Network Segregation

Introduction

In late 2014, a steel mill in Germany suffered massive damage as the result of a cyber attack that required advanced hacking skills, applied industrial control knowledge and endurance. No one has claimed responsibility for the attack and the lack of attribution and clear objective emphasize that not only are the threats very real, but anyone can become a target without apparent reason. This incident is only one of many examples of successful cyber attacks of industrial automation and control systems, and together with Stuxnet manifests the reality of cyber threats.

Unless you have experienced a serious security incident first hand, it is easy to believe such attacks only happens to someone else, but not having experienced such an incident however, does not mean you haven't been compromised. In fact, according to a report from KPMG from 2014, it is more likely than not that information is being exfiltrated by malware without your knowledge from your office networks. 14 companies were studied, and data was actively stolen from 10 of them without their knowledge.

The report is yet another strong indication traditional, best-practice defense like anti-virus, perimeter firewalls and network intrusion detection systems based on signatures are easily avoided. It also indicates insufficient organizational readiness as no action was taken even when malicious code was detected.

Westermo present a series of five basic applications assets owners can apply in their own networks to improve the security posture in a sustainable way.

Network Segregation

In large enterprise networks, it is common to create subnets to separate for example regular office assets from line-of-business (LOB) servers, and even separate networks for different geographic locations to reduce maintenance, increase availability and troubleshooting. It is however also common to still treat the internal networks as one single trusted zone where all assets are free to communicate with any other asset. To mitigate the risk of infection propagation, office assets are often equipped with host based network filters instead of introducing filters between the network segments.

Using WeOS switching routers, it is possible to add network filters in the segment boundaries, creating security zones with additional layers of protection and detection.

Protection

The fundamental idea with security zones is to ensure only the expected data flows between zones, which for industrial applications could be for example industrial protocols like MODBUS TCP, DNP3 or IEC 60870-5-104. Everything else should be filtered out reporting all known bad and unknown data to an intrusion detection system.

Using a highly simplified view of a water purification process illustrated in *Figure 1 - Network segregation into security zones*, as example, there are at least three obvious areas with assets with similar purpose, function and criticality; water disinfection (Zone A), water clarification (Zone B) and control system (Zone ICS). These zones can all be grouped into one bigger zone, the Production Zone, to create a layered zone model (onion) where one layer must be compromised before lower layers can be attacked from the network.

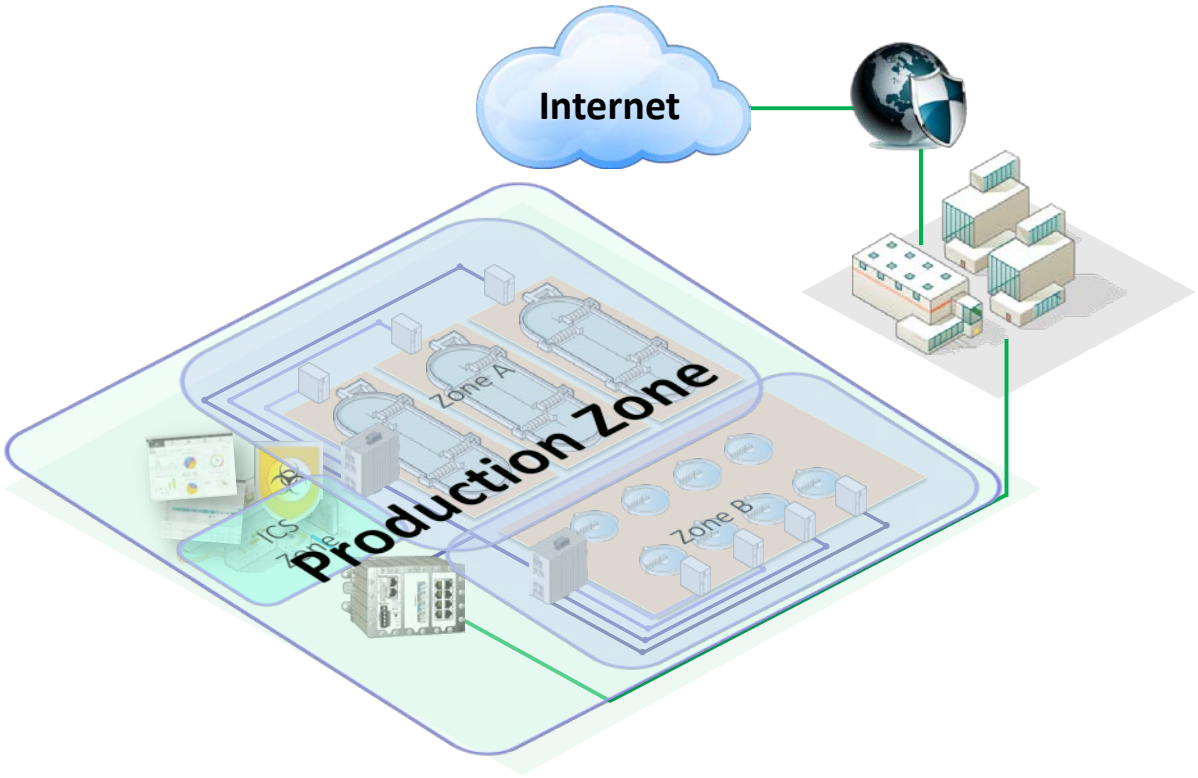

Figure 1 - Network segregation into security zones

Zones does not necessarily provide any protection by themselves, but by activating technologies like stateful packet inspection (what we call our firewall) and NAT, it is possible to limit the data flowing between zones.

There are major benefits with WeOS based devices;

- All WeOS devices have these capabilities
- Every access port on the device can be a separate security zone

- They work in resilient network architectures (ring topologies) as illustrated in *Figure 2 - Resilient network segregation*

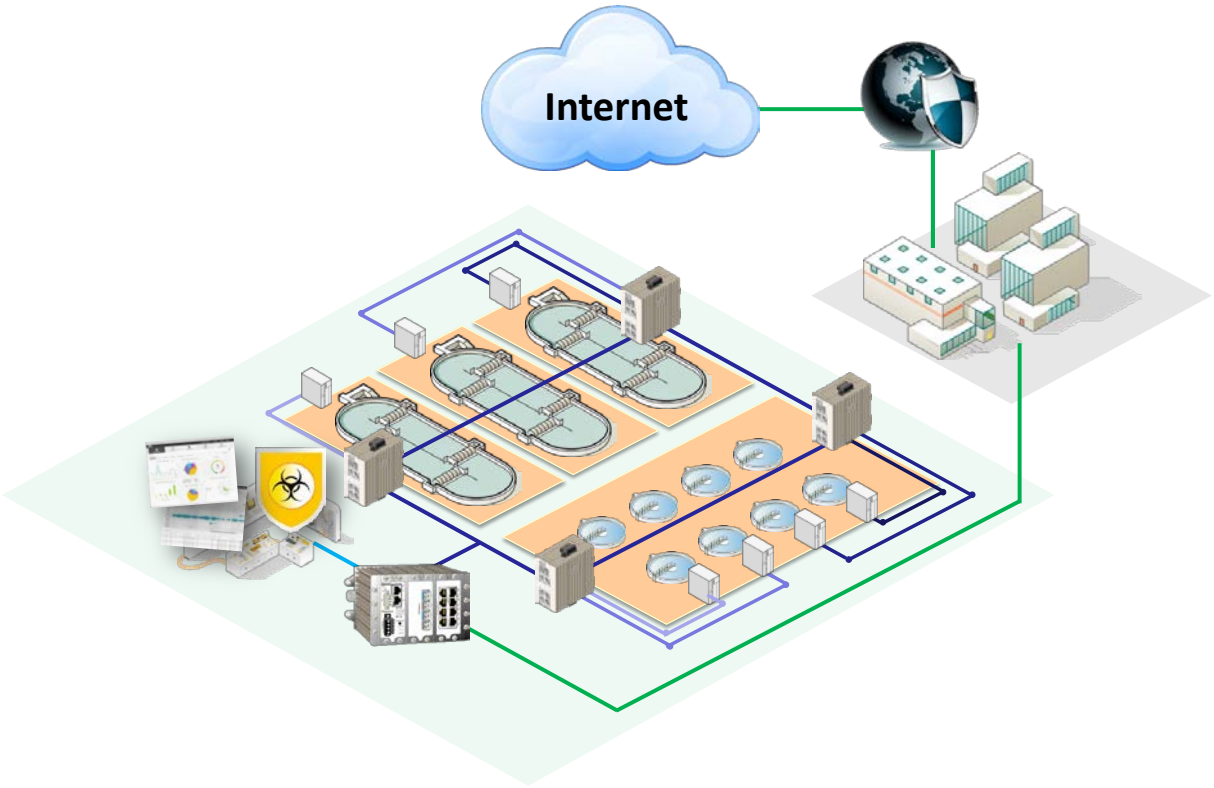

Figure 2 - Resilient network segregation

The example in *Figure 2 - Resilient network segregation* illustrate how each PLC becomes an individual security zone in resilient topology, making it possible to control all data flows in the network. With this technique, it is possible to establish IP level traffic whitelisting, meaning only approve data flows are allowed between assets (assuming there is only one asset per zone).

Detection

Most malware, including Stuxnet, access the network in distinct ways which makes it possible to detect their presence unless they are using covert channels (secretly using an already existing and approved channel). For example, knowing that only one asset in a system have a share printer connected makes it possible to detect all unsolicited requests to shared printers in other assets.

In the example above, it would be possible to detect malware calling home in the Production Zone perimeter as the malware is likely to make requests to unapproved resources, like DNS, external addresses or ports.