SDN/NFV Whitepaper

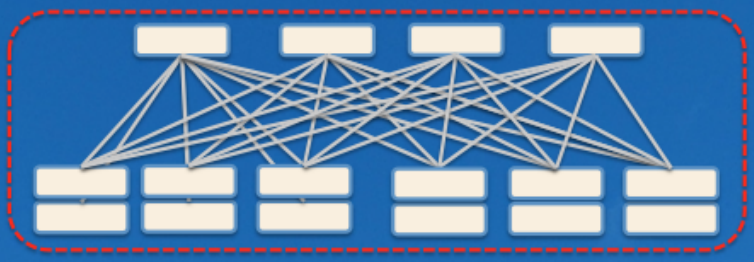

CORD Fabric

Proof-of-concept demonstration at Open Networking Summit, June 2015

A Whitepaper by:

ONF OPEN NETWORKING FOUNDATION

ON.LAB

at&t

onos Open Network Operating System

In collaboration with:

DELL

inMon

Introduction

Many carriers in both the telecom and cable fields are espousing next generation architectures that borrow heavily from web scale data centers and cloud providers. With this they use SDN and NFV to transform their carrier functions into workloads that are hosted on this common infrastructure. We've coined the phrase *Central Office Re-architected as a Datacenter* (CORD) to refer to this approach of providing infrastructure as a service and networking services as tenant applications for that infrastructure. This paper focuses on the fabric within CORD and a companion paper focuses on applications that can support broadband access by virtualizing Optical Line Termination equipment (vOLT) as well as CPE (vCPE). The fabric is the common re-usable interconnect for all the applications and hardware, and can be thought about like the backplane of multi-card chassis devices that are typically found today in central offices and cable head-ends.

CORD[1] Fabric

The typical legacy data center is designed as a three-layer hierarchy with the devices getting larger, more complex and more expensive moving from the top-of-rack (ToR) to the Core. In effect, this design has a spanning-tree with three layers, ToR, Aggregation and Core switches, with significant oversubscription at each layer.

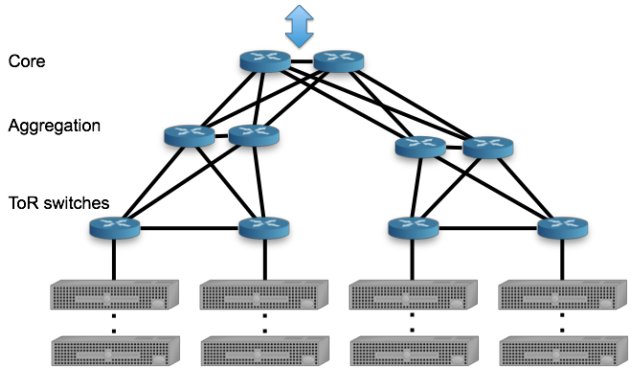

Figure 1. Legacy Datacenter Architecture

Instead of this oversubscribed and hierarchical design, CORD moves towards a next-gen flat fabric which is conceptually one giant distributed high performance switch/router with 'line cards" and "backplanes" instantiated by bare-metal switches and optimized for east-west traffic. This fabric uses the well known folded-Clos leaf-spine architecture to create a flat, homogenous network that interconnects compute, storage and IO.

[1] CORD is an acronym for Central Office Re-architected as Datacenter

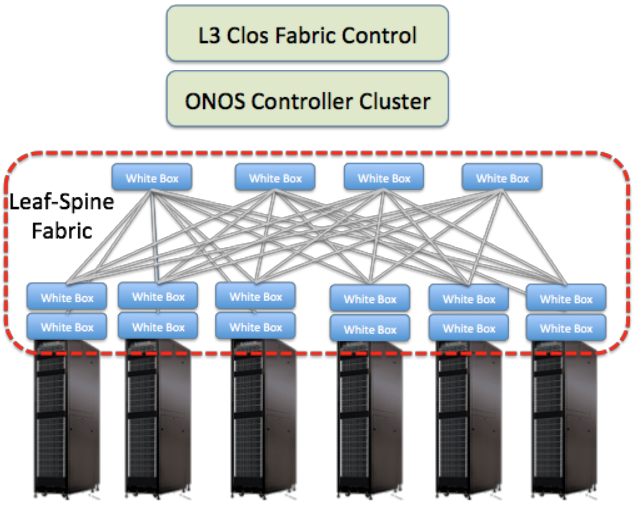

Figure 2. CORD Fabric architecture

We've developed the CORD fabric to follow the webscale approach of using open-source software and commodity hardware built on merchant Silicon ASICs. It is based on SDN principles which afford simplicity of control and reduce the complexity of the software required in the switches. SDN allows the fabric to be treated as a system, instead of a collection of individual boxes. It can be configured as a whole, and managed with a single-pane-of-glass visibility via a scalable, highly available ONOS controller cluster.

In this "one giant router" fabric, the line cards are the leaf switches that are standard Top-of-Rack switches in data-centers (with 48X10G ports and 6X40G uplinks). The backplane of the distributed router is comprised of spine switches that are typically 32X40G bare-metal switches. Internally, the fabric uses segment routing technology to apply MPLS labels and use ECMP shortest paths to route between ToRs that have their uplinks connected to multiple spine switches. This approach supports hashing traffic across the spine in order to gain availability and capacity. In this design we've chosen to use IP to the ToR, and Ethernet within the rack.

The CORD fabric automatically performs traffic-engineering. It allows external analytics to detect elephant flows that cause congestion in the uplinks, and take immediate, direct action by moving the flows to under-utilized links via the fabric control application running on the ONOS cluster.

ONF, AT&T, ON.Lab and ONOS project, in collaboration with Dell and InMon, will demonstrate a CORD Fabric proof-of-concept at Open Networking Summit on June 15-18, 2015. This whitepaper describes the CORD Fabric demo at ONS.

CORD Fabric Demo at ONS

CORD Demo Highlights

CORD demo at ONS will showcase:

- leaf spine architecture with following software and hardware building blocks-
  - SDN control using ONOS and segment routed fabric control app
  - hardware comprised of Dell Switches
  - analytics from InMon for automatic traffic engineering
- fabric capabilities including-
  - L2 switching within a rack
  - L3 routing across racks
  - ECMP with different paths taken through fabric between same pair of racks
  - Internet connectivity through the fabric
  - Recovery from dataplane and control plane failures
  - Analytics using InMon and automated traffic engineering

CORD Fabric demo set-up

The CORD demo puts together all the above software and hardware building blocks as shown below:

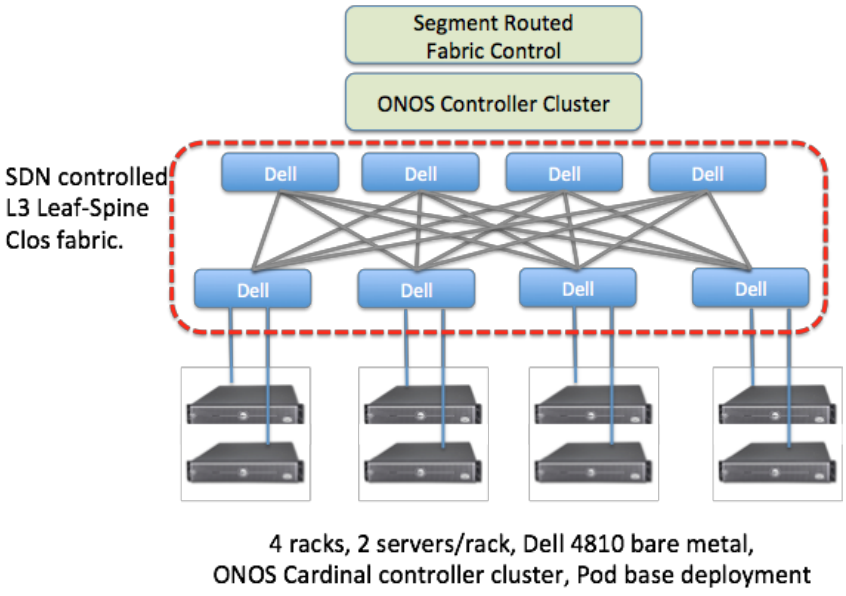

Figure 3. CORD Demo Set-up

The CORD Fabric demo hardware comprises of Dell switches. The SDN control plane comprises of ONOS with Segment Routed Fabric Control app.

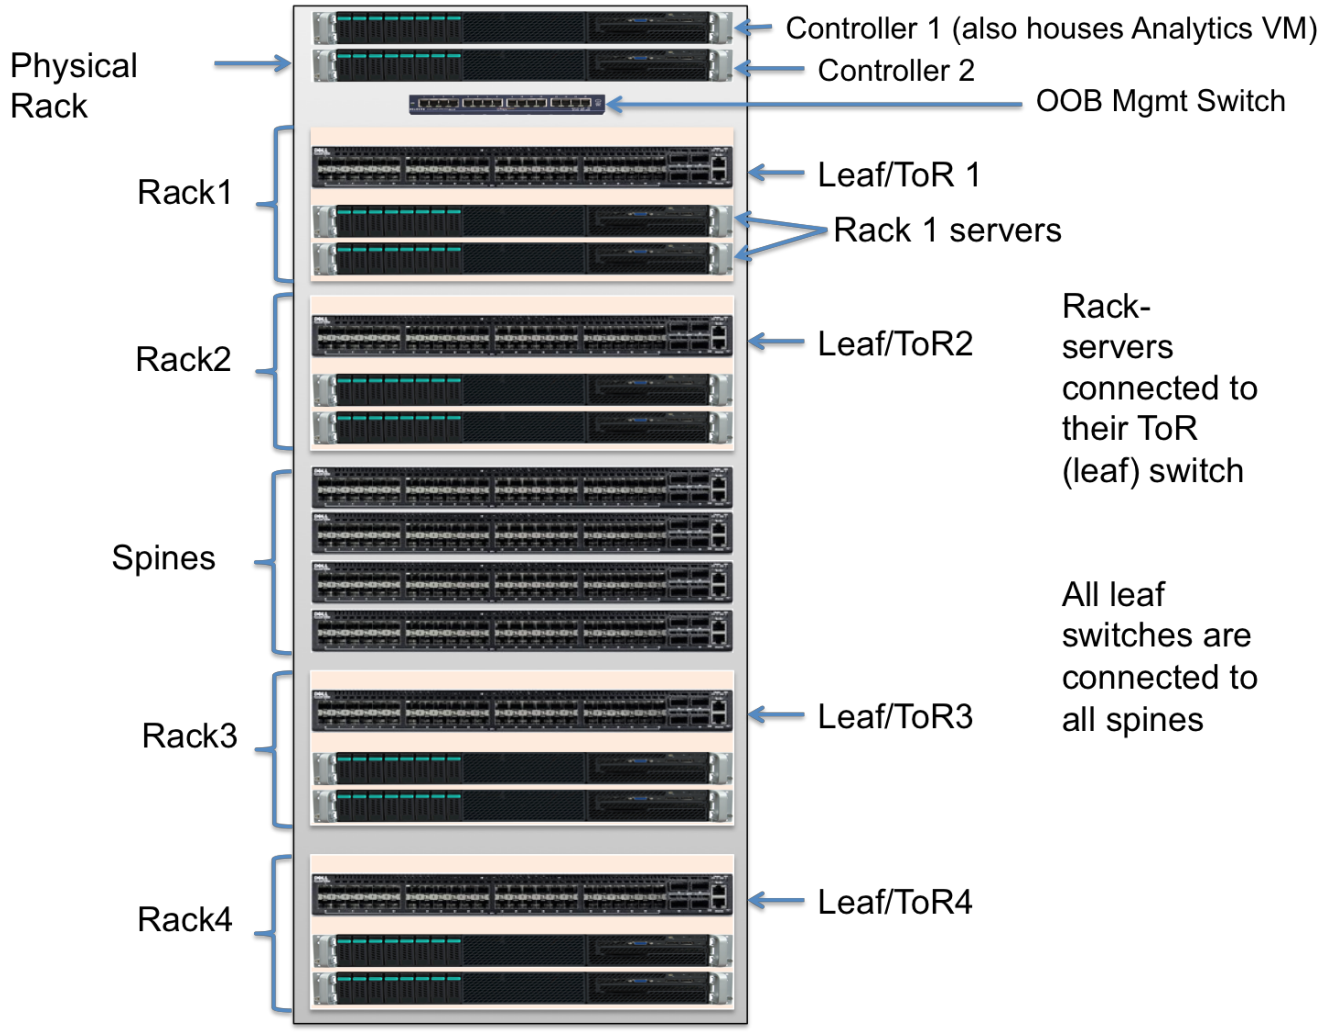

Figure 4. CORD Fabric demo hardware rack

External analytics from InMon are used in the demo to perform automatic traffic engineering.

Figure 5. CORD Fabric Analytics and Traffic engineering

CORD Fabric demo in action

Come to see the CORD Fabric demo live at Open Networking Summit from June 15th-18th, 2015. Additional details here.

Next Steps

It's our intention to build upon this CORD Fabric proof-of-concept, to expand the types of scale-out networking functions that can be supported, and to mature the work throughout the rest of 2015. Some initial steps may include application of this work to the other CORD PoCs, including vOLT, vCPE and NFaaS. Collaborative participants are encouraged to contact onos-dev@onosproject.org