Process Solutions

Honeywell

White Paper

Mitigating Cyber Security Risks in Legacy Process Control Systems

Executive Summary

The term “legacy process control system” has different connotations for different people. To many, it refers to proprietary systems from a past era. To others, the term may imply the new generation DCS that have been founded on open technology, or systems using no-longer-supported Microsoft operating systems. These systems have fundamentally different architectures and present different risks.

The continuous evolution of the DCS enabled organizations to protect the investment in equipment and control strategies over long periods of time. However, interfacing decades-old controllers with current technology also makes this equipment indirectly vulnerable to attack. All these systems have one common denominator: they experience gaps in support. This makes them more vulnerable than contemporary systems. This white paper will discuss various techniques for protecting legacy systems, the problems surrounding these techniques, and new methods for analyzing security.

Table of Contents

Introduction

The term “legacy process control system” has different connotations for different people. To many, it refers to proprietary systems from a past era, such as the Honeywell TDC 2000 and TDC 3000 Distributed Control Systems (DCS). To other people, the term may imply the new generation DCS that have been founded on open technology, such as the Honeywell TPS and PlantScape systems introduced in the 1990s. To some, it may suggest the generations of Experion PKS systems using no- longer-supported Microsoft operating systems (e.g. Windows 2000 and soon Windows XP). These systems have fundamentally different architectures and present different risks.

It’s important to note that proprietary systems were not impervious to security risks. The mechanisms of attack were different in those days, generally relying on physical access and inside knowledge in the absence of external network connections. The mechanisms to attack today’s process control systems allow attackers to target assets across vast geographical distances.

Different attack scenarios have been developed and executed over the years. An attacker can sneak through the network, from firewall to firewall, to penetrate into a process control network. Today’s client-side attacks that could start anywhere, including:

- Within internal corporate networks
- From communication with external parties
- From your home computer

The discovery of the Stuxnet attack had a significant impact on the process control world. Suddenly everyone became aware that targeted, client-side attacks on process control systems could occur, and that specifically the national critical infrastructure could be attacked.

Security researchers focused on SCADA systems and discovered many new vulnerabilities, placing more pressure on both vendor and owner/operator organizations to secure their process control environments. Internet sites allowed for the scanning of SCADA systems directly connected to the Internet. Some of these systems still used default vendor passwords. While system vendors improved product design processes and security for new products, there are still many legacy systems left that have known vulnerabilities. These legacy systems have vulnerabilities in their operating systems, communication protocols, control applications and computer equipment.

The continuous evolution of the DCS enabled organizations to protect the investment in equipment and control strategies over long periods of time. However, interfacing decades-old controllers with current technology also makes this equipment indirectly vulnerable to attack. All these systems have one common denominator: they experience gaps in support. This makes them more vulnerable than contemporary systems. This gap can be caused by:

- Unavailability of security patches
- Loss of skills and knowledge
- Outdated system support documentation
- Inability to deploy specific security countermeasures on the older software and equipment

Security is not equivalent to making every system component resistant to attack. Security is an approach used to create multiple layers of defense around the production process to protect it. It is architecture, a chain with many links, where the weakest link breaks first. This weakest link in the chain is often a server or station with an outdated operating system, missing security patches, or a critical application for which software updates are no longer available. Each system component should contribute to overall resilience against attacks, but these components alone are not sufficient to protect the entire system. They need to be embedded in an architecture with multiple layers of protection. There will always be areas in a system (such as areas that require real-time performance) that need to rely on the protection layers surrounding them for their resilience against attacks.

This white paper will discuss various techniques for protecting legacy systems, the problems surrounding these techniques, and new methods for analyzing security.

How to Prevent or Delay a System from Becoming a Legacy System

Time converts state-of-the-art technology into legacy technology. Time also changes secure systems into vulnerable systems. This is why security must be viewed as a process, not a product. Security is a set of continuously-evolving strategies to counter attackers, who are constantly finding new ways to reach the target. To remain secure in the long term takes effort and requires investments. But there are ways to reduce the costs, such as preventing systems from becoming legacy systems.

Prevention Measures

Preventing systems from becoming legacy systems involves several aspects:

- Performing periodic system refreshes to maintain a system's serviceability
- Maintaining security patches to keep the system up-to-date with the latest vulnerability fixes
- Maintaining documentation so you can know your system and document the assets, network traffic flows and security controls

Obviously, time can't be stopped, so new developments in the process control environment will impact the security stance of the process control system.

Twenty years ago, control systems were built with proprietary technology with a lifecycle of at least 15 years. Today's open technology world has other rules. Operating systems undergo major changes at least once every three years, while hardware platforms change even more rapidly. Changes in CPU, memory, and storage technology enforce changes in operating systems to support this new technology. However, this evolution simultaneously creates gaps in the serviceability of the older systems. Legacy operating systems do not provide the system software to support the new technology, and legacy hardware platforms do not provide the performance and technology to support the new operating systems.

A five-year refresh cycle seems to be a reasonable compromise between a reliable and serviceable system with a high availability and the return on investment for the new software and equipment. Changing the hardware and operating system will impact the DCS software. Migrating to a higher DCS release supporting the new operating system becomes unavoidable.

How does this affect security? Remaining secure requires the installation of a continuous flow of new security patches, which contain software fixes for vulnerabilities. Vendors of operating systems have a limited support window for security fixes; once the product is no longer sold, the support is generally limited to a three-to-five-year period. After this period, no more security patches will become available, resulting in a rapid degradation of the product's security. Software that protects the process control system, such as anti-virus and whitelisting applications, also has support limitations pertaining to legacy platforms. Therefore legacy systems can suffer from the unavailability of security patches as well as the unavailability of security protection software.

Knowing your system is essential when building a secure system. In order to keep your process control system secure, you must have a good overview of its applications, configuration and communications. You need to know which protection layers exist and on which security controls these protection layers depend.

A network drawing showing a firewall doesn't explain how this firewall contributes to the overall protection. Zone and conduit channel diagrams are a better method to document the security architecture because they show how assets are grouped in security zones, the sequence of defenses protecting these zones and the interdependencies of defenses. These diagrams show how legacy systems are separated from other system components and how their vulnerabilities are remediated.

Security relies on proper configuration management processes and up-to-date documentation. If you aren't sure which components are in your system, you can't determine whether they are vulnerable. You must know the assets and network traffic flows to design effective security architecture.

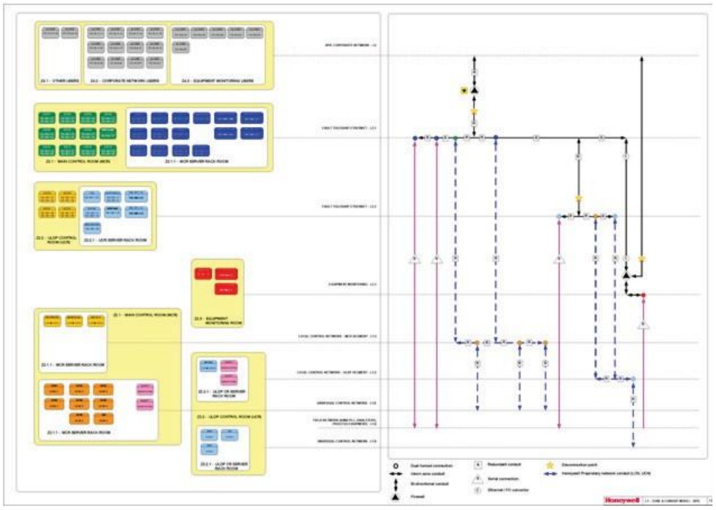

Figure 1: Zone and conduit channel diagrams

Delaying Measures

Can a system's transformation into a legacy system be delayed? Can the serviceability of a system be extended? Yes, but maintaining accurate documentation is imperative for the serviceability of any system. Apart from out-of-date documentation, other factors contribute to the loss of serviceability, such as the aging of software and hardware, and having a backlog of security patches.

As previously discussed, operating systems must support new hardware, different storage technology, faster CPUs, better graphics and larger memory. Isolating hardware platform changes from their impact on the operating system by using server virtualization can help delay the aging process. The virtualization software layer separates the new technology of a new server platform from the legacy operating system, allowing the operating system to interact with the new technology as if it was still the old server platform.

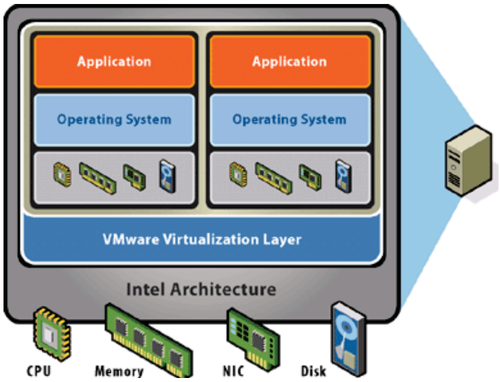

Figure 2: Virtualization diagram

Software such as VMware creates a layer between the operating system and hardware. This layer, the hypervisor, simulates the server or station hardware and provides a Virtual Machine (VM). This virtual layer allows for the application of new hardware and storage technology. The result? A legacy Windows 2000 operating system can exist in a VM running in a brand new server. This extends the lifecycle of a system to at least the software lifecycle—which could be up to eight years. The capability to run multiple VMs on the same hardware platform also provides additional advantages, such as a smaller system footprint and reduced hardware cost.

How does virtualization support security? Virtualization runs legacy software on more capable hardware; it adds security functions between the VM and the hypervisor[1]. Security functions can be implemented as a shield around the VM to do their job without being installed on the legacy software residing in the VM. Organizations can implement anti-virus and virtual patching solutions[2] without leaving a footprint on the VM that runs the legacy software. Additionally, the way the VM communicates with the hypervisor reduces the attack surface of the virtualized function. Several types of attacks, such as a number of buffer overflow attacks, will fail because of the hypervisor limitations.

Another form of virtualization is application virtualization. This method combines the application and the operating system in one package, making the application less dependent on the server's operating system. Such "sandbox" architecture also reduces security risk and makes the application more tolerant of a changing environment.

But how should organizations address the unavailability of security patches or the inability to install them? There are two security controls that can be used: Application White Listing (AWL) and Virtual Patching (VP). Each method has advantages and disadvantages.

Application White Listing

AWL is a method that prevents unauthorized executable code from running. Malware can install malicious software on the machine to carry out attacks. When this malicious code is initiated, AWL will stop it, acting during application execution time.

Virtual Patching

Virtual Patching (VP) is another security control technique that acts on network level. This can be either the physical wired network, or it can be within the virtual network of a virtualization solution. VP inspects the traffic and monitors the traffic for the exploitation of a particular vulnerability. It does this with vulnerability filters, which inspect the attack activity rather than monitor a particular bit pattern in the traffic. Vulnerability filters do not use static signatures.

AWL and VP work differently. AWL leaves a footprint on the legacy node. It needs to be installed on this node, so it must be compatible with both the operating system and the application. VP is network resident, so it doesn't leave any system footprints in the legacy node or the network.[3]

[1] The hypervisor is the software layer that creates a layer between the operating system in the "Virtual Machine" (VM) and the server platform.

[2] Various security controls will be discussed later in this white paper.

[3] No IP or MAC address.

For a legacy node that supports AWL, the implementation can offer additional security when security patches are no longer available. For legacy nodes where AWL cannot be installed, VP is still an option. Both methods reduce the risk of having legacy nodes in your system and therefore delay your system from becoming vulnerable to attacks.

How to Protect Legacy Systems

Security Architecture

As previously discussed, knowing your system is essential. The security posture of a specific system component is also important to know, but it's not essential for securing this component. As an analogy, a person doesn't need to be bulletproof to be protected from being shot – he could instead wear a bulletproof vest. Similarly, the security architecture as a whole provides the layers of defense to protect the critical assets of the process control system, even if individual components differ in level of vulnerability. The security architecture will surround it and place the most vulnerable components in the internal layers of the architecture.

Security by design is a methodology used to structure such a layered defense. The Honeywell Security by Design process formulates five steps[4] for the design process. These steps help you make the right security design decisions.

Step 1—What are the assets you are trying to protect? How do systems interact, and what network traffic flows do you need to protect? This question might seem basic, but it is often ignored. For example, the following pose all different security problems that require different solutions:

- Securing a network connected to a camera network—Cameras create network access outside the physical security boundaries of the control room and server rack rooms.
- Securing a wide area network connecting multiple production systems – Connecting multiple systems to a common network create the risk of a security incident (such as a malware infection) impacting connected systems.
- Interfacing a process control network with a wireless network—Wireless networks don't stop at the physical boundaries of a plant like a fence and increase the risk of unauthorized access if not properly protected.
- Interfacing a DCS with a safety system—DCS and safety systems have distinct roles in an industrial control system. A security impact on the DCS should not impact the safety system at the same time.
- Interfacing with terminal servers—Terminal servers connect the process control network to many different support systems. A security breach in one of those systems can impact the main control system if not properly protected.

Step 2—What are the risks to these assets? Consider the need for security. Are you protecting exclusively against unintentional attacks, or are you including intentional attacks? Answering these questions involves understanding what is being defended. National critical infrastructure obviously requires more security than a production process manufacturing soap. Considering the consequences of a successful attack, who wants to attack, which methods are available, and why they want to attack are important to determine the level of security required.

Step 3—How well does a particular security solution mitigate a risk? If a security solution doesn't solve an issue, it is no good. Examples of this can include demilitarized zones (DMZ) that do not provide any barriers to access, and network filters that are easily bypassed. These "solutions" create more insecurity than security.

Step 4—What other risks or unintended consequences does the security solution cause? For example, you can implement an anti-virus solution, but then you must also secure the daily update of the signature files. Security solutions often have ripple effects and can cause new security problems. It's important that these new problems are smaller than the older ones.

Step 5—What are the costs and trade-offs of the security solution? Every security system has them. There could be investment required, less user convenience, or an impact on overall system resilience and availability. Installing security patches can also induce labor cost, risk of unavailability, and loss of functionality due software reboots.

[4] The five step process was initially formulated by Bruce Schneier in his book "Beyond Fear" and modified for process control systems.

These five steps by themselves do not lead to a secure system, but together, the steps provide tools to evaluate and analyze a design. This five-step process may seem obvious when stated in the abstract form, but applying the steps to real situations is hard work. It requires detailed information about all of the components that make up a process control system. It also requires experience with the techniques used, such as:

- Threat modeling
- Security zone, conduit, and channel modeling
- Defining security patterns for the authentication and authorization processes

A well-designed architecture that uses multiple layers of defense can protect legacy systems. Even if vulnerable to many attacks, a successive layering of protection mechanisms reduces risk.

Can You Air Gap Legacy Systems to Protect Them?

Locking the doors at home and heading for the cellar is not security—security requires constant vigilance and responding to new threats, rather than hiding from them. Similarly, removing all external access to a process control system will not protect the control system. It will reduce what is called the attack surface, but only from server-side attacks. Popular client-side attacks are not stopped by air gaps. The Stuxnet attack was a client-side attack, and it successfully targeted an air gapped system. Of course, Stuxnet was an exception; this type of targeted attack is difficult to stop without also seriously impacting the business processes. Security remains a trade-off between being secure and doing business. However, similar attacks launched since, such as Duqu, had a much bigger impact. This type of attack often propagates over removable media such as USB devices. The most common cause of malware infections in process control systems is the use of an infected USB stick, CD/DVD or connecting an infected PC to the network.

Exchange of data for various supply chain and reporting processes is a crucial function for the business. Security remains a trade-off between the business benefits offered by the process control system and security counter measures restraining these benefits. Isolating legacy systems as an answer to security threats has a limited effect. Today’s biggest threat is the client-side attack, since it has many mechanisms to propagate, including via the isolated network once it has entered into the system.

Application Whitelisting

Application White Listing (AWL) is probably the most natural protection strategy. AWL allows what is explicitly authorized and blocks everything else by default. It is the opposite of traditional anti-virus programs, which allow everything that is not explicitly blocked. This means AWL can cause false negatives, and anti-virus (blacklisting) can cause false positives. The advantage to using AWL with legacy systems is that it protects the status quo and blocks new unexpected actions, which is exactly the kind of behavior that can occur if malware exploits a particular vulnerability and downloads or drops executable code into the computer. Anti-virus would only offer protection against known exploits that have a recognizable signature to detect. This means AWL offers better protection for legacy systems, since they become more vulnerable for new (zero day) attacks over time due to lack of fixes for newly-disclosed vulnerabilities.

However, AWL has its limitations, both in detecting attacks as well as applying the technology to legacy systems. AWL has difficulty intercepting attacks that are fully memory resident and attacks that are exploiting interpreted code, such as used for mobile code (JavaScript, Pearl) and web-based applications. Another limitation of AWL in relation to legacy systems is that the software needs to be installed in the system. This software might not be tested for compliance with the legacy system and could also overload the system performance. Therefore AWL always leaves a distinct footprint in the system and cannot always be used. An alternative is virtual patching (VP).

Virtual Patching

Virtual Patching (VP) is relatively new technology. VP is fully network-based, which makes it have no impact on the legacy system. It is an appliance in the network with input and output ports. It is called a “bump in the wire”. Honeywell’s VP solution has no IP or MAC address, making it fully transparent to the network. It protects the system using vulnerability filters, which monitor the network activity and intercept the traffic if a particular known vulnerability is exploited. It doesn’t use a static bit pattern signature as filter, but instead uses an activity filter evaluating every protocol step.

VP is in line with the network and works as an Intrusion Prevention System (IPS). The manufacturer of the Honeywell solution creates a filter as soon as a new vulnerability is discovered by their research institute or disclosed by others. Since one vulnerability is often exploited in various ways, a vulnerability filter can stop many exploits, including those that are brand new. VP is valuable when protecting legacy systems because it doesn't require any software to be installed on the legacy system and doesn't impact performance[5].

The limitations of VP are that the technology is restricted to the network. It will stop attacks over the network, but it can't stop attacks using removable media (e.g. USB drives) or file sharing as propagation methods. A combination of VP with anti-virus is the most logical solution for remediating this security gap.

How to Choose Which Security Control is Right for Your System

The answer depends on various characteristics of the particular legacy system. Is it a single node? Is it a system comprised of multiple nodes? Is the problem the operating system? Can you replace the hardware? Is the application the bottleneck? Is it a "ghost" system without documentation? There are many different situations and variables to consider when selecting an optimal mitigation strategy.

Security strategy examples include:

Method	Characteristics
Application White Listing (AWL)	• AWL protects the legacy system if new security patches are no longer available. • AWL protects against unauthorized execution of executables at node level. • AWL requires that both the software platform (compatibility) and the hardware platform (performance) support the solution. • AWL protects against malware infections originating from removable media and file sharing. • AWL can be used for a particular node or at system level.
Virtual Patching (VP)	• VP protects the legacy system if new security patches are no longer available. •VP leaves no footprint on the network or on the legacy node, so it can always be applied. •VP intercepts attacks on network level using vulnerability filters. Vulnerability filters monitor the network activity and block illegal activity. •VP would typically be used to protect multiple nodes. •VP is a "bump in the wire".
Virtualization in combination with VP	• When used in combination with virtual machines, VP can protect one or multiple virtual machines. •Runs in its own VM and communicates with the VP device on the network. •Can be combined with VP "on the wire".
Creating an air gap	• Air gaps would protect against all server-side attacks. • Air gaps impact the business functioning, isolating the process control system. It eliminates the possibility for real-time integration of various business functions. • Air gaps offer no protection against client-side attacks. • Air gapped systems still require anti-virus, which would require an update mechanism to stay effective.

A single node can be best protected by installing AWL, but if for some reason AWL is not an option (due to an unsupported or slow hardware platform), then using VP is an alternative security strategy because its characteristics of being fully network resident matches the requirement better. Multiple nodes can be protected by both AWL and VP in combination with AV. Using only AWL or only VP wouldn't provide sufficient protection. AV is supplemental for both solutions.

VP can also be used in combination with virtualization, where a specific VM protects the communication in the virtual server environment and the wired environment. In this way, VP can protect one or multiple VMs.

[5] The added network latency is less than one millisecond and therefore negligible.

The diagram shows the various protection solutions and where they are active. In principle, these solutions should be supplemented with each other, as each technology has its weak and strong points. AWL struggles with attacks that fully remain in memory and replace authorized code, as well as attacks based upon interpreted software code. AV also struggles with memory-based attacks and is vulnerable to changing malware signatures. Both AV and AWL do not protect against a denial of service (DDoS) attack. With VP, protection is limited to network traffic. Legacy systems will not always support all protection options. Therefore often methods should be combined.

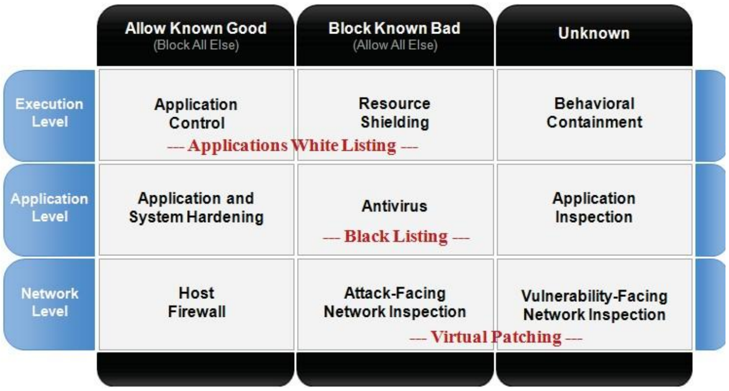

	Allow Known Good (Block All Else)	Block Known Bad (Allow All Else)	Unknown
Execution Level	Application Control	Resource Shielding	Behavioral Containment
Application Level	Application and System Hardening	Antivirus	Application Inspection
Network Level	Host Firewall	Attack-Facing Network Inspection	Vulnerability-Facing Network Inspection

Figure 3: Various protection solutions and where they are active

Conclusions

Legacy systems form a weak link in the security chain. If a legacy system gets compromised, chances are that the other parts of the system are impacted. When protecting legacy systems, consider:

- Replacement or upgrade. The logical choice is to remove the vulnerable system. However, this requires budget and might not be opportune because of the impact on the continuity of the production process.
- Hardening. Every system should be hardened, including legacy systems.
- Apply the “least privilege” principle. Various legacy systems have limited, role-based access control functionality, providing users with more authority than needed. Try to restrict authorizations as much as possible.
- Apply strong passwords. Legacy systems often allow the use of weak passwords. Sometimes, even default passwords are used. Correct this. If available, use domain authentication rather than workgroups.
- Apply one of the discussed methods (AWL, VP) to compensate for the absence of security fixes.
- Consider virtualization to extend the lifecycle of a system.
- Apply a defense-in-depth security defense. Use the security by design methodology to evaluate the various design decisions.
- Make certain that there are sufficient spare parts available.
- Make certain you have a backup that can be restored.
- Maintain the skill level. Often plants are confronted with legacy systems when one employee changes jobs or retires and the skills to support the system are lost.
- Maintain the system documentation.

If you can apply the above recommendations, it is a good start for protecting your systems against threats and extending the lifecycle of legacy equipment.

For More Information
Learn more about how Honeywell's solutions
for legacy systems, visit our website
www.becybersecure.com or contact
your Honeywell account manager.

Honeywell Process Solutions
Honeywell
1250 West Sam Houston Parkway South
Houston, TX 77042

Honeywell House, Arlington Business Park
Bracknell, Berkshire, England RG12 1EB

Shanghai City Centre, 100 Junyi Road
Shanghai, China 20051

www.honeywellprocess.com

WP-12-12-ENG
November 2014