DEFINING THE 21ST CENTURY CYBERSECURITY PROTECTION PLATFORM FOR ICS

November 2014

About the Authors

Mario Chiock, CISSP, CISM, CISA

Cybersecurity & Disruptive Technology Executive Adviser

Mario Chiock possesses over 33 years of experience in Oil Field operations, IT, Security, Risk, Privacy and Auditing. Prior to his current role as executive advisor, Mario was the CISO at Schlumberger where he was responsible for developing the company's worldwide, long-term cybersecurity strategy. He is recognized for his leadership and management in all aspects of cybersecurity throughout the company as well as within the community.

Through his vision, he successfully transitioned Schlumberger from legacy firewalls to a more robust infrastructure based on next-generation firewalls. His experience in successfully deploying advanced technologies and approaches also spans Incident Response, Advanced Threat Prevention, Digital Rights Management, and RightsWatch to watermark and fingerprint sensitive documents. He also implemented federation services to minimize 3rd party risk and created the extended security team to foster collaboration with other IT groups. Mario is also known for applying his Lean Six Sigma expertise for measuring performance and creating dashboards that have led to more simple operations and reduction of waste.

Outside of Schlumberger, Mario has been an active member of the Information Systems Security Association (ISSA) for over 14 years; he has held numerous board positions in the Austin, Capital of Texas Chapter, as well as the South Texas Chapter in Houston. He was president of the South Texas Chapter in 2007, bringing in the "Chapter of the Year" award. He continues to serve on the board. Mario is also an active volunteer trainer for the Certified Information Systems Security Professional (CISSP) certification, and has mentored and helped many to obtain their CISSP certification.

Mario was recognized as one of the top 25 out of more than 10,000 security executives in the ExecRank 2013 Security Executive Rankings, he also won the 2012 Central Information Security Executive (ISE) "People Choice Award", in 2014 he is a recipient of the CSO40 – 2014 award, named "ISSA Fellow", won ISC2 Americas Information Security Leadership Awards (ISLA) and won the "ISSA Honor Roll" award.

He is an active member of the Houston Security community and gives security talks, training and volunteers his IT security expertise to local non-profit organizations. He is currently Vice-President of the Houston InfraGard Chapter, has served in executive and technical Advisory boards of many security companies such as WatchFire (Now IBM), McAfee, ISS (now IBM), Qualys, Solera Networks (now Blue Coat) and currsently is active on the Palo Alto Networks® Executive advisory board, Onapsis advisory Board and Watchful Software Board.

Mario has a CISSP, CISM & CISA Certifications, and is currently the chair for the American Petroleum Institute Information (API) Security Sub-Committee and was involved in the creation of the Oil & Gas ISAC.

Del Rodillas, MSEE, MBA

Senior Manager, SCADA and Industrial Controls Cybersecurity

Del Rodillas is responsible for Palo Alto Networks' SCADA and Industrial Control Systems industry initiative. He works with a range of asset owners in Utilities, Oil & Gas, Transportation, Mining, Manufacturing and similar industrial market sectors to increase awareness and education on the need for improving control systems cybersecurity as well as what technologies and best practices can be deployed to address the security challenges.

Del has extensive technology industry experience spanning Cybersecurity, Networking, Aerospace/Defense, and Semiconductors with roles in strategic marketing and engineering. He holds a Masters in Electrical Engineering degree from Santa Clara University and an MBA from the Wharton School of the University of Pennsylvania. Del participates in several industry consortia including EnergySec, the ICS-ISAC, and the San Francisco Bay Area Chapter of Infragard.

Executive Summary

SCADA and Industrial Control Systems found in critical infrastructure and manufacturing industries have enjoyed unprecedented levels of agility, speed, and cost savings with the pervasive adoption of information technology and increased connectivity to supporting networks. Even an Industrial Internet of Things (IIOT) is starting to emerge, enabling new capabilities such as predictive maintenance and even new business models. However, with this modernization have also come undesired IT vulnerabilities and other threat vectors which are increasingly being exploited by malicious actors such as nation states, cybercriminals and malicious insiders. Recent years have shown a concerning rise not only in the number but also the sophistication of attacks specifically targeting critical infrastructure and manufacturing asset owners. Real-world cases have shown disruption of critical processes and even destruction of ICS equipment. The need for improved security in ICS has never been higher and has become a board-level issue for many organizations.

While IT administrators have been quick to deploy the latest and greatest technologies and practices to secure corporate environments, operational technology (OT) administrators have not been as aggressive. The extreme sensitivity to ensuring availability and performance of the industrial process has led to a more conservative and rigorous approach to how security is deployed and maintained. For example, to minimize disruptions to the process and sometimes because of the very nature of a process, it is not uncommon to have systems with maintenance cycles in excess of 12 months with some even running multi-year cycles. Within this window software cannot be patched and AV signatures cannot be updated. Other times, administrators avoid inline protections such as network IPS or AV because of the concern over accidental blocking or performance degradation. These products are put in detection-only mode or thrown out altogether. Even practices commonplace in IT environments such as vulnerability scanning can cause malfunctions or denial of service scenarios in industrial controllers which were not designed to deal with such events. These constraints makes securing SCADA/ICS environments both unique and difficult.

The result is that many organizations are therefore still working with a mixed-bag of antiquated security technologies which operate in silos, are difficult to manage, provide limited situational awareness and do not provide the kind of preventive security required. Hence such organizations become prime targets for attackers who are likely using similar environments as quality assurance test beds for their sophisticated attacks. The bigger gap is the inability to address the constantly evolving threats which utilize never before seen attacks. Actions must be taken to build up the right capabilities to better protect ICS.

A new kind of platform—a 21st century ICS security platform—is required to properly secure control systems from the new threat landscape. This platform consolidates different core technologies in a way that ensures prevention even against advanced attacks. The integration must also allow interfaces to alert & perform security actions in an automated way, not only with its own services but also with other supporting technologies. It must also facilitate information sharing within the organization as well as with peer organizations. In the same way that the bad guys collaborate to develop targeted attacks, so too must the good guys.

In this paper we take a look at the nine core capabilities that define this 21st century security platform for industrial control systems.

1. Integrates network and endpoint security with a threat intelligence core
2. Classifies traffic based on applications and users, not ports and IP
3. Supports granular network segmentation including role-based access
4. Natively blocks known threats
5. Detects and prevents attacks by unknown malware
6. Stops Zero Day attacks to the endpoints
7. Provides central management and reporting
8. Secure use of mobility and virtualization technologies
9. Powerful API and industry-standard management interfaces

With these capabilities in place, organizations are better able to deter advanced threats and adapt and scale as the threats evolve.

Introduction

The Evolution of Industrial Control Systems

The automation systems used to monitor and control industrial processes in factory floors and critical infrastructure such as electric substations and oil rigs have many names: Industrial Control Systems (ICS), Supervisory Control and Data Acquisition (SCADA), and Distributed Control Systems (DCS) to name a few. These systems which are holistically referred to in this paper as "ICS" have evolved dramatically over recent decades. What once were isolated, proprietary systems interconnected by serial communications technologies are now highly interconnected and geographically-distributed systems which utilize commercial-off-the-shelf (COTS) products and the Internet Protocol (IP). This merging of Information Technology (IT) and Operational Technology (OT) has allowed many critical infrastructure and manufacturing asset owners to enjoy tremendous productivity and cost savings. Further efficiencies are anticipated with the deployment of mobility, virtualization and even cloud-based components.

The New Cyberthreat Landscape in ICS

Economically, IT-OT integration has been very advantageous for many asset owners. However, along with these benefits has come a wider exposure to a variety of cyberthreats which could compromise availability, integrity and confidentiality. Organizations are more than ever being called out to revisit their control systems security posture to assess just how capable it is at preventing cyberincidents.

Some of these threats are unique to ICS components and others relevant to both IT and OT products. In addition, these threats could originate from within the ICS or originate from extraneous locations. Finally, cyberthreats to ICS could be malicious or accidental in nature. Figure 1 lists the most concerning ICS threat vectors as identified by the respondents in the SANS Institutes 2014 Survey on Industrial Control Systems [1].

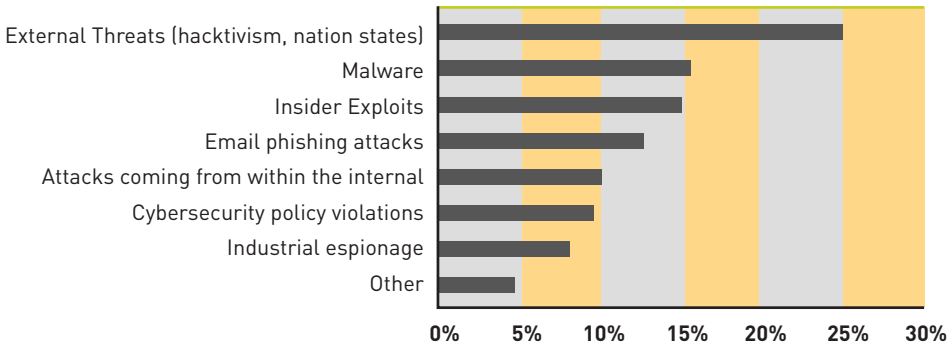

Figure 1: Top threat vectors in ICS (SANS ICS Survey 2014).

In terms of the first category of external threats, Stuxnet was the first publicly disclosed cyberattack targeting ICS specifically. It exploited applications, files and vulnerabilities in COTS and ICS software to achieve its goal of disrupting Iran's uranium enrichment program. In this case, ICS equipment, the centrifuges themselves, were reported to have been damaged, highlighting the reality of cyber-physical consequences. Since then, we have seen an increased sophistication of the techniques used by targeted attacks to ICS. Reports for the recent Energetic Bear campaign revealed two new and concerning techniques used by APTs to attack ICS [2]. First it employed malware hidden in the ICS software packages downloaded from vendor websites. It also utilizes ICS protocols to learn about the affected organization's environment. The risk of these protocols being used to control or disrupt the process cannot be dismissed.

The next category of attacks are spearphishing attacks, or to be more general, social engineering techniques which include watering hole attacks or something as simple as leaving malware-infected USB thumb drives in target organization parking lots as bait. Virtually all targeted attacks involve the compromise of the endpoint at some point via social engineering. For example Stuxnet purportedly utilized USB drives to infect laptops used by support personnel and Energetic Bear applied a combination of spearphishing, watering hole, and trojanized malware attacks. The concentrated spearphishing campaign to Norwegian oil and gas companies around August 2014 [3], where there were 50 confirmed attacks to organizations including StatOil, highlights how this has become a staple method for hacker organizations in trying to breach networks of critical infrastructure operators.

After APTs and spearphishing, the next biggest concern is the introduction of malware into the ICS. This is often done by accident through infected mobile computing devices or removable media used by personnel with access to the ICS. Worms could also sneak into the ICS via "trusted" vendor and partner networks. Whether malicious or unintentional in nature, malware could lead to costly downtime and potential safety issues. These often do not make the news, but can result in multi-million dollars in losses due to lost production, remediation costs, and perhaps legal feels when incidents involve injury, death, or environmental damage.

Insider exploits are also of a concern. A good publicly disclosed example of that is the Maroochy Shire incident [4]. In this case, a disgruntled employee of an ICS vendor supporting the Shire's sewage treatment works took vengeance after an employment-related disagreement. Using his deep knowledge of the shire's sewage treatment control systems including an unsecured wireless network, he released 800,000 liters of sewage into the local parks, rivers and hotel grounds causing significant environmental damage.

While not a cyberthreat perse, the risk of failing a regulatory audit has increased the pressure on security organizations. Several countries have adopted regulations around controls systems security that could lead to severe fines for non-compliant organizations. In the U.S. and Canada, the NERC CIP standard has been adopted in the Electric Utilities industry. For chemical facilities in the US, the CFATS standard serves a similar purpose of enforcing cybersecurity compliance. The NIST Cybersecurity framework (NIST CSF) is a more recent standard which government agencies must be compliant to while serving as a best practices reference for other ICS asset owners.

Is Your Organization Equipped to Deal With These Threats?

The discussion on the different types of cyberthreats raises some very important questions organizations need to ask themselves in terms their ability to defend their ICS from the range of cyberthreats to ICS.

- Do you have the right level of traffic visibility to validate proper use of the ICS and more importantly quickly detect improper use? How easy is it to extract this information?
- Can you enforce sufficient access controls that align to business policies and effectively limit extraneous and internal attack vectors, while meeting stringent performance requirements? How easy is it to deploy and maintain the controls?
- Are your unpatched or possibly unpatchable legacy systems protected from exploits and malware? Could you further reduce downtime due to cyberincidents or patch maintenance?
- If faced with a targeted cyber attack which utilizes methods and malware never before seen in the wild, would you be able to prevent the attack?
- Do your network and endpoint security solutions work together to support the goal of prevention or are they disjointed?
- Do your security solutions facilitate or increase the burden of meeting regulatory standards?
- If you are using advanced technologies such as virtualization or mobile devices is your security implementation consistent or might these serve as "chinks in the armor"?

The ICS Security Gap

Most critical infrastructure and manufacturing organizations have some level of cybersecurity today. However there is a still a good portion of organizations, especially in less regulated industries, who have legacy technologies which are inadequate at addressing modern day ICS cybersecurity challenges. Figure 2 below shows a typical legacy security configuration found in many ICSs today:

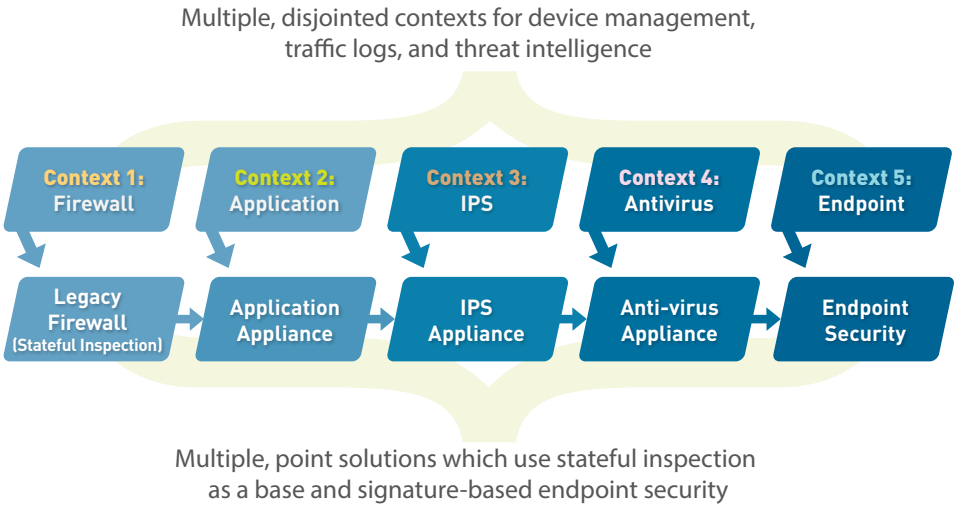

Figure 2: Typical Legacy Security.

From a network security standpoint, legacy solutions are often based on stateful inspection firewalls which do not provide the layer-7 visibility and user-based access controls required to effectively detect anomalies and minimize exposure. Organizations try to fill the gap by deploying multiple disjointed solutions such as application, IPS and AV appliances, but this typically results in increased risk of misconfiguration, silos of uncorrelated information, performance degradation and increased costs in terms of capital and operational expenditures. To add to the difficulties, existing endpoint security products operate separately from the network security, but more concerning is how they are only able to address threats which have known signatures, strings, and behaviors. They are unable to prevent attacks which use never-before-seen exploits and malware. Given the high risks involved in protecting ICSs, security solutions must be able to prevent attacks, even Zero Day attacks. Furthermore, the proliferation of point solutions has increased the load on organizations and made the job of securing ICS very complex.

21st Century Platform Security for ICS

Organizations can no longer rely on disjointed and ineffective legacy point solutions to defend critical infrastructure. The stakes are just too high. They need a 21st century cybersecurity platform that has the complete and tightly-coupled set of capabilities to prevent threats and reduce the burden on organizations in deploying and maintaining security. When selecting a platform, one must look for the following nine must-have capabilities:

1. Integrates network and endpoint security with a threat intelligence core
2. Classifies traffic based on applications and users, not ports and IP
3. Supports granular network segmentation including role-based access
4. Natively blocks known threats
5. Detects and prevents attacks by unknown malware
6. Stops Zero Day attacks to the endpoints

7. Provides central management and reporting
8. Secure use of mobility and virtualization technologies
9. Powerful API and industry-standard management interfaces

1. Integrates network and endpoint security with a threat intelligence core

The process as we have already discussed can be compromised directly at hosts such as HMIs or automation servers or they could pivot from other hosts such as a 3rd party vendors laptop or adjacent networks. It is very clear today that advanced threats will exploit weaknesses on both the network and hosts in a highly-orchestrated fashion to achieve their agenda. Organizations must be aware of this and have provisions that can stop threats across the attack life-cycle, the so-called "kill-chain" model where preventive methods are employed across both network-based and endpoint-based attacks. Besides just having this preventive toolbox as individual parts, the new requirement is that the preventive mechanisms must work collaboratively and share threat intelligence amongst each other.

There are at least two very powerful benefits organizations can derive by adopting such a platform. The knowledge gained about threats collected at endpoints can now be correlated with the knowledge learned about threats traversing the network. Organizations can better understand and respond to cyberattacks especially targeted attacks when they are able to view these repositories of threat intelligence with shared context. Secondly, protections for threats discovered at endpoints can be quickly sent to the network to prevent threats from propagating and protections for threats discovered at the network can also be distributed to stop attacks at endpoints. The two work hand and hand to prevent threats and the threat intelligence core provides centralized, automated intelligence. The challenge with most existing solutions is that endpoint and network security operate in separate silos. Actually this has been the paradigm for most organizations, but newer technologies have begun to tightly integrate these capabilities along with a shared threat intelligence core or cloud that automates threat intelligence and dissemination of protections. [5]

Figure 3 below shows the concept of integration of endpoint, network and a threat intelligence core to protect the process. The red arrows represent areas where advanced threats could initiate their attacks and the blue arrows capture the interaction of the network and endpoint security with the threat intelligence core or cloud. This is the basis of the 21st century security platform. Make sure to select a platform that tightly integrates these components as described above.

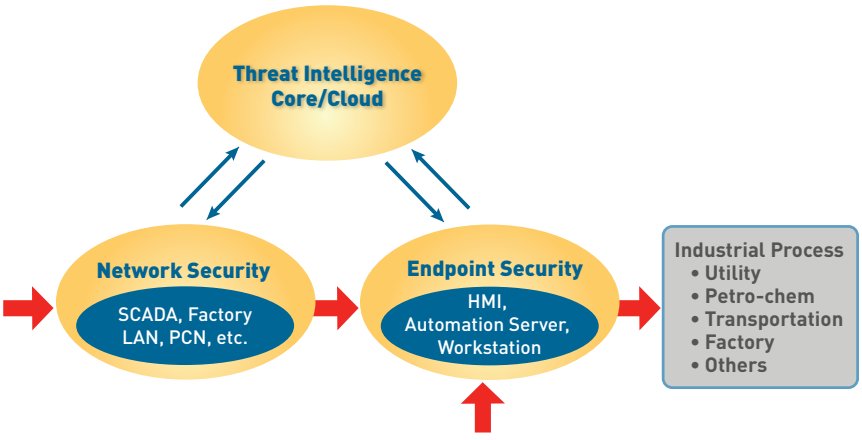

Figure 3: Concept of Endpoint, Network and Threat Intelligence Core as the basis for the 21st century security platform.

2. Classifies traffic based on applications and user, not ports and IP

Drilling down into requirements for network security, often discussed capabilities are layer 7 (application) and layer 8 (user) visibility. While these may have been nice-to-haves in the past, both are now essential to enable the level of visibility required to detect anomalous use. And rather than being add-ons, they need to be intrinsic in the technology and done with high performance. Here are several reasons why these are essential capabilities.

1. **Advanced threats cleverly use open ports**—targeted attacks are likely to exploit ports open in the ICS environment for example as channels for command and control. By classifying traffic at the application layer, one will be able to better distinguish between expected traffic and potential malicious traffic exploiting open ports.

2. **Applications port hop**—Aside from the malicious traffic, other risky or bandwidth-hogging applications are introduced by personnel to improve productivity or provide leisure. These applications often port hop to avoid detection. By identifying applications, one is better able to detect such undesired applications that could impact process availability.

3. **Some protocol functions are more "interesting" than others**—ICS protocols like Modbus and DNP3 have read and write functional variants. As write commands have the ability to alter a PLCs state and potentially take down the process, one will want to have visibility to functional variants to increase intelligence on the nature of control systems traffic.

4. **Adding user/user group visibility increases intelligence**—In all three scenarios adding the context of user or user-group to the application and protocol traffic allows one to detect anomalies based on role.

When selecting your platform, make sure that application visibility are natively available. Also make sure that the platform supports integration of user-IP mapping repositories which can then be contextually linked to application/protocol traffic.

3. Supports granular network segmentation, including role-based access

A common misconception is that perimeter based security is enough for protecting control systems from cyberthreats. As shown in our discussion of threats, separating SCADA from the business network, for example, is not enough to stop cyberthreats. Targeted attacks will find a way into the ICS, malicious insiders leverage their inside knowledge, well-meaning engineers unintentionally introduce risks, and malware jumps from "trusted" partner/vendor networks. More granular security zones interconnected by more intelligent segmentation gateways or "conduits" are required to ensure that users are given just enough access to do their jobs. This is the so called least-privilege approach described in ISA 62443 [6] or "zero-trust" model developed by Forrester Research [7].

A next-generation ICS security platform should not only be able to see the traffic with high granularity at the application and user level, it should also be able to apply these parameters in policy. By having the capability to segment the network by applications and users, the concept of role-based access control can now be applied between security zones. The platform must make it very easy and intuitive to apply application/protocol control by user and not require multiple policies to realize the desired access control. This reduces the administrative burden and reduces the likelihood of mistakes. Figure 3 below shows this important first step of reducing one's attack footprint by controlling the use of protocol, applications and other potential vectors.

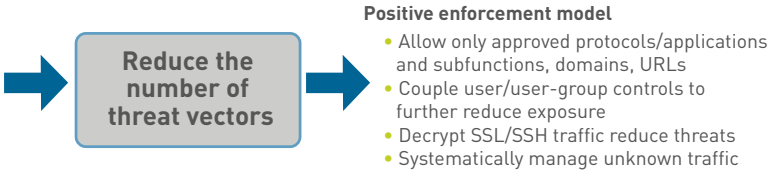

Figure 4: Applying segmentation to reduce the number of threat vectors.

Some powerful use cases can now be considered with this capability to not only whitelist protocols and applications, but also whitelist how individual users or user groups use them:

- Allow 3rd party vendors access to Modbus reads from a DMZ into a PLC zone. If they need to make any write changes, a ticket would need to be opened to allow a Modbus programming commands
- Allow only certain business users access to a Historian server in the process control network which may be using OSIsoft PI or other ERP/database applications
- Limit the use of administrative applications such as SSH, RPC, RDP only to approved SCADA admins who understand the risks involved with maintaining critical infrastructure assets

It is interesting to note that some technologies, like data diodes, were created to address the access control and security limitations of stateful inspection firewalls. Data diodes limit traffic to one direction at the IT-OT perimeter, for example to allow data flow only from the ICS environment to the business network. However many applications still require bidirectional communications leading organizations to have a pair of data diodes. With the advanced segmentation and access controls as described above, one can use the same device that is used to provide fine-grained micro-segmentation within the control systems to manage the perimeter.

4. Natively blocks known threats

Also part of the network security discussion is the capability to stop known threats. There is a large universe of known threats to ICS including:

- Exploits to ICS-specific products such as the controllers (PLCs, RTUs, IEDs) or SCADA software packages
- Exploits to IT products used in ICS such as Operating Systems, browsers, and even specific modules such as OpenSSL and the Unix BASH shell
- Protocol functions such as DNP3 warm restarts which although are normal features are so risky that they warrant treatment as exploits
- Run-of-the-mill viruses that even if introduced by accident could still take down hosts and cause downtime.
- Known bad domains/URLS used by malware for command and control and watering hole attacks

Despite knowledge of these risks, many organizations leave devices in IC unpatched and unprotected against these threats for extended periods. This could be because the product vendor has not yet made a patch available or because the operator needs to wait for a maintenance window before taking the device offline for patching. It is not uncommon for the systems to never be patched due to product end of life. It is therefore critical to have compensating protections for these devices while they are left vulnerable.

Preventing known threats is part and parcel of being a 21st century security platform. Unlike legacy systems where threat information is analyzed separately from application or user information, threat information must be analyzed at the same time as the application and user information not only to improve performance but also so that there is shared context between the repositories. Shared context allows increased intelligence in terms of recognizing the nature of a threat with respect to the originating application and users as well as the impacted assets and users. Furthermore, it allows much easier creation of policies which are simultaneously based on a more effective approach of allowing legitimate applications/protocols between security zone based on role then blocking known threats may be using that channel to propagate. Figure 5 below shows the concept of blocking the known threats which may have come in via whitelisted traffic.

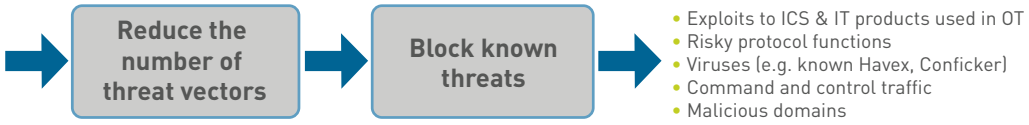

Figure 5: Block known threats natively after reducing potential attack vectors.

5. Detects and prevents attacks by unknown malware

The question then arises of how to deal with unknown malware that passes through the network. This is where the threat intelligence core comes in. A 21st century platform must be able to isolate suspicious network-borne payloads and send them to the threat intelligence core for rapid and automated analysis and dissemination of protections such as anti-virus, exploit, and command and control signatures. While solutions exist for this functionality as a standalone sandboxing appliance, this capability must be native to the platform so that the protections can be quickly and automatically provided to the enforcing device, the firewall. Detection is helpful, but there must be a closed loop to ensure prevention in a timely manner. Furthermore, because of the shared context, more intelligence could be collected in terms of the relationship of the Zero Day malware with the application and user information. When selecting your platform, be cautious of standalone sandboxing solutions that only tell you that you have a problem and do nothing in terms of providing you protections. Figure 6 shows the concept of adding the ability to quickly detect unknown threats and quickly converting them to known threats which can be stopped. Also for asset owners who might have sensitivity to sharing files outside of the organization, make sure to look for solutions that support local sandboxing and creation of protective signatures.

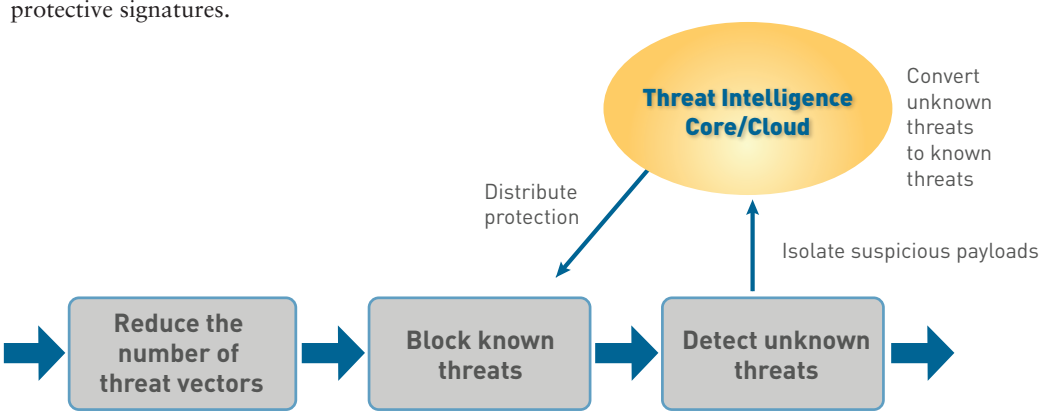

Figure 6: Native sandboxing capabilities to detect known threats and quickly convert them to preventable known threats.

6. Stops Zero Day attacks to the endpoints

Let's now shift our attention to the protection of hosts such as HMIs, automation servers, workstations, and even hosts belonging to managers/admins with privileged access. These systems run software which have remotely exploitable vulnerabilities. Furthermore, users can be tricked to run malware directly. Traditional endpoint security has been focused on looking at known malicious signatures, strings and behaviors to stop such endpoint-based attacks. This approach has been proven to be both ineffective and operationally burdensome. It only stops known threats and the moment you put in place a signature to block a known threat, several more new and then unstoppable exploits or malware variants crop up. Control systems endpoints are highly exposed not only to Zero Day attacks, but even known attacks for which the endpoints may not be protected given the long patching/update cycles which could span several months or even years.

So far the approach based on known signatures has been the only available option, but technologies have emerged to prevent even unknown attacks by stopping the underlying attack techniques used by exploits and malware. [8] This approach makes perfect sense given that the number of new exploit and malware techniques discovered each year are in the 2-4 and 10s to 100s range respectively. This is a much more manageable set versus the large universe of exploits and malware discovered each year. Most exploits and malware actually use more than one technique in their attack sequence, but one simply has to stop one of them to prevent the attack from completing. By focusing on stopping techniques rather than known signatures the endpoint security solution is more effective at preventing attacks. Such technology could also be used to validate installation packages from software vendors to ensure that there are no Trojan malware per attacks such as Energetic Bear. It can also be used to prevent and receive notifications for unauthorized installation of applications. These may be applications that actually have valid business use, but it allows organizations to have a more regimented and auditable approach to application deployment at endpoints.

In addition to utilizing this disruptive technique based approach to stopping Zero Day exploits and malware, the endpoint solution, as with the network security component, must also interact with the threat intelligence cloud to make use of and contribute to centralized and automated threat intelligence. Figure 7 below shows the concept of advanced endpoint protection and its interaction with the threat intelligence cloud.

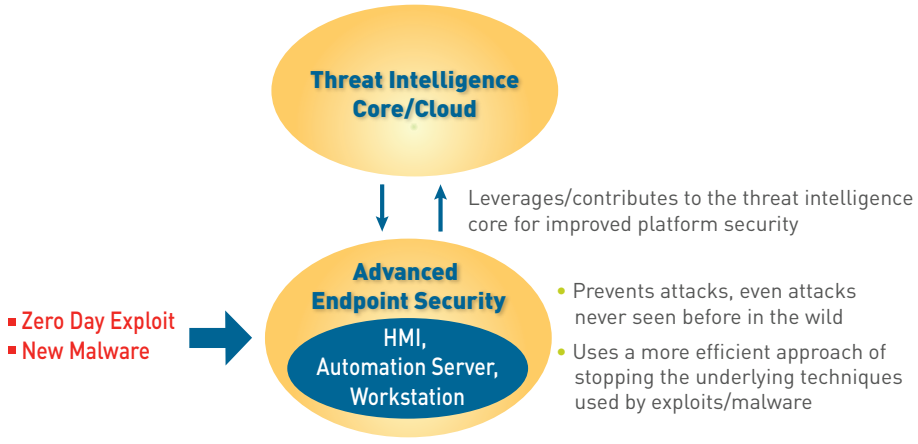

Figure 7: Advanced endpoint security stops attacks to host, interacts w/ threat intelligence cloud.

7. Provides central management and reporting

Given how highly distributed ICS tends to be whether across multiple plants on a campus or across multiple remote facilities such as substations or production facilities, a 21st century platform must provide a means for centrally managing the platform. Rather than having a separate central management device for application policy, threat prevention, URL and other functions, these should all be provided via a single management device. Furthermore, the platform must be able to efficiently aggregate the local information from each of the remote devices to create a consolidated view of the operations which in some cases may be a global operation. This capability helps dramatically in terms of being able to perform forensics and is indispensible when creating supporting documents required in regulatory audits.

8. Secures the use of mobility and virtualization technologies

While the use of mobile devices and virtualized operational datacenters are not yet very common and often intentionally avoided in ICS environments, there are already some front running organizations leveraging these technologies to further improve operational efficiencies and reduce costs. For example, some organizations have started to deploy mobile HMIs on tablet devices used in the field and on the factory floor. A next generation platform must be able to ensure that security policies are consistently enforced on such mobile devices to ensure safe use even outside the walls of the control center or plant.

Also, some organizations have started to consolidate physical servers running historian, SCADA, and other application servers onto several virtual machines residing on the same hypervisor. A great majority of ICS asset owners are hesitant to adopt such technologies and prefer fixed assets and non-virtualized servers. This is justified as there are new security considerations with virtualized environments such as securing east-west traffic between machines and ensuring consistency of security when moving virtual machines around. Whatever their current position on virtualization technology, organizations must consider that strong economic drivers often compel the migration to advanced technologies especially once the technologies are more proven. To set themselves up for the future, it would therefore be wise for organizations to select a security platform that also supports securing virtualized environments.

9. Powerful API and industry-standard management interfaces

A platform which has the above core features will cover a lot of the bases, but there needs to be a provision to accommodate additional products which may address unforeseen need or bring in additional value-add functionality. To achieve this, the platform must support industry-standard management interfaces and an open application programming interface (API). Together these capabilities allow integration with third party solutions which for example improve policy/configuration management, log analysis, reporting, and other important security functions. Security Information and Event Management (SIEM) devices are in particular very powerful platforms for aggregating data from many sources including network, server, database, and of course security.

Alignment with Industry Standards

Several cybersecurity standards specifically focused on critical infrastructure and industrial control systems have been developed in recent years. Some like NERC CIP and CFATS are regulated while others such as ISA 62443 and NIST Special Publication 800-82 serve as guidelines. The NIST Cybersecurity Framework (CSF) is a more recent standard that calls for mandatory compliance for US government agencies and serves as a best-practice reference for other critical infrastructure asset owners [9]. With a 21st century platform in place, organizations should be able to better address the requirements set forth by these standards and respond more efficiently during cyber incidents and compliance audits. While an exhaustive mapping of capabilities to the NIST CSF is beyond the scope of this paper, it useful to discuss at a high level how some of the capabilities described above map back to the NIST CSF's functional areas as follows:

NIST CSF Functional Area	Supporting Capabilities of the 21st Century Security Platform

IDENTIFY	▪ Identify network traffic and usage at very granular levels
	  ▪ Applications, ICS protocols, protocol functions
	  ▪ Users and user groups, IP Address, countries
	  ▪ Files, data strings, URLs, domains

PROTECT	▪ Reduce the number of attack vectors including applications, protocols, Domains/URLs, user, and other segmentation ▪ Protect unpatched systems from zero day exploits and never-before-seen malware ▪ Prevent malicious use of ICS protocols ▪ Secure mobile and virtualized environments ▪ Prevent data exfiltration
DETECT	▪ Detect unauthorized use whether malicious or non-malicious ▪ Decrypt encrypted traffic to identify stealthy malicious traffic ▪ Detect known threats and unknown threats which have never been seen before in the wild (IPS, AV, malicious domains/URL, command and control, "Son of Stuxnet" attacks) ▪ Detection can be performed at the network or at endpoints
RESPOND	▪ Shared context between application, user and threat/content information increases intelligence which simplifies forensics process ▪ Threat intelligence cloud provides automated threat analysis and protections for both endpoints and the network ▪ Integration with other security devices such as real-time SIEMs enriches the analytics
RECOVER	▪ Protections from threat intelligence cloud are automatically disseminated to endpoints to prevent attacks ▪ Knowledge of any impacted devices are provided back to centralized management and can be remediated ▪ Easy deployment of any additional policies/segmentation to improve security posture

Summary

In summary, the threat landscape to Industrial Control Systems has escalated to a point where legacy technologies are no longer effective in stopping the different types of cyberthreats. A new kind security platform—a 21st century security platform- is required to effectively combat cyberthreats and achieve the all-important goal of keeping the process available. The platform must combine the benefits of network security and endpoint security while leveraging a threat intelligence cloud to ensure attacks are prevented wherever they originate. Furthermore it must provide granular visibility and control at the application and user levels to allow for network segmentation that better aligns with the business needs. The platform must do more than just detect threats and attacks, it must prevent them, even attacks never before seen in the wild. The risks involved in securing critical infrastructure and manufacturing assets are just too high to allow anything but prevention as far as treatment of cyberattacks. Finally, the platform must be easy to deploy and maintain and should be able to interoperate with other security products. With all these in place, organizations will have the right set of capabilities required to secure modern day ICS while keep uptime and operational efficiency high.

References

[1] Report "2014 Survey on Industrial Control Systems Security", SANS Institute

[2] Article "Havex Hunts for ICS/SCADA Systems", F-Secure, http://www.f-secure.com/weblog/archives/00002718.html

[3] Article "Oil Industry Under Attack by Hackers", NewsInEnglish.no, http://www.newsinenglish.no/2014/08/27/oil-industry-under-attack-by-hackers/

[4] Report "Malicious Control System Cyber Security Attack Case Study—Maroochy Water Services, Australia", Joe Weiss and Marshall Abrams

[5] Video "Enterprise Security Platform", Palo Alto Networks https://www.paloaltonetworks.com/products/platforms.html

[6] Standard "ISA-62443-1-1 Security for Industrial Automation and Control Systems Models and Concepts", ISA99

[7] Report "Developing a Framework to Improve Critical Infrastructure Cybersecurity", Forrester Research

[8] Brief "Traps: Advanced Endpoint Protection", Palo Alto Networks, https://www.paloaltonetworks.com/content/dam/paloaltonetworks-com/en_US/assets/pdf/datasheets/Endpoint/endpoint-protection.pdf

[9] Standard "Framework for Improving Critical Infrastructure Cybersecurity", National Institute of Standards and Technology

Glossary

API	Application Programming Interface
APT	Advanced Persistent Threat
AV	Anti-Virus
CFATS	Chemical Facility Anti-Terrorism Standard
COTS	Commercial-Off-the-Shelf
DCS	Distributed Control Systems
DMZ	De-militarized Zone
ERP	Enterprise Resource Planning
HMI	Human-Machine Interface
ICS	Industrial Control Systems
IED	Intelligent Electronic Device
IP	Internet Protocol
IPS	Intrusion Prevention System
IIOT	Industrial Internet of Things
IT	Information Technology
NERC CIP	North American Electric Reliability Corporation Critical Infrastructure Protection
NIST CSF	National Institute of Standards and Technology Cybersecurity Framework
OPC	Open Platform Communication
OT	Operational Technology
PLC	Programmable Logic Controller
RDP	Remote Desktop Protocol
RPC	Remote Procedure Call
RTU	Remote Terminal Unit
SCADA	Supervisory Control and Data Acquisition
SIEM	Security Information and Event Management
SSH	Secure Shell
URL	Universal Resource Locator
USB	Universal Serial Bus

paloalto NETWORKS®
the enterprise security company™

4401 Great America Parkway
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com