AC-Hunter™ CTF User Guide

Greetings and welcome! This condensed user guide will walk you through the process of connecting to and using AC-Hunter for the CTF.

(Note: This is not the complete AC-Hunter software User Guide)

Logging in to AC-Hunter

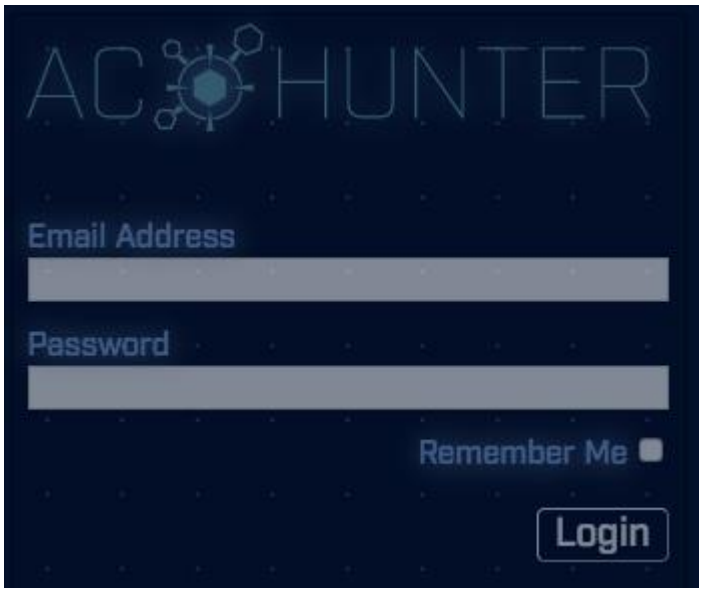

Dataset Selection

When you login to AC-Hunter for the first time, there may not be a default dataset selected. This will cause the Dataset Selection window to appear (you can always go here manually by clicking the gear icon on the Dashboard tab).

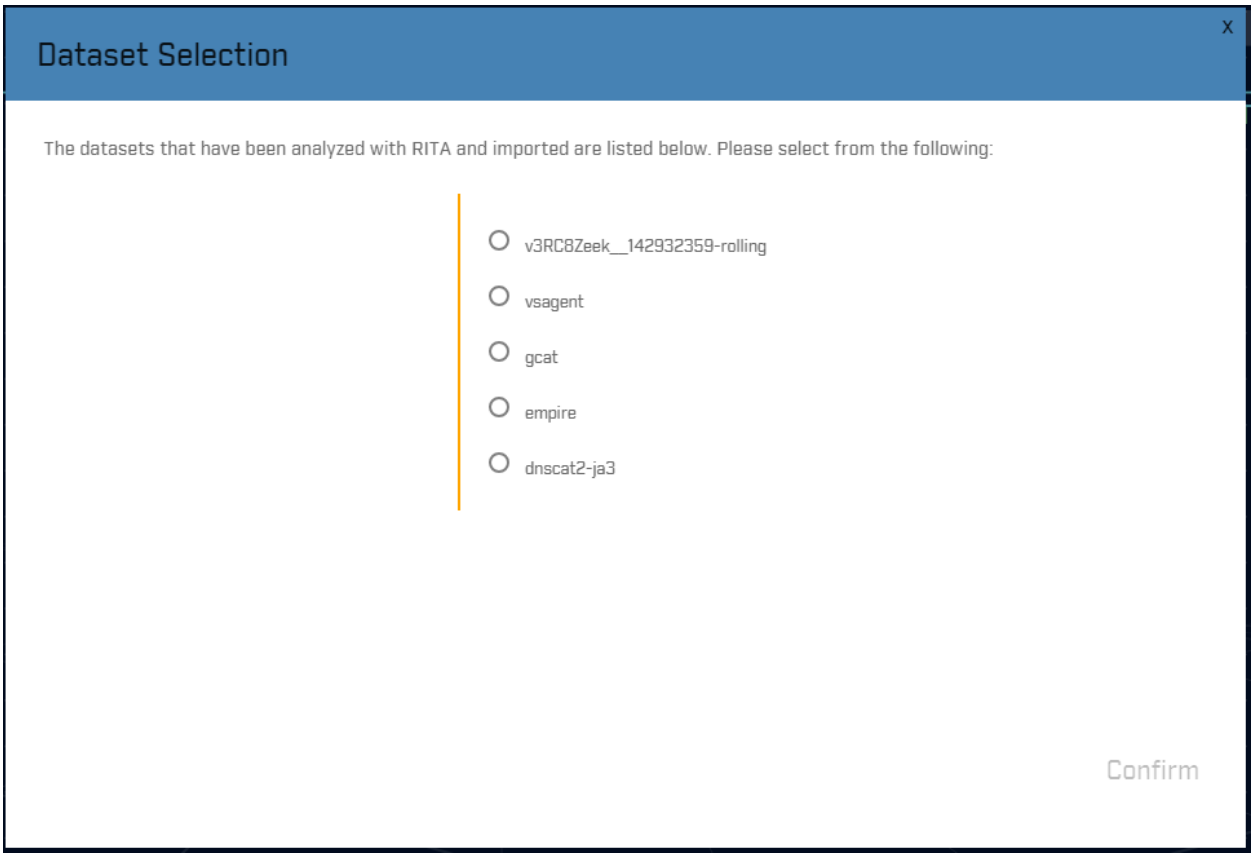

The datasets used for the CTF challenges are **DB1**, **DB2** and **DB3**.

The other sample datasets in the list are examples of command and control traffic. This is sample data that will give you an opportunity to explore within AC-Hunter. Click the radio button to the left of the first dataset you want to work with, then click the "Confirm" button on the bottom right.

This will return you to the Dashboard. You can reach the Dashboard at any time by clicking the Dashboard button in the bottom left.

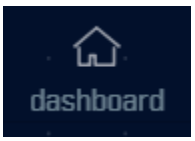

If you later want to change the dataset you are working with, return to the Dashboard and then click the gear icon in the top right.

Dashboard

The Dashboard identifies which systems are most likely to be compromised and why.

The panel on the left side of the screen (Hosts) lists suspicious IP addresses, sorted by threat score, with the highest score listed at the top. You can scroll the list and click other IP addresses to analyze them as well.

The panel on the right-hand side of the screen (Threat Activity) identifies which threat vectors were detected and how they were combined to create the threat score for the selected IP address.

If you want to perform a deeper analysis on any of the threat activity, you can click the listed threat activity item on the right side of the screen. This will load whatever module can be used to analyze the threat. For example, clicking on "beacon score" will automatically load the beacon module. The search function will be set to the IP address being analyzed so you can focus on just that system.

Beacon Analysis

To check for internal systems that are exhibiting beaconing behavior, use the "beacons" module.

Beacons by IP or hostname

AC-Hunter has two ways to identify Beacons - regular traffic between one of your systems and an external IP address, or regular traffic between one of your systems and a hostname. The second approach, newly added in version 5.1.0 allows you to identify Beacons even when the remote IP keeps changing by looking at the hostname used instead of just the IP.

To see how these differ, look at the beacons and beacons fqdn tabs.

There are two ways you can enter the module, and each will present the data slightly differently. If you are on the dashboard and click the "beacon Score" link under "Threat Activity", the beacon module will load with a filter options set in the Results Feed (see below) to only show beacons from this source IP address. This will appear in the top left of the screen.

If however, you click the "beacons" icon on the bottom of the screen and enter the module that way, no filter will be implemented. The result is that the top listed IP address will be the source IP that generated the highest beacon score, regardless of its final threat score. This is useful when you are only interested in reviewing beacon activity.

Beacon results are displayed based on IP pairs. It is possible for a source IP to be listed multiple times if beacon activity was detected with multiple destination IP addresses on the Internet.

Results Feed

The results feed identifies all internal IP addresses where beacon activity was detected. As mentioned above, it’s possible for a source IP address to be listed multiple times if beacon activity was detected with multiple destination IP addresses on the Internet.

Each identified connection is preceded by a colored icon. These can be red, orange, yellow, blue or green. The color code provides a quick visual indicator of the likelihood of the system exhibiting beaconing behavior. Red is most likely while green is least likely.

The ">50 v" immediately under "RESULTS" allows you to set a minimum threshold; you will only see entries whose likelihood is greater than this cutoff. Click on it to change the threshold:

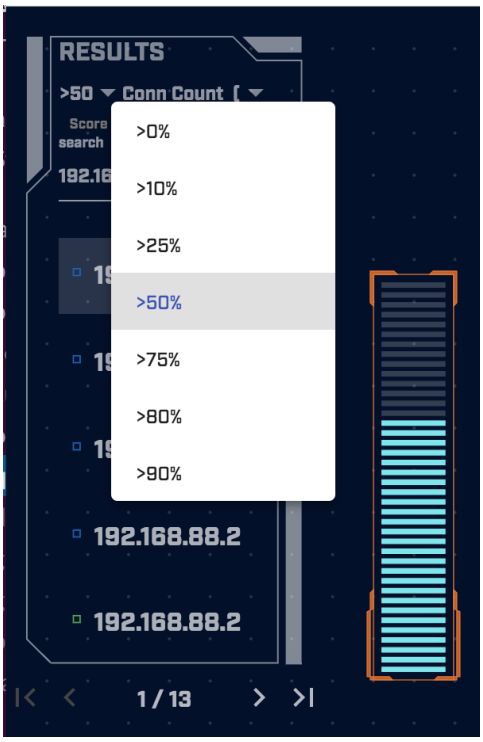

By clicking the "Conn Count [^] v" at the top of the IP list, you can change the way the IP addresses are sorted. When you change the sort order, the label changes to show how the records are sorted.

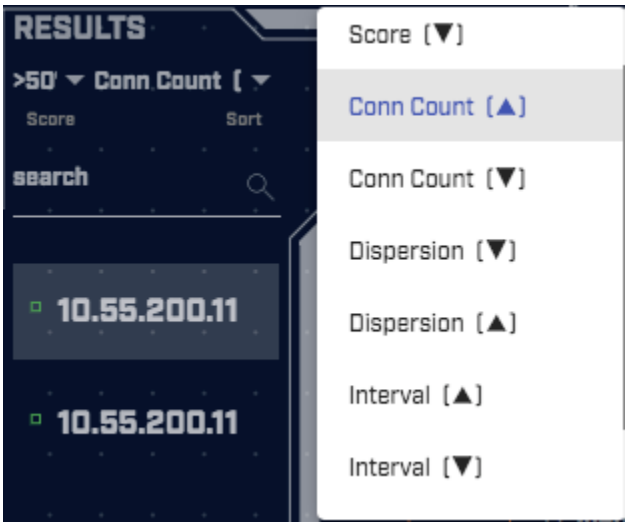

The default is to sort the list based on the score from highest to lowest. This will automatically place the systems most likely to be beaconing at the top of the list. Another useful sort option is by connection count as compromised systems tend to generate thousands of beacon signals. Once you select a sort option, you will be returned to the results feed.

Searching Results

If there is a particular IP address you are interested in analyzing, you can search for it by entering the IP address in the search prompt just above the list of IP addresses. You can enter a partial IP addresses to match on entire ranges:

Source and Destination Analysis

Just to the right of the results feed is the source and destination IP address identified in each connection. To the right of each of the addresses are two symbols, a minus sign and a filter icon.

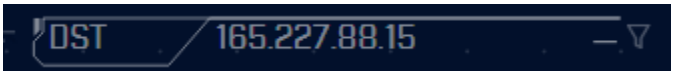

Clicking an IP address will open a menu with one or more options, commonly "AbuseIPDB", "VirusTotal", "copy to clipboard, and "deep dive". The first two are external websites you can

use to research the destination IP address, the third puts the IP address on the clipboard for ease of pasting to other tools, and the last brings you directly to the deep dive module:

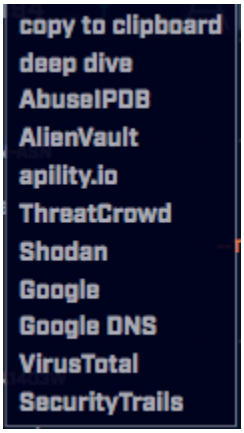

Clicking the minus sign collapses the details regarding the IP address. If the minus sign is replaced with a plus sign, the details are already collapsed, and you can expand them by clicking the plus symbol.

The details include useful information regarding the IP address, such as its Autonomous System Number (ASN), as well as the organization responsible for that ASN. You can see the IP address's range assignment as well as geolocation information.

Note the Fully Qualified Domain Name (FQDN) is displayed. It is important to note that this is *the name the source IP queried* via DNS which brought it to this destination IP address (A or AAAA record query). This is not the PTR record associated with the IP address. This is important because it can provide additional insight as to the intent of the connection. Knowing that an IP is part of a public cloud hosting provider is unremarkable. However, knowing the actual host name the system queried can be a valuable data point.

Below the FQDN is the "Comm" line. This shows what protocols were used in communications between the two systems. The syntax is:

(protocol / port; service)

So for example:

(tcp/443;ssl)

Would be traffic going to port TCP/443 that included an SSL or TLS negotiation at the beginning of the session. Sometimes the service may not be identified. There are a number of reasons this can occur:

- Zeek does not have a decoder for this specific service

- The capture data missed the initial packets that included the service headers
- An unknown service is being run over the port that Zeek does not recognize

If multiple protocols were used in communications between the source and destination IP address pair, multiple entries may appear on this line, up to a total of five. If more than five protocols were observed, the Comm data will end with "+++". If you need to review all communications, you can do so within the deep dive module.

Whitelisting

Whitelisting editing has been disabled for the CTF.

Timeline Analysis

The graph along the bottom of the screen identifies how frequently the communication connection was observed within the data set being analyzed. Here's an example:

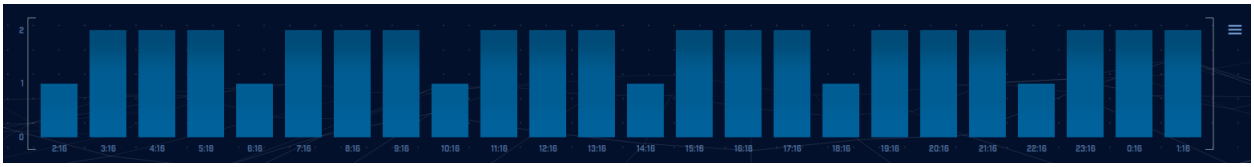

Each bar represents the number of communication connections seen within a specified period of time. The default is one hour. So in this example, the communication connection being observed consistently fires off twice per hour, but occasionally one time every hour (the six shorter bars in the graph). The fact that communication takes place so consistently makes the communication connection suspect.

You can increase the resolution by decreasing the time value represented by each bar in the graph. This is done by clicking the bar icon on the top right portion of the graph:

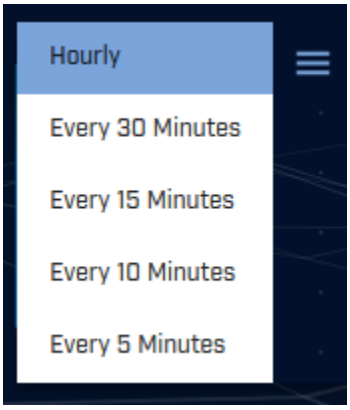

Note that increasing the resolution requires increased processing of the data, so you may see the screen refresh more slowly.

Connection Frequency Chart

The connection frequency chart can be found on the top right side of the screen. It provides a visual representation of the number of observed connections that exhibit the same timing behavior.

The x-axis identifies the timing between connections in seconds. The y-axis identifies the number of connections between the two specified IP addresses that exhibited the timing shown along the x-axis. For example, the bar on the right is identifying that 10 connections were separated by 2,049 seconds, or 34.15 minutes. This is approximately twice per hour. The second bar (on the left) is showing that 30 connections were separated by 2,048 seconds, or 34.13 minutes.

Scoring Chart

The scoring chart appears just to the right of the results feed and will look similar to the following:

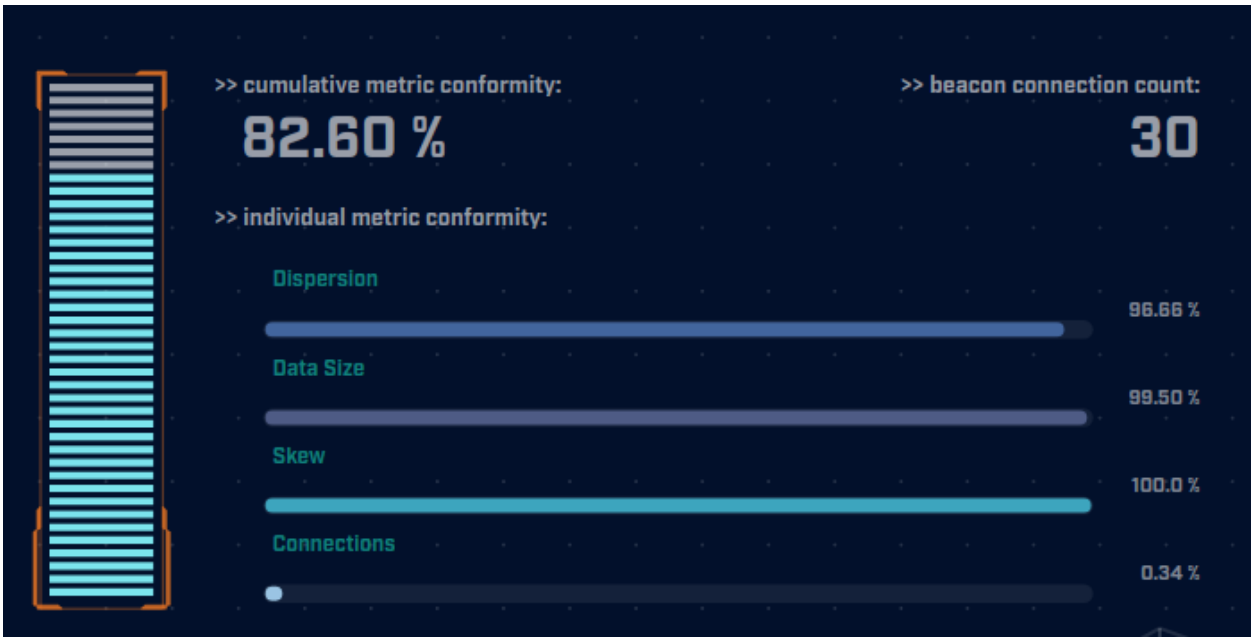

The chart visually displays the results of the timing analysis for the specified IP source and destination address pair. The right side of the chart displays the results from analyzing various timing attributes, while the left-hand side shows the overall score.

The most important value is the overall score on the left-hand side of the chart. In the above example this value is 82.60%. Any connection pair that scores higher than 98% should be considered highly suspect. Scores below 80% are mostly likely not an indication of beaconing. Scores between 80% and 98% require a deeper analysis, with priority given to higher scores.

The right side of the score chart shows the results from all of the timing attributes that were analyzed to generate the overall score.

Data Size Analysis

Up until now we have been analyzing the connection timing attributes within the dataset. We can also perform a similar analysis based on data size. This can be accessed via the "Views" radio button on the right hand side of the screen (right beneath the Report icon).

View "1" permits you to analyze connection timing attributes. By clicking the radio button to the left of the number "2", we can analyze data size attributes.

When you click the number "2" radio button, you will see the screen refresh. However, the layout will remain similar to when we were performing a timing based analysis.

Some of the charts and features described in the previous section are available here as well. The following charts and features are unchanged between views.

- Results Feed, however the results are prioritized based on an analysis of data size instead of connection timing.
- Source and Destination IP
- Whitelisting functionality
- Timeline Analysis

The remaining charts are modified to be more appropriate for performing a data size analysis.

Data Size Frequency

In view 2, the chart furthest to the right changes from analyzing timing dispersion to analyzing consistency of data size. Here's an example:

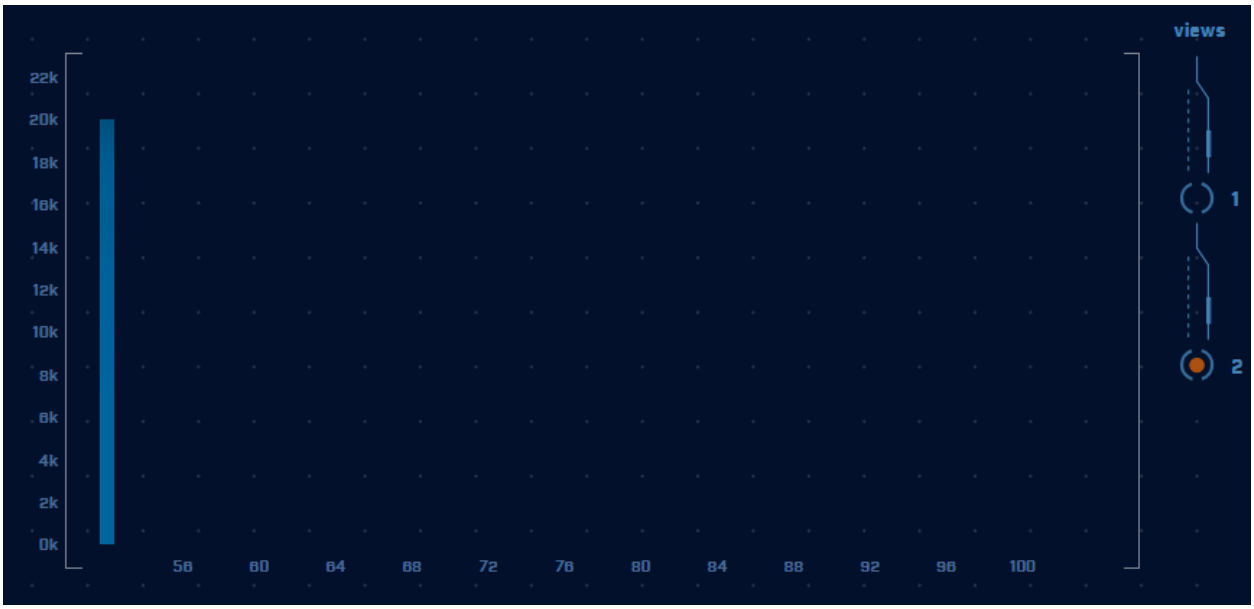

The x-axis identifies data size, while the y-axis identifies quantity of connections. Each bar represents the number of connections seen sending the specified amount of data. Note our example is extremely consistent. All 20,000 or so packets were approximately 50 bytes in size. This is useful information, as it tells us that the attacker never activated the compromised system. If they did, we would see this reflected in the packet size analysis as the additional information would cause larger data exchanges. We would see some number of larger data sessions mixed in as well.

Scoring Chart

The scoring chart has changed to show data size related attributes.

Strobes Analysis

Strobes are similar to beacons in that they are repeated connections between two IP addresses. However, unlike a beacon which may try and hide its signaling, a strobe makes no attempt at being stealthy. A signal that triggers two or three times a second is an excellent example of a strobe. To analyze strobes, click the strobe icon on the menu.

This will produce the strobes analysis screen, however there are no strobe examples in the CTF datasets.

Long Connections Analysis

One way attackers attempt to evade beacon analysis is by creating persistent connections. In other words, they attempt to leave the connection active for as long as possible. This creates fewer firewall log entries, and thus is indicative of more advanced malware. To analyze these long connections, click the long connections icon at the bottom of the screen.

This will produce a screen similar to the following:

In the top left you can select a threshold for how long a connection should stay active before it appears in this output. The default is one minute, but you may wish to set this to a larger value. By default, connections are displayed from longest to shortest. You can reverse this sort order if you choose.

The output is fairly self-explanatory. You see the IP address of the system that initiated the session. You also see the destination IP address and port number. The "service" is an application layer analysis of the protocol being used to communicate over this port. For example, "ssl" indicates that a normal SSL/TLS handshake was detected. Be suspicious of applications using non-standard ports or communications to standard ports that do not follow the associated protocol (example: traffic to TCP/80 but http is not detected as the service).

To only see long connections to and from a particular IP address, enter that address in the search box in the upper left:

Clicking an IP address will open a menu with one or more options, commonly "AbuseIPDB", "AlienVault", "copy to clipboard", and "deep dive". The first two are external websites you can use to research the destination IP address, the third puts the IP address on the clipboard for ease of pasting to other tools, and the last brings you directly to the deep dive module:

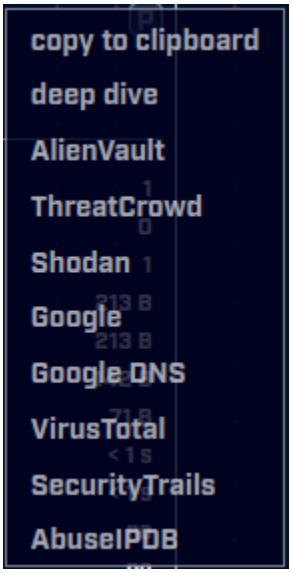

Threat Intel Analysis

To see if any connections occurred with systems that appear on one or more threat intel feeds, click the "threat intel" button at the bottom of the screen.

View 1 (the default view) shows connections between internal systems and external IP addresses that appear on one or more threat intel feeds.

The layout of the threat intel screen is a bit different than the screens discussed previously.

Results Feed

On previous screens, the Results Feed on the left hand side of the screen identified the source of the connection being reported. However, on view 1 of the threat intel screen the Results Feed is displaying all threat IP addresses that have had connections with internal IPs, whether the threat IP initiated the connection or received the connection. This makes sense when you consider that we are trying to analyze which of the target IP addresses have appeared on a threat intel feed.

The functionality of the results feed is similar to that which was discussed on previous screens. By default, higher priority addresses are listed at the top. The default sort order is "Total Bytes Exchanged". You can change the sort order by clicking the “Total Bytes Exchanged” label near the top of the Results panel. Below this is a dialog box you can use to search for occurrences of specific IP addresses.

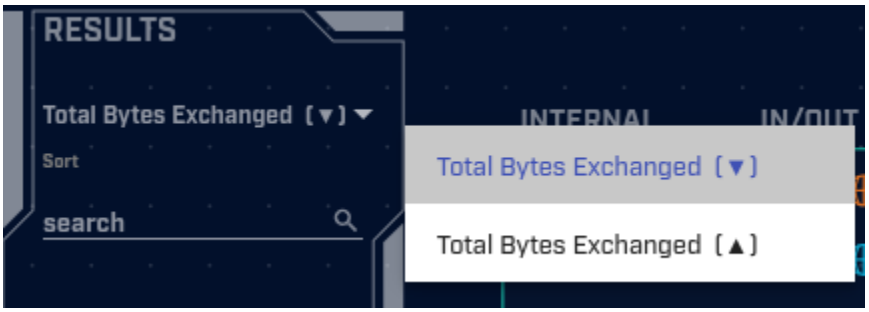

DNS Analysis

DNS can be used as both a covert communication channel, as well as a way to exfiltrate data out of a network. Because DNS is such a noisy protocol, it tends to have minimal logging

enabled. Combine that with the fact that most environments permit DNS out of their environment, and it makes a good choice for hiding suspect traffic patterns in plain sight.

To analyze domain name traffic observed on your network, click the "dns" button on the bottom of the screen.

This will produce the DNS A and AAAA record analysis screen, which will appear similar to the following:

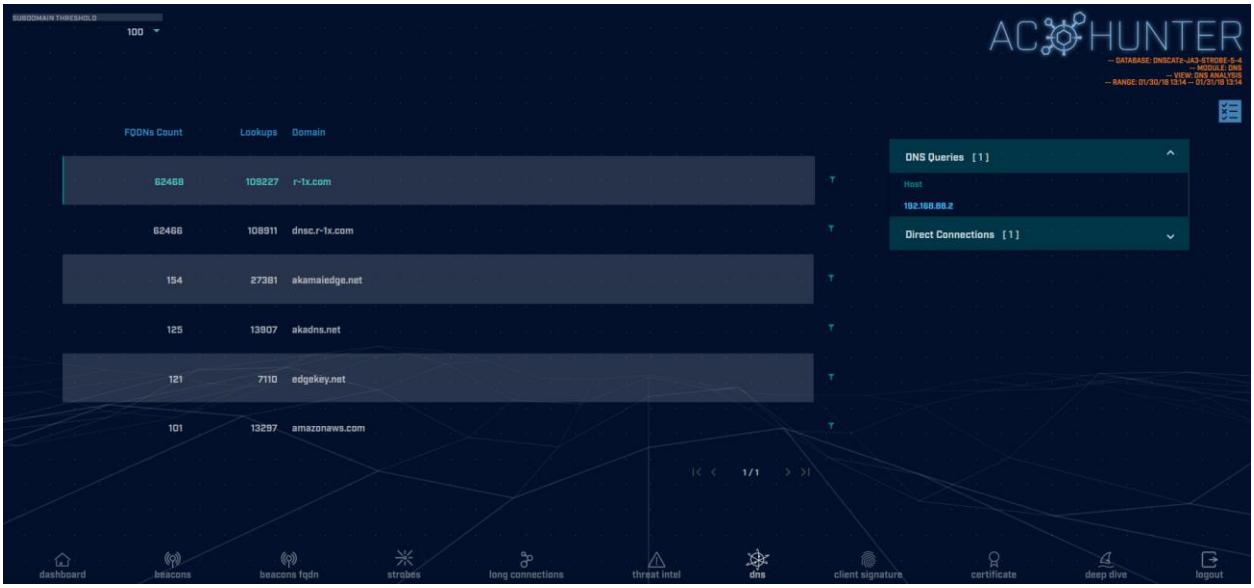

This screen reports summary information regarding the external hosts being contacted by your internal systems. For example, in the sample screen capture above there were 62,468 subdomains of r-1x.com.

This view is a powerful tool for routing out covert communications taking place over the DNS protocol. For example, consider the data in the above example for the "r-1x.com" domain.

In this output, we see that DNS lookups were performed on 62,468 different host names within the "r-1x.com" domain. While it is normal to see a few dozen lookups within most domains, and even 500-800 in extremely popular and well-known domains such as "microsoft.com" or "amazon.com", this is an excessively high number of hostname lookups, especially for a somewhat obscure domain name.

Compromised systems need to "call home" in order to get their marching orders. Some strains of malware perform this task by generating DNS lookups. The remote name servers (in this case the name servers for "r-1x.com") are actually acting as command and control (C&C) servers. The FQDN being queried is actually an encoded message to the C&C servers. To ensure the local resolver forwards the request and does not hand back cached information, the compromised system will vary the FQDN being queried with each request. This results in an excessive number of FQDNs being queried within a domain, as seen in this example.

The "Subdomains Threshold" option in the top left corner can be used to set a minimum display threshold. For example, you could choose to display only domains that have 100 or more FQDNs associated with them.

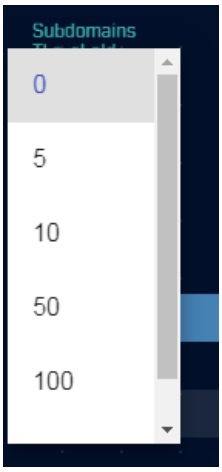

Client Signature

The client signature module is used to identify systems on your network that communicate in a unique fashion. While a unique signature is not always a telltale sign of a C&C channel, it occurs often enough that it is worth verifying. To launch the client signature module, click the client signature icon at the bottom of the screen.

View 1 - User Agent Strings

"User agent" names are sent as part of HTTP requests; they identify the browser or tool making the web request. Many environments maintain standards for both server and end user computer configurations. This will cause the user agent field to be consistent across those platforms. Unique user agents can be interesting in that they may identify a non-standard tool or browser being used on the network.

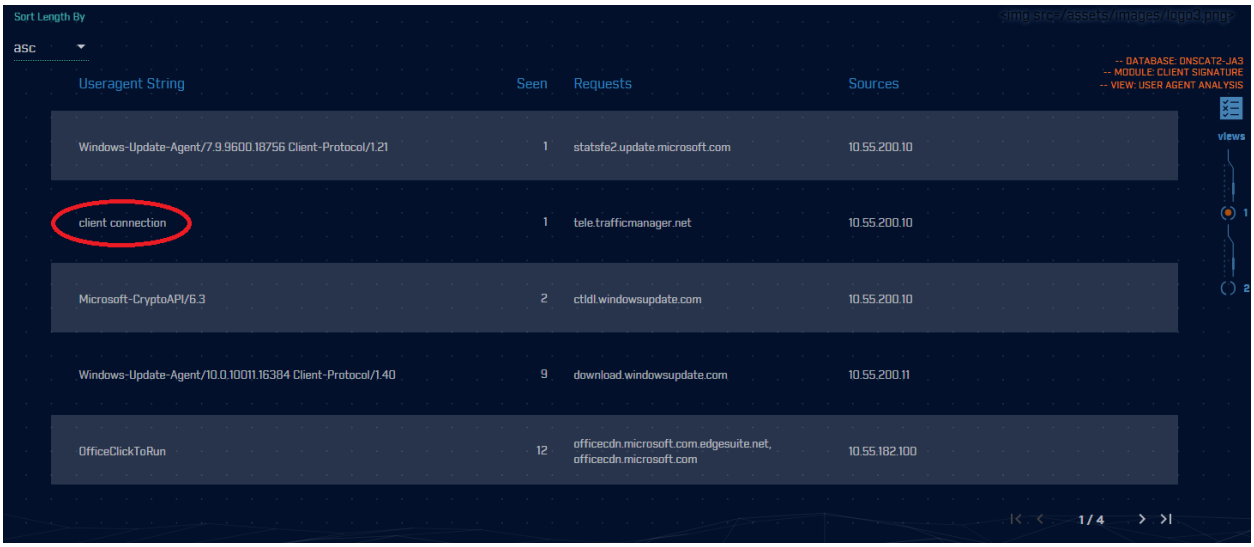

The default sort option shows the most unique user agents first. The "Sort By" option in the top left can be triggered so that the least unique appears first.

The left-hand column displays the user agent identifier. In addition to showing the type of tool that made the request, it can include hints to the operating system, processor architecture, and particular program versions being used.

The "Seen" column identifies the number of unique source IP addresses that have been observed using the specified user agent. The "Requests" column shows where these requests were sent. Finally, the "Sources" column shows the internal IP address from which the requests came.

View 2 - SSL/TLS Hash

Switching to View 2 under client signate will bring up the Ja3 SSL/TLS Hash analysis screen. This provides a similar analysis as the user agent view, except it is used for HTTPS connections rather than HTTP.

SSL/TLS Hash	Seen	Requests	Sources
5e573c9c9f8ba720ef9b18e9fce2e2f7	1	clientservices.googleapis.com	10.55.182.100
bc6c386f480ee97b9d9e52d472b772d8	2	clients4.google.com, 556-emw-319.mktoresp.com	10.55.182.100
f3405aa9ca597089a55cf8c62754de84	2	builds.cdn.getgo.com	10.55.182.100
28a2c9bd18a11de089ef85a160da29e4	2	mediaredirect.microsoft.com	10.55.100.105, 10.55.182.100
08bf94d7f3200a537b5e3b76b06e02a2	4	files01.netgate.com	192.168.88.2

HTTPS uses SSL or TLS to both authenticate and encrypt data sessions. These start with the client sending an SSL hello packet to the server, which identifies which authentication and encryption methods the client prefers to use. Similar to the user agent field, if you are running standardized software within your environment, these SSL hello packets should be consistent across all of your system. As an example, if you have standardized on using Mac OS X with a Chrome browser, a hash of the SSL client hello from every system should be identical. So again, similar to the user agent field, unique systems may be an indicator of non-standard software and should be investigated.

The sort option functions identically between views 1 and 2. The additional displayed data fields are also identical.

To only see client signatures for a particular IP address, enter that address in the search box in the upper left:

Certificate

The certificate module checks the digital certificate being issued by servers during outbound HTTPS connections. To launch this module, click the "certificate" icon at the bottom of the screen.

Currently the certificate module does a simple check to see if the certificate is valid. We will be expanding these checks over the next few releases. While an invalid certificate is not a

guarantee that a C&C channel has been established. If it is observed in combination with beaconing or long connection traffic it should be judged suspiciously.

To only see certificate issues for a particular IP address, enter that address in the search box in the upper left:

Deep Dive

While the other AC-Hunter modules focus on a specific threat vector (beaconing, long connections, DNS tunneling, etc.), the deep dive module is designed to help assess the threat of a specific system. Let's say that while you are reviewing one of the other modules, you identify an internal system that is acting suspiciously, but you are unsure if the system is safe or a threat. The deep dive module will show you all communications associated with that system so that you can make a more informed threat assessment. Further, let's assume that you detect suspicious activity from an internal system to an external IP address, and you want to quickly assess if any other internal systems have been in contact with that external IP address. The deep dive module can quickly summarize all internal systems that have communicated with that external IP address.

To access the deep dive module, click the deep dive icon at the bottom of the screen:

You will first be prompted for the IP address you wish to investigate. As you type in the IP address, a drop down list of possible IP addresses will be presented. You can click the IP address you wish to investigate at any time, rather than having to type in the full address.

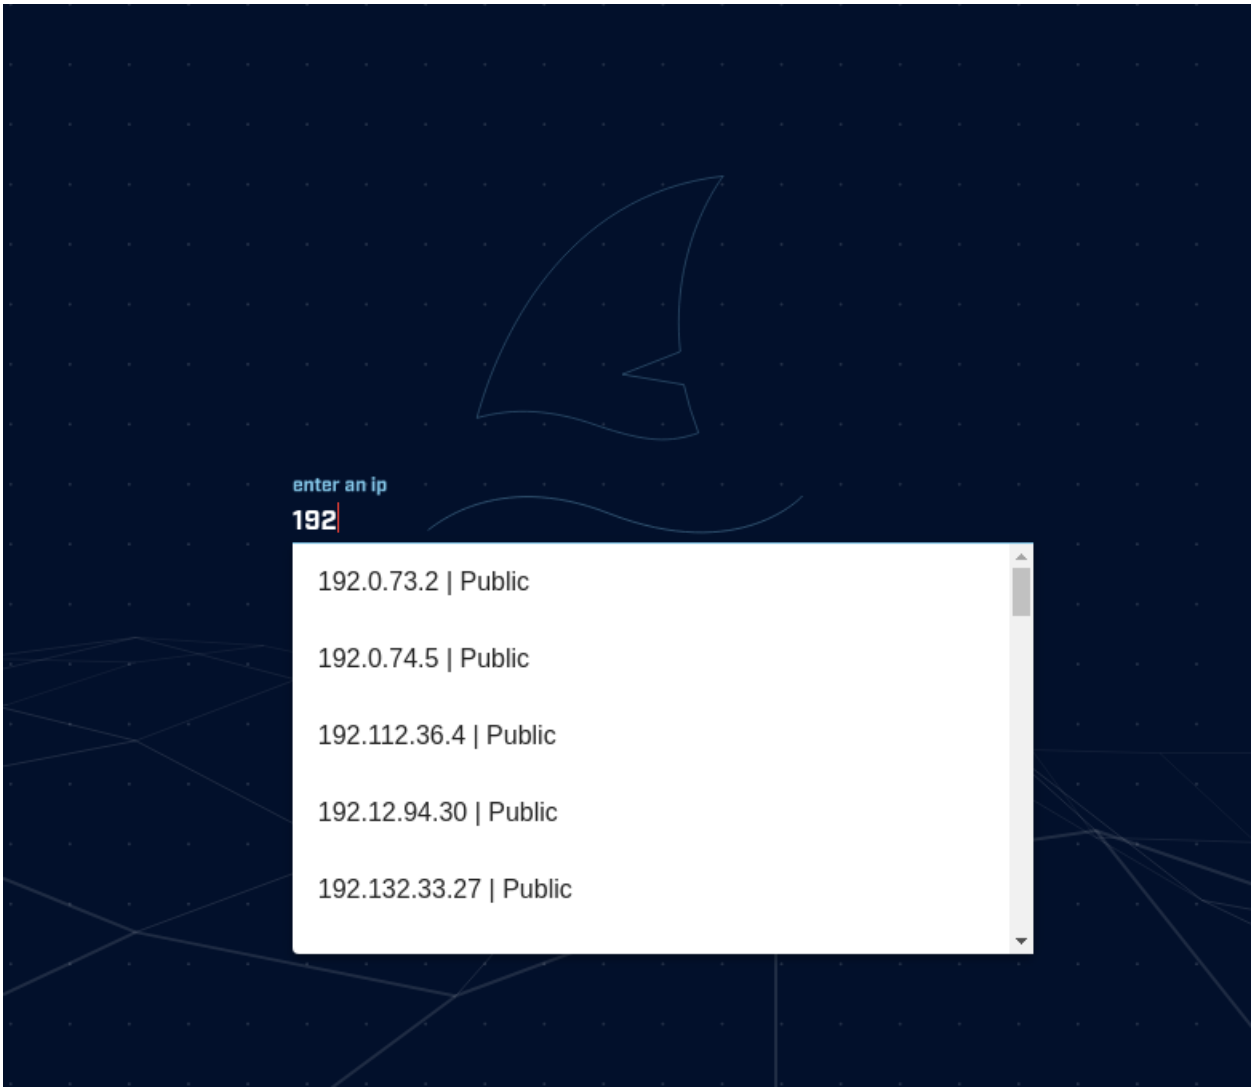

Once you select an IP address, the deep dive main screen will be loaded.

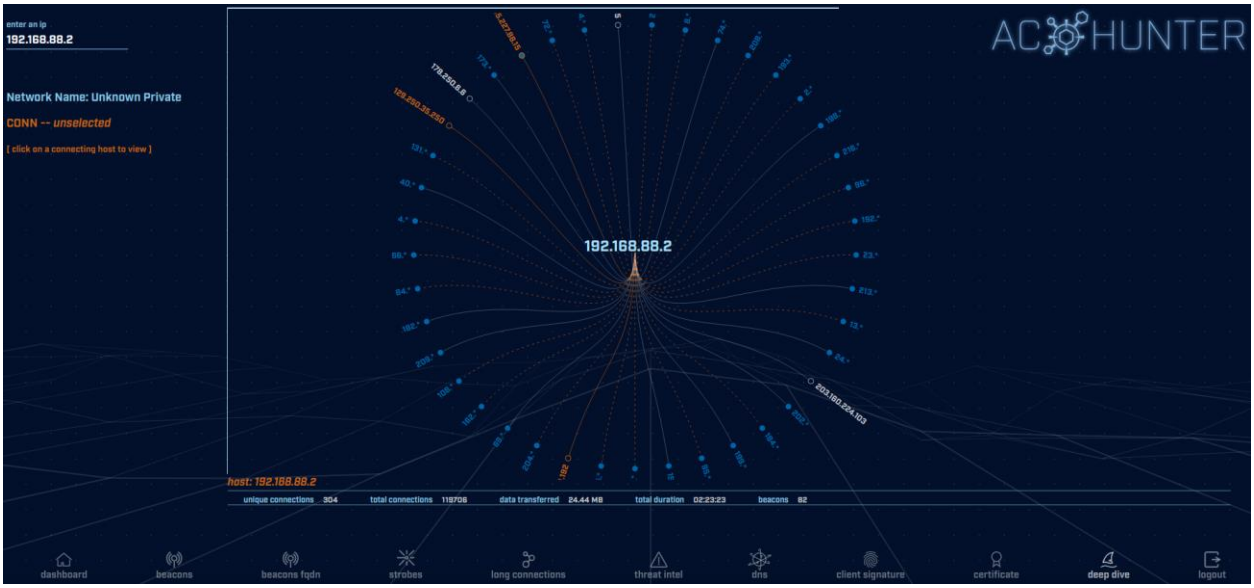

The IP address you entered will appear in the middle of the graphic. All of the connections protruding out from it are systems with which it communicated.

Just below the graphic, you can see overall summary information regarding the system being investigated.

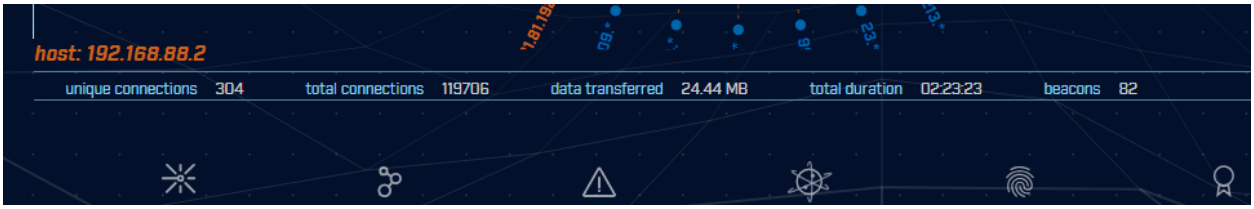

This includes data points such as the total number of connections the system created, the amount of data transferred, and the time spent with active connections open.

You can manipulate the graphic by clicking your mouse and dragging it around the screen. You can also zoom in and out by using the mouse wheel.

You will notice that a number of patterns are used to express the connections taking place.

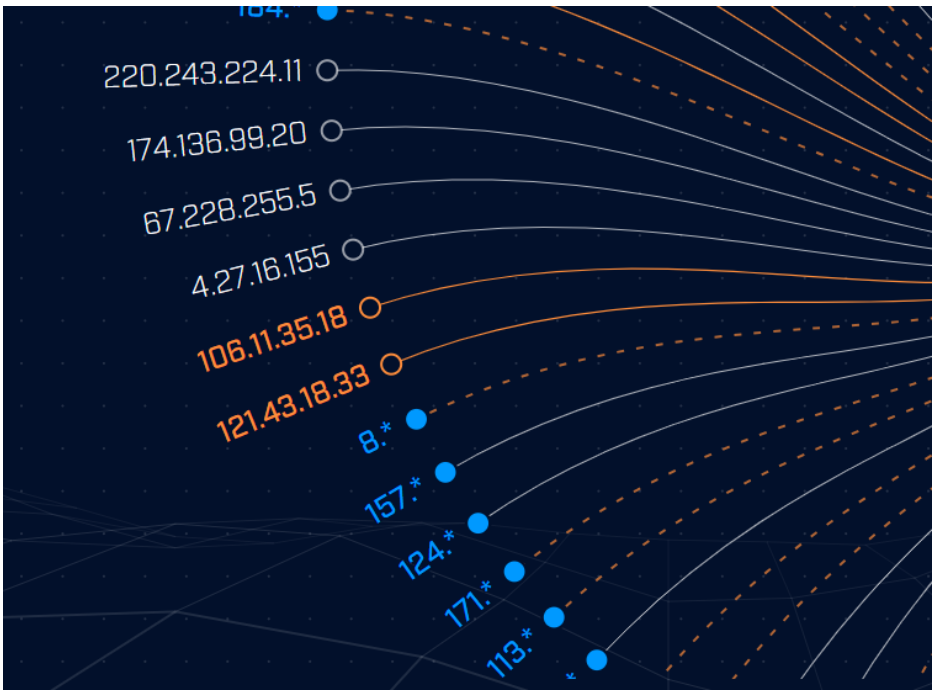

A white line indicates that no suspect activity was detected within communications with the specified target. An orange line indicates that one or more beacons were detected. A dashed orange line indicates that there is a mix of normal and beacon activity. The circle at the end of the line tells you if there is additional data available when clicking through. A filled in circle indicates you can click through for more data. An empty circle indicates this is the final data point. For example, if I click on "8.*" in the above example, my screen expands out any additional information:

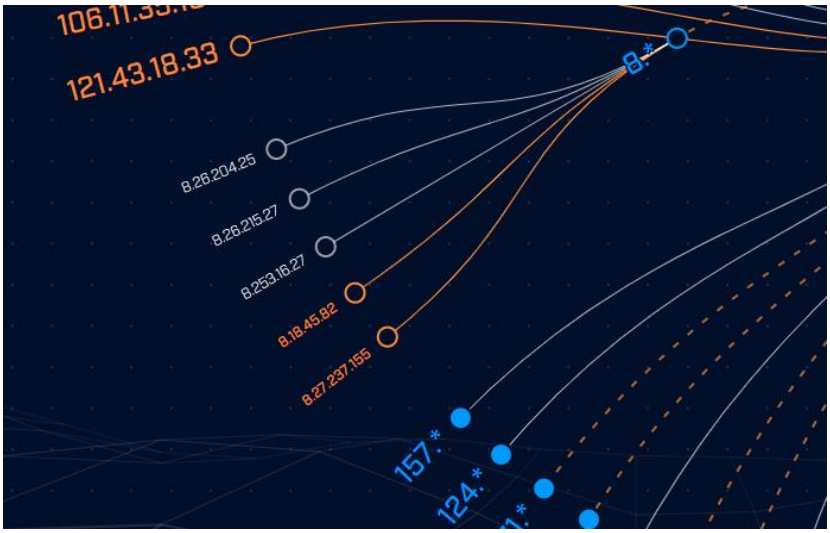

When an IP address is clicked, the left hand side of the screen updates with information regarding communications with this host:

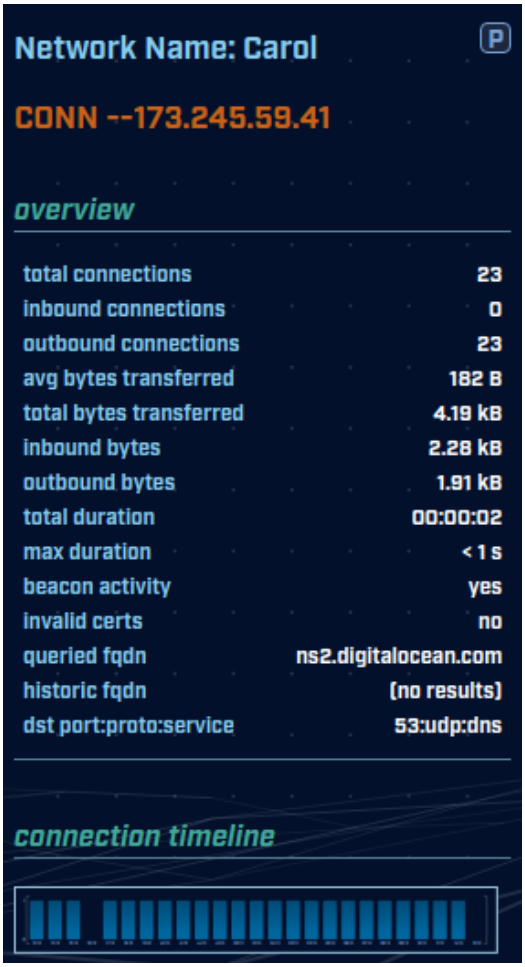

The "overview" section is a summary of all communications between the IP address I am researching and the IP address I just clicked. I can see the number of connections, bytes transferred, duration of all sessions combined, and whether beacon or invalid digital certificates were observed.

The connection timeline is a graph of the number of connections that took place each hour over a 24 hour period of time. If I mouse over the timeline it will expand across the screen so that it’s easier to read the values on the X and Y axis.

If this conversation looks malicious, then it's worth investigating the external IP address (165.227.88.15) some more. The first question: Did that system communicate with any other systems at my end? We make that easy; click on the "P" (Pivot) icon to the right of the IP address. Deep dive immediately switches view from the original where we see all conversations with 192.168.88.2 to one where we see all conversations with 165.227.88.15. As you might expect we still have the conversation with 192.168.88.2, but thankfully there aren't any other machines that have talked to this external IP address:

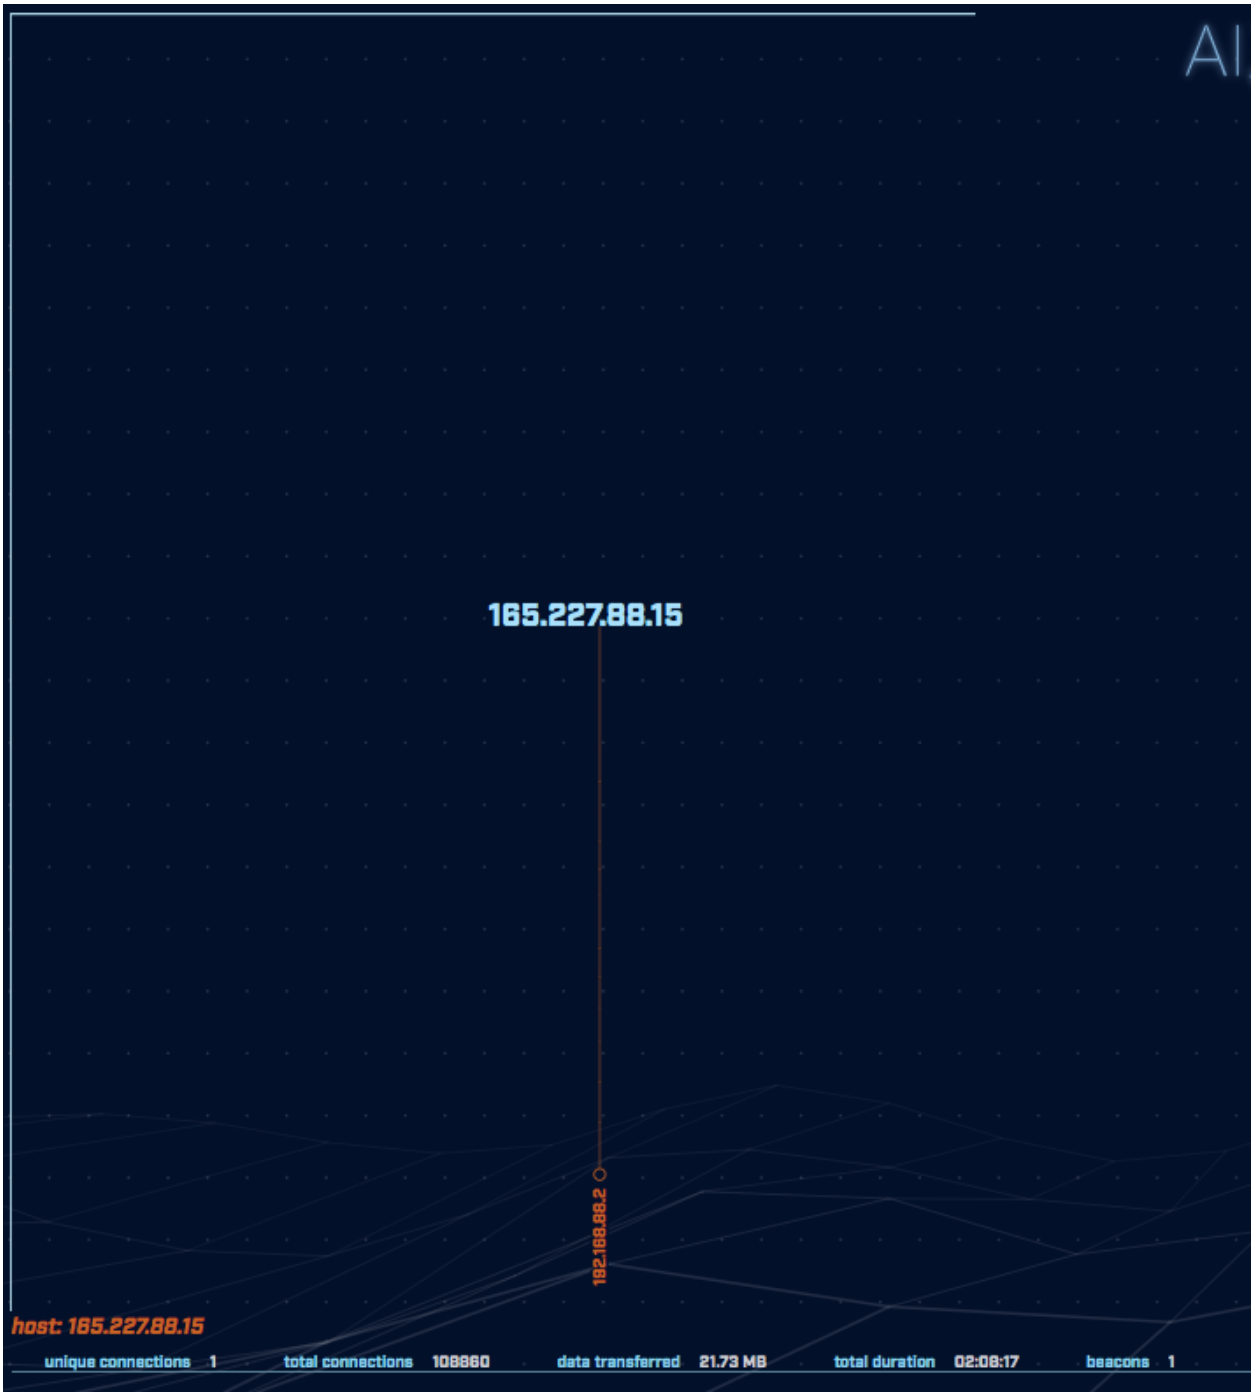

Note that if the system being evaluated has connected to a lot of systems, it is possible that this graphic can get pretty busy. Here's an example:

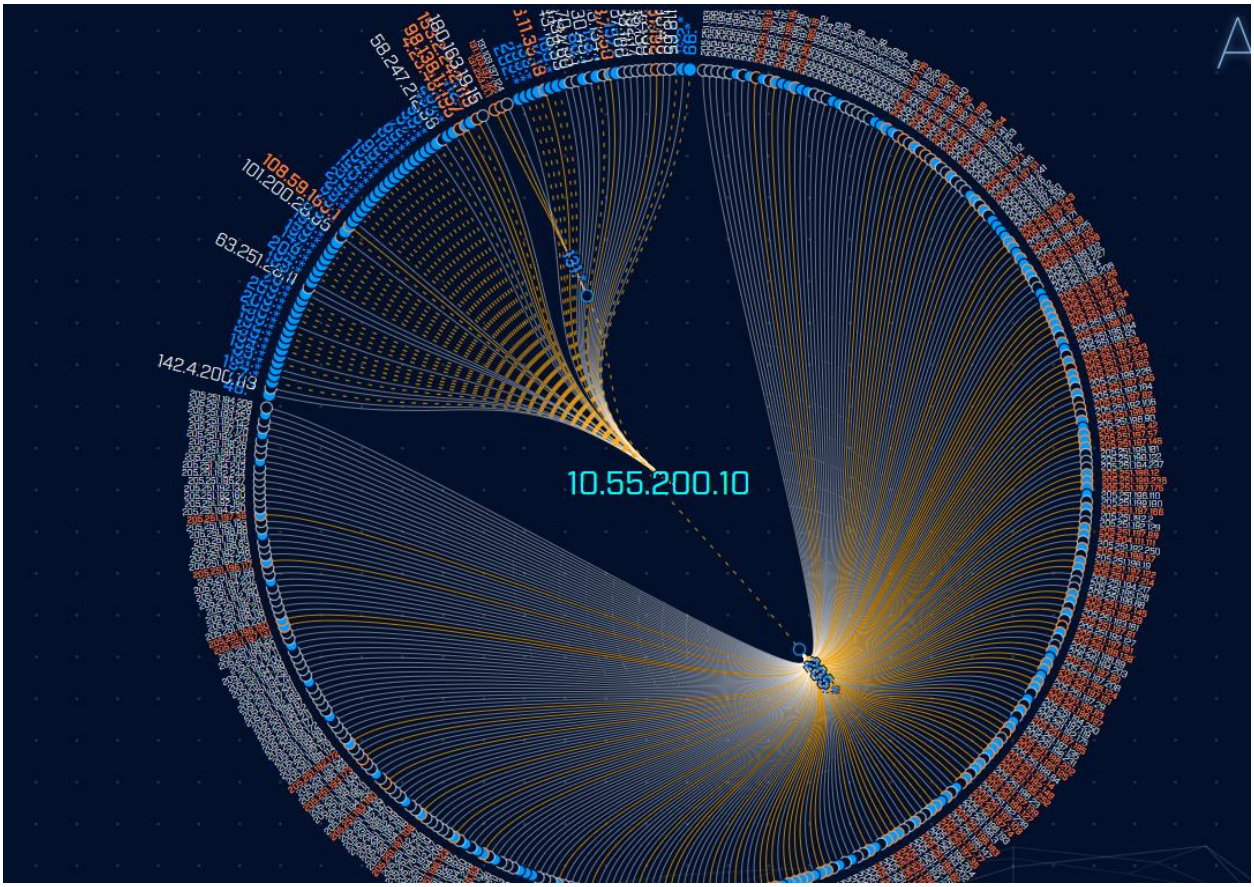

Note that when the blue "205.*" address was clicked, many IP addresses appeared below it. The graphic updated to move all unassociated IP addresses into a confined space, so that maximum space was available for the subnet under review. This ensures that all possible targets are fully visible.

Changing the display Theme

In the same Gear menu (Settings) you can pick which display theme to use. The default theme "Game Mode (dark)" places light text and graphics on a dark background.

We also offer "daVinci Mode (light)", which switches to a light background with darker text and graphics when you press Confirm. Give both a try and see which you find more readable.

Logout

The logout button is used to disconnect from AC-Hunter when you are done using the system. Clicking this button will return you to the login screen.

https://www.activecountermeasures.com

ctf@activecountermeasures.com