

Chapter 5

APPLYING TRUSTED NETWORK TECHNOLOGY TO PROCESS CONTROL SYSTEMS

Hamed Okhravi and David Nicol

Abstract Interconnections between process control networks and enterprise networks expose instrumentation and control systems and the critical infrastructure components they operate to a variety of cyber attacks. Several architectural standards and security best practices have been proposed for industrial control systems. However, they are based on older architectures and do not leverage the latest hardware and software technologies. This paper describes new technologies that can be applied to the design of next generation security architectures for industrial control systems. The technologies are discussed along with their security benefits and design trade-offs.

Keywords: Process control systems, trusted networks, security architectures

1. Introduction

The increased interconnectivity of industrial control networks and enterprise networks has resulted in the proliferation of standard communication protocols in industrial control systems. Legacy SCADA protocols are often encapsulated in TCP/IP packets for reasons of efficiency and cost, which blurs the network layer distinction between control traffic and enterprise traffic. The interconnection of industrial control networks and enterprise networks using commodity protocols exposes instrumentation and control systems and the critical infrastructure components they operate to a variety of cyber attacks.

Security surveys reveal significant increases in external attacks that target critical infrastructure assets [2]. The entry points in most of the incidents were corporate WANs, business networks, wireless access points, modems and the Internet.

Several government agencies and industry associations have proposed standards and security best practices for industrial control systems [11, 12, 17, 19].

However, these efforts are based on older technologies and security architectures that rely on the differentiation and separation of enterprise and control traffic. While the efforts are, no doubt, important, the underlying security philosophy exposes industrial control systems to attacks that exploit misconfigurations, out-of-band connectivity and blind trust in the identities of traffic sources.

However, new technologies are emerging that provide more pervasive security within networks [10]. These technologies push security from perimeter devices such as firewalls to the networked devices themselves. This paper reviews technologies that can be applied to designing the next generation of secure industrial control systems. The technologies are discussed along with their security benefits and design trade-offs.

2. Control System Security Recommendations

Industrial control systems (ICSs) are highly distributed networks used for controlling operations in water distribution and treatment plants, electric power systems, oil and gas refineries, manufacturing facilities and chemical plants. ICSs include supervisory control and data acquisition (SCADA) systems and distributed control systems [19]. The main components of an ICS are the control server or master terminal unit (MTU), remote terminal units (RTUs), intelligent electronic devices (IEDs), programmable logic controllers (PLCs), operator consoles or human-machine interfaces (HMIs), and data historians. Generally, an ICS comprises two distinct networks: a process control network (PCN) containing controllers, switches, actuators and low-level control devices, and an enterprise network (EN) incorporating high-level supervisory nodes and corporate computers.

The National Institute of Standards and Technology (NIST), Institute of Electrical and Electronics Engineers (IEEE), Instrumentation Systems and Automation (ISA) Society, International Electrotechnical Commission (IEC) and Industrial Automation Open Networking Association (IAONA) have specified guidelines for securing ICSs (see, e.g., [11, 12, 19]). In fact, most security best practices recommend the segregation of PCNs and ENs.

Firewalls are often used to segregate PCNs and ENs [1, 18, 19]. A firewall can be configured to block unnecessary services, protocols and ports, thereby providing a higher degree of segregation between a PCN and EN. A router may be positioned in front of the firewall to perform simple packet filtering, leaving the firewall to perform more sophisticated tasks such as stateful filtering and acting as a proxy.

Using a single firewall between a PCN and EN has a serious drawback. This is because the firewall must allow the data historian to have a wide range of access to the PCN. Essentially, each service needs a “hole” in the firewall to operate correctly. Configuring too many holes in the firewall reduces PCN-EN segregation and opens the PCN to a slew of attacks. This problem is typically addressed by creating a “demilitarized zone” (DMZ) [1, 18, 19].

An architecture deploying a DMZ has three zones: an outside zone containing the EN, an inside zone containing the PCN, and a DMZ containing the data

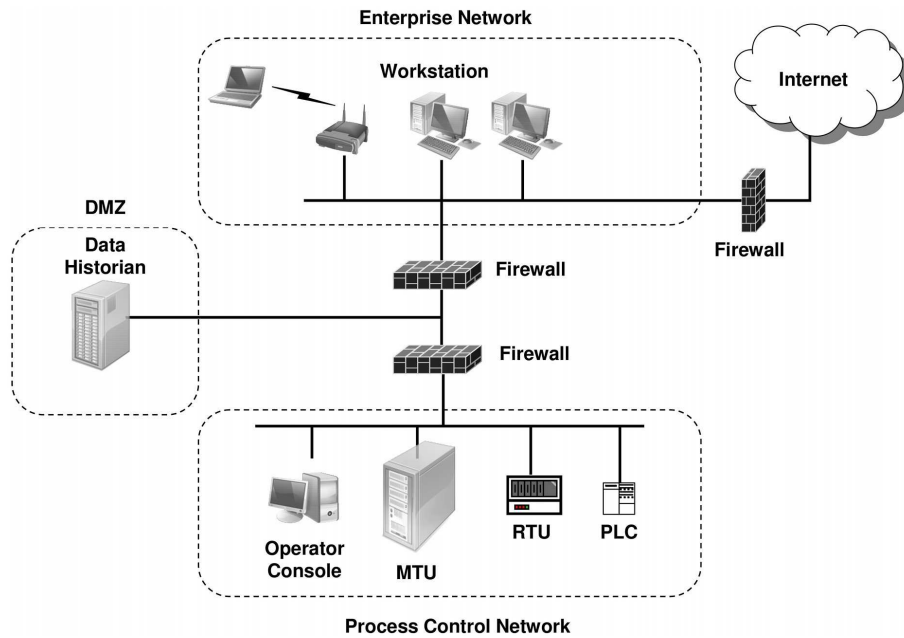


Figure 1. Paired firewall PCN architecture.

historian. Firewall rules are crafted to make the DMZ historian the sole point of contact between the PCN and EN. The historian can access PCN services that provide it data; in turn, the EN is allowed access to the historian. Firewall rules block PCN access by all other devices. Most attacks originating in (or passing through) the EN and targeting the historian will not affect the control systems; at worst, they would corrupt the historian's data (a redundant copy of this data is stored elsewhere).

A PCN architecture deploying paired firewalls separated by a DMZ [18, 19] is shown in Figure 1. It simplifies the firewall rules and achieves a clear separation of responsibility as the PCN-side firewall can be managed by the control group and the EN-side firewall by the IT group [18, 19]. This architecture is highly recommended for ICSs, and best practices have been identified for configuring the firewalls (see, e.g., [1, 11, 12, 19, 21]).

3. Security Challenges

Firewall configuration errors can lead to security vulnerabilities. One problem is that firewalls often have large rule sets that are difficult to verify. According to a study by Wool [21], firewall rule sets may have as many as 2,600 rules with 5,800 objects, and a significant correlation exists between rule set complexity and the number of configuration errors. A second problem is that firewalls are usually the main line of defense. Configuration errors enable at-

tackers to exploit holes in a firewall and target the otherwise defenseless devices inside the network.

Wool [21] notes that 80% of rule sets allow “any” service on inbound traffic and insecure access to firewalls. He emphasizes that “the analysis of real configuration data shows that corporate firewalls are often enforcing rule sets that violate well-established security guidelines.” The Wool study and others demonstrate that firewall configuration errors pose a real threat to ICS security.

Even properly configured firewalls can be bypassed [3]. This occurs, for example, when a vendor creates a direct (e.g., dial-up) connection to a device for maintenance, or when unsecured wireless access points exist behind a firewall. Firewalls can also be thwarted by tunneling attack traffic using legitimate means (e.g., via a corporate VPN) or by using encryption (firewalls do not inspect encrypted packets). A widely-reported firewall breach occurred in January 2003, when the MS SQL Server 2000 worm infected systems at the Davis-Besse nuclear power plant in Oak Harbor, Ohio [16].

Vulnerable devices are typically secured by patching their services, updating software or deploying the latest versions of the devices. However, manual patch/update/version management are difficult and costly tasks, especially when careless users introduce vulnerable (wireless) devices into an industrial control network that establish new entry points for attackers.

Unsecured physical access also exposes ICSs to serious security threats. Open wireless access points and Ethernet ports on office walls enable attackers to enter ICS networks and target critical assets. Nothing in the traditional ICS architecture prevents suspect devices from connecting to the network; thus, serious threats are posed by devices whose hardware, operating systems, executables and/or configurations have been tampered with by attackers.

Many ICS vulnerabilities admit malware such as worms, viruses, Trojan horses and rootkits [15, 16]. ICS security trends reveal that external malware attacks are becoming increasingly common [2]. Finally, rogue users (insiders) are an ever-present threat to ICSs.

4. Trusted Process Control Networks

In a traditional network access control model, access is granted to a user without considering the security state of the user’s machine. Likewise, firewall access control is agnostic about the security status of the device that sends traffic. A port on a machine is opened or not opened to traffic based entirely on the identity of the source.

A trusted network architecture uses information about the hardware and software states of devices in admission and access control decisions. When a device first “joins” the network, its hardware and software are checked; based on these checks, the appropriate access control rules are applied dynamically to the user, device and traffic. The same principle can be applied to process control architectures. This section discuss technologies that support this concept and their application to ICSs.

4.1 Trusted Networks

A trusted network (TN) architecture uses existing standards, protocols and hardware devices to implement “trust.” TNs provide important security services such as user authentication, comprehensive network device admission control, end-device status checks, policy-based access control, traffic filtering, automated remediation of non-compliant devices and auditing.

The Trusted Computing Group (TCG) has promulgated industry standards for TNs [20]. Several commercial TN technologies have been developed, including Cisco TrustSec [6], Cisco CleanAccess [7] (formerly known as Cisco Network Admission Control (NAC) [5, 8]), and Microsoft Network Access Protection (NAP) [13]. Cisco NAC is interoperable with Microsoft NAP; details about their interoperation can be found in [9].

4.1.1 Trusted Network Components. TN component vendors use a variety of names to describe their products. We use generic terms with a bias towards those adopted by Cisco CleanAccess.

A TN has the following components:

- **Client Device:** Every client device must be evaluated prior to admission to a TN.
- **Network Access Device:** All connectivity to a TN is implemented via a network access device (NAD), which enforces policy. NAD functionality may exist in devices such as switches, routers, VPN concentrators and wireless access points.
- **Authentication, Authorization and Access Control Server:** The authentication, authorization and access control (AAA) server maintains the policy and provides rules to NADs based on the results of authentication and posture validation.
- **Posture Validation Servers:** Posture validation servers (PVSs) evaluate the compliance of a client before it can join a TN. A PVS is typically a specialization for one client attribute (e.g., operating system version and patch or virus signature release).
- **Posture Remediation Servers:** These servers provide remediation options to a client device in case of non-compliance. For example, a server may maintain the latest virus signatures and require a non-compliant client device to load the signatures before joining a TN.
- **Directory Server:** This server authenticates client devices based on their identities or roles.
- **Other Servers:** These include trusted versions of Audit, DNS, DHCP and VPN servers [5, 7, 8].

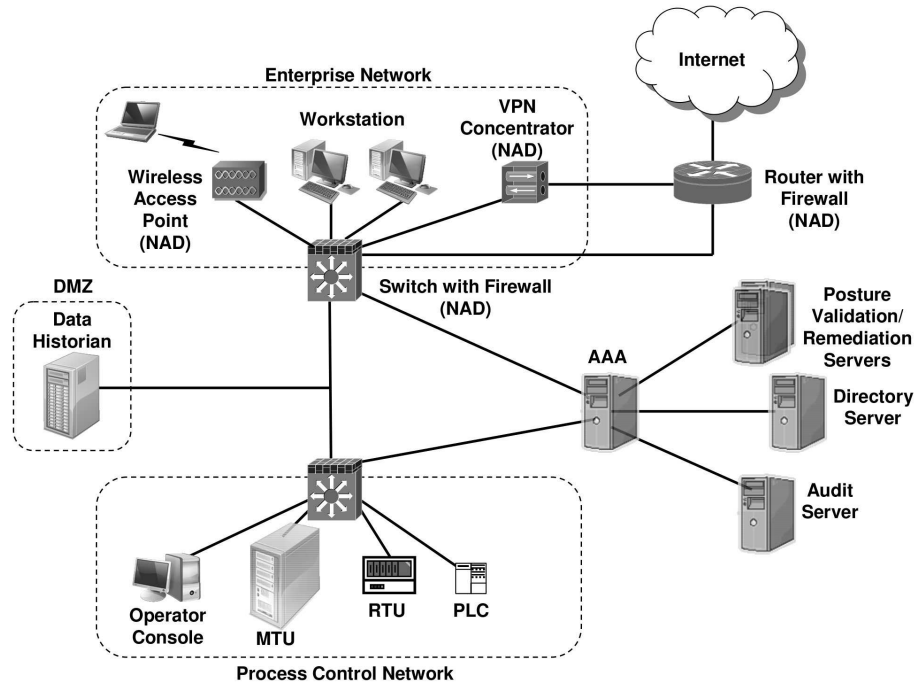


Figure 2. Trusted process control network.

4.1.2 Trusted Network Protocols. TNs leverage existing standards and protocols to implement the required security functionality; this reduces the cost of building TNs. Protocols used in TNs include IPsec for hardening communications [7], EAP and 802.1x for authentication [5, 6], RADIUS/LDAP/Kerberos for directory services and authentication [5, 7], HCAP for compliance communication [5], and GAME for communications between AAA and audit servers [4].

4.2 TPCN Architecture

A trusted process control network (TPCN) architecture is presented in Figure 2. A client device intending to join the network communicates its request to the NAD. The NAD establishes the client device's identity using EAP over the 802.1x protocol and sends the results to the AAA server using the RADIUS protocol. The AAA server returns a list of posture validation requirements and the addresses of the appropriate PVSs.

The client then validates its posture with each of the PVSs. If the client is in compliance, the results are sent to the AAA server using the HCAP protocol. On the other hand, if the client lacks one or more requirements, the appropriate posture remediation servers suggest remediation actions to the client.

The directory server determines the client's group or role. Given all the results from the PVSs and the directory server, the AAA server determines the set of rules that apply to the client's access and traffic and sends them to the NAD for enforcement. From this point on, the client is permitted to communicate via the NAD and all its activities are monitored for policy compliance. Interested readers are referred to [5, 7, 8] for additional details.

The policy held by the AAA server is in the form of an authentication requirement and a list of posture validation requirements. For example, token-based authentication may be required and postures must be validated with the anti-virus server, patch management server and driver validation server. When a client device joins the network, a NAD communicates with an AAA server on behalf of the device. The AAA server authenticates the device and provides rules based on the device's security postures to the NAD. From this point on, the NAD enforces the policy on all ingress and egress traffic to/from the device. For example, an RTU with valid firmware is allowed to communicate with the historian; all other traffic is blocked. The two examples below further clarify the workings of a TPCN.

Example 1. Consider a scenario where an analyst on a workstation intends to connect wirelessly to the PCN to access historical data about plant operations. The workstation connects to a wireless access point (AP) in the enterprise network with NAD functionality. The AP applies the default policy, which is to block all traffic except what is needed to establish trust. The workstation then authenticates with the AP using EAP over the 802.1x protocol to send a stored certificate. The AP uses RADIUS to send the workstation's identity to the AAA server. The AAA server then sends the user's identity to the directory server, which knows the user's role ("analyst"). The AAA server uses RADIUS to send the workstation a list of posture requirements (anti-virus version number and OS patch history). The workstation uses a trusted platform module (TPM) chip to sign and send the posture values to the relevant PVSs, which proceed to validate these values. The patch management PVS discovers that the workstation OS has a missing patch and coordinates with the remediation server to have the appropriate patch sent to the workstation. The PVSs transmit the results back to the AAA server using the HCAP protocol. If the workstation is compliant, the AAA sends a rule set to the AP for enforcement. Since the user role is "analyst," the rule set allows TCP connections to the historian but blocks access to all other devices.

Example 2. Consider a scenario where an RTU intends to join the PCN. The RTU connects to a switch on the factory floor via a network cable; the switch has NAD functionality. The protocols used are the same as in Example 1, so we avoid repetition. The switch authenticates the RTU using the RTU's stored token. The AAA server requires the RTU to validate its configuration with a configuration management server. The RTU sends its configuration to the configuration management server, which returns the successful result to the

AAA server. The AAA server, in turn, sends the appropriate rule set for the compliant RTU to the switch for enforcement. The RTU may now communicate with other RTUs, the MTU and the historian; the switch blocks all other traffic.

4.3 TPCN Requirements

For added security and separation of duty, a TPCN requires at least two NADs (switches with firewalls) and a AAA server (Figure 2). An enterprise can add as many PVSs as required, e.g., an anti-virus validation server to ensure that devices have up-to-date virus protection, a patch management server to check that devices have the correct patches and a software validation server to verify the authenticity of embedded device firmware. Incorporating multiple PVSs adds to the cost of a TPCN, but enhances security.

All NADs (switches, routers, wireless access points, etc.) must support trusted network functionality. Many vendors offer products with trusted network functionality. Therefore, if an enterprise is already using new equipment, implementing a TPCN may be very cost-effective. Older systems would likely involve significant upgrades, which can be costly. Note that in a TPCN architecture the firewall functionality is integrated in NADs.

Client devices may need software and firmware upgrades to support trusted network functionality. A trusted network client is required for authentication with the AAA server and for sending posture values. For secure applications, TPM chips can be used to verify configurations and obtain posture signatures. Devices such as RTUs and PLCs do not usually have TPMs; however, as some RTUs already come with built-in web servers, adding TPM to these devices is feasible, especially if government regulations mandate the implementation of trusted ICS architectures.

The administrator applies system updates by imposing new requirements in the AAA and PVSs. The AAA server informs devices of the new policy. If the devices have the update, they verify this fact with a PVS and remain in the network. Otherwise, the appropriate server provides them with the required patches (or installs the patches automatically), upon which they can enter the network.

TPCNs have the same availability issues as traditional PCNs – applying patches can cause components to crash. Therefore, every patch or update must be tested thoroughly before being placed on the AAA server. Exact replicas of TPCN components should be used for testing. If concerns exist after testing, a backup device may be placed in the TPCN. In such a situation, the AAA server holds two different policies for the device. One policy is associated with the actual role and the other policy with the backup role. The backup policy does not enforce the new requirement on the backup device until the actual device is verified to function correctly with the patch. It is only then that the administrator applies the requirement to the backup device as well. Note that if the actual device is affected by the patch, the backup device can function correctly since it is not required by its policy to have the patch in order to connect to the network. TPCNs do not positively or negatively affect

system availability; they merely enforce the requirements. It is the testing phase, before the specification of a requirement, that determines whether or not system availability is affected.

5. TPCN Evaluation

The benefits of a TPCN are best seen in the light of how it addresses the security issues that impact traditional networks. A TPCN addresses the following security issues either partially or completely.

- **Firewall Configuration Errors (Partial):** A TPCN breaks the set of firewall rules into smaller rule sets associated with each access control group or role. These rule sets are sent by the AAA server to the NADs for enforcement upon completion of the authentication phase. According to Wool [21], the number of configuration errors decreases logarithmically as the rule set complexity decreases. Because a TPCN has smaller rule sets, the potential for firewall configuration errors is correspondingly lower. Moreover, access rules in a TPCN are defined based on groups or roles, not just IP addresses; this helps reduce confusion and, consequently, configuration errors. Note that configuration errors will never be completely eliminated; therefore, TPCN only provides a partial solution to the problem.
- **Bypassing Firewalls (Complete):** TPCNs explicitly address this issue by securing all NADs and requiring them to establish trust relationships with client devices before forwarding traffic (including wireless traffic and VPN traffic). Furthermore, the access control and traffic rules are applied at every access point. It is not possible to bypass the rules by hooking a line behind a firewall; this is because the line's switch (access point) enforces the rules.
- **Vulnerable Devices (Partial):** In a traditional network architecture, patch/update/version/configuration management is performed manually by the network administrator. This is an extremely difficult task for remote and mobile devices. As a result, it may be done less frequently than recommended or it may be simply ignored. In a TPCN, the state of a device is checked automatically before it can join the network. Moreover, its behavior is continuously monitored upon entry and status checks can be performed at the desired frequency. Consequently, a TPCN is less vulnerable to known attacks. Note, however, that a TPCN is still vulnerable to zero-day attacks.
- **Unsecured Physical Access (Complete):** TPCNs again address this problem by enforcing security policies on NAD ports. This is sometimes referred to as "port-based access control." Thus, a malicious or careless user cannot hook a device to an open Ethernet port and gain entry into the network. Note also that ports on TPCN switches and wireless access

points do not forward traffic until trust relationships are established with the communicating entities.

- **Malware (Partial):** The compliance rules enforced on devices before and after joining a TPCN reduce the likelihood of infections by malware. A SCADA security study [2] notes that “the majority of worm events occurred months or years after the worm was widely known in IT world and patches were available.” This implies that the majority of incidents can be prevented by enforcing compliance rules before a node joins a network. Since nearly 78% of the (external) SCADA security incidents are caused by malware [2], TPCN incidents are reduced dramatically. Nevertheless, a TPCN remains vulnerable to zero-day attacks.
- **Untrusted Devices (Complete):** TPCNs address this problem explicitly by verifying the signatures of the critical components of a device using the TPM chip and also checking the device status. Note that if the TPM chip is trusted, the device can attest its identity.
- **Untrusted Users (Partial):** By using stronger authentication methods and clearly defining user roles, TPCNs prevent attacks such as password cracking/stealing, access violations and impersonation. Also, by blocking all unnecessary accesses, TPCNs partially prevent accidents caused by careless insiders that account for more than 30% of all security incidents [2].

We employed the Common Attack Pattern Enumeration and Classification (CAPEC) database [14] to further compare the TPCN architecture with traditional PCN designs. CAPEC contains twelve attack categories along with their descriptions, prerequisites, methods, consequences and mitigation strategies. We consider nine attack categories (with 31 attack patterns), which we believe are meaningful in the ICS context and showcase the differences between TPCNs and traditional PCNs. For example, while buffer overflow attacks are effective against software applications, they are not relevant when evaluating network designs.

Tables 1 and 2 present the results of the comparison. The descriptor H (high) means that an attack is performed with little effort and cost; M (medium) implies that an attack is still possible but requires expert knowledge and is costly; L (low) indicates that an attack is highly unlikely or involves enormous effort, time and/or cost. The last column in Tables 1 and 2 shows the security controls provided by a TPCN to address the attack (if any).

Considering the 31 total attack patterns, a PCN is vulnerable to nineteen (61.3%) high, nine (29%) medium, and three (9.7%) low feasibility attacks. On the other hand, a TPCN is vulnerable to only two (6.5%) high feasibility attacks along with nine (29%) medium and twenty (64.5%) low feasibility attacks. Note that this is a qualitative comparison of the two architectures; the quantitative assessment of network architectures based on security metrics is an open research problem and is beyond the scope of this paper.

Table 1. Feasibility of attack patterns.

Category	Attack Pattern	PCN	TPCN	TPCN SC
Abuse of Functionality	Inducing Account Lock-out	H	L	Strong Authentication
	Exploiting Password Recovery	H	L	Strong Authentication
	Trying Common Application Switches and Options	H	L	Configuration Verification
	Exploiting Incorrectly Configured SSL Security Levels	H	L	Configuration Verification
Spoofing	Faking the Data Source	M	L	Message Authentication
	Spoofing the Principal	H	L	Strong Authentication
	Man-in-the-Middle Attack	H	L	Device Authentication
	Creating a Malicious Client	M	L	Accounting
	External Entity Attack	H	L	VPN Access Control
Probabilistic Techniques	Brute Forcing Passwords	L	L	Strong Authentication
	Brute Forcing Encryption	L	L	N/A
	Rainbow Table Password Cracking	L	L	Strong Authentication
	Manipulating Opaque Client-Based Data Tokens	M	M	N/A
Exploiting Authentication	Exploiting Session Variables, Resource IDs and Other Credentials	M	M	Software Verification
	Reflection Attack on Authentication Protocol	H	H	N/A
	Bypassing Authentication	H	L	Port-Based Access Control

Table 2. Feasibility of attack patterns (continued).

Category	Attack Pattern	PCN	TPCN	TPCN SC
Resource Depletion	Denying Service via Resource Depletion	H	M	Compliance Verification
	Depleting Resource via Flooding	H	M	Traffic Filtering
Exploitation of Privilege or Trust	Lifting Credentials/Key Material Embedded in Client Distributions	M	L	Software Verification
	Lifting Cached, Sensitive Data Embedded in Client Distributions	M	L	Software Verification
	Accessing Functionality Improperly Constrained by ACLs	H	M	Small Rule Sets
	Exploiting Incorrectly Configured Access Control Security Levels	H	M	Role-Based Access Control
	Manipulating Writeable Configuration Files	H	L	Configuration Verification
Injection	LDAP Injection	H	H	N/A
	Sniffing Information on Public Networks	M	M	IPSec
	Manipulating User-Controlled Variables	H	L	Configuration Verification
	Manipulating Audit Log	H	L	Audit Verification
	Poisoning DNS Cache	H	L	Trusted DNS
Protocol Manipulation	Manipulating Inter-Component Protocol	M	M	N/A
	Manipulating Data Interchange Protocol	M	M	N/A
Time and State	Manipulating User State	H	L	Configuration Verification

6. Conclusions

Trusted network technology can help address the challenges involved in securing industrial control systems that are vital to operating critical infrastruc-

ture assets. Adding trust to industrial control networks eliminates security problems posed by inadequate controls, non-compliant devices and malicious users. It dramatically reduces vulnerabilities to malware attacks that constitute the majority of external attacks. Also, the likelihood of internal attacks is reduced via compliance verification, port-based access control, device and user authentication, and role-based access control. Implementation and maintenance costs are major issues, especially when deploying security solutions for industrial control networks containing modern and legacy systems.

References

- [1] E. Byres, B. Chauvin, J. Karsch, D. Hoffman and N. Kube, The special needs of SCADA/PCN firewalls: Architectures and test results, *Proceedings of the Tenth IEEE Conference on Emerging Technologies and Factory Automation*, 2005.
- [2] E. Byres, D. Leversage and N. Kube, Security incident and trends in SCADA and process industries: A statistical review of the Industrial Security Incident Database (ISID), White Paper, Symantec Corporation, Cupertino, California, 2007.
- [3] E. Byres and J. Lowe, The myths and facts behind cyber security risks for industrial control systems, *Proceedings of the VDE Congress*, pp. 213–218, 2004.
- [4] D. Capite, *Self-Defending Networks: The Next Generation of Network Security*, Cisco Press, Indianapolis, Indiana, 2006.
- [5] Cisco Systems, Implementing Network Admission Control – Phase One Configuration and Deployment, Version 1.1, San Jose, California, 2005.
- [6] Cisco Systems, Cisco TrustSec: Enabling switch security services, San Jose, California (www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns147/ns774/net_implementation_white_paper0900aecd80716abd.pdf), 2007.
- [7] Cisco Systems, Cisco NAC Appliance – Clean Access Manager Installation and Configuration Guide, Release 4.1(3), San Jose, California (www.cisco.com/en/US/docs/security/nac/appliance/configuration_guide/413/cam/cam413ug.pdf), 2008.
- [8] Cisco Systems, Getting started with Cisco NAC network modules in Cisco access routers, San Jose, California (www.cisco.com/en/US/docs/security/nac/appliance/installation_guide/netmodule/nacnmsg.pdf), 2008.
- [9] Cisco Systems and Microsoft Corporation, Cisco Network Admission Control and Microsoft Network Access Protection Interoperability Architecture, Redmond, Washington (www.microsoft.com/presspass/events/ssc/docs/CiscoMSNACWP.pdf), 2006.
- [10] M. Franz and D. Miller, Industrial Ethernet security: Threats and counter measures (www.threatmind.net/papers/franz-miller-industrial-ethernet-security-03.pdf), 2003.

- [11] Industrial Automation Open Networking Association, The IAONA Handbook for Network Security, Version 1.3, Magdeburg, Germany (www.iaona.org/pictures/files/1122888138-IAONA_HNS_1.3-reduced_050725.pdf), 2005.
- [12] Instrumentation, Systems and Automation Society, Integrating Electronic Security into the Manufacturing and Control Systems Environment, ANSI/ISA Technical Report TR99.00.02-2004, Research Triangle Park, North Carolina, 2004.
- [13] Microsoft Corporation, Network access protection platform architecture, Redmond, Washington (www.microsoft.com/technet/network/nap/naparch.mspix), 2004.
- [14] MITRE Corporation, CAPEC: Common Attack Pattern Enumeration and Classification, Bedford, Massachusetts (capec.mitre.org).
- [15] North American Electric Reliability Council, SQL slammer worm lessons learned for consideration by the electricity sector, Princeton, New Jersey (www.esisac.com/publicdocs/SQL_Slammer_2003.pdf), 2003.
- [16] Office of Nuclear Reactor Regulation, Potential vulnerability of plant computer network to worm infection, NRC Information Notice 2003-14, Nuclear Regulatory Commission, Washington, DC (www.nrc.gov/reading-rm/doc-collections/gen-comm/info-notices/2003/in200314.pdf), 2003.
- [17] R. Ross, S. Katzke, A. Johnson, M. Swanson, G. Stoneburner and G. Rogers, Recommended Security Controls for Federal Information Systems, NIST Special Publication 800-53, National Institute of Standards and Technology, Gaithersburg, Maryland, 2005.
- [18] M. Sopko and K. Winegardner, Process control network security concerns and remedies, *IEEE Cement Industry Technical Conference Record*, pp. 26–37, 2007.
- [19] K. Stouffer, J. Falco and K. Scarfone, Guide to Industrial Control Systems Security, Second Public Draft, NIST Special Publication 800-82, National Institute of Standards and Technology, Gaithersburg, Maryland, 2007.
- [20] Trusted Computing Group, Trusted network connect to ensure endpoint integrity, Beaverton, Oregon (www.trustedcomputinggroup.org/groups/network/TNC_NI_collateral_10_may.pdf), 2005.
- [21] A. Wool, A quantitative study of firewall configuration errors, *IEEE Computer*, vol. 37(6), pp. 62–67, 2004.