



White Paper

Improving Workflow Efficiency - The Challenges of True End-to-End Control and Monitoring

Maurice Snell,
Snell Ltd

Written for presentation at the CABSAT 2011 exhibition.

Abstract

The traditional way for broadcasters to achieve an integrated system covering their entire workflow was to invest heavily in customized development. Only the wealthiest broadcasters could afford to do this, and therefore benefit from the workflow efficiencies that such an integrated system brings - giving operators a user-friendly view and control of all aspects of the operation. Today's financial environment drives all broadcasters, large and small, to maximize the efficient use of their assets, giving their staff the tools to do more with less. This increases the requirement for off-the-shelf packages that can approach the level of integration that previously necessitated custom development. This presentation will discuss the technical challenges involved in using a combination of generic protocols such as SNMP, and vendor-specific protocols, to provide cost-effective userconfigurable end-to-end control and monitoring systems.

Keywords

Control, Monitoring, SNMP, Infrastructure, Modular, Efficiency, Workflow, End-to-End, Systems, Integration, GUI.

Introduction

Throughout the history of television broadcast, it has been desirable to present operational staff with integrated user interfaces for remotely monitoring and controlling the equipment that delivers the services the operators are responsible for. The environment in which these operators work is inevitably a high-pressure one, with potentially millions of media consumers affected by any issues that arise and, in the case of premium commercial channels, millions of dollars of advertising revenue and other commercial contracts at risk, not to mention the individual and corporate reputation that could be publicly damaged if on-air services are affected. Clearly it would be beneficial to give these operators tools to help them rapidly identify and resolve issues that arise, whilst performing day-to-day tasks as efficiently as possible. The key requirements of these tools are to enable the remote control and monitoring of equipment, defined as follows:

Remote monitoring: provide rapid notification of any change of status, and allow access to the full details of the status of the equipment, without needing to walk through the various equipment racks and check device status locally, e.g. from individual device front panel or card edge indicators.

Remote control: enable the adjustment of configuration and operational parameters provided by the individual devices, without needing to access the local devices' front panels, card edge controls, etc.

With modern network-based technologies, both control and monitoring can apply to equipment relevant to the needs of the operator, wherever it is located - whether in the same room, same building, or even on a different continent to the operators themselves.

In the past, best practice for remote control and monitoring, (C&M), was typified by these examples: Figure 1 shows a typical custom-made hardware control panel, and Figure 2 shows a typical custom-made software GUI.




Figure 1: Typical custom-made hardware panel for remote control and monitoring of a range of equipment. This panel is still in use at the BBC, London.

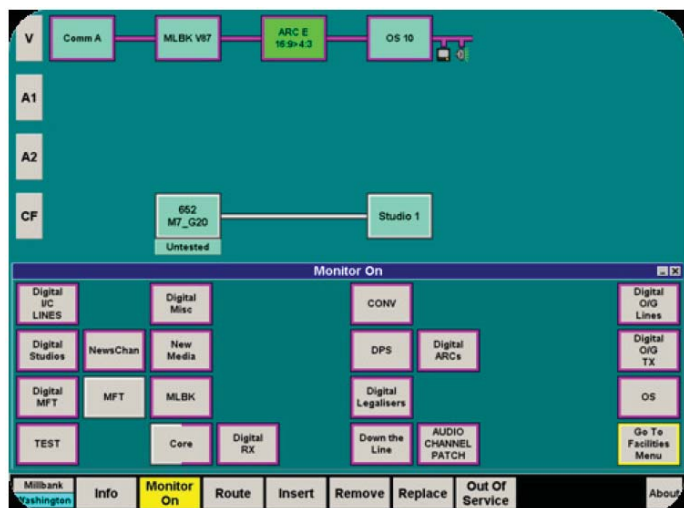


Figure 2: Typical late-20th century custom-made GUI for remote control and monitoring of a range of equipment. This is the Sprint system built for ITN, London, providing centralized intuitive control of a complex system of multi-format signal routers.

Both of these achieved some of the key requirements for such systems: giving operators simplified and consistent remote access to monitoring status and control of parameters from a variety of devices throughout the broadcast chain, removing the need for the operator to know the location of and visit the individual devices, and reducing human error by preventing access to other engineering or configuration controls from those units. However, systems like these were expensive to produce, because they required customized engineering development work to build each specific hardware or software panel. Where a broadcast chain remained unchanged for many years, it could be viable for such panels to be built, although only the most well funded broadcasters could afford the in-house or contract software effort required to produce a dedicated GUI for even a small part of their operation. Hardware panels were typically cheaper to build, but the resulting functionality was far more limited than a custom-built software system. Therefore most systems in the broadcast industry tended to make do with a sub-optimal combination of disparate sub-systems: often vendor-specific tools that provided a

good interface to that vendor's own products, without any substantive integration with the broader chain. This has led to one of the biggest practical issues facing system designers today: how to fit in enough PC screens to show all the different user interfaces required by a typical operator.

In recent times, many changes have affected the industry, such as the consolidation of broadcast operations through acquisitions and mergers, and a trend towards contracted service provision. Driven by a decrease in revenue per channel, almost all broadcasters now have to produce an increasing number of channels using a stable or shrinking engineering workforce. At the same time, modern digital and file-based workflows involve far more complex underlying technologies to deliver each channel. Therefore less budget is available to develop C&M solutions, but there is an increasing requirement to make efficient use of personnel by giving them highly customized tools that help perform their daily operations, alert them to problems that arise, and assist them in solving these problems. In addition, these systems need to be responsive to the changing requirements of the users, and allow rapid redeployment, expansion, and adjustment as the business responds to market transitions.

Some of the Challenges

Over the past 15 years, we have been working on C&M, initially for our own products and latterly on vendor-agnostic tools. In the process, we have confronted many issues affecting the delivery of real and usable C&M systems. The experience gained has been invaluable in understanding how to achieve tangible efficiency gains for our customers. Some of the commonly-occurring issues are described below.

Not all standards are created equal.

SNMP – Simple Network Management Protocol, standardized by the Internet Engineering Task Force (IETF) – has been widely used in the IT and telecommunications industries since the 1990s, and has become increasingly common in broadcast products in the last decade. However, as a standard, SNMP leaves a lot to be desired – ultimately because the standard is too loose, and leaves too much to the discretion of the individual implementer. This applies at different levels: at the low level, there are wide variations in protocol implementation between different vendors, for example in the use of array indices, in the mapping of low level values in to the different data types supported by SNMP, and in the relationship between the data objects used for the different classes of SNMP communication. At the higher level, there is no standardization of parameters within broadcast products so, for example, to adjust the video gain of one device from one vendor requires different commands than adjusting the same gain parameter from another device from the same vendor, let alone a device from a different vendor. Within the IT industry, there is a base line of functionality that has been standardized: for example, it is possible to extract the status of an ethernet port from any managed switch using common commands, however accessing hardware status such as PSU, fans, temperatures etc. is, as with broadcast, unique to each switch vendor and even product line from the same vendor.



There have been attempts to standardize some subsets of broadcast functionality in SNMP, for example ETSI Technical Specification 102 032, for DVB test and measurement devices, published in 2002.

Unfortunately, in our experience, few if any devices actually meet this standard – all devices implement vendor-specific MIBs, (Management Information Base – the SNMP documentation format for describing the parameters available from a unit).

Whilst SNMP is indeed an open standard, not all broadcast vendors have an open process for distributing their MIBs. Universally in the IT world, and for some broadcast vendors, MIBs are published freely on web and/or ftp sites. Other broadcast vendors consider their MIBs private and confidential, restricting access, and thereby complicating the use of this supposedly-open standard by their customers and partners.

In addition to these variations between implementations, there are a number of technical limitations of the SNMP protocol, for example:

- There is no standardized way to request resending any previously-sent traps, (the unsolicited messages sent by the device to the remote manager software to indicate a change of status such as an alarm condition occurring), and therefore no consistent way to initialize a monitoring system and determine the current status of devices that send traps.
- SNMP is normally implemented over UDP/IP, which is a best-effort non-guaranteed delivery mechanism. There is no standardized way to determine if a trap was or was not received by the manager software, and determine appropriate retries if required, for example to recover from a temporary network interruption.
- SNMP does not support registering for notification of changes to controllable values. A typical device is able to send trap notifications for changes to alarm conditions, (e.g. PSU has just failed), but is not able to send notifications on changes of controllable values such as video gain. This makes it unsuitable for applications where multiple users may need to adjust parameters interactively and see the changes made by other users in real-time. Where SNMP must be used in such situations, a high rate of polling can give tolerable performance at the cost of high network and CPU utilization.

Proprietary workarounds for each of these limitations are provided by some vendors, exacerbating the variation seen in the field from products that all meet the same “standard”.

Over the years, we have had to integrate with many “challenging” implementations of SNMP, for example:

- A video playout server that crashed, rebooted, and stopped playing video when it was polled for SNMP data.
- A multiviewer that sent thousands of non-informative traps per hour although no status was changing.
- A modular frame that claimed to monitor air temperature, but always reported it as -128 Celsius.
- Various devices sold as SNMP-compliant, which actually did not support SNMP at all.
- Systems whose SNMP parameters depended on the serial number of the hardware unit, and thus changed unexpectedly when a failed unit was replaced with an apparently-identical one.

- Devices that sent SNMP traps but did not allow polling of values, thus providing no way to find out the current status and initialize the monitoring system.
- Devices that used proprietary data packing algorithms to put many pieces of data within a single SNMP data type, thus requiring non-standard down-stream handling of the data.
- Devices that require the C&M system to perform logical calculations across multiple data values to extract simple status such as whether a particular PSU has failed.

Due to these limitations and inconsistencies, implementing real-world control and/or monitoring systems using SNMP is usually much harder, (i.e. more expensive in man-days for software development and commissioning), than might be expected by a new-comer to the field. However, the widespread adoption of SNMP, at least for monitoring purposes, has made it now a de-facto standard in the broadcast industry. For many purchasers, the presence of SNMP is considered a default requirement for any device, whether or not it will be immediately used.

We have had to invest many man-years of development to get an SNMP manager solution capable of working with the whole spectrum of devices used by broadcasters in a sufficiently flexible way to make a robust and usable system.

Control is different to monitoring

The technical requirements are quite different to achieve generic configurable alarm and status monitoring, versus generic configurable device control. If you take a C&M product optimized for SNMP monitoring, and another one optimized for generic device control, the underlying architectures are likely to be quite different. For example, the SNMP protocol is widely implemented in broadcast and IT equipment, usually for monitoring, and sometimes also for control. However, the limitations of the protocol make it unsuitable for some kinds of control, for example live interactive multi-user control. Therefore it may be necessary to use different connection methods to the same device simultaneously for different C&M purposes.

There is also a major potential security benefit in allowing the control mechanism to be different from the monitoring mechanism. Unfortunately, this is not generally possible with SNMP: once you have opened up access to the UDP/IP port that permits retrieving, for example, the status of a PSU for monitoring purposes, you have the potential for device control to be sent through the same port, and this may be seen as an unacceptable risk in some environments. With some other vendor-specific protocols, such as Snell’s RollCall, it is possible to strictly separate control from monitoring. This allows the user to choose to distribute fully-functioning monitoring workstations more widely within their organization, including low-security areas such as offices and remote access laptops, where it might not be safe or desirable to permit use of software and IP ports that had the potential to control live equipment.



Implementing vendor-specific protocols can be time-consuming

For the reasons described above, i.e. the limitations of SNMP, and the differing requirements of control as opposed to monitoring, in order to deliver a fully-featured system it is often necessary to interface directly with a third-party vendor-specific protocol. However, this can be very time-consuming, and thus expensive to the end user, because the software engineer typically needs to start from the ground up with each new protocol. One technical solution that has served us well over the years is to develop generic configurable translators, which can be scripted to work with a wide range of different protocols. There is a trade-off: a scripted solution using a generic translator will typically require an order of magnitude less development time than a fully-custom coded solution, but the scripting system may impose some limitations, preventing access to particular vendor-specific functionality from some protocols. Therefore the ideal C&M toolkit needs to allow for both methods: scriptable translators to be used where possible and appropriate, with the ability to fall back to full custom development if the system requires it.

Not all vendor-specific protocols are created equal

One modular supplier uses a proprietary protocol where network broadcast packets are sent by each device to communicate with the C&M system. As the number of devices grows, so does the network bandwidth consumed by these broadcasts which, unfortunately, the ethernet switches are obliged to duplicate to every connected device. Some other devices can not cope with this flood of unhelpful broadcast packets and, in a large system, the devices transmitting the broadcasts can not reliably function either. Therefore additional complexity has to be added to the network infrastructure, to segregate small groups of the offending units in to isolated virtual networks.

The impracticality of human monitoring of multi-channel audio

In the early days of television broadcasting, there would have been at least one engineer able to monitor each and every television channel, watching a screen and listening to the mono or later stereo audio via headphones or loud speakers. Today however, many broadcast organizations now manage 10s or 100s of television channels, with typical ratios of between 5 to 50 channels being monitored by a single person. It is feasible for a person to "watch" many simultaneous video channels, through video walls with many discrete screens or virtual screens presented through multiviewer technology on to large display screens. However, none of these solutions allow the human operator to listen to more than one audio stream, (be that mono, stereo, or surround sound), and each video stream shown on the video wall may have multiple audio streams, typically for different languages.

Therefore a modern monitoring solution must offer a variety of tools for automating the monitoring of audio, and indeed video, to allow a high quality of service to be delivered across increasing numbers of channels per operator.

False alarms – the nail in the coffin of many attempted monitoring systems

In our experience of visiting many customers attempting to use a variety of software packages for monitoring alarms from broadcast equipment, one of the most common problems is the occurrence of

false alarms – i.e. SNMP traps or other messages being received by the system and presented to the operator, despite there being no actual fault in the broadcast chain. Whilst these are an annoyance and a waste of time to the operator, the more serious problem is that operators quickly learn to ignore what the monitoring system is telling them. Inevitably this will lead to a genuine alarm being lost amongst the flood of false alarms, followed by a backlash against the monitoring system which, in practice, failed to help the operators deal with the genuine problem. There are sadly many broadcast facilities where audible alarms are being fed to loud speakers with the volume set to zero, or thousands of emails are collecting in an inbox that nobody bothers to look at, or a screen in the corner is flashing red with numerous error reports that have the occasional genuine fault obscured by thousands of unhelpful false alarms.

To make a truly usable monitoring system, a variety of tools must be provided to allow alarms to be masked, filtered, and sorted. These have to be intuitive and easy enough to use so that the daily changes in the use of the broadcast system can be mirrored in the filtering of alarms presented to the users.

Advanced C&M solutions can go beyond notifying the user, and perform automated self-healing actions: for example, having detected frozen video on the primary chain but valid moving video on the backup chain, the system could automatically drive a router crosspoint or bypass switch to put the backup chain to air. This can avoid the delay and potential human error that could occur if the system only warns the operator and requires them to take the corrective action. For this kind of control to be trusted to work automatically without user intervention, it is even more important that the monitoring can be trusted, false alarms suppressed, etc.

One-stop-shop solution?

Since most broadcast hardware suppliers provide some level of C&M software to work with their own products, perhaps the simplest way to achieve a complete end-to-end software system would be by choosing a single supplier for the entire hardware chain. Unfortunately this is rarely possible, with few if any suppliers truly covering the entire value chain, e.g. from camera to transmitter for a typical broadcaster. Most systems contain equipment from a wide variety of manufacturers, either as a result of a best-of-breed purchasing policy, or an attempt to get the cheapest possible price for each component, without necessarily considering the higher cost of integrating these diverse devices.

There is a variety of C&M tools used in the broadcast industry. Most of these are specific to one vendor's products, sometimes with extensions to limited third-party control, or limited third-party monitoring. There are also more generic tools optimized for generic multi-vendor control, and other tools optimized for generic SNMP-based monitoring. However, there have been few tools able to competently cover both the monitoring and control of a wide range of equipment.

What defines the ideal C&M solution?

To meet the needs of today's broadcast and related industries, a C&M product should aspire to be able to provide a single system that delivers everything the operator needs to perform their tasks – both the scheduled and planned tasks, and unplanned or emergency tasks that arise.

The following essential qualities need to be designed in throughout the solution:

- A very high level of reliability, allowing trust to be placed in the tool as the primary method for interacting with the equipment, being alerted to problems, and being able to resolve them. Ten to fifteen years ago, most customers were reluctant to rely on Microsoft Windows platforms for such essential systems, however there has been a gradual acceptance and widespread adoption of Windows across all aspects of the industry, such that Windows-based clients are now the de-facto standard.
- Scalability, in pricing and technology, to work cost-effectively from the smallest simple system up to a large multinational enterprise-wide solution.
- Flexibility to integrate with any broadcast, IT, or other device critical to the customer's workflow. This may require any combination of old-fashioned closing-contact GPI I/Os, RS422/232 serial ports and ethernet interfaces with a combination of generic and proprietary protocols.

A services and support organization is required, able to assist with any or all of the following: capturing the customer requirements, designing a system to meet these requirements, making commercial proposals for different options such as

- Only supply of configurable products, with the end-user to perform all configuration and installation. Subject to availability of local staff, this option can give the most economical final solution, and enables rapid response to required changes.
- A hybrid solution where some partial configuration and installation is performed, perhaps including formal training, with the end user continuing the configuration and on-going maintenance of the system.
- A complete turn-key system where the supplier does all installation and configuration to meet the customer's requirements, providing a ready-to-use system.

Whilst every system is different, the following features are commonly required:

- Advanced tools to automate the process of monitoring and aggregating alarm conditions from large numbers of devices.
- A range of methods allowing the user to manage the suppression of temporary false alarms, (e.g. a device removed for maintenance), and permanent false alarms, (e.g. a device reporting no signal present on input 2, when this input is not wired and never used).
- Collection and logging of information to a central database, from where it can be queried to analyze the current and historic state of the system, ascertain trends, plot graphs, and derive a variety of business intelligence from the underlying data.
- User-configurable screens, able to display alarm conditions, and offer intuitive controls that map to any device in the system. Figure 3 shows an example configured screen.

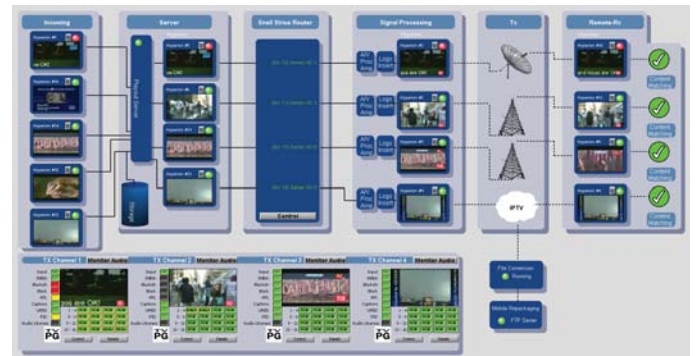


Figure 3 : Example of a user-configurable screen presenting a graphically-intuitive high-level overview display summarizing alarm conditions from a typical multi-channel broadcast system.

- Optional hardware panels, to allow hands-on control that can be operated blind, e.g. allowing the user to adjust video gain while looking at a video analyzer. This kind of operation is difficult to perform with a touch-screen or mouse-operated GUI. Hardware panels can also be operated in a much smaller rack space than a PC-based solution. As with the user-configurable GUIs, the end user must be able to easily reconfigure the panels, to give the operator the required controls from products such as routers, modules, and multiviewers, presented as a combined workflow. Figure 4 shows a range of configurable hardware panels.




Figure 4 : Examples of user-configurable hardware panels, able to offer integrated access to multiple devices, simplifying the user's workflow.

- A range of tools to automatically check the quality of the video and audio content and metadata at specific points in the broadcast chain, which can include local or remote signals, return and off-air feeds, incoming and out-going lines, and content being ingested. To aid remote analysis by an operator who does not have access to the original media, (e.g. a remote SDI feed), it should be possible to stream video and audio proxies to the operator, in addition to the metadata measured from the remote signals. To enable remote access through commodity communication networks, low-bandwidth streaming modes should be supported. Some alarm conditions are straightforward technical errors, e.g. SDI signal not present, while others are subjective measurements to help guide the operator to potential as well as actual problems, for example if the level of motion in the video is suspiciously low for a particular genre of program content expected for this channel or program. Examples of the parameters that should be monitored include:

Video alarms:

- Video carrier missing
- Video wrong format, (e.g. SD when HD required)
- Black video
- Frozen video
- Video too dark
- Level of motion too low
- CRC / EDH errors
- Average picture level too high or too low
- Luma or chroma values too high or too low
- Video bit depth too low, (e.g. 10-bit data cropped to 8 bits)

Audio alarms (per channel):

- Audio presence and type
- Audio bit depth too low, (cropped)
Audio silent, or quiet, or too loud, or overflowing the transmission format/clipping
- Audio previously clipped and attenuated
- Phase errors
- Unexpected correlation: mono, stereo, dual-mono
- Dolby guard-band and alignment errors

Metadata alarms:

- Presence and contents of closed captions or subtitles, in various formats including CEA608, CEA708, etc.
- Presence and value of AFD information, (aspect ratios etc.), in various formats such as SMPTE 2016, WSS, and VI
- Presence and value of content advisory, e.g. XDS, V-chip
- Presence, value, and continuity of timecode in ANC or VITC
- Presence and value of UMIDs as per SMPTE 330M, embedded in SDI as per SMPTE RP223
- The ability to interface to external systems that can provide additional information about what control and monitoring should be applied to different channels at different times. For example, a playout automation schedule may have information about the expected genre of program content, (number of type of audio channels, aspect ratio, presence and location of logo, etc.), and ideally this should be fed in to the C&M system, so that the correct settings for processing devices can be applied and the optimal set of monitoring can be enabled at each point in the schedule. This can significantly reduce the presence of false alarms, whilst giving a more thorough and faster notification of genuine problems. A famously difficult example is the classic motion picture 2001, which contains a sequence where the audio is entirely silent for more than 2 minutes. Almost any audio silence detector, configured to suppress alarms on “typical” material but raise alarms to indicate a technical fault, will raise a false alarm from this content. Alternatively, if the detector is configured with a time threshold such that it ignores silent audio for at least 2 minutes before raising an alarm, then many typical failures will go undetected for much longer than is desirable. If the automation schedule can contain sufficient metadata about each item in the schedule, and this is communicated to the monitoring system, then the audio silence detector could be disabled or configured with longer time thresholds when playing this particular item, or even just for the necessary segment within the motion picture.
- User-configurable rules engine, able to perform automated actions in response to simple or complex combinations of status conditions. For example, self-healing systems could be constructed

where backup chains are automatically routed in and reconfigured to replace a main channel where failure has been detected. Figure 5 shows an example of a rule being graphically edited.




Figure 5 : Example of a rule in a graphical editor, allowing the user to define logical combinations of status conditions that trigger automated corrective actions.

- User configurable alarm notification systems, e.g. sending emails and SMS text messages containing information about alarms that have occurred – with sophisticated handling of multiple or repeated alarms without flooding the users with too many messages.
- User-configurable scriptable interfacing to third-party devices through GPI, SNMP, web services, and generic serial or ethernet protocols.

C&M tools are starting to appear that have this broad range of functionality, such as Centra from Snell, allowing systems to be user-configured for any scope, up to a fully-featured end-to-end enterprise-wide control and monitoring system.

Conclusion

This paper has described the challenges of building real-world integrated systems for the monitoring and control of broadcast infrastructure. It has also detailed the requirements for the cost-effective creation of integrated GUIs and hardware panels for broadcast operators, giving them intuitive access to monitor and control all underlying equipment contributing to the broadcast services they are responsible for. Technological improvements allow this to be economically viable for any scale of operation, and therefore contribute to the ongoing efficiency improvements demanded of all organizations.

Acknowledgements

I would like to thank my colleagues at Snell for their comments and contributions, and the directors of Snell Ltd for permission to publish this paper.