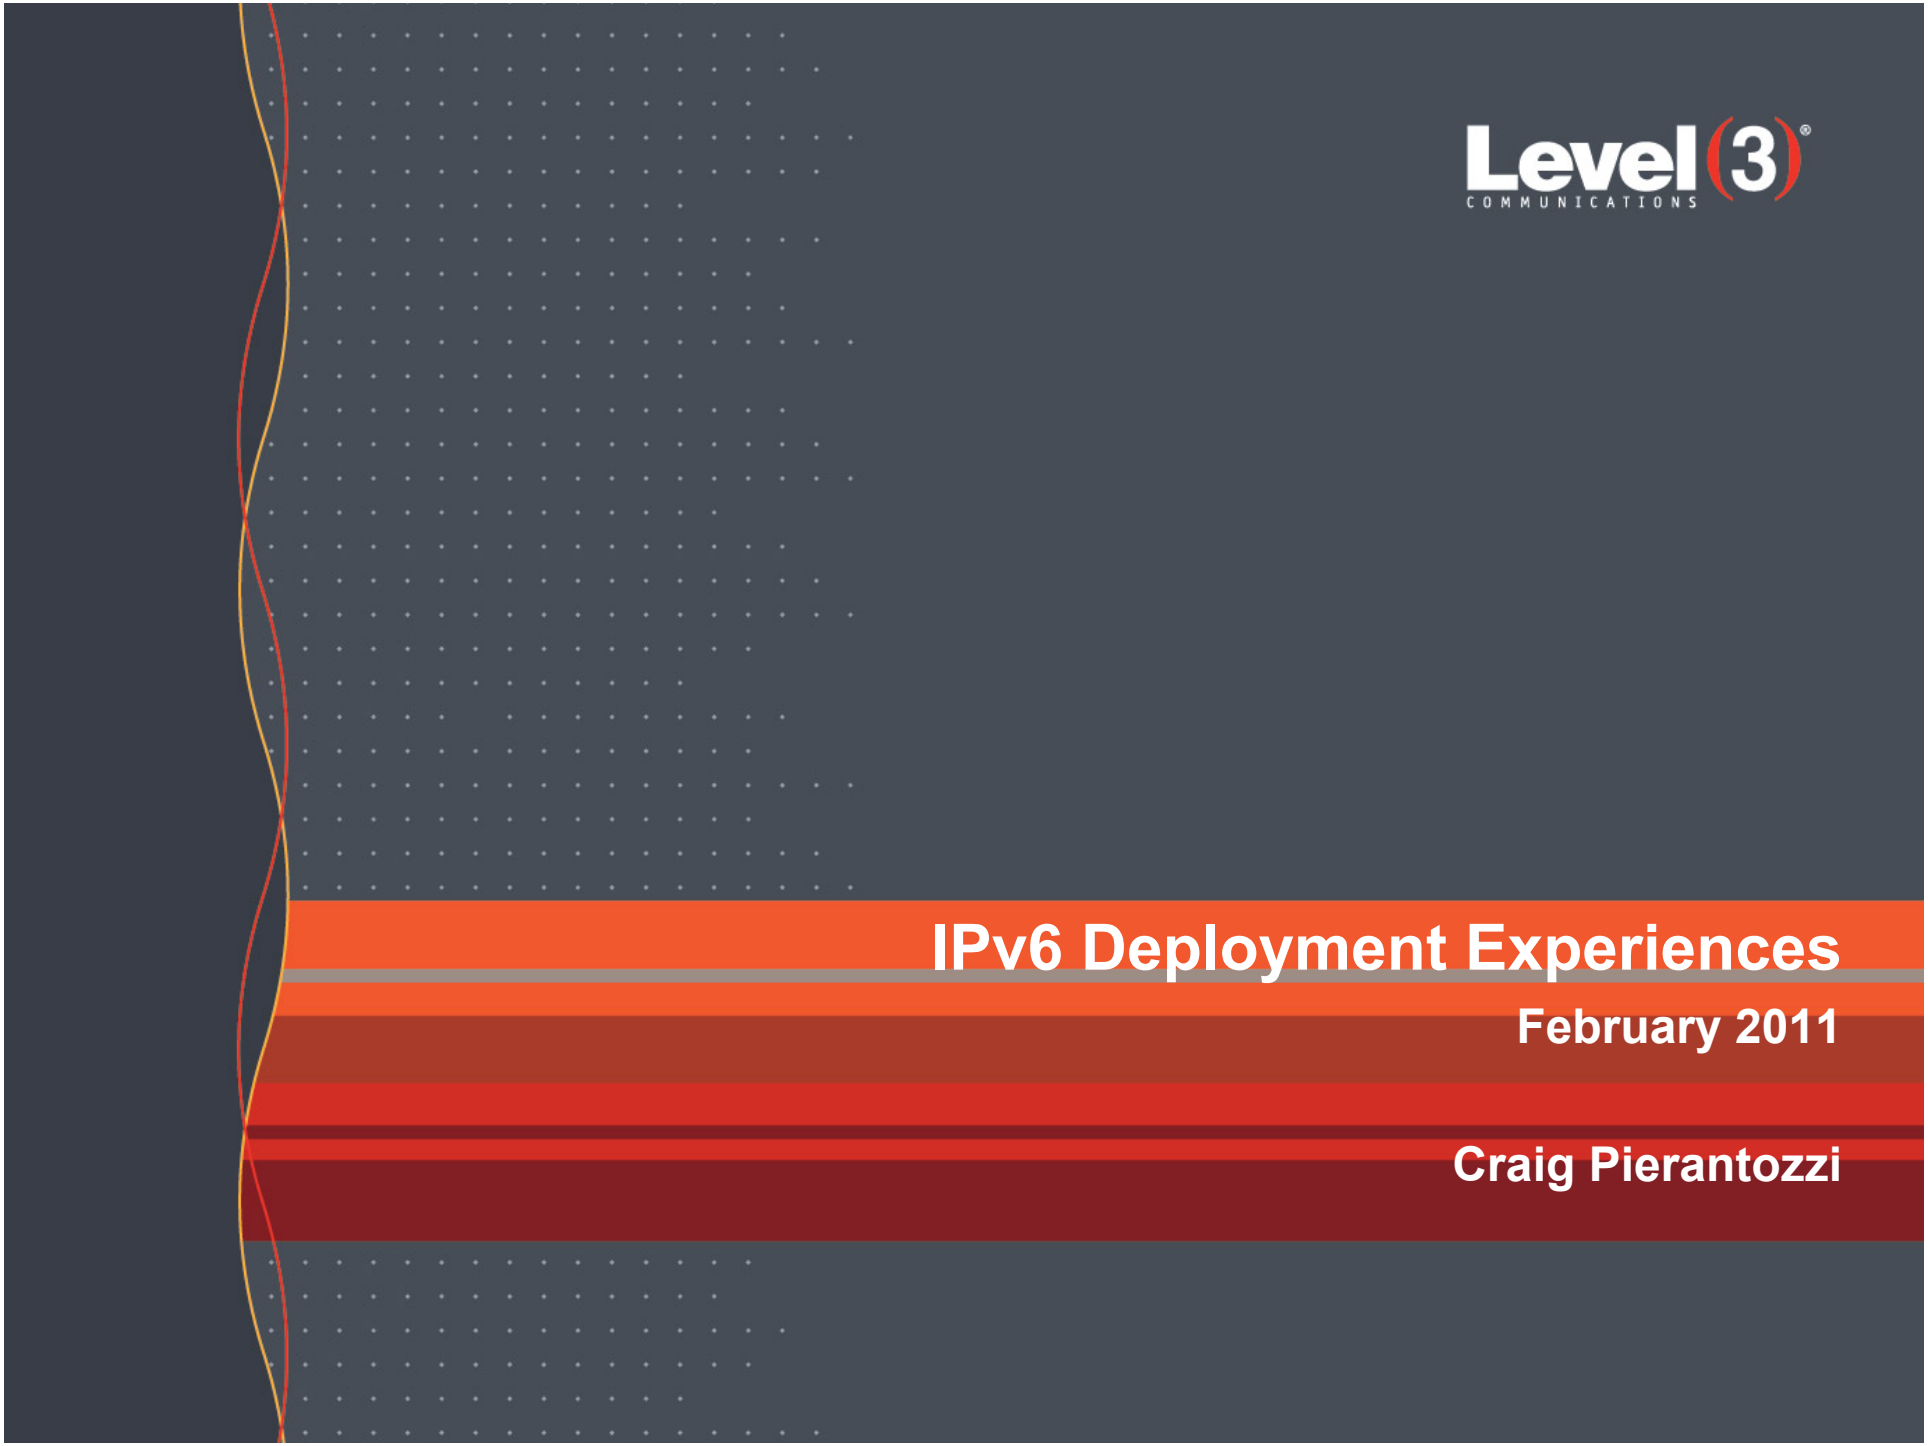

IPv6 Deployment Experiences

February 2011

Craig Pierantozzi

Agenda

- Level 3 IPv6 history
- IPv6 Design Goals and Implementation
- Issues and Observations
- Takeaways

History

- Started offering tunnelled IPv6 service in May 2005.
  - EU-based 6PE solution.
  - Limited number of selected customers and peers using tunnels or @LINX using low speed ports.
- 2005 - 2008 the solution grew and grew…
  - Tunnel boxes deployed in North America for better geographic coverage.
  - Routing options reached parity with IPv4 (i.e. community based TE).
  - Larger number of customers and peers. AMSIX peerings also established.
  - All connections standardized to manually configured tunnels. No longer 6PE.
- 2009 started offering dual-stack services in EU.
  - Some customers started to genuinely “demand” dual-stack.
  - Based on DPC tunnel service on Juniper MX960 in limited locations and still connected back to tunnel based service.
- Today the network is now fully dual stack enabled, GA was July 1, 2010 and the routing ecosystem is ‘evolving’.

Goals and Decisions

“Thou shall not break the IPv4 network”

- Tunnel solution was difficult to scale due to hardware limitations and in 2009 demand for a full IPv6 solution reached the tipping point.
- Issue (as it is for us all) was keeping core network costs down.
- IPv6 over VLAN to provide economic and technical efficiencies.
  - Allows us to provide “full speed” dual-stack IPv6 everywhere.
- Separate IGP for IPv6 for safety & stability reasons.
  - OSPFv3 chosen over existing IS-IS.
- Moved route reflection off tunnel devices onto standard route reflectors.
- IPv6 enabled natively on all existing LANs, alongside IPv4.
- Make the overall design as similar to IPv4 as possible for Operations

IPv6 Design

- PE pairs function as ‘core’ devices
  - LANs run IPv6 & OSPFv3
  - Ptp VLANs for intercity connectivity
  - No segregation of IPv4 and IPv6 traffic on LAN broadcast domain.
- Existing route-reflectors stood up for IPv6 reflection.
- Allows mapping IPv6 forwarding-path for intercity transit to follow same path as IPv4 if built correctly.
- Future design will go back to IPv6 over an MPLS core

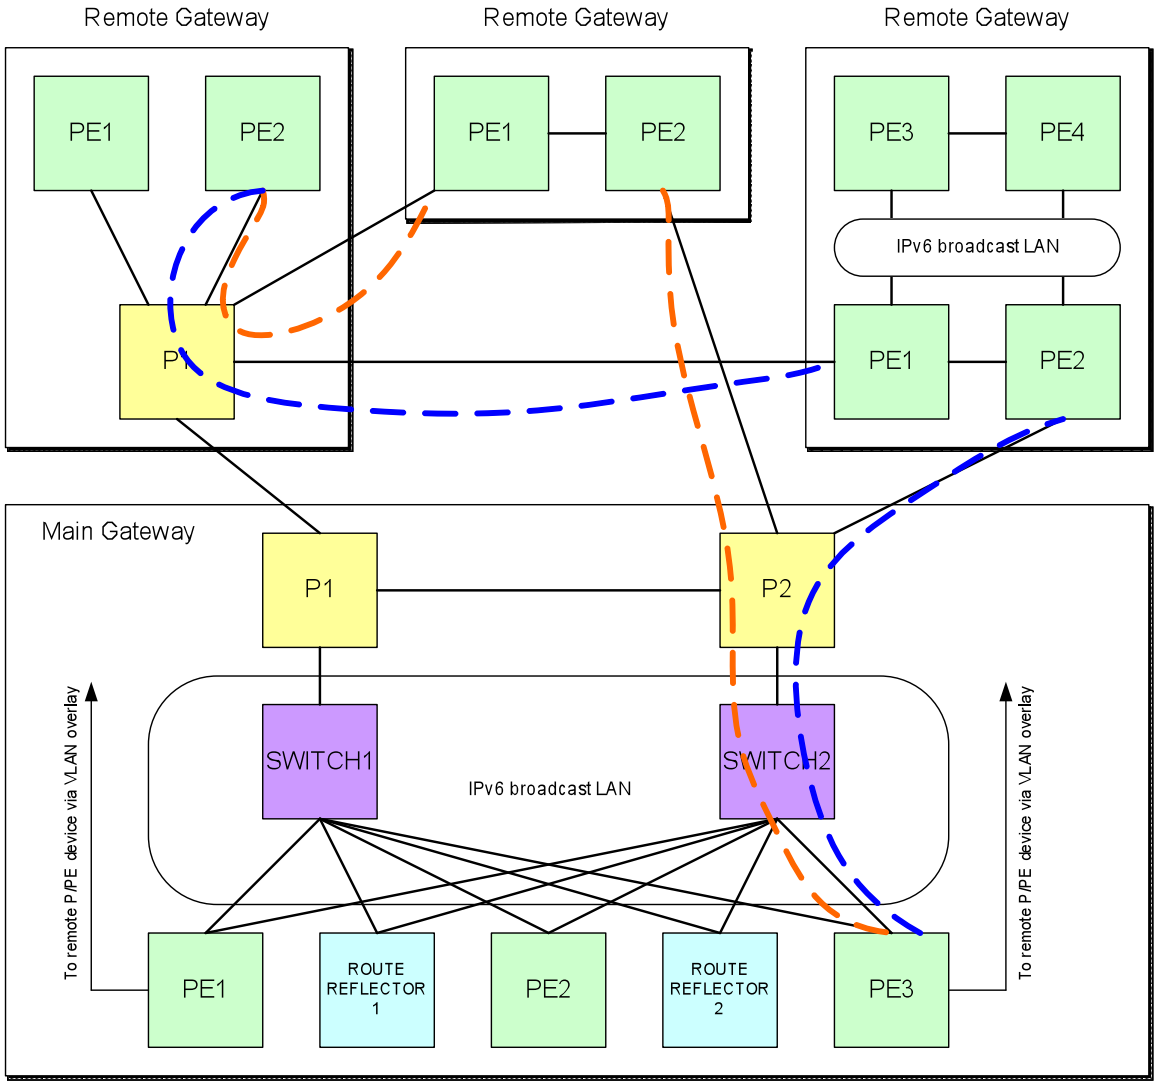

Supporting IPv6

- You can chose a number of ways to support IPv6 on your network.
  - Transition – Tunnels and such which customers use to connect in some way, over your existing IPv4 network. If you like that sort of thing…
  - Dual-stack – Actually enable IPv6 on your existing network, at least at the edge.
- Is your existing kit capable of supporting IPv6?
  - There's still a LOT of very old kit out there and some of it doesn't support IPv6 at all.
  - Some of it supports IPv6 but only in software.
  - Some of it supports IPv6, but only with specific software versions.
  - Some of it supports IPv6, but TCAM issues prohibit its use.
  - Some of it supports IPv6, but load-balancing isn't good.
- You need to understand what your customers need from you before you can make this decision.
- If they're going to leave if you don't support dual-stack, you need to have a plan in place one way or the other!

Platform Issues

- Most modern network hardware supports IPv6 in some way.
- Check that your kit includes (full) IPv6 support.
  - But do you have to update all your configs? ADDRESS-FAMILIES!!!
- TCAM issues:
  - On many platforms, TCAM is used for fast route table lookups.
  - The way TCAM is ‘carved up’ means it can’t be shared in many different ways (ACLs, QoS, MAC, IPv4 and IPv6) allocated in ‘profiles’.
  - Allocating a reasonable amount of TCAM space to IPv6 can mean you have a very limited amount of space left for IPv4. You may run out sooner than you think! How much would this cost you?
  - Have a look at profiles and think what would be reasonable for you for now and you may come up with different answers for your core and your edge.

Differentiating IPv4/IPv6 Traffic

- One common problem is how to create differentiated data sets and graphs for IPv4 and IPv6 traffic on your network.
  - Most existing systems just SNMP poll IF-MIB which just counts octets in and out. It doesn't care what protocol the traffic is and this can cause support and planning issues for you and your customers.

- No easy solution, but here are two things commonly done:
  - On the 'J' vendor, have an inet6 traffic filter inbound and outbound to count IPv6 traffic in and out of interfaces. You can collect this data by SNMP or CLI, although you'll probably have to do some special stuff to make it work in your regular package.
  - On most vendors, you could have different VLANs on a physical interface for IPv4 and IPv6 traffic. This allows you to measure per VLAN and therefore do differentiation. However this isn't *real* dual-stack, introduces other problems and often isn't an acceptable solution for customers.

- You could use NetFlow, but this requires v9 and a good support.

Other Important Factors to Consider

- Supporting IPv6 on your network are **NOT** the biggest issues you'll encounter.
- Much bigger are issues include:
  - **Training your staff**
    - Technical and non-technical staff need to be trained on IPv6. They need to ensure they understand it so they can sell it correctly, and support it correctly.
    - Who is going to do your IPv6 training? Do you have an in-house expert or are you going to outsource your IPv6 training requirements?
    - Once the staff are trained, how will you keep their skills up to date?
  - **Supporting your back-office systems**
    - This can be a large number of new and old systems.
    - This is often non-trivial and can take huge amounts of time and expense.
    - Are you sure your billing system is going to be happy with IPv6?
  - **Supporting your network management platform(s)**
    - Does your existing NMS support IPv6? How about load balancers, firewalls, IDS etc?

Takeaways (1)

- Good news is that most modern router and server kits do support IPv6 properly…BUT code still has bugs because it hasn't been stressed and people haven't cared much to this point when it's failed.
- Many monitoring systems etc. do not support IPv6, at least not as well as IPv4.
- Security is likely to be a messy. How many of the firewalls, AV, load balancers, VPN and such support IPv6 properly?
- Skill base for IPv6 is still VERY low.

Takeaways (2)

- If you don't already, you really need to have a plan in place for IPv6.
  - Many networks are now complete, so how are you going to compete?
- Don't assume that activating IPv6 on your network is the only task. You may be pulling along other organizations!
  - Think about training.
    - Even to the point of (re)training everyone in hex!
    - Differentiate your training.
  - Think about back-office systems.
  - Think about monitoring and supporting IPv6.

Thank You

Questions?

Craig Pierantozzi
Sr. Director, Network Engineering
craigp@level3.com