Flattening the Network

Intro to BoF Session in Denver, June

Neil Farquharson

Birds of a Feather in Denver

Please sign up for our

BoF Session

At NewNOG52, June

Data traffic predictions

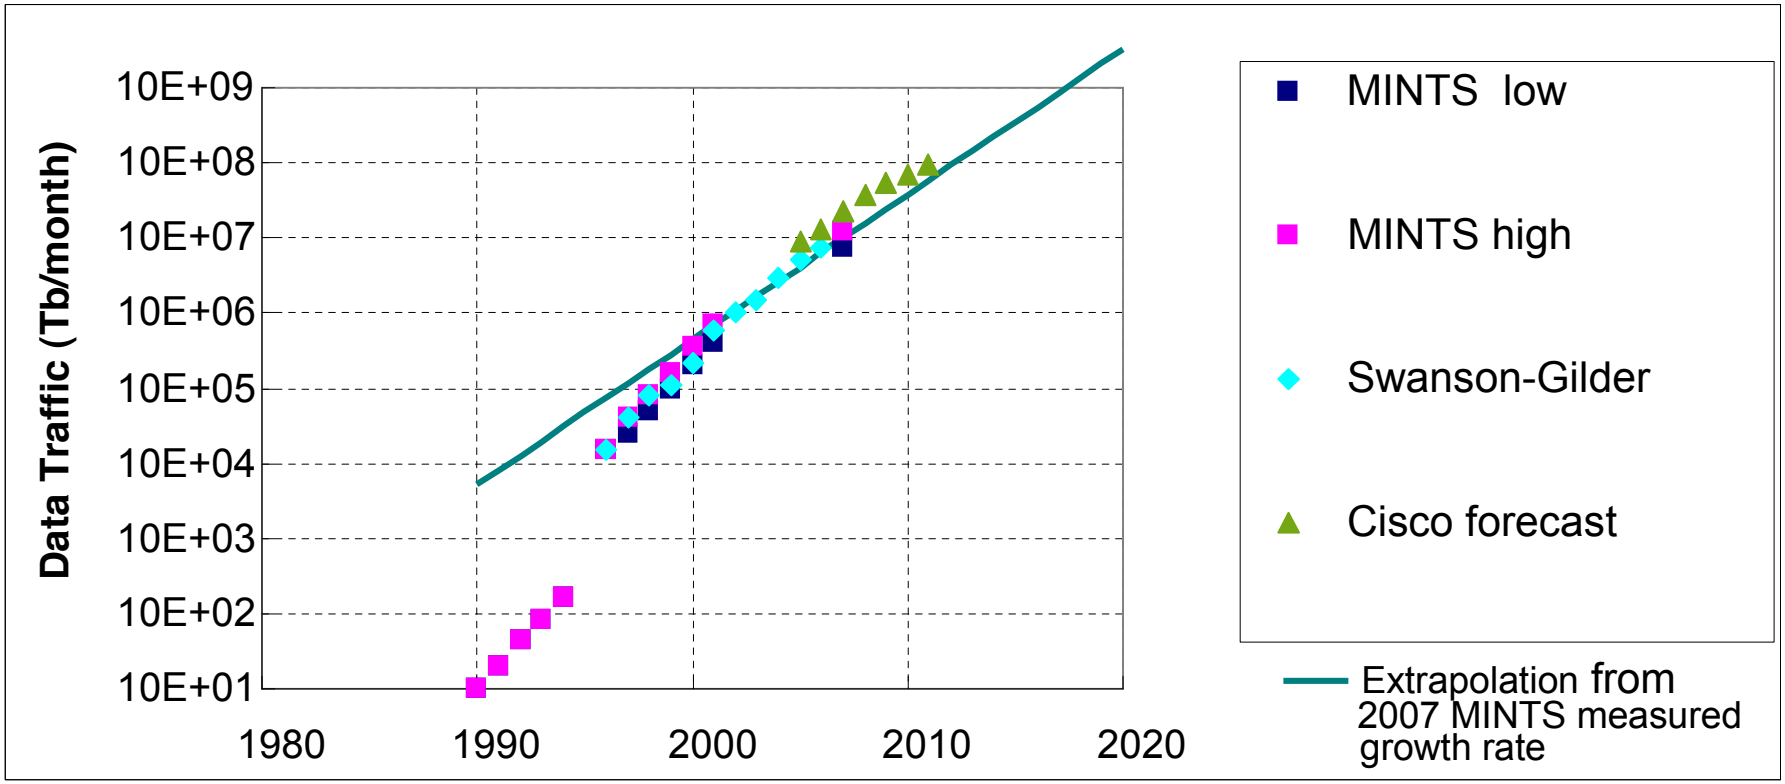

MINTS – **Minnesota Internet Traffic Study, http://www.dtc.umn.edu/mints/home.php**

Swanson-Gilder – ***Estimating the Exaflood*, 2008.**

Cisco Forecast – ***Global IP Forecast and Methodology*, 2007.**

Traffic Predictions

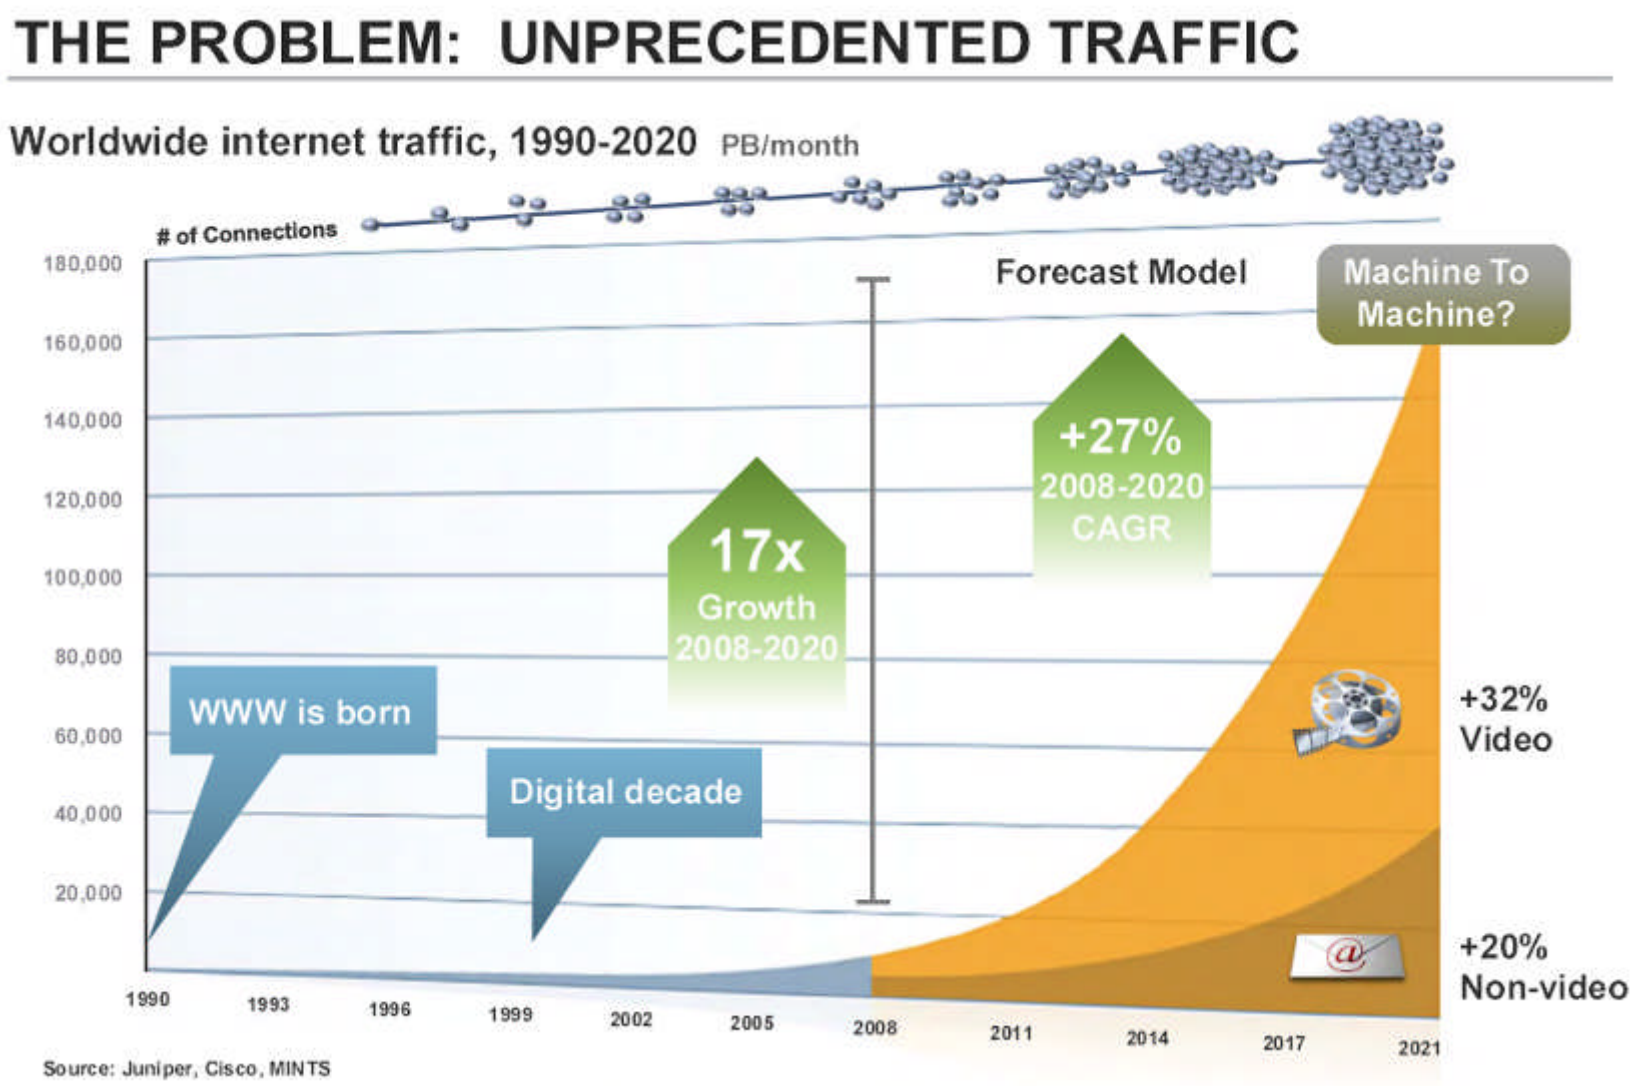

Alcatel·Lucent


“Star Trek” copyright of Paramount

Drivers for change

- Traffic growth leads to more core routers being deployed, as all traffic passes through a meshed router core
- Traditional delineation of Transport and IP networks results in limited operational visibility across domains, thus the cost of delivering services is not optimized
- Core architectures are increasingly expensive and power hungry
- Carriers looking for higher speeds (40G, 100G) and dramatic cost efficiencies

From

- High volumes of video & internet traffic negatively impacting profitability
- Growing meshes of core routers to handle ever-increasing traffic
- Separate IP & Transport domains with minimal cross-domain operational visibility

To

- A new architecture that can accommodate profitable growth for internet and content distribution
- Leveraging the efficiencies of integrating flexible OTN and of powerful IP services
- Integration of the Optical domain into the IP domain with cross-Layer visibility and automation

OTN is the key convergence technology uniting current technologies with robust OAM and FEC for reliable transport that is very high bandwidth at low cost/bit

Switching at Lambda or Sub-lambda (Port, Sub-port) Level

- **Complete transparency** for clients in flexible containers
- **ODU container sizes** well-suited for Ethernet transport

- **ODU Termination** (G.709 OAM) guarantees a clear demarcation
- **Intermediate Monitoring** can be electronic and/or photonic

- **HO-ODU** networking used when client throughput does not need further aggregation within a λ
- **LO-ODU** networking used when sub-λ multiplexing is needed (no stranded BW)

- **Switching** can be electronic and/or photonic

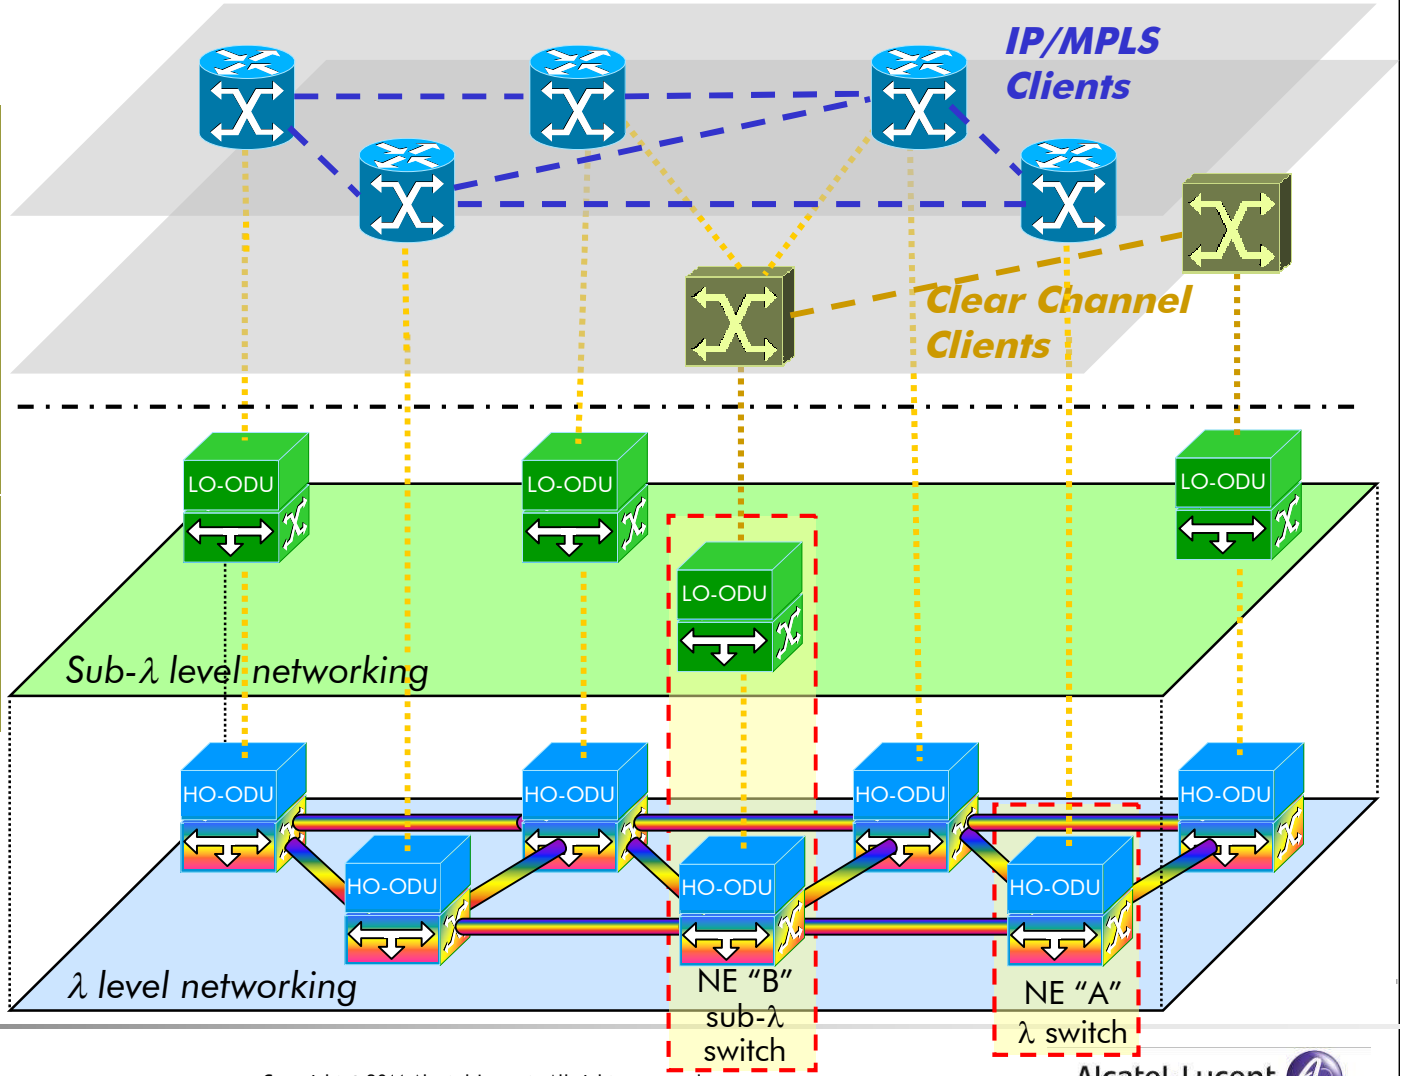

ODU(flex) switch

ROADM

Alcatel·Lucent

More flexibility for OTN

ODUflex – The Flexible Data Unit

- Provides for connection-oriented transport of any variable bit-rate, leveraging its **flexible-sized** container (1.25 Gbps increments)
- Decouples client physical interface from the network transport entity
- Maintains a 1:1 relation between client interfaces and transport entities, reducing the # of transport entities and simplifying operations
- Supports transport of variable rate packet streams (for instance VLANs within a 100GE port)
- Similar to VCAT, but avoids differential delay problem of VCAT
- Enables full utilization of network resources, **avoiding capacity waste**

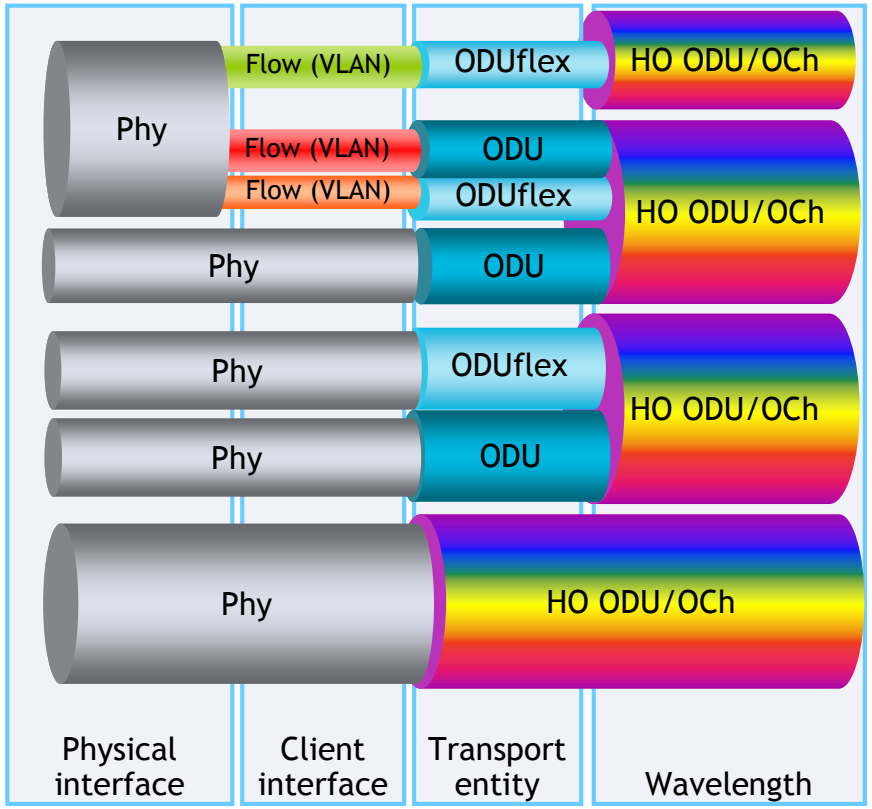

ODUflex maps any bit-rate client to a single, right-sized transport container

More flexibility for OTN
Flexible grooming + OTN

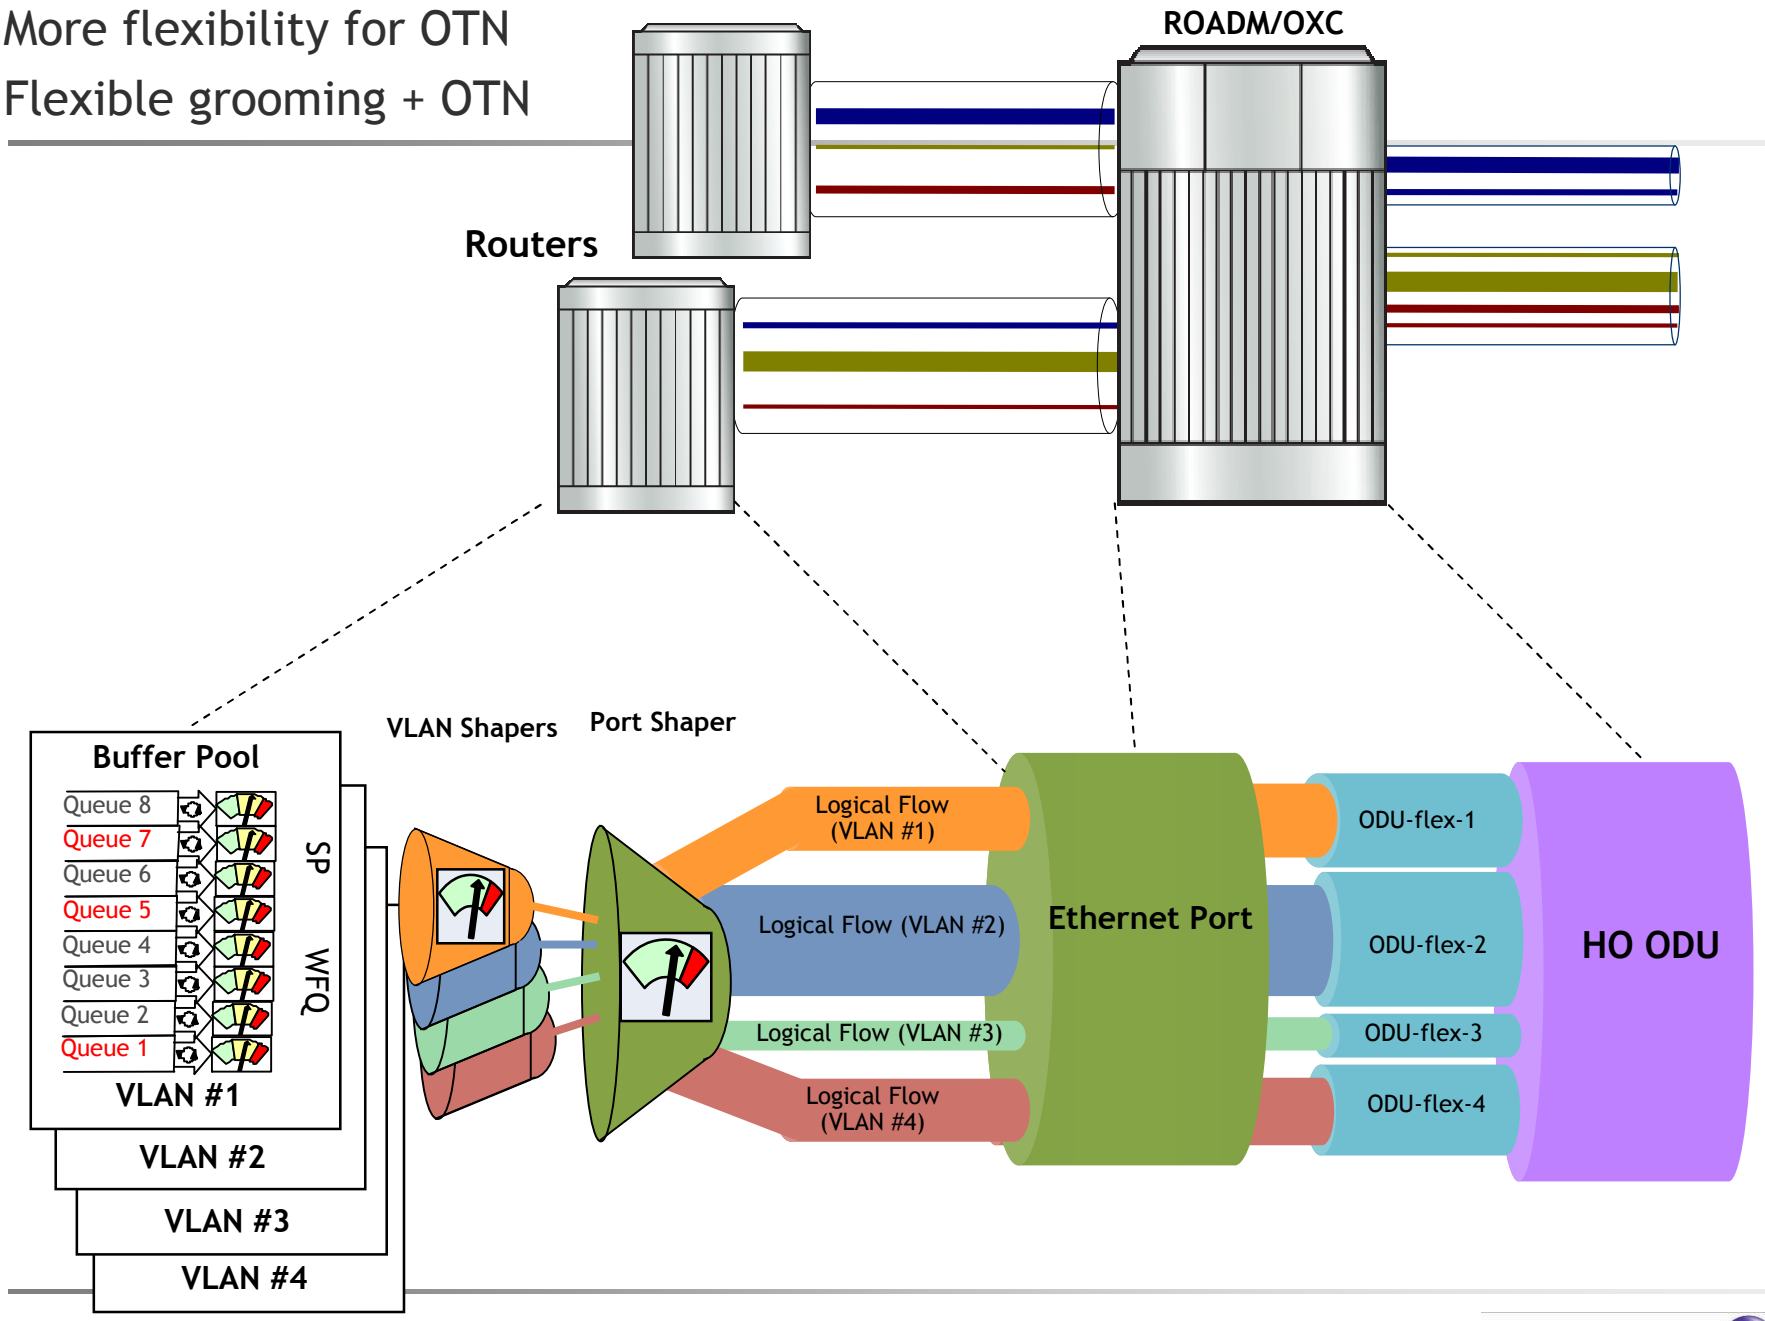

Before Flattening - Present Mode of Operation

Separate and independent IP and optical domains

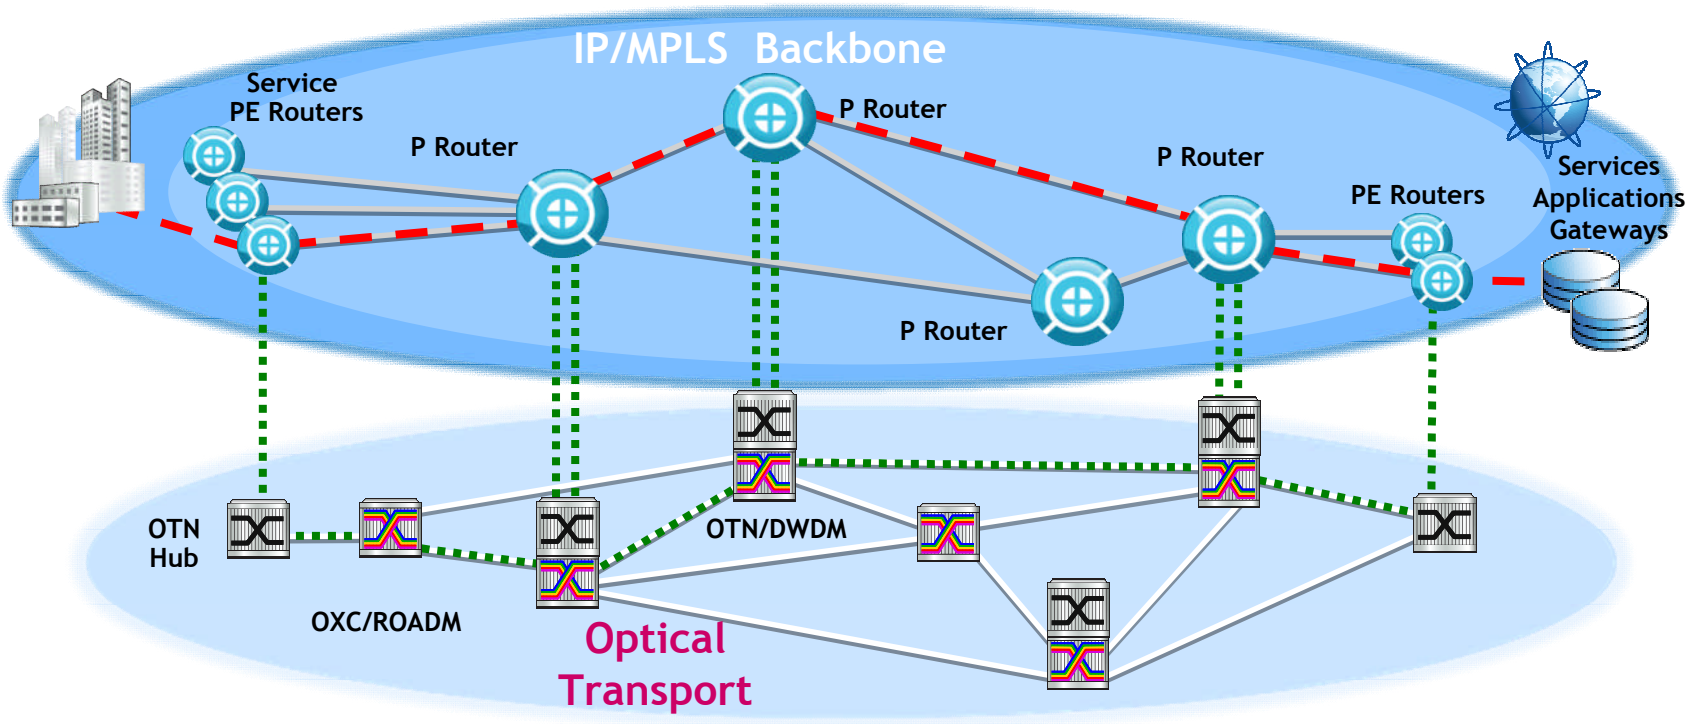

After Flattening, using a unified control plane

Mix of technologies to give best overall performance

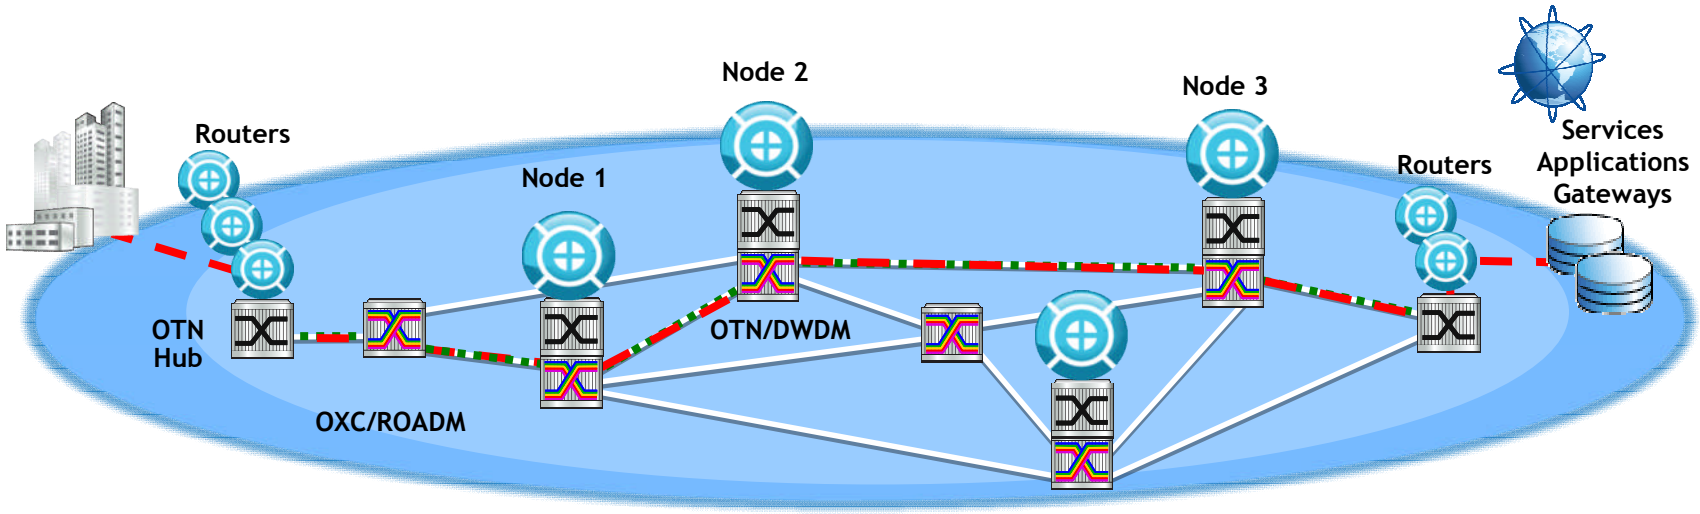

On <u>average</u> equates to:

- lower latency
- higher bandwidth
- lower cost per bit

After Flattening, using a unified control plane

Example of one traffic stream

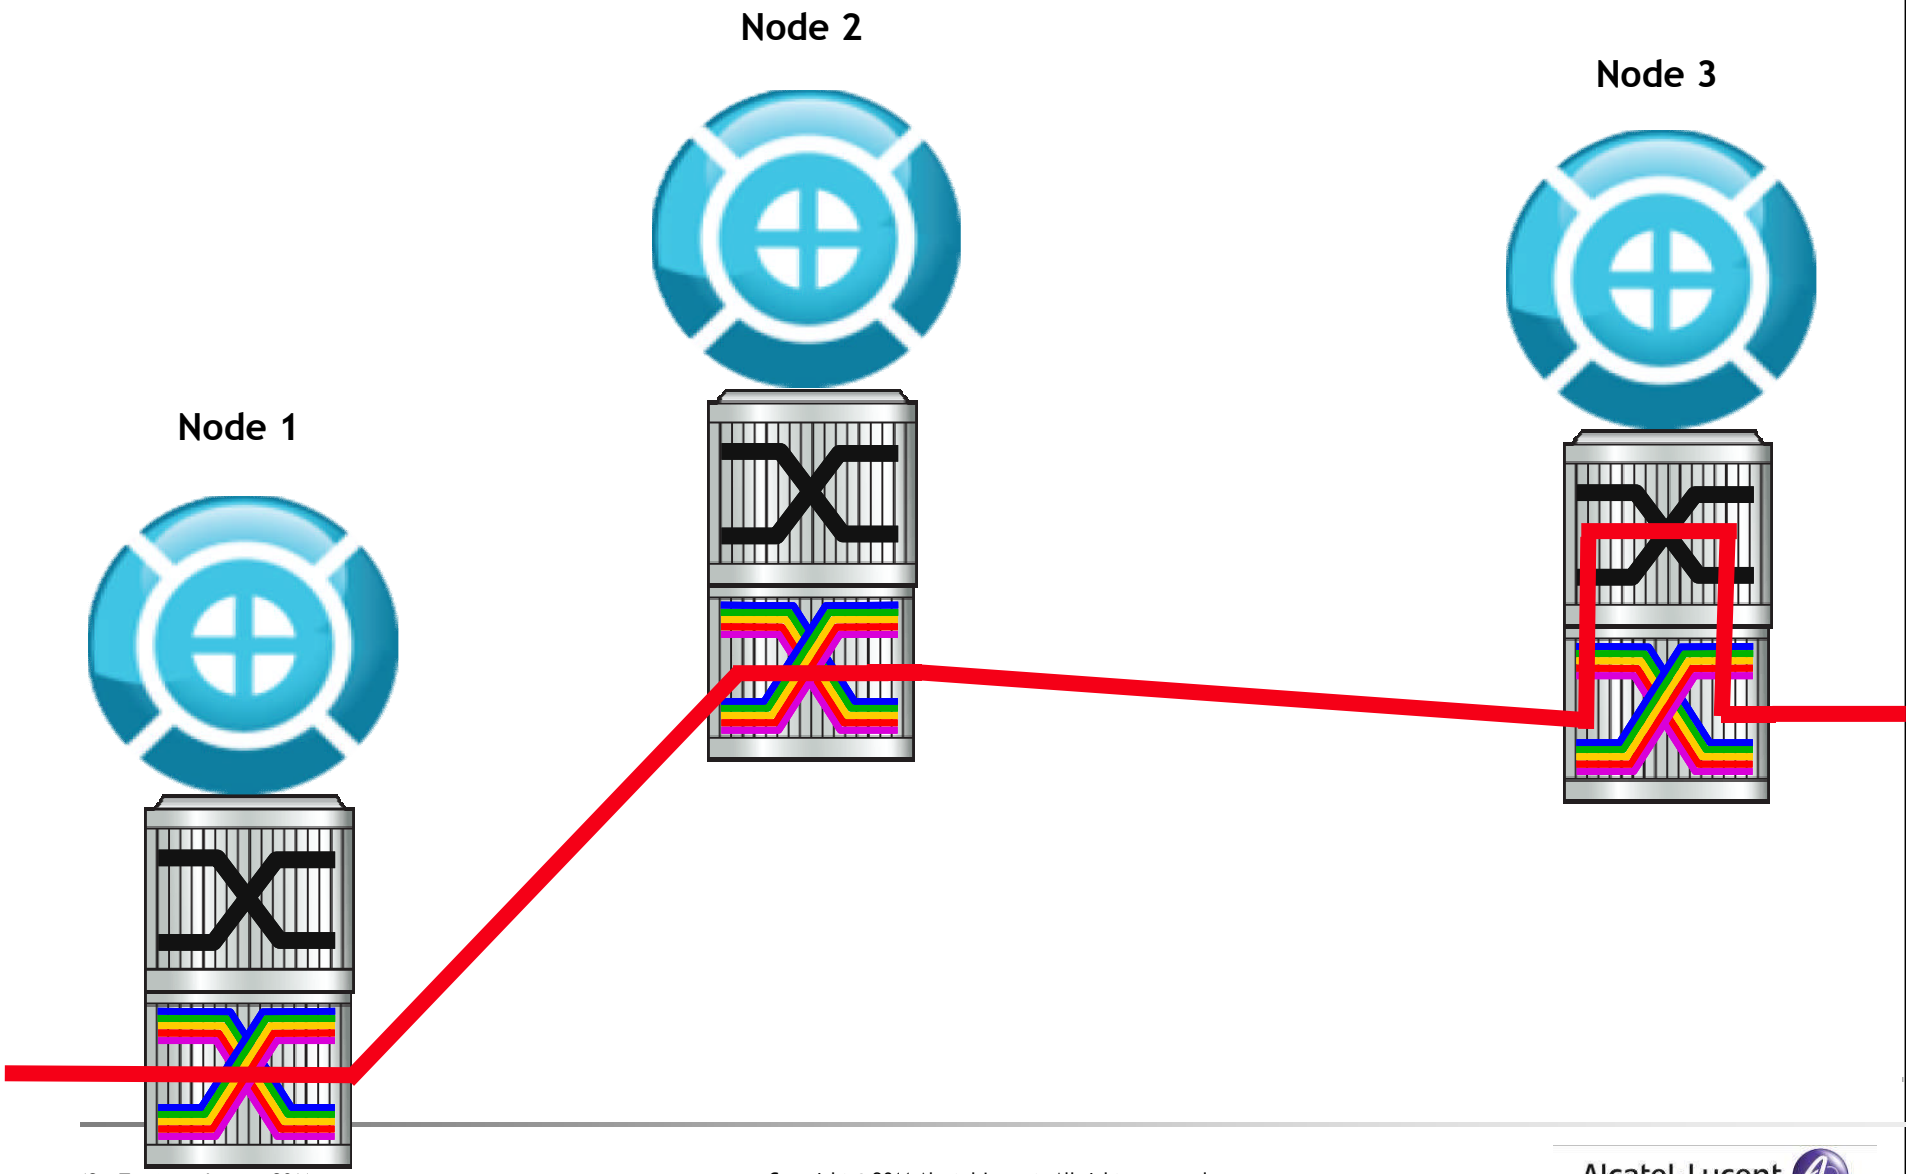

After Flattening, using a unified control plane

Example of another traffic stream

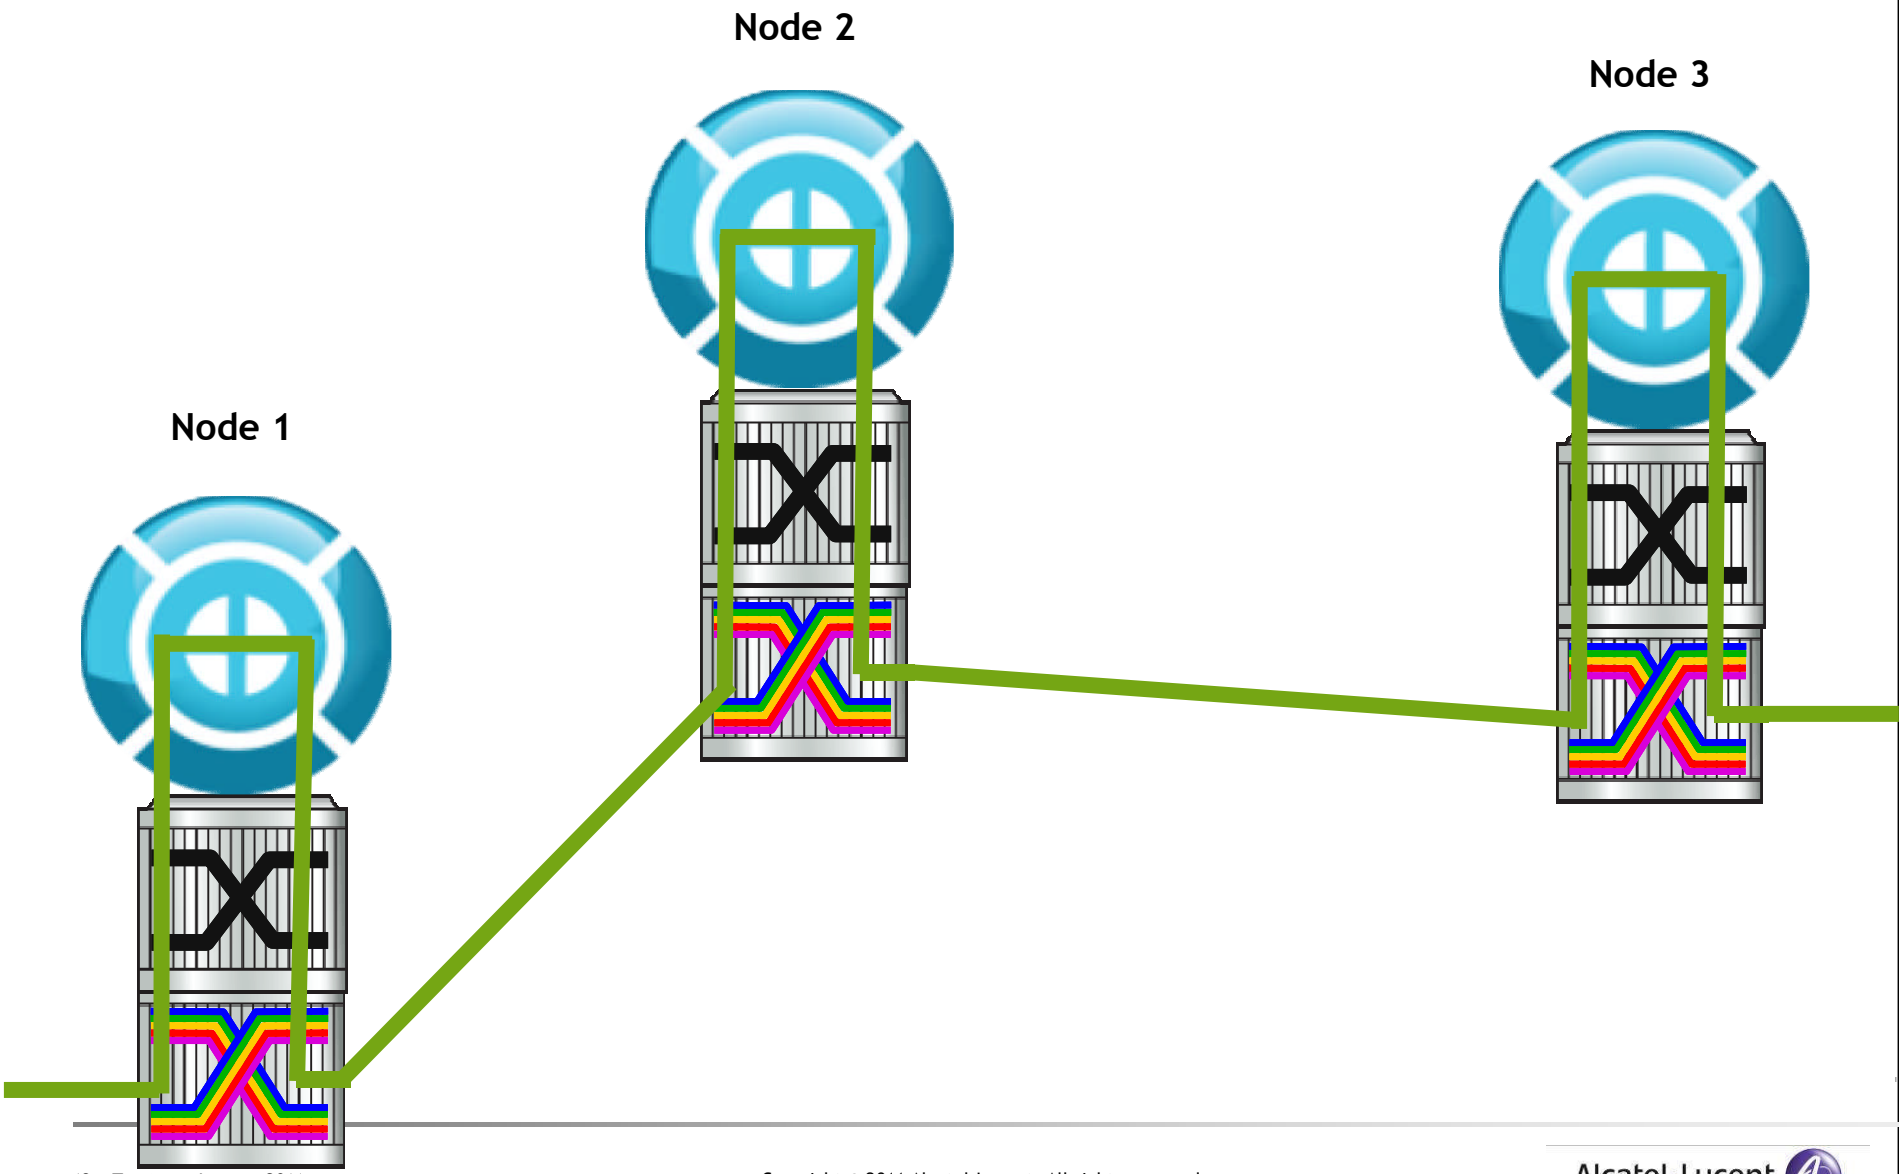

After Flattening, using a unified control plane

Example of a third traffic stream

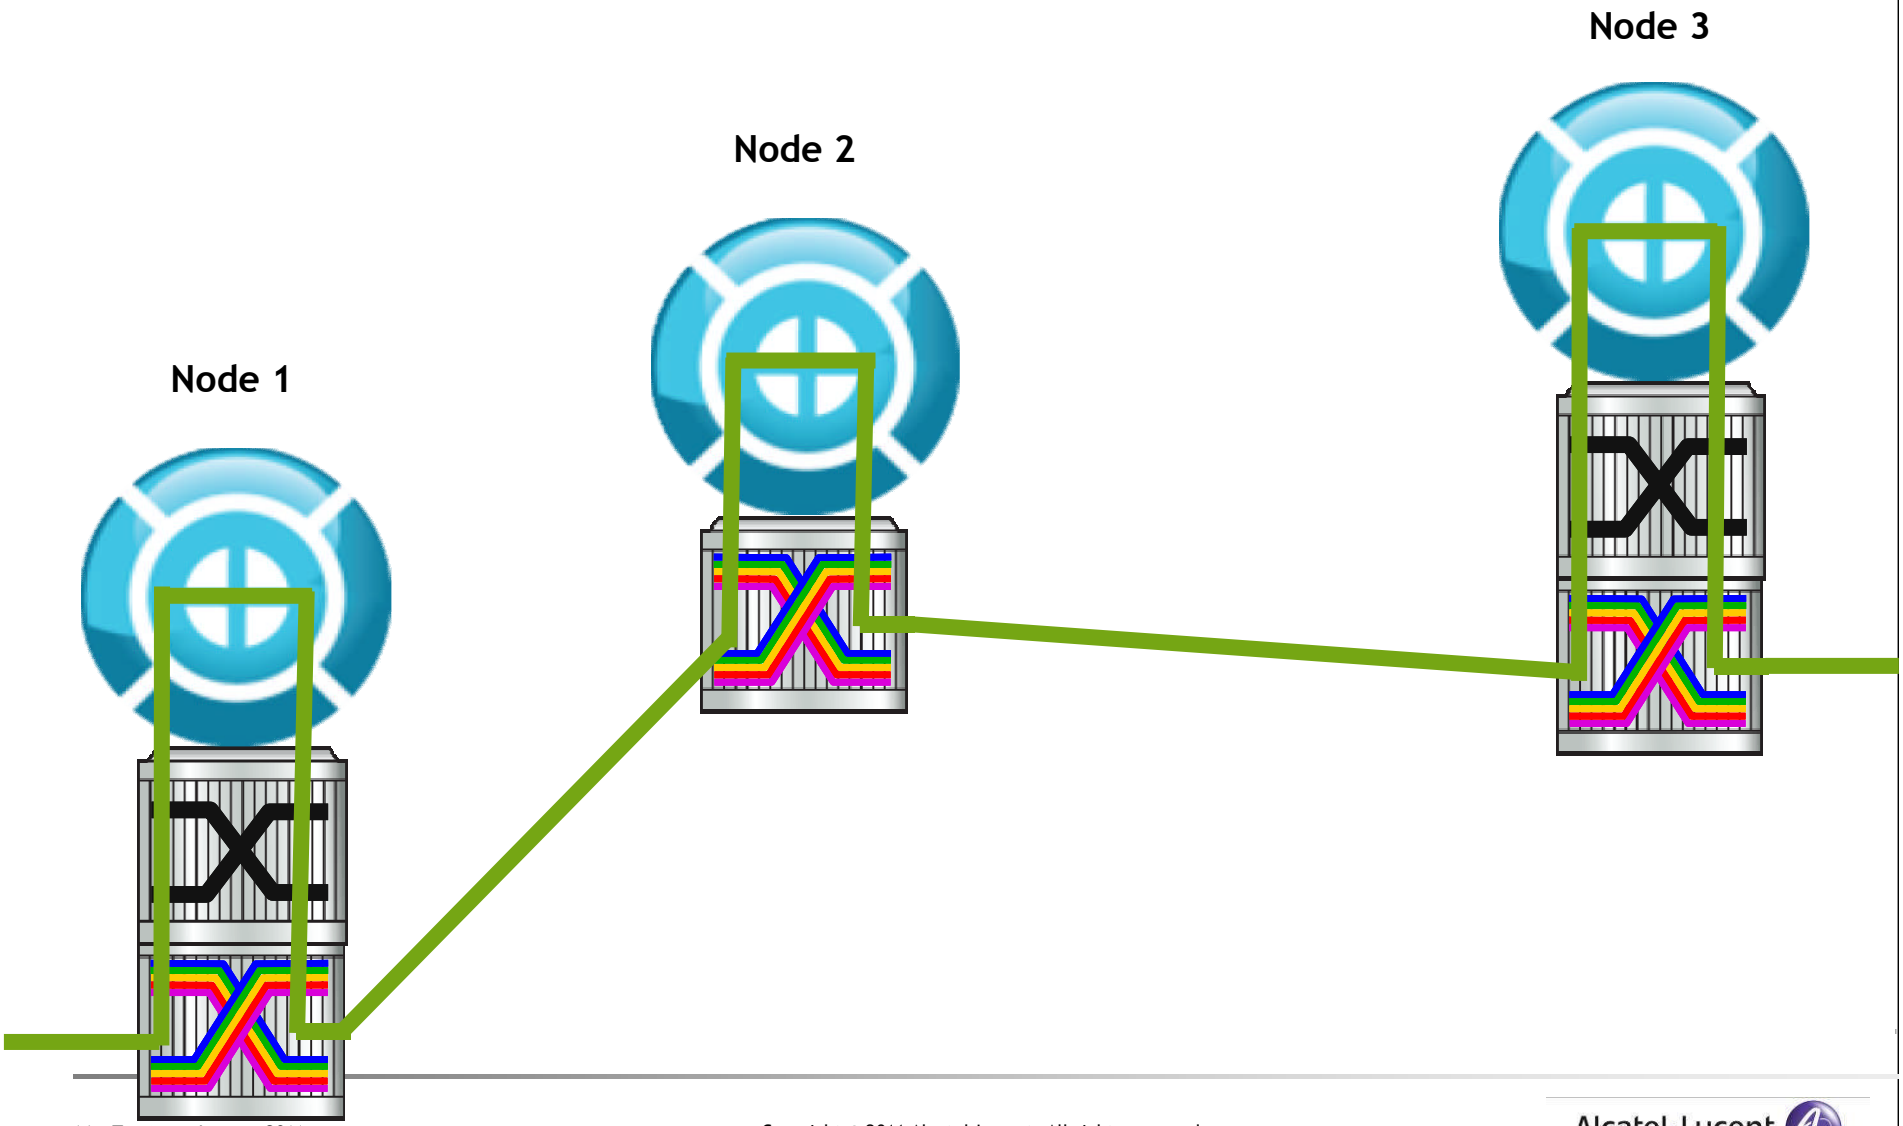

Engrained beliefs to be overcome

- Belief #1: Network Flattening (NF) should be done at WDM layer, not ODU layer (via IPoDWDM with tunables on routers) to reduce complexity and the # of OXCs.
- Belief #2: You don't need to protect both IP and optical layers – optical protection works for link failures, not router failures. Emphasis should be on IP layer protection.
- Belief #3: As you make an optical device more dynamic, it tends towards becoming a router; with comparable capex, complexity and opex challenges.
- Belief #4: Network Flattening requires more router resources for protection purposes. Also, you need > 100% spare to account for potential types of failures.
- Belief #5: Router-OTN (mapping onto ODU layer) handoff is expensive and complex; and unlike the optical layer, MPLS LSPs can grow and shrink as traffic dictates.
- Belief #6: The adage of "switch where you can, route where you must" is still pertinent: NF leads to much worse scalability problems since meshing of PE-routers becomes an $O(N^2)$ problem rather than O(N). Also, P-routers are able to stat-mux traffic better than PE-routers. Consequently the reliance on PE-routers as aggregators requires overprovisioning of the PE-routers.

Birds of a Feather in Denver

- Utilize P-Routers for what they do best
- Optimize resource sharing and distribution
- Increase network capacity fast
- Will save a darn lot of money – once we get it to work

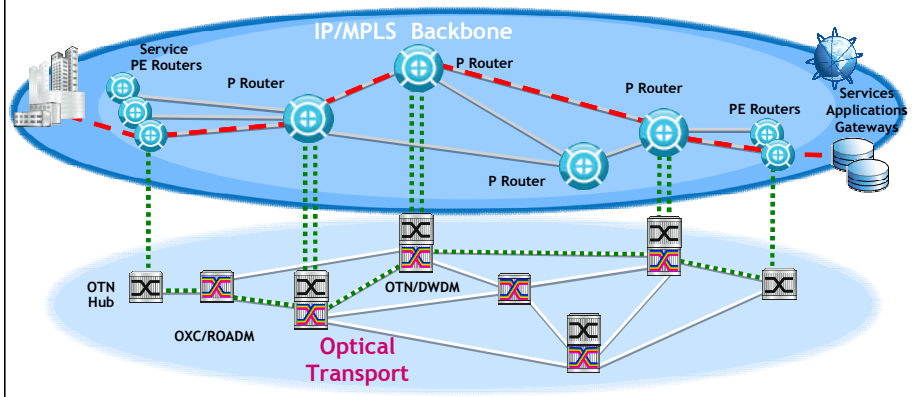

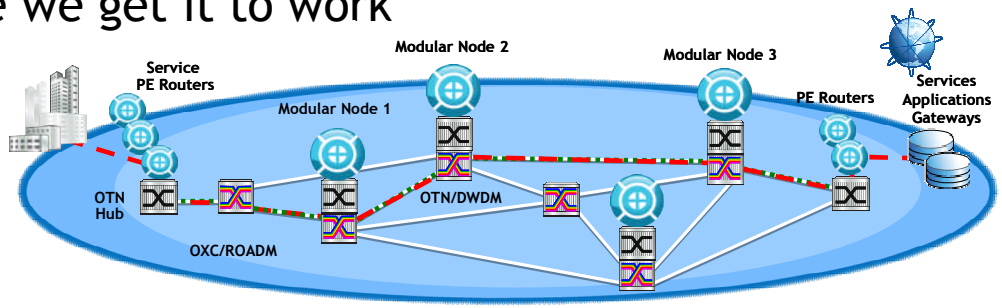

Please join us for our BoF Session, NewNOG52, June

Alcatel·Lucent

Replies to Slide 15. Not included in origin slide deck.

Belief 1: There are situations where IPoWDM (lambda grooming) makes sense, especially for hub & spoke flows. The tradeoffs are (1) the provider needs integrated IP/optics operations groups, (2) providers who need a clear demarc between IP and optical domains prefer black & white interfaces from the router, (3) IPoWDM is typically limited to metro and regional networks, since router wavelengths put into the DWDMs must conform to rigorous engineering rules, (4) the wavelengths may be underutilized. Depends upon the design.

Belief 2: I think that possibly this claim assumes that optical protection would be slow since it's at lambda layers. If so this is a fallacy because protection is much faster the further down the Protocol stack one goes. However, I do believe two protections may be the way to go – it is something we should discuss in June.

Belief 3: Increased complexity, Capex and Opex yes. But NOT comparable complexity, CAPEX and OPEX. If one looks at the total network (IP+optical) cost, complexity and Opex, the benefits of dynamic control at the optical layer outweigh these added costs. More so if we are able to move to the unified control system we're discussing here.

Belief 4: Again, this is not true if one properly engineers the resiliency mechanisms and coordinate between layers to avoid overlaps and conflicts. A main benefit of mesh topology is to provide resiliency without 100% redundancy. And if some % of traffic bypasses a core router, that traffic does not even need to be protected by that core router.

Belief 5: This only holds true if one assumes a fixed 1-1 mapping between MPLS LSPs, VLANs, router ports and TDM containers. However if one includes per-VLAN shaping and the flexibility of ODU-flex's granular pipe sizes (multiples of 1.25 Gbps) then the handoff is not expensive. With the solution I'm talking about we can put multiple VLANs into a single port and ensure QoS.

Belief 6: The proposed optical paths between PE-routers do create new adjacencies that impact PE routing tables and – potentially - convergence time. This requires the PE-router to expand routing tables, while maintaining convergence time. Further, these paths are selectively created by us, avoiding a full PE-router mesh and the resulting O(N2) scalability problem. Also, my understanding is that P-routers do not stat-mux traffic better than PE-routers, particularly when the PE-routers are coupled with intelligent optical switches to support flexible grooming.