30th Space Symposium, Technical Track, Colorado Springs, Colorado, United States of America
Presented on May 21, 2014

AUTOMATING REMOTE SYSTEMS

Roger Franklin
Crystal Solutions, Roger.Franklin@crystalcc.com

ABSTRACT

Operators of networks – data, video, communications, pipeline, transportation, and others – find themselves responsible for an ever increasing number of remote sites worldwide. For many, it is not practical to man all of those sites, given the resources and investment that would be needed. Further complicating matters is the fact that many of those sites are not easily accessible, so the practicality of manning those sites diminishes further. Those exact features that make a site impractical to man, however, make it suitable for a satellite communication connection.

Therefore, there is now a vast array of unmanned sites connected via satellites. However vital, unmanned systems create their own set of issues. Many of these are not built to react properly or even shut down when there are problems, thereby compounding issues as they continue to operate in a partially failed state. If there are no personnel on site, it can be very difficult to ascertain the quality of input data and output data for the system at the remote site. Having trained personnel on site also helps to warn of other issues, either before they happen or as soon as they do. But in the absence of personnel on site, you need to replace site awareness and knowledge with some semblance of automated intelligence.

This paper discusses how automation across a network can give operators the flexibility to monitor, command, and control systems and equipment at any of their locations worldwide. It will cover potential environmental issues, equipment malfunction, continued operations during communication interruptions, and disaster recovery, as well as how automation can improve satellite bandwidth usage efficiency.

SITES THAT REQUIRE MANAGEMENT – UNMANNED / INACCESSIBLE SITES

Data collection, in general, has gotten much easier. As the internet of things starts to explode, data will be available from just about any object you buy from shoes to litmus strips to milk jugs. In order to make intelligent business and operations decisions, we want access to as much data as possible. It's logical to deduce that more and more remote locations require systems to collect data with the expectation the data will be useful, somehow, to decision makers.

We want to know what the flow rate of product through pipelines are; determine how much leakage there is at pumps and junctions; calculate our fuel usage cost for transporting goods; understand the data throughput on switching stations; track the number of rounds fired by each rifle on the battle field; capture images which might show faults and cracks; enable data translation from one format to another; and more.

The sheer number and type of remote sites will continue to increase for all types of businesses and operations. For many of these remote sites, it is not practical to man all of the sites, given the resources and investment that would be needed. The result is that there is a growing array of unmanned sites worldwide.

Further complicating matters is the fact that the features of many of these sites that make them impractical for manned operations (geographic isolation, low real estate costs, etc.) also make them difficult to use traditional terrestrial communication methods. This is just one reason for the explosion of satellite communication networks.

The fact that satellite-based data connections are used presents competing requirements for the people and companies that own and operate the remote sites. The first requirement is to maintain operational efficiency and resiliency, which requires data from the remote sites; the second requirement is to minimize satellite bandwidth demand, which requires minimal data transmission to and from the remote sites.

In order to maintain operational efficiency and resiliency you need access to data from the systems at remote sites. And the world is telling you that you need as much data as you can possibly get because if you analyze the big data that will be available then you'll be able to make better decisions. There may be a flaw in that logic, and the law of diminishing returns probably applies, as well. The other reality is that you pay per bit of data that flows over the satellite communication network and you need to use that bandwidth smartly and efficiently.

The easy, expensive, and impractical answer to ensuring a resilient and efficient remote system is to just put a trained person at each site and let them handle everything. We can learn from that scenario by analyzing what would cause a human to get uncomfortable or nervous about a remote system's operation – waiting for equipment failure indicators to become active isn't good enough. We take the analysis a human would conduct, build the data collection capability, and then add decision making algorithms into an *intelligent control* system that can take automated actions – all contained at the remote site.

Some network-wide management systems are designed such that all components of a secondary site are directly controlled and monitored by servers and personnel at a central location. While this can work with a single, very simple secondary site, this strategy does not scale well.

Having an *intelligent server* at the secondary site offers several advantages: First, it allows relatively complex switching *rules* to be implemented in the case where equipment failures call for redundancy switching or ceasing of operations. Software based decisions can cover a multitude of 'what if' cases.

Second, it allows for the recording of high bandwidth diagnostic data (such as digitized spectrum traces) to occur *without burdening the connection channel* between the primary control site and the secondary sites.

Finally, it allows for monitoring (and responses to problems) to occur in a manner that *is not dependent upon the connection* between the primary and secondary site.

The ideal architecture would have a primary Monitor and Control server at the central location, fed by secondary intelligent servers at each remote site, the latter of which perform all of the low level 'bit twiddling' involved in interfacing with the remote system equipment.

THE ANALYSIS – POTENTIAL ISSUES WITH REMOTE SYSTEMS

The analysis required to maintain an efficient and resilient operation starts with predicting and understanding what could wrong. What are the failure possibilities and what could cause those failures to occur. So, we'll first look at a few examples of issues that can arise at remote systems.

Environment Issues

The weather, external and internal, can have a huge impact on all types of systems and it can often be overlooked. If you have remote, unmanned sites it is extremely challenging to be aware of changing environmental conditions, such as a rise in temperature, humidity, or heavy rain. Equally, forgetting to change an air filter in an

air conditioner can cause a rise in internal temperature that will ultimately affect equipment. And, of course, there's lightning. If you aren't monitoring the environment then issues can occur without warning and before anyone is aware the operations of the remote system could be severely impacted.

Equipment Malfunction

In a remote system, equipment malfunction can go unnoticed for a large period of time. Naturally, faulty backup equipment is even less noticeable than online equipment and any failure may result in major service malfunction if backup equipment must become active and fails to operate. Any and all equipment should be monitored for failure indications. When identified early, repair of faulty equipment can be cost-effectively scheduled while minimizing the risk to the remote system's operations.

GIGO

It's not unusual to have individual equipment at remote sites functioning properly, but the system as a whole doesn't. A primary cause is due to bad input which results in bad output – Garbage In, Garbage Out. It is also important to configure the right operational parameters to avoid cascading deterioration in operational quality. The output of one system may be within spec, but if it doesn't quite meet the input of the next system, then the initial spec needs to be adjusted. A few relatively minor issues at one site can cause much bigger issues further down the line by the time the product or transmission has been through other devices or other remote sites. This is especially true with IP video.

DATA COLLECTION – MONITORING THE RIGHT PARAMETERS

Once the analysis of possible remote system issues has been completed, a monitoring system needs to gather the necessary data, and store it at the remote location, so that periodic analysis can be conducted. The task of data collection, and subsequent control, is where the rubber meets the road when automating remote systems that must be resilient. Any monitoring and intelligent control system must have top-notch and reliable interfaces to the equipment that makes up the remote system.

There is a challenge in determining the right data to be collected. Once a monitoring system is installed to collect data from a remote system's equipment and environmental sensors, it is easy for that system to collect a significant amount of data that may (or may not) turn out to be useful in the future. The general rule of thumb is to collect as much as you can afford to collect.

The other challenge lies in not wasting precious communication bandwidth to transmit the massive amounts of collected data to a central location, unless it is absolutely necessary. The key lies in a system that does more than simply monitors and collects data or issue the occasional command. The system must be intelligent enough to determine what data is to be communicated back to a central location and when.

DECISION MAKING ALGORITHMS – INTELLIGENT CONTROL

An intelligent control system that monitors, manages, and controls a remote system should be able to filter data, conduct analysis, and make automated decisions at a remote site, in addition to collecting and storing equipment and environmental parameters.

It's not only possible, but desirable for an intelligent control system at a remote site to create concise summary reports of information learned, actions taken, and maintenance required for a managed remote system. The management system located at the central location aggregates the reports and summaries and provides further network-wide analysis as needed. In general, no news is good news. However, a minimal amount of

available bandwidth should be dedicated to period health and status summary packets to ensure that the communication link is functional.

INTELLIGENT CONTROL SYSTEMS AT REMOTE SITES

Remote sites can be extremely challenging to manage, but if you set up the automation system correctly, these issues (and many more) are manageable. If you are transmitting high power RF, you are obligated to be able to shut it off remotely, but in all cases it makes sense for your operation to have that control. With the correct automation algorithms, you can configure the system to attempt to resolve issues, and eventually shut down if unsuccessful. Many sites are not currently built to shut down when there is a problem, thereby compounding problems as they continue to operate. This can lead to satellite interference, amongst other things.

The number, type and monitoring requirements of these remote sites are unique to each network, so there is no 'one size fits all' solution. While it would be naïve for anyone to think they can envision all the possible automation scenarios when setting up a system, there are a number of important 'ground rules' that can be followed which will handle most types of problems.

Rule 1 – Move Intelligence 'to the edge'

Any management and control system to be considered must have the ability to distributed decision making capability. Intelligence at the edge of a management and control system is imperative in order to execute the decision making algorithms that are required. Further intelligence is required to store a significant amount of data that might be useful for analyzing future unanticipated events. Even though satellite communications are some of the most reliable available, the remote site must be able to operate and handle issues when the satellite communication link is not functioning as expected.

A key component of intelligent control capability at remote sites is for the control system to be able to monitor itself. The intelligent server is only effective when it is operating normally. If its power, computing, memory, or storage capacities aren't within reasonable limits then effective operation is in jeopardy. The analyze-potential-issues, collect-data, and make-decision thought process applies to the very system that collects the data and makes decisions.

Rule 2 – Have Secure Connectivity to each site

While the open Internet could be used to connect a central location to remote sites, one must be mindful of several potential problems. First, is the bandwidth sufficient and *guaranteed*? Second, does your IT infrastructure prevent hacking? If not, your status data at the remote site is exposed to the outside world, but what's worse is that the equipment at the remote site is ripe for *spoofing*. An ill-willed person, with just a little knowledge, could actually commandeer control of remote equipment and reprogram the settings for maximum bad effect.

If the Internet is employed, firewalls should be in place at a minimum. Additionally a VPN could be set up between the central location and the remote site. If this is done, the *scope* of the VPN should only include the IP segment at the central location where the main control and monitor server connects to the outside world – it should <u>not</u> include the entire LAN at the central location to minimize the risk of attacks from within the central location.

It may be a good idea to have multiple communication paths between the central location and each remote site, if possible. Internet connections using two distinct service providers would be a good idea (so long as they

don't both share the same physical copper or fiber to a site). One might also consider a permanent, contracted, non-Internet connection to the sites (with an Internet connection as a backup).

As discussed earlier, satellite connections are often used for remote sites. They are secure and reliable, and also support VPN tunnels. But for companies that don't own their own teleport or earth stations, there are usually one or more terrestrial connections between a teleport and the central location.

Rule 3 – Use Positive Acknowledgement Throughout

As a certain President once said, "Trust, but Verify!"

No part of a control and monitoring system should be 'fire and forget'. That is, every commanded action should be accompanied by a confirmation signal from the affected system or piece of equipment. Likewise, **status** information should never be given *on failure only*. The monitoring system should periodically *poll* each piece of equipment or system for health and status, and a non-reply should be treated as a failure. Of course, equipment should also assert error reports and conditions as soon as they occur, without being asked for it! But, relying just on notifications, or traps, from equipment is not sufficient.

Implementation of this requirement can be harder than it sounds. For example, just how long should one wait for a return signal from a piece of equipment once it is asked for? The delay time can be significant based on equipment 'busy-ness' and network traffic. Often there is a trade-off between the timeliness of response to legitimate problems versus monitoring 'over aggressiveness' resulting in false alarms.

In any case, when setting up a monitor and control system, the designer should ask, "If this piece of equipment or communication channel fails, will an error condition be asserted?" The answer needs to be 'YES' for all segments and equipment.

RESULTS

Managing remote sites can be challenging. But if some simple design rules are followed, major stumbling blocks can be avoided. The result is a Network Operations Center that seamlessly monitors and controls remote *sites just as intimately and effectively* as they do the central facilities, while minimizing bandwidth requirements to those remote sites.

