# Headend Re-architected as a Data Center (HERD)

## Adaptation of Central Office Re-architected as a Data Center (CORD) to the MSO market, evolving the headend infrastructure to become a data center

A Technical Paper prepared for SCTE/ISBE by

Soumen Chatterjee
Director, Product Line Management, Blue Planet
Ciena
3939 N. First Street
San Jose, CA 95134
(408) 219-6125
schatter@ciena.com

Jim Kleinsmith
Director, MSO Business Development
Ciena
130 Conifer Lane
Walnut Creek, CA 94598
925-457-2451
jkleinsm@ciena.com

 

# Table of Contents



# List of Figures

# Introduction

MSO Head Ends have evolved from the early days of offering only cable TV to providing voice, video and data services. The services landscape is evolving rapidly and bandwidth requirements have increased from 10 Mbps in DOCSIS 1.0 to 1Gbps and 10Gbps in DOCSIS 3.1. Cable MSOs are under competitive pressure to cost effectively offer new services as they compete against over-the- top and new service offerings from other service providers.

Central Office Re-architected as a Datacenter (CORD) is a carrier-grade solution led by the ON.Lab and AT&T to use the Open Source Operating System (ONOS) to bring datacenter economics of scale and cloud-style agility to service provider networks by applying the relatively new Software Defined Networking (SDN) and Network Functions Virtualization (NFV) technologies using a commodity infrastructure.

We will present "Head Ends Re-Architected as a Data Center" (HERD), the adaptation of CORD to the MSO market, by evolving the headend infrastructure to become a data center. In the HERD concept, SDN, NFV, commodity off the shelf servers and networking infrastructure virtualize the network by disaggregating network functions from hardware into Virtualized Network Functions or VNFs. VNFs running in virtual machines (VMs) or containers can be chained together to create new orchestrated services, offering global and local functionality.

The HERD software architecture includes Multi-Domain Service Orchestration (MDSO), which manages end-to- end service provisioning and life cycle management, Openstack, which acts as the virtual server infrastructure manager and ONOS, used to program the white box (generic, off-the-shelf switching are routing hardware) based switches and host control plane applications. The session includes HERD benefits and applications such as Access as a Service (ACCaaS), Subscriber as a Service (SUBaaS), Internet as a Service (INTaaS) and Content Distribution Network (CDN).

## 1. Challenges

Cable MSOs are facing significant challenges as the demand for more bandwidth is increasing year after year at unprecedented pace, stressing the current infrastructure. In addition, the need to deploy new services to more effectively compete against competitors as well as over-the-top (OTT) service providers is proving to be a daunting task using traditional methods of headend design. As a diverse set of services are deployed, the smooth integration of such services with existing ones such as voice and video over existing headend architecture and networks is equally crucial. Therefore, a new approach to design a next generation headend infrastructure that is more scalable, agile and programmable is in the interest of many MSOs.

Traditional MSO networks, including the headend, is architected and implemented using proprietary devices that were designed for a time where the number of services were limited, the subscription model and traffic flows were more static and bandwidth requirements were significantly lower than today's expectations. This reflects a similar situation that the traditional network operators face. 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 ultrafast 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).

 

The Telco Central Office (CO) is analogous to the headend, in that it contains a diverse collection of purpose built devices, with little coherent or unifying architecture. For this reasons we will present how Software-Defined Networking (SDN) and Network Functions Virtualization (NFV) along with white-box appliances such as switches and servers are used to transform the CO to an agile, elastic, cloud-like infrastructure running on commodity hardware, achieving a cost-effective solution to address tomorrow's MSO requirements.

## 2. CORD

Central Office Re-architected as a Datacenter aims to transform the traditional Head End using two related trends and technologies:

- SDN is used to separate the network's control and data planes making the control plane programmable, simplifying the forwarding devices, resulting in lower cost white-box switches to be deployed.
- NFV is used to then move the data plane from running on dedicated, proprietary hardware to running as virtual machines on common off the shelf (COTS) servers. Decoupling the function from the proprietary hardware and running it as software on regular servers reduces CAPEX as high-margin devices are replaced with commodity servers. In addition, and more importantly, this move also reduces OPEX as orchestration software automates the complete end-to-end service, removing human error, reducing service provisioning time and improving operator agility with the added benefit of creating an opportunity for innovation.

The goal with CORD is to not only replace today's purpose-built hardware devices with their more agile software-based counterparts, but also make the Head End an integral part of every MSOs larger cloud strategy, enabling them to offer more valuable services. By utilizing a general software architecture to orchestrate the aforementioned infrastructure, MSO's can offer a wide range of services that include applications such as Access as a Service (ACCaaS), Subscriber as a Service (SUBaaS), Internet as a Service (INTaaS) and Content Distribution Network (CDN).

### 2.1. Commodity Hardware

Figure 1 illustrates the leaf-spine based network fabric consisting of mostly white box Top-of-Rack (ToR) switches connecting the servers in a highly redundant fashion, thereby creating ample bandwidth for east-west traffic that is commonly seen in scale-out type cloud applications.

As illustrated in figure 1, all hardware consists of COTS servers except for phy-specific devices.

In the CORD Passive Optical Network (PON) use case, there is a server based virtual Optical Line Terminal (vOLT) and phy-specific devices, the OLT Media Access Control (MACs), are separate from the server.

A simple CORD POD (a modular infrastructure) consists of Open Compute Project (OCP)[1] qualified servers with 128GB of RAM, 2x300GB HDDs, and 40GE dual port NICs connected with OCP-qualified, OpenFlow-enabled switches deployed in a 32x40GE port configuration.

 


**Figure 1 - Target hardware built from commodity servers, I/O Blades and switches**

## 2.2. The CORD Software

Figure 2 illustrates the major software building blocks deployed in the CORD Pod.

**OpenStack [3]** is used as the virtualized infrastructure manager (VIM) and is responsible for creating and provisioning virtual machines and virtual networks.

**Open Network Operating System (ONOS) [**4] is the software used to manage the underlying white-box switch fabric. ONOS is more than an SDN controller and can also host many control applications that determine how the network fabric should be configured.

**XOS [5]** is an open source orchestrator for assembling and composing services.

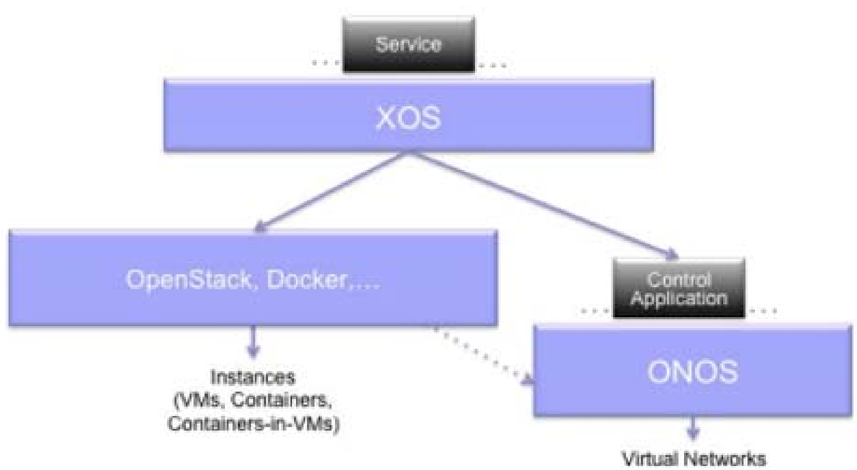


**Figure 2 - Open source software components used to build CORD**

Given the current hardware and software components, the first step is to transform the central office into a data center by mapping traditional hardware based, proprietary devices to the data center architecture created using the CORD POD.

 

# 3. Mapping Hardware Based Head End to CORD Based Infrastructure

The following Passive Optical Network (PON) use case shows how the PON architecture is mapped to a CORD POD. A possible use case for DOCSIS and CCAP functions is discussed later in this document.

**CORD Passive Optical Network Use Case**

In the CORD PON implementation, optical line terminals are mapped to an I/O blade with the PON OLT MAC, which uses an Open Compute Project co-developed 1RU pizza box GPON MAC device that nicely fits into the server-rack form factor. This particular box is controlled via OpenFlow. The virtual OLT (vOLT) runs on top of ONOS, and delivers the functions commonly found in traditional OLT devices. The vOLT implements subscriber authentication, VLAN management and various other control plane functions.

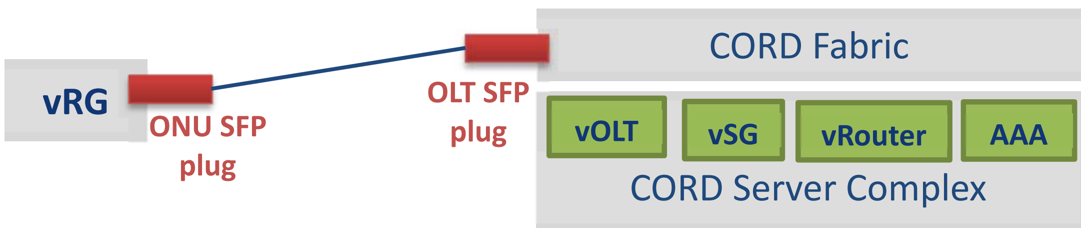


**Figure 3 - CORD PON Architecture**

Next is the CPE, which is sometimes called a "residential gateway or RG" that is traditionally located in the customer premises. Services provided by such CPE devices may include DHCP, NAT, Firewall and so on. One of the trending use cases for SDN and NFV is the concept of virtualized CPE, in which as many functions as possible are moved to the "cloud" with the CPE device being reduced to the bare essentials, sufficient to provide connectivity and maybe crude tests for troubleshooting. The benefits are clear; fewer truck rolls and the ability to offer value-added services that are traditionally considerably costlier to do with traditional, hardware based, proprietary CPEs. With CORD, virtual CPE applications can be implemented using a full VM, lightweight containers or a chain of such containers.

The last function we will cover is the broadband network gateway (BNG), which is more complex than the aforementioned devices. The BNG enables subscribers to connect to the public internet. Depending on the sophistication of the traditional BNG device used, it is not uncommon to see various functions such as VPNs, GRE and MPLS tunneling to be also supported. In the CORD implementation, the virtualized BNG sometimes known as the vRouter is implemented as a VM hosted on the ONOS controller, managing the flow of traffic through the leaf-spine architecture based switch fabric.

Once the minimum CORD infrastructure is mapped to the new CORD design, additional value added services can be added for differentiation and better competition. This is where CORD's *Everything-as-a-Service (XaaS)* principle comes in to play.

- **vOLT** implements Access-as-a-Service, where each tenant corresponds to a subscriber VLAN.
- **vSG** implements Subscriber-as-a-Service, where each tenant corresponds to a subscriber.
- **vRouter** implements Internet-as-a-Service, where each tenants corresponds to a subnet.

# 4. CORD Applied To The Headend

## 4.1. Generic Head End Architecture

The Cable Head End architecture has not changed much for the last 10 years. Access routers connect the Primary Head End via the backbone to the Media Centers and the Internet. Metro Head Ends connect to the Primary Head Ends via DWDM fiber optic rings. User routers in the Head Ends connect to Cable Modem Termination Systems (CMTS). CMTS connect to Optical Nodes. Optical nodes connect to cable modems via coax.

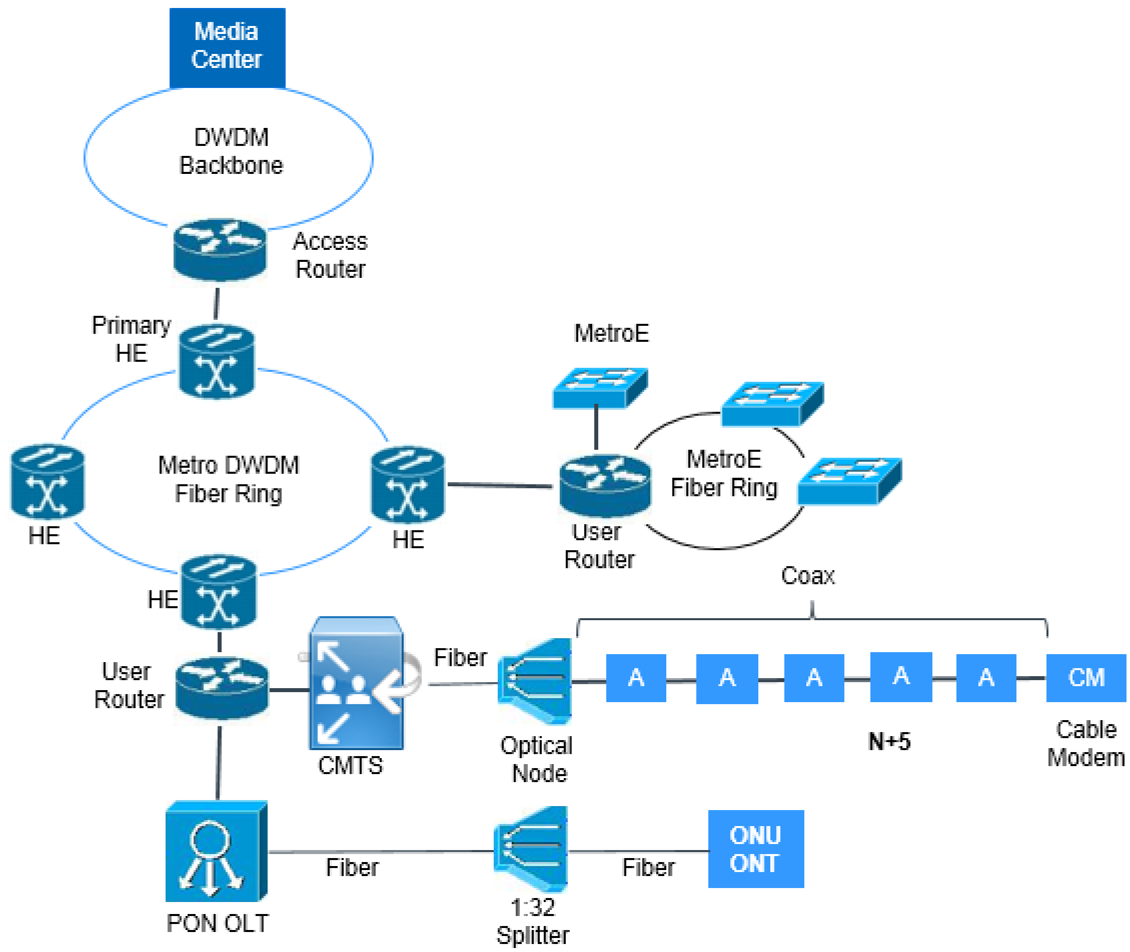


**Figure 4 - Generic MSO Head End and Access Network Example**

Metro Ethernet business services connect enterprise customers to the user router ports via fiber.

In the conventional Head End architecture, each device is a proprietary implementation running on a vendor's purpose built hardware.

 

## 4.2. HERD Architetcure

Head Re-architected as a Datacenter (HERD) uses software defined networking, merchant silicon in common off the shelf switches and servers and virtual network functions.

A HERD POD can include any number x86 Compute & Storage and Merchant Ethernet Switches.

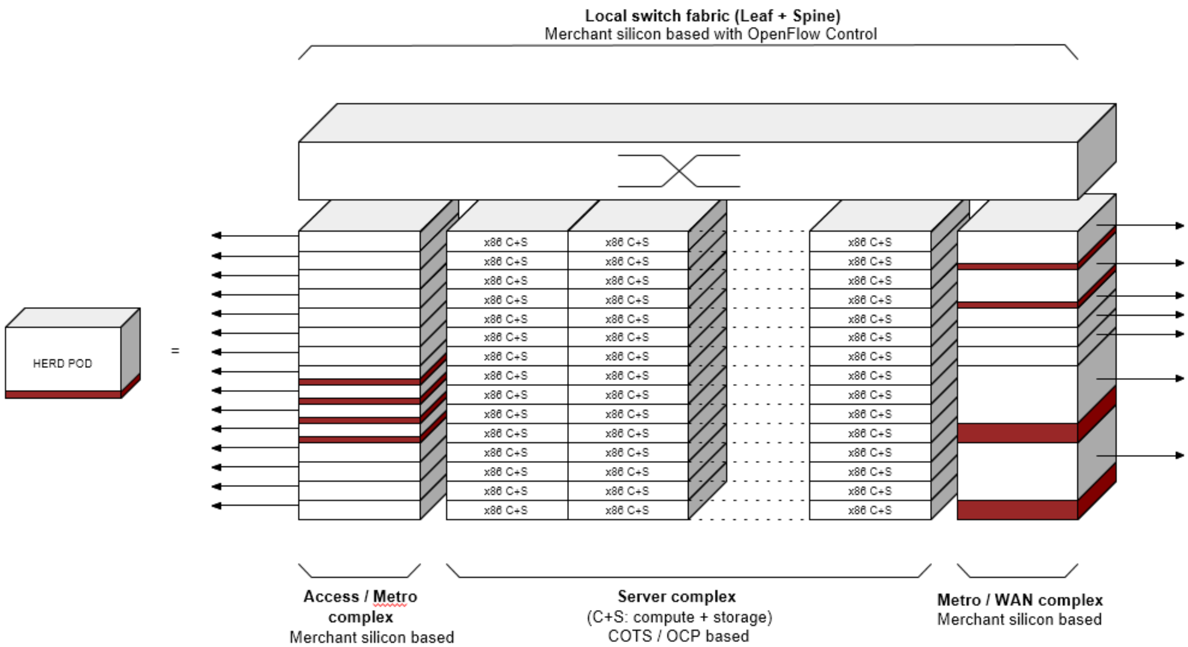


**Figure 5 - HERD POD Example**

As conventional systems become disaggregated, the virtual network functionality can be assigned to Common Off The Shelf (COTS) x86 servers on an as needed basis. These virtual network functions running on servers eliminate the different vendors' equipment with the accompanying delivery intervals and maintenance.

The User Routers are replaced by vRouters running on servers. A new vRouter can be configured and be running in minutes.

As discussed in the CORD PON example, virtual OLTs and virtual Subscriber Gateways can be deployed on x86 servers in the HERD POD.

## 4.3. Virtual CCAP Use Case

As the CCAP devices become more disaggregated, virtual CCAP (vCCAP) functionality including the Cable Control Plane, Subscriber Management, MAC, Service Flow Engine, DOCSIS Policy Engine and DOCSIS Provisioning System could be run on the servers in the HERD POD.  As subscriber capacity increases, new virtual CCAP servers could be configured and new Remote PHY devices deployed to accommodate the increased bandwidth and subscribers.

 

Once the virtual CCAP functions become available, the virtual CCAP POD connects via 10 gigabit Ethernet fiber to the Remote PHY Devices (RPD) which connects to the Cable Modems (CM) over coax.

Vendors are now delivering Remote PHY Devices (RPDs) Devices that can be deployed closer to the user locations using optical fiber based distributed access architectures.

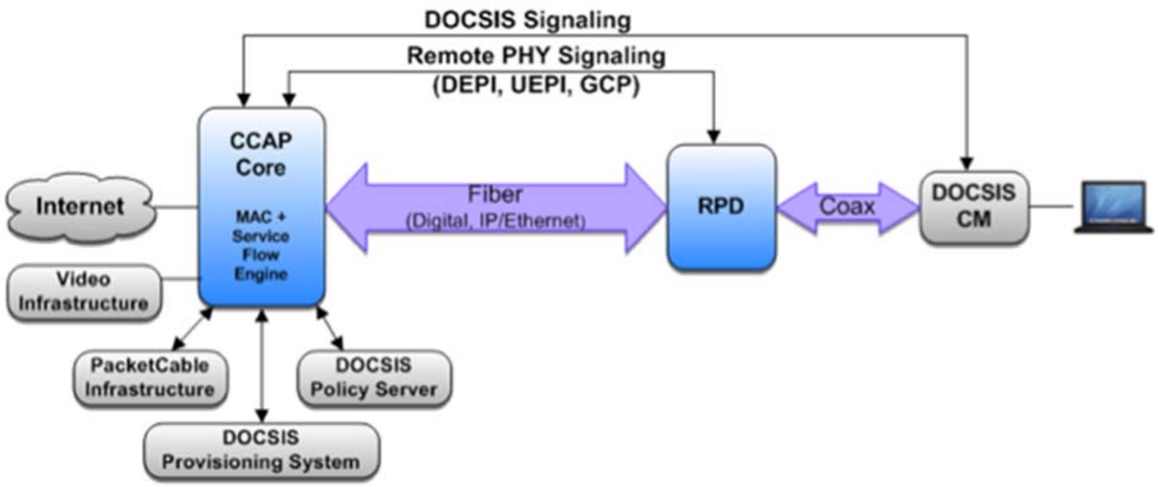


**Figure 6 - CCAP Remote PHY architecture**

## 4.4. CableLabs CCAP Specifications

CableLabs has completed many specifications defining a common CCAP architecture and functional definitions.

1. CableLabs Converged Cable Access Platform Architecture Technical Report
2. CableLabs Operations Support System Interface Specification (OSSI) defines the requirements necessary for the Configuration, Fault Management, and Performance Management of the Converged Cable Access Platform (CCAP) system
3. CableLabs MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.0
4. The CableLabs Remote PHY technology is detailed by six specifications including:
    - The System Specification that describes System level requirements such as initialization sequences and security.
    - The R-DEPI and R-UEPI specifications that describe the downstream and upstream pseudowires and the L2TPv3 control plane.
    - The General Control Protocol (GCP) specification that defines a protocol used for configuration of Remote PHY Devices (RPD).
    - The R-DTI specification that defines the timing interface between the CCAP-Core and RPD.
    - The R-OOB specification that defines support for the SCTE55-1 and 55-2 out of band data for video applications.

 

### 4.5. Webscale Head End Architecture

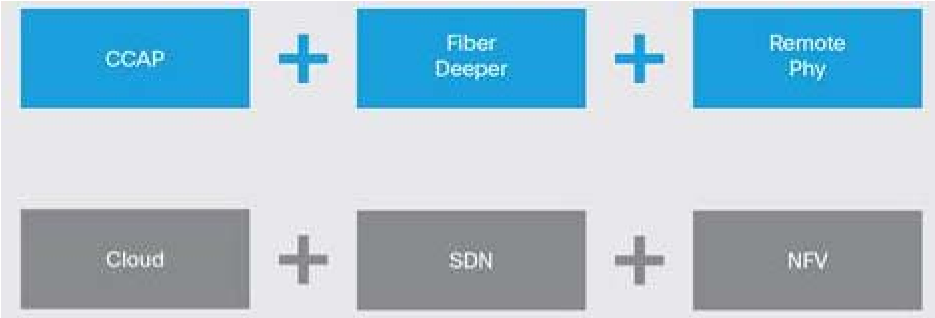


**Figure 7 - Webscale Head End Architecture**

Virtualizing of CCAP onto HERD PODs will help move the Head End into the Webscale world bringing datacenter economies and fast service deployment.

# Conclusion

HERD is a significant milestone in bringing cost effectiveness and agility to the Head End.

HERD is designed to replace today's purpose-built hardware devices with their more agile software-based counterparts.

HERD provides a new network architecture enabling:

- Reduction of CAPEX/OPEX and Total Cost of Ownership (TCO)
- Driving complexity out of the network
- Improving network agility
- Increasing modularity

Data Center virtualization brought in the Cloud Computing era. HERD is the beginning of the Elastic Cloud Network era.

# Abbreviations

| | |
|---|---|
| 10G EPON | 10Gbps Ethernet Passive Optical Network |
| ACCaaS | Access as a Service |
| BNG | Broadband Network Gateway |
| CAPEX | Capital Expense |
| CCAP | Converged Cable Access Platform |
| CDN | Content Delivery Network |
| CMTS | Cable Modem Termination System |
| CORD | Central Office Re-architected as a Datacenter |
| COTS | Common Off The Shelf |
| CPE | Customer Premise Equipment |
| DOCSIS | Date Over Cable Service Interface Specification |
| DWDM | Dense Wave Division Multiplexing |
| EPON | Ethernet Passive Optical Network |
| GCP | General Control Protocol |
| GPON | Gigabit Passive Optical Network |
| GRE | Generic Routing Encapsulation |
| HERD | Head End Re-architected as a Datacenter |
| HFC | hybrid fiber-coax |
| INTaaS | Internet as a Service |
| MAC | Media Access Control |
| NFV | Network Functions Virtualization |
| NIC | Network Interface Card |
| OCP | Open Compute Project |
| OLT | Optical Line Terminal (PON) |
| ONOS | Open Network Operating System from ON.LAB |
| ONU | Optical Network Unit (PON) |
| OPEX | Operations Expense |
| R-DEPI | Remote Downstream External PHY Interface |
| R-DTI | Remote DOCSIS Timing Interface |
| R-OOB | Remote Out of Band |
| R-UEPI | Remote Upstream External PHY Interface |
| RPD | Remote PHY Device |
| SCTE | Society of Cable Telecommunications Engineers |
| SDN | Software Defined Network |
| SUBaaS | Subscriber as a Service |
| vCPE | virtual Customer Premise Equipment |
| VIM | Virtual Infrastructure Manager |
| VM | Virtual Machine |
| VNF | Virtual Network Function |

| | |
|---|---|
| vOLT | virtual OLT vnf |
| VPN | Virtual Private Network |
| vRouter | virtual Router vnf |
| vSG | virtual Subscriber Gateway vnf |
| XaaS | Everything as a Service |

# Bibliography & References


1. Delivering a Software-based Network Infrastructure. Krish Prabhu. AT&T Labs (October 2015).
2. Atrium: A Complete SDN Distribution from ONF. https://github.com/onfsdn/atrium-docs/wiki (2016).
3. OpenStack: Open Source Cloud Computing Software. https://www.openstack.org/ (2016).
4. ONOS: Towards an Open, Distributed SDN OS. HotSDN 2014. P. Berde, M. Gerola, J. Hart, Y. Higuchi, M. Kobayashi, T. Koide, B. Lantz, B. O'Connor, P. Radoslavov, W. Snow, G. Parulkar (August 2014).
5. XOS: An Extensible Cloud Operating System. ACM BigSystems 2015. L. Peterson, S. Baker, A. Bavier, S. Bhatia J. Nelson, M. Wawrzoniak, M. De Leeneer, and J. Hartman (June 2015).
6. CORD: Re-inventing Central Offices for Efficiency and Agility. http://opencord.org (2016).
7. Network Functions Virtualization—An Introductory White Paper. SDN and OpenFlow World Congress (October 2012).
8. CORD: Central Office Re-Architected as a Datacenter, Larry Peterson, Open Networking Lab, https://wiki.opencord.org/pages/viewpage.action?pageId=1278047
9. CCAP Operations Support System Interface Specification, http://www.cablelabs.com/specification/ccap-operations-support-system-interface-specification/
10. CCAP OSSI Specification, http://www.cablelabs.com/specification/ccap-ossi-specification/
11. Distributed CCAP Architectures Overview Technical Report, http://www.cablelabs.com/specification/distributed-ccap-architectures-overview-technical-report/
12. CCAP Architecture Technical Report, http://www.cablelabs.com/specification/ccap-architecture-technical-report/
13. Remote PHY Specification, http://www.cablelabs.com/specification/remote-phy-specification/
14. Remote PHY OSS Interface Specification, http://www.cablelabs.com/specification/remote-phy-oss-interface-specification/