SDN As A Matchmaker For Remote Phy Architecture

Alon Bernstein, Sangeeta Ramakrishnan
Cisco Systems

ABSTRACT

The Remote Phy (RPHY) architecture facilitates the separation of MAC (Media Access Control) resources from Phy (Physical layer) resources. The follow-up question is how do we pair them up?
The most straightforward method is to view the Remote Phy Device (RPD) in the RPHY architecture as a fixed "satellite" of the MAC and use static configuration for pairing them up. But this would only be a starting point.
This paper will outline how a dynamic and flexible separation of these resources facilitates new options for OPEX savings as well as new load sharing and availability models.
The paper will also outline an SDN architecture to guide the implementation of these new capabilities.

OVERVIEW

The traditional view of RPHY is that it's satellite architecture. This view sees the RPD as an extension of a CCAP core (not an independent entity) with a static association of the RPD to CCAP core.

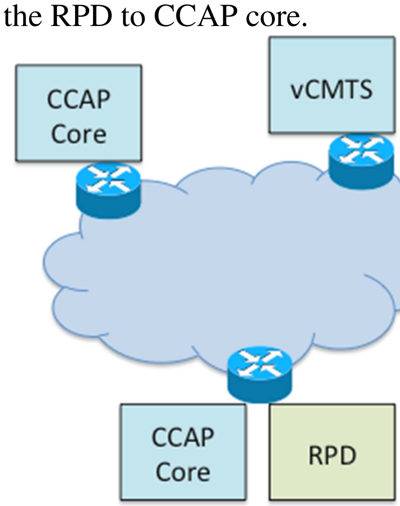

Figure 1 Dynamic association on MAC to RPD

Our point of reference is that we have a pool of Phy resources and a pool of DOCSIS MAC resources that can be paired up in any way and based on any kind of policy. Figure 1 shows a number of resources distributed at various sites in the operator network. The RPD refers to a Remote PHY node or a Remote PHY shelf. (Throughout the paper we will continue to use the term RPD to apply generally to either a RPHY node or a RPHY shelf). The CCAP-core refers to an entity that contains all CCAP functionality but the PHY as defined in the Remote Phy Architecture Document (see ref[4]). The vCMTS refers to a virtualized version of CCAP-core where the MAC functions, including scheduling and MAC level processing, have been virtualized.
As shown in Figure 1 the RPD can be paired with a MAC port on the co-located CCAP core but it can also be paired with a CCAP core that is remote or with a virtual CMTS (vCMTS) in a remote or local data center. Because the MAC can be implemented on either a CCAP or vCMTS we will refer to the term "MAC" throughout this document as the resource that implements the DOCSIS MAC functions) instead of referring to a CCAP core/CMTS/vCMTS.
Since various instances of MAC functions may be available and may reside in different locations, a method to determine how to pair a RPD with a MAC resource is required.
The most obvious pairing is distance based, i.e. connect the RPD to the nearest MAC resource. As this paper will demonstrate there are other options and choices that can be made. The following sections will outline a few use cases for an intelligent mapping approach.
The actual "pairing" itself is done by means of controlling the L2TPv3 tunnel endpoints that are used for the DEPI and UEPI protocols.

PHYSICAL CMTS AND VCMTS

From an RPHY architecture perspective the vCMTS is a MAC resource just like the physical CCAP and the choice between the two is a policy-based decision. For example the physical CMTS ports can be paired first and only if the core runs out of ports then vCMTSs can be spun up and dynamically connected to the RPD.

LOAD BALANCING

A common issue with capacity planning is that enough equipment has to be installed to support a worst-case or near worst-case traffic load. With flexible mapping of resources its possible to optimize the network resource usage.

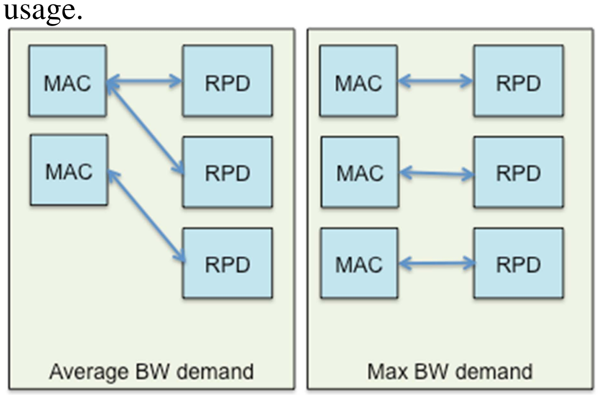

Figure 2 Balancing load

Figure 2 depicts a case where a single MAC resource can handle the load of two RPD during average traffic loads, but only one RPD when traffic peaks. A new MAC can either be a vCMTS that is spun to accommodate the load or a centralized physical CCAP that is designed to handle the overflow traffic from several remote nodes.

AVAILABILITY

The RPD can be connected to an active and a standby MAC resource. While this is a standard mode of operation, the concept of “matchmaking” can be used to have the redundant port in a different box than the active one. In fact, the redundant instance may be intentionally placed in a different data center to allow for geo-redundancy. See ref [2] for a general discussion of redundancy for NFV that is applicable to RPD as well.

PLACEMENT MODULE

The location of the RPD in relation to the MAC is an important piece of information when deciding how to pair the two. There are several methods that can be used to identify the location of MAC and RPD:
1. Pre-configuration: as part of the installation process of a CCAP device or an RPD its location is stored in a database. In the case of vCMTS the location of the data center is sufficient.
2. DHCP servers usually assign IP address pools based on physical location. Inspecting the IP address assigned can be mapped to a physical location.
3. Both (1) and (2) can be combined and cases of mismatch can be flagged. This will provide an extra layer of assurance that the mapping is correct.

Once the location of the endpoints is known the placement module can use the location information in combination with other data and policies to decide which RPD and MAC to pair. As mentioned in the use cases above the additional data can be the role of the MAC port (active/standby), bandwidth load and more.

LOCATION OF SERVICES

Placement of MAC and RPD resources need to take into account the location of other services such as Deep Packet Inspection engines, carrier grade NAT and more. The reason is that we don’t want to create “tromboning” of data through the network. Consider the following example with a caching service:

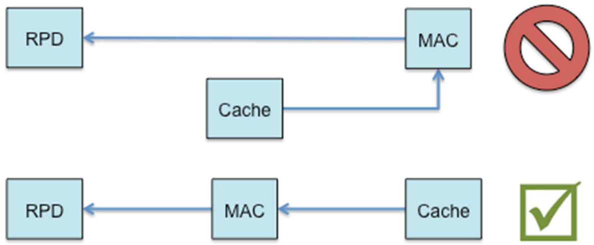

Figure 3 Caching service example

If the cache is placed near the RPD and the MAC is remote (e.g. a distributed cache and a centralized CCAP core) then traffic from the cache will flow all the way up to the MAC only to go down to the RPD as depicted in the top of `Figure 3`. Clearly the placement of the MAC functions needs to take into account the caching service and placed after the cache (the bottom of `Figure 3`).

TRAFFIC ENGINEERING

When the RPHY and the MAC are co-located there is little need for traffic engineering (TE). When the two are remote its critical that the path between them is engineered to have minimal delay and jitter.
Traffic engineering is well understood and the same engineering that is used for other use cases can be used for engineering the path between the MAC and the RPD.
Over the years traffic engineering has been perceived as relatively complex and not easy to manage. Fortunately the introduction of new TE methods (such as Segment Routing - see ref [3] for a discussion of use cases) in combination with SDN greatly simplifies and automates TE creation.

SDN ARCHITECTURE

Previous sections outlined the main components and functions needed for matchmaking: RPD, MAC, Placement, services and TE. The next step is to bind them all together in an SDN framework

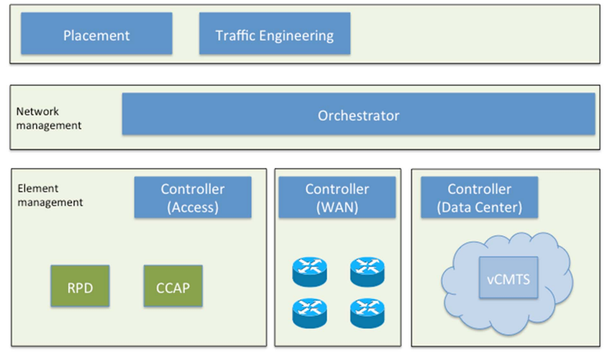

Figure 4 SDN framework for matchmaking RPD and MAC

As depicted in Figure 4 we can fit the placement and TE modules on top of a general purpose SDN/orchestration framework. Furthermore, <u>even the placement and TE modules are fairly generic and can apply to many other use cases.</u>
Let's explore this base architecture and how it relates to other use cases that require placement and TE.
A typical network is segmented into domains. These domains are typically segmented on organizational as well technological boundaries (for example the access, WAN and data center domains in Figure 4).
While this division to domains helps in "dividing and conquering" the complexity of a service provider network it leads to siloes and lack of simple end-to-end provisioning.
The orchestrator, or as some prefer to call it the "controller of controllers" solves this problem. It can manage a complex distributed transaction across all these domains in order to provide end-to-end functions. Note that even the domains themselves may be fairly complex networks and might require their internal orchestration, but for the sake of simplicity we assume a simple two level hierarchy. The orchestrator takes upon itself the traditional TMN (Telecommunications Management Network, see ref [1]) role of "network management" and the controllers assume the role of element managers. This can be confusing as sometimes controllers are branded differently for different domains, but they might all still be built from the same building blocks and similar concepts. Its also

worth noting that while the "purist" SDN view as represented by OpenFlow may be relevant in the data center, it is less so in other areas of the network and irrelevant to pairing the MAC and RPD since OpenFlow is very limited to creating network paths.

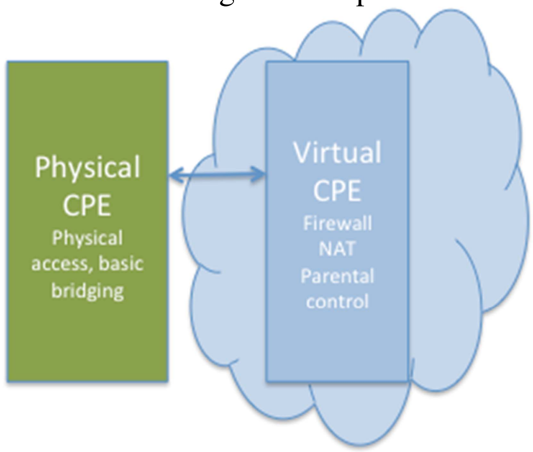

Figure 5 Virtual CPE use case

What other use cases can use a similar infrastructure? Consider the case of a virtual CPE. In a nutshell a vCPE refers to a system where the customer has a very simple physical CPE in the premise and all the advanced services (e.g. NAT, parental controls, firewall) are implemented in the cloud. Without getting into too many details on vCPE we can observe similar themes to the RPD/MAC use case since both have a relation between a physical device and virtual one. The issue of placement appears in the vCPE use case as well because the same policies that govern the placement of a vCMTS relative to an RPD apply here as well - whether distance based, or availability based or load based. The same is true for Traffic Engineering as well. The CPE may be a business CPE with a service level agreement on QoS that requires TE. The same TE solution that applies to the CPE case can be used for engineering a path between RPD and a MAC resource.

There will be customizations needed for each use case. As an example, the CPE use case may use GRE or IPSEC to connect the virtual instance to the physical one. RPD and MAC are connected via L2TPv3, but one can see how the majority of functionality remains the same.

Finally this is meant only as an example and other use cases can leverage this infrastructure as well.

In summary the solution for pairing RPD and MAC resources does not need to be a highly unique solution to solve only the RPD/MAC use case. It can build upon a foundation that is relevant to other use cases and could integration in general.

CONCLUSIONS

RPD can be dynamically and in a highly flexible way be assigned to various MAC resources, whether a local CCAP device in the hub, a remotely located CCAP or a vCMTS. The choice can be policy based and can be changed along with changing network dynamics.
Furthermore, the SDN infrastructure used to achieve this agility does not need to be built from the ground up for the specific use case of pairing a MAC resource with an RPD resource. Instead it builds upon the same SDN infrastructure that is used to enable a host of other network use cases where pairing of resources is needed.

REFERENCES

1. "Telecommunications Management Network". ITU-T Recommendation series M.3000
2. Availability for NFV, Alon Bernstein, INTX 2015
3. https://tools.ietf.org/html/draft-filsfils-spring-segment-routing-use-cases-01
4. Remote Phy Architecture, cablelabs doczone