

## **Keeping "S" In Your SNMP and "\$" In Your Pocket**

*By Orest Holyk, Product Development Manager and Jeff St. Denis, VistaLINK™ Specialist, Evertz Microsystems Ltd.*

*During the past several years, we have seen a number of papers and presentations introducing, describing, discussing, and prophesizing the merits of SNMP monitoring in broadcast facilities. Alternatively, there have been an almost equal number of reports highlighting some negative aspects of SNMP implementations. The question remains – is SNMP an effective and efficient signal monitoring tool and can it also be implemented as an equipment configuration solution within a broadcast environment? With reference to recent articles and existing facility SNMP case studies, this paper provides an overview of SNMP usage, dispels some common myths and addresses some common questions about ROI.*

It is quite common now to read about SNMP monitoring implementations within broadcasting. In a recent article by Jack Verner, VP and CTO of Digital System Technology, entitled “Effective Workflow Strategies for Digital Conversion”, SNMP is growing in popularity with the most practical implementation being new facilities. Further, the ability to display when and where equipment fails allows the engineer to pinpoint problems quickly and potentially protect against revenue loss, continuing that the investment in SNMP monitoring makes the most sense for larger facilities, where the possibilities for signal loss are more complex.

However, in most cases, a new facility is not in the capital budget, yet there is a growing need to implement SNMP monitoring on existing or expanding services, possibly consisting of both new and legacy equipment. In addition, the need to easily configure modules from a central, networked point vs. the traditional serial or card-edge interface is also desired. Subsequently, a cost-effective solution for both monitoring and configuration through SNMP is possible with easy-to-use and intuitive GUI interfacing, acceptable “reaction” times and minimal network traffic congestion, if any.



**Figure 1: Traditional signal monitoring has been further enhanced with multi-display virtual monitor walls and SNMP.**

Yes, it is still true that "seeing is believing", and with increased video, audio and data content, watching every picture across dozens of displays is no longer practical for any one operator. Traditional monitoring still exists and in some cases will not change, but as equipment costs continue to decrease and technology improves, small and large, older and newer facilities will consider new signal monitoring options. Even since a previously, SMPTE-published article entitled "Considerations for Multi-Channel Signal-Integrity Monitoring" (November, 2001)<sup>1</sup>, there has been considerable progress in technology, price and monitoring implementations. At that time, the conclusion identified SNMP as an effective tool for monitoring both incoming and departing signals from local signal monitoring equipment located at strategic points throughout a video enterprise network and

through this protocol, the operator effectively had additional monitoring "eyes and ears" in an ever-expanding digital television universe! This, along with configuration via SNMP is still true today.

Since that time, equipment monitoring through SNMP has unofficially evolved to provide "health" and/or "status" monitoring, whether it is situated at a local, central facility or at a remote site. Also, the scope of SNMP has expanded to include configuration of compatible gear all using the same protocol. In general,

<sup>1</sup> This paper introduces the basics and components of SNMP. It addresses security and interoperability of SNMP implementations and reviews signal monitoring goals in general.

SNMP is generally a “simple” protocol that can be used for both monitoring and configuration. But the primary advantage of SNMP is its interoperability among 3<sup>rd</sup>-party equipment manufacturers and software developers – it virtually bypasses the need to deal with proprietary protocols. However, along with SNMP proponents, there are some concerns as well as those who propose anything but SNMP since it is deemed to be an ineffective tool for broadcast facilities, preferring “tried and true” solutions instead.

**Introducing the Myths**

During recent customer visits, a number of common SNMP-related remarks were noted and are described here. Most feedback revolves around the “need for dedicated bandwidth” required to deploy a SNMP-enabled network, while others are unsubstantiated claims, which add to network “fear-mongering”. Some common quips about SNMP based solutions include:

- SNMP takes up too much bandwidth
- SNMP is OK for monitoring but not good for configuration due to latency and lost packets, and it also increases network traffic; there is no reliability for transmission of critical information. It should be referred to UTP – Unreliable Transport Protocol instead of UDP
- SNMP is too complicated and too expensive since dedicated hardware and customized software is needed with programmers on staff to make changes when I need them; SNMP software usually



**Figure 2: VistaLINK PRO PLUS – configuration, monitoring, 3<sup>rd</sup>-party interface, thumbnails, streaming on the same network and at your fingertips!**

ends up being site specific. There are many vendor-specific software (NMS) solutions, and I don't want to run another software GUI or dedicate another PC to the application

- MIBs are far from standard across vendors
- SNMP is insufficient for 3<sup>rd</sup> party, non-networked, legacy equipment

With several thousand VistaLINK (Evertz SNMP Monitoring and Configuration Solution) network nodes worldwide, monitoring tens of thousands video, audio and data signals, there is a vast amount of implementation experience to draw from. Subsequently the following sections attempt to demystify these common SNMP myths.

### **Bandwidth Requirements for SNMP Monitoring and Configuration**

'What size of pipe (i.e. how much bandwidth) do I need to dedicate to an SNMP monitoring solution?' is a commonly asked question. The simple answer is "it depends". More specifically, the amount of bandwidth that is needed depends on the volume and frequency of monitoring, and the "type" of monitoring. Configuration via SNMP tends to have little effect on network congestion since most parameters such as video or audio proc are adjusted occasionally, then left alone. For reference, Table 1 identifies typical byte sizes for SNMP configuration commands – GET, SET, and TRAPs.

<b>SNMP Operation</b>	<b>Size (bytes)</b>
GET Request (on integer)	89
GET Response (on integer)	90
GET Request (on Octet String of 15 chars)	89
GET Response (on Octet String of 15 chars):	105
SET Request (on integer)	91
SET Response (on integer)	91
TRAP (without any "varbindings") Note: with varbindings the size can vary depending on the number/size of each varbinding.	111

**Table 1: Typical 'byte-size' SNMP packets**

Regarding bandwidth, SNMP packets do not use up too much bandwidth as shown in the following simple calculation:

Equipment:	1	7700FR-C Frame/Chassis (Evertz)
	1	7700FC Frame Controller (Evertz)
	14	7761AVM2-DC Dual Channel Audio/Video Monitoring Modules (Evertz) <sup>2</sup>

Description: One frame (ref. "test frame") contains 14 dual channel modules, monitoring 28 video signals each with associated discrete balanced analog audio. Each channel is capable of monitoring 23 parameters (consequently, the dual channel provides up to 46 TRAPs).

Scenario: All TRAPs are enabled and the frame has encountered a complete signal failure, as all cards are now reporting all faults to the network monitoring system.

Results: Fourteen (14) modules provide 644 TRAPs (46 TRAPs x 14 modules), with each TRAP being 111 bytes for a total of 71,484 bytes

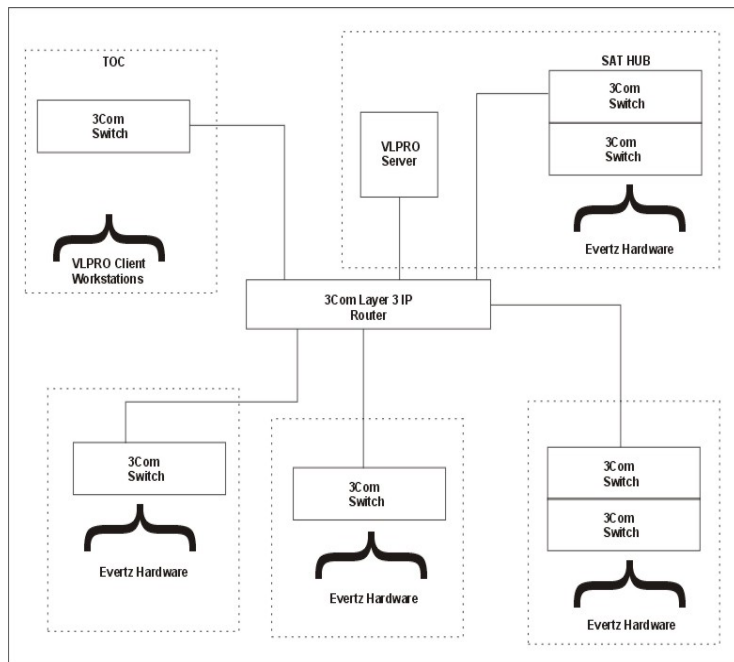
Subsequently, if the test frame issues every trap possible, approximately 69.8 kbytes are transferred. It would therefore require 183 fully loaded test frames, all simultaneously issuing every TRAP possible, to saturate a 100Mbit "NON-SWITCHED" network.

TRAP monitoring via SNMP introduces two philosophies that can affect network congestion: "push" vs. "pull". In the "push" case, a NMS listens for TRAPs from equipment or signal "sniffers" and upon receiving such a packet, extends a notification to the clients or drives a pre-configured event. The "push" philosophy is considered in the above TRAP bandwidth calculation example. Conversely, a "pull" approach based on polling configuration, proactively queries equipment at a given rate and expects information packets to be returned with appropriate health status updates. Depending on the polling rate, this approach can greatly increase network congestion and potentially bring the network to a standstill.

---

<sup>2</sup> For specific frame and module details, visit [www.evertz.com](http://www.evertz.com)

To expand on this push vs. pull concept with an example, consider again the previously introduced test frame, fully loaded with dual channel monitoring modules. The “pull” concept is employed by the NMS with a query rate set to poll each of the 23 fault parameters per channel. In addition, presuming that there are 500 channels (250 modules in 18 frames) being monitored, equating to 11500 polls and an equal 11500 responses. If the polling rate is set to a 1 second interval then, referring to Table 1 and presuming that there is an equal number of integer and text string type parameters (average packet size of 100 bytes) for each request and response (200 bytes total), a total of 2.3 Mbytes/sec (or 18.3 Mbits/sec) are transmitted and received. Of course, reducing the query rate in half, doubles the amount of traffic, so it is easy to see how network congestion can add up, without even considering any other network traffic.



**Figure 3: Simplified block diagram of complex network architecture used for facility wide equipment health and signal status monitoring.**

Alternatively, if bandwidth is still a concern, the network administrator can consider “manager management” techniques, namely reducing the polling rate or the number of parameters to query, or using a combination of push (signal monitoring and report only when something is at fault) and pull (equipment status) concepts. According to Jeffery Case, Founder and CTO of SNMP.com, a rule of thumb for equipment health monitoring should be somewhere

between 2 packets per second per node (a general health poll every second with 1 request and 1 response) and 2 packets per node every 10 - 30 minutes. How often this is maintained depends on a) how fast you want to detect that something is not healthy, b) how smooth you want your hourly, daily, and

weekly report graphs to look, and c) how many event notifications you want to tell you proactively when something is unhealthy or about to go bad. Then, if something interesting is found, then there might be a flurry of activity to detect what is going on and why ... perhaps in the 10s of packets per second...then back to the slow background polling.

### **SNMP Configuration Reliability**

In general, there are two “reliability” considerations: first, no network protocol is 100% reliable, and secondly, SNMP implementation reliability begins with properly designed network architecture. Reliability can further be divided into protocol (SNMP) reliability and system (redundancy) reliability. The latter is easily explained with either hardware that can be hot swappable and/or contains redundant power-supplies to prevent unscheduled power cycles on key equipment, or primary and back-up servers that collect data and distribute information and notifications to attached clients. In the server case, there is network-health monitoring with the capability of detecting that the primary server is down, forcing the back-up server to step up to the plate and re-routing the attached clients to communicate directly with the back-up location, thereby eliminating the single-point of failure. Alternatively, multiple servers can be used in an implementation to further improve on reliability.

The notion that “SNMP is unreliable” is by definition, true (SNMP, which is part of the UDP network protocol branch is a non-guaranteed protocol). However, protocols in general are all inherently unreliable to a varying degree and the difference among them is that they offer different (more or less) probabilities for losing a message along with tools to improve reliability. Within an SNMP framework there are a few tools such as trap counters, packet prioritization or SNMP v3 (with trap acknowledgement) that exist to further improve reliability in networks with traffic congestion issues. Similarly, SNMP configuration reliability via the “SET” command can also be easily implemented by doing a programmed, systematic, follow-up “GET” on that same parameter, again confirming complete message transfer. But clearly the most important factor, after a well-designed network, is to minimize network traffic or congestion, as this is

what potentially causes messages over UDP to “get lost” and never reach their destination. One method of not overloading the network is to ensure that the system is not flooded with useless TRAPs, which add to network congestion. Ultimately, reliability on the network still “depends” on a number of factors.

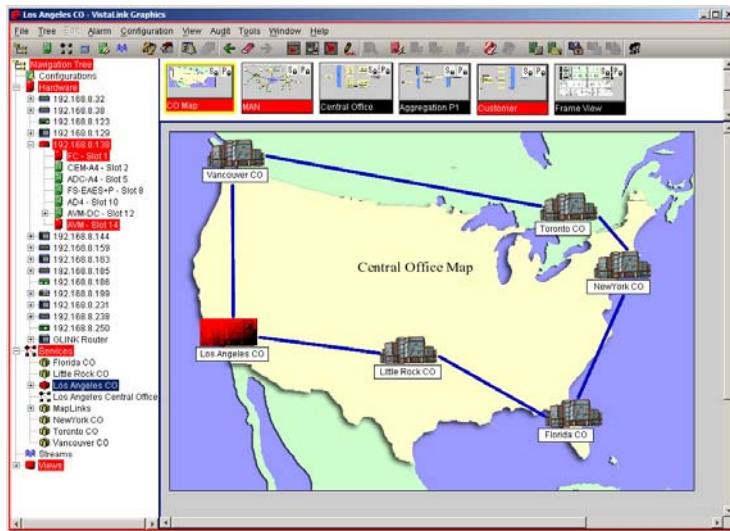
### **Yet Another Vendor-specific GUI**

A common statement when discussing network monitoring and configuration solutions is “not another software GUI”. At first thought, this seems odd as software unlocks the interface to the latest product-specific features, should be operating system agnostic (but most use Windows® anyway), is meant to be user-friendly, and usually can be the differentiating feature between rival solutions. However, additional software means additional equipment and training headaches not to mention interconnectivity and interoperability issues. As previously reported, competitors to gain the upper hand with end-users always use product differentiation, but when considering a facility monitoring system, ideally one manager should be interoperable with all equipment in one form or other. Although still quite a ways to go, there are certain developments to keep an eye on.

Most equipment manufacturers offer some form of installable configuration software or web interface. For advanced, large, complex installations, the usual route is to choose a larger NMS Software Developer, specializing in control system configuration and monitoring GUIs. However, end-users are still in a quandary for several understandable reasons. Namely, they may be unsure which NMS offers the most robust solution, which interfaces the best with all third-party equipment - either legacy and/or new, who is responsible for fixing interface and/or equipment problems when they arise, and which Manager developer will continue to support the product for the long-term.

However, a single-source NMS provides a seamless integration and a single direction to fix problems when faults occur – in effect a new meaning to “one-stop servicing”. Furthermore, upgrades and maintenance are easier to manage for the end customer. However, this service does not come free, and

additional costs may be incurred for customization work to communicate with and support other vendor's equipment. A few years ago, using HP OpenView seemed like an acceptable alternative NMS solution as it offered agnostic control of 3<sup>rd</sup> party equipment, built-in monitoring and enough horsepower for future expansion needs. This tool also allowed facilities to test potential user interfaces not only for problem-solving engineers and technicians, who required in-depth information and statistics about monitored signals, but also for front-line operators who need simple user-friendly GUIs for immediate problem identification and solution implementation. This software was implemented in a number of dedicated broadcast facilities, but for the most part did not fit the bill for others.



**Figure 4: Evertz VistaLINK PRO PLUS with graphics using SNMP for monitoring and configuration of Evertz equipment and 3<sup>rd</sup> party gear.**

Similarly, vendor-specific equipment interoperability efforts among third-party equipment manufacturers that were initiated previously did not develop as expected, and the only interoperability/networking effort really underway currently is through SMPTE at Sony's initiation of a enterprise-wide MIB interoperability effort. Sony is trying to add to the widely accepted SNMP MIB structure and standardize it within

SMPTE – this effort is meant to make the MIB uniform across vendors and does not change the true nature of a MIB. Instead, it actually adds a small branch to the existing MIB allowing vendors to enter product name, serial number, and firmware version etc. fields to further promote both equipment interoperability and vendor-specific GUIs.

A recent integration effort gaining popularity sees higher-end NMS Control Systems/Managers operating as a “collectors of everything”, meaning a repository for traps and a notification tool for operators,

monitoring both health status and signal status through a combination of push-pull philosophies using 3<sup>rd</sup>-party equipment. Upon receipt of a fault, the data is processed at the central point allowing the operator to access it through the NMS or get directed to the equipment vendor's specific GUI for additional configuration and troubleshooting tools through built-in 'launchable' applications, sparing the customer from needless software re-development costs. In the process, solutions are making use of the available tools at hand. In effect, this means that it is also up to the vendor, not the NMS to keep files up to date, and so the end customer is better off.

### **Return on Investment Considerations**

For service providers, the ROI from a signal monitoring implementation is increasingly important. Operators watching a signal leave the facility as well as customers watching the signal comfortably at home usually handle traditional signal monitoring. The viewer could detect problems at home and, before the operator is drawn to it, contacts the call center to complain. Each call on the toll-free line, along with the added costs of used space and personnel, equates to roughly U\$5.00. Clearly if the problem is far-reaching, and many subsequent calls are fielded by Customer Service, dollars start adding up. With the implementation of a signal monitoring system, it is anticipated that problems are detected before the viewers see them thereby reducing call center costs. Further, with the use of various trending tools, data collected within the monitoring system and over a specified period can be parsed further to provide significant details about certain outages, channels, services, providers, etc. in an attempt to provide cost credibility.

### **Summary**

With the future poised for expansions of channels, services, coverage area and equipment, investment growth in SNMP monitoring and configuration solutions is also expected. As needs and trust in network monitoring grow, expect to see more monitoring by exception (report of problems only when they occur) and automated reasoning (automated event control due to a detected fault/problem) solutions to surface, offering even more interoperability, convenience and ease of use.