

Embracing Service Delivery Changes with Machine Learning

Change-Driven Segment Identification and Scoring Can Increase Operator Confidence in Network and Device Modifications and Upgrades

A Technical Paper prepared for SCTE•ISBE by

Andrew Sundelin

Director, Product Management & Cable Innovation

Guavus, Inc.

San Jose, CA

Andrew.Sundelin@guavus.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction	3
1. The Change Imperative	3
1.1. Software, Devices and Combinations.....	3
1.2. Change Initiatives and Concerns.....	4
2. Change-Driven Segmentation.....	4
2.1. Timing Before and After States.....	4
2.2. Segmentation and Attributes	5
3. Change Scoring	6
3.1. Direct Measures.....	6
3.2. Indirect Measures	7
4. Operationalizing Change Metrics	8
4.1. STB Software Upgrades.....	9
4.2. Node Split Use Case	9
4.3. All IP Video	10
Conclusion	10
Abbreviations	10
Bibliography & References	11

List of Figures

Title	Page Number
Figure 1- Attribute-driven Subscriber Identification.....	6

List of Tables

Title	Page Number
Table 1 – Before and After Attribute Sharing.....	5
Table 2 – Sample Data Value Analysis for Customer Experience Metric.....	8
Table 3 – Node-Split Change Events	9
Table 4 – Node-Split Deployment Schedule.....	9

Introduction

Cable operators face significant challenges in launching and improving services with agility and velocity within ever more complex service delivery architectures. The dilemma is clear: Improvement requires change; but change drives performance incidents. So how can operators best test improvements and new services while minimizing unintended consequences?

Evidence-based methodologies, such as A/B testing, that are transforming other industries provide a good answer to the *what* needs to be done. The more challenging question is *how* can operators meaningfully apply these techniques to their operations? Creating a formal evaluation program around every change occurring in MSO access networks would be both labor-intensive and cost-prohibitive, if even possible.

Two techniques that operators can use to improve upon existing ways of tracking and maintaining customer satisfaction before and after changes in their network are the following:

Change-Driven Segment Identification. This is an automated technique for identifying populations and micro-populations of subscribers within the operator network whose service, CPE or service delivery network has experienced the same change.

Change Scoring. Once groups with like changes are identified, the quality of a given upgrade needs to be evaluated. There are many ways to score changes, from straight customer experience measures to automatic feature selection models. Big data and machine learning provide operators with options for systemically evaluating potentially service-impacting change.

The goal of this paper is to describe ways that operators can use machine learning in combination with well-designed operational practices to quickly differentiate between upgrades that improve service quality and those that make it worse – thereby allowing operators to embrace change, not fear it.

1. The Change Imperative

1.1. Software, Devices and Combinations

Not long ago, industry leaders talked about handling one or maybe two major initiatives in a year. While there are still major capital expenditures and large-scale projects, overall change has become more incremental and rapid, less exclusively waterfall and more agile, to use the terms that originated with software developers. That language is appropriate, because of all the modifications occurring in operator networks, software changes are the most common.

In the home, software upgrades and planned improvements can occur to the CM, the STB, the MTA, the home security keypad – even the video camera or other IoT element. In the network, software changes also occur with growing frequency. The CMTS, the CCAP (either CCAP core or even the RPD), the DPoE system, the video controller can all have software changes that potentially impact service.

Combinations of CPE and network changes can also occur with a number of potential permutations that make it impossible to fully validate prior to production deployment. Furthermore, capacity-driven node splits, fiber-deep initiatives, redistribution of CPE across access equipment (or subscribers across shared video controllers) are also common changes to the actual network/service topology, all of which can negatively impact service performance and customer satisfaction.

1.2. Change Initiatives and Concerns

Change creates anxiety in many industries. In a 2015 study of some 300 clients in the area of IT Operations Analytics, Gartner reported that 85 percent of performance incidents are traceable to changes. [Cappelli] The quicker cadence and more expansive scope of change in the cable industry is reflected in initiatives and operations involving several areas of technology, including:

Set-top boxes. Legacy devices remain in the field, but the market is no longer characterized by stodgy single-purpose devices with quarterly software release cycles. The RDK, or “platform behind the platform,” now in more than 40 million set-tops and gateways around the world, is a good illustration of how fast-paced, agile techniques are transforming the industry. Last year an industry executive stated that the RDK open-source consortium was shipping more than new 70 features per month. [Ismail]

DOCSIS infrastructure. As DOCSIS and WiFi networking standards evolve, bringing greater speed and new architectures, both CPE and network hardware will continue to change rapidly. On the CPE side, new firmware for traditional devices such as CMs and MTAs is more frequent, and the types of IoT devices in the customer premises (most with remotely upgradeable software) are exploding. On the other side, access network devices are also a constant work in progress. Traditional CMTSs, CCAPs and DPoE systems all have software upgrades; and with distributed access architectures (DAAs), decomposed CCAPs are likely to increase this rate of change as the industry undergoes a significant transformation.

Fiber optic nodes. Changes in fiber nodes can have many potential motivations. An operator could be switching to a DAA architecture; “splitting a node” for capacity reasons; or transitioning to an n+0 architecture and going “fiber deep.” Comcast launched a major fiber-deep initiative two years ago, and last year one executive said that after this multi-year ramp-up is complete, the MSO could end up with eight to ten times as many nodes. [Breznick]

Video delivery. No more service may be undergoing more change and more difficult to manage than video. Adaptive bit-rate (ABR) technologies help preserve IP video quality, but an operator has little visibility into those streams. Being one-way broadcast, QAM-based video has no feedback, apart from subscriber complaints or telemetry from expensive probes, making it a necessary to find predictive metrics elsewhere. Changing architectures, such as cloud-based guides, can add new components to the service delivery architecture. Changes to many video delivery components can impact service in ways that can be difficult to detect or to test in the lab (due to both scalability and permutation challenges).

Apart from those four areas, there are other network elements that can also impact service delivery. The common thread connecting many of the techniques described in this paper is the ability to tie these devices (and, thus, the attributes of these devices, such as manufacturer, model and software version) to individual subscribers. That is the key to creating an anxiety-reducing change management system that is not only dynamic and extensible, but also service-centric and customer-centric.

2. Change-Driven Segmentation

2.1. Timing Before and After States

The idea behind segmentation of operational data is to identify groups of subscribers with identical service changes so as to best assess and compare differences across groups or across time (e.g. before/after a change).

Change-driven segmentation has a number of fundamental concepts. A change-driven segment has a common “before” state (i.e. the value of a subscriber attribute before a change), but possibly different

“after” states. Change occurs in different portions of the network at different times and at different scales. Each change-driven segment has an associated time period for that change, typically, a day.

For example, 1 million subscribers might have a Cisco DPC3010 running software version 1.2.3. The operator may be migrating those subscribers from version 1.2.3 to version 1.2.4. The operator may upgrade 100,000 subscribers on January 1; then 250,000 subs on January 8; and finally, 650,000 subs on January 15. However, these 1 million subs, since they share a “before” and “after” state would be in the same “before” segment and the same “after” segment, regardless of the timing of this change for any individual subscriber in that group.

Alternatively, an operator might have the scenario seen in Table 1. Here some subscribers share the before state (e.g. pre-node split) and some share the same after state (e.g. post-node split or merge). But the change-driven segment is uniquely identified by a single attribute sharing the same before and after attribute value regardless of the timing.

Table 1 – Before and After Attribute Sharing

Subscribers	Before Attribute	After Attribute	Change Date	Change-Driven Segment #
100	Node = ABC	Node = ABC-1	August 1st	1
100	Node = ABC	Node = ABC-2	August 1st	2
100	Node = ABC	Node = ABC-3	August 1st	3
100	Node = ABC	Node = ABC-4	August 1st	4
23	Node = An34q	Node = ABC-4	August 12 th	5

2.2. Segmentation and Attributes

This type of segmentation can be driven by any number of service attributes:

CPE Attributes: One segment of subscribers might be those with the same make and model of CM or STB, being upgraded from software version X to software version X+1. In this case, the software version is the service attribute which is changing and driving their membership in that group.

Topological Attributes: Another segment of subscribers might be created when a service group is being divided by the splitting of a particular node. Or when a new CMTS/CCAP is deployed and subscribers are migrated to that CMTS. Here it is a common change in a subscriber’s network topology that is changing and driving their membership in that group.

Billing/Service Attributes: During service changes (e.g. going from QAM-based video delivery to IP-based video delivery) an operator might change a subscriber’s rate code. Or the rate limits on a particular HSD speed tier might get changes. Again, these common changes in subscriber attributes – whether they happened simultaneously or not – allow for change-driven segmentation of subscribers with “like” changes from a “before” state to an “after” state.

Today these selections are today happening on an ad-hoc, manual basis, and so an exercise such as examining a fiber-deep initiative with speed and precision is difficult. Software changes are also happening to more devices with greater frequency. Thus, programmatic change evaluation programs are labor-intensive and expensive. The amount of change in the network makes them all the more difficult.

A lack of automation in identifying groups who have experienced change is one reason to fear rather than embrace improvement initiatives. Change-driven segment identification that leverages big-data analytics provides a way to automatically identify subscribers who have experienced change. (See Figure 1.)

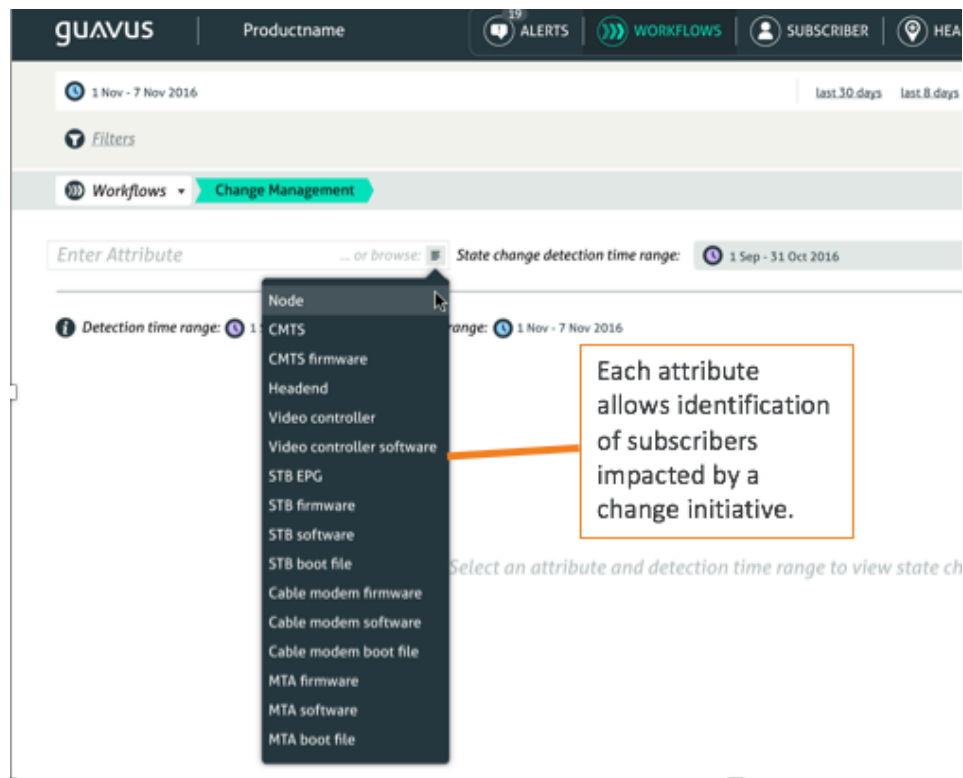


Figure 1- Attribute-driven Subscriber Identification

3. Change Scoring

Once operators can automatically identify groups of subscribers with identical changes, the next step is to evaluate those changes and determine which changes improved customer experience, and which changes made it worse.

3.1. Direct Measures

One way to score an upgrade is simply to look at direct measures of customer experience. To assess a change initiative or other operation, an operator would assemble data before and after to determine whether it yielded a positive or negative result. Every subscriber on the network has a set of attributes. As shown above, those can be device, service, topology, billing or other per-subscriber attributes.

It is common for operators to track all of the care events (e.g. technical support calls, customer-reported trouble tickets and scheduled truck rolls) from individual subscribers. This becomes even more powerful when operators leverage a big-data system to not just track this per subscriber, but also per subscriber attribute. Then operators can know how many subscribers with a given CPE are calling, how many associated with a given CMTS, how many with a specific rate code, and so on.

Thus, the simplest way to determine whether a change leads to a better or worse customer experience is to take a segment of subscribers identified by automatic, change-driven segmentation and look at the rate (by population) of direct measures of customer experience.

3.2. Indirect Measures

But there's another way to score an upgrade. Care events as a direct measure of customer experience are not ideal, because they indicate that a group of customers has already experienced enough pain to take the trouble to contact their operator. There are many other indirect measures of customer experience. These can be service-layer data (e.g. STB errors/reconnects, speed tester data, video probe data, CM reboots, ABR metrics, buffer size, etc.) or network-layer data (e.g. correctable codewords, uncorrectable codewords, RX power, TX power, dropped packets, etc.) None of these measures involve a human picking up a phone and calling the operator, but many of them can be predictive of that behavior.

There is a key challenge to leveraging these indirect measures of customer experience. At some level they do not negatively impact customer experience, but then at some point, the customer experience becomes so bad that customers begin to call, get tickets and instigate truck rolls. How can we determine when these various indirect measures (or features) start to impact customer experience?

The traditional way to determine at what level various network performance measures are contributing to negative customer experience is to convene a group of internal and external multi-discipline subject matter experts to perform a study. They then monitor any number of indirect measures of customer experience and develop a set of thresholds, which, when exceeded, lead to bad customer experience. However, these experts are expensive, and historically, the data has been difficult to accumulate for a holistic view of the entire network.

Perhaps more importantly, these studies generally have limitations. Because they are expensive, they often happen over a limited portion of the network and over a limited period of time, which means that the resulting thresholds may not be universally accurate across the operator's network and will likely become inaccurate as they age and customer behavior, expectations and applications all change.

A better way to determine when these various indirect measures/features start to impact customer experience is through machine learning. This is a two-stage process.

The training stage begins when the machine-learning algorithm discovers the relationship between features and target variable(s), such as care events. Note that not all attributes are created equally; CPE/UE, subscriber health and network elements tend to be more highly predictive. (See Table 2). At the end of training comes validation, or confirmation through another set of feature data, not included in the training data, that the model is truly predictive of the target variable.

Table 2 – Sample Data Value Analysis for Customer Experience Metric

Attribute	Value	Notes
CPE/UE	High	Actionable, predictive
Plant/telemetry health	High	Directional, actionable
Usage patterns	Med	Noisy, requires proper handling
Subscriber health	High	Directional, actionable, predictive
Device characteristics	High	Actionable
Service tiers	Low	Noisy
Tenure of CPE/UE	Med	Temporally directional
Installation type	Med	Temporally directional
Network elements	High	Actionable, predictive
Plant capacity	Low	Varies based on service type

In the second stage of machine learning-driven analytics, live data is put into the model. The model then outputs a score. That value indicates how predictive of the target (e.g., care events) the input data set is.

For example, in the case of a node split, a machine-learning model could be developed where the features are indirect measures of customer experience (such as, correctable codewords, uncorrectable codewords, RX power, TX power and SNR) while the target variable is technical support calls. By monitoring a time series of feature values before and after the change, the probability of technical support calls can be predicted (before and after the change). If the probability of support calls goes up, then this change made things worse for customers.

The machine-learning approach is better than a one-time study because the applications that subscribers use change fairly quickly over time, and different applications have different sensitivities to different types of impairments. Machine learning models can be periodically retrained to adjust to these application and behavior changes. This approach is called continuous learning.

To recap this discussion of change scoring, the obvious way to evaluate the user experience before and after a change is to look at direct measures of customer experience and to see if the changed attribute makes the experience of subscribers with that shared change better or worse. The problem with this approach is that it depends upon customer impact and input.

The second way, involving indirect measures of customer experience, can be driven by machine learning. Through a model that predicts the likelihood of poor customer experience, machine learning can output scoring on a time series of indirect measures before and after the change, creating scores that can be used to evaluate whether the change is likely to make the customer experience better or worse.

4. Operationalizing Change Metrics

The vision of this paper is that systems be developed to (a) automatically identify changes to a wide variety of per subscriber attributes; and (b) score changes so as to readily identify those that worsen customer experience or are predicted to do so.

The scoring aspect of the system could be a daily report that lists all of the change-driven segments identified in the network sorted by the difference in the score for that segment before/after the change.

(See Figure 2.) Thus, a team could easily identify bad changes to the service or network in order to more quickly roll them back.

Table 3 – Node-Split Change Events

<input type="checkbox"/>	Oldest Change	Newest Change	Changed Subs	Old State	New State	Impact (Customer Experience Score)
<input checked="" type="checkbox"/>	9-Oct	9-Oct	10	AAA	BAAA1	+10.5
<input checked="" type="checkbox"/>	9-Oct	9-Oct	10	AAA	BAAA2	0
<input checked="" type="checkbox"/>	9-Oct	9-Oct	9	AAA	BAAA3	-2.5
<input type="checkbox"/>	7-Sep	7-Sep	46	PDQ	PDQ1	-3.7
<input type="checkbox"/>	12-Oct	12-Oct	23	XYZ	XYZ1	0
<input type="checkbox"/>	12-Oct	12-Oct	16	XYZ	XYZ2	-1

4.1. STB Software Upgrades

At some operators, this basic model is being followed for modern STB software upgrades already. However, this usually requires a complex program of phased upgrades where larger and larger groups of subscribers receive the same upgrade over time. It also requires an equally complex method of evaluation, where a priori knowledge of the groups being upgraded at specific times is used to create a group to be monitored. This entails tracking which devices are being upgraded when and where, as well as specific monitoring to see if the upgrade is good or has unintended consequences.

Operationalizing this approach also permits operators to apply the web technique of A/B testing of alternative software upgrades, allowing them to readily determine which of two implementation options is better. Thus, change analytics can increase agility, enabling a business to “lean in” to quick learnings.

To summarize, phased software upgrades remain a best practice, but a continuous and comprehensive approach to evaluating changes across many different attributes removes the need to specifically track when and on what specific devices these phased changes are being made. Plus, it brings these benefits to other types of changes, beyond software upgrades.

4.2. Node Split Use Case

The industry’s hunger for fiber in the access plant is growing, yet with no common mechanism today to automatically identify “bad” node splits, there is no quick and easy way to know how well these initiatives are going. A pace of 25,000 or 50,000 per year leads to a nightly deployment rate that is too numerous to track manually. (See Table 2.)

Table 4 – Node-Split Deployment Schedule

Per year	Per week	Per night
25,000	481	120
50,000	962	240

Suppose there are bad splits. A scenario requiring an additional truck roll for each split could add \$3 million in the case of 50,000 deployments per year. (\$3m/year = 50,000 x 1 x \$62.) If only one in six is

bad, that’s another \$516,000. Adding in service calls that might follow a defective node upgrade could easily and significantly increase that amount.

Given the high stakes, operators need to identify bad nodes as quickly as possible. Without scoring, however, there is no baseline for subscribers following a split, and with no baseline, you’re looking at up to three weeks to begin detecting anomalies. No scoring mechanism also means no before/after metrics, or some way of understanding how customer experience on the original nodes differs, if at all, from customer experience on the new or fiber-deep nodes.

4.3. All IP Video

Many changes to the network are designed to escape subscriber notice. They may not even be visible when looking at subscriber-related attributes such as STB make/model/software version.

If a subscriber has a STB which is capable of both IP- and QAM-based video, then device-oriented attributes will not allow an operator to differentiate between subscribers’ being delivered all-IP video and those being delivered primarily QAM-based video. However, they may be visible when looking at the billing system, or other subscriber attributes. Some larger operators are using a different rate code for their subscribers on all-IP video, and this is one way they can be identified on a per-subscriber basis.

Thus, operators can leverage Change-Driven Segmentation to identify subscriber populations with like changes to a vast array of subscriber-related attributes. This can either be part of a formal phased-roll out process or more ad hoc, but the technique works equally well in either situation – even when the like changes happen at different times. Once these subscriber populations are automatically identified their before/after experience can be scored (Change Scoring) such that changes which make customer experience worse can be quickly backed out or otherwise mitigated.

Conclusion

Network and device upgrades and other change initiatives cause concern among technical operations and customer care personnel. This is the case among service providers and tech companies of all stripes, but it is especially so in cable as the pace of change and complexity of the industry’s service delivery platforms increases. But evaluating every change occurring in MSO access networks using traditional means is a non-starter. The costs are too high, resources too limited and workload overwhelming.

An attractive, efficient and cost-effective alternative is to leverage change-driven segmentation techniques and machine learning-based change scoring. The benefits include a reduction of uncertainty and anxiety surrounding change; greater clarity about which changes promote and which detract from customer experience; and advancement of the long-sought goal of transforming a network-centric view of operations into one focused more on services and the subscriber.

Abbreviations

ABR	adaptive bit rate
CCAP	converged cable access platform
CM	cable modem
CMTS	cable modem termination system
CPE	customer premises equipment

DOCSIS	DOCSIS over Cable Service Interface Specification
DAA	distributed access architecture
DPoE	DOCSIS Provisioning over Ethernet
HSD	high-speed data
IP	Internet protocol
IoT	internet of things
MSO	multiple systems operator
MTA	multimedia terminal adapter
QAM	quadrature amplitude modulation
RDK	Reference Design Kit
RPD	Remote PHY Device
RX	receive level
SNR	signal-to-noise ratio
STB	set-top box
TX	transmit level
UE	user experience

Bibliography & References

Breznick, A. Why Cable’s Feasting on Fiber. Light Reading, Nov. 9. 2017.

Cappelli, W. Causal Analysis Makes Availability and Performance Data Actionable. Gartner, Oct. 7, 2015.

Cheikhrouhou, A., and J. Davenport, A. Kelkar. The Imperative of Customer-Centric Operations. SCTE/ISBE, 2017.

Cunha, G. Approaches for Proactively Managing Customer Experience and Reducing OPEX in a Cable Operations Environment. SCTE-ISBE 2017.

Dorairaj, S, and C. Basian, B. Burg, N. Pinkernell. Simplifying Field Operations Using Machine Learning., SCTE/ISBE, 2017.

Ismail, L. RDK Product Roadmap Update, Inside RDK. Summary of session offered at European RDK Summit, Sept 17, 2017.

Sundaresan, K. and N. Metts, G. White, A. Cabellos-Aparicio. Applications of ML in Cable Access Networks, NCTA, 2016.

Sundaresan, K., and J. Zhu. Access Network Data Analytics. SCTE/ISBE, 2017.

Sundelin, A. Leveraging Machine Intelligence and Operational Analytics to Assure Virtualized Networks and Services, SCTE/ISBE, 2017.