Visibility, Monitoring, and Critical Infrastructure Security

domaintools.com/resources/blog/visibility-monitoring-and-critical-infrastructure-security

Introduction

On 08 February 2021, officials from Pinellas County, Florida announced an unknown entity accessed water treatment operations for the city of Oldsmar. In addition to technical analysis based upon limited details, multiple media outlets responded to the incident with immediate reporting lacking significant additional details. At this time, while the general nature of this incident is somewhat known, many questions remain, especially concerning what entity was responsible for the incident and what their precise intentions were in attempting to modify water treatment operations.

While further investigation is warranted, available details allow us to reach some preliminary conclusions on the incident itself and its likely implications. Furthermore, based on what we know with respect to this event and past Industrial Control System (ICS) intrusions, we can formulate an understanding of this incident's maturity. Finally, the event in Oldsmar provides sufficient information to provide defensive guidance to detect, mitigate, or prevent similar scenarios in the future.

Overview of the Oldsmar Incident

On 05 February 2021, operators at the municipal water treatment facility serving the small city of Oldsmar, Florida noticed strange activity on the systems used to monitor and control operations at the plant. Initial reporting from Reuters indicated the facility used TeamViewer remote access software for remote monitoring and management, which was subsequently confirmed in follow-on reporting and interviews conducted by Wired. While plant operators have since removed the software, at the time an unknown entity identified the TeamViewer instance and managed to authenticate to the system.

Identification of critical infrastructure systems exposed to the internet is hardly a new phenomenon. As previously documented by Kim Zetter in several articles from the early 2010s, various tools exist that enable researchers (or less scrupulous entities) to search for and identify ICS devices or similar equipment. Prior to the Oldsmar incident, researchers identified several instances of likely malicious entities remotely accessing control system equipment in the water sector, with the following standing out as most interesting:

- **2013**: Intruders, assessed to be linked to Iranian entities, accessed control systems for the small Bowman Avenue Dam. Although a relatively minor structure and resulting in no disruptive consequences, analysts theorize the intruders may have intended to target the much larger Arthur R. Bowman Dam in Oregon, instead.
- **2016**: In its annual data breach report, Verizon reported several water utilities—combined into the single, pseudonymous organization "Kemuri Water"— experienced breaches of varying severity. In a few instances, the unknown intruders appear to have manipulated water treatment controls in a haphazard way, causing operational disruption but no harm or destruction.
- **2018**: An unknown entity utilized VPNFilter malware to attempt an unspecified attack on a chlorine production plant in Ukraine. Although not directly targeting water treatment operations, the incident would have significantly impacted the sector had a disruption occurred at the targeted site.
- **2020**: Unknown entities, although tentatively linked to Iranian interests, accessed and performed minor modifications to multiple water pumping and treatment devices in Israel in April and July 2020. Based on analysis of available information, affected devices were externally accessible with minimal or no authentication preventing the intruder from accessing control systems. While there is some speculation that this event may be linked to the Oldsmar incident, no evidence exists connecting the two and all similarities appear circumstantial.

In all of the above cases, external access to control systems resulted in either no or very limited disruption to physical operations. While the same is roughly true of the Oldsmar incident in terms of ultimate impact, the unknown intruder's actions in the environment are concerning. Specifically, the entity utilized remote access to ICS equipment to manipulate sodium hydroxide levels in the treatment plant. While normal operations run at 100 Parts Per Million (PPM) of sodium hydroxide in the treatment environment, the unknown intruder attempted to increase the amount to 11,100 PPM.

While attempting to make the above alteration, personnel monitoring the equipment, likely a Human Machine Interface (HMI) in the plant environment, noticed the action and reversed it. Had the action not been caught in progress during standard working hours, officials indicate it would take 24-36 hours for the change to be reflected downstream among the population served by the district, and that automated testing and similar safeguards would have detected the physical process change.

Based on all available evidence and statements from local authorities, the intrusion and manipulations to the ICS environment were prevented through operator attentiveness and interaction, with further engineering controls providing additional layers of safety. Although the incident resulted in neither significant disruption nor outright damage, the simple fact that some unknown entity attempted the above action is deeply concerning—reflecting either callousness given the potential harm, or ignorance as to what the attempted change might have produced in the serviced population.

Process Visibility and Intruder Maturity

Although deeply concerning, the intrusion (if not "attack") scenario described above shows multiple immaturities. Particularly:

1. Events took place during normal operational hours where personnel were on-hand and available to quickly respond.
2. The intruder did not attempt to hide or mask their activity through interaction with or overwrite of HMI systems or spoofing of sensor data.
3. The modification to sodium hydroxide levels was so extreme as to almost certainly trigger engineering or other non-ICS controls or alarms within the environment.

These three factors, taken together, indicate an attack that was either immature, rushed, or potentially unintentional following access to the controlling HMI. To better understand how these items function, especially in light of a potential integrity-targeting ICS incident, a quick review of historical ICS incidents is helpful.

Stuxnet Incident

Although well-documented, especially through resources such as Kim Zetter's Countdown to Zero Day and Symantec's Stuxnet Dossier, certain elements of Stuxnet are frequently misunderstood or overlooked in general discussion. While reviewed in other sources in depth, the critical item enabling Stuxnet's success was the malware's ability to induce a general loss or denial of view condition in the victim environment. In this specific case, the malware recorded "normal" plant operations then played these recordings back to monitoring systems during physical attack sequences to mask events from plant operators. Absent this critical step, operators would have been able to detect anomalous operations in the plant environment enabling intervention and process diagnosis.

2015 Ukraine Power Event

In 2015, entities linked to Russian military intelligence (GRU) penetrated multiple electric distribution centers in Ukraine. In a coordinated operation in late December of that year, the attackers disrupted distribution operations inducing blackout conditions through a combination of rogue control devices and logging on to user workstations to disconnect equipment.

Yet for such operations to succeed, personnel had to be locked out of their workstations to prevent operator intervention during the initial phases of the attack. Subsequent activity in the victim environment resulted in the use of wiper malware to remove remote operational control, followed by a malicious firmware update to serial-to-ethernet converters which made communication to equipment impossible. Overall, these steps amount to a coordinated effort to induce a loss or denial of control condition that enabled a sustained, widespread impact to Ukrainian electric utility operations.

2016 Ukraine Industroyer/CRASHOVERRIDE Incident

In 2016, Ukraine again witnessed an electric power incident in December linked to Russia's intelligence services, this time targeting a single transmission substation. Referred to as the Industroyer or CRASHOVERRIDE event, the incident again wiped control systems to induce loss of control—but also likely aimed at a loss of view condition as well to enable a potentially destructive (if failed) physical damage scenario. In this particular case, removing operator logical control (to force manual operations) combined with loss of logical view into the health and status of the system was used in sequence to enable a process protection-focused attack scenario. Absent these conditions, it would be highly unlikely for the sequence of events required to restore operations in an unprotected, unsafe state (enabling possible destruction) would materialize.

2017 TRITON/TRISIS Event

In 2017, a petrochemical plant in Saudi Arabia experienced multiple unexpected plant shutdowns due to the plant's Safety Instrumented Systems (SIS) tripping for then-unknown reasons. Subsequent investigation identified a purpose-built malware variant, referred to as TRITON or TRISIS, as responsible for the disruption. Further investigation and analysis indicated that rather than a direct attack on plant safety equipment, the malware's purpose was to enable undetected, arbitrary modification of SIS parameters. Combined with access elsewhere in the plant environment, an attacker could remove or alter safety controls to induce physical damage. Yet to succeed, the attacker needed to ensure not only access to modify safety parameters, but also the ability to alter such parameters without operators knowing such changes took place.

Implications for Oldsmar

Overall, these four examples of high-profile, technically complex ICS attack scenarios emphasize a critical barrier to adversary success: the ability to evade, influence, or outright deny operator visibility into and control over ICS environments. In all four examples, the attacks required some mechanism to hide from operators or deny their ability to correct or mitigate changes made to operating parameters.

In the case of the Oldsmar treatment plant incident, the intruder failed to attempt any such action based on information currently available. Had the unknown entity spoofed or otherwise interfered with HMI display parameters or sensor data, the operator on duty would be less likely to notice the incident as it took place, resulting in an attack moving on to engineering and process controls for potential mitigation or detection. Not only did the intruder fail to limit or manipulate process view in the environment, they executed the event during primary working hours on a weekday, almost ensuring that such activity would be quickly noticed (and mitigated).

Based on these observations and in light of past ICS incidents, we can therefore make a reasonably confident claim with available evidence that this was not an especially complex or savvy “attack”. As described in multiple sources, the intruder appears to have merely taken advantage of weakly secured, accessible remote access mechanisms to connect to plant equipment controls, followed by either deliberate or potentially inadvertent manipulation of the environment. That such an attempt occurred at all is certainly concerning, but the overwhelming evidence given event timing and execution indicate that there were only slight possibilities for this event to produce significant damage or harm.

Defensive Countermeasures and Attack Surface Reduction

While this particular incident did not result in any damage or even notable operational disruption, events at the Oldsmar water treatment facility highlight the real risks and dangers associated with remote access to critical infrastructure systems. While the knee-jerk answer to such issues would be to remove or curtail access to such systems as much as possible, this is unrealistic and inactionable in modern operational environments. For reasons ranging from centralized control over geographically distant control systems to vendor requirements to system access for telemetry and maintenance purposes to the limitations placed on personnel by COVID-19 restrictions, remote accessibility cannot simply be removed or shut off. However, the ways in which such connectivity are implemented can ensure that such operations are done in a securable, defensible fashion.

First and foremost, while precise details on the system in question are not available, available evidence indicates that if the HMI controlling sodium hydroxide levels was not directly accessible via the TeamViewer instance at the facility, then access to such a system was easily gained from that initial access point. While such direct access may be convenient, it is insecure and undesirable. Instead, having a purpose-built bastion or “jumphost” can provide a single, hardened point for remote access and monitoring. By using

different sets of credentials for the bastion host to internal network and system authentication, security can be increased further as password brute forcing or credential capture for the bastion will not enable immediate follow-on access to other systems in the network.

Network segmentation, access controls, and sound network engineering can work in concert to reduce the overall attack surface to a limited number of defensible nodes (such as the bastion), while also facilitating monitoring of activity to a smaller set of devices. When applied with even more robust security controls, such as the implementation of robust Multi-Factor Authentication (MFA) schema for remote login activity, exposed attack surface can be reduced even further.

While certain controls such as complex passwords or hardware token MFA may be undesirable in certain ICS environments due to operational overhead and similar considerations, applying such defensive measures and similar controls to external facing network access points is critical. The profusion of scanning and indexing tools for remotely accessible services combined with the ability of adversaries to either brute force or potentially capture user credentials mean these accessible systems must be hardened.

Remote Access Monitoring and Indicator Enrichment

Once networks are appropriately hardened and segmented, network operators and defenders can then proceed to Network Security Monitoring (NSM) and traffic analysis. In well-designed environments with only a few externally-communicating or -accessible bastions, NSM operations are simplified and manageable, allowing for potentially powerful security analysis and response.

With proper monitoring, defenders can begin asking a number of questions or formulating hunting hypotheses relative to their network posture. Particularly, operators can utilize near real-time enrichment of network observables such as IP addresses and domain names to build an intelligence picture of traffic flows and communications.

At the most basic level, this may constitute little more than geographic enrichment of IP addresses to identify odd or anomalous connections. This methodology can certainly produce errors—such as the case of an Illinois water pump station where a worker on vacation remotely accessed a device prompting calls of Russian critical infrastructure hacking. Typically, observing remote authentication attempts from non-local address space for a municipal utility network can form a good starting point for follow-on investigation.

In more advanced use-cases, enrichment of external infrastructure using third-party intelligence sources can enable complex, higher-confidence queries for suspicious or known malicious activity. Examples would include being able to correlate source infrastructure for connection attempts to anonymizing infrastructure such as TOR, or

identifying connections from Internet Service Providers (ISPs) or Autonomous System Numbers (ASNs) either highly correlated with known-malicious activity, or never previously associated with any known-good, legitimate operation.

As documented in previous items from DomainTools researchers, identifying indicators in general, and network indicators in particular, as composite objects yields a number of possibilities for enrichment and analysis. In this particular case, defenders and network operators can leverage enrichment of network data to identify suspicious remote access activity or other signs of initial intrusion, potentially enabling operator response and mitigation actions before such intrusions proceed to the manipulation of controls in a critical infrastructure environment. From the opposite view, and as previously documented in the context of the SUNBURST campaign, near real-time enrichment of outbound traffic can potentially identify stealthy, otherwise difficult to detect intrusion scenarios by flagging items related to Command and Control (C2) activity.

Conclusion

The Oldsmar water treatment plant intrusion raises many concerns: first, that such an intrusion even happened in the first place indicates a certain maliciousness or lack of caution by the entity in question; second, while this specific instance appears relatively simplistic and immature, the same initial access vectors used in this event could be leveraged by more operationally savvy entities to produce a disruptive or dangerous impact. Yet while the incident is concerning, network operators and defenders have many options available to fight back against such events. Through a combination of network hardening, attack surface reduction, network segmentation, and NSM with indicator enrichment, defenders can dramatically reduce the likelihood of such events, significantly reduce their efficacy, or increase the likelihood of identifying such activity at relatively early stages.

Of course, for many smaller organizations, such as municipal water treatment entities, some of the security suggestions offered above may remain out of reach for budgetary or technical reasons. While a “100% solution” may not be feasible for such entities, certain preliminary steps such as attack surface analysis and reduction remain within reach. Irrespective of maturity and capability, given that various entities are actively probing and attempting to interact with connected critical infrastructure systems, all organizations responsible for operating such equipment must take all available steps within reason to secure these networks as best as possible.

Overall, and as documented by various entities, malicious activity against critical infrastructure networks in general and ICS networks in particular appear to be increasing. Through a combination of proactive defensive hardening and enhanced visibility or network monitoring, asset owners, operators, and defenders can better position themselves against

such intrusions. Failure to apply these lessons now mean potential adversaries at varying levels of complexity and capability will continue to find vital networks unprepared for future intrusion scenarios.