


# BURSTING DATA BETWEEN DATA CENTERS – CASE FOR TRANSPORT SDN

**Abhinava Sadasivarao, Sharfuddin Syed**, Ping Pan, Chris Liou (Infinera)  
Inder Monga, Andrew Lake, Chin Guok – Energy Sciences Network (ESnet)

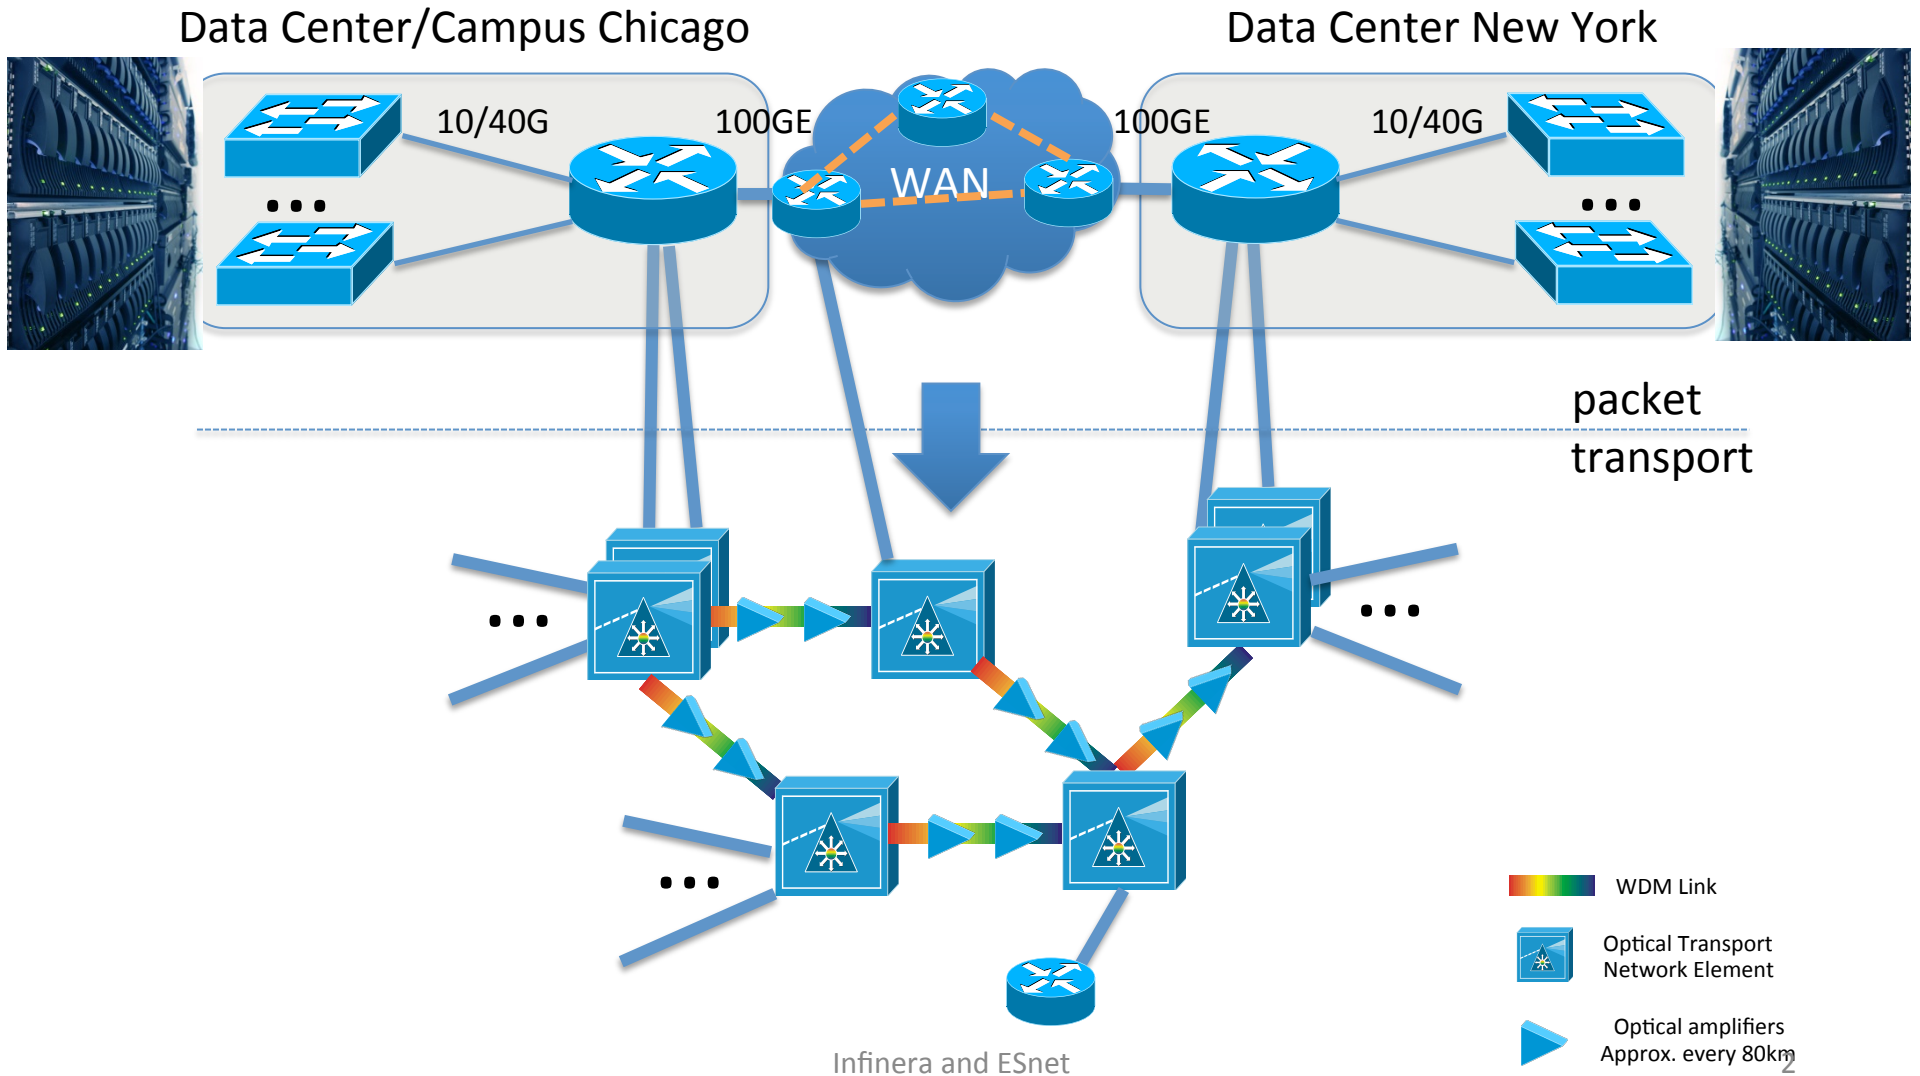
IEEE Hot Interconnects, August 2013



Infinera and ESnet

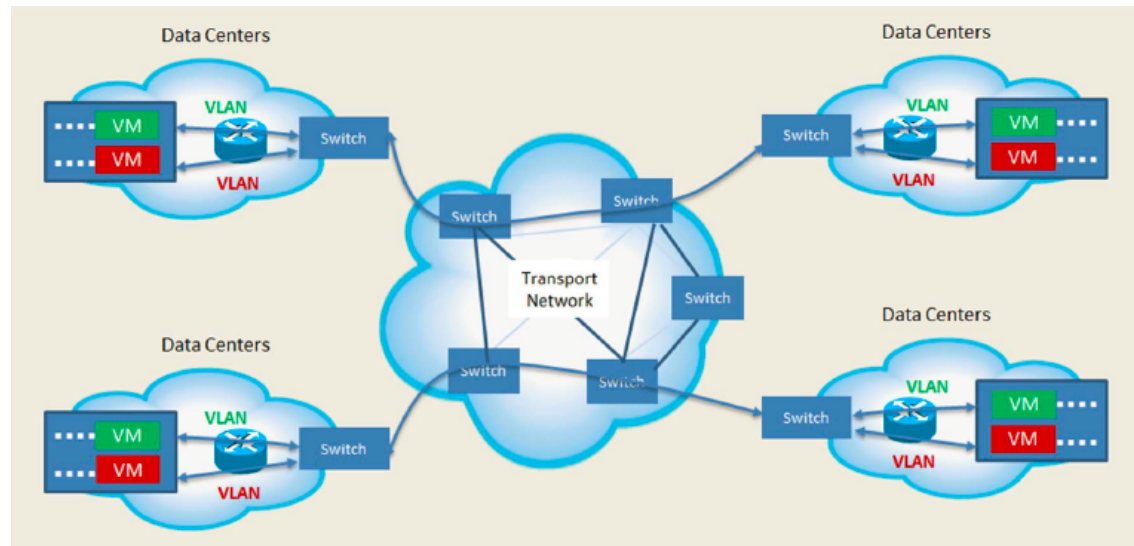


# A multi-layer network



# In Particular, Data Centers...

- Massive pool of compute and storage
- Network
  - Intra-DC to allow VM migration, HA, ...
  - Inter-DC over WAN for replication, delivery, ...



Infinera and ESnet

# Data Center – Packet World

- Intra-DC programmability
  - Lot of prior work done here [Portland][VL2]
  - OpenFlow making datacenter switches programmable easier
  - Centralized control logic can make these switches cheaper through simpler hardware (\$\$\$)
    - Decoupled data and control plane along with switches built with cheaper GbE/FCoE ports
- Inter-DC programmability
  - DC edge routers do L2/L3 tunneling
    - Packet layers are programmable – OpenFlow is an example
  - Google's B4
  - Microsoft's SWAN

[PortLand] Radhika Niranjana Mysore, Andreas Pamboris, Nathan Farrington, Nelson Huang, Pardis Miri, Sivasankar Radhakrishnan, Vikram Subramanya, and Amin Vahdat. 2009. PortLand: a scalable fault-tolerant layer 2 data center network fabric. In *Proceedings of the ACM SIGCOMM 2009*

[VL2] Albert Greenberg, James R. Hamilton, Navendu Jain, Srikanth Kandula, Changhoon Kim, Parantap Lahiri, David A. Maltz, Parveen Patel, and Sudipta Sengupta. 2009. VL2: a scalable and flexible data center network. In *Proceedings of the ACM SIGCOMM 2009*

# Data Center Interconnects – Transport View

- Providing high-capacity links for the edge DC routers
- However (traditional service providers)
  - If the inter-DC router links are overprovisioned, equally overprovisioned transport capacity required
  - These links are circuit-switched, statically allocated ahead of time
  - Operational are isolated by network layers
    - SP manage L3/L2 and L1 infrastructure separately - Multiple set of people to operate them – Engineering, planning, provisioning
    - L2/L3 control planes don't talk to L1. Can't do end-to-end multi-layer TE (Multi-vendor L1 control plane also is "in-theory")
- How can we make networks end-to-end programmable in a multi-layer, multi-vendor scenario?

# Transport Paradigm is Different!

## Packet World

- Connectionless
- Dynamic flows
- Inline control plane (NMS independent)
- Distributed CP solutions with numerous protocols


## Transport World

- Connection (circuit) oriented
- Static pipes/configuration
- EMS/NMS + Cross-connect paradigm
- Nascent distributed CP (GMPLS), not inline

**Historically, transport networks have been programmable by Centralized NMS/OSS.**

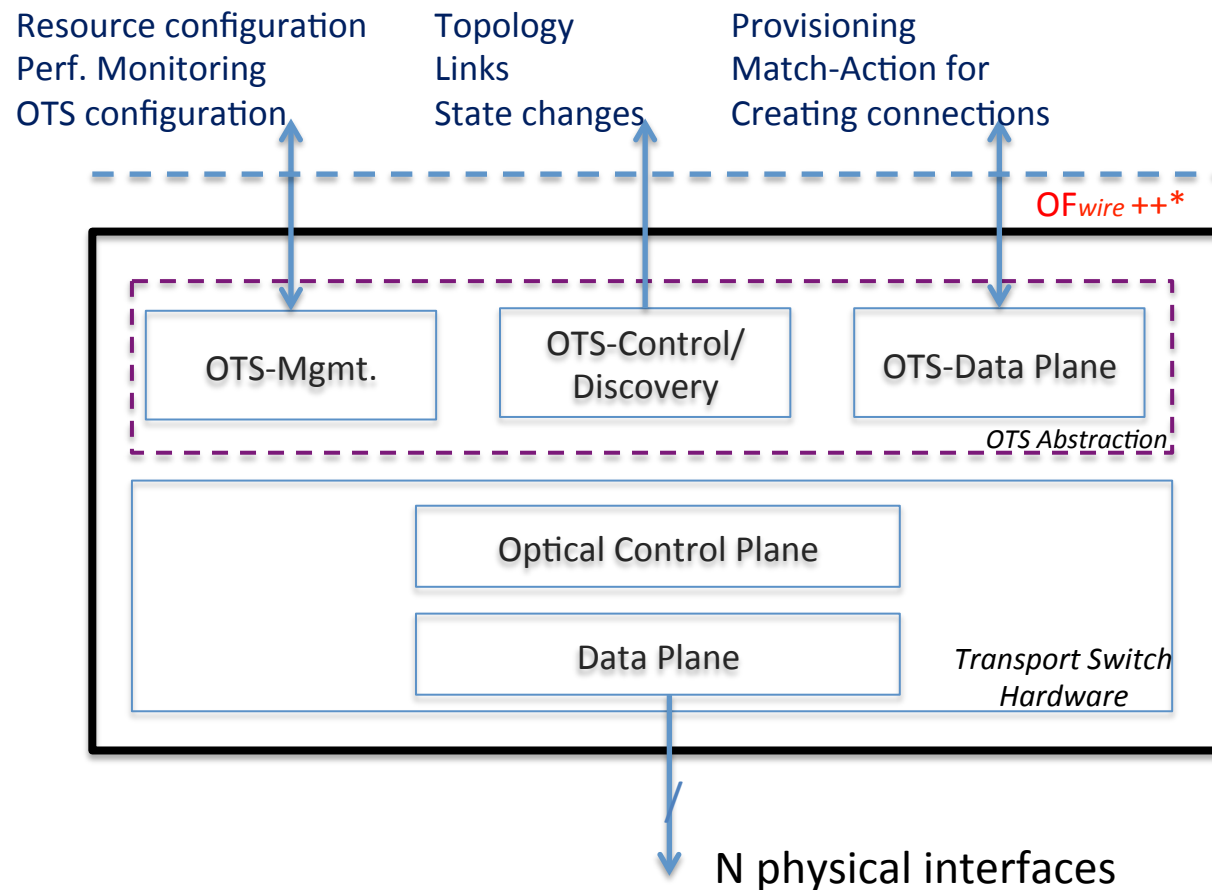
# Architectural Approach

- Abstract programmability between layers
  - Virtual overlay of Layer3/Layer2/Layer1
  - Central controller layer operates on this overlay
  - Program the “network fabric” than on a per-element basis



- Open Transport Switch (OTS)
  - OpenFlow-enabled, lightweight virtual switch representing a Transport Network Element
  - Capabilities exposed by OTS depends on the capabilities of the Transport element
  - Provides all the information needed to provision, control and monitor

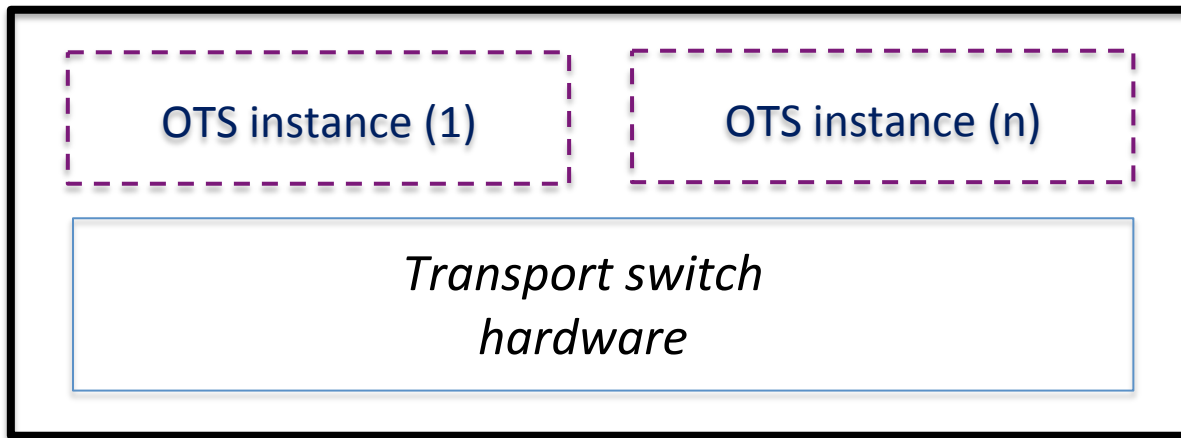
# OTS Subsystems: A high-level view



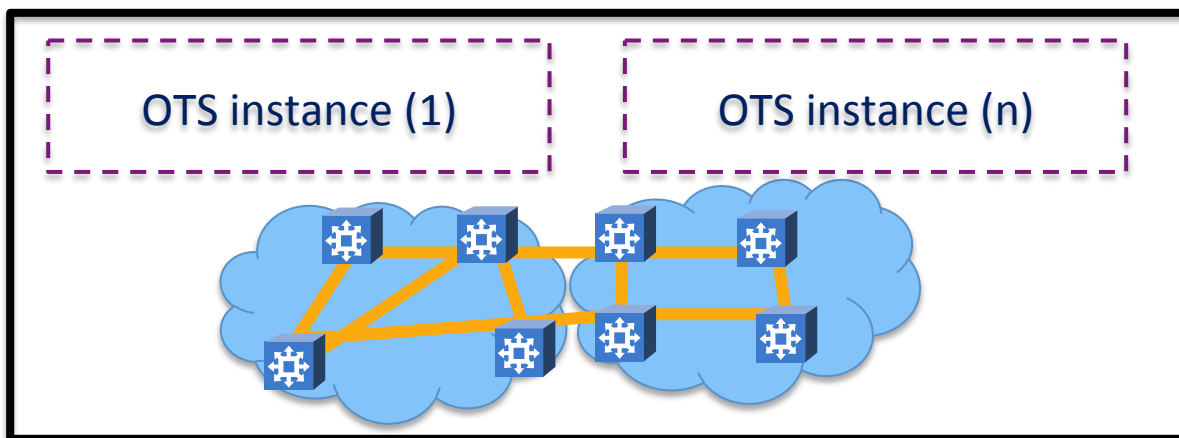
\* Recent architectural discussions have decided to not change OFw protocol. Infinera and ESnet



# OTS Building Blocks (contd.)

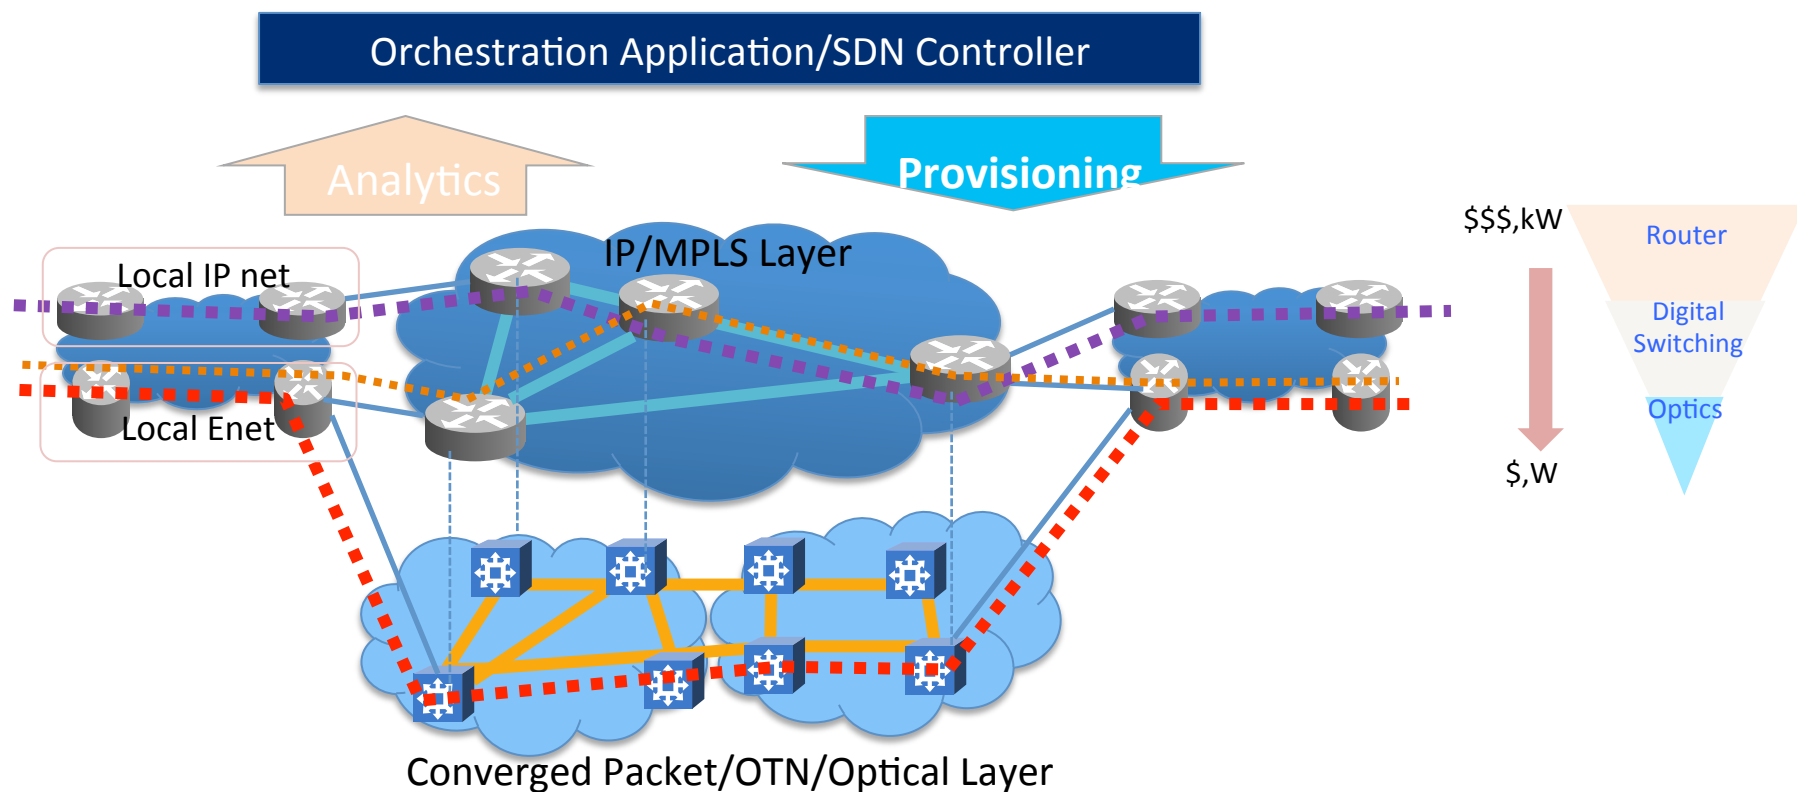


*Partition a hardware device*





*Virtualize a multi-domain transport network*

# Use Case: Multi-Layer Optimization



- Next-gen networks will drive need for multi-layer representation, topology computation & provisioning
- SDN approach facilitates orchestration across layers & domains

# To Summarize...

- Advantages 
  - Programmability of Transport Networks
  - Virtualization of Transport Networks
- Challenges 
  - Migration
  - Co-existence with control plane
  - Multi-vendor interface

# THANK YOU!

asadasivarao at infinera dot com

ssyed at infinera dot com