Operational Focus on Network Intelligence and

Neighborhood Health for Advanced Service Assurance

Presented at
SCTE Cable-Tec Expo$^{®}$ 2009
Denver, Colorado

By

Richard Berthold, CTO
Proxilliant Systems Corp.
One Clock Tower Place, Suite 350
Maynard, MA 01754
978.823.8100 phone / 978.823.8107 fax
Richard.Berthold@Proxilliant.com

and

Keith Hayes, VP-Network Operations and Engineering Services
Charter Communications
6399 S. Fiddler's Green Circle
Greenwood Village, CO 80111
303-323-1347 phone / 303-323-1319 fax
Keith.Hayes@CharterCom.com

Introduction and Overview

Cable operators continue to wage the battle of maintaining service quality and bandwidth while offering a myriad of technologies and advanced two-way services. Some of these technologies and services are well managed, well behaved and predictable, but others are loosely managed, dynamic, complex, and extremely hard to plan for. As services have evolved in their sophistication and unpredictability, managing and assuring the quality of the networks and services require a new level of intelligence – not just passive monitoring – but built directly *into* the network. This new intelligence must span both network intelligence and service intelligence, crossing the lines of RF and IP, to tell us how the network is performing and being utilized as well as how each of the disparate services is performing over that network.

The recent article “Redefining MTTR” in *Communications Technology Magazine* (July 2009) illustrates a new, more accurate understanding of a nine-step MTTR and the comprehensive costs associated with assuring service quality and bandwidth in an increasingly complicated operations environment. After articulating the four categories of intelligence gaps that bloat MTTR, the article launched a discussion about how intelligence is beginning to be built into the cable network to bridge those gaps and dramatically improve quality and costs.

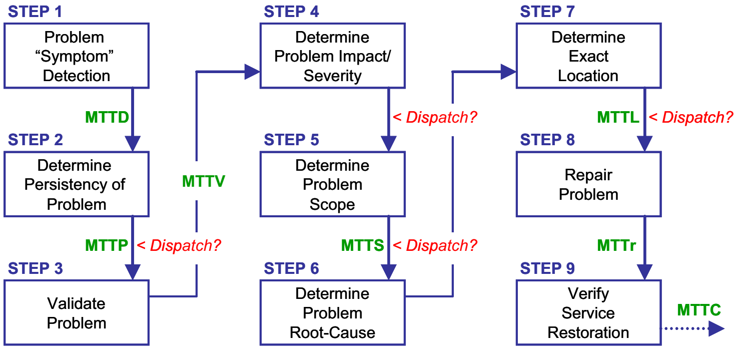

Key Service Assurance Operations Process Steps

A new equation to better understand and discuss MTTR (Mean-Time-To-Repair) is defined as:

MTTR = MTTD+ MTTP+ MTTV+ MTTS+ MTTL+ MTTr (+MTTC)

Where Mean-Time-To-Detect (MTTD) is defined in the article as the time to detect that a problem exists (from the actual initial onset of a problem). Mean-Time-To-Persist (MTTP) is defined as the persistency time to determine if the problem is persistent or just a “glitch”. Mean-Time-To-Validate (MTTV) is defined as the time spent to validate that the problem is real rather than a false alarm or symptom of a different problem. Mean-Time-To-Severity/Scope (MTTS) is then defined as the time spent determining the severity and impact of the problem, such as what and how customers and services are affected. Next, Mean-Time-To-Cause (and location) or MTTL is defined as the time to determine the root cause and location of the problem. And finally MTTr (little ‘r’) represents the actual repair time. An optional additional component is Mean-Time-To-Close (MTTC) which is the time to verify service restoration and close out any ticket or notification.

As a result, the article also defines a subsequent MCTR (Mean-Cost-To-Repair) equation, built to better understand and discuss the *cost* of repair (or service maintainability). This is defined as:

MCTR = MCTD+ MCTP+ MCTV+ MCTS+ MCTL+ MCTr (+MCTC)

Where the “T” for time is replaced with a “C” for the cost spent on each sub-task.

With the newly defined MTTR as a conceptual foundation, this paper will investigate the next-generation access network architecture – a network with intelligence built into the service delivery infrastructure – to illustrate how cable operators can cost effectively wage this battle of maintaining service quality and bandwidth while offering a myriad of technologies and services. In addition, the paper will discuss how this network intelligence provides the keys to operational change in the focus and manner by which advanced broadband services are managed. It will illustrate a deeper focus of operations, down to the neighborhood (or last active) level, to both deliver consistently reliable services and reduce MTTR and costs in today’s complex environment.

Business Objectives for Next-Generation, Intelligent Network

In the fourth quarter of 2008, Charter launched the initial strategic deployment, integration and evaluation of a new intelligence-based system with an operations focus on the neighborhood in its Georgia market. Charter's business objectives for the initial deployment of a service health management system centered on three major areas.

Supporting the Next Wave of Differentiated Service Offerings

Charter continues its focus on providing new, differentiated services within competitive service packages to its customers. In 2008, Charter launched its most aggressive HDTV (high-definition television) offering ever, with channel counts up to 40. The HD offering was made available to all Key Market Areas (KMAs) throughout Charter's network.

This HDTV launch follows the successful roll-outs of VoIP (voice over Internet protocol) and VOD (video on demand) services in recent years. By the end of 2008, Charter had deployed VoIP services to more than 1 million subscribers, achieving double-digit penetration by aggressively promoting its triple-play service bundle. Charter's VOD is deployed in every major market, with more than 8 million streams played each week.

The current economic climate notwithstanding, Charter is focused on the next wave of technological advances to ensure a robust network platform for the continued delivery of additional differentiated services, allowing the operator to maintain and extend its market leadership in the face of competition from Verizon and others. Two technologies are especially important in fueling Charter's next phase of service offerings and customer growth – SDV (switched digital video) and DOCSIS 3.0. SDV delivers expanded video programming capacity and DOCSIS 3.0 expands two-way data capacity, making both attractive investments. In combination they make possible dynamic video switching and dynamic high-speed two-way bandwidth delivery, representing an even more compelling entry into a new generation of interactive services to customers. Ensuring the success of these two enabling technologies is important.

Delivering the Highest Service Reliability

Charter is committed to providing the highest service reliability with the highest service availability for its customers. Along with the advances in services and the technologies that enable them, Charter is highly focused on driving down service calls, improving network availability, and increasing service performance across all services. For data services, outside of maximum bandwidth flow, this means nearing 100% network availability, representing anytime rapid file/video loads and uninterrupted streams. For video, this means reduced video tiling or other artifacts from video broadcasts, and reduced unavailable, interrupted, or lost VOD sessions. And for VoIP, this means zero dropped calls, minimized echo/jittery voice, and the delivery of 911-life-line availability.

As services have evolved in their sophistication and unpredictability and the myriad of enabling technologies has accumulated layer upon layer, managing and assuring the reliability of networks and services now require an approach that is proactive and automated. Charter is aggressive in exploring new, more effective ways of building intelligence directly into the network to ensure the highest service reliability.

Containing Costs and Driving Profits amidst Significant Growth

While the first two objectives are key drivers for revenue generation, maintaining operations costs while achieving them is the challenge for Charter – and all cable operators. Simply speaking, Charter's business equation reads like this:

- Deploy enhanced service offerings AND
- Deploy competitively priced product bundles AND
- Grow customer base AND
- Increase service reliability AND
- Contain or reduce operations cost

So, the question is how to make this equation work with a relatively stable number of technical operations staff. With increasingly complex services, the challenge is clear, particularly when, before you can dispatch or simply detect a perceived problem currently, service has already severely degraded, or worse yet, gone off-line. In this economic climate, progressive operators like Charter are taking the challenge one step further. They are envisioning operating plans that accommodate significant growth without increasing support resources – simply using advanced network intelligence to correctly focus precious human intelligence. As such, Charter is among the leading operators who are exploring innovative approaches that break out of existing paradigms to reinvent technical operations.

Key Obstacles in Cutting MTTR and MCTR

Many obstacles exist in moving toward the goal of cost-effective, continuously managed service and network health, where MTTR approaches MTTr. These challenges fall into four categories:

1. Lack Of Continuous Behavior Awareness
2. Lack Of Data Synchronicity
3. Lack Of Network Topological Relationships
4. Lack Of Automated Compensation or Containment Systems

1. Lack of Continuous Behavior Awareness

The lack of continuous behavior awareness is represented as gaps in data and time. Only some health, status, and service performance information is available at any given time. These gaps may reflect information not being available from all systems simultaneously or continuously. Unless we could poll all DOCSIS data from every CM all the time, or make test calls and video measurements (at the CPE) on a continuous basis, which is not possible (or practical), there remain gaps to the complete set of information and the complete timeline over which they occur. In addition, all customer transactions (use and experiences) are not reported through customer calls or service performance records. Instead, we get snapshots in time from a subset of locations, and we *infer* continuity.

Unfortunately, today's sophisticated cable network and the dynamic services that utilize it are anything but continuous, predicable or of uniform scope. While there are general trends around peak usage and "quiet" (noise-wise) times of the day, service-affecting problems of varying severity and scope can occur at any time, last as short as two minutes or as long as two days, effect 1000 customer or 0, and cause anything from a blip on the TV to dropping a 911 call.

As an example, the plot below shows a time-based graph of high-intensity impulse noise at a given end amplifier leg location. This is a result of continuous monitoring of the number of noise events from an intelligent device physically at that location. In the example graph below, ingress/noise is depicted as both "short" duration and broad (commonly CPD). It is interesting to note, however, that the detection (and impact) of this noise lasts for 1 to 2 hours and only twice per day. If a system or operations person were to sample the noise effects every four hours, on the hour as depicted by the dashed lines, there would be no noise detected. The plant would be deemed and recorded in reports as "clean" (noise free).

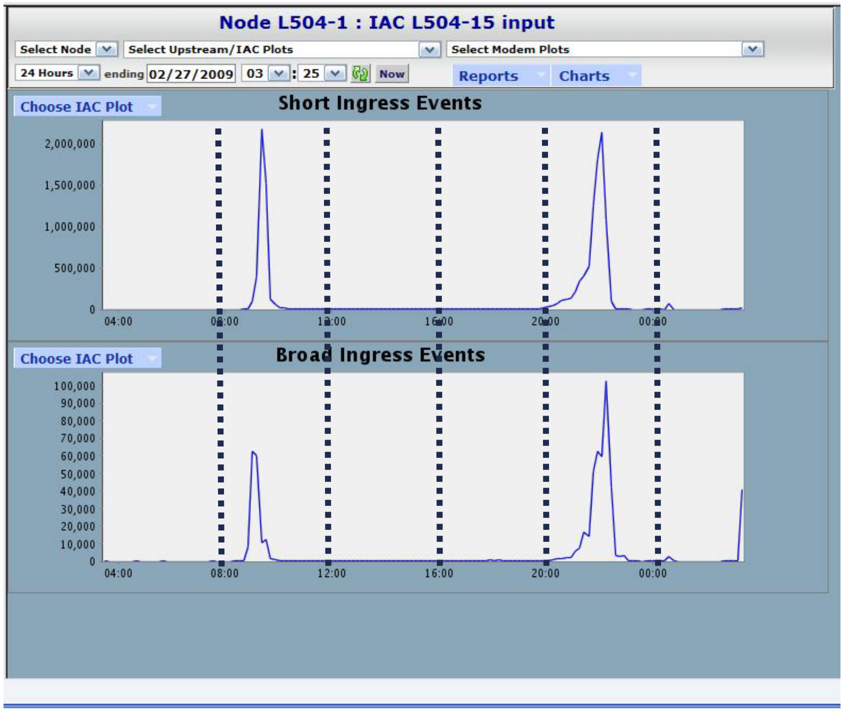

Intermittent Impulse Noise at End Amplifier Location

Without continuous behavior awareness, we must infer continuity, which is often, as with this example, a false assumption. Can we really continuously monitor all data, status, devices, and systems continuously? The answer, of course, is no. And, even if we could, it would take so much time and computer power to analyze the results that the problem (symptoms) may have gone away or otherwise be undetectable (and therefore in most cases un-actionable).

The key to solving this dilemma is that we need *some* continuity in our behavior awareness, at the least cost/impact in performance, to be used as our "trigger" to investigate further and deeper. This approach will allow us to dispatch an automated, highly focused health determination system, which can be complete, continuous, and comprehensive in a limited focus area.

2. Lack of Data Synchronicity

The lack of data synchronicity is a very subtle but critical obstacle. A comprehensive health management system is composed of information from different sources, of different formats, and unfortunately, at different sampling rates. Making a diagnosis of a problem relies on accurate correlation of *time synchronous* information. At the doctor's office, we may be asked whether the chest pain was triggered by physical exertion or after eating a bowl of chili. You certainly would not want to choose open heart surgery when you should have simply taken a Tums.

In today's cable plant health management, lack of data synchronicity is most prominent across RF, IP and QoE layers. Several independent monitoring systems may provide snapshots of each. One system may gain transmit power and SNR parameters and IP performance from cable modems at one rate. Another system may measure noise levels (intensity and frequency) at the headend. Yet another system may supply feedback on power supplies, reporting battery status and other parameters. Video and telephone logs may include previous 24 or 48 hours worth of data. Customer calls and ad-hoc field tests again provide another source of unsynchronized (but potentially coincidental) performance and health measurements.

Rather, we need to answer at what exact instance did we detect high-power ingress isolated in one plant segment that within 2 minutes manifested itself in modem drop-off, high SNR drop, or excessive packet loss from *only* the devices in that plant segment. The best we can do without synchronized data is to make a *general* assessment of health. With this approach, the one who sneezes the most is most likely to have the cold.

As an example, using signal level stability, in the plot below we see a results of monitoring the cable modems upstream transmit power from all cable modems (and only those) common to an end amplifier. In other words, these represent only those in the same "RF domain". This plot also illustrates the previous point that behavior cannot be assumed to be continuous. As we see, if we sampled the system at various times every fours hours (out of skew), as indicated by the vertical dashed lines, we would assume that the upstream transmit power is fairly stable and varies by no more than 2 dB over a 24-hour period. Everything (falsely) looks pretty good.

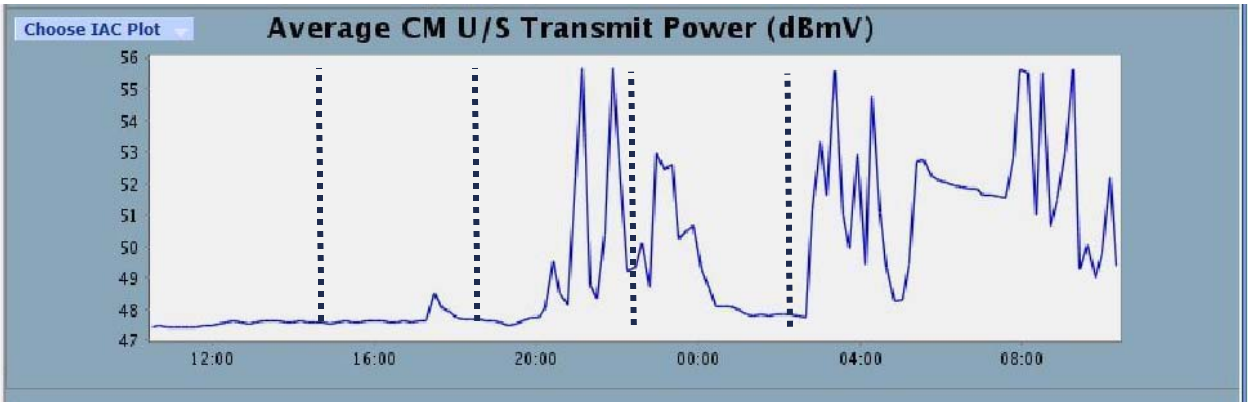

In the plot below we see results from monitoring an independent upstream pilot carrier that is expected to around +2 at the headend. Given a margin of ±3 for normal "breathing" of the plant, our thresholds of acceptability are set at -1 to +5 (indicated by the red horizontal lines). However, if this pilot power monitoring system was sampled every four hours (out of skew) as indicated by the vertical dashed lines, the power levels would be deemed in range (and not worth any future, deeper analysis).

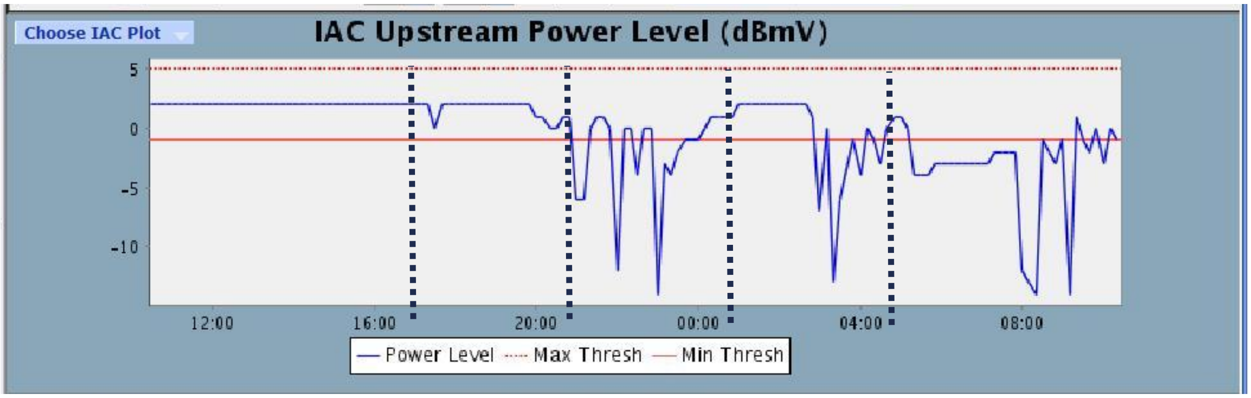

Now, if you took the every-four-hour sample data from the modems and the every four hour sample data from the pilot power, plotted them together (and assumed continuity), you would infer that two these two independent systems had their independent results correlated to an end result of an undeniable conclusion that the signal levels are stable. Now, let's say the modem's transmit power monitoring system's sample rates were as indicated, but the other was skewed such that the pilot carrier was frequently detected out of range. Now you would conclude that there is an attenuation affecting the upstream pilot carrier (e.g. at an upstream amplifier from the modem), but not the modems. How do we explain this? Perhaps the pilot monitoring system had a failure or reported false data. This illustrates that independent monitoring of a single system, as well as correlation of non-synchronized, non-continuous data from multiple sources leads to false conclusions.

On the contrary, if one continuous monitoring system (e.g., the pilot) was used to trigger and synch the monitoring of a second measuring system (modem's transmit) and then compared in the exact same time domain, then in fact a very different conclusion could be made. As the synchronized data plot indicates below, a repeated, chronic change (attenuation) of the upstream pilot, located at an upstream amplifier from the modems, correlates precisely with the change in average transmit power of every modem attached to that amplifier. In the example graph, this is shown to be a swing of over 7 db several times a day, a great cause for concern and a worthwhile investigation for repair.

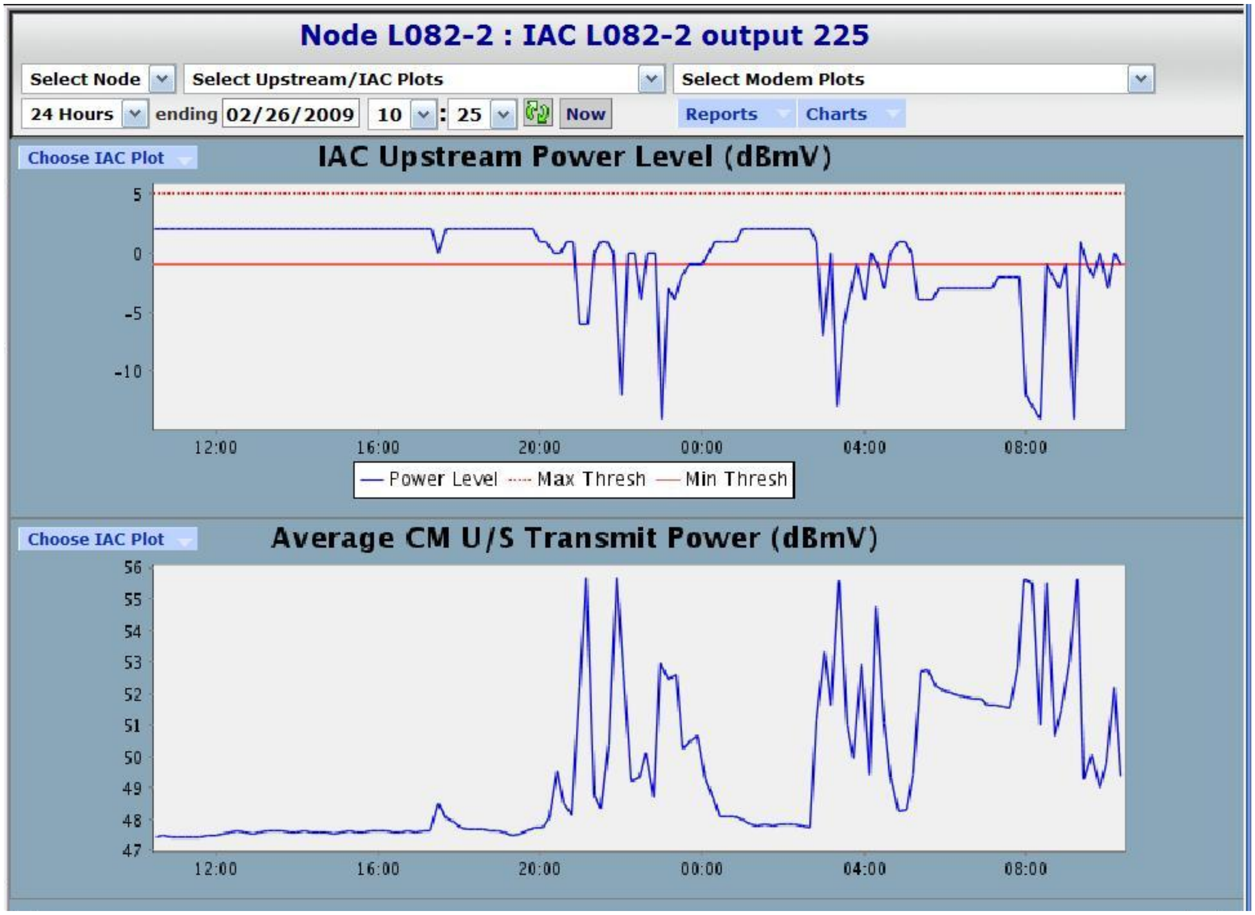

Without cross-system data synchronicity, plant and service performance diagnostics become extremely suspect when it comes to the reliable determination of service impact, root-cause, and precise location of a problem.

3. Lack of Network Topological Relationships

The lack of plant topology relationship between the various reporting systems severely impacts the operations team's ability in both locating the problem as well as correlating and synchronizing the behavior. Even if we could continually monitor every modem, phone call, set-top and customer experience, it is extremely difficult – if not impossible – to reliably correlate the behavior of various data elements within the coaxial plant, fiber, amplifier, etc. without a physical plant (network topology) relationship.

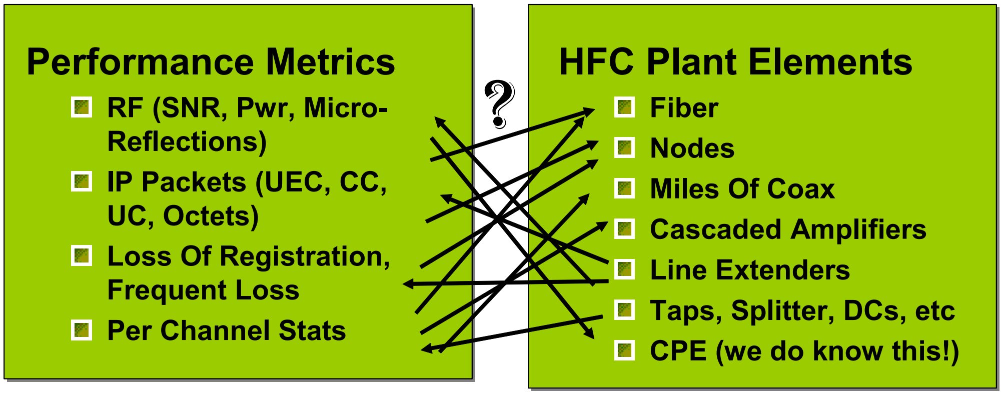

Performance Metrics

- **RF (SNR, Pwr, Micro-Reflections)**
- **IP Packets (UEC, CC, UC, Octets)**
- **Loss Of Registration, Frequent Loss**
- **Per Channel Stats**

HFC Plant Elements

- **Fiber**
- **Nodes**
- **Miles Of Coax**
- **Cascaded Amplifiers**
- **Line Extenders**
- **Taps, Splitter, DCs, etc**
- **CPE (we do know this!)**

A topological relationship must be made for accurate and automated problem isolation

Despite this lack of physical plant relationship, we still need to start somewhere. Today, operations teams start at the node and work their way down. Thanks to billing records with node names recorded with house-keys upon plant build-out, we have some notion of which nodes the customers are on. And, by using the street addresses, we can see approximately where the troubled customers are – but not necessarily where the *trouble* is. Also, since billing records are of varying integrity, many false indicators are considered and skew the analysis.

The lack of physical plant relationship raises the question: "Since two problem-customers have a geographical relationship (e.g., same street name, within 500 feet from another) and both have problems, it must a common problem, correct?" Not necessarily, as is evidenced by a typical example plant map below. Two of the highlighted areas (A and B), most likely have homes (modems) back to back. Therefore, they have a geographic relationship. However, from a network topology they have very little in common. Since they are on separate ports out of the fiber node, the only common intersection is the node itself. So, coincidental power level changes between the two – unless the problem is at the node, fiber, or headend – is just that, coincidental. Correlation of behavior in this scenario will mislead the diagnostics, location, root-cause, and most likely waste a lot of technician time.

In another scenario involving homes in groups B and C, here, again, there is a geographic relationship. Most likely, though, these problem homes are on two entirely different nodes. That means coincidental power, noise, SNR, packet loss changes between the two, unless the problem is now at headend (pads, CMTS, router), has little value and if combined will create not just misleading but also erroneous conclusions.

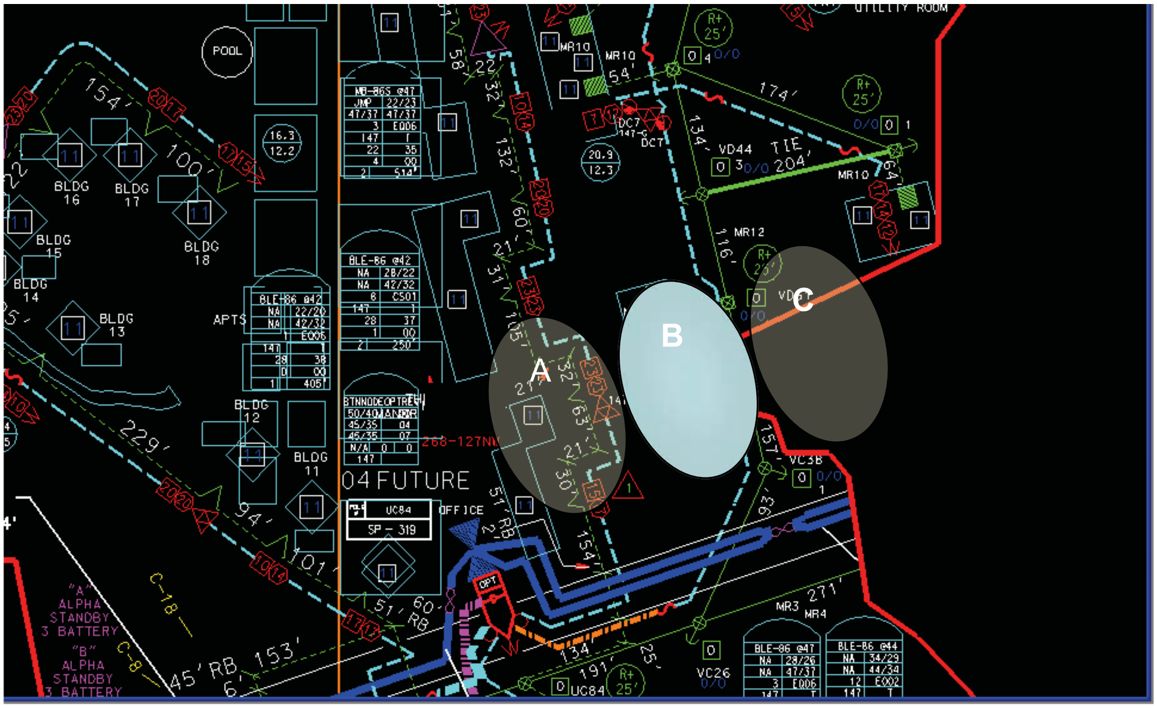

Geographic Relationships vs. Network Relationships

4. Lack of Automated Compensation or Containment Systems

The lack of automated compensation or containment systems is the final category of obstacles in moving toward the goal of cost-effective, continuously managed service and network health where MTTR approaches MTTr. The objective with such automated techniques is to kick in upon recognition to contain and mitigate the offending issue as much as possible until the problem can be fully resolved with human involvement. To put in human terms, such a technique would operate as a pacemaker, not purely a defibrillator.

Some techniques have been developed for headend systems to be agile in the face of noisy conditions. They include frequency hopping and forward error correction. These techniques do allow compensation in the light of unpredictable faults and are most effective during broad and constant noise situations, such as AGWN. During transient impulse noise, however, they have little effect (no predictability). Many amplifiers have built-in temperature and tilt compensation systems to accommodate seasonal temperature changes. However, these lack any real-time situational intelligence to detect and react to partial failures or degradations.

A New Network Intelligence

A new intelligence layer built into the cable infrastructure, deep into the HFC plant, is necessary to fill the gaps highlighted above and to ultimately achieve the industry's network availability and operations efficiency goals. This new intelligence must deliver:

1. Scalable, continuous behavior awareness
2. Auto-generation and maintenance of a topological relationship of information
3. Real-time data synchronicity across multiple layers
4. Automation all of STEPs 1-7 of the operations process (and STEP 9)
5. Provide automated compensation and containment (parts of STEP 8)
6. Operations focus at the neighborhood level

A deployment of this new intelligence layer recently completed by Charter Communications offers a view into the extraordinary potential of a new intelligence layer. This solution enabled the operator to fill the service assurance gaps and to experience the operational gains and expense savings associated with continuous, real-time service and network health assurance.

The system and architecture would be designed around the unique needs of the cable industry's HFC networks. It would feature real-time intelligence deployed neighborhood-deep into the access network that works in concert with sophisticated service health management software deployed in the headend and regional data center.

This system would include coordinated management of new intelligent service health management devices, generally referred to as a Health Management Devices (HMDs). (These are also sometimes referred to as an intelligent access controller (IAC)). These HMDs would be deployed deep in the HFC plant – typically at amplifier segments of 50-60-homes-passed clusters. Together with the centralized service health management software, this next-generation system would deliver intelligent management of real-time monitoring of critical fault and performance indicators, including continuous path and upstream ingress monitoring and analysis, with synchronized correlation of RF, IP and QoE level information. This intelligent management of the HFC network and services would be combined with managed ingress suppression technology – to eliminate the propagation of quality-eroding ingress.

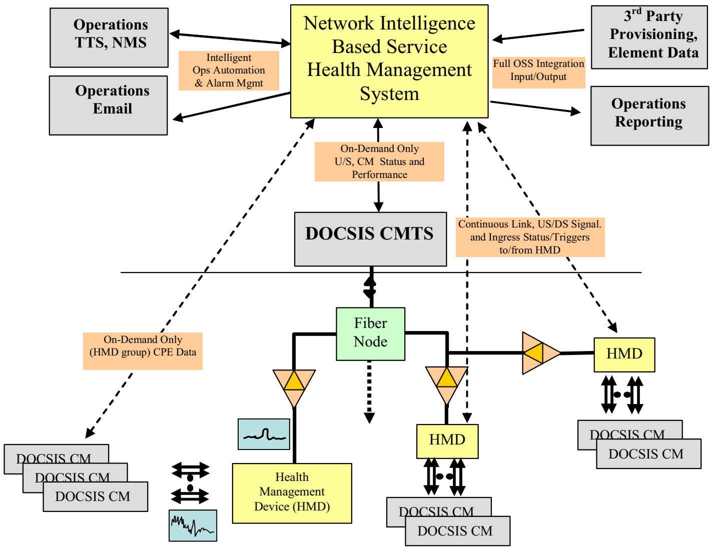
Operations TTS, NMS
Operations Email
Intelligent Ops Automation & Alarm Mgmt
Network Intelligence Based Service Health Management System
3rd Party Provisioning, Element Data
Full OSS Integration Input/Output
Operations Reporting
On-Demand Only U/S, CM Status and Performance
DOCSIS CMTS
Continuous Link, US/DS Signal. and Ingress Status/Triggers to/from HMD
On-Demand Only (HMD group) CPE Data
Fiber Node
HMD
HMD
Health Management Device (HMD)
DOCSIS CM
DOCSIS CM
DOCSIS CM
DOCSIS CM
DOCSIS CM
DOCSIS CM
DOCSIS CM

High-Level System Architecture

A Detailed View of the Intelligent HFC Plant Architecture

Remote, distributed intelligent devices are at the heart of this next-generation HFC plant architecture. As in the Charter deployment, HMDs would typically be installed deep into the cable access network just after each end-amplifier. This placement results in an average coverage area of 50 homes-passed for each HMD. One key function of an HMD would be to block ingress and noise coming from the cable network and at the same time allow valid signals, such as modem signals, to pass through. An HMD would largely prevent ingress and noise from entering the network and propagating to the headend during valid signal transmission. This ingress suppression capability isolates the return path coming from each cable network's amplifier leg segment until a valid signal is detected. Once a valid signal is detected, the HMD would instantaneously connect the return path from (and downstream from) the amplifier leg location to the upstream and headend network as long as the signal is present.

An HMD would communicate both upstream and downstream, using a low-bandwidth FSK coded channels (or other means) with a headend base RF modem. Using this communication channel, it would be possible to continuously control and monitor each HMD in the system. Each headend modem would communicate with several HMDs spanning several fiber nodes.

The HMD would be specifically intended to be element and location agnostic. It could be embedded with an amplifier, tap, or as a stand-alone in-line device, placed anywhere in the HFC plant, based on the specific network architecture. This flexibility is required in order to allow intelligence to be placed universally, deep and wide into the plant, providing precise, real-time service health management within a given neighborhood (or plant segment, such as amp leg).

This new intelligent architecture would provide complete, continuous coverage across the entire plant footprint, primarily covering each end amplifier (or amp leg) region. It would provide a platform to deliver real-time, un-ambiguous plant health intelligence, while providing active problem detection, isolation, compensation and containment within a fraction of a node's network geography. Coupled with real-time intelligence to manage and control the distributed HMDs, this system would allow Charter's operations to increase service availability, reduce operations cost while increasing service performance. Key elements of this intelligent architecture are highlighted below.

1. Delivering Scalable Continuous Behavior Awareness

The intelligent HFC plant leverages its deep presence of HMDs to provide a continuous behavior awareness of the HFC plant and the services running over it. Because of its distributed nature and autonomous and imbedded intelligence, each HMD would allow scalability to deliver health management automation with accuracy and efficiency.

This system would provide three types of continuously monitored base health indicators:

a. Loss of Communications (to/from headend and HMD)
b. Upstream and Downstream Signal Levels and Attenuations
c. Upstream Noise Detections and Characterizations

These three indicators offer the needed low-impact and continuous monitoring of vital statistics and a very high-speed, low bandwidth "trigger" (STEP 1 from the figure on page 1) for advanced operations automation for STEPs 2-7, as well as STEP 9, through other advanced features of the intelligent HFC plant.

a. Detection of Loss of Communications

Using independent (e.g. FSK) telemetry channels, an HMD would maintain a continuous heartbeat of communications to a headend controller. Its status, its location and what it “sees” is known at all times by a service health manager.

A loss of communications represents an immediate trigger for a potential plant power failure or line break detection. However, this is only a trigger or symptom of a potential issue that must be validated in the standard operations processes. If un-validated and without a persistency test, such as that of traditional transponders, this issue could result in many false alarms, residential power outages (not a plant issue) or intermittent “no trouble found”. Through the automated discovery and maintenance of plant topology (e.g., cable modems, set tops, etc.) below an HMD (and amplifier), as discussed in the next section, the centrally based service health management system would automatically qualify these event triggers into a validated service-affecting outage condition, At Charter, this automated discovery would be instrumental in significantly increasing the reliability of plant outage detection and significantly cutting the MTTR for outages.

b. Upstream and Downstream Signal Level Monitoring

The intelligent network elements would also use the telemetry communications as a power level pilot, where the upstream and downstream signal levels can be continuously monitored. When each of the HMDs is configured for placement into the plant, it would be set for an expected (designed) DOCSIS return level hitting the HMD. This expected level and respective configuration, will also affect the power level of the upstream transmission. In conjunction with a headend combining plan, it is typically set to reach the headend modem at approx 0 dBmV (the same as DOCSIS at the CMTS). Through the service health manager, we can then continuously evaluate the current power levels across a given node, relative to all peer HMDs, at any given time. This view gives the operations team an unambiguous comparison of upstream balancing and overall performance of gain/attenuation against plant design. The top chart below shows power levels from each of eight HMD devices within a given node. From this chart, it is apparent that there is a wide range of upstream values, which are typically expected to be fairly consistent around +2. This suggests a probable balancing issue at the two extremes (e.g., +9 and -3) and should be verified for proper balance.

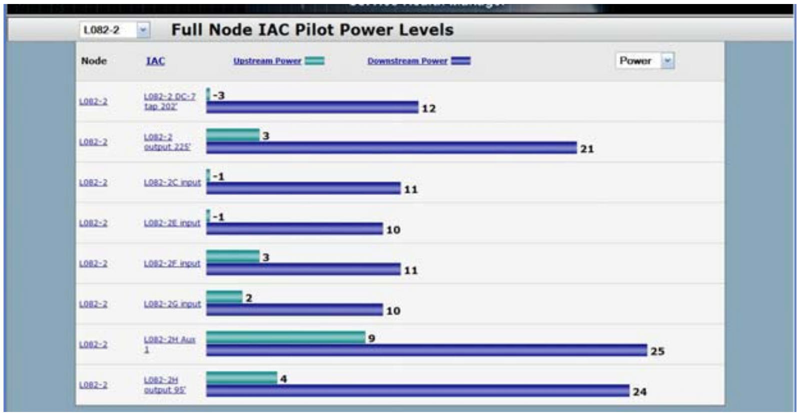

In the next example illustrated in the second chart, we see that a node is viewed to be “balanced” relative to one another (and at the CMTS). However, they are universally low, at -10 from expected. Instead of indicating a balancing issue at one or two locations, this chart suggests a more serious issue affecting the entire node, potentially at the node itself, fiber, or at the headend.

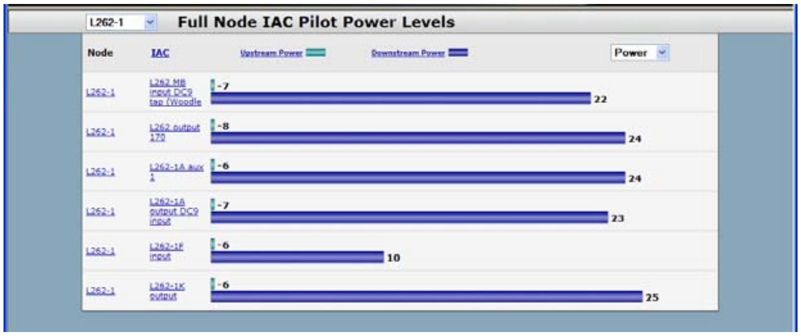

Once a node is deemed balanced, the current values can be viewed as a baseline against which major changes can be used as a second problem symptom detection (in STEP 1) as well as serve to prompt re-visitation for re-balance when variances are detected. And, through the CM discovery process and real-time correlation of performance levels, this change can be automatically validated as service impacting or not (given it is already located), thus implementing STEPs 2-7 of the process *before* ticket/dispatch would be considered.

c. Continuous Noise Detections and Characterizations

While an HMD would provide intelligent and automated suppression of upstream noise and ingress, which will be discussed in detail later, it also would provide continuous capabilities to detect, rank (i.e., compare) and precisely locate sources of noise and ingress. In addition, an HMD would be able to identify the type of ingress (e.g., short duration, long duration, or broad noise, potentially Common Path Distortion or CPD) and its exact time of occurrence and intensity in real-time. The plot below shows periodic intermittent short ingress, predominantly in one location, off amplifier leg AA23B main output. The predominance of broad ingress indicates the source most likely is CPD.

These ingress values are continuously monitored, recorded and transmitted using the standard FSK channel telemetry, with no additional bandwidth for polling required. Through the continual awareness of ingress and noise, thresholds can be set to be used as a third problem symptom trigger at the initial onset of sudden ingress (STEP 1), again with real-time detection and no latency or impact on usable bandwidth.

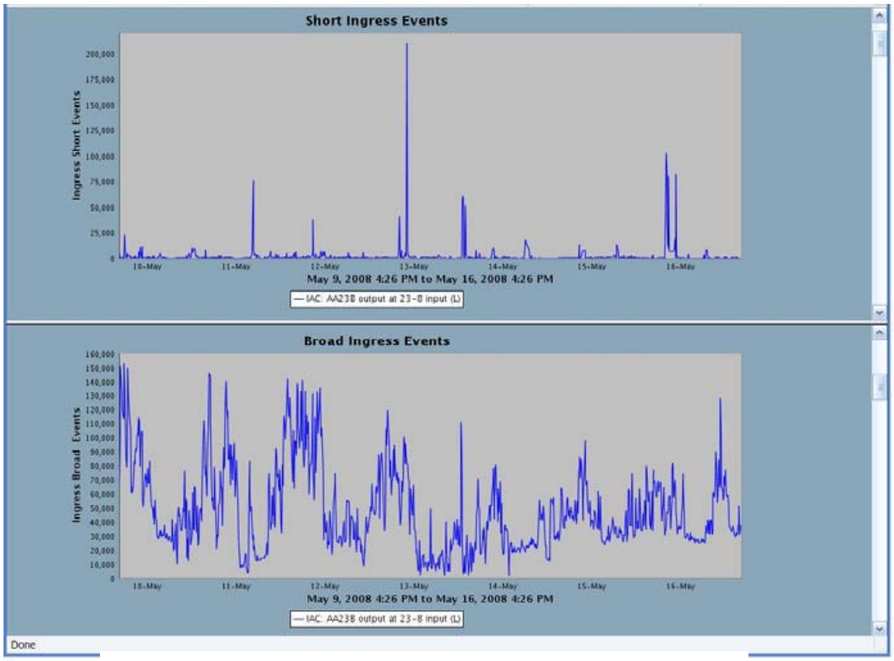

Continuous Ingress and Noise Monitoring

This figure below shows an initial rise in detections of high-intensity ingress. These rising changes in ingress counts are used as initial ingress and noise problem triggers. With this trigger, utilizing the per-IHMD-to-Cable Modem discovery, the 10-20 modems within this amplifier leg are then targeted for precise and high-rate polling (e.g. 2-5 minute rate over 24 to 48 hours) to automatically correlate both RF and IP performance response in light of the on-set of severe RF impairments. This approach is a radical departure from standard massive periodic "event-less" polling, and represents a completely scalable, automated, and latency-free determination of problem validation, impact, and scope within the isolated location. More importantly, correlated time and synchronous evaluation of *just* the devices experiencing this same RF impairment is vital to accurate assessment of root cause and true impact, which will be further illustrated in several examples to follow. Through direct, real-time correlation of an independent RF condition with modem power, noise, and packet performance, service-impacting criteria can be set to prioritize and filter against all current conditions. Therefore, through this process, STEPs 2-7 of the process are

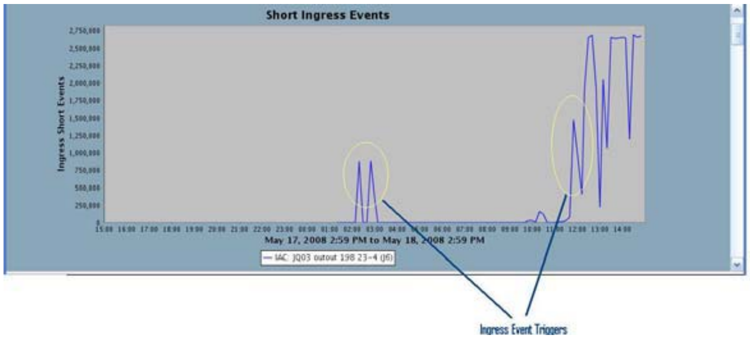

Using Start of Ingress/Noise as Event Trigger automated with highest level integrity before a ticket/dispatch would be considered and ultimately dispatched against the highest service-impacting priority.

2. Auto-Generation and Maintenance of Topological Relationship

To effectively implement a scalable system that uses a deeply placed, independent symptom trigger device, understanding which other network elements should be considered for synchronized polling and analysis and which ones should not is an absolute requirement. Knowledge of just which modems, phones, set-tops, power supply transponders, etc., are within the network segment of concern allows scalable, full qualification of the service impact of any RF- or IP-triggered anomaly. Further, the inclusion of other devices not in this segment for high-rate analysis leads to diminished scalability in both monitoring and reporting. More importantly, performance data from devices outside the RF condition will mislead, dilute analysis, and most-likely misdirect maintenance activity.

For example, assume packet loss data is polled and then calculated from 10 modems on a given amplifier leg during a given severe RF condition. It is determined that these 10 modems transmitted 1000 codewords and lost 40 (or 4%) over the event-triggered analysis period. This example suggests a pretty severe condition that would wreak havoc on phone call quality. Let us now expand the example to include a second set of 10 modems in a second amplifier leg that is not experiencing the same RF condition – with one modem uploading several large files to MySpace or YouTube – and as a group transmitted 9,000 codewords and lost only 10 (or approximately 0.1 %). The combined calculation for packet loss rate would be 50/10,000 or just 0.5 %. While this result is not great, it would be deemed acceptable – despite 10 homes (50% of sample) having severely unacceptable voice performance. Now, by including the 300 or so modems from the remainder of the node in the analysis, the quality problem for these 10 customers are completely masked. Furthermore, when adding in the dimension of time (our non-continuity problem), if the polling rate for packet loss is every 4 hours, with an intense RF condition occurring for two of those hours, then 2 hours of uploads (unaffected by RF condition) would be incorporated into the analysis. With 2 hours of near-perfect data, the reporting is again diluted and may be masking a bona fide condition that is severely impacting 10 customers.

The above example shows that a data time-synchronicity and topology relationship is essential for reliable and effective service health management. There have been many attempts over the years to convert geographical plant maps from CAD drawings or to manually recreate paper-based plant maps into database-driven systems and then to instill strict change management practices to maintain the accuracy of the maps through all plant changes. This approach has achieved mixed results, since it requires institutional change in plant modification policies and well-trained and skilled CAD/drawing technicians permanently on staff to rapidly transfer all changes into the maintained database system. Unfortunately, it is also prone to manual errors in either reporting the changes or properly entering those changes. Which in most cases are only to be detected, perhaps, when the time arrives to reference them (for problem analysis), and come to find out we are directed to the wrong location, and perhaps for the wrong problem. Without constant audit, these maps tend to accumulate inaccuracy.

To solve this difficult industry problem and enable the power of real-time, precise, and time-synchronous data correlation, a key part of the embedded intelligent devices are to automatically generate a network topological relationship of a given HMD and all network devices connected upstream and downstream from them. This includes its co-location at the input or output side and an end-amplifier. In this newly intelligent network, this critical function needs to be accomplished not through geographic inference or data-mining of potentially faulty billing data records, but through active validation of device connectivity through each HMD. Unfortunately there are currently no standards for auto-discovery of HFC elements and network topology today, but would be beneficial for future standard considerations.

A commitment to this approach and the results of network awareness enables an unambiguous, validated topological relationship based on actual behavior of all customer premises equipment (CPE) or any other elements, such as DOCSIS power supply transponders.

The notional relationship of node and amplifier associated with the HMD's location and the HMD's location itself result in an accurate, continuously maintained network topology of the customer premises device heath and status information. This relationship can now include all service performance information – including customer complaints – and the actual localized physical cable infrastructure (e.g., cable, taps, amps, DCs, splitters). Unlike a converted plant map, the network topology integrity is constantly audited and re-discovered through each automated discovery process, eliminating any chance of slow integrity degradation through plant changes. Only with such data integrity can any system be reliable and accurate enough to automate intelligent decision logic for problem validation, scope, service impact, isolation, and root cause analysis. This ultimately includes automatic verification of recovery after plant repair against the exact devices and services operating though that physical plant.

This plot from the Charter deployment includes a list of all HMD locations within a given node, FA, designated by the amplifier name and either input or output (Aux1, Aux2) as well as tap type and location. Within a given amplifier leg, FA16A output, at the first tap location (23-4), 15 devices were discovered to be downstream of this location. Using the discovered MAC addresses, the billing address is easily correlated and displayed. On the example plot shown, the rapid decline of average SNR from *just* these 15 cable modems can be clearly seen. Through focused analysis of these modems, a direct correlation of eventual major packet loss occurring for six hours (4:00 to 10:00pm) is observed.

Therefore, with this topological relationship and high-rate detection of the problem "triggers" (e.g., signal, noise, connectivity, throughput) discussed earlier, the system would initiate focused, real-time correlation of service performance metrics, to give the most complete and accurate assessment of problem validation, scope, and service impact, which will be illustrated in more detail in the following pages.

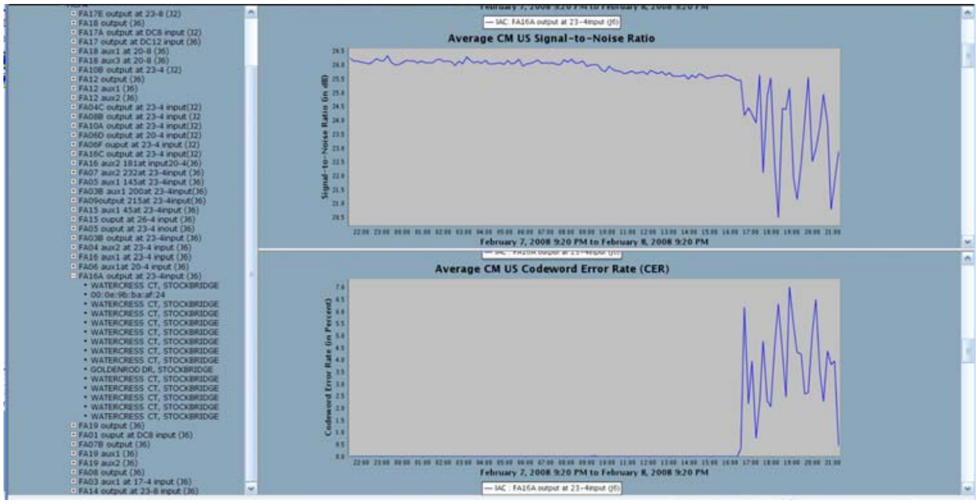
FA17E output at 23-8 (32)
FA18 output (36)
FA17A output at DC8 input (32)
FA17 output at DC12 input (36)
FA18 aux1 at 20-8 (36)
FA18 aux3 at 20-8 (36)
FA10B output at 23-4 (32)
FA12 output (36)
FA12 aux1 (36)
FA12 aux2 (36)
FA04C output at 23-4 input(32)
FA08B output at 23-4 input (32
FA10A output at 23-4 input(32)
FA06D output at 20-4 input(32)
FA06F ouput at 23-4 input (32)
FA16C output at 23-4 input(32)
FA16 aux2 181at input20-4(36)
FA07 aux2 232at 23-4input (36)
FA05 aux1 145at 23-4input (36)
FA03B aux1 200at 23-4input(36)
FA09output 215at 23-4input(36)
FA15 aux1 45at 23-4input (36)
FA15 ouput at 26-4 input (36)
FA05 ouput at 23-4 inout (36)
FA03B output at 23-4input (36)
FA04 aux2 at 23-4 input (36)
FA16 aux1 at 23-4 input (36)
FA06 aux1at 20-4 input (36)
FA16A output at 23-4input (36)
WATERCRESS CT, STOCKBRIDGE
00:0e:9b:ba:af:24
WATERCRESS CT, STOCKBRIDGE
WATERCRESS CT, STOCKBRIDGE
WATERCRESS CT, STOCKBRIDGE
WATERCRESS CT, STOCKBRIDGE
WATERCRESS CT, STOCKBRIDGE
WATERCRESS CT, STOCKBRIDGE
WATERCRESS CT, STOCKBRIDGE
GOLDENROD DR, STOCKBRIDGE
WATERCRESS CT, STOCKBRIDGE
WATERCRESS CT, STOCKBRIDGE
WATERCRESS CT, STOCKBRIDGE
WATERCRESS CT, STOCKBRIDGE
WATERCRESS CT, STOCKBRIDGE
FA19 output (36)
FA01 ouput at DC8 input (36)
FA07B output (36)
FA19 aux1 (36)
FA19 aux2 (36)
FA08 output (36)
FA03 aux1 at 17-4 input (36)
FA14 output at 23-8 input (36)
IAC: FA16A output at 23-4input (J6)
Average CM US Signal-to-Noise Ratio
Signal-to-Noise Ratio (in dB)
22:00 23:00 00:00 01:00 02:00 03:00 04:00 05:00 06:00 07:00 08:00 09:00 10:00 11:00 12:00 13:00 14:00 15:00 16:00 17:00 18:00 19:00 20:00 21:00
February 7, 2008 9:20 PM to February 8, 2008 9:20 PM
Average CM US Codeword Error Rate (CER)
Codeword Error Rate (in Percent)
7.0
6.5
6.0
5.5
5.0
4.5
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
22:00 23:00 00:00 01:00 02:00 03:00 04:00 05:00 06:00 07:00 08:00 09:00 10:00 11:00 12:00 13:00 14:00 15:00 16:00 17:00 18:00 19:00 20:00 21:00
February 7, 2008 9:20 PM to February 8, 2008 9:20 PM
IAC : FA16A output at 23-4input (J6)

Topological Relationship of Performance

3. Real-Time Data Synchronicity Across Multiple Layers

The network-based intelligent system no longer provides one silo or layer of information, but cross layer, time synchronized analysis of RF, IP, QoS, and the customer experience. This system provides direct correlation of aggregate ingress/noise and its effect – or not – on upstream SNR, on packet performance, i.e., correctable codeword rate (CCR) or uncorrectable codeword error rate (CER), or other metrics such as Mean Opinion Score (MOS) for measurement of voice quality. These and other key correlations provide direct relationships of the RF characteristics affecting RF performance, detected by two independent means (HMDs and CMs), how they are related, and what impact they have on service performance, at the time they are used.

When this data is viewed as an aggregate from a given upstream channel (which most current systems do), it includes performance data from the entire DOCSIS upstream (i.e., all connected plant and devices) as well as combined data from all HMDs within a given upstream scope, such as a node. This perspective gives only the top-most view of the upstream channel performance and service health.

In the example upstream performance plots from the Charter deployment shown here, a user can view a time-synchronous correlation of the upstream channel's composite CER for all modems communicating on upstream channel with the total upstream codeword traffic. An operations team can see that the CER reaches a level of 11% during a total traffic of nearly 4,000,000 unerrored codewords (or approx 400,000 codewords lost). This view allows for qualitative analysis of actual customer experience (i.e., severity and time of day) as an aggregate across the entire upstream channel's utilization.

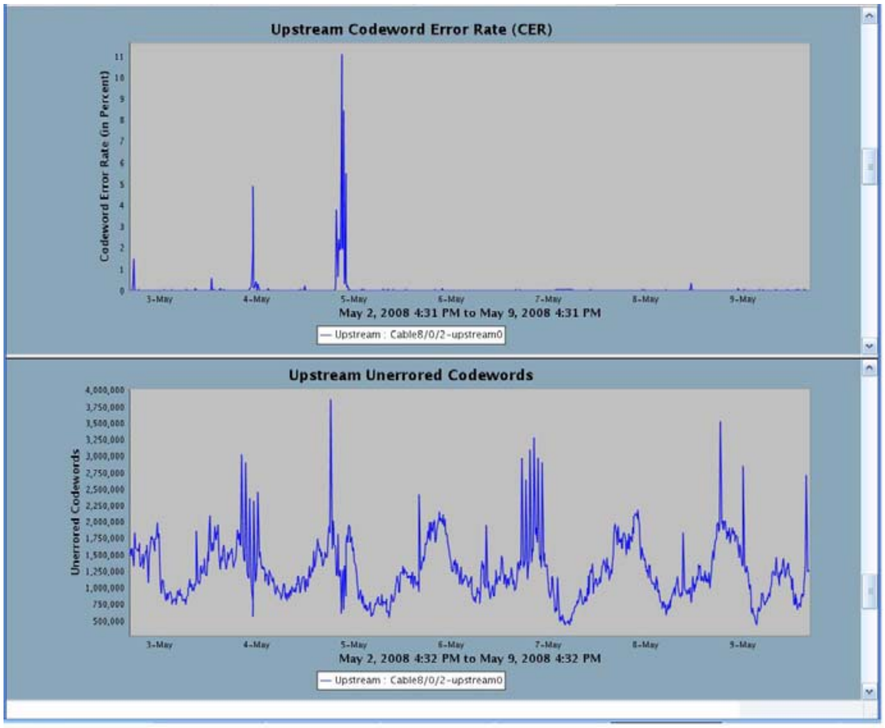

Intermittent Severe Pack Loss on Upstream

In the plots to the right, operations can also view a time-synchronous correlation of the upstream channel's same composite CER with the composite upstream Signal-to-Noise Ratio. As illustrated, when the CER reaches a level of 11%, the overall upstream SNR dropped to 19 dB. It is also clear that, although a very severe condition, it occurred for only an hour or two over a seven-day period.

However, from a service outage perspective, an hour or two is a very long time. This example highlights that if we were not polling fast enough or during that period of adverse condition, we would likely reach an incorrect conclusion that this upstream looks good with continuous SNR near 32 dB.

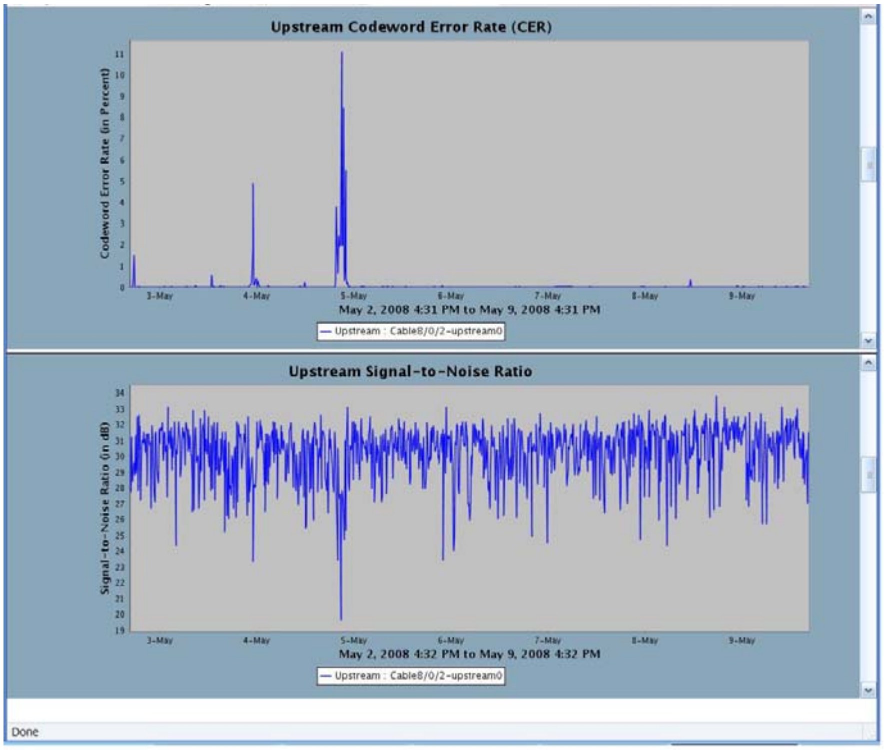

Intermittent U/S Pack Loss and SNR

In the next example, the upstream performance plots show a time-synchronous correlation of the upstream channel's aggregate broad ingress events from the several distributed intelligent network devices compared to the aggregate upstream uncorrectable packets.

The total given is the sum of events reported from all HMDs deployed in the same node/US. The user can see a direct correlation of the broad ingress and its effect on severe codeword loss for more than a given two-hour period. Through this type of view, we can now see a direct correlation of sudden impulse noise occurring on the system (from an RF perspective) and then the severity of this noise, not in terms of RF (such as amplitude or count) but in IP (packet) terms of service impact and customer experience.

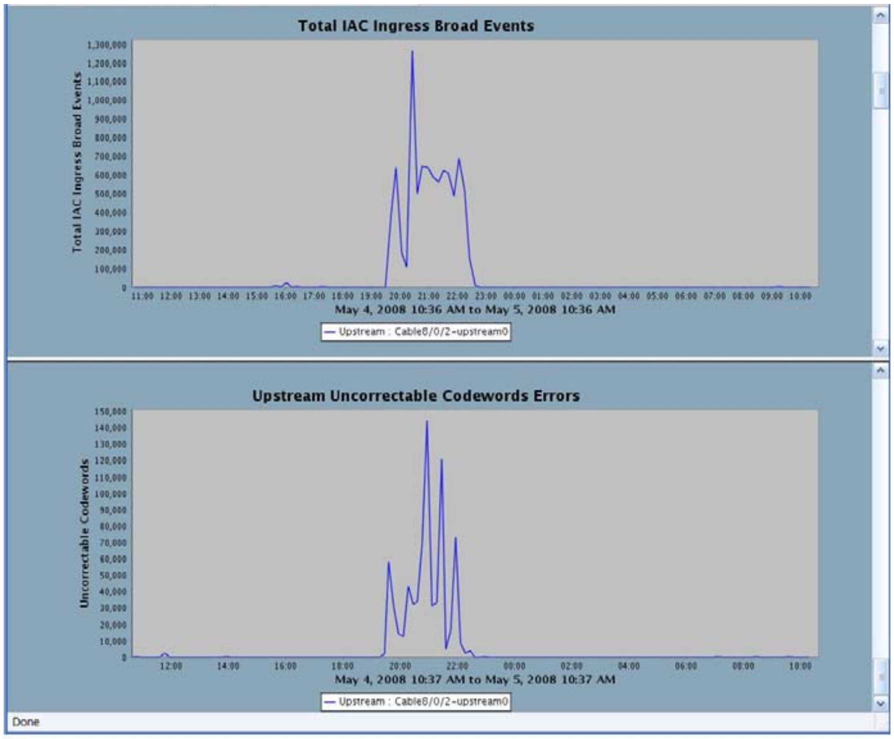

Precision Broad Ingress and U/S Pack Loss Analysis

With this precision correlation, we see a time-synchronous correlation of the upstream channel's aggregate broad ingress events and uncorrectable codewords over time. In this example, we see that this broad ingress condition is a chronic, cyclic problem that can cause short duration, but major impact on service performance – and, certainly, a top priority target for investigation and repair.

Top-level aggregate upstream comparisons are valuable for detected large problems manifesting in a full node effect. They do very little to help focus our attention (and operations effort), detecting and repairing where most of the problems lie, in each plant segment. Although upstream level analysis is common among standard DOCSIS monitoring tools, they typically are not correlated with independent RF monitoring systems, as was the case at Charter. They only provide comparisons across other DOCSIS values and therefore provide only an overall, top-level indication of HFC plant and service health. They may indicate a generally healthy or problematic service performance, but they provide very little information as to what and where the problems are, and many times dilute (and therefore, hide) localized plant regions exhibiting extremely bad performance. A general "healthy" upstream- (or downstream-) based performance metric, while severe conditions exist in an amp leg region, is what we call the "whole upstream illusion".

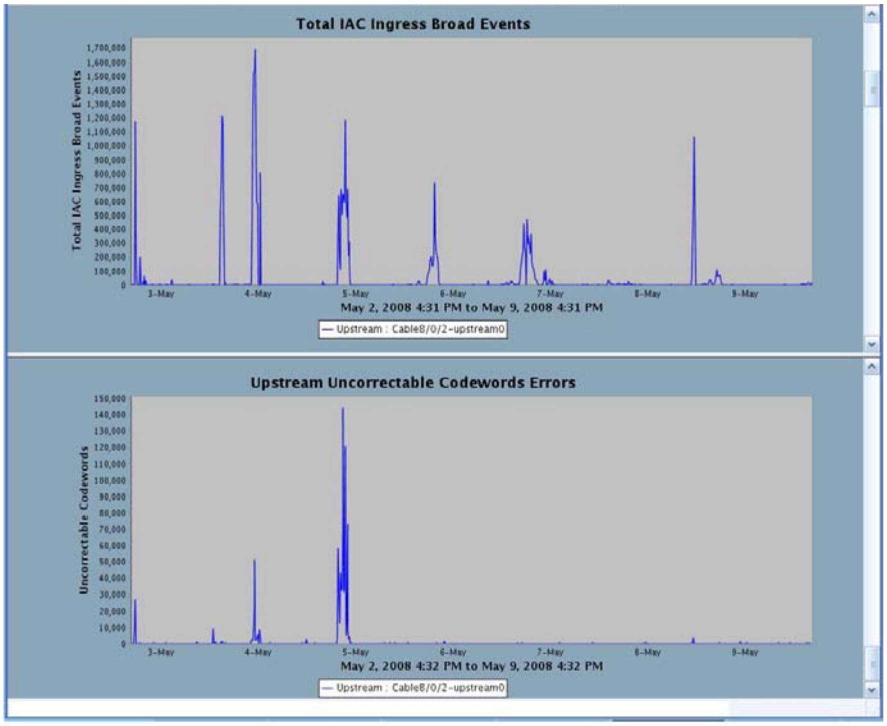

Repeated Broad Ingress and U/S Packet Loss

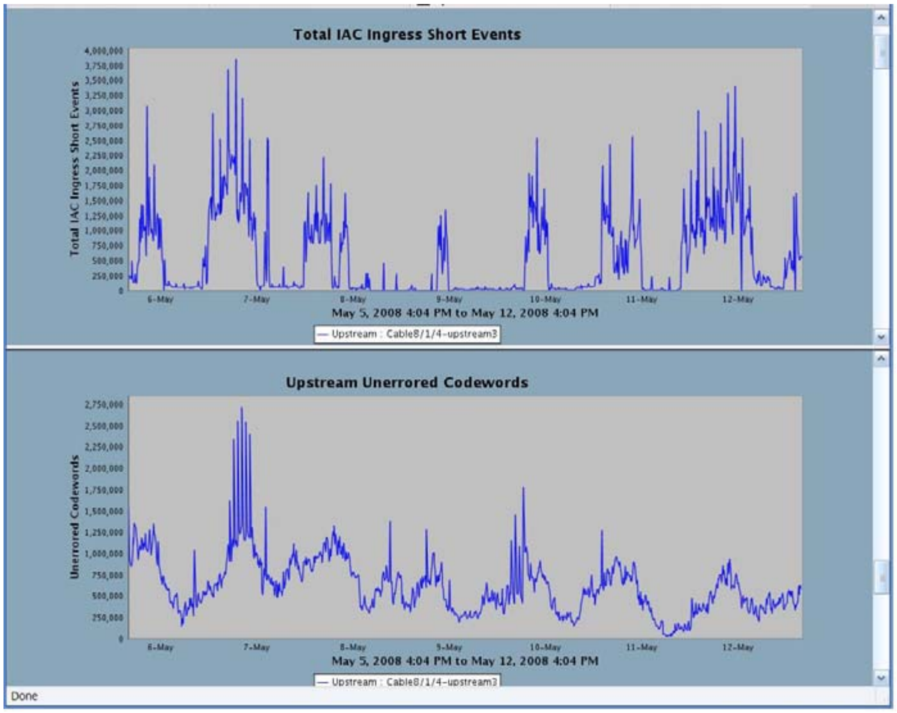

Chronic Short Ingress Aggregated at All Nodes HMDs

As an example, let us look at the following aggregate performance plot at the entire upstream level. From a top-level view, we at least can see that severe chronic short duration ingress (< 50 µs pulses) is occurring over a highly utilized upstream channel.

However, looking at the service impact using the upstream aggregate CER measurements in the next plot, we see very moderate impact, typically running at .025% (1 codeword in every 4000).

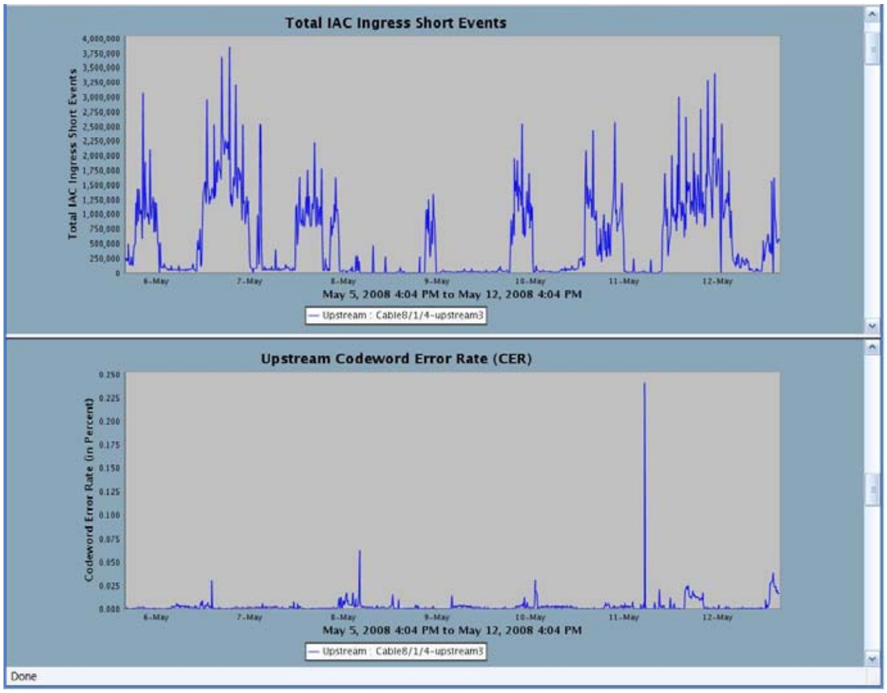

Chronic Short Ingress With No Measurable Upstream Impact

Also, analyzing service impact (or symptoms thereof) using the upstream's aggregate SNR measurement, we do see some overall variance, but not any reason to be concerned about or to dispatch for. In fact, a mean of the SNR shows a reasonably healthy SNR of 32-33 dB.

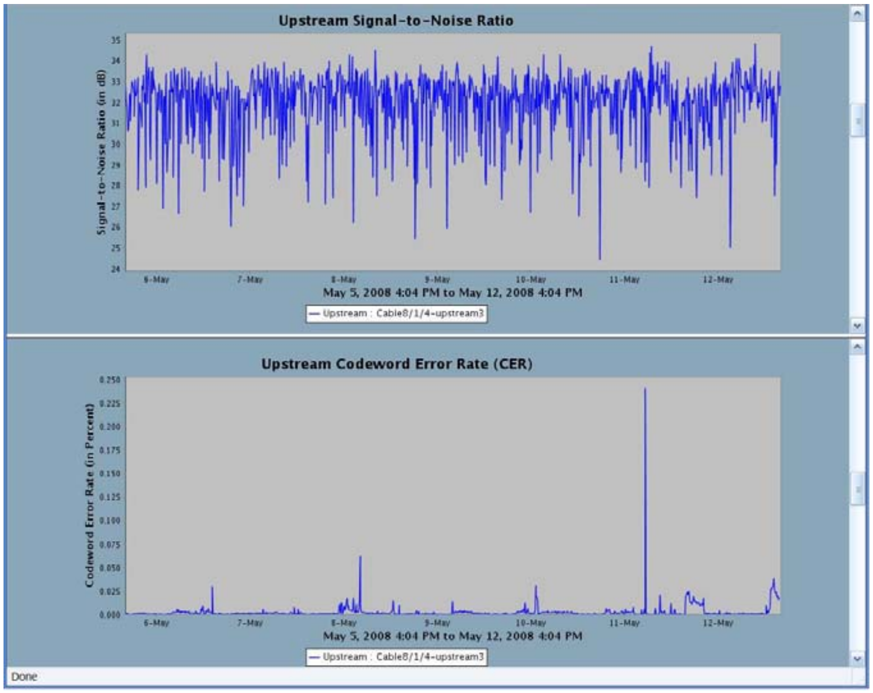

No Measurable Upstream Impact on SNR

We see a completely different story, however, by digging deep and analyzing service impact (or symptoms thereof) using the localized HMD regional noise indications as well as HMD region's modem SNR and CER measurements. This view is enabled by the HMDs deep deployment, automated CM discovery, and event-triggered, high-rate modem data correlation within an HMD region.

We see, that during the same exact period as that measured at the whole upstream (the upstream channel allocated to node JQ), that one primary location, when analyzed only over the HMD region, had repeated bad performance of up to 6% codeword loss. This means pockets of 20 customers, under amplifier leg "JQ-03 output," are experiencing repeated, unacceptable service performance, given that CER > 1% generally causes a measurably garbled phone call.

Without this deep network intelligence, this degradation is undetectable when measured at the entire upstream/node level.

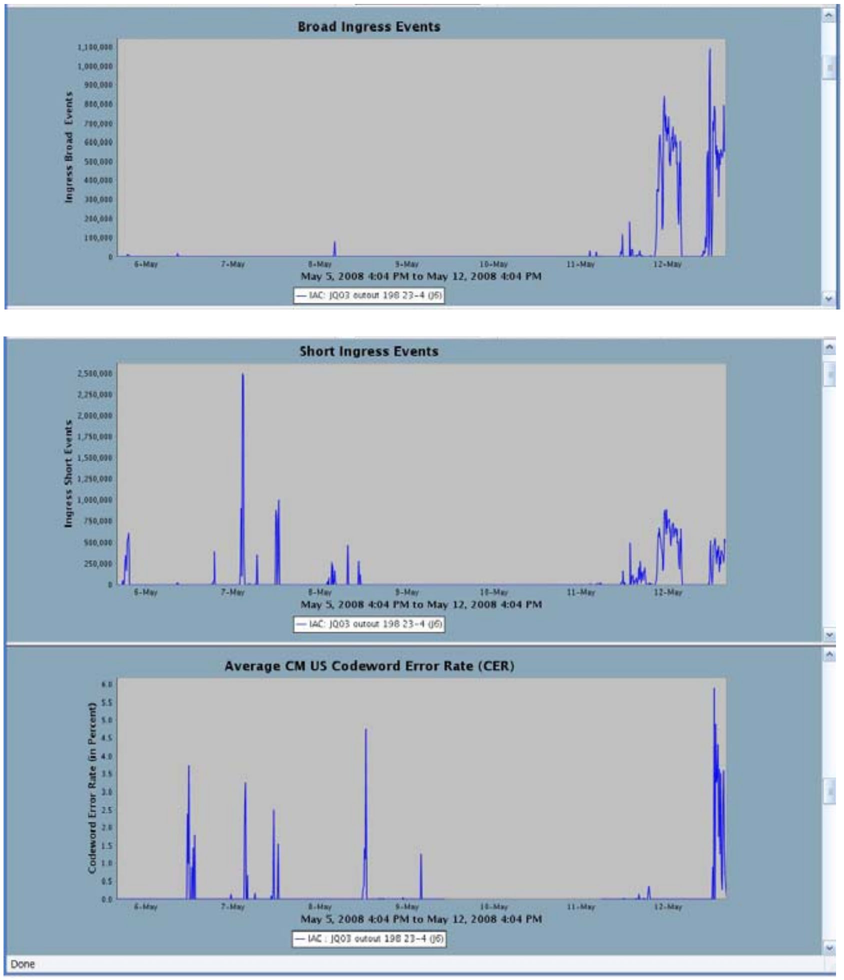

Localized Ingress With Serious Impact on CER

Also, with deep analysis of average SNR of just those modems experiencing the same repeated high-level ingress, the operations team can see that the average CM SNR is running more than 3 dB less that indicated at the entire upstream.

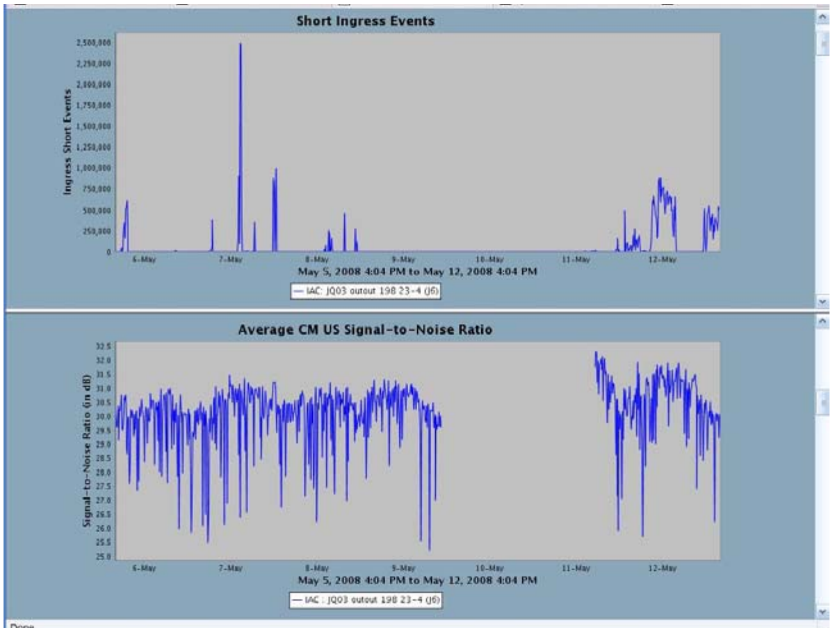

Localized Ingress With Moderate Impact on CER

The examples above illustrate that to effectively detect – without random, ad-hoc and manual testing – and to prioritize, locate, and repair or improve service performance, it is imperative to look deep into the HFC plant. Localizing and synchronizing the analysis by network segment is essential – right down to the individual amplifier or amplifier leg. The objective is to focus on the performance across only a cluster of relevant devices, all with direct network topological relationship, i.e., on common bus leg, taps ports, through common amplifiers, off common power supplies.

A Neighborhood View into Real-Time Network and Service Performance

The subsequent example plots from the Charter deployment zoom into those topological HFC regions to provide a precise view of RF and service performance across individual HMD pockets on common actives and passives of 50-60 homes passed. For each selected HMD region (i.e., the segment of plant defined by an HMD of a given amp leg and the physical plant connecting its discovered modems), the HMD's specific detected short and broad ingress, and upstream and downstream levels are plotted. In addition, several key aggregate modem performance data (i.e., composite over all modems within a HMD region) is provided across the exact time scale to that of the signal and noise plots. This modem performance data includes average modem group US transmit power, average modem group US SNR, and average modem group CCR and CER. Also, data synchronized include average modem downstream receive power, average modem downstream SNR, and total unerrored codewords (throughput) that is transmitted through the HMD by attached modems.

By selecting individual plot pairs, direct time correlation of key cross-layer metrics can now be viewed. For example, direct correlation of aggregate ingress/noise reported at a single location, demonstrated its effect (or not) on average upstream SNR variance, average modem transmit power variance (indicating a common attenuation change in a common path), as well as its effect (or not) on packet performance (CCR, CER).

These key correlations provide direct relationships of the RF characteristics affecting RF performance, detected by two independent means and how they are related. This gives a very precise view of the IP/service Performance correlation as it relates directly to the common RF behavior experienced within this HMD region of the HFC network. This also gives an unambiguous determination of service impact for any (and all) RF layer anomalies, qualifying the relative priority for action (dispatch).

Equally important, this evaluation allows us to know where the problems are *not*. Most of the time, looking at each individual or aggregate metrics at any given time, only one or two regions (or clusters) indicate a disproportionate weight of problematic metrics, as an experienced cable technician would expect. But the questions remain as to where to start, how severe and how to prioritize among other problem hot spots. With this system, continuous and built-in localization would be provided across the entire operations footprint. And, with a ranking analysis of key metrics, field operations staff is focused only in the plant segment with problems – at the exact time they are occurring.

In the HMD region performance plots shown here, we see synchronous correlation of the HMD detected short ingress events with the average upstream transmit power for all modems communicating through the HMD/amp leg. By examining the *change in average* transmit power rather than simply *absolute* transmit power (which can be different from modem to modem), we can detect common attenuations that affect all modems connected to a given amplifier. By synchronously correlating with independent noise or attenuation experienced at this location, we can discover if the transmit variance is a result of induced noise or change in signal gain or attenuation. In this example, the average CM group upstream transmit power level fell 4 dB in the same time period of HMD pilot power attenuation at Amp JQ03 main output. In the time-synchronous analysis, you can see that nearly 3-6 out of every 100 codewords were lost persistently over this 24 hour period.

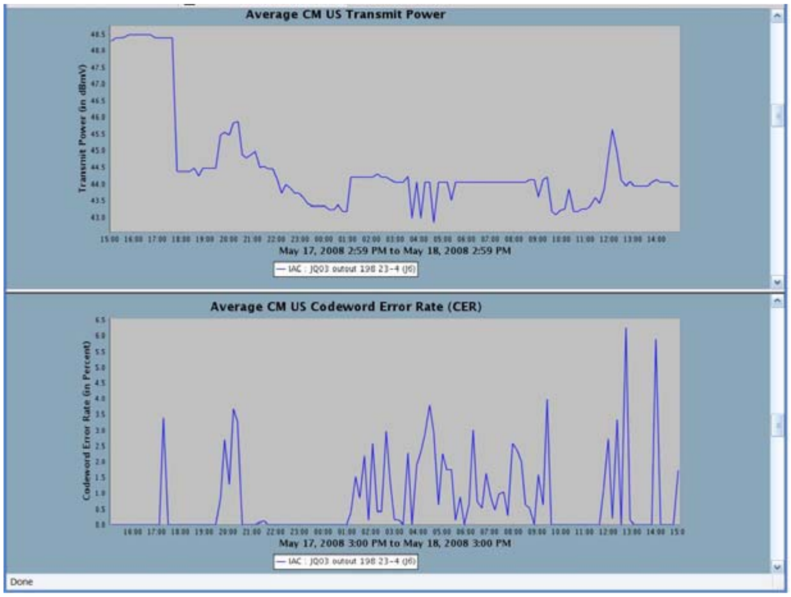

24-Hour Zoom of HMD Region Transmit Power and CER

As shown in earlier sections, the system would allow continuous analysis of its upstream and downstream pilot power levels. Using changes from these pilots as triggers, it would automate the time-synchronous polling and analysis of the change in the HMD modem group's upstream transmit power and downstream receive power accordingly. It also will incorporate high-rate analysis of the potential effects of customer experienced performance through packet loss of this same modem group. This method provides a very accurate means to detect and validate true attenuations – ones that affect *both* the HMD pilots and the CM group levels. By using only the CMs within the given amp leg (in this case, JQ03 main output) and then by calculating change in average transmit instead of just looking for out-of-range modems (e.g. > 60 dBmV), an accurate assessment of true attenuation can be detected, validated, located, and scope of service impact determined as shown below.

Moreover, if other modems were considered – ones that are not on a common amplifier – then problems would either be masked or a false problem determination would be made.

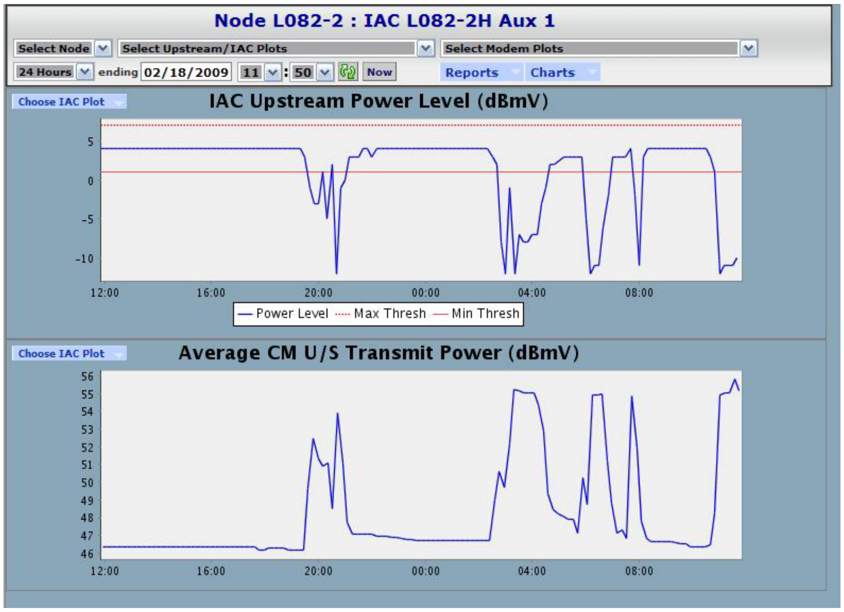

Continuous Cross System Signal Level Correlation

For example, let us say that 400 modems on given node's transmit power ranged from 40 to 55 and averaged 50 dBmV. In a given modem group of 20 modems, the average fell from 48 to 44. In that scenario, the overall node's CM upstream transmit average would range from 50 before an attenuation, and then (20*44 + 50*380)/400 = 49.7 dBmV after. In other words, this situation would be undetectable (at a change of 0.3) and therefore never analyzed for service impact. Furthermore, it would most likely not be resolved until during scheduled maintenance or after customer complaints.

Alternatively, if another modem group, under a different node port or amplifier was also experiencing a different attenuation effect, the two problems may wrongly indicate a full node problem, rather than two isolated amp leg issues.

In performance plots shown here for another HMD region at Charter, we see a time-synchronous correlation of the average US SNR for modems only within a given segment of the network with the total unerrored upstream codewords transmitted from those same modems.

In this example plot, we clearly see large dips in SNR, on average, from those modems on Amp FA05 main output leg, while viewing the upstream traffic flow through this amp leg.

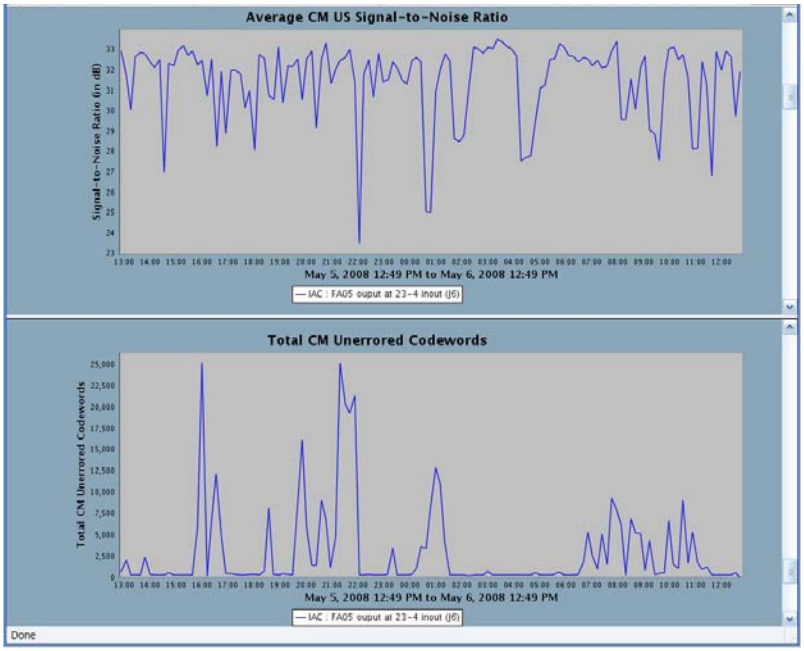

24-Hour Zoom of HMD Region Transmit SNR and Throughput

In the HMD region's performance plots shown here, we see a time-synchronous correlation of the HMD's detected short ingress events at that location, with the average CM codeword error rate (for all modems communicating through the HMD/Amp leg). In this example plot, chronic and cyclic short pulse width noise directly correlates with chronic, repeated packet loss. This pattern has occurred over a seven-day period, affecting on average all modems within that plant segment (i.e., Amp JQ03's main output). This plot provides a direct "service impact" view of this detected and persistent noise and on the service quality received by customers on that leg – regularly, a near 6% loss for a few hours a day.

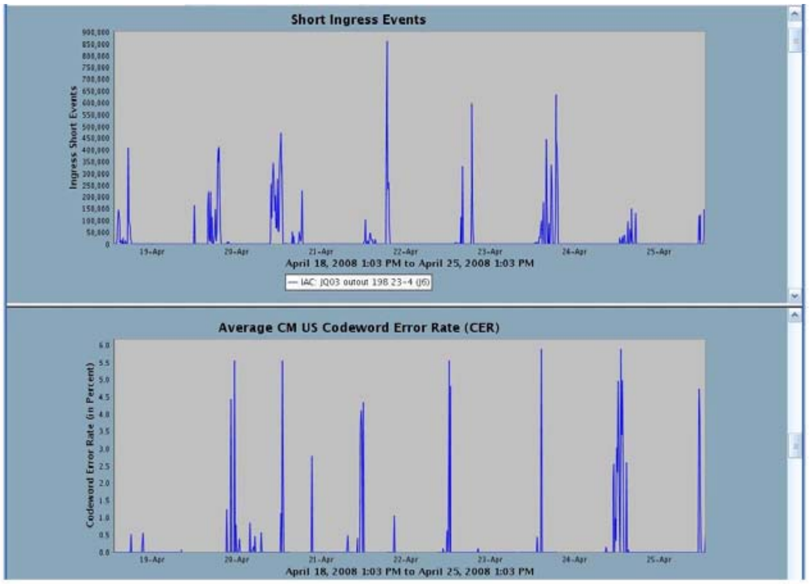

Chronic Noise With Repeated Regional Service Impact

In the next few examples, we can see a 24-hour zoom of major degradations in service and plant performance. Most compelling to note is that each of these problems occurred primarily in one of 35 HMD locations on Charter's FA node, and more importantly, it seriously affected service for 2 to 3 hours. Due to its relative short duration, it was imperative to perform symptom event-triggered analysis to capture these problems, their severity and scope at the times they are occurring. Without this sort of implementation, i.e., with simple time-based periodic (e.g. 4 hour) status-polling from various systems, it would be near impossible to detect these situations, much less understand their true service impact.

In the first example to the right, we see high-intensity broad noise affecting one plant segment FA09 amp main output. Nominally, we see the SNR mean range near 31 dB, where as during the broad ingress detection, we see clearly the 10 dB drop as a result. Keep in mind that this average SNR plot is from only the modems off this same amp leg and therefore represents a true measure of impact.

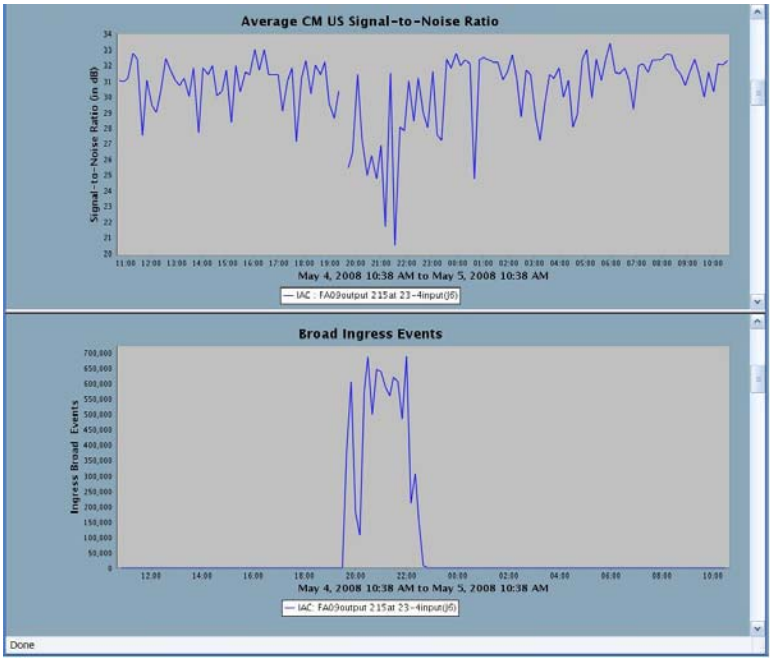

3-Hour Duration of Broad-Ingress and Ave CM SNR

Any other time-based analysis of SNR alone (even at a very high rate) that includes many more modems in the node would not detect this truly service-impacting problem. To illustrate, node FA has approximately 400 subscribers. So, if 20 modems have approximately 20 dB upstream SNR and the other 380 are at 31 dB, then the node's average of all CMs would go from 31 dB (before ingress effect) and (20*20+31*390)/400= 30.45 dB during, and then 31 after. This does not represent detectable variance (outside normal variation), and therefore undetectable by operations, and therefore un-worked. And as what was shown in the previous plot, that of repeated chronic, short duration (e.g. 2 hour) noise, this may then persist for days, weeks, or longer.

In the next plot to the right, we see a similar SNR drop from the modem group, also synched with sudden ingress (not shown) under FA amp leg FA08 output. Here, we see not just the effect on SNR, but also on packet loss. Similar to the analysis discussed on calculating SNR and CER on the entire upstream, without analyzing these effects by incorporating only packet counts at a high-rate from the modems experiencing the same noise effects under FA08 output during the time of noise occurrence, would lead to an undetectable (masked) service-impacting problem affecting all customers in a given neighborhood

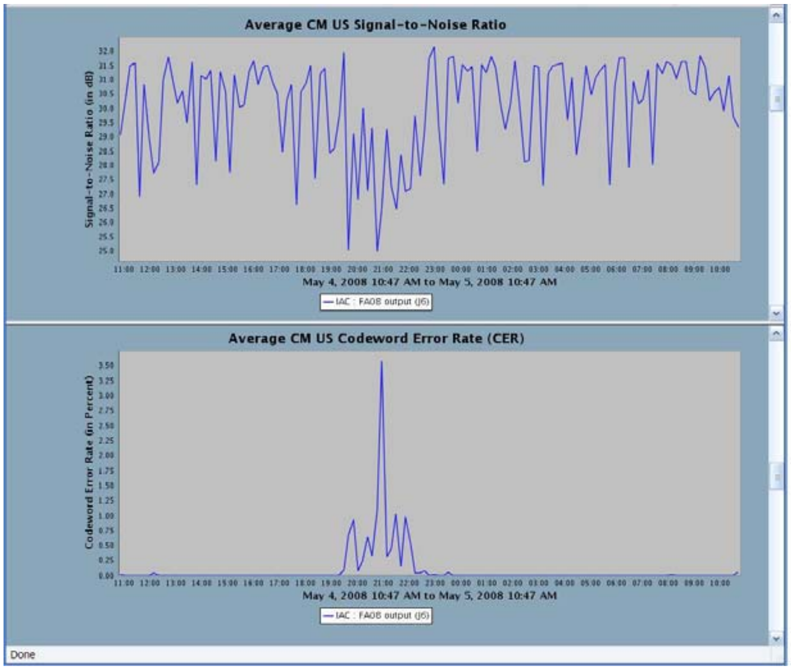

3-Hour Duration of Broad-Ingress and Ave CM CER

In some cases, noise and ingress can be correctable by the CMTS. Detecting situations of high noise-generated correctable packets and dispatching accordingly is an important part of a proactive maintenance plan. With this analysis, this system allows Charter to target major areas of concern, their location, and time of occurrence, to quickly repair plant problems *before* they impact the customers.

In the HMD performance plots shown here, synchronous correlation of the HMD-detected broad ingress events with the total errored (but corrected) codewords transmitted are shown (CCR of nearly 10%). As we have discussed, broad ingress typically is a result of CPD. However, in this case, the noise condition, although severe and chronic, is successfully being corrected at the CMTS. However, this would be an early warning that this noise condition would intensify to the point that error correction could not be performed, resulting in major packet loss as shown in the previous plots. This segment (Amp JQ01A main output) would be a priority target for proactive maintenance to repair the problem's cause.

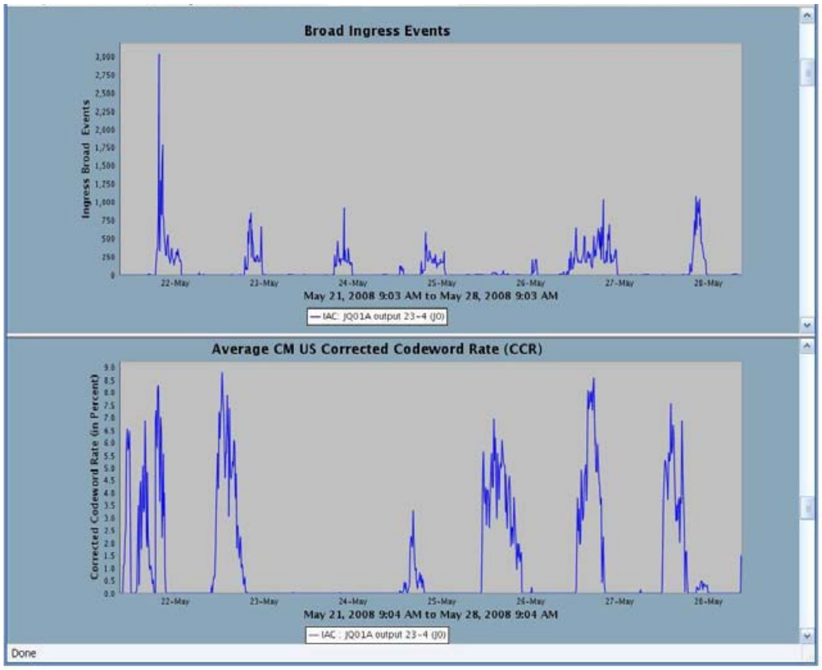

Cyclic Broad Ingress With Repeated High-Rate CCR

This system deployment at Charter realized:

- Increased efficiency in determining plant-related vs. premises-related problems when reported through customer service calls, with significant savings in resources spent per customer call

4. Automation of the Technical Operations Process with Complete OSS Integration

The previous sections in this paper have illustrated how intelligence built into the network enables a paradigm shift in cable technical operations. With precise, real-time and automated detection of symptom triggers, the system would provide a uniquely scalable, non-latent response to potential line, signal, and noise conditions (STEP 1).

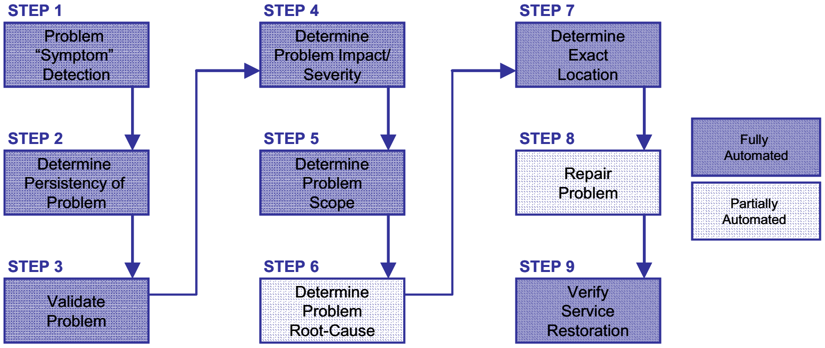

Now Automated: Key Service Assurance Operations Process

With continuous monitoring and analysis of symptom triggers, the system would provide automated, high fidelity, and highly reliable measure of persistency (STEP 2), ensuring that transient problems do not occupy technical operations staff.

Coupled with other powerful and unique features of this intelligent network, such as the automated discovery of IP devices and network topology and trigger-based high-rate analysis of service impacts, the service health management system would fully automate the many time-intensive steps of the process shown in the figure above, including:

- problem validation (STEP 3) through real-time, time-synchronized correlation system and independent health and status metrics,
- precise trouble location to precise amp leg (STEP 7) and scope of problem (STEP 5), easily and reliably determined, through automated and continually maintained network topology, and
- automation of the natural implementation of configurable technical operations logic based on operations policies or service levels, to determine problem impact and severity (STEP 4),.

The automation of STEP 4 has a two-fold operations benefit. First, all problems are not treated equally. Unacceptable service level impacting effects, based on accepted operations and business standards, are prioritized for immediate action. Given limited technical resources, the highest service-impacting conditions can quickly be determined and worked first. Secondly, it enables the technical operations team to progress with gradually tightening service level settings and resulting higher service delivery to the point of being mostly dispatched for proactive service management operations before a customer experiences a problem.

The system would allow threshold settings both on the initial base trigger events as well as the service level impact on affected customers before a dispatch alarm report (and/or ticket) is delivered. These were discussed in previous sections in detail.

- Loss of communications (from head-end to HMD)
- Upstream and downstream signal levels and attenuations
- Upstream noise detections and characterizations

Based on these initial triggers, real-time high-rate analysis of service performance would be initiated and then automatically analyzed against the time-synchronous HMD and modem data. The result analysis and report are then "pushed" to external 3rd party OSS systems in the form of SNMP traps, email, trouble tickets and general purpose XML messages (e.g. for integration into real-time mapping systems). In general, it is imperative that any system not have to be constantly viewed to detect these problem conditions, nor should it be polled in order to distribute real-time health information. Rather, it should automate all correlation logic, data analysis and reporting and push this complete package out to appropriate operations and dispatched personnel.

For a Loss of Communications event, key operations settings would include (at a minimum):

- Persistency duration
- Size of HMD plant region (number of modems)
- % of off-line CPE devices

When the HMD communication loss is detected, persisted, and the discovered cable modems are offline, a Service Impacting Outage report would be created and distributed. Below is an example alarm report for an outage event at Charter.

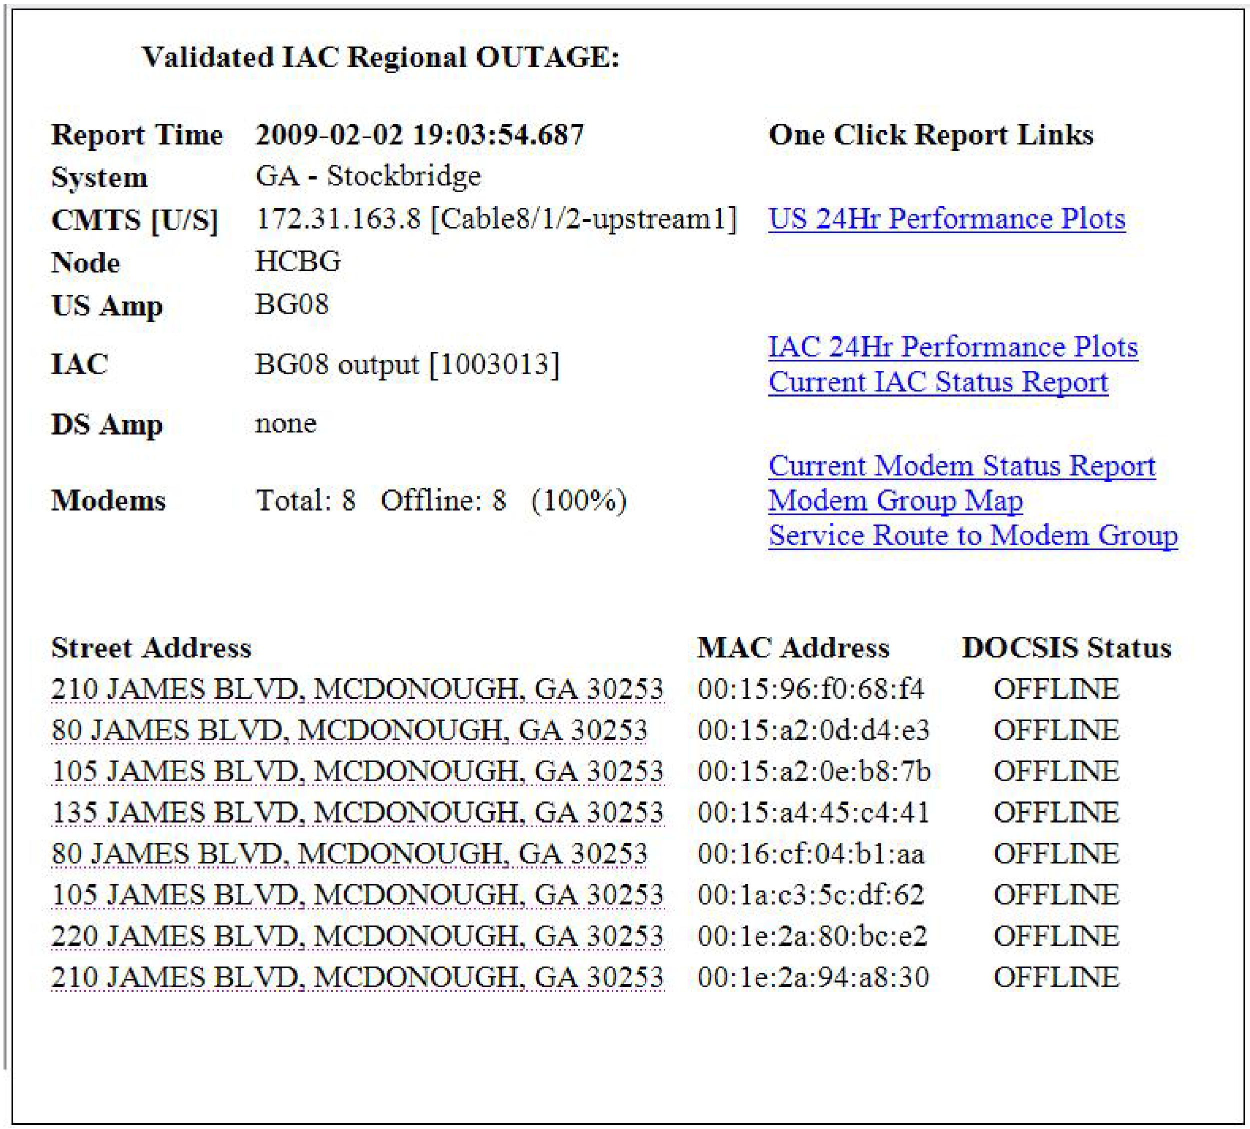

Validated IAC Regional OUTAGE:

Report Time	**2009-02-02 19:03:54.687**	**One Click Report Links**
System	GA - Stockbridge	
CMTS [U/S]	172.31.163.8 [Cable8/1/2-upstream1]	US 24Hr Performance Plots
Node	HCBG	
US Amp	BG08	
IAC	BG08 output [1003013]	IAC 24Hr Performance Plots Current IAC Status Report
DS Amp	none	
Modems	Total: 8 Offline: 8 (100%)	Current Modem Status Report Modem Group Map Service Route to Modem Group

Street Address	MAC Address	DOCSIS Status
210 JAMES BLVD, MCDONOUGH, GA 30253	00:15:96:f0:68:f4	OFFLINE
80 JAMES BLVD, MCDONOUGH, GA 30253	00:15:a2:0d:d4:e3	OFFLINE
105 JAMES BLVD, MCDONOUGH, GA 30253	00:15:a2:0e:b8:7b	OFFLINE
135 JAMES BLVD, MCDONOUGH, GA 30253	00:15:a4:45:c4:41	OFFLINE
80 JAMES BLVD, MCDONOUGH, GA 30253	00:16:cf:04:b1:aa	OFFLINE
105 JAMES BLVD, MCDONOUGH, GA 30253	00:1a:c3:5c:df:62	OFFLINE
220 JAMES BLVD, MCDONOUGH, GA 30253	00:1e:2a:80:bc:e2	OFFLINE
210 JAMES BLVD, MCDONOUGH, GA 30253	00:1e:2a:94:a8:30	OFFLINE

In the above example, the HMDs loss-of-communications event triggered a persistency test as well as synchronous analysis of the ping status of the modem group. In this case there were eight modems downstream from this HMD. All of them were reported as off-line. Since the HMD would be a cable line powered device, and assuming the use of backup power supplies, a loss of communications will not be the result of a residential power outage. In addition, due to the correlation of two independent sources, the HMD and the CMs, discounts any possible reporting fault by either one or the other. Thus, a reported outage represents a true, validated cable problem and can be dispatched immediately. Not only due to the deep scope of an HMD (focused on a small segment of the total node space), many smaller outages, only affecting a single amplifier leg are detected, that otherwise would go un-noticed until customers start calling in.

Specifically, the deployment at Charter realized:

- An average reduction of 50% in MTTR for plant outages

For Signal Level events, key operations settings would include (at a minimum):

- Base HMD signal level variance (from post balancing baseline)
- Persistency duration (of level variance)
- Size of HMD network region (number of modems)
- Resulting service impact on HMD group modem performance over exact analysis period (base list, other levels are also provided).
  - Average CM US and DS modem SNR variance
  - Average CM US transmit and DS receive power variance
  - Worst average CM CER and CM CCR %

When the HMD pilot level threshold violation is detected, persisted, and the discovered cable modems performance metrics exceed acceptable service levels, a Service Impacting Signal report would be created and distributed. This process would be completely automated. Below is an example alarm for a signal-triggered event at Charter.

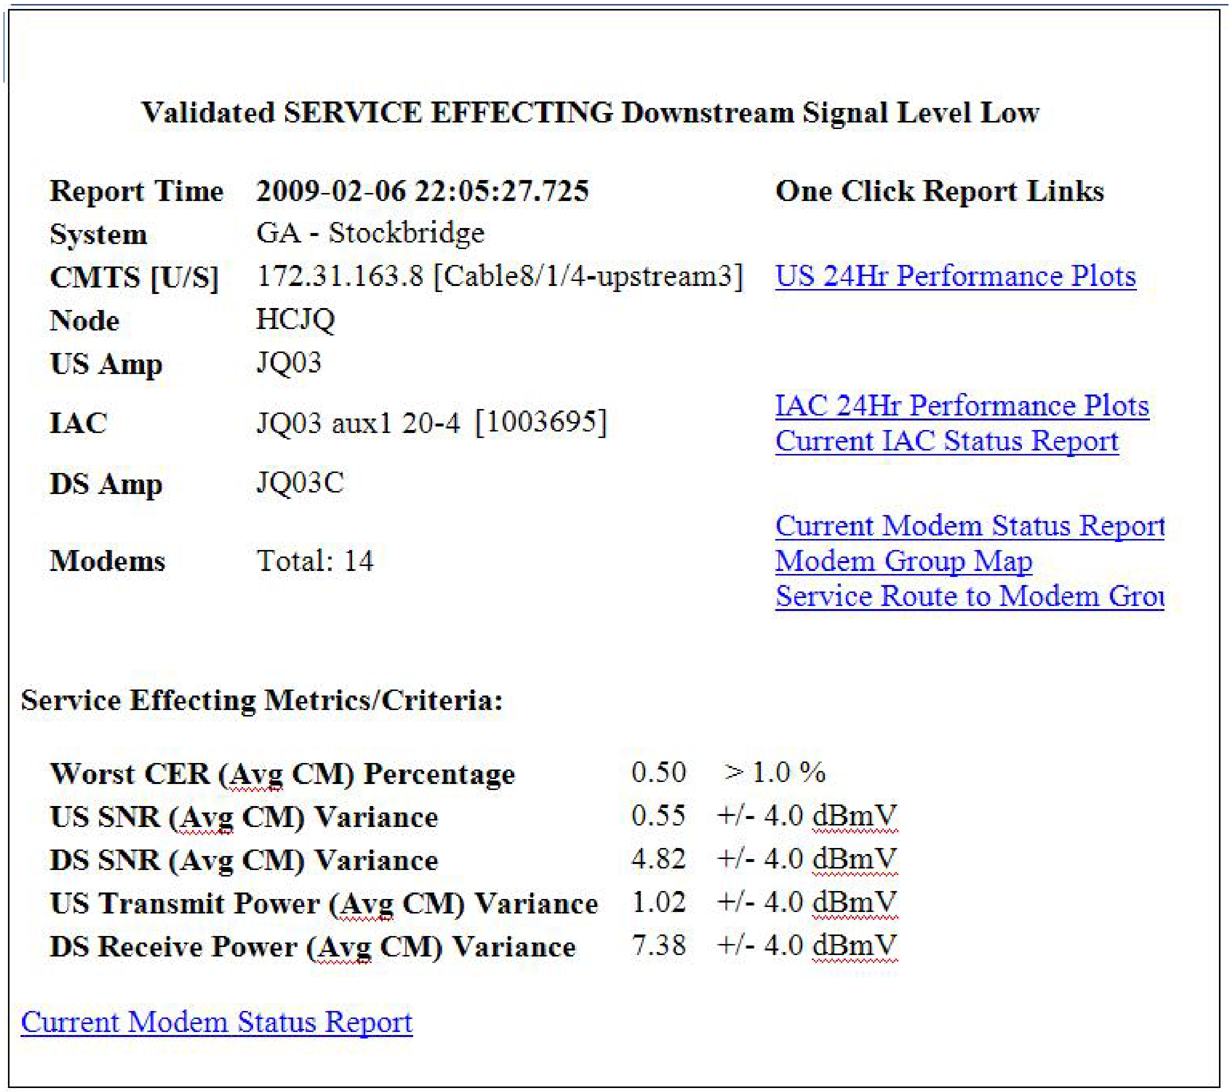

Validated SERVICE EFFECTING Downstream Signal Level Low

Report Time	**2009-02-06 22:05:27.725**	**One Click Report Links**
System	GA - Stockbridge	
CMTS [U/S]	172.31.163.8 [Cable8/1/4-upstream3]	US 24Hr Performance Plots
Node	HCJQ	
US Amp	JQ03	
IAC	JQ03 aux1 20-4 [1003695]	IAC 24Hr Performance Plots Current IAC Status Report
DS Amp	JQ03C	
Modems	Total: 14	Current Modem Status Report Modem Group Map Service Route to Modem Grou

Service Effecting Metrics/Criteria:

Worst CER (Avg CM) Percentage	0.50	> 1.0 %
US SNR (Avg CM) Variance	0.55	+/- 4.0 dBmV
DS SNR (Avg CM) Variance	4.82	+/- 4.0 dBmV
US Transmit Power (Avg CM) Variance	1.02	+/- 4.0 dBmV
DS Receive Power (Avg CM) Variance	7.38	+/- 4.0 dBmV

Current Modem Status Report

In the above example, an initial trigger event was created from the HMD's downstream pilot power level dropping below the given threshold (set at ±3 dB). The resulting automated analysis included a persistency test (e.g. 2 minutes) and then a subsequent service impact test against the discovered modem group. Using the operational defined Service Level criteria (an effect of 1.0 % CER or 4 dBmV on power or SNR), it is shown, that this signal level event experienced by the HMD, is synchronous with an almost 5 dB change in average downstream SNR (of all 14 modems) and a 7 dB change in receive power.

Similarly for Ingress/Noise events, key operations settings would include (at a minimum):

- Base HMD ingress rate (sudden occurrence of high-rate noise)
- Persistency duration (of high ingress/noise)
- Size of HMD plant region (number of modems)
- Resulting service impact on HMD group modem performance over exact analysis period (base list, other levels are also provided).
    - Average CM US and DS modem SNR variance
    - Average CM US Transmit and DS receive power variance
    - Worst average CM CER and CM CCR %

When the HMD ingress/noise level threshold violation is detected, persisted, and the discovered cable modems performance metrics exceed acceptable service levels, a Service Impacting Signal report would be created and distributed. Again, this process is completely automated. Below is an example alarm for an ingress-triggered event at Charter.

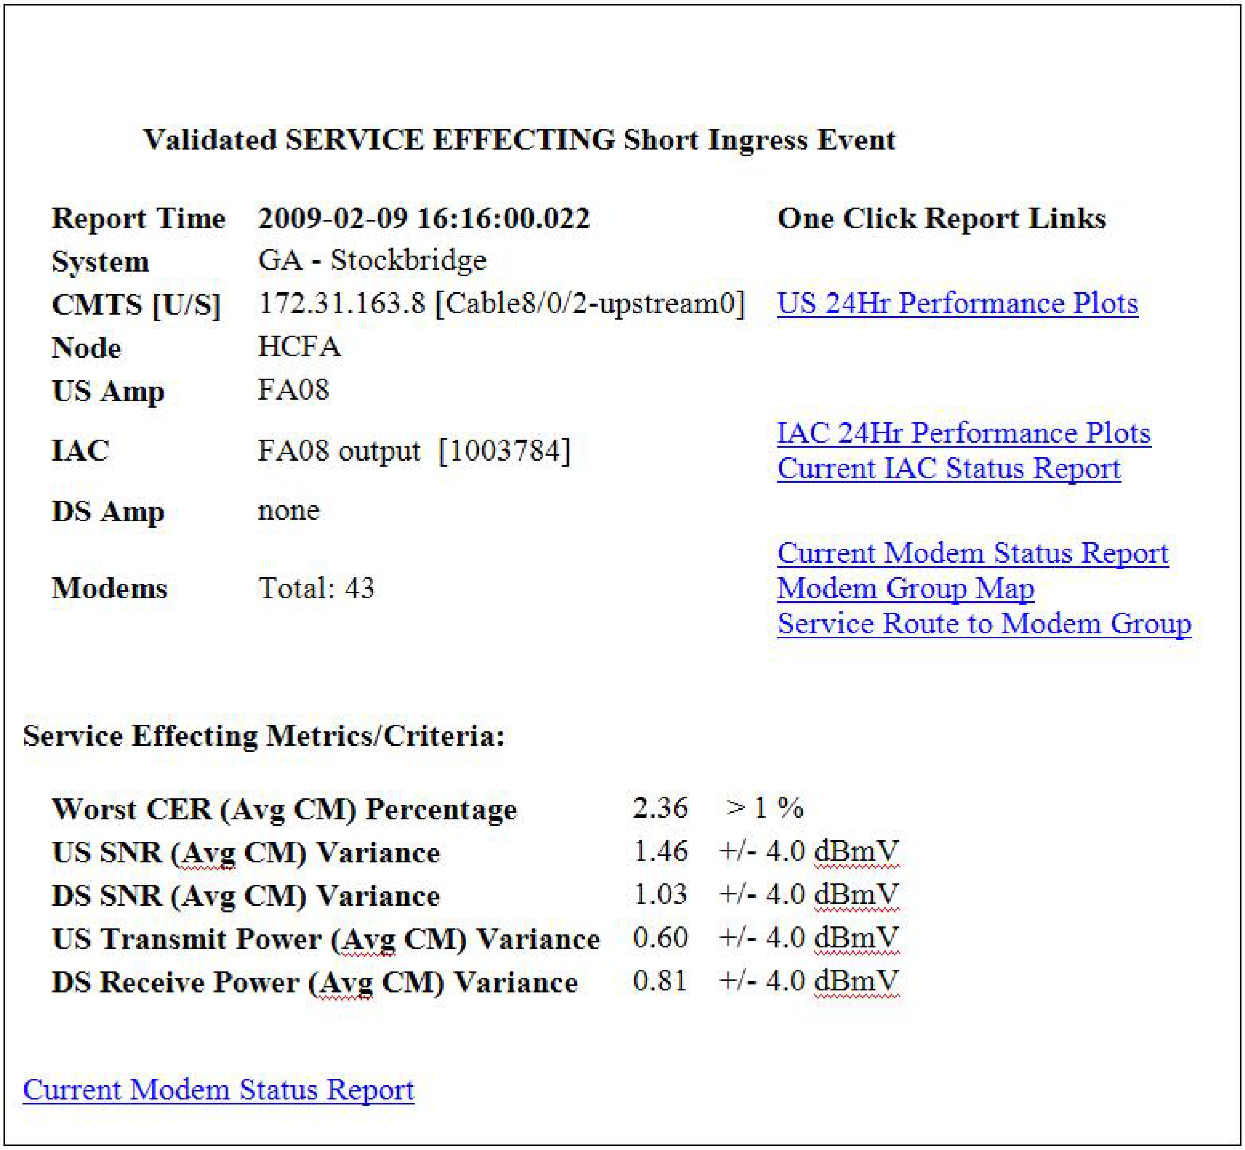

Validated SERVICE EFFECTING Short Ingress Event

Report Time	**2009-02-09 16:16:00.022**	**One Click Report Links**
System	GA - Stockbridge	
CMTS [U/S]	172.31.163.8 [Cable8/0/2-upstream0]	US 24Hr Performance Plots
Node	HCFA	
US Amp	FA08	
IAC	FA08 output [1003784]	IAC 24Hr Performance Plots Current IAC Status Report
DS Amp	none	
Modems	Total: 43	Current Modem Status Report Modem Group Map Service Route to Modem Group

Service Effecting Metrics/Criteria:

Worst CER (Avg CM) Percentage	2.36	> 1 %
US SNR (Avg CM) Variance	1.46	+/- 4.0 dBmV
DS SNR (Avg CM) Variance	1.03	+/- 4.0 dBmV
US Transmit Power (Avg CM) Variance	0.60	+/- 4.0 dBmV
DS Receive Power (Avg CM) Variance	0.81	+/- 4.0 dBmV

Current Modem Status Report

In the above example, an initial trigger event was created from the HMD's detection of the sudden on-set of persistent short duration (less that 50 μs) ingress. The resulting automated analysis included a subsequent service impact test against the discovered modem group of 43 modems. Using the operational defined Service Level criteria (an effect of 1.0 % CER or 4 dBmV on power or SNR), it is shown that this ingress detection event experienced by the HMD is synchronous with an average of 2+ percent in packet loss across the entire 43 device modem group under amplifier leg FA08 main output.

With this automated, service impact awareness, technical operations are quickly "tuned-in" to a *current* service-impacting condition. By looking at the current real-time measurements of noise focused at the precise map leg location, we know where and when (and when not) to dispatch techs, as illustrated below.

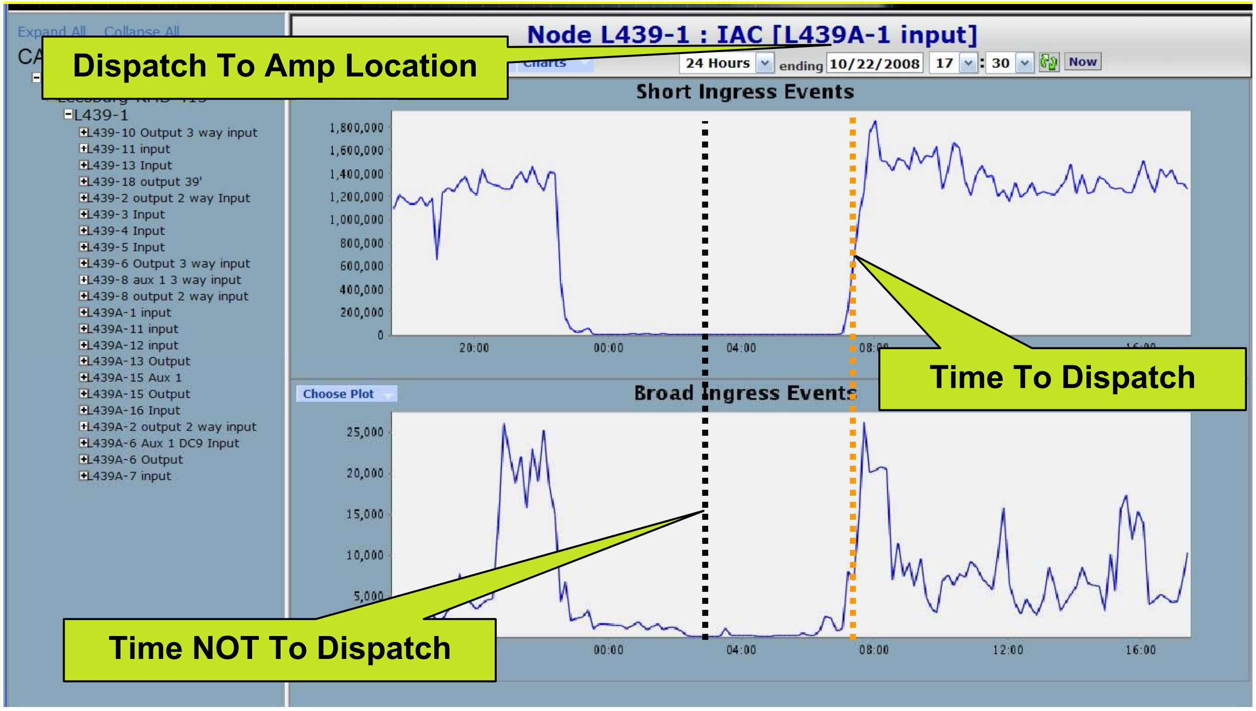

Dispatch To Amp Location

Time To Dispatch

Time NOT To Dispatch

On top of the automation of STEPs 1-7 as illustrated above and triggered *only* on the qualification of STEPs 1-7, the system would continue automated high-rate post alarm (and repair) analysis of the service effecting criteria for a configurable period of time. In this way, not only are the detection, propagation, and scope captured, but so is validation of the service restoration (STEP 9), which then a notification (i.e., clear of service impact) can be sent out accordingly.

The system would represent the industry's first system that automates the service impact analysis, correlation, and dispatch to a precise segment of the plant. Only through this analysis can problem symptoms, whether detected from the RF side or the IP side, provide a true measure of relative qualification from one segment to another, providing the utmost in efficiency in operations resources. The key is being able to qualify all problems' symptom conditions, using localized and time synchronous service impact measurements.

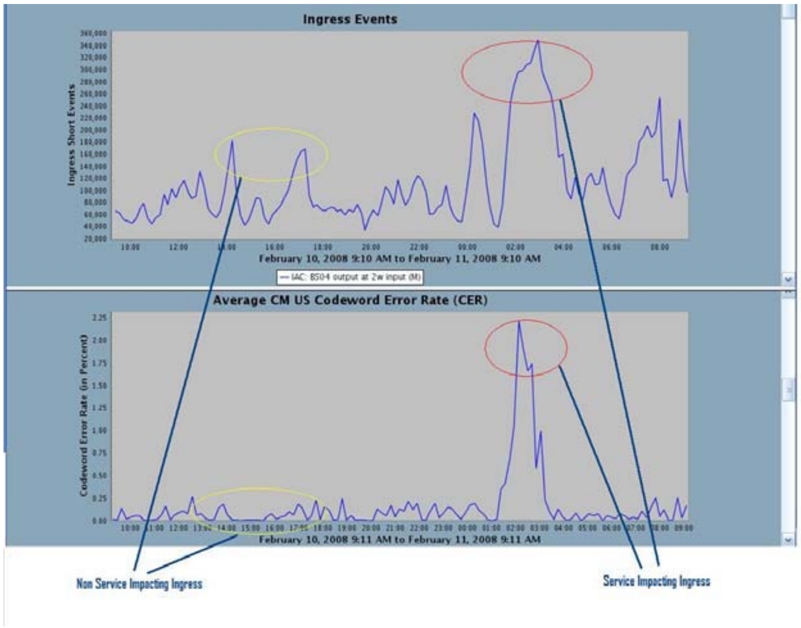

For example, in the plots to the right, we see that through the service-affecting criteria correlation, that although high-rate ingress is occurring earlier (to the left), it is not considered service impacting at a 1% CER service level threshold. However, as the noise is detected at higher rate with greater intensity, a definitive service impact is seen. The correlated metric could also have been CER, CCR or SNR. However what this illustrates is that changes in RF noise and attenuations are a regular, common fact in a cable system, but simply do not provide the necessary information of service impairment alone. Conversely as we have shown, periodic sampling of SNR from DOCSIS devices alone does not prove a service-impacting condition – even if we could poll all the modems all the time!

Single-dimensional systems – whether they are stand-alone RF monitoring, headend based spectrum analysis systems, DOCSIS MIB pollers or even technicians with the latest hand-held devices – are inadequate for cost-effective and continuous service health management. In summary, a system that employs deep real-time intelligence, automates nearly the entire operations process, and implements automated logic set by configurable operations settings, truly delivers an MTTR nearing MTTr with dramatically reduced costs.

The deployment at Charter realized:

- Implementation of a cost-effective, priority-driven, proactive maintenance program, focusing maintenance efforts to top service-impacting regions by amplifier leg, no longer focusing simply and grossly on full node performance data.

- A 75% reduction of resources required for identification, location, and repair of network-related plant degradations.

5. Automated Compensation and Containment

On top of removing the obstacles to reducing MTTR, a comprehensive service health management system should maintain the highest level of service performance during sudden and unpredictable severe degradations, through automated compensation and containment. This is a critical aspect (and quite possibly a minimum requirement) of any active service health management system. As broadband-based service operators deploy more and more complex service offerings, requiring continuous higher modulation and higher bandwidth transmissions, this will become imperative. These new services create a demand to utilize the full upstream spectrum for multiple DOCSIS channels or to run one or two channels at 64 QAM. Today, this is itself a fairly challenging task. What becomes more daunting is when the problem is compounded with the utilization of upstream channel bonding, available in DOCSIS 3.0.

The promise of DOCSIS 3.0 involves the ability to fully utilize a wide, multi-channel (bonded) bandwidth spectrum, on-demand, and on behalf of business-driven objectives. Such capabilities not only open the door for faster downloads (e.g., for push content), but also enable bandwidth availability for high-quality upstream video for video conferencing, home and business security monitoring, and other advanced services not practical before.

The challenge with channel bonding, unlike that of single channel higher modulation or session-based QoS, is that it requires multiple, continuous, clean bandwidth channels across the entire usable spectrum simultaneously. However, unpredictable impulse noise and other impairments in the spectrum have unpredictable impacts on high-bandwidth two-way services. Many upstream DOCSIS channel utilizations today are allocated high in the sweet spot, above 30 MHz. In most cases, the channel can be maintained at a high enough signal-to-noise ratio to assure reliable 16 QAM transmission without performance dropouts. However, achieving the required margin for 64 QAM becomes somewhat challenging and elusive, even at this upper spectrum. If achieved, it may offer only the ability to add more devices per channel (i.e., capacity), but not necessarily a major increase in any one customer’s bandwidth (i.e., throughput).

In order to reliably deliver dynamically allocated, high-throughput, continuous bandwidth with high-quality (i.e., low packet loss) through channel bonding, the margin of carrier-to-noise (C/N) must be continuously maintained. And, it must be maintained across multiple, simultaneous channels using the full spectrum, including that traditionally shied away from or completely avoided between 15 and 30 MHz. Changes in noise and signal levels and their effect on service needs monitored, minimized, and if too severe, quickly located and repaired. Traditional methods to maintain this reliability through refined plant-quality maintenance procedures have provided good results in some cases. However, given the HFC plant is such a diverse and dynamic network, such methods are extremely difficult and costly – if not impossible – to maintain and institutionalize.

The following plot of a sample node in the Charter deployment shows a spectrum analyzer trace of the DOCSIS upstream. Taken from the headend, the plot shows graduated noise amplitude increasing toward the lower spectrum. This trace illustrates a single DOCSIS transmission and appears to be achieving required C/N for acceptable error free transmission. In addition, it appears that to the left (higher in frequency) from the transmitted DOCSIS channel, a second channel could be transmitted, again with acceptable SNR margin.

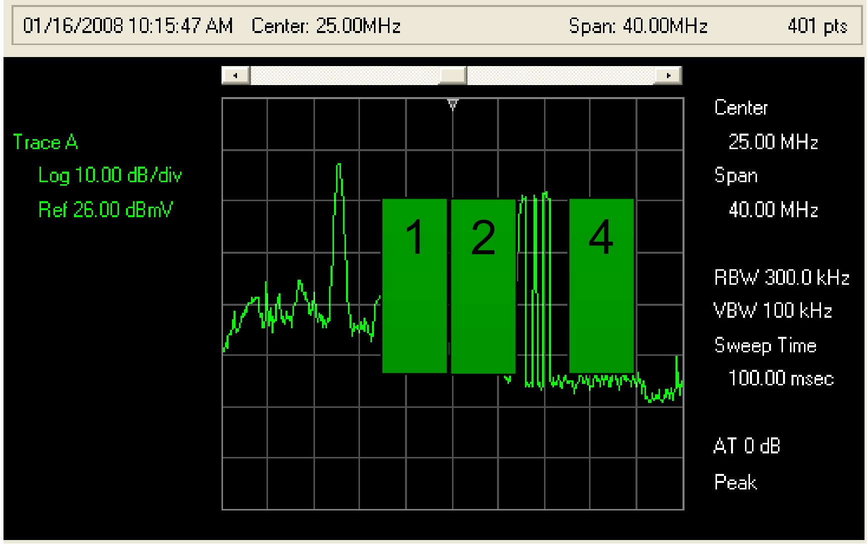
01/16/2008 10:15:47 AM Center: 25.00MHz Span: 40.00MHz 401 pts

Trace A
Log 10.00 dB/div
Ref 26.00 dBmV

Center 25.00 MHz
Span 40.00 MHz
RBW 300.0 kHz
VBW 100 kHz
Sweep Time 100.00 msec
AT 0 dB
Peak

Multi-Channel Bonding Not Feasible With Ingress

To the right (i.e., lower in the spectrum), however, we see that with an attempt to run two additional channels (i.e., to form a four-channel bonded transmission as illustrated), a C/N margin is insufficient. Thus, effective loss-less transmission will not be provided for these channels. Therefore, high-bandwidth transmission would not be achieved at this given time. Due to the packet multiplexing nature of channel bonding, where a given stream's packets are divided across each of the four channels, even if two channels are perfect, significant loss in any one channel will proportionately degrade the entire stream's performance. In this case shown above, under a DOCSIS 3.0 four-channel bonded transmission, there might be anywhere from 5-25% packet loss in aggregate. And, since the onset of noise and ingress – especially impulse noise – is not predicable or constant, the effective and continuous delivery of high-bandwidth, low-error rate could not be reliably delivered or guaranteed.

Here, built-in features of a more intelligent deep and self-managed access network system become critical: active ingress suppression technology. This ingress suppression capability continuously analyzes the (RF) and disconnects parts of the return path where there is no traffic. It would operate similar to an Ethernet switch, connecting cable modems (paths), when needed, and allowing transmission to the headend/CMTS while the other modems/paths remain disconnected.

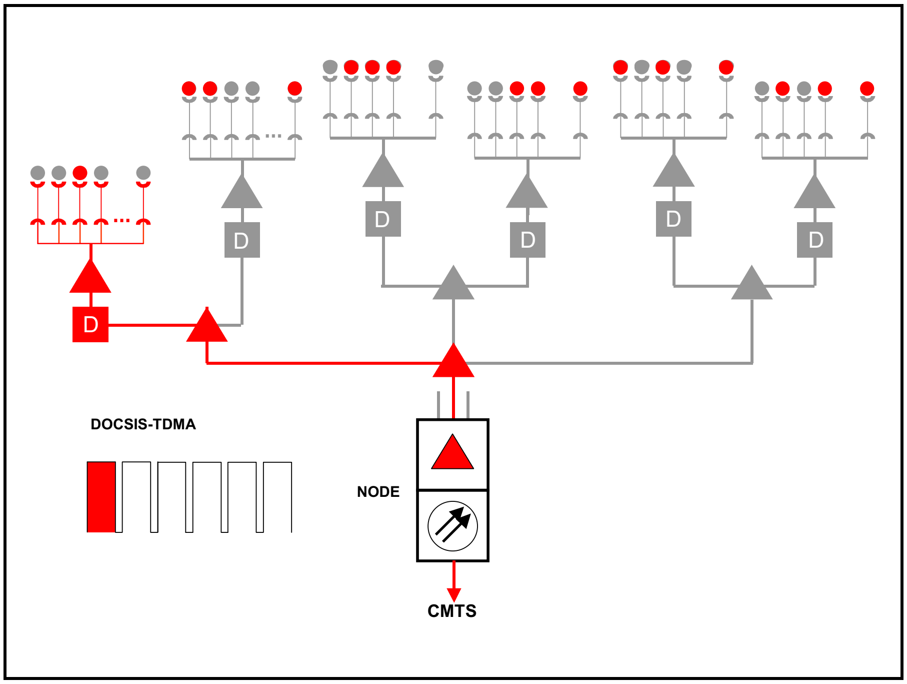

One DOCSIS Carrier/Device Connects One Network Leg

The ingress suppression would be accomplished by HMD electronics. This would provide exclusive fast-switching access of the return path, one RF segment at a time. As such, it enables a general reduction of contributing energy and noise from n-1 RF network segments while one segment is transmitting (where n is the number of independent locations of an HMD within a given or combined fiber node region).

With continuous ingress reduction, C/N improvement in a typical cable network can be significant. For lower frequencies (<25 MHz), the noise has a more transient behavior (i.e., short spikes with high amplitude). For higher frequencies (>25 MHz), the noise has a more thermal and continuous behavior. Noise added from a large number of homes, however, will behave more like the high frequency noise from a single home. The reason is that low frequency spikes will add up in numbers, but not so much in amplitude (given several spikes will not often coincide in time). The result will be a more continuous low frequency noise floor. The high frequency noise (or pick up from transmitters) will simply add in power and also appear as a constant noise floor. For a large node, the “tilt” of the noise power as a function of frequency could be about 10-15 dB, meaning that low frequency noise of 5-15 MHz is about 10-15 dB stronger than noise at 40-65 MHz (e.g. for Europe).

Through continuous noise floor and impulse noise reduction across the entire upstream spectrum, an active service health management system would enable the continuous C/N margin required to run higher modulations, such as 64-QAM. In addition, when spanning the spectrum with multiple upstream bonded channels for DOCSIS 3.0, the system allows uninterrupted, highly reliable utilization to achieve the resulting expansion in bandwidth delivery, as illustrated in the figure below.

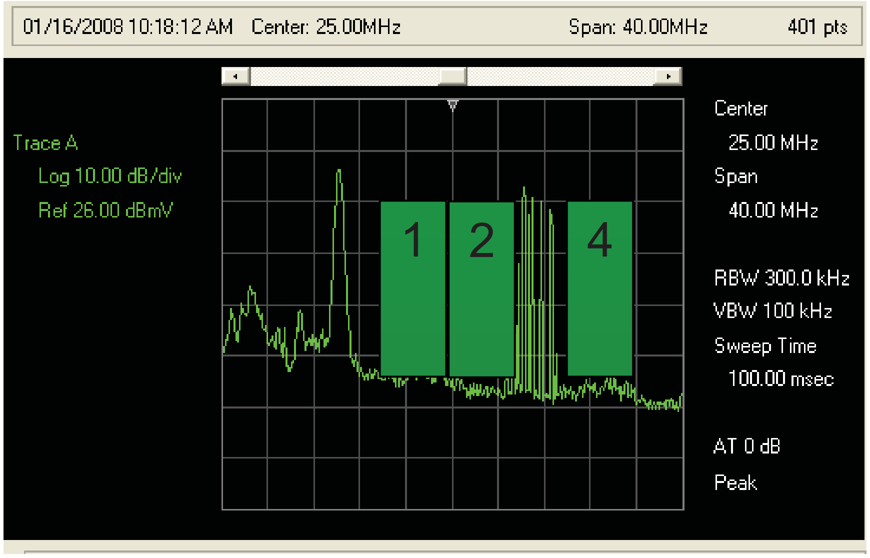

01/16/2008 10:18:12 AM Center: 25.00MHz Span: 40.00MHz 401 pts

Trace A
Log 10.00 dB/div
Ref 26.00 dBmV

Center
25.00 MHz
Span
40.00 MHz

RBW 300.0 kHz
VBW 100 kHz
Sweep Time
100.00 msec

AT 0 dB
Peak

DOCSIS 3.0 Multi-Channel Bonding Assured With Ingress Suppression

Without active ingress suppression and active ingress detection and repair through the same processes defined for reducing MTTR, a DOCSIS 3.0 bonded channel set will ultimately becomes vulnerable to it weakest link / channel.

What this also means is that unmanaged ingress and noise effects will now take the form of a service outage (unusable performance). What this also means is that operations processes and tools need to evolve accordingly. This real-time service health management capability would provide continuous bandwidth assurance for the delivery of high upstream, even in the face of the inevitable and unpredictable noise conditions that are part of the everyday and on-going life of a RF network. Active bandwidth assurance and containment of sudden impairments, coupled with a program of rapid identification, prioritization, and repair are essential to delivering and maintaining the dramatic improvement in bandwidth capacity enabled by DOCSIS 3.0.

Specifically, the deployment at Charter realized:

- Instantaneous and continuous maintainability of at least 2 dB and as much as 7 dB improvement in upstream CNR across four or more upstream channels.
- Delivery of a continuously maintained clear path, even in the face of sudden and unpredictable RF conditions, assuring reliable, high-upstream bandwidth service delivery for future deployment of DOCSIS 3.0.

6. Neighborhood Health Versus Node Health: The New Operations Paradigm

Charter's deployment provided convincing proof of the power of proactive monitoring and management for system maintenance at the neighborhood level. The system generated the efficiencies in workforce management by isolating the type of problem to at least an amplifier level location, but in most cases to an individual distribution leg. By viewing the network as a collection of "neighborhoods" rather than nodes, it illustrated for Charter the tremendous improvements in productivity and time to repair.

Like most cable operators, Charter has traditionally utilized a method of problem aggregation to create a "score card" to rate and rank physical fiber node segments of various systems. In this way, performance as a whole can be evaluated and compared from one node to another, one system to another, as well as one division to another. Known as Node Health score or similar names at individual operators, this method allows both the evaluation of the network's aggregate performance as well as the operations group that manages them. It involves taking various statistics, concentrated within a given fiber node region and the customers in it, and assigning a "weight" to each individual statistical category. In other words, customer reported trouble tickets may have a higher weight against a healthy node score than say, a modem's upstream transmit power being above acceptable ranges.

Weighted Aggregate Per Node (μ/S)			
Network Performance Status (DOCSIS)	Voice Service Related Performance Stats	Video Service Related Performance Stats	Customer Reported Trouble Stats

Classic Node Health Score Metrics

A node health information "aggregator" serves as a data miner of sorts, periodically (say, every 4-24 hours), polling various independent monitoring systems, including customer trouble call databases, voice performance metrics, video performance metrics, and out-of-range or exception conditions from the standard DOCSIS modem and CMTS pollers. A combination of all of these attributes, relative to the total size of the node (i.e., subscribers), determines an aggregate Node Health score. Ranking these scores against each other can therefore determine the worst performing (least healthy) nodes. Combining averages over a given system can determine the worst performing system, and so on. This method has also been used to evaluate technicians' and supervisors' performance levels, and even become part of an individual's performance bonus objectives.

The good news is that this approach created a universal measuring stick for management to measure and reward operations performance. In addition, it helped proactive maintenance to move away from a pure scheduled maintenance process to prioritized plant maintenance process focused on the worst 10% of the nodes in any operations group's domain of responsibility.

The bad news, however, is that often *only* these 10% of the nodes with the worst health scores actually get worked, beyond the ones with major customer reported trouble tickets. As detailed previously, due to the dilution of metrics resulting from too large of a sample set (i.e., across the whole node), many isolated, poorly performing neighborhoods are rarely, if ever, being worked. This is due to the fact that a given pocket of 20 customers off a severely degraded amplifier leg that is part of an otherwise healthy node of another 180 customers, would never rank as one of the worst top 10% as a node

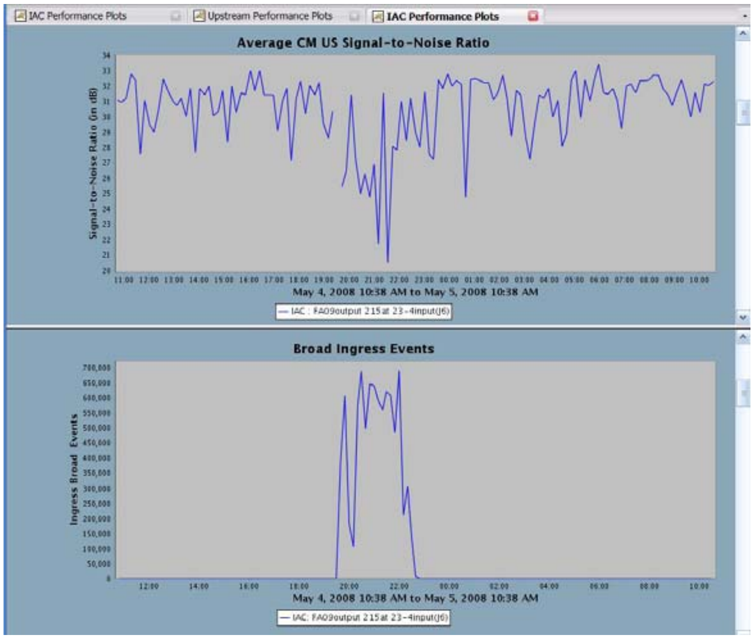

Example of Worst-Amp Segment Previously Undetected aggregate score – even though among amplifier segments it would rank as by far the worst segment *within* the node and possibly among the worst over all node segments. An example of this situation is captured in the screenshots below.

To complicate the situation further, the Node Health Score does very little to help field technicians locate and repair a problem. Cable operators may know who calls in, which modems have had recent symptomatic data, but have no ability to know "when, where, and what" based on Node Health Score alone. Technicians must begin the frustrating manual search and test process with hand-held devices, starting at the node, to locate and repair the root causes. All the while, those 20-home clusters on the "moderate" nodes experience chronically bad service.

The deployment by Charter illustrated a new and more accurate approach, which would calibrate the scoring at the same level as the HFC plant's elements (i.e., amp and amp leg). In this way, the evaluation fits more closely the actual physical architecture and has a direct connection to the small geographical (topological) area upon which a technician would focus their efforts.

If we take the same individual metrics gathered for node health (e.g., specific customer calls, performance reports, and customer device status information) and then use the discovered network topology information, we could sort this data by these individual amplifier CPE groups (cable modems, MTAs, set-tops). By also combining other metrics specifically targeted at the amp leg, such as high noise and out of range signal levels provided by real-time deep intelligence provide by Health Management Devices, we create a measure of health that is much more precise and actionable – an amplifier *neighborhood* health.

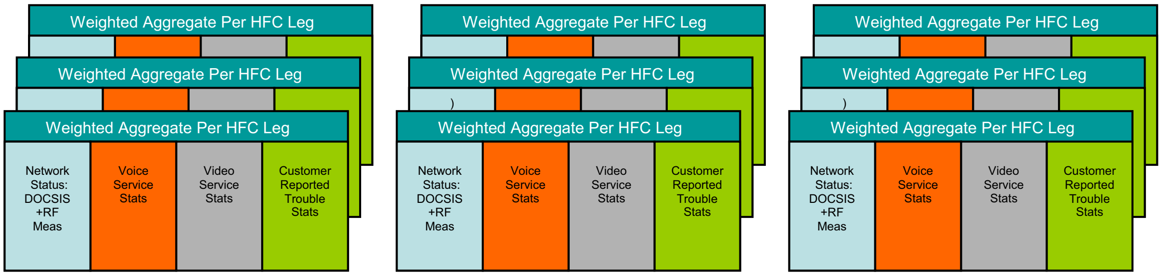
Weighted Aggregate Per HFC Leg
Weighted Aggregate Per HFC Leg
Weighted Aggregate Per HFC Leg
Network Status: DOCSIS +RF Meas
Voice Service Stats
Video Service Stats
Customer Reported Trouble Stats
Weighted Aggregate Per HFC Leg
Weighted Aggregate Per HFC Leg
)
Weighted Aggregate Per HFC Leg
Network Status: DOCSIS +RF Meas
Voice Service Stats
Video Service Stats
Customer Reported Trouble Stats
Weighted Aggregate Per HFC Leg
Weighted Aggregate Per HFC Leg
)
Weighted Aggregate Per HFC Leg
Network Status: DOCSIS +RF Meas
Voice Service Stats
Video Service Stats
Customer Reported Trouble Stats

New Neighborhood Health Score Driven Process

From this weighted score, we can now rank the amplifier segment regional performance, from one to another. Instead of then working the worst 10% of the nodes in a given 200-node system, cable operators can prioritize maintenance activity in the 20 worst *neighborhoods* across what may be 2,000 individual node (amp) segments. With this approach, efforts are focused exactly in the locations where the problems are, with resolution occurring in a more timely fashion and the entire plant operating at higher performance levels because the metrics are not obscuring problem areas.

Summary Results of Re-Inventing Technical Operations With Built-In Intelligence

A deployment of a new intelligence-based system with operational focus at the neighborhood level by Charter showed remarkable results in improving operational efficiency, as presented throughout this paper:

- An average reduction of 50% of in MTTR for plant outages
- A 75% reduction of resources required for identification, location, and repair of network-related plant degradations
- Increased efficiency in determining plant-related vs. premises-related problems when reported through customer service calls, with significant savings in resources spent per customer call
- Implementation of a cost-effective, priority-driven, proactive maintenance program, focusing maintenance efforts to top service impacting regions by amplifier leg, no longer focusing simply and grossly on full node performance data
- Instantaneous and continuous maintainability of at least 2 dB and as much as 7 dB improvement in upstream CNR across four or more upstream channels
- Delivery of a continuously maintained clear path, even in the face of sudden and unpredictable RF conditions, assuring reliable, high-upstream bandwidth service delivery for future DOCSIS 3.0 deployment.

Appendix: Acronyms

CCR:	Correctable Codeword Rate
CER:	Uncorrectable Codeword Error Rate
CM:	Cable Modem
CMTS:	Cable Modem Termination System
C/N:	Carrier-to-Noise Ratio
CPD:	Common Path Distortion
CPE:	Customer Premise Equipment
DS:	Downstream
FSK:	Frequency Shift Keying
HDTV:	High-Definition Television
HMD:	Health Management Device
IP:	Internet Protocol
KMA:	Key Market Area
MCTR:	Mean Cost to Repair
MTTC:	Mean Time to Close
MTTD:	Mean Time to Detect
MTTL:	Mean Time to Location/Cause
MTTP:	Mean Time to Persist
MTTr:	Mean Time to (Implement Identified) Repair
MTTR:	Mean Time to Repair
MTTS:	Mean Time to Scope/Severity
MTTV:	Mean Time to Validate
RF:	Radio Frequency
QoE:	Quality of Experience
QoS:	Quality of Service
SNR:	Signal-to-Noise Ratio
US:	Upstream
VOD:	Video on Demand
VOiP:	Voice Over Internet Protocol