

AN EMBEDDED SYSTEMS APPROACH TO IP-BASED BROADCAST FACILITY REMOTE CONTROL

Stephen Dinkel
Burk Technology
Lee's Summit, MO

Nathan Burk
Burk Technology
Littleton, MA

Jonathan Burk
Burk Technology
Littleton, MA

ABSTRACT

HD Radio, DTV and multicasting have combined with the realities of downsized engineering departments to raise the expectations for broadcast facility remote control. At the same time, LAN/WAN connectivity has invited vast improvements in the way broadcasters collect, distribute and manage data. An embedded systems approach to IP-based broadcast facility remote control is the outcome of these dynamics. Combining IP architecture with a self-contained, embedded design offers broadcasters major advancements in alarm propagation, alarm aggregation and multi-site operability. Exploring this approach explains how to adapt broadcast facility management to today's technologies while satisfying reliability requirements in ways where PC and server-based solutions have fallen short.

DESIGNING FOR REMOTE DATA NETWORKS

Consolidation presents a major challenge to traditional remote control. Bringing digital buildouts into the mix further underscores the need for scalability and a broader approach to managing remote sites. With engineering attention divided among more stations and increased need for high level corporate overview, the point-to-point approach to site management had to be expanded while still maintaining the reliability of a closed system.

The ability of new solutions to meet the need for scalability is directly linked to the availability of faster, more efficient communication paths to remote sites. Solutions are available that leverage existing broadcast infrastructure to lower upfront costs. Satellite IP and cellular GSM/GPRS links have also gained acceptance in the broadcast environment, and new technologies such as WiMAX are gaining ground. Moreover, many groups have already invested in infrastructure for email, mobile messaging, corporate LANs, VPNs, Blackberries, etc. Bringing remote control onto the same canvas adds to the return on such an investment. The challenge lies in delivering these solutions to the realm of RF engineering, demonstrating broadcast-grade reliability and proving future readiness.

Navigating the IP Landscape

For all the advancements in IP deployment, migrating to an IP-based system for remote broadcast facility management can be more complicated than the service providers, equipment manufacturers and end users would like it to be. Even if a high speed link exists between studio and transmitter, program needs must be met first, often leaving little bandwidth for other applications. Dropped packets, network downtime and latency all challenge a broadcaster's confidence in using IP for broadcast remote control.

However, these obstacles to using IP for remote facility management can be overcome. Broadcasters understand this as one more extension of the marriage of RF engineering and IT engineering. For manufacturers, the challenge is to understand that IP-based equipment needs to live in the reality of broadcast networking, not the "universal broadband" model of the software industry.

MAKING LOW BANDWIDTH WORK

This distinction between developing IP-based products for the broadcast industry and designing for the corporate/IT environment is critical, and is what led us to the embedded system model for IP-based broadcast facility remote control.

In a networked environment, it is easy to think of bandwidth as unlimited. However, there are two problems with this assumption. First, networks are shared by many devices that compete for bandwidth. Second, bandwidth at remote sites is far more limited than it is in the office. Any real-time control device running over an IP network cannot assume that 100Mbps are available for its use.

In designing an IP-based remote facility management system, the realities of bandwidth limitations required us to design a protocol that allowed maximum data content in each packet. Comparing this approach to SNMP showed that OID strings generated larger packet sizes, consumed more processing power, and required substantially more bandwidth to achieve scalability. Fig. 1 compares the embedded system to an SNMP-based system to show that even if the pipe is narrowed to modem or ISDN speeds, the embedded system

remains exceptionally scalable. Calculations at DSL speeds and faster revealed similar results: over a T1, the embedded system showed more than six times the site capacity of an SNMP system.

TCP/IP, although reliable, requires too much processing overhead to allow for ample scalability. We opted to build a design that utilizes UDP/IP, taking advantage of its leaner connectionless model. The downside is that UDP packets are not acknowledged by the receiving network node. To overcome this obstacle, we added application layer acknowledgements to the protocol to ensure commands and data updates are received by their intended recipients. This results in lower overhead, while still safeguarding against dropped packets.

Using a peer-to-peer architecture further economizes communications. Unlike an approach where passive interfaces report telemetry back to a central server, a peer-to-peer network takes the remote server out of the equation, removes a point of failure and eliminates the extra “hop” in the data path that a server creates. This model places additional burden on protocol efficiency because without a central server, there needs to be a way for multiple clients to access the same data simultaneously. This includes software applications,

web browsers, or other units in the system. To address this requirement, a subscription model is used, where the embedded system adds to the data stream only those parameters requested by the clients. For example, if a user at the front panel selects a new site for display, the system terminates the subscription to the previous site’s data and provides the data for the newly selected site instead. Data is not added to the stream if it is not needed for display at that moment. This is how an embedded system, without the aid of PCs or servers, can operate at modem speeds and still encompass more than 3,000 discrete monitoring and control parameters.

The efficiency of such a design supports site-to-site control over low speeds, which is important when considering system operability over existing data links.

IP Over Serial Links

Where IP connectivity is not available, a serial link may be used with the addition of a learning serial-to-Ethernet bridge. The serial-to-Ethernet bridge extends a LAN to a remote location over a serial link, at speeds as low as 4800 baud. Unlike a serial-to-Ethernet converter, which locally transforms serial data into TCP/IP and requires an existing IP path, a serial-to-Ethernet bridge connects two LANs and operates on a serial link such as one provided by a digital STL or dedicated telco

Scalability at Low Bandwidth: Proprietary UDP/IP Protocol vs. SNMP




Fig. 1: Embedded System Scalability. Scalability at various modem/ISDN speeds of an embedded system with proprietary UDP/IP protocol compared to SNMP monitoring. Calculations assume 32 channels of analog metering updating every second.

link. To optimize throughput, the bridge “learns” the location of each MAC address, and transports only that network traffic which is destined to the other side of the bridge. With such a solution in place, the convenience of IP addressability at the remote site is achieved without the expense of broadband connectivity.

Integrating SNMP

With so much emphasis on using custom protocol for bandwidth economy, it was important for us to keep the door open for integration with standard protocols like SNMP. Even though the larger packet size and slower data polling make SNMP less suited for real-time monitoring and control of the broadcast network, SNMP bridges the IT and broadcast engineering worlds and provides access to a broader range of equipment than before. The main difference is that a real-time embedded system can sample data very quickly, which allows it to handle data that naturally fluctuates, such as power output, line voltage, temperature, etc. With typical polling rates in minutes instead of milliseconds, SNMP is often used for monitoring network infrastructure and reporting performance metrics where data granularity is less critical.

Because of the differences in SNMP monitoring and real-time reporting, combining the two models onboard a single platform was a unique challenge. It ultimately led to the development of an SNMP manager embedded in the client software, allowing SNMP monitoring alongside real-time broadcast data without inflating the native protocol or requiring greater bandwidth.

EMBEDDED SYSTEM FOR STABILITY

An embedded system offers advantages beyond protocol and bandwidth efficiency. Because the system is self-contained and independent from a PC, disruptions induced by changes in operating system

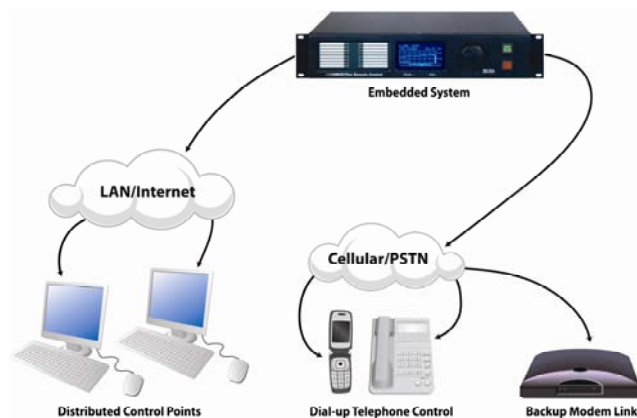


Fig. 2: Server-Client Topology. An embedded system remains connected to the outside world without requiring a central server or computer for data management. Client-side access to the system is flexible.

versions, service packs, security updates, etc. are not an issue. Because a PC is used for multiple purposes, there are often conflicts between applications and competition for CPU time. Desktop operating systems like Windows and Linux are non-deterministic, which further challenges implementation of real-time, mission critical monitoring and control functions. PC solutions also present the problem of inconsistent performance from one installation to the next because each PC is configured with different computing resources and a different complement of software. Finally, because the PC relies on moving parts, its MTBF is lower than an embedded device.

To avoid these pitfalls, one of our primary design requirements was to manage alarms, macros and system configuration onboard the system hardware instead of depending on PCs. We looked to an industrial real-time operating system that offers several advantages over an embedded version of a consumer OS. The result is that the system only requires sufficient memory and processing power to handle its specific tasks. There is also very little OS overhead, which is not the case with an embedded Windows OS. It also means that it was possible to precisely control process priorities to prevent, for example, an urgent alarm notification from being queued behind a configuration update.

Running a remote control system in an embedded device does not, however, leave it disconnected from the PC-driven world, as shown in Fig. 2. In the IP environment, software and web browsers can connect to the system using a client-server model. The hardware at the remote site acts as a server, leaving the desktop PCs to the purpose for which they are best suited: handling the user interface features as a remote client. The upshot is if a service pack is released for the PCs operating system, if a hard drive fails, or if the computer is reallocated for a new use, the integrity of the broadcast facility remote control system remains uninterrupted. This is in contrast to a system that depends on remote commands from a PC or server for routine events like tower light monitoring, AC power management, backup transmitter switching, etc.

PUTTING SCALABILITY TO WORK

A system built on IP architecture breaks the barrier of point-to-point communications, enabling site-to-site connectivity on a far greater scale than before. Being able to monitor and control one remote site from any other fits perfectly into the model of consolidation, where there are more sites to manage and fewer staff to share the responsibility. The ability to automatically (or manually) issue commands from one site directly to another brings the entire broadcast network within easy reach of any location and resolves problems of “orphaned” sites due to traditional communication limitations.

A logical follow-on to the site-to-site capability is being able to integrate unattended studios in the broadcast facility management system. A trouble condition at either location can result in an off-the-air event. Especially with trends toward centralcasting and using the Network Operations Center model for site management, the ability of IP to handle point-to-multipoint communication provides an effective means to expand the scope of broadcast facility management to include studios.

Simplifying Data Presentation

With a system designed to transport larger volumes of data, the need to present that data to the broadcaster in a manageable format becomes more important. All of the advantages of scalability only compound the problem of workload management unless there is a system in place to reduce the volume of data to a reasonable size, deliver it using an effective medium, and apply user-assigned rules to ensure information is directed to the appropriate personnel.

Alarm notification is the area where this concept is most critical. A real-time system can report events quickly, but the alarm report is of little value if it does

not prompt an effective resolution. Being able to direct alarm notifications to only the appropriate personnel is a major improvement toward facilitating an effective response to the alarm. By routing transmission system alarms to the RF engineer, audio alarms to the PD, and so on, each system user has a more clearly defined role and there is less potential for alarm conditions to remain unresolved due to communication errors. All operators can still receive notifications via email while, for example, the personnel with primary responsibility receive a pager alert or telephone call from the remote management system. This approach, which maximizes data sharing within the workgroup while directing data distribution based on area of responsibility, is shown in Fig. 3. Client software, as always, can log and track all data for future review.

An extension of selective alarm notification is the capability to analyze multiple out-of-tolerance events so that system reports a small number of primary alarm reports instead of every alarm that results as a consequence of the primary problem. As shown in Fig. 4, a traditional reporting scheme results in numerous alarm calls to the engineer, even though the alarms are the result of a single event. If the system instead collects the out-of-tolerance data over the span of a few




Fig. 3: Alarm Propagation. Selective alarm notification, combined with flexible vehicles for incident reporting, results in an effective means of sharing data among stakeholders and resolving alarm conditions efficiently.

seconds and applies aggregation logic to the data, a more informed alarm notification can be made.

An embedded system is well positioned to handle alarm aggregation because the extra processes involved can occur in virtually the same moment as the raw data enters the system. A remote server is never tasked with receiving the data, processing rules, and determining how to report the alarm.

FROM CONCEPT TO REALITY

In bringing all of these concepts together, a guiding requirement was to achieve relevance to real-world broadcast applications. A solution that leverages a station or group's investment in networking communications is of little use if it simply leaves the broadcaster with a new set of problems to solve.

In an IP-based system for broadcast facility management, the major concern is over redundancy in the communication path. The embedded approach is designed to keep processing power at the remote site to protect the system's integrity if there is LAN/WAN failure. However, a means to access the site during that failure remains a critical requirement. Maintaining a dial-up modem link and dial-up telephone control to the site is good practice even with the most robust IP link. A system with a protocol that operates solely in an IP environment precludes this type of redundancy. This was a consideration that led us to support dial-up in addition to IP.

Outlining a Reasonable Migration Path

The transition to IP has tempted an "out-with-the-old" attitude on the design side, but on the implementation side, some level of backwards compatibility with older equipment and older link types is needed to ensure an easy transition. This consideration played a major role in development and resulted in a smooth integration with existing remote control equipment, as shown in Fig. 5.



Fig. 5: Backwards Compatibility. Two generations of broadcast facility remote control equipment at a central studio facility in Hooksett, NH. Photo courtesy of Dirk Nadon, Director of Engineering NH, Nassau Broadcasting.

The advancements in network technologies have moved broadcast forward in impressive ways. When we began work on a broadcast facility remote control system to take advantage of these advancements, it was evident that the technology would only provide value if the system would work in real-world applications using tools that broadcasters can access now. Wrapping the functionality into an embedded system that thrives on the LAN/WAN removed the burden of OS maintenance and software conflict resolution from the broadcaster and provided a system that would deliver stability in an unattended environment over a lifespan that meets the expectations of broadcast.

Alarm Aggregation: Main AC Failure	
Out of tolerance conditions	Alarms reported
MAIN AC OFF	MAIN AC OFF
XMTR OFF AIR	XMTR OFF AIR
SUPPLY VOLTAGE LOW	
LINE VOLTAGE LOW	
LINE VOLTAGE OFF	
XMTR POWER LOW	
HVAC OFF	
STL SIGNAL LOW	
SILENCE SENSOR ALARM	
EXCITER POWER LOW	
EXCITER OFF	

Fig. 4: Alarm aggregation. As the scalability of a broadcast facility management system increases, so does the need for more intelligent alarm reporting. Alarm aggregation generates alarms on actionable conditions, not the conditions that occur as the result of a primary failure.