ArcSight and Malware Beacon Detection

Prepared By: John Bradshaw and Damian Skeeles

Date: Wednesday, February 17, 2010

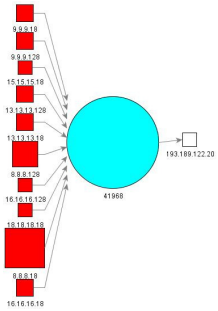

ArcSight, Inc.
5 Results Way, Cupertino, CA 95014, USA
www.arcsight.com
info@arcsight.com

Corporate Headquarters: 1-888-415-ARST
EMEA Headquarters: +44 870 351 6510
Asia Pac Headquarters: 852 2166 8302

Table of Contents

MalWare Beacon Overview

Malware has evolved over the years into sophisticated code that incorporates error detection, stealth capabilities, as well as distributed command and control capabilities. While security vendors constantly search for methods to identify and detect malware before it can infect a system, there is always the threat that a newer, more sophisticated method will bypass initial detection. One of the biggest threats facing customers is the unknown sleeper agent awaiting instructions from a master controller to execute its payload.

In order for this command and control structure to work, there must be some form of communication that occurs between the zombie system(s) and the master controller. The regular checking in of a zombie with its master controller (MC) is commonly referred to as ***Malware Beaconing.*** The purpose of this ArcSight Use Case is to document methods the ArcSight Enterprise Security Manager (ESM) correlation engine can assist security analysts in detecting these Malware Beacons.

Beacon Traffic Types

As with malware payloads themselves, beacon communications has grown more sophisticated. Initial beacons were designed to let the MC know the zombie was up and active. This could be achieved with simple beacons that had the following characteristics:

- Single UDP / TCP packet
- Set interval when the beacon occurs (hourly, daily, etc.)
- Small packet size – generally carried only identifying information

Now with increased sophistication, malware communications have evolved to incorporate additional capabilities designed to make them harder to detect:

- Full two-way communication (TCP) over common well-known ports (ie: 80, 443)
- Adherence to protocol RFCs to escape packet anomaly detection (ie: full HTTP GET requests and HTML delivery)
- Randomized time intervals for communication
- Heartbeat with version updating of malware code

With the communications now looking more and more like every day traffic, what detection methods are available to help combat this threat?

ArcSight Solutions and Malware Beacon Detection

ArcSight solutions provide abilities that assist analysts in detecting and responding to malware beacon events. The following capabilities outline embedded technologies in the ArcSight product lines:

ESM Asset Modeling

ArcSight Administrators can classify assets based on importance to business operations. Asset criticality rankings are used by the Threat Priority Formula to determine overall incident priority. While this is one of the strongest capabilities of ArcSight, it is also one that is usually not setup to its full potential.

The first basic step to detecting any external threat is to know what's yours. ArcSight categorization provides a simple label that can be applied to all assets within the organization regardless of asset priority and use. That label is the "Site Asset Categories\Address Space\Protected" tag:

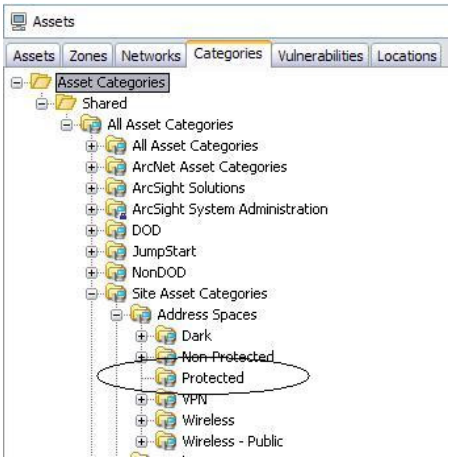

The first step of correlation can now be defined with a simple condition:

Any communication from a Protected Labeled asset to anywhere else.

IP Geo-spatial Location

ArcSight ESM includes publicly licensed geo-spatial information for IP addresses. This allows security analysts to prioritize events based on a source/destination's physical location. If traffic at your site wouldn't normally include traffic to foreign countries, then certainly *any* traffic detected to a foreign country would become suspect. You also have the ability to identify

countries you deem more suspect than others (see Active List Technologies below) or IP Address ranges of international business partners you wish to exclude.

Active List Technologies

The ability to check information against a static or dynamic list of entries if another key capability of ArcSight ESM. For example, we may wish to isolate our search to only a list of ports that are commonly open for outbound traffic through a firewall (such as 80 and 443), or as noted above, we may have a list of foreign countries where any traffic seen would create a much higher priority incident.

You may also wish to import information from third-party security threat intelligence subscriptions or use the US CERT Watchlist to monitor for any signs of activity going to known-bad sites.

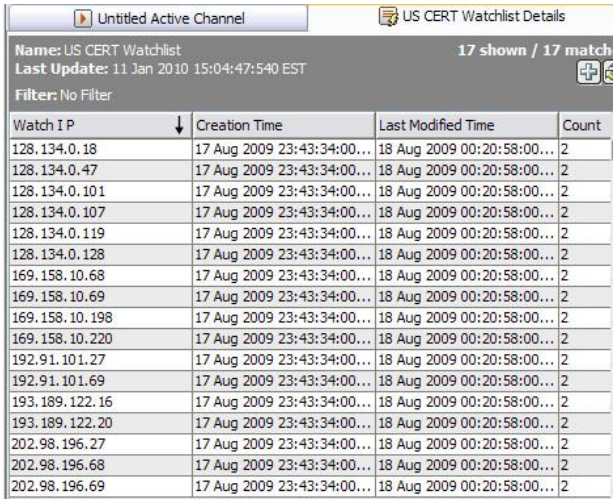
Untitled Active Channel | US CERT Watchlist Details

Name: US CERT Watchlist — 17 shown / 17 match
Last Update: 11 Jan 2010 15:04:47:540 EST
Filter: No Filter

Watch I P	Creation Time	Last Modified Time	Count
128.134.0.18	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
128.134.0.47	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
128.134.0.101	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
128.134.0.107	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
128.134.0.119	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
128.134.0.128	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
169.158.10.68	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
169.158.10.69	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
169.158.10.198	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
169.158.10.220	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
192.91.101.27	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
192.91.101.69	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
193.189.122.16	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
193.189.122.20	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
202.98.196.27	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
202.98.196.68	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2
202.98.196.69	17 Aug 2009 23:43:34:00...	18 Aug 2009 00:20:58:00...	2

The key advantage here is that you now have a dynamic list and if that list were to contain known MC sites, then any system communicating with that site is now highly suspect of having a malware infection.

Required Event Feeds

Ensuring the right information is being captured for analysis is another important step in the detection of any threat. For malware beacons, we will focus on the egress points to the Internet. There are several requirements that must be fulfilled in order to properly evaluate the information for potential beaconing. The more of this information that can be gathered, the better our detection model will operate:

- Source and Destination IP address and port information
- # of bytes transferred out
- # of packets transmitted in the session
- Protocol type (UDP or TCP)

There are three likely event source candidates that may provide some, or all, of the information:

- Network traffic flows (ie: NetFlow)
- Firewalls
- Network Intrusion Detection/Prevention Systems (NIDS/NIPS)

ArcSight and NetFlow

For the purposes of this paper, we will be using Cisco NetFlow events to discuss content creation and capabilities. Depending on the event information available, you will be able to create all, or a subset, of these capabilities in your ESM solution. ArcSight currently supports Cisco NetFlow v5 and v9. These event feeds contain all the relevant information listed above:

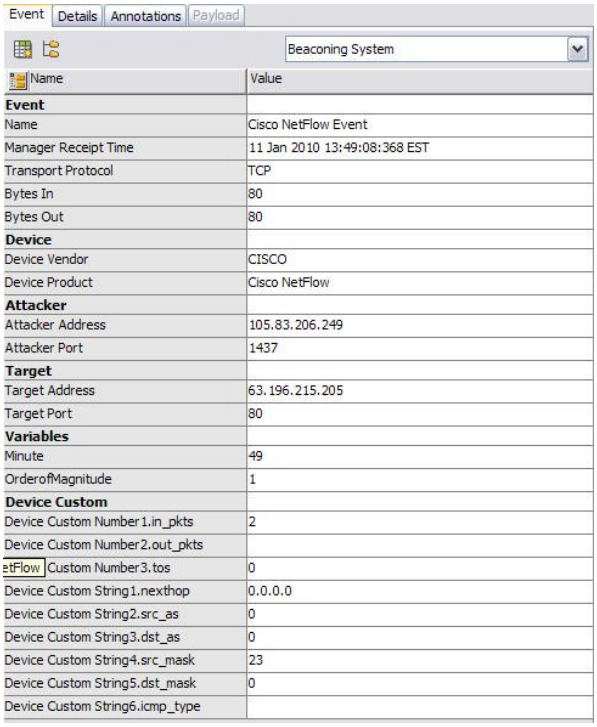

Using ArcSight Correlation To Increase Beacon Detection Probability

Now that we've classified what systems we're protecting and what event information will be helpful in beacon detection, we can now start defining scenarios that indicate possible beaconing activity. Remember, our assumption at this point is that traditional malware detection mechanisms (host-based IDS, anti-virus and NIDS/NIPS) missed the malware installing itself on the target system and it is now operating as the attacker planned.

The way we're going to assist security analysts in identifying beaconing zombies is by an increasing threat profile based on patterns of activity we might expect to see from a beaconing

system. We will never reach 100% positive identification, but certain patterns of activity are far more suggestive of beacon traffic than others. Let's discuss some possibilities.

Traffic to Foreign Countries

What if you saw activity from one of your systems described as follows:

- Single packet transfer
- Well-known port (ie: 80, 443, 53)
- Destination is to a foreign country
- Occurred during off-hours

Would you say the likelihood of this being a beacon is very high? This scenario is very easy to detect in ArcSight as follows:

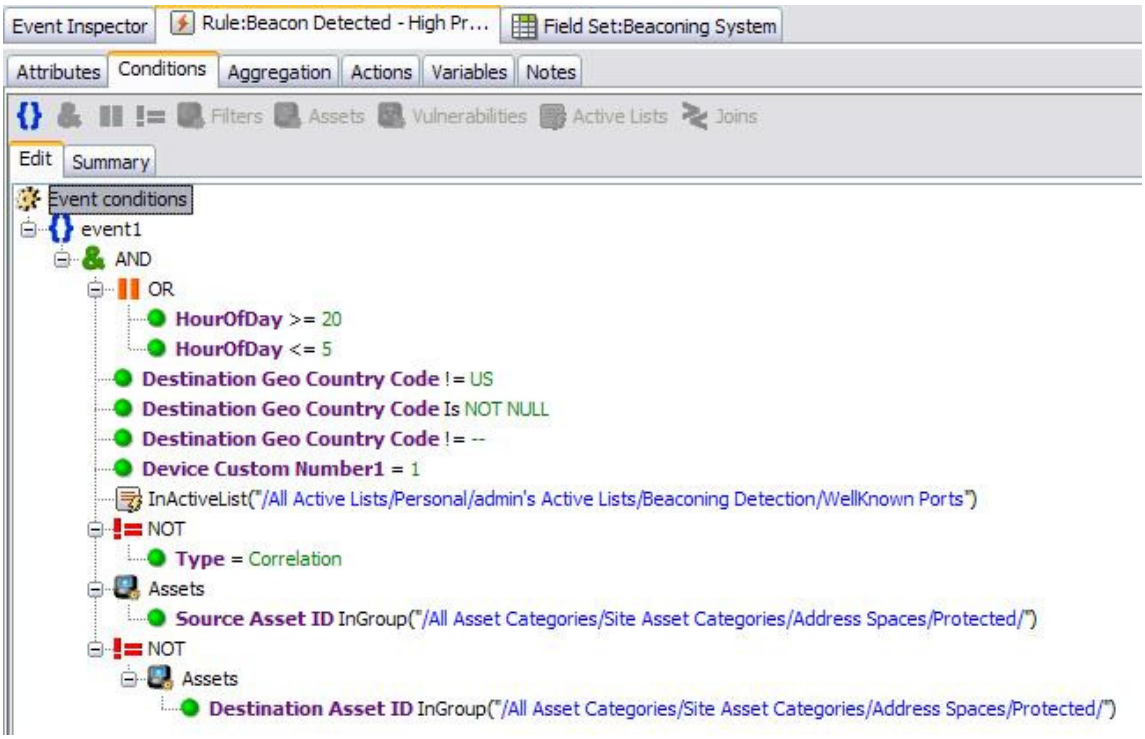

This rule states: Correlate any single-packet transfer over a well-known port between the hours of 8:00PM and 5:00AM from a protected system to a system outside of our protection and outside of the United States.

Traffic to Destination on Watch List

What if you saw activity from one of your systems described as follows:

- Low # of packet transfers
- Low # of bytes transferred
- Well-known port (ie: 80, 443, 53)
- Destination is to an IP or subnet on one of your watch lists
- Occurring anytime of the day

What are the odds that this traffic is a beacon for malware (or at the very least something that merits additional investigation). Again, ArcSight Active Lists and asset categorization help us quickly isolate this pattern of activity:

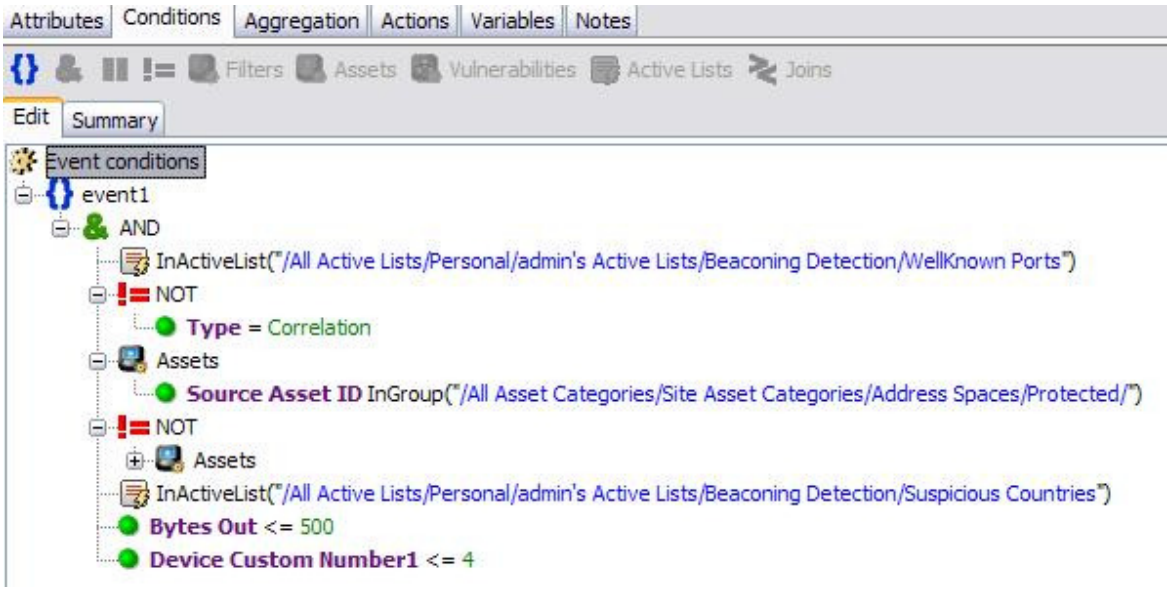

This rule states: Correlate any traffic where the number of packets is four or less with less than 501 bytes transferred out and where the source is a protected asset and the destination is to one of the countries on my suspicious watch list.

Putting It All Together

Once we've mapped out the various correlation rules that define patterns of activity that could be beaconing and we've assigned a relative priority to each pattern, we can then begin construction of an appropriate dashboard that will visually display this information as it happens.

The following dashboard illustrates ArcSight's ability to alert on potential beaconing systems

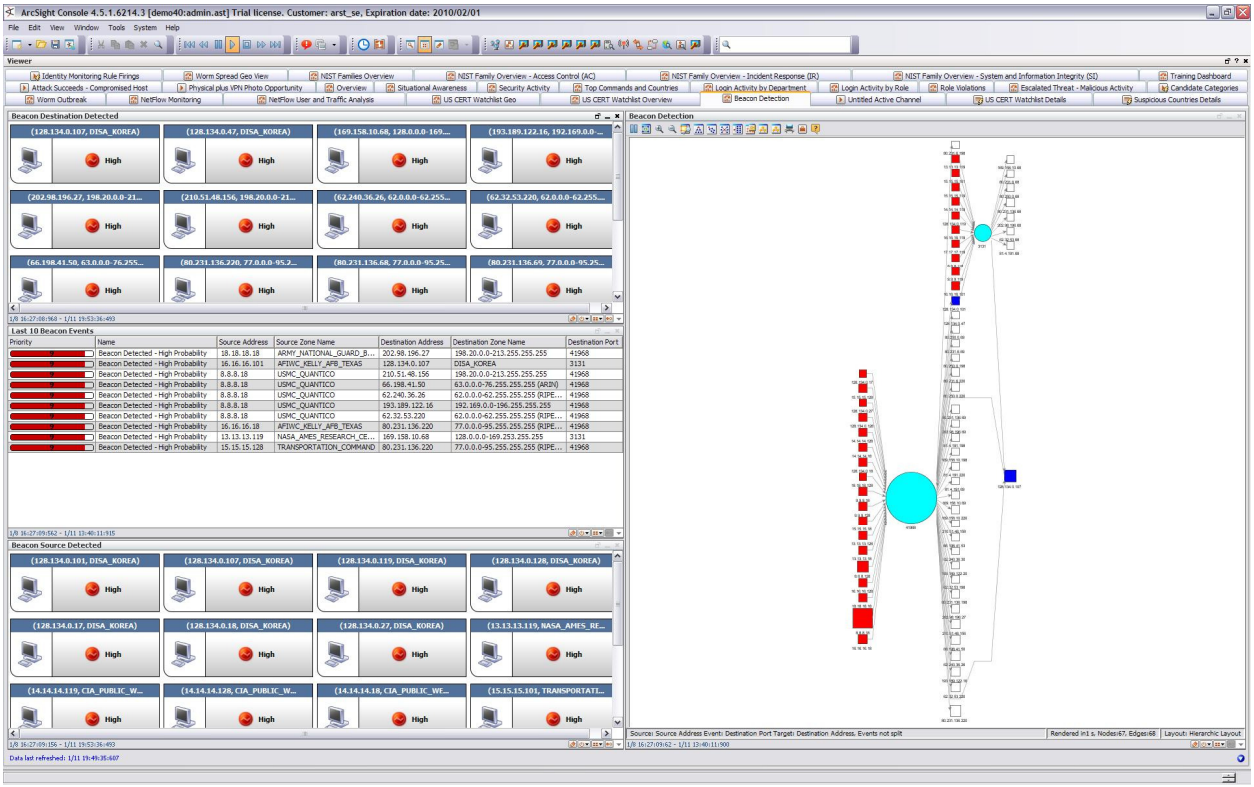

Last State Data Monitors

ArcSight Last State Data Monitors analyze incoming event data and populate the dashboard with status mappings based on event field values. For example, the overall priority of the correlated beacon detection event is used to map the criticality symbol in the monitor:

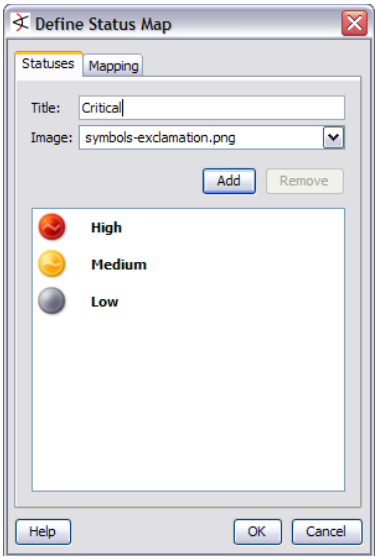

Red represents high-probability of beaconing malware, yellow represents a medium probability and green the lowest probability. Based on the patterns of activity discussed earlier, analysts can set the data monitor to prioritize detection probability rankings based on their site's environment.

Event Graph Data Monitor

There is nothing that catches the eye more than an event graph as it starts to populate with relevant information. The event graph setup for this dashboard maps the source IP address, target port and target IP address of any potential beaconing system and its MC.

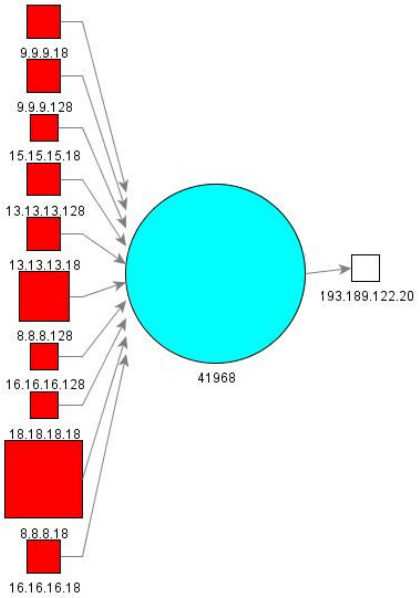

Worked Example: Malware in an MSSP

The following is an illustration of the above concepts being put into practice, when malware was detected by an ArcSight ESM installation during a five-day Proof-of-Concept at a large MSSP in Europe. This example illustrates how Pattern Discovery can detect such activity, even with short timescales and minimal configuration.

Defining the Pattern Boundaries

As the customer had well-defined security policies, it was possible to use Outbound Firewall Denies as an initial filter for selecting events for pattern discovery. This is on the basis that legitimate outbound traffic should pass through established proxies or ports, and hence outbound traffic which is dropped by the firewalls is likely to be due to misconfigurations or illegitimate traffic. Once this has been passed, our attention can then turn to legitimate traffic.

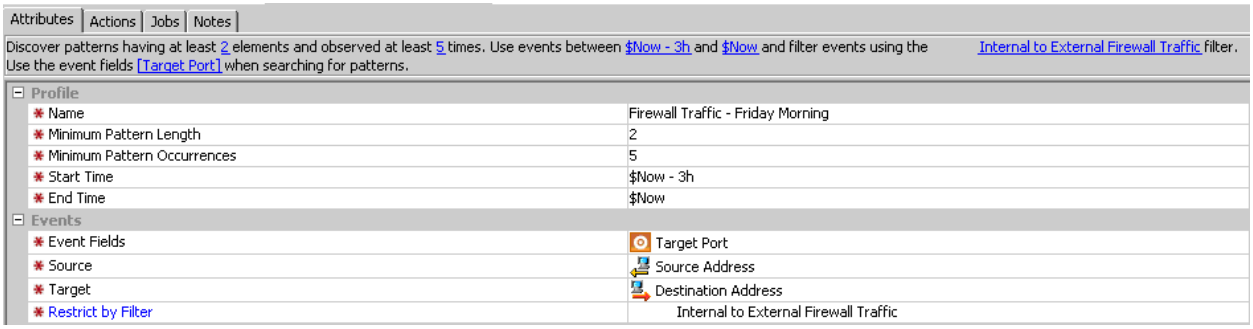

Within this filter, and a sample 3-hour data set, Pattern Discovery is instructed to identify instances of failed attempts by an internal IP to access an external IP, on the same combination of ports each time. The precise ports and IPs are not specified – this is for Pattern Discovery to identify.

Running the Analysis

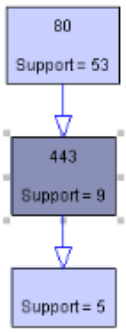

Pattern Discovery takes a snapshot of the filtered events, runs an analysis, and then returns the results of identified patterns, displayed as a flowchart. As can be seen, it has detected nine instances of systems attempting and failing to contact an outside IP address using ports **80,443**, and of these, five attempted to use **80,443,** and then a third port.

Double-clicking this graph brings up a detailed analysis of the pattern.

In this screen, we can see the Source and Destination IPs of that pattern, for each occurrence. The Source IP is always the same, but the Destination IP is very different. This suggests a single attacker, and that this is likely to be malware.

Pattern Discovery also details the time-spread of the events for each individual sequence, and can give a statistical breakdown of the sequences as a whole. For example, how often the Attacker attempts port 80 first, against how often it uses port 443 first. Pattern Discovery does not require events to be in the same order each time

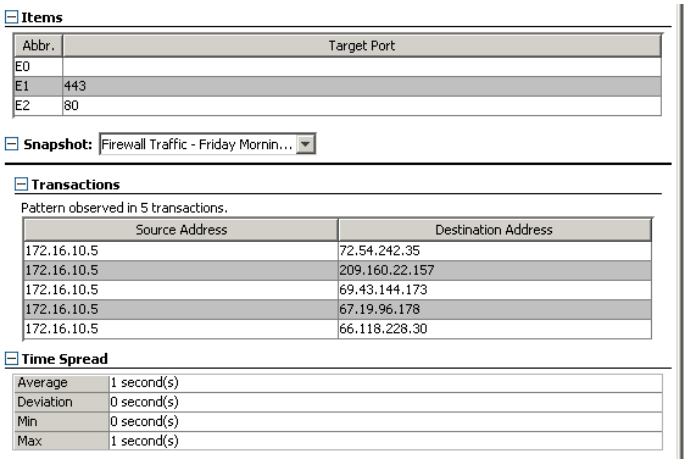

Items

Abbr.	Target Port
E0	
E1	443
E2	80

Snapshot: Firewall Traffic - Friday Mornin...

Transactions

Pattern observed in 5 transactions.

Source Address	Destination Address
172.16.10.5	72.54.242.35
172.16.10.5	209.160.22.157
172.16.10.5	69.43.144.173
172.16.10.5	67.19.96.178
172.16.10.5	66.118.228.30

Time Spread

Average	1 second(s)
Deviation	0 second(s)
Min	0 second(s)
Max	1 second(s)

This is all more clearly illustrated as an event graph, as seen here.

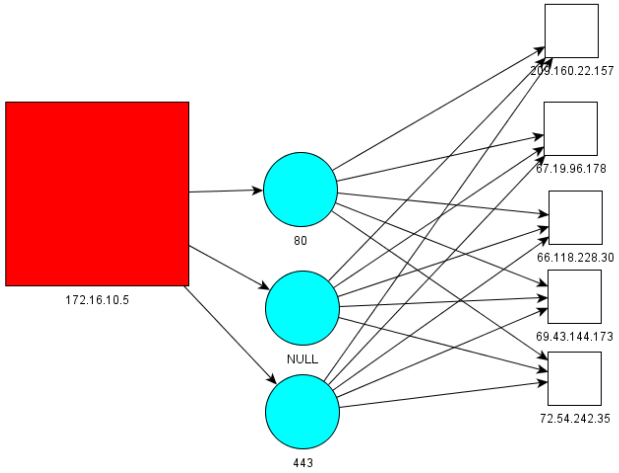

Investigating the Attacker

With the attacker identified, it is then possible to right-click the IP Address and investigate **all** activity originating from that IP address, including traffic passed successfully by the firewall, OS activity, IDS alerts, Netflow traffic,and so on.

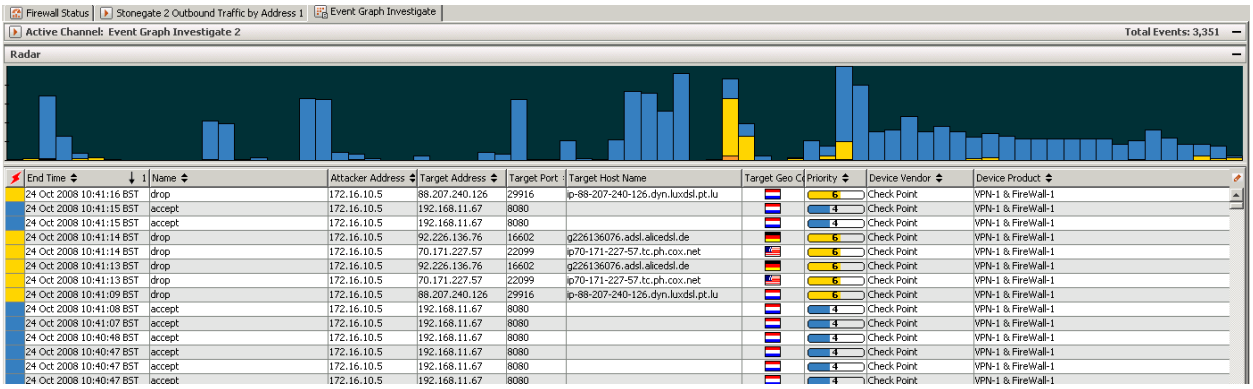

End Time	Name	Attacker Address	Target Address	Target Port	Target Host Name	Target Geo C	Priority	Device Vendor	Device Product
24 Oct 2008 10:41:16 BST	drop	172.16.10.5	88.207.240.126	29916	ip-88-207-240-126.dyn.luxdsl.pt.lu		6	Check Point	VPN-1 & FireWall-1
24 Oct 2008 10:41:15 BST	accept	172.16.10.5	192.168.11.67	8080			4	Check Point	VPN-1 & FireWall-1
24 Oct 2008 10:41:15 BST	accept	172.16.10.5	192.168.11.67	8080			4	Check Point	VPN-1 & FireWall-1
24 Oct 2008 10:41:14 BST	drop	172.16.10.5	92.226.136.76	16602	g226136076.adsl.alicedsl.de		6	Check Point	VPN-1 & FireWall-1
24 Oct 2008 10:41:14 BST	drop	172.16.10.5	70.171.227.57	22099	ip70-171-227-57.tc.ph.cox.net		6	Check Point	VPN-1 & FireWall-1
24 Oct 2008 10:41:13 BST	drop	172.16.10.5	92.226.136.76	16602	g226136076.adsl.alicedsl.de		6	Check Point	VPN-1 & FireWall-1
24 Oct 2008 10:41:13 BST	drop	172.16.10.5	70.171.227.57	22099	ip70-171-227-57.tc.ph.cox.net		6	Check Point	VPN-1 & FireWall-1
24 Oct 2008 10:41:09 BST	drop	172.16.10.5	88.207.240.126	29916	ip-88-207-240-126.dyn.luxdsl.pt.lu		6	Check Point	VPN-1 & FireWall-1
24 Oct 2008 10:41:08 BST	accept	172.16.10.5	192.168.11.67	8080			4	Check Point	VPN-1 & FireWall-1
24 Oct 2008 10:41:07 BST	accept	172.16.10.5	192.168.11.67	8080			4	Check Point	VPN-1 & FireWall-1
24 Oct 2008 10:40:48 BST	accept	172.16.10.5	192.168.11.67	8080			4	Check Point	VPN-1 & FireWall-1
24 Oct 2008 10:40:47 BST	accept	172.16.10.5	192.168.11.67	8080			4	Check Point	VPN-1 & FireWall-1
24 Oct 2008 10:40:47 BST	accept	172.16.10.5	192.168.11.67	8080			4	Check Point	VPN-1 & FireWall-1
24 Oct 2008 10:40:47 BST	accept	172.16.10.5	192.168.11.67	8080			4	Check Point	VPN-1 & FireWall-1

It is also possible to use Event Graphs to gain an insight into the full scale of the malware's activity. The 'nest' on the top right are the individual denied attempts to access external systems by the malware.

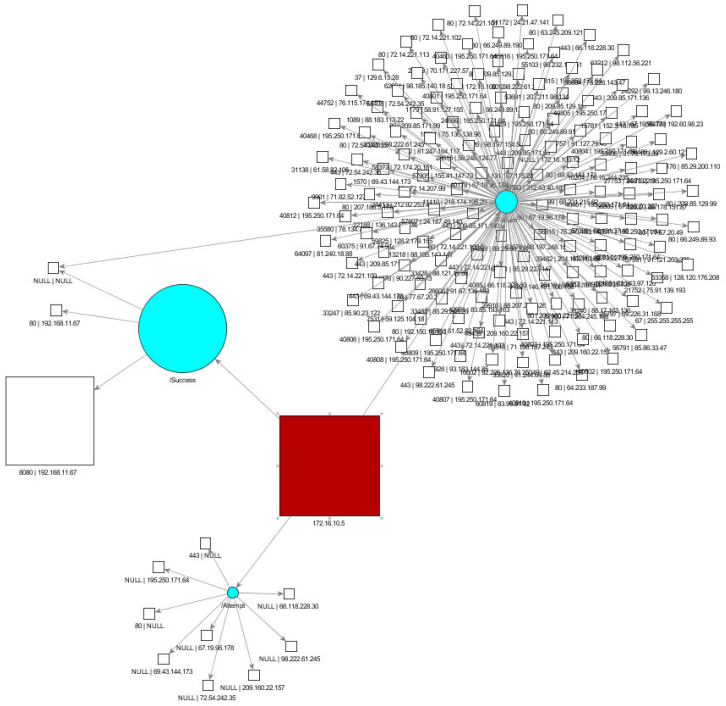

This can also be portrayed geographically on dashboards, or on Google Earth, to understand any geopolitical context to the activity.

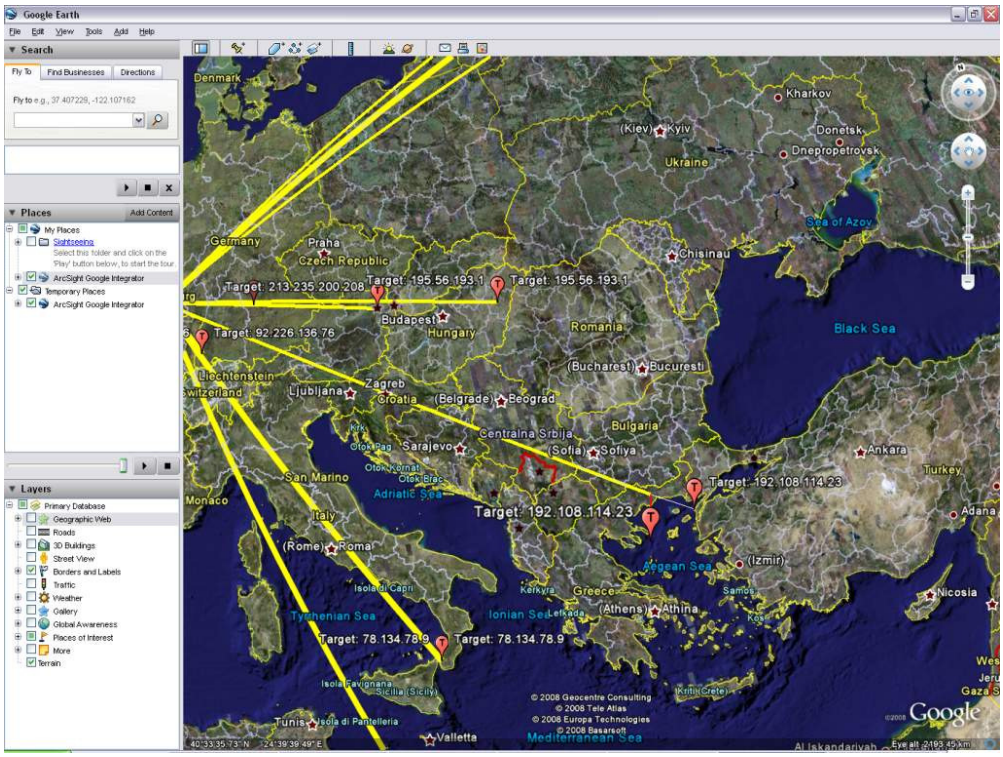

Clean-up

This malware was traced by the security team to a Security Analyst's laptop, running in the main office. Given this took place in a major European MSSP, which was the dominant provider of managed security services to Financial Institutions in their country, this was an alarming breach. Prior to the ESM POC, they were unaware of the malware's existence, and did not know how long it had been in place, or whether it had succeeded in exporting data out of the organization.

This company has since become an ArcSight ESM customer.

Solution Overview

Malware is constantly being evolved by hackers in an effort to evade detection by anti-virus and intrusion detection systems. Security Analysts require tools that will allow them to detect patterns of activity that can be associated with communications methods used by malware to control the code's operations.

ArcSight's state of the art correlation engine provides methods that can define these patterns of activity, prioritize them and visually display them in a manner that allows analysts to take meaningful actions in the course of their investigation.