Leveraging Machine Intelligence and Operations Analytics to Assure Virtualized Networks and Services

A Technical Paper prepared for SCTE/ISBE by

Andrew Sundelin
Director, Product Management -- Cable Innovation
Guavus
1800 Gateway Drive, Suite 160
San Mateo, CA 94404
303-883-1226
Andrew.sundelin@guavus.com

Table of Contents

List of Figures

Abstract

Operators will benefit from the flexibility that comes with the ability to apply IT virtualization technology to CPE and network functions. However, virtualization introduces additional complexity, creating the need for an orchestration layer with more sophisticated assurance capabilities, including data center and network analysis capabilities, encompassing physical and virtual resources in real time.

This paper focuses on a valuable use-case, which leverages operations analytics (OA) and machine intelligence (MI) to drive a variety of resource allocations within virtualized networks. This resource allocation applies both within the “cloud” and in the access network.

Within the cloud, OA and MI can be utilized to predict when additional capacity is needed for services such as additional compute to maintain quality of experience (QoE) for a cloud-based guide or additional storage for a cloud-based DVR system. These technologies can also be utilized to predict unexpectedly popular content and pre-position it optimally within content distribution networks. Within the access network, OA and MI can be utilized for congestion prediction and service optimization. For example, trends in PHY performance parameters can be leveraged to predict the need to tweak OFDM/OFDMA profiles for maximal efficiency.

This paper also explores an emerging concept: “Software-Defined Operations (SDO),” which applies Software-Defined Networking’s (SDN’s) separation of data and control planes to key pieces of modern care and operations equipment. For example, if one views technical support calls as the data then the Interactive voice response (IVR) system is the equivalent of a router (and, thus, part of the data plane) and programming that IVR system is, thus, part of the control plane. With that concept established, the paper looks at the power SDO can have when combined with Operations Analytics.

Virtualization creates benefits and challenges, both of which are driven by the new dynamicity of network services, topology, inventory and hardware resources. The combined application of OA technologies and MI can help operators to assure virtualization by providing an automated, real-time approach that learns from data and autonomously adapts to new information, intuiting connections and relationships to proactively detect anomalies and prescribe solutions.

Introduction

Networks are once again in transition. First there is ongoing development of existing cable technologies. Then there is the adoption of innovation from outside the industry. This has happened in the past, for instance, with fiber and high-speed data. It is beginning to happen with cloud computing and virtualization. MI is another promising and timely import.

Operators will benefit from a new level of flexibility that comes with the ability to apply IT virtualization technology to customer premises equipment (CPE) and network functions. However, virtualization introduces additional complexity, creating the need for an orchestration layer with more sophisticated assurance capabilities, including data center and network analysis, encompassing physical and virtual resources in real time.

Combined with OA, MI can play a key role in this emerging era. After a brief background discussion, this paper addresses the applicability of MI/OA through several use cases. Of special relevance are solutions to resource allocation problems within the cloud (nDVR, guide) and access (DOCSIS) portions of the cable network. The emerging category of "Software Defined Operations" holds additional promise.

Advanced analytics work apart from virtualization. But OA technologies and MI are especially well suited to help operators assure virtualized networks and services by providing an automated, real-time approach that learns from data and autonomously adapts to new information, intuiting connections and relationships to proactively detect anomalies and prescribe solutions.

Clouds and Virtualization

Unveiled about five years ago, notably at a SDN gathering, and directed toward telecom service providers, network functions virtualization (NFV) was aimed at a number of goals, including cost reduction, speed, agility, innovation, and improved services. [1] While having appeared first, SDN complements NFV. The separation of control and data planes that SDN delivers can enhance the infrastructure enabled by NFV.

At NFV's debut, MSOs already were assessing the relevance of SDN. [2] They were engaged in cloud initiatives, such as network DVR and cloud-based guides. Industry leaders also agreed that web technologies could apply to networking. [3] In terms of deployment, however, neither SDN nor NFV are in wide-scale production. For many operators, virtual CCAP (vCCAP), among other technologies, fits that description. At best, virtualization falls in what one industry leader calls the upcoming "second wave." [4]

To date, there have been limited implementations of SDN-style control planes in DOCSIS networks, [5] but industry leaders are aligning with the broader NFV initiative. Comcast, for instance, serves on the board of the Open Network Operating System (ONOS) project, according to which central offices (and, by inference, head-ends) are being 're-architected' as data centers. [See the related Central Office Re-architected as a Datacenter (CORD) and Headends Re-architected as Datacenters (HERD) concepts.] [6] [7]. Where data centers go, clouds are sure to follow.

Machine Learning, Machine Intelligence & Operations Analytics

Figure 1 shows the evolution of systems from legacy database systems on the left which are generally fairly siloed in the data they contain (e.g. operated by a single department and only containing that department's data) to big data systems (for example, data lakes) which federate disparate data and often provide wider-scale to the organization via cloud-based access. However, one common complaint about data lakes is that operators don't know how to derive actionable insights from all the data that they have centralized – this is where big data analytics come into play. The right side of the figure illustrates real-time streaming analytics – leveraging machine intelligence and integrated with an orchestration layer – more realizing the potential of the data lake (for data at rest) while both accelerating and automating operations decision making incorporating data in motion and machine intelligent algorithms.

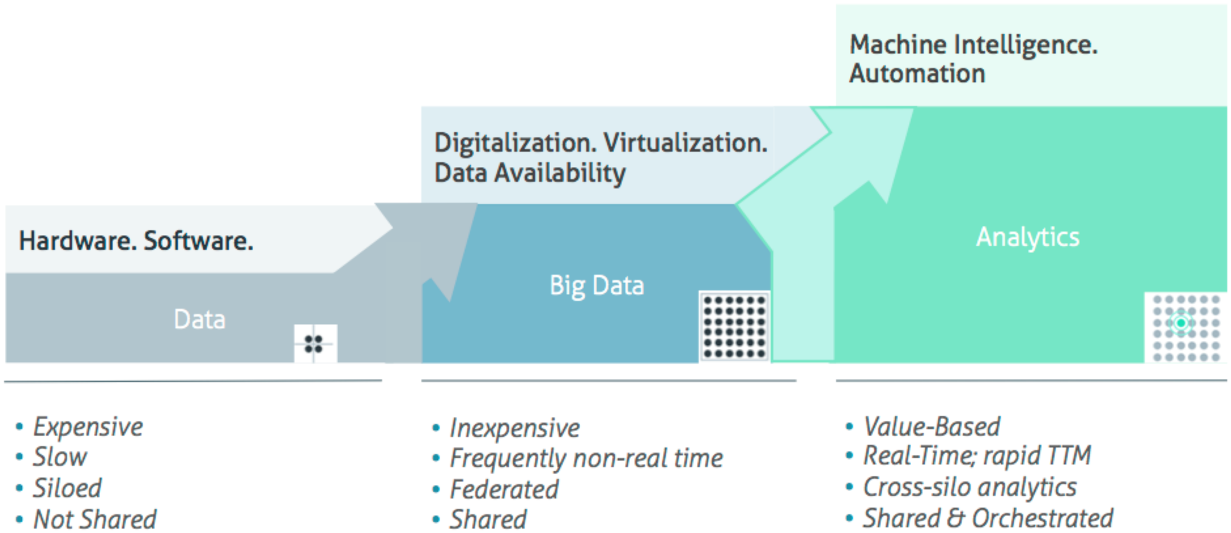

Figure 1 - The Evolution of Analytics

Machine learning (ML) is a relatively mature branch of analytics in which machines improve their ability to recognize patterns as they continue to be trained with additional examples – without having to be programmed to handle each new example or pattern. MI is a somewhat newer field which improves upon ML by adding the ability to reason. Machine intelligent systems are capable of forming hypotheses from raw, disparate data to develop new, valid information that is not a direct result of data in the original data set.

MI is a branch of the broader field of analytics, which involves studying past historical data to identify and interpret patterns. OA applies analytics techniques to the operational realm. More specifically, OA relates to historical and real-time business processes, including resource planning, service monitoring and service diagnostics.

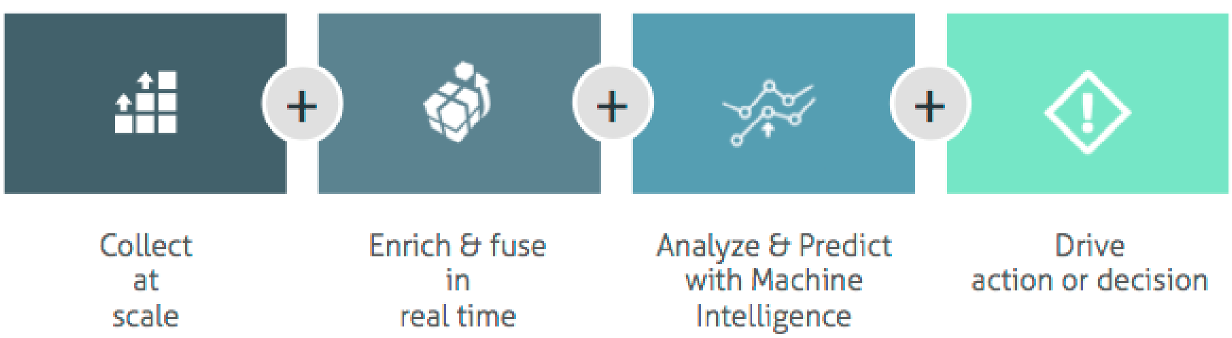

Figure 2 - Operations Analytics Functions

This paper looks at Operations Analytics systems which can address virtualization use cases in real time by leveraging a number of broad functions:

- Collection at Scale: collect event data in real-time and at massive scale from a variety of sources leveraging a big data engine;

- Real-time Enrichment: enrich and fuse cross-functional data in real-time with other events and reference data – combining data in motion with data at rest;
- Analyze & Predict with MI: monitor millions of event time-series, apply machine learning for baselining, anomaly detection and real-time prediction with models built via machine-intelligent algorithms for actionable intelligence; and
- Drive Action or Decision: prescribe actions and integrate with downstream systems to perform a network or operations function.

OA can have numerous applications to the service provider domain, but this paper focuses on those applications which relate to virtualization.

How does MI/OA Apply to Virtualization?

From the start, NFV was envisioned as interacting with legacy OSS and network management system (NMS), especially given the requirements of service provisioning and management, as well as with new elements, such as network controllers and cloud managers. [1] [8]. Within an SDN/NFV framework, the orchestration layer handles those complex tasks bridging physical and virtual resources.

For SDN/NFV to reach its potential, automation became a clear requirement. The logic is compelling: Programmable networks generate considerable amounts of data, which creates the potential for more intelligent, closed-loop decision-making. And as services can be spun up (or torn down) in an NFV environment, what was once challenging enough to manage becomes even more so. Operations analytics – which adapt to circumstances and draw new connections and insights through real-time data processing – therefore become integral to orchestration and next-gen OSS. [9]

The link between analytics and the programmable or virtualized networks does not occur in a vacuum. In practice, hybrid implementations draw upon legacy and new network elements for their purposes, such as anomaly detection and prescriptive decisions. [10] Of special value, as noted in some of the following use cases, is OA’s ability to be more accurate and timely than legacy systems. Typically involving long-term and quickly outdated studies, the status quo approach yields static metrics that are applied universally, often misaligned with particular circumstances.

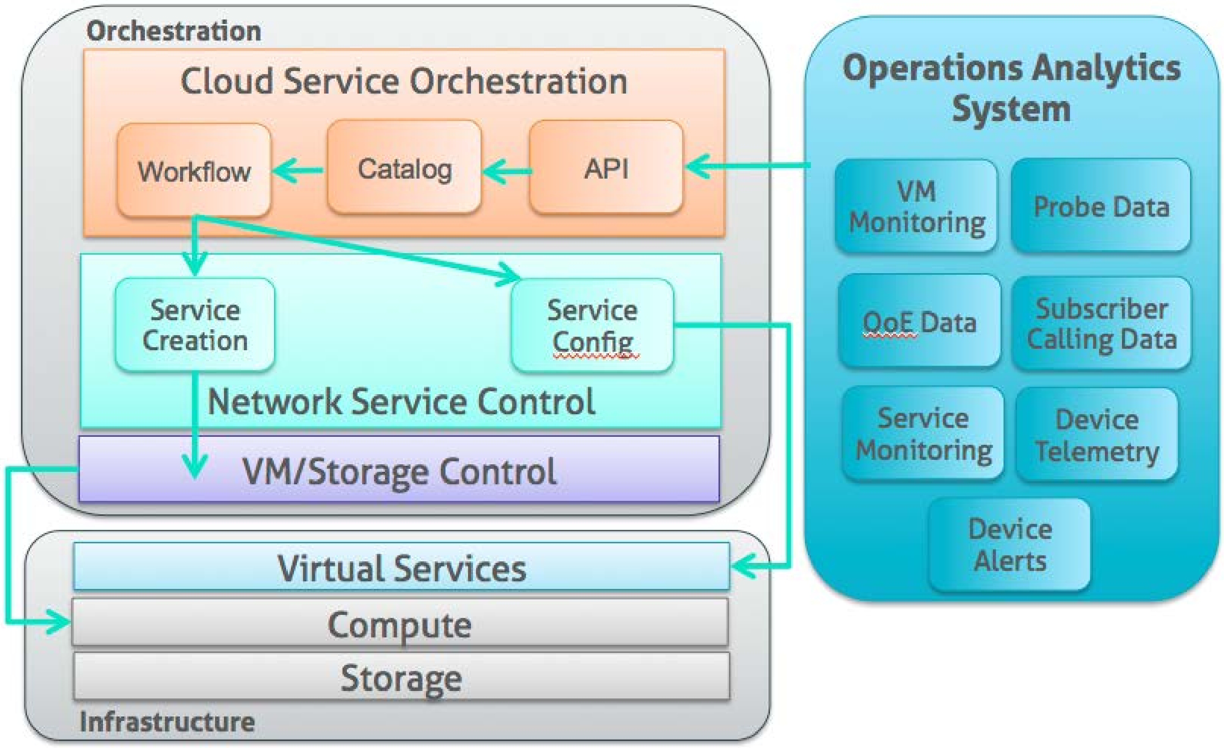

Figure 3 - Operations Analytics and Network Function Virtualization

Operations Analytics Use Cases

These first three use cases apply to cloud computing; the next two to the access network. The final case is an emerging category. They all drive a variety of resource allocation optimizations in networks that are or could be virtualized. Each case reviews the service or scenario, considers the status quo approach, and indicates what the application of OA could deliver.

1. Cloud Guide

A cloud-based, media-rich user interface no longer depends upon limited compute resources of the set-top box. Yet cloud-based guides have limits, too. A system overwhelmed by requests will generate denials for access. The worst approach is simply to wait until those denials are issued and customer complaints begin hitting the care staff. A less naive approach involves fixed thresholds, such as waiting until the resources hit a certain percentage, and then adding CPU capacity. But how can it be assured that the threshold chosen is the correct one? Or that it is correct for all parts of the network?

The ideal approach would be right-sizing a dynamic threshold for given portions of the network with the goal of optimizing resource allocation and only adding virtual machine (VM) capacity when it was truly needed, but also performing this optimization with an eye towards real subscriber QoE measures. Thus, meeting the dual goals of maximizing resource utilization and QoE at the same time. One way to achieve these goals is to leverage automatic feature selection and develop a predictive model for a given target variable. When building a predictive model, different "features" (or data attributes) are evaluated for their level of correlation with a target variable. For example, one might use a target variable of technical support calls about video issues from cloud-based guide users and use CPU utilization as a feature; with this approach one can build a predictive model to determine at what level of CPU utilization do video-related calls from this population begin to deviate significantly from their baseline rate (i.e. become anomalous). Similarly, one might use a target variable of set-top box retry messages and CPU utilization as a feature to determine at what level of CPU utilization does the rate of retry messages begin to be anomalously high.

By performing anomaly detection on the rate of set-top retry messages, an operations analytics engine could determine the level of CPU utilization when retries on the guide begin, or perhaps when retries reach a certain percentage of all requests within a given market. Sensitivity to slow-downs or retries in different regions of the network could vary, depending upon various demographic factors – and could vary over time as STB software or end-user expectations change. Thus, by utilizing a machine-intelligence-based approach to determine when to add additional capacity to a cloud-based guide system an operator creates a closed-loop system that provides optimized, just-in-time capacity all the while maintaining QoE and, yet, adapting to both localized and time-varying user perception.

Figure 3 shows a variety of data which might be applied to determining when to add VMs to a cloud-based guide service (e.g. probe data, QoE data, subscriber technical support calls related to cloud-based guide, device alerts [such as retry messages], service monitoring data, etc.). It also illustrates how an operations analytics system can utilize an orchestration layer to spin up/down VM instances as indicated by the feature-driven model.

2. Network DVR

Leveraging the popularity of CPE-based DVRs, operators have deployed cloud or network DVR service to reduce costs and provide more value to subscribers related to greater mobility, tuning options and storage capacity. Business models vary. In a standard scenario, there may be a static threshold, at which point a subscriber receives more space or an offer to purchase more. Or the business rules may call for deleting content after a certain number of days.

A more dynamic, MI-driven approach, however, will look at the trends in the recordings made by subscribers. This kind of system assesses the need for additional capacity on the basis of not only virtual disk capacity, but also variables such as file-size distribution, delete history and recording frequency. At the macro level, operations analytics could also guide operators on the allocation of resources within clusters of VMs. A large network storage array, for instance, is likely serving not only the network DVR, but also the guide, billing and other needs. Operations analytics could determine which services get how many CPUs allocated to them, and at what times of the day, week, month or year this allocation takes place.

3. CDN Placements

MSOs first engaged with streaming servers and content libraries over QAM VOD. With the introduction of IP VOD and the migration toward all IP, many operators now run their own content delivery network (CDNs). In any case, for efficient transport and better user experience, placing content as close as possible to end-users ("content at the edge") is a best practice. But methods differ. Whereas the status quo approach may use Nielsen ratings or a human expert to make a determination of what is normally popular in a given region, MI/OA will detect more granular patterns, which allows for better QoE while minimizing resource consumption.

With a real-time understanding of actual user behavior in some regions, an operator with a closed-loop operations analytics platform can make timely predictions about user behavior in other regions. For example, as a baseline such a system would prepopulate content based on historic local popularity across all time zones. At the same time, real-time OA would use MI to predict what is likely to become popular – in specific locations within later/western time zones – using actual viewership and demographic data from earlier/Eastern time zones. This approach minimizes latency and maximizes QoE while also minimizing WAN capacity requirements (and avoiding potential WAN congestion) ensuring that both normally popular content and surprisingly popular content are optimally positioned within the CDN.

4. DOCSIS Channel Licenses

With the emergence of the CCAP and massive channel densities, most CCAP vendors have moved from a pure hardware-based model for selling DOCSIS channels (e.g. a 5x20 card where you get 5 DS QAMs and 20 US channels solely by buying the card) to a license-based approach where the hardware itself supports hundreds, if not thousands of channels, but these channels are only usable with a valid license. CCAP vendors have also allowed operators to pool licenses across their networks, decoupling them from a single physical device. This allows operators flexibility in deployment of DOCSIS channel licenses, but with that flexibility comes an optimization challenge – where are these licenses best deployed?

This situation presents an opportunity for MI/OA and, essentially, license-based SDN within the access network. By way of smart license allocation, an operator can optimize license allocations such that licenses from under-utilized portions of the network can be reallocated to “hot spots” in other portions of the network.

Part of the challenge here is essentially a classic capacity planning challenge – but applied in an SDN-oriented fashion. This MI-based approach allows operators to transcend simple, static “one size fits all” notions of capacity thresholding, where rules of thumb might say that any DOCSIS channels are congested at, say, 75 percent utilization. By combining QoE measures such as speed testing probes, customer calls about “slow speeds”, etc., MI and OA can be applied to capacity planning, empowering operators to right-size capacity based on the specific sensitivities of the local subscriber population. For example, some populations may be very sensitive to congestion and start to “feel” the impact of congestion at, say, 72 percent utilization. Other populations may be less sensitive to congestion (either due to different expectations or different traffic types) and not “feel” congestion until 80 percent. This variation presents an opportunity to minimize CAPEX by optimally assigning DOCSIS channel licenses.

Whereas a status quo approach to making these licensing decisions might rely upon one-time historic generalizations, operations analytics can identify real-time variations from narrower sub-populations and prescribe precise and ongoing resource reallocation solutions with much greater efficiency.

5. DOCSIS 3.1 Profiles

As the deployment of DOCSIS 3.1 gains speed, it is useful to recall how its use of Orthogonal Frequency Division Multiplexing (OFDM) enables thousands of QAM sub-carriers per channel, each with its own profile, or modulation value and amplitude. This is the property that enables OFDM to “fit a downstream transmission path like a glove." [11] However, there are only 16 OFDM profiles per downstream channel and each CM has its own unique RF characteristics. Thus, optimizing this limited set of profiles for a large number of CMs per channel can be a difficult task. Configuration of CCAP modulation profiles is relatively manual in early deployments, but over time this process will become more optimized. While embedded processing within the CCAP itself could perform this optimization, this computationally intensive task is better handled externally.

Using the SDN strategy of separating control from data planes, a CableLabs working group proposed and trialed a remote application to manage these profiles, as reported on last year. [12]. The overall effort to use SDN to optimize DOCSIS 3.1 turns on the applicability of k-means clustering, a type of unsupervised machine learning algorithm. [13]. This effort remains a leading example of the application of ML/OA within a software-defined framework for cable. These applications take advantage of the lower price and higher performance of standard computing platforms as well as much longer data retention capabilities of standard computing platforms. Bringing MI/OA to the table allows for not only looking at current conditions when optimizing OFDM profiles, but also predicting future conditions. Operators can do so by leveraging current and past performance and RF conditions. By looking back at a longer history, they can then incorporate seasonality and other longer-term factors in the profile optimization.

6. An Emerging Concept: Software-Defined Operations

A new case involves the extension of SDN into the OSS arena. As the "data plane", for instance, an IVR system is being considered. In effect, it is a network element making decisions about routing calls. The external systems that program the IVR represent the "control plane." In this framework, it can be argued that SDN concepts can be applied much more broadly within an operator's infrastructure than for just intelligently routing data and optimizing network performance. For lack of a better term, this emerging trend could be called "Software-Defined Operations." Many of the leading MSOs already are applying SDN concepts more broadly within their operations.

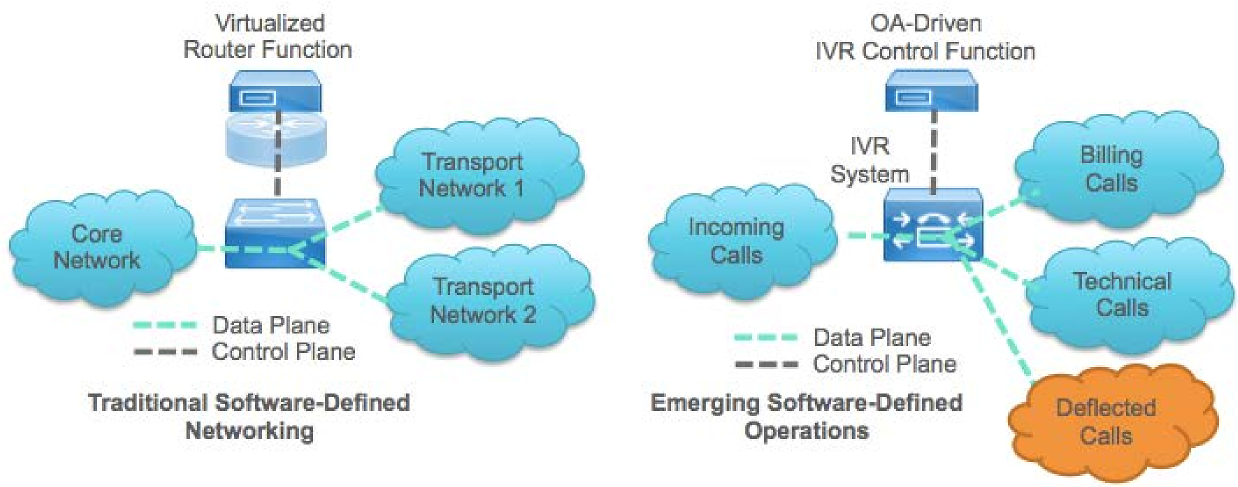

Figure 4 - Software-Defined Networking & Software-Defined Operations

MI/OA is a natural fit in this case, and there is a growing base of operator experience with this approach. One MSO's customer experience analytics (CEA) engine was built to detect network anomalies and automatically trigger a call deflection within the IVR to inform subscribers of a known or likely outage.

The results were impressive, leading to $6.7 million in savings from call deflections and elimination of unnecessary truck rolls. [14] While any outage inflicts damage, proactive and accurate customer engagement driven by "Software-Defined Operations" can not only reduce operations costs but also limit the impact of outages and other service issues on a net promoter score (NPS). With customer satisfaction levels continuing to fall, there is some urgency to finding applications in that domain. [15]

Other potential applications of Software-Defined Operations include:

- Auto-generating tickets in trouble-ticketing systems using MI//OA on network alarms that are predicted to lead to incidents/outages;
- Auto-cancelling booked truck rolls for subscribers which MI/OA has determined are unnecessary; in these cases, the true root issue of their problem is an outside-plant (or other multi-user issue) and not a single subscriber issue, as originally diagnosed by the care agent.

Conclusion

The application of MI and OA is not dependent upon virtualization. As discussed in a paper presented at last year's SCTE Cable-Tec Expo 2016, MSOs are deploying advanced analytics today, to good effect. [16] With other results showing that analytics can reduce truck rolls by 30 percent, the potential for savings in customer care and network maintenance is tremendous.

MSOs can realize those benefits now, apart from SDN or virtualized network functions. Indeed, now is a good time to explore such applications of MI and OA. But the time is coming – possibly in the next technology "wave" – when deploying these smart technologies to assure more fully-fledged virtualized platforms will become critical. Already cases can be sketched out where the application of MI/OA is a good fit for both cloud environments (guide, nDVR) and access networks (DOCSIS). Another promising area for active exploration and experimentation is the emerging category of "Software-Defined Operations."

Abbreviations

CCAP	Converged cable access platform
CEA	Customer experience analytics
CORD	Central office re-architected as a datacenter
CDN	Content delivery network
CPE	Customer premises equipment
DOCSIS	Data over cable service interface specification
HERD	Headends re-architected as a datacenter
IVR	Interactive voice response
MI	Machine intelligence
ML	Machine learning
MTTU	Meant time to understand
nDVR	Network personal video recorder
NFV	Network functions virtualization
NMS	Network management system
NPS	Net promoter score

OA	Operations analytics
OFDM	Orthogonal frequency division multiplexing
ONOS	Open network operating system
OSS	Operations support system
PCMM	PacketCable multimedia
QoE	Quality of experience
SDN	Software-defined networking
SDO	Software-defined operations
vCCAP	Virtual CCAP
VM	Virtual machines

Bibliography & References

Network Functions Virtualization: An Introduction, Benefits, Enablers, Challenges & Call for Action. SDN and OpenFlow World Congress, 2012.

G. White. Can DOCSIS Networks Leverage SDN? NCTA, 2013.

B. Field. Applying Web Principles to the Network. NTCA, 2014.

R. Howald. Aboard the Technology Wave: Surf Conditions Report. Cable-Tec Expo, 2016.

D. Early. A Practical Guide to Implementing Software Defined DOCSIS. Cable-Tec Expo, 2016.

N. Nandiraju. Distributed Access Architecture - Goals and Methods of Virtualizing Cable Access. Cable-Tec Expo, 2016.

S. Chatterjee. Headend Re-architected as a Data Center (HERD). Cable-Tec Expo, 2016.

Telcos eye servers & software to meet networking needs. (Interview with Don Clark, technical manager of NFV industry specification group.) Gazettabyte, April 1, 2013.

CEO Chat with Anukool Lakhina, Guavas, LightReading, July 20, 2015.

B. Lynch. Anukool Lakhina. Implement Closed-Loop Network Decisioning Now with Big Data Analytics and Fuel Future-State SDN Use Cases Through a Common Platform Deployment. NCTA, 2014.

J. Chapman. The Power of DOCSIS 3.1 Downstream Profiles. NCTA, 2013.

S. Rahman. J. Solomon. J. Schnitzer. D. Early. DOCSIS 3.1 Overdrive: Dynamic Optimization Using a Programmable Physical Layer. NCTA, 2016.

J. Schnitzer. DOCSIS 3.1 Profile Optimization Using SDN. Cable-Tec Expo 2016.

A. Sundelin. J.P. Goyet. Using Analytics to Extract Operational Intelligence and Redefine Customer Experience Management. Cable-Tec Expo, 2016.

Wireless Competition Boosts Customer Satisfaction, While Subscription TV and ISPs Face Problems. ACSI Telecommunications Report 2017.

C. Menier. Contextualizing Data to Gain a Real-Time Integrated View of the Network and Improve Customer Experience. Cable-Tex Expo, 2016.