Defense-in-depth, or:

how to secure Industrial Control System Critical Infrastructure

Author
Andreas Agostin

Industrial Network Sales Specialist
Eaton's Crouse-Hinds Business

23.02.2016

Introduction

Stuxnet was an eye opener for many: a malicious piece of software (malware) that infiltrated the Iranian Uranium Enrichment plant and causes physical damage to centrifuges by spinning the motors up and down repeatedly until they failed [1]. Shortly after the incident, several other malware were detected which remained hidden for a long time prior to their detection: Flame, Gauss, and Duqu [2].

Stuxnet made people aware that governments (even though it remains unproven which government was responsible) have the ability to infiltrate critical infrastructure and cause damage.

But what about the average plant? They are safe, right?

The average plant is far from safe. Even if nothing happened so far, the lack of security measures leaves the doors wide open. There are several databases collecting security incidents. RISI, the repository for Industrial Security Incidents, specializes on ICS. It lists over 240 incidents [3] for various industries. The VERIS community database [4] which started in 2013, lists more than 4700 incidents, out of which approximately 200 are industrial type [5].

It needs to be considered that such databases have to rely on publicly available reports. The number of unreported cases is unknown. Which company loves to read their name in news associated with negative headlines?

Incidents

Independent of which database you look at, a large number (in the magnitude of 80%) [3] are unintentional, and are caused by malware infection, software flaw, human error, and such.
Out of the 20% of incidents which are intentional, half of them are caused internally, for example by a disgruntled employee, the other half remotely.
While this probably is surprising for many, the damage that these incidents cause is probably even more surprising. While malware accounts for the majority of reported incidents, the damage caused by it is often minor. Sabotage and accidental incidents cause the more significant and sometimes devastating damage (see figure 1).

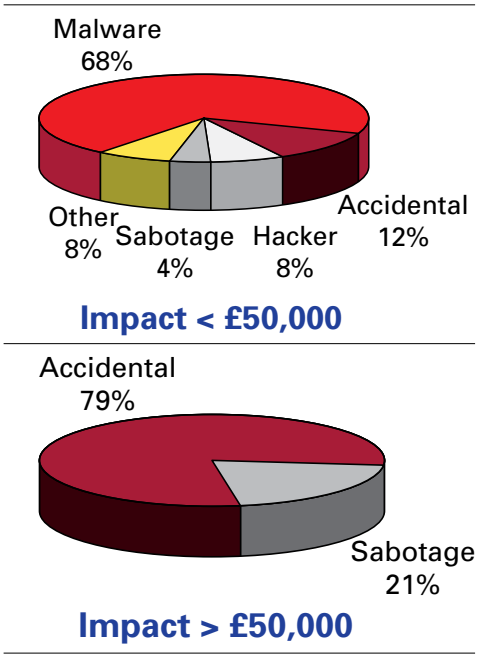

Figure 1: Analysis of industrial incidents for years 2002 to 2005 from the ISID (now RISI)

Powering Business Worldwide

Whether small damage but frequent nuisance, or rare occurrence with potentially devastating damage, it becomes important to prevent such incidents in the first place. Since we are dealing with software, the risk can probably not be eliminated, but at least significantly reduced.

I hear you: you have a firewall to the Internet. And some of you even have another one between the office LAN and the plant network.

That's a good start. But you will soon understand why this leaves your plant at great risk.

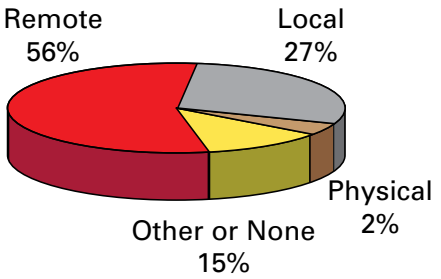

Figure 2: Analysis of industrial incidents for years 2002 to 2005 from the ISID (now RISI)

Risks

Looking at ISID data in a different way, we find that over 50% of incidents are initiated remotely.

In other words, in 44% of incidents, a firewall to the internet (or between office and plant network) is probably not of any help at all.

In case you are puzzled by the figure 15% for "Other or None": a typical representative of this category is malfunctioning equipment and equipment failure.

So what are the risks? There are various potential causes for an incident:

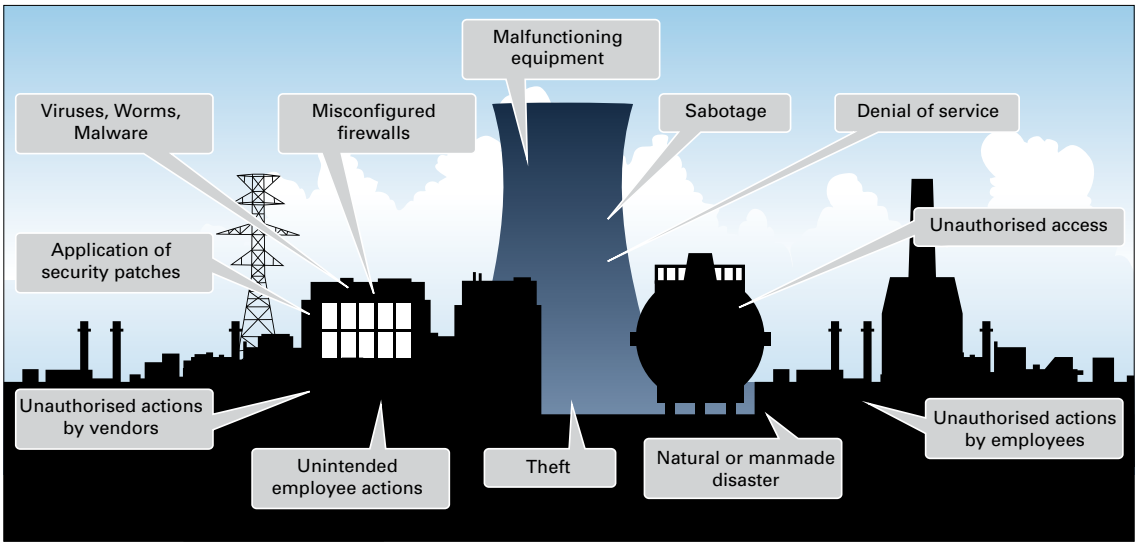

Figure 3. Possible causes for industrial incidents

There are also equally many potential points of entry, which makes it nearly impossible to define a clear perimeter:

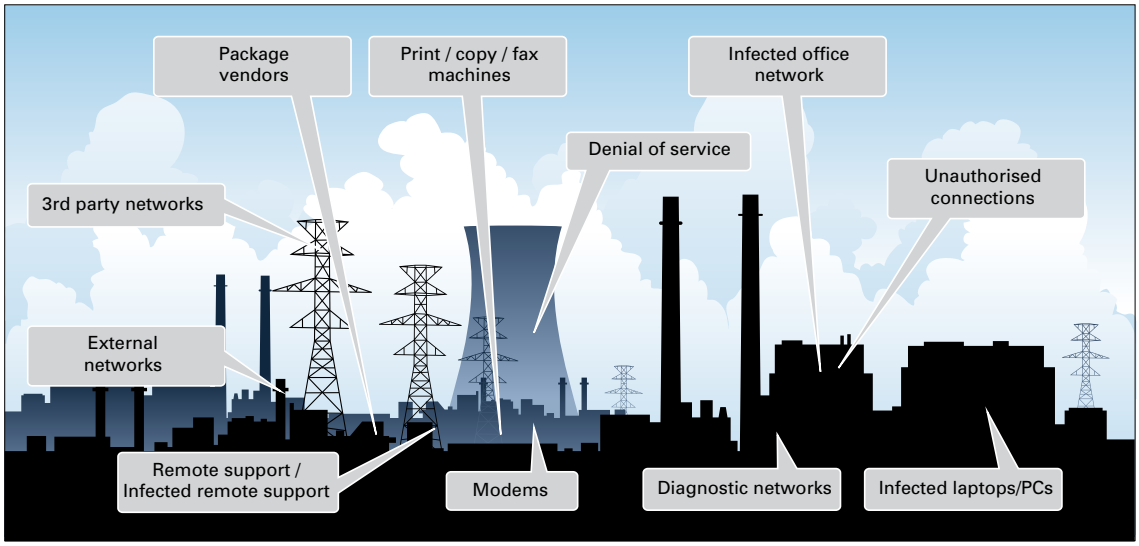

Figure 4. Possible points of entry

A plant operator may consider one or several of above possible points of entry as a particular threat. Considering that most damage is caused by accidents and sabotage, it will probably be difficult to predict where these incidents will happen.

A comprehensive approach is unavoidable. So what options are available?

The Three Approaches to Security

1. Air gap

An air gap is a complete isolation of one system from everything else. Obvious by definition, an air gap is most complicated in its implementation. How so?

First of all, it is highly difficult to keep systems isolated. At a hearing before the Subcommittee on National Security, Homeland Defense and Foreign Operations on May 25, 2011, Sean McGurk, Director National Cybersecurity & Communications Integration Center, U.S. Department of Homeland Security, said: *"In our experience, in conducting hundreds of vulnerability assessments in the private sector, in no case have we ever found the operations network, the SCADA system or energy management system separated from the Enterprise network. On average, we see 11 direct connections between those networks and in some extreme cases, we have identified up to 250 connections between the actual producing network and the enterprise environment." [6].*

Secondly, information from the plant floor is often needed by the management. How much is the plant output? What is our efficiency? How much losses is the unscheduled downtime costing us?

There is disagreement among experts whether air gaps are more harm than cure [7]. In theory, they can completely eliminate the risk of externally-caused incidents. In practice, the risk appears to even increase, as air gaps are frequently not implemented properly [6], and even if, they frequently result in the implementation of "sneakernets", unofficial (or even official) trafficking of data using a USB stick for example, which often bypasses all security measures in place [8].

2. Bastion model

Building a wall around the city core to protect the city against intruders was a frequently adopted security measure during the middle ages and beyond. The perimeter is clear: everything inside is protected from the outside.

Today, this model is implemented with firewalls, to protect the beautiful, sunny and friendly internal network from the ugly, dark and nasty internet. But is everything this black and white?

In August 2006, a malfunction in a redundant network caused both reactor recirculation pumps to fail. It is believed that a PLC malfunctioned and flooded the Ethernet with spurious traffic ("traffic storm"), disabling the VFD controllers controlling the pumps [9].

Firewalls are a frequent item of concern, too. A 2004 study revealed that 80% of firewalls are misconfigured [10]. Another study from 2012 revealed several "common firewall flaws" [11]:

- passing Microsoft Windows networking packets
- passing rservices (rlogin, rsh, and rexec)
- having trusted hosts on the business LAN
- Most common: not providing outbound data rules. This may allow an attacker who can sneak a payload onto any control system machine to call back out of the control system LAN to the business LAN or the Internet

With sufficient evidence available, the overall risk reduction achieved by the Bastion model is more than questionable.

3. Defense-in-depth

A security strategy based on layered defenses has proven significantly more effective, no matter whether for military, commercial or industrial applications. What makes the security measures more effective is the layered approach: if one layer fails, there is still another that can protect the property more firmly.

How many security measures are in a bank? Guards, security glass, steel – lots of it, safes, time locks, access control, and lots more. The advantages of adopting a layered defense are an enormous risk reduction of a cyber security incident.

So how do we layer the defenses?

Defense in depth is a concept standardized in ISA-99 (all new revisions will be released under the new name ISA-62443) and adopted in IEC 62443 in agreement with ISA [12]. Applying defense in depth means to divide a plant into Zones which represent functional units. The idea behind is: if one functional unit (Zone) is affected, it should not be possible to spread to other Zones.

Zones will be completely isolated from each other if they do not have to communicate with each other. Otherwise, they will be linked via "conduits". Such concept implements multiple layers of security, so that if one layer (or Zone) is penetrated (affected), the other layers (or Zones) remain unaffected.

The conduits linking Zones must fulfill certain security requirements. First and foremost, they are used only between communicating Zones. Zones that are not communicating with each other must be isolated from each other. Next, conduits must allow only the traffic required between these two Zones they are linking. Lastly, all requirements must be understood, and traffic must be reduced to the required minimum. Unused or unnecessary traffic must not be let through.

It appears that conduits can be easily implemented with firewalls, but after reading the section about the Bastion model, one might think that firewalls are not a good idea. The traffic through the conduit is however stripped to the minimum; in most cases, such as a PLC communicating with a remote I/O system, only one protocol will be required; in some rare cases, the number of protocols may extend to a handful. This reduces the risk of misconfigured firewalls significantly. And since the zones are small but plenty, the probability that a fault in one Zone spreads to another is significantly reduced as well.

So the remaining question is: how do we determine the Zones?

One could take the approach of segregating into "Office" Zone and "Plant" Zone, and, BANG, we are back to the Bastion model.

One could take the approach and put each and every piece of equipment into its own Zone. Maximum security, but is this efficient?

Let us take a gas-fired power plant as example as to how reasonably divide a plant into Zones.

Our power plant operates under the following assumptions:

- It has more than 1 turbine (the actual number does not really matter, as long as it is greater than 1)
- Each turbine has their own associated controller (e.g. PLC)

I suggest the following approach in order to determine the Zone: if there are two pieces of equipment, A and B (no matter what this is), and in case A fails, it does not matter whether B fails as well, then A and B can be in the same Zone. If A fails and B should not fail (under any circumstance), then A and B should be in separate Zones.

So allow me to apply this philosophy to the turbines:

- Turbine 1 would have one or several PLCs, some transmitters, gas analyzers, valves, etc. If any one piece of this set of equipment fails, it is very likely that the whole turbine cannot be operated any longer. As such, all related equipment should be in the same Zone.
- Turbine 2 however runs independently. If turbine 1 is down, it is not acceptable that Turbine 2 goes down as well. As such, segregate each turbine into an individual Zone.
- How about the HMI? If the HMI is down, you may not be able to visualize the process, but the process should probably continue nevertheless. If so, put it into its own Zone.
- Does Turbine 1 communicate with Turbine 2? No. So don't put them onto the same switch.
- Does Turbine 1 communicate with HMI? Yes. Put a conduit, consisting of a defense-in-depth firewall.
- Does Turbine 2 communicate with HMI? Yes. Put a conduit, consisting of a defense-in-depth firewall.

Continue this thought and segregation process until all equipment is distributed and Zones and Conduits are in place.

The Tofino defense-in-depth firewall

To implement Conduits, a large powerful firewall appears not to be the right product, particularly due to the tremendous costs associated with it.

A smaller, less powerful and hence less costly product appears more reasonable. On top of that, since the product would be distributed throughout the plant floor among the process control equipment, a product that does not require the skills of an IT expert and that understands the usual process control protocols would likely be beneficial.

Eaton's Tofino Security appliance is such a product. It is specifically designed to implement defense-in-depth, with an industrial design, redundant power feed, alarm contact, DIN rail mount and other typical industrial features.

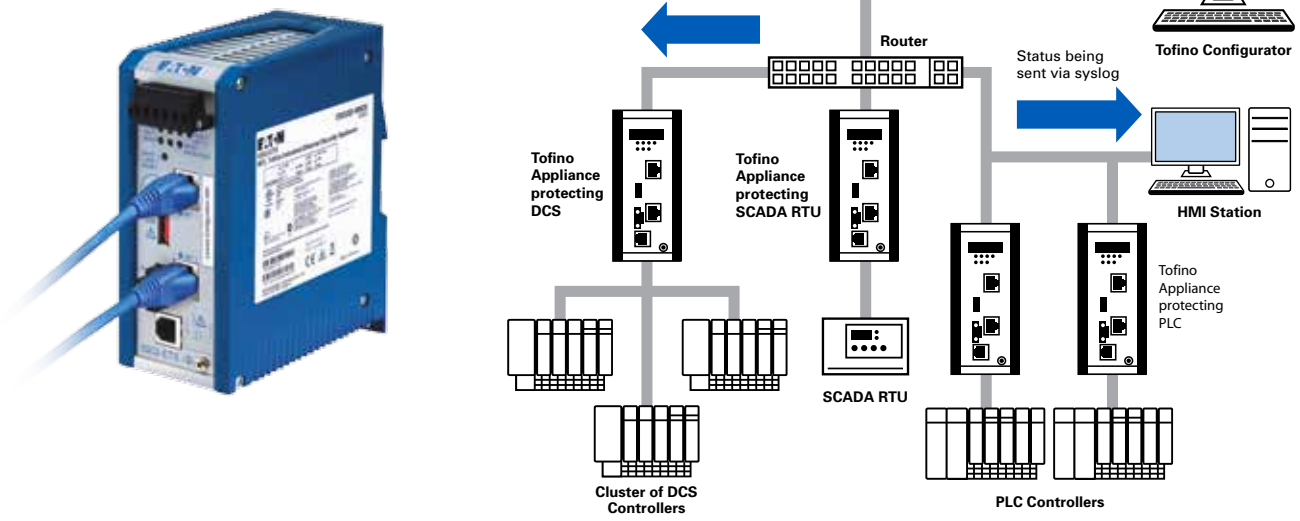

Figure 5. Implementing Defense-in-depth with Eaton's Tofino Security Appliance

Conclusion

Among the various approaches to Industrial Control System (ICS) security, the Bastion model using a central firewall provides only marginal risk reduction.

While air gaps theoretically provide a reasonable level of protection, they fail to achieve any significant risk reduction due to practical difficulties in its implementation.

Defense-in-depth is internationally standardized as IEC 62443 and the only concept that can achieve the maximum possible level of risk reduction, while minimizing potential flaws in its implementation at the same time. The level of risk reduction however highly depends on the level of segregation into Zones, and the proper implementation of defense-in-depth firewalls. In practice, the exercise is often implemented as a cost-reduction rather than security-enhancement measure.

References

[1] Source: http://www.jpost.com/IranianThreat/News/Article.aspx?id=199475, accessed 24.02.2016

[2] http://www.cnet.com/news/a-whos-who-of-mideast-targeted-malware/, accessed 24.02.2016

[3] RISI, http://www.risidata.com/Database, accessed 24.02.2016

[4] VERIS, https://github.com/vz-risk/VCDB, accessed 24.02.2016

[5] VERIS, http://vcdb.org/explore.html, accessed 24.02.2016

[6] Sean McGurk, Director National Cybersecurity & Communications Integration Center, U.S. Department of Homeland Security, at the Hearing before the Subcommittee on National Security, Homeland Defense and Foreign Operations , May 25, 2011, http://www.gpo.gov/fdsys/pkg/CHRG-112hhrg70676/html/CHRG-112hhrg70676.htm, accessed 24.02.2016

[7] #1 ICS and SCADA Security Myth: Protection by Air Gap, https://www.tofinosecurity.com/blog/1-ics-and-scada-security-myth-protection-air-gap, accessed 24.02.2016

[8] Schneier on Security, https://www.schneier.com/blog/archives/2013/10/air_gaps.html, accessed 24.02.2016

[9] Data storm caused nuclear plant shutdown, http://www.securityfocus.com/news/11465, accessed 24.02.2016

[10] Avishai Wool, "A quantitative study of firewall configuration errors", IEEE Computer Magazine, IEEE Computer Society, June 2004, http://csdl.computer.org/comp/mags/co/2004/12/rz074.pdf, accessed 24.08.2012

[11] US-CERT, http://www.us-cert.gov/control_systems/csvuls.html, accessed 24.08.2012

[12] IEC 62443, https://webstore.iec.ch/searchform&q=IEC%2062443, accessed 24.08.2012

Powering Business Worldwide

Eaton Electric Limited,
Great Marlings, Butterfield, Luton
Beds, LU2 8DL, UK.
Tel: + 44 (0)1582 723633 Fax: + 44 (0)1582 422283
E-mail: mtlenquiry@eaton.com
www.mtl-inst.com

Publication No. **WP Defense-in-depth 171016**
October 2016

EUROPE (EMEA):
+44 (0)1582 723633
mtlenquiry@eaton.com

THE AMERICAS:
+1 800 835 7075
mtl-us-inf o@eaton.com

ASIA-PACIFIC:
+65 6 645 9888
sales.mtlsing@eaton.com