



Industrial Defender, Inc.
16 Chestnut Street · Suite 300
Foxborough · MA · USA · 02035
T: +1-508-718-6700
F: +1-508-718-6701

Experience with Network Anomaly Detection on Industrial Networks

Abstract

Industrial networks are routinely described as smaller and more stable than typical enterprise networks and so should be good candidates for anomaly-based intrusion detection. This article describes simple network anomaly detectors attached to industrial networks. The anomaly detector looks only at IP addresses and TCP/UDP ports. At each site, a manual calibration/learning process is undertaken to identify network traffic that complies with the site's security policies. All other traffic triggers alarms.

This article summarizes experience with anomaly detection at a number of sites. A surprising finding is the degree of customer interest in understanding and reviewing traffic flows identified by the anomaly detection system. Many control systems are small enough to make such manual review feasible with the anomaly detection tool. The article describes traffic which surprised site personnel, and describes remediations which were initiated as a result of the observed traffic. The article concludes that some industrial networks described are large enough to be difficult to characterize manually. Some automatic learning or characterization is desirable for such networks, but only if such characterization is amenable to manual review.

Andrew Ginter, ISP, CIPS, CISSP
Chief Security Officer
Industrial Defender

Based on a presentation at the Department of Homeland Security
Industrial Control Systems Joint Working Group (ICSJWG) Spring Meeting
San Antonio, Texas, USA

April 7, 2010

Table of Contents

Abstract.....	1
1: Introduction.....	3
1.1: Whitelisting vs. Blacklisting.....	3
1.2: Anomaly Detection and Control Systems.....	3
1.3: Experience with Anomaly Detection.....	4
2: Virtual Segmentation	4
3: Host Intrusion Detection Agents.....	5
4: Firewall Retrofits.....	5
5: The Anomaly Detector.....	7
6: Experience with Real Networks.....	9
6.1: SP99 Level 3/Level 4 Firewalls	9
6.2: SP99 Level2/Level 3 Firewalls	9
6.3: Anomaly Detection within Control Networks	9
7: Security Incidents and Remediations	11
7.1: Windows XP Automatic Updates.....	11
7.2: Network Driver “Phoning Home”.....	11
7.3: Corporate IT Scanning Control Equipment	11
7.4: Unauthorized Historian Connections.....	11
8: Conclusions	12
References	13
About the Author.....	13
About Industrial Defender.....	13

1: Introduction

Anomaly detection is a network intrusion detection technique in which expected or approved network traffic is either “learned” or described by rules, and unexpected or unapproved network traffic triggers alerts. Anomaly detection technology typically tracks characteristics of communications which include:

- communications protocols,
- source & target IP addresses,
- source & target TCP and UDP ports,
- the direction of initiation of communications,
- the volume of different kinds of traffic, and
- packet formats and content.

In addition, traffic characteristics may be charted against time, time of day, and calendar effects like weekends and holidays. Advanced products may look deeper into traffic streams and track the characteristics of different kinds of traffic payloads. COTS anomaly detection products include the CISCO Anomaly Guard Module^[1] and the IBM Proventia Anomaly Detection System^[2].

1.1: Whitelisting vs. Blacklisting

Anomaly detection can be thought of as “network whitelisting”. Conventional “blacklisting” network intrusion detection products such as those bundled with modern firewalls or unified threat managers have rules or “signatures” describing dangerous or unwanted network traffic. When traffic matching a rule is detected, an intrusion detection system raises an alert and an intrusion prevention system prevents the traffic from passing into another part of the network. In contrast, network anomaly detection products describe normal or allowed traffic and raise an alert when traffic deviates significantly from the normal profile.

Just as blacklisting network technology has both detection and prevention aspects, whitelisting anomaly detection can have both detection and prevention aspects. Conventional firewalls provide a simple kind of whitelisting / anomaly-based intrusion prevention. “Allow” rules in a firewall can describe allowed traffic, and a “deny all” rule prevents traffic from passing through the firewall when that traffic has not matched a specific “allow” rule.

Whitelisting technologies, be they network-based or host-based, have advantages over blacklisting technologies. Whitelisting technologies address high volume, low volume and targeted threats equally effectively, because they do not depend on a third party lab or vendor having observed an attack and created a signature for it. Whitelisting technologies address zero-day and zoo-list attacks equally effectively, again because there is no dependence on signatures. But of course, whitelisting technologies are no silver bullet – whitelisting is blind to malicious content that is allowed by one or more whitelisting rules.

1.2: Anomaly Detection and Control Systems

Network anomaly detection has been criticized in enterprise networks because of the great variety of allowed traffic in most enterprises. Such variety makes rules-based or learning-based characterization of “normal” or “allowed” traffic difficult, and leads to false positives. Further, traffic patterns in an enterprise normally evolve fairly rapidly as a result of changing technologies and changing business needs. This makes a “continuous learning” mode the only practical mode for most networks, because of the cost and complexity of trying to track changing traffic patterns with a process of manually defining rules. But continuous learning has been shown to be vulnerable to

“stealth” attacks, which change normal communications patterns so slowly that they are mistaken for natural change by the traffic learning algorithm.

In contrast, control systems tend to be simpler than enterprise networks and tend to have traffic patterns that change much more slowly than do communications patterns in enterprise networks. As a result, network anomaly detection components have been described as good fits for industrial networks^[3], but little has been reported about the application of such technologies to live industrial networks.

1.3: Experience with Anomaly Detection

This paper describes experience with anomaly detection on industrial networks. Given that control networks are simpler than enterprise networks, we at Industrial Defender thought we'd apply a simple anomaly detector to a number of live industrial networks to see if those networks were simple enough to make it practical to manually characterize high level communications patterns using whitelisting rules. We planned to review the characterization with the operations team responsible for the network and see if unauthorized communications could be observed in the ruleset.

We also planned to review with our security services personnel their experience with anomaly detection and with firewall-based whitelisting and host intrusion detection (HIDS) calibration. Firewall-based whitelisting is a technique used by Industrial Defender personnel routinely, to retrofit firewall technology into existing industrial networks. The retrofit methodology uses firewall rules and session logs as a simple form of anomaly detection. In addition, HIDS agents on control system hosts report network communications and support “accept” type rules in a way that yields information equivalent to a simple network anomaly detector as well. From all these experiences, this paper describes security findings and the ease/difficulty of carrying out such traffic characterization and rules creation.

2: Virtual Segmentation

An early, well-documented example of simple anomaly detection techniques is the section on “Policy Based IDS” in the “Snort 2.0: Intrusion Detection”^[4] text. Most of the text describes the most common use of Snort – defining rules which describe dangerous packets and which raise alerts when those packets are detected. The policy-based section describes a technique of defining rules which describe allowed traffic and raise alerts on all other traffic. For example:

```
pass udp 192.168.31.2 53 -> any any          # DNS server
pass udp 192.168.31.2 any -> any 53
pass udp 192.168.31.0/24 any -> 192.168.31.2 53 # DNS clients
pass tcp 192.168.31.0/24 any -> any 443      # HTTP/S clients
pass tcp 192.168.31.0/24 any -> any 80
Alert any any any -> any any (msg:"unauthorized traffic")
```

will raise an alert for every packet Snort sees that does not match the “pass” rules. This is noisier than is optimal – ideally operations administrators would like a small number of notifications of each kind of unauthorized traffic.

This approach to anomaly detection is used infrequently by Industrial Defender services personnel as a “virtual firewall”. When a set of operational equipment is very sensitive, but the process and control system are under very tight change control, it may not be feasible to insert a conventional firewall into the network to control communications to the sensitive equipment. In

this case, Industrial Defender recommends a “virtual firewall” - a passive intrusion detection Network Sensor configured as above, to alert on unauthorized communication with the sensitive equipment.

In practice our teams encounter such circumstances, but not often.

3: Host Intrusion Detection Agents

The Industrial Defender product suite includes software agents that run on industrial hosts and report network communications patterns employed by those hosts. The agents report network sessions – protocol, IP addresses and ports – unless an “ignore” rule exists matching the session. The session logs convey similar information to the output of the simple anomaly detection script described below, but of course the agent has access only to communications involving the host on which the agent runs.

The Industrial Defender Security Event Management system (SEM) makes it possible to capture groups of rules describing different kinds of hosts and applications' communication patterns. Industrial Defender services personnel report that when such rule sets exist already for the types of applications found on an industrial network, rule creation proceeds very quickly, since the session/alert logs contain only a comparatively small number of new kinds of communications to investigate and verify. When such rule sets do not yet exist, creating a set of rules for a new kind of host or application can take several hours per rules set.

Industrial Defender customers report that characterizing communications using the HIDS tools, and monitoring hosts over time with such tools routinely reveals unauthorized communications patterns.

4: Firewall Retrofits

Industrial Defender services personnel routinely retrofit firewalls into already-deployed, continuously-running industrial networks, without shutting down the control systems which rely on those networks. The technique they use for creating firewall rules is equivalent to a simple anomaly detector.

The existing network tends to be a mix of SP-99 level 2 and level 3 security zones, all on the same IP address range and sometimes on the same switch.

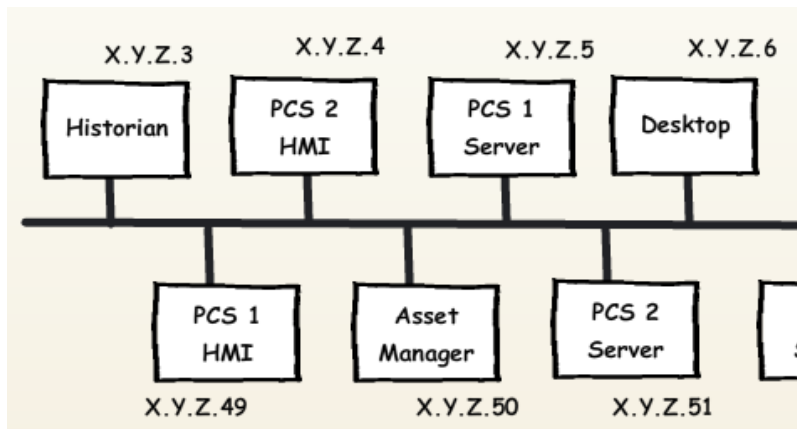


Figure 1: Original Network Architecture

The target architecture generally one closer to an SP-99 model, with separate level 2 control systems each in their own security zone, and the level 3 equipment in its own security zone.

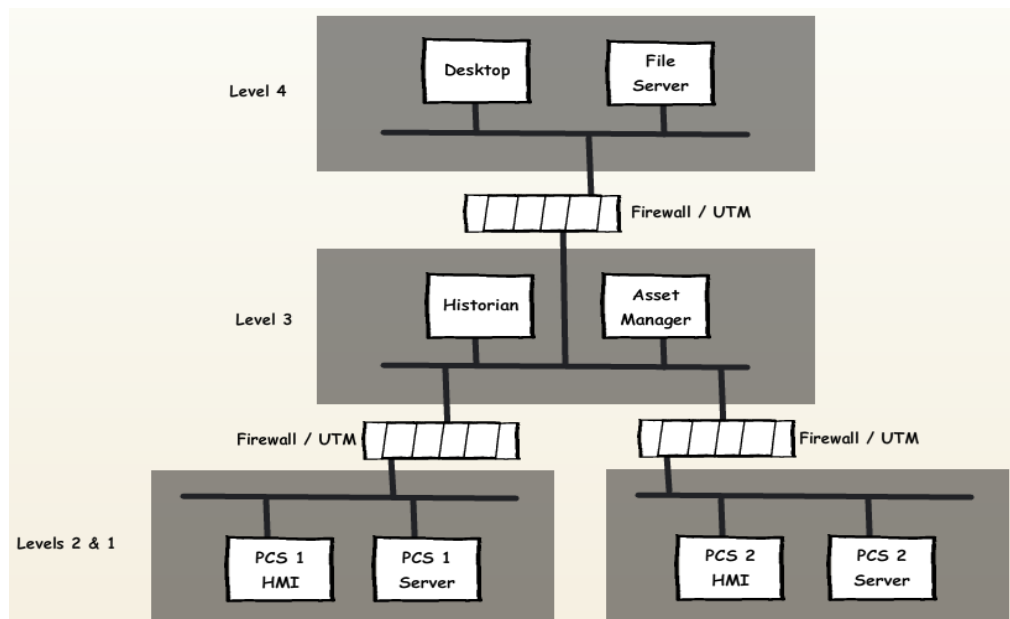


Figure 2: Desired Network Architecture

The problem is that static IP addresses are often more or less randomly assigned amongst existing equipment. For example, a desktop workstation and a control system server may have adjacent IP addresses. This makes it impossible to use a conventional “layer 3” firewall/router to segment the network, since such equipment needs easily identifiable subnets of IP addresses routeable via each network interface.

To segment this homogenous network then, our people use a security appliance in “transparent” mode, which other vendors may call “bridging” or “layer 2 routing” mode. The equipment on the network is gradually redistributed between layer 2 switches or VLANs so that each switch or VLAN is connected to only one security zone. The switches or VLANs for different zones are themselves connected by transparent-mode security appliances, initially configured to themselves act as switches.

To configure the firewall to act as a switch, our people add an “accept all” rule for every pair of interfaces. This allows traffic to flow freely between interfaces, and the interfaces auto-detect what IP addresses are on which interfaces. Once the physical or VLAN partitioning into zones and switches is complete, session logging is enabled on each of these “accept all” rules. After a period of time – typically one or two days – the session log is reviewed.

The session log looks very much like the output of the simple anomaly detector described in the next section, the two main differences being that:

- the session log records less data than an anomaly detector, because the “accept” rule sees only data passing between zones, and
- the session log is rather noisier than the anomaly detector about the data that it does record because each session is logged. There is no “compression” of output by logging each kind of connection only once, or by flagging server ports and summarizing future connections to those ports as “CLIENT” connections.

The session logs are reviewed with site personnel to ensure that the existing communications flows are in fact authorized and understood, and rules are created for such authorized flows. As

each “accept” rule goes in, the volume of new session logs decreases, since session logging is not enabled on the new “accept” rules.

After enough “accept” rules have been defined, the remaining session logs show only unauthorized or undesirable communications. The rules on the new firewalls are compared to rules that have proven sufficient to put a similar control system test bed through enough testing to be confident of the firewall not impairing the correct operation of the control system. At this point the site runs the new firewalls for a period of time long enough to be confident that they have been configured correctly. Anomalous communications show up in the “accept all” session log. Once the site is confident of correct operation, the “accept all” can be replaced with a “deny all” rule. The individual whitelisting rules continue to allow legitimate traffic to pass, and the “deny all” default rule segments the security zones.

5: The Anomaly Detector

The anomaly detector prototype is illustrated in Figure 1. The detector is a short script that uses tcpdump under the hood to capture packets. The detector is designed for deployment on Industrial Defender Network Sensors, but should work on any system with tcpdump, bash and tclsh. As the script uses tcpdump, it can either receive packets from a live network interface, or replay packets from a packet capture (pcap) file. The tcpdump command captures only packet header information.




Figure 3: Anomaly detector architecture

The session cache module turns the packet stream into a stream of sessions with unique combinations of:

- protocol (TCP, UDP or ICMP),
- endpoint IP addresses, and
- TCP or UDP ports associated with the endpoint IP addresses.



The pattern matcher then applies star-name match patterns to the stream of sessions. If a session matches a rule, the session is suppressed. If there is no match on a rule, an alert is issued for the session. A sample of rules follows:

```
# High-volume connections
A udp 192.168.31.2 53 *.*.*.* # DNS server
A udp 192.168.31.2 *.*.*.* 53
A udp 192.168.31.*.* 192.168.31.* 53 # DNS clients
A tcp 192.168.31.*.*.*.*.* 443 # HTTP/S comms with world
A tcp 192.168.31.*.*.*.*.* 80
```

The anomaly detector running on a network sensor would typically have access to either a mirror port on one or more switches, or to a tap replicating traffic across one or more important points in the network.

A sample of anomaly output looks like¹:

```
192.168.31.191:50799 17.8.8.2:993 tcp (/imaps)
239.255.255.250:1900 192.168.31.8:1024 udp (/)
192.168.31.198:50114 173.8.8.12:993 tcp (/imaps)
192.168.91.31:58683 192.168.31.39:80 tcp (/www http)
192.168.91.31:58684 192.168.31.39:80 tcp (/www http)
192.168.91.31:CLIENT 192.168.31.39:80 tcp (/www http)
192.168.95.11:34840 192.168.31.53:443 tcp (/https)
192.168.95.11:34841 192.168.31.53:443 tcp (/https)
192.168.95.11:CLIENT 192.168.31.53:443 tcp (/https)
```

The detector does not track which IP addresses initiated a connection. The detector prints the IP address with the highest-numbered port first. Where a port is well known in the `/etc/services` file, the name of the port is shown on the right of the output. When three different high-numbered ports on one machine communicate with the same port number on another machine, a CLIENT indicator is printed on the left and no further sessions for that pair of IP addresses and the right-hand port are printed, since client side port numbers are semi-randomly assigned and convey little useful knowledge to a person reading the output.

The rules and output sample shown above were taken from an enterprise network running in a small branch office. The security policies for the office were “relaxed” by industrial standards. Employees in the office had a legitimate need to routinely visit a wide variety of internet websites and they used a variety of messaging techniques to communicate with customers all over the world. These employees also had legitimate need to routinely stage and tear down internal web servers for testing purposes. As a result, many anomaly rules were phrased as “*.*.*.*” indicating that any IP address was a legitimate target of common internet communication protocols. In short, using the simple anomaly detector to scan the communications for IP address/port anomalies revealed very little because pretty much all communications were authorized. To detect unauthorized communications would have required digging deeper, as commercial anomaly detectors do, into black-listed IP addresses, payloads, traffic volume patterns and so on.

1 Names and addresses have been changed

6: Experience with Real Networks

6.1: SP99 Level 3/Level 4 Firewalls

Industrial Defender services personnel report that firewalls retrofit between the SP-99 level 3 networks and the level 4 enterprise network tend to be straightforward to configure via session discovery. There tends to be very few kinds of communications allowed across such firewalls and discovering and interpreting those communications is straightforward. Furthermore, most control systems can run correctly for many hours if not much longer without any communications with enterprise networks, and so there is plenty of time to correct rules if they are found to be impairing legitimate communications after you switch over to the “deny all” rule.

Our services personnel report that the average L3/L4 segmentation results in 50-75 firewall “accept” rules, and that the interpretation of session logs and creation of rules generally takes less than one day per firewall.

6.2: SP99 Level2/Level 3 Firewalls

Industrial Defender services personnel report that retrofitting firewalls to segment existing control systems between SP-99 level 2 and level 3 networks, and within existing level 3 or level 2 networks is considerably more difficult than retrofitting into the enterprise perimeter of the control system. There are more kinds of communications typically authorized between level 2 and level 3 networks than there are between control and enterprise networks. Further, internal segmentation errors are much more likely to lead to control system failures. As a result, sites typically require a long test period of monitoring the session logs with accept/log rules flagging anomalies, before activating the “deny all” rule.

Our services personnel report that currently, internal segmentation is still unusual, but is increasing in frequency – most firewall retrofit engagements are still at the control/enterprise perimeter. A typical L2/L3 firewall might have two or three times as many rules as a L3/L4 segmentation effort, and might take up to two days to interpret session logs and configure those rules.

6.3: Anomaly Detection within Control Networks

We were provided with packet capture data from Industrial Defender Network Sensor products monitoring switches in eight DCS networks at level 2 of the SP-99 architecture, an SP-99 level 3 aggregation point for those networks and an EMS QA testbed. Packet data from switch span ports exposes all internal communications, not just communications between security zones. The size of the packet captures from live systems was capped at 50MB to minimize the WAN impact of transferring the files to the analysis site. The 50MB limit meant that on the busiest switches we captured less than five minutes traffic. As a result, the analysis below should be interpreted as a lower bound of what might be found at those sites over a longer period of time. The QA testbed packet capture was larger and represents one full day's data.

The results are summarized below. The summary describes:

- Sessions – the number of session descriptions the anomaly tool produced. Note that this is typically less than the number of sessions observed, since “CLIENT” session outputs can represent many sessions from the same client IP address to the same server port.
- TCP Ports – the number of TCP server ports observed in the session outputs. A server port is a port numbered less than 1024, or one to which more than one connection was observed.

- UDP Ports – the number of UDP server ports observed in the session outputs. A server port is a port numbered less than 1024, or one to which more than one connection was observed.
- Unique IP addresses – the number of unique IP addresses observed in the anomaly detector output.
- TCP/UDP/ICMP – percentages of session outputs for each protocol.

Site	Sessions	TCP ports	UDP ports	IP addresses	TCP/UDP/ICMP
1	465	27	16	93	42/34/17
2	1177	32	33	144	62/33/4
3	708	41	18	102	62/26/12
4	569	9	11	60	52/31/18
5	168	38	25	111	59/35/6
6	566	44	18	98	59/25/17
7	224	13	15	74	51/33/16
8	643	49	19	111	55/33/12
L3	4259	29	26	618	42/53/4
QA	382	101	20	40	78/20/3

The data shows that we observed between 60 and 144 IP addresses in each L2 network for which we had packet data, and that observed between 224 and 1177 communications sessions for each such network. The QA network was comparable to the L2 networks, but one would expect a QA network to contain less equipment and so fewer types of communications sessions than a production network. The L3 network was roughly four times as complex as the most complex L2 network we observed, and roughly ten times as complex as the simplest L2 network.

Experience with the tool indicates that networks the size of the L2 networks above can be characterized manually with “accept” rules with a few hours work. Our attempt to work with the site personnel to characterize the L3 network failed – there was too much session data for the one day we had allotted to the task. Progress on the task suggested that it might take quite a number of days for a manual characterization, and that our confidence in the resulting characterization might be limited because of the sheer volume of data to track manually.

It seems clear that additional automation is needed to undertake analysis of a network of the size of the L3 site. What seemed to be indicated was one or more of:

1. a way to visualize and browse the communications patterns in the network – there are many such tools already available, and/or
2. a way to automatically group IP addresses by similarity in behavior, and suggest rules automatically to account for that group behavior..

7: Security Incidents and Remediations

While there are no incidents reported here from the L2 networks because of lack of time to complete the traffic characterization & rules development at the experimental sites, there is a long history of anomalous traffic having been discovered with HIDS agents and during firewall retrofits. This section describes some examples.

7.1: Windows XP Automatic Updates

A number of Windows XP systems were found to be attempting communications with the Microsoft website. Investigation showed that the communications were attempts to retrieve information about operating systems updates. This alarmed operations staff at the site because none of the equipment on the control network was authorized for automatic updates – all software updates must be tested and approved before deployment.

Further investigation showed that automatic updates were in fact disabled on all equipment on the control network, but that on the offending equipment it was only disabled. The “Automatic Updates” service was still running. Disabling the “Automatic Updates” service stopped the anomalous communications.

7.2: Network Driver “Phoning Home”

A number of systems were found to be attempting communications with a site on the open internet. Investigation showed that the communications were due to a management application supplied by the vendor of a network card in use on that equipment. The site did not investigate further – they simply disabled the communication path for that card through the operations firewall.

7.3: Corporate IT Scanning Control Equipment

Industrial Defender personnel at site doing anomaly-based firewall rules development *frequently* find corporate IT personnel using 'nmap' or other scanners to do port scans on operations equipment. Investigations are often undertaken, not so much because there is any question as to which team is doing the port scans or why, but more often to find the individuals responsible and educate them as to the difference between control equipment and enterprise equipment.

Remediation is generally to design the final firewall rules set to block all such scans.

7.4: Unauthorized Historian Connections

In the course of a firewall retrofit, repeated communications with the plant historian client listen port were detected from an office on another continent. Plant personnel had prepared what they thought was a complete list of authorized users and the IP addresses from which they connected to the historian. The user list included local users and corporate users responsible for company-wide initiatives that required access to historical data. None of these users or IP addresses was from the continent in question.

No investigation was undertaken. Remedial action was to omit the anomalous communications from the firewall rules set.

8: Conclusions

Customers involved in this prototype expressed interest in a couple of aspects of anomaly detection:

- Unauthorized communications, and attempts at communication, with enterprise IP addresses or internet IP addresses were of great concern to everyone who offered an opinion.
- Monitoring internal switches for new types of communications was of interest; especially to operators of the largest, most complex networks, but only if the exercise did not result in a lot of false positives.

In addition, completely characterizing communications on a network by representing such communications in a set of rules small enough to review manually intrigued operations and security staff. More understanding of critical control systems is better than less and this tool seemed to offer staff an insight they do not currently have with other approaches.

The experience of trying to apply the anomaly detection to operations networks produced a number of learnings:

- Communications on small control network segments, containing 50-100 hosts, can be completely characterized with 100-200 rules, using the described rules syntax.
- Characterization of traffic on such small networks takes roughly a full-time afternoon for someone already familiar in large part with the function of equipment on the network.
- Characterization of traffic on a network of 850 IP addresses using the manual methods is impractical. Additional automation or visualization is necessary.
- The existing rules language is not powerful enough. Describing allowed communications between groups of addresses – such as the DNS traffic in the example above – required many lines of rules. A more concise language makes producing rules faster, and makes manual review and approval of the resulting communications patterns more practical.

Firewall installation personnel had the following observations of the anomaly detection tool:

- The session-tracking output of the tool is comparable to and somewhat more compact than the session logging capabilities of our firewall platform.
- The firewall platform only sees communications through the perimeter of a network segment. Anomaly reports across the perimeter are the most useful kind of anomaly reports, but reports of new or unauthorized communications within the control network are also of value.
- As a rule, the firewall platform is not used as an anomaly detection tool. At self-managed customers, firewall session logging on the default “deny” rule is either disabled, or if left enabled is not given much attention by personnel at site. At the customers we manage remotely, we do leave denied packet / session logging turned on, but we have automation to summarize the large volumes of denied packets.

In summary, industrial networks are simpler than enterprise networks and operations staff would like to be able to understand communications on those networks, especially if those communications patterns can be expressed accurately and concisely. There is great interest in alerting on new kinds of communications attempts through the operations perimeter, especially new kinds of egress attempts. Given the simplicity of operations networks, the tools described here are close to the mark, but need additional automation to be practical for larger operations networks.

References

[1] *CISCO Anomaly Guard Module*, 2009, CISCO Systems,
http://www.cisco.com/en/US/prod/collateral/modules/ps2706/ps6235/product_data_sheet0900aecd80220a7c.html

[2] *IBM Proventia Anomaly Detection System helps improve the user experience and adds network profiling to network behavior analysis*, 2008, IBM Corporation2
http://www.ibm.com/common/ssi/rep_ca/1/897/ENUS608-021/ENUS608021.PDF

[3] *Modeling Flow Information and Other Control System Behavior to Detect Anomalies*, 2008, Brian Moran and Rick Belisle, Proceedings of the SCADA Security Scientific Symposium (S4), 2008, S4 Digital Bond Press, Sunrise, FL.

[4] *Snort 2.0: Intrusion Detection*, 2003, Brian Caswell, Jay Beale, James C. Foster, and Jeffrey Posluns, Syngress Publishing Inc., Rockland, MA.

###

About the Author

Andrew Ginter is Chief Security Officer of Industrial Defender, Inc. Andrew is responsible for technology direction, product architecture and corporate IT for the company. Ginter has more than 25 years in security systems and software development, and brings critical technical leadership to Industrial Defender. He previously held various technical management roles at Agilent and Hewlett-Packard. A central theme in Ginter's career has been the development of high quality, reliable, and secure systems software - including compilers, real-time operating systems, network software, industrial control systems, message-passing middleware and alert management systems. Ginter holds a B.Sc. in Applied Mathematics and a M.Sc. in Computer Science from the University of Calgary. His professional background also includes an Information Systems Professional (ISP) accreditation from the Canadian Information Processing Society (CIPS), and a Certified Information Systems Security Professional (CISSP) from the International Information Systems Security Certification Consortium (ISC2).

About Industrial Defender

Industrial Defender, Inc., the global leader in Cyber Risk Protection™, is the first company to offer a completely integrated cyber security solution designed to protect the real-time process control and SCADA environment in a flexible and cost effective platform. This comprehensive Cyber Risk Protection™ Lifecycle solution enables the efficient assessment, mitigation and management of cyber security risk within the critical infrastructure network domain. Industrial Defender is a privately held company with over 18 years of real-time process control and SCADA industry experience and more than 7 years of industrial cyber security experience. Industrial Defender has completed more than 100 process control / SCADA cyber security assessments, more than 10,000 global technology deployments in securing critical infrastructure systems, more than 3,000 mission critical SCADA deployments and provides managed security services for 170 process control plants in 21 countries.