Delivering Ultra-Fast Broadband

WHITE PAPER

Contents

Executive Summary

This white paper describes technologies, deployment methodologies, and management practices for delivering *ultra-fast broadband* with downstream transmission speeds of 100 Mbps or more. Its content is aimed at wireline service providers with significant copper infrastructure, who are investigating options for a transition to ultra-fast broadband.

The possible technology choices include:

- Vectored VDSL
- G.now™
- G.fast
- GPON

Vectored VDSL delivers downstream speeds up to around 100 Mbps, is most suitable for deployment from a node, and is the most economical in terms of required capital expenditure. Services with yet higher speeds can be delivered by combining Vectored VDSL with shared WiFi, which can deliver peak speeds in the range of 100s of Mbps by aggregating multiple WiFi streams and using Vectored VDSL as backhaul. Vectored VDSL management functions must include service qualification, coexistence with legacy lines, and automatic optimization for line stabilization against residual noise.

G.now delivers speeds of 100 Mbps to 1 Gbps, and is most suitable for fiber-to-the-street or fiber-to-the-basement deployments. It requires a higher investment than Vectored VDSL, but remains much more economical than GPON for brownfields. G.now requires management for crosstalk mitigation, and for line optimization against interference sources in the home environment.

G.fast will deliver speeds of 100 Mbps to 1 Gbps, and is expected to be widely available between 2016 and 2018. Otherwise, G.fast has many similarities to G.now in terms of deployment scenarios and costs. G.fast also will require management functions for power control, for mitigating interference sources in the home environment, and for coexistence with legacy VDSLs.

GPON has a downstream capacity of 2.5 Gbps, typically shared among 32 or 64 subscribers. The ideal deployment scenarios for GPON are new construction and high take-rate areas with aerial plant; for other scenarios, GPON requires the highest capital expenditure among the technologies discussed in this paper. Important GPON management functions include service installation verification, customer care recommendations, technician guidance, bandwidth allocation and capacity planning.

Service providers must consider a multi-technology approach for delivering ultra-fast broadband. The technology choice for each geographical area must be based on an assessment of capabilities and deficiencies of the existing network. Service providers also must deploy a unified management solution capable of delivering diagnostics, analytics, recommendations and optimization for any of the deployed access technologies. Finally, service providers must anticipate virtualization of the access network, either in the form of access node functionality migrating to cloud-based software, or in the form of multiple providers sharing a common infrastructure but controlling a virtualized sub-network containing their respective customers. For both of these scenarios, a unified management solution delivers essential functionality for network control and virtualization.

1. Introduction

The bandwidth race is on, and access service providers are facing the pressure of increased demands from broadband consumers. Factors such as the increasing popularity of streaming video (especially high-definition and ultra-high-definition), cloud-based services, and the proliferation of Internet devices at home (associated with the emergence of the Internet-of-things) are creating new requirements for both higher speeds and higher reliability of Internet access services. Additionally, several governments are adopting stricter definitions of basic broadband services [1][2].

This paper explains technologies, deployment methodologies, and management practices that network operators can use for delivering *ultra-fast broadband*. Ultra-fast broadband is defined in this paper as *achieving data transmission rates of 100 Mbps or more in the downstream direction*. Ultra-fast broadband encompasses super-fast broadband, and next-generation-access.

The technology choices that are described in this paper include:

1. Vectored Very high-speed Digital Subscriber Lines; also known as Vectored VDSL, G.vector, or G.993.5 [3]
2. G.now™ [4]
3. Fast Access to Subscriber Terminals; also known as G.fast [5]
4. Gigabit Passive Optical Networks (GPON); also known as G.984.x [6]

The above technologies are already commercially available, with the exception of G.fast, which is expected to be widely available by 2016. Technologies for delivering ultra-fast broadband over coaxial cable are not included, because this paper is aimed at network operators with significant copper-based infrastructure.

The paper starts in Section 2 by describing the technologies in terms of their technical characteristics, their deployment environment, and their management functions. Section 3 contains a comparison of the technologies, and Section 4 concludes with a set of recommendations for service providers.

2. Technologies

This section presents the most important aspects of Vectored VDSL, G.now, G.fast, and GPON. For each, there is a simple description of the technology, a presentation of the deployment architecture, and an explanation of necessary management functions.

One way to distinguish these technologies is from the fiber deployment needed to support them:

- Vectored VDSL brings fiber to a cross-box or node within about 1 km (3300 feet) of the subscribers.
- G.now and G.fast bring fiber all the way to a distribution point (dp) or drop-wire terminal within about 200 meters (700 feet) of the subscribers.
- GPON brings fiber all the way to the customer premises.

2.1. Vectored VDSL

2.1.1. Vectored VDSL Systems

Vectored VDSL was originally proposed in 2001 as "Dynamic Spectrum Management – Level 3" [7]. It was later standardized by the ITU-T as Recommendation G.993.5 [8]. Vectoring uses physical layer signal processing to enable cancellation of the crosstalk between all the lines that terminate on a single DSLAM. This has the effect of as much as doubling VDSL speeds on very short lines, with diminishing speed increases on longer lines where the crosstalk is weaker.

Figure 1 shows speeds achievable by vectored VDSL, using profile ADE17, with 20 dB of crosstalk reduction in a vectored group of 12 lines, on 0.4 mm cable, with 10% rate overhead. Speeds reach more than 100 Mbps in the downstream directions for loops shorter than 550 meters (1800 feet). The sum of upstream and

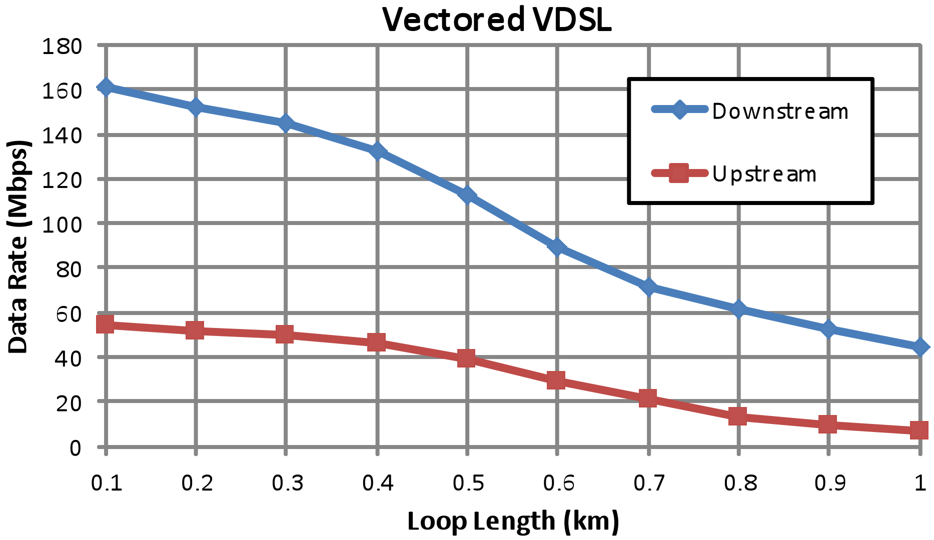

Figure 1: Typical achievable speeds of Vectored VDSL.

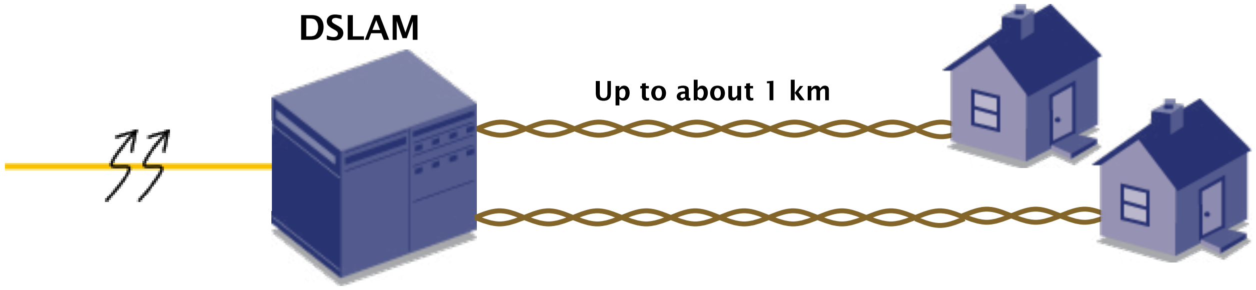

Figure 2: Vectored VDSL deployment; DSLAM located at a node in the Outside Plant (OSP).

downstream rates exceeds 200 Mbps for loops shorter than 200 meters (700 feet).

2.1.2. Vectored VDSL Deployment

As Figure 2 shows, Vectored VDSL typically brings fiber to within about 1 km (3300 feet) of subscribers. Bringing fiber within 500 meters (1600 feet) guarantees all subscribers downstream speeds of 100 Mbps. This may often be closer than legacy cabinet installations and can require plant re-arrangement; however Vectored VDSL needs the least amount of plant changes among the technologies considered in this paper.

Care needs to be exercised to ensure that all the lines emanating from the location, or node, of the Vectored VDSL DSLAM are either in the same *vector group*, or are managed to avoid interference (see Section 2.1.3). Node-level vectoring allows large vector group sizes (e.g. 96 to 384), so that every line from a given node can be vectored. Board-level vectoring uses smaller vector group sizes (e.g. 16 to 64), and is appropriate for small nodes or for FTTB deployment.

Vectored VDSL can deliver more than 100 Mbps to each home on a dedicated line. For most broadband consumers, this is ample bandwidth to support anticipated services for many years. But even higher speeds can be delivered with a deployment architecture that combines Vectored VDSL with WiFi sharing (see Figure 3). This architecture uses IP Layer Bonding for aggregating a collection of Vectored VDSL (or other types of) broadband links and for making the combined bandwidth available to consumer devices over WiFi. For example, three living units in a dense urban environment can share their Vectored VDSL links; any device in these units has access to a peak rate of 3 x 100 Mbps or 300 Mbps. IP layer bonding can be managed to group the multiple physical links into a single virtual link for the user, with individual services mapped into appropriate QoS levels as provided by the different links.

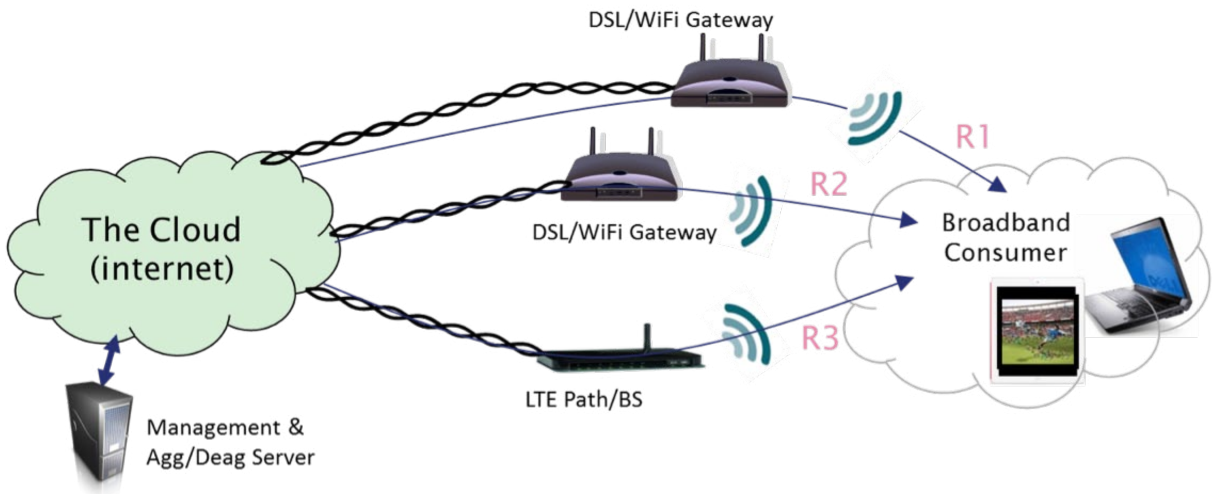

Figure 3: Vectored VDSL-WiFi deployment; IP Layer Bonding aggregates broadband connections.

2.1.3. Vectored VDSL Management

Service providers deploying Vectored VDSL must accurately qualify customers and networks for vectored services. By planning a gradual rollout to the most promising parts of the network, costs can be contained, while services can be delivered to those customers most likely to purchase. The ability to predict the services that can be offered after the upgrade is very important for maximizing the return on vectoring investment.

During rollout, there is the need to manage the coexistence of Vectored VDSL with legacy (non-vectored) VDSL. The problem is that existing non-vectored VDSL lines usually have non-vectored modems on the customer end of the line, and so these lines remain non-vectored even after the network-end is upgraded to a new vectored DSLAM. This may be a long-term problem, as some current users may refuse to upgrade their modems. Hardware-based solutions offered by DSLAM vendors appear to only work in the downstream direction, and to have poor performance. Dynamic Spectrum Management (DSM) is a comprehensive solution that manages the vectored and non-vectored lines to ensure the desirable tradeoffs, and allows super-fast speeds on the vectored lines while the non-vectored lines successfully continue to deliver their legacy service levels. DSM is a software-based technique requiring zero hardware upgrades. DSM enables service providers to deploy vectoring across networks with hardware equipment from multiple vendors and to upgrade the network gradually, managing vectored and non-vectored lines as needed.

When Vectored VDSL lines are operating and crosstalk is eliminated, a range of other noise sources can seriously affect performance: radio-frequency ingress, impulse noise, power-line interference, etc. Many of these sources are time varying; and so they can have even worse effects than slowing the line speed, which is to cause instability and frequent re-trains. A full retrain of a vectored group can take several minutes, which can be very damaging to low-delay IPTV services. Managing Vectored VDSL to ensure stability is crucial, because it can greatly improve customer satisfaction. Automatic optimization and interference management of all lines are essential to prevent dramatic OpEx costs, which would otherwise result from a higher volume of trouble calls and dispatches.

2.2. G.now™

2.2.1. G.now Systems

G.now is a MARVELL brand for a broadband access platform based on G.hn technology [4]. G.hn is a family of ITU-T Recommendations defining home networking over phone lines, power lines and coaxial cables [9][10][11]. G.now systems use standard G.hn chips and apply system-level software enhancements to deliver an access solution over phone lines of length of 200 meters (700 feet) or less. G.now can also operate over coaxial cabling.

G.hn and consequently G.now share many of the features of older DSL, but include many enhancements and adaptations. Just like ADSL and VDSL, G.now makes use of advanced Discrete Multi-Tone (DMT) modulation. However, G.now operates up to frequencies of 100 MHz, well beyond VDSL maximum frequencies of 17.7 or 30 MHz. G.now also uses Time Division Duplexing (TDD) for separation of downstream and upstream transmissions, which enables a programmable downstream to upstream asymmetry ratio. Transceivers operating in the same bundle use the same TDD ratio, and TDD frames are synchronized to a common clock to eliminate Near-End Crosstalk (NEXT).

G.now also has a number of advanced features such as automatic retransmission, and forward error correction based on LDPC codes with programmable FEC rates (1/2, 2/3, 5/6, 16/18, 20/21). Bit loading is dynamically recomputed based on real-time signal-to-noise ratio (SNR), channel frequency response and statistics from errors and LDPC decoder iterations. This means that bit-loading and the rates of each line can change within seconds without having to wait for a long initialization cycle to complete.

G.now performance results (downstream throughput rates) are shown in Figure 4 for a split ratio of 80:20. Results are shown for scenarios with one line, 8 lines and 16 lines. The effects of crosstalk for these specific cable types appear to be weak. Marvell reports that for a single line, the aggregate upstream and downstream PHY rate is 850 Mbps over 100 meters (300 feet) of phone-line; the corresponding throughput rate is 700 Mbps. The maximum aggregate PHY rate is 1 Gbps.

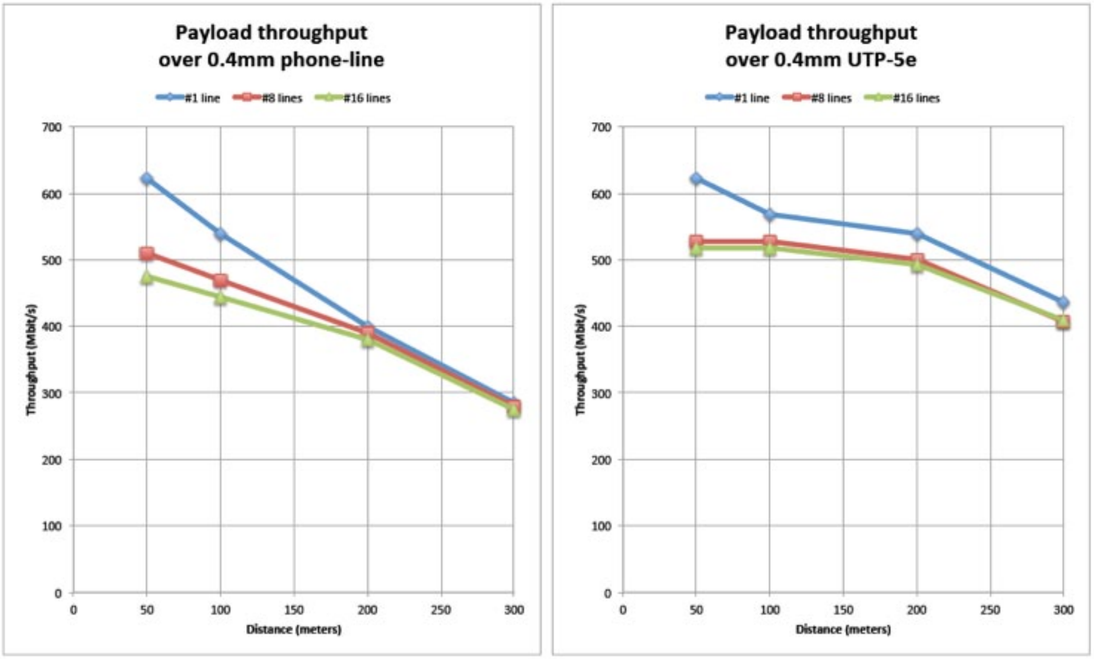

Figure 4: G.now downstream performance examples (from [12]).

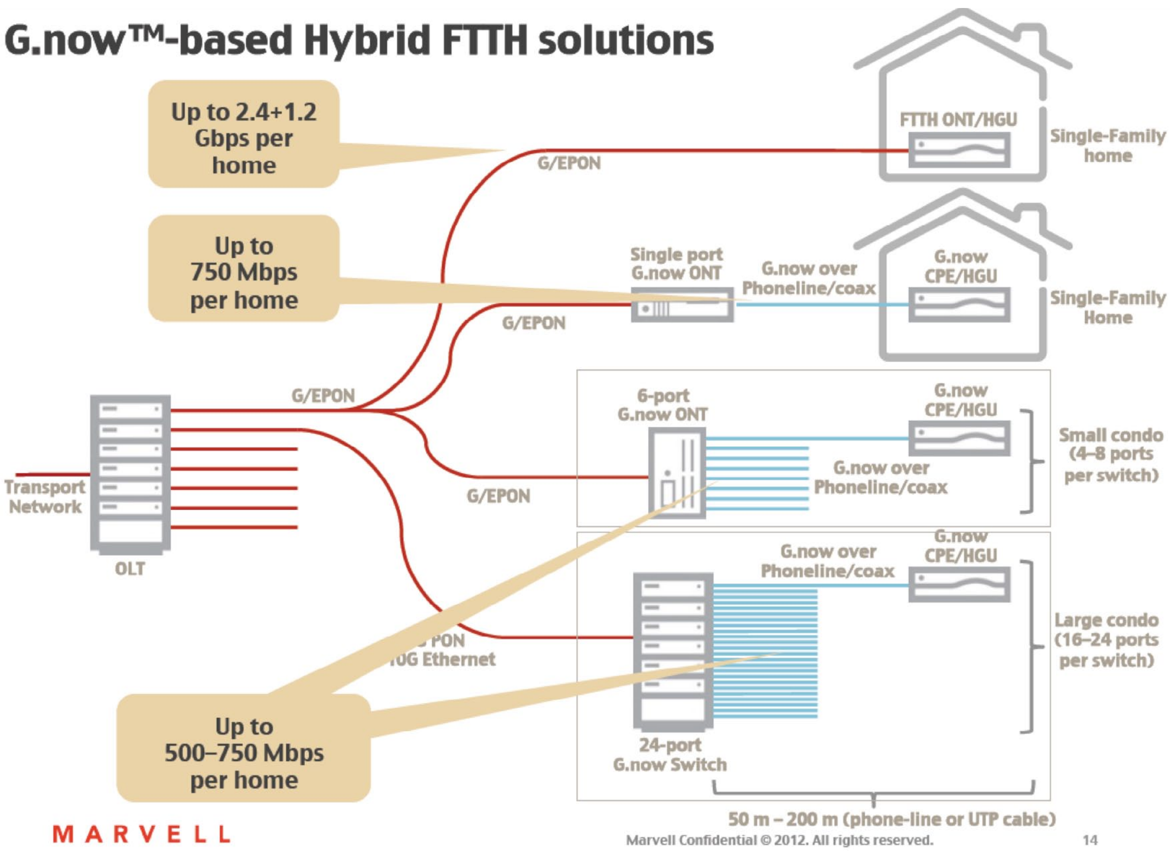

Figure 5: G.now deployment scenarios (from [12]).

2.2.2. G.now Deployment

Figure 5 shows possible deployment scenarios for G.now: These include a single-port G.now system suitable for a single-family home or apartment; and multi-port G.now systems for small or large multiple-dwelling units. The G.now systems are fed by GPON, EPON or active Ethernet; G.now can use either phone lines (up to 200 meters or 700 feet) or coaxial cabling.

2.2.3. G.now Management

Similarly to Vectored VDSL, service providers deploying G.now can benefit greatly from accurate service prediction. Observation and analysis of the existing ADSL or VDSL connections deliver estimates of the rates and services that can be supported after an upgrade to G.now is completed. These estimates provide essential input for planning for a transition to G.now in a way that maximizes the return on investment.

When G.now nodes are sharing copper cabling with VDSL nodes, the performance of both systems can be affected. When G.now and VDSL are transmitting over partially overlapping frequency bands, they will each cause crosstalk onto each other and degrade the achievable rates. If G.now and VDSL are configured to transmit over non-overlapping frequency bands, then this leads to a loss of data rate performance on one or both systems. The best practice is to apply DSM to dynamically select the spectrum configuration that achieves coexistence for the two systems while meeting the desired target data rates.

As mentioned earlier, G.now does not include a vectoring engine for crosstalk cancellation. In scenarios with strong crosstalk, the solution is to deploy a cloud-based crosstalk management system as shown in Figure 6. The system mitigates crosstalk by optimizing the assignment of power and time-slots based on the real-time bandwidth demands of each line. The system executes this optimization using knowledge of the crosstalk environment for the specific G.now switch.

2.3. G.fast

2.3.1. G.fast Systems

G.fast aims to provide ultra-high speeds over copper twisted pairs, up to and sometimes even exceeding speeds of 1 Gbps. The planned loop lengths for G.fast are from 50 to 250 meters (150 to 750 feet). G.fast is being standardized as ITU-T Recommendation G.9701 [5]. Similar to vectored VDSL, G.fast supports vectoring, which reduces crosstalk that is found in multi-pair cables and at higher frequencies. The first version of G.fast operates over frequencies of up to 106 MHz, and uses linear vector pre-coding to eliminate crosstalk in the downstream direction. A future version of G.fast may operate over frequencies of up to 212 MHz, and may support higher-performance non-linear pre-coding to allow for even higher speeds, as shown in Figure 7.

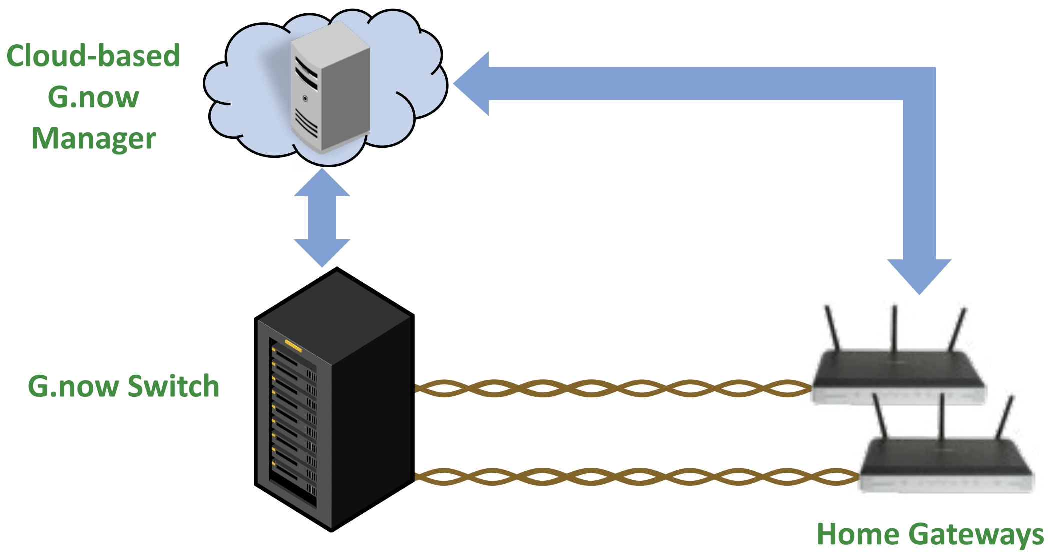

Figure 6: Cloud-managed G.now.

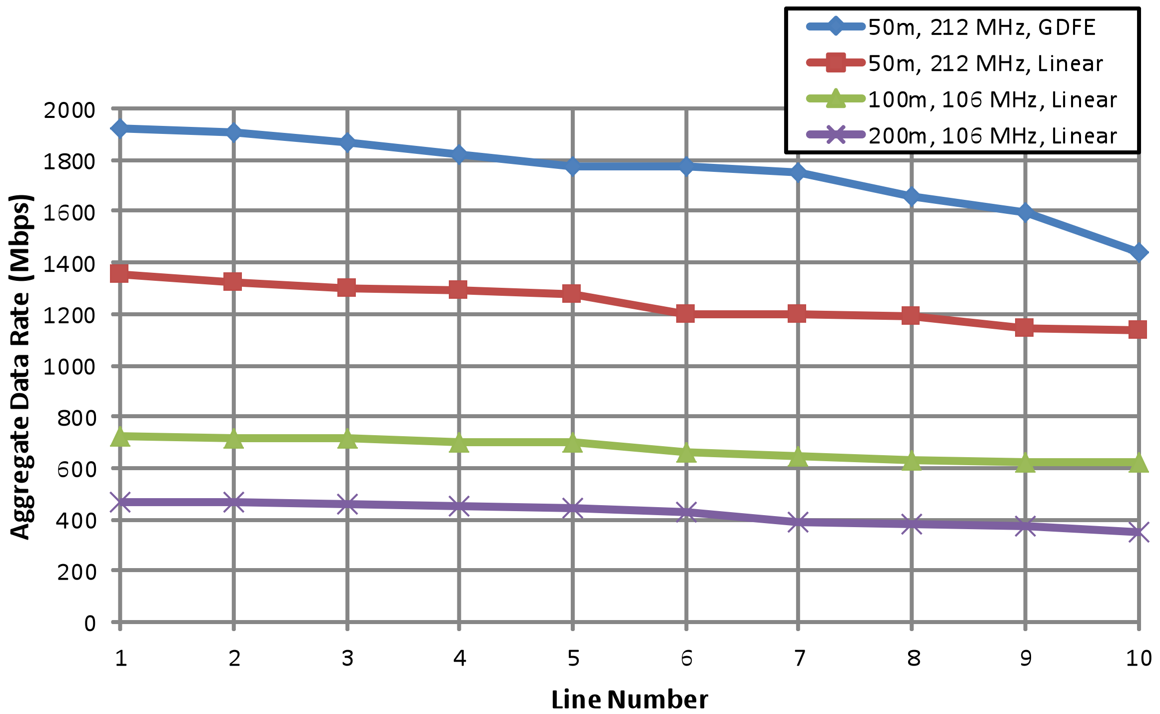

Figure 7: G.fast speeds, sum of upstream and downstream rates. "Linear" refers to the use of linear pre-coding techniques. "GDFE" is an advanced non-linear pre-coding technique delivering performance gains when higher frequencies are used.

Unlike prior DSL technologies which used Frequency-Division Duplexing (FDD), and similarly to G.now, G.fast uses Time-Division Duplexing (TDD). With TDD, the system transmits only downstream signals for a fraction of time, and transmits only upstream signals for the remaining time. TDD allows the speed asymmetry to be varied at will among all the lines emanating from the same Distribution Point Unit (DPU, which is the name for a G.fast DSLAM). This allows some areas to be served with business-class symmetric service, while other areas can be served with asymmetric service that best addresses consumer needs.

G.fast is amenable to low-cost self-install deployment, similar to ADSL. G.fast supports new On-Line Reconfiguration (OLR) techniques, including Fast bit-Rate Adaptation (FRA), to adapt to changes in the transmission environment and overcome the harsh home wiring environment. G.fast supports reverse-powering, which sends power from the customers' CPE to the DPU, thus eliminating the need for network power. Reverse powering needs to provide approximately 10 Watts, which is sufficient because the G.fast transceivers are very close to the subscribers and require little transmitter power. Reverse powering, and low-power modes planned for G.fast, are expected to lower G.fast deployment costs by eliminating the need for costly network powering and battery back-up of remote DPUs.

Figure 7 shows aggregate (upstream plus downstream) per-line data rates of G.fast lines for four scenarios. The data rates were calculated by simulation, with background noise of -140dBm/Hz. These conditions are characteristic of environments with no external interference sources, and are not representative of all situations that may be encountered in the field. There are 10 lines, and results are shown with both linear pre-coding and Generalized Decision Feedback Equalization (GDFE), which is a non-linear pre-coding technique. All systems in the figure have transmission frequencies starting at 23 MHz (above the VDSL frequencies) except for the 200m lines that start at 2.5 MHz (above the ADSL2+ frequencies). The transmission band extends to 106 MHz or 212 MHz, as noted in the figure legend. Note that the first generation of G.fast [5] will use linear pre-coding, and frequencies up to 106 MHz, thus it can reach 1 Gbps only over very-short loops.

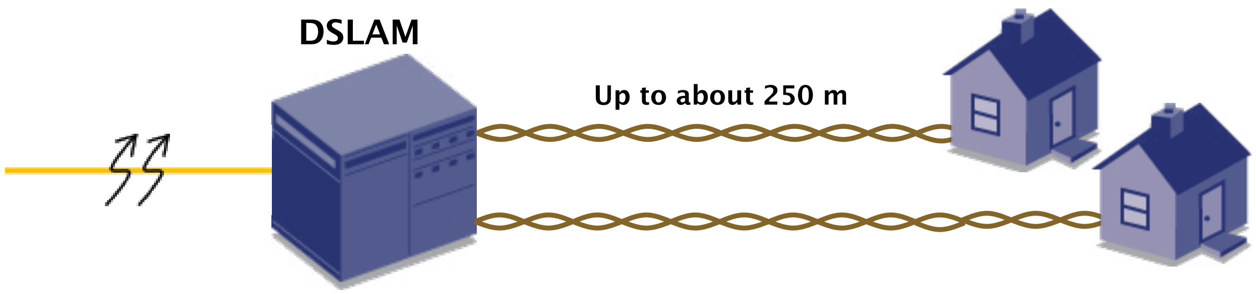

Figure 8: G.fast deployment; Distribution Point Unit (DPU) located at a Distribution Point (terminal) in the Outside Plant (OSP).

2.3.2. G.fast Deployment

G.fast is expected to be deployed in a Fiber-To-The-distribution point (FTTdp) architecture as shown in Figure 8. The "dp" may also be called the "terminal" or "drop-wire terminal," and is where the Distribution-Point-Unit (DPU) is located. Fiber is fed to the terminal, and from there, very short copper cables and drops, up to about 250 meters (750 feet) long, serve subscribers.

Some service providers will require G.fast DPU equipment to be backwards compatible with vectored VDSL. Such G.fast transceivers can fallback to support VDSL, albeit with lower performance compared to a dedicated VDSL transceiver. This is useful for initial installation of G.fast; an existing VDSL customer can be cut-over to the new G.fast DPU and have her VDSL-capable modem sync up. At a later time, a customer can receive a G.fast-capable modem, which will automatically sync up in G.fast mode.

In addition to higher speeds, FTTdp supports other features that appeal to operators, such as customer self-install, variable asymmetry, reverse-powering, fast On-Line Reconfiguration (OLR), and Remote Copper Re-configuration (RCR). Remote Copper Re-configuration (RCR) means that that after initial installation, it is possible to provision service for any customer connected to the DPU with no need for a technician dispatch. RCR is typically enabled by a switching matrix that connects subscriber lines either to G.fast transceiver ports or back to exchange or cabinet-based legacy services.

G.fast can coexist with ADSL and VDSL on adjacent pairs of copper wires by using frequencies above these technologies. However, operation above typical VDSL frequencies implies a minimum G.fast frequency of about 23 MHz, and this incurs a significant performance penalty on longer G.fast loops, beyond about 100m.

Both G.now and G.fast are candidates for Fiber-To-The-Building/Basement (FTTB) deployments, and perhaps for feeding small cells as they proliferate. For other types of deployments, G.now and G.fast entail a very high number of active electronic boxes in the Outside Plant (OSP), which raises concerns with regard to operational costs. G.now is available today, while G.fast is currently progressing through the standards. The first version of the ITU-T G.fast standard [5] should be completed in December 2014, with some interoperability testing occurring in 2015 under the auspices of the Broadband Forum.

2.3.3. G.fast Management

G.fast shares most of the management needs of G.now that were described in Section 2.2.3. The main exception is that G.fast equipment contains its own vectoring engine instead of using a cloud-based crosstalk manager.

G.fast management interfaces are currently undergoing standardization in the ITU-T [13] and Broadband Forum [14]. Features of G.fast such as Time Division Duplexing, Fast Rate Adaptation, Seamless Rate Adaptation, On-Line Reconfiguration, Vectoring and discontinuous operation will have their own controls. This provides additional room for automatic line optimization for purposes of protecting against impulsive and time-varying noise sources in the home environment.

The use of reverse powering means that the DPU may lose power if all the CPE connected to it are turned off and no longer provide reverse power feed. The solution is to employ a Persistent Management Agent (PMA) located either in a continuously-powered part of the network, or preferably virtualized in the cloud. The PMA will store diagnostics data from the DPU so these are available after the DPU is powered down. The PMA can also accept configuration changes and apply them after the DPU regains power.

Finally, management of Remote Copper Reconfiguration (RCR) can avoid truck-rolls to the network equipment for installation and activation of new or upgraded broadband service.

2.4. GPON

2.4.1. GPON Systems

Gigabit-capable Passive Optical Network (GPON) extends fiber all the way to the home or premises, and uses an entirely passive OSP (with the exception of Optical Network Terminals, ONTs, that are sometimes located outside of homes). While the point-to-point links of Vectored VDSL, G.now or G.fast are entirely separate until they are aggregated by Ethernet switching at a DSLAM or DPU, GPON shares the fiber medium among multiple subscribers. This sharing is performed with Time-Division Multiple Access (TDMA) under the scheduling control of the Optical Line Terminal (OLT). In this way, multiple (typically 32) subscriber lines are combined into a single fiber running into an exchange and terminating on an OLT. Thus, there are relatively few ports on the network-end active equipment.

GPON is standardized by the ITU-T G.984 series of Recommendations and typically supports aggregate line rates of 2,488 Mbps in the downstream direction and 1,244 Mbps in the upstream direction on two separate wavelengths. GPON uses the GigaPON Transmission Convergence sub-layer (GTC). GTC defines several different types of containers for scheduling Time-Division-Multiplexing (TDM) services with guaranteed QoS. Under this scheme, portions of each time frame are dedicated to services, and the OLT allocates downstream transmission opportunities. Also, the OLT performs Dynamic Bandwidth Allocation (DBA) to control the allocation of upstream transmission opportunities, which is assisted by a feedback mechanism from the ONTs to the OLT for reporting buffer-fill. DBA can be further assisted by inputs from policy management systems. The GTC protocol allows fragmentation and is efficient, losing only a few percent of line rate to overhead, unless there are an unusually high number of service streams.

While the speeds of GPON may seem ample, they are shared across many users and may at some times be exhausted, particularly for high-bandwidth services that are not amenable to concentration such as unicast video streaming during prime-time. GPON can be upgraded by splitting nodes, e.g., serving 32 users per OLT port instead of 64. Or, GPON can be upgraded to new, higher-speed systems with no change to the fiber components of the OSP. XG-PON supports 10 Gbps down and 2.5 Gbps up using TDMA similar to GPON, but at a faster line speed. XG-PON has been standardized, and equipment is now becoming available.

Beyond that, NG-PON2 runs each of multiple pairs of wavelengths as a single TDMA PON, each carrying up to 10 Gbps down and 2.5 Gbps up; this is called Time and Wavelength Division Multiplexed (TWDM) PON. NG-PON2 can use up to 8 wavelengths for each direction, increasing overall speed by up to a factor of 8 compared to XG-PON. In addition, NG-PON2 can optionally support point-to-point (PtP) Wavelength-Division-Multiplexing (WDM) PON for even higher per-user speeds. NG-PON2 offers a way to eliminate the need for point-to-point fiber (also known as active Ethernet), because a GigE or 10GigE link can be dedicated to each subscriber on separate WDM wavelengths. NG-PON2 standards are nearing completion in the ITU-T.

These next-generation PON systems should be able to work with existing PON OSP, and are compatible with GPON and each other on the same glass plant due to their wavelength assignments. An upgrade should require no change in the OSP, however, the OLTs and ONTs would have to be replaced and upgraded.

2.4.2. GPON Deployment

The GPON physical architecture is shown in Figure 9. ONTs terminate the GPON on the subscriber end of the line; these a re traditionally installed on or near the outside of a home for ready access by fiber installers. More recently, bendable fiber which can be run inside has made "customer self-install" possible for GPON.

The main cost of GPON is installation of fiber all the way to the home or apartment; the fiber-drop and ONT installation are particularly costly. Greenfield deployments are clear winners for GPON, however they represent roughly only 1% of subscribers per year. In brownfield deployments, aerial plant is easier to upgrade to fiber than buried plant. Generally speaking, only a subset of locations is cost-effective for GPON installation based purely on return-on-investment criteria. An alternative model for deployment is to rely

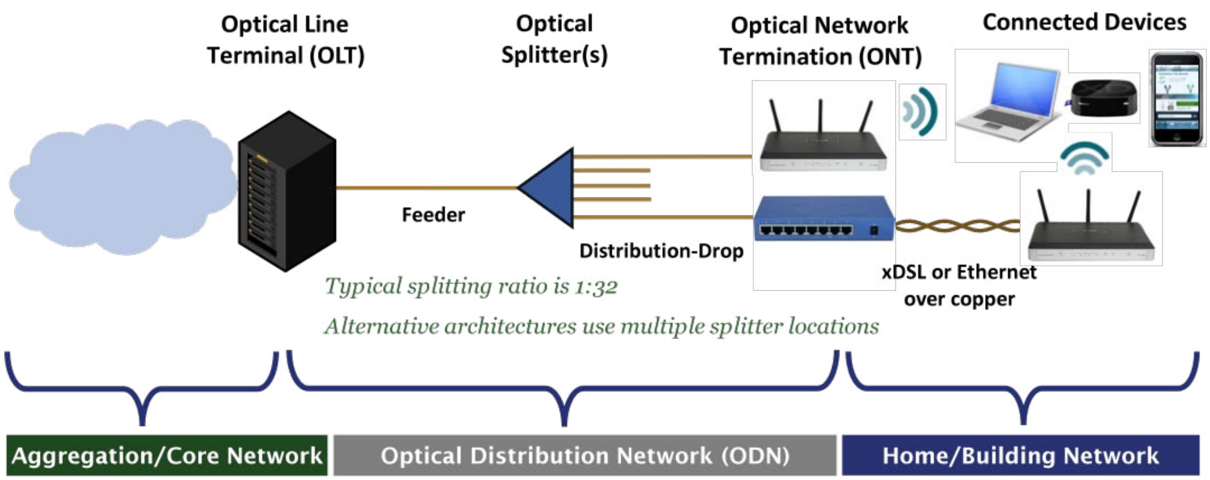

Figure 9: GPON physical architecture, with the OLT typically in an exchange.

on "pull" demand from subscribers, whereby GPON installation takes place only after a certain number of customers have committed to purchasing the service.

2.4.3. GPON Management

A key requirement for GPON management is the "turn-up test" verification of correct installation for a newly activated service. Providers with GPON deployment experience report that the majority of faults appearing in passive optical networks can be traced to installation flaws such as connector contamination. A management function is necessary for assuring installers that that new service is correctly installed, and that neighboring services have not been inadvertently disrupted.

Service providers require proactive monitoring of the GPON network for the purpose of having a complete view and being able to react to degradations or disruptions. Connectors, fiber cables, and splitter housings can degrade, for example if they get wet over time. Tight bends, dirty connectors, and in-home cabling can also cause a poor optical signal. A full characterization of the lines in terms of historical and current quality of service, and in terms of identified faults of the active equipment or of the passive elements, can drive analytics algorithms for determining actions by maintenance teams. Combined with long-term data, Optical Time-Domain Reflectometry (OTDR) outputs [15] can feed analyses to spot faults and see trends in degradation before they become service affecting. OSP faults that occur between the OLT and the splitters affect multiple subscribers, faults after the splitters generally only affect a single subscriber; fault correlation sectionalizes these types of faults and minimizes redundant dispatches.

Expert recommendations are necessary to assist customer care agents with identifying access troubles, isolating the root causes, and responding accordingly. Expert recommendations can also be directed to field technicians to identify fault types and locations, and to verify performance after fixes are applied.

Finally, a very important need for GPON operation relates to capacity management. Analysis of traffic and of congestion periods is used to identify and report emerging capacity issues in the shared optical distribution network. With proper capacity planning and with tracking of congestion trends, a balance can be struck between avoiding overbuilding and being ready to grow capacity as bandwidth usage grows.

3. Comparisons

This section provides a summary comparison of the technologies that were described in the previous sections. This comparison is made along the lines of speed, capital expenditure (CapEx), operational expenditure (OpEx), and availability.

Speed

Vectored VDSL offers downstream speeds up to about 100 Mbps, which can be further boosted with the use of bonding for networks with spare copper pairs, or with IP Layer Bonding. G.now and G.fast can deliver rates approaching 1 Gbps over sufficiently short loops. Each of these three technologies provides a dedicated link. GPON is sharing capacity of 2.5 Gbps in the downstream direction typically among 32 to 64 lines. At the extreme cases, GPON may deliver as high as 2.5 Gbps to a single line, or as low as 40 Mbps if 64 lines are sharing the available bandwidth.

Technology	Speed	CapEx	Availability
Vectored VDSL	50-100 Mbps[1]	$300-500 per line	Today
G.now	100 Mbps to 1 Gbps	$1100 per line	Today
G.fast	100 Mbps to 1 Gbps	$1400 per line	2016-2018
GPON	2.5 Gbps shared over 32 to 64 lines	$2500-5000 per line	Today

[1]Speeds exceeding 100 Mbps are possible by combining Vectored VDSL with shared WiFi.

Table 1: Comparison of Vectored VDSL, G.now, G.fast and GPON.

CapEx

Overall CapEx comparisons are difficult to make, because they are heavily dependent not only on the technology and on equipment costs, but also on the geography, on the existing infrastructure, on the city environment and on the labor costs. The cost estimates that are given next for vectored VDSL, G.fast and GPON are obtained from [16], and are representative of developed countries. Vectored VDSL is acknowledged as the most cost-effective option, because it requires the least amount of plant reconfiguration and active electronics. An average cost of $300-500 per line has been reported by service providers deploying Vectored VDSL. G.now and G.fast are estimated to have higher costs than vectored VDSL, because they require the addition of a large number of new nodes in the network (or in MDUs). G.now is expected to have a lower price for its hardware, because of the lack of an expensive vectoring engine, and because of volume synergies with G.hn production. For these reasons, the G.now cost is estimated to be $1100 per line, while G.fast is estimated to be $1400 per line. Finally, GPON has the highest costs, since it requires fiber installations to extend to the premises. An often cited cost concern about GPON is installing fiber in the drop segment, which requires consent and coordination with the homeowner or landlord.

OpEx

OpEx data for access networks are difficult to obtain. An attempt to collect such data was made in [17], but the resulting report does not produce a direct OpEx comparison. The general expectation is that GPON, which is essentially a new network build, and which contains few active electronics in the OSP, should have lower costs to operate than older networks that contain copper segments. However, management practices can very significantly affect OpEx of copper networks; for example, through software-based automatic optimization of underperforming lines, or through tools for guiding customer care agents and field technicians (these factors are acknowledged in [17]). For this reason, this paper makes no attempt to draw a conclusion on which technology offers OpEx advantages.

Availability

Vectored VDSL is available today, and is already being deployed by a number of service providers. G.now is also immediately available, and has a growing number of lines in the field. GPON is considered a mature product, and is being deployed by tens of service providers. The first version of the G.fast standard is expected to be complete by the end of 2014. Considering past trends of technology evolution from standards to deployment, and even assuming a higher speed of evolution, it is reasonable to assume that systems fully supporting G.fast will be available in early 2016. Lab and field trials, and interoperability, typically require two years of further work, so mass deployments of G.fast may be possible by 2018.

A summary of these comparisons is presented in Table 1.

4. Recommendations for Service Providers

This section provides a set of recommendations for service providers that are planning to deliver ultra-fast broadband services:

- **Choose the technology that is the best fit in both technical and economic terms for each geographical area.** An accurate assessment of the capabilities and deficiencies of the existing network is necessary to guide these technology choices. Multiple access technologies give service providers the flexibility to balance network investment with service demand. A provider's decision to only invest in GPON may delay for a very long time ultra-fast broadband delivery to a very large percentage of customers.

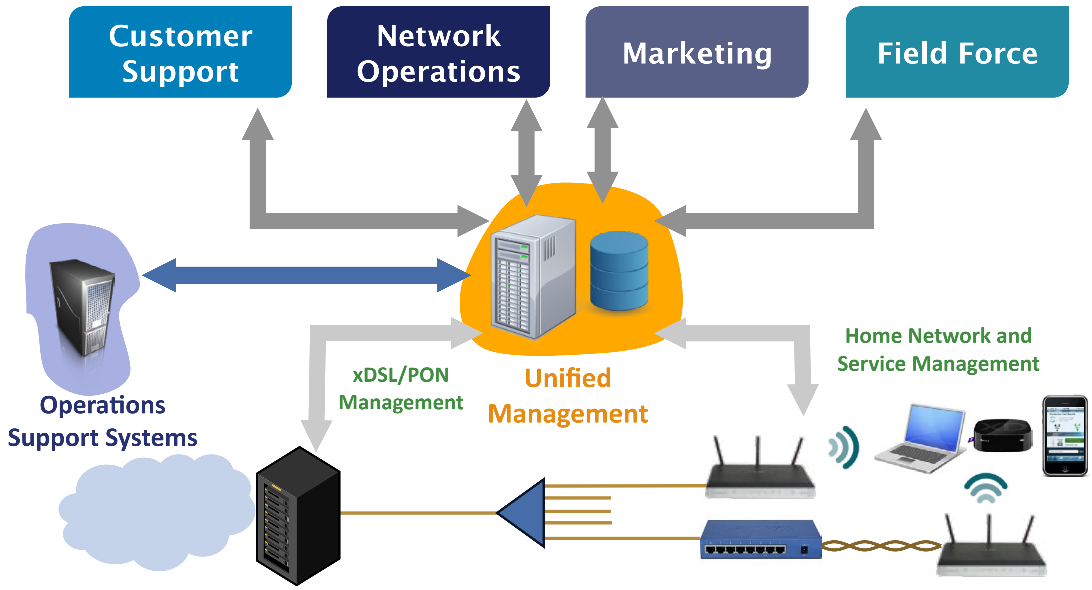

Figure 10: Architecture for deploying a unified management solution.

- **Deploy a unified solution for diagnostics, analytics, recommendations and optimization of all access technologies.** A software-based management solution must support all deployed technologies, and must provide the framework for both reactive and proactive strategies for service assurance. Automatic optimization should be used to minimize the need for labor-intensive interventions. Expert recommendations should be used by customer care professionals (at the lowest tier possible) to identify and respond to issues. Technicians should have tools for real-time guidance during installation or remediation. Network planning and marketing departments should be equipped with knowledge of capacity bottlenecks and of upsell opportunities. Figure 10 is a graphical depiction of how a unified management solution should be deployed.

- **Virtualize functions of the access network; move network intelligence from hardware to software systems.** Virtualization for access networks is favored by technology trends, such as the commoditization of access node hardware, and the desire for low-cost and non-continuously-powered units for G.now and G.fast. Access network virtualization is also supported by business trends, such as the desire to share physical network infrastructure among competing providers, with each provider having a virtual view of the respective sub-network. More broadly, virtualization fever is high, because operators stand to gain agility and lower costs with logically centralized touch points, and with the vast computing and storage resources available in the cloud. Service providers must assume that hardware-vendor-neutral unified management solutions (as in Figure 10) are the brains of ultra-fast broadband networks. Such solutions (for so-called Software-Defined Access Networking [18]) enable virtualization with logically centralized and agile network control. They also enable virtual network unbundling that lets network operators efficiently manage wholesale to retail relationships.

For more information, visit www.assia-inc.com or email us at info@assia-inc.com.

References

[1] "The FCC may consider a stricter definition of broadband in the Netflix age," *Washington Post*, May 30th, 2014; retrieved from: http://www.washingtonpost.com/blogs/the-switch/wp/2014/05/30/the-fcc-may-consider-a-stricter-definition-of-broadband-in-the-netflix-age/

[2] *Goals of the Digital Agenda for Europe*; retrieved from: http://ec.europa.eu/digital-agenda/en/about-our-goals

[3] ITU-T Recommendation G.993.5, Self-FEXT cancellation (vectoring) for use with VDSL2 transceivers

[4] "Marvell Unveils Game Changing G.now Technology Providing Fiber-to-the-Home (FTTH) Class Gigabit Broadband Services for the 'Smart Life and Smart Lifestyle'", Marvell press release, February 24, 2014; retrieved from: http://www.marvell.com/company/news/pressDetail.do?releaseID=5077

[5] ITU-T draft Recommendation G.9701, Fast Access to Subscriber Terminals (FAST)

[6] ITU-T Recommendation G.984.x, Gigabit-capable Passive Optical Networks

[7] J. Cioffi, "Proposed Outline and Text for DSM," ATIS Standards Contribution T1E1.4/2001-189, August 20, 2001

[8] An Overview of G.993.5 Vectoring – Issue 2, Marketing Report, Broadband Forum MR-257i2, 2014

[9] ITU-T Recommendation G.9960, Unified high-speed wireline-based home networking transceivers - System architecture and physical layer specification

[10] ITU-T Recommendation G.9961, Unified high-speed wire-line based home networking transceivers - Data link layer specification

[11] V. Oksman and S. Galli, "G.hn: The New ITU-T Home Networking Standard", IEEE Communications Magazine, October 2009

[12] "Providing Gigabit broadband access in MDUs using G.hn over existing copper wires in Korea," K.Y.Park – Ubiquoss, Chano Gomez – Marvell Semiconductors, TNO DSL Seminar, June, 2014

[13] ITU-T draft Recommendation G.997.2, Physical Layer Management for FAST Transceivers

[14] Broadband Forum draft WT-318, Management Architecture and Requirements for FTTdp

[15] Broadband Forum TR-287, PON Optical-Layer Management

[16] J. M. Cioffi and G. Ginis "Giga Feast or Fast," G.FAST Summit May 21, 2014

[17] "NGA Operational Expenditure: A Comparison of OpEx on FTTH, VDSL and ADSL Networks", Analysys Mason, August 22, 2013; retrieved from: http://www.analysysmason.com/Research/Content/Viewpoints/NGA-opex-Aug2013-RDTW0/

[18] K. Kerpez, J. Cioffi, G. Ginis, M. Goldburg, S. Galli, and P. Silverman, "Software-Defined Access Networks," *IEEE Communications Magazine*, vol. 52, no. 9, Sep. 2014; retrieved from: http://www.assia-inc.com/technology/knowledge-center/#white-papers

Acronyms and Abbreviations

ADSL	Asymmetric Digital Subscriber Line
CapEx	Capital Expenses
CPE	Customer Premises Equipment
DBA	Dynamic Bandwidth Allocation
DMT	Discrete Multi-Tone
DSM	Dynamic Spectrum Management
dp	Distribution Point (also known as terminal)
DPU	Distribution Point Unit
DSL	Digital Subscriber Line
DSLAM	DSL Access Multiplexer
EPON	Ethernet Passive Optical Network
FDD	Frequency Division Duplexing
FRA	Fast bit-Rate Adaptation
FTTB	Fiber-To-The-Building or Fiber-To-The-Basement
FTTdp	Fiber-To-The-distribution-point
FTTH	Fiber-To-The-Home
GDFE	Generalized Decision Feedback Equalizer
G.fast	ITU-T Recommendation G.9701, Fast Access to Subscriber Terminals
GPON	Gigabit-capable Passive Optical Network
IPTV	Internet-protocol Television
LDPC	Low-Density Parity Check
MDU	Multi-Dwelling Unit
NEXT	Near-End Crosstalk
NFV	Network Function Virtualization
OLR	On-Line Reconfiguration
OLT	Optical Line Terminal
ONT	Optical Network Terminal
OpEx	Operations Expenses
OSP	Outside Plant
OSS	Operations Support System
OTDR	Optical Time-Domain Reflectometry
PMA	Persistent Management Agent
PtP	Point-to-Point
QoS	Quality of Service
RCR	Remote Copper Re-configuration
TDD	Time Division Duplexing
TDM	Time-Division Multiplexing
TDMA	Time-Division Multiple Access
TWDM	Time Wavelength Division Multiplexing
VDSL	Very-high-speed Digital Subscriber Line
VULA	Virtual Unbundled Local Access
WDM	Wavelength Division Multiplexing
WiFi	802.11 family of IEEE standards

United States
333 Twin Dolphin Drive, Redwood City, CA 94065
Tel: 1-650-654-3400 • Fax: 1-650-654-3404

Europe
Calle Maria Tubau 3, Madrid 28050, Spain
Tel: +34 914842940 • Fax: +34 913446182

China
Suite 470, F/4, Beijing Sunflower Tower No.37, Maizidian Street
Chaoyang District, Beijing 100125, P. R.China
Tel: +86 10 85276788 • Fax: +86 10 85276488

Rev 101014