Automating Service Creation and Provisioning for Creative Ethernet Services

Delivering Ethernet Service at Scale

A Technical Paper prepared for SCTE/ISBE by

John Hawkins
Product and Solutions Marketing
Ciena
1185 Sanctuary Parkway
Alpharetta, GA
678-867-3331
jhawkins@ciena.com

Table of Contents

List of Figures

Introduction

Ethernet service uptake continues on the rise, especially among the small and medium business segments. The increasing interest in high-reliability and low-cost options drive operational challenges with deployment and activation headaches that can only be addressed by comprehensive automation of the network infrastructure in an intelligent and operationally sound model.

The MEF has proposed a Lifecycle Services Orchestration (LSO) architecture that allows for such comprehensive automation when coupled with technologies such as Software Defined Networks (SDN), Network Function Virtualization (NFV) and service orchestration. The approach allows for flexibility in service definition to meet varying qualities of service yet is cost-effective in terms of CapEx and (more pointedly) OpEx.

However, the approach breaks down at scale unless specific considerations are accounted for from the outset. The rise of popularity of the software defined WAN (SD-WAN) is a case in point, where an application drives the adoption of a promising technology whose business case depends on automated deployment and provisioning. Here we explore the technology components that make this possible in the cable operator context.

The Challenge of Scale

1. Unrelenting demand for data

The chairman of India’s Reliance Industries conglomerate and the world’s richest man recently opined that “data is the new oil.” His perceptive remark points out that the world runs on data. Proof of this can be found in many forms. Consider the fact that the largest companies in the world today (by market capitalization) are Apple, Alphabet, Amazon, Microsoft and Facebook. Ten years ago, only one of these (Microsoft) made that list alongside the likes of Exxon Mobil, Petrochina and Royal Dutch Shell. It is no coincidence that these newcomers are also the top cloud data players in the market today and serves as proof that data, and quick access to it, drives not only business and commerce, but modern-day life as we know it.

The prevalence of cloud-based commerce makes for a challenge operators haven’t experienced since the early days of telephony (or perhaps mass electrification): how to scale the network to meet the demand for instant, secure and reliable access to all that data. In the case of telephony or electricity, the rollout of those services was taken up country by country, region by region, whereas in this instance everyone wants access to their data, and want it now irrespective of locale or other context. In both business service and mobile backhaul situations, Ethernet and IP-based services have become the instruments of choice to allow for private network builds as well as access to the public Internet for cloud based application support.

Of course, making the business case for the massive infrastructure shift required is not simple. Observe that Operational Expense (OpEx) exceeds Capital Expense (CapEx) in most operator environments (see Figure 1), and owes this to the care and feeding of the growing infrastructure. New technologies have been pioneered to address this fact given that operational efficiencies deliver the best chance for increasing margins and justifying the new investment in infrastructure.

CFOs have mixed feelings about the financial implications. While CapEx allows for better up-front planning of costs and amortization of the investment over an extended period of time, the relatively short technology refresh cycle (as compared to former telephony or electrical technologies) tends to work against predictable budgeting as technology moves faster than operators can digest. In an OpEx environment, at least a portion of the technology refresh cycle can be outsourced to the vendor, who collects a usage-based fee for the latest and greatest technology and sees to it that it maintains currency with the state-of-the-art. At least that's the theory.

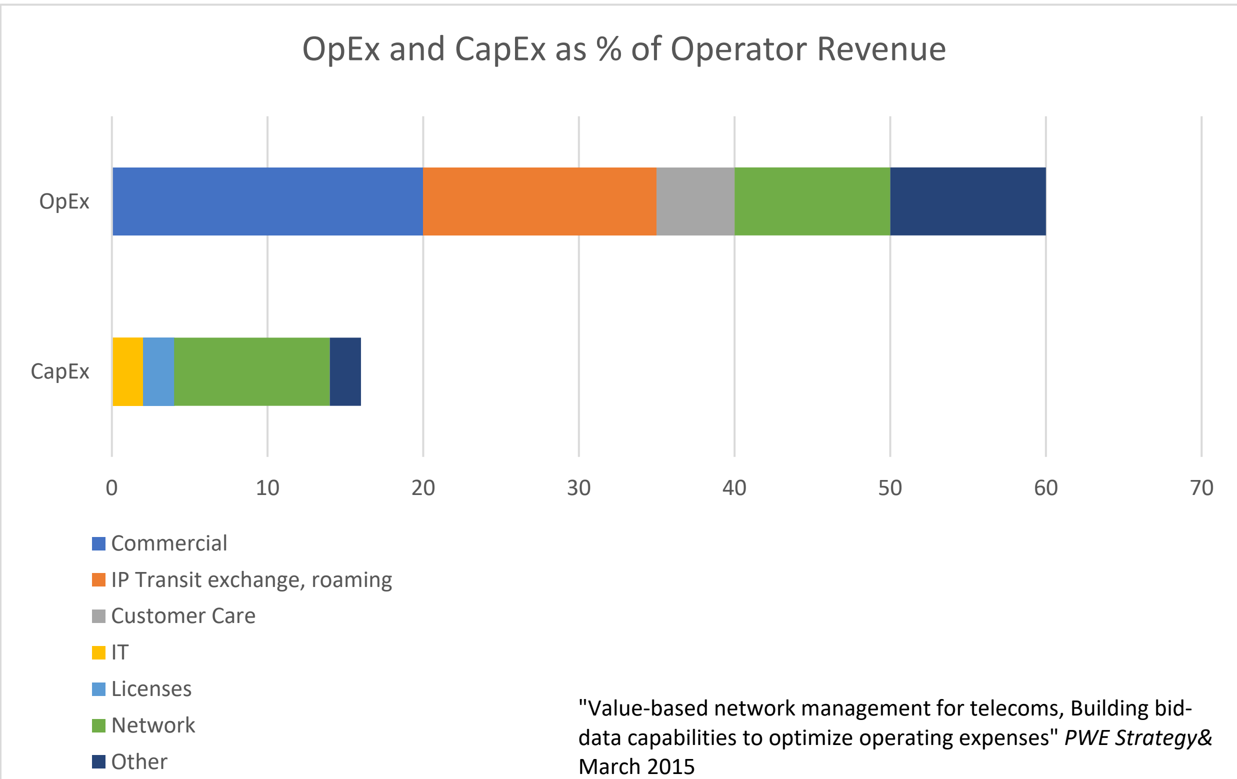

Figure 1 - Cost Breakdown for Telecom Operators

But shifting costs from CapEx to OpEx (or from operator to vendor) doesn't solve the fundamental underlying challenges of complexity and scale.

Into this complex mix, technology has been brought to bear to solve the problems it may have itself created. In fact, this challenge of "scaling out" the network is at the root of most of today's technology advances. Fundamentally, Software Defined Networks (SDN), Network Function Virtualization (NFV), Lifecycle Services Orchestration (LSO) all purport to address the high cost of operationalizing an expanding network. Automation is the bottom line and reducing OpEx the main goal.

2. Modern deployment and service commissioning

The processes and procedures employed to initially start up a new service at a new customer site are some of the most labor intensive and costly activities operators face. Thousands of end points must be brought “on-net” and services provisioned properly and quickly before revenue can flow to the operator’s books. Inefficiency and error can delay revenues, increase customer complaints and lead to competitive disadvantage. The conventional approach involves scheduling several tasks:

- Customer premise equipment (CPE) must be delivered, installed in an appropriate environment, powered up, connected to the network, configured with the appropriate software, commissioned for the appropriate services and authorized to use the appropriate network resources.
- Live communication between the field technician and the network operation center (NOC) personnel is required to ensure both ends are communicating.
- The NOC or technician then initiates a series of service activation tests such as line rate testing, latency, packet loss, etc. to produce a benchmark for the service.
- The NOC then enables the end-to-end service per the order once the back-office tasks of performance measurement and billing have been set up.

One or more customer visits or “truck rolls” are required for all this to come together, and still the process remains fraught with the potential for error and further delay leading to grumbling customers.

To date several techniques have been perfected to address these service startup challenges.

2.1. Zero Touch techniques

Zero touch provisioning (ZTP), or low touch provisioning, has been used for over a decade to auto-provision CPE gear allowing for rapid service turn-up at newly installed locations. While fiber availability or installation is a pre-requisite for any service activation, turning up the service once the media is in place can also be a significant bottleneck if not addressed.

The process involves automating the provisioning tasks listed above. This might look something like Figure 2 and involve the following steps

1. An order arrives from the customer and is entered into the order flow.
2. The NOC confirms the order (payment terms, timeframes, etc) and generates a work order to the warehouse for CPE.
3. The NOC generates the appropriate service profile(s) and a software load with the appropriate configuration script(s). These are pre-loaded to a network server and awaits the device installation.
4. Meanwhile the appropriate device is shipped to the customer or alternatively delivered via truck roll with technician.
5. Once powered and connected to the network, the CPE device requests an IP address via a call to the DHCP server on the private VLAN. Once the device authenticates, the software and configuration scripts are downloaded to the device. Alternatively, a technician can scan a barcode on the device indicating its unique ID, allowing the NOC to initiate the transfer from the server once the correct device ID is confirmed.
6. Service templates are applied once the device is discovered by the NOC.
7. Line-rate service activation testing is completed and a “birth certificate” report can be sent to the customer indicating the service is up and running. Further “report cards” can be issued per the service level agreement (SLA).

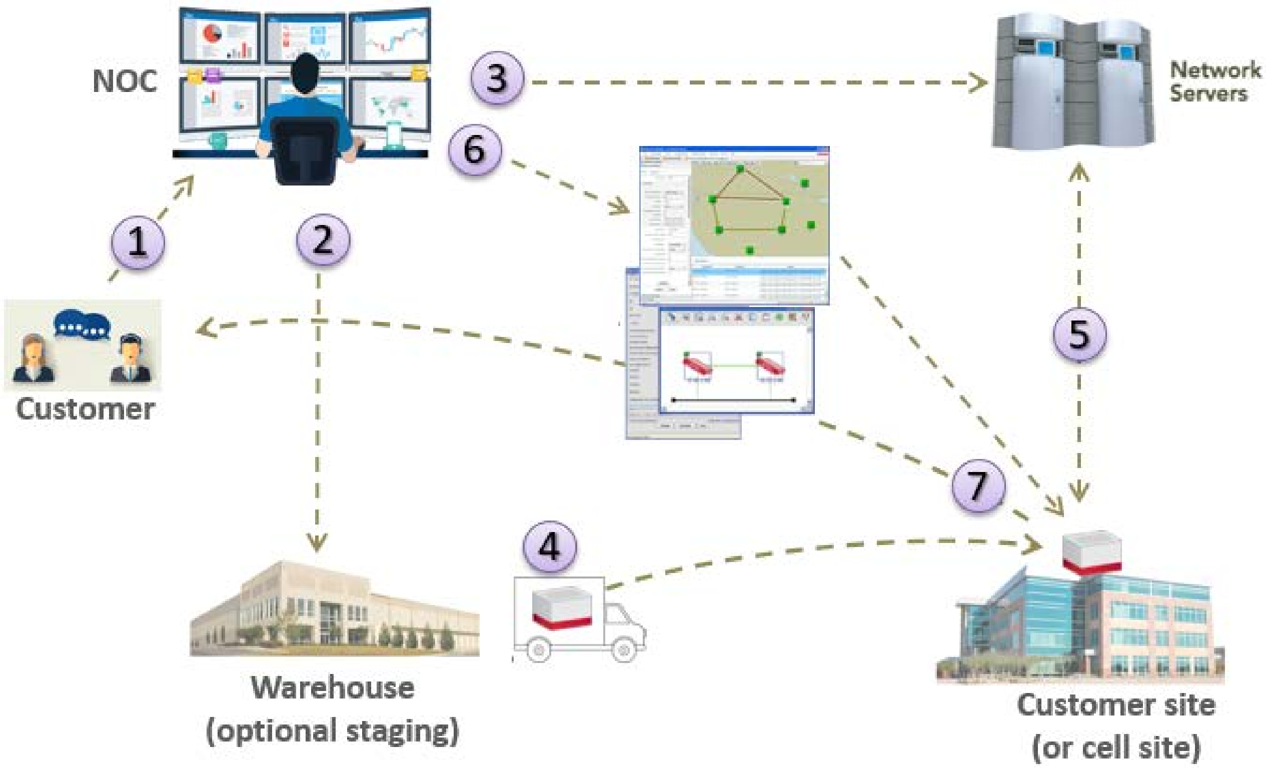

Figure 2 - Zero Touch Technique

Automating these service deployment steps results in faster time to market, increased accuracy of provisioning, higher customer satisfaction, minimized training costs for field personnel, highly scalable growth in service turn-up, and ultimately lower OpEx.

2.2. Remote maintenance and upgrades

Maintaining the ongoing health and performance of the service also drives significant OpEx throughout the lifecycle of the service. Given it is far more expensive to capture a new customer than keep a satisfied one, it is critical that these expenses too are minimized by maintaining a high level of customer satisfaction. Maintenance upgrades and feature enhancements can be difficult to schedule and may result in service interruption. While these cannot be completely eliminated via automation, consistent and accurate software patch/upgrade distribution can deliver similar benefits to those from automated turn-up, namely lower incidence of human error, faster deployment and higher customer satisfaction. Similarly, automated system configuration backup and restoration can minimize downtime when errors or network failures do occur and the network state needs to be quickly restored.

To the extent these processes can be triggered, executed, and verified remotely (and admittedly, not all can) the OpEx challenge is minimized and downtime is minimized.

2.3. Orchestration

As packet services have become more complex involving multiple operators, often crossing service territories and frequently involving dynamic behavior such as route selection, virtualized function deployment and variable SLAs, orchestrating all of the moving pieces has been a critical component to facilitate mass rollouts (see Figure 3). Of course with such flexibility of services comes complexity, and complexity usually breeds expense.

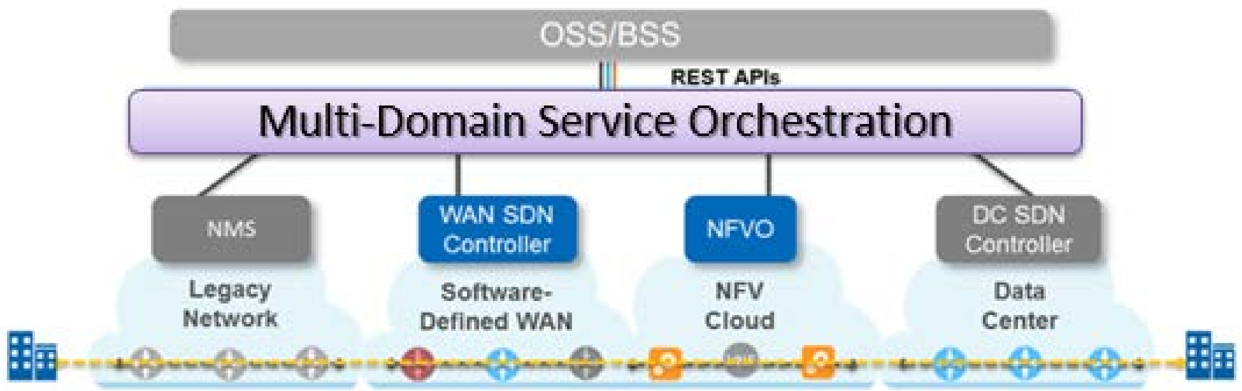

Figure 3 - Multi-Domain Service Orchestration Concept

Network management system (NMS), Operations Support System (OSS) and Business Support System (BSS) coordination is one of the benefits of orchestration and much of the reason the technology has captured the industry's attention and generated a significant amount of hype. Automating NMS functions can enable significant cost savings, as well as adding considerable capabilities to operators' OSS/BSS systems, but should be regarded as a journey, not a destination. So far, large scale deployments have been restrained and lessons learned are being carefully watched by the industry as a whole. MSOs are among the most aggressive proponents of the orchestrated service approach.

The orchestration approach relies on a single point from which devices can be provisioned, monitored, and troubleshot. The benefit of visualizing the entire network from a single point of administration cannot be overstated. By collecting the network state in this manner, the network is better documented, making troubleshooting tasks shorter and more straightforward. This ultimately impacts metrics like Mean Time to Insight (the time it takes to correctly diagnose and triage new issues) and Mean Time to Repair.

Not only is the visibility to network bottlenecks and trouble-spots improved by this approach, configuration files and scripts can be greatly simplified, minimizing the possibility of human error which is the single largest source of network downtime.

2.4. Comprehensive model for automation

Evolving from existing and legacy networks to automated networks is among the primary hurdles many organizations are currently addressing. Few have the luxury of deploying technology in a greenfield environment. However, most are at least evaluating automation via orchestration in limited environments. A select few have implemented orchestration in production environments. Those early adopters are beginning to realize the benefits, as well the challenges involved when doing so at scale.

So far we have only implied the use of today's hottest network technologies: Software Defined Networks (SDN) and Network Function Virtualization (NFV). These are natural corollaries to the orchestrated network as they provide the ability to control the behavior and functionality of the network from a logically centralized controller allowing for considerable flexibility and cost-containment. Individual devices (whether core, edge or CPE) need not be as expensive and high-functioning as before. End-to-end service attributes can be pre-staged, verified against network policies, and pushed to the network as needed.

Ultimately the nirvana state for this virtualized, software-centric infrastructure is the so-called "self-driving network." The vision is one in which high-level software entities can address the network directly via well-defined Application Programming Interfaces (APIs) that allow access to key network attributes and functions. These applications are thereby empowered to request network resources (capacity, end-points, service types, performance tiers, etc.) and receive responses from the network as to whether the request can be fulfilled and at what cost. These interactions can be between consumers and service providers (in the case of a retail transaction) or between operators (in the case of a wholesale transaction). In combination, service providers can thereby "stitch together" end-to-end services using a combination of on-net and off-net infrastructure. While our focus here has been on the effect of automation on OpEx and CapEx, this has the effect of adding considerable value to the service provider's product offerings and hence revenue streams.

3. The MEF LSO framework

In this environment of rapid technology evolution, standards organizations such as the MEF (formerly Metro Ethernet Forum) and the TMF (formerly Tele-management Forum) have stepped in to offer a framework for network automation (see Figure 4). As we stated earlier, automation is not a new concept, but corralling the various industry efforts revolving around SDN and NFV including proprietary and open-source projects has been a challenge so far. The MEF Lifecycle Service Orchestration (LSO) reference model (MEF 55) is particularly useful and forms the basis for our discussion here.

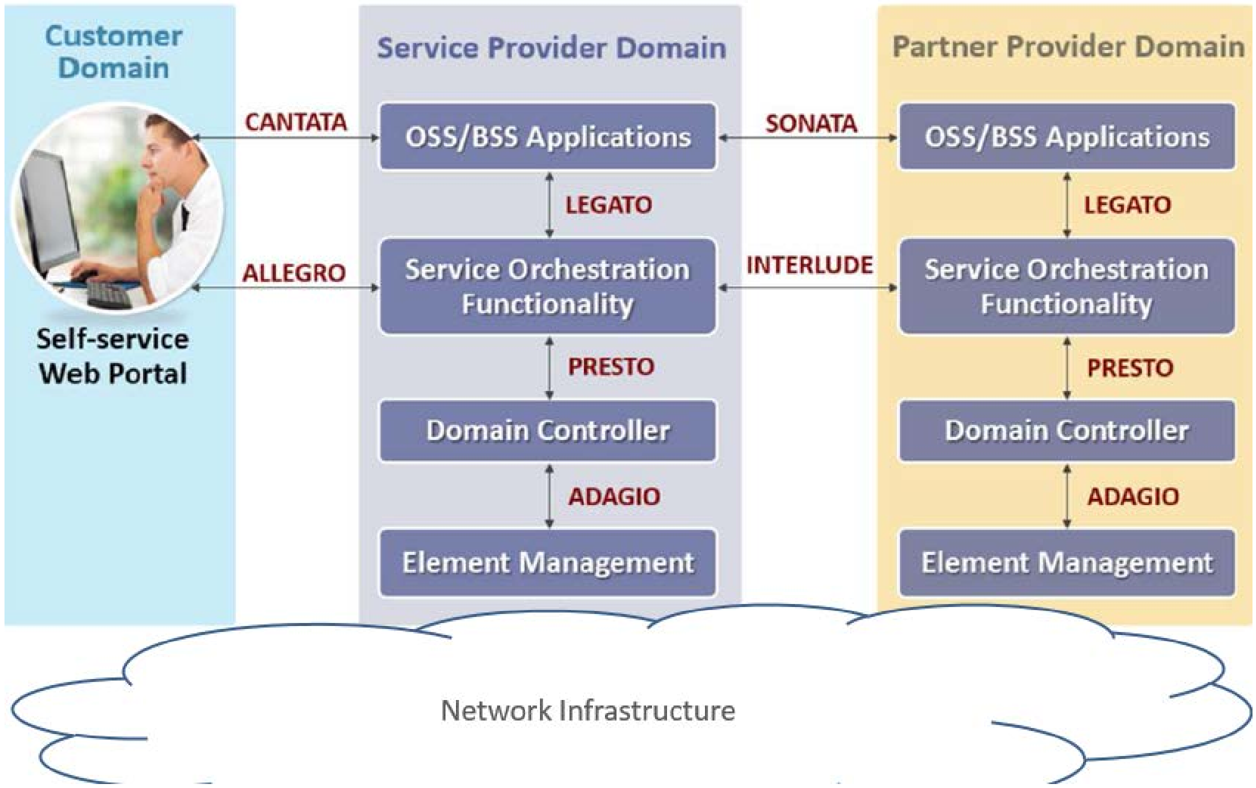

Figure 4 - MEF LSO Reference Architecture

3.1. Standards-based interfaces and operational processes

In order to facilitate interactions between customers, service providers and partner providers, the model defines each of these as individual domains, each with specific interfaces to allow a request/response between them. These interfaces (cleverly named for musical terminology) along with the APIs that give them their functionality, are in early stages of development within the Technical and Operations Committee of the MEF, but already useful early implementations have been floated in the industry via proof of concept and "hackathon" activities involving multiple operators and equipment vendors. Specifically, multiple demonstrations of the Presto interface between service orchestration and multiple domain-specific controllers have been demonstrated, creating MEF-compliant Carrier Ethernet services across more than one operator (also using the Interlude interface).

The ultimate goal is to allow for certified MEF services to be created using these interfaces along with value-added virtual functions. Layering additional managed services via NFV and orchestrating these such that they can be deployed, updated, monitored, and managed by the service provider opens entire new avenues of revenue than can compensate for the commoditization of the bandwidth-only service. Several use cases have been proposed by the MEF and are reviewed in section 4.

3.2. Standards-based connectivity services

The basis for any managed service is the connectivity service itself and the MEF service types continue to enjoy good market penetration worldwide. E-Line, E-LAN, E-Tree, E-Access, and E-Transit service types are well understood by operators and a growing number of enterprises who increasingly demand them. These provide options of native Ethernet connections, either private or virtual that can address capacity, QoS, and resilience, requirements for access to data centers, branch offices, and the Internet itself. Given the challenges of access and metro network connectivity, along with the comparative expense of IP/MPLS services, these Ethernet based services have become the workhorses for those building networks to access cloud-based data.

Their value lies in the fact that these are standards-based, well-understood and increasingly widely-available services. In the past, a TDM-based service such as ATM or Frame Relay would have been the mainstay of enterprise connectivity. They were well understood, consistent in behavior and performance, and very cost-effective. Today, they are challenged to deliver the bandwidth required by modern applications and have largely been replaced with IP-based alternatives such as IP-VPNs based on MPLS technology. But further growth in bandwidth requirements and increasing complexity of those protocols have led to a resurgent interest in Carrier Ethernet services.

4. Ethernet business service use cases (a sampling)

4.1. Connectivity +

Operators are interested in introducing new value-added services that can command higher revenues than simple commoditized connectivity. Doing so at scale can lead back to complexity, and therefore requires careful planning and considerable automation to avoid jumping from the pan back into the fire. There is a need to rapidly introduce new services to market, test their viability (both economic and technical) and then scale out the network.

NFV based services allows for a software-centric model of doing just that. This “Connectivity +” or Virtual Network Function as a Service (VNFaaS) business model is in its infancy and is made possible by the holistic view of the lifecycle of the service now well captured in MEF standard MEF 55. Concepts such as “branch-in-a-box” are being tested in various markets with connectivity services bundled along with firewalls, encryption, virus scanners, WAN optimizers and the like. With this distributed NFV model, service providers can provide ongoing, subscription-based updates to the software, and avoid separate and costly appliances otherwise required for each service. The generic processor complex required for NFV also allows a future opportunity for added value services to be upsold for added revenue opportunities.

4.2. Wholesale Network as a Service (NaaS)

Figure 5 shows an example of an MEF E-Access service using two operator networks to access a datacenter in an “out-of-footprint” scenario. While the service provider cannot offer direct access to the desired datacenter, an E-Access service from the wholesale operator can complete the end-to-end connection. These wholesale arrangements are not uncommon. The reality, however, is that such arrangements are labor-intensive given both business and technical perspectives. Automating the interactions based on pre-arranged business-driven agreement can drastically reduce the time to provision and turn-up the service using the MEF Sonata interface and associated API. Reducing provisioning from weeks to hours, perhaps minutes makes such a service viable at scale and highly valuable to end-users large and small.

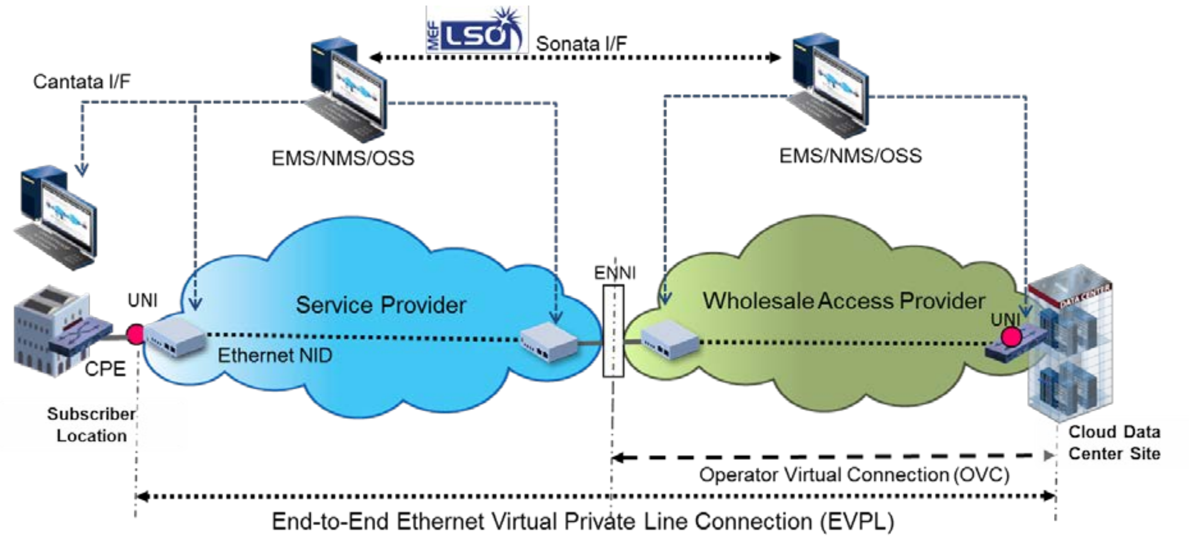

Figure 5 - Multi-operator Datacenter Access Service

4.3. SD-WAN

The software defined WAN (SD-WAN) is thought to be an even more compelling use case for enterprises wanting to connect branch locations to their vital sources of data. SD-WAN is intended to optimize the use of multiple links to a branch office based on policy and best-practice. For example, web access might be provided using a low cost Internet access links, whereas more vital IP telephony or video conferencing traffic might be steered towards a fiber-based, protected and private E-Line service. Wireless connections may also be leveraged in the event of poor performance or outages of landline services. Policy may also dictate that different connections be used at differing times of day to take advantage of cost savings or congestion avoidance.

A more ambitions implementation of SD-WAN envisions a centralized controller using SDN techniques selecting actual routes to minimize latency or jitter performance of a given link for a given performance target. This would involve active and dynamic management of the operator's available resources and would require sophisticated automation. From the point of view of the operator, path selection is already in use to route around network hot-spots (or downed links/nodes) so extending the value to the customer is an additional opportunity for incremental revenues.

4.4. Cloud exchange

Leveraging multiple services from multiple cloud providers offers another use case whereby an operator becomes a broker of cloud services. Similar to SD-WAN in application, multiple cloud services (such as Amazon, or Microsoft's cloud offers) can be provided via a single connectivity service via such a brokerage arrangement. While the network architecture for such connectivity (see Figure 6) is facilitated by MEF service constructs (end-to-end E-Line connectivity stitched together from multiple operator virtual connections), the orchestration of the services in an automated fashion relies on interactions between operators via the MEF Sonata API.

It is important to define methods for assembling services across multiple administrative domains in this model. Not only is interconnecting networks for transiting traffic critical, but also the interconnection business resources to support the network services. In this way, the assemblage of resources appears to the end user as a single, end-to-end service, dedicated to his use and billed by a single entity, the service provider.

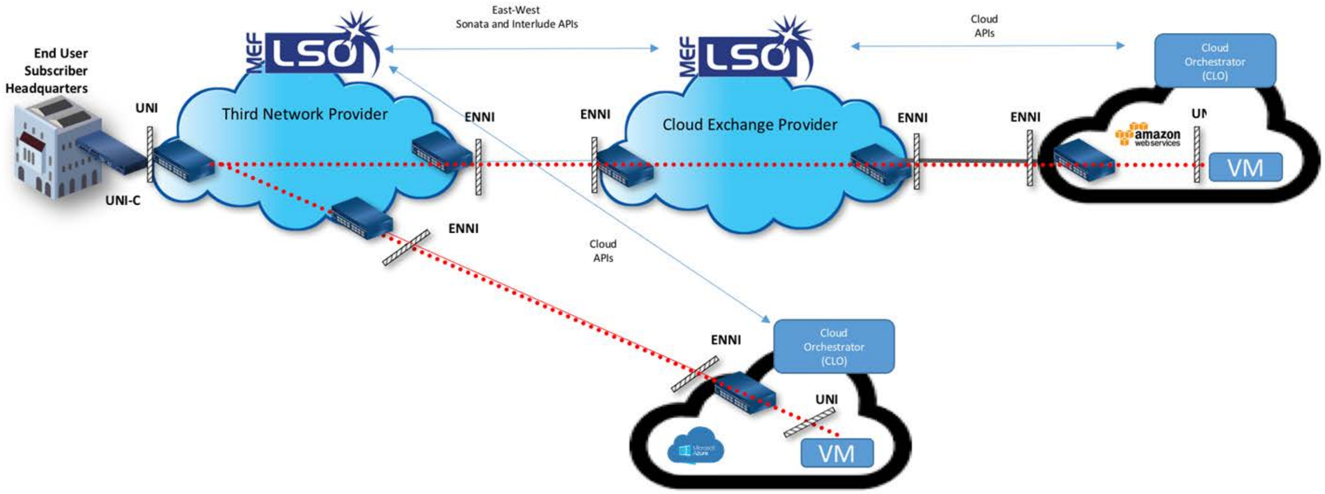

Figure 6 - Cloud Exchange Access

Conclusion

A seismic shift is underway in today's networking market. Connectivity services, while vital to the efficient operation of modern business, is becoming a commodity. Layering services in addition to connectivity is seen as a viable and valuable approach to deliver what end-users really want, however the opportunity for simply layering complexity onto an already complex environment constrains its rollout at scale. Techniques such as zero-touch-provisioning, service scripting, SDN/NVF, and lifecycle service orchestration address this complexity so these value-added services can be rolled out en-masse without breaking the service provider's business case.

Abbreviations

BSS	Business Support System
CaPex	Capital Expense
CE	Customer Edge
CEN	Carrier Ethernet Network
CPE	Customer Premise Equipment
EMS	Element Management System
EPL	Ethernet Private Line
EVPL	Ethernet Virtual Private Line
LAN	Local Area Network
OpEx	Operational Expense
OSS	Operations Support System
MAC	Media Access Control
MEF	MEF Forum (formerly, Metro Ethernet Forum)
NFV	Network Function Virtualization
NFV-O	Network Function Virtualization Orchestrator
NID	Network Interface Device
NMS	Network Management System
SDN	Software Defined Network
SLA	Service Level Agreement
UNI	User-Network Interface
VLAN	Virtual Local Area Network
VM	Virtual Machine
WAN	Wide Area Network

Bibliography & References

Shane, D. (2017, July 31). Reliance's Big Data Play. Barron's. Retrieved July 31, 2017, from http://www.barrons.com/articles/reliance-makes-big-data-play-with-cheap-mobile-1501303073M.

Walker, "A Growth Opportunity for Vendors: Telco Opex," OVUM, Oct. 31, 2012.

"Central Office Re-architected as a Datacenter (CORD)." (n.d.). Accessed June 01, 2017, from https://wiki.opencord.org/pages/viewpage.action?pageId=1278047

"Network Functions Virtualization — An Introductory White Paper," *SDN and OpenFlow World Congress, Oct. 2012*.

Berde, P., M. Gerola, J. Hart, Y. Higuchi, M. Kobayashi, T. Koide, B. Lantz, B. O'Connor, P. Radoslavov, W. Snow, and G. Parulkar "ONOS: Towards an Open, Distributed SDN OS." *ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking (HotSDN 2014).* August 2014.

MEF Forum, *Service Operations Specification: Lifecycle Service Orchestration (LSO) Reference Architecture and Framework, MEF 55*, March 2016

2007.Mehta, Kumar. “The Evolution of SD-WAN to SD-Branch.” *SDx Central*. July 14, 2017, Accessed July 15, 2017, from https://www.sdxcentral.com/articles/contributed/evolution-sd-wan-sd-branch/2017/07/

Tarplin, Johnathan. "Can The Tech Giants be Stopped?" *Wall Street Journal*, July 15, 2017, Review sec.

Whelan, Greg. "How the CORD Project Will Change the Economics of Broadband." *Linux.com*. August 10, 2016. Accessed June 1, 2017, from https://www.linux.com/blog/how-cord-project-will-change-economics-broadband

Aitor Gutierrez, et al March 31, 2015. “Value-based network management for telecoms: Building big-data capabilities to optimize capital and operating expenses.” Accessed June 01, 2017, from https://www.strategyand.pwc.com/reports/value-network-management-telecom

