Broadband Access Transformation: Challenges and Architectures

Mark H. Mortensen
Rick Talbot

Table of Contents

Executive Overview

The access network is evolving to such an extent that it is more proper to say that it is undergoing a fundamental transformation in economics, architecture, and operations management. This paper outlines these three areas and provides the background, fundamental understanding, and projections of these changes over the next five years.

The changes are driven by several forces:

- **Economics:** Government and private equity investments. Similar to the rural electrification and interstate transportation government programs of the past century, the US and many European governments are collectively spending over $100,000,000,000 over the next decade to bring broadband services, mostly via fiber to the home, to their entire populations, driven by:
  - The societal desire to bridge the digital divide, especially in rural areas in the US and poorer households.
  - The potential economic benefits of implementing enhanced rural broadband to underserved and unserved areas.

  These government funds also often require matching funds from private equity or other sources.

  This area is covered in a companion document: *Broadband Access Transformation: Drivers and Funding.*

- **Technology:** FTTH access, middle mile, coherent pluggable optics in routers, and distributed intelligence. Technology advances in optical transport, more efficient and distributed routing technologies, and software control allow new, more efficient interconnection of the metro networks to the customers.
  - Considerable advances have been made in hardware technologies, including sophisticated passive optical networks (PONs) for fiber to the home (FTTH), coherent pluggable optics in routers, and equipment intelligence distributed closer to the customers offer new options for efficient, high-bandwidth access services.
  - The faster FTTH and high-bandwidth 5G mobile services and more distributed mobile network architectures are requiring significant upgrades in the access network, advances that are leading to new "middle mile" architectures that feed the metropolitan networks.

- **Operations:** Autonomous networks and distributed intelligence require new approaches for managing the increasingly complex networks, with distributed software.
  - The goal of moving toward autonomous networks that provides the three zero-Xs: zero touch, zero wait, and zero trouble in a network architecture that follows a philosophy of single-domain autonomy, multidomain orchestration.
  - The integration of the IP and optical transport domains.
  - The movement of network functionality and complex networking closer to the customer.

This paper is divided into sections, providing:

- An introduction to the paper.
- An overall vision of the new, transformed, local broadband network architecture.

- A review of the technology options for the enhanced access network.
- The operations architecture of the transformed local broadband network.

These sections represent the collective views of the authors, each an expert in their area.

Introduction

The digital divide is a topic of conversation in much of the world, as high-speed digital communications become an increasingly important part of everyday life. Governments in the US and Europe, in particular, have made commitments to close the digital divide by large government funded programs targeting both availability and affordability.

This drive toward broadband for all is happening at the same time as the access network is undergoing a fundamental transformation in technologies, economics, architecture, and operations management.

This paper provides an overview of the technologies and operations underpinning this transformation. It's companion document, *Broadband Access Transformation: Drivers and Funding*, covers the funding aspects.

Part 1: The New Broadband Network Architecture

The Need for Transformation

As previously addressed, an unprecedented amount of funding is being made available to deploy broadband network infrastructure to underserved, and particularly unserved, locations. However, government subsidies are contingent upon several requirements of the funding recipients. A fundamental requirement among these is that the deployed infrastructure must provide subscribers with a download speed of at least 100 Mb/s and an upload speed of at least 20 Mb/s. This requirement is beyond the capabilities of most deployed network access systems and even some systems currently offered by network equipment vendors to provide broadband access.

Simply meeting the requirements of government subsidies does not guarantee that the deployed broadband infrastructure will meet the network operator's needs. Once new broadband services of at least 100 Mb/s are offered, subscribers will undoubtedly increase their data demand with streaming and video-based security services, large file transfers, and online gaming. This tsunami of data demand per subscriber will likely exceed 100 Mb/s in the coming years. Since service providers will need their newly deployed broadband network infrastructure to suffice for several years, probably out to 2028, they will seek totally new solutions to transform their broadband networks.

Segmentation of the Broadband Network

As network operators have been planning new infrastructure to enable true broadband services, much of the industry's focus has been on the broadband access network. This focus is natural since most of the spending will be in that segment of the overall network. However, all services require the reliable operation of the total end-to-end network, not just access. Specifically, an enhancement of the access network that enables higher data speeds will stimulate higher traffic throughput that will be seen throughout the network. As depicted in Figure 1, the traffic transported by the access network is passed on to a middle mile network and then on to the metro network and finally the core network.

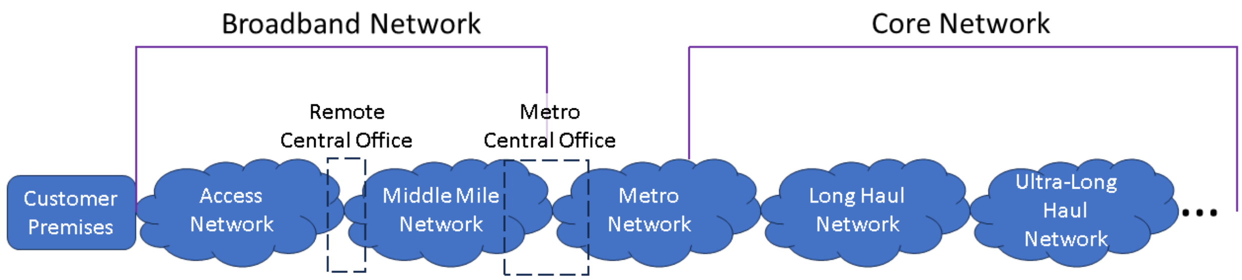

Figure 1. Broadband Network Segments in Overall Network

Of these various network segments, the middle-mile network is the most sensitive to increases in traffic from the access network; this is the segment that must be enhanced with any associated enhancement of the access network. This association of the broadband access network with the middle-mile network results in a combined segment called, simply, the broadband network.

The Broadband Access Network

The fundamental role of the access network in telecommunications is well established: it aggregates all traffic from end-users in the access area and delivers the traffic to (and/or distributes digital content from) a central office[1] that routes that traffic to (and/or obtains digital content from) various destinations over local and long-haul networks. Neither the Infrastructure Investment and Jobs Act (the Act which established the BEAD program) nor the Notice of Funding Opportunity (NOFO, which was written by the NTIA to provide rules for the BEAD program) mentions the broadband access network or broadband access infrastructure. For the Act and the NOFO, the access addressed is simply internet access. The concept of a network is assumed to be whatever it takes to provide broadband internet access (coinciding with the

[1] Not all broadband access operators will term their broadband access network termination site a central office. MSOs often refer to headends; Internet Service Providers (ISPs) often use the term Point of Presence (POP); fixed wireless operators refer to their antenna tower site; and operators that will launch their access networks directly from data centers will simply refer to the data center (if not POP). For simplicity, this report will continue to use the term central office unless the other term is explicitly called for.

concept of the broadband network of this paper). Thus, the rules laid out for the entire broadband network by the BEAD program apply equally to the broadband access network:

- The network must provide reliable broadband service with speeds no less than 100 Mb/s downstream and 20 Mb/s upstream.
- 95% of latency measurements of the network must read below 100 msec round trip delay.
- Outages on the network are not to exceed, on average, 48 hours over a 365-day period.
- The costs of the network must be low enough to enable the network's operator to meet the following requirements:
    - Offer at least one low-cost broadband service option,
    - Participate in the Affordable Connectivity Program (ACP) and
    - Refrain from implementing data usage caps for subscribers.
- The network operator must permit other broadband service providers to interconnect with middle-mile network facilities on a just, reasonable, and nondiscriminatory basis.

The federal funding requirement for broadband service speed may be 100 Mb/s downstream and 20 Mb/s upstream, but technical capability and commercial factors are driving higher speeds. Currently available FTTH solutions easily provide symmetric speeds above 100 Mb/s, and an FCC study confirms that network operators are offering significantly higher speeds[2]. Though subscriber data demand is currently under 100 Mb/s[3], that demand is growing at approximately 20% annually[4]. At this growth rate, an average broadband download speed will reach 69 Mb/s by 2028. Thus, competition will drive offered speeds while the overall network will need to support approximately 70 Mb/s per subscriber.

The BEAD NOFO's 100 ms latency limitation is easy to achieve for most broadband access technologies, but commercial issues, once again, will be the controlling driver for latency limitation. Many consumer applications require latencies of considerably less than 100 ms. Further, if 5G Radio Access Network (RAN) mid-haul is to be supported, the latency must be less than 5 ms, and the end-to-end 5G latency is limited to 30 ms.

The Middle-Mile Network

The Act also defines and funds the middle-mile (network) infrastructure. The purpose of the middle mile, as articulated in the Act, is:

---

[2] The FCC's Twelfth Measuring Broadband America Fixed Broadband Report found that:

- "The weighted average advertised download speed of the participating ISPs was 307.73 Mbps…."
- "For most of the major broadband providers that were tested, measured download speeds were 100% or better than advertised speeds during the peak hours (7 p.m. to 11 p.m. local time).

[3] Assuming a current 25 Mb/s household demand, considering a 25 Mb/s for Ultra HD 4K Video and multiple active devices per home.

[4] According to quarterly OpenVault Broadband Industry Reports since 2019

(A) to encourage the expansion and extension of middle mile infrastructure to reduce the cost of connecting unserved and underserved areas to the backbone of the internet (commonly referred to as the ''last mile''); and

(B) to promote broadband connection resiliency through the creation of alternative network connection paths that can be designed to prevent single points of failure on a broadband network.

The definition of middle-mile infrastructure that is provided by the Act, however, is rather nebulous. The Act states that middle-mile infrastructure "means any broadband infrastructure that does not connect directly to an end-user location, including an anchor institution…," and then the Act lists a multitude of infrastructure types that are included in this definition. From a network perspective, the infrastructure that does "connect directly to an end-user location" is the broadband access network. The middle-mile infrastructure, then, connects the broadband access network to the internet.

Broadband access networks are terminated in subscribers' local (sometimes called serving) central offices. The connection to the internet, in turn, is made at an edge router. In the case of broadband services, the administration of the subscribers' services, which is required regardless of the broadband access technology employed, takes place at this edge router. For the purposes of this paper, a PON network example is used.

The edge router used to administer a PON network is called a Broadband Network Gateway (BNG, which will be described in another section). The middle-mile network connects one or more PON networks to the BNG. In the past, the BNG was centralized, often located at the interface of the metro with the core network. However, as internet traffic increased, the BNGs migrated toward the edge of the metro network, ultimately resulting in the BNG residing in the serving central office of the access area.

As network operators begin to deploy new broadband access infrastructure, the targeted unserved and underserved areas will be served by central offices that lack a BNG and will often be outside of the metro network (and increasingly in outright rural areas). Metro central offices that receive new broadband access systems either already contain an edge router that can be equipped with the BNG function or will be supplied one. However, small central offices located in remote areas that receive the new broadband access systems will likely continue to lack those BNGs. The connection between a broadband access system located in a remote central office and the BNG located in a metro central office is the middle-mile network. From a transport perspective, the middle mile is simply a backhaul connection linking a remote central office to a metro central office. A closer look at the connected broadband access network and middle-mile network, depicted in Figure 2, reveals the interfaces of the middle mile network.

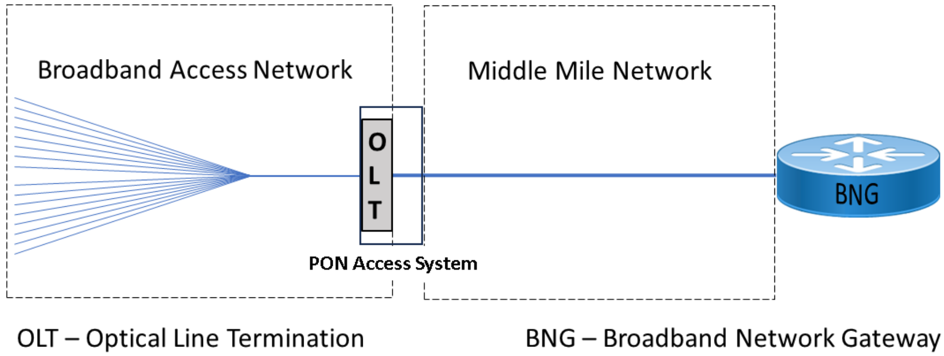

Figure 2. Current Broadband Network Segmentation

The broadband access used in this example is a PON. The optical line termination (OLT) is a module that resides in a PON access system (which supports multiple OLTs). Note that the middle-mile network interfaces the PON access system (in the remote central office) on one end and the BNG (in the metro central office) at the other end. Similarly, the middle-mile network includes the access-facing ports of an edge router that serves as the BNG.

IP Functions in a Middle-Mile Network

Though the middle mile is currently a transport network, enhancements to the broadband network will expand the domain of the middle mile. IP functionality will be added to the broadband access systems (addressed later, in the Broadband Access Technology section), which will enable coordination between the broadband access network and the BNG as well as with the connection between the two (addressed in the Management and Control section). In addition, coherent optical modules will be introduced to both the broadband access system and the BNG (addressed in the Middle-Mile Technology section). The resulting next-generation middle-mile network is depicted in Figure 3.

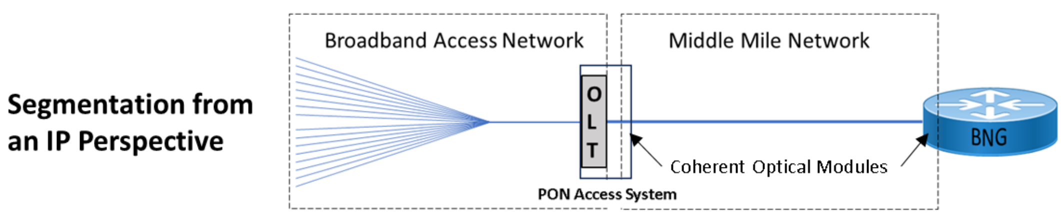

Figure 3. Next-Generation Broadband Network Segmentation

The coherent optical modules, which provide the transponders for the middle-mile optical systems, are now part of the middle-mile network.

The Middle Mile as a Multi-Service Network

The middle-mile network does more than provide high-capacity transport of the aggregated broadband access network traffic between the broadband access system and the BNG. It also provides transport for private lines and Radio Access Network (RAN) mid-haul and backhaul connections to the metro central office. These connections are depicted in Figure 4. In this figure, only one remote office is served by the optical ring, but in practice the ring would serve multiple remote offices. The mid-haul lines are shown connecting directly to the Central Unit (CU) in the remote central office. In practice, they could be connected via an optical transport system or coherent optical modules in the Distributed Units (DU) to a CU located in either the remote central office or in the metro central office. Further, wavelength services could connect to the optical ring either directly or through the remote central office.

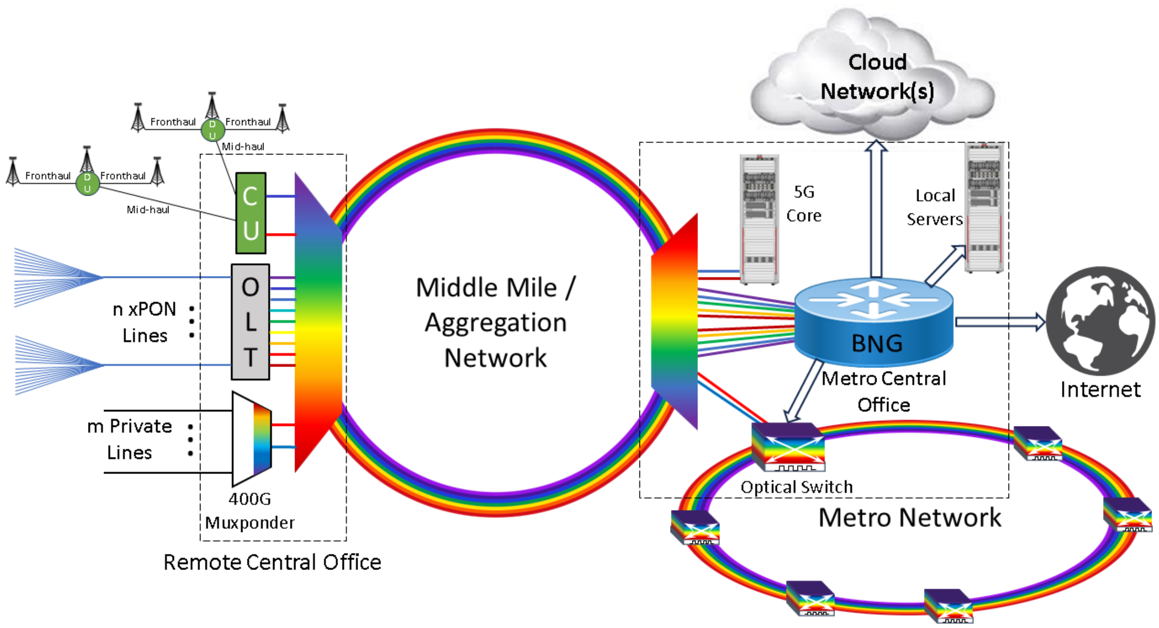

Figure 4. Multi-Service Middle-Mile Infrastructure

On the left of the figure, the remote central office contains:

- A CU for a 5G RAN to the upper left of the central office. The RAN consists of two DUs that are connected to the CU via mid-haul and unpictured remote units (RUs) at six cell towers that are connected to the two DUs via fronthaul.
- A PON access system (depicted as OLT, representing the OLTs that terminate the xPON lines).
- A 400G muxponder that aggregates the private lines.

The CU, PON access system, and 400G muxponder all employ coherent optical modules to present 400G wavelengths to a DWDM multiplexer that aggregates them into an optical line system. In this figure, the optical system is configured as a ring that connects the remote central office to the metro central office (on the right of the figure) that, itself, is connected to the metro optical network (depicted as another ring) via an optical switch.

In the metro central office:

- The wavelengths from the CU (which is RAN backhaul) are connected to the 5G core,
- The wavelengths from the PON access system are connected to the BNG, and
- The wavelengths from the 400G muxponder are connected to the optical switch.

The 5G core, BNG, and optical switch all employ coherent optical modules to terminate the 400G wavelengths from the remote central office. The BNG connects subscribers served by the PON access system to:

- One or more cloud networks,
- Servers located in the metro central office,
- The internet, and
- The optical switch (provides connections for any private lines served by the PON access system).

Scaling the Middle-Mile Network

The required capacity of broadband access systems is determined by how many 70 Mb/s subscribers are supported by each system. Until BNGs are placed in remote central offices, all broadband service traffic aggregated in the central office will be sent via the middle-mile network to the BNG of the system since broadband service access to the internet is through the BNG. The required capacity of the middle-mile network to support a single broadband central office is therefore determined by multiplying the number of broadband subscriber lines served by the central office by the traffic of each line.

This paper uses a figure of 12,000 lines for the total lines terminated in suburban central offices[5] and 660 lines for the total lines terminated in rural central offices[6]. Not all end users in the service area of the central office will subscribe to a broadband access operator's service:

- 14% of households will not subscribe to broadband service[7];
- Some areas will be served by multiple broadband providers, which will capture some of the subscribers[8].

Combining these factors yields three values for the amount of broadband traffic that must be supported by a middle-mile network operator:

- From a suburban central office in which there is broadband competition, 423 Gb/s, requiring two 400G DWDM wavelengths for transport,

---

[5] *FCC Report on Major LECs as of 2000*. These figures produce an average number of lines per central office (of the major local exchange carriers of approximately 16,000 lines. As subscribers have cut the cord, the figure of 16,000 lines per central office in 2000 have declined. For the purposes of this paper it is assumed that this decline is 25%, yielding a current average number of lines per central office of 12,000.

[6] Universal Service Fund 2022 Submission of 2021 Study Results by the National Exchange Carrier Association, Inc. (NECA).

[7] December 22, 2022 LRG press release: 90% of U.S. Households Get an Internet Service at Home:

- 90% of households use a laptop or desktop computer.
- Of those that use a laptop or desktop computer at home, 96% have an internet service at home.
- Broadband accounts for 99% of households with an internet service at home.

[8] Assume the following competitive effects:

- In a suburban central office in which there is broadband competition, that competition takes 40% of the broadband service subscribers,
- Some suburban central offices will have no broadband competition, and
- There is no broadband competition in a rural central office.

- From a suburban central office in which there is no broadband competition, 705 Gb/s, requiring two 400G DWDM wavelengths for transport, and
- From a rural central office, 39 Gb/s, requiring one 100G DWDM wavelength for transport.

Because part of the stated purpose of the middle-mile infrastructure is "to promote broadband connection resiliency," the middle mile is likely to be deployed as a ring, collecting traffic from multiple central offices.

Part 2: Technologies Used for the New Broadband Network

Access Network Technology Options

Active Ethernet is the use of standard Ethernet (GbE, 10GbE, 25GbE, 50GbE and 100GbE) on a fiber strand or pair to each subscriber. The technology is simple and easy to manage, but it requires a separate Ethernet port at the serving central office for each subscriber, and there is no standard method of providing subscriber management. Because this technology employs separate fiber pairs for each subscriber and separate fiber ports at the central location, this solution is significantly less cost-effective than PON. This added expense likely prevents the service provider from meeting the service affordability requirement.

PON is a single fiber shared by a number (16, 32, 64, 128 or 256) of subscribers. The termination of the PON is an OLT located at the serving central office. The OLT institutes and administers a number of time slots on the PON for use by the subscribers, identifying each subscriber by its time slot. Because of the generous bandwidth of the fiber and advancements in metro optical networking, the ultimate capacity of PON systems is only limited by the speed, power and space requirements of the electronics.

There are two types of PON systems: Ethernet-based PON (EPON) and ITU-based Gigabit PON (GPON). The speeds of the variants of EPON are specified by the total system speeds (not the speed available to a subscriber). These are GbE-PON, 10G-EPON, 25G-EPON, 50G-EPON, all of which support symmetrical subscriber service speeds. These indicated speeds are not exactly the capacity shared by the subscribers. Ethernet employs a total bitrate that is higher than the indicated speed; for example, GbE actually runs at approximately 1.25 Gb/s to support overhead bits of the GbE. Some of the time slots of the system are used for monitoring, controlling, and synchronizing the system, limiting the delivered subscriber service speeds. EPON systems can be implemented on a pair of fibers (transmit on one and receive on the other) or a single fiber (one wavelength for transmit and another wavelength for receive).

The speeds of the variants of GPON are also specified by the total speeds of their shared lines. These are currently GPON, XGS-PON (10 Gb/s symmetrical PON) and 25G-PON; 50G-PON and 100G-PON have been demonstrated and supposedly will be ready when needed. As does EPON, these systems all provide symmetrical speeds. Whereas GPON systems run at their designated speeds (unlike EPON), they, like EPON systems, use some of the time slots of the system for monitoring, controlling and synchronizing the system. For example, a GPON user that subscribes to a Gigabit service cannot receive quite the full 1,000 Mb/s.

A list of the vendors of PON solutions includes:

- Adtran

- Calix
- Ciena
- DZS
- Nokia

Hybrid fiber coax (HFC) uses fiber from the cable system distribution facility (headend or hub) to an optical node that translates the optical signal to radio frequency (RF) and sends it over coaxial cable lines to subscribers. Cable operators primarily use it. Cable operators use a protocol Data Over Cable Service Interface Specifications (DOCSIS) to enable high-speed bandwidth data transfer over the existing coaxial cable systems that were originally designed for cable TV signals. DOCSIS is in its 4th generation (DOCSIS 4.0) with each generation delivering more throughput. DOCSIS 4.0 combined with bringing optical nodes closer and closer to subscribers and distributing the access network will enable the cable industry to reach its 10G goal to deliver 10Gb/s symmetrical with low latency.

Fixed wireless access (FWA) uses radio links between two fixed points to provide wireless broadband. It consists of a base station connected to a fixed network. FWA has recently emerged as a major technology alternative for broadband deployments, particularly in areas where the cost of deploying fiber is prohibitive. The introduction of C-band 5G spectrum at the start of 2022 opened the door for faster connections. Recent innovations are also mitigating some of the inherent challenges of FWA such as interference and low throughput. Today, it is possible to deliver 100Mbps downstream and upstream with newer FWA technologies.

One issue to consider regarding FWA is that, because its spectrum is shared, it only supports a fixed number of subscribers for a given service speed. This limitation is unlikely be a problem since:

- For areas with high subscriber densities, such as in the metro, the upper spectrum used limits the reach of the signal, resulting in relatively small coverage areas and thereby limiting the number of subscribers. For these FWA systems there will be many towers in the overall coverage area, requiring a large number of middle-mile network connections.
- The use of FWA in rural areas is ideal. The low subscriber densities in these areas will allow for the use of lower spectrum radios that have significantly increased reach. FWA in these applications may be particularly cost-effective, particularly compared to terrestrial solutions (such as PON and HFC), which would be the broadband alternatives to FWA.

To date, FWA has largely been focused on urban and suburban areas, in part due to the availability of midband 5G spectrum in those areas. Going forward, the growth of FWA will come from suburban and rural areas. Wells Fargo's analysts predict FWA, in general, will capture around 12−13% of the overall US broadband market and fully 10% of the residential broadband market by 2025[9].

[9] https://www.lightreading.com/broadband/fixed-wireless-access-(fwa)/fwa-to-remain-biggest-disruptor-through-2024/d/d-id/785484

Satellite Systems

These broadcast and receive RF signals between a satellite and subscribers in a relatively wide area. Satellite system types are based on the type of orbit the satellite uses:

- Low Earth Orbit (LEO): This type of satellite orbits at an altitude of up to 1,200 miles. The round-trip delay of signals to and from a LEO satellite would meet the latency requirement, but the reliability and cost-effectiveness of the solution are questionable.
- Medium Earth Orbit (MEO): This type of satellite orbits at an altitude between 1,200 miles (the maximum LEO altitude) and 22,236 miles of a Geosynchronous Orbit. Most of this region, that above 4,650 miles in altitude, would cause unacceptable round-trip delay. The placement of satellites in this orbit would be expensive because of the additional rocket boost requirements.
- Geosynchronous Orbit (GSO): At an altitude of 22,236 miles, a satellite in GSO has a 24 hour orbital period, matching the rotation of the Earth and thereby remains over a single spot on the Earth's surface, but its round-trip delay of 478 ms is far too great to meet the requirements of the government funding programs.

Twisted Pair Copper

- Digital Subscriber Line (DSL): Several signal modulation techniques are used (hence xDSL) to achieve speeds of 6 Mb/s (Asynchronous DSL–ADSL) up to 100 Mb/s (very-high-bit-rate DSL–VDSL). However, the quality of older copper cables can become degraded, slowing DSL speeds, and interference from other DSL services in nearby cable pairs can limit DSL speeds. The speed of the service also decreases with the length of the copper pair, with VDSL only operated over a very short pair length. These factors render DSL incapable of supporting federally-subsidized broadband.
- G.fast is an ITU standard DSL modulation technique that can provide from 100 Mb/s to 1 Gb/s speeds on very short lengths of copper pairs, less than 500 meters, making it impractical and costly.

Middle-Mile Network Technology

Broadband Network Gateway (BNG)

As noted previously, the administration of subscribers' services is required regardless of the broadband access technology employed, and for the purposes of this paper, a PON network example is used. This section describes the subscriber service administrative element of a PON system, the BNG. Beyond its service administrative function, the BNG is the edge router of the internet for the broadband infrastructure, connecting subscribers to the internet, which is the central purpose of the BEAD program. The BEAD requirements apply to an edge router that serves as a BNG. The underlying router requirements for the BNG are practically universal for edge routers:

- **Routing Protocols**: The edge router must support a variety of routing protocols, including OSPF, BGP, RIP, IS-IS, MPLS, MPLS SRv4 and v6, SRv6 and EIGRP.

- **Firewall Features**: The edge router must have a powerful firewall with features such as NAT, IPSec, DoS protection and URL filtering.
- **Load Balancing Features**: The edge router must be able to load balance traffic between multiple WAN links, providing redundancy and improved performance.
- **VPN Features**: The edge router must be able to create VPN tunnels for remote access and site-to-site connectivity.
- **Security Features**: The edge router must, itself, be secure and have the aforementioned firewall protection as well as having intrusion detection and DDoS protection.
- **High-Availability Features**: The edge router must be highly available, having features such as hot standby and failover.
- **Scalability Features**: The edge router must be extremely scalable to be able to handle the increasing traffic demands of the network, both the termination of numerous broadband access systems as well as all the business traffic arising from the access areas it serves.
- **Management Features**: The edge router must be easy to manage and configure, offering web-based GUI, CLI, and SNMP support.

BNG Functions

Added to the requirements of the underlying edge router of the BNG are a set of features that specifically support the broadband services and are distinct to the BNG. These features include:

- **PPPoE Support**: The BNG router must be able to support PPPoE for subscriber authentication and traffic management.
- **DHCP Support**: The BNG router must be able to provide DHCP services to subscribers.
- **NAT Support**: The BNG router must be able to provide NAT for subscribers, allowing them to access the internet with private IP addresses.
- **CGNAT Support**: It is recommended that the BNG router implement CGNAT, which can help to conserve public IP addresses and increase the scalability of the BNG.
- **Traffic Shaping Support**: Because the broadband service providers will offer services with a variety of service speeds, it is recommended that the BNG router be used to shape traffic, ensuring that subscribers receive the bandwidth for which they are paying.
- **QoS Support**: Similarly, because the broadband service providers will offer services with different priorities, it is recommended that the BNG router be used to provide QoS for subscribers, ensuring that critical traffic is given priority.
- **High-Availability Features**: The BNG router must be highly available and have features such as hot standby and failover.

A list of the vendors of BNG solutions includes:

- Cisco Systems
- Juniper Networks
- Nokia

BNG Disaggregation

Because BNG functionality is implemented in software, this BNG software can be disaggregated from the edge router, itself, and run at a separate location (Figure 5). In a move aimed at simplifying the network at most locations, particularly remote locations, as well as saving space and cost, the industry has begun disaggregating many network functions, including BNG functions, which had traditionally been centralized. This separation of the function's user plane from its control plane enables network operators to cost-effectively implement these functions anywhere in the network, leaving the user plane to use smaller, simpler, and more cost-optimized platforms such as simple aggregation routers or switches. These scale-out architectures also allow operators to:

- Build their networks incrementally based on demand,
- Improve network performance and resiliency, and
- Provide a natural evolutionary path to the next-generation architectures of the future.

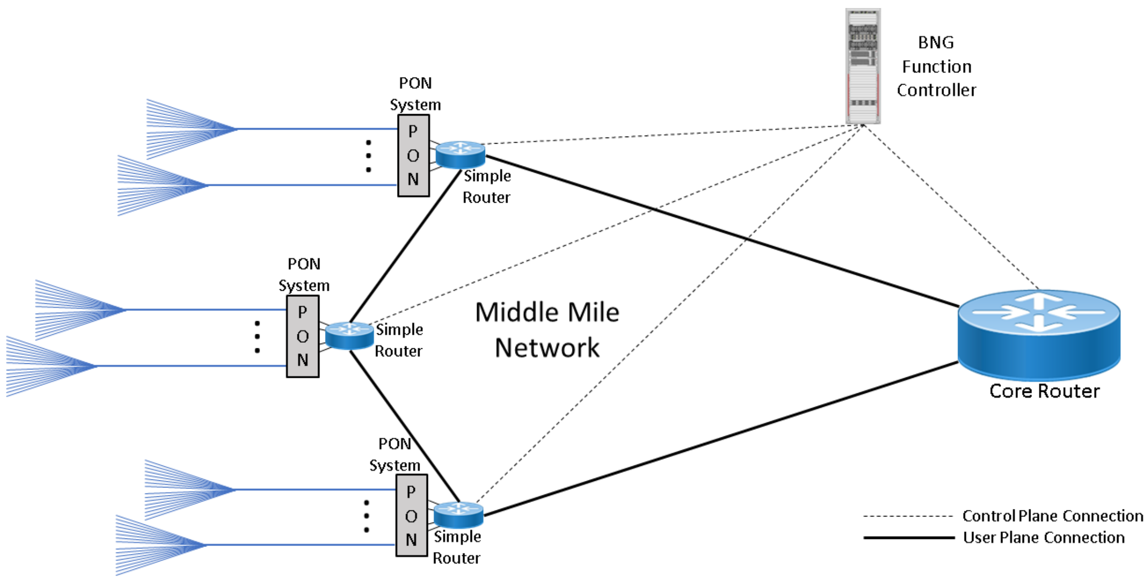

Figure 5. Disaggregated BNG Function

This emerging broadband network architecture introduces new models of network operation:

- **Distribution**: Making smaller ongoing investments when adding capacity, and incrementally scaling BNG servers in step with demand;
- **Disaggregation**: Using BNG Control / User Plane Separation (BNG CUPS), splitting BNG control and user plane functions with control plane functions remaining centralized while the user plane moves closer to subscribers to meet growing capacity and performance demands more readily. Similarly, operators are increasingly distributing routing functions closer to users, employing an IP services fabric composed of smaller, more compact, and more cost-optimized platforms.
- **Virtualization**: Using servers rather than complex routers or other specialized equipment at the centralized location to run the BNG software creates virtual BNG. In addition, router functions, themselves, can be run by a separate server, resulting in virtual routing.

- **Convergence:** BNG CUPS provides a future-ready architecture that makes it easier to introduce next-generation innovations as well as controlling multiple service networks from a common core. For example, operators will be able to control wireline and wireless services from a single network core, achieving Wireless Wireline Convergence.

The importance of the disaggregated BNG to the broadband network is fundamental. Routing can now be performed at even remote central offices, and the BNG functions controlling that routing hardware enables the broadband access system to be connected to a virtual BNG at the local central office. That BNG would remove the overhead required on middle-mile connections as well as filter traffic bound for the core of the network, reducing the load on connections with the metro network as well as enabling subscribers to access the cloud at the extreme edge of the network.

A list of the vendors of disaggregated and virtual BNG solutions includes:

- Casa Systems
- Ciena (Benu Networks)
- Cisco
- Juniper
- netElastic
- RtBrick
- Sanctum Networks
- UfiSpace
- Waystream AB

New Transport Network Technology

Coherent Optical Modules

Improvements in the electronics of optical transponders have enabled the increase of optical transport performance over the years, from 100 Gb/s to 400 Gb/s to 600 Gb/s to the current 800 Gb/s (and soon 1.2 Tb/s). Simultaneous improvements in optical components have helped (and in cases were necessary), but the reduced architectural scale of several of the electronic components (such as the DSP, DAC and ADC) were the fundamental enablers of the increased optical performance. The reduction of electronic component scale from 40 nm to 5 nm has increased the baud rate from 30 GBaud to 130 GBaud. Of similar importance is the reduction of space and power enabled by these smaller electronic components. Once the scale reached 7 nm, the power consumption of these components could be reduced to the point of producing a 400 Gb/s transponder module that could fit in the QSFP-DD 400G ports of a switch/router.

The Optical Internetworking Forum has leveraged this development to issue an implementation agreement for a standard 400G coherent optical pluggable module called the 400ZR. The first application for the 400ZR was data center interconnection, which seldom required a reach of more than 120 km, the reach of the 400ZR module. This 120 km reach is also likely to be sufficient for most middle-mile network connections, eliminating the need for separate transponder-based optical networking systems, greatly reducing cost, space and power consumption. These reductions are particularly important in remote central offices.

A list of the vendors of coherent optical pluggable modules include:

- Acacia/Cisco
- Ciena
- Infinera
- Nokia
- Numerous component manufacturers that sell directly to hyperscale data center operators

IPoDWDM

As the network becomes predominantly a set of interconnections between routers and/or switches, and this is the case for middle mile networks, the need for optical transport systems is fading into irrelevancy. The new network architecture that emerges is based on wavelength connections between the coherent optical pluggable modules housed in the interconnected routers/switches (Figure 6).

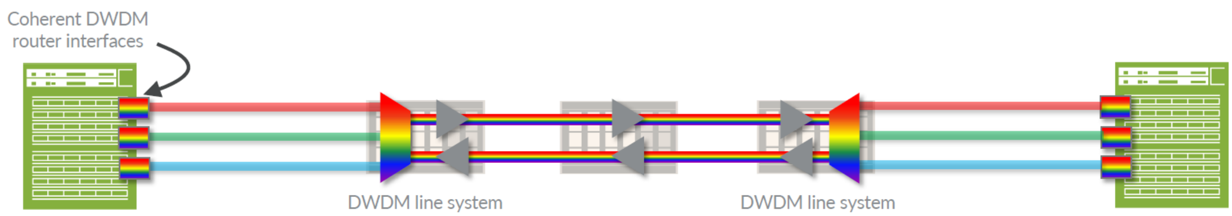

Figure 6. IPoDWDM Using Coherent Optical Pluggable Modules

Network architects foresaw this architecture more than a decade ago. These architects named the architecture IP over DWDM (IPoDWDM). However, at that time DWDM transponders were large and dissipated significant power, so much so that if they were plugged into a router, the module would take up an inordinate amount of space on the router, which was very expensive at that time. Since the router cost has declined, and now that the coherent optical pluggable modules are becoming available, the IPoDWDM model has become practical, at least in most networks. Granted, the term IPoDWDM has so fallen out of favor that the marketing departments of the routing vendors claim that the interconnection of routers equipped with coherent optical pluggable modules is something new and better than IPoDWDM, but that is only semantics. In reality, this is IPoDWDM.

Part 3: The New Intelligent Local Broadband Operations Architecture

The access network used to be considered merely a way of connecting the customer's device to the functions provided in the local central office and beyond to a metropolitan network. This architecture is changing as access is increasingly moved to fully digital PON technology that is connected to a middle mile to the metropolitan network. This is creating a new, distributed intelligent local broadband network architecture for both mobile and fixed networks that includes substantial functionality beyond mere connectivity. This includes several capabilities:

- Enhanced end-user services supported by devices such as MEC and BNG platforms.
- Enhanced flexibility in the deployment and operations of the intelligent local broadband network.

- Movement of packet and computing technologies closer to the customer.
- Tighter coupling between optical and packet technologies.
- Disaggregation of operations control via a control-/user plane separation (CUPS) architecture for several classes of network elements.
- Slicing of the local broadband network (and beyond) to support end-to-end QoS service guarantees.

However, this trend of distributing intelligence is pushing against another technological trend: less expensive, higher bandwidth regional networking from the introduction of coherent plug-ins for optical transport. This will allow less expensive backhauling to more centralized network intelligence, the more classic arrangement. However, it is inevitable that the more distributed, intelligent, local broadband network architecture will prevail.

In this section, we provide our answers to the following questions:

- How does the growing concept of autonomous networks play in the transformed local broadband network operations?
- What are the new operations requirements of the transformed local broadband network?
- What is the status of vertical element disaggregation of operations control (CUPS architecture), and how will the CUPS architecture be operationally implemented?
- What is the vision for operations control architectures that will make operations easier, faster, less expensive?

The Promise of Autonomous Networks

Autonomous networks (AN) is a concept from the TM Forum, very similar, and converging with, the ETSI concept of zero-touch networks. It posits a network (with its supporting software systems) that can configure and manage itself and the services that it supports. The CSPs who have signed onto the AN program at the TM Forum represent 38% of the worldwide telecoms market, meaning that this concept has moved from the early adopters to the early majority stage of acceptance.

Autonomous networks have three main characteristics:

- Zero-wait as changes to the network to account for a new service or service instance or adaptation to changing network and traffic conditions are done immediately,
- Zero-touch where the network requires no human intervention,
- Zero-trouble where any problems are either proactively predicted and mitigation efforts automatically applied or faults are fixed so fast that no trouble is detected by the users.

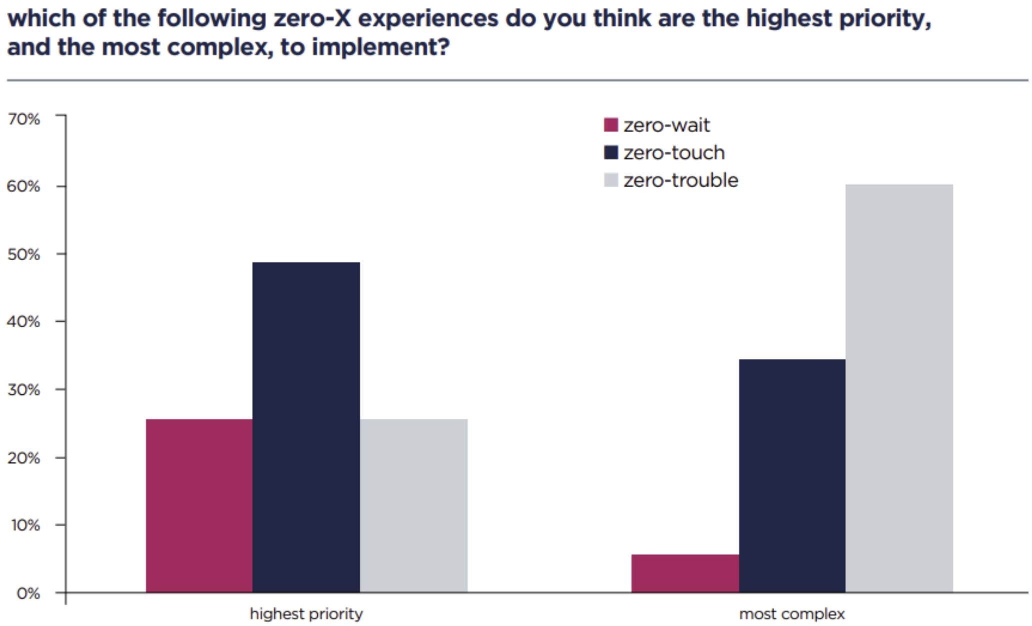

Figure 7. Characteristics of Autonomous Networks

Zero-touch is considered by CSPs to be the most important characteristic as indicated in a survey conducted by the TM Forum[10]. When asked which capability is the most complex to implement, 60% of respondents said zero-trouble.

The AN is also software architectural framework of autonomous, cross-orchestrated domains to implement this vision. The basic concept is to divide the network resources, be they physical or virtual, into domains that serve as autonomous but interconnected building blocks for the network with the operations following those same boundaries. These are depicted as resource closed loop (four in Figure 8).

With resources divided into the autonomous domains, there is a need for a coordination layer across the domains (3), the service operations closed loop, identified as **cross-domain network orchestration**. Business operations are another layer above, with, again, closed loop operations (2). This is **cross domain business orchestration**. Topping it all off, self-directed networking defined as **business intent** connects the user to the entire stack to give a **user closed loop operation** (1).

[10] Mortensen, Mark H., *Leveling up: achieving Level 3 autonomous networks and beyond*, September 2023. https://inform.tmforum.org/research-and-analysis/reports/leveling-up-achieving-level-3-autonomous-networks-and-beyond

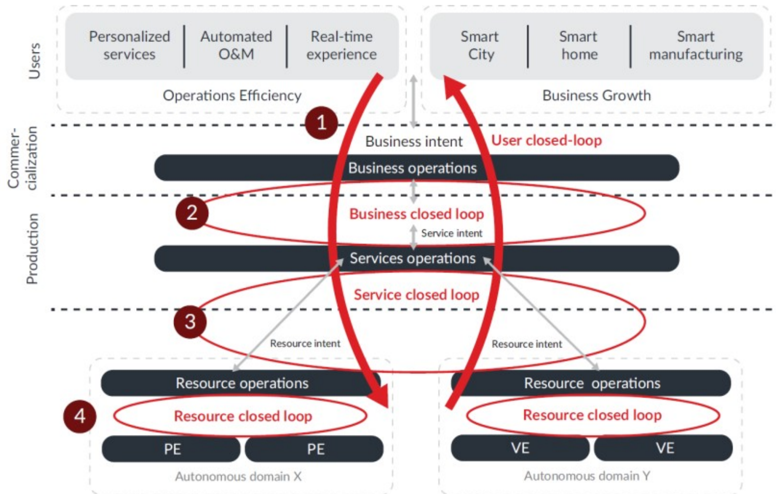

Figure 8. AN Framework (Source: TMForum, 2020)

The net effect is to have an overall closed-loop operation that extends from the top of the operations structure through the business operations of the CSP, to the network resources, through the hierarchical layers, based on the concept of autonomous domain-based operations at the bottom of the stack using a philosophy of **single-domain autonomy, multidomain orchestration**.

In the context of the new local broadband architecture, the access (PON) represents one domain, while the middle mile is usually broken into an IP/MPLS and optical operational domains. Each domain will operate autonomously, with cross-domain coordination.

PON and Wireless Access Networks Require an Updated Intelligent Access Operations Architecture

The access network used to be considered merely a way of connecting the customer device to the functions provided in the local central office and beyond. However, a distributed local broadband network operations architecture is developing in mobile and fixed networks that includes substantial functionality, beyond mere connectivity. This includes several capabilities:

- Enhanced end-user services supported by devices such as MEC and BNG platforms.
- Enhanced flexibility in the deployment and operations of the access and middle mile domains.
- Movement of packet and computing technologies closer to the customer.
- Tighter coupling between optical and packet technologies.
- Disaggregation of operations control via a CUPS architecture for several classes of network elements.
- Slicing of the local broadband network to support end-to-end QoS service guarantees.

Middle Mile Creates a New Management Domain

The control and orchestration of the transport domain in PON and middle mile scenario will need to undergo significant enhancement to consider:

- Disaggregation of the PON and middle mile transport equipment into multiple connected boxes.
- Virtualization of some of the transport equipment.

- Possibility of multiple vendors' equipment involved in the transport PON and middle mile.
- New low latency requirements of some 5G services.
- Introduction of multi-domain slicing of the 5G network into separate slices for services with different characteristics and perhaps for individual customers also. This has the potential to swamp the operations team without substantial automation.

The implications are still being understood by CSPs and vendors. But some of the major ones are shown in the table:

OPERATIONS FUNCTION	TECHNICAL IMPLICATIONS
Planning the Network Capacity & Configuration	• Planning systems for hybrid PNF/VNF PON and middle-mile transport. • Planning of slicing across PON and middle-mile and other domains. • New functions for planning for latency & slicing, probably involving simulations of black-swan events in addition to the normal algorithmic methods. • Planning systems will be more likely to come from independent software vendors (ISVs) and systems integrators (SI) rather than network equipment manufacturers, since multivendor PON and middle mile requires the cooperation of vendors in providing specifications.
Creating & Instantiating VNFs	• Standard processes and systems will be used, currently Kubernetes. • May be different systems for different vendors unless home-grown or other multi-vendor systems are employed.
Configuring the PNFs and VNFs	• Resource provisioning must be enhanced to cover new slicing, latency, and perhaps other requirements. • Disaggregated network elements will require more coordination and configuration.
Measuring Health & Performance (Service Assurance)	• Latency service quality metrics required • Multi-vendor configuration management a requirement for full automation (either as a single, multi-vendor system or sub-domain management systems with a manager-of-managers at the domain level).
Managing the Domain	• Diagnosing subtle failures and sub-par performance in a multi-vendor system will require substantial new functionality, (or manual effort), especially for latency requirements.

	• Probes may well have to be an important part of the architecture.
Cross-Domain Orchestration	• Northbound interface to cross-domain orchestration systems required for multi-domain orchestration. • Including provisioning and service assurance.

Table 2. Technical Specifications

Operations Architecture Implications

The key operations architectural decision will be how a CSP will break down the network into the sub-domain, domain, and cross-domain functions and who will provide these functions:

- Build their own domain management and orchestrations system(s) or implement or buy SaaS functionality from an ISV or SI that would work with the multiple vendors,
- Depend upon an equipment manufacturer to provide multi-vendor domain management,
- Depend upon an ISV or SI to provide the systems for their equipment as a sub-domain and implement a cross-subdomain orchestration function (home grown, SI or ISV as on-prem, SaaS or managed service) for the access and middle-mile domains.

Market Implications

Traditionally, the equipment manufacturers have provided the bulk of the domain management functions via their element management (EMS) and network management (NMS) systems with limited multi-vendor functionality for all but the simplest functions. As the PON and middle-mile arrangements move to include more vendors in various complex arrangements, this will require greater adherence to standards (and deeper standards) by equipment manufacturers and require more inter-vendor cooperation. This cooperation has been hard to achieve over the last 50 years, leading to sub-optimal operations processes. In the future, ACG Research expects there to be:

- Greater CSP dependence on ISVs and SIs for the domain management function, as they represent neutral third parties that can work with the various vendors.
- Greater cooperation among the vendor winners in the PON and middle-mile arrangements as they are pressured by the CSPs to work together. This may lead to tactical or even strategic alliances between the larger major and the specialty vendors (especially NFV vendors) that manage to find a propitious niche.

Coherent Optical Pluggables Integrate Formerly Separate Optical and IP Domains

When digital transmission was first introduced into the circuit switched network in the early 1980s, all 24 switched analog trunk circuits were each converted using external A/D converters in a DS1 multiplexer to DS0 digital channels that were multiplexed into a single DS1 (in the US). As the switches became digital themselves, the extra A/D and D/A conversions were unnecessary, and the switched digital channels were sent directly into a multiplexer embedded in the switch. This created an operational problem; however, there were separate organizations responsible for the switching and transmission equipment monitoring and maintenance. The solution was to have the switch send the normal alarms and performance information to both the switching organization and the transmission organization, to a transmission service assurance OSS. Similarly, the individual organizations configured their respective equipment.

The situation is similar now that coherent pluggable optics capable of high rates (100, 400, and 800 Gbps) and long distances (in excess of 1,000 kilometers) are available in a number of vendors' routers. The operational problem of two different organizations—in this case, packet routing and optical transport— is being solved in a similar way, with the information being sent through the administrative channels of the router, then sent to the appropriate optical transport OSS. When SDN control is available, it goes through the router control channel but is configured by the optical transport organization using its own tools.

However, most equipment vendors have upgraded their DC systems to offer an integrated transport domain controller that handles both the router and the optical information. Cisco has its Router Optical Network and its Network Service Orchestrator. Nokia uses its Network Service Platform to control its Coherent Routing offering. Some, such as Juniper, have gone further and added the full range of optical transport management capabilities for optical multiplexing and switching to create a fully functional transport domain controller with the different needs of what are usually two separate operational organizations covered by role-based user interfaces and ubiquitous process automations.

Horizontally Disaggregated and Virtualized Equipment Architectures

As previously mentioned, certain network functions have been virtualized (usually those with large compute requirements such as firewalls, CDNs, security policy enforcement points, and BNGs). These can be placed on metro or edge compute platforms closer to the customer. Some, such as BNGs, can also be run as applications distributed across a geography. These network element architectures require the supporting domain controllers (or cross-domain orchestrators) handle the more complex arrangement in both network provisioning and assurance, including multivendor support. The standards for horizontal disaggregation are underway in several standards development organizations. ACG Research believes that such disaggregated standards will cause the commodity equipment market to lag the more proprietary offerings by 18 to 24 months and represent a minority of the market for the foreseeable future.[11]

Vertically Disaggregated CUPS Architectures

Disaggregation is happening vertically, with the control functions of the network elements being vertically disaggregated from the underlying hardware. This is called control-/user plane separation (CUPS), a term originally from the 5G mobile core architecture but used as a general term in the industry. The control functions are either virtualized and run on a standard computing platform or when real-time requirements are tight, run on a standardized computing environment specific to that application. The interface from the control function and the network element function is standardized, as may be the application environment in the control. Beyond 5G, CUPS is being used as part of the O-RAN specification. The purpose is to provide software-focused providers (or even CSPs to offer innovative software packages. The O-RAN CUPS concept is still untested in a commercial sense, but already well established in the market in other areas, such as the distributed BNG.

[11] As a point of reference, we note that over a decade after virtualized routers were introduced in the data center market, that it has gained only 40% market share. We believe that the same order of magnitude will be found in the horizontally disaggregated market, including distributed BNGs and O-RAN for the next five years.

Overall Network Provisioning and Service Assurance

The increasing distributed nature of the network functions (both vertically and horizontally) and the growing complexity of the local broadband network is bringing more complexity to network planning, installation, management, and assurance.

- The number of equipment options for the access and middle-mile architectures is growing, giving greater complexity to the equipment purchase decisions.
- Vendors' management solutions of the complex network functions are increasingly a part of the decision of which vendors' solution to buy and deploy. As seen before, the willingness to invest in network automation is growing.
- CSPs are increasingly looking to equipment providers for guidance on engineering their solutions, including the user plane traffic as well as the control and management plane connections.

Network Slicing Requirements and End-to-End Orchestration

Network slicing has been a well-discussed topic, with many vendors offering their versions of **slicing managers**. All of these slicing managers contain the full set of functions (provisioning, assurance, inventory and design, PAID functions). The slicing managers usually work with existing domain controllers in their respective domains.

Operations Summary

The operations requirements of the local broadband network are changing rapidly, with new equipment types, new software technologies and enhanced operations standards. The concept of a layered control architecture, as articulated in the TM Forum's Autonomous Network project, is well-established.

This paper is sponsored by Juniper Networks