Capital Technology Management Hub

IPv6 Implementation Lessons Learned and Motivation in the United States

A Panel Discussion
April 13, 2010

Qwest IPv6 Implementation Experience

Shawn Carroll

Previous Qwest Implementation Work

- Obtained 6bone Pseudo Next Level Aggregator (pNLA) from Abilene in 1999
- Obtained production Sub Top Level Aggregator (sTLA) 2001:428::/35 in 2000 (now /32)
- Built IPv6 test network in 2000
  - Overlay
  - Native IPv6 across OC3s and Generic Routing Encapsulation (GRE) over IPv4 OC48s
  - Cisco 7507s and 12008s
  - 9 Total PoPs across the country
  - Alpha customers connected via GRE over existing IPv4 circuits
- Built to gain experience with operating a native IPv6 network
  - Gauge customer interest
  - Maintained v6 peering connectivity

Qwest IP Networks => IPv6 Networks?

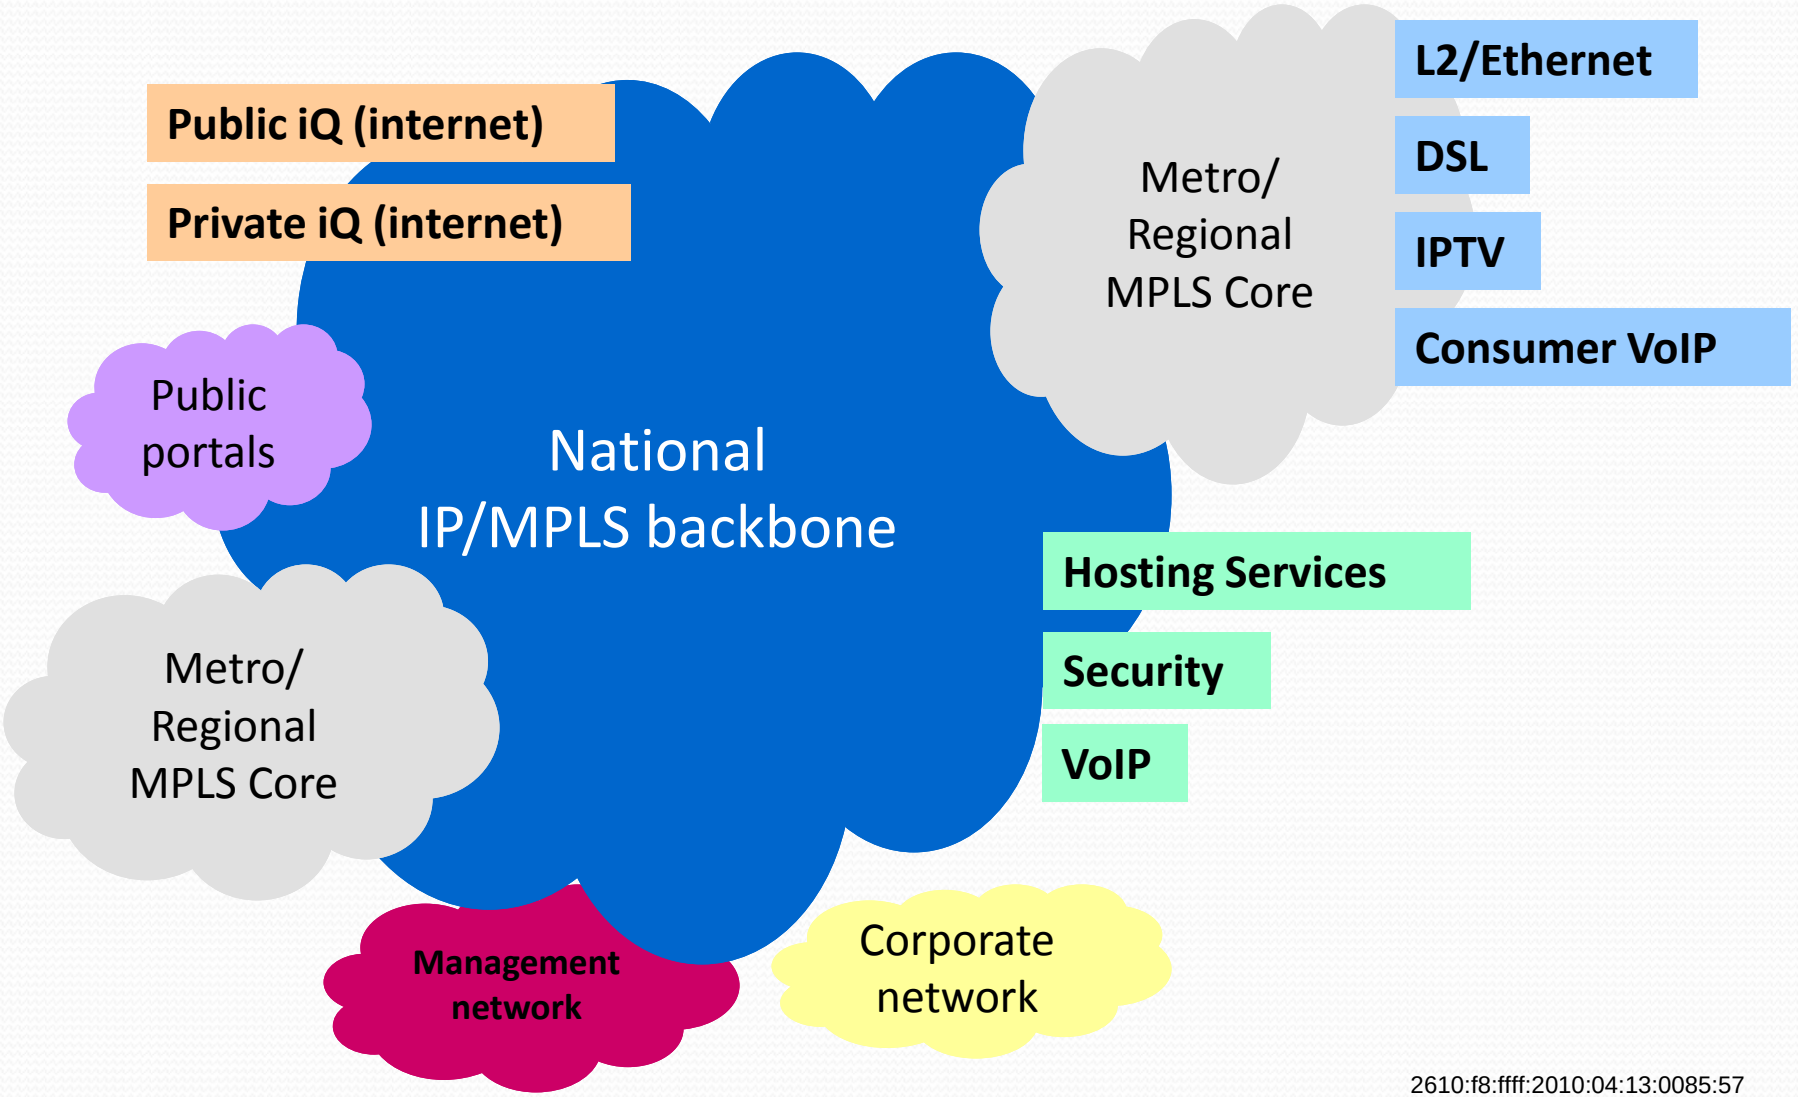

General Environment

- Overall transition plan is a phased approach
  - Need to evolve the systems and the network
  - Work from the IP Core out toward edge
  - Work from lower layers (L3) to application layers (VoIP/CDN/etc)
- Challenges
  - Resource contention
    - If you wait until 2011, you have a significant challenge
    - But hard to justify applying scarce resources today when they can be applied to other projects with superior financial metrics
  - Business case
    - Customers seem to expect IPv6 to be free
  - Integration with ongoing projects
    - Hard to integrate without stalling the product pipeline
- Seek balance
  - Ensure that IPv6 work is performed in a measured way
  - Transition networks that need it first/derive most benefit
  - Make sure that the network has been assessed, regardless of transition time

First Phase Implementation?

- Target international backbone first
  - Need to get your core working before it makes sense to work on clients of core
- Implement public peering
- Implement basic public and private IP services
- Make systems support IPv6
  - Inventory, ordering, provisioning
- Seek a totally integrated solution

IPv6 Service Objectives

- Target specific services that need IPv6 first
- Enable IPv6 equivalents of existing IPv4 iQ services
  - Public port – Connect to the public internet
    - Options include customer static routes, BGP with customer, Qwest vs. customer address space
  - Private port – L3VPN product
  - Enhanced port – Mix public and private services on a single interface toward the customer
    - Only support L2 separation initially
- Enable full mixing of IPv4 and IPv6 on the same physical port
  - In full complexity, an enhanced port would offer public access and private L3 VPN for both v4 and v6 on a single customer interface
  - Across all interface types (Ethernet, POS/TDM, ATM, FR)

Overall Design Considerations

- Does IPv6 warrant a new/different network design?
  - Architecture review determined:
    - Existing architecture supports IPv6 well
      - Same network fundamentals are required as in IPv4
    - All existing protocols support the IPv6 address family
      - ISIS, BGP, MPLS, RSVP
    - MPLS 6PE/6VPE meshes well with existing network
      - Provides a bridge until native IPv6 implementations of MPLS are available from suppliers
- Implement via overlay?
  - Doesn't scale operationally
    - Need to manage two networks
  - Makes dual stack customer ports difficult to implement
    - A significant percentage of existing customers will add IPv6 to their existing service
    - Commercial growth rate makes scale a near term problem

Hardware Assessment

- Which hardware:
  - Existing inventory partitioned into three classes:
    - Must support: New/required
    - Never support: Old/hopeless/too expensive to fix
    - Maybe support: Older but support IPv6, near technology refresh
  - Work queued for “musts” first, then “maybes”
- Certification work:
  - Five different network roles
    - Public edge, private edge, agg, P, border
  - Ten different element types
    - Roughly a set of 200 carrier card+daughtercard combinations to certify
  - New ACLs and policy maps to be developed
- Types of testing
  - Redo throughput on most IO cards
  - Scale testing: CPU testing, memory (RIB/FIB)
  - Must verify both IPv4 and IPv6 performance

System Assessment: Network OSS

- Scope: Around 18 systems need to be touched
  - Plus around 30 network scripts used by operations
- IP addresses are used in many of the systems
  - The interface IP address is used by many systems to identify the service
- Significant upgrades:
  - Inventory
  - Alarming
  - Performance
- Each different system must have the entire connectivity path audited to ensure any v6 required upgrades
- Not all 3rd party systems had IPv6 components ready

System Analysis: Business OSS

- Sales:
  - Customer service forms and systems
  - Sales engineering processes
- Billing:
  - Flow addresses to bill?
  - Charge for IPv6?
- Customer portals
  - Display IPv6 information, performance
- Ordering
  - OSS backplane that interconnects the above systems.
- Result:
  - Almost all systems had IP addressing that had to be touched.
  - Took significant time to discover where addressing was already implemented

IGP Considerations

- Two choices for support of IPv6
  - Single topology ISIS:
    - Assumes that the IPv6 and IPv4 topology are the same.
  - Multi-topology ISIS:
    - Different topologies for IPv4 and IPv6. Two SPF runs.
- ISISv6 support affects all elements in IGP domain, including elements that aren't taking part in IPv6 services.
  - Elements are supposed to ignore TLVs they aren't using or don't understand.

IGP Implementation

- Single topology ISIS seemed to be best choice
  - At least one significant element in the network that needed to support IPv6 had a MT-ISIS hardware limitation
  - Better match for single session iBGP
  - Less operations work short term
    - Link bounces for MT-ISIS
  - More generic vendor support
- Push suppliers to support “transition mode” to make MT-ISIS transition in the future simpler.

BGP

- Dual stack can be implemented over one or two sessions:
  - iBGP: one session over IPv4, two families in the session
    - Tracks that MPLS and iBGP topology derived from IPv4 topology
  - eBGP: two session: IPv4 family over IPv4, IPv6 over IPv6
    - Tracks inter-carrier topology can diverge for families

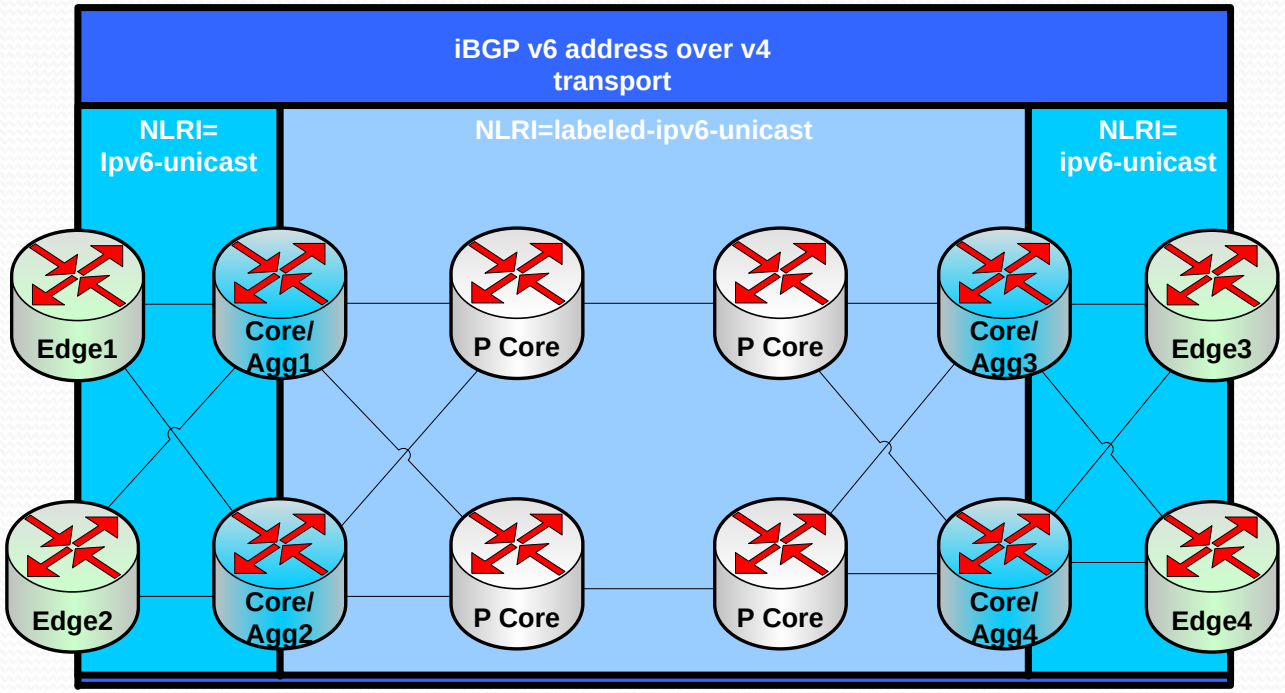

Peering

- Target: dual stack existing ports with IPv6
  - Many key peers only supported tunneled connections
  - Tunneling added
    - GRE and 6in4 both supported

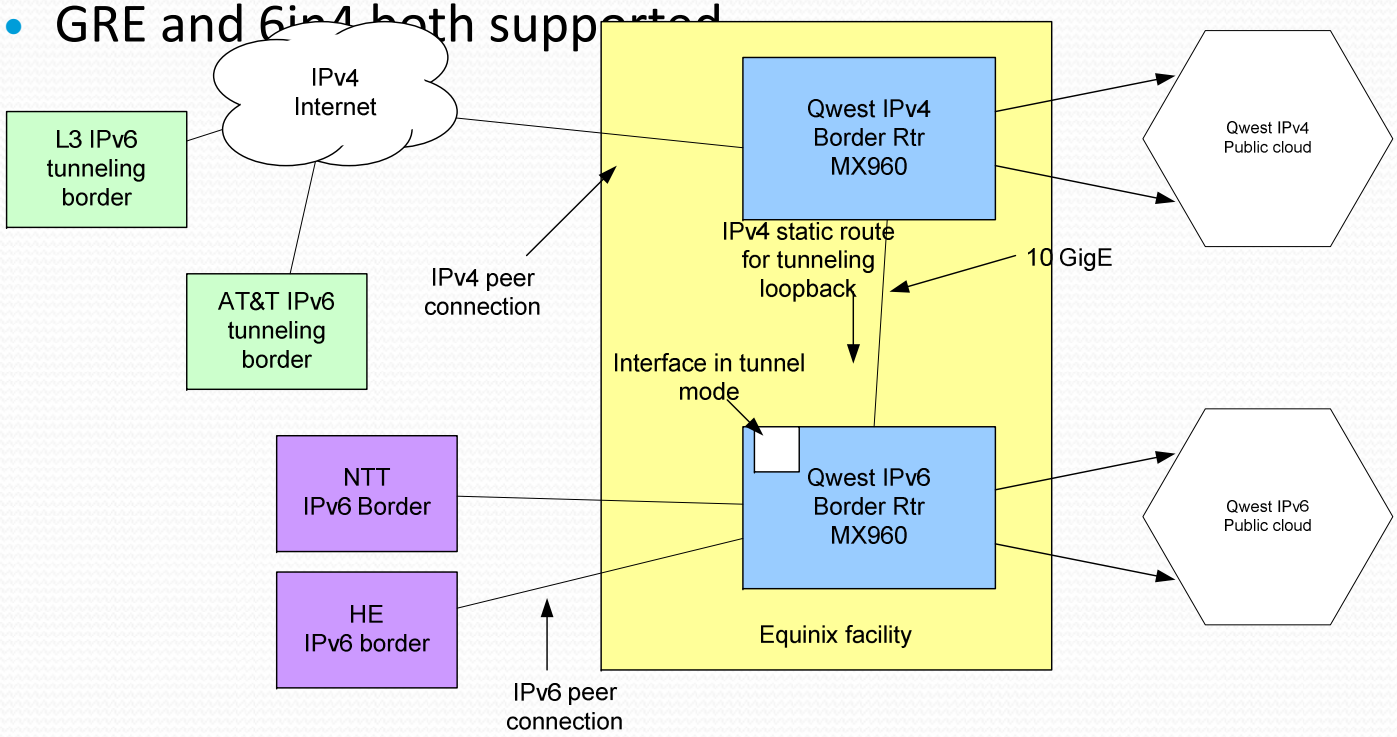

DNS

- DNS queries can be made over IPv4 or IPv6
  - A query can return records that contain both IPv4 records (A) and IPv6 records (AAAA)
- A small number of new DNS resolvers were added that support queries over IPv6 transport
  - Distributed throughout country to maintain a service latency objective
  - DNS systems were upgraded to record customer IPv6 DNS service components
- Servers are cheap, not worth the effort to dual stack existing servers until there is significant DNS query traffic

Actual Implementation Timeline

- Systems and Certification delivered on time [June 08]
- Issues on rollout:
  - Being cautious, chose to initially deploy 2 “IPv6 only” border routers
    - Within a week of turn-up, uncovered a “rare” RIB tree lookup bug that causes the router to reboot
  - Wait for new software release to be certified and deployed [~5 month delay]
    - New release is stable, dual stack policy rolled out to all border routers
- Beta tests in 1H2009

Issues Unique to IPv6

- Most suppliers HW support regular IPv6 data forwarding just fine
- But carriers typically need more than just data forwarding
  - Vendor support IPv6 unicast RPF
    - Not all suppliers support it yet
  - Vendor support for sFlow / netFlow
    - Used as a packet accounting mechanism and security mechanism
    - Little support for IPv6
    - Consequently, 3rd party support not available till mid/late 2008
  - Differentiated IPv4 and IPv6 stats
    - Useful on combined interfaces and shared LSPs
  - Inconsistent implementation of address representation
    - :: zero collapse

Lessons Learned

- Most of the network elements required two certification runs
  - First to identify bugs, wait for supplier to fix, and final certification run
- Hardware in general is still not at feature parity or with IPv4
  - sFlow / netFlow, and the tools that support the associated analysis are still not carrier grade
  - SNMP MIBS are non-standard between vendors
  - Counters: Should be able to answer how much IPv4 vs. IPv6 traffic is flowing over a dual stack interface
- Third party system support is extremely slow to develop
- Training so far better than expected (BGP is BGP)
- Dual stack/integration can slow down rollout, as other services drive release dates when IPv6 specific fixes are required

Conclusion

- Most Qwest IP networks have been aligned into some form of IPv6 transition plan
- Core network transition is complete
- Many of the more complex customer facing networks will take several years to transition