

SECURE-IP SITE MONITORING & CONTROL OVER THE INTERNET

Alan Brown¹, Rob Spiers¹, John Ahern² and Andrew T. Mulrooney²

¹Radica Broadcast Systems Ltd., UK and ²Comlab Inc., Canada.

ABSTRACT

Monitoring of transmission sites and remote locations has come a long way since the advent of broadcasting last century, but how can the move from point-to-point communications to more readily accessible, and therefore more vulnerable, communications methods such as the Internet be made more secure through the use of modern technology?

This paper examines the growing requirement for site monitoring and control at broadcast transmission sites, and how the Internet can be used without compromising the security of communications and risking unauthorised users gaining access to equipment.

The significant cost-benefit advantages of using telemetry systems at transmission sites are described in terms of reduced transmitter downtime and fewer emergency site visits by an engineer. This has led to the gradual proliferation of Remote Terminal Units (RTUs) at these sites and the increasing need to communicate securely with them. Historically, communications security has not been a problem as most early communications methods have been via RS232 over a radio link or by dial-up telephone. Some of the more-recent RTUs have more than one means of communications available to them, and IP connectivity is becoming more popular. Using IP overcomes the inherent delays found in conventional telephone technology, making the polling of multiple sites much easier and quicker. Once all of the RTUs and associated management software are located on an IP network the ability to create a 'virtual site' becomes a possibility. The use of IP does, however, introduce potentially serious security issues. It is proposed that these can be overcome by using the widely available public domain information on the Advanced Encryption Standard (AES) to implement added data security using encrypted communications. In using these methods, similar to those used by high-security operations such as banks and government departments, it is possible to exploit the benefits of IP data communications systems without exposing the system to potential abuse.

THE REQUIREMENT FOR SITE MONITORING & CONTROL.

When radio and television broadcasting started in the early part of the last century, transmitters were huge, unreliable pieces of equipment which required 24 hour supervision and continual adjustment to keep them working properly. Over the years, transmitters have become smaller and much more reliable, so the need for a fully staffed broadcast site has diminished, although given the importance of the transmitter for a radio or TV station it is still very essential to build in redundancy for the inevitable "breakdown". Indeed, in some countries; most notably the United States of America, legislation still requires the transmitter to be monitored and for the broadcaster to have the ability to shut down the system in the event of a fault occurring.

Even in countries where there is no legislative need for monitoring, some organisations have discovered the benefits of knowing what is happening at their transmitter sites without having to go there, and without having to staff the site 24 hours a day. When faults do occur, being able to pinpoint the problem can save multiple trips to site, substantially decrease transmitter down-time and reduce engineering costs by allowing the engineer to take the necessary replacement parts to site on the first visit.

Remote telemetry i.e. the ability to meter equipment at a distance, allows you to know what is happening, but remote control can help you to solve the problem without even having to go to the site. Transmitters nowadays often have built-in redundancy and the remote control system can automatically or manually switch to the reserve chain whilst an engineer is dispatched. One obvious advantage of this lies in the significant reduction of transmitter downtime, but other advantages might include allowing the engineer to attend during normal working hours; an option which is generally less expensive, and for the correct spare parts to be taken first time.

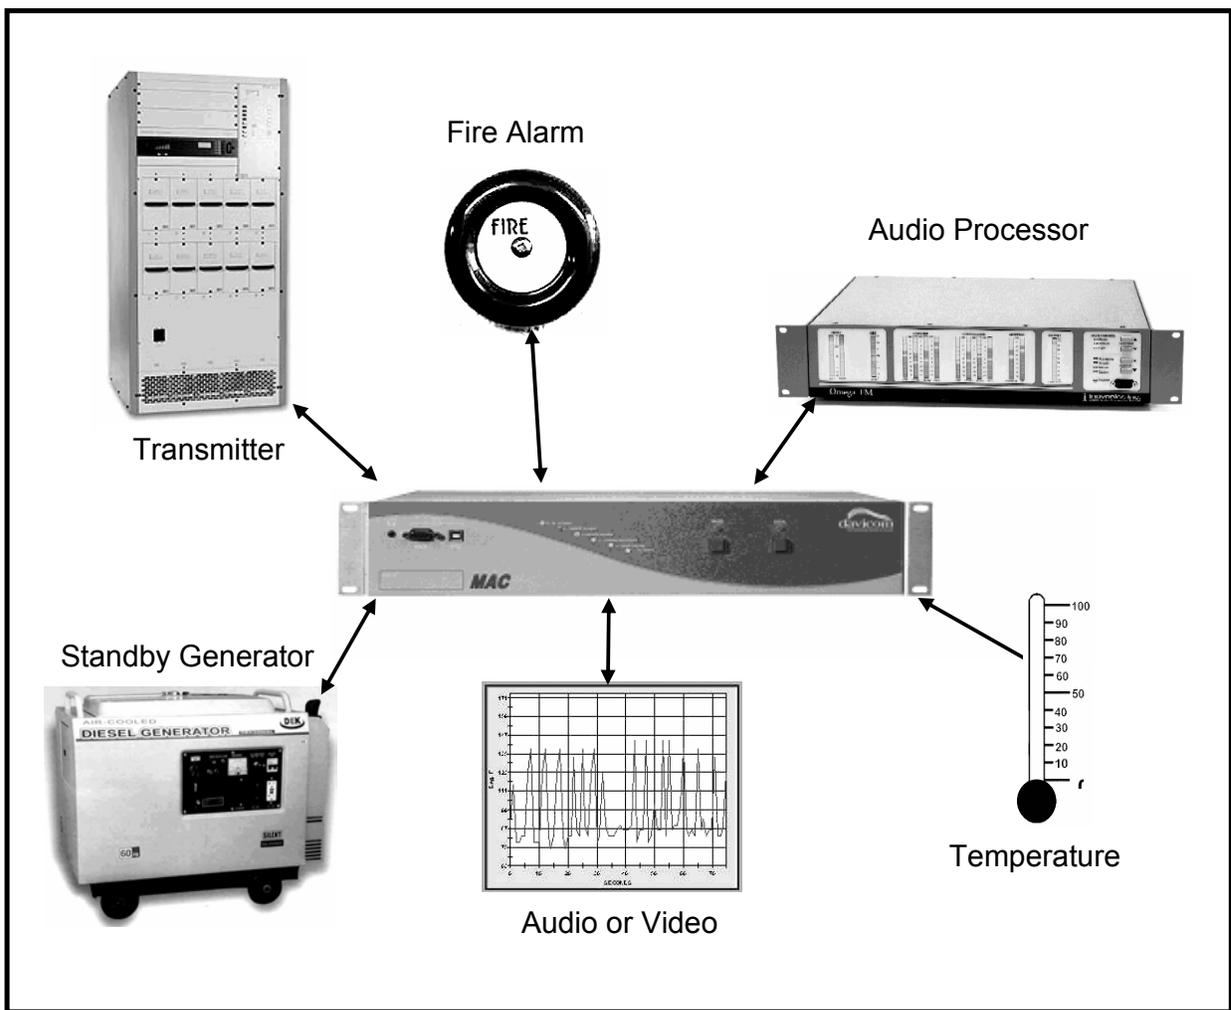


Figure1 An RTU with Associated Equipment

The remote control and telemetry unit is often referred to as an RTU (Remote Terminal Unit). An RTU is a general class of monitoring, and in some cases, control instrument. Typically, it would have a mixture of both analogue and digital input/outputs with one or more communications ports. The RTU inputs might allow a centrally-controlled computer system to monitor and control input

conditions, while some RTUs might even have the ability to perform control logic and could therefore execute their own control operations. Historically, most of the development of RTU equipment has been driven from the monitoring and control requirements of utility installations such as electricity, gas and water plants; very few RTU developers exist that concentrate on the specific telemetry requirements of the broadcast market.

More specifically, with respect to transmission sites, the RTU would usually be located in, or near to, the transmitter rack. Inputs can either be taken directly from equipment that has built-in sensors, or alternatively, external sensors can be added to provide readings for the RTU. These inputs can be Status Inputs such as an on/off relay, or a high/low voltage; these are often rather confusingly referred to as digital inputs. It is also possible to have Metering Inputs which comprise a voltage that is proportional to a particular reading, such as transmitter power or audio, for example. These are often called analogue inputs. An audio input would provide the ability to listen remotely to an analogue audio signal which would enable the reporting engineer to determine subjectively how a particular site is behaving, rather than relying solely on discrete, measured values.

Site parameters monitored by the RTU can include the following:

- Transmitter forward power
- Transmitter reflected power
- Power amplifier status lines
- Audio or video inputs
- Mains power presence
- Ancillary equipment status (e.g. audio processors, RDS generators, etc.)
- Intruder or fire detectors
- External door switch
- Internal and/or external temperature
- Standby generator

A good, intelligent RTU should have the ability to monitor a variety of different types of inputs and status conditions, together with the ability to switch to back-up systems once a pre-set threshold level has been met for a given period of time. It should also have the ability to prioritise alarms into primary, secondary, major and minor conditions. A suitable hierarchy should also be defined for the alarms, which can suppress the secondary failure reports and thus prevent multiple reporting of alarm conditions. A more comprehensive RTU might include internal timers that would work with specific, periodic functions, such as enabling mast lights, triggering daytime and night-time equipment settings, or performing scheduled checks of peripheral equipment such as stand-by generators and HVAC systems. Ideally, multiple RTU installations should have the ability to be centrally interrogated and controlled using proprietary software that would enable local, regional, national or global monitoring strategies to be implemented by the organisation.

MEANS OF COMMUNICATIONS

Having defined our RTU, with inputs and outputs, monitoring and controlling our transmitter, we need to establish a means of communications between the RTU and an operator. If the RTU switches to a back-up system to keep the transmitter on air, we need to be alerted of the fact so that a technician or engineer can assess what level of action needs to be taken. Technical staff must also have the means of calling the RTU to see the current state and interrogate the history of the unit.

Traditionally, the two most popular methods were by RS232 communications over a radio link, or by dial-up telephone. The first required that the studio be not too far from the transmitter or that several RF relays be used. The second meant that a telephone line had to be installed to site; something that is not always easy when the transmitter is at the top of a mountain or an equally remote area.

More recently GSM and other mobile phone systems have enabled dial-up connections to be implemented without running telephone wires to the site. This can present major cost savings, but mobile signals can often be difficult to receive at transmitter sites.

Some of the more recent RTUs will have more than one means of communications available to it, and becoming increasingly popular is IP connectivity. Larger transmission and network operators have started to install IP-based communications systems on their sites. These have generally been serviced by microwave or landline connections and enable IP control of equipment as well as performing other IP-based functions such as audio and/or video streaming, over private Wide Area Networks (WANs). IP connectivity is also available via dial-up or 'always on' connections such as DSL or cable and now directly by satellite. These however, generally go directly out to the Internet, and consequently, raise serious security issues previously not present when using point-to-point communications, such as dial-up, radio links or private WANs.

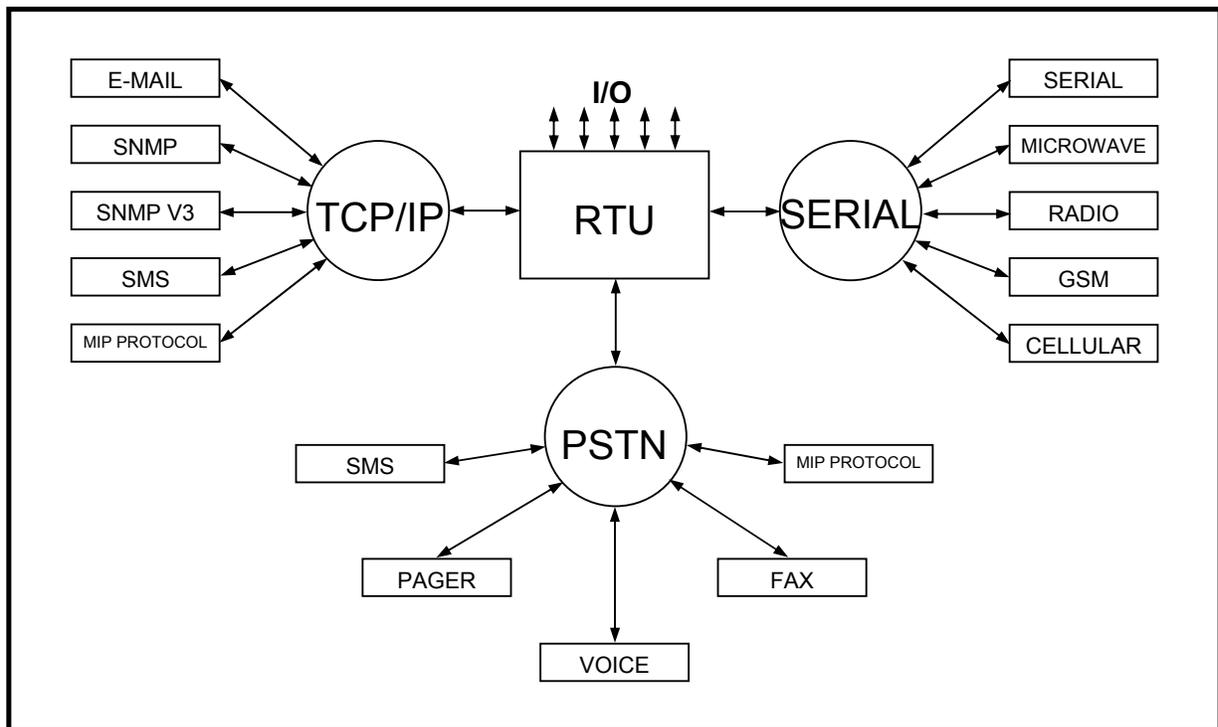


Figure 2 Connectivity Options for an RTU

THE MOVE TO IP

As computers become more ubiquitous, the advantages of IP-based communications become more apparent. Almost everyone has some familiarity with IP because of the widespread use and proliferation of both the Internet and in-home networks. IP has therefore presented itself as a standardised means of communications between many different types of equipment in a variety of different locations, and as such makes it an ideal connectivity choice when implementing network-wide control and monitoring strategies. This becomes especially important and beneficial when IP-enabled Transmit and Receive equipment are being installed into the sites of the near future.

By using IP, many of the limitations of dial-up or serial technology are overcome. In systems with more than one site, it is far easier for central management software to poll sites without the delays associated with normal telephone technology. For example, calling a site via a telephone line can typically take 30 seconds before a full connection is made even though the transfer of status may only take a few seconds. With IP, not only are connection speeds generally higher, therefore meaning faster transfer of data, but connection times are substantially reduced, meaning sites can be polled more frequently.

Alarms, too, can be delivered faster and by different methods. Simple Network Management Protocol (SNMP) is an industry standard for delivering alarms, or traps, and this is easily implemented within an IP framework, enabling RTUs which have traditionally been tied to proprietary software to be integrated into third-party systems. E-mails can also be sent and the opportunity of converting these into other forms of alarms, such as paging or SMS text messages, are now being offered by many companies on the Internet.

The fact that the central management software resides on an Internet-facing computer also allows operators to access the information remotely from any computer that has Internet access, anywhere in the world. Disaster recovery back-up locations can easily mirror the main software and allow a level of redundancy that would have been very difficult only a few years ago.

Once all of the RTUs and the management software are located on an IP network, the ability to create a "virtual site" becomes a possibility. Alarms from one site could be qualified by alarms from a different site, such that if a main transmitter fails, its relays are prevented from sending alarms to the management software. This helps an operator pinpoint the problem without being inundated with alarms because of a central failure.

The inherent security problems with IP connectivity are also becoming easier to overcome. Together with third party assistance, and using widely available public domain information on the Advanced Encryption Standard (AES), it has become increasingly possible to implement added security using encrypted communications.

Using point-to-point communications, such as serial-over-microwave or dial-up connections, has required little in the way of security, as access has been restricted to those who have knowledge of the installation or telephone number; usually a username and password was sufficient in these cases. Even a private WAN offers some level of security as 'outsiders' do not normally have access. However, opening up the RTU to the world by placing it on the Internet has given an unprecedented level of access to almost anyone.

ENCRYPTION

Generically speaking, increasing the level of security over IP using encryption algorithms means that the publicly transmitted data between the site and the RTU is effectively scrambled in such a way that only the RTU and the monitoring addressee can understand. Put simply, the encryption/decryption process between the RTU and the interrogating platform consists of passing the data exchanged between the two sites through an encryption/decryption 'engine'. Thus, all data communications via IP will have been transformed from ordinary ASCII-coded data to encrypted data. The standard for this encryption process has been set by the Federal Information Processing Standards 197 (FIPS-197) and adopted as the Advanced Encryption Standard (AES). A detailed explanation of this encryption and decryption process would fall outside the scope of this paper. However, this standard specifies the use of the Rijndael algorithm. This is a symmetric block cipher that can process data blocks of 128 bits, using cipher keys with lengths of 128, 192, and 256 bits.

Rijndael was also designed to handle additional block sizes and key lengths, however they are not adopted in this standard.

In the context of remote monitoring and control systems and encryption over IP, the bi-directional sequence of events is detailed in Figure 3. Essentially, the raw application data is generated by the RTU at the transmitter site and encrypted prior to being packaged into the transport layer for delivery via IP. Using the example of the Davicom MAC RTU system, the application data from the RTU is generated via the Mac Internet Protocol (MIP) and then encoded using the AES algorithm. The encrypted data is now in a secure form for transport over IP and only the RTU and the addressee know the encryption key. Using 128 bit encryption for example, there would be 2^{128} i.e. 3,402,823,669,209,384,634,633,746,074,300,000,000,000,000,000,000,000,000,000,000 possible combinations. At the other end, this process is essentially reversed; the data remains encrypted and 'safe' until it is decrypted immediately before entering into the MIP layer of the process and fed back into the application.

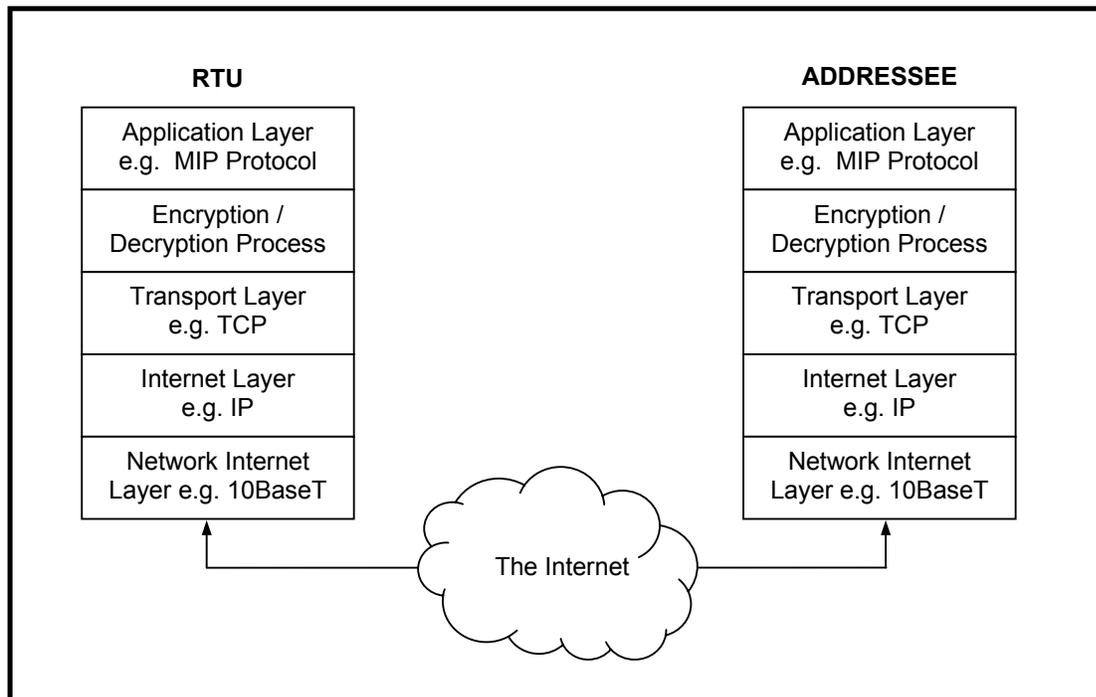


Figure 3 Secure Data Transfer over the Internet

CONCLUSION

Site monitoring and control continues to be a popular option for transmitter operators wanting to reduce both costs and downtime. RTUs are now available that are so sophisticated, they can be programmed to handle most situations at a transmitter site; switching to reserve equipment as and when necessary to maintain continuity of transmission. However, all this is of limited value unless the information is relayed to the site operator as speedily, efficiently and as securely as possible.

The emerging opportunity of IP connectivity, particularly over the Internet, makes data retrieval and exchange faster, more reliable, and more diverse than traditional methods of communications, but at the price of increased vulnerability to interception and attack. However, by using methods developed for high-security operations, such as financial institutions and government departments, secure encryption can be implemented which will allow the benefits of IP to be exploited without opening the system up to potential abuse.

The simple addition of an encryption/decryption engine inserted into the transport stream can render the communications unreadable to anyone who does not have the necessary key to unlock the code, and implementation can be accomplished easily using readily available technology. This level of security is essential to a professional, trouble-free installation, and in many ways it even improves on the level provided by point-to-point communications.

As long as the pitfalls of IP connectivity, particularly over the Internet, are addressed during the development stage, the advantages far outweigh the disadvantages and the move should be embraced as the step forward that it is.