Remote PHY 2.0

The Next Steps For Remote PHY Technology

A Technical Paper prepared for SCTE•ISBE by Pawel Sowinski, Andy Smith and Tong Liu.

Pawel Sowinski
Principal Engineer
Cisco Systems Inc.
psowinsk@cisco.com

Andy Smith
Principal Architect
Cisco Systems Inc.
andsmit@cisco.com

Tong Liu
Principal Engineer
Cisco Systems Inc.
tonliu@cisco.com

© 2019 SCTE•ISBE and NCTA. All rights reserved.

Table of Contents

List of Figures

© 2019 SCTE•ISBE and NCTA. All rights reserved.

Introduction

In the period between the initiation of the Remote PHY program and the deployment of phase one significant changes have taken place in cable operator services, and network deployments and also in the technologies that have become available to build and operate systems such as Remote PHY. It is appropriate to consider the next generation requirements R-PHY will need to solve and what are the best tools to use to do this.

As the first phase of Remote PHY is moving towards technological maturity, the time is right to plan the next steps in its evolution. The paper proposes a strategy for enhancing the Remote PHY architecture by examining a set of its real and perceived issues, and by suggesting a how to effectively tackle them.

The paper focuses on the following technical issues:

- **Cloud APIs and Automation**: How to transition the R-PHY control plane towards mature, open model driven network management protocols such as NETCONF, RESTCONF or gNMI.
- **Manageability**: A recommendation for model driven telemetry for Remote PHY.
- **Latency**: How to eliminate latency issues resulting from the physical distances between R-PHY system components by incorporating a DOCSIS upstream bandwidth scheduler into the RPD.
- **Operation in Multi-Service Networks**: The application of modern network standards, such MPLS and Segment Routing to R-PHY data plane transport.

The paper presents an in-depth technical analysis and discusses the utility and economic value of the proposed enhancements. The paper demonstrates that the proposed functional advancements, taken together represent the next generation of R-PHY architecture, Remote PHY 2.0.

The paper does not explain the details of R-PHY architecture. The paper is written with the assumption that the reader has at minimum a rudimentary familiarity with Remote PHY. The necessary background information can be found in [RPHYTR] and [RPHY] .

Content

1. The Opportunity

Remote PHY technology has finally entered the phase of wide scale deployments. At the time of this paper's writing, several cable operators are providing commercial service based on R-PHY technology to hundreds of thousands of subscribers. The multi-year Remote PHY standardization efforts led by CableLabs are drawing to a conclusion. Soon, the R-PHY project at CableLabs will enter the maintenance phase. The working group's focus will shift towards fixing specification bugs rather than the definition of new functionality.

The confluence of these events creates a perfect opportunity to take step back, take a critical review of the R-PHY technology, assess its weaknesses and gaps and devise a strategy to best address these issues. This paper presents a menu of options for several selected new R-PHY features. Our intent is to initiate a conversation within the industry about the future direction of R-PHY technology. Therefore, the list of R-PHY 2.0 features discussed within this paper is open to further additions and changes.

2. Introducing Remote PHY 2.0

In this paper we refer to the existing R-PHY technology and specifications as R-PHY 1.0 or simply 1.0.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

"Remote PHY 2.0" is nothing more than a convenient name chosen as a common label applied to the set of new R-PHY architecture options proposed in this paper. We don't claim that any of these features cannot be added to existing specifications and products without such a label. We believe however, that there are tangible benefits and a convincing argument to be made for separating these options from current R-PHY technology and packaging them under a new version label. The primary concern is the ability of the existing products to support these new options. The following factors also need to be considered:

- The proposed features are interdependent. For example, the proposed model driven telemetry relies on RPD supporting the data driven control plane.
- The proposed technical solutions do not constitute incremental development. They offer replacement for currently utilized techniques and may not provide backward compatibility.
- It is beneficial to logically separate these options because of the large scope of changes to the involved software infrastructure.

3. Cloud Friendly Control Plane

In this section we propose a strategy to replace the main control protocol deployed in R-PHY 1.0 architecture. First, we describe the existing R-PHY control protocol and analyze its strengths and weaknesses. Later we detail the approach to upgrade the control protocol and how to minimize the transition impact on the existing R-PHY system. Finally, we explain the technical and business benefits of the proposed transition.

3.1. What is R-PHY Control Protocol?

In a Remote PHY Architecture, the integrated Converged Cable Access Platform (CCAP) is separated into two distinct components. The first component is the CCAP Core. The second component is the Remote PHY Device (RPD). The CCAP Core inherits all I-CCAP functions except for the PHY layer which is implemented in the RPD. The CCAP Core and the RPD communicate over a permanent IP connection.

The relationship between the CCAP Core and the RPD resembles a master-slave communication model. The direction of control is from the CCAP Core to the RPD. The CCAP Core remotely controls the functions of the RPD through a protocol which we refer to as the R-PHY Control Protocol.

The R-PHY control protocol incorporates all elements of the FCAPS (Fault, Configuration, Administration, Performance, Security) management framework. In this context the CCAP Core acts as the Network Management System and the RPD acts as the Managed System. There is however a number of important differences in requirements for a typical FCAPS operation and for a R-PHY control protocol with the Core and RPD having a much tighter coupling than a typical FCAPS manager and client. In many instances, the CCAP Core and the RPD operate with a common set of configuration parameters and state information. Whenever the operator, or internal processes in the CCAP Core impose changes to the values of these parameters or state variables, the control protocol needs to coordinate them between the systems, sometimes with tight real-time constraints. For example, when the configuration of a downstream profile changes on an OFDM channel, the change needs to be enacted in both systems by a detailed procedure prescribed by the control protocol.

The R-PHY architecture incorporates a great deal of flexibility in how the CCAP can be functionally decomposed into a set of independent CCAP Cores. For this reason, each RPD is required to provide service to multiple (from one up to 10) CCAP Cores. Serving multiple masters is hard. The R-PHY

© 2019 SCTE•ISBE and NCTA. All rights reserved.

control protocol solves this problem by subdividing the set of managed resources into isolated information silos. Each silo is controlled by a single Core.

For example, a selected CCAP Core, referred to as the Principal CCAP Core, is designated to provide the central management functions such as the initial configuration of the RPD-CCAP Core pairing, the division of managed resource between the Cores, the RPD general configuration, the fault handling, the control over device software upgrades, etc. The Principal Core does not handle any video or data signals. Other CCAP Cores can provide individual CCAP data services, e.g. DOCSIS, or SCTE 55-2 out-of-band service, and only manage the RPD resources dedicated to these services.

3.2. Control Protocol in Remote PHY 1.0

In R-PHY 1.0, virtually all aspects of the master-slave relationship are managed with a protocol commonly referred to as Generic Control Plane/R-PHY Control Protocol (GCP/RCP). R-PHY also relies on several other protocols for narrower purposes, such as Layer 2 Tunneling Protocol Version 3 (L2TPv3) control protocol. In this section we focus solely on the GCP/RCP protocol.

Figure 1 shows an example of GCP connections between an RPD and a set of four CCAP Cores.

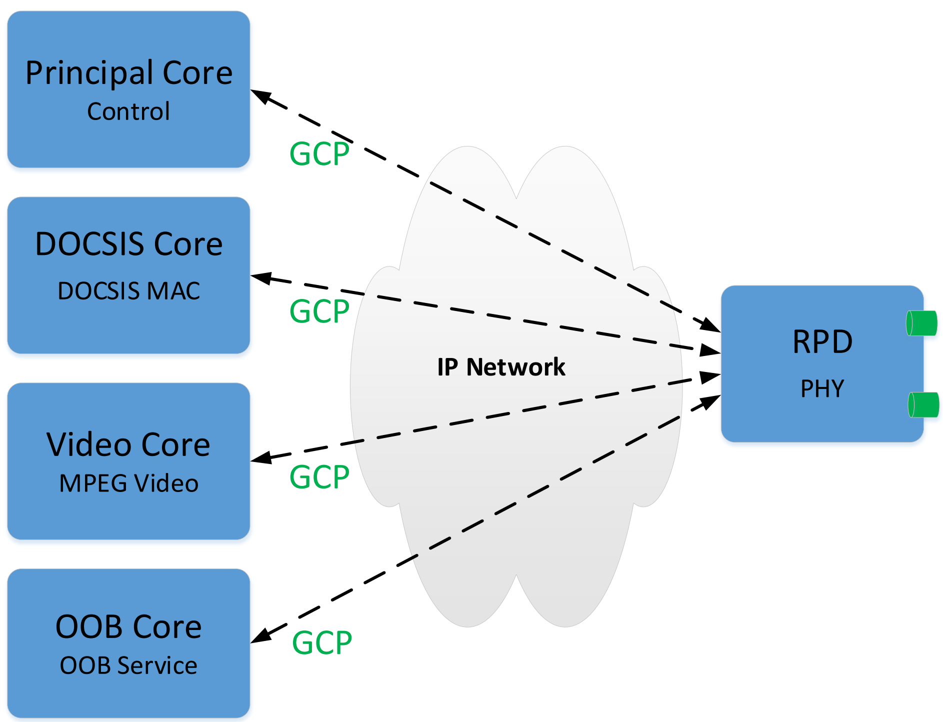

Figure 1 – R-PHY Control with multiple RPD and CCAP Cores.

In the example, as mentioned earlier, the Principal Core provides general management of the RPD. The DOCSIS, Video and OOB Cores control distinct sets of RPD resources associated with the services they respectively provide.

3.2.1. GCP Protocol Stack

Figure 2 presents a typical control protocol stack of an RPD. The GCP protocol relies on TCP for reliable transport and on the IPsec suite for security protection. The formats of messages and their exchange rules

© 2019 SCTE•ISBE and NCTA. All rights reserved.

are specified in [GCP] . The upper layers implemented by the RCP are documented in [RPHY] and in [RPHYOSS].

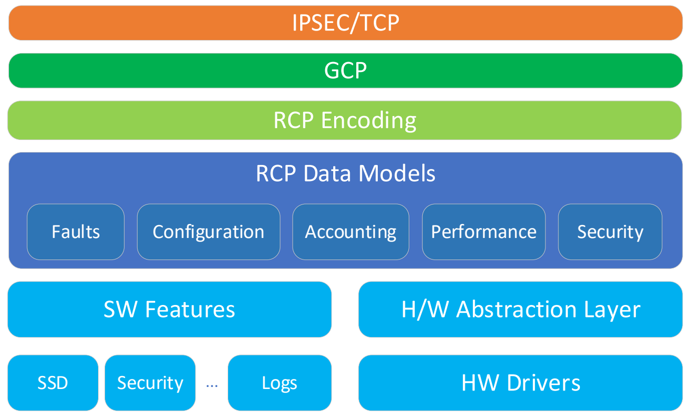

Figure 2 – RPD 1.0 control protocol stack.

The RCP protocol design was intended to mirror the DOCSIS protocol deployed for communication between the CMTS and Cable Modems. RCP operates as an abstraction layer over the foundation provided by the GCP protocol. Just like the DOCSIS protocol, RCP carries data in hierarchically organized Type-Length-Value (TLV) tuples.

Unlike DOCSIS, RCP defines rules for structured access to an RPD data model represented by a hierarchy of application specific object-TLVs. Object-TLVs form a tree in which each node has a type and either a value or a set of child nodes. The detailed specifications of objects-TLVs serve two roles, the formal definitions of the elements of the information models and how to encode the information in exchanged messages. The definition of RCP object-TLVs includes semantics and syntax, default values, attribute units as well as range and access constraints. The specifications render selected object classes as UML diagrams to provide an informal, visual representation of the information models of object-TLVs constructs. An extensive set of distinct objects describes RPD capabilities. Lastly, RCP also directly incorporates selected DOCSIS messages for configuration of specific DOCSIS channels' operational parameters.

The RCP provides the CCAP Core with the ability to remotely manage properties of modelled objects such as RF channels, RF ports, performance variables, etc., maintained by the RPD. The relative complexity of R-PHY information models, in our opinion, reaches at best modest levels.

3.2.2. Transactions

Just like most network management protocols, RCP allows the CCAP Core to perform CRUD (create, read, update, and delete) operations on object classes and individual attributes. RCP configuration operations are transactional within sets of grouped objects. The protocol also provides a means for the RPD to send asynchronous notifications to CCAP Cores to inform them about defined events, e.g. when the state of a modelled attribute changes or to report errors.

3.2.3. Extensibility

The RCP mirrors the DOCSIS protocol extensibility mechanism through vendor-proprietary extensions at the top of the TLV hierarchy. Some implementations also support proprietary extensions at lower levels in the hierarchy.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

3.2.4. Performance Requirements

The volume of traffic exchanged over the GCP/RCP connection is highest during the initial RPD configuration. The size of a configuration set for a typical RPD is modest. It can be measured in 10s of kilobytes. After the initialization, the volume of GCP traffic varies; it depends primarily on the level of status information retrieved from the RPD by the CCAP Core(s). The specifications do not impose limitations on the traffic volume. We can assume that the volume of GCP traffic is not an issue we need to be concerned with.

The RCP operates with few real-time constraints. The majority of RCP protocol interactions require a response within one second. The most stringent real-time requirements are imposed by procedures for dynamic updates to DOCSIS channel parameters such as DOCSIS OFDM profiles or DOCSIS upstream channel parameters. During these procedures, the GCP/RCP transport carries selected DOCSIS messages from the DOCSIS Core to the RPD. The procedures require coordination with parallel procedures conducted between the CMTS Core and Cable Modems. For example, during upstream channel parameters change, the RPD is mandated to process received UCD message in less than 50 msec.

To summarize, the RCP defines abstract information models and a set of protocol rules for CRUD operations on objects from these models. The RCP closely follows the operational principles of network management protocols and is subject to few performance or stringent real-time requirements.

3.3. R-PHY 1.0 Control Protocol Status

At the functional level, few complaints can be made about the R-PHY 1.0 control protocol. The protocol has demonstrated sufficient flexibility to enable effective replication of I-CCAP features in the distributed R-PHY environment without compromising functionality or performance. RCP 1.0 also offers very compact encoding.

Remote PHY 1.0 specifications certainly meet the criteria of an open standard. The specifications have been developed by a working group open to any willing participants. CableLabs processes make the specifications available, royalty free for download to the public. R-PHY 1.0 specifications, including the control protocol have been developed with the goal of ensuring multivendor interoperability. Successful deployments of R-PHY systems with interoperable components from several vendors have proven that the R-PHY 1.0 control protocol has successfully achieved this goal.

The R-PHY 1.0 control protocol is deployed just in one, relatively narrow application field, the Remote PHY architecture. While the development of the R-PHY 1.0 control protocol for a system does not represent a high technological barrier, the number of existing implementations is limited to a handful of equipment vendors and cable operators. The ecosystem of applications and support tools available to test GCP/RCP is scarce or non-existent.

The R-PHY 1.0 control protocol was designed to implement the RPD control functions efficiently and to meet certain DOCIS real-time requirements. The reliance on a purposely developed and narrowly deployed protocol as well as the inherited real-time constraints made the R-PHY 1.0 control protocol unique.

This uniqueness is also the Achilles heel of the R-PHY 1.0 control protocol because it translates into a set of business issues such as the difficulty in testing and validation, the slower adoption curve and in the end into higher OPEX for the operators.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

Without a doubt, the availability and the maturity of the development and test ecosystems has been a contributing factor to some of the interoperability problems encountered during initial deployments. The root causes of interoperability problems can be notoriously complex. This is especially true in a multivendor environment, where vendors cooperate developing on ever-changing specifications but also fiercely compete against each other in the marketplace. The causal analysis of the interoperability issues lies beyond the scope of this paper.

Despite the protocol design based on open CableLabs specifications, the developer community has currently no Open Source code project at its disposal to stop reinventing software each time they need to build a new product.

Another issue is the maturity of the 1.0 protocol definition. We are referring here to the 1.0 control protocol rules, not the information models on which the protocol operates. Design issues are uncovered, new protocol rules are added, existing rules are modified or clarified in each release of the R-PHY specifications. Even though the rate at which the protocol changes are introduced is consistently decreasing, these processes are likely to continue for some time.

Finally, we need to consider the significant changes in the cable operator infrastructure environment in which the R-PHY is implemented. The R-PHY 1.0 protocol was developed with requirements of coherent integration with the physical CCAP infrastructure. As operators transition towards a cloud-based infrastructure, the requirements for the protocol shift as well. The R-PHY control protocol needs to be reimagined to better conform to the new cloud environment.

3.4. Remote PHY 2.0 Control Protocol

In this section we describe the proposal for the control protocol for R-PHY 2.0. We explain the goals, the methodology, discuss the feasibility and demonstrate how the proposal addresses the issues with R-PHY 1.0 control protocol identified in the preceding section. Finally, we discuss the options for a transition from 1.0 to 2.0 control protocol.

We propose to transition the R-PHY control protocol away from GCP/RCP towards modern, YANG model driven protocols. YANG, which stands for Yet Another Next Generation, is simply a better choice, especially for cloud APIs. This is not a new idea even in the domain of Distributed CCAP Architectures. The first example is CCAP Config. CableLabs specifications have been publishing YANG based APIs for CCAP configuration for almost a decade. Most recently, the Flexible MAC Architecture (FMA) group has embraced a similar approach using YANG for the management APIs of the FMA MAC Network Element.

YANG is a data-modeling language used to describe network device configuration and operational data developed by the Internet Engineering Task Force (IETF). YANG models the hierarchical organization of data as a tree in which each node has a name and either a value or a set of child nodes. YANG provides clear and concise descriptions of the nodes and of the interaction between them. Details about YANG can be found in **Error! Reference source not found.** and **Error! Reference source not found.**.

Moving to a YANG models will address the R-PHY 1.0 control plane issues identified previously, enable cloud friendly tools and automation, and improve system manageability, testability and multivendor interoperability.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

3.4.1. Why YANG?

Over the past decade, YANG became a universally adopted standard for modeling of APIs for management of physical and virtual network elements. YANG emerged as the default choice for network management APIs when automation, agility, and scaling are the key requirements. The proliferation of YANG based programmatic interfaces extends into the Internet of Things (IoT) space, and outside of networking, into fields such as medical, vehicular and even aeronautical technology.

The IETF has been developing standards with YANG as the data modeling language for all elements of FCAPS framework. To date, within the IETF, about a hundred of YANG modules have been adopted for standard track and hundreds more are circulating as drafts.

The wide availability of mature development ecosystems, including Open Source code libraries, toolchains and applications are the key factors driving YANG's adoption. The development organizations can pick from dozens of Open Source rooted tools and commercial systems specifically developed to help with YANG model creation, validation and testing. The functionality supported by tools includes conversion from other modeling methods and even automatic API source code generation in modern programming languages.

3.4.2. Proposed Methodology

3.4.2.1. Data Model Translation

The first step towards R-PHY 2.0 control protocol is the formalization of RPD data models in YANG. The RPD data models can be translated from the current representation in RCP object-TLVs to equivalent YANG modules. Such a translation, in our opinion, is feasible, if not completely straightforward or even somehow mechanical in nature. The existing hierarchy of RCP object-TLVs, and their constraints can be directly replicated into YANG. The translation does not need to result in a perfect mirroring of the existing RCP models. Where optimizations are appropriate, the formalization process can include a desired level of refactoring. Selected YANG modules developed by the CableLabs FMA working group could be adopted for reuse in R-PHY 2.0. The product of the translation will be a set of YANG modules representing the same managed objects of an RPD as those embedded in RCP 1.0.

3.4.2.2. Relaxation of Real Time Requirements

To eliminate the most stringent real-time requirements on the control protocol, we propose to remove DOCSIS MAC Management messages from the R-PHY 2.0 control protocol. The UCD, OCD and DPD messages are sent from the CMTS to Cable Modems in-band, as packets embedded in downstream data streams. In R-PHY 1.0 these packets simply pass transparently through the RPD. [MULPI] specifies precise procedures by which DOCSIS CMs operate on these messages. An R-PHY 2.0 compatible RPD can snoop the messages from the data plane stream sent to CMs and participate in the channel change procedures just like Cable Modems.

3.4.2.3. Protocol Selection

The next step is the selection of the protocol over which the CCAP Core and the RPDs exchange information. Three protocols have emerged as the favorite choices for operation with YANG defined APIs. These protocols are NETCONF, RESTCONF and gNMI.

- The Network Configuration Protocol (NETCONF) defines a mechanism for manipulating configuration data and for retrieving operational data. NETCONF carries configuration and

© 2019 SCTE•ISBE and NCTA. All rights reserved.

operational data encoded in XML over a reliable transport. NETCONF's definition can be found in **Error! Reference source not found.**.

- RESTCONF stands for Representational State Transfer Configuration Protocol. RESTCONF is a REST-like protocol which relies on HTTP protocol and methodology. Request and response data can be represented in XML or JSON format. RESTCONF is described in **Error! Reference source not found.**.
- gNMI is a network management protocol developed primarily by Google. gNMI provides the mechanisms to manage the configuration of network devices, and also to retrieve operational data. gNMI typically relies on Thrift or Google Protocol Buffers (GPB) for data encoding and serialization. gNMI is specified in [GNMI-SPEC].

It's outside of the scope of our paper to analyze the technical differences between these protocols. A good example of such analysis between NETCONF and RESTCONF can be found in [CLAISE]. We will however examine the common properties of these protocols to consider their suitability as the replacement for GCP/RCP 1.0.

All of these protocols have been adopted by major network equipment providers and gained strong industry support in both physical and virtualized applications. A rich ecosystem of tooling and test equipment benefits from participation by a wide vendor community. Open source code libraries for the client and the server side are available to accelerate development, lower costs and ensure seamless interoperability.

We don't claim that there are not any differences between these protocols. On the contrary, the design of each protocol allows it to perform certain tasks better while having disadvantages in other areas. They have been developed to solve different issues with a different set of goals. However, the following properties are common to all the protocols. Each protocol supports a superset of the primitives offered by RCP 1.0. All protocols offer security protection features that can be seamlessly integrated into the CableLabs public key infrastructure security defined for R-PHY 1.0. Mirroring a capability of other networking devices, a R-PHY 2.0 compliant RPD could even have the flexibility to simultaneously operate all three protocols on top of common YANG models.

The multi-protocol RPD stack is shown on Figure 3.

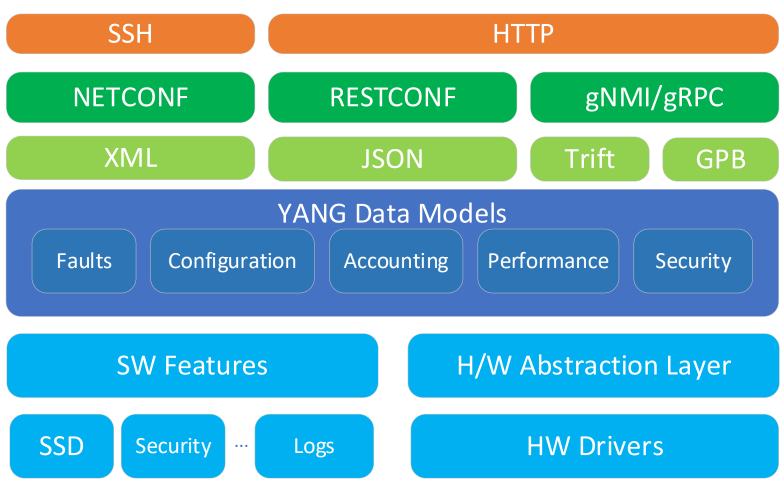

Figure 3 – RPD 2.0 control protocol stacks.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

The choice of which protocol is enabled on RPDs in any particular deployment could be left to operators and driven by the requirements of the CCAP Cores and other OSS infrastructure systems deployed in their networks. Another choice is to narrow the RPD mandate to just one, carefully selected protocol, to maintain the low development costs and reduce the complexity of the RPD.

3.5. 1.0 to 2.0 Transition

By applying the translation methodology explained in the previous sections, the R-PHY specifications can establish a precise correspondence between object classes and individual attributes represented in YANG and in RCP object-TLVs. For example, a definition of each YANG leaf attribute could include a cross-reference to an RCP object-TLV. With such a mapping, a 2.0 RPD can support dual 1.0 and 2.0 control protocols. The protocol stack of such an RPD is shown on Figure 4.

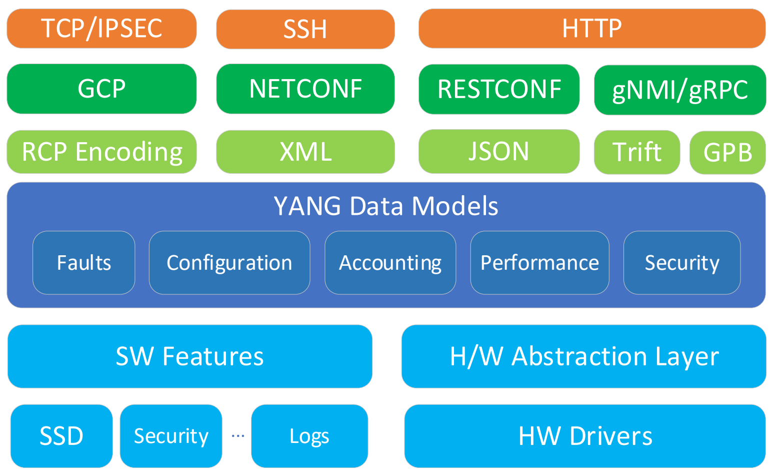

Figure 4 – RPD 1.0 and 2.0 control protocol stacks.

It is worth noting that an RPD with a dual protocol stack may not only allow backward compatible operation with R-PHY 1.0 but can also inter-operate with 1.0 and 2.0 CCAP Cores simultaneously. The operators can utilize this feature as a part of their transitional strategy to R-PHY 2.0. It offers the flexibility to selectively and incrementally upgrade CCAP Cores to version 2.0 protocols. Further, certain CCAP Cores supporting narrow functionality, e.g. SCTE 55-2 OOB, may continue to operate in R-PHY 2.0 environment at 1.0 level until they reach end-of-life.

3.6. Enabling Automation

Studies show that networks are evolving faster today than they have in the previous decades while their OPEX and CAPEX are being continually reduced. The key evolutionary drivers are automation and virtualization. An R-PHY control protocol transition to depend on YANG models and widely deployed, standard-based protocols will align the R-PHY with modern cloud-native technologies. It will also help in addressing all of the RCP 1.0 issues explained earlier in the paper. Few automation tools support GCP/RCP protocol. Many existing cloud automation tools are available, and their APIs are YANG based. Thus, transitioning to a 2.0 control protocol will be the necessary step to more easily integrate with cloud-native CCAP Core systems and automated OSS systems. The results will be the enablement of automation, the acceleration of the network evolution and significant reduction of the total cost of ownership for cable operators.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

4. Model Driven Telemetry

Model-Driven Telemetry (MDT) is a modern technique for monitoring in which operational data is streamed from network devices continuously using a push model. Applications can subscribe to selected elements of YANG data models over a standards protocol such as NETCONF, RESTCONF or gNMI/gRPC. Model driven streaming telemetry allows monitored data to be pushed from the monitored device, e.g. the RPD, to an external collector at a higher frequency than polling, as well as to push data only when a change is recorded. A periodic collection method, when a device pushes data at a defined interval, is better suited to monitoring of frequently changing metrics, e.g. data plane statistical counters. An on-change collection method is a better fit for monitoring infrequently changing data such as state objects, faults or error counters. Through a combination of these methods, MDT provides a highly flexible, efficient communication process for automatic near real-time access to operational data.

In order to stream data from the device the application, or the collector, establishes a subscription to a data set which can be any subset of a device's YANG model. A subscription is a contract between subscription service and a collector that defines the data set to be pushed and the collection methods. Subscription allows clients to subscribe to modeled data. The device pushes the data to the collector as per agreed contract.

Figure 5 shows how MTD could be integrated with the proposed R-PHY 2.0 protocol stack.

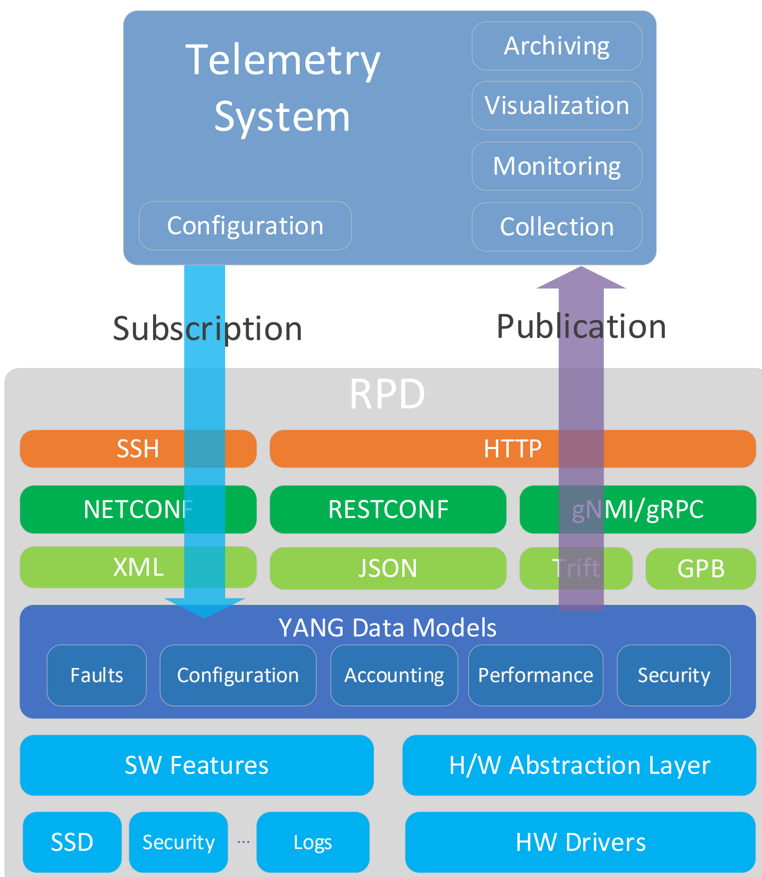

Figure 5 – Model Driven Telemetry with R-PHY 2.0 Protocol Stack.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

The expanding popularity of MDT is driven by many factors, including its simplicity, the reliance on open standards and YANG models, the broad availability of commercial solutions and Open Source software for all elements of the development ecosystems. For example, well-known Open Source components such as the Apache Kafka messaging bus and the ELK stack (Elasticsearch, Logstash, and Kibana), can be used to build a reliable MDT infrastructure and automated systems for processing and visualization of received information.

The cable industry is already quite familiar with model driven telemetry. MDT originated as a data collection technology for cloud-based infrastructure, but it is also implemented on modern CCAP hardware-based platforms. Several cable operators already deploy MDT collection and monitoring systems in their networks. Including MTD in the R-PHY 2.0 feature set will provide a simple, yet extremely powerful technology for pushing useful metrics from where they are generated to where they are consumed, fitting well into the operators' publish/subscribe (PUB/SUB) model. MTD will become a foundation for modern monitoring of the real-time health of the RPD population as well as of the services it provides.

Finally, we examine a simple example which shows how MTD technology could be applied to monitor a vital metric of the health of the R-PHY data plane. In R-PHY, the user data is transported over L2TPv3 pseudowires. In each consecutively transmitted packet, the L2TPv3 transmitter increments a sequence number embedded the packet header. By examining the continuity of the sequence numbers, the receiver can detect when the network drops packets in-between the transmitter and the receiver. In such a case the receiver (i.e. the RPD) increments a statistical counter of sequence errors for the corresponding pseudowire. Any change to the values of the sequence error counters provides an immediate indication of a potential issue with the health or the performance of the network. The RPD telemetry agent could be configured to monitor changes to the values of the sequence error counters and stream the counters' values to the MTD collector whenever they change. An application within the MTD collection system could then in real-time analyze the received data and alert a network manager about on-going network problems.

5. R-PHY with Remote Upstream Scheduler

The R-PHY remote upstream scheduler is an architectural option that moves the real-time DOCSIS upstream scheduling function together with the PHY element to the RPD. It is suitable for providing low latency DOCSIS transport over long distance R-PHY deployments.

The location of the upstream scheduler has been part of the R-PHY architecture consideration since the very beginning of the R-PHY development. The optimum location choice depends on both business and technical reasons.

For R-PHY 1.0, the primary goal is to enable DCA by replacing the analog optical link between the CCAP and the Node with a digital link. Just with this initial step, cable operators would be able to get better SNR performance, pull the fiber deeper, rebuild the plant and cut a large service group into much smaller ones. All these can be achieved by simply moving the PHY element out of the CCAP Core, while keeping all MAC elements including the DOCSIS upstream scheduler centralized. This also allows the operators to leverage the existing CCAP MAC functions to simultaneously support both integrated PHY and Remote PHY for a quick and smooth transition to DCA.

The main technical reason for applying centralized upstream scheduling to R-PHY 1.0 is because with a 2 msec MAP interval, the CIN delay is not a dominant latency factor within the 100-mile I-CCAP HFC

© 2019 SCTE•ISBE and NCTA. All rights reserved.

reference range. In this case R-PHY 1.0 is equivalent to an I-CCAP in terms of the upstream request-grant (REQ-GNT) latency.

As the network keeps transitioning to DCA, there are, however, reported cases where the CIN is stretched beyond the 100-mile mark, for reasons such as hub-site consolidation that relocates a CCAP Core to the central headend or a regional data center. Meanwhile, driven by the new low latency applications, such as cloud gaming and mobile xhaul, the DOCSIS REQ-GNT protocol is being tightened to shorter MAP intervals, such as 1 millisecond on DOCSIS 3.1 OFDMA channels. In such circumstances, the CIN delay could be exposed as a significant factor affecting the REQ-GNT latency. This problem can be solved by simply moving the REQ-GNT handling to the RPD, which will effectively cut the CIN out of the REQ-GNT loop as shown in the Figure 6 below.

The remote upstream scheduler will be proposed as a R-PHY 2.0 feature for low latency support. The specification will focus on the remote upstream scheduling interface definition to allow the CCAP Core to work with remote upstream schedulers from different RPD vendors. The remote upstream scheduling management interface will be a data model driven, taking advantage of the new R-PHY 2.0 control plane infrastructure.

The overall remote upstream scheduling definition will provide R-PHY 2.0 with the flexibility for load balancing the upstream scheduling tasks between the CCAP Core and the RPD, enabling backward compatibility with 1.0 RPDs, and ultimately the ability to maximize both the centralized and distributed computation resource to achieve low-latency at high system efficiency.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

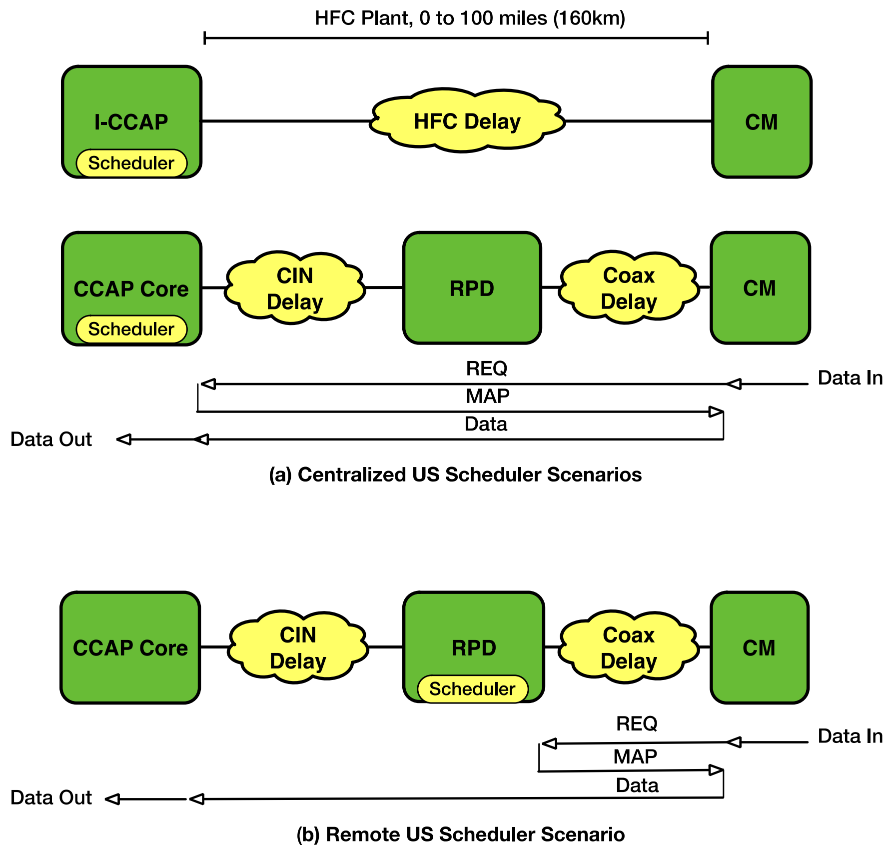

Figure 6 – Centralized vs. remote R-PHY upstream scheduling options for R-PHY1.0 and R-PHY 2.0.

Additional details about the Remote Scheduler can be found in [RemoteScheduler].

6. Data Plane Improvements

Since the early days of M-CMTS, the MAC and PHY split has been enabled through the application of L2TPv3 tunneling technology. The (DEPI) tunneling scheme utilized L2TPv3 over IPv4 or IPv6 as a simple, lightweight and standards-based encapsulation to enable scalable connectivity between the CMTS and the edge QAM device.

L2TPv3 was utilized again with the advent of Remote PHY, to enable a MAC and PHY split. DEPI was used on the downstream and the upstream direction (called Upstream External PHY Interface or UEPI) was added. In the Remote PHY architecture, the L2TPv3 UEPI tunnels are unicast (point to point) from the RPD to the CCAP Core. In the downstream direction, the L2TPv3 DEPI tunnels are either unicast from the Core to a single RPD or multicast to multiple RPDs. Multicast DEPI tunnels permit an efficient allocation of CCAP Core resources across many DOCSIS service groups and are ideal when adapting existing centralized hardware CCAP devices to a Remote PHY deployment. The network that DEPI and UEPI tunnels transit is called the CIN (Converged Interconnect Network).

While the existing L2TPv3 DEPI and UEPI tunneling schemes have served both M-CMTS and Remote PHY well, as cable networks and CCAP software evolve it might be prudent to re-examine the tunneling architecture of a MAC/PHY split in cable. The current architecture leaves a few things to be desired:

© 2019 SCTE•ISBE and NCTA. All rights reserved.

- It is difficult to traffic engineer IP tunnels without resorting to another, additional encapsulation. Specifically, traffic engineering refers to redistributing traffic loads across different paths.
- ECMP (Equal Cost Multi Path) load balancing in the CIN can be a challenge, primarily due to the lack of decipherable entropy in the payload of the packet. ECMP is a hop-by-hop algorithmic load balancing mechanism that depends on sufficient input to the algorithm, known as 'entropy', to make an efficient decision. Because DEPI and UEPI tunnels carry encrypted DOCSIS traffic and have a common set of IP addresses, a standard router has limited visibility in to how to best make a load balancing decision.
- A high caliber network architecture is required. It is incumbent that the CIN be engineered akin to circuits, not paths. This is a packet transport network, not an internet routing network. Remote PHY architecture requires that packets arrive in order, no high priority packets be dropped and symmetric latency (round trip times) be maintained.
- L2TPv3, despite its versatility, is still a niche tunneling protocol in the industry. Service provider networks have generally embraced Multi-Protocol Label Switching (MPLS) as the tunneling technology of choice. If you look at the protocol diagrams for R-PHY 1.0, it allowed for an expansion to include MPLS.

The state of cable access architecture is evolving quickly to encompass software, cloud native technology and multi-modal access methods. Software implementations of CMTS infrastructure permit a horizontal scaling of DOCSIS resources, which eliminates the need for IP multicast in DEPI tunnels for DOCSIS. Many cable operators would like to build out one CIN or Ethernet aggregation network, of which DOCSIS technology is but one method for last mile connectivity. Also, service provider control and data planes, automation methods, telemetry retrieval and system programmability have all evolved significantly since M-CMTS and Remote PHY were first proposed.

This confluence of events and technology means it's a good time to revisit the CIN architecture and the way in which the Remote PHY system interconnects the CCAP Core with RPDs. This paper explores two options. In both cases, MPLS technology will play a much larger role in the transport of Packet Streaming Protocol (PSP). Option 1 is to simply encapsulate the existing DEPI and UEPI unicast into an MPLS LSP (Label Switch Path). Option 2 is to eliminate the L2TPv3 tunneling layer altogether, and tunnel PSP natively over MPLS.

Option 1: PSP over L2TPv3 over MPLS

In this option, the existing Remote PHY DEPI and UEPI tunnels are further encapsulated into MPLS, typically by the first hop router in the CIN. Once in an MPLS LSP, Remote PHY traffic can be merged and integrated into a multi-purpose CIN.

This option has the benefit of maintaining backward compatibility with existing CCAP and RPDs, including the L2TPv3 control channel. While this is the easiest option from a Remote PHY architecture change perspective, and it pushes any MPLS integration work to a CIN engineering exercise, it is suboptimal from a life-of-a-packet perspective. It is generally unadvisable to deploy networks with multiple layers of encapsulation, especially when the final encapsulation (in this case MPLS) is capable of carrying the ultimate payload (in this case Remote PHY PSP). L2TPv3 in this case is simply redundant and adds unnecessary complexity, MTU size, and administrative complications into the architecture.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

Option 2: PSP over MPLS

In this option, R-PHY PSP is natively transported over MPLS. This is a unified tunneling approach, in that one tunneling mechanism transports not only Remote PHY traffic but also any other traffic the CIN may be called upon to transport. MPLS provides a payload agnostic method of transporting any type of data across a packet switched infrastructure, and Remote PHY will be one of any number of services.

Benefit	Drawback
RPD traffic is under policy control with XTC (PCE). TE is a function of network management.	RPD and CCAP participate in control plane – must engineer correctly for scale and resiliency
MPLS enables a multi-purpose network	MPLS skillset development needed
MPLS enables scale and automation, alignment with progressive SP network directions	More software running on the RPD and CCAP (control plane only – ISIS, BGP, PCEP)
CIN is free of any remote phy state (labels only)	Software and standards changes for Remote PHY architecture

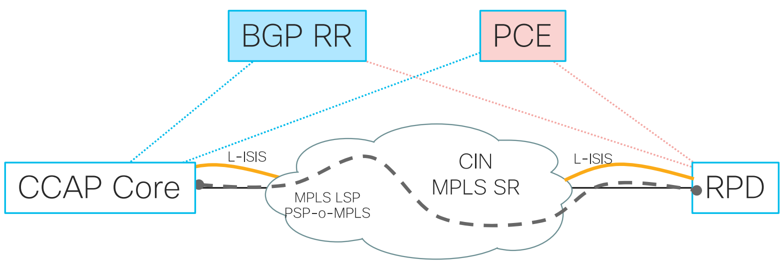

Figure 7 – MPLS Transport for R-PHY.

It is worth noting that the RPD has legacy DOCSIS traffic that does not use PSP as well as video traffic that is based on an MPEG-TS. All this legacy traffic could also be placed natively on an MPLS infrastructure or over-laid with old encapsulation over new encapsulation.

Key to the use of MPLS in this architecture is creating the equivalent of a circuit for PSP packet transport. In Remote PHY it is critical that between the CCAP Core and the RPD, no packets should be dropped, no packets have asymmetric latency, and no packets arrive out of order. Traffic Engineering (TE) permits this type of network to be built by specifying path or path constraints that MPLS encapsulated traffic must follow. Modern approaches to TE, such as Segment Routing, permit a lightweight and scalable approach to delivering the type of network Remote PHY performs best in.

7. Other Functional Improvements

In addition to the four architectural options outlined above, Remote PHY 2.0 as described can provide a foundation for a much broader set of functional improvements. Here we list several such options without describing them in detail.

- **Extended Spectrum DOCSIS** (ESD), a part of DOCSIS 4.0. The process of standardization of ESD has only just begun at CableLabs but it will undoubtedly require changes to R-PHY specifications.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

- **NetFlow** agent in the RPD. NetFlow is a valuable tool which can help in debugging end-to-end data plane issues.
- **Broadband Digital Forward and Broadband Digital Return**. Building on Moore's law progress these techniques push existing R-PHY features, such as Narrowband Digital Return (NDR) and Narrowband Digital Forward (NDF) into digitization of wider spectrum blocks. DCA effectively replaces the legacy technologies.
- **Advanced Power Management** of the RPD system and its RF module may allow for significant reduction of the power consumption of the HFC network.

R-PHY 2.0 can be also used to more completely specify those data unit formats which are kept as vendor proprietary in R-PHY 1.0 specifications.

Conclusion

R-PHY 1.0 provides a valuable addition to the toolset available to operators s as they continue to extend service offerings and provide ever increasing bandwidth while reducing capital and operational costs. The paper has described a number of issues faced by the 1.0 version which may limit the utility of the R-PHY system going forward.

The paper describes a menu of potential architectural improvements for Remote PHY technology and describes four of these in some detail. Individually, each one of the proposed options solves a different problem and offers valuable business and operational benefits to the cable operators. The paper demonstrates that each one of these options is worthy of the consideration by the cable operators in its own right and that when combined they can provide value which is greater than the sum of the parts. Taken together, these technical improvements constitute a new generation, Remote PHY 2.0.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

Acknowledgments

The authors would like to sincerely thank John Chapman and Gerry White for their valuable comments and contributions to the paper.

Abbreviations

BGP	Border Gateway Protocol
CAPEX	Capital Expenditure
CCAP	Converged Cable Access Platform
CM	Cable Modem
CMTS	Cable Modem Termination System
DCA	Distributed CCAP Architectures
DEPI	Downstream External PHY Interface
DOCSIS	Data over Cable System Interface Specification
ECMP	Equal-Cost Multi-Path
ESD	Extended Spectrum DOCSIS
FCAPS	Faults, Configuration, Accounting (or Administration), Performance, Security
FMA	Flexible MAC Architecture
HFC	Hybrid Fiber Coax
I-CCAP	Integrated CCAP
IoT	Internet of Things
ISYS	Intermediate System to Intermediate System
GCP	Generic Control Plane
GPB	Google Protocol Buffers
L2TP	Layer 2 Transport Protocol
L2TPv3	Layer 2 Transport Protocol version 3
LSP	Label Switching Path
MDT	Model Driven Telemetry
MPLS	Multi-Protocol Label Switching
msec	millisecond
NETCONF	Network Configuration Protocol
OPEX	Operational Expenditure
PCE	Path Computation Element
PSP	Packet Streaming Protocol
RCP	R-PHY Control Protocol
RPD	Remote PHY Device
RR	Route Reflector
R-PHY	Remote PHY
SCTE	Society of Cable Telecommunications Engineers
SR	Segment Routing
SSH	Secure Shell
TE	Traffic Engineering
TLS	Transport Layer Security
UEPI	Upstream External PHY Interface
XTC	XR Traffic Controller
YANG	Yet Another Next Generation

© 2019 SCTE•ISBE and NCTA. All rights reserved.

Bibliography & References

[CLAISE]	https://www.claise.be/2017/10/netconf-versus-restconf-capabilitity-comparisons-for-data-model-driven-management-2/, Benoit Claise
[GCP]	DOCSIS DCA-MHAv2, Generic Control Plane Specification CM-SP-GCP-I04-180509, CableLabs
[GNMI-SPEC]	"gRPC Network Management Interface (gNMI) v0.6.0", OpenConfig operator working group https://github.com/openconfig/reference/blob/master/rpc/gnmi/gnmi-specification.md
[MULPI]	DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I18-190422, CableLabs
[NetProg]	Network Programmability with YANG, Benoit Claise, Joe Clarke, Jan Lindblad, ISBN-13: 978-0135180396
[RemoteScheduler]	SCTE-Expo 2019, R-PHY with Remote Scheduler, Tong Liu
[RFC3931]	https://tools.ietf.org/html/rfc3931
[RFC6020]	https://tools.ietf.org/html/rfc6020
[RFC6241]	https://tools.ietf.org/html/rfc6241
[RFC7950]	https://tools.ietf.org/html/rfc7950
[RFC8040]	https://tools.ietf.org/html/rfc8040
[RPHY]	DOCSIS DCA-MHAv2, Remote PHY Specification, CM-SP-R-PHY-I12-190307, CableLabs
[RPHYOSS]	DOCSIS DCA-MHAv2, Remote PHY OSS Interface Specification, CM-SP-R-OSSI-I12-190510, CableLabs
[RPHYTR]	DOCSIS DCA - MHAv2, Modular Headend Architecture v2 Technical Report, CM-TR-MHAv2-V01-150615, CableLabs
[R-DEPI]	DOCSIS DCA-MHAv2, Remote Downstream External PHY Interface Specification, CM-SP-R-DEPI-I12-190307, CableLabs
[R-UEPI]	DOCSIS DCA-MHAv2, Remote Upstream External PHY Interface Specification, CM-SP-R-UEPI-I10-190307, CableLabs

© 2019 SCTE•ISBE and NCTA. All rights reserved.