Homeland Security

September 2011

ICSJWG Quarterly Newsletter

Register for the ICSJWG 2011 Fall Conference!

The Industrial Control Systems Joint Working Group (ICSJWG) 2011 Fall Conference is around the corner, October 24-27, 2011, at the Westin Long Beach in Long Beach, California, so don't forget to register at http://www.regonline.com/Register/Checkin.aspx?EventID=990303!

Some agenda highlights include:

- *Unclassified Threat Briefing (plenary session)*
- *Panels on Vulnerability Disclosure and Regulation*
- *Presentations on Standards, Technology, and Case Studies*
- *ICSJWG Subgroup Meetings*
- *Intermediate Training*

Check out these features and other conference activities in the working agenda: http://www.us-cert.gov/control_systems/icsjwg/conference.html

CSSP Rolls Out CSET Version 4.0

The Control Systems Security Program (CSSP) has released Version 4.0 of the Cyber Security Evaluation Tool (CSET). This new version of the tool can be downloaded from the CSSP website - http://us-cert.gov/control_systems/satool.html. The Version 4.0 release incorporates several new standards, such as North American Electric Reliability Corporation (NERC) Critical Infrastructure Protection (CIP) Revision 3, Nuclear Regulatory Commission (NRC) Regulatory Guide 5.71, a new key requirements set, and Version 7 of the DHS "Catalog of Security Requirements: Recommendations for Standards Developers." The new CSET tool also includes a fully revised report suite with complete gap rankings, new diagramming functionality, and a new resource library. The updated tool supports evaluations of business and industrial control systems.

About the ICSJWG

The ICSJWG is a collaborative and coordinating body operating under the Critical Infrastructure Partnership Advisory Council (CIPAC). The ICSJWG provides a vehicle for communicating and partnering across all critical infrastructure and key resources (CIKR) sectors between federal agencies and departments as well as private asset owner/operators of industrial control systems. The goal of the ICSJWG is to continue and enhance the facilitation and collaboration of the industrial control systems stakeholder community in securing CIKR.

For more information, visit http://www.us-cert.gov/control_systems/icsjwg/

Table of Contents

ICSJWG Subgroup Status

Below is an update on the progress of the ICSJWG subgroups. If you would like to become a member of any of the subgroups, send an email with your contact information to icsjwg@dhs.gov or contact the co-chairs directly.

- **Roadmap to Secure Industrial Control Systems Subgroup**
  GCC Co-Chair: Perry Pederson (Perry.Pederson@nrc.gov)
  SCC Co-Chair: Tim Roxey (Tim.Roxey@nerc.net)

  The Roadmap subgroup completed the initial Cross-Sector Roadmap document and the revised Charter. They have finalized both documents and forwarded them for final review and approval to the ICSJWG Government Coordinating Council (GCC)/Sector Coordinating Council (SCC).

- **Vendor Subgroup**
  GCC Co-Chair: Marty Edwards (Marty.Edwards@dhs.gov)
  SCC Co-Chair: Eric Cosman (ECCosman@dow.com)

  The Vendor subgroup is busy with two sub-committees committed to developing white papers. First, the Cross-Vendor Position Paper outlines the direction that the ICS community should take to improve security and the importance of owners/operators, vendors, and system integrators collaborating to design, implement, and maintain ICS security. Second, the ICS Common Vulnerability Disclosure Framework paper is intended to provide a consensus-based foundation for ICS vendors and integrators working to develop a vulnerability disclosure policy. The paper provides recommended ranges and formats for different aspects of the disclosure process.

- **Workforce Development Subgroup**
  GCC Co-Chair: Keri Nusbaum (Keri.Nusbaum@dhs.gov)
  SCC Co-Chair: Michael Glover (M.Glover@prime-controls.com)

  The Workforce Development subgroup is currently reviewing its charter and assessing the possible addition of two functional roles to the National Initiative for Cybersecurity Education (NICE) Specialty Area Framework.

- **Research & Development Subgroup**
  GCC Co-Chair: Dr. Douglas Maughan (Douglas.Maughan@dhs.gov)
  SCC Co-Chair: VACANT

  The R&D subgroup is seeking applicants for SCC co-chair. If anyone is interested in applying for the position, please send an email to icsjwg@dhs.gov. The subgroup also is looking for subject matter experts who have the time and resources to dedicate to the position.

Homeland Security Information Network

HSIN is the information sharing tool used by ICSJWG subgroup members. All subgroup members can stay abreast of upcoming meetings through the calendars and subgroup reference materials in HSIN (e.g., charters, meeting minutes, agendas, etc.).

In addition, the "Alert Me" feature notifies users of changes to the portal, which eliminates the need for users to constantly log in to find out if updates have been made. Alerts can be sent immediately, daily, or weekly. To sign up for alerts, click on the "Alert Me" link on the left-hand side of the ICSJWG homepage and choose your delivery option. ICSJWG subgroup members who still need access to HSIN can send an email to icsjwg@dhs.gov to request an account.

- **If you do not currently have a HSIN account**, please provide your name, company, contact information, critical infrastructure sector, and ICSJWG subgroup affiliations to icsjwg@dhs.gov.

At this time, DHS is not able to grant non-U.S. citizens or those residing outside of the U.S. and its territories access to the HSIN portal. The owners of the HSIN portal are reviewing sharing agreements concerning information posted to the site. Until that process is complete, international user accounts will be on hold. ICSJWG Communications will contact all international members immediately if there are new developments.

Participation is Key!

Your participation and input is **critical** to the success of these subgroups and to the overall mission of the ICSJWG in coordinating cybersecurity efforts to secure industrial control systems across the nation's critical infrastructure. Please email the co-chairs or icsjwg@dhs.gov to get involved with one or more of the subgroups.

Advanced Training Events Scheduled for Fiscal Year (FY) 2012

CSSP is currently offering advanced cybersecurity training sessions at the Control Systems Analysis Center located in Idaho Falls, Idaho. These sessions provide intensive hands-on training in protecting and securing control systems from cyber attacks, including a realistic Red Team/Blue Team exercise that is conducted within an actual control systems environment. It also provides an opportunity for attendees to network and collaborate with other colleagues involved in operating and protecting control systems networks.

- **Day 1:** Welcome, overview of DHS CSSP, a brief review of cybersecurity for industrial control systems, a demonstration showing how a control system can be attacked from the internet, and hands-on classroom training on Network Discovery techniques and practices.

- **Day 2:** Hands-on classroom training on Network Discovery, instruction for using Metasploit, and separation into Red and Blue Teams.

- **Day 3:** Hands-on classroom training on Network Exploitation, Network Defense techniques and practices, and Red and Blue Team strategy meetings.

- **Day 4:** A 12-hour exercise where participants are either attacking (Red Team) or defending (Blue Team). The Blue Team is tasked with providing the cyber defense for a corporate environment and with maintaining operations to a batch-mixing plant and an electrical distribution Supervisory Control and Data Acquisition (SCADA) system.

- **Day 5:** Red Team/Blue Team lessons learned and roundtable discussion.

The following advanced training events have already been scheduled for FY 2012:

- **October 10-14:** International (Reserved)
- **November 7-11:** Reserved
- **December 5-9:** Industry Partners
- **January 16-20:** Industry Partners

There is no cost to attend the training; however, travel expenses and accommodations are the responsibility of each participant.

Additional offerings are being planned and will be announced once dates are finalized. As scheduled advanced training gets closer, an invitation along with a link to register for the course will be sent out and posted to the following website - http://www.us-cert.gov/control_systems/cscalendar.html. Please monitor the site periodically, as this schedule is updated as new courses are confirmed.

Register by clicking on the link provided on our webpage - http://www.us-cert.gov/control_systems/cscalendar.html. Registration is open approximately 2 months before the start of a class. Due to high demand, class size is limited to approximately 35 people with a maximum of 2 individuals per company per event. Classes fill quickly, so early registration is encouraged. Notification of cancellation is appreciated, with as much advance notice as possible so that others who wish to take the course can do so.

Industrial Control Systems Contributed Content

ICSJWG is now accepting contributions from the community pertaining to control systems security for the December Quarterly Newsletter. If you want to submit an article for the December Newsletter, please email icsjwg@dhs.gov, and we will take your submission into consideration for publication. The deadline for submissions for the December Newsletter is **November 25, 2011**.

Past ICSJWG newsletters are located on the CSSP website http://www.us-cert.gov/control_systems/icsjwg/index.html and in HSIN https://cs.hsin.gov/C10/C1/ICSJWG/Document%20Library/Forms/AllItems.aspx?RootFolder=%2fC10%2fC1%2fICSJWG%2fDocument%20Library%2fICSJWG%20Newsletters%2fICSJWG%20Quarterly%20Newsletter&View=%7b6F252F6A%2d18EB%2d447A%2d96D4%2d106024729AB9%7d.

Also, thank you to all members who contributed content for the September Quarterly Newsletter! The following content was submitted by members of the ICSJWG for publication and distribution to the ICSJWG community. Content and opinions are those of the authors and do not represent DHS opinions, endorsements, or recommendations. The advice and instructions provided in the contributed content should be confirmed and tested prior to implementation.

Software Assurance Events

From the DHS Software Assurance Program

SwA Working Group Sessions - Winter 2011

Nov 28-Dec 2, 2011 at MITRE-1, 7525 Colshire Drive, McLean, VA 22102-7539

The Software Assurance (SwA) Program of the Department of Homeland Security's National Cyber Security Division co-sponsors the Software Assurance Working Group's sessions to provide venues for public-private collaboration in advancing software assurance initiatives. Status updates from the SwA Working Groups are presented during SwA Forums and to other relevant stakeholder groups.
https://buildsecurityin.us-cert.gov/bsi/1292-BSI.html?branch=1&language=1

SwA Forum - Spring 2012

March 26-30, 2012 at MITRE-1, 7525 Colshire Drive, McLean, VA 22102-7539

The Software Assurance Program of the Department of Homeland Security's National Cyber Security Division co-sponsors SwA Forums semi-annually with organizations in the Department of Defense and the National Institute for Standards and Technology. The purpose of the forums is to bring together members of government, industry, and academia with vested interests in software assurance to discuss and promote integrity, security, and reliability in software.
https://buildsecurityin.us-cert.gov/bsi/1293-BSI.html?branch=1&language=1

Nuclear RoadMap

From the Nuclear Sector-Specific Agency (SSA)

The Nuclear Information Technology Strategic Leadership (NITSL) Workshop is a premier industry workshop for nuclear power reactor cyber security professionals, attended by both public and private sector partners, including nuclear utility owner-operators, cyber system vendors, cyber security experts, industry regulators, and other government stakeholders.

The Nuclear Sector-Specific Agency (SSA) has been invited for the past three years to provide an update on Nuclear Joint Cyber Subcouncil activities and initiatives. On June 13, 2011, Ms. Rachel Liang and Mr. David Martin from the Nuclear SSA were joined by Mr. Bill Gross, Cyber Program Manager with the Nuclear Energy Institute (NEI), and the Co-Chair of the Nuclear Joint Cyber Subcouncil who formally introduced the Roadmap to Enhance Cyber Systems Security in the Nuclear Sector.

The Roadmap to Enhance Control Systems Security in the Nuclear Sector describes coordinated activities to improve cyber systems security in the Nuclear Sector. It provides nuclear control and cyber systems vendors, asset owners and operators, and relevant government agencies with a common vision, goals, and objectives for cyber systems security in the sector. It also provides milestones to focus specific efforts and activities for achieving the vision, goals, and objectives over the next 10 to 15 years, addressing the Nuclear Sector's most urgent challenges, as well as its longer-term needs to reduce the cyber security risk to nuclear industrial cyber systems.

NITSL was a unique opportunity to not only discuss the Roadmap, but to interface directly with the wider Nuclear Sector community.

The New Paradigm for Utility Information Security: Assume Your Security System Has Already Been Breached

(Note: This was also featured on Asian Power Magazine and is submitted to ICSJWG with permission.)

By Ernie Hayden CISSP CEH, Managing Principal – Energy Security, Verizon Global Energy & Utility Practice

In 1990, Steven Covey published the very popular book *The Seven Habits of Highly Effective People*. With the publication of that book I am certain, Dr. Covey introduced the term "paradigm" into the popular business lexicon and it has been persistently used ever since.

As a reminder, a paradigm is defined as a worldview underlying the theories and methodology of a particular scientific subject. However, what does that have to do with information security practices of enterprises and energy companies? Below you'll see that the paradigm for system security is getting flipped on its head.

The Old Paradigm: The Fortress

Basically, there has been a standard practice if you will for many years where the "fortress" approach was the norm—or paradigm—for enterprise and energy company security. This applied to physical security and cyber security. The fortress concept included a strict perimeter—usually defined by gates, guards, and firewalls.

In this approach, the assumption was that all the attackers were on the outside of the perimeter and that the strong perimeter would prevent the attacker from not only entering the walls but they could not access the crown jewels (aka data) because it was housed within layers of more security barriers that included more walls, more guards, and more firewalls and maybe a moat.

The insider threat—that is the threat of an attacker from the inside of the perimeter—was viewed as very unlikely and not a "real" threat.

Using this approach, when the attacks became stronger and bolder, the captains of the fortress added more walls, more guards, and more firewalls with some extra intrusion detection systems (perhaps vats of boiling oil?) and security management tools.

Frankly, this was how I was trained as a security professional. But there are new ideas surfacing that turn this model upside down.

The New Paradigm: Assume Security System Breach

If you have been following the news these past few months, there have been some large cyber security hacks resulting in huge breaches of data and personal information. For example, one company's gaming system was shut down for a considerable time resulting in lost revenue and investigation/mitigation costs. While writing this article, another separate company announced that their gaming system had been hacked causing them to shut down their system for a time and reset all the user's passwords. And the recent notice of breach at a major defense contractor certainly raised many eyebrows. Don't forget there have been many other hacks in the recent past including a major credit card processor resulting in thousands of credit card numbers being stolen.

All these companies have employed the Fortress Paradigm. They have employed walls, fences, gates, guards, firewalls, intrusion detection systems, two-factor authentication systems, etc. and their systems still were breached by presumably outside miscreants or nation-states.

This changes the security defense posture considerably! The security paradigm needs to be adjusted.

As I attend different security conferences and read new thought leadership on the subject of security of companies, I'm noticing a new theme surfacing. That theme is you should assume your security systems are breached. You should assume that you can, ***and will be breached***.

Who is saying this?

For instance, I heard Mr. Kris Herrin, Chief Technology Officer of Heartland Payment Systems, make a speech at The Source Security Conference in June in Seattle. Kris said that the new approach by Heartland is to take all possible and practical steps to protect the data but they will assume the security systems and data can and have been breached.

In December 2010, Deborah Plunkett, the head of the U.S. National Security Agency's (NSA) Information Assurance Directorate, announced that computer systems must be built with the assumption that the adversaries will get in. She even stated that the most sophisticated attackers are going to go unnoticed on the NSA's networks. With these new paradigms, the focus will be on assuming that all components of the system are not safe and to make sure their practices, policies, procedures, and mitigation schemes are adjusted accordingly.[1]

This same theme of assumption of breach was also echoed in a PriceWaterhouseCoopers white paper called, "Are You Compromised But Don't Know It? A New Philosophy for Cybersecurity." Here, they go on to reinforce the new paradigm—assume you have been or will be breached and protect your systems and data accordingly. They advocate that this approach is more realistic and can allow you to be more flexible in protection of your high-value assets.

Lastly, my friend and mentor, Mr. Kirk Bailey, Chief Information Security Officer of the University of Washington in Seattle, has always been an advocate for assumption of breach. He has maintained this philosophy for as long as I have known him, and he steadfastly keeps his cyber mind aware that University systems and data stores could be hacked at anytime by sophisticated attackers.

Even in the Verizon Data Breach Investigations Report for 2011, there is a demonstrated increase in data breaches caused by external agents. In other words, these external entities need to somehow breach the security systems to gain access to the information. And statistically, the report goes on to show that the data breaches occurred by:

- *50% utilized some form of hacking,*
- *49% incorporated some sort of malware,*
- *and*
- *11% employed social engineering tactics.*

[1] Reuters Canada, December 16, 2010 http://ca.reuters.com/article/technologyNews/idCATRE6BF6BZ20101216

What Do You Do?

With this new paradigm, the fortress will still remain but you need to realize that the standard "signature-based" defensive measures do not necessarily work to identify and stop the more sophisticated attacks. You also have to realize that even the smallest hole in your perimeter could be compromised. Don't forget, that is all the attacker needs.

For example, there is a barrage of targeted attacks being made on various energy companies. However, these attacks are not like a cloud of arrows attacking your facility but instead a "rifle shot" aimed at an executive in your organization. This shot is usually a targeted phishing attack where a single email is sent to the executive that looks innocent enough but has a single URL or attachment that when opened will take advantage of vulnerabilities in computer programs such as your PDF reader. This then allows the hacker to install "back doors" in the corporate computer systems for future malware injects, network reconnaissance, and data retrieval. Hence, the assumed security of systems you rely upon is not necessarily effective enough.

A signature-based system would not stop this attack but education of the executive might.

Please recognize that this challenge to the assumption of a secure perimeter is not just the failure of employees to not open phishing emails. Often the factors include highly complex software, new attack methodologies, and the ever-crumbling perimeter caused in part by constant detection of vulnerabilities by security researchers and organized criminals.

So, what to do? Mr. Kirk Bailey is offering 10 Key Practices he is implementing under the philosophy of assumed breach. They are listed below—but please realize that these are not easily implemented, they are fraught with pushback from traditional security professionals, and each one could be described more thoroughly than this article can accommodate. That said, they include the following:

1. Implement a Risk Management Framework for reporting. Have a structure that is repeatable and readily demonstrates trends.

2. Conduct asset profiling and inventory. Know where your "crown jewels" of data are and separate out the data that can be lost with minimal impact.

3. Prioritize assets and related risk-mitigation efforts. Focus on protecting the "crown jewels." Continue to use and implement traditional layers of defense but only to recognize that they will not be 100% effective.

4. Clearly define roles and communication plans. Know who your trusted contacts are for incident response.

5. Implement aggressive risk transfer programs through detailed contracts and insurance underwriting.

6. Establish and sustain active and strategic alliances to allow for effective and trusted cross-communication about threats, mitigation schemes, and lessons learned.

7. Implement a business intelligence (aka “warfare intelligence”) program that includes effective situational awareness features.

8. Establish “advanced” incident response and management capabilities. Think outside of the normal cyber incident response practices to include incorporation of trusted contacts, stealthy communications, and attacker evidence.

9. Develop an active response capability.

10. Practice strategic isolation for your key executives, scientists, and knowledge workers. Limit presence on social networks that can be used by attackers for targeted hacks.

Kirk has reiterated that the above are not a checklist and are not adequately described in one or two sentences; however, he sees that the new paradigm including assumption of breach will require new thinking and for security programs to be built upon a “flexible fabric.”

In any risk discussion, the notion that there will always be some percentage of risk that cannot be eliminated is always present. You have to assume that this risk will always be there, and often it's due to things way beyond your control. Examples include human misbehavior, fundamental flaws with networking protocols, ditto for software, hidden back-doors, and design flaws in third-party hardware and software that you've bought and installed. That said, even with a “zero risk” mentality, you still need to realize that no security system is 100% effective.

Conclusion

This is a new shift in security thinking and many of my peers are still in disbelief. However, this new approach may allow you to be more effective in implementing layered security systems, protecting the high-value data, and being flexible enough to think like a cyber criminal and stop the attacks or at least mitigate their damage early in the theft. Basically, be ever vigilant. Constantly monitor and inspect your security systems, inspect your “crown jewels” and look for suspicious activity or minute changes that cannot be explained, and look at your logs and egress filters for stealthy communications to and from these systems.

About the Author

Ernie is currently a Managing Principal for Verizon Business with extensive experience in the power utility industry, critical infrastructure protection/information security, and cybercrime and cyberwarfare. His primary focus is on supporting customer projects regarding smart grid security, energy supply security, and electric grid security with special emphasis on NERC Critical Infrastructure Protection (CIP) standards. He travels extensively and speaks at many security and energy conferences and has been a previous contributor to the ICSJWG Quarterly Newsletter. You can contact him at ernie.hayden@verizon.com.

Industrial Automation Control Systems Standard Project Moves Forward

By Mike Ahmadi

The results were tallied in mid-July, 2011, and on July 29th, 2011, the International Electrotechnical Consortium (IEC) informed the standard making world that the proposal referenced as 65/482/NP has been approved and has now been re-designated as 65/487/RVN.

This may certainly seem confusing to those not intimately familiar with the workings of IEC (I know it confused me for a while), but suffice it to say that this is a huge step towards development and eventual ratification of one of the most promising standards to have ever surfaced in the Industrial Automation Control Systems (IACS) world. I am talking about IEC 62443-2-4 "Security for industrial process measurement and control - Network and system security Part 2-4: Certification of IACS supplier security policies and practices."

So why am I making such a lofty claim? What makes this particular standard so promising?

Well, to put it succinctly, it is end-user driven. Yes, that is what makes all the difference in the world, if you ask me.

Let me explain...

Let's start with a bit of history. The International Instrument Users' Association, more commonly known as the WIB (www.wib.nl), was founded in 1962 by a group of companies whose goal was to share information about process instrumentation. The members of this newly formed organization quickly realized that it was indeed beneficial to collaborate and share information and soon generated interest from companies throughout the world. In 1968, the WIB was officially registered as a non-profit organization known as the "Working Party on Instrument Behaviour".

Since its inception, the WIB has served as an invaluable resource to the member organizations that have participated in it. Despite the fact that many WIB members are direct competitors, the collaboration has proven extremely beneficial beyond any competitive concerns. The sharing of process control information among experts meant that organizations could leverage each other's knowledge at a fraction of the cost of managing research and development projects entirely in house. As these organizations continued to grow, the WIB continued to serve as a fantastic resource to help address the ever-changing process control landscape, and the collaborative process led to consistency, which ultimately translates to agility as well as enormous cost savings for both the end users and their vendors.

That brings us to cybersecurity…

As we are (hopefully) all aware by now, as technology has grown, so have cybersecurity challenges. In fact, that is probably why you are reading this article right now. As it turns out, organizations that have had to deal directly with cybersecurity issues in the process control space quickly came to realize that the cost of failure to protect against security threats is, as one person in the oil and gas industry once explained to me "As serious a threat as a fire." Knowing full well what a fire means to someone in the gas and oil industry, it was patently clear that this is a serious concern.

What several WIB members soon discovered was that they were all facing similar challenges in addressing cybersecurity issues, and it was indeed prudent to take on a collaborative project to pro-actively edge against these threats. The WIB leadership enlisted the assistance of cybersecurity experts to develop and build what became known as the WIB PROCESS CONTROL DOMAIN-SECURITY REQUIREMENTS FOR VENDORS. Now at version 2.0, this document consists of a set of requirements that is meant to serve as a normative set of baseline requirements that vendors can certify their IACS products to, and thereby conclusively demonstrate what security capabilities they can provide to end users.

But why vendors?

Simply put, an organization can (ostensibly) exercise some level of consistent control over their environment and thereby achieve some level of understanding of their inherent risk. Unfortunately, vendors represent a "wildcard" in the equation. Some vendors do a better job than others in addressing cybersecurity, and nearly all vendors will have a tendency to, at times, overstate their security posture. Moreover, try as they may, vendors rarely understand the challenges and concerns of end users as well as the end user does. Because of this situation, both end users and vendors must engage in costly Factory Acceptance Testing (FAT) and Site Acceptance Testing (SAT) activities for each and every engagement. What this ultimately leads to is the need for both end users and vendors to recoup these expenditures and that translates to higher costs for those whom the principal serves.

The WIB security requirements were created without vendor input. In fact, vendors were explicitly not permitted to participate in the process. The intention of excluding vendors was to ensure that the focus remained on end user needs, without the potential of vendor objections. While many vendors recognize the need for security in IACS products today, having vendors participate in setting requirements for their own product lines introduces the potential for conflicts of interest, since upper level management in such companies (especially publicly traded companies) are driven by mandates to maximize shareholder interests. Since the addition of pro-active security is so difficult to justify from an ROI perspective, vendors who actively push for the additional costs associated with building in security without an immediately recognizable financial justification face the potential of board-level scrutiny and investor backlash. However, if the vendors are removed from the decision-making process, this particular difficulty is essentially eliminated. Now I must admit, this was not explained to me as the reason why this was done, but it is a consequence I derived as I thought about this, and if it is indeed one of the reasons for the WIB choosing to exclude vendors, it is quite brilliant.

Almost immediately upon finalization of the requirements, a certification program was developed to test for conformity to the requirements, and the end users began mandating that their vendors certify to the requirements. This was initially met with quite a bit of objection by some vendors (which is certainly understandable), but customer mandates generally prevail, and the certification process was underway.

Despite the healthy adoption of the WIB security requirements by several large industry end users and certification by associated vendors, the fact remained that this was an industry group-based standard and not an international standards development organization (SDO)-based standard. While such standards can (and often do) become de-facto standards for many industries, the additional assurance and clout a true SDO standard carries is undeniable. Standards Development Organizations exist with the sole purpose of building and maintaining standards on a global scale. They have well defined (and in many cases onerous) processes in place to assure that the way standards are developed and maintained are consistent.

These SDOs also have a global reach that generally surpasses the reach of industry groups, which allows them to procure the needed resources to build standards. Additionally, in many countries, SDO-based standards have the force of law behind them, which (at times) makes it easier to ensure compliance.

The import of this was not lost on the proponents of the WIB, and in early 2011, The Netherlands introduced a new proposal for a work item to the IEC (one of the largest and most respected SDOs globally) as part of the vaunted IEC 62443 series, with the designation of IEC 62443-2-4. The IEC 62443 series is comprised of multiple standards and began life as the ISA99 series. The other parts of the series have been under development by IEC TC65 (Technical Committee 65) working group 10 for some time now, and the WIB authors felt that, due to the tight relationship between the WIB requirements and many of the various parts of IEC 62443 (which is a set of security standards for Industrial Control Systems), it was a natural fit. Additionally, many of the same companies who created the WIB security requirement and vendors who were certifying to the requirements were deeply involved in the IEC 62443 series.

Today, I am a US expert for TC65, and we now have an impressive international contingency of technical experts who are in the process of resolving comments on the working draft of IEC 62443-2-4 (which is essentially WIB 2.0). The project leadership has set a rather aggressive timeline for ratification of the standard, but it shows great promise due to the fact that adoption of the standard is virtually guaranteed due to the interest shown by some very large end user organizations. Additionally, the NIST CSWG (Cyber Security Working Group) formed an IEC 62443-2-4 task force whose mission is to harmonize the NISTIR 7628 requirements with the proposed standard, and it is currently one of the most active groups in the NIST CSWG. The UCAIug OpenSG Security Conformity Task Force is currently completing a draft charter, which seeks to define how certification labs for the IEC 62443-2-4 should operate (not by developing criteria, but assuring that certification labs operate under a consistent set of SDO and industry defined guidelines). Last, but certainly not least, there are many members of the ICSJWG who are actively participating in the standard development process (many who cross over into the aforementioned groups).

Interest is indeed quite high, and the IEC 62443-2-4 project has led to an increased interest in the other parts of the IEC 62443 standards project. This is extremely promising since the other parts of the series dive quite deeply into specific areas of the control system, and also define requirements that are intended for end users as well as the vendors. In other words, the entire series “completes the picture”.

For anyone who is interested, the IEC 62443 project is always in need of volunteers, and the extraordinary interest that the IEC 62443-2-4 process has now brought to the series means that we are going to need more bodies to help move the process forward. Cybersecurity standards are no longer an option in the control systems world, and standardizing requirements is a good way to make sure the discussions remain productive and that the requirements are actionable.

Planting a stake in the ground is never an easy task, but once it is there, it perpetuates motion, and I will take that over stagnation any time.

Consistency Under Deception Implies Integrity

From Fred Cohen & Associates - Analyst Report and Newsletter

Consistency analysis has been found useful in detecting corruptions of all sorts, ranging from accidental bit flips (i.e., parity checking) in the 1960s, to multiple bit error detection (i.e., cyclical redundancy checks), to malicious alteration detection (i.e., cryptographic checksums)[1] in the 1980s, and has been used in digital forensics since at least the 1990s[2]. All of these are in the digital space. But as the digital and analog spaces increasingly converge in industrial control systems, complex corruptions of the combined digital and analog spaces are being used to induce harmful physical effects through exploitation of the combined systems. As an approach to defeating limited attempts to alter control system, effector, and/or sensor signals, the notion of consistency checking again comes to the fore.

Computational leverage

Notionally, the use of cryptographic checksums for detection of intentional alteration gains its utility from the computational leverage of detection over forgery. The forger not holding the cryptographic key to generate a true cryptographic checksum for any desired bit sequence is unable to systematically forge sequences in which the cryptographic checksum is consistent with the content it covers. While replay attacks and similar methods may function in systems not well designed to defeat them, such systems can be and have been designed and are successful in mitigating attacks, up to the point where the attacker is able to determine the cryptographic key, at which point forgery becomes feasible and inexpensive. The selection of the cryptographic checksum method is intended to be such that this takes a long time, and thus without infeasible computational capacity, the attacker cannot systematically forge for a long enough time to make the system secure for the intended use. Or at least this is how it can be designed.

Notionally, consistency checking in digital forensics also gains its utility from computational advantage. In this case, normal operation of computer systems produces redundant traces, and these traces can be compared for consistency. While trivial forgeries can function against detectives who are unfamiliar with the consistency methods available today, the complexity of creating a forgery that cannot be detected in the larger overall system is thought to be so high that it is infeasible in almost all cases.[3] Again, the forger cannot anticipate and alter enough traces to defeat all feasible consistency checks, and the alteration of these traces introduces potential inconsistencies with still other traces that are also subject to detection.

Thus, consistency checking leverages computational advantages of defenders over attackers. In doing so, it refutes the common but false assumption that the defender has to protect against all possible attacks in detail while the attacker only has to find one attack the defender failed to defend against. Note also that caught attackers don't always get to keep trying.

[1] F. Cohen, "A Cryptographic Checksum for Integrity Protection", IFIP-TC11 "Computers and Security", V6#6 (Dec. 1987), pp 505-810.
[2] F. Cohen, “A Note on Detecting Tampering with Audit Trails”', 1995
[3] F. Cohen, “Digital Forensic Evidence Examination”, ASP Press, 2009-2011

Detection Under Normal and Altered Operating Conditions

Under normal operating conditions, a failure in a sensor, effector, or other system component will be reflected in altered signals indicative of a move away from the weighted "center" of the control envelope (i.e., the set point). The response from the automated control system will be to compensate by altering effectors so as to move the system back toward the center of the control envelope. Of course there may be some losses (e.g., leaked fluids) resulting from the actual fault (e.g., a hole in a pipe), and the control system will continue to compensate as well as it can, and potentially alert operators to the fault condition (e.g., fluids are running low).

If an attacker alters a sensor to produce false information, the response from the automated control system will be to compensate by altering the effectors in response under the assumption that the information is true. Thus, a reflexive control attack is realized as the reflexes of the control system react to the information available. To some tolerance, the changes will be within the control envelope of the system and stability will be retained, even if some other performance effects may occur (e.g., undetected theft of fluids or excess line voltage). As tolerances are exceeded, the control envelope may be exceeded and the system may become unstable, may collapse, and may, as a side effect, destroy physical components.

However, if there are multiple sensors, to the extent that unaltered sensors are effected by control changes based on false information, the unaltered sensors may generate signals inconsistent with the altered sensor signals, supporting inconsistency detection. If redundant sensors are separate and different, common mode failures may also be avoided and better diagnosis supported. The same is true of an altered effector or an altered physical system, and to the extent that there are redundant control systems, to an altered control system. Thus, the potential exists to use the control system to detect inconsistencies between sensors, perhaps diagnose the most likely bad sensor(s) and overall situation, and in an advanced system, perhaps compensate for and reduce the trust and dependence on the false signals.

At some level of induction and/or suppression of signals (i.e., alteration), so many signals may be altered that an entirely false picture that is itself consistent may result. If every system and component is taken over by an attacker, the system may not detect or report anything. But even short of this, the deception may be of sufficient quality as to mimic the legitimate control system and deceive the operator and the systems they depend upon.

Induction and suppression of signals (i.e., deception) for detection

To compensate for this class of attack, an alternative approach is to intentionally induce and suppress the normal control signals that would be used in a systematic way so as to produce systematic changes in the overall system that (1) remain within the safety margins of the control envelope and (2) produce time variant effects within and throughout the system under control. By doing so, the sophisticated control system may induce changes that ripple through the system as a form of diagnostic test, creating sequences of alterations that remain safe and relatively efficient while inducing feedback that reveals attempts to circumvent normal controls. By doing so in an externally unpredictable sequence, the malicious actor wishing to alter the control system may be detected if they are unable to predict the proper control signals in time to reflect a globally consistent system state and variance, even though they have control of most of the sensors and effectors. Thus, computational advantage is used by the control system designer to detect and potentially diagnose malicious alteration.

Limitations of these methods

There are two cases to consider; a quasi-static case in which the system is assumed to be over damped and fed by more or less constant volumes relative to measurement time frames, and a dynamic case in which the system is under-damped and constantly changing, so that waves normally build upon each other.

The quasi-static case:

In the quasi-static case, subject to small-time delays associated with the propagation of change, conservation of mass dictates that the volume in a repository is equal to the sum of all the flows into that repository. Because the system is quasi-static, the conservation rule can be applied to within the measurement precision of sensors, and variations detected as inconsistent if they exceed the variances in precision, again subject to the relatively small propagation delay. The detection time for a leak is dictated by the precision of measurements, so that a device that measures a water tank to the nearest 1000 liters will detect leaks totaling no more than 2000 liters essentially as it happens.

For a system of tanks and pipes, the height of each stage in the system should match the flows in and out, and an attempt to provide false signals for any single sensor will result in a detection as soon as the leaks create a mismatch. By altering multiple sensors, it is feasible to shift apparent usage from one area to another only to the extent that it doesn't create a downstream inconsistency. Thus, at downstream endpoints, a malicious actor could steal water and place the blame on another party, but theft in the middle of a system would be detected by the downstream loss, and in order to avoid such detection forged signals from downstream sensors would be required in the proper combinations so as to create a consistent resulting overall system.

Time to detection depends on the volumes flowing and sensitivity, so that for 2000 liters of loss in a system flowing 200 liters per second, 10 seconds plus propagation time plus sensor scan time limits the time till detection. However, noise factors such as rain and wind may effect volumes and sensor precision. For example, if it rains, the system will gain water that could be stolen as it is gained.

For passive detection, consistency analysis allows forgeries by parties who understand and can model the system reasonably well, perhaps including taking sensor readings from other (downstream) sensors and forging multiple values so as to retain consistency of the overall system in near-real-time. By adding active alterations to the system through actuator changes, the system can further resist modeling by malicious actors. For example, instead of having a fully predictable control system that seeks to keep water in tanks at constant set points, suppose the system intentionally changed set points over time. Now the attacker seeking to model the system has to take into account the changes in set points in order to create a consistent forgery. Instead of simply forging sensor readings to cover up a loss for a time, the attacker has to calculate the proper values for the entire downstream system taking the changes in all set points into account, or the forged values will be internally consistent, but inconsistent with the changes set points. The control system’s model of the overall system is unchanged, and detection remains the same problem, while forgery becomes far harder and requires access to and analysis of all of the changes set points for effect.

The dynamic case:

Unlike the quasi-static case, an underdamped system with changes that don't have time to propagate to stability produces a far more complex challenge for both attacker and defender. Conservation remains true, of course, but measurement becomes far harder. A sensor measuring a wave form may vary significantly from the average value, and integration does not produce a reliable mean value in an underdamped system over a short time frame. Creating additional dynamics is far more problematic in that, without complete knowledge of the system state over time, induction of changes may force the system out of stability and generate positive feedback, ultimately resulting in catastrophic failure. Sensors may be at nodes in the system providing incorrect feedback, and perhaps more importantly, the utility gained in detection may be nullified by the large variance in sensor values in different operating modes.

On the other hand, to the extent that these problems can be solved, the situation for the attacker is potentially far more complex. Now they have to launch their attack in an environment where their changes may produce system instabilities that make their changes obvious quickly, they have to compute values upstream and downstream, and the computation of effects of alterations may take too long for real-time analysis. The defender has the advantage of proposing a change, calculating expected values with time through a complex analytical process, and then using actuators in combinations and sequences so as to produce predictable dynamics. The attacker without advanced knowledge of the planned changes is forced to try to do real-time analysis of the effects of the changes and produce solutions in time to forge sensor data to within the fidelity required to fool the defender's predictive system.

Limits based on relative volumes:

A further limit is worth pointing out. Intentional variations in flows for detection may be problematic in systems such as power infrastructure, where back forces and phase shifts may result from attempts to alter flows. While at distribution points, such changes are relatively straight forward with smart meter technologies and force levels are relatively low, if many such changes are made in concert, the combined forces may be enormous. Thus, scheduling of changes may become a serious challenge for the overall system. In transmission, shifting enough power to produce changes in excess of detection thresholds may require so much energy that the system becomes less stable, back forces on generation may be problematic and damaging, interactions of wave forms may cause nodes in the system in excess of allowable tolerances, and compensation for rapid changes in load that occur all the time may force the detection thresholds to be set so high as to make large-scale dynamic detection infeasible without risking system stability. Computation of changes and detectable effects may be so complex as to make dynamics infeasible.

Conclusions:

Consistency analysis appears to be a feasible method for detecting intentional acts altering control systems, and intentional deception in the form of induction and suppression of signals can be used to gain computational leverage over attackers. In the quasi-static case this is now feasible, while in the dynamic case it is far more complex and potentially dangerous, but also potentially far more advantageous to the defender.

CSSP Contact Information

If you would like to contact the ICSJWG to ask a question or inquire about participation, please send an e-mail to icsjwg@dhs.gov.

The CSSP and Industrial Control Systems Cyber Emergency Response Team (ICS-CERT) encourage you to report suspicious cyber activity, incidents, and vulnerabilities affecting critical infrastructure control systems. Online reporting forms are available at https://forms.us-cert.gov/report/.

In addition, the ICS-CERT Monthly Monitors are published on HSIN as appendices to the ICSJWG newsletter and can be found here http://www.us-cert.gov/control_systems/ics-cert/.

Other important contact information:
Website Address: http://www.us-cert.gov/control_systems/
ICS-CERT Email: ics-cert@dhs.gov
Phone: 1-877-776-7585
CSSP Email: cssp@dhs.gov