Fixed Access Network Sharing

Kenneth J. Kerpez, John M. Cioffi, Peter J. Silverman, Bruno Cornaglia, and Gavin Young

Abstract

Deep fiber deployments for ultrafast broadband are both technically and economically challenging. Fixed access network sharing (FANS) offers a highly enhanced form of virtual unbundling of broadband access networks, enabling cost sharing and the dynamics of a competitive landscape. Data, control, and management interfaces are automated and harmonized among wholesale infrastructure providers and retail virtual network operators. Standardized FANS interfaces can greatly decrease OPEX while increasing customer satisfaction. When FANS is implemented, network slicing logically partitions and isolates network resources among the VNOs, and such sharing can split the cost of network upgrades among several entities. FANS works with virtualization, where control plane functions are migrated from dedicated network equipment into software running on commodity hardware, with FANS providing network as a service. This article describes the drivers behind FANS, several architectures supporting FANS, and emerging standards supporting FANS.

Introduction

Fixed access network sharing (FANS) applies when an infrastructure provider (InP) controls a physical access network that supports virtual unbundling to virtual network operators (VNOs) [1]. FANS provides interfaces [2, 3] that allow diagnostics and status data to be disseminated from the InP to VNOs, and FANS allows a VNO to request or perform changes in network configuration and control their own virtual network. These FANS interfaces enable automated operations spanning the InP and VNO domains. An objective of FANS is to enable VNOs to perform operations with virtual unbundling similar to their operations with physical unbundling [4].

Not only is the physical access network provided by the InP and shared with FANS, but management data and configuration are also shared. With FANS, network management is a shared responsibility, with some management being performed by the InP and some by VNOs. There may be different levels of data sharing, distinguished by both the data shared between VNO and InP, and by the resolution and accuracy of the shared data.

Figure 1 shows how a single network and its equipment is administered by an InP, and a network sharing system logically divides the physical network resources between a number of VNOs, each of which has their own virtual network slice. The access node is generalized terminology for a digital subscriber line access multiplexer (DSLAM), optical line terminal (OLT), cable modem termination system (CMTS), and so on. The aggregation node aggregates traffic.

As fiber penetrates deeper, physical unbundling gives way to virtual unbundling, and FANS provides the means to continue a vibrant competitive ecosystem supporting economically viable, differentiated, and innovative services. FANS is synergistic with the software-defined access network (SDAN) concept [5, 6].

Competition

Current Multi-Operator Landscape

Competition today on copper infrastructure is often via facilities-based physical unbundling, with competitive providers deploying their own access nodes and leasing copper loops from InPs. Physical unbundling allows a retail operator to essentially run all their own operations and independently offer services. Physical unbundling is very popular in some areas (e.g., Europe) and has, arguably, led to very low-cost and very high-speed service to consumers.

There is also some "virtual unbundling," which has two flavors: bitstream and virtual unbundled local access (VULA). Bitstream is generally provided by giving each VNO access to their customers at the IP layer and is simple resale of the service provided by the InP. Bitstream usually is offered in a way that does not support differentiation of services between VNOs. VULA is generally provided at the Ethernet layer, and can enable layer 2 class of service differentiation and sometimes multicast. Generally, the management, backhaul, and other operations are opaque to the VNO with current virtual unbundling, although there are limited exceptions [7].

Fiber-Deep Ultrafast Broadband

Vectored VDSL2 and G.fast coordinate the signals across all the copper pairs emanating from an access node. Therefore, physical unbundling is technically unattractive with fiber to the node (FTTN) using vectored very high rate DSL 2 (VDSL2), and with fiber to the distribution point (FTTdp) using G.fast because multi-operator vectoring is not yet standardized. Physical unbundling with fiber to the premises (FTTP)/fiber to the home (FTTH) using passive optical network (PON) technologies could be possible using an overlay network or separate wavelengths. However, for all these new ultrafast broadband deployments, physical unbundling is generally

Digital Object Identifier:
10.1109/MCOMSTD.2017.1600001ST

Kenneth J. Kerpez, John M.. Cioffi, and Peter J. Silverman are with ASSIA Inc.; Bruno Cornaglia and Gavin Young are with Vodafone

ISSN 2471-2825/17/$25.00 © 2017 IEEE

economically unattractive. This is because physical unbundling would require many operators to each run fiber in the outside plant to each of many small nodes, install all these small nodes, and deploy equipment at each of these small nodes. Much of this fiber and equipment would be redundant with other operators' deployments and be underutilized.

The existence of multi-operator environments will depend on virtual unbundling. Basic resale is insufficient; an ecosystem supporting innovative service offerings from multiple operators depends on a platform that allows retailers to perform operations in a way that is nearly indistinguishable from physical unbundling. For example, FANS can enable multicast or real-time service variations such as turbo-boost, bandwidth reservation for video streams, and real-time charging. Moving forward, FANS is especially desirable in its ability to encourage such competitive innovation and differentiation, and to drive economic growth of broadband services at lower cost and lower operational complexity.

FANS BENEFITS

Deploying superfast broadband networks is quite costly in terms of both upfront investment and resources needed for design and implementation. Multiple parallel networks covering the same areas can lead to low take rate and long return on investment. Operators should therefore share their own infrastructure in order to reduce upfront investments and operational resources. The International Telecommunication Union (ITU) Commission for Sustainable Growth and the European Commission have recognized that sharing infrastructure can speed broadband rollouts. In some countries, including Portugal, Spain, and the United Kingdom, regulatory bodies have taken into account this consideration and have created rules for sharing existing passive infrastructures. Operators also create their own agreements for sharing the access network.

Two types of sharing are possible, as shown in Fig. 2: passive and active sharing. With passive sharing, only the passive infrastructure is shared, including ducts, poles, and cabling; and each operator installs its own access node equipment with a fiber tree connecting it to the optical distribution frame (ODF) in the central office (CO). For PON, passive sharing leads to a cost increase above a single infrastructure because multiple parallel connection trees have to be built. In the case of active sharing, both passive infrastructure and access node equipment are shared; the drawback is that the VNO connects via bitstream. Active sharing with bitstream has limited customer management and service differentiation capabilities, which sometimes make active sharing unattractive. Passive sharing can give around 30 percent savings, while active sharing can lead up to 40 percent savings, compared to deploying multiple parallel networks.

To overcome this situation, active sharing with FANS uses VULA and virtualization concepts to allow slicing the network elements so as to assign dedicated resources to each operator. In this case the cost of the passive infrastructure remains the same, but new slicing capabilities have to be developed.

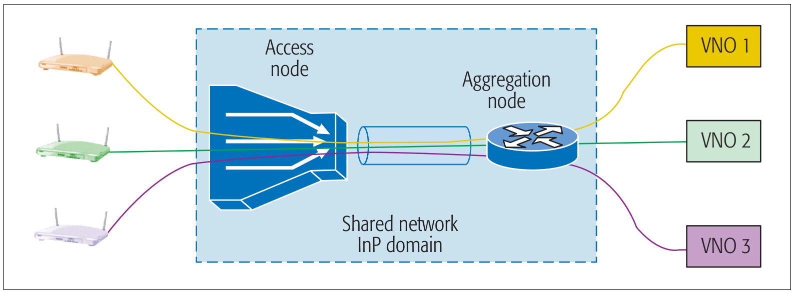

FIGURE 1. Network sharing data plane, showing network slices in different colors.

Using an active sharing approach is more convenient than the passive one, mainly because of the extra cost in the passive approach for each operator to install and connect its own equipment. On the other hand, active sharing without using FANS limits the capability of the operator to differentiate its offer from others, so active sharing can be less attractive. FANS extends active sharing to encompass management and control functions. With FANS the solution implemented is a "virtual unbundling," so each operator has similar capability to implement its own access network, which is why FANS is an appealing option.

USE CASES

With current bitstream or VULA, operations "interfaces" between InP and VNO are often manual. In conjunction with appropriate business arrangements, FANS can assist in automating operations interactions between InPs and VNOs; including fault, configuration, performance monitoring, and optimization. In particular, standardized FANS interfaces allow efficient data exchange between all parties. This can lower operation expenditure (OPEX) costs, improve customer satisfaction by enabling rapid response times, and increase the number of customers.

FANS can enhance competitiveness with competing broadband media. Multiple companies can share the costs of infrastructure and its upgrades. The InP can offer enhanced product offerings to the VNOs, with enhanced value for providing access to automated data and control interfaces. FANS use cases include:

Performance monitoring and optimization is enabled. Each VNO can have automated real-time access to performance and fault monitoring data. This is useful for VNOs' network monitoring as well as repair and troubleshooting operations. Customer relationship management (CRM) can improve. VNOs can optimize the configurations of their lines. Further, multi-line optimizations can be performed across multiple operators' lines by a centralized sharing system, which can increase performance of all lines [4].

Fault correlation can be performed using shared data to correlate multiple faults across multiple lines and multiple service providers; and this can further be used to help coordinate dispatches. For example, pooled data across multiple VNOs and/or InPs can be used to identify a fault that occurs in a single shared cable section. A single dispatch to fix that cable section is much better than dispatching to each troubled line separately.

Configuration can be automated, so a VNO can rapidly turn up service. VNOs can offer ser-

The centralized sharing system may be administered by the InP, a VNO, or a third party. The centralized management system itself could be provided by multiple parties either via shared or open source software, or via APIs between proprietary systems. With shared or open source software, the cost of the software is split across multiple parties.

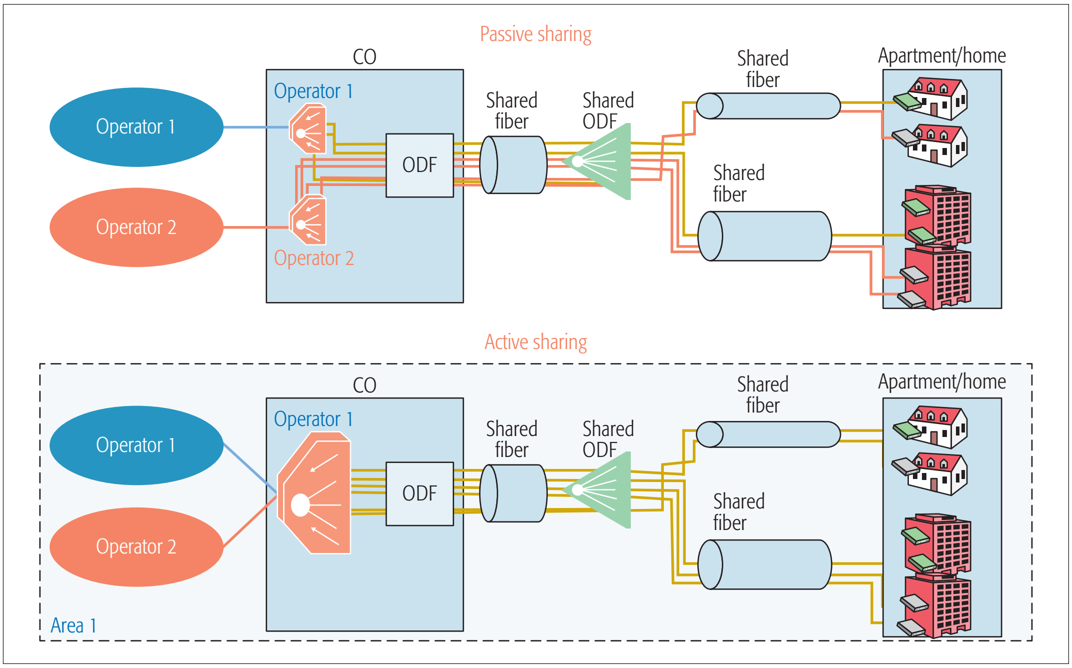

FIGURE 2. Passive sharing and active sharing approaches, and their relative 10-year TCO per operator.

vices with different quality of service (QoS) levels, for example, to maximize speed or stability, or to minimize delay or power usage. Data sharing can help ensure that the lines can support the necessary QoS attributes.

Services innovation can flourish; for example, VNOs can provide different types of assured services, business class services, sponsored services, and so on. This can build on basic broadband offerings to grow the pie for all entities.

Network planning can be enhanced with network topology and capacity data. Lines in a small geographic area should all have about the same performance, so neighborhood data can be anonymized by a sharing system to provide line performance projections to VNOs.

A **single operator** can exploit FANS for their company and segment their network among different operation teams for business customers, consumer broadband, and/or mobile backhaul connectivity. Moreover, FANS can also facilitate mergers and acquisitions because operations can initially be run separately but with an integrated network.

FANS COMPONENTS

ACTORS

FANS defines interfaces between an InP (aka wholesaler) and multiple VNOs (aka retailers). The InP operates the physical network, while the VNOs interface to broadband consumers.

The InP is responsible for deploying and managing the physical network; the InP:

- Enables physical resource slicing and carries out the slicing
- Provides an interface to the VNO for data and control
- Gets revenue from resource leasing

The VNO leases resources from the InP, and the main VNO functions are:

- Operates, controls, and manages its own virtual area networks
- Runs and redesigns customized systems in its own virtual network, such as diagnostics and optimization and DSM/DLM systems [10]
- Provides specific and customized service through its own area network
- Utilizes the network resources provided by the InP to provide its services
- Obtains revenue by selling services to end users

CENTRALIZED SHARING SYSTEM

FANS is generally instantiated around a centralized sharing system, as shown in Fig. 3. The centralized sharing system has southbound interfaces to equipment and northbound interfaces to VNO systems. The centralized system performs certain functions, such as authentication, authorization, and accounting (AAA), and arbitration of requests for resources, data, and control. The centralized sharing system can perform operations such as diagnostics, configuration, and optimization for the VNO, as shown in Fig. 3 for VNO A; or the external interface to the centralized sharing system can enable these functions to be performed by a VNO itself, which is the case for VNO B in Fig. 3. The centralized sharing system may be administered by the InP, a VNO, or a third party. The centralized management system itself could be provided by multiple parties via either shared or open source software, or application programming interfaces (APIs) between proprietary systems. With shared or open source software, the cost of the software is split across multiple parties.

The centralized sharing system implements multi-tenancy, although the functionality may be distributed among multiple systems or locations. Aspects of FANS could be implemented without a centralized sharing system, for example, with a distributed architecture simply having interfaces between the various players. In any event, the centralized sharing system conceptually glues FANS together.

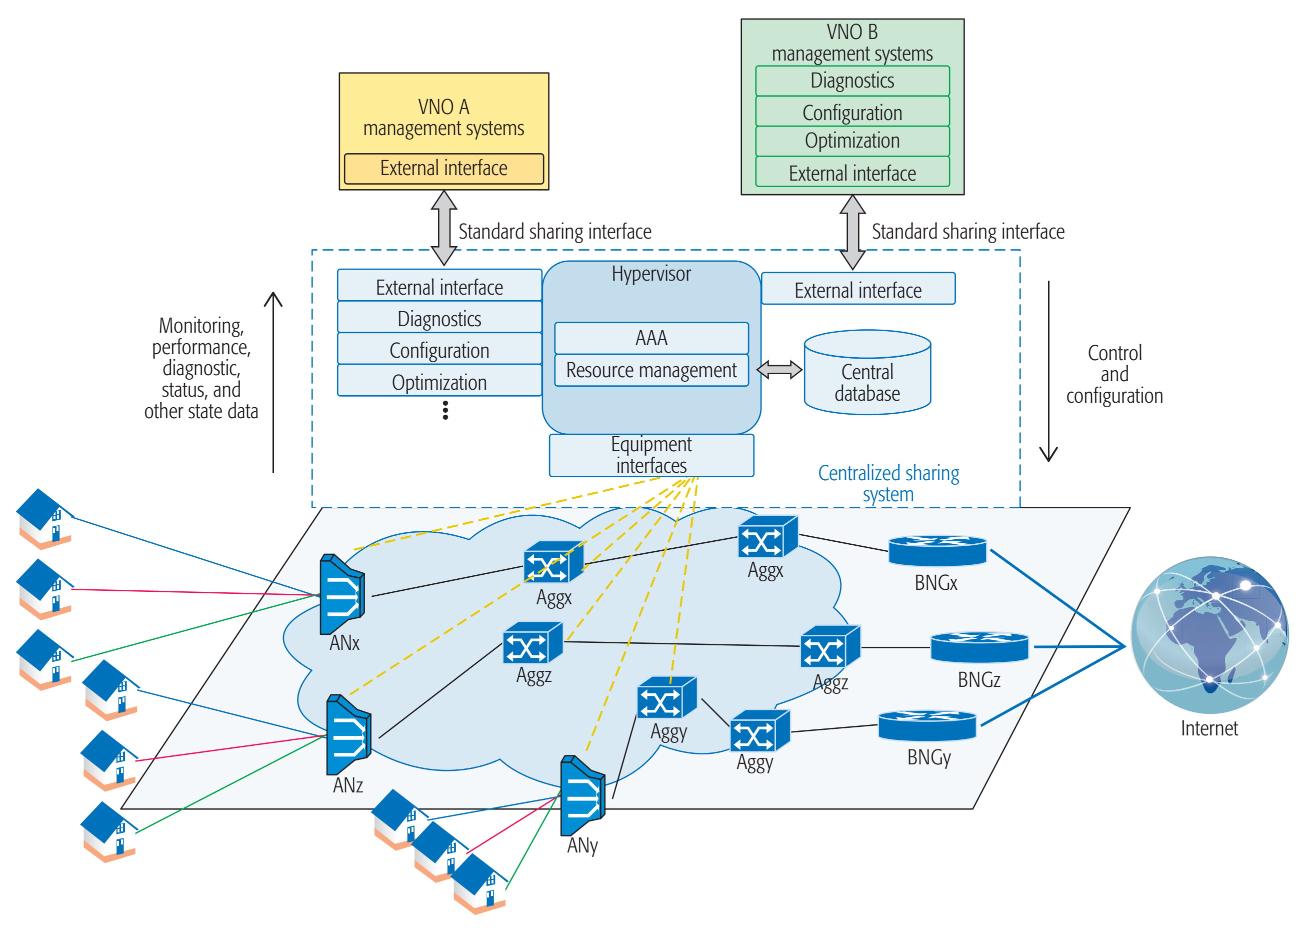

FIGURE 3. Centralized sharing system.

The components in Fig. 3 are part of the infrastructure that supports FANS:

- The "hypervisor" manages authentication and authorization to use the centralized sharing system, and manages resources. The hypervisor here oversees the virtualization or slicing of the network, not the compute infrastructure.
- AAA: Authentication, authorization, and accounting verifies user credentials, admits requests and limits access, and maintains transactional records for billing and other purposes.
- Resources management includes assignment of network bandwidth, equipment interfaces, equipment computational resources, and management interface bandwidth/frequency of admissible requests.
- Southbound equipment management interfaces typically utilize the existing management capabilities of network elements or their management systems.
- The northbound sharing interface should be standardized. This interface provides monitoring, performance, diagnostics, status, and other state information from the network to the VNOs, and also relays or interprets control and configuration requests from VNOs to the network.

There are additional interfaces between InPs and VNOs that may or may not involve the centralized sharing system:

- The business support system (BSS) interface supports ordering and billing functions between InP and VNO. The BSS includes a catalog of offers and resources, an ordering interface, inventory (of physical and virtual components, equipment, outside plant, admissible configuration settings, etc.), and revenue. The resource catalog should include physical resources, virtual resources, and services.
- Test and diagnostics interface. This invokes specific test actions to provide diagnostics data beyond what the equipment can provide during normal operation. It provides outside plant and network test data, and may provide analyses of these data. Test may be abstracted to hide details of the network equipment capabilities and to provide results in terms of the virtual services provided. The interface may issue commands for requesting tests.
- Logical inventory. This includes equipment, interfaces, virtual ports and their assignments, equipment configuration settings, and virtual functions inventory.

Messaging across the southbound interface from the centralized sharing system may be implemented by an abstraction layer or an adaptation layer. An abstraction layer hides the details of equipment interfaces to present a simplified interface toward management systems. An adaptation layer directly translates signals from one format to another format, and usually has a different adapter for each type of interface.

RESOURCES MANAGEMENT AND SECURITY

Resources must be assigned carefully, to control access permissions, arbitrate conflicts, ensure correct or fair resource utilization, and guarantee reliability for the underlying physical infrastructure. Resources need to be assigned, with data access and control separated, for the access network, equipment, and computing infrastructure. A VNO cannot be allowed to access private data about another VNO's customers. Sharing of resources must be managed to ensure that resources are properly allocated among the competing VNOs, and that any particular VNO cannot either impair another VNO's service.

Computing resources, including CPU, memory, and virtual network, can be shared between any or all of the actors. Managing computing infrastructure resources is particularly important for virtual functions.

BACKHAUL

The backhaul network extends from the access node to whatever point the traffic is handed off to a VNO or service. Traffic on the backhaul broadband network needs to be segregated between different operators and services. Backhaul segregation can use VLANs, MPLS tunnels, or software defined networking (SDN) approaches. An interesting new VLAN approach is to use a new "operator VLAN" (O-VLAN) tag, which may be a third VLAN tag in addition to C-VLAN and S-VLAN tags of IEEE 802.1ad Q-in-Q. This allows the VNO to manage two levels of VLANs (S+C VLAN) for its service configurations, while the InP only assigns the O-Tag for each operator.

A VNO may wish to use its own backhaul network. In this case, data traffic may be handed off from InP to VNO at various reference points, including the V-interface to the regional broadband network or the A10 interface to service provider network(s) as defined in TR-101 [8]. The handoff could be at a broadband network gateway (BNG), in the aggregation network, or even in the outside plant at a cabinet location.

FANS ARCHITECTURES

Two types of virtualization are emerging [1]:

- Equipment slicing, where network resources are virtualized with slicing to support multi-tenancy. Separate tenants perform separate functions on logically separate parts of the equipment. This is most aligned with virtual node sharing.
- Full virtualization, where network functions migrate from equipment to being hosted on cloud infrastructure platforms. This is most aligned with management-system-based sharing.

MANAGEMENT-SYSTEM-BASED SHARING

The centralized sharing system in Fig. 3 is the core of management-system-based sharing. With this sharing technique, a management system performs the network slicing at the management system level and not directly in the equipment itself. The management could be virtualized and hosted in the cloud or at other operator locations. The management system supports multi-tenancy, where each VNO is a separate tenant.

Management system sharing separates the management plane from the data plane, with sharing and network slicing performed by the management systems. The data plane can remain unchanged, and data-plane functions such as packet forwarding continue to be performed in the network elements. Aspects of the control plane may also support sharing and network slicing functions.

Management system sharing separates the management plane from the data plane, with sharing and network slicing performed by the management systems. The data plane can remain unchanged, and data-plane functions such as packet forwarding continue to be performed in the network elements. Aspects of the control plane may also support sharing and network slicing functions.

The centralized management system could provide functionalities that include the following:

- Security, which includes authentication to verify user credentials, authorization to admit requests and limit access, and accounting to maintain transactional records for billing and other purposes
- Fault correlation, particularly for faults that occur on lines or equipment that impact multiple VNOs
- Inventory maintenance of the physical plant and equipment, as well as the virtual assignment of resources
- Data maintenance needed to access VNOs and equipment such as addressing
- Support of an automated data clearinghouse that allows automated operations
- Providing data to assist VNOs with network planning and to assist in development of innovative services and differentiated services
- Multi-line optimization across multiple VNOs

Management system sharing allows a VNO to choose to perform the following operations, among others:

- Services provisioning
- Fault and performance management
- Configuration of the network elements
- Testing and gathering of diagnostic data
- Line optimization
- Call center operations to answer trouble calls

Unlike virtual node sharing, management-system-based sharing can be implemented with currently deployed equipment.

VIRTUAL NODE SHARING

Virtual node sharing is based on the concept of equipment slicing. Virtual node sharing is performed within equipment, including access nodes, aggregation nodes, and virtual port mappers. This may require the equipment to host a common execution environment for sharing, such as a segmented space running a version of Linux. A hypervisor controls the life cycle and resources of virtual machines (VMs). Equipment slicing allows interchangeable functions to be hosted similar to the way a data center can host virtual network functions (VNFs). This concept is part of the virtual OLT (vOLT) in the Central Office Re-architected as a Datacenter (CORD) initiative. FANS can be considered a use case of CORD.

The virtual access node model performs equipment slicing on physical access nodes to abstract them into multiple virtual access nodes, where each VNO accesses a logically separate virtual access node. Separate access node functions can be sliced independently, and as shown in Fig. 4 some functions may also be fully virtualized and hosted on cloud virtualization infrastructure.

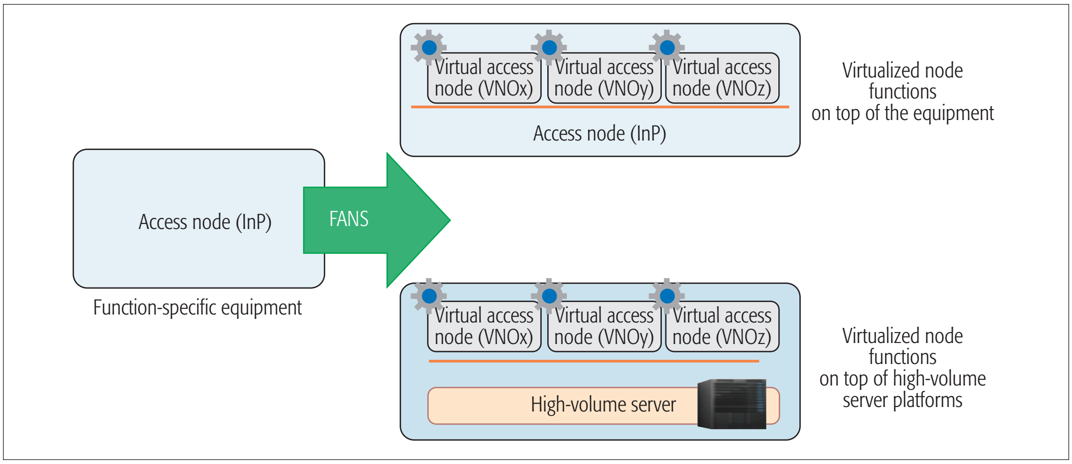

FIGURE 4. Deployment scenarios for virtual access node functions.

SDN controllers for the access network can also be virtualized. There may be a hierarchy of SDN controllers, with an end-to-end SDN controller on top of domain-specific SDN controllers. Network SDN control can be divided between the InP, having an infrastructure controller, and the VNOs, having virtual controllers.

Some functions may also be fully virtualized, as shown in the bottom half of Fig. 4. VNFs implement some of the functions that traditionally reside in the access node or the BNG.

A virtual access node element represents the whole set of characteristics of a physical access node. A centralized sharing system is still present and is involved as part of virtual node sharing. The centralized sharing system described earlier performs orchestration, and monitors and scales virtualized and physical network resources.

Other network nodes can similarly use virtual node sharing, including virtual aggregation nodes (e.g. Ethernet aggregation switch, multiprotocol label switching [MPLS] router, SDN switch) [1].

The port mapper concept is part of virtual node sharing. A port mapper maps a disparate set of physical ports into a logical set of ports assigned to each VNO. The port mapper is a virtual entity used to map logical ports over the host physical ports. The virtual ports are identified through virtual port IDs. For example, as shown in Fig. 5, an access node may assign each user-facing port to a separate VNO, and then the VNO references the port through its virtual port ID. The port mapper may be combined with a virtual switch, which can intelligently forward data by inspecting packets before passing them on, ensuring traffic isolation. The port mapper moreover facilitates the customer migration as the end customer maintains the same physical ID and only changes the virtual port ID when moving from one operator to another.

FULL VIRTUALIZATION

Full virtualization moves functions into cloud platforms/data centers. Virtualization generally follows the architecture of the European Telecommunications Standards Institute (ETSI) Industry Specification Group (ISG) for Network Functions Virtualization (NFV). Here, the NFV infrastructure (NFVI) runs a system such as OpenStack or CloudStack, and the VNFs run on this infrastructure. Management and orchestration (MANO) are controlled by a system such as Open Platform NFV (OPNFV), Open Source MANO (OSM), Open Orchestrator (Open-O), or a proprietary system.

FANS with full virtualization can extend management-system-based sharing to include control functions and other VNFs and network services (NS); this can be thought of as fixed access network as a service (NaaS). Here, the NFVI and MANO support multi-tenancy, where each VNO is a tenant, and they are logically separated. For a given function, each VNO would have separate VNFs, NSs, and VMs, thereby using the underlying NFV components to allocate resources, ensure privacy, perform life cycle management, and so on.

SDN controllers for the access network can also be virtualized. There may be a hierarchy of SDN controllers, with an end-to-end SDN controller on top of domain-specific SDN controllers. Network SDN control can be divided between the InP, having an infrastructure controller, and the VNOs,a having virtual controllers.

The following are among the fixed access network functions that may be virtualized:

- VLAN translation/addition/removal: The access node would focus on basic connectivity, whereas additional VLAN tagging could be performed in the NFVI.
- Virtual inventory management performs assignments of physical assets.
- Per subscriber QoS enforcement (e.g., policing or shaping) enforces QoS policy, and allocates QoS and class of service (CoS) levels.
- Port-based access control/authentication is performed, for example, by using a centralized 802.1x agent
- Traffic is managed, filteried, and shaped, and flow is controlled.
- Forwarding, traffic steering, load balancing, and SDN control are virtualized.
- Application awareness, deep packet inspection (DPI), and services-aware networking are virtualized
- There can be virtual gateway functions, residential and business.
- Control and configuration: Each VNO controls and configures their own virtual access node dataset of configuration objects.
- Diagnostics and state information: Each VNO accesses virtual functions providing test, diagnostic, performance, and status information.
- Dynamic rate allocation (DRA): This function controls traffic scheduling, such as dynami-

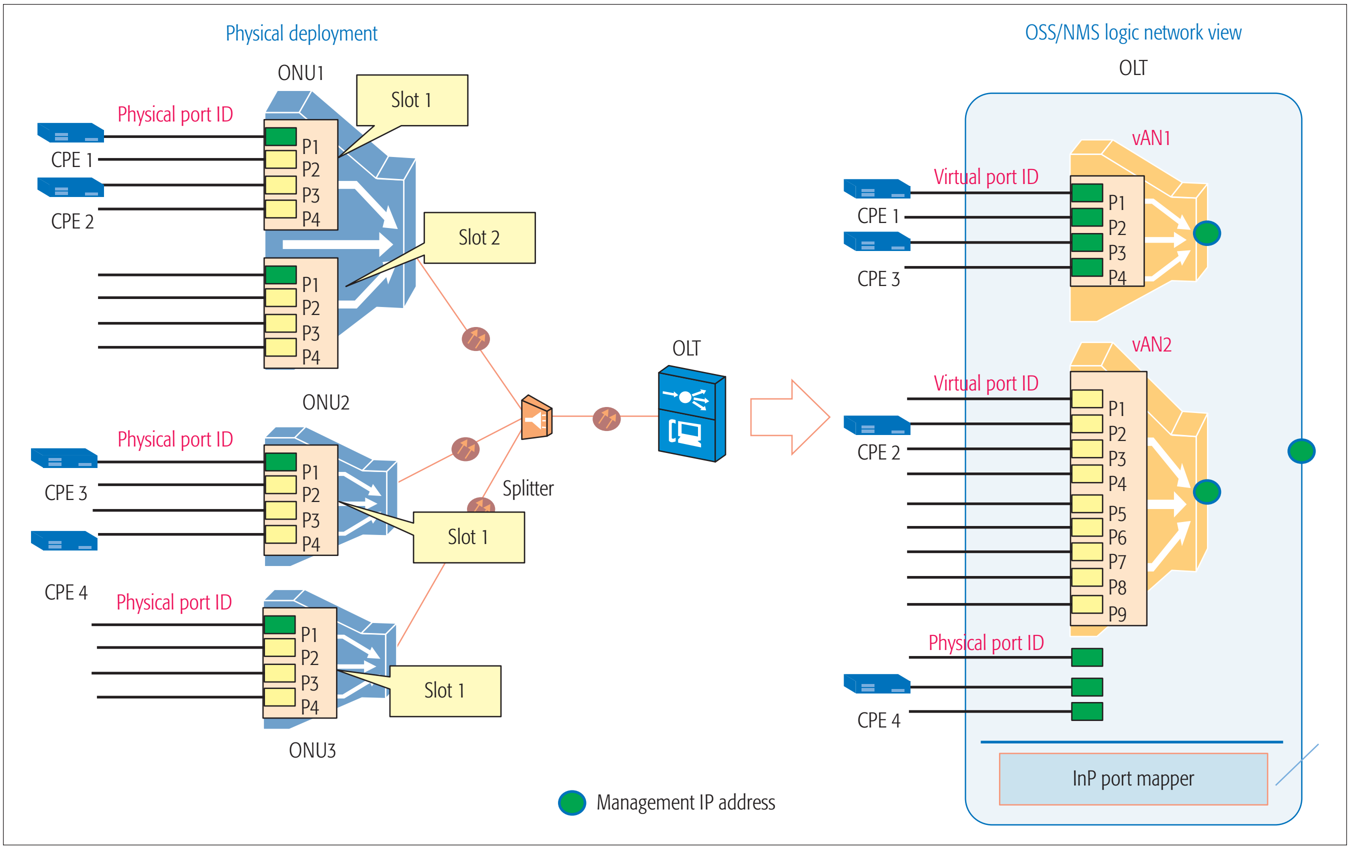

FIGURE 5. Port mapper [1].

cally varying the G.fast asymmetry ratio, or configuring PON dynamic bandwidth allocation (DBA).

- Dynamic resource assignment can be virtualized (e.g., access and backhaul bandwidth assignment).
- There can be virtualized dynamic spectrum management (DSM) and dynamic line management (DLM) [10].
- Power control entity (PCE), cross-layer low-power mode control, for G.fast: There are a number of thresholds and other settings that can be varied to configure low-power mode on individual transceivers, and these settings and primitives can be determined in a virtualized power control entity and communicated to the transceivers.
- VDSL/G.fast vectoring control and management: Virtualized functions can control part of the vectoring configuration, and could even calculate vectoring coefficients [9].

STANDARDS

An initial effort at standardizing part of FANS for DSL was undertaken in the U.K. Network Interfaces Coordinating Committee (NICC) [4]. This then fed into the definition of an interface for DSL data sharing in the Broadband Forum [3], and now a larger project on FANS is underway in the BBF [1, 2] to define FANS architectures, requirements, interfaces, and so on.

NICC ND1518, DATA SHARING FOR DSM

The NICC specification ND1518 [4] describes the use of sharing data between operators for the purpose of DSM in DSL environments. Sharing data on cable plant and DSL configuration and performance allows DSM level 2 and 3 multi-line optimizations and DSM level 1/DLM single-line optimizations [10] to enhance the performance of all lines.

Figure 6, from ND1518, shows a somewhat complicated view of data sharing interfaces, including interfaces for DSM/DLM, BSS/ordering, plant inventory, and plant test and repair. Three entities comprise the InP: two access node operators (ANOs), and a transmission path facility (TPF) provider.

BROADBAND FORUM DSL DATA SHARING

Broadband Forum TR-349 [3] describes and defines data sharing for managing DSL. DSL data sharing architectures are described, with centralized and distributed architectures presented. High-level use cases are defined and requirements presented for these use cases. The data and control parameters applicable to each use case are identified, and most parameters are common to all use cases. Two types of DSL data sharing interface are defined in detail: profile-level (general) and parameter-level (specific). The parameter syntax is defined in YANG data models for G.fast and for VDSL in Broadband Forum TR-355.

BROADBAND FORUM FIXED ACCESS NETWORK SHARING

The Broadband Forum has embarked on a program to specify FANS. This is to investigate technical aspects associated with FANS that involve the access network, including access nodes and aggregation nodes. The Broadband Forum FANS project should generate several specifications, the first of which is FANS — Architecture and Nodal Requirements [1]. This document identifies architectures and interface points for FANS. Both management-system-based sharing and virtual-access-node-based sharing are included. Requirements are included, and operation, administration, and maintenance (OAM) and other operations are also described.

A project has also been started in the Broadband Forum to specify FANS access network sharing interfaces [2]. Further specifications, which have not yet been started, include access network virtualization and SDN-enabled FANS.

SUMMARY

Ultrafast fiber-deep broadband deployments are changing the competitive landscape. With current bitstream or VULA virtual unbundling, operations interfaces between InPs and VNOs are often not real time, and VNOs provide simple resale with little or no differentiation. Sharing network resources and management interfaces will allow virtual unbundling to be economically and operationally efficient, and enable vibrant competition based on differentiation of offered services between providers.

FANS extends virtual unbundling to unbundle management and control functions, with network sharing and data sharing. FANS opens up management and control interfaces such that VNOs can perform the same operations as they would with physical unbundling, where they own and operate their own network elements. Much as the Internet offers a platform for innovative applications, so FANS can offer a platform for innovative broadband services.

Standardized interfaces and a central management system are keys to enabling FANS in the near term. In the longer term, virtual nodes, SDN control, and full virtualization will all feed into enabling FANS. Resource control, AAA, security, and configuration control must be carefully administered with FANS to ensure privacy and avoid harm to the network.

FANS offers many benefits, both to InPs and VNOs:

- Automated interfaces lower operational costs relative to manual interfaces, both for the InP and the VNOs.
- Fault correlation across multiple operators' lines is enabled, again lowering operations costs.
- Multi-line, multi-operator optimizations are enabled, which increases performance of all lines.
- Multiple companies can share the costs of network upgrades to superfast broadband.
- It enables enhanced service levels, services differentiation, and innovation.
- The InP can offer FANS as an enhanced service to the VNOs.
- The VNOs can offer enhanced services to the broadband customers (e.g., enterprises).

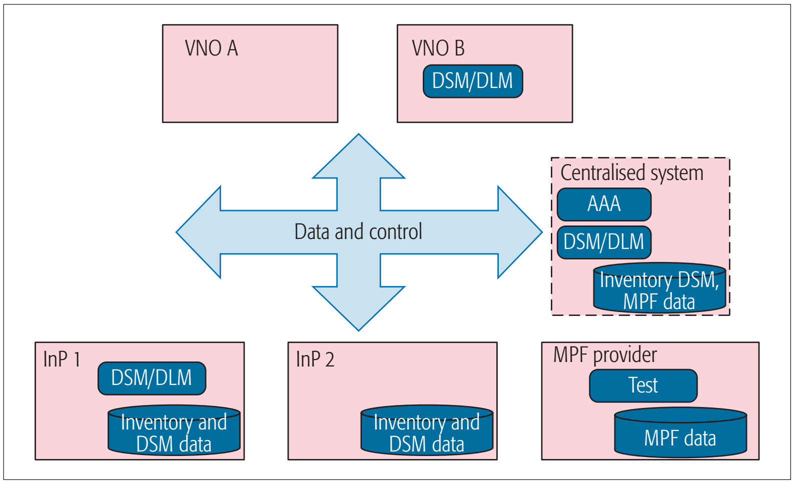

FIGURE 6. Simplified view of data sharing [4].

REFERENCES

[1] Broadband Forum WT-370, "Fixed Access Network Sharing — Architecture and Nodal Requirements."
[2] Broadband Forum WT-386, "FANS Access Network Sharing Interfaces."
[3] Broadband Forum TR-349, "DSL Data Sharing," 2016.
[4] NICC ND1518, "Data Sharing for DSM," 2015.
[5] K. Kerpez *et al.*, "Software-Defined Access Networks," *IEEE Commun. Mag.*, vol. 52, no. 9, Sept. 2014, pp. 152–59.
[6] K. Kerpez and G. Ginis, "Software-Defined Access Network (SDAN)," *CISS 2014: The 48th Annual Conf. Info. Sciences and Systems*, Mar. 19, 2014.
[7] W. Schramm, "NGA Regulation in Austria," TNO DSL Seminar, June 2013.
[8] Broadband Forum TR-101i2, "Migration to Ethernet Based DSL Aggregation," 2011
[9] J. Cioffi, "FULLY Unbundled Vectored DSLs: It Is Absolutely Possible!" G.fast Summit May 2016.
[10] Broadband Forum TR-197i2, "DQS: DSL Quality Management Techniques and Nomenclature," 2014.

BIOGRAPHIES

KENNETH J. KERPEZ [F'04] (kkerpez@assia-inc.com) received his Ph.D. from Cornell University in 1989. He worked at Bellcore and Telcordia for 20 years, and he now works at ASSIA. He became an IEEE Fellow in 2004 for his contributions to DSL technology and standards. He has many years of experience working on networks of all sorts, including DSL, fiber access, home networks, wireless systems, Wi-Fi, broadband service assurance, IPTV, IP QoS, triple-play services, and virtualization.

JOHN M. CIOFFI (jcioffi@assia-inc.com) received his B.S.in electrical engineering in 1978 from the University of Illinois, and his Ph.D. in electrical engineering in 1984 from Stanford. He was with Bell Laboratories, 1978–1984; IBM Research, 1984–1986; and has been a professor of electrical engineering at Stanford, 1986–present, now Emeritus. He founded Amati Com. Corp in 1991 and was officer/director from 1991 to 1997. He currently is chairman and CEO of ASSIA, Inc. His specific interests are in the area of high-performance digital transmission. He has been the recipient of numerous highly prestigious awards, has published over 600 papers, and holds over 100 patents.

Peter Silverman (psilverman@assia-inc.com) is director, Standards and Technical Marketing at ASSIA Inc. Prior to employment at ASSIA he was employed at Bell Laboratories, Ameritech, 3Com, and Valo Inc. before taking his current position at ASSIA in 2005. He has edited numerous international telecommunications standards, and is a co-author of two books, *Understanding Digital Subscriber Line Technologies* (Prentice-Hall, 1999) and *DSL Advances* (Prentice-Hall, 2003) and 10 patents.

BRUNO CORNAGLIA (bruno.cornaglia@vodafone.com) received his Degree in electronic engineering from Politecnico di Torino in 1990. He worked at CSELT, the research center of Telecom Italia Group, Omnitel, which then became Vodafone Italy, and now at Vodafone Group Services. His main areas spread from microwaves, where he developed the concept and solution of microwave Ethernet with adaptive modulation;, satellite, where he worked on defining new SDH payloads to be used in satellite links; and now fixed broadband, including copper, cable, and fiber technologies, where he is working on different innovative approaches, like distributed architecture for cable, NG-PON2 in fiber, and full access virtualization. He is an author or co-author of more than 20 papers.

GAVIN YOUNG [M] (gavin.young2@vodafone.com) received his

B.Sc in electrical and electronic engineering from the University of Wales, Swansea, in 1985 and his M.Sc. in communications engineering from the University of London (Imperial College of Science and Technology) in 1986. He is head of the Fixed Access Centre of Excellence within Vodafone, where his interests include technology strategy, architecture, innovation, and performance of broadband networks spanning copper, cable, and fiber technologies. He is a Fellow of the IET and Distinguished Fellow of the Broadband Forum.