IPv6 @ Comcast

Managing 100+ Million IP Addresses

Alain Durand

Office of the CTO
Director – IPv6 Architect
Alain_Durand@cable.comcast.com

comcast

Agenda

- Background
- Comcast Approach
- Where We Are
- What We Learned

Background

When Net 10 is Too Small….

- Unlike DSL modems, Cable Modems are managed and consume an IP address.
- Until recently, Comcast was using Net 10 (RFC1918) for managing the cable modems:
  - That space was exhausted in 2005.
  - Since then, Comcast was allocated a very large bloc of public IPv4 address space for device management
- In the control plane, all devices need to be remotely managed, so NAT isn't going to help us, nor is federated Net 10 islands…
  IPv6 is the clear solution for us.

Simple View of the Scope of Comcast IP problem

Set-Top boxes moving from proprietary management to Docsis-based IP management (provisioning, EPG,…)

- 20 Million video customers
- 2.5 set-top box per customer
- 2 IP addresses per set-top box

- Total: 100 Million IP addresses

Note: This is not IP-TV

And we have not yet talked about High Speed Data…
nor Comcast Digital Voice…
nor merger/acquisition…

Triple Play Effect on the Use of IP Addresses

	2005 HSD only	2006+ Triple Play
Cable Modem (CM)	1 (private only)	1
Home Computer / Router	1	1
eMTA (Voice adaptor)	0	1 – 2
Set Top Box (STB)	0	2
Total number of IP addresses (assume 2.5 STB per household)	**1 – 2**	**8 – 9**

IP Addresses: Natural Growth vs New Services
(in the coming years)

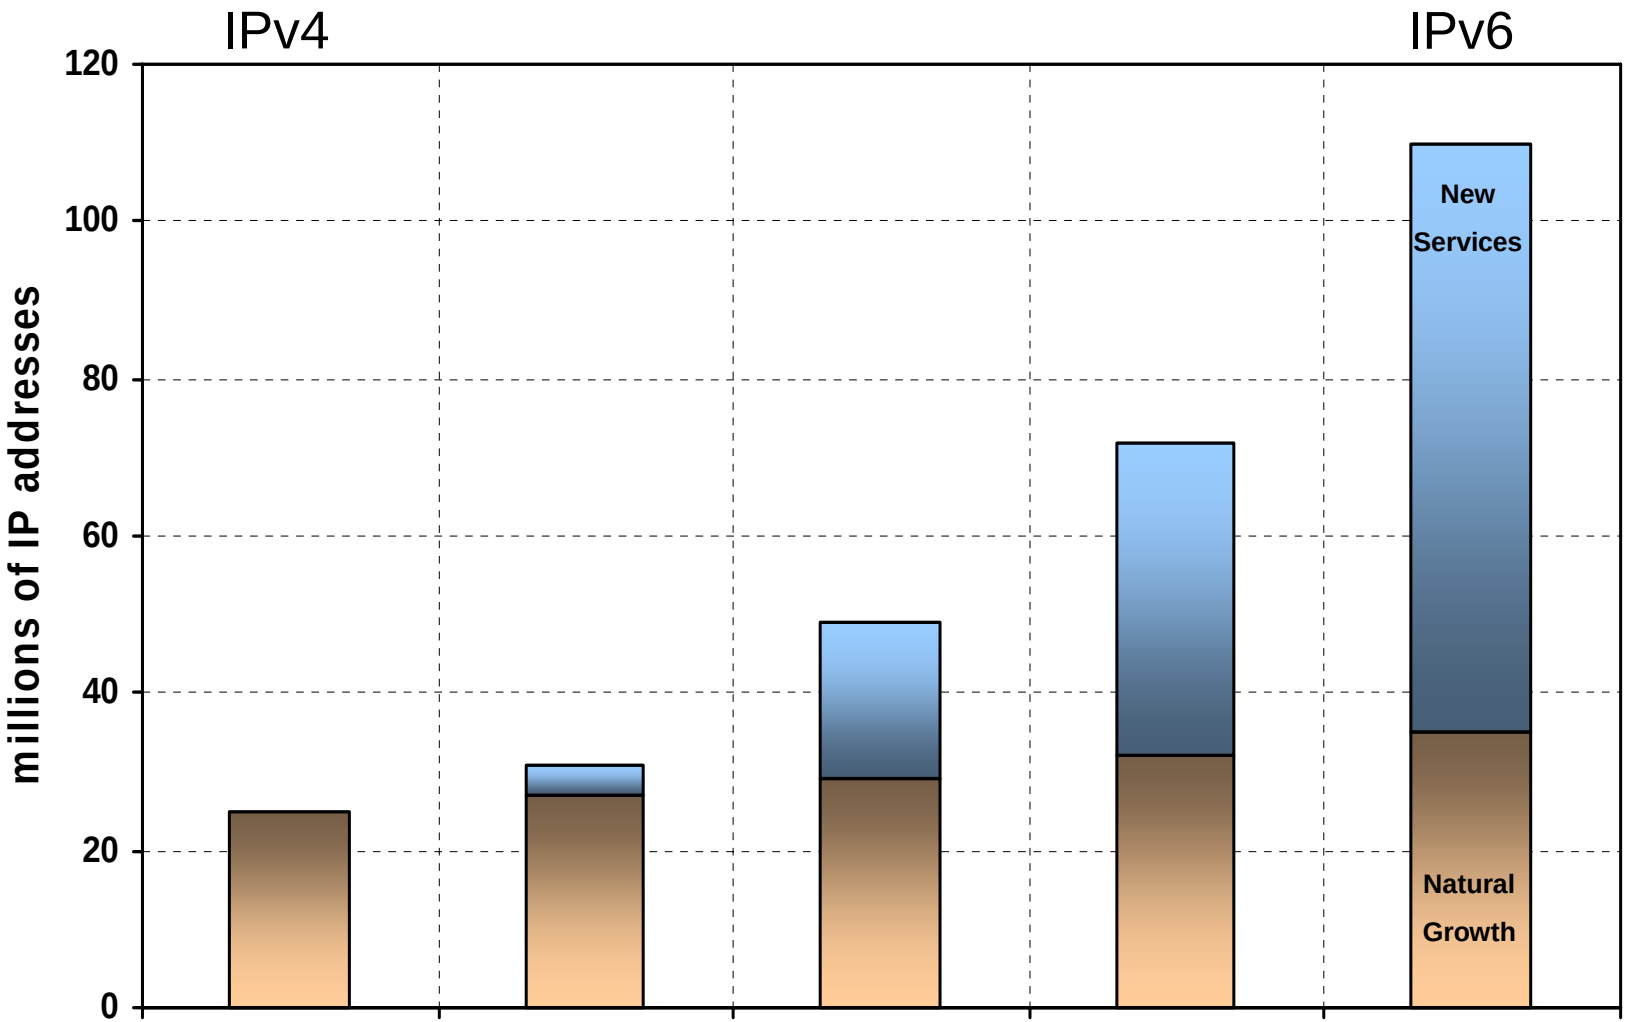

Note: this graph shows trends, not actual data

Comcast Approach

Comcast IPv6 Strategy

- ***Start early***
  - Deployment plans have started back in 2005
- Deploy IPv6 ***initially*** on the ***Control Plane*** for the ***Management*** and ***Operation*** of the ***Edge Devices*** we manage
  - Docsis CM, Set Top boxes, PacketCable eMTA (Voice),…
- Be ready to offer our customers new services that take advantage of IPv6

IPv6 Deployment: Principles and Approach

- Deploy IPv6 ***only where it is absolutely necessary***
  - *and nowhere else!*
- Architecture: ***dual-stack at the core, v6-only at the edges***
- Deployment approach: ***from the core to the edges***
  - Backbone -> Regional Networks -> CMTS -> Devices
  - This is an incremental deployment; existing deployments will be unaffected in the beginning.
- Follow same operational model as with IPv4
  - DHCP-based provisioning and access control

Modems and “Single IP Version” Mode of Operation

- New CM will be IPv6 ready (dual-stack capable)
- On an IPv4-only CMTS, CM will be provisioned with IPv4
- On IPv6-enable CMTS, CM will be provisioned with IPv6
  - **CM will never have both IPv4 & IPv6 addresses at the same time** (If we could give both an IPv4 and an IPv6 address at the same time, we will not need IPv6 in the first place!)

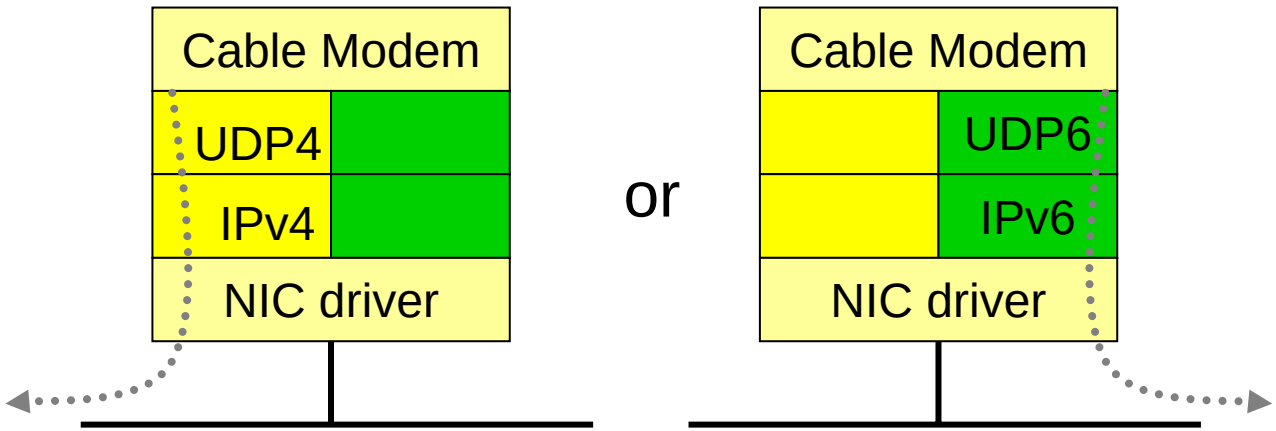

Note: the modem is a bridge, it will forward IPv4 packets even if it is configured in an IPv6-only mode

IPv6 Architecture for Devices & Services

- High Speed Internet
  - customer service remains IPv4 for now
    - May add IPv6 service in later phase
  - IPv6-ready modems will be manage only with IPv6

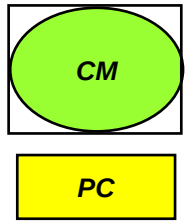

- Comcast Digital Voice
  - The embedded modem of eMTA will be managed only with IPv6
  - The MTA itself and the soft-switches will remain IPv4 for now

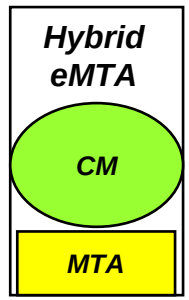

- Video
  - New Set-Top boxes will have their embedded CM and their host stack managed only with IPv6.

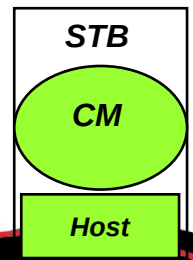

IPv6 Architecture for Applications: Provisioning, Monitoring, Back-Office

- Mostly a software upgrade problem
  - Similar to the Y2K problem,
    Fields need to be bigger in database & web scripts
- Should “application X” be upgraded for IPv6?
  - Transport questions: Does “application X” communicate with devices that are potentially IPv6-only (e.g. CM)?
  - Payload questions: Does “application X” manipulate IP data (store, input or display) that could be IPv6?
- Comcast inventory analysis: about 100 “applications”
  - 10 need major updates for transport
  - 30 need minor updates for display/storage

IPv6 Architecture for Back Office

- Back-office systems that do communicate directly with the CM or STB migrate to dual stack
- The other back-office systems keep using IPv4
- However back office systems may need to be modified to display/input/store IPv6 related data (CM IPv6 addresses)

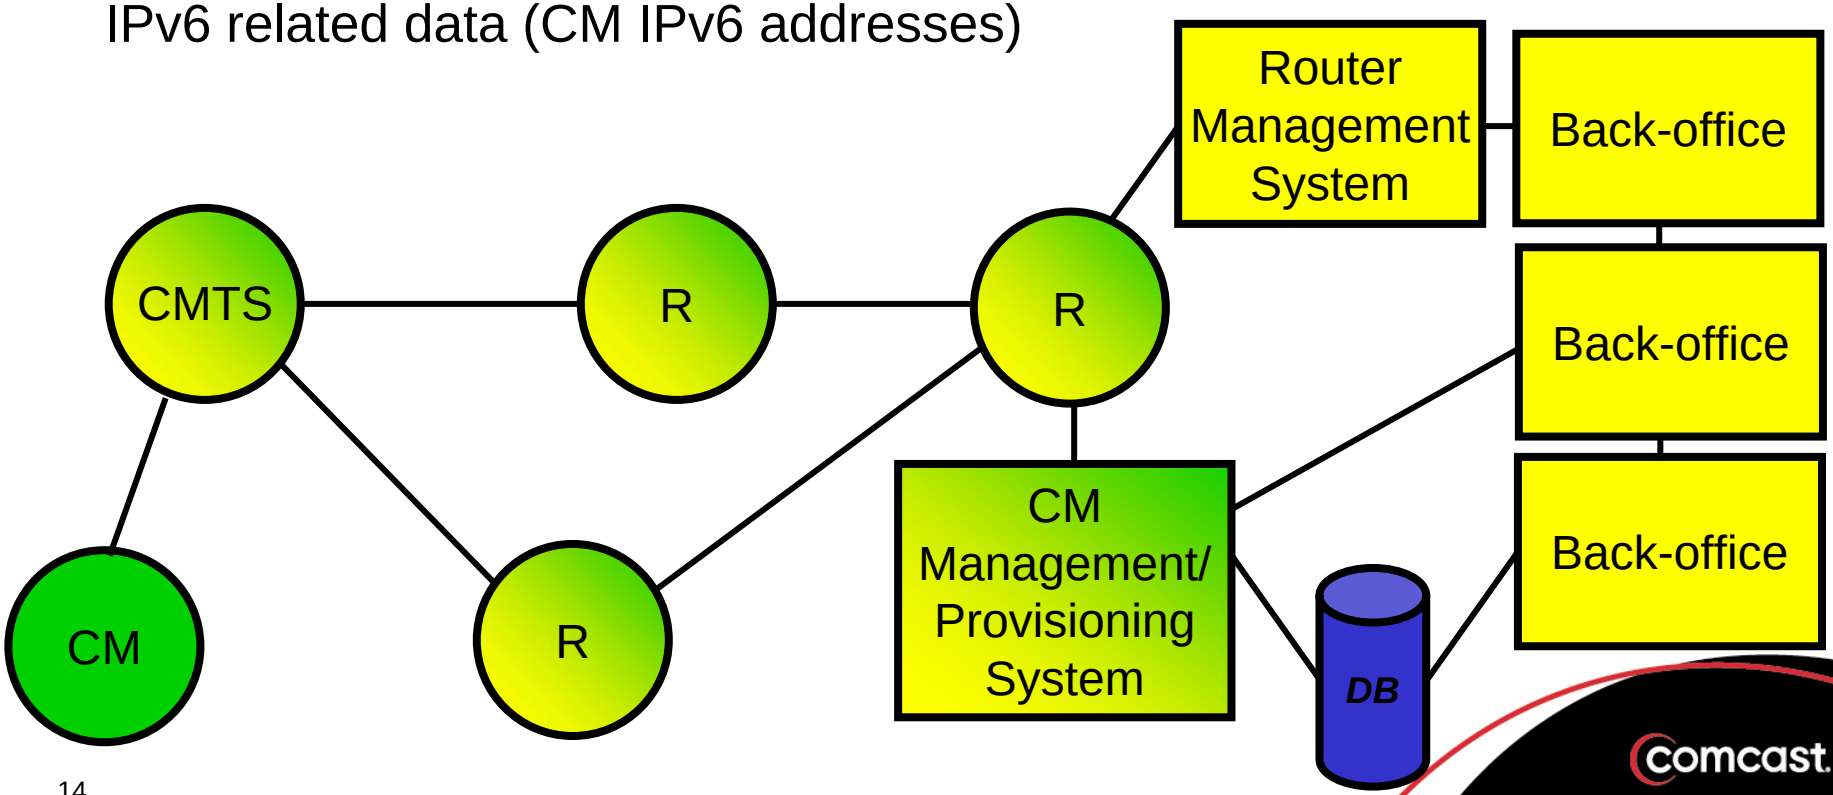

Where We Are

IPv6 Scope & Timeline

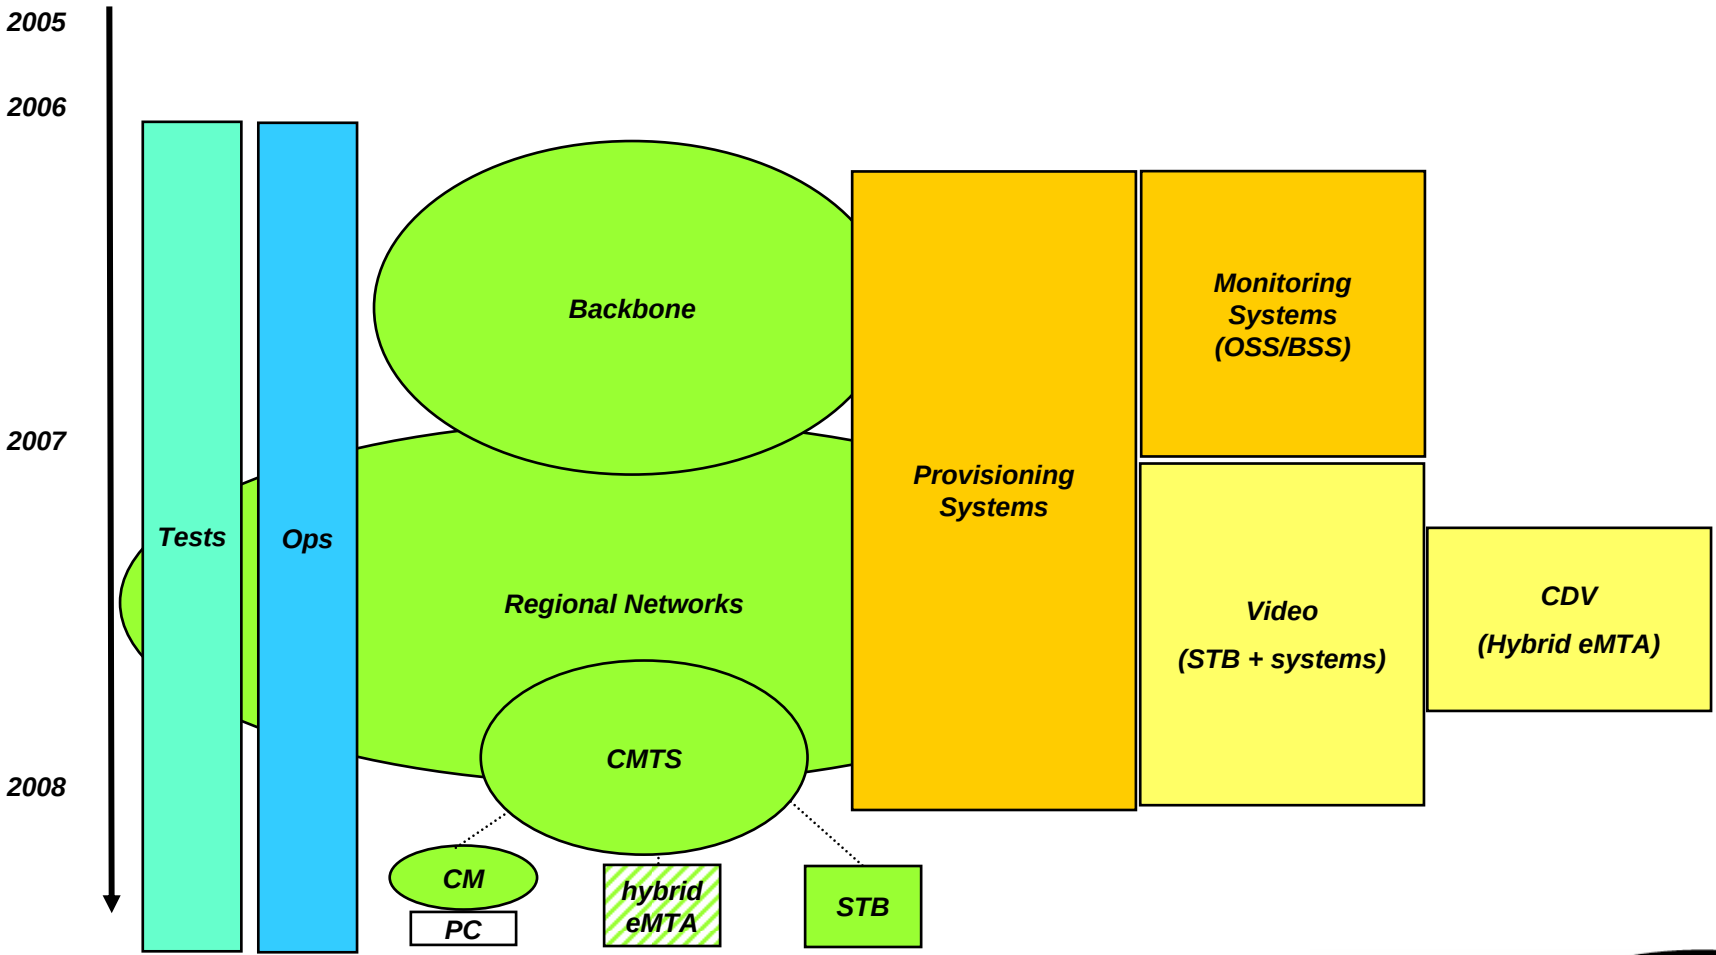

June 2006

All routers on Comcast IP backbones are IPv6 enable

First PING on our 10GE **production** backbone:

```
ping ipv6 2001:558:0:f501::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:558:0:f501::1, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 3/5/14 ms

Jun  2 09:31:49.589, len: 162, hits: 1, i/p i/f: TenGigE0/7/0/0
00146a7d 29810014 6ac4dd08 86dd6000 0000006c 3a3c2001 05580000 f5010000
00000000 00022001 05580000 f5010000 00000000 00018000 50fae0da 00004480
3e53000f 062b0809 0a0b0c0d 0e0f1011 …

Jun  2 09:31:53.533, len: 162, hits: 1, i/p i/f: TenGigE0/0/0/0
00146ac4 dd080014 6a7d2981 86dd6000 0000006c 3a402001 05580000 f5010000
00000000 00012001 05580000 f5010000 00000000 00028100 4ffae0da 00004480
3e53000f 062b0809 0a0b0c0d 0e0f1011 …
```

November 2006

- All routers on Comcast IP backbones route IPv6

Traceroute coast to coast

```
1  2001:558:0:f511::1 24 msec 17 msec 17 msec
2  2001:558:0:f510::1 26 msec 24 msec 24 msec
3  2001:558:0:f50f::1 40 msec 39 msec 39 msec
4  2001:558:0:f50e::1 49 msec 47 msec 47 msec
5  2001:558:0:fe0b::1 73 msec 72 msec 73 msec
```

What We Learned

Docsis 3.0

- Docsis 3.0 standard has been published
- Equipment under tests

Provisioning

- Cable motto: If you can't provision it, don't build it…
- Two components:
  - DHCPv6 server
    - DHCPv6 solutions are getting there
      - Open source code under development (ISC)
      - Successful DHCPv6 bake-off held at RIPE-NCC on 3/14-3/16
  - Back office glue with billing system
    - Where the logic of the system is
    - The hard part to get right

OSS Tools

- Router monitoring tools are in place
- Rest of OSS stack is actively tested
- Will be 100% ready way before roll out

- MIBs
  - Confusion between ‘old’ IPv4 & IPv6 MIBs and ‘combined’ IPv4/IPv6 MIBs:
    - Some only implement the ‘old’ MIBs, some only populate the IPv6 branch of the combined MIB…
    - Very little implementations of the TCP/UDP MIBs

Routers

- Core routers
  - No major issues (we run fairly recent hardware platforms)
  - Routing protocols
    - We chose IS-IS for IPv6 after long discussions
      - Extensions to IS-IS still needed
    - Still running (for now) OSPFv2 for IPv4
- Access routers (CMTS)
  - Code still under development
  - Mileage may vary with vendors

Labs

- We decided NOT to create a big IPv6 lab
  - already too many labs anyway
- We pushed each lab to develop its own IPv6 story
- Part of the “Make IPv6 part of Comcast DNA” effort

Operations

- If you can't operate it, don't deploy it…
  - Getting early buy-in from Ops people is important
  - Need to demonstrate that "IPv6 is not the end of the world as we know it"

- We chose the "baby steps" approach
  - Deploy one thing at a time
    - Apply appropriate training
    - Start with little, non critical traffic
    - Show nothing breaks; loop

Training

- Training is key for Ops people
  - Traditional “3-5 days generic IPv6 off-the-shelf training” is not adapted to our population
  - Early information to combat FUD factor
  - Just-in-time training before roll-out

- Tons of slides exists on the web
  - Quality is not always there
  - Usually very generic content
  - Focus mostly on Layer 3 issues
    - but the problems are elsewhere…

Future Headache: SIP / IMS / PacketCable 2.0

- Inter-networking IPv4-only and IPv6-only devices on a SIP network is a difficult thing
- 3 translations need to happen:
  - SIP message
  - SIP SDP headers
  - Data

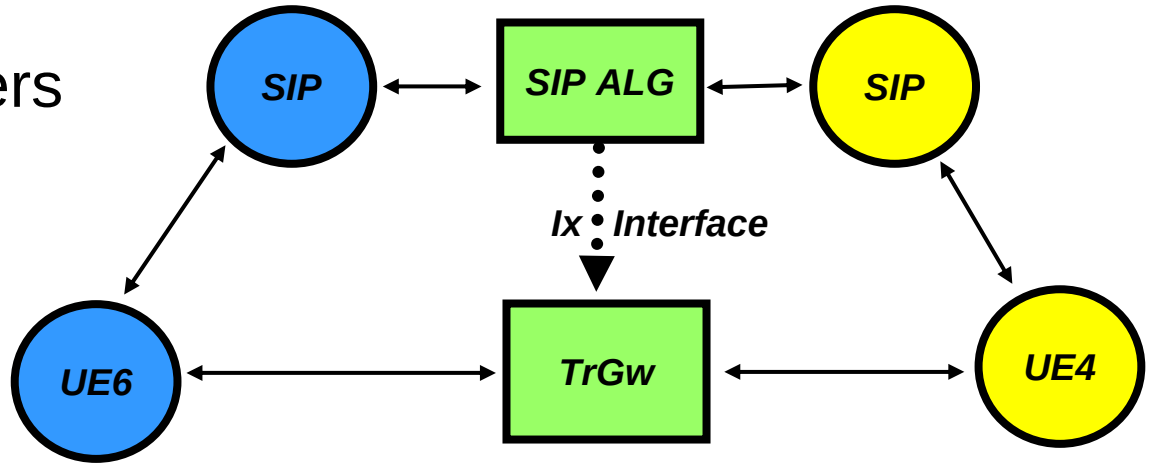

- Specs are not ready
- **Sounds difficult to engineer to scale**

Conclusion

- 96 more bits, no magic
- Deployment problems are not at layer 3 but at layer 7 & 8
- Engaging vendors early helps
- Training & early Operation engagement are important steps