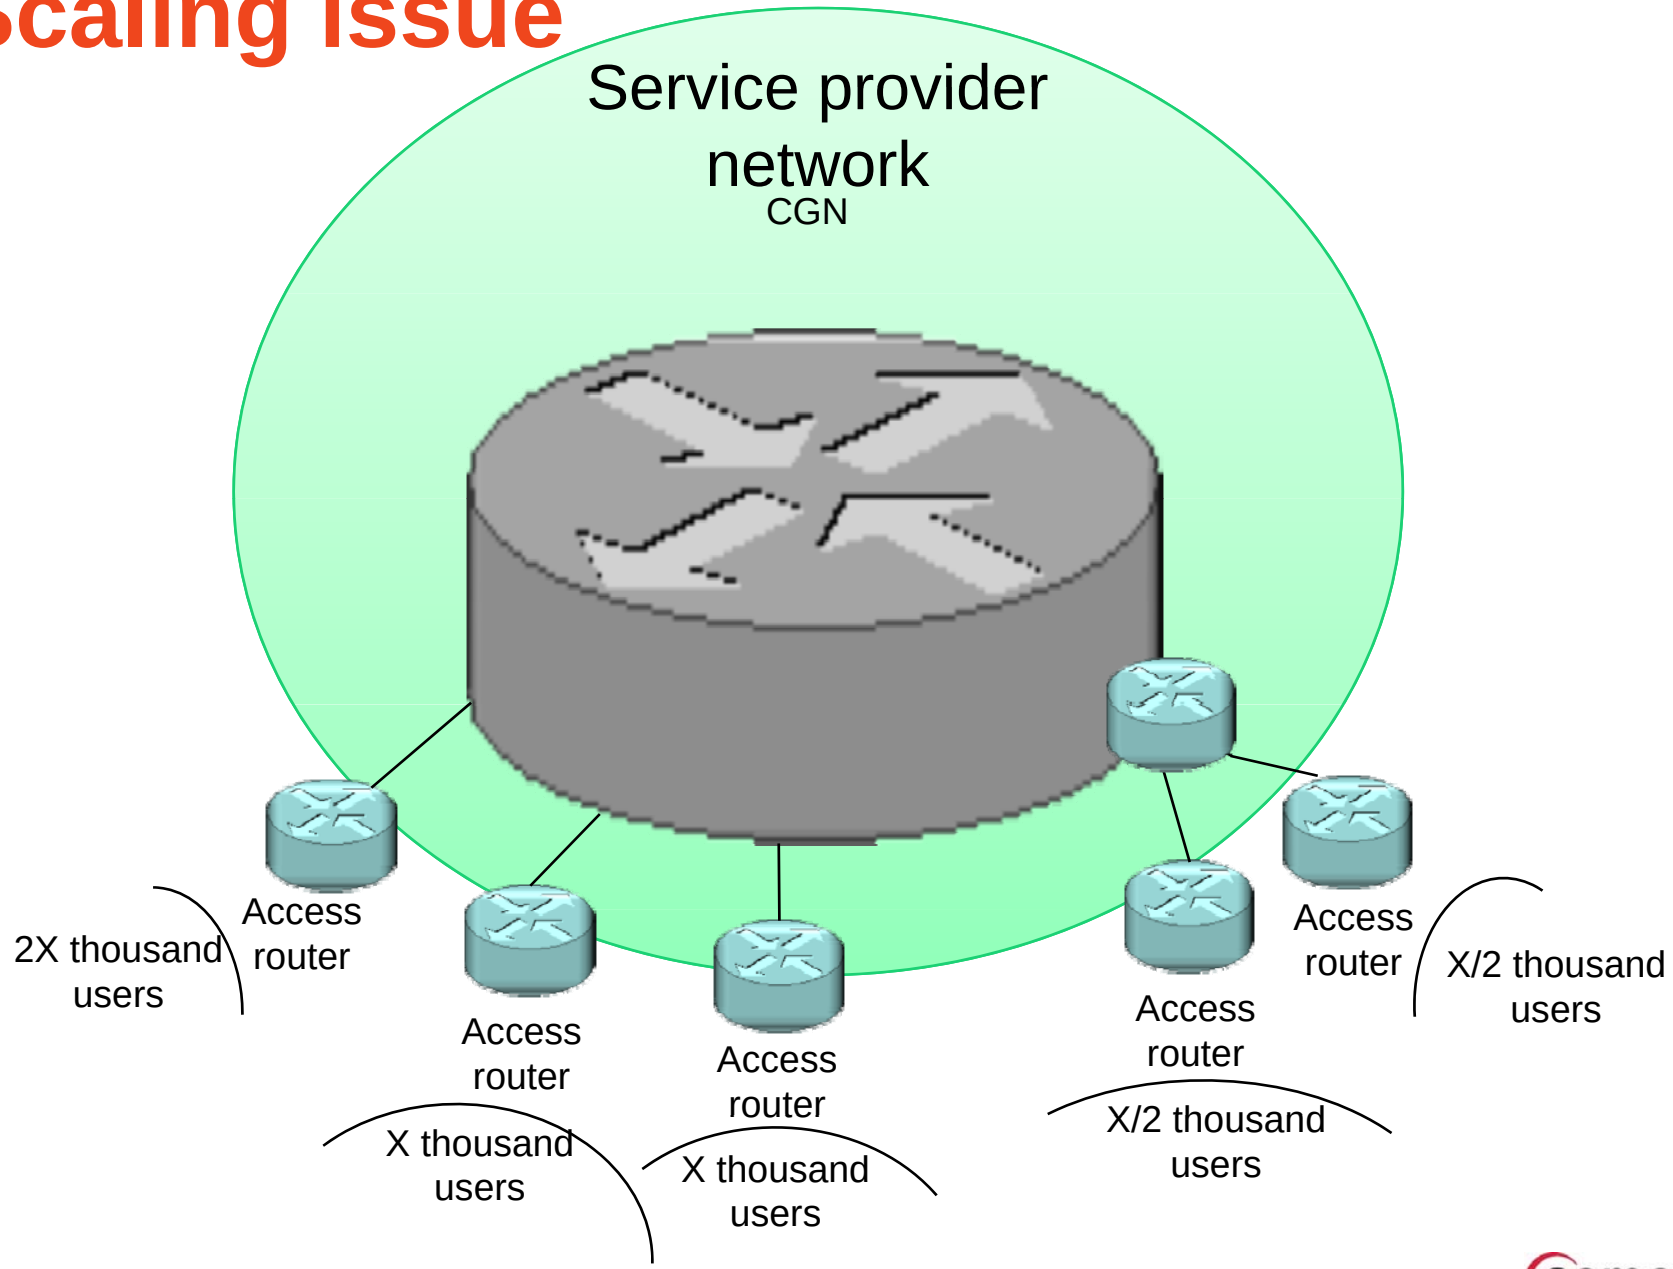

comcast

Horizontal scaling

- DHCPv6 option to configure tunnel end-point
- Enable sending the traffic to as many CGNs as necessary

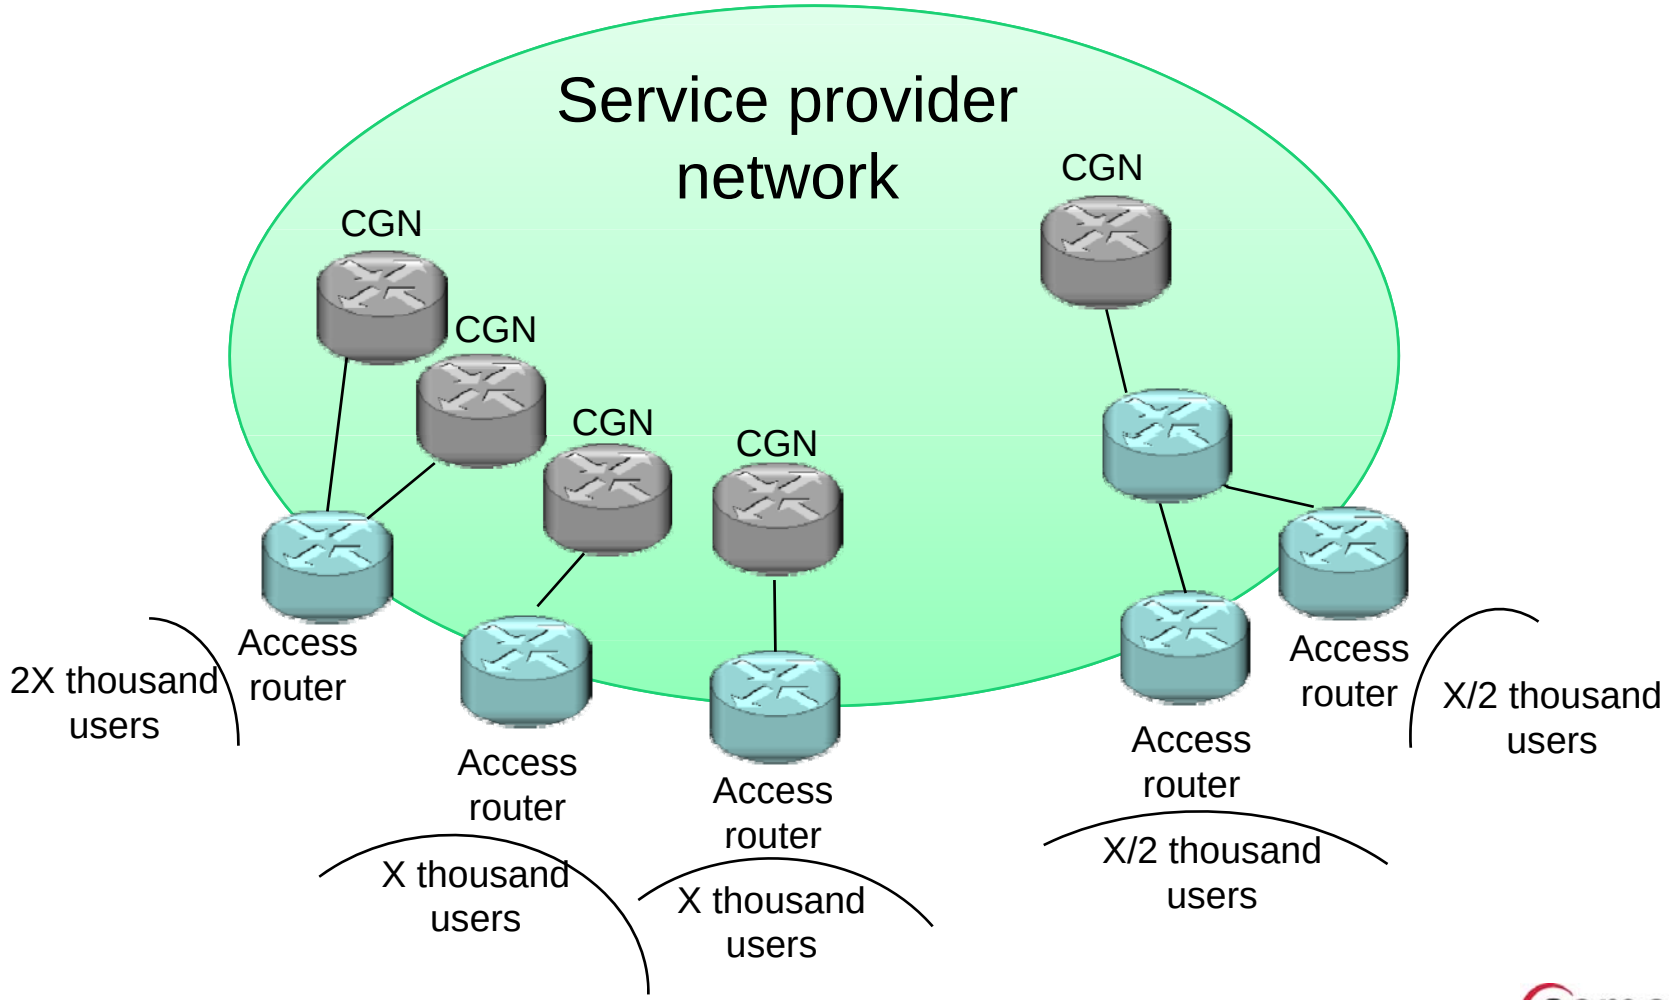

comcast

DS-lite demo

Thanks to:
Yiu Lee, Carl Williams, Anthony Veiga
ISC
ARIN

comcast

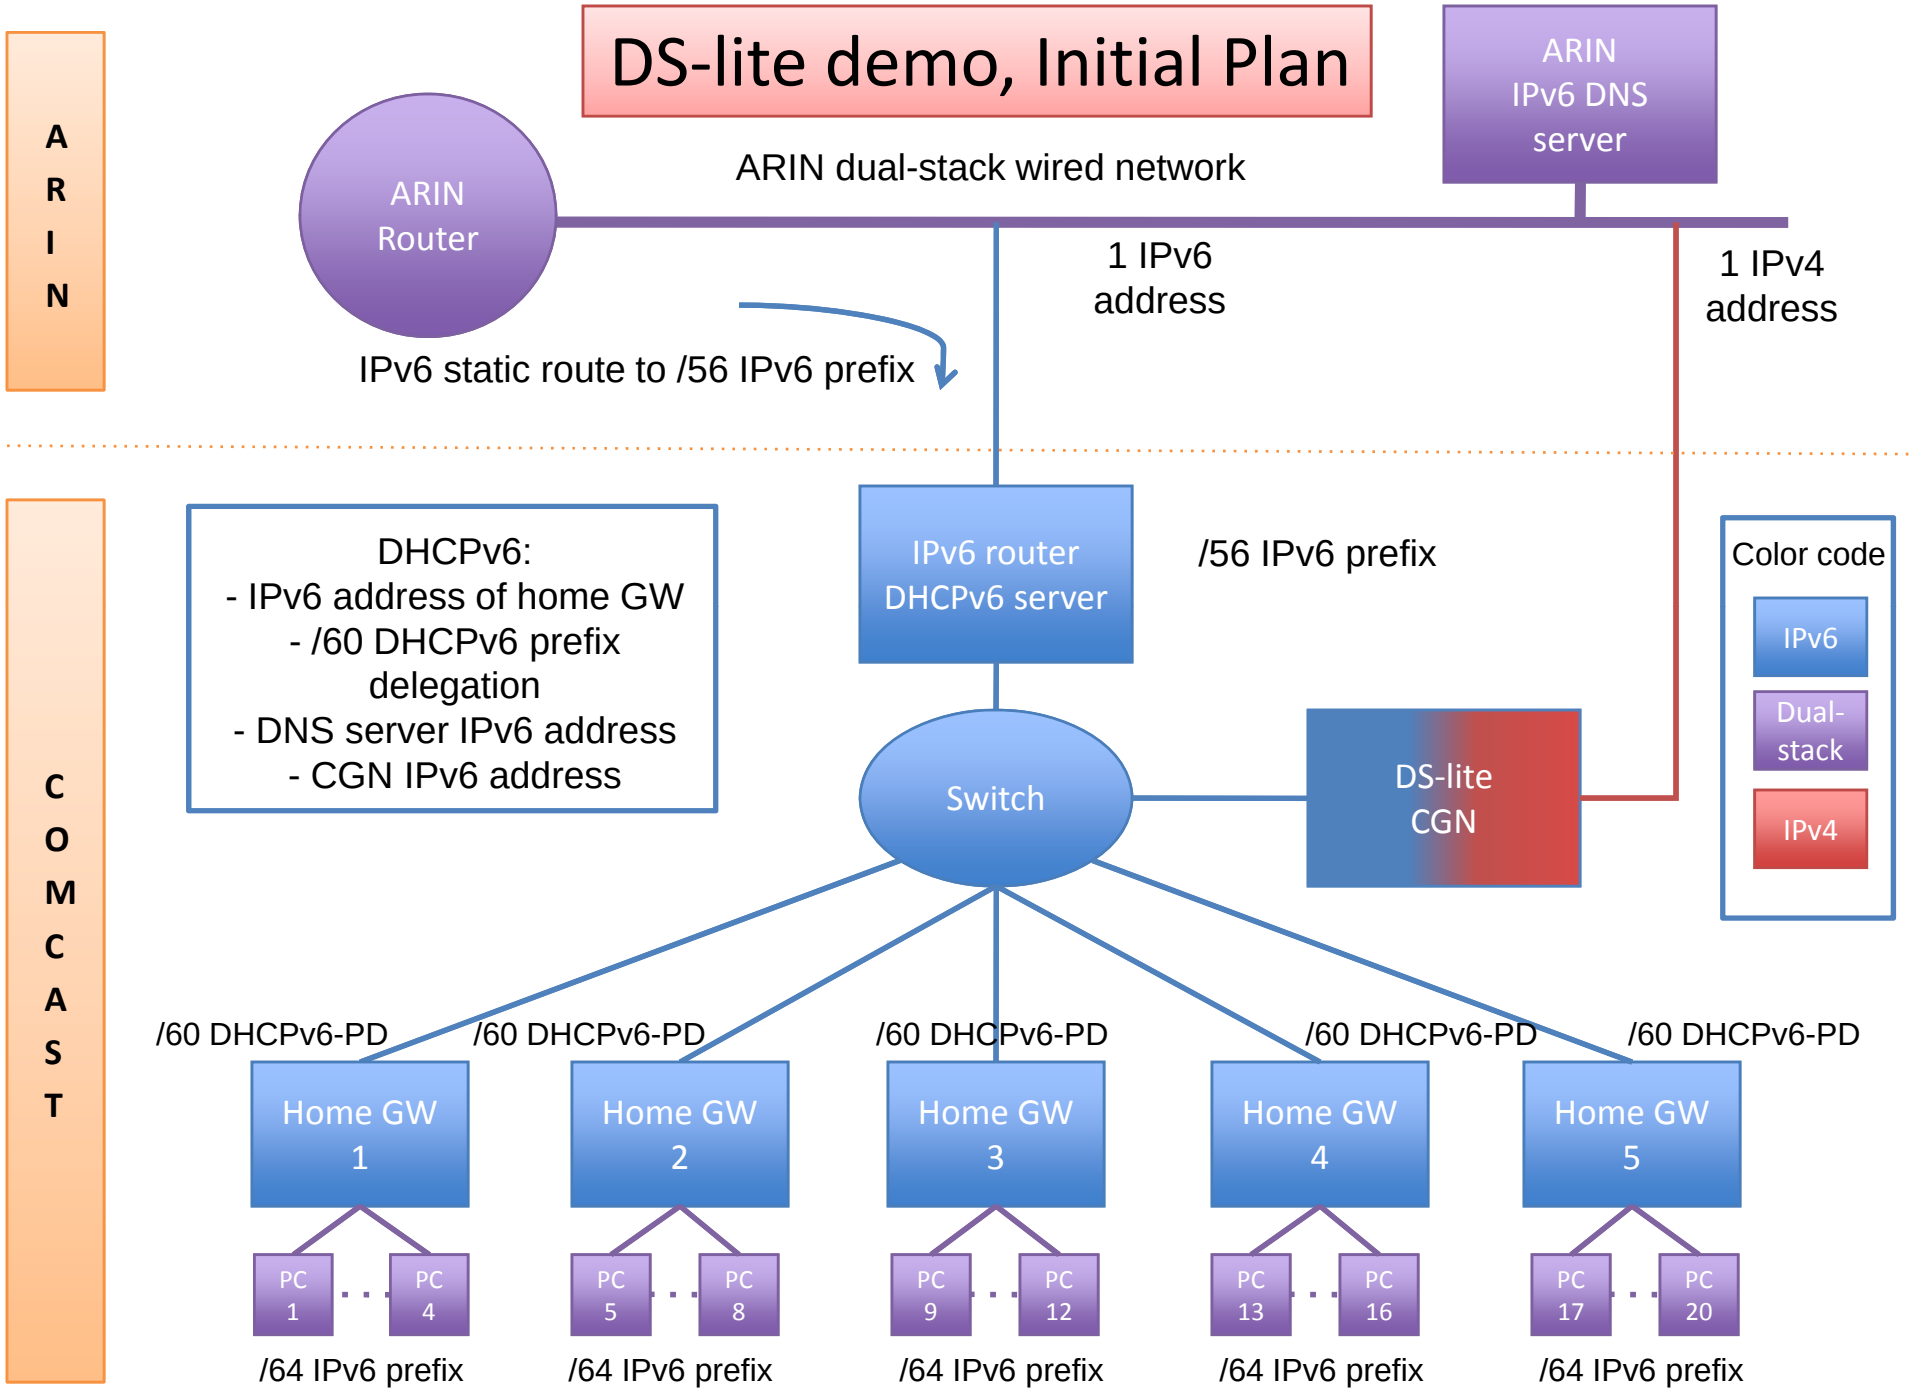
DS-lite demo, Initial Plan
ARIN
COMCAST
ARIN Router
ARIN dual-stack wired network
ARIN IPv6 DNS server
1 IPv6 address
1 IPv4 address
IPv6 static route to /56 IPv6 prefix
DHCPv6:
- IPv6 address of home GW
- /60 DHCPv6 prefix delegation
- DNS server IPv6 address
- CGN IPv6 address
IPv6 router DHCPv6 server
/56 IPv6 prefix
Switch
DS-lite CGN
Color code
IPv6
Dual-stack
IPv4
/60 DHCPv6-PD
/60 DHCPv6-PD
/60 DHCPv6-PD
/60 DHCPv6-PD
/60 DHCPv6-PD
Home GW 1
Home GW 2
Home GW 3
Home GW 4
Home GW 5
PC 1
PC 4
PC 5
PC 8
PC 9
PC 12
PC 13
PC 16
PC 17
PC 20
/64 IPv6 prefix
/64 IPv6 prefix
/64 IPv6 prefix
/64 IPv6 prefix
/64 IPv6 prefix

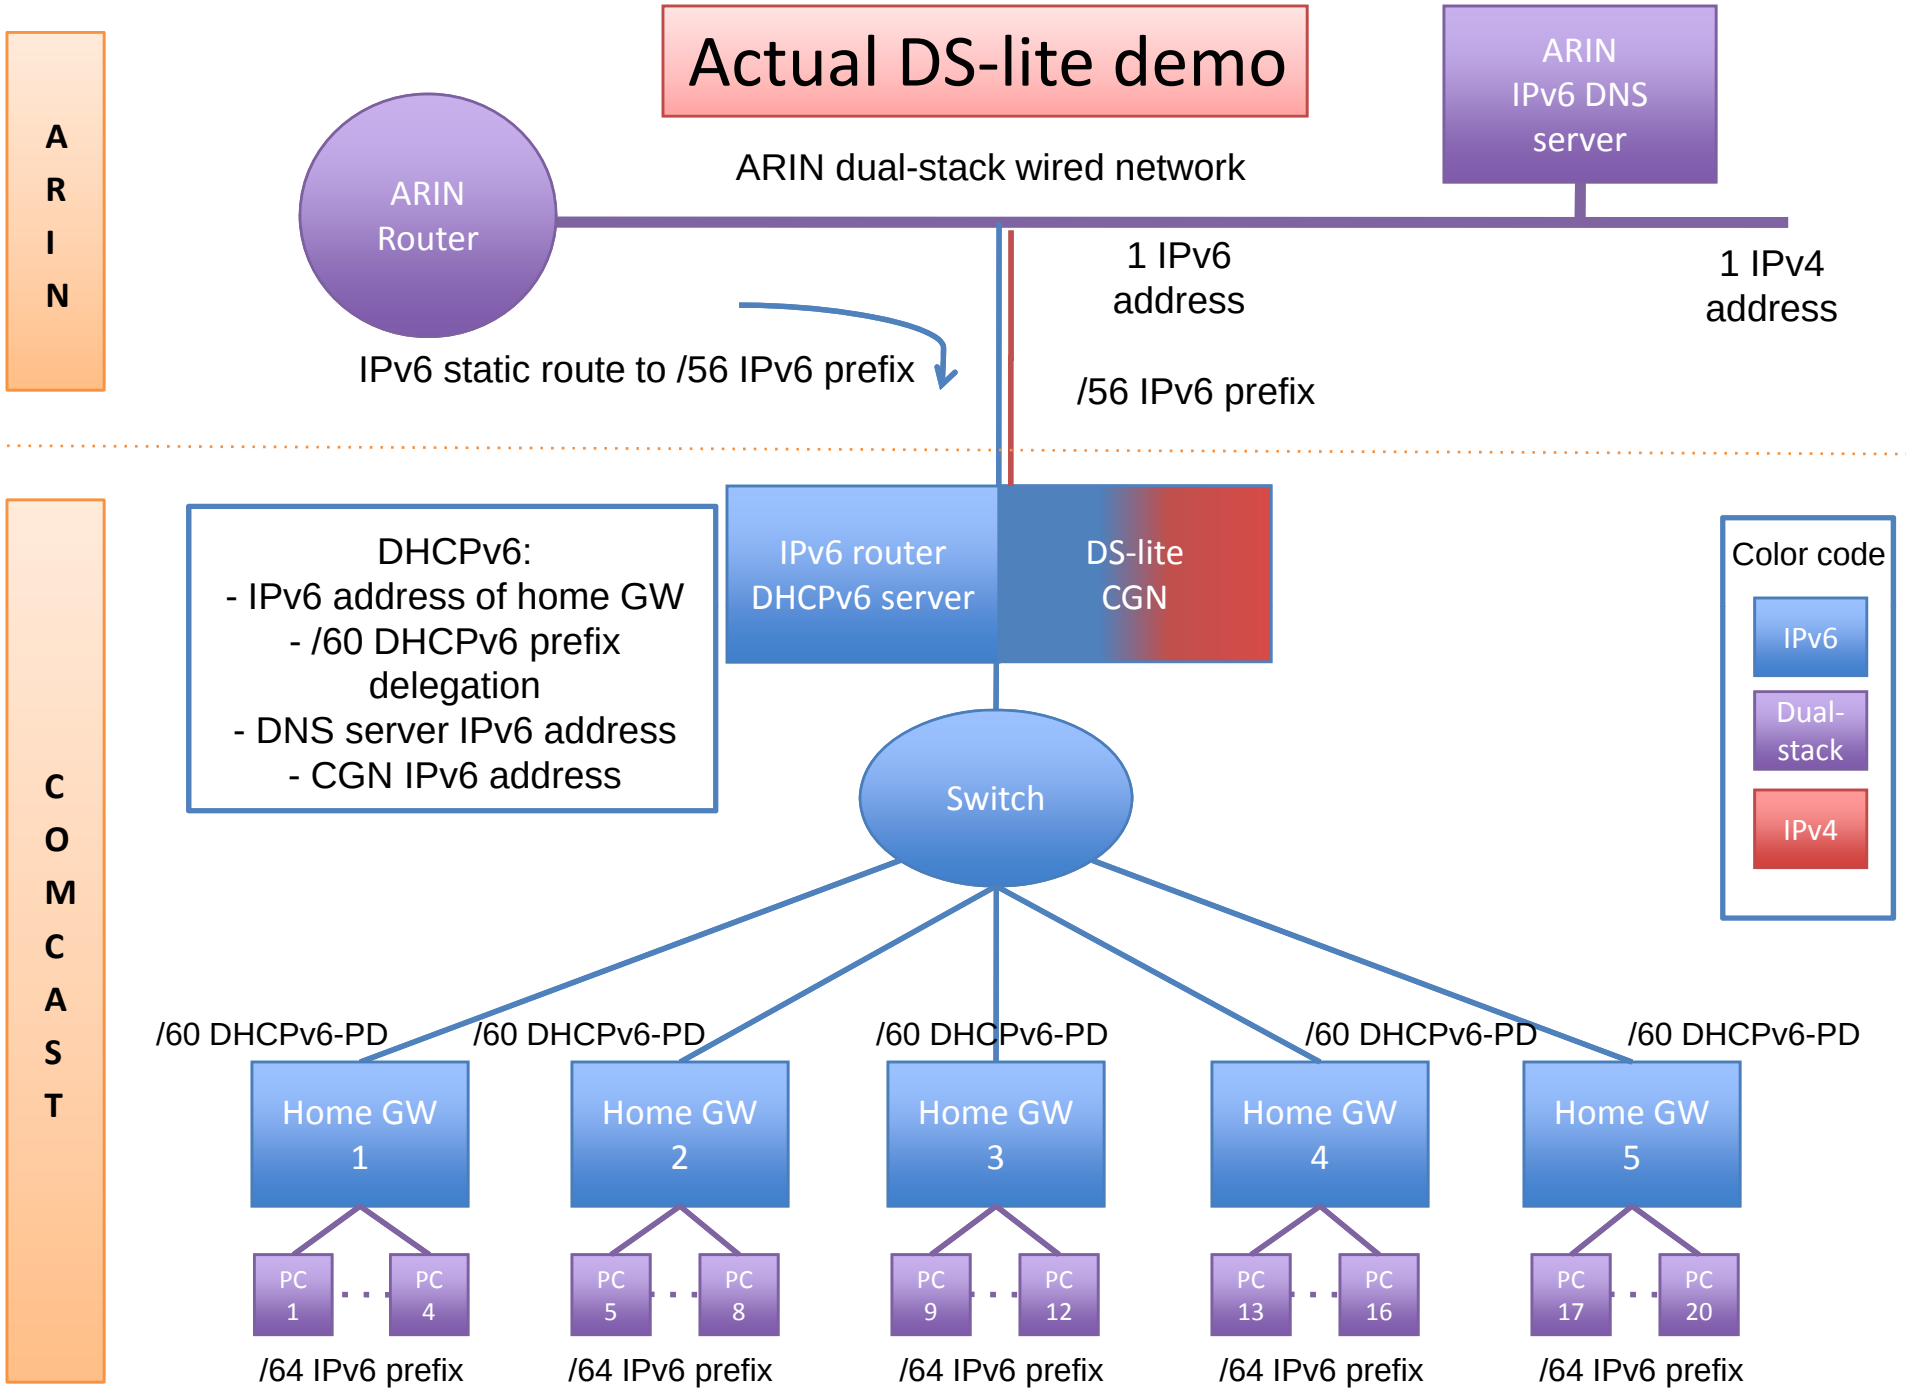
Actual DS-lite demo
ARIN
COMCAST
ARIN Router
ARIN dual-stack wired network
ARIN IPv6 DNS server
1 IPv6 address
1 IPv4 address
IPv6 static route to /56 IPv6 prefix
/56 IPv6 prefix
DHCPv6:
- IPv6 address of home GW
- /60 DHCPv6 prefix delegation
- DNS server IPv6 address
- CGN IPv6 address
IPv6 router DHCPv6 server
DS-lite CGN
Color code
IPv6
Dual-stack
IPv4
Switch
/60 DHCPv6-PD
/60 DHCPv6-PD
/60 DHCPv6-PD
/60 DHCPv6-PD
/60 DHCPv6-PD
Home GW 1
Home GW 2
Home GW 3
Home GW 4
Home GW 5
PC 1
PC 4
PC 5
PC 8
PC 9
PC 12
PC 13
PC 16
PC 17
PC 20
/64 IPv6 prefix
/64 IPv6 prefix
/64 IPv6 prefix
/64 IPv6 prefix
/64 IPv6 prefix