

Contents

Introduction

Control System Vulnerabilities and Exploits

Control System Attack Vectors and Examples: Field Site and Corporate Network 1

Threat Modeling NetDDE Vulnerabilities in Control Systems 2

Security Metrics for Calculating Risk in Control Systems

Measurable Control System Security through Ideal Driven Technical Metrics 3

SCADA Threat Modeling Using Attack Scenarios 4

Security Assurance Levels: A SIL Approach to Security 5

Detecting Attacks on Control Systems

Detecting Attacks in Power Plant Interfacing Substations through Probabilistic Validation of Attack-Effect Bindings 6

Modeling Flow Information and Other Control System Behavior to Detect Anomalies 7

Control System Protocols

Key Management and Cryptography for Advanced Metering Infrastructures (AMI) and Other Large, Low Power Networks 8

Introduction

Digital Bond created the SCADA Security Scientific Symposium (S4) in January 2007 to provide a venue to encourage control system security research and to develop a collaborative research community. S4 2008 is the second occurrence of the symposium.

Consider this Introduction our view of the state of control system security research, and we can provide our view succinctly. The quality and impact of control system security research is improving, but the quantity of control system security research is still poor and has shown no sign of growth over the past two years.

Our assessment of the quality of research is based on the S4 papers received this year and published in this Proceedings book. This work is by the far the best we have seen published in a variety of areas such as:

- control system security-related anomaly detection
- control system security metrics and risk theory
- vulnerabilities and exploits of control system applications that show the fallacy of many widely held beliefs on security perimeters and default permissions

The papers are digging deeper into the protocols, application software and unique elements of control systems, and we congratulate all of the authors on their outstanding work.

This is not meant to degrade last years work, which was the best available to that date. It is common that researchers build on past efforts and papers to move the overall body of research forward. In fact sharing research results in technical papers and fomenting an open control system security research community was a prime reason Digital Bond created S4. Little progress on control system security was made from 2001 to 2006, and we believe one reason for this was research was hidden behind closed doors.

While we were encouraged by the progress in the quality of control system security research, we were highly disappointed by the quantity of true control system security research being produced in 2007. We had hoped that the S4 selection team would face difficult decisions on what papers to include in the 2008 event. This was not the case. We chased down every possible research project we were aware of to fill the program. This uncovered some gems from previously unheard of authors, which is another S4 goal, but even with that effort of hundreds of telephone calls and emails there were not scores of papers of merit that did not make the cut.

Given the amount of money spent on research, the increase in control system research programs at Universities, and the amount of technical talent in asset owners and vendors, the community should be doing better. We are at a loss as to why more progress is not occurring but suggest two possible reasons. First, perhaps too much time is being spent rehashing old arguments on how the control system and IT cultures are different, responsible disclosure and information sharing, and competing control system security standards efforts. These may all be worthwhile discussions, but they are really only a diversion to the researcher.

Second, the nascent state of the control system security research community and limited control system knowledge by the researchers results in “research” efforts that are rehashing well known results from years past. The community does not need another research project showing Modbus can be spoofed because it lacks any source or data authentication, or that Digital Bond’s SCADA IDS signatures will identify selected attacks. The new researcher is not as much to blame for this as the research supervisor who should guide the researcher throughout the project. Admittedly there is a lag time between research starting and research results, and we are cautiously optimistic that at least some of the research institutions did a better job in 2007 of focusing precious research time and dollars on worthwhile efforts.

We addressed another major concern in our view of the state of the control system security research community at S4 2008 - - that is little or no security processes integrated into the control system vendors’ software development lifecycles. This is a ticking time bomb as most businesses and individuals have experienced with the frequent vulnerabilities and patches required for enterprise operating systems and applications.

Steve Lipner of Microsoft, who was a primary creator of Microsoft’s Security Development Lifecycle and co-author of a book with that same title, was the Day One Keynote Speaker. We highly recommend everyone read his book, and then go ask your control system vendor to describe and provide evidence of how security is integrated into their software development lifecycle. The results of this investigation will scare you.

Dave Aitel of Immunity, Inc. provided a different view to the vulnerability problem. What would an elite attacker do to identify new vulnerabilities in a control system application he had no previous experience with? Hacking has become a for profit business so it is not a stretch to foresee a circumstance where a bad actor hired a black or gray hat hacker to do just that. Dave’s keynote was further emphasized by two S4 papers that showed how minimal research efforts identified easily exploited vulnerabilities in popular control system applications.

To make matters even worse, the expected lifetime of control system applications is five, ten or even twenty years. The problem needs to be addressed now to hopefully have secure control systems in the 2015 / 2020 timeframe. Unfortunately it may take widespread exploited vulnerabilities like those that occurred with the worms in 2001 / 2003 before the control system community wakes up to software quality issues. However there is always hope that well meaning asset owners, vendors and researchers can raise awareness and force the issue, and we encourage readers to do just that.

So control system security researchers, including the team at Digital Bond, need to step it up in 2008! Our hope is the S4 2009 call for papers will be swamped with high quality research requiring difficult decisions on what papers to select.

Dale Peterson
Editor and S4 Chair
Digital Bond, Inc.

Control System Vulnerabilities and Exploits

Public discussion on control system application security to date has focused on missing security features and security patches that could not be applied without breaking the application's functionality. This is the proverbial low hanging fruit for an attacker. A control system protocol or application that is unable to authenticate the source and data integrity of a message is easily misused. Similarly, a control system application running on a host with many easily exploitable published vulnerabilities in an operating system, database or third party component, such as the Java Runtime Environment (JRE), can be easily compromised by even a low skilled attacker using one of the many open source or commercial exploit platforms. So the initial focus by the control system community on this low hanging fruit is warranted.

History has shown that software has bugs and some bugs will result in vulnerabilities. Reducing the number of bugs and resultant vulnerabilities is highly dependent on integrating security into the software development lifecycle. Unfortunately, most of the control system vendors have largely ignored the security mistakes and subsequent improvements made in the commercial software sector, and therefore likely have a bug / vulnerability problem in the magnitude of what was seen in Microsoft and other vendors circa 2000. This is obviously a bold statement that has yet to be proven true, but Digital Bond believes this is merely due to lack of research effort allocated to finding vulnerabilities in control system software rather than any inherent security. In fact, many of the classic attacks that were successful against web servers, databases, media players and other commercial software are likely to be successful against control system applications due to their ignoring lessons learned and reasonable software development practices.

S4 2008 highlighted the need to promote secure software development practices in the two keynote addresses. Steve Lipner of Microsoft discussed their Security Development Lifecycle and the impact this has had on reducing vulnerabilities in the Day One keynote. Dave Aitel of Immunity, Inc. approached the issue from the opposite side by explaining how an elite hacker would look for new vulnerabilities in control system applications in the Day Two keynote.

This section includes two papers that are just the beginning of an effort to research vulnerabilities in control system applications. The first paper shows two exploits, including a heap overflow exploit, that highlight particularly dangerous vulnerabilities because they can be exploited from outside the security perimeter. The second paper shows how a default configuration setting that is largely hidden from the asset owner can be exploited to own the host and launch attacks on other systems in the security zone.

Control System Attack Vectors and Examples: Field Site and Corporate Network

Eyal Udassin
C4 Security
Ramat Hasharon 47114, Israel
eyal.udassin@c4-security.com

Abstract: SCADA systems directly influence the lives and wellbeing of all civilians in almost any modernized country. The best site for an attacker to compromise in order to cause maximum damage is the control center. Much like Aikido, an attacker can use your strengths (centralized management of assets, multiple control applications) to his benefit.

In the whitepaper we will show two possible attack scenarios on the control center:

1. Attack from the field – how to study the weaknesses and exploit the control server software
2. Attack from the corporate network – research and vulnerability assessment of Proficy software to use it to take over the only authorized connection to the control center

As the aim of the whitepaper is to be as realistic as possible, we will demonstrate all three vectors on a very common target – an oil distribution company running GE Fanuc software: Simplicity 6.1 for their control system and Proficy Information Portal 2.6 for reports.

Keywords: Application Security, Heap Overflow, Historian, Field Site, Fuzzing Protocol Analysis

1 Introduction

The SCADA domain is a high-profile target for terrorist groups aiming their attacks at the civilian population of developed countries. The implication of this from the security aspect is that the attacker is not your ordinary "arbitrary 14 year old internet hacker", but a well-funded, well equipped and highly motivated group of computer and control experts.

According to Dr. Boaz Ganor, director of the Institute for Counter-Terrorism (ICT), the definition of terrorism is "the deliberate use of violence aimed against civilians in order to achieve political ends" [1]. SCADA systems directly influence the lives and wellbeing of all civilians in almost any modernized country, and the best site for an attacker to control in order to cause maximum damage is the control center. Much like Aikido, an attacker can use your strengths (centralized management of assets, multiple control applications) to his benefit.

Threat Modeling NetDDE Vulnerabilities in Control Systems

Jason Holcomb, Charles Perine
Digital Bond, Inc.
Sunrise, Florida
holcomb@digitalbond.com, perine@digitalbond.com

Xavier Panadero, Lluís Mora
Neutralbit
Barcelona, Spain
xpanadero@neutralbit.com, llmora@neutralbit.com

Abstract: Microsoft's NetDDE protocol is used in many control system applications to exchange data between two disparate systems, such as populating an Excel spreadsheet from data on a historian. NetDDE clients and servers are found in asset owner written programs and free or low cost utility programs. NetDDE interfaces are also still found in popular SCADA and DCS systems.

This paper provides a brief overview of NetDDE server configuration with emphasis on the access control features available in NetDDE shares. With that background, a vulnerability resulting from poor NetDDE share configuration in Wonderware's InTouch HMI version 8.0 default installation is described. A tool, *nbDDE*, is introduced and demonstrates how an attacker could exploit the InTouch vulnerability and other misconfigured NetDDE shares in a variety of methods. The paper also includes a discussion on how the NetDDE shares could be modified to reduce the risk to prevent or limit the exploit.

The paper concludes with a discussion of how integrating security, particularly threat modeling, into the software development lifecycle could have identified and addressed this vulnerability prior to the release of vulnerable code.

Keywords: NetDDE, DCOM, Wonderware, Threat Modeling

1 NetDDE Overview and Use in Control Systems

Dynamic Data Exchange (DDE) is a protocol designed to share data between programs that run on Microsoft Windows. The protocol defines a set of messages and guidelines for exchanging data using shared memory. The data exchange can either be a one-time transaction, or a continuous exchange as new data becomes available.

Network DDE, or NetDDE, manages the network communication necessary for DDE communication between different computers. DDE and NetDDE are used extensively in SCADA applications. In fact, NetDDE was developed by SCADA software vendor Wonderware. Wikipedia describes the origins of NetDDE:

Security Metrics for Calculating Risk in Control Systems

Security metrics was an area of high interest at S4 2007, and in the past year excellent work has been done to move this critical research area forward. Why is it critical? C-level executives will only spend money on security when they are convinced it is addressing a real risk; this is a wise decision by the executives. The allocation of money and other resources in a security program should follow a risk management approach. The lack of specific, credible information on the threat, vulnerability, and impact of security controls in SCADA and DCS is necessary to make the business case for security and efficiently apply security resources.

There are three very different and innovative papers in this section. The first paper is the most comprehensive and statistically sound work to date on measuring control system security by a set of technical metrics. The second paper veers away from trying to develop threat metrics and instead takes an attack scenario approach to characterize the threat. The third paper begins to tackle the often discussed question of why can't Security Assurance Levels (SAL's) be developed and products tested and certified to these SAL's in a similar manner to what is happening today in the safety world with Safety Integrity Levels (SIL's).

Measurable Control System Security through Ideal Driven Technical Metrics

Miles McQueen, Wayne Boyer, Sean McBride
Idaho National Laboratory
{miles.mcqueen, wayne.boyer, sean.mcbride}@inl.gov

Marie Farrar
Securicon, LLC
marie.farrar@Securicon.com

Zachary Tudor
George Mason University
ztudor@gmu.edu

Abstract: The Department of Homeland Security National Cyber Security Division supported the development of a small set of security ideals as a framework to establish measurable control systems security. Based on these ideals, a draft set of proposed technical metrics was developed to allow control systems owner-operators to track improvements or degradations in their individual control systems security posture. The technical metrics development effort included review and evaluation of over thirty metrics-related documents. On the bases of complexity, ambiguity, or misleading and distorting effects the metrics identified during the reviews were determined to be weaker than necessary to aid defense against the myriad threats posed by cyber-terrorism to human safety, as well as to economic prosperity. Using the results of our metrics review and the set of security ideals as a starting point for metrics development, we identified thirteen potential technical metrics - with at least one metric supporting each ideal.

Two case study applications of the ideals and thirteen metrics to control systems were then performed to establish potential difficulties in applying both the ideals and the metrics. The case studies resulted in no changes to the ideals, and only a few deletions and refinements to the thirteen potential metrics. This led to a final proposed set of ten core technical metrics. To further validate the security ideals, the modifications made to the original thirteen potential metrics, and the final proposed set of ten core metrics, seven separate control systems security assessments performed over the past three years were reviewed for findings and recommended mitigations. These findings and mitigations were then mapped to the security ideals and metrics to assess gaps in their coverage. The mappings indicated that there are no gaps in the security ideals and that the ten core technical metrics provide significant coverage of standard security issues with 87% coverage.

Based on the two case studies and evaluation of the seven assessments, the security ideals demonstrated their value in guiding security thinking. Further, the final set of

core technical metrics has been demonstrated to be both usable in the control system environment and provide significant coverage of standard security issues.

Keywords: Cyber Security Metrics, Control System Security

1 Introduction

Electronic control systems that operate much of a nation's critical infrastructure are increasingly connected to public networks. Therefore, control systems and the associated critical infrastructure are at risk from cyber attacks. Meaningful metrics are needed to make informed decisions that affect system security. The Department of Homeland Security National Cyber Security Division supported development of a small set of security ideals and associated metrics to provide control system owners/operators guidance in managing their system security.

A metric is a standard of measurement [1]. The scope of this paper is limited to quantitative technical metrics. A cyber security technical metric is the security relevant output from an explicit mathematical model that makes use of objective measurements of a technical object. Other types of metrics (such as operational and organizational metrics, and metrics that are qualitative such as "low impact" or "highly unlikely") can provide insights about security but are beyond the scope of this work.

An important use of technical metrics is in the estimation of risk where risk is defined as the probability of an event times the consequence of the event. The risk we would like to measure is the expected value of the loss from cyber attacks per unit time. Risk is usually measured in dollars or lives. The estimation of risk could provide the ability to weigh the benefits versus costs of security counter measures. However, a credible estimation of cyber security risk in real world control systems is not currently feasible because the problem involves an unpredictable intelligent adversary and very complex systems. Previous work [2] proposed "mean time-to-compromise" as a security metric and proposed a simple method for calculating it as a function of the number of known vulnerabilities. A method was also proposed for estimating risk reduction for a simple control system using the mean time-to-compromise metric [3]. Unfortunately, those methods require simplifying assumptions that are not valid in general.

In our opinion a good set of metrics should support the concept of risk estimation within the practical constraints of what is currently objectively measurable and under the control of the defender. A good set of metrics should have the following attributes: The number of metrics should be small (less than 20)¹ to be manageable; the metrics should be easy to understand, measurable and objective; the metrics should be directly related to security risk; and the set of metrics should represent the most important measurable security attributes of the system. Previous work [4] introduced the concept of security metrics based on seven ideals of security and proposed a set of metrics intended to meet the above criteria. This paper is an extension of that work.

¹ NIST 800-55 [SBS03] recommends that to keep the set of metrics manageable, the number of metrics should be about 10 and no more than 20.

SCADA Threat Modeling Using Attack Scenarios

Ralph Langner
Langner Communications AG
rl@langner.com

Bryan L. Singer
Wurldtech Security Technologies
BSinger@wurldtech.com

Abstract: The threat of cyber security failures in industrial automation and SCADA is ever present and even increasing, but there is often little understanding as to what the threat actually may be. While engineers and security professionals can often measure the vulnerabilities in a system, and can apply quantitative logic to the impact analysis, one of the key components to understanding risk often goes ignored. The challenge is that the threat is a key component of risk.

Understanding the threat is tantamount to creating awareness, without which a business case and justification for security expenditures cannot be developed. Most organizations exist happily in the thought that, “this could never happen to us,” but to those informed about security, it is not a matter of if, but when. Without sufficient threat data, however, developing threat models and attack trees are a matter of academic interest with often low perceived value to the business community.

This paper takes an alternate approach to building threat data. Leveraging past known cases of cyber security failures as well as recent research and development into the areas of industrial cyber security, this paper adopts an approach of creating scenarios that explain the why, how, and consequences of various attack vectors and compromise scenarios. The paper begins with expressing the traditional based approaches to understanding risk, the various strengths and weaknesses in ascertaining risk, and then tours various threat scenarios that can be used or extended by the industrial automation and SCADA community to help create an accurate picture of risk and utilize in business case development for security countermeasures.

Keywords: Threat Modeling, Cyber Terrorism, Risk Assessment

1 No Threat, No Risk: Towards a Scenario-Based Risk Model

If there is a terra incognita in the SCADA security landscape, it's threat country. We know quite a lot about vulnerabilities, about impact, and we even have a good selection of countermeasures available. But when it comes to determining threat, we know little,

Security Assurance Levels: A SIL Approach to Security

Dr. Nate Kube, Bryan L. Singer
Wurldtech Security Technologies
Vancouver BC Canada
nkube@wurldtech.com, bsinger@wurldtech.com

Abstract: A Safety Integrity Level (SIL) is a statistical representation of the reliability of the Safety Instrumented System (SIS) when a process demand occurs. SIL's are correlated to the probability of failure of demand (PFD), which is equivalent to the unavailability of a system at the time of a process demand. Given the de facto acceptance of SIL and the widely recognized interrelationships and interdependencies between safety and security, many have pushed for the adoption of a security level concept similar to safety in the security environment. Several efforts have been directed at this including security assurance levels defined in documents such as NIST 800-53. This paper takes a critical look at the SIL concept, its overall strengths and weaknesses as applied to security, and proposes general models for use within the security arena.

Keywords: Security Assurance Level (SAL), Safety Integrity Level (SIL)

1 Introduction

The continued deployment of Ethernet enabled devices in the industrial automation and control systems community has exposed a number of new potential risks previously unrealized in both IT and process control safety and security. Industry can draw upon previous history to exact a similar model. The Safety Integrated Levels (SIL) as defined in ISA – 84 [5] and internationalized in ISO 61508 [3] and IEC 61511 defined a four layer model for dealing with process safety requirements in two separate categories: hardware safety integrity and systematic safety integrity. This probabilistic model utilizes Failure Model and Effects Analysis to project risk and damages during a system failure, and provides a clear model by which systems can be assessed or implemented to achieve the desired level of safety risk reduction. While many draw parallel between security and safety, this is one intersection in which the lessons are clear and provide an effective model for security as well.

The following paper outlines the history of SIL, testing rigors as applied to control devices, models by which security can be evaluated similar to SIL, and implications for implementing this model both from an asset owner's and a vendor's perspective. The authors' intent is not to invalidate any existing safety models, but rather to show how extensions upon the SIL concept, whether applied as a separate standard or as a basic extension, provide an excellent framework upon which vendors can design, test, and market more resilient components. Asset owners are then free to determine and meet

Detecting Attacks on Control Systems

Control systems are much more uniform than a typical enterprise network. There is a limited amount of authorized use of the network, and this use is restricted to functions essential to the monitoring and control of a process. The communication flow often varies little from hour to hour or day to day. There are requests followed by responses, and even the requests and responses issued to individual points or groups of points is typically limited and ordered.

This structure offers tremendous opportunities for anomaly detection algorithms to identify attacks on control systems. Anomaly detection was discussed in two papers at S4 2007, and the two papers in this section build on that work. The first paper studies the IEC 61850 protocol and uses structural equations to probabilistically characterize the legitimacy and abnormality of IEC 61850 traffic. The second paper shows how flow information available from most routers can be used to identify attacks including changes in the volume of traffic between various points in the control system network.

Detecting Attacks in Power Plant Interfacing Substations through Probabilistic Validation of Attack-Effect Bindings

Julian L. Rrushi and Roy H. Campbell
Department of Computer Science
University of Illinois at Urbana-Champaign
201 N. Goodwin Avenue
Urbana, IL 61801
{jrrushi, rhc}@uiuc.edu

Abstract: In this paper we provide a mathematical approach to detection of attacks on relays in electrical substations speaking IEC 61850, an abstract industrial protocol devised by the Technical Committee 57 of the International Electrotechnical Commission as a standard for substation communications. Our contribution regards those electrical transmission substations which interface with the generators of a power plant through step-up transformers. In this paper we take as an instance power plants which use nuclear reactors as a source of energy.

The basis of the proposed approach is formed by structural equations which semantically model the relations between operational variables of substation and nuclear power plant components as monitored by the respective control systems. Causality relations investigated via structural equations are reflected on Bayesian belief networks to probabilistically characterize the legitimacy and abnormality of IEC 61850 traffic. We then employ the stochastic activity network formalism to construct composed models of substation operation from which we derive intrusion detection rules.

Keywords: Anomaly Intrusion Detection, IEC 61850, Stochastic Activity Networks, Structural Equations, Bayesian Belief Networks

1 Introduction

In this paper we discuss an intrusion detection approach for IEC 61850 [5] which builds upon statistical and probabilistic methods, namely structural equations modeling (SEM) [17], Bayesian belief networks (BBN) [16], and stochastic activity networks (SAN) [10, 19]. The work discussed here is part of a larger intrusion detection framework that we have devised for operation in nuclear power plants. This framework is based on SAN modeling of a fusion of control protocols used in nuclear power plants, for example Modbus [13], protocols used in electrical substations, for example IEC 61850 [5], and the operation of physical components of a nuclear power plant and an electrical substation, respectively. The electrical substations considered in this research are those

Modeling Flow Information and Other Control System Behavior to Detect Anomalies

Brian Moran, Rick Belisle
IBM Internet Security Systems
6303 Barfield Road
Atlanta, GA 30328
bmoran@us.ibm.com, rbelisle@us.ibm.com

Abstract: The stereotypical behavior of most control systems lends itself to modeling the system and detecting anomalies that could be cyber attacks or other dangerous network activity. In this paper we will discuss and demonstrate how the flow information available from routers and other network devices can be used in an anomaly detection system (ADS) to detect attacks from the communication flows from a source / destination / service perspective. We will also use this flow information to identify anomalies in the volume of communication on the network as a whole or between any two hosts. Finally, the paper will discuss the potential to model control system protocol usage to detect anomalies caused by cyber attacks.

Keywords: Netflow, network behavior analysis, network anomaly detection, SCADA

1 Introduction

Most SCADA and Process Control Systems in use today were developed years ago, long before public and private networks or desktop computing were a common part of business operations. As a result, the need for network security measures within these systems was not anticipated. At the time, good security for SCADA systems meant limiting and securing the physical access to the network and the consoles that controlled the systems. Planners rationalized that if they were suitably isolated from any physical entryways, and if access was limited to authorized personnel only, the systems were fully secure and unlikely to be compromised.

The increasingly networked and linked infrastructure of modern SCADA systems has rendered those early security plans obsolete. As companies have added new applications, remote access points and links to other control systems, they have introduced serious network risks and vulnerabilities that cannot be addressed by their physical control policies. Often, these risks are underestimated due to the complexity of the network architecture, the lack of formal network security guidelines and assumptions about the privacy of the network. Organizations are now realizing the security of these systems means more than physically separating the system and the components they control and monitor.

By analyzing network traffic flow and modeling behavior of the control system, commercially available security systems can offer the much needed protection for

Control System Protocols

The control system environment, particularly a SCADA environment, introduces a number of challenges to developing secure protocols - - even new protocols for new systems without any legacy issues. The field site equipment is likely to be low power, low cost, low bandwidth and in a physically exposed environment. Advanced metering infrastructures (AMI) are a prime example of these limitations with networks potentially in the millions of devices with unit prices as low as \$50. The paper in this section discusses the challenges in developing security protocols and implementations for AMI and then offers key management and cryptography proposals for AMI.

A second control system protocol paper was presented at S4 on OPC Unified Architecture (UA). This paper is available to subscribers on Digital Bond's website and will also be published as a chapter in a Digital Bond Press book titled OPC UA Exposed in 2008.

Key Management and Cryptography for Advanced Metering Infrastructures (AMI) and Other Large, Low Power Networks

Grant Gilchrist and Darren Highfill
EnerNex Corporation
170C Market Place Blvd. Knoxville TN 37922-2337
grant@enernex.com, darren@enernex.com

Abstract: Advanced Metering Infrastructures (AMI's) have limitations that require unique information security approaches. Some of the characteristics addressed in this paper include the mixture of one-way broadcast and two-way communication, the limited processing power found in organizational issues in managing cryptographic credentials, the sheer number of endpoints (in the millions), and the expected type and frequency of message traffic to be addressed.

This paper describes in detail the currently proposed authentication mechanism for broadcasting controls to Programmable Communicating Thermostats (PCT's) in the state of California. PCT's are a key component of the advanced metering and demand response infrastructure being deployed there. It discusses the alternatives that were considered in developing this mechanism.

This paper goes on to discuss some of the ideas currently under consideration for securing the two-way AMI networks, and how they might relate to this PCT solution.

Keywords: AMI, AMR, PCT

1 Introduction

The Programmable Communicating Thermostat (PCT) system is envisioned as a mechanism for reducing the amount of electrical power in use by the state of California during critical periods up to several hours in length. The intent of this system is to permit a statewide organization (as yet undefined) to send out a broadcast message that causes all thermostats in the state to automatically adjust themselves so the household uses less power. In this way the state can avoid blackouts or power system instability.

The broadcast message may be an emergency event, requesting an immediate curtailment of load, or it may be a price signal indicating that the cost of power will increase drastically. A price signal would permit some consumers to choose to continue using power at higher cost if they felt they could not curtail their power usage.

The state is in the process of mandating that all new homes built starting in 2009 will include a PCT. The legislation used to make this requirement is known as "Title 24" [1] [2]. This legislation will affect a fairly small number of homes relative to the size of the population, approximately 60,000 per year.