SOLUTION BRIEF

The Cable Access Node of the Future is Intelligent, Modular, and Scalable

Deploy any access technology using a module that conforms to specifications of the Generic Access Platform (GAP).

Figure 1. Hybrid Fiber-Coaxial (HFC) Node Hangs on Cable Strand

"Today, there are operational complexities supporting a multitude of custom node housings and designs. GAP is a means to streamline node deployments, and future proof and simplify our outside plant."

John Williams
VP, OSP Engineering and Architecture
Charter Communications*

Charter Communications is currently the second largest cable provider in the United States. The company is at the intersection of technology and entertainment, facilitating essential communications that connect more than 26 million residential and business customers in 41 states.

Access Network Complexity

With network requirements and delivery models constantly evolving, multiple-system operators (MSOs) need to deploy a mix of ever-changing access technology. For example, a DOCSIS node may, in the future, require a remote-PHY and passive optical network (PON) to serve a given neighborhood. In the coming years, another location might implement full-duplex (FDX) DOCSIS together with 5G to better serve an enterprise.

Keeping up with this changing landscape necessitates many different node SKUs, leading to high deployment, inventory, development, and maintenance costs. This is because today's hybrid fiber-coaxial (HFC) nodes, amplifiers, and other outside plant gear are usually purpose-built to a pre-defined set of functions and cannot be easily changed after installation. Therefore, operators would have to swap out an entire HFC node (Figure 1) in order to make the previously described upgrades, which may be cost and time prohibitive. Further exacerbating management complexity, operators are driving fiber deeper in the network, resulting in up to a ten-fold increase in the number of HFC nodes.[1]

What is needed is a set of standards that helps increase the agility and manageability of HFC nodes. This paper discusses how Intel, Charter Communications*, and others in the industry are pioneering a new concept to address these industry challenges with the Generic Access Platform (GAP) project. The platform also eases the migration to multi-access edge computing (MEC), and HFC node connectivity to Citizens Broadband Radio Service (CBRS) or other 5G networks- two key advancements in the networking and communications industries.

GAP Overview

HFC nodes connect fiber optic cables to coax cables:

- Fiber originating from a cable headend site, and
- Coax terminating at cable modems located in subscriber homes.

While GAP is not yet fully defined, it is likely to include a variant similar in size to today's strand mount nodes, which are roughly 50 pounds and 12 in. x 12 in. x 24 in., as depicted in Figure 2. It either hangs on a fiber/cable strand (i.e., strand mount) or is bolted inside an enclosure that is on the ground, in a pedestal, or underground in a vault. A major feature of a GAP enclosure is it can be upgraded while installed in the field either with a remote software upgrade or with a simple module swap instead of an entire housing swap as is typical today.

GAP Specifications

GAP specifies a standards-based, modular enclosure with well-defined slots that can house any module conforming to GAP's physical, thermal, mechanical, and electrical specifications. Thus, GAP-compliant modules are able to coexist within a GAP-compliant enclosure. Akin to the Peripheral Component Interconnect (PCI) specification, GAP provides enough information to allow vendors (e.g., OEMs) to independently develop interoperable modules.

In order to take into account a variety of environments, the plan is to specify a number of GAP housing sizes and footprints that can host the same modules.

GAP Objectives

Today, an OEM is typically responsible for the entire HFC node, including all internal electronic components as well as the enclosure (usually cast aluminum). A goal of GAP is to eliminate the need for each OEM to develop its own housings; and instead, devote these engineering resources to developing service-generating modules.

GAP Benefits

Unlike most fixed-function HFC nodes today, MSOs will be able to make in-field updates of the access technologies in GAP-compliant HFC nodes. With the ability to swap GAP-compliant modules, MSOs can cost-effectively add new functionality to HFC nodes, which ultimately increases their return on investment (ROI). This future-proofing of HFC nodes enables MSO's to quickly upgrade networks with emerging capabilities, like deep fiber, distributed architectures (Remote-MAC/-PHY/-MACPHY), full duplex, 5G, and fiber to the home.

Some of the other benefits enabled by GAP include:

- Reducing inventory: MSOs can address different access mixes at the module level, which costs less than carrying complete HFC nodes in inventory for each access mix.

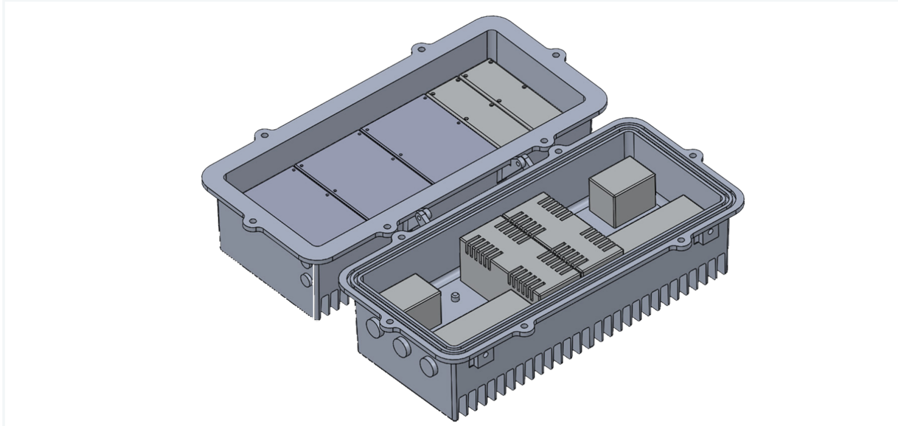

Figure 2. Mockup of a GAP-Compliant HFC Node

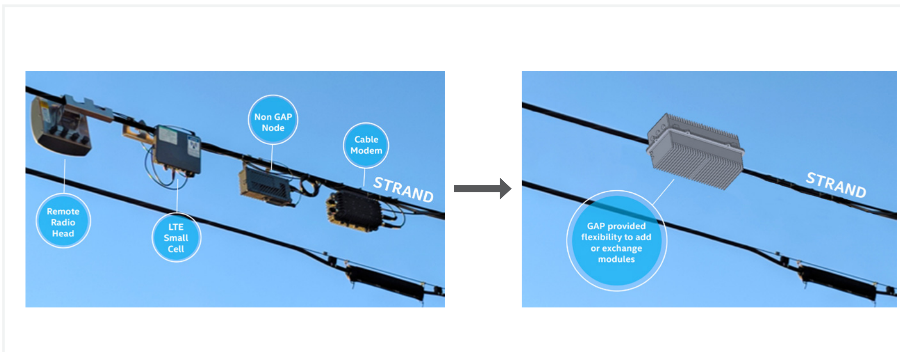

Figure 3. GAP Consolidates Access Technologies

- Using best-of-breed vendors: OEMs no longer have to build the entire node and can focus their investment in areas of expertise. MSOs are not tied to a single OEM and can choose module vendors that best fit their needs.
- Facilitating edge computing: Compute capabilities can be more easily moved to the far edges of the network with GAP-compliant modules equipped with server-grade microprocessors.
- Minimizing strand sprawl: A variety of access technologies can be consolidated into a single GAP-compliant HFC node, as shown in Figure 3.

Making GAP Nodes Intelligent

Network functions virtualization (NFV) is a pervasive theme within the telecom industry, producing in a major transformational change the network. The underlying technologies create a platform for rapid service deployment, software upgrades, improved resource utilization and future proofing. MSOs can also take advantage of readily-available software from a large ecosystem to address new business trends.

Using the tenets of NFV, GAP nodes can be architected to extend the benefits of NFV to the network's edge. This requires making GAP nodes more intelligent and software-based by adding high-performance compute modules into GAP-compliant HFC nodes. These modules can also implement multi-access edge computing (MEC),

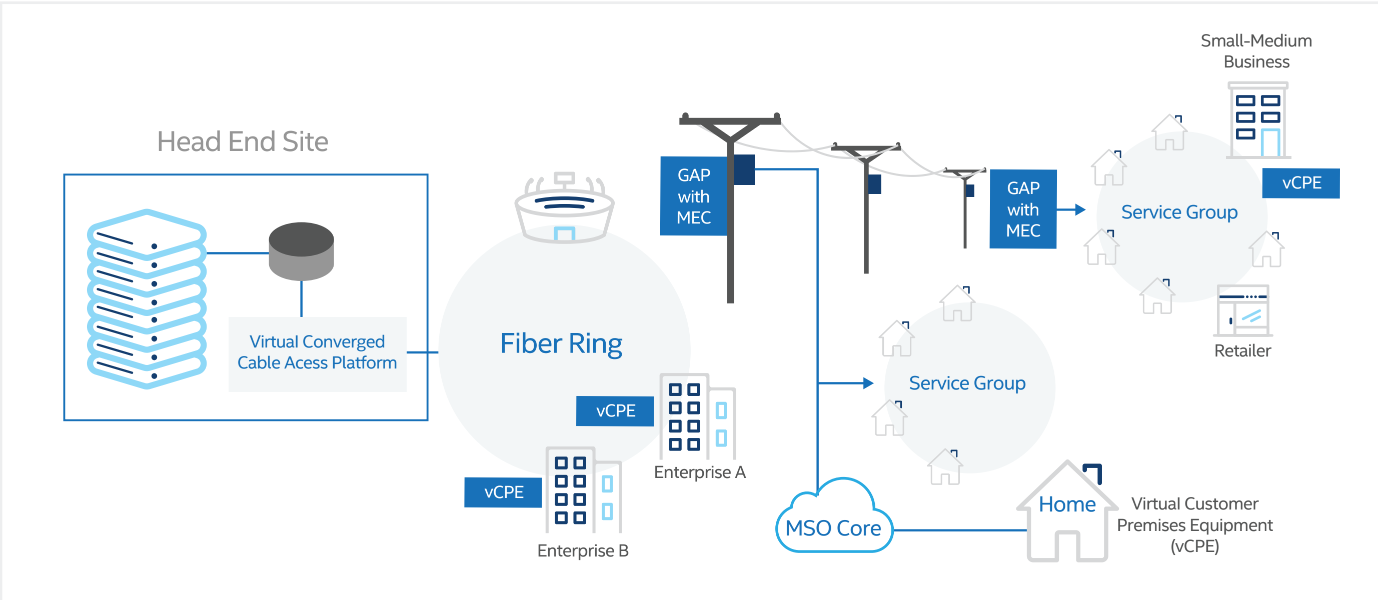

Figure 4. MEC in GAP Is Located Close to Subscribers

an emerging industry strategy for:

- Lowering operating expenses (OpEx)
- Reducing application latency
- Deploying new revenue-generating applications

The underlying concept is to deploy data center or network-class compute and storage resources in close proximity of subscribers (e.g., HFC nodes, base stations), as depicted by the GAP nodes in Figure 4. As a result, MSOs can migrate select service applications from their headend site to the network's edge. GAP facilitates MEC by defining an enclosure capable of housing high-performance compute and storage modules.

Recognizing the growing interest in MEC, the European Telecommunications Standards Institute (ETSI) has initiated the development of standards and definitions for MEC to enable real-time services and alleviate core network congestion.[3]

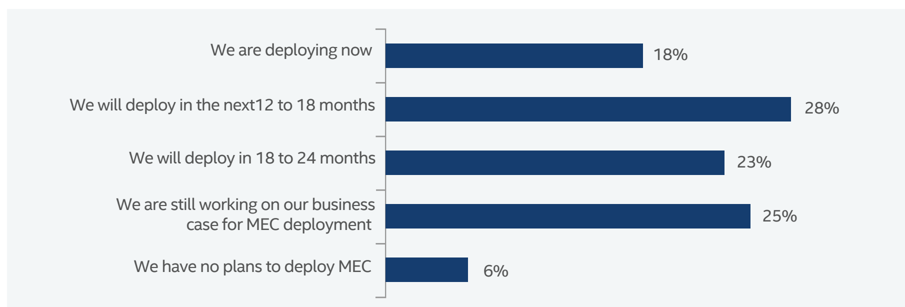

Figure 5. Survey Indicates Strong Adoption of MEC

Question: When does your company plan to start deploying MEC?

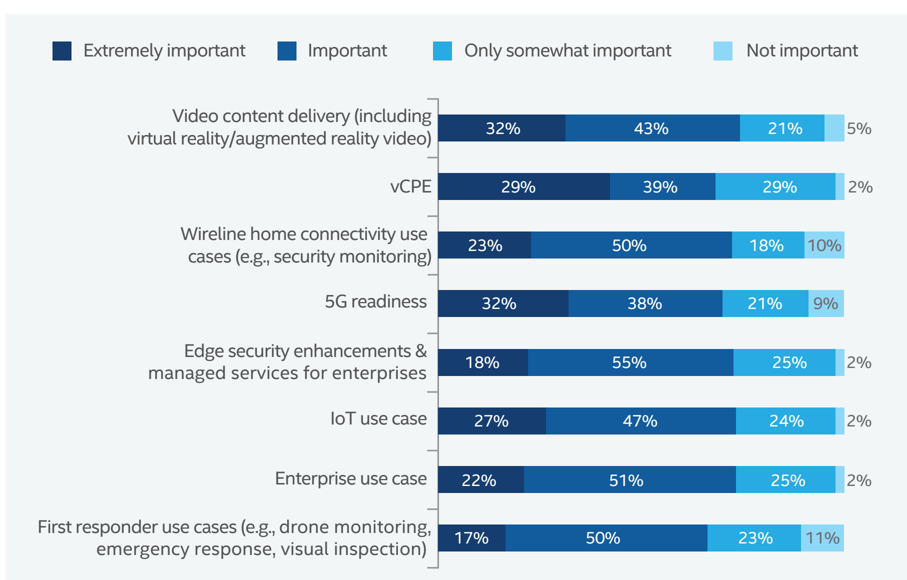

Figure 6. MEC Supports Mobile, Fixed, and Enterprise Use Cases

Question: Please rate the importance of the following use cases in driving your company's MEC deployment strategy. (N=102-107)

MEC Deployment and Use Case Survey

Heavy Reading, in conjunction with Intel, conducted a global survey of communications service providers (CoSPs) to assess MEC adoption trends. The findings in Figure 5 indicate there is a very strong commitment to MEC, with 18% currently deploying, 28% planning to deploy in the next 12 to 18 months, and 23% planning to implement thereafter.[4]

Although MEC first emerged in a mobile context, respondents indicate the importance of MEC for fixed and enterprise use cases as well. The momentum behind MEC can be attributed to lower latency, improved programmability, and the distributed spirit of the cloud. Figure 6 shows the six top-rated use cases: 5G readiness and video content delivery (both 32%), virtual customer premises equipment (vCPE) (29%), Internet of Things (IoT) use cases (27%), wireline home connectivity (23%), and enterprise use cases (22%).[5]

Services Enabled by MEC

MEC can be used to cache, store, and manipulate data that would otherwise need to be uploaded-to/download-from an MSO's core network, thus saving valuable bandwidth. Response times can be perceptibly improved when applications are run in close proximity of subscribers. Applications that can benefit from MEC include:

- Video content delivery (i.e., morning shows)
- IoT applications (i.e., machine communications)
- Wireline home connectivity (i.e., home monitoring: surveillance, temperature, appliances)
- Virtual customer premises equipment (vCPE) (e.g., Cable Modem)
- 5G readiness (i.e., low-latency services)
- Enterprise business models (e.g., extreme localization, augmented reality)

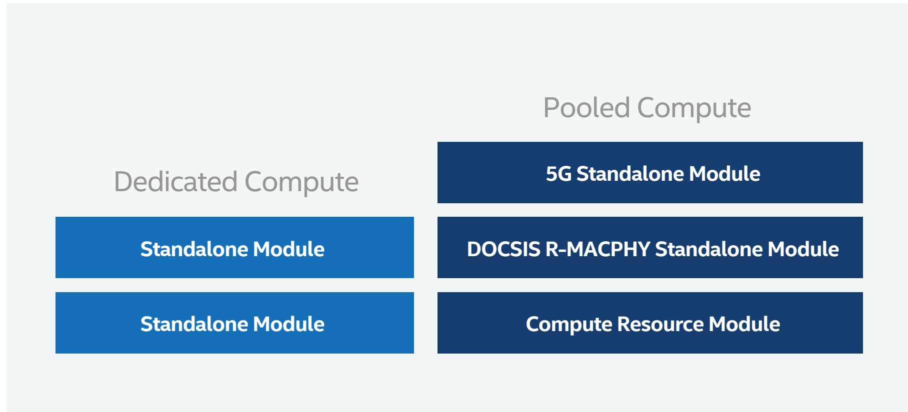

Figure 7. Two Ways to implement Compute Resources in a GAP-Compliant HFC Node

GAP Compute Module Platform

Designers of GAP-compliant HFC nodes can add high levels of compute performance with the Intel® Xeon® D-2100 processor, a system-on-a-chip (SoC) designed for low-power, high-density solutions. It integrates essential network, security, and acceleration capabilities in a form factor optimized for flexible and scalable MEC solutions.

With a range of 4 to 18 cores and up to 512 GB of addressable memory, this SoC contains a platform controller hub (PCH), high-speed I/O, and up to four 10 gigabit Intel Ethernet ports – all within a thermal design point of 20 to 110 watts. It runs the same instruction sets as Intel® Xeon® Scalable processors to provide software consistency and scalability from the data center to the network edge. The Intel Xeon D-2100 processor and other Intel® processors are designed to run in harsh environments, serving applications such as automotive and wireless.

Platform Ingredients

Intel also provides GAP module designers with other components needed to build out the infrastructure to support diverse access networks. These include high performance and reliable field-programmable gate arrays (FPGAs), Ethernet controllers, solid-state drives, and hardware accelerators. These ingredients enable MEC to perform high speed packet processing as well as low latency content delivery.

Compute Resource Design

When designing a GAP-compliant HFC node supporting multiple access technologies (e.g., 5G and DOCSIS R-MACPHY), designers have two fundamental choices for designing an access module, as depicted in Figure 7.

Pooled compute: Add a compute module and share its resources among the node's access technology modules. Compute module resources can be appropriately sized for the peak traffic of the combined modules and not oversized to simultaneously handle the peak traffic of all the modules. For example, 4G / LTE wireless traffic tends to be high during the day, while cable traffic is high in the evening.

Standalone Module: Integrate dedicated resources on each access technology module. Access technology modules are standalone and operate relatively independent of other modules. This approach is best for less intelligent modules that do not need to communicate with other modules.

Intel Proof-of-Concept

Intel developed a GAP proof-of-concept compute module using an Intel Xeon D-2100 processor and an Intel® Arria® 10 FPGA to consolidate wired and wireless access technologies on a shared, software-configurable platform, shown in Figure 8. The module can run a wide range of network functions and new services within stringent power, space, and cost constraints.

The module streams a video feed on a virtual baseband unit (vBBU) and sends/receives digital cable modem signals on a virtual cable modem termination system.

Intel's Role in GAP

Intel is now working with ecosystem partners to develop a GAP Reference Architecture that employs Intel® processors and Intel® FPGAs to create compute modules based on the most recent GAP specification. The purpose of this effort is to demonstrate the ability of a GAP node to house a powerful compute element. In addition, Intel will share design guidelines to give module vendors a higher level of confidence when developing their own compute modules.

Intel has decades of experience in bringing ecosystem components together to enable architectural transformation in the communications and IT industries. Now, Intel is applying its expertise from cloud computing to the edge, in close cooperation with GAP project members of SCTE. These efforts are expected to significantly increase a MSOs' agility and manageability of multi-technology access networks.

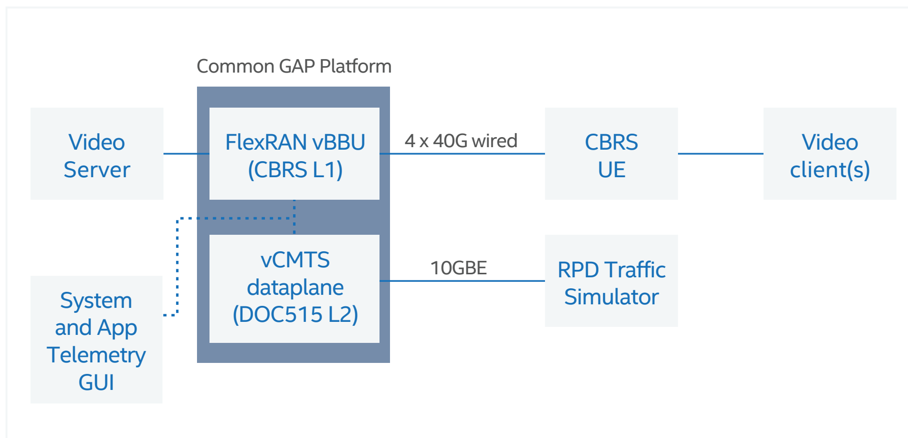

Figure 8. Intel® Gap Compute Module: Proof-of-Concept

For more Intel solutions for networking and communications, **visit https://www.intel.com/content/www/us/en/communications/communications-overview.html.**

[1] Todd Loeffelholz, "Fiber Deep Networks and The Lessons Learned From The Field," 2017, page 4, https://www.nctatechnicalpapers.com/Paper/2017/2017-fiber-deep-networks-and-the-lessons-learned-from-the-field/download.

[2] Joe Madagan, "SCTE•ISBE Standards Program 'Minds the Gap,' Creates Generic Access Platform Working Group to Promote Technical Innovation," August 6, 2018, https://www.scte.org/SCTE/News/SCTE/SCTE_News_Archive.aspx?hkey=56b678e9-98f6-4eca-afbd-4a0b6197c9a2#/publication/5b684656dfcf720004759eca/56d0ca5d54e2f90c0087312f?&sh=false.

[3] White paper, "Intel® Network Builders Program Goes to the Edge," April 2018, https://networkbuilders.intel.com/docs/intel_network_builders_program_goes_to_the_edge.pdf.

[4] Jim Hodges of Heavy Reading, "Transforming the Edge: The Rise of MEC," Q4, 2017, https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/the-rise-of-multi-access-edge-computing-paper.pdf

[5] Jim Hodges of Heavy Reading, "Transforming the Edge: The Rise of MEC," Q4, 2017, https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/the-rise-of-multi-access-edge-computing-paper.pdf

Printed in USA 1018/BB/ICMCSW/PDF Please Recycle 338327-001US