Converging User Planes in Modern Networks: A Story of GUPpies

A Technical Paper prepared for SCTE by

Jennifer Andreoli-Fang, PhD
Senior Manager, Cloud Architecture
AWS
nextjen@amazon.com

John T. Chapman
CTO Broadband Technologies, Cisco Fellow
Cisco
jchapman@cisco.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

1. Introduction

Many modern broadband networks today are built with three areas of enhancement:

1. containerized implementation,
2. serviced-based architecture (SBA), and
3. control and user plane separation (CUPS).

5G networks are an example network that have all three traits. DOCSIS networks, on the other hand, have yet to define and benefit from SBA and CUPS. This paper focuses on CUPS. We discuss why it is useful and how to implement it. We also cover practical use cases that CUPS can enable.

While we focus on DOCSIS, PON, and 5G networks, the ideas presented in this paper are extendable to Wi-Fi and fixed wireless access (FWA). These ideas build upon the authors' previous white paper on cable and mobile convergence [1].

2. Separating the Control and User Planes

Communication service providers in our industry operate many access networks:

- DOCSIS networks, where the cable modem termination system (CMTS) is the subscriber termination system
- 5G networks, where the 5G core is the subscriber termination system
- Fiber-based passive optical networks (PON), where the BNG (broadband network gateway) is the subscriber termination system
- Wi-Fi and fixed wireless access (FWA), where the BNG is also the subscriber termination system

DOCSIS and 5G were designed by different standards groups. Some of the key differences are: 5G needs to provide mobility management, an unnecessary functionality in DOCSIS. 5G establishes and tears down user sessions, whereas DOCSIS maintains an always-active primary service flow. Most DOCSIS secondary service flows are also active even though the specification allows the ability to dynamically establish, change, and delete them.

The user authentication processes differ significantly. Different terminologies are also adopted. For example, the term "user plane" in 5G is equivalent to "data plane" in DOCSIS. While most of the 5G implementations are containerized, DOCSIS implementations are on the path to containerization.

DOCSIS and 5G are both access technologies, one over the wire, and one over air. While different, there are also many similarities. In the following, we will start by providing a short overview of the 5G network architecture while pointing out similarities to DOCSIS to help the reader orientate.

2.1. 5G Network Architecture

5G has a flat and highly disaggregated architecture. The core network consists of a set of building blocks, or network functions (NFs), offering services needed to execute the functionality of a mobile network. These NFs are interconnected via the "service-based architecture" (SBA) [1].

Prior to 5G, network functions talked to each other via point-to-point interfaces standardized by the 3GPP. The idea behind SBA is to migrate from telecom-style protocol-leaden interfaces to standard web-based application programming interface (API). With SBA, network functions can register themselves and subscribe to each other's services. Doing so allows operators to deploy new services without waiting for standards bodies to define and vendors to develop new protocols. 5G control plane functions are connected via HTTP/2 from the IETF, with JSON as data format, while user plane functions use the classic point-to-point protocols.

Key functional blocks for 5G, as shown in Figure 1, are:

- Access and mobility management function (AMF) – controls the radio access network (RAN) and the connections to the user equipment (UE), as well as UE mobility management
- Session management function (SMF) – manages user sessions and user plane functions
- User plane function (UPF) – handles packet processing, traffic aggregation, and tunneling
- Unified data management (UDM) – manages and provides network user data to the services that request it
- Authentication server function (AUSF) – handles authentication process for a subscriber
- Policy control function (PCF) – works with the SMF and supports policy and charging
- Radio access network (RAN) – last mile wireless network to the user

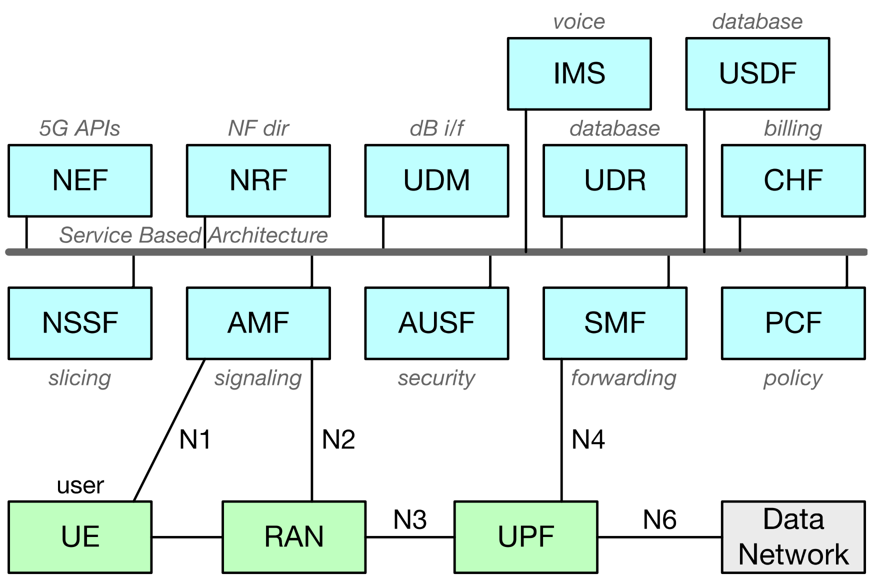

Figure 1 – Key 5G Core Network Functions

2.2. 5G CUPS

Control and user plane separation, or CUPS, was first introduced in 4G / LTE and defined in 3GPP Rel-14. The idea was to separate the packet gateway into control and user planes to allow for resource sharing, flexible deployment, and independent scaling to achieve benefits in CapEx and OpEx. CUPS is now required for 5G standalone architecture.

Figure 2 shows the 5G network architecture simplified with CUPS. The control point to access the 5G user plane across N4, the UPF, is the SMF. Multiple SMFs can share the same UPF and a SMF can interface with multiple UPFs. More importantly, the UPF is mostly protocol agnostic – it just needs to support a 3GPP packet tunneling protocol (N3) and a control protocol (N4) that it uses to talk to the SMF. The SMF is the network function that interfaces with other 5G control plane functions.

The 5G control plane functions including the SMF support diverse and complex set of functions all related to the essential functions of the mobile network and subscribers, while leaving the user plane to handle tasks solely related to packet handling. The UPF specializes in packet processing that is not unique to 5G, but also to other communication networks. This opens the possibility for the UPF to be reused in a multi-access network scenario. We will expand on this point later in the paper.

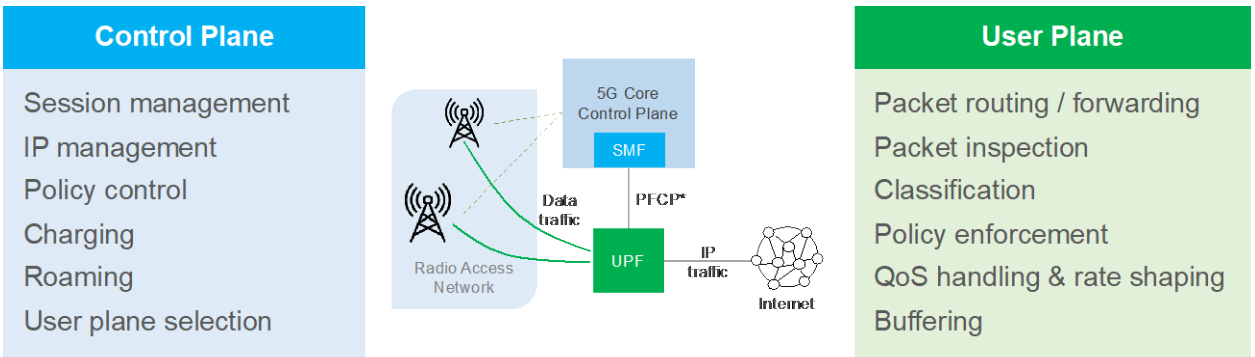

Figure 2 – 5G CUPS Simplified

2.3. CUPS in Other Access Networks

Many cable operators also operate passive optical network (PON) for their fiber network. The broadband network gateway (BNG) is the subscriber termination system in PON. BNG is specified by the Broadband Forum (BBF), and has gone through several iterations, from monolithic architecture, to virtualized BNG that leverages commodity hardware in lieu of proprietary hardware, to disaggregated BNG (DBNG) that is based on the CUPS principle [3].

Figure 3 shows DBNG architecture. The DBNG control plane (CP) provides subscriber management functions including authentication, authorization, and accounting (AAA), IP address management, and policy. The DBNG CP interfaces with external management services with the northbound control interfaces. The CP is also responsible for programming data traffic forwarding rules for the DBNG UP. CUPS enables the option of having a CP function to control multiple user plane (UP) functions. This allows the CP can have the ability to steer a subscriber session to a particular UP. The steering function enables this functionality and is currently being defined by the BBF.

There are strong similarities between the DBNG and the 5G core architectures. The steering function is a combination of the 5G's AMF and SMF. The information flow between the DBNG CP and UP adopted and extended two protocols standardized by the 3GPP: packet forwarding control protocol (PFCP) and general packet radio services (GPRS) tunneling protocol user plane (GTP-U)[1]. The BNG architecture is

[1] Note that GTP-U in DBNG architecture is used to tunnel the control packets between the CP and CPE with the UP as the intermediary. In 5G, the control packets between the CP (AMF) and CPE (UE) are carried on the N1 interface.

also set to evolve into a new network function, the aggregation gateway function, or AGF (to be covered later) to support wireless-wireline convergence.

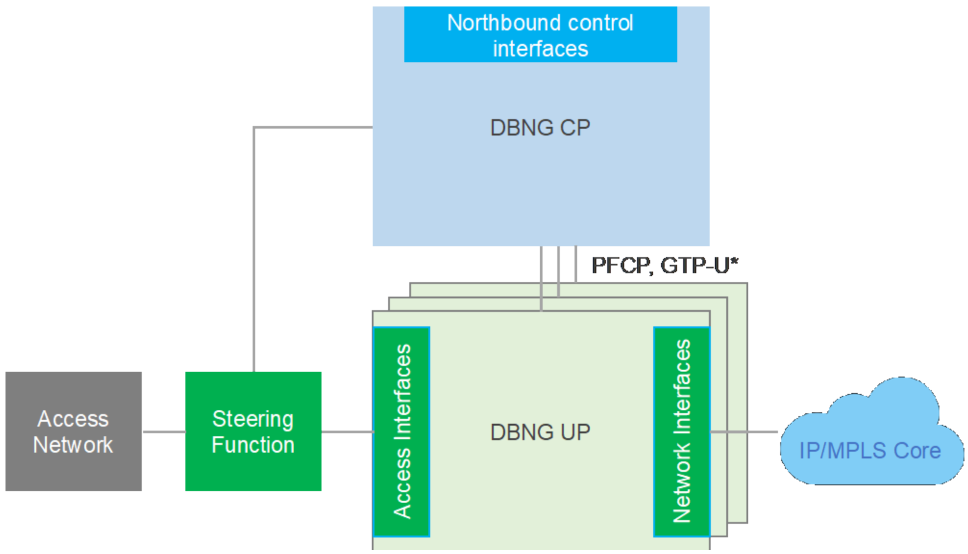

Figure 3 – Disaggregated BNG

2.4. Motivation for CUPS

Many argues that a key benefit of 5G CUPS is that it allows the network operator to distribute resources throughout the network, with the control plane (CP) sitting in a centralized location and the user plane (UP) closer to the application.

More importantly, CUPS architecture provides several benefits in operating a modern network.

Operational considerations: CUPS allows control and user plane software to be updated independently. Considering the complexity of the control plane and the multitude of functionalities it supports, the CP should be updated frequently. Operators can leverage DevOps techniques to update the CP. On the other hand, user plane software updates potentially create network downtime and outages, and ideally should be updated infrequently. By decoupling the control and user planes, operators will have more system uptime and results in better high availability.

Implementation and deployment optimization: The Remote PHY architecture is essentially a DOCSIS radio in the fiber node with the CMTS core software in the hub. With CUPS, the DOCSIS control plane can be deployed in the hub or in the cloud, while the DOCSIS data plane can remain in the hub. The DOCSIS data plane can further be implemented as dedicated central processor unit (CPU) or network processor unit (NPU) cores.

Enabling rapid adjustment of compute and network resources: Unplanned outages or expected capacity bursting require operators to react quickly to deal with massive surge of control and/or user plane traffic. An area outage can cause a large number of cable modems (CM) to reboot and re-register at the same time. In the case of a registration storm, only the control plane is impacted. CUPS enables the independently scaling of the CP – in this case, the registration function only, without the need to scale the UP. Another example is busy hour or holiday-related uptick in capacity demand that requires the increase in UP capacity without necessarily increasing the capacity of the CP.

3. The Protocols Behind CUPS

3.1. PFCP

The packet forwarding control protocol (PFCP) is the session layer protocol that runs on the 5G's N4 interface between the SMF and UPF, and is used by SMF to associate with and configure the UPF. An SMF in 5G can control multiple user plane functions. The SMF needs to select and associate with a UPF. The SMF configures a user session by sending packet detection rules (PDRs) for each session. Figure 4 illustrates what happens when a GTP-U encapsulated packet is received by the UPF and eventually egresses as an IP packet to the data network.

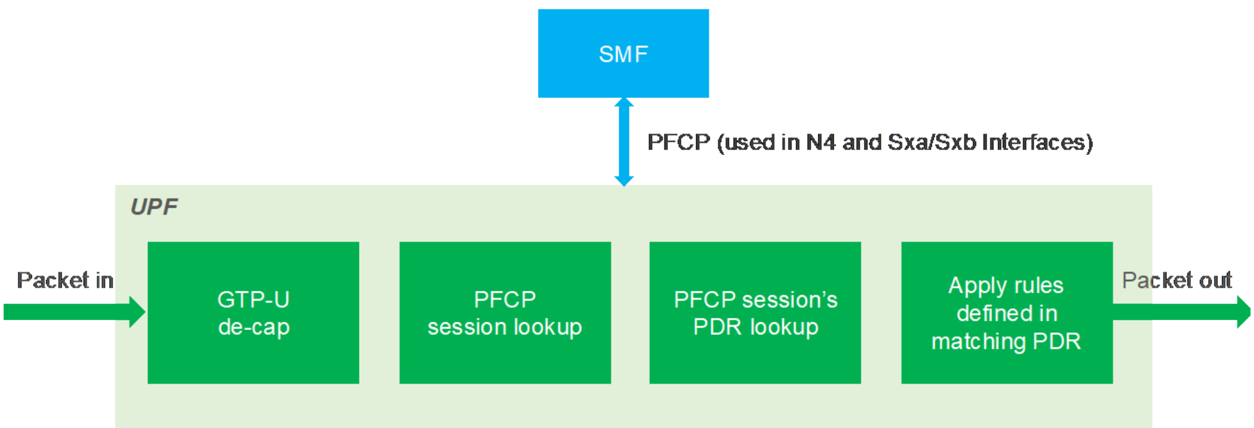

Figure 4 – Packet Processing Workflow in UPF

3.2. Remote PHY vs. Mobile RAN

The Remote PHY architecture was invented by John Chapman of Cisco in 2001 [4] [10] [11] [12] and was originally part of the Modular Headend Architecture (MHA). The main protocols of Remote PHY are:

- **R-DEPI**, the remote downstream external PHY interface, a pseudowire that connects the CMTS Core to the RPD [5],
- **R-UEPI**, the remote upstream external PHY interface, a pseudowire that connects the RPD upstream to the CMTS Core [6],
- **GCP**, the generic control plane, that is the management protocol for the RPD [7], and
- **R-DTI**, the remote DOCSIS Timing Interface, that uses IEEE 1588 to manage the DOCSIS timestamp between the CMTS Core and the RPD [8].

DEPI, UEPI, and DTI were written in 2004. GCP was written in 2012 at a Pete's Coffee shop. DEPI and DTI were standardized as part of MHAv1 at CableLabs in 2005 [13]. UEPI and GCP, along with a "Remote" prefix, were standardized at CableLabs as part of CDOCSIS in 2014 [14] and MHAv2 in 2015 [15].

Remote PHY is a fundamental enabler of the DOCSIS CUPS architecture. Remote PHY disaggregated the integrated CMTS into MAC and PHY components. The MAC is placed into a CMTS Core and can be virtualized [9]. The PHY is placed into a remote PHY device (RPD) that is located in the HFC fiber node.

The RPD and node together act as a “DOCSIS radio” and are analogous to the mobile radio access network (RAN). In fact, the Remote PHY protocol was the influence for the development of the Small Cell Forum (SCF) network functional application programming interface (nFAPI) split 6 protocol [16].

In 5G, the RAN is managed by the AMF. In Remote PHY, the RPD is managed by the principal core. There is much more to both stories, but the analogy holds true. The vCMTS contains the Remote PHY tunnel generation and termination as well as the DOCSIS MAC. As a containerized virtual implementation, the DP and CP should be separate containers and could be located on different servers, or even as diverse as moving the CP to the regional or far edge cloud.

Today, the protocols between the DOCSIS DP and CP are proprietary and internal to each CMTS manufacturer. PFCP should make an excellent baseline as a protocol, just as it has for mobile (4G and 5G) and PON (BNG). There will likely be extensions required for DEPI and UEPI, and for any DOCSIS functionality that does not map directly into existing PFCP.

Note that the DOCSIS upstream scheduler is a real-time function that would likely stay with DOCSIS DP or stay close by. This is also true in mobile and PON where the upstream scheduler stays close to the RAN.

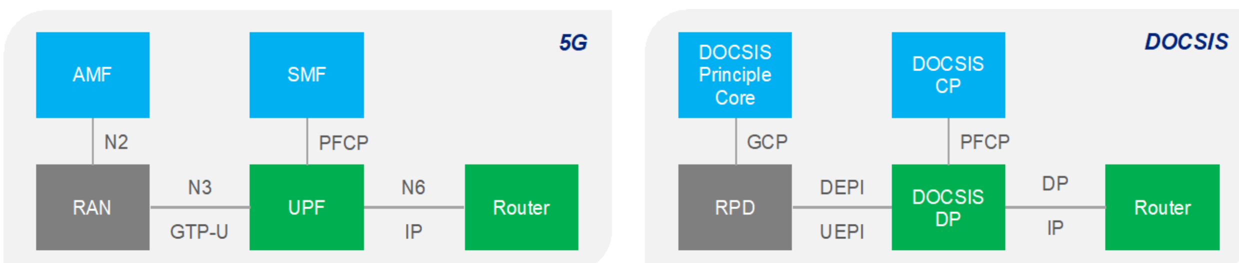

Figure 5 – Remote PHY vs. Mobile RAN

3.3. Converging the User Planes

If DOCSIS can be implemented with a CUPS architecture that is very similar to 5G and PON/BNG, then could the networks be converged? Convergence does not mean that all the protocols must be changed, but it does imply reuse of major blocks, with the intent of achieving a common management and provisioning system.

One early example of convergence was the use of DOCSIS Provisioning over EPON (DPoE) [17] that reused a DOCSIS provision system to provision a passive optical network (PON) system. DPoE specified a signaling gateway that converted between the two worlds.

The 5G architecture defines the access gateway function (AGF) that connects a wireline termination system to a 5G core. Jennifer Andreoli-Fang, while at CableLabs, defined a cable version of this called the cable fixed access gateway function (CFAGF) in the 3GPP technical report on wireless-wireline convergence [20]. CFAGF was later renamed as wireline AGF (W-AGF) in the specification stage [21].

CableLabs adopted the W-AGF terminology in the converged core technical report (TR) [22], where the W-AGF was shown to connect an integrated CMTS to a 5G core. The role of W-AGF in a Remote PHY system was not discussed in the CableLabs TR. However, just as in the I-CMTS case, W-AGF can connect

R-PHY to a 5G system with an additional functionality of converting between DEPI/UEPI and GTP-U. This is all shown in Figure 6.

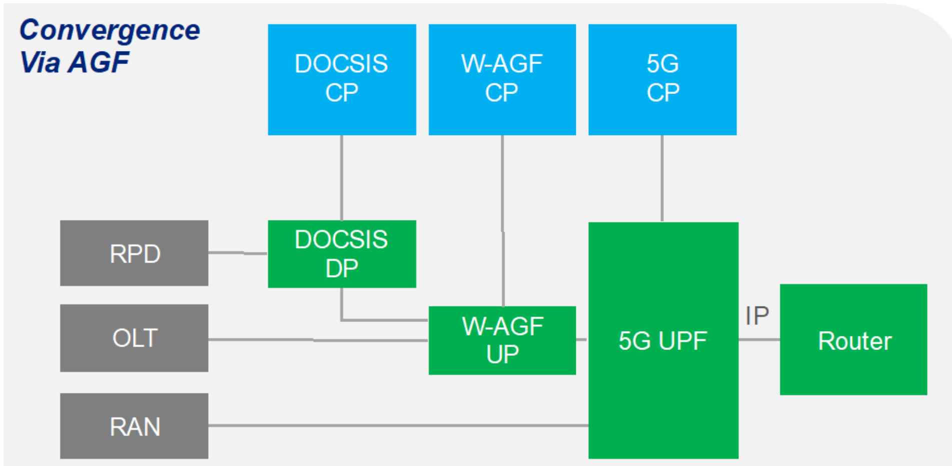

Figure 6 – Convergence via W-AGF

Care must be taken in the interpretation and implementation of this functionality. The W-AGF as defined in 5G connects the DOCSIS and 5G worlds together though a gateway. Without any optimization, the result may end up with two completely separate data planes that are now interconnected. Such an interconnection couples the two data planes and can make the combination redundant and more complicated that two separate systems, without achieving cost reduction and simplicity goals. For example, there are now two quality of Service (QoS) and service level agreement SLA control points and data planes.

3.4. GUP – Generic User Plane

A more pragmatic implementation would be something described in Figure 7. The DOCSIS RPD, the PON OLT, Federated Wi-Fi, and the 5G RAN would all send their traffic to a generic user plane (GUP). The GUP would have a pseudowire function (PWF) that terminated and generated pseudowires of interest, specifically:

- GTP-U for mobile
- VLAN over L2TP/PPPoE for PON
- DEPI/UEPI (L2TPv3) and DOCSIS Framer for DOCSIS
- L2TP (or equivalent) for Federated Wi-Fi

It is worth noting that Remote PHY uses a pseudowire called L2TPv3 that is derived from the BNG pseudowire L2TP. L2TPv3 is the base protocol for DEPI and UEPI.

The interface between the pseudowire function (PWF) and the packet services function (PSF) would be an IP over Ethernet packet coupled with context information containing metadata like a generic service identifier (GSID) and port interface IDs.

It would then apply a common set of subscriber management processes and QoS tools. In DOCSIS, this would be packet classification into service flows and the associated rate shaping. There are equivalently defined operations in mobile and PON. The key is to have this all done once, on the same hardware and software user plane, not a separate operation per service domain.

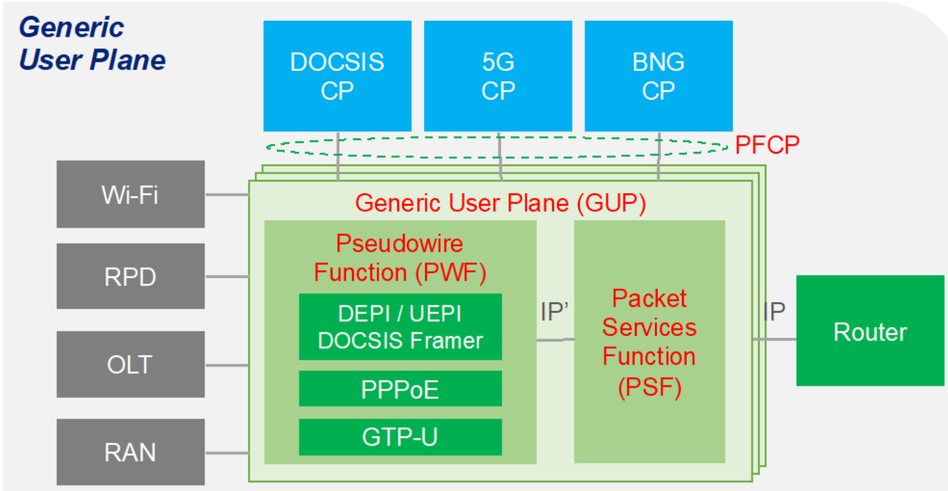

Figure 7 – Truly Converged User Planes

The GUP function could be a larger centralized model or a distributed model. And yes, a distributed GUP system would be a bunch of GUPpies.

We have spoken strictly about the user plane in this study. There are also incredible opportunities in the control plane and management planes to converge workloads even though the access protocols may differ. But this is for another whitepaper. Meanwhile, let's go down one more level on the user plane implementation.

3.5. P4

P4 stands for "Programming Protocol-Independent Packet Processors". P4 is a programming API for programming a network processor unit (NPU).

What is an NPU? Let's compare to what we know. Central Processing Units (CPU) are used to run generic software and are programmed with a machine specific assembly language. Graphical Processing Units (GPU) are massively parallel processors that specialize in vector operations and can be used for graphical rendering, and more recently bitcoin mining and artificial intelligence. They also have a machine specific programming interface. An NPU is used for switching packets. Routers have NPUs. An NPU can be implemented in a hardware application-specific integrated circuit (ASIC) or as a software core on an CPU. Hardware NPUs typically had a proprietary machine specific programming language. P4 changes that by standardizing the interface to an NPU.

PFCP is a layer 3 control plane protocol that is intended to be run between two software stacks on two different network elements. The GUP software stack would convert between PFCP and P4. This is shown in Figure 8.

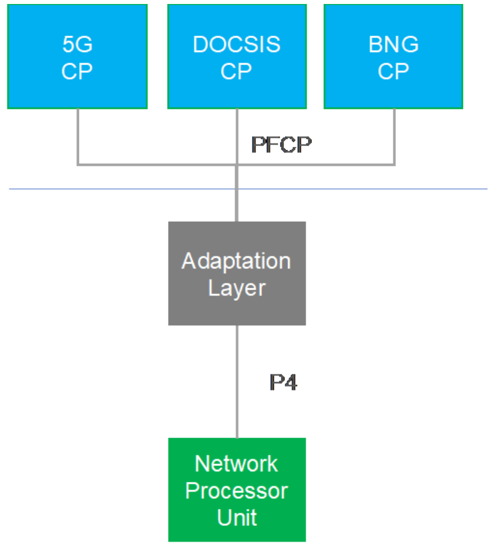

Figure 8 – NPU Programming Model

P4 is agnostic to all protocols. P4 works on bytes and tables. P4 allows fields to be defined within a packet. Fields can be extracted and inserted, and those fields can used as indexes into tables and be retrieved from tables.

For example, in P4, you could define a packet type called "Ethernet". The first six bytes would be the "Destination Address", the next six bytes would be the "Source Address", and the next two bytes would be the "Type". For layer 2 forwarding:

- a forwarding operation would be performed by telling the hardware to use the Destination Address as a lookup into the forwarding table.
- a learning operation would be performed by storing the "Source Address" as a new forwarding table entry along with the port number the packet came in on.

A router would use the Type field to understand what the next protocol is, such as an Internet Protocol (IP) packet, and P4 would have a description of what to do. When a layer 3 packet is forwarded, P4 would have a task to rewrite the Ethernet Destination Address.

Why is all this important?

It is important that all the functionality that is defined by PFCP must be implementable by P4.

And why is that also important?

Because once P4 can explain everything to an NPU in agnostic terms, then a common user plane becomes realistic. And once that user plane is common, then it could be converted into an ASIC to drastically increase performance, decrease costs, and distribute user planes across the IP network.

4. Practical Use Cases

As a bonus, the CUPS architecture enables the partial or full offload of the CMTS core to the cloud. We will start with a primer on CMTS disaggregation and the cloud, before describing the use cases.

4.1. CMTS Disaggregation

Virtual CMTS functions can be split into service domains based on timing requirements. Figure 9 is adopted from [19] and illustrates the functional blocks in the DOCSIS management plane, control plane, and data plane.

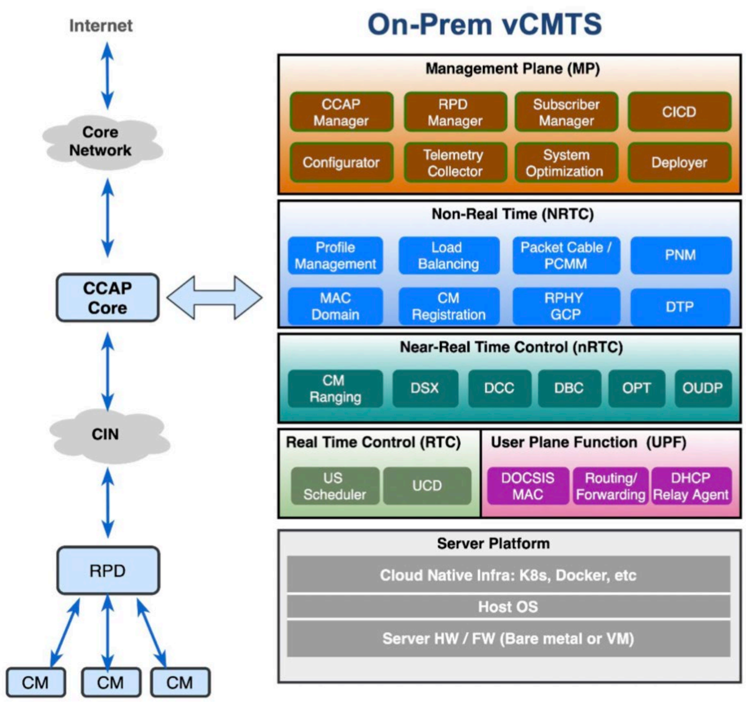

Figure 9 – vCMTS Functional Blocks

4.2. Cloud Primer

This section provides a primer on general cloud terms and concepts. To describe the use cases with sufficient detail, we will use some Amazon Web Services (AWS) specific terms as examples, although we will keep the terminologies agnostic whenever possible. Other cloud providers may have similar concepts and naming conventions.

Every communication service provider (CSP) is likely familiar with what cloud providers call the "region." It's best to relate a cloud region with a CSP's regional data center. Each region is typically consisted of multiple, isolated locations where data centers sit to provide a resilient infrastructure. But the cloud infrastructure is much more.

As shown in Figure 10, the cloud infrastructure extends from cloud region to cloud edge and to the cloud far edge to support applications with different timing constraints or local residency requirements. Cloud edge is typically located in major metro centers and the cloud far edge can sit in the communication service providers (CSP) facilities. In AWS, region, edge, and far edge are inter-connected and have the same "look and feel" for cloud infrastructure services such as networking, compute, storage, and deployment automation.

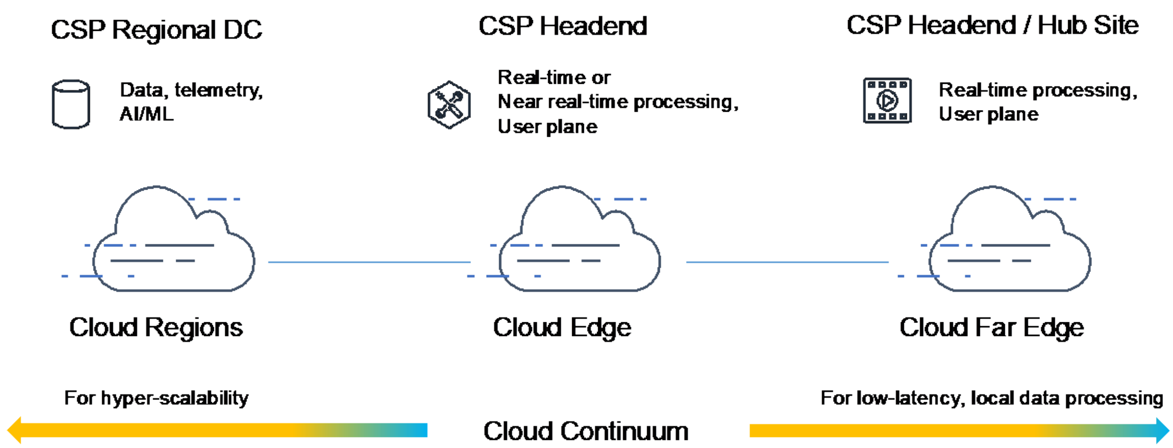

Figure 10 – Cloud Infrastructure

Putting containerized applications in the cloud allows an operator to leverage some of the fundamental benefits that come with the cloud.

- **Resiliency**: Due to its massive scale and design, the cloud infrastructure is inherently more resilient compared to a CSP's on-prem data centers.
- **Pay-what-you-use model**: With on-prem data centers, CSPs need time, CapEx, and OpEx to plan, build, and maintain data centers for peak capacity. Large amounts of resources planned for peak usage sit idle during normal operating conditions. With cloud, CSPs only pay for normal network capacity and pay for peak capacity only on the occasions when capacity demand surges.
- **Horizontal scaling**: Modern software applications are architected to enable horizontal scaling, or scaling out, which involves spinning up additional containers used for a functionality when demand arises. (In contrast, vertical scaling, or scaling up, refers to scaling by adding more power like CPU or memory.) Horizontal scaling and cloud go hand-in-hand. With cloud's near-infinite compute resources, applications can automatically scale out to deal with a surge in demand.

4.3. CUPS and the Cloud

This section focuses on DOCSIS CUPS use cases, although the same concept is extendable to BNG and mobile. DOCSIS management plane applications such as subscriber management or remote PHY orchestrator are well-suited to be deployed in the cloud due to their traffic and latency characteristics. CUPS enables a service provider to partially offload the vCMTS application to the cloud. As shown in Figure 11, the vCMTS control plane functions can be deployed in the region. When containerized, the CP can leverage the cloud's native ability for horizontal scaling to scale out when required in cases such as registration storm. CUPS allows the DOCSIS data plane to be deployed on-prem. A dedicated connection between CSP on-premise infrastructure and cloud region (shown in Figure 11 as AWS Direct Connect) carries control plane traffic, which can be PFCP.

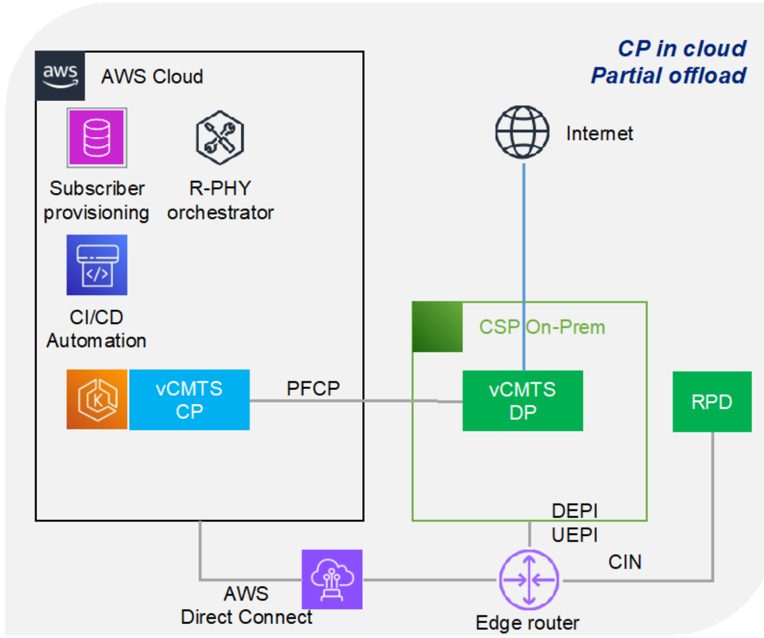

Figure 11 – Partial Offload – CP in the Cloud

In a temporary network outage, deploying the DOCSIS DP in the cloud can accelerate disaster recovery. Figure 12 shows the basic architecture where both the DOCSIS CP and DP run in the cloud. The DP may be run in the cloud edge such as a metro center to minimize latency from the outside plant. The dedicated connection between cloud region and on-prem would now carry DEPI and UEPI traffic. The ingress and egress point of DOCSIS data plane traffic to the internet is directly through cloud.

Another use case not shown is cloud-based capacity bursting. During normal operations, DOCSIS DP runs on CSP premise. When capacity demand increases, additional DP can spin up in the cloud. This way, CSPs don't need to build and operate their data center for peak capacity, where a large number of resources sit idle for most of the time.

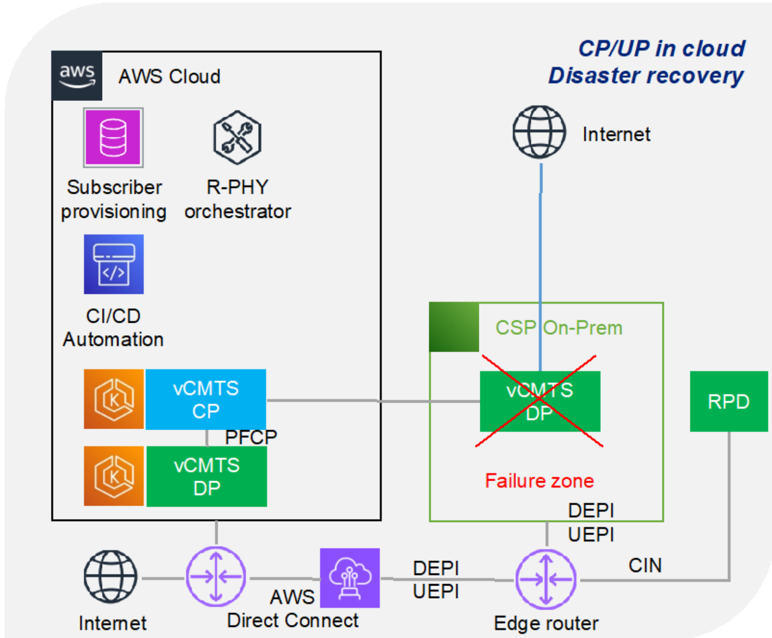

Figure 12 – Full Offload – Disaster Recovery in the Cloud

5. Business Considerations

As with all business decisions for adopting of new technologies, there needs to be a strong business case or the product may be technically brilliant but fail in the marketplace. There must be a return on investment (ROI). This means either an increase in revenue or a decrease of capital and/or operational costs for the operator.

The pro-side of convergence is a single system that does everything. There could be less redundant spend across multiple systems. There could be one team to manage that one system. The customer can seamlessly move between the different access systems.

The con-side of convergence is that by combining the systems together, the entire system becomes more complex. It is more complex in implementation, more complex to test, and more complex to update software. It requires cross domain knowledge just to operate in one of the access domains which make it more complex to staff.

There is a further deployment risk. The 5G architecture seems to be the most attractive to converge to, but it is also the least deployed architecture compared to PON and DOCSIS. 5G systems also have the least bandwidth requirements when compared to 10G/25G/100G links for PON and 10G links for DOCSIS. 5G also has the most complex feature set when it comes to roaming, hand-offs, deep packet inspection for flow discovery, and more.

A practical move to convergence should prioritize a move to simplicity, more common features, and less legacy access-specific features.

In this white paper, we took a simple approach of just converging the user plane. All user planes classify and forward IP over Ethernet packets that are contained in an access-specific encapsulation. This is a contained problem. An optimized user plane that could process packets at higher speeds and be usable across all access market segments should also be lower cost per bit.

In addition, a common user plane technology could be removed from the centralized access systems and moved into the network to allow for distributed user plane termination. This would then allow distribution of edge processing functions like media and application caches.

6. Conclusion

A generic user plane, GUP, enables scale. Scale of the business case. Scale of the actual NPU. A larger market for a common device. An option of a software-based or hardware-based user plane. The opportunity to move the user plane from gigabits to terabits and petabits of data throughput.

Convergence may not be an all-or-none proposal. It may be just one step at a time, and the data/user plane looks to be one of the first steps.

Abbreviations

AAA	authentication, authorization, and accounting
AGF	access gateway function
AMF	Access and mobility management function
API	Application programming interface
ASIC	Application-specific integrated circuit
AUSF	Authentication server function
BBF	Broadband network Forum
BNG	Broadband network gateway
CFAGF	cable fixed access gateway function
CM	Cable modems
CMTS	Cable modem termination system
CP	Control plane
CPU	Central processing unit
CSP	Communication service provider
CUPS	control and user plane separation
DevOps	Development and operations
DEPI	Downstream external PHY interface
DOCSIS	Data over cable system interface specification
DBNG	Disaggregated BNG
DPoE	DOCSIS provisioning of EPON
EVPN	Ethernet Virtual Private Network
FWA	Fixed wireless access
GCP	Generic control plane protocol
GPRS	General packet radio services
GTP	GPRS tunneling protocol
GTP-U	GTP user plane
GUP	Generic user plane (first defined in this paper)
IP	Internet protocol
LAN	Local area network
NF	Network function
NPU	Network processing unit

PCF	Policy control function
PDF	Packet detection rules
PON	Passive optical network
PSF	Packet services function (first defined in this paper)
PWF	Pseudowire function (first defined in this paper)
QoS	Quality of service
RAN	Radio access network
RPD	Remote PHY Device
SBA	service-based architecture
SMF	Session management function
TR	Technical report
UDM	Unified data management
UEPI	Upstream external PHY interface
UP	User plane
UPF	User plane function
VLAN	Virtual LAN
W-AGF	Wireline AGF

References

[1] Jennifer Andreoli-Fang, John Chapman, et. al, "Cable and Mobile Convergence: A Vision from the Cable Communities Around the World", *SCTE Cable-Tec Expo Fall Technical Forum*, Denver, Oct, 2020. [link]

[2] "System architecture for the 5G System (5GS)", 3GPP TS 23.501.

[3] "Multi-Service Disaggregated BNG with CUPS", BBF TR-459, 2020.

[4] John T Chapman, "Packet Fiber Node", US Patent 7,209,442, filed Jun 27, 2001, issued April 24, 2007. [link]

[5] John T Chapman, "Downstream Remote Physical Interface for Modular CMTS (DEPI)", US Patent 7,688,828, filed May 17, 2005, issued March 30, 2010. [link]

[6] John T Chapman, "Upstream physical interface for modular cable modem termination system (UEPI), US Patent 7,639,617, filed May 17, 2005, issued Dec 29, 2009. [link]

[7] John T. Chapman, "Generic Control Protocol (GCP)", US Patent 9246700, filed Sept 12, 2012, issued Jan 26, 2016. [link]

[8] John T. Chapman et. al, "Managing time offset and frequency drift in asynchronous docsis remote phy network environments (DTI)", US Patent 9,722,739, filed April 15, 2015, issued Aug 1, 2017. [link]

[9] John T Chapman, Alon Berstein, et. al., "Virtual Cable Modem Termination System", US Patent 9,369,339, filed Sept 30, 2013, issued Jun 14, 2016. [link]

[10] John T Chapman, "Modular CMTS Architecture", NCTA Spring Technical Forum, 2005. [link]

[11] John T. Chapman, "DOCSIS Remote PHY", *SCTE Cable-Tec Expo Fall Technical Forum*, Atlanta, GA, Oct 2013.

[12] John T. Chapman, "Remote PHY for Converged DOCSIS, Video and OOB", *NCTA/SCTE Technical Forum*, Los Angeles, Jun, 2014. [link]

[13] CableLabs, Modular Headend Architecture Specifications, 2005. [link]

[14] CableLabs, CDOCSIS Specifications. [link]

[15] CableLabs, Remote PHY Specifications (MHAv2). [link]

[16] "nFAPI Specifications", Small Cell Forum. [link]

[17] CableLabs, DPoE Specifications, 2012. [link]

[18] R. MacDavid et al., "A P4-based 5G User Plane Function", Proceedings of the ACM SIGCOMM Symposium on SDN Research (SOSR), Oct 2021. [link]

[19] John T. Chapman, Tong Liu, "Unleash the Power of Cloud Computing for CMTS", SCTE Expo 2021. [link]

[20] "Study on the Wireless and Wireline Convergence for the 5G system architecture", 3GPP TR 23.716.

[21] "Wireless and Wireline Convergence Access Support for the 5G System," 3GPP TS 23.316.

[22] "5G Wireless Wireline Converged Core Architecture Technical Report", CableLabs, June 2020.

