Carrier Grade NAT - Observations and Recommendations

Chris Grundemann | North American IPv6 Summit | 11 April 2012

CableLabs®

...Revolutionizing Cable Technology®

Agenda

- CGN Technology
- CGN Challenges
- CGN Architectures
- Conclusions

Starting with the Basics

CGN TECHNOLOGY

Address sharing needed when IPv6 is not available

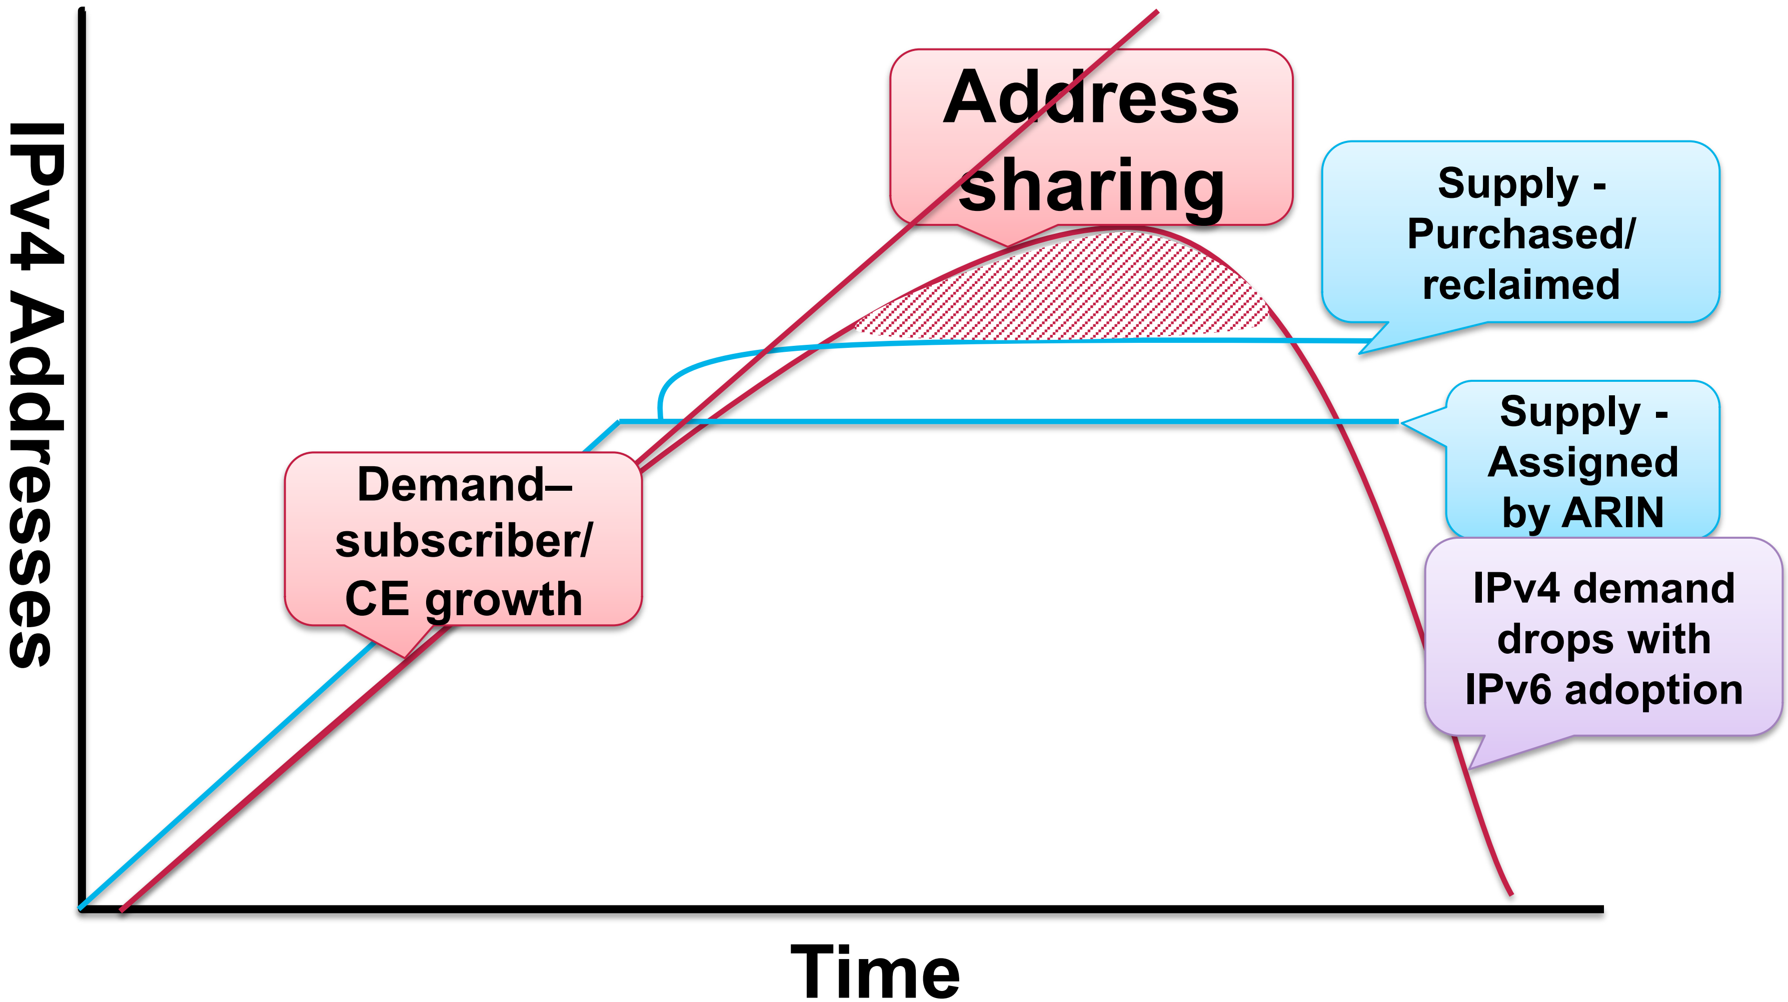

Network Address Translation (NAT)

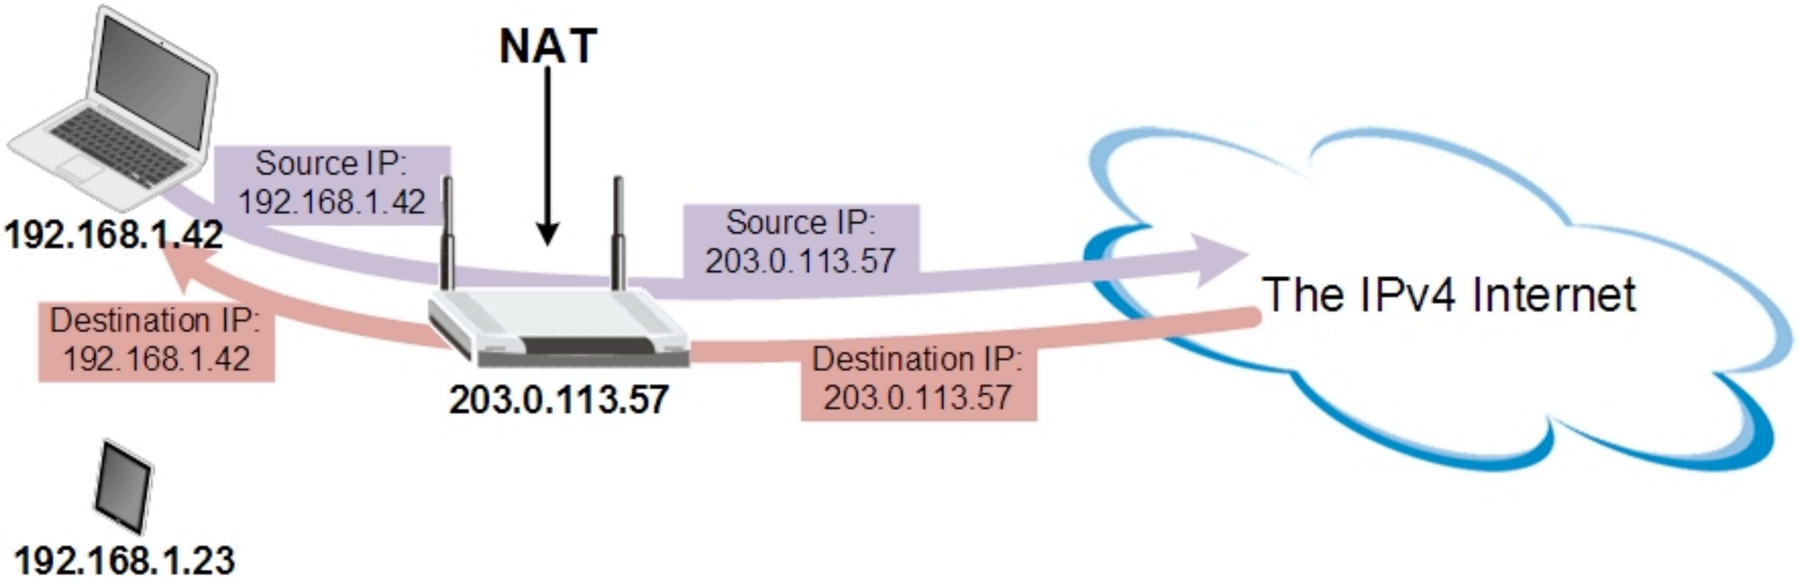

PAT and Address Overloading

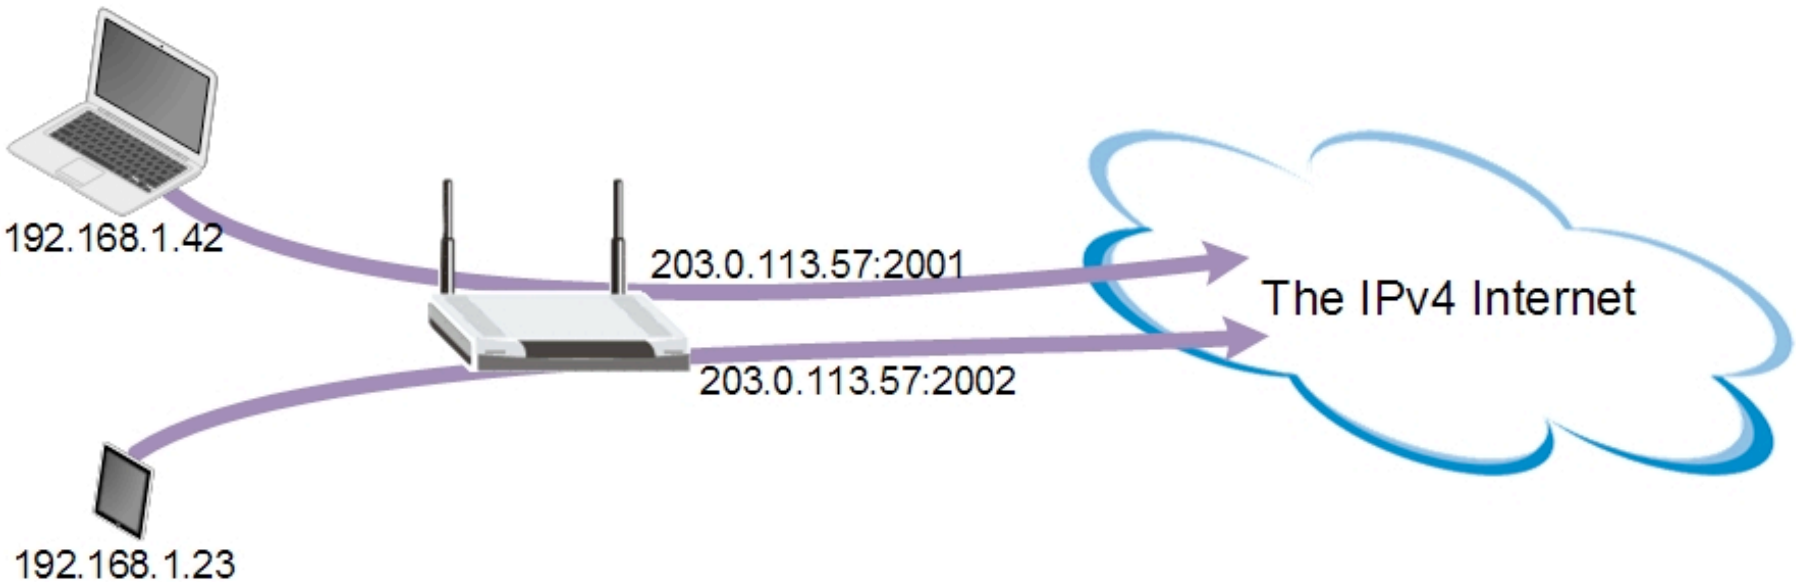

NAT and the End to End Principle

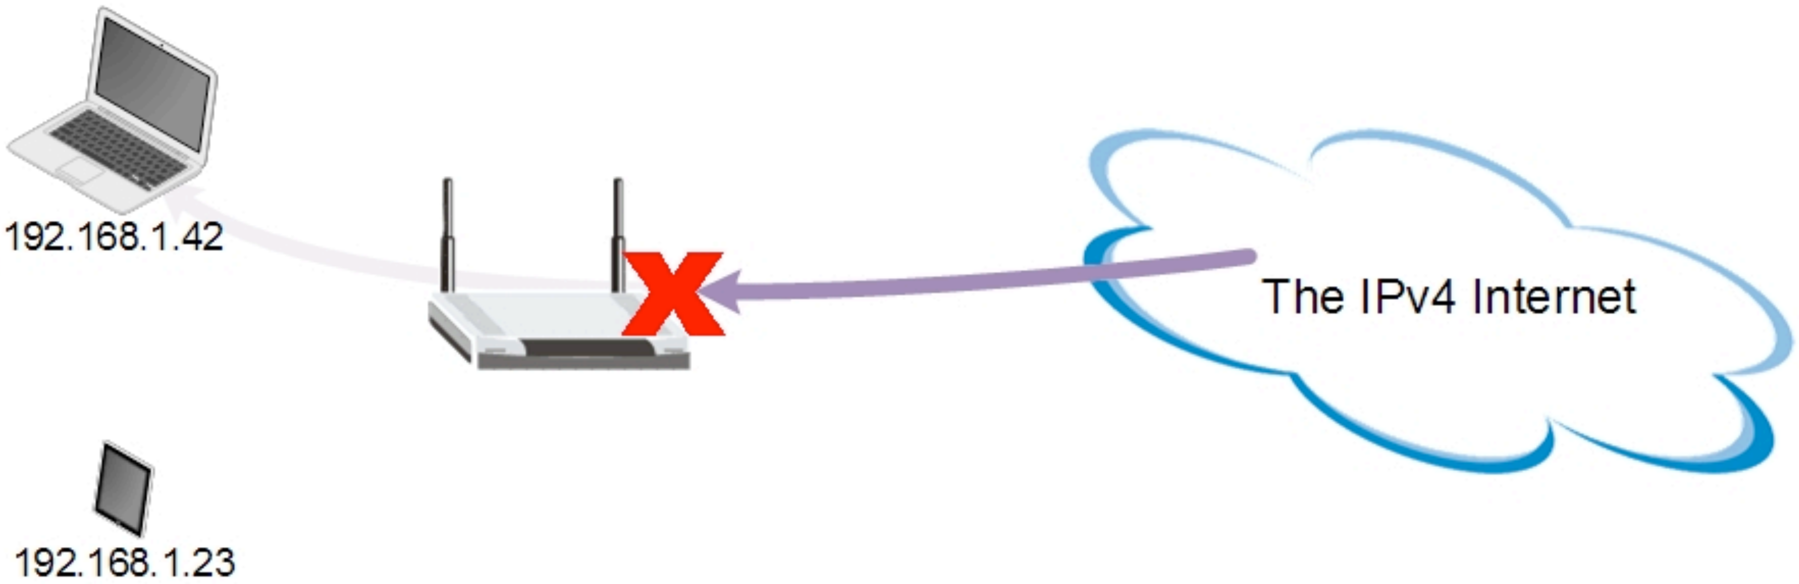

NAT444

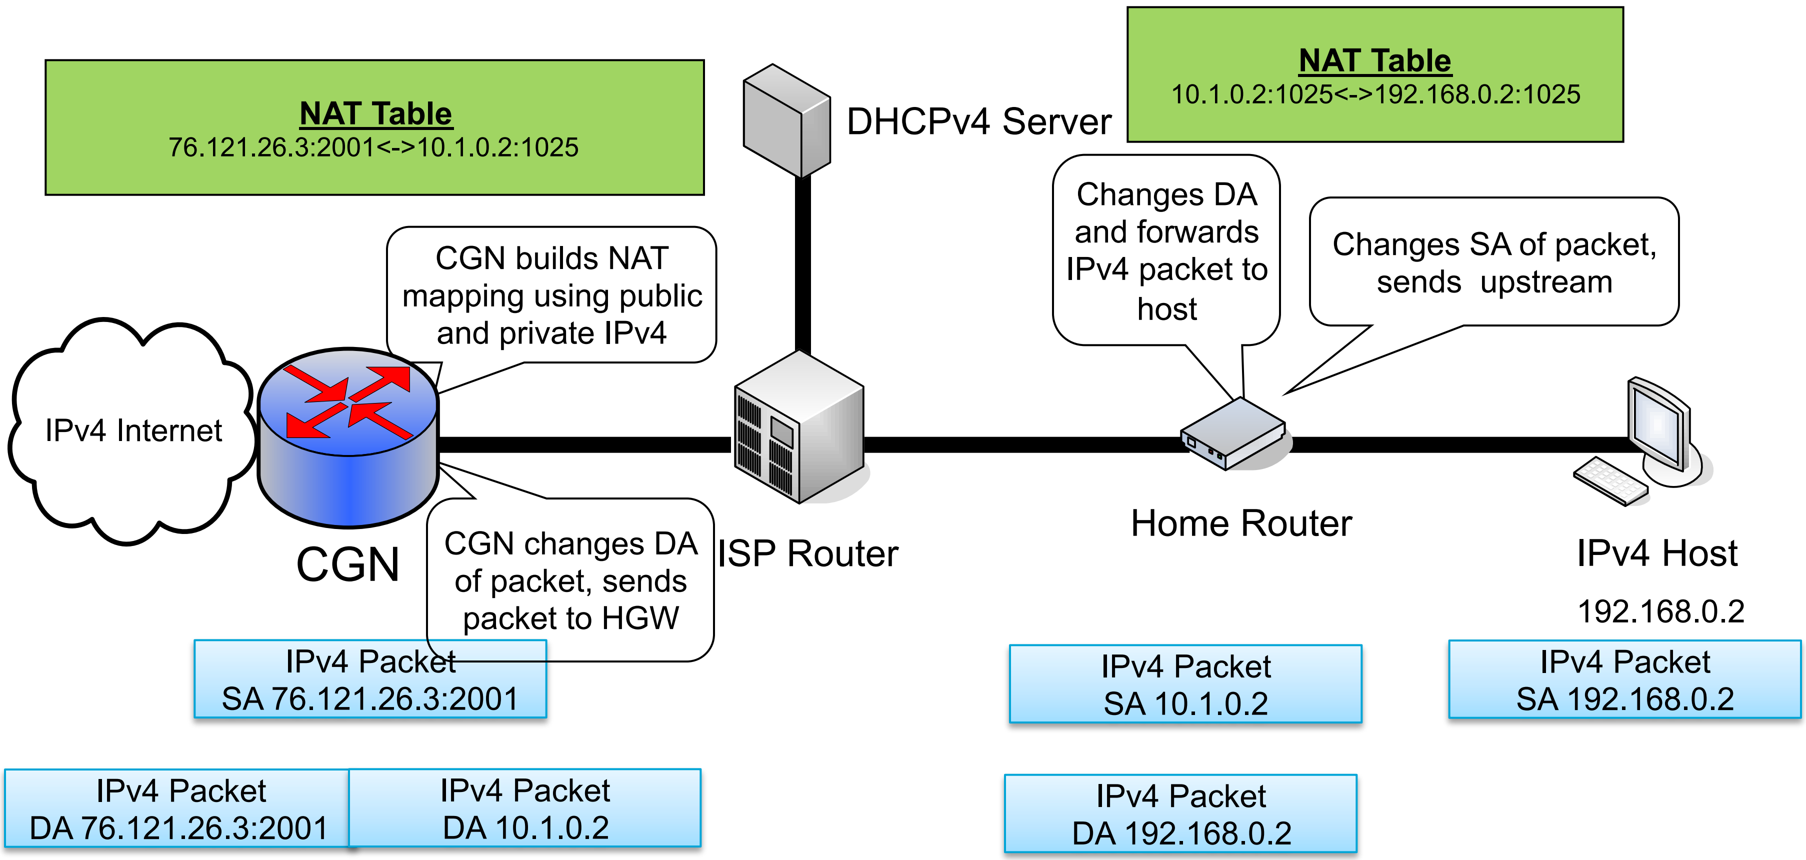
NAT Table
76.121.26.3:2001<->10.1.0.2:1025
DHCPv4 Server
NAT Table
10.1.0.2:1025<->192.168.0.2:1025
Changes DA and forwards IPv4 packet to host
Changes SA of packet, sends upstream
CGN builds NAT mapping using public and private IPv4
IPv4 Internet
CGN
CGN changes DA of packet, sends packet to HGW
ISP Router
Home Router
IPv4 Host
192.168.0.2
IPv4 Packet
SA 76.121.26.3:2001
IPv4 Packet
SA 10.1.0.2
IPv4 Packet
SA 192.168.0.2
IPv4 Packet
DA 76.121.26.3:2001
IPv4 Packet
DA 10.1.0.2
IPv4 Packet
DA 192.168.0.2

Dual-Stack Lite

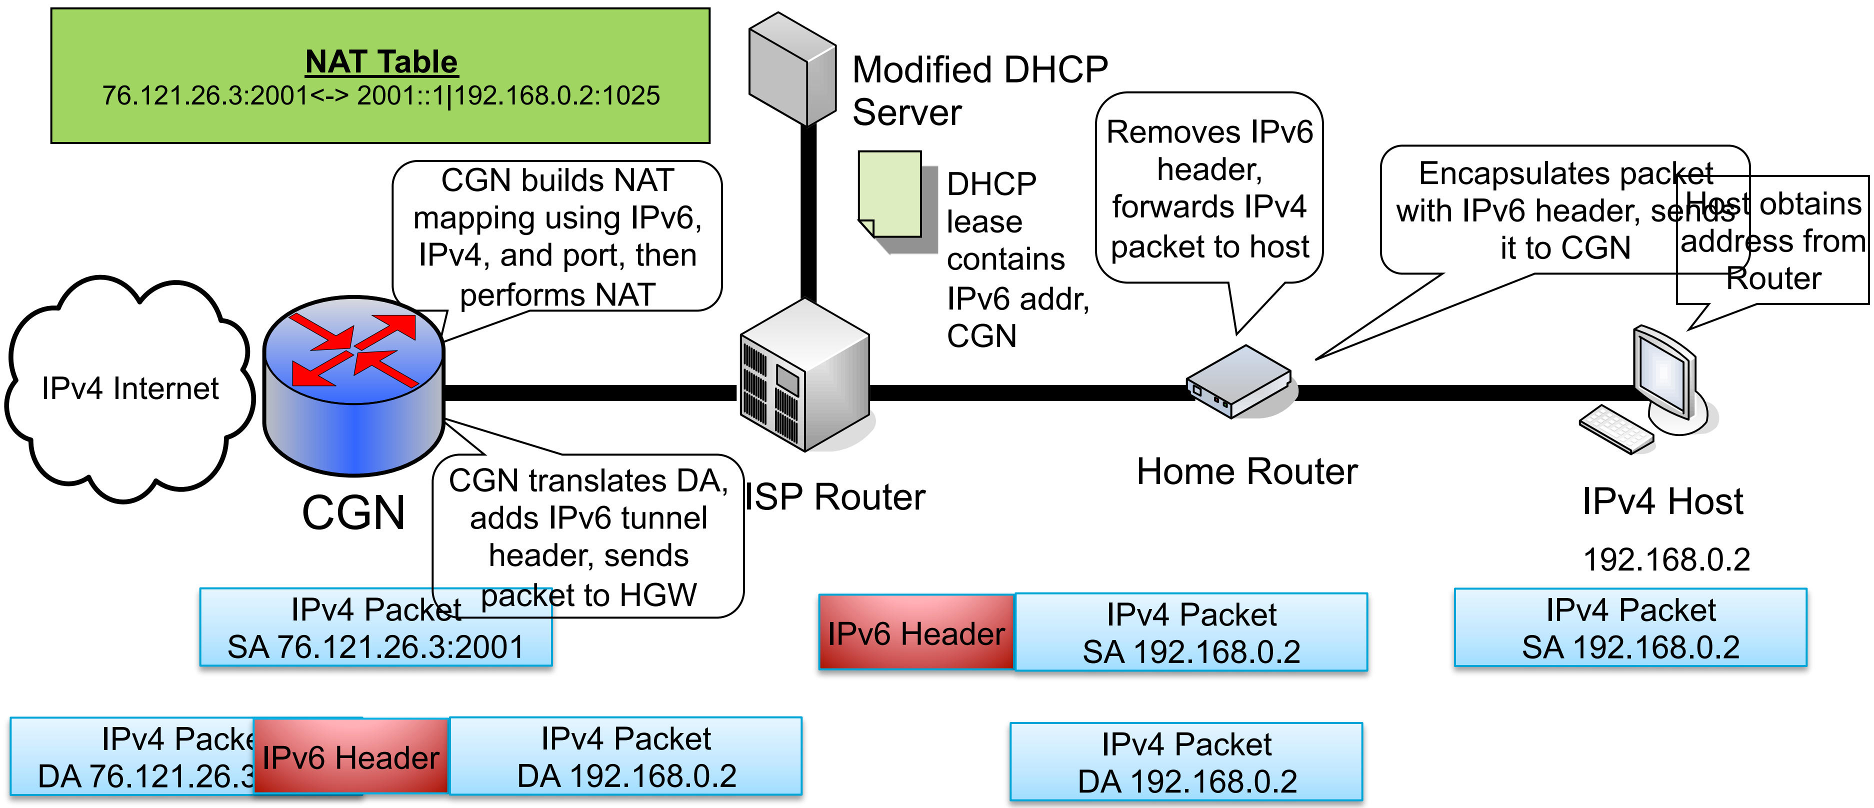

   4/5/12

CableLabs
...Revolutionizing Cable Technology

Typical Access Technology Transition Timeline

Connectivity Type	Time →			
IPv4	Native	NAT444	NAT444	DS-Lite
IPv6	None	6RD	Native	Native

The Evil in Necessary Evil

CGN CHALLENGES

CGN Testing Background

- CableLabs first conducted CGN testing in 2010
  - NAT444 only
- Second round June – Sep, 2011
  - Both NAT444 and DS-Lite
- Additional CGN testing in IPv6 interop events
  - About one a quarter

Overview of test scenarios

- Single and dual ISP networks with one or more users on multiple home networks
- Test applications include
  - Video services – e.g. Netflix, YouTube, iClips, Silverlight
  - Audio streaming – e.g. Pandora, Internet Archive
  - Peer-to-peer – e.g. on line gaming, uTorrent
  - FTP – large file transfers
  - SIP calls – e.g. X-Lite, Skype
  - Video chat – e.g. Skype, OoVoo
  - Social networking – e.g. Facebook, Webkinz
  - Web conferencing – e.g. GoToMeeting

Client devices and gateways used for testing

- Laptops running Vista, Win 7 and MAC OS
- Gaming consoles
- Tablet devices
- iPhone and Android smartphones
- CE devices
  - Blu Ray players, Smart TVs
- CPE routers
  - Most vendors represented

Observations

- The following types of applications behaved erratically or had the potential to break:
  - Video streaming, e.g. Netflix, YouTube
  - Peer-to-peer, e.g. uTorrent, Bittorent, Limewire
  - On line gaming, e.g. X-box
  - FTP file transfer
- Performance dependent on home gateway
  - Different NAT types (full cone, partial cone) perform differently
- Observed behaviors were exacerbated when multiple users or multiple home networks were involved
- User experience further degraded when crossing ISPs and when “hairpinning” through the same CGN

Log volumes

150 - 450 bytes/connection

* 33k - 216k connections per sub per day

----

5 - 96 MB / user / day

That's potentially over 1 PB per 1M subs per month

It's also over 20Mbps for just the log stream…

CGN Challenges

- Poor quality of experience for advanced services
  - Peer-to-peer, video streaming, gaming, etc.

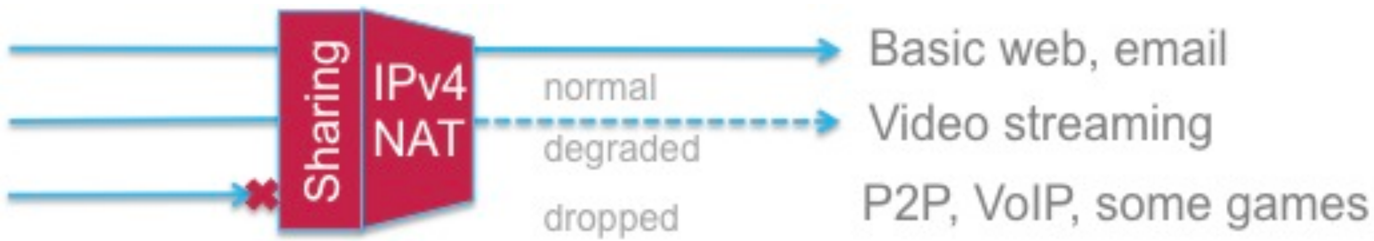

- Negative impact to targeted advertising/geo-location
- Logging requirements for lawful intercept
  - Petabytes of data

CableLabs®
...Revolutionizing Cable Technology®

Workarounds

Challenges	Workarounds
P2P SIP (cannot initiate/ receive calls) uTorrent (seeding does not work)	Use Proxies for Peer to Peer applications Port Control Protocol
P2P Gaming	Software Upgrade from Manufacturer Port Control Protocol
Degraded experience for services such as Netflix, video streaming	Deploy tested home-routers from an approved list
Slower Download rates (some clients)	No known workarounds (Try larger MTU)
Negative impact to targeted advertising/geo-location	Distributed CGN, Regional IP and Port assignments
Logging requirements for lawful intercept	Deterministic NAT, Data compression, Bulk port assignment
Overlapping Addressing / NAT Zones	Large enough shared transition space
Impacts to traffic engineering	Distributed CGN, VRF (MPLS/VPN)

Port Control Protocol (PCP)

- PCP is an IETF protocol
  - Expected to be an RFC soon
  - Allows an IPv6 or IPv4 host/router to control how incoming IPv6 or IPv4 packets are translated and forwarded by a network address translator (NAT) or simple firewall
  - PCP can solve a number problems identified with CGN
- Challenges
  - Requires CPE Router and CGN support
  - Requires that trust boundary be extended to subscriber for port assignment

Summary

- Significant improvement year over year
  - CGN improvements
  - Content provider updates (X-Box live, Netflix)
  - Application updates (X-Lite, uTorrent)
- CGN experience not as good as un-NATed IPv4
  - Degradations in P2P, streaming applications
- DS-Lite and NAT444 perform similarly
  - Additional impacts to hairpinned DS-Lite connections
- Troubleshooting issues will be difficult

- More: **draft-donley-nat444-impacts**

Looking for Answers

CGN ARCHITECTURES

21 Cable Television Laboratories, Inc. 2011. All Rights Reserved. Proprietary/Confidential. 4/5/12

CableLabs®
...Revolutionizing Cable Technology®

Architectural Constraints

- Relative deployment cost (day 1 cost)
  - Ease of implementation
- Impact on routing: Changes required in current routing infrastructure
  - Traffic Engineering: Allows MSO to distribute/route traffic
  - Load Balancing: Sharing load between different devices
- Scalability: Response to increased traffic/subscriber growth
- Subscriber IP addressing
  - Size of Private Subnet needed
  - Number of Public Addresses used
- Geo-location: Granularity of geolocation information obtained
- On-net server deployment: Ease of placement of various servers

22 Cable Television Laboratories, Inc. 2011. All Rights Reserved. Proprietary/Confidential. 4/5/12

Architecture – Centralized

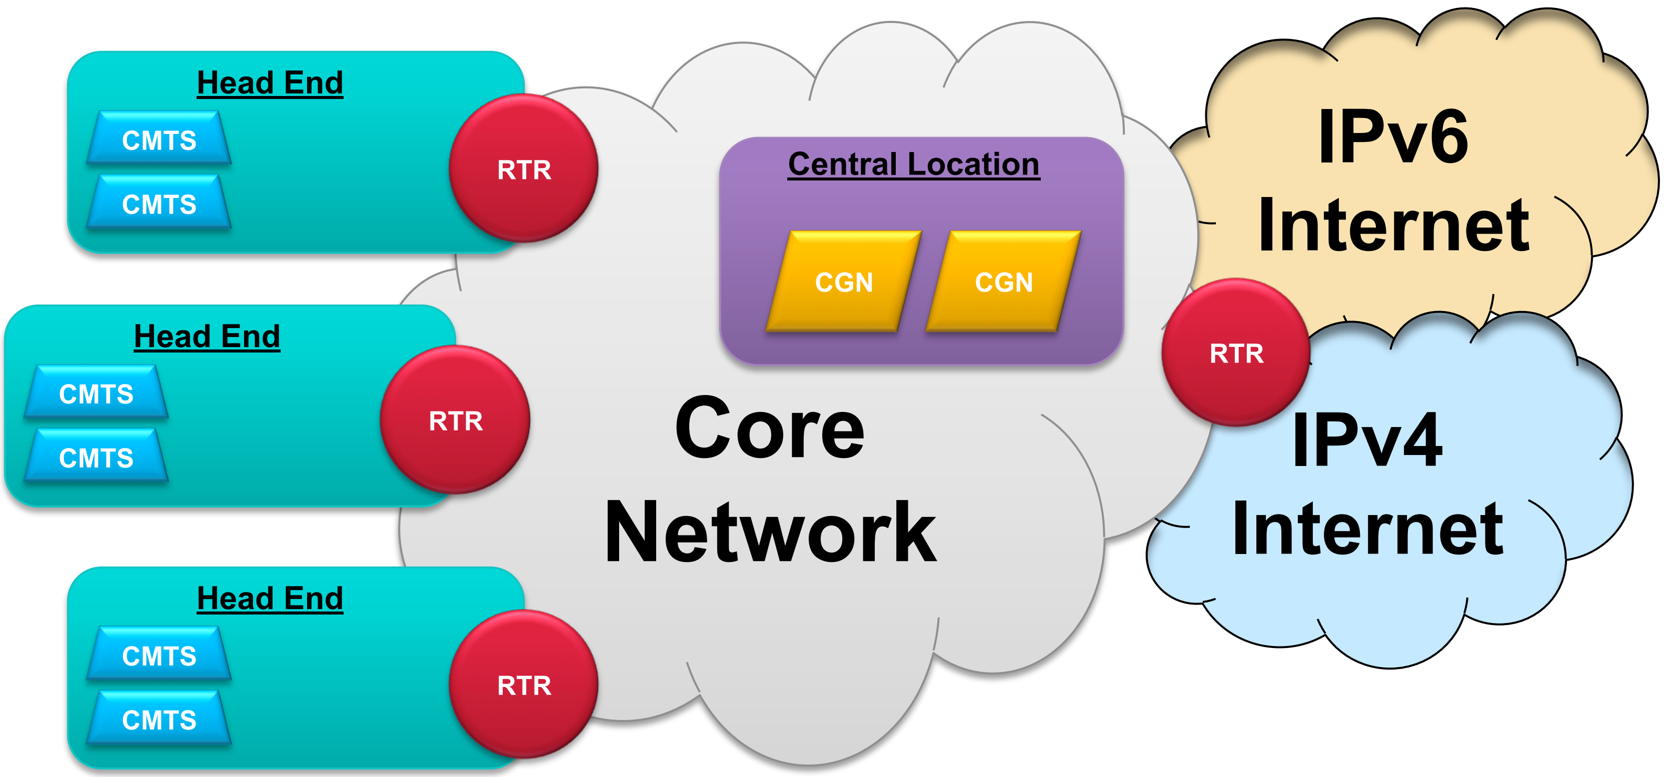

23 Cable Television Laboratories, Inc. 2011. All Rights Reserved. Proprietary/Confidential. 4/5/12

Architecture - Distributed

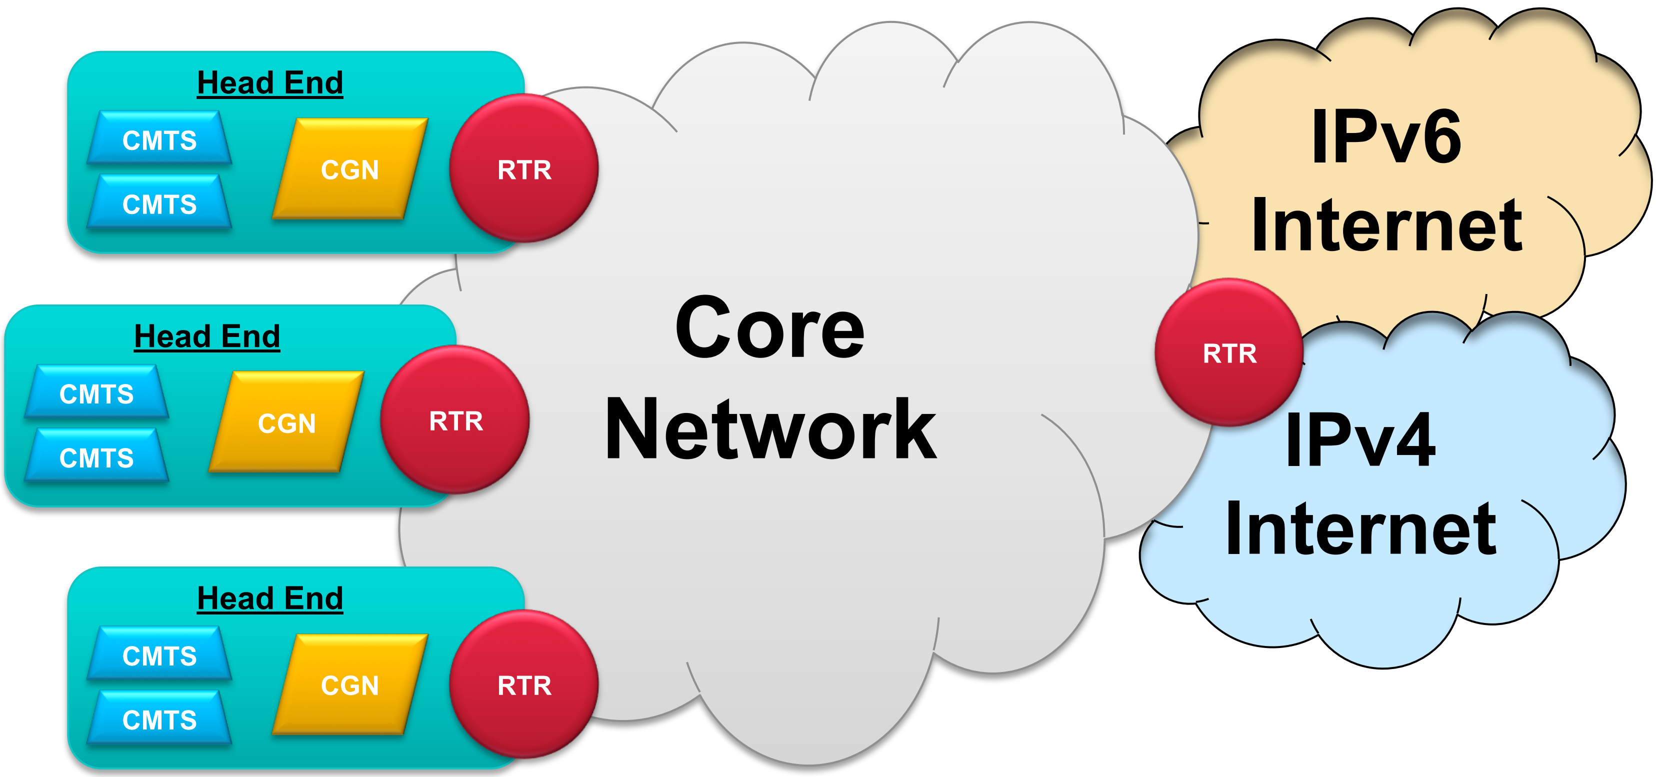

24 Cable Television Laboratories, Inc. 2011. All Rights Reserved. Proprietary/Confidential. 4/5/12

Architecture – Hybrid (Phased approach)

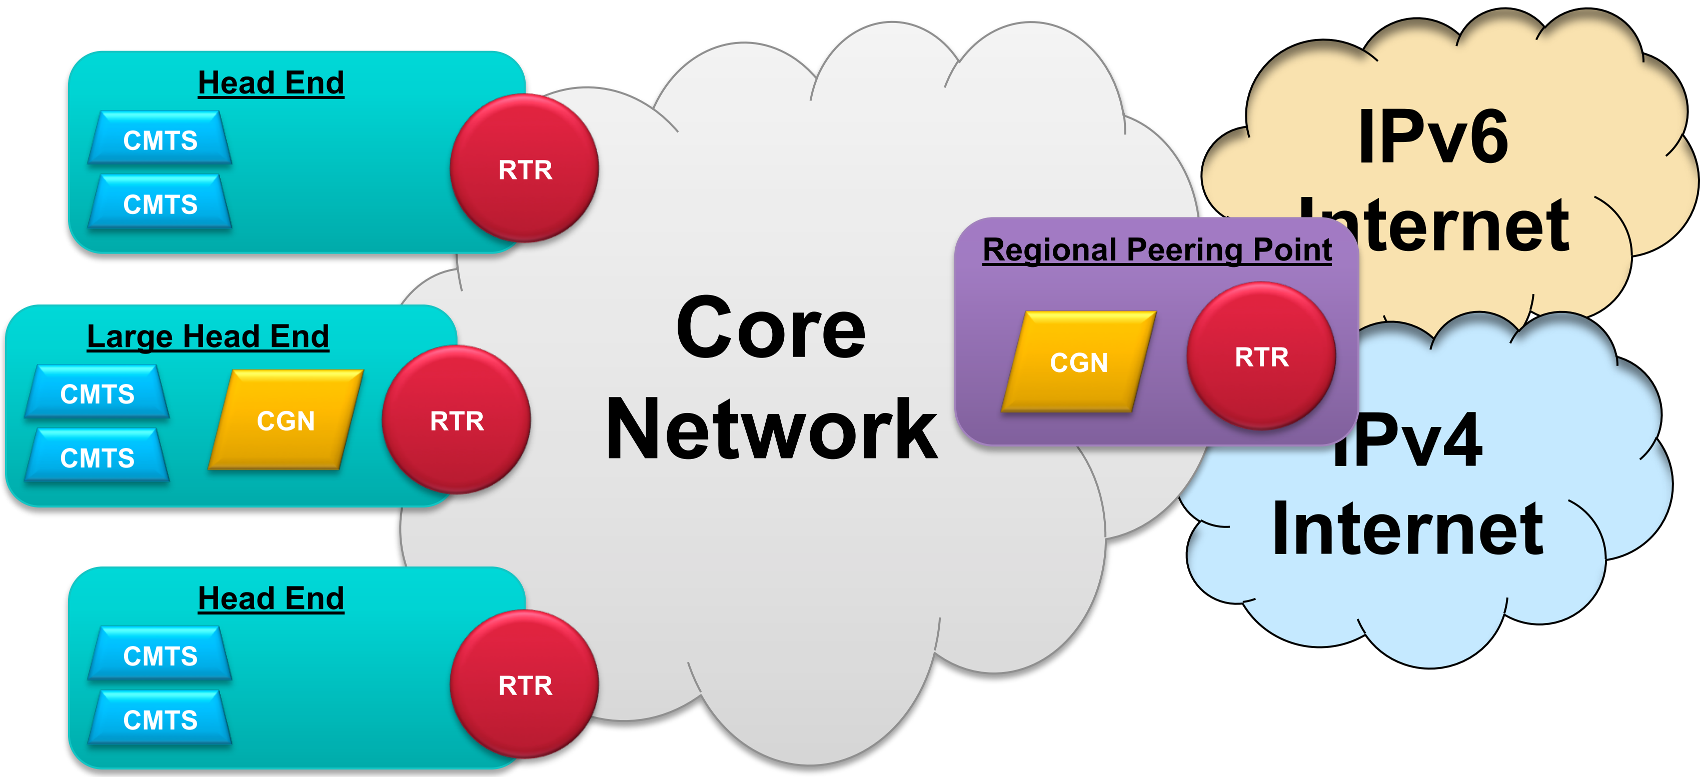

25 Cable Television Laboratories, Inc. 2011. All Rights Reserved. Proprietary/Confidential. 4/5/12

Recommendation

- A phased hybrid approach is recommended
  - Start with Regionalized CGNs
  - Add CGNs as needed locally as the CGN user base grows
- Rationale
  - Offers ISPs easy starting point and wide reach
  - Low impact to routing and traffic engineering
  - Offers the most flexible scalability over time

Further Considerations

- Subscriber differentiation
- Routing CGN Traffic
- Redundancy
- Load balancing & Scalability
- Server location & NAT bypass
- IP Addressing
- Geo location
- Logging
- Security
- Address Reputation

NAT Bypass and Server Location

- Goal: Optimizing local traffic and subscriber access to advanced services
- Server Location (in a NAT444 environment)
  - Any internal (e.g. voice, video) or 3rd party (e.g. CDN) application servers that are placed inside the CGN should offer better performance
  - This is less important for basic services such as web and email
- NAT444 CGN Bypass
  - Don’t send traffic through the CGN unnecessarily
  - Use native dynamic routing to reach servers inside the CGN
  - Add servers to CGN VPN, if in use
- DS-Lite: Enable IPv6 on all servers (all IPv4 goes through CGN)

IP Addressing: Public "outside" addressing

- Number of addresses required determined by number of CGN subscribers and compression algorithm
  - Start low; ~8x
- Where to get addresses?
  - Re-purposing existing addresses
    - Renumber infrastructure to IPv6 or private IPv4
    - Renumber customers to inside CGN addresses
  - Acquire new addresses – transfer market?
  - Reserve addresses now
  - Does not need to be contiguous space
- Port restrictions
  - Should not be an issue at low compression ratios

IP Addressing: Inside Addressing

- NAT444: Use a single network-wide pool of inside addresses
  - 100.64.0.0/10 Shared Transition Space
  - Assign local (per site) blocks out of larger pool for operational clarity, logging, the ability to insert local CGNs, and potential geo-location benefits
- DS-Lite: Any addresses are acceptable and can be reused per tunnel

Geolocation

- Local (per-site) CGNs will offer roughly equal granularity to what is available today
- Regional CGNs will dilute geo-location data
- One idea to minimize this dilution is to use separate outside pools of addresses which correspond to the per-site private subnets
  - These public pools should be loose, to borrow from the next pool if needed
    - Either borrow from an adjacent pool, or higher level pool

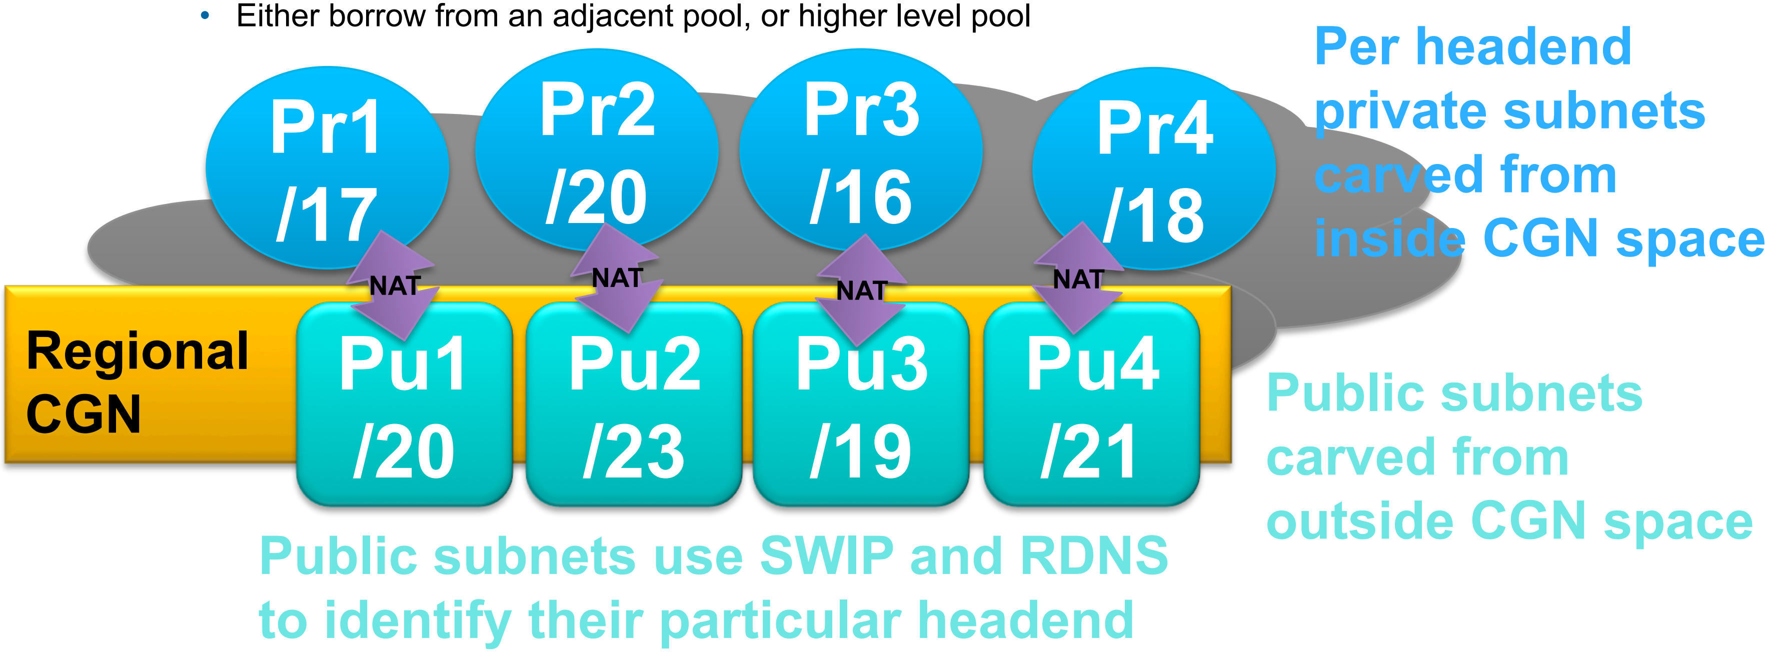

Log Reduction Strategies

- Port block reservations
  - Reduce logging up to 100x
- Log compression
  - Reduces volume, but not search time
- Deterministic reservation
  - See next slide…

Proposal: Deterministic Port Reservation

- [draft-donley-behave-deterministic-cgn]
- Collect inside range, outside range, compression ratio
  - Compression ratio ≥ inside/outside
  - Inside range/compression ratio = ports/user
  - Set aside well-known ports (<1024) & dynamic overflow range
  - Pre-reserve port ranges for each internal IP address
  - Allow dynamic reservation above that threshold
    - Remote logging only required for dynamic reservations
    - Still need state logging locally for every active connection
- Limitations:
  - Requires low compression ratios
  - Requires configuration change control process

Deterministic NAT Illustrated

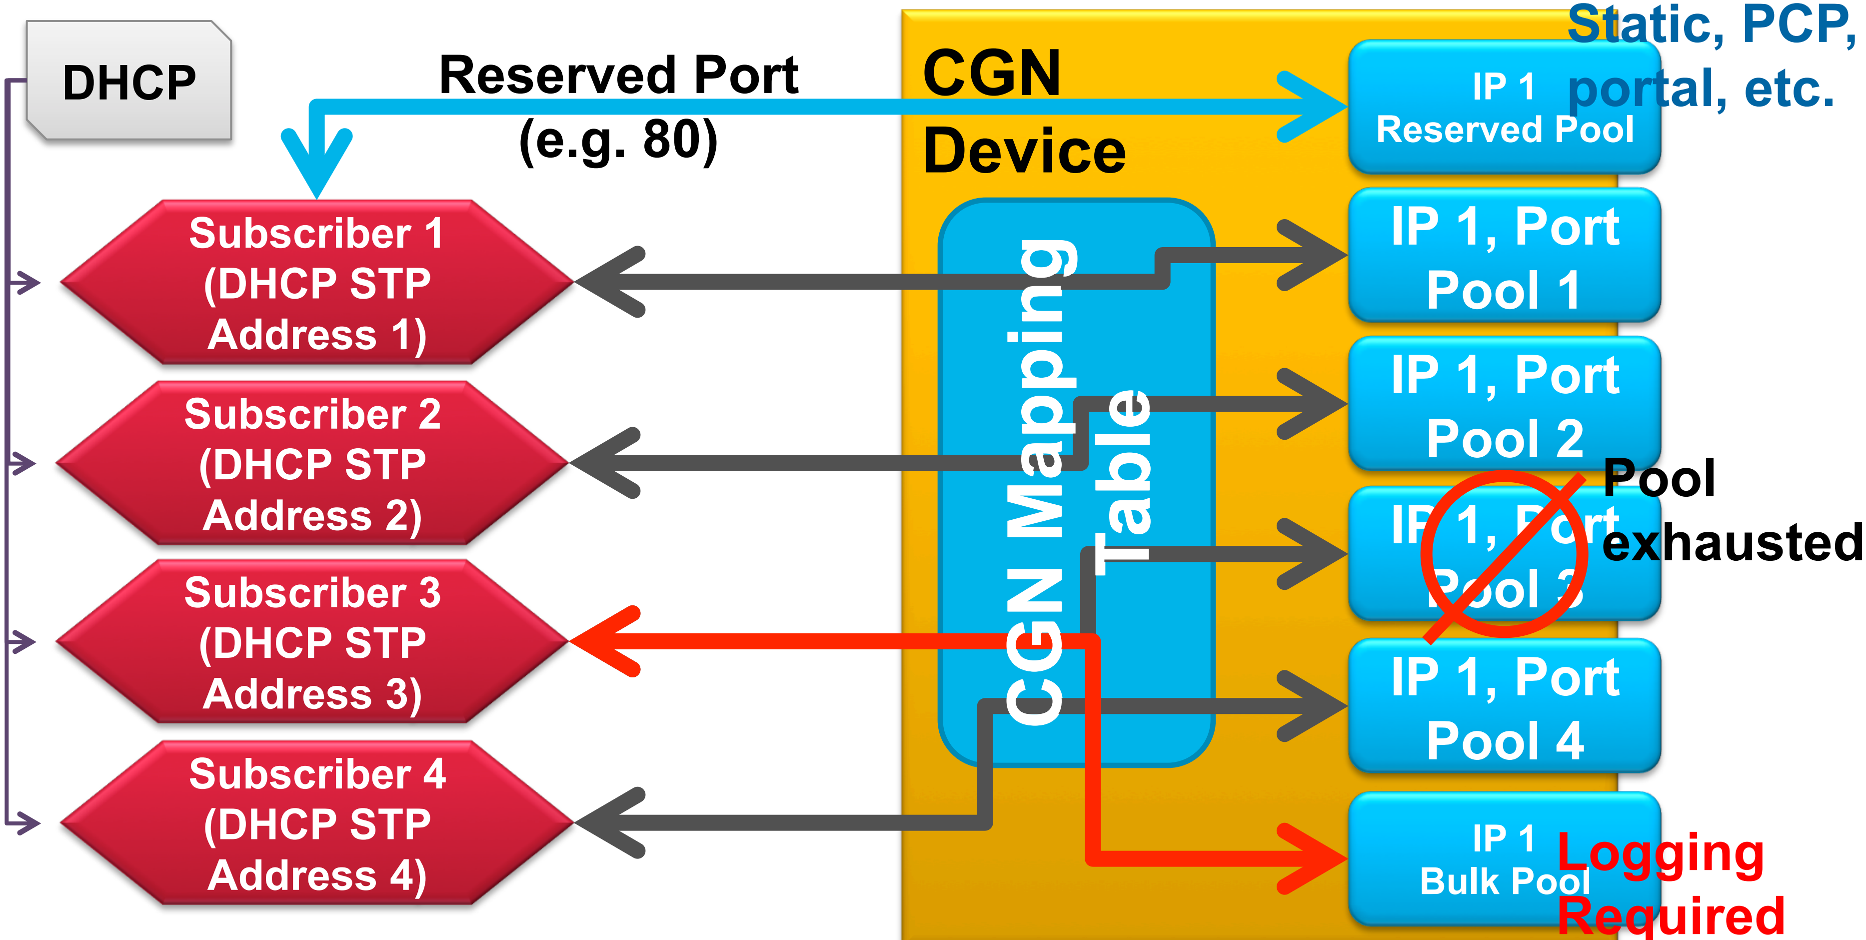

Security Considerations

- CGN Inside IP Space Filtering:
  - Block CGN routes from being advertised to and from peers
  - Block traffic with CGN source or destination IPs at borders
  - This filtering likely does not happen on the CGN device
- DOS Mitigation at the CGN:
  - CGN device becomes target for DOS and other IP-focused attacks from outside your network
  - CGN device is also bottleneck for attacks sourced from CGN subscriber networks

IP Address Reputation

- IP blacklisting is more problematic with multiple subscribers behind a single outside IP
    - All subs behind that IP are affected
    - Any sub behind that IP can cause the listing
- Examples:
    - Secure transactions (Banking, Storefronts, etc.)
    - Email spam lists (Spamhaus, etc.)
    - Individual website blocking (comment spam, etc.)
- Difficult to troubleshoot
    - Requires CGN logging

The Big Picture and What's Next

CONCLUSIONS

CableLabs®
...Revolutionizing Cable Technology®

IPv6 Offers a Better Experience than Shared IPv4

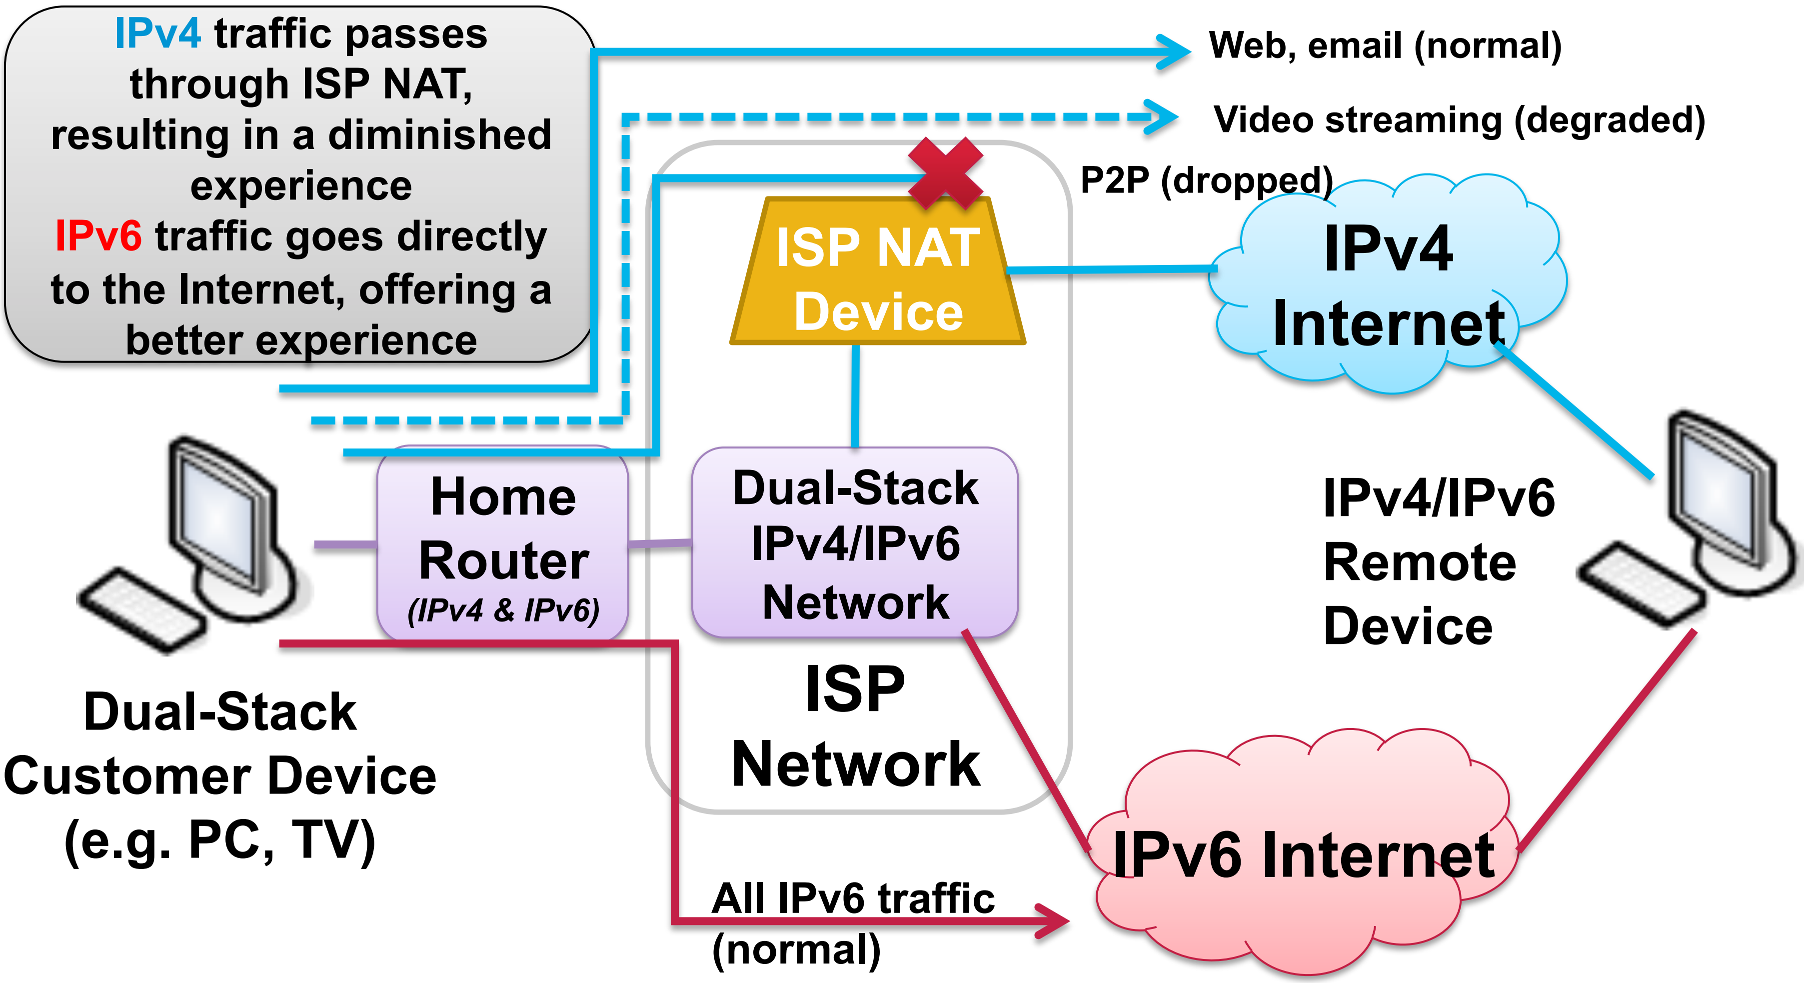

Requires Dual-Stack (IPv4 & IPv6) PC and Home Gateway

CableLabs®
...Revolutionizing Cable Technology®

We still have a lot of work to do!

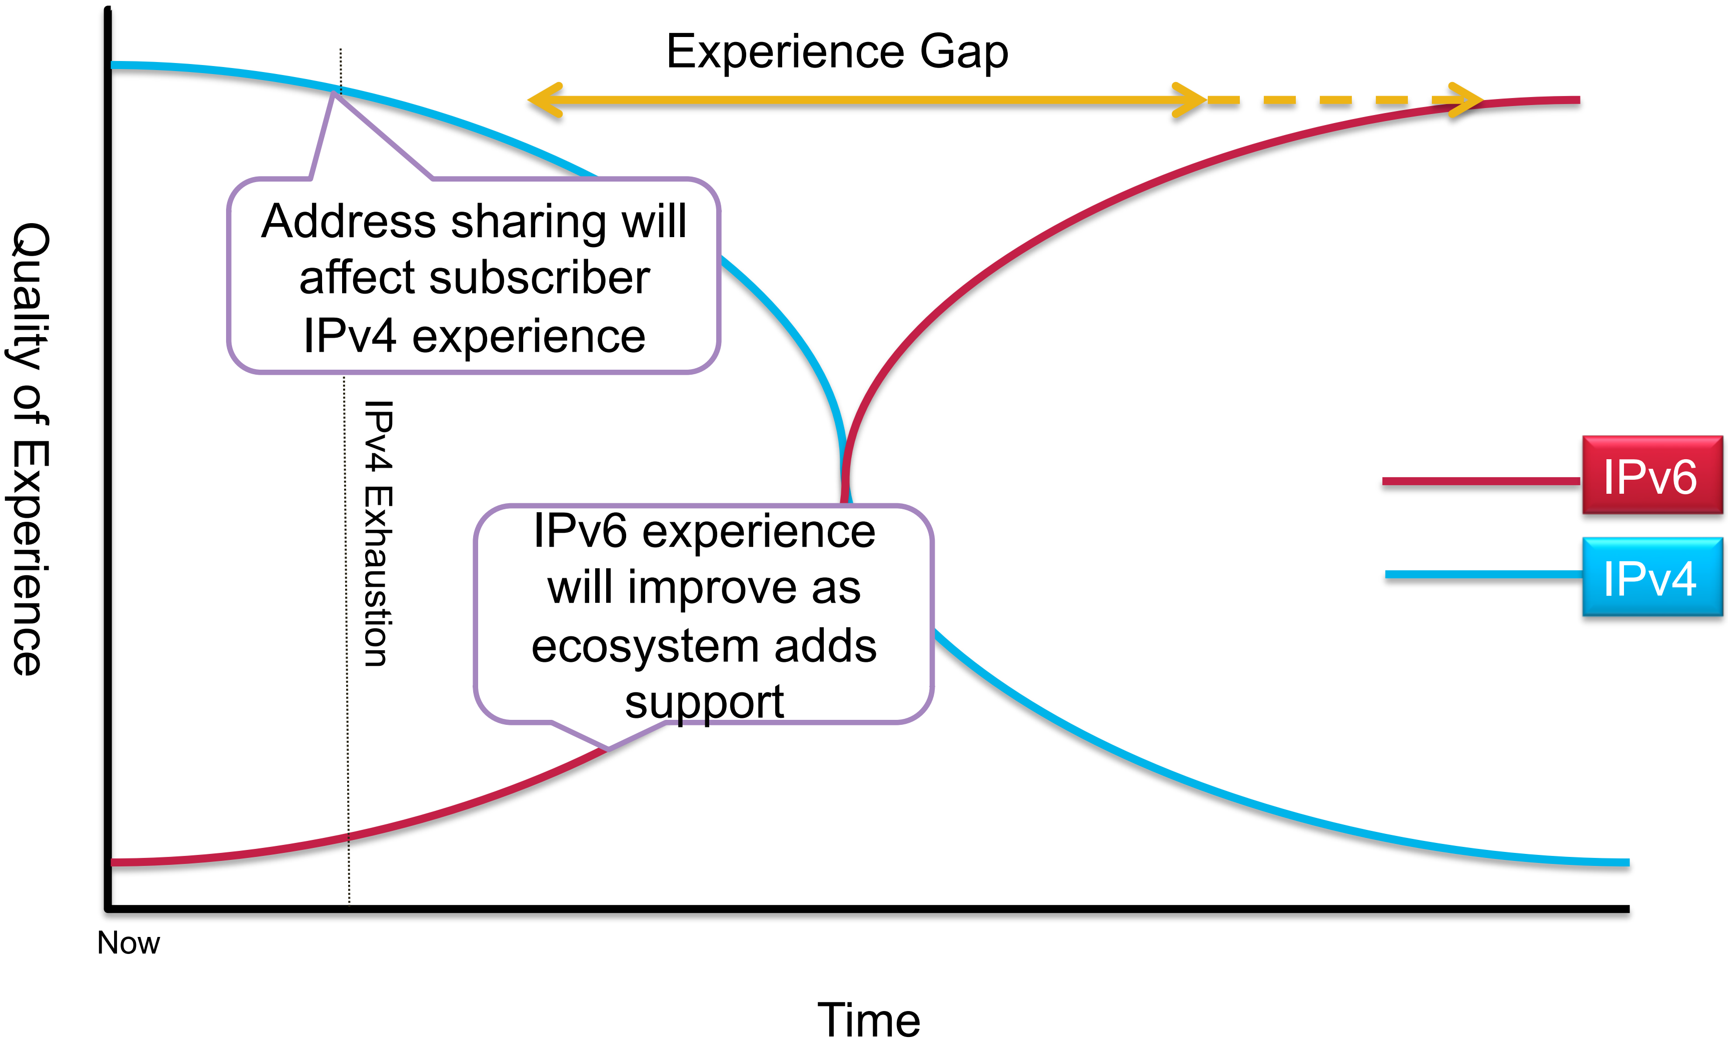

  4/6/12

In Short

- IPv6 is the answer to IPv4 address exhaustion
- CGN can support legacy IPv4 systems for some time
- Deploying CGN **will** impact your customers
  - P2P, VoIP, gaming, video, streaming & geolocation, etc.
  - For many, a necessary evil to maintain IPv4 service
- A properly designed architecture can help
  - Optimize routing, latency and jitter
  - Reduce logging requirements
  - Improve targeted advertising results
  - Mitigate the impact on your customers

Questions?

Chris Grundemann
c.grundemann@cablelabs.com
http://chrisgrundemann.com

