NON-INVASIVE NETWORK MANAGEMENT

Bruce F. Bahlmann
Alopa Networks

Abstract

Broadband operators have extremely limited visibility into the health of Hybrid-Fiber Coax (HFC). Their limited visibility is due to several factors including the complexity of their transport medium, type of network, and sophistication of their back office. To navigate these factors, a more non-invasive approach to managing HFC is needed. This paper will present such an approach that is completely scalable, offers sophisticated location logic that enable one to quickly locate the most common types of HFC outages, utilizes a mere fraction of the bandwidth required by other network management efforts, and will fully integrate into existing top-level network management systems without additional stand alone status monitors.

Introduction

Broadband operators increasingly find themselves in a tough spot. Which is trying to figure out how to support and manage new services that are nearing deployment or have already been deployed. HFC being a complex animal as it is but now having to mange multiple services running on-top HFC leaves most broadband operators scratching their heads.

The fast paced deployment of new services is usually managed with a brand new Network Operations Center (NOC). These NOCs are quickly staffed with as many people (most internal) as it takes to provide 24-hour coverage 7 days a week. The result of promoting installers and plant operations personnel provides the broadband operator with increasing challenges in terms of organizing people and equipment to perform the minimalist amount of monitoring of the HFC and its increasing number of services – training is a huge issue.

The scarcities of technical employees skilled in network operations to design, build, and run broadband operator NOCs contribute to increasingly ugly statistics on building successful network management organizations. Over 70% of the attempts to initiate network management resort in failure.

Considering the cards stacked against broadband operators, they generally succeed in providing a minimal level of network management over HFC. However, the areas they face the most challenge are providing in-depth multi-service network management support, expanding their visibility beyond their backbone and hubs – down to their End of Lines (EOL), and developing more advanced associations (correlations) between related events. While most broadband operators would say they would like to explore these areas there is this problem regarding their obvious lack of commercially available tools as well as experienced network operations people to use them.

Background

Managing HFC is far from an exact science. In fact, HFC is so crammed with various types of technology, transmission media, and content that it is extremely tough to keep operational. It is also difficult for broadband operators to have visibility all the way down to their EOL – the point(s) at which each HFC node terminate. As a result of much of the HFC being invisible and complex, Broadband operators end up picking and choosing which components of HFC represent the most critical and monitor them with whatever tools are available. Basically the rule of thumb is the more customers that share

the same HFC component and/or transport the more critical it becomes (thus their current focus of the backbone and hubs).

Unfortunately, there are still not many tools available for managing HFC (the actual fiber nodes) – certainly not many that don't represent a completely new and often stand alone system. Since the last thing broadband operators are looking for is yet another monitoring system to drop into their already over crowded NOC, not many companies have gained traction with broadband operators for their HFC management products. This has left much of the market of managing HFC open. In fact, broadband operators to this day still do not have a cost effective way to manage all the way down to their EOL. The key here is cost effective – as some vendors have provided solutions for EOL monitoring but they are extremely cost prohibitive and not all together realistic to deploy operationally.

The focus of this article is to explore ways that broadband operators can gain more visibility into their HFC nodes including an example EOL monitoring system.

HFC Monitoring Challenges

The reliability of any network management system is directly dependent on the extent that it reaches out to all of its network elements. Customer Premise Equipment (CPE), Cable Modems (CM), Set Top Boxes (STB), and Media Terminal Adapters (MTA) represent attractive customer located network elements for NOCs to reach out and verify HFC health and availability. These installed customer network elements are placed throughout the network, represent no additional cost, and provide increasingly useful status and health information to a traditional Network Management System (NMS). However, traditional NMS rely on actively polling network elements to collect operational status of the environment where these network elements reside. Thus, current efforts to use customer network elements along with a traditional NMS fall short of the mark because they rely on active polling of these network elements. The method of active polling suffers from several issues that will be explained.

Scalability is one of the most obvious issues. Essentially, the sheer numbers of customer network elements can reach a point where it impacts the frequency that a single application can poll them on a regular basis within a timeframe that is worthwhile. As a result, the frequency that network elements would be polled by a traditional NMS would be increased so as to allow all these elements to be polled without impacting the network performance that is trying to be measured and monitored.

Because there are so many network elements to poll this also prevents one from being able to obtain much of any detail from each network element. The more information obtained from each element the greater the time it takes to collect this information, the greater each request impacts the usable bandwidth of the network it seeks to monitor, and the greater the impact on the network element performance it is attempting to use to monitor it.

Not being able to poll frequently presents another problem if using a traditional NMS. This is because the reason one polls network elements is to determine their current status and look for potentially service degrading behavior. This does not work well with HFC as it changes constantly. While much of these changes are tolerable (hardly noticeable operationally) the changes that can indicate much more serious problems are brewing also happen sporadically. Since a traditional NMS can only poll periodically it is likely to surmise that it will not be able to capture (or detect) these sudden changes and thus be unreliable in determining much more than trivial (on/off) status of the HFC.

Another problem with active polling is that unless similar network elements are polled

together (or within a reasonable time frame) the information gathered is useless across all similar network elements. For example, all network elements on a network can span several HFC nodes (i.e. they are combined). Unless all network elements are polled by-HFC-node and within a reasonable timeframe the information gathered may only indicated that something is potentially wrong with one of the HFC nodes. However, since none of this information can be collated by-node, the resulting data is unreliable and only marginally useful. NOCs that actually poll network elements on various nodes usually resort to managing subsets of network elements, if at all, so they can only look at very small-controlled samples of the customers.

Traditional NMS applications are best suited to manage network elements with static Internet Protocol (IP) addresses and are not capable of managing network elements with dynamic IP addresses. NOCs that do manage customer network elements must re-map the IP addresses of these network elements with every renumbering of the network. Since NOC resources are at a premium, this often results in fewer instances of monitoring customer network elements on each HFC node. This practice only leads to increasingly less monitoring of the HFC.

Monitoring customer network elements using a traditional NMS pays a heavy price on ones network because it uses Simple Network Management Protocol (SNMP) get. An SNMP get requires a NMS to send a question to a network element somewhere on the network. Each question in SNMP terminology corresponds to a specific Management Information Base (MIB) located in the network element's operating environment. Essentially, each network element maintains a wealth of MIBs each of which corresponds to some configuration or operational data stored locally on the network element. Depending on what is asked (i.e. which MIB(s) is/are requested in the get) the network element responds to the request by determining all the answers to these questions and then sends back a reply. As a result of this transaction, the network between the NMS and the network element pays twice for this transaction – once for the get and again for the reply.

Since not many customer network elements (e.g. a CM) actually have a need to communicate with the outside world they often fall off most routing tables. What this means is when some other system decides that what information they have is all of the sudden useful it must blaze a trail to each element from a networking perspective. This process of blazing a trail involves creating routing table entries for each network element. Routers that provide connectivity for each network element outside its network must re-learn about these network elements before communications can flow between the NMS and the network element. Each routing table entry requires the network element's router to Address Resolution Protocol (ARP) its physical address. ARP allows one to determine the mapping of IP address to Media Access Control (MAC) address (also called physical address). Once this table entry is created in the router it is able to relay packets to the network element. Albeit, this process is extremely quick (even in network time) this latency across all subscribers only further contributes to the inefficiency of using traditional NMS on network elements.

There is also this issue of determining the correlation between the network element and its associated real world information. One can usually derive certain things given pieces of information. For example, knowing ones IP address and subnet mask other information about the network can be derived. Likewise, knowing ones phone number or first and last name one can determine where one lives. However it is impossible to derive relationships between dissimilar or unassociated things. While there exists ways to complete these relationships they currently go beyond the capability of traditional NMS. For example, typical billing and customer care systems can

associate network elements with real world customer information. However, NMS do not provide hooks into such systems. Instead, these associations must be built manually – a very tedious and unmanageable process.

The NMS is also the wrong tool to manage HFC because it must monitor customer controlled network elements. That's right, all these network elements belong to the customer (or at least an increasing number of them do with the advent of retail CMs, etc.). So many of its requests to these network elements will not go through – such is the unpredictable nature of a customer-controlled device. As a result, it must actually ignore many of its responses because they will come up empty (or non-responding). In the early stages of deployment of CMs (or any other new technology) only a handful (if any) of network elements may exist on each node or Cable Modem Termination System (CMTS). What this means is that most (if not all) network elements may be down and this could actually represent 'normal' operating conditions.

Focus Areas

There are two general areas one needs to focus on when monitoring HFC. These general areas are:

- Monitoring HFC health
- Monitoring HFC EOL

There are "solutions" that claim to address both of these areas using one technology, but to do this right these areas actually represent two drastically different approaches that no one system can sufficiently achieve.

For example one system uploads tables of information specifically from their CMTS to provide some visibility into the HFC. However solutions like these are highly proprietary and will not work in a multi-vendor environment. They also provide only limited health monitoring, can impact the performance of their CMTS, and do not address EOL monitoring. Keep that in mind when looking at network management systems that claim to address both of these areas.

Since there are many reasons (several were discussed previously) why traditional NMSs are not up to the task of performing reliable HFC monitoring one must explore uncharted territory to achieve the visibility needed to manage/maintain high service quality.

Uncharted Territory

The most promising technology that can facilitate scalable visibility into HFC is actually quite old and extremely well established – the use of SNMP traps. SNMP traps provide a non-invasive way of monitoring the health of ones HFC plant because they are non-solicited. Network elements capable of SNMP traps are first configured to look for certain events/conditions (e.g. some threshold is reached), and then inform their configured trap host when these traps conditions are met. SNMP traps provide the following attractive features:

Completely scalable – Since SNMP traps use individual network elements to capture its information there is no need for any one application to perform direct polling on network elements. Once more, SNMP traps can be easily directed to any number of applications further distributing the load of handling all the traps across multiple servers. The distribution of this load can be easily added over time and more importantly in conjunction with the number of network elements.

Minimized network element performance impact – SNMP traps actually minimize the performance impact as compared with that of direct polling. That is because the network element only communicates an SNMP trap if a threshold is met rather than continually respond to SNMP requests. In the mean time the network element merely examines these thresholds along side its normal function.

Network elements monitor more states – Unlike direct polling which must carefully optimize what it request from network elements, SNMP traps can look at wider variety of MIBs. MIBs whose data only seldom changes are not good candidates for direct polling. However these make terrific SNMP traps as they send extremely valuable information about the network element at the time it occurs. This optimizes the collection of SNMP trap responses and allows applications receiving these traps to be overall more responsive to the needs of these network elements.

Network elements report independently – SNMP traps allow network elements to report their information independently rather than as the result of being polled directly. Reporting independently allows network elements to report the moment their thresholds have been met. This also allows applications receiving these traps to correlate multiple responses so as to determine the extent and severity of these traps.

Tolerant of dynamic IP addresses – Since SNMP traps come from network elements they are inherently tolerant to changes in the IP address on the network element. While this does force the application receiving the traps to be cognizant of each network elements' current IP address, this problem is much more manageable than the challenges that would exist to provide this functionality in a traditional NMS.

Half the bandwidth cost of SNMP gets – SNMP traps (like SNMP gets) communicate over something called the User Datagram Protocol (UDP) in networking world. However the SNMP trap does not require an acknowledgment from its destination. Therefore once a network element sends a SNMP trap it is done – no waiting around for some type of reply. From a networking perspective the use of UDP is an extremely efficient means of communicating. Note however that UDP requires a fairly reliable network to operate properly because if the network drops packets (i.e. it is unreliable) it may very well drop the SNMP trap message. In this case it is a self-fulfilling prophecy in that SNMP traps are used to make the network that much more reliable.

Tolerant of delays in routing – SNMP traps don't care about how long they take to reach their destination. Likewise, their originating network elements also don't care how long they take to reach their destination. As discussed previously, once they are fired out of the network element the communication is over as far as the network element is concerned – which goes on about its business (no matter what is left for the SNMP trap to negotiate to reach its destination).

Address correlation between element & customer – SNMP traps do not in themselves provide any real correlation between the network element and the customer. However what they do provide is a means of translating these traps at the application responsible for receiving the traps. At this point network information and real world information can easily meet. This particular area is where focus is needed to build the necessary relationships between the network element and the customer it represents. When these two pieces of information meet the result is an extremely powerful database capable of advanced reporting, troubleshooting, and modeling.

Tolerant of being customer controlled – SNMP traps can easily withstand having a customer connecting and disconnecting as well as powering up and down the network element. In the event the network element can communicate over the broadband medium its information is transmitted along side others. Should it be shut down or disconnected, it will not participate in the collective monitoring of the network health. In this way it makes no difference what so ever if the network element is on or off. Having it on would be great but if it is shut down it doesn't break anything or cause any false alarms.

SNMP traps form the basis for providing an excellent HFC health monitoring system. Combine this with an intelligent trap collection application that can be distributed and you have a fairly cheap means of monitoring HFC health.

Monitoring HFC Health

Monitoring HFC health is an evolutionary process. Just remember that managing HFC is far from an exact science. To do this right one needs many different sources to accurately track its health. One source can come from CMs, one from MTAs, yet another from STBs, and so on. It is important that when one allocates spectrum to these services that it strategically selects frequencies across the entire spectrum. Sticking to only those frequencies that are known to be good defeats the purpose of using customer network elements to monitor the health of ones HFC. If various customer network elements are not positioned across the available spectrum the use of SNMP traps will not be able to provide sufficient sampling to determine overall HFC health and should therefore not be used. Essentially this would merely provide visibility to a narrow portion of the overall spectrum. One actually needs several reference points (at least 3) across the entire spectrum to provide any kind of reliable HFC health monitoring.

Several ancillary benefits may be achieved out of a comprehensive use of SNMP traps to monitor the health of ones HFC. Some of these include:

Less reactionary network operations – Given the NOC now has visibility to significant changes on the HFC it can direct resources to make repairs before these changes become overly noticeable to customers.

Significant individuals can be more closely observed. Very important people (VIP) as well as past trouble makers can be more carefully watched without drawing attention as would be the case if the were modeled within an NMS.

Exploring more extensive use of SNMP traps will elevate their importance in the eyes of standards bodies and CableLabs. This will result in more specific trap requirements in future network elements geared more specifically towards monitoring network health.

Monitoring HFC EOL

Unfortunately, monitoring HFC EOL represents a totally different process than monitoring the health of HFC and is relatively void of cost effective solutions. This is because EOL monitoring concerns itself with connectivity and availability where as monitoring HFC health concerns itself mainly with reliability. Albeit these are somewhat related, connectivity and availability are distinctly different from reliability. Essentially monitoring HFC EOL insures the entire physical plant is available (no breaks, outages, etc.) all the way down to its EOL. Consequently the types of messages sent by EOL network elements are quite different than those sent by HFC health type network elements. This actually represents a relatively new area for broadband operators in terms of deploying/using extensive EOL monitoring. In fact, there are not many commercially available EOL monitors that will provide a cost effective solution to this problem.

Today's EOL monitors are extremely expensive (around $200 US), proprietary, and require some type of line voltage where they connect to the HFC. These requirements force many broadband operators away from the technology (even though they really would like to have it). Instead broadband operators seek to use other means of observing the HFC but none of which provide them with the same kind of visibility. The requirement of an available line voltage is also highly restrictive as this is not always available near EOLs – certainly not every EOL.

One cost effective EOL means of monitoring may be through the use of a miniaturized CM. The concept here is to drive chip technology to reduce the footprint of a CM down to something

that would fit in an enclosure no bigger than a line filter. In this case, a slimmed down CM is placed in a line filter. It is then powered by the 90v square wave that flows down the HFC for telephony. As a result one can have an operational standalone CM capable of sending and receiving information on the network with one slight exception. Instead of responding as a normal CM, these EOL CMs will provide some additional functionality. This includes the ability to perform predictable chatter. Chattering is a process where by the EOL network element periodically talks to its host. This predictable chatter allows the system to determine possible outages.

Any network element that is past due (not heard from) would need to be followed up using some basic form of direct polling. In this case direct polling is used to speed the resolution of the event. In other words, is the network element down or is some link down. Once this information is obtained one can forward this information to the broadband operator's NOC for analysis and resource assignment.

However new the concept of monitoring EOL is to the broadband operator its benefits outweigh the barriers needed to make this component a necessary part of the overall network management system. Yet broadband operators' general lack of interest in this component is likely the result its current lofty price tag. Until these components can get within the $20-30 US range (or lower) there use will not enter the main stream.

Bruce Bahlmann
Alopa Networks
bahlmann@bigfoot.com