ON.LAB

SDN Control of Packet-over-Optical Networks

Marc De Leenheer — ON.Lab
Guru Parulkar — Stanford University & ON.Lab
Tom Tofigh — AT&T

Acknowledgements

People and companies that were instrumental

- Bob Doverspike, Tom Anschutz, Martin Birk, Mehran Esfandiari (AT&T)
- Prajakta Joshi (ON.Lab)
- Praseed Balakrishnan (Fujitsu), Eiichi Kabaya (NEC), Patrick Liu, Wei Wei, Hongtao Yin (Huawei)
- Saurav Das, et. al., "Packet and circuit network convergence with OpenFlow," OFC 2010, OTuG1 (ONF)
- Infoblox

Outline

- SDN for Service Providers
  - Background
  - Use cases
- Packet/Optical Use Case
  - Problem statement and conceptual solution
  - Implementation
  - Demonstration
- State of the Industry & Future Work

Explosive Growth

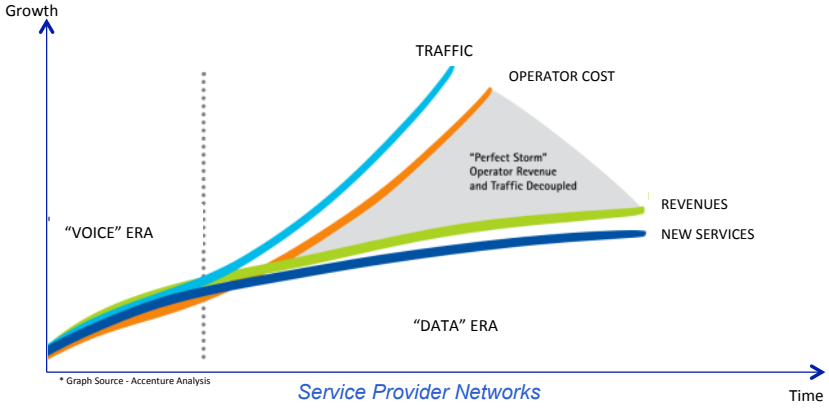

Unprecedented Traffic Growth	Orders of magnitude increase in users, devices, apps	Video, Mobile traffic exploding	CAPEX continues to rise
2016 traffic = triple of 2011	More mobile devices than people	IP Video: 79% of all IP traffic in 2018	AT&T spends $20 Billion per year on CAPEX

Turning Growth into Opportunity

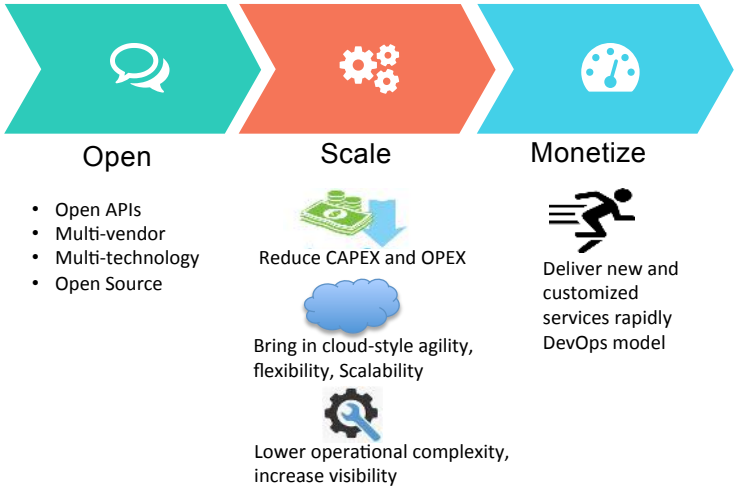

Key Enabler: Software Defined Networking

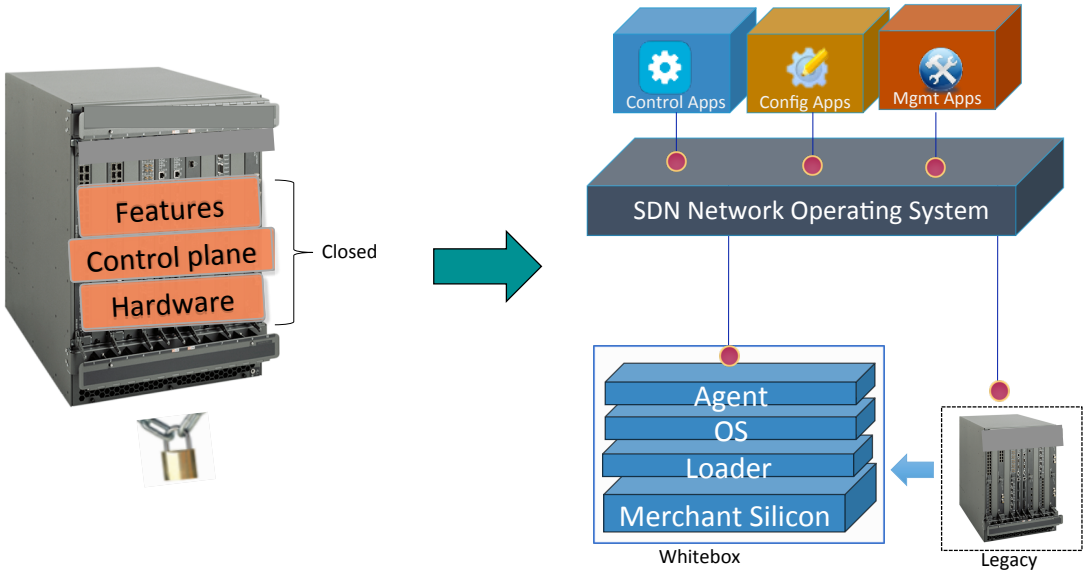

Service Provider Networks

- WAN core backbone
  - Multi-Protocol Label Switching (MPLS) with Traffic Engineering (TE)
  - *200-500 routers, 5-10K ports*
- Metro Networks
  - Metro cores for access networks
  - *10-50K routers, 2-3M ports*
- Cellular Access Networks
  - LTE for a metro area
  - *20-100K devices, 100K-100M ports*
- Wired access / aggregation
  - Access network for homes; DSL/Cable
  - *10-50K devices, 100K-1M ports*

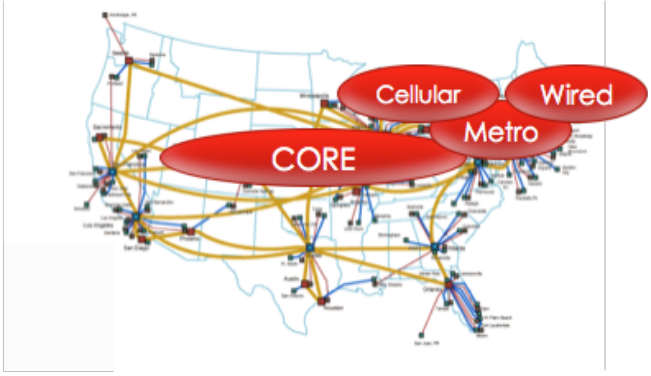

Service Provider Network of the Future

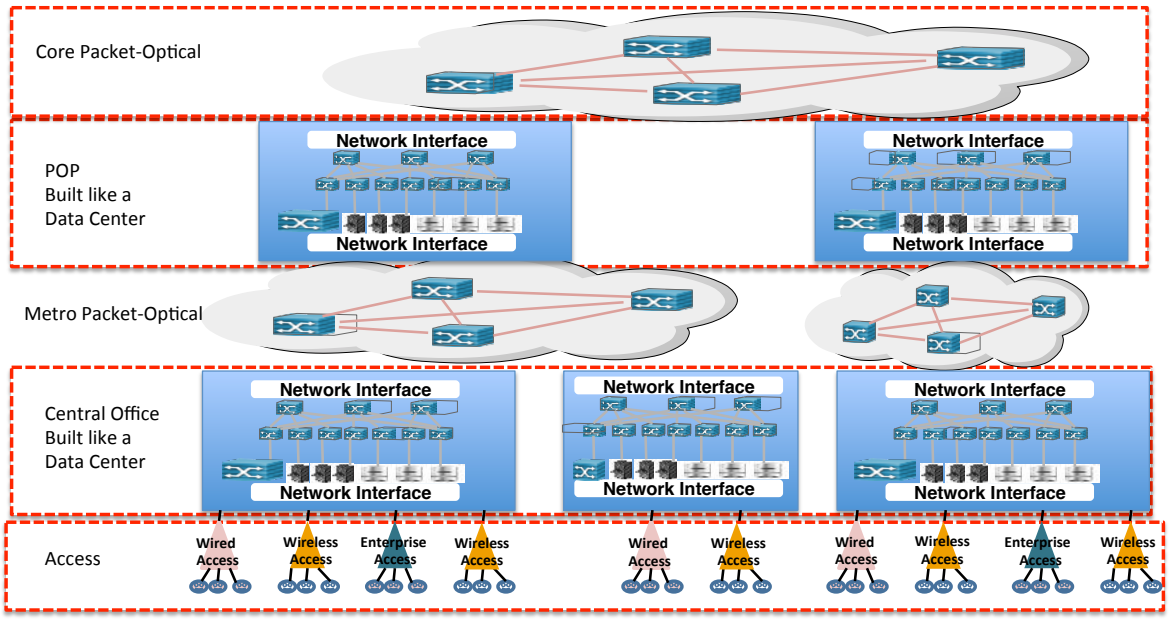

SDN Control Plane: Key Performance Requirements

- High Throughput:
  - ~500K-1M paths setups / second
  - ~3-6M network state ops / second
- High Volume:
  - ~500GB-1TB of network state data
- Difficult challenge!

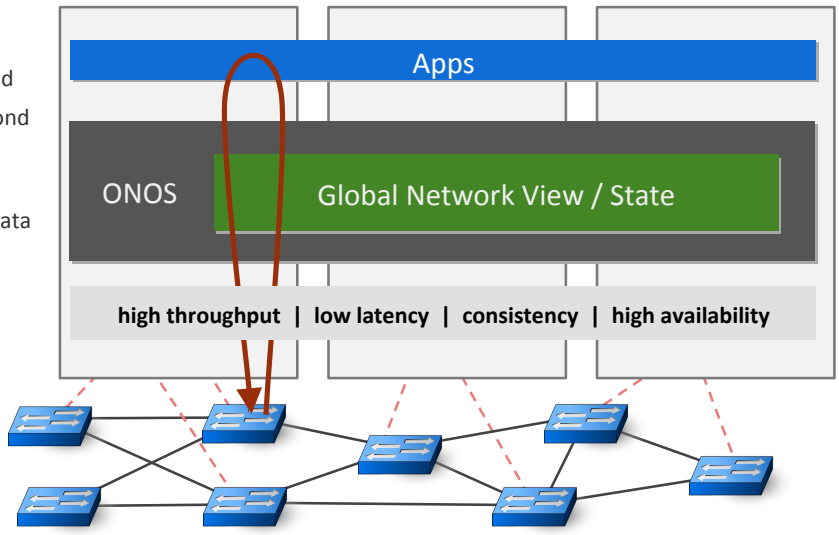

ONOS: SDN Network OS for Service Providers

- Scalability, High Availability & Performance
- Northbound & Southbound Abstractions
- Modularity

ONOS: Distributed Network OS

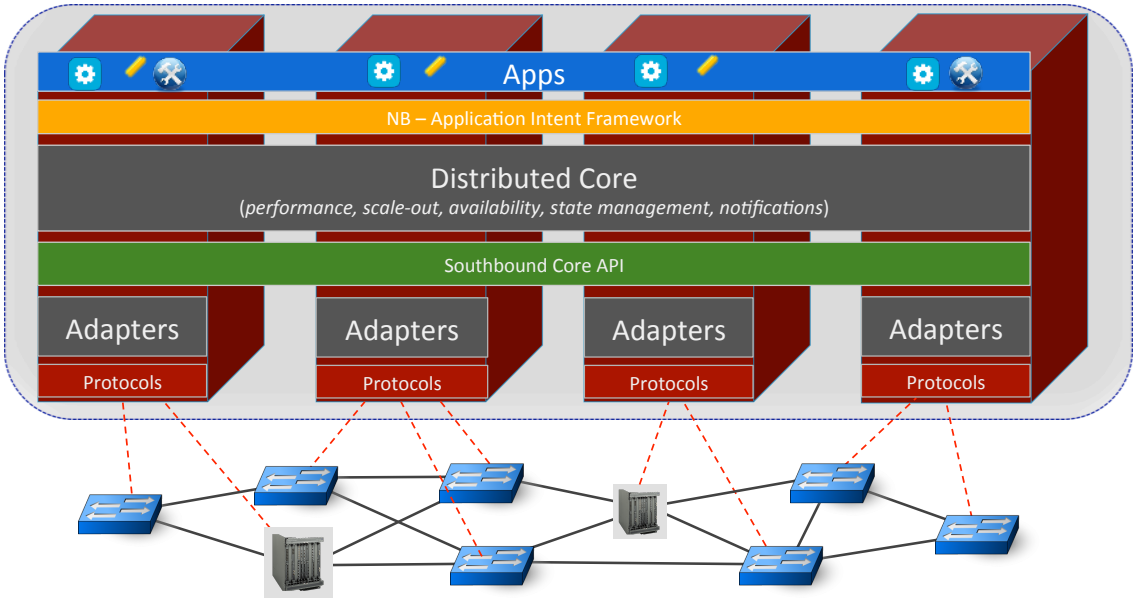

Application Intent Framework

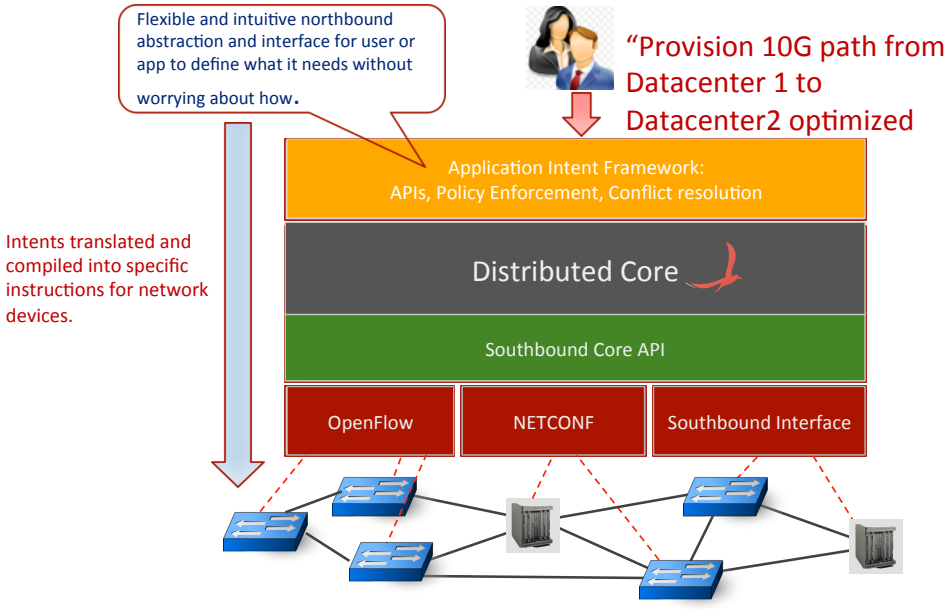

SDN Use Cases for Service Providers

- Converged multi-layer packet/optical networks
- Seamless SDN and IP peering with SDN-IP
- Segment routing with SDN control
- Central Office Re-architected as a Datacenter (CORD)
  - Network Functions as a Service (NFaaS)
  - vCPE and vOLT
- And many more
  - Mobile backhaul (IP RAN)
  - IP multicast
  - …

Seamless Peering: SDN-IP

SDN-IP enables communication between:

- **SDN network and external IP networks**
- **external networks across SDN island**

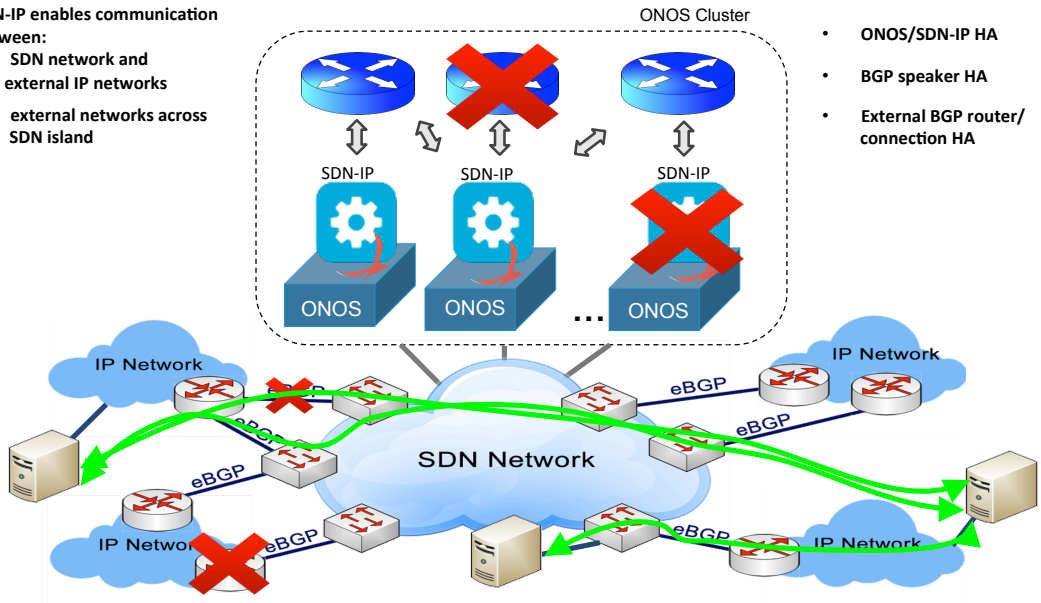

- **ONOS/SDN-IP HA**
- **BGP speaker HA**
- **External BGP router/ connection HA**

SDN-IP Deployments

Internet 2: Provide L3 connectivity between 6 universities around US

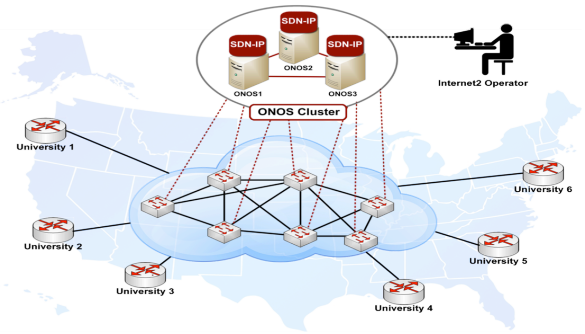

Peering Router on Corsa hardware ~14,000 routes

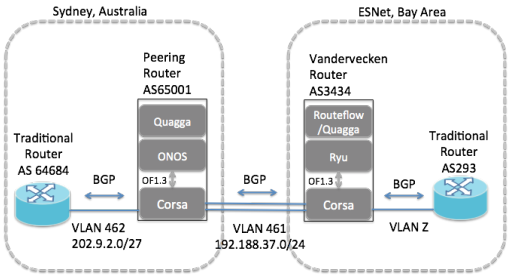

Seamless peering of SDN islands with existing networks
=
Migration strategy for real networks

Segment Routing

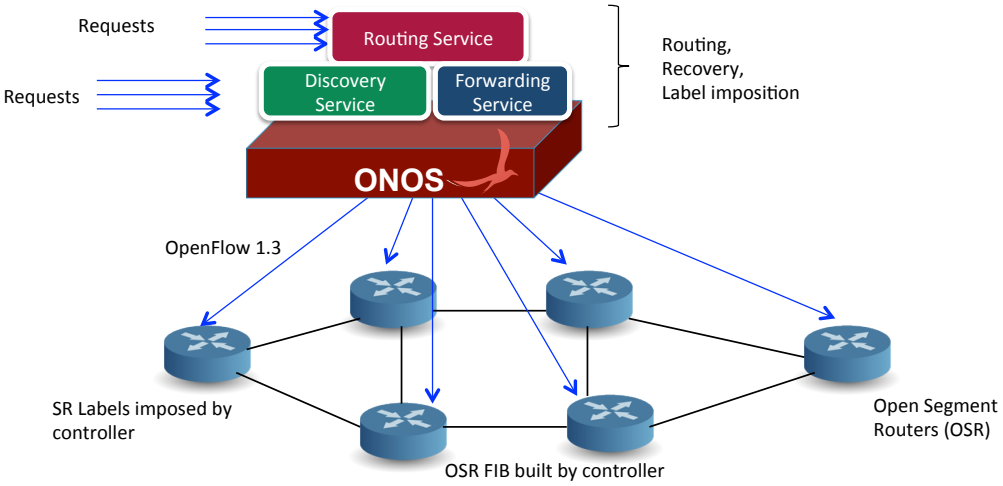

Outline

- SDN for Service Providers
  - Background
  - Use cases
- Packet/Optical Use Case
  - Problem statement and conceptual solution
  - Implementation
  - Demonstration
- State of the Industry & Future Work

Problem Statement

- Today IP packet and transport networks are separate.
- They are planned, designed and operated separately by different teams.
- This leads to significant inefficiencies.
- They are subject to under-utilized networks with significant pre-planning and highly over-provisioned for worst case.
- A lot of the path planning in these networks is off-line.

> Given these considerations, **WAN links are typically provisioned to 30-40% average utilization**. This allows the network service provider to mask virtually all link or router failures from clients. Such overprovisioning delivers admirable reliability at the very real costs of **2-3x bandwidth over-provisioning** and high-end routing gear.
>
> *S. Jain, et. al., "B4: Experience with a Globally-Deployed Software Defined WAN," SIGCOMM 2013.*

Multi-Layer Network without Converged Control Plane

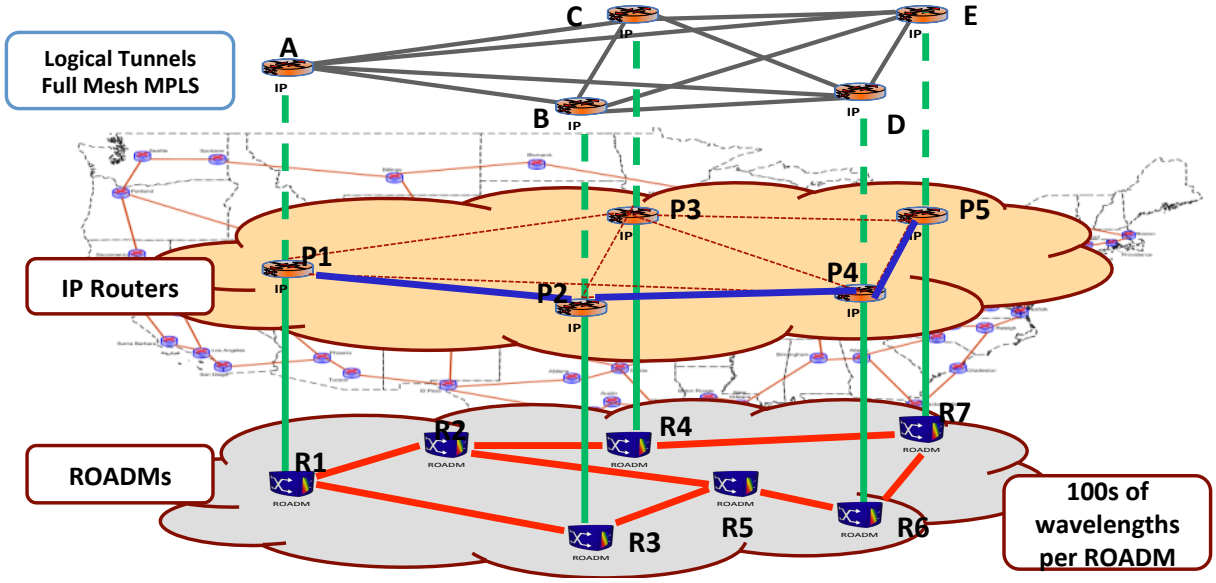

Multi-Layer Network without Converged Control Plane

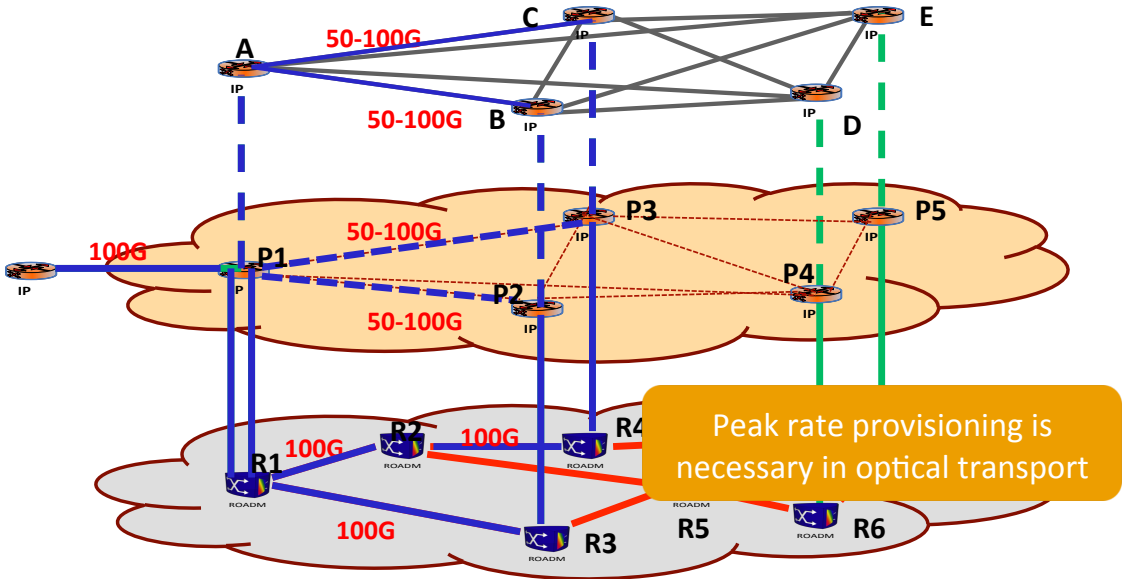

Multi-Layer Network without Converged Control Plane

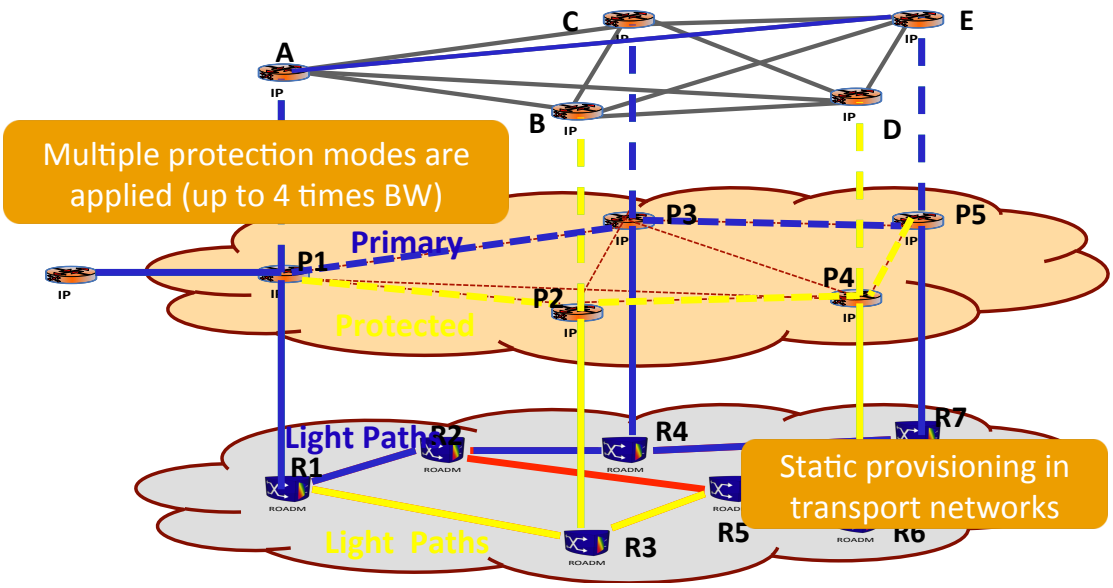

Conceptual Solution: Multi-Layer SDN Control

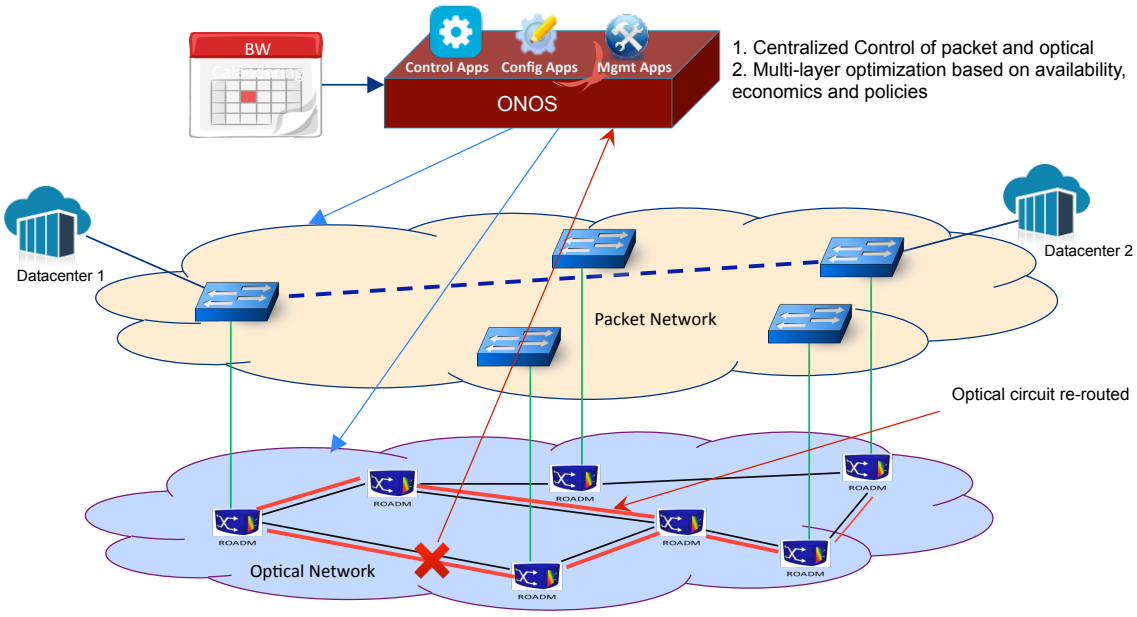

Benefits of Converged Control Plane

- Much faster bandwidth provisioning
- Drastically improve network utilization
- Perform dynamic restorations in response to packet and transport network failures
- Agile development and rapid deployment of new services

Implementation

- Code is king
  - Less is more
  - Vendor neutral
  - Scalability, high availability, performance
- Work focused on the three SDN layers
  - Data plane
  - Control plane
  - Applications

Implementation – Data Plane

Packet Switches

- Open and standardized interface to forwarding plane?
- Reality
  - OpenFlow
  - Available today in many products

ROADMs

- Open and standardized interface to forwarding plane?
- Reality
  - Legacy protocols such as TL1
  - Vendor specific

Built an optical emulation platform LINC-OE with our partner Infoblox

https://github.com/FlowForwarding/LINC-Switch

ROADM Emulation Basics

- Emulates optical layer topology from predefined table
- Includes characteristics of optical cross connect and Packet to Optical Link Interface (Add/Drop)
- Ports, links and switches are remotely reconfigurable by Mininet
- Supports OpenFlow 1.3+ Optical Add/Drop match actions
- Supports failure scenarios of links, ports, and ROADM
- Work in progress
  - Emulates channel signal/power measurement
  - Regenerator support

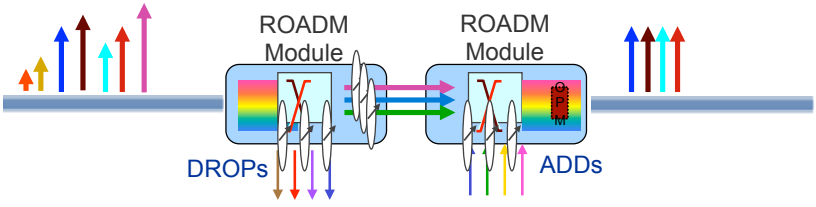

Forwarding Model for ROADMs

- Match/action abstraction for ROADMs
  - ROADM has three functions: add, drop, and forward
  - Match is really about wavelength provisioning

Add	
Match	Packet port and traffic type
Action	Transponder uses lambda and output to optical port

Forwarding	
Match	Optical port and lambda
Action	Output to optical port (easy to extend when considering regenerators)

Drop	
Match	Optical port and lambda
Action	Transponder uses lambda and output to packet port

Forwarding Model for Packet and ROADM Layer

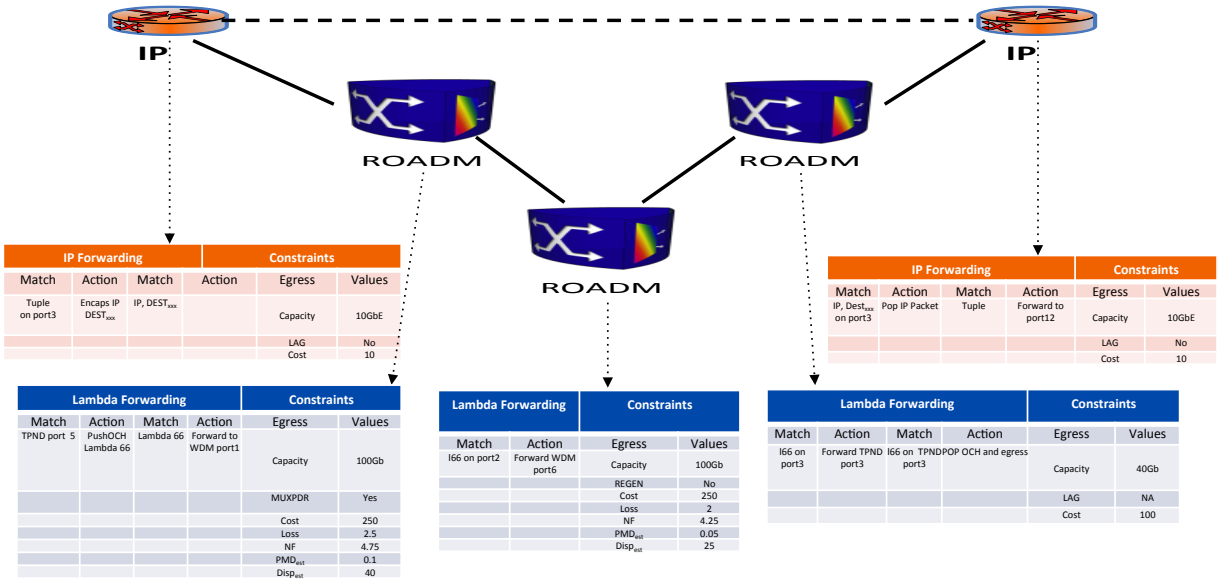

IP Forwarding				Constraints	
Match	Action	Match	Action	Egress	Values
Tuple on port3	Encaps IP $DEST_{xxx}$	IP, $DEST_{xxx}$		Capacity	10GbE
				LAG	No
				Cost	10

Lambda Forwarding				Constraints	
Match	Action	Match	Action	Egress	Values
TPND port 5	PushOCH Lambda 66	Lambda 66	Forward to WDM port1	Capacity	100Gb
				MUXPDR	Yes
				Cost	250
				Loss	2.5
				NF	4.75
				PMD_{est}	0.1
				$Disp_{est}$	40

Lambda Forwarding		Constraints	
Match	Action	Egress	Values
I66 on port2	Forward WDM port6	Capacity	100Gb
		REGEN	No
		Cost	250
		Loss	2
		NF	4.25
		PMD_{est}	0.05
		$Disp_{est}$	25

IP Forwarding				Constraints	
Match	Action	Match	Action	Egress	Values
IP, $Dest_{xxx}$ on port3	Pop IP Packet	Tuple	Forward to port12	Capacity	10GbE
				LAG	No
				Cost	10

Lambda Forwarding				Constraints	
Match	Action	Match	Action	Egress	Values
I66 on port3	Forward TPND port3	I66 on port3	TPNDPOP OCH and egress	Capacity	40Gb
				LAG	NA
				Cost	100

Transport Network Metering Model

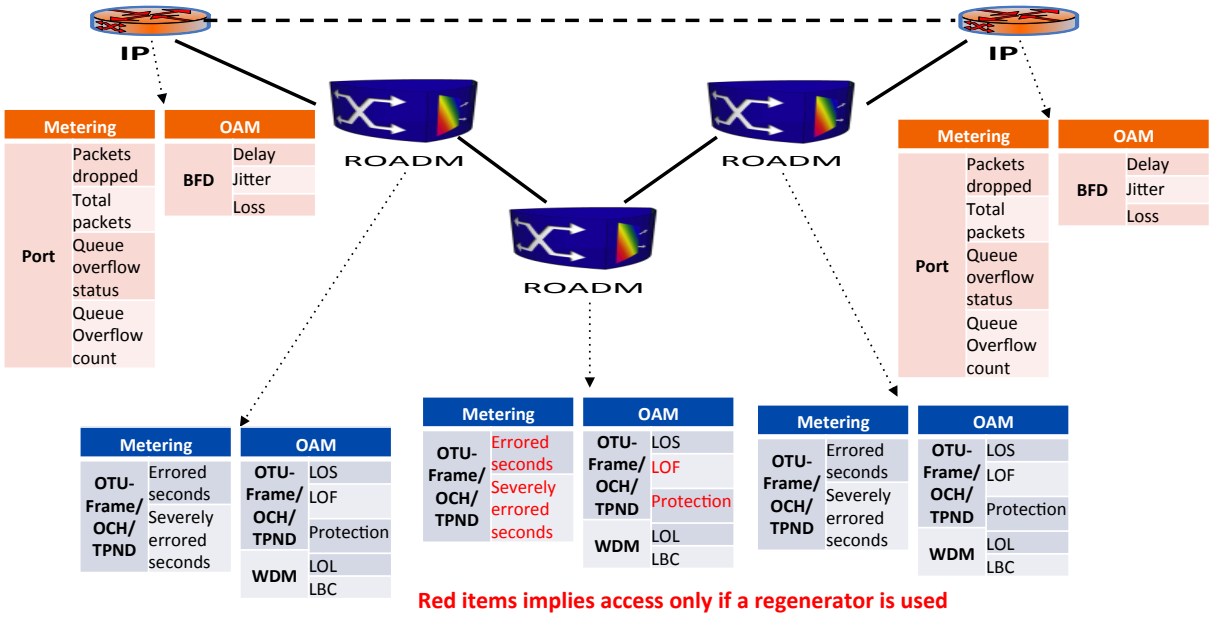

Red items implies access only if a regenerator is used

Implementation – Control Plane

- Southbound protocol for ROADMs
  - ONF Optical Transport Working Group
  - OpenFlow 1.3+ experimenter messages
  - Southbound abstractions simplify adding new protocols
- Discovery
  - Automatic L3 topology discovery (LLDP)
  - Static configuration of L0 topology
  - L0 discovery work in progress

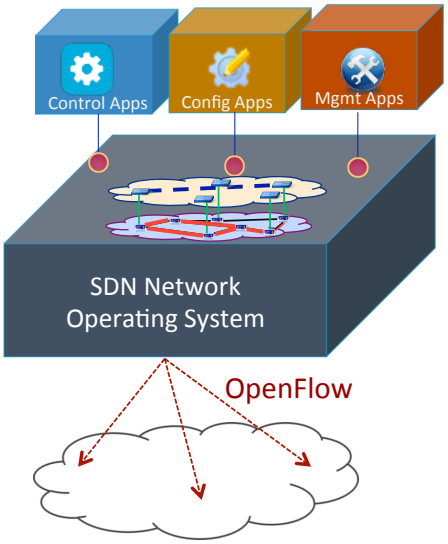

Implementation – Control Plane

- Converged topology
  - Control both packet and optical layers
  - Allows adding additional layers, e.g., OTN
- Path calculation takes place on the multi-layer graph
  - Constraints and resource management
  - Wavelength continuity, bandwidth, latency, …
- Packet and optical layer restoration
  - First try packet layer, then optical layer
  - Easily add multi-layer protection and restoration mechanisms

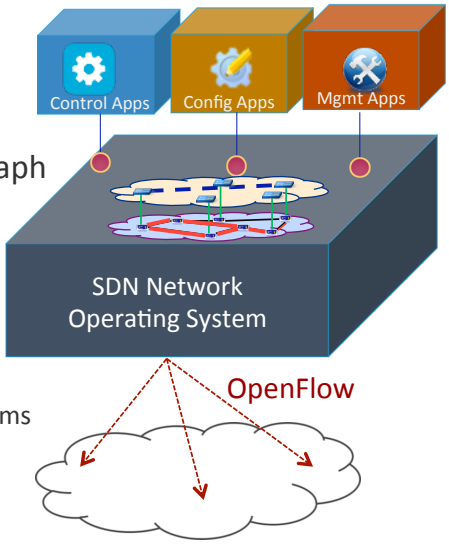

Implementation – Applications

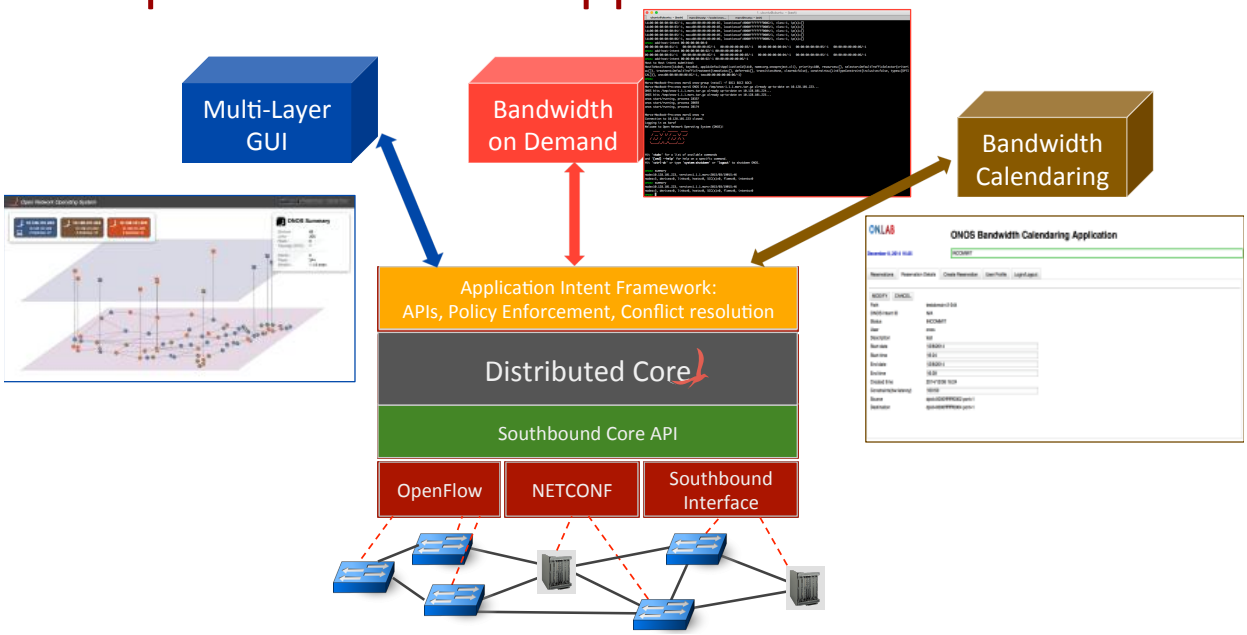

Multi-Layer Network Reference Platform

- "Packet/Optical Network in a Box"
- Open Source components
  - Data plane
    - LINC-OE & OVS
    - Mininet
  - ONOS
  - Pre-packaged apps
- Benefits
  - Rapid prototyping, agile
  - Scalability testing
  - Control plane interoperability between vendors

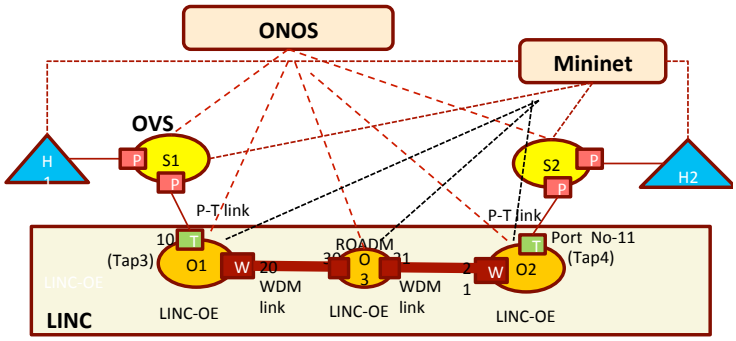

Demo GUI

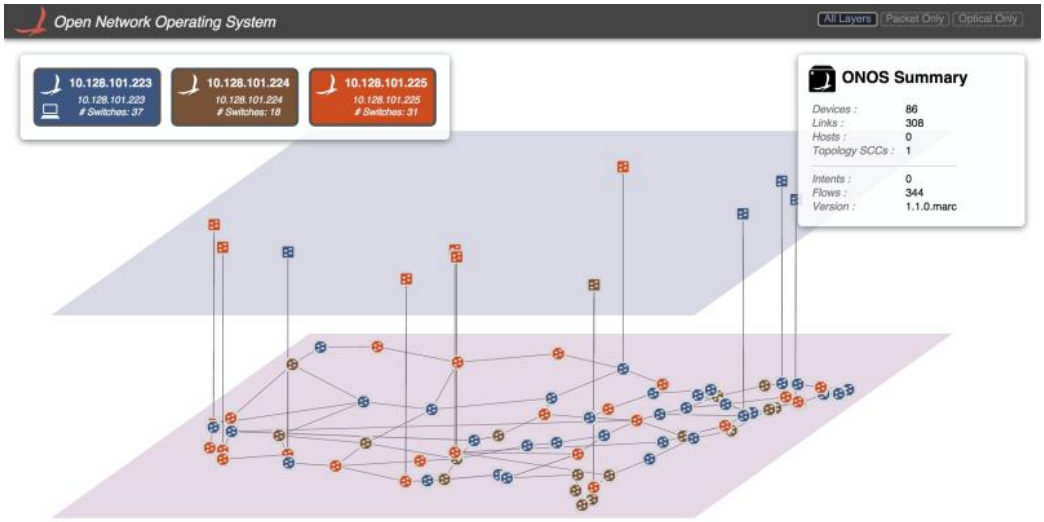

Try it yourself

https://wiki.onosproject.org/display/ONOS/Packet+Optical

Demo

https://www.youtube.com/watch?v=QA9ECsKpSug

Lessons Learned

- Feasibility
  - Converged packet optical control plane is possible
  - Offers scalability, HA, and performance
- Benefits
  - Significant improvement in network utilization
  - Drastic reduction in CAPEX and OPEX
  - DevOps model for transport networks
- Deeper insights
  - OpenFlow packet switches commercially available, resistance from L0 vendors
  - Abstractions are critical: intent framework, multi-layer graph

Customer Facing Services/Apps

- Bandwidth on-demand, bandwidth calendaring
- Customer portals for zero touch service provisioning
- On-demand and dynamic virtual private networks
- Elastic bandwidth services
- And so on…

Outline

- SDN for Service Providers
  - Background
  - Use cases
- Packet/Optical Use Case
  - Problem statement and conceptual solution
  - Implementation
  - Demonstration
- State of the Industry & Future Work

Vertical Integration: Packet Switches

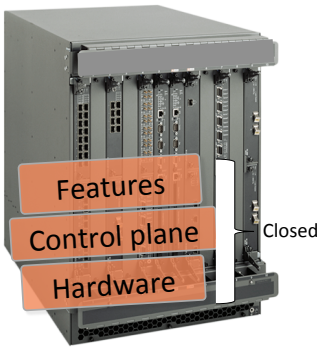

Packet switches are undergoing this transformation right now!

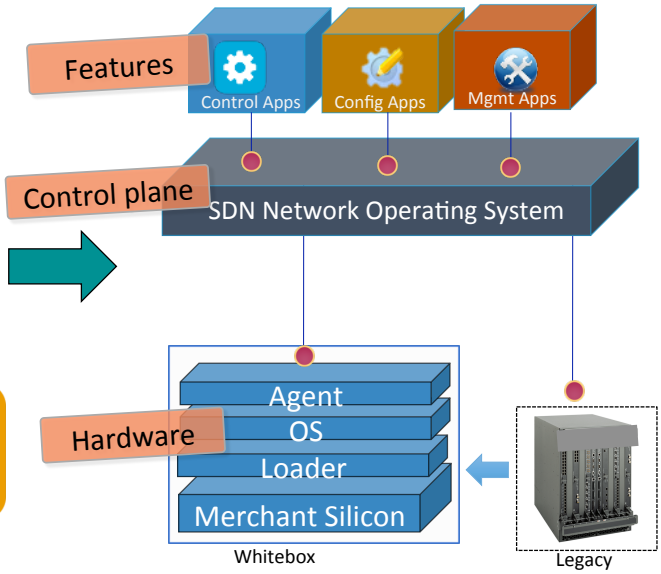

Vertical Integration: ROADMs

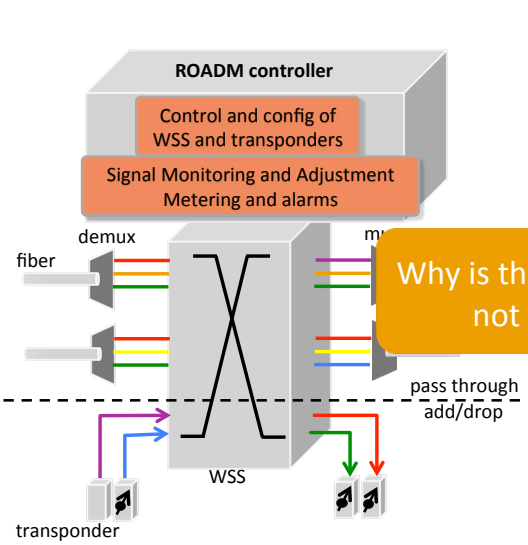

Why is this de-aggregation not happening?

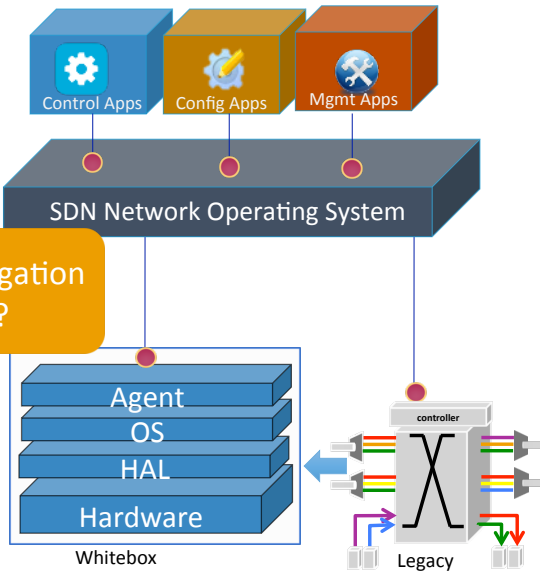

What Makes Optical Devices Different?

- "We need specialized mix of L0, L1, and L2 functions"
- "Physical impairments are too complex to monitor and manage externally"
- "Our analog transmission system is custom designed"
- "Every vendor has his own DSP which is proprietary and without programmable dynamics "
- "It's impossible to control all configuration and forwarding at scale"
- "You can't achieve sub-50ms failovers"
- And so on…

None of this is fundamental!

De-aggregation is inevitable

Open Optical Hardware

- Hardware Abstraction Layer
  - Hides optical impairments, thermal instability, power balancing, etc.
  - Can autonomously fix problems or perform maintenance
- OS
  - Server-like environment for switches
  - Manages various hardware sensors
  - Boot loader, utils, switch management, etc.

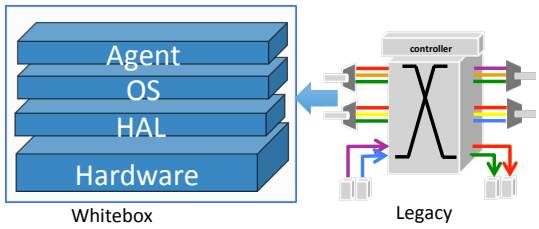

- Agent
  - Open and standardized interface for forwarding, configuration, and observability

Inviting all vendors to join us!

It's Happening Now!

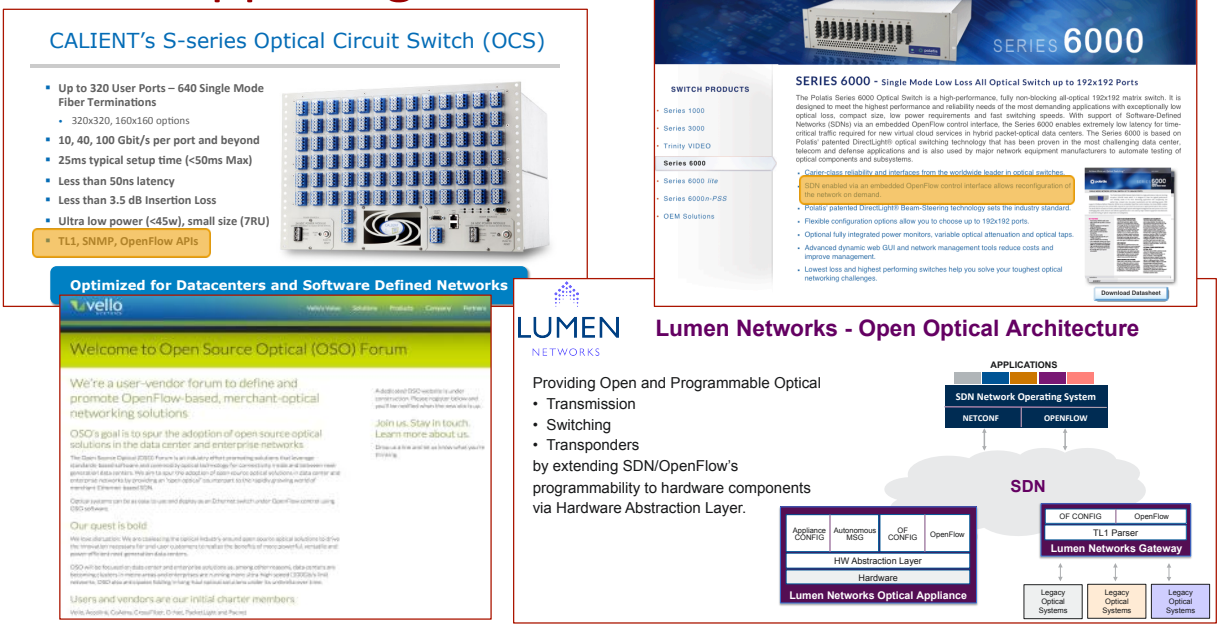

Vendor-Specific Domains

- Second problem with Optical Transport Industry
- Transport networks suffer from vendor lock-in
  - Domain consists of equipment from a single vendor
  - Each domain requires vendor-specific NMS/EMS
  - No data plane interoperability
- Profound impact on service providers
  - Complex management & orchestration tools
  - Problem identification & resolution
  - Expensive
- Is this fundamental?

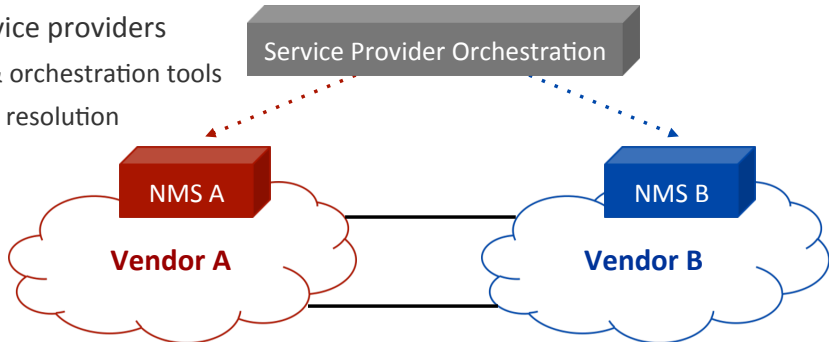

Why Vendor-Specific Domains?

- "We monitor network state and performance in NMS"
- "We built intelligent alarm and event handling between boxes and EMS"
- "Our EMS is the only system that can control our transmission"
- "Failures are handled faster and more efficiently by our NMS"
- And so on…

None of this is fundamental!

Vendor-specific domains will disappear

Vendor-Neutral Domains

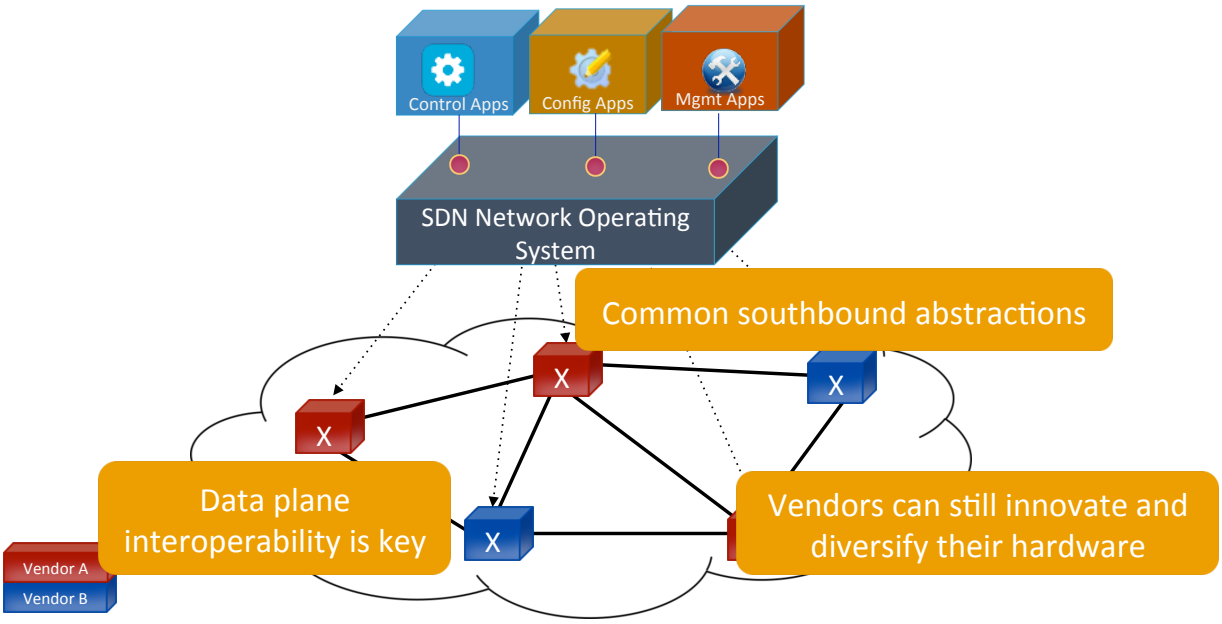

Future Work

- Looking to work with vendors that offer OpenFlow support
  - Something better than vendor-specific TL1
- Experiments on data plane interoperability
- Drive adoption of DevOps model for transport networks
- Hardware deployments

Proof of Concept

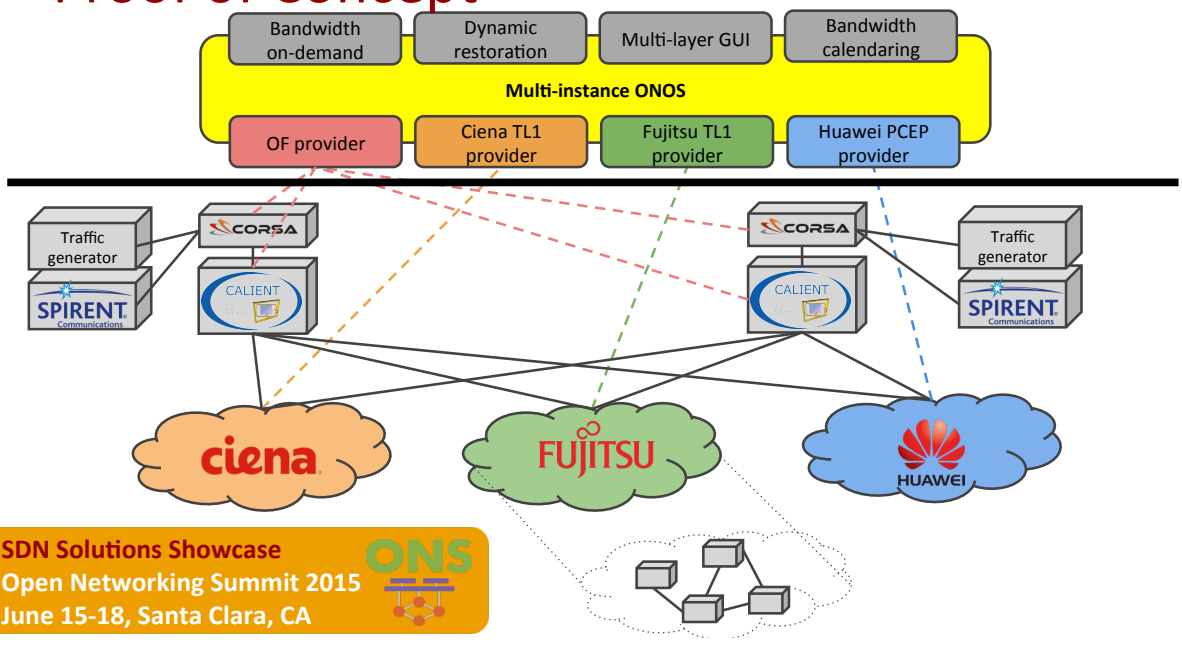

SDN Solutions Showcase
Open Networking Summit 2015
June 15-18, Santa Clara, CA

Cap-Grow-Drain Strategy

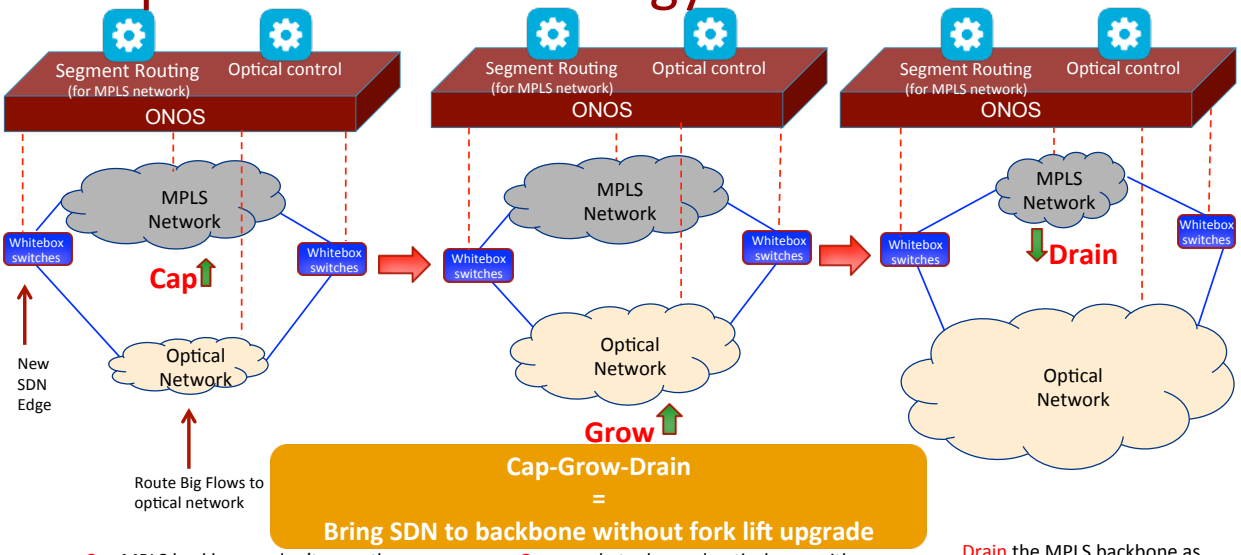

Cap MPLS backbone – don't grow the legacy MPLS backbone of proprietary routers

Grow packet edge and optical core with SDN control plane and make the best use of packet-optical technologies

Drain the MPLS backbone as most traffic transitions to new packet edge and optical core network

Summary

- Demonstrated converged packet/optical control plane for service providers
  - Scalability, HA, performance
  - Potential to dramatically decrease CAPEX & OPEX
  - Innovative services using DevOps model
- Need the right abstractions
  - Intent framework
  - Multi-layer graph

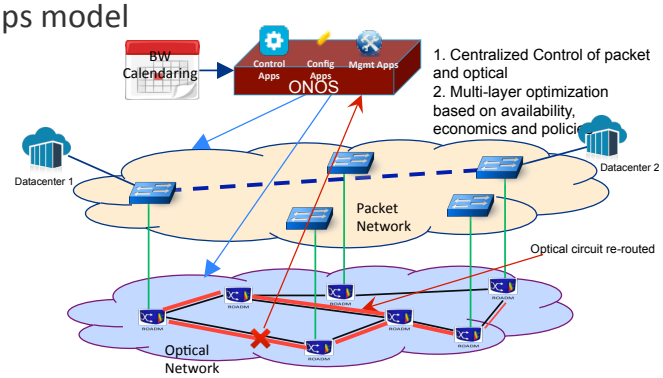

Call to Action

- Open and standardize hardware interfaces
  - Achieve control plane interoperability
- Eliminate vendor-specific domains
  - Achieve L0 data plane interoperability
  - Remove vendor-specific approaches (EMS & NMS)

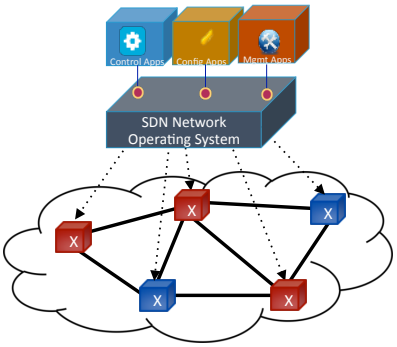

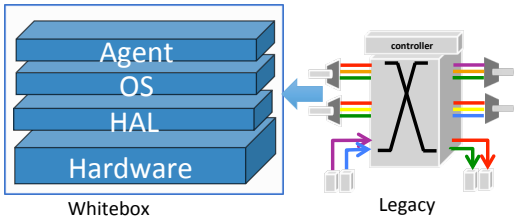

Whitebox Legacy

- If existing vendors don't take action, others will step in!