Importance of this topic

Network operators face significant challenges supporting ever-increasing bandwidth demands and ever-increasing service expectations. For example, AT&T has seen data traffic increase by 100,000 percent in the last eight years, and plans are now underway to roll out ultra-fast fiber and access to 100 cities across the US [1]. At the same time, introducing a new feature often takes months (waiting for the next vendor product release) and sometimes years (waiting for the standardization process to run its course).

In response to these challenges, network operators are looking for ways to benefit from both the economies of scale (infrastructure constructed from a few commodity building blocks) and the agility (the ability to rapidly deploy and elastically scale services) that commodity cloud providers enjoy today.

Cloud economies and agility are especially needed at the edge of the operator network—in the Telco *Central Office (CO)*—which contains a diverse collection of purpose-built devices, assembled over fifty years, with little coherent or unifying architecture. For example, AT&T currently operates 4700 Central Offices, some of which contain up to 300 unique hardware appliances. This makes them a source of significant CAPEX and OPEX, as well as a barrier to rapid innovation.

This tutorial describes *CORD*, an architecture for the Telco Central Office that combines *Software-Defined Networking (SDN)*, *Network Functions Virtualization (NFV)*, and elastic cloud services—all running on commodity hardware—to build cost-effective, agile networks with significantly lower CAPEX/OPEX and to enable rapid service creation and monetization.

The Central Office Re-architected as a Data Center (CORD) Tutorial

- **Architectural Overview** (1hr) -- An overview of the tutorial and a brief introduction to the CORD architecture. Includes a description of the hardware and software building blocks, plans for a market trial.

- **Software Organization** (1 hr) -- An in-depth look at CORD's software organization, including the role XOS plays in creating, naming, operationalizing, managing, and composing micro-services (VNFs).

- **Sizing and Configuration Options** (1 hr) -- A discussion of various options of how CORD can be sized (hardware) and configured (software), including where the access equipment and CORD compute complex is located, and plans to augment the CORD service portfolio with micro-services that can play a role in VNF disaggregation.

- **ONOS** (1 hr) -- An in-depth look at CORD's use of SDN, including the role ONOS plays in controlling the white-box switches and OLT pizza boxes. The focus is on ONOS's scalable and highly available design, the northbound interface it exposes to control applications and the southbound interface by which it interacts with various switching devices.

- **Hardware Building Blocks** (30 min) -- An in-depth look at CORD's OCP-based hardware infrastructure, including servers, OLT "pizza boxes", and white-box switches. The discuss includes both an overview of the underlying merchant silicon, and the software stack (e.g., Atrium) running on the white-box switches.

- **Switching Fabric** (30 min) -- An in-depth look at the CORD switching fabric, including the use of segment routing and the implementation of virtual networks in support of service composition. The focus is on how these elements are implemented as control applications running on ONOS.

- **VNF-as-a-Service: Part 1** (1 hr) -- An in-depth look at the how CORD's core VNFs (micro-services) -- starting with vOLT -- are implemented, including the relationship between these VNFs and XOS, ONOS, OpenStack, and the underlying hardware.

- **VNF-as-a-Service: Part 2** (1 hr) -- An in-depth look at the how CORD's core VNFs (micro-services) -- continuing with vSG and vRouter -- are implemented, including the relationship between these VNFs and XOS, ONOS, OpenStack, and the underlying hardware.

- **Wrapup** (1 hr)-- Discussion of generalizations of CORD for Mobile and Enterprise deployments, and a discussion of how CORD can be sized and configured for different scenarios.

Speakers

Larry Peterson, Chief Architect, Open Networking Lab

Larry is the Chief Architect at Open Networking Lab and is the lead architect for the CORD project. Larry Peterson comes to ON.Lab from Princeton, where he was the Robert E. Kahn Professor of Computer Science and Director of the Princeton-hosted PlanetLab Consortium. He served as chair of the CS Department from 2003-2009. In 2007, Peterson co-founded CoBlitz LLC to commercialize CDN technology developed on PlanetLab. CoBlitz was acquired by Veriue Inc. in 2010, and subsequently by Akamai in 2012. Peterson is co-author of the best-selling networking textbook Computer Networks: A Systems Approach (5e), and chaired the initial planning efforts that led to NSF's GENI Initiative. His research focuses on the design and implementation of networked systems. Some of his recent projects and papers can be found here. Peterson is a former Editor-in-Chief of the ACM Transactions on Computer Systems, was on the Editorial Board for the IEEE/ACM Transactions on Networking and the IEEE Journal on Select Areas in Communication, and served as program chair for SOSP, NSDI, and HotNets. Peterson is a member of the National Academy of Engineering, a Fellow of the ACM and the IEEE, and the 2010 recipient of the IEEE Kobayashi Computer and Communication Award. He received his Ph.D. degree from Purdue University. Check out www.cs.princeton.edu/~llp for more information.

Ali Al-Shabibi, CORD Architect, Open Networking Lab

Ali Al-Shabibi is a software architect and the lead developer for CORD and cord developer of ONOS. He was a founding engineer with ON.Lab and has also been the lead engineer and maintainer of OVX, and FlowVisor, two network hypervisors, at the Open Networking Laboratory. Previously, he was a post-doc at Stanford University researching OpenFlow and SDNs in Nick McKeown's group. He received his Ph.D from the University of Heidelberg in Germany in 2011 after performing his doctoral research at CERN (European Centre for Nuclear Research) in the ATLAS (A Toroidal Lhc ApparatuS) Networking group, where he contributed to the design and development of the TDAQ (Trigger and Data Acquisition) Network. Ali Al-Shabibi brings vast knowledge of flow models and congestion avoidance protocols. He comes to ON.Lab from Stanford and CERN, where he analyzed large, mountainous systems, such as 'Portes du Soleil' and 'Argentiere.' Urban wannabe, unrelenting espresso consumer and dedicated traveler, Ali Al-Shabibi is an avid soccer and table tennis player. While he was born in Baghdad, he grew up in Geneva, Switzerland, where he attended the Swiss Federal Institute of Technology (EPFL) for his BSc and MSc degrees. Nowadays, Al-Shabibi can be found theorizing and philosophizing about SDNs or dreaming up cool networking applications with the ONRC crew over many coffees.