Integrating Remotely Located Substations Into SCADA Systems: A Case Study Using Commercially Available Satellite Internet Service Providers for SCADA Communications

James W. Rice
Plumas-Sierra Rural Electric Cooperative

Nicholas C. Seeley
Schweitzer Engineering Laboratories, Inc.

Presented at the
DistribuTECH Conference
San Diego, California
February 4–6, 2007

Originally presented at the
8th Annual Western Power Delivery Automation Conference, April 2006

Integrating Remotely Located Substations Into SCADA Systems: A Case Study Using Commercially Available Satellite Internet Service Providers for SCADA Communications

James W. Rice, *Plumas-Sierra Rural Electric Cooperative*
Nicholas C. Seeley, *Schweitzer Engineering Laboratories, Inc.*

Abstract—**This paper presents observations and analysis experienced by a utility end user.**

As SCADA systems become a crucial part of daily operations for utilities, finding low-cost, reliable communications for substations in remote geographic locations often presents a challenge. Current SCADA communications methods, including frame-relay, T1, and fiber are often unreasonably expensive, in terms of installation costs and length of service contracts. Consequently, this can make integrating a utility's entire system into a SCADA network economically unfeasible or difficult to justify. In an effort to overcome these obstacles, Plumas-Sierra Rural Electric Cooperative (PSREC) recently installed redundant, commercially available satellite Internet connections, from two separate providers, at two of its substations. These satellite Internet connections serve as redundant communications paths for the substations and enable their integration into the SCADA system.

This paper discusses the architecture, security, implementation, results, and lessons learned from installing commercial satellite Internet services as a means of communications for SCADA systems, as well as the economic reality of using such a service over typical, more common SCADA communications methods.

I. Introduction

In an effort to keep up with the growing demand of their customers, PSREC, in northeastern California, began several studies to identify ways to increase system capacity and system reliability. Among the suggestions for increasing capacity was 1) acting as an interconnection node between Pacific Gas and Electric and Sierra Pacific utilities and 2) adding generation to their system, ideally through the construction of a wind farm. Both options required that PSREC make significant improvements to their existing SCADA system to increase their system reliability. Because PSRECs system would be used as an inter-tie, bridging two larger utilities, the ability for operators to monitor and control various devices at substations throughout the system is extremely important. Such control would allow conscious decisions to be made to restore or cut off power to areas as required.

The development of a SCADA system to integrate all of the substations in the PSREC system was to be divided into several stages, each of which would be designed, installed, and commissioned at various times over the course of several years. In March 2005, PSREC completed their initial stage where they integrated 5 of their 13 substations into the new SCADA system. This initial integration saw many challenges, the main challenge being how to economically integrate two remote substations into the SCADA system.

The main considerations of the new SCADA system design included the following:

- The inclusion of control center system monitoring with HMI control capability
- Engineering access to the relays at each substation
- Adequate functionality (integration of as many substation devices as possible)
- Transmission of as much useful information as possible from each substation to the control center
- Availability of all useful and necessary control functions
- The ability to remotely upload new relay settings, modify existing settings, or retrieve Sequential Events Reports (SER)—an immensely useful tool

After identifying the SCADA system criteria, PSREC began the logistical planning of its implementation, focusing on the required communications infrastructure, the common link in all the SCADA design considerations.

II. Communications Infrastructure

When designing a SCADA system, communications requirements are critical, namely: direction, latency, throughput, security, and cost. Taking each one of these factors into consideration, PSREC arrived at the following conclusions:

1. Communication must be two way; controls must be able to be sent and metering and target information must be able to be received.
2. An engineering access point must be available so that designated engineers can make settings changes to the relays remotely or download system event reports.
3. System functionality should take precedence over the overall speed of the system.

   Provided that data gathering capabilities, necessary to analyze and describe system events, could be implemented, the speed with which the SCADA system needs to respond was of secondary importance. Quick operation and update rates were desired, but multiple-second delays would be tolerable because real-time operation of the SCADA system was not required or desired.

4. System throughput is more important than latency.

   The system should be able to transmit large amounts of information, such as setting files, and large amounts of substation information.

5. Security must be implemented and possible areas of accessibility to intrusion must be addressed in light of the security regulations soon to be released by FERC regarding substation communications.

6. Cost must be proportional to the functionality that PSREC would be gaining.

A. Leased Lines

After consideration of these issues, PSREC decided to lease lines from the local telephone service provider, run a frame-relay Ethernet line to each station, and network these stations to PSREC's main office. The frame-relay option provided high throughput capabilities, a platform to ensure adequate security could be installed to prevent unauthorized access and electronic intrusion, as well as excellent speed.

The weak point of this option came down to cost. Three of the five substations that were to be integrated in the initial phase, Marble, Beckwourth, and Quincy, were located in more populated regions of PSREC's service area. The phone company already had the existing infrastructure to run lines to these stations, thereby putting the cost of monthly service in an acceptable range. The other two stations, Patton and Leavitt, were located in more remote regions, putting the monthly service cost for the Ethernet connection at around $1,500 per month, per station. In addition, the phone company required a three-year contract before the service could be installed. The total cost for integrating Patton and Leavitt stations into the frame-relay Ethernet network would amount to almost $60,000 over three years, including installation costs. This cost was not justifiable; considering the "non-critical" status of both stations, another solution needed to be implemented.

B. Satellite Internet

In addition to PSREC being an electric power provider in the northeastern region of California, they are also the local satellite Internet Service Provider (ISP) for the area. Recently, in a partnership with several other businesses, they launched their own satellite for the purposes of providing satellite Internet service to an even wider geographic area. Given this, PSREC began to ask if it would be possible to put a satellite dish at the Patton and Leavitt substations and run communications over the Internet.

A few weeks later in Portola, California, PSREC began testing this communications option. Initial tests consisted of polling a single relay for basic target and metering information and sending a control command using both OPC and DNP/IP protocols. Both protocols, after some adjustments to delay timeouts, worked extremely well. Initial cost estimates also made this idea look favorable. Monthly subscription to the satellite service would cost $75 per month, as compared to the $1,500 a month estimate for the frame-relay service. This accounts for an approximate savings of $50,000 over the course of 3 years.

The option of running SCADA communications over an Internet network, via a satellite connection, was selected for the Patton and Leavitt substations. However, before this decision was made, the consequences of communicating via satellite and communicating over the Internet were taken into full consideration.

III. Satellite Basics

In order to understand the pros and cons of using satellite communications, it may be helpful to provide some background on the basics of satellite communications, and specifically cite the performance capabilities of the satellite used in PSREC's application.

The capabilities of these two satellite-linked substations are very similar to that of its three frame-relay counterparts and include such capabilities as:

- Automation, control, and testing
- Control center SCADA and HMI
- Engineering access
- Remote data collection and analysis
- Communications via DNP/IP protocol

Satellites used for Internet service follow a geostationary orbit approximately 22,500 miles above the earth at the equator. Given the distance from point A—to satellite—to point B is roughly 45,000 miles, a round trip transmission of data at the speed of light takes nearly 500 ms, not including delays introduced by hardware and software (Fig. 1). While an extra half-second delay in the receipt of most metering and target information from a substation is no cause for alarm, this latency tends to cause problems for ordinary TCP/IP networks.

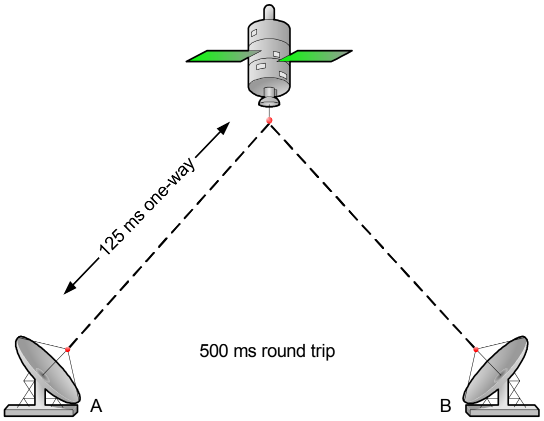

Fig. 1. Basic Satellite Latency

IV. TCP/IP Basics

The latency of a basic land-based Ethernet network falls somewhere between 40–100 ms (point-to-point). When a TCP session starts, information is sent out and the sender waits for its intended recipient to acknowledge it has received the data. If the recipient does not send acknowledgment in a given length of time, the sender assumes the data have been lost, resends the information, and will not transmit more information until the receipt of that information has been acknowledged. As noted earlier, the delay inherent to a satellite network is considerably longer than that of its land-based counterpart, thus, the sender does not receive the acknowledgements from the receiving end in a timely enough manner. The sender takes this lack of acknowledgement and/or delayed acknowledgement from the recipient as a sign that the network is congested (when, in actuality, the long response time is due to the satellite latency) and begins slowing down the rate of transmission to compensate for the network "congestion" and minimize the need for retransmissions.

On a normal, land-based network, the TCP sessions begin slowly, referred to as slow-start, and gradually increase in speed as the rate of acknowledgements confirms the ability of the network to accept larger amounts of data. This only hampers the already slow performance of the satellite network. TCP will never fully ramp up to maximize its speed and throughput because the acknowledgements are never received quickly enough.

V. Maximizing TCP/IP Performance Over Satellite

As long as Einstein's theory of special relativity holds true, the 500 ms delay will never be improved, making it essential to tackle the problem in another arena. In order to address the latency problems of satellite networks, a few solutions have been implemented that provide a work-around, making satellite TCP communications more practical. The current trend in satellite Internet technology is to employ techniques called "IP Spoofing" and "IP Acceleration."

As shown in Fig. 2, spoofing equipment, sometimes referred to as Performance Enhancing Proxies (PEPs), is placed at ground locations. This equipment, which may actually be part of the satellite modem and not an external piece of hardware, acts as an intermediary between the satellite hub and the workstation or remote site. When this spoofing equipment receives any Internet traffic addressed to the remote site, it automatically and immediately sends an acknowledgment response to the sender. Receiving this acknowledgement, the sender is fooled into thinking the acknowledgment came from the remote site and begins to send more packets. TCP, in turn, ramps up in speed and eventually levels out at the highest practical speed it can obtain, all the while the effects of the latency on the sending end are never felt because the acknowledgments are not coming from the remote site, but from a spoofing device that is accessible to the sender over the terrestrial network.

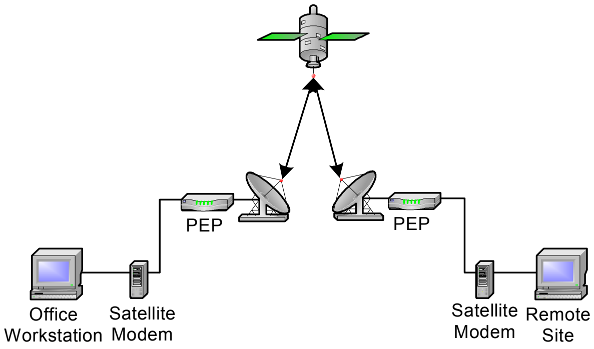

Fig. 2. Satellite Network With PEPs

These PEPs operate by decoding the TCP packet and using the TCP header information to spoof the remote site. With this information the PEP can act as the remote site and perform the necessary handshaking required by TCP, mimic the destination port, provide the proper sequence numbers, and discard the remote site acknowledgments, thereby fooling the sender. Therefore, it is of utmost importance that the PEP is able to read the TCP header information in order for it to successfully spoof the remote device and allow for maximum utilization of the TCP network.

Should the need arise for packet retransmission, due to actual communication interruptions and not merely perceived interruptions because of latency issues, the proxy itself will resend the data. In most cases, the PEP will cache the data segments within each TCP packet passing through it on the network. When the acknowledgement is received for a particular packet, those data are purged from the cache and the cycle continues. It is through this process that the PEP, using the TCP acknowledgements and/or timeouts, decides when and where to retransmit corrupted or lost packets.

VI. Security, VPNs, IPSec, and PEPs

Because PSREC opted for a commercially available satellite ISP for its SCADA communications, the issue of system security arises and with good reason. While the satellite itself provides some measure of security, the following explanation provides an excellent reason why more robust security should be implemented:

> "In absence of strong encryption, an attacker can purchase a VSAT terminal and with a basic knowledge of the data link protocol, hack the terminal to listen in on data intended for others. In order to listen, the hacker would need to reverse engineer the VSAT's embedded code, and command the terminal to tune to different frequencies and timeslots (for TDMA based systems) in order to receive transmissions from the satellite intended for other terminals. This is no kiddy script exploit, but once compromised the VSAT can act as a powerful packet sniffer… A simple brute force attack of the weakly encrypted data will yield its contents." [1]

The most common and widely accepted method of Internet security is Virtual Private Networks (VPN) with IPSec policies. While this paper will not go into extensive detail on the interworkings of VPNs using IPSec, some basic knowledge is necessary to understand how it will affect a satellite network or any other network using PEPs.

Basically, IPSec accomplishes securing TCP packets by encrypting the entire packet and sending it over the network. When it reaches its destination, the decryption process takes place using an Internet Key Exchange, which allows the destination site to decrypt the information back to its original TCP form.

Here lies the inherent problem of using IPSec over a satellite network. As described above, IPSec encrypts the entire packet, meaning the TCP header gets encrypted, as well as everything else. As we discussed earlier, in order for PEPs to maximize the performance of the satellite network, they must be able to decipher the TCP header in order to spoof the destination site. So, by adding IPSec encryption to provide data security, the PEPs are essentially useless because they cannot make sense of the encrypted TCP packet. This results in the same latency problems that we had before we introduced PEPs.

VII. Security Versus Latency

While VPNs using IPSec will work over most satellite connections, the question becomes whether or not the network latency can be tolerated. Such latency, while not rendering the SCADA system useless, may become a nuisance as polling rates may have to be slowed down to five- or six-second intervals and control commands will take up to five seconds to reach their intended destination. However, it should be noted that in most instances, latency can be greatly improved merely by selecting a satellite ISP who is using the latest technology. For example, PSREC noted a 50% decrease in latency when they switched to a provider with updated technology, including a newly launched satellite. However, if no such provider exists, options are available, such as Selective Layer Encryption (SLE) that can be configured to encrypt all but the TCP header, allowing PEPs to do their job and maximize the satellite connection. Leaving the header information unencrypted introduces a drawback because would-be intruders can gain valuable information in the header, such as the sending and receiving of device IP addresses.

VIII. The Security Decision

After investigating the available options, and some basic testing with their SCADA system, PSREC made a decision on which security method made the most sense for their application and needs. The basic architecture of the PSREC SCADA system is shown in Fig. 3. Notice that two stations are transmitting information over the Internet, thus the concern for security and need for additional research concerning security. As with any SCADA system, the exact architecture of the system should remain company confidential, because any security information that is made public only provides more information to those that could use the information for nefarious purposes.

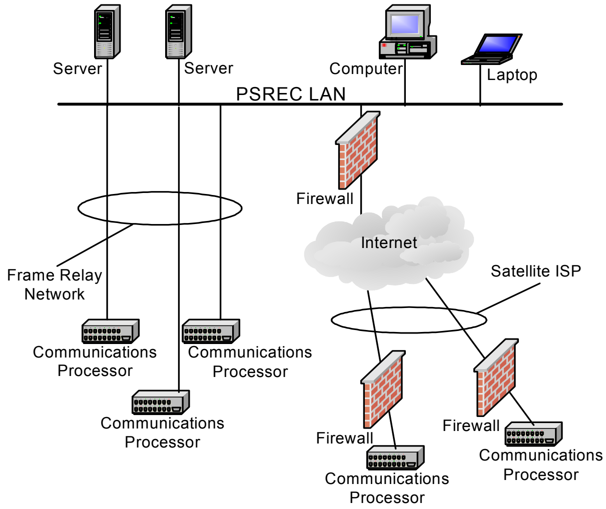

Fig. 3. PSREC SCADA System Architecture

IX. PSREC Results

After tracking communications for nearly 180,000 minutes, recording disruptions in the communications link, recording the causes of such interruptions when possible, some basic analysis was done to provide PSREC with details on the strengths and weaknesses of using satellite internet communications. Using a worst-case scenario approach, every minute that the satellite communications were not functioning was recorded and figured into the total downtime. This included time the satellite was down due to software upgrades within the satellite itself, routine satellite maintenance (late Sunday evening/early Monday morning outages), and, more significantly, downtime due to operators not being available to restore communications manually.

The vast majority of downtime was due to the unavailability of an operator to manually restore the connection after the satellite momentarily lost communications. These instances were a direct result of the security implemented within the network. More specifically, a result of the chosen system architecture was a momentary loss in communications, which could cause the system to close down completely until the operator manually starts communications again. Such an instance, while not terribly common for PSREC, did result in extended periods of downtime for stations when the operator was absent and not available due to vacation time, illness, or elsewhere on site to restore the connection.

PSREC has since remedied much of the problem by merely switching to a statically assigned IP address, versus the DHCP that they were using previously. This relieved much of the problem, because it was the dynamically assigned IP addresses that created problems within their security architecture.

A total of 179,970 minutes were monitored at both the Patton and Leavitt substations. The total downtime at Patton was 23,763 minutes and the total downtime at Leavitt was 16,247 minutes. The total uptime can then be calculated at 87% and 91%, respectively. While this number may seem somewhat disappointing, it is important to note that if we remove the downtime resulting from the operators inability to manually restart communications in a timely manner, we can subtract 21,695 minutes of downtime from Patton and 14,686 minutes from Leavitt. Recalculating the numbers, we get revised uptimes of 98.86% and 99.2% at Patton and Leavitt, respectively. More analysis has to be done regarding the recent system optimization using static IP addresses; however, the initial results look promising.

Other causes of signal loss are largely related to weather. The biggest enemy PSREC has seen so far is heavy, wet snowfall. Snow collects on the dish and on the head of the Transmit Receive Integrated Assembly (TRIA). This will cause the signal to be attenuated beyond use until the snow falls off or is cleared away. PSREC recently installed “unapproved” heaters (i.e., not recommended for the satellite dish by their ISP) on the dishes and TRIA assemblies to help speed the melting of the snow and hopefully avoid outages (due to snow) all together. As the winter progresses, PSREC will learn more and adapt to minimize signal loss.

The dishes mounted at both sites have been exposed to the heaviest rainfall the area has seen in nearly 50 years along with hurricane-force winds. Data show that under these conditions, the signal interruption lasted only a matter of seconds, and in a few cases minutes, but amounted to no more than a minor nuisance, rather than a serious problem for PSREC.

X. Conclusion

With a few initial setbacks resulting in less than desirable communications, PSREC has since implemented a successful communications scheme using commercially available satellite ISPs. After a few network modifications and optimizations, their satellite communications times are comparable to, although slightly slower than, their landline counterparts at a fraction of the cost. The result is an acceptable communications alternative for remotely located substations within a utility's system. However, words of caution need to be expressed. While this method of communications has been proven successful for PSREC, the results seen thus far are still variable enough to discount this method from being used as the primary means of SCADA communications for "system critical" stations. While the dependability of all methods of communications are subject to unanticipated events, satellite-based communications seem to have a higher degree of unknowns to account for, thereby making them a higher risk. However, for stations that are not critical to a utility's reliable operation, satellite communications offer a very economical alternative to other methods of communications.

XI. References

[1] Gregory Totsline, "Issues When Using IPSec Over Geosynchronous Satellite Links," p. 4, SANS Institute, August 12, 2002.

XII. Biographies

James W. Rice graduated from Cal State Northridge in 1975 with a Bachelor of Science Degree. After graduation he spent the next thirteen years as Captain of Engine Companies and Helicopter Operations for the local Fire Department. He went to work for Plumas-Sierra Rural Electric Cooperative nineteen years ago as a Meter Reader and has progressed to Manager of Technical Services in charge of the Meter Department and Substations.

Nicholas C. Seeley graduated from the University of Akron in 2002 with a B.S. degree in Electrical Engineering. After graduation Nic began working at American Electric Power in Columbus, Ohio for the Station Projects Engineering group where he focused on substation design work. In June 2004, Nic was hired at Schweitzer Engineering Laboratories, Inc. in the Systems and Services Division as an Associate Automation Engineer where he has been involved in the development, design, implementation, and commissioning of numerous automation-based projects.

20060413 • TP6233-01