

Contextualizing Data to Gain a Real-Time Integrated View of the Network and Improve Customer Experience

A Technical Paper prepared for SCTE/ISBE by

Chris Menier
VP Products and Strategy
Guavus
1800 Gateway Drive, Suite 160, San Mateo, CA 94404
650-243-3400
Chris.menier@guavus.com

Table of Contents

Title	Page Number
Introduction_____	3
The Quality of Experience Mandate_____	4
The Quality of Experience Challenge_____	5
Defining the Capabilities of a Cable Operational Intelligence Application _____	7
1. Comprehensive Use of Data Resources _____	7
2. Tasks Performed by the Operational Intelligence Framework _____	8
A Tier 1 MSO's Use of Operational Analytics_____	10
Conclusion _____	12
Abbreviations _____	13
Bibliography & References_____	13

List of Figures

Title	Page Number
Figure 1: Historical representation of customer information, which was locked in siloed systems	6
Figure 2: Integrated view of the customer experience	6
Figure 3: Operational intelligence framework	7
Figure 4: Service anomalies for an example MSO by category	9
Figure 5: Monitoring the trending service assurance issues for an MSO	10

Introduction

Never has maintaining consistent quality of experience been more important to cable operators in residential and commercial services alike, and never has achieving this goal been more challenging.

In a competitive environment where consumers and businesses have more providers to choose from than ever, cable companies' ability to capitalize on superior networking technology and service capabilities can only be realized to its full potential if customers can be assured the quality of experience they obtain from relying on cable operators is second to none. In this sense, the customer experience and overall quality of experience has become cable's most important product.

To meet this mandate MSOs must be able to identify and resolve issues independent of reliance on customer input. In an environment where, based on tabulations from work with leading MSOs, over half the customers who experience service problems never report them¹. Equally important, operators must find a way to cut operations costs that takes the guesswork out of finding the root causes of problems, thereby reducing mean-time-to-repair (MTTR), identify what alerts to prioritize and eliminate a significant proportion of unnecessary truck rolls to customer premises and outside plant locations.

Meeting these goals is more challenging than ever in a marketplace characterized by multiple developments unfolding on parallel tracks, including:

- An expanding array of residential and business services;
- Multiple types of core, edge and network infrastructure;
- Proliferation of the types of devices connecting to operators' services;
- An IP-based transformation touching on all facets of operations;
- Consolidation of headends and NOCs (network operations centers) that brings with it the need to maintain operational control over ever larger swathes of service territory; and
- The growing complexity of networks and consumer data fueled by industry mergers.

The only way to meet the customer experience and operational cost challenges posed by these new developments is to reduce the impact of service impairments through preventive action and more rapid responses to issues that are directly affecting customers. This requires implementation of a comprehensive, highly automated and very granular approach to gathering, correlating and analyzing data from every component that contributes to overall performance, whether in the core, the edge, transport links, customer premise equipment (CPE) or personal devices used to connect to services on the go.

Fortunately, the facilities, network and device technologies underlying cable's move into this new service era, including the expansive capabilities of the DOCSIS OSSI (OSS Interface) platform, have been designed to accommodate the need to generate useful performance data, ensuring that operators have plenty of data to aggregate and correlate in the analytical process. As a result, MSOs are now dealing with a Big Data environment that requires industrial-scale advanced analytics to provide actionable insight into everything pertinent to assessing service performance and ensuring superior customer experience.

Such cable-optimized, use case-driven applications must be able to look across all relevant points of functionality tied to any given service in real time and put this vast amount of data into context in order to diagnose the causes behind emerging as well as customer-perceived problems with minimal human

¹ The Problem with Customer Service, [ConsumerReports](#), July 2015

intervention. And this new kind of customer aware analytics application, which provides operational intelligence, must be able to present easy-to-understand graphic renderings of what's happening, providing the most general trend views to highly detailed views of performance at the local and customer premises levels of activity.

Armed with this new operational intelligence, cable operators are able to incorporate into their workloads a number of tasks that once seemed out of reach, including:

- Issue isolation and machine learning root cause analysis;
- Care and field operations optimization;
- Events and alarms reduction, correlation and prioritization;
- Proactive maintenance; and
- New service introduction impact analysis.

Such capabilities are in increasingly wide use by cable MSOs who understand that reducing service maintenance processes by even a few minutes can save millions of dollars while greatly enhancing customer satisfaction. By applying customer aware analytics to continuous streams of data from network elements, end user devices and sensors, call center interactive voice response (IVR) and stored data from CRM systems, billing servers and catalogs, operators can meet the quality of experience (QoE) challenge as never before, notwithstanding the greater complexities of today's operational environment.

In the discussion that follows we detail what operators can expect from a customer aware analytics designed specifically for their current and future needs, including key ingredients for success such as making sure advanced tooling is integrated with home-grown tools, balancing the speed of data versus data enrichment and ensuring that end data is deterministically relevant and actionable. To help illustrate operational intelligence in action, we reference one use case where an MSO is applying advanced data analytics to differentiate between plant and CPE-based issues, resulting in significant reductions in truck rolls and customer complaints.

The Quality of Experience Mandate

The well-documented reality cable companies have confronted for years is that the industry consistently ranks at the bottom in measures of consumer satisfaction across multiple industry sectors. For example, in the latest update to the multi-industry American Customer Satisfaction Index consumer surveys conducted under the auspices of the University of Michigan, the American Society for Quality and CFI Group, the ISP and subscription TV sectors rank lowest on a 100 point scale across 41 industrial sectors with ratings of 64 and 65, respectively, barely higher than the 63.9 points scored by the federal government.² Ratings for eight cable companies included in the subscription TV survey ranged between 54 and 66, all below the ratings scored by satellite and telco providers of pay TV services, which ranged between 67 and 70.

In another widely followed series of reports, J.D. Power & Associates' latest set of four regional customer satisfaction surveys involving over 30,000 pay TV customers found that telephone and satellite companies outranked cable TV providers everywhere, with MSOs consistently falling below the sector

² American Customer Satisfaction Index, [Telecom Report](#), June 2016

averages in all four regions.³ Cable companies fared better in J.D. Power's Internet and telephony surveys, but in all regions the top ranking in both categories went to a telco provider.

Meanwhile, as cable companies battle these traditional competitors for market share they face intensifying competition for viewers' time from an expanding array of over-the-top (OTT) providers, now including individual broadcasters and "skinny" bundlers like Sling TV and Sony Vue as well as long-time SVOD providers like Netflix, Hulu and Amazon Prime. While MSOs' subscriber losses in 2015 represented a 72 percent improvement over the year before, cable operators continued to take the greatest hit in net subscriber attrition with losses totaling 345,000 in 2015, compared to 125,000 telco TV subscriber losses and a gain of 86,000 for DBS providers, according to Leichtman Research Group.⁴

On the OTT side of the ledger, viewing continues to surge. In the U.S., Parks Associates says 57 percent of consumers in broadband households now subscribe to OTT video services with payments averaging \$9 per month.⁵ Globally, as forecast by Digital TV Research, SVOD homes will increase from 117 million at year's end 2015 to 249 million by 2020.⁶ OTT SVOD revenues are projected to hit \$21.6 billion in 2020, up from \$7.6 billion in 2014. Ten countries will have SVOD penetration in excess of 33 percent by 2020, the researcher said.

The Quality of Experience Challenge

Cable companies have made great progress combatting these headwinds, as evidenced by the improvement in pay TV subscriber attrition, the industry's dominance in broadband connectivity, a growing Wi-Fi footprint and rapid growth in commercial services. But, as the ongoing lag in consumer satisfaction performance attests, there's room for significant improvement.

One important contributor to customer satisfaction is a user experience that consistently measures up to or exceeds customer expectations across all service categories. This means operators not only must be able to quickly diagnose through highly automated processes the causes of problems that have produced customer complaints; they must be able to view in real time signs of impending problems that can be redressed through preventive measures before customers are impacted.

Such capabilities require processing of data from every aspect of cable operations, including OSS, billing, customer and network data, which traditionally have sat in data silos as well as more recently data stored in data lakes. At the basic PHY level, the types of facilities over which operators must maintain vigilant control has increased to include multiple types of long-haul and metro networks, their various peering points with the Internet and, at the access level, a wide range of HFC architectures, several types of fiber access networks and a growing Wi-Fi infrastructure, where the ability to parse interactions between devices and access points is vital to sustaining solid performance.

³ J.D. Power & Associates, [Regional Surveys](#), September 2015

⁴ Leichtman Research Group, [Cable, DBS, & Telcos: Competing for Customers](#), March 2016

⁵ Parks Associates, [press release](#), April 2015

⁶ Digital TV Research, [Global OTT TV & Video Forecast](#), June 2015

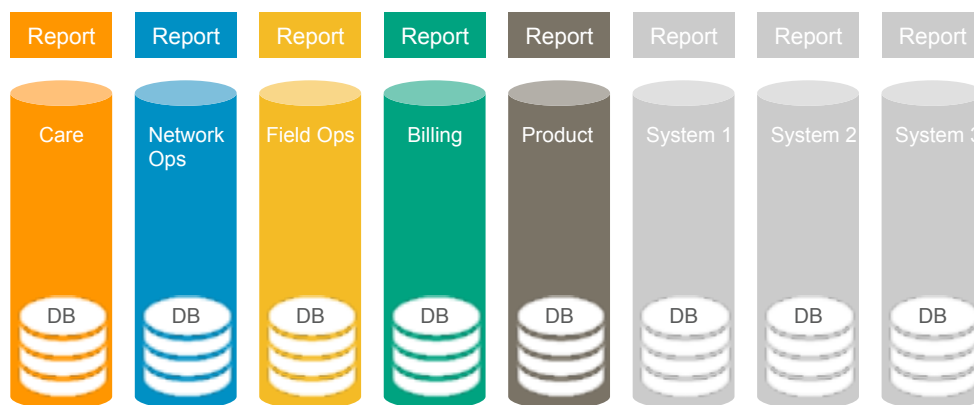


Figure 1: Historical representation of customer information, which was locked in siloed systems

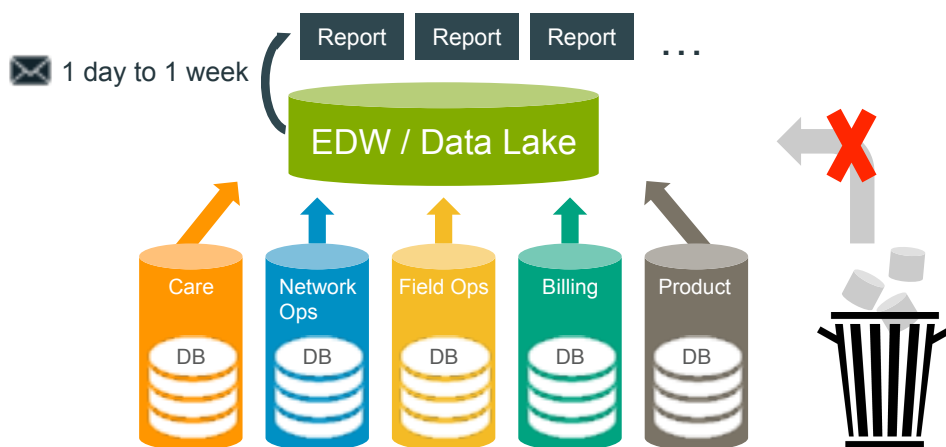


Figure 2: Integrated view of the customer experience

Premises networks pose new challenges as well with operators' growing reliance on Wi-Fi and secondary and tertiary wireless connectivity over ZigBee, Z-Wave and Bluetooth links. And with new means of connectivity on premises comes the need to maintain awareness of a growing volume of devices as well as cloud-based functions that increasingly play a role in how set-tops and gateways operate as the industry moves to virtualization of residential and business CPE.

Higher up in the operational protocol stack, an increasing variety of services and the growing role of IP-centric modes of processing and distributing service payloads to residential and business customers have greatly expanded the number of potential trouble spots to be factored into performance analysis. Not only

do operators need to track the newer IP and virtualization components of headend operations; they need to be aware of what's happening on third-party content delivery networks (CDNs) that can impact the quality of user experience with OTT-delivered content.

All of these challenges are magnified further by the scale of the cable service footprint under management of any given NOC. Consolidation of operations to regional and even national levels has greatly expanded the volume of data relating to all the areas of expanding operational responsibilities that must be brought into a coordinated holistic approach to service maintenance. Moreover, with the shift to core headends linked to automated local master hubs that are replacing traditional headends, centrally located operations personnel must be able to continue maintaining close scrutiny over all the edge functionalities involved in transferring multiple categories of downstream and upstream service flows into and out of multiple access networks.

Defining the Capabilities of a Cable Operational Intelligence Application

1. Comprehensive Use of Data Resources

To meet all these challenges, cable operators need customer aware analytics built on top of an operational intelligence framework that leverage data mining processes that can aggregate useful data from every available source in real time and at massive scale. The OI framework must be equipped with interfaces that support integration with all existing data platforms, including HFC telemetry, change management events, element management system (EMS), network management system (NMS) alarms and care interactions and others outside the operator-owned facilities domain.

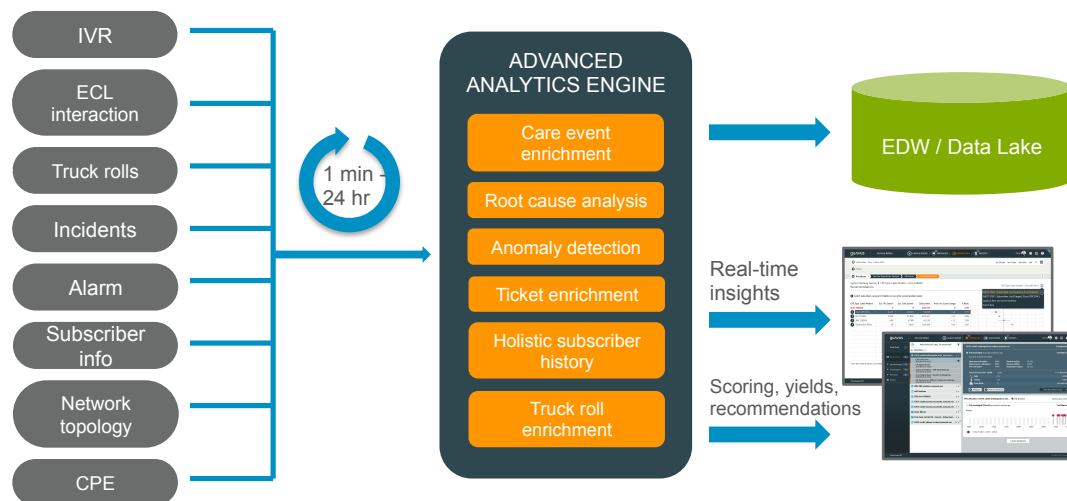


Figure 3: Operational intelligence framework

Such deployments provide voluminous performance and usage data from cable modems, set-tops and other CPE. Employing streaming data analytics, the operational intelligence application can collect, analyze and mediate data on millions of devices, subscribers and network elements. Fusing data which is often siloed, such as calling data, network topology data and truck roll data with CPE firmware information, the platform can apply the analytics processes from the vantage of system-wide visibility into networks, devices, services and subscribers' overall experiences.

With the ability to use all the resources discussed in this section for aggregation of performance data across the end-to-end cable infrastructure, an operational intelligence application can provide operators everything they need to capitalize on the cost-saving and QoE benefits of pre-emptive operational efficiency and reduced response times to outages and systemic issues.

2. Tasks Performed by the Operational Intelligence Framework

Process engines in the Operational Intelligence Framework are essential to achieving these ends and can be categorized as follows:

Large-scale real-time ingest – The magnitude of the collection layer tied to all these data-generating resources is far beyond what is often claimed for cable-oriented OSS systems. It must be capable of collecting massive amounts of network streaming and enterprise stored data that can scale well beyond typical enterprise platforms. A hint of what's involved beyond traditional quality of service (QoS) analytics can be seen in the scope of information associated with the performance of each “chunk” in an Adaptive Bitrate (ABR) streaming session amid communications between servers and devices, often involving just-in-time manifest manipulations tied to advertising, time-shifting and other advanced features.

At any given moment there might be millions of ABR sessions underway, all with a direct impact on customer satisfaction. In a centralized cable operations management environment there are additional millions of simultaneous processes occurring across the infrastructure that must be closely monitored and analyzed to arrive at a comprehensive understanding of events impacting customer experience.

Data enrichment and correlation – In preparation for applying analytics processes all the incoming data must be instantly categorized and correlated as it applies to any use of any service underway at any given moment. The vast dynamically shifting pool of data thus becomes a reservoir of constantly changing compilations of all the events that are germane to what each user is experiencing across the network.

Anomaly detection – With customer aware analytics, operators can establish reference baselines of aggregate system performance from the very local service area level to regional and national levels. Baselines can be configured to reflect everything happening over time or more narrowly defined to track what happens in conjunction with specific operational events. Some examples include CPE reboots with firmware upgrades or the ability to encompass seasonal timeframes such as when people travel to second homes in the winter.

With such baselines in place, operators have a multi-layered reference foundation against which the analytics engine can register any anomalies that may signal emerging or immediate alarm-inducing problems. This anomaly detection can be calibrated to reflect the significance of impacts in terms of statistically significant deviations from the baseline, numbers of subscribers affected, duration, performance of devices newly introduced to a specific service, levels of disruption to subscriber experience or other parameters.

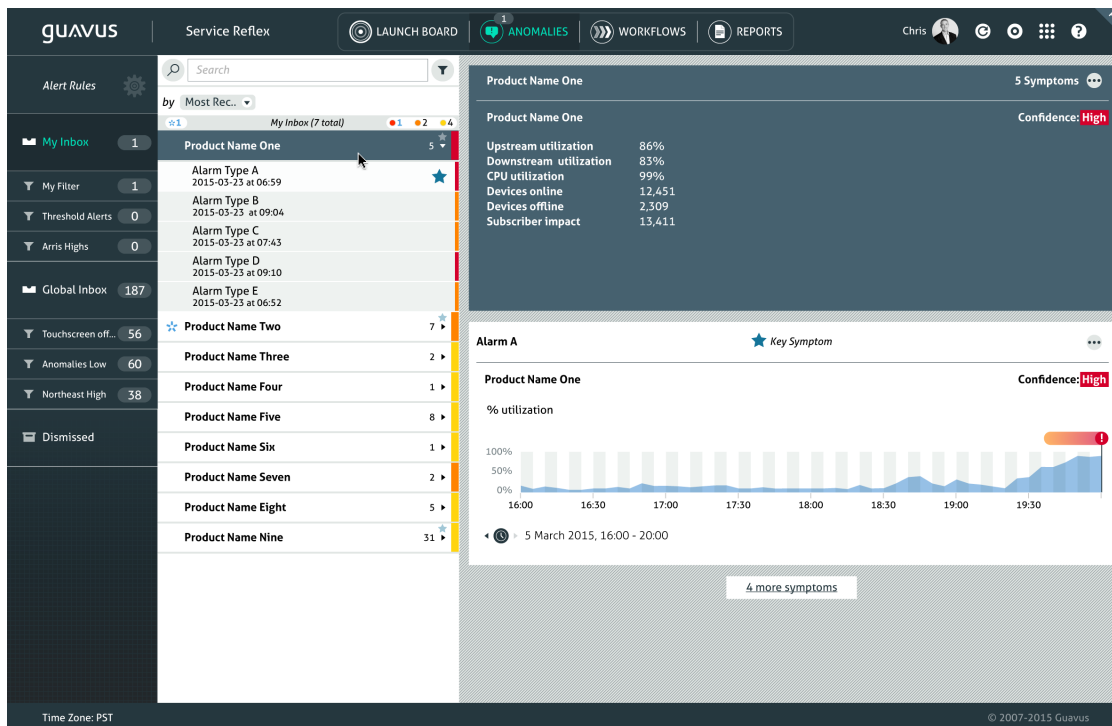


Figure 4: Service anomalies for an example MSO by category

Diagnostics – Leveraging all the information generated through the above processes, the analytics application continually performs analysis of all anomalies categorized as meriting either preventive or alarm-responsive action. By virtue of its system-wide view, the analytics engine can immediately sort through all processes that are engaged in enabling any sub-par or potentially sub-par service experience to identify single or multiple points of contribution to the anomaly.

At the same time, the diagnostics process has to filter out alarms falsely associated with a given anomalous event, either eliminating them from the analysis or re-categorizing them in cases where they apply to some other event.

With potentially tens or hundreds of millions of combinations to be analyzed, the analytics engine is able to generate results and trigger action in accord with whatever policies are set for any given anomalous event. The engine can perform this analysis hierarchically so that, for example, it can identify a clue common to multiple locations, such as the fact that this particular anomaly is occurring wherever a particular brand of cable modem is interacting with certain types of CMTSs.

The analytics applications must be able to integrate machine-to-machine learnings at the click of a mouse showing baseline charts and anomalies, diagnostic findings, results of remedial action and other aggregations of data as required by operators. Operators should be able to dynamically direct use of the platform along with how results are presented to fit their most pressing needs.

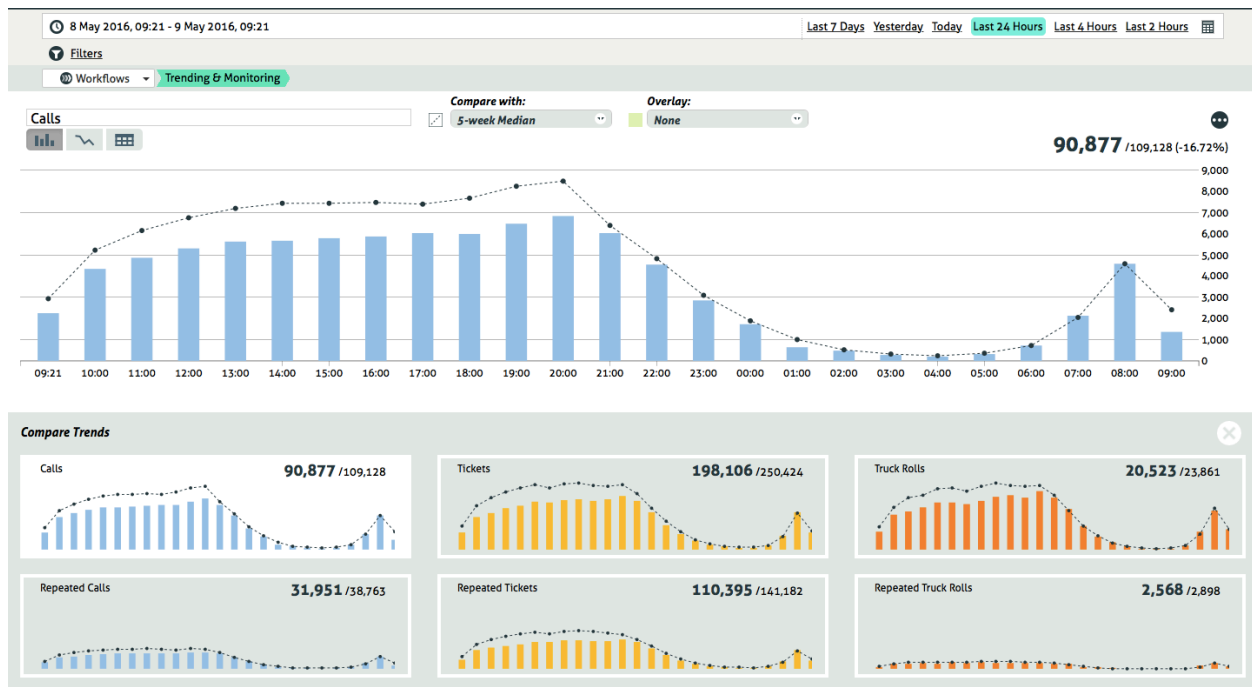


Figure 5: Monitoring the trending service assurance issues for an MSO

A Tier 1 MSO's Use of Operational Analytics

One example of how customer aware analytics solution is being used to help an operator improve customer experience can be seen in the strategy of a Tier 1 MSO in North America that wanted to differentiate between plant and CPE-based issues to help minimize truck rolls, which cost approximately \$80 each time a truck is deployed. In this case, the application has been directed to maintain a real-time view of baselines and anomalies across the MSO's footprint in all its regions of operation.

Rather than relying on customer calls to care centers to establish patterns that would help to narrow the range of possible causes, the operator is able to instantly isolate on the issue before the call volume intensifies. Once any given anomaly crosses a machine-learned threshold (rather than a pre-set threshold) in terms of numbers of people affected, duration of the issue or other parameters, the platform can zoom in on the data pouring in over a selected period of time from the affected area to perform the correlations and analysis essential to identifying the root cause.

The machine-learning algorithm interprets a customer's state to make real-time prescriptions for optimizing her experience. In doing so, it derives associations between the events it is seeing, the relevant history of the associated customer and the knowledge (e.g., care issues) it has accumulated from other customers and systems. Ontologies are dynamically grown and used to manifest the knowledge.

As anomalies in quality of performance on user sessions appear in any part of the service territories, the analytics application automatically looks for patterns associating specific users with specific network elements, for example CMTSs. Using the operator's OSS data that associates users with their CPE, the technology addresses questions such as whether the problem is occurring with only certain types of consumer premise equipment for example, modems, or is randomly spread across multiple modem types accessing a given CMTS or whether the problem is common to a certain brand or model of CMTS in association with certain CPE models. At the same time, the application looks at all the status monitoring data flowing in from the nodes and coaxial plant serving the affected areas to determine whether those elements rather than CMTSs or CPE might be causing the problem.

Avoiding the need for manual intervention, the applications automatically send the analysis to the operator's incident management system to trigger whatever maintenance activities are appropriate. The application also rates the analysis as to level of certainty, so that if there's a high degree of confidence in the diagnosis assignment of remedial tasks can be more automated than if manual intervention is required to direct farther investigation into the root cause.

The operator can also direct the analytics application, using standard APIs, to automatically interact with IVR systems to create messages that will be heard by any affected subscribers who call in to customer service. Additionally, the framework can take actions on specific devices in the home, for example, it can interact with the ACS to push new firmware, reboot the device, reset factory defaults, and change the Wi-Fi channel, etc. The subscriber can be immediately informed that the operator is taking remedial action, what the nature of the problem is and when it should be fixed.

With this highly automated, contextually driven approach to dealing with CMTS and CPE issues in place, the operator is saving immense amounts of money on truck rolls that might have gone out unnecessarily to fix CPE problems that are not a factor. Based on field experiences of this MSO and others, once the operational intelligence framework is in place it turns out that about 30 percent of truck rolls ordered with previous procedures are unnecessary. The total in annual cost savings was \$70 Million in the first year alone.

The analytics application has been integrated with the MSO's field management and other back-office systems so that all unnecessary truck rolls can be automatically canceled without waiting for action by dispatchers. In addition, by adjusting the IVR messaging to affected subscribers, the operator avoids the costs of CSR time that would have been consumed talking to those subscribers.

In another use case, the MSO was able to prioritize and cut a budget for replacing legacy video infrastructure. Initially, the team forecasted a budget of more than \$50 Million dollars to replace everything older than a certain number of years, since it was assumed that equipment older than the threshold would most likely cause problems and failures. However, using customer aware analytics, the MSO was able to cut the budget to approximately \$10 Million by correlating subscribers to each piece of equipment, correlating subscriber impact of each equipment failure and examining failure rates over time. This enabled the MSO to prioritize the 8 different categories of equipment that it should replace, and within each of the categories, prioritize the specific devices that have higher-than-normal impact when

they do fail. This process is now part of the MSO's annual capital expenditures planning process to help the MSO rationalize and prioritize large capital spends.

Analytics applications also assist the MSO's efforts to improve customer satisfaction in other ways, starting with providing information essential to guiding preventive action. In general, the majority of anomaly spikes that the application detects are linked to customer-impacting care issues within micro-segments of the subscriber population (micro-spikes). These previously undetectable issues can comprise as much as 80 percent of the baseline volume. Traditional approaches have mainly targeted the remaining 20 percent of call volume (macro-spikes). In order to truly drive down the baseline, a combined approach of attacking both the micro and macro spikes is required.

The question then becomes, how should preventive maintenance be prioritized based on which issues are most likely to have the greatest impact on customer experience and which of those are closest to crossing that threshold? Here, again, using threshold parameters set by the operator around such vectors as the intensity, frequency and scale of anomaly spikes, the applications can provide guidance for the best use of resources at hand.

In addition, the MSO is able to prioritize its reactions to issues, ensuring that its most valuable customers are the ones that are least affected by service performance issues. The operator is able to set rules that factor such considerations into the matrix of other prioritization parameters to ensure the most efficient and effective use of maintenance resources. By prioritizing preventive maintenance to ensure minimal impact of issues on an ongoing basis, the operator is able to minimize personnel and other maintenance costs while staying ahead of problems that could have costly ramifications. This type of actionable intelligence can predict with 90 percent accuracy what alarms would convert into care incidents.

The MSO has also been able to bring new QoS monitoring data generated by other suppliers into the analytics process. For example, the operator can leverage a premises Wi-Fi system supplier's metrics to get a fuller view of causes behind issues that have been associated with customers' CPE.

Conclusion

MSOs have access to a growing volume of performance-monitoring data streams from every point in the service infrastructure that impacts subscriber experience. Utilizing customer aware analytics applications that provide operational intelligence, they can now apply this data to achieve unprecedented levels of automation and cost-saving efficiency in their efforts to maintain consistently high levels of customer experience across their service territories.

By quickly identifying and addressing root causes to issues that are having a tangible impact on customer experience in any location, operators can lower costs of truck rolls and call center operations while minimizing the impact of performance lapses on customer experience. With a highly automated approach to identifying, prioritizing and proactively addressing emerging problems that have the potential to affect customer experience, operators can manage maintenance operations to maximum affect in their efforts to sustain and drive improvements in customer satisfaction.

Growing numbers of MSOs are putting such capabilities to use, allowing them to centralize real-time processing and analysis of data representing millions of simultaneous service-enabling functions from the

core to every end point in their residential and commercial services markets. As a result, operators are able to continue building new services delivered over multiple access facilities with the flexibility to migrate to new IP-oriented modes of operation in complete confidence that they can do a better job than ever at containing costs and maintaining high levels of customer experience.

Abbreviations

ACS	Auto configuration server
CCAP	Converged cable access platform
CDN	Content delivery network
CMTS	Cable modem termination systems
CPE	Customer premise equipment
CWMP	CPE wide area network management protocol
DHCP	Dynamic host configuration protocol
DOCSIS OSSI	OSS interface
EDW	Enterprise data warehouse
EMS	Element management system
IoT	Internet of things
IPDR	Internet protocol detail records
IPDRdocs	Open-ended streaming of data records
IVR	Interactive voice response
MTTR	Mean-time-to-repair
NAT	Network address translation
NMS	Network management system
NOC	Network operations center
OIA	Operational intelligence applications
OTT	Over-the-top
QoE	Quality of experience
QoS	Quality of service
RDKit	Reference Design Kit
SDV	Switched digital video
XML	Extensible markup language

Bibliography & References

American Customer Satisfaction Index, *Telecom Report*, ACSI Telecom Report Shows Competitive Industries Have Higher Customer Satisfaction, June 2016

J.D. Power & Associates, *Regional Surveys, Improved Network Performance and Reliability Drive Satisfaction with Residential TV and ISP Providers, as More Households Connect Digitally*, September 2015

Leichtman Research Group survey, Cable, DBS, & Telcos: Competing for Customers, March 2016

Parks Associates, press release, *Over 50% of consumers in U.S. broadband households subscribe to an OTT video service*, April 2015

Digital TV Research, Global OTT TV & Video Forecast, June 2015