



Table of Contents

[Table of Contents](#)

[First Time Use](#)

[Logging in to AC-Hunter](#)

[Dataset Selection](#)

[Datasets that are Being Filtered](#)

[Dashboard](#)

[Beacons Analysis](#)

[Results Feed](#)

[Searching Results](#)

[Source and Destination Analysis](#)

[BeaKer Integration](#)

[Safelisting](#)

[Timeline Analysis](#)

[Connection Frequency Chart](#)

[Interval Scoring Chart](#)

[Data Size Analysis](#)

[Data Size Frequency](#)

[Data Size Scoring Chart](#)

[Beacons web Analysis](#)

[Beacons Proxy Analysis](#)

[Strobes Analysis](#)

[Long Connections Analysis](#)

[Threat Intel Analysis](#)

[Results Feed](#)

[Timeline Chart](#)

[Host Chart](#)

[Threat Intel Feeds Used](#)

[DNS Analysis](#)

[Client Signature](#)

[View 1 - User Agent Strings](#)

- [View 2 - SSL/TLS Hash](#)
- [Cyber Deception](#)
- [Deep Dive](#)
- [Active-Flow Databases](#)
- [Espy Databases](#)
 - [Differences Between Zeek-sourced and Espy/Netflow-sourced Data](#)
- [Active-Flow Installation](#)
- [Espy Installation](#)
- [Investigation Sources](#)
- [Canary Token Setup for Cyber Deception](#)
- [BeaKer Installation](#)
- [Managing Databases](#)
 - [Importing Packets from a PCAP File](#)
- [Managing Safelists](#)
- [Changing the Display Theme](#)
- [Modifying the Sensor Name](#)
- [Configuring User Accounts for the Web Interface](#)
 - [Managing Internal AC-Hunter User Accounts](#)
- [System Maintenance](#)
 - [Log Maintenance](#)
 - [Deleting Zeek Logs](#)
 - [Deleting RITA Logs/Databases](#)
 - [More information on log and database management](#)
- [Logout](#)
- [References](#)

Greetings and welcome to the AC-Hunter™ User Guide. This guide will walk you through the process of connecting to and using AC-Hunter for the first time. Please refer to the Install Guide if you have not yet installed the software. If you have any questions or problems, please do not hesitate to contact support@activecountermeasures.com.

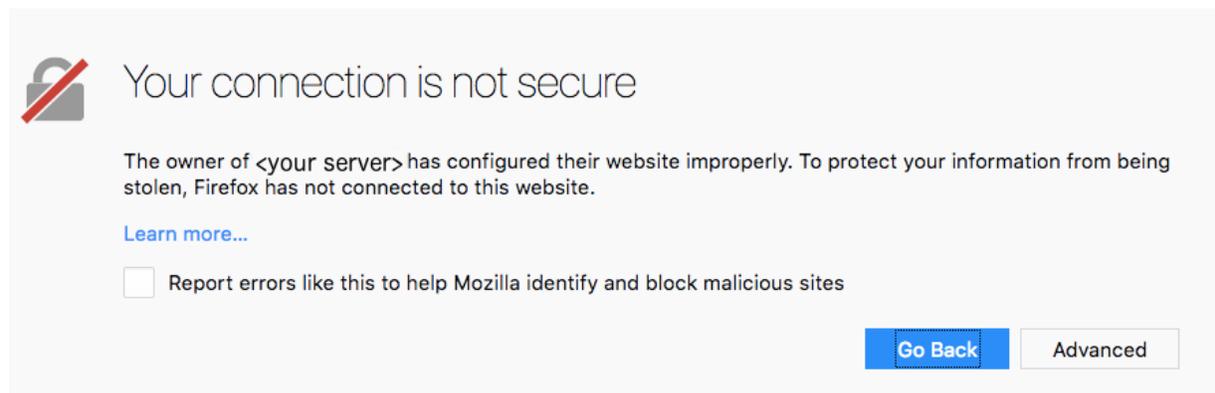
First Time Use

The following is a brief introduction to using AC-Hunter. It is not designed to be an in-depth explanation on how to perform threat hunting. This guide will introduce you to the AC-Hunter interface and describe how to interpret the data being presented. For a deeper look at using AC-Hunter for threat hunting, please see the various [blog entries](#) and [videos](#) we have created.

Logging in to AC-Hunter

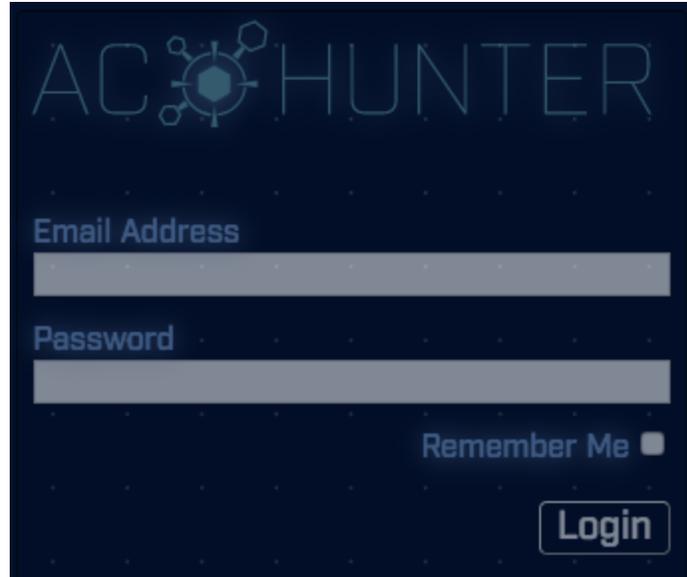
Connect to the AC-Hunter server using your Web browser (we recommend Chrome):
`https://<AC-Hunter server name or IP address>`

When you first connect to the system, you may see an error similar to the following:



By default AC-Hunter uses a self signed digital certificate. Since the validity of the certificate cannot be checked by your web browser, an error similar to the above example is produced. If the AC-Hunter server is connected to your internal network, and is not directly accessible from the Internet at large, it is safe to create an exception for this server. If you prefer, you can have a third party certificate authority generate a digital certificate for the system which will remove the above error. Our [FAQ on replacing this certificate](#) covers the steps in more detail.

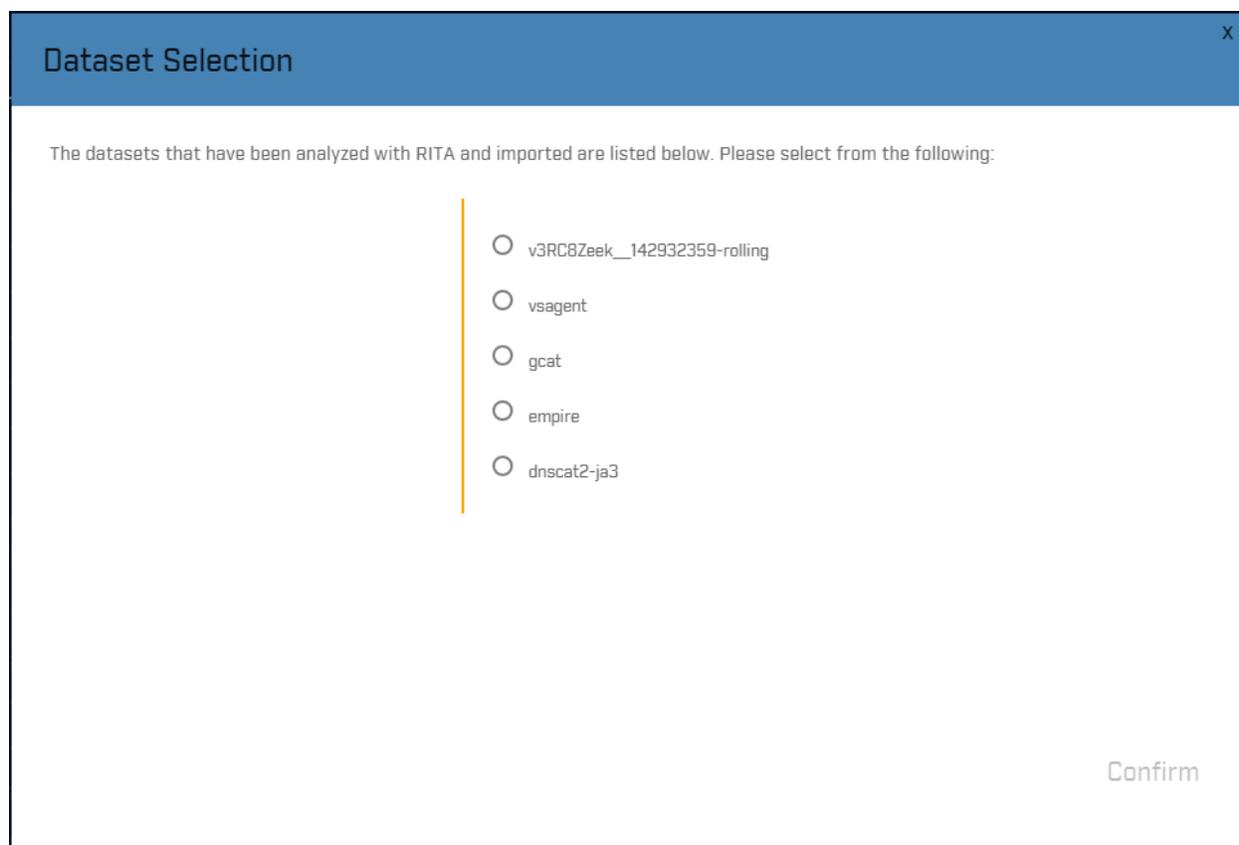
Once you clear the certificate error screen you will see the AC-Hunter login screen.



Enter the login name and password you created while installing AC-Hunter. Press the Login button to authenticate to the system.

Dataset Selection

When you login to AC-Hunter for the first time, there will be no default dataset selected. This will cause the Dataset Selection window to appear (you can go here manually by clicking the gear icon on the dashboard tab).



Right after install, you will have two types of datasets available. The first will have the word "rolling" at the end of its name. This is the dataset that will always include the last 24 hours worth of data. If you just finished installing AC-Hunter, and you try to select this dataset, you may see the loading icon spinning but not data being presented. This is normal as the system needs to run for a minimum of two hours (longer preferred) before any data can be presented. Permit the system to run for a while and come back to check this dataset later.

The other datasets are examples of command and control traffic. This is sample data that will give you a chance to experiment with AC-Hunter while the system does the first threat hunt of your network. 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.



Live data from your network will appear on the system after a few hours. After 24 hours, AC-Hunter will begin archiving the data it has analyzed. This permits you to go back in time to review systems as required. The naming convention for the archive datasets includes the name of the system that submitted the data and the date stamp associated with the submission. This makes it easy to pinpoint the exact dataset you need to review.

Datasets that are Being Filtered

When you are picking a data set, you may see a note below the database name that says "Safelisting in progress." In the background, AC-Hunter is applying safelist entries to the database. You're still welcome to select it and work with it just like any other database, but be aware that recent safelist changes may not have been applied yet (you may see records that should have been filtered by a recent safelist addition, or you may not see records that should come back to view after a recent safelist removal).

The warning note may show up in multiple places:





Dashboard

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

The dashboard features a central 'THREAT RATING' bar at the top, currently set to 100. Below this, there are two main data panels:

- Hosts [15]:** A table listing suspicious IP addresses and their threat scores.
- Threat Activity:** A table detailing various threat indicators and their associated values and points.

Hosts [15]	Network Name	Threat Score
192.168.88.2	Unknown Private	332.01
10.55.100.111	Unknown Private	140.57
10.55.100.105	Unknown Private	131.00
10.55.100.100	Unknown Private	128.45
10.55.100.107	Unknown Private	128.45
10.55.100.110	Unknown Private	127.43
...

Threat Activity	Value	Points
Beacon Score [strongest signal seen]	99.90%	99.90
Beacon FQDN Score [strongest signal seen]	82.90%	82.90
Beacon Web Score [strongest signal seen]	0.00%	0.00
Beacon Proxy Score [strongest signal seen]	0.00%	0.00
Longest Connection	02:06:34	2.11
Threat Intel Connections [Outgoing]	1	10.00
Threat Intel Connections [Incoming]	1	10.00
Threat Intel Average Bytes	199.82 B	0.00
Too Many FQDNs Per Domain	62468	100.00
Unexpected Protocol on Well Known Port	0	0.00
Total		332.01

The panel on the left side of the screen lists suspicious IP addresses, sorted by threat score, with the highest score listed at the top. This is effectively your action item list, as the system most in need of a deeper investigation will be 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 identifies which threat vectors were detected and how they were combined to create the threat score for the selected IP address. For example, an internal system generating a strong beacon signal will be scored lower than a system with a strong beacon signal to an IP address that has appeared on a threat intel feed and appears to be moving a lot of data.

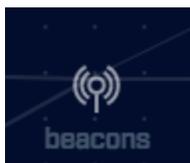
If you want to perform a deeper analysis on any of the threat activity, you can click the listed threat activity on the right side of the screen. This will load whatever module can be used to analyze the threat. For example, clicking on "beacon score" will automatically load the beacon module. The search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

Beacons Analysis

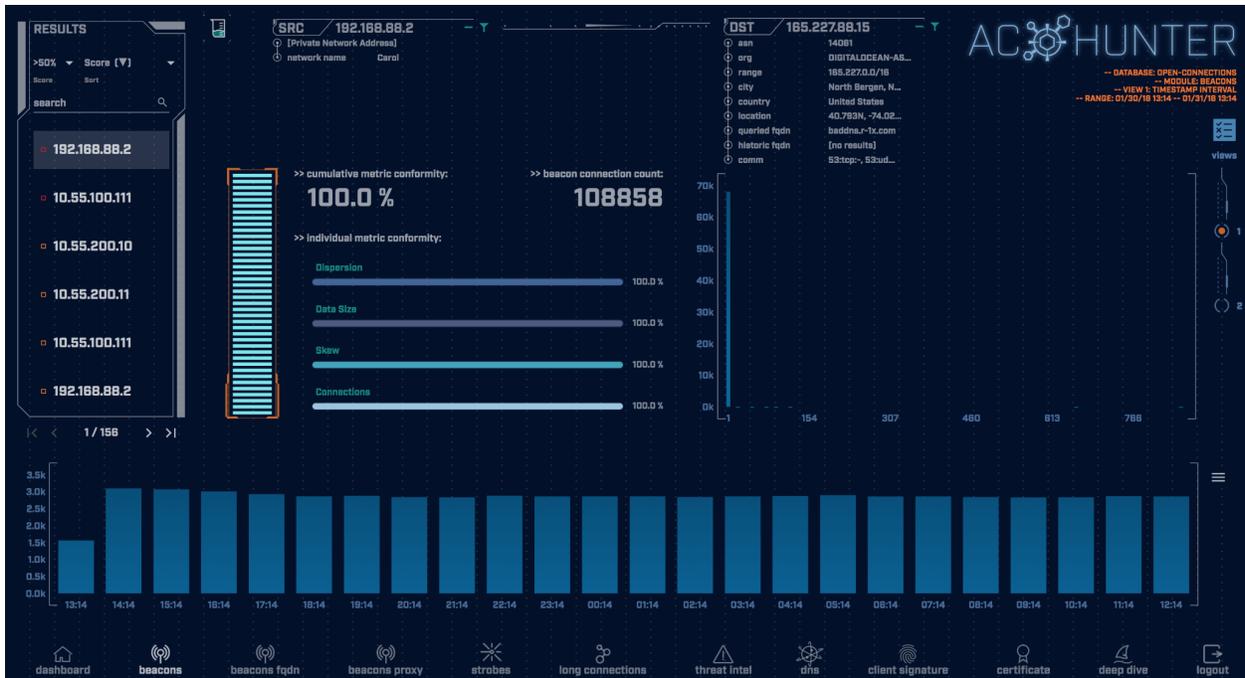
To check for internal systems that are exhibiting beaconing behavior, use the "beacons" module.

AC-Hunter has three ways to identify Beacons - regular traffic between one of your systems and an external IP address, regular traffic between one of your systems and a hostname, or regular HTTP traffic between one of your systems and the Internet that's funneled through a proxy. 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.

This section of the document will focus specifically on the "Beacons" module, which analyzes traffic from a source IP to an external IP. Many of the features overlap between each beaconing module. Please refer to the other beaconing module sections below for key differences in the other modules.



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.



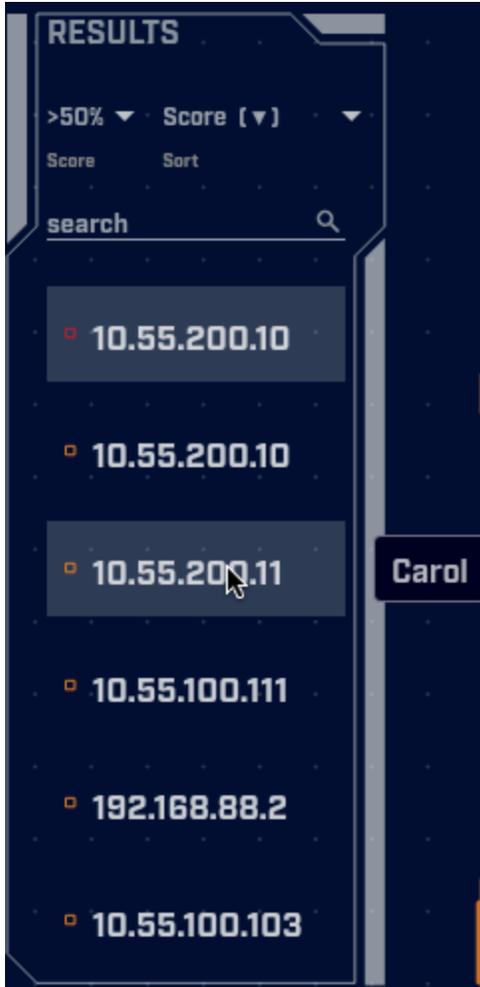
If however, you click the "beacons" icon on the bottom of the screen and enter the module that way, no filter will be implemented. The result is that the top listed IP address will be the source IP that generated the highest beacon score, regardless of its final threat score. This is useful when you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

Beacon results are displayed based on IP pairs. It's 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

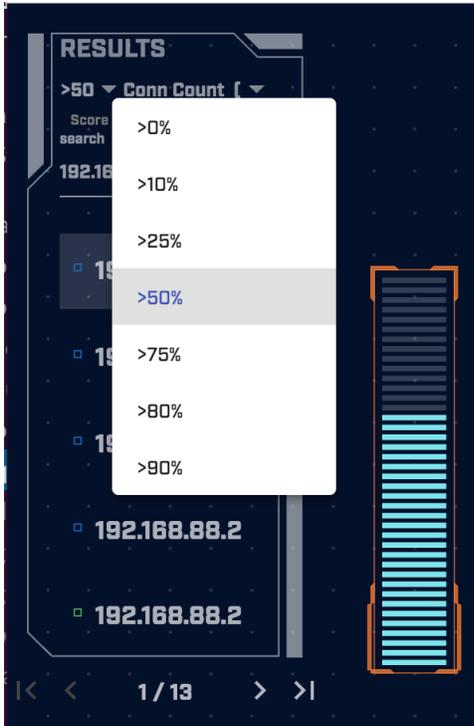
The results feed identifies all internal IP addresses where beacon activity was detected. As mentioned above, it's possible for a source IP address to be listed multiple times if beacon activity was detected with multiple destination IP addresses on the Internet.

When hovering over each ip, that host's name will appear in a tooltip. This makes it easier to identify which internal device is using said IP address..

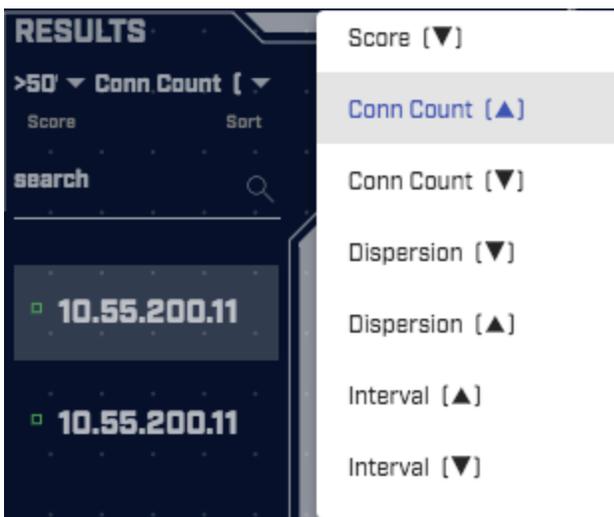


Each identified connection is preceded by a colored icon. These can be red, orange, yellow, blue or green. The color code provides a quick visual indicator of the likelihood of the system exhibiting beaconing behavior. Red is most likely while green is least likely.

The ">50 v" immediately under "RESULTS" allows you to set a minimum threshold; you'll 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.

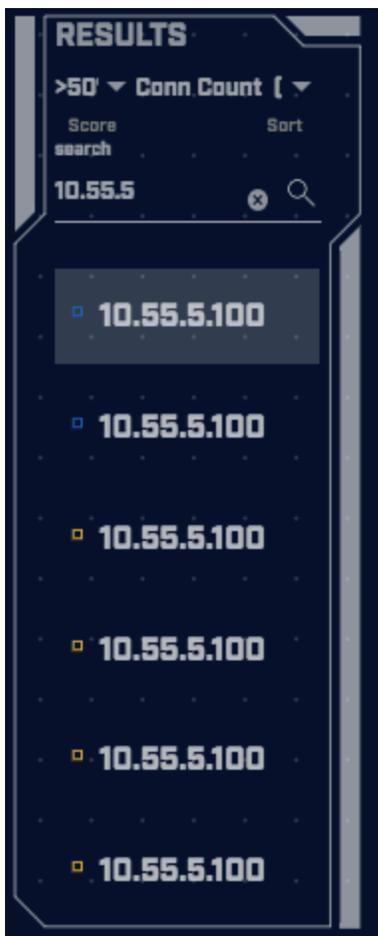


The default is to sort the list based on the score (more on this later) from highest to lowest. This will automatically place the systems most likely to be beaconing at the top of the list. Another useful sort option is by connection count as compromised systems tend to generate thousands of beacon signals.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

When the results feed displays more than one connection, clicking on an IP address in the list will update the other portions of the screen so that information regarding that connection can be analyzed. You may see the same source IP address appear multiple times in the results feed. This is because AC-Hunter analyzes source and destination IP connection pairs. While each entry may have the same source IP address, the destination IP addresses will be different.

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:



On the Beacons Web, Beacons FQDN, Beacons Proxy, and Client Signature tabs you can search by IP address or hostname. On the DNS tab you can search by hostname.

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:



Please see [Customizing Investigation Sources](#) for information on how to add your preferred lookup tools to this drop-down list.

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.

In some cases, Fully Qualified Domain Name (FQDN) information is displayed. The FQDN information can appear on one of two lines: the "Queried fqdn" line, or the "Historic fqdn" line. If an FQDN appears in the "Queried fqdn" line, then it means this source made either an A or AAAA query for this FQDN and received the destination IP address as a response. 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 the actual host name the system queried can be a valuable data point. The “Historic fqdn” line shows FQDNs that were queried by other systems in the dataset. The other systems received the destination IP on this page in response to their FQDN queries.

Clicking an FQDN in either line will open a menu with one or more options. If the row you clicked has more than one FQDN, then the menu will open with all of the available FQDNs. Clicking on one of the FQDNs will open a menu with investigation sources that you can use to investigate the FQDN with:



Several investigation sources may appear in the menu, commonly “whois”, “URLScan”, and “copy to clipboard”. The first two are external websites you can use to research the hostname, and the last copies the hostname to the clipboard (as shown above).

Below the FQDN lines 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.

BeaKer Integration

You will see the BeaKer icon between the results feed and the source and destination details.



If you have a BeaKer install, then clicking the icon will open a new browser tab (similar to how the investigation menu works). This new tab will load details for the connections between the selected source and destination IP addresses in the BeaKer dashboard. This information will include the processes that generated the connections. You can also find a quick overview on interacting with BeaKer [here](#).

Note: If you do not have a BeaKer install then clicking the BeaKer icon will have no effect.

Safelisting

Let's say in the course of your analysis, you determine that the beaconing behavior is actually normal. For example, the connection may simply be a system verifying the current time with a known NTP server. AC-Hunter includes the ability to "safelist" communications in order to ensure they no longer appear in the results feed.

To create a safelist, click the filter icon to the right of the IP address you wish to safelist.



This will produce the safelist management window.

Safelist this Entry?

SRC	Safelist by IP Address
DST	View/edit your full safelist in Home > Settings > Safelist.
PAIR	<p>Whitelist ...</p> <p><input checked="" type="radio"/> 165.227.88.15</p> <p><input type="radio"/> 165.227.88.0/24</p> <p><input type="radio"/> 165.227.0.0/16</p> <p><input type="radio"/> 165.0.0.0/8</p> <p><input type="radio"/> Organization Name [DIGITALOCEAN-ASN]</p> <p><input type="radio"/> Organization ASN [14061]</p>
DOMAIN	<p>For all connections where it is the ...</p> <p><input type="radio"/> Source IP [initiator of connection]</p> <p><input checked="" type="radio"/> Destination IP [receiver of connection]</p> <p><input type="radio"/> Either</p>
	<p>Comment</p> <hr/>
	<p>Cancel Safelist</p>

Based on your selection, AC-Hunter will default to a number of specific values which you can choose to override. For example you can choose to apply the safelist to just that specific IP address, or to all IP addresses within a class A, B or C range. This is helpful when you may have a group of systems all on the same subnet performing the same function. You can choose to apply the filter to all IP addresses that are part of that specific organization (Amazon, Google, Microsoft, etc.), or to just the IP addresses that are part of the current Autonomous System Number (ASN). Filtering by ASN is handy when you only want to trust a portion of an organization's network. For example, you may want to safelist Microsoft patching servers but not systems that are part of Azure.

At the bottom of the screen, you have the option to record comments with this safelist. We highly recommend doing so, as this can simplify management down the road. For example, a year from now you may not remember why this exception was created, or who created it. By documenting this information in the comments section, you always have that data handy.

By selecting the "PAIR" tab, you can also safelist all conversations between these specific IP addresses:

Safelist this Entry?

SRC

DST

PAIR

DOMAIN

Safelist by pair of IP addresses

View/edit your full safelist in Home > Settings > Safelist.

Safelist Connections Between the Following IP Pair:

192.168.88.2--165.227.88.15

192.168.88.0/24--165.227.88.0/24

Comment

Cancel Safelist

This is useful when you have two systems whose normal traffic shows up as beacons, long connections, or any other type of alert in AC-Hunter. Note that safelisting an IP pair safelists *all* traffic back and forth between the systems, not just the specific traffic you were looking at.

The second radio button allows you to safelist between two IP *networks*, the Class C (256 address) networks around each endpoint. As with all network safelisting, you should consider whether the hosts around the endpoints are equally trusted; this would not be the case at a cloud provider, for example.

Once you make a selection, the safelisting will automatically apply in the background. However, the larger the block you safelisted, the longer this background processing will take. It could take several minutes for the safelisting to apply if you select a large organization such as "Amazon, Inc" or "Microsoft Corporation". This may slow down the update speed of the screen.

When a safelist is applied to a dataset, you will see a large filter icon appear between the Results list and the source IP address:

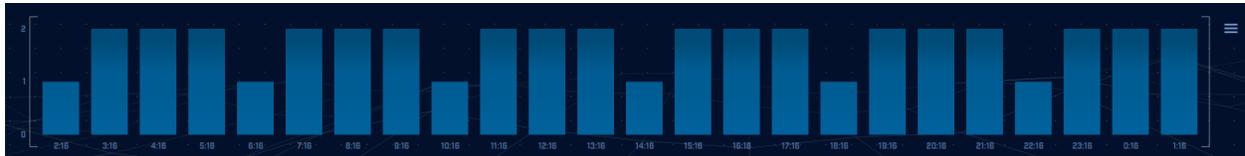


This icon is to remind you that you are only looking at a subset of the actual data for this module, as some entries are being safelisted or blocked from view.

You can learn more about managing safelists in the "[Managing Safelists](#)" section of this document.

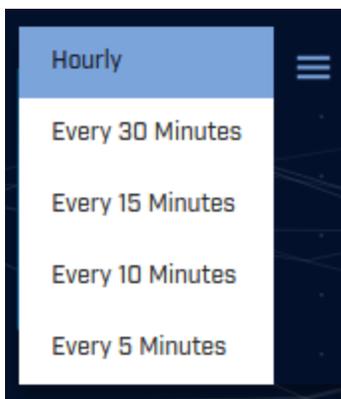
Timeline Analysis

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. This may sound confusing, but it's actually pretty straight forward. Let's look at an example.

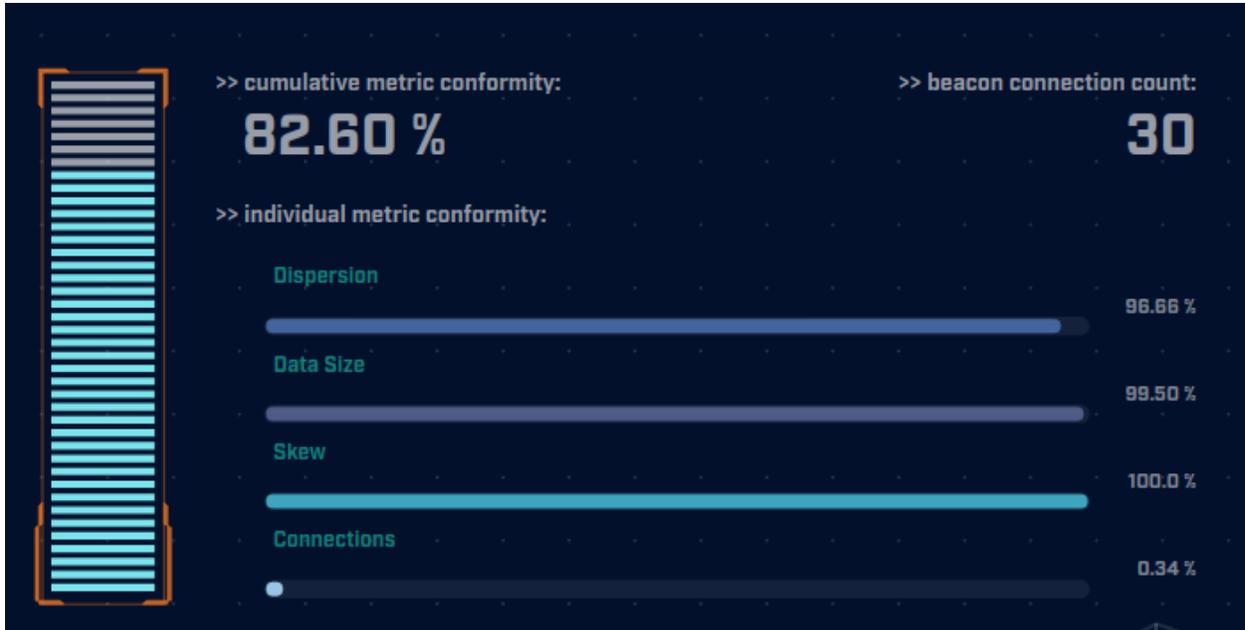


The x-axis identifies the timing between connections in seconds. The y-axis identifies the number of connections between the two specified IP addresses that exhibited the timing shown along the x-axis. For example, the bar on the right is identifying that 10 connections were separated by 2,049 seconds, or 34.15 minutes. This is approximately twice per hour, which was the same timing behavior we observed in the time analysis chat earlier in this document. The second bar (on the left) is showing that 30 connections were separated by 2,048 seconds, or 34.13 minutes. Again, this syncs with what we saw earlier.

Also of interest is that the timing between the connections all range from 2,048 to 2,049 seconds. So even though all of the connections do not show the exact same timing separation, the dispersion between all connections within the data set is very tight.

Interval 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 most 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.

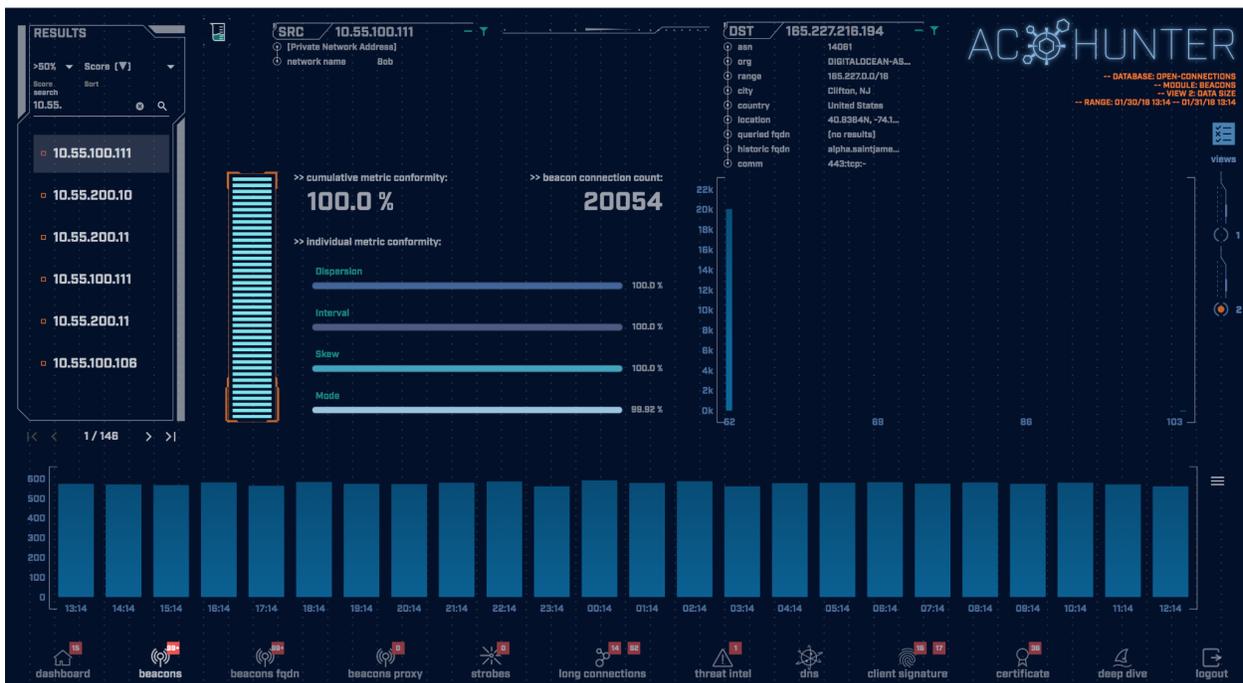
Data Size Analysis

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.



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.



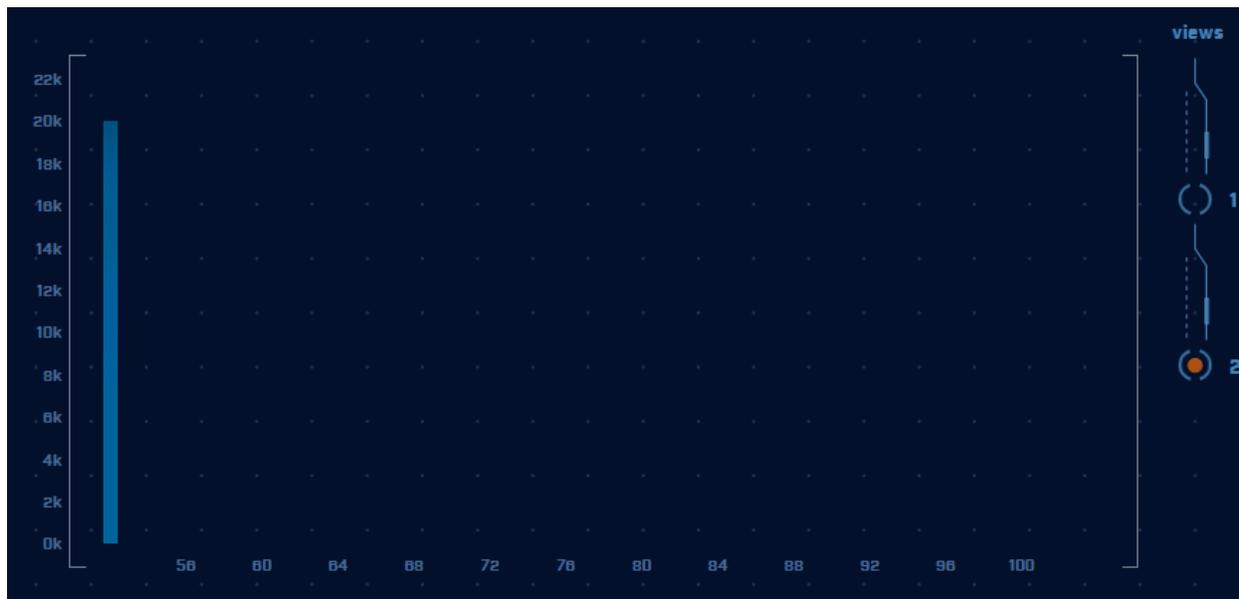
Some of the charts and features described in the previous section are available here as well. The following charts and features are unchanged between views.

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

Data Size Scoring Chart

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



The left hand side of the chart shows the overall data size score for this IP address pair. The right hand side of the chart displays the individual attributes that were analyzed when generating the overall score.

Like the time interval scoring chart, the most important value is the overall score. In fact, the overall score is the same value between the two views because it takes both data size and connection frequency into account. Connections scoring higher than 98% should be considered highly suspect.

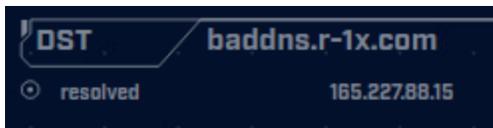
Beacons web Analysis

The “Beacons web” module is useful for cases where an internal host is beaconding out to an external host through the use of an FQDN (extracted from the HTTP or TLS conversation). In these cases, the external IP address to which the internal host is communicating might change over a short period of time.

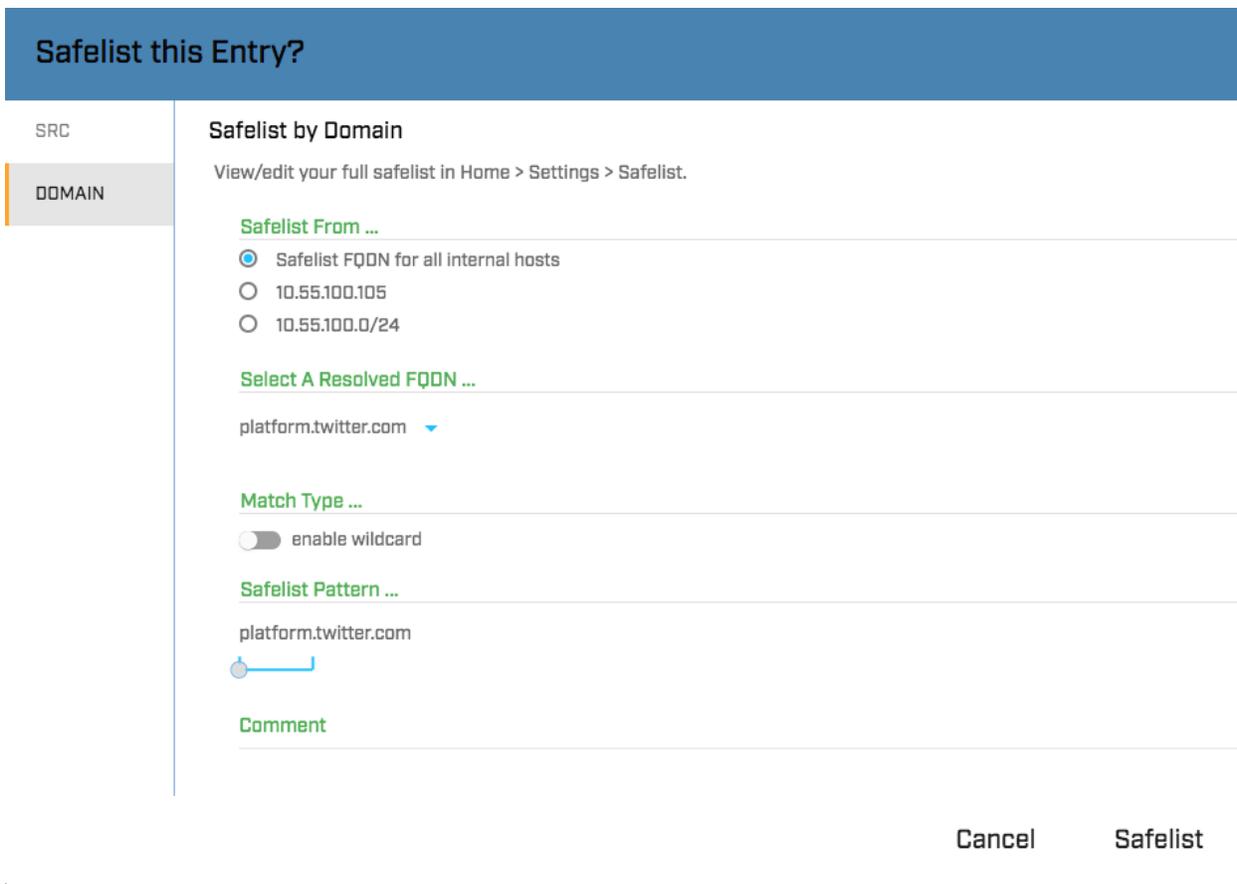


Nearly every aspect of the “Beacons web” module is the same as what is present in the “Beacons” module. One main difference is the destination information that is displayed. On the “Beacons web” module, an FQDN is displayed for the destination. Under the FQDN, a “Resolved” line is present. The “Resolved” line will display any IP addresses that were found to be returned in response to recent DNS queries for the FQDN. The queries could have either

been made by the source host shown on the page or by another system in the dataset. Clicking either the FQDN destination or a resolved IP will bring up the same investigation menus that were discussed in the “Beacons Analysis” section of this document.



Clicking the filter icon next to the destination will also result in slightly different behavior than what was seen for “Beacons Analysis”. The menu will start on the “Domain” tab rather than the “Dst” tab. The menu will allow for filtering based on the entire FQDN, a higher-level domain of the FQDN, or a wildcard match based on the FQDN that will result in subdomains being filtered.



In addition to defining a target hostname to safelist you can also decide to limit this to a single local IP address, the 256 addresses around it, or all internal IP addresses with the "Safelist From..." radio buttons.

The “Dst” and “Pair” tabs are intentionally not present for the filtering dialog of the “Beacon FQDN” module.

Beacons Proxy Analysis

The “Beacons Proxy” module can be useful for some environments in which hosts communicate to the internet through one or more proxy servers. Please note that the functionality of the “Beacons Proxy” module will depend heavily upon the proxy configuration of the environment and the placement of data sensors. The “Beacons Proxy” module currently only supports analysis on HTTP requests. Specifically, the requests must have used the HTTP CONNECT method to initiate the connection.



Nearly all aspects of the “Beacons Proxy” module are the same as the “Beacons” module. One main difference is the destination information that is displayed. On the “Beacons Proxy” module, an FQDN is displayed for the destination. Under the FQDN, a “Proxy IP” line is present. The “Proxy IP” line will display the IP address of the proxy host that was used to communicate with the FQDN. Clicking either the FQDN destination or the Proxy IP will bring up the same investigation menus that were discussed in the “Beacons Analysis” section of this document.



Clicking the filter icon next to the destination will also result in slightly different behavior than what was seen for “Beacons Analysis”. The menu will start on the “Domain” tab rather than the “Dst” tab. The menu will allow for filtering based on the entire FQDN, a higher-level domain of the FQDN, or a wildcard match based on the FQDN that will result in subdomains being filtered.

Safelist this Entry?

SRC	Safelist by Domain View/edit your full safelist in Home > Settings > Safelist. Safelist From ... <input checked="" type="radio"/> 10.55.100.111 <input type="radio"/> 10.55.100.0/24 <input type="radio"/> Safelist FQDN for all internal hosts Select A Resolved FQDN ... alpha.saintjameschurch.org ▼ Match Type ... <input type="checkbox"/> enable wildcard Safelist Pattern ... alpha.saintjameschurch.org <input type="text"/> Comment
DOMAIN	

Cancel **Safelist**

In addition to defining a target hostname to safelist you can also decide to limit this to a single local IP address, the 256 addresses around it, or all internal IP addresses with the "Safelist From..." radio buttons.

The "Dst" and "Pair" tabs are intentionally not present for the filtering dialog of the "Beacon Proxy" module.

Another major difference for the "Beacons Proxy" module is that it does not have a view 2 since the relevant data size information cannot be reliably obtained.

One final difference for the "Beacons Proxy" module is that it does not allow for BeaKer integration.

Something to note for the "Beacons Proxy" module is that having chosen to filter the IP address for a proxy server will not remove results involving that proxy server in the "Beacons Proxy"

module. For instance, in the screenshot above you can see a “Proxy IP” address of 142.93.18.199. If the IP address of 142.93.18.199 had been filtered in another portion of the program, the entry seen in the screenshot above would still appear in the “Beacons Proxy” module. This behavior is intentional. In the event that you have false-positives in other modules that arise from systems proxying through the proxy server, those results can still be filtered from the other modules without potentially removing true-positives from the “Beacons Proxy” module.

Strobes Analysis

Strobes are similar to beacons in that they are repeated connections between two IP addresses. However, unlike a beacon which may try to 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. These could easily be spotted just by watching a packet sniffer scroll through packets. To analyze strobes, click the strobe icon on the menu.



This will produce the strobes analysis screen:

Src	Dst	Connection Count	Dst FQDN
10.55.100.106	199.38.164.54	14	app.lime.com, a.rfihub.com, 20741079p.rfihub.com, p.rfihub.com, 20753598p.rfihub.com
10.55.100.108	162.212.41.11	12	app.lime.com, www.shutterstock.com
10.55.100.105	199.166.0.200	8	app.lime.com, sc.lasds01.com, anycast.sc.lasds01.com
10.55.100.107	52.27.226.96	8	app.lime.com, dcs-edge-usw2-620097651.us-west-2.elb.amazonaws.com
74.217.250.101	23.217.149.21	4	app.lime.com, e4343.g.akamaiedge.net

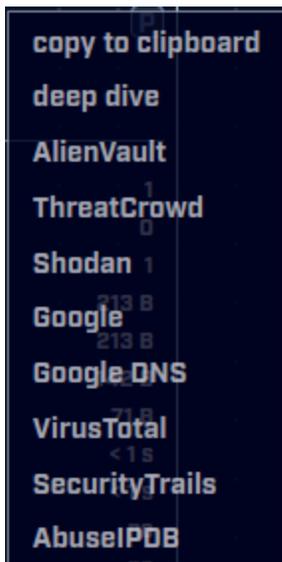
Because strobes communicate so frequently, an in depth analysis is not required to identify that regular communications are being maintained. For example, in the first line above we see that 10.183.20.144 initiated communications with 192.168.117.138 a total of 3,043,773 times over the 24 hour period under review. That’s approximately once every 3 milliseconds.

Because strobos are so noisy, you may find that the source is due to poorly written code or a misconfiguration rather than actual malware. Most malware will at least attempt to be stealthy.

To only see strobos 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:



Please see [Customizing Investigation Sources](#) for information on how to add your preferred lookup tools to this drop-down list.

Clicking on a fqdn will open a menu with one or more options. If the row you clicked has more than one fqdn, then the menu will open with all of the available fqdns. Clicking on one of the fqdns will open a menu with investigation sources that you can use to investigate the fqdn with:



These sources can be one or more, commonly “whois”, “URLScan”, and “copy to clipboard”. The first two are external websites you can use to research the hostname, and the last copies the hostname to the clipboard (as shown above.)

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.

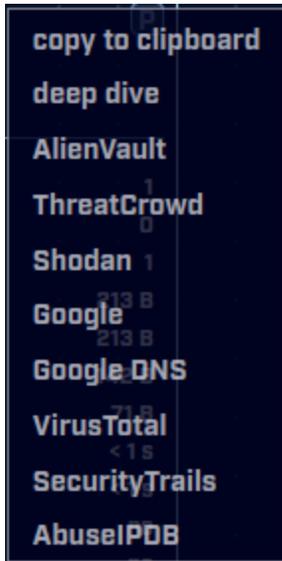
The output is fairly self explanatory. You see the IP address of the system that initiated the session. You also see the destination IP address and port number. The "service" is an application layer analysis of the protocol being used to communicate over this port. For example, "ssl" indicates that a 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).

AC-Hunter analyzes both "closed" connections (where the connection was shut down before the end of this time period) and "open" connections (where the connection was still going at the end of the time period). Note that you need to be using the Zeek sensor supplied with version 5.3.0 or higher to see the open connections.

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:



Please see [Customizing Investigation Sources](#) for information on how to add your preferred lookup tools to this drop-down list.

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.



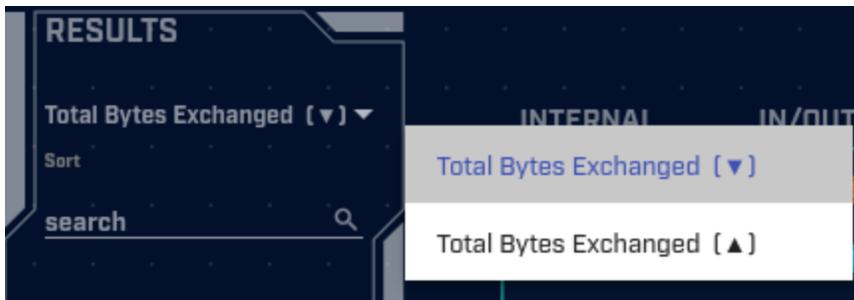
The layout of the threat intel screen is a bit different than the screens discussed previously.

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.

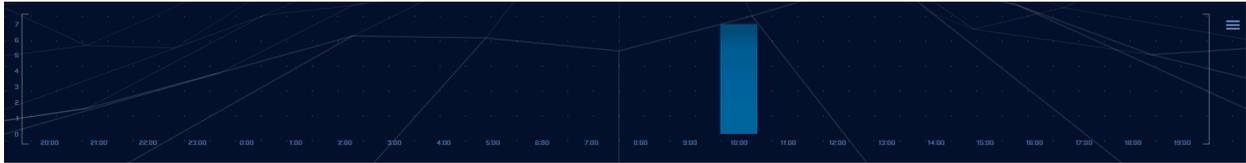


The functionality of the results feed is similar to that which was discussed on previous screens. By default, higher priority addresses are listed at the top. The default sort order is "Total Bytes Exchanged". You can change the sort order by clicking the "Total Bytes Exchanged" label near the top of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



Timeline Chart

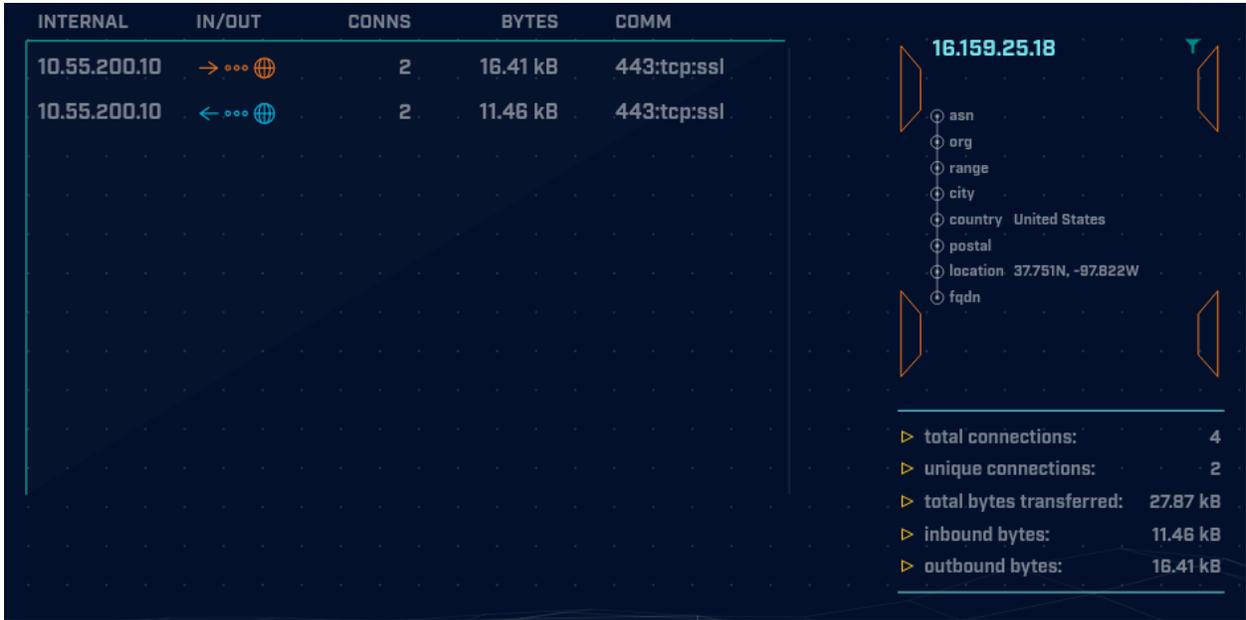
The timeline chart identifies the number of connections observed to the IP address that is on at least one threat intel feed over the time range within the dataset being analyzed.



Each bar represents the number of connection requests seen within a specified period of time. The default resolution is one hour per bar. 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.

Host Chart

The host chart provides detailed information regarding the destination IP address being analyzed.

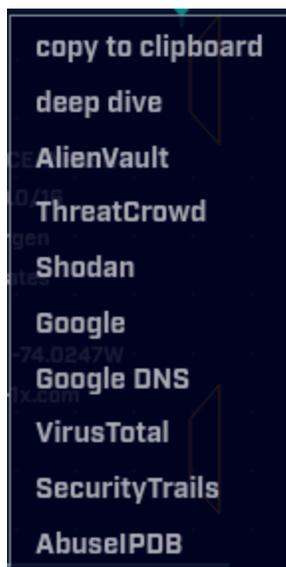


The left hand side of the chart provides summary information regarding which internal systems contacted the IP address recognized from at least one threat intel feed. The In/Out column contains an icon that indicates whether the threat connection is inbound or outbound. Notice that the first IP in the table (orange) is an outbound connection, as the arrow is pointing away from the internal IP and out towards the internet. The last IP in the table (blue) is an inbound connection, as the arrow is pointing from the internet towards the internal IP. If an internal IP has both inbound and outbound connections, each connection will display in it's own row (as shown above) since each connection can have different connection counts, bytes transferred, and comm ports.

The top right side of the chart shows identification information regarding the IP address recognized from at least one threat intel feed. Clicking on the filter icon brings up the safelisting screen. Safelists take precedence over threat intel feeds, so creating a safelist entry will remove the IP from being presented when analyzing this data set.

The importance of this information depends on where the Zeek system collects packets on your network. If the device is outside of your firewall, then most connection attempts listed here will (hopefully) get filtered out by your firewall and can be safely ignored. If, however, you are collecting packets inside of your firewall filtering, results listed here are worthy of a deep dive, especially if it appears data was transferred.

Clicking an IP address will open a menu with one or more options, including "AbuseIPDB", "AlienVault", "copy to clipboard", and "deep dive". The first two are external websites you can use to research the 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:



Please see [Customizing Investigation Sources](#) for information on how to add your preferred lookup tools to this drop-down list.

Threat Intel Feeds Used

RITA and AC-Hunter leverage the following threat intel feed from the following source.

- feodotracker.abuse.ch

You also have the ability to create your own Threat Intel list. For details and step-by-step instructions, see https://portal.activecountermeasures.com/support/faq/?Display_FAQ=1890.

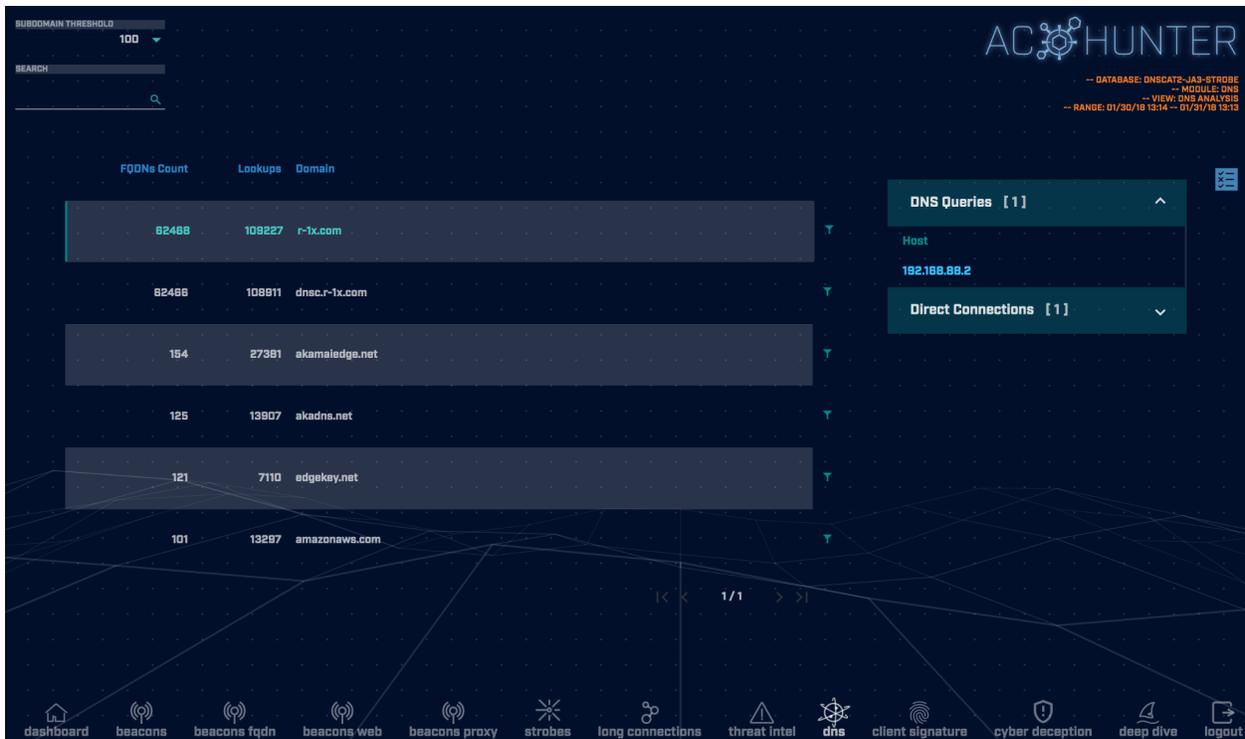
DNS Analysis

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:



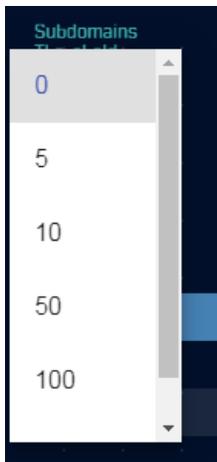
This screen reports summary information regarding the external hosts being contacted by your internal systems. For example, in the sample screen capture above there were 62,468 subdomains of r-1x.com.

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.



You can also search for specific hostnames or domains by using the "Search" box in the upper left:



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.

Useragent String	Seen	Requests	Sources
Windows-Update-Agent/7.9.9800.18758 Client-Protocol/1.21	1	statsfe2.update.microsoft.com	10.55.200.10
client connection	1	tele.trafficmanager.net	10.55.200.10
Microsoft-CryptoAPI/8.3	2	ctldl.windowsupdate.com	10.55.200.10
Windows-Update-Agent/10.0.10011.18384 Client-Protocol/1.40	8	download.windowsupdate.com	10.55.200.11
OfficeClickToRun	12	officecdn.microsoft.com.edgesuite.net, officecdn.microsoft.com	10.55.182.100
Microsoft BITS/7.9	25	7au.download.windowsupdate.com, au.download.windowsupdate.com	10.55.200.11

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 "Search" field allows you to search by either IP address, hostname, or domain name.

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.



The screenshot shows a table with the following columns: SSL/TLS Hash, Seen, Requests, and Sources. The data is as follows:

SSL/TLS Hash	Seen	Requests	Sources
5e573c9c9f8ba720ef9b18e9fce2e2f7	1	clientservices.googleapis.com	10.55.182.100
bc6c386f480ee97b9d9e52d472b772d8	2	clients4.google.com, 558-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
08bf94d7f320da537b5e9b76b06e02a2	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, hostname, or domain name, enter that 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:



Please see [Customizing Investigation Sources](#) for information on how to add your preferred lookup tools to this drop-down list.

Cyber Deception

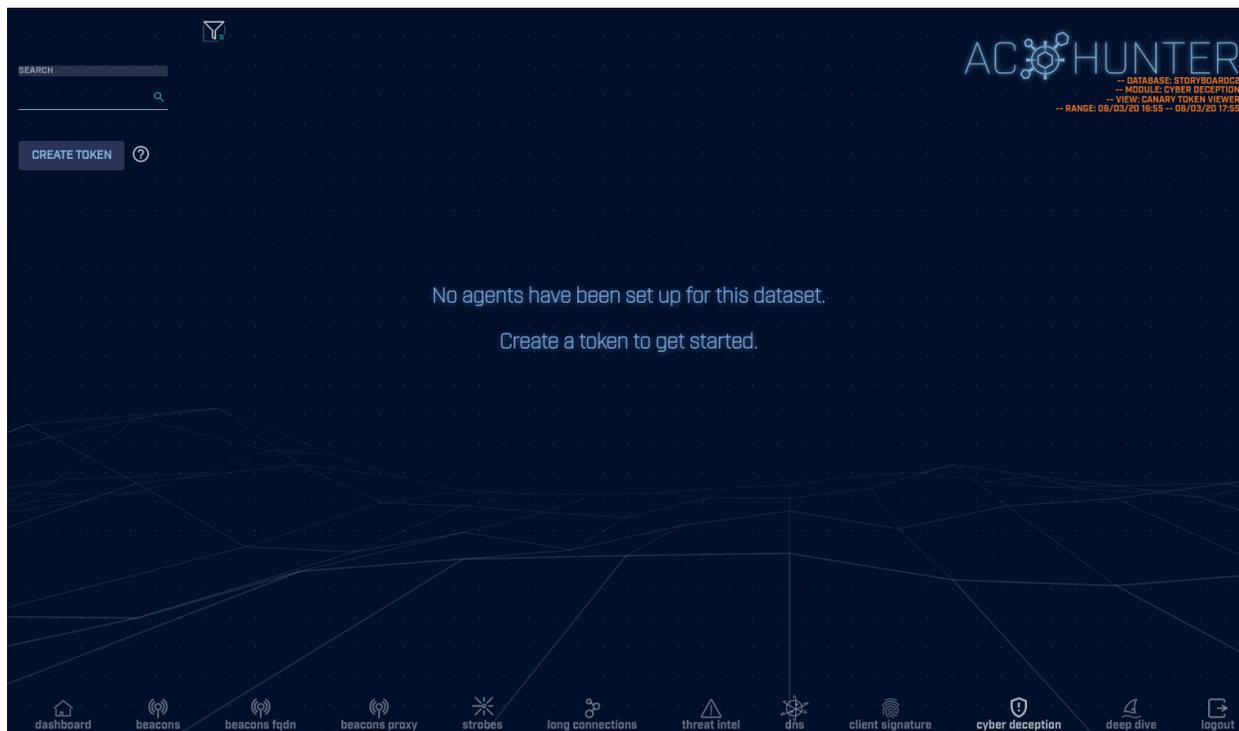
The cyber deception module allows for the creation and monitoring of canary tokens. There are two types of canary tokens available. File access tokens will generate an alert when a designated file has been accessed. User-access tokens will generate an alert when an authentication attempt is made against a monitored user or a Kerberos ticket is requested for that user.

This module requires listening agents to be registered on your Domain Controller, which will in turn allow for file and/or user account monitoring. Once these agents and monitored resources are properly set up, triggered deception events will begin to appear within the “Triggered Events” table in the cyber deception module for any alerts that are generated from the tokens.

To view the cyber deception module, click the cyber deception icon on the bottom of the screen.



Upon first use, no information will be available.



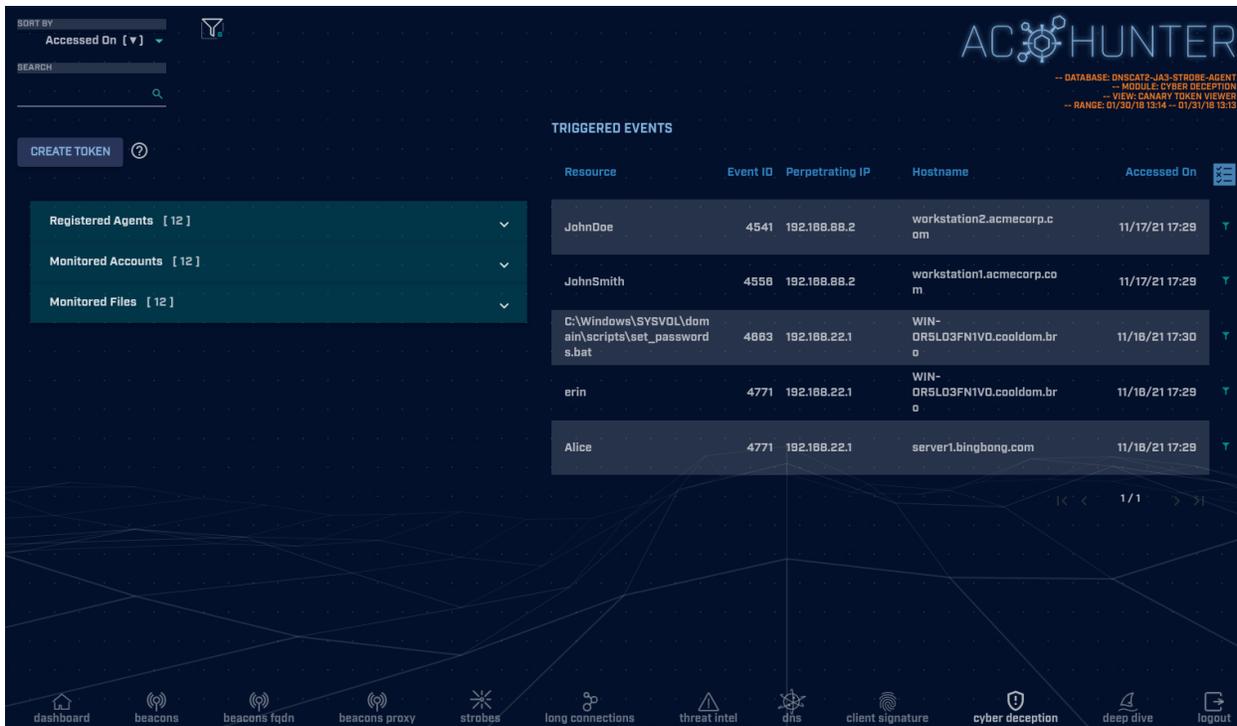
Clicking on the “CREATE TOKEN” button in the top left hand corner will download a .zip file containing a PowerShell script and other files that will assist in setting up registered agents and monitoring files/accounts.



If for any reason you need a reminder, clicking the  icon will bring up a brief overview of the tokens and instructions on how to set them up.

Once agents are set up, resources are being monitored, and the tokens have been triggered with deception events, the cyber deception module will display a split view.

For more information on setting up agents and resource monitoring, see Canary Token Setup for Cyber Deception.



On the left hand side, any registered agents, monitored files, and monitored accounts *set up for the current dataset* will be shown. Simply expand each section to view their details.



On the right hand side, any triggered deception events will be displayed. The Resource column represents the resource that is being monitored and has been triggered (either a file or user account.) The Windows Event ID for that event is shown as well. The Perpetrating IP represents the IP address of the system that accessed the file or attempted to log in to the shown account.

Clicking on a row will display the entire Windows Event Log for that triggered event, as well as some extra information.

Event Details

Path: C:\Windows\SYSVOL\domain\scripts\set_passwords.bat

Logon ID: 0x42013cd1

Process ID: 0x1764

Process Name: C:\Windows\explorer.exe

Windows Event Log 

```

<Event xmlns='http://schemas.microsoft.com/win/2004/08/events/event'>
  <System>
    <Provider Name='Microsoft-Windows-Security-Auditing' Guid='{54849625-5478-4994-a5ba-3e3b0328c30d}' />
    <EventID>4663
  </EventID>
  <Version>1
  </Version>
  <Level>0
  </Level>
  <Task>12800
  </Task>
  <Opcode>0
  </Opcode>
  <Keywords>0x8020000000000000
  </Keywords>
  <TimeCreated SystemTime='2021-11-16T22:30:23.536032200Z' />
  <EventRecordID>1069215
  </EventRecordID>
  <Correlation/>
  <Execution ProcessID='4' ThreadID='520' />
  <Channel>Security
  </Channel>
  <Computer>WIN-0R5L03FN1V0.cooldom.bro

```

Close

Clicking on the 'copy' icon next to Windows Event Log will copy the entire event log to your clipboard.

Windows Event Log 

Triggered deception events can be safelisted by the perpetrating IP address. If there is a system that should be allowed to access the file or account, safelisting by the IP address of that system will remove any triggered events by that system from the event list. It will also remove any events generated from that system from impacting threat scores.

Safelisting a perpetrating IP will NOT safelist that IP in any other modules but cyber deception.

Safelist this Entry?

DECEPTION

Safelist Triggered Events by IP Address

View/edit your full safelist in Home > Settings > Safelist.

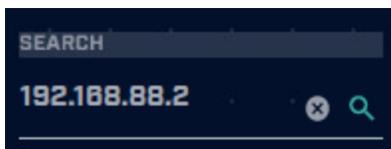
Whitelist ...

- 192.168.88.2
- 192.168.88.0/24
- 192.168.0.0/16
- 192.0.0.0/8

Comment

Cancel Safelist

To view events triggered by a specific perpetrating IP, enter the whole or partial IP address in the search bar in the top left hand corner.



Clicking on a perpetrating IP will bring up an investigation menu.



*Note that since deepdive requires the investigated host to be watched by Zeek, results may not be available if the perpetrating IP has not been seen by Zeek.

*Note that if the host is internal on your network and being monitored with Espy, it also will not have any results for deepdive.

*Note that deception entries, by default, will only be visible in a rolling dataset for 1 day. The number of days to keep entries present in the rolling dataset can be adjusted with the DeceptionManagement->DaysKeepEntries setting in the /etc/AC-Hunter/config.yaml file.

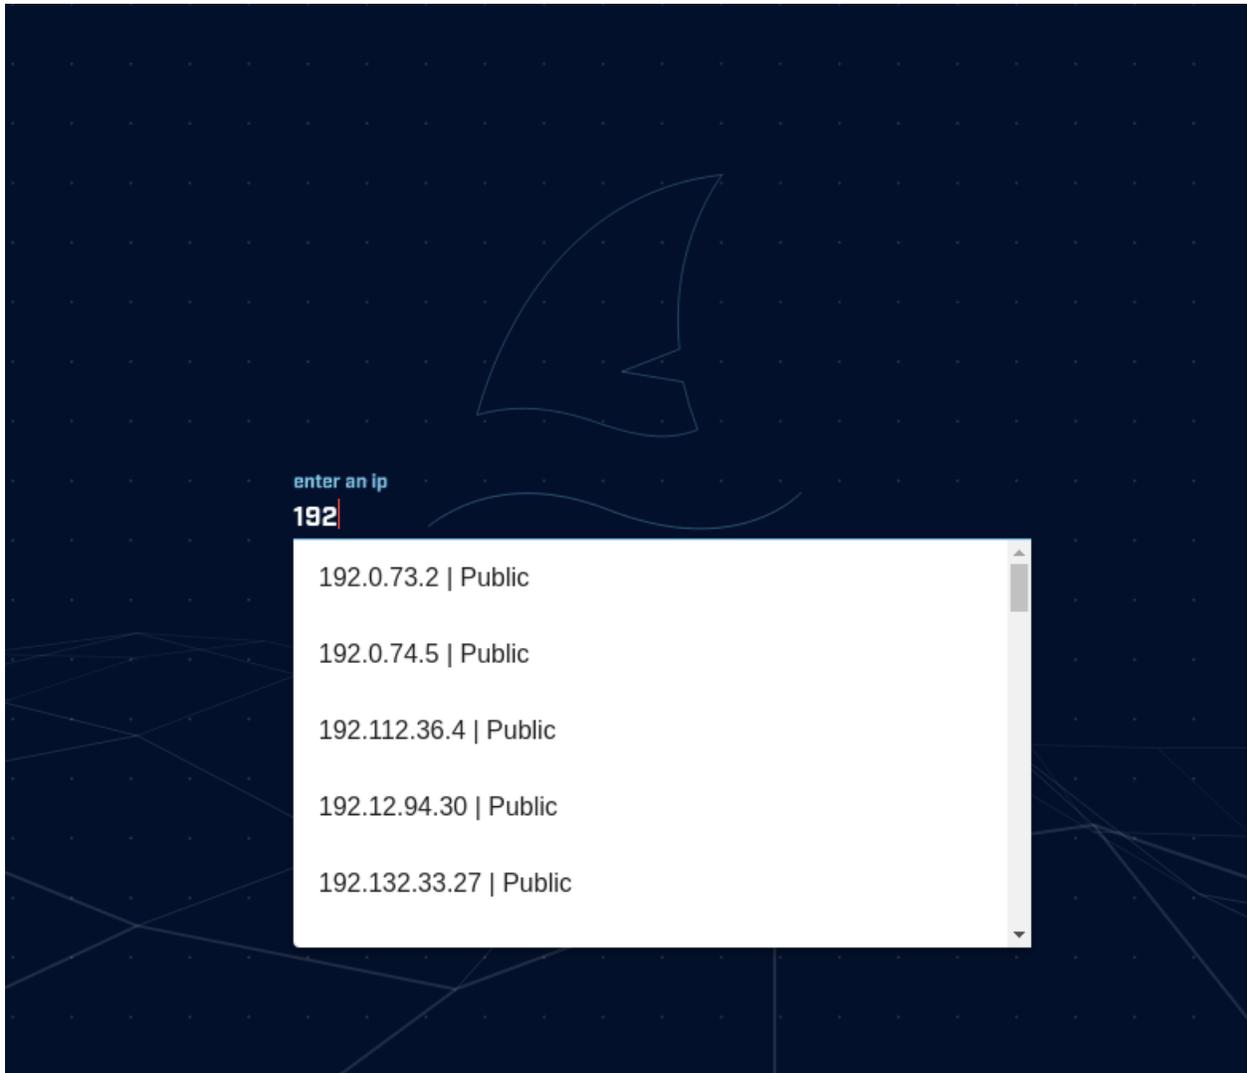
Deep Dive

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.

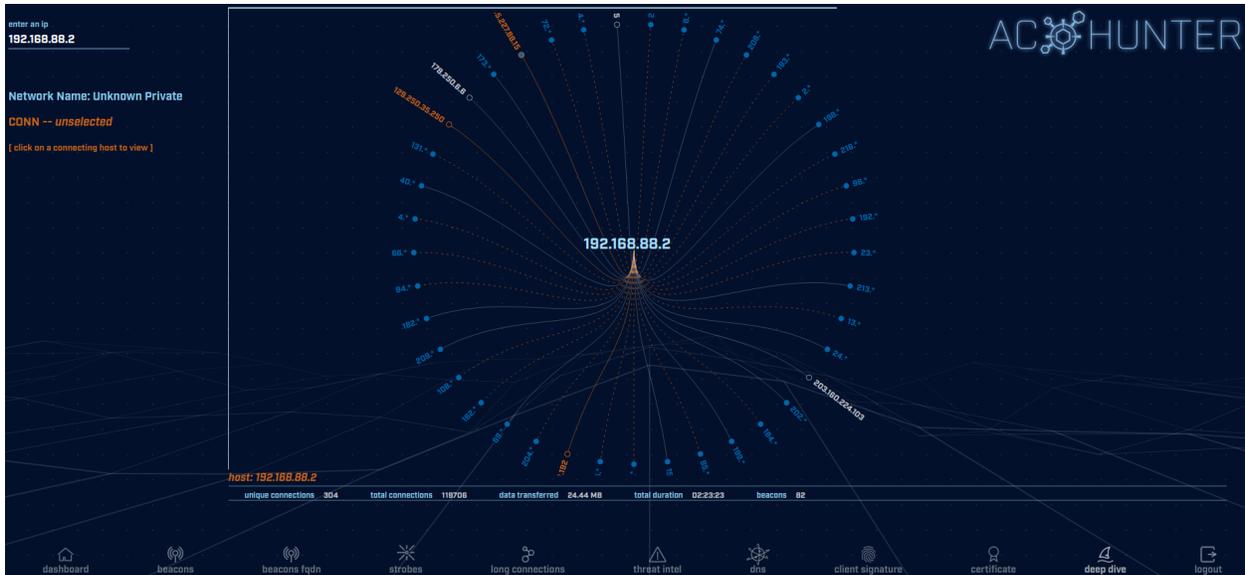
To access the deep dive module, click the deep dive icon at the bottom of the screen:



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.

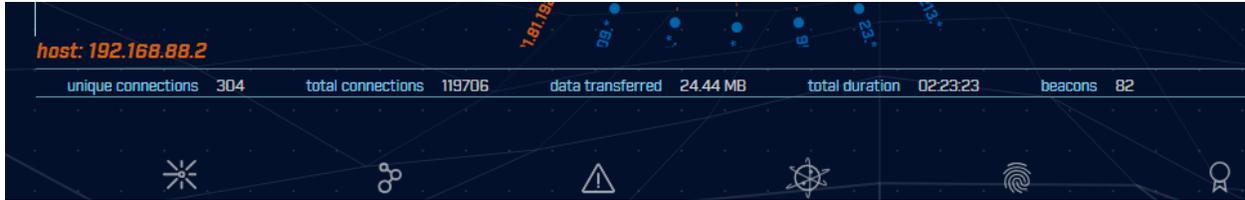


Once you select an IP address, the deep dive main screen will be loaded.



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.

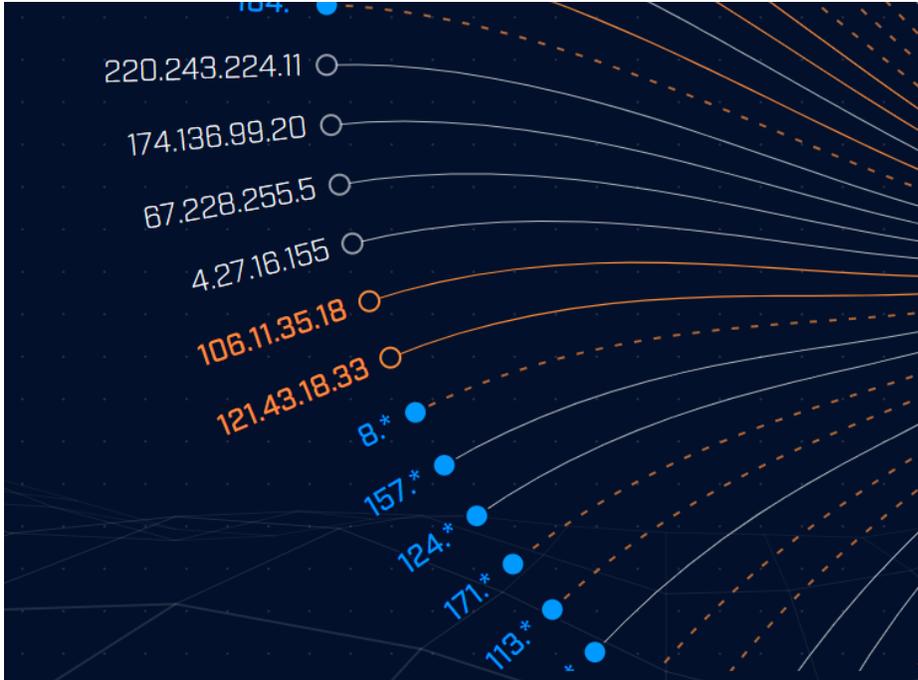
Just below the graphic, you can see overall summary information regarding the system being investigated.



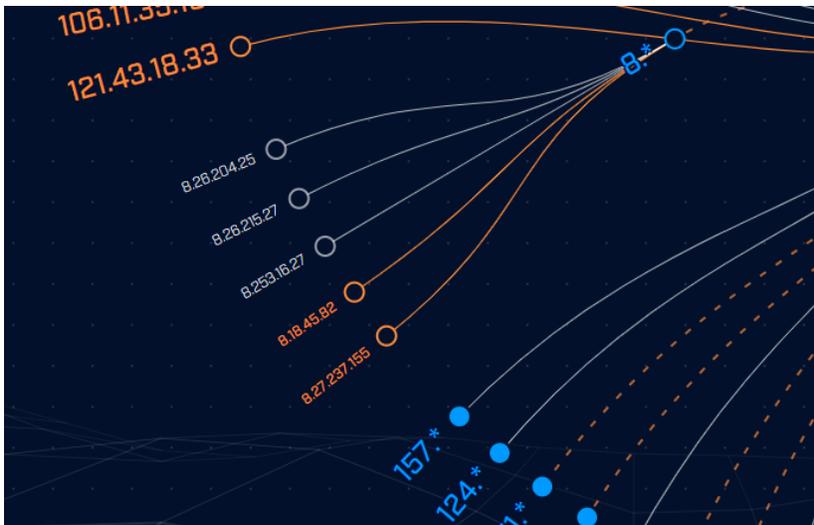
This includes data points such as the total number of connections the system created, the amount of data transferred and the time spent with active connections open.

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:

Network Name: Carol P

CONN --173.245.59.41

overview

total connections	23
inbound connections	0
outbound connections	23
avg bytes transferred	182 B
total bytes transferred	4.19 kB
inbound bytes	2.28 kB
outbound bytes	1.91 kB
total duration	00:00:02
max duration	< 1 s
beacon activity	yes
invalid certs	no
state	closed
queried fqdn	ns2.digitalocean.com
historic fqdn	[no results]
dst port:proto:service	53:udp:dns

connection timeline

The "overview" section is a summary of all communications between the IP address I'm 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.

Clicking the "CONN--" IP address will open a menu with one or more options, commonly "AbuseIPDB", "AlienVault", and "copy to clipboard". 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.

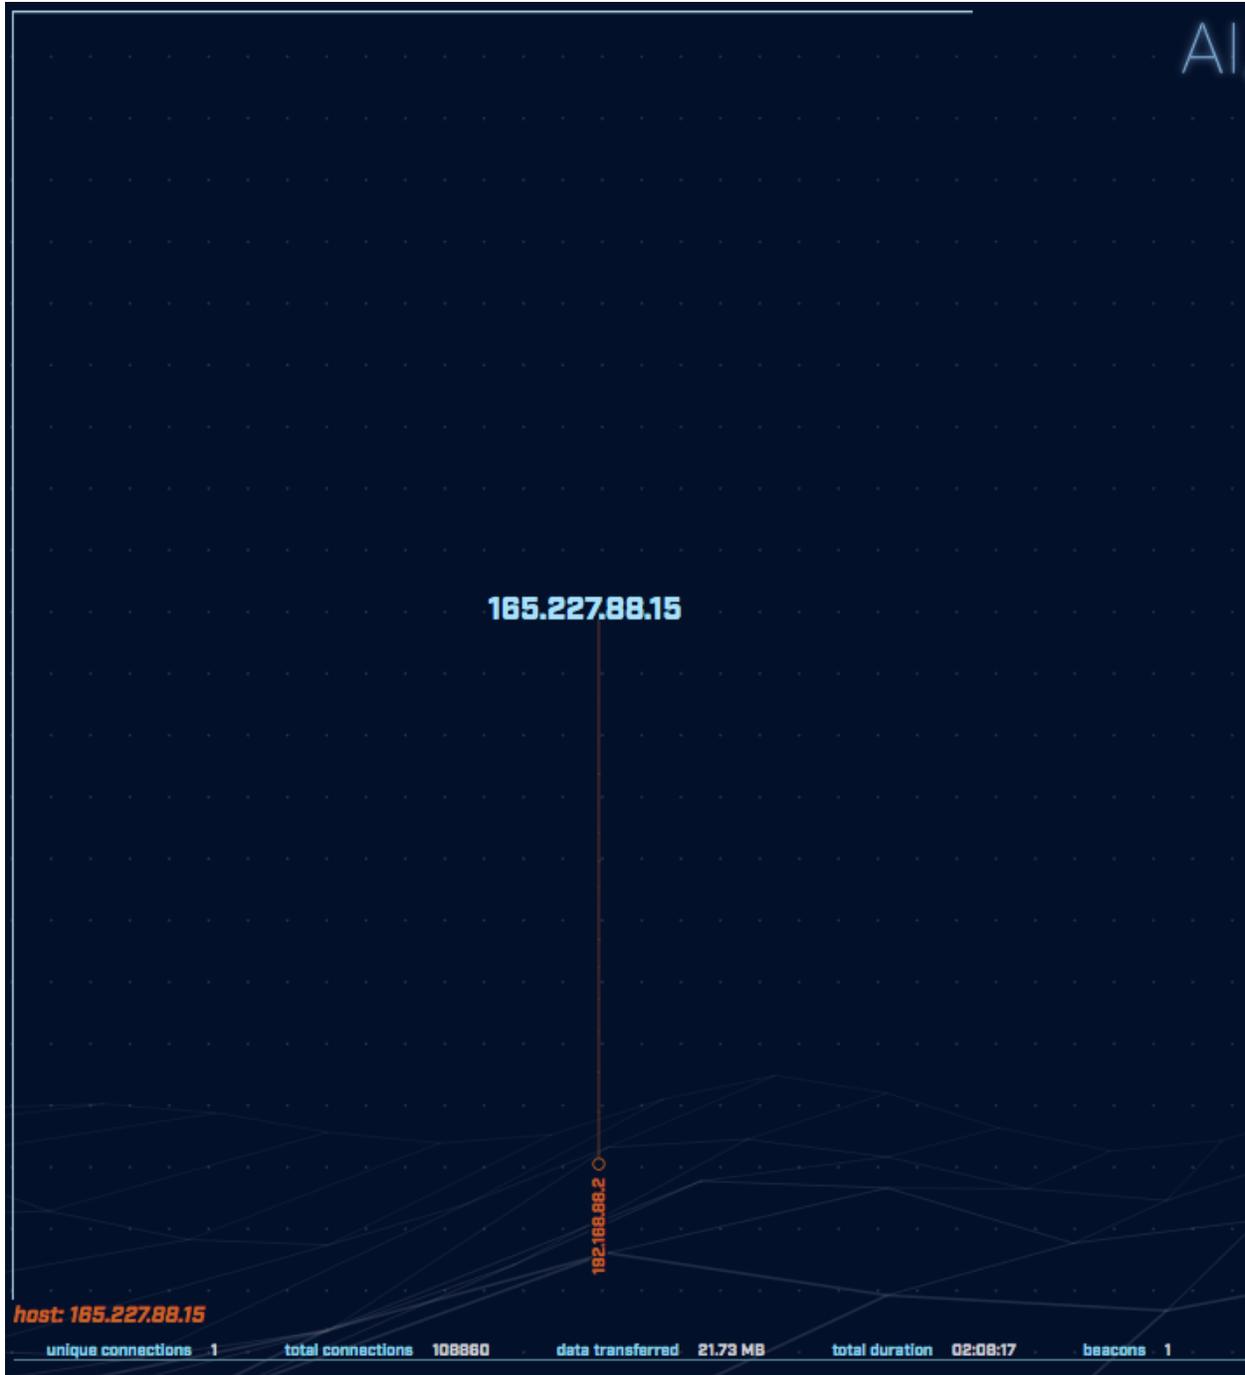


When available, a fqdn will be visible in the overview pane. Clicking on the fqdn will open a menu with one or more options, including “whois”, “URLScan”, and “copy to clipboard”. The first two are external websites you can use to research the fqdn, and the last copies the fqdn to the clipboard:

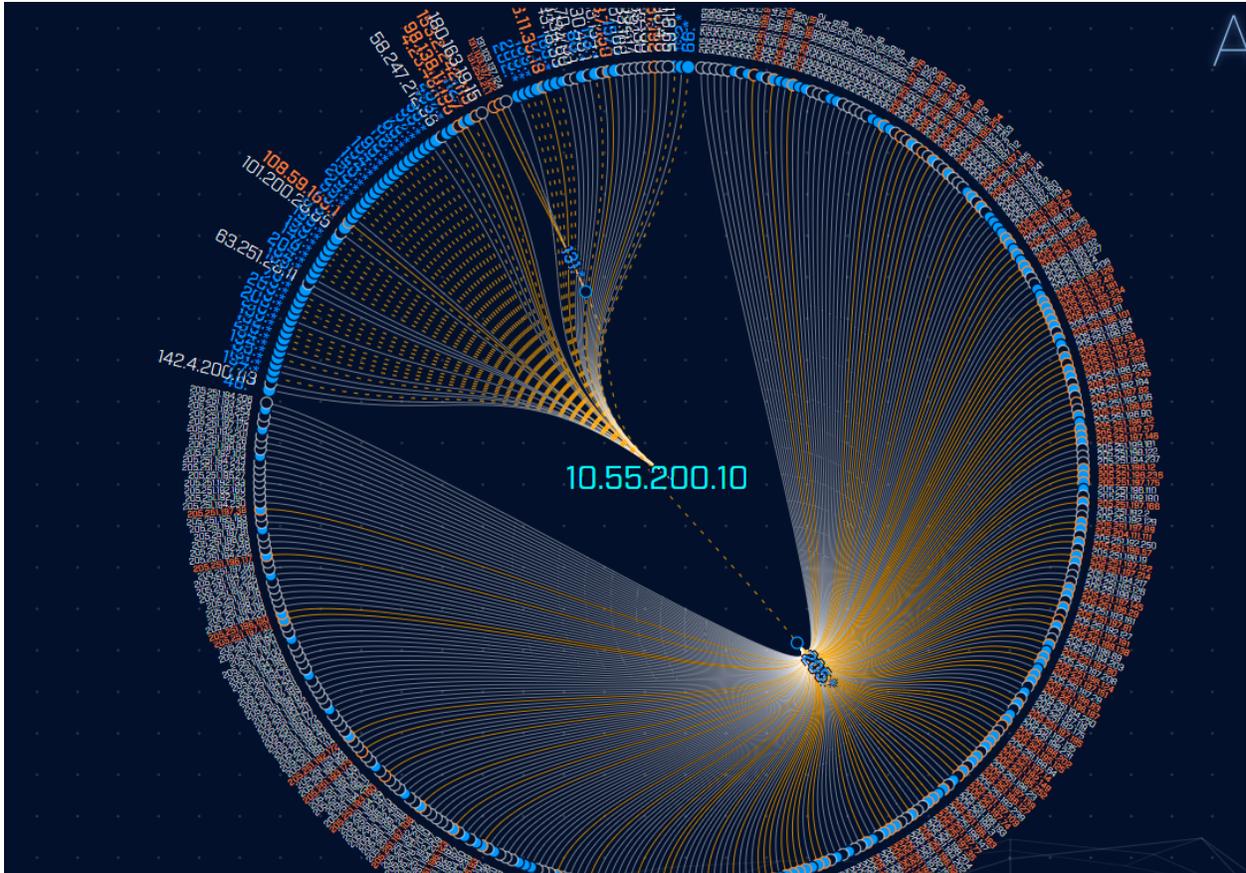


The connection timeline is a graph of the number of connections that took place each hour over a 24 hour period of time. If I mouse over the timeline it will expand across the screen so that it's easier to read the values on the X and Y axis.

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's 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.

Active-Flow Databases

In some network environments it's not possible or not practical to run Zeek to capture the raw data needed by RITA and AC-Hunter. Packet sniffers might not be allowed by policy, the network might generate more traffic than Zeek can ingest, or there might not be a place to put the sniffer to get a good picture of the network.

As an alternative, if the routers in that network can send summaries of those conversations over Netflow V9, RITA and AC-Hunter can use them instead of Zeek logs. Because the Netflow summaries are more limited than the summaries provided by Zeek, there are some limitations that we'll cover below.

The Active-Flow module converts Netflow data provided by routers into a format that AC-Hunter can read and import.

Espy Databases

Espy provides a second alternative for environments where it's not possible or practical to run Zeek. Where Active-Flow learns about network connections from Netflow packets provided by routers, Espy learns about network connections from monitors running on Windows systems.

Differences Between Zeek-sourced and Espy/Netflow-sourced Data

If you take this approach of importing Netflow or sysmon/Espy data to use in AC-Hunter, you should be aware of some contrasts between Threat Hunting based on Zeek logs and Threat Hunting based on Netflow/Espy data:

- Netflow and Espy records are a high level summary that shows a conversation took place between two systems on a given port. They do not include any of the content of that conversation. (It's like knowing that a phone call took place in a given time period between two numbers but having no audio of the two people speaking.)
- For that reason, you will not have any data on the DNS, Client Signature, or Certificate tabs.
- On the Beacons screen we use Zeek's ability to classify a conversation not only by port but also by the application protocol used (for example, "443:TCP:TLS"). Since that data is not provided outside of Zeek, this classification is limited to the port ("443:TCP:-"). This also means that in a non-Zeek-sourced database, "Unexpected protocol on a well-known port" will not have valid results on the dashboard tab.
- Because there is not enough information recorded in Netflow records to reliably show which machine started the conversation, the Threat Intel lines on the dashboard are duplicated between Outgoing and Incoming. For the same reason you may see double entries on the Beacons tab; one for external to internal, one for internal to external. It's best to focus on internal IP to external IP beacons.
- Netflow records are commonly generated every N minutes (perhaps every 5 or 15 minutes, sometimes configurable). This means that the Long Connections tab will report fragments of connections that are far shorter than the overall connection.
- Espy data does not include connection duration so there will be no results in the Long Connections tab.
- Espy data does not include the amount of data transferred in a connection. This means that the Beacons score for data size will always be a constant value (100%) but the interval portion of the score will vary as usual. It also means that in Deep Dive and Threat Intel there will be no data on bytes transferred.

With the exception of the above notes, the Dashboard, Beacons, Strobes, Long Conns, Threat Intel, and Deep Dive tabs should work much the same as they do with Zeek-sourced databases.

Active-Flow Installation

The easiest way to get set up with Active-Flow is to use the official installer, which will connect Active-Flow to your AC-Hunter install automatically. This will install the Active-Flow server on the system you specify (which can be the AC-Hunter system as long as that machine is not also running Espy or Zeek).

Once the server is running you'll need to send Netflow V9 records to it. You'll need to configure a router or switch on your network to send these records to the hostname (or IP address) of the Active-Flow server.

Active-Flow expects to see certain fields in these records; to see the minimum list and how to enable these on a Cisco ISR, see https://portal.activecountermeasures.com/support/faq/?Display_FAQ=2726 . If you run into any problems with Active-Flow, please see the troubleshooting FAQ at https://portal.activecountermeasures.com/support/faq/?Display_FAQ=2722 .

Within a few hours of when you have the Netflow data flowing correctly you should see a new "localhost__1000-rolling" database in AC-Hunter's database list which contains the Netflow connection data.

Espy Installation

The easiest way to get set up with Espy is to use the official installer and install Espy to the same system as AC-Hunter (as long as neither Zeek nor Active-Flow is also running on AC-Hunter).

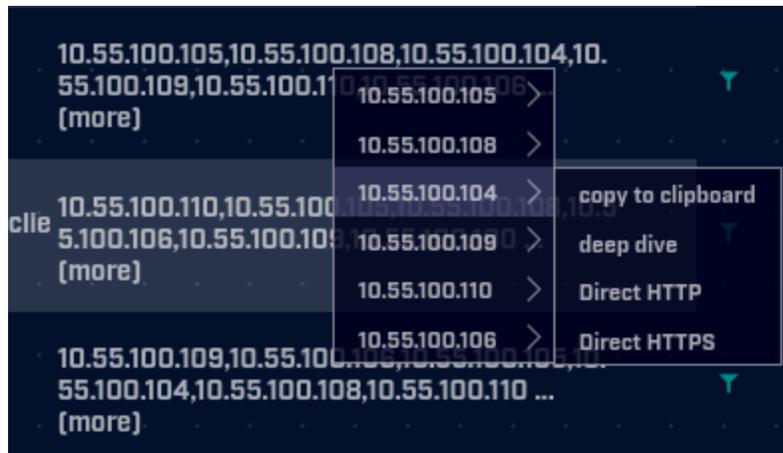
Once the server is installed, you'll need to install the Espy client on each Windows system you wish to monitor. To do this, see the "Automated Install: Espy Agent" section of <https://github.com/activecm/espy/> . Within a few hours you should see a new

"localhost__1000-rolling" database in AC-Hunter's database list which contains the Espy connection data.

Investigation Sources

AC-Hunter comes with several Investigation Sources pre-loaded. You can see these in the investigation drop down menu that appears when clicking an IP address, FQDN, ASN, user agent string, or JA3 hash. The options presented will differ for IPs based on whether the IP selected is defined as an external IP address or an internal IP address. All investigation menus have the "copy to clipboard" feature enabled, allowing you to copy the item clicked to the clipboard.

For instances where there is more than one investigatable item available, the menu will display all available items (FQDNs, IPs, etc) with their respective investigation sources as a submenu.



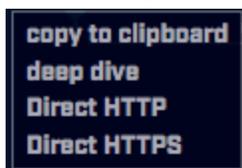
External IP Investigation Sources



The default sources are described below:

- AbuseIPDB is a crowd-sourced collection of IPs associated with malicious activity
- AlienVault Open Threat Exchange provides IOCs and threat intel
- ThreatCrowd presents AlienVault OTX's data in a different format
- Shodan records historical open ports and services on the host
- Google searches the web for an IP
- Google DNS uses Google's DNS over HTTPS service to do a PTR lookup on the IP
- VirusTotal's passive DNS service displays historical DNS records for an IP
- SecurityTrails' passive DNS service displays historical DNS records for an IP
- copy to clipboard puts the IP address on the clipboard for ease of pasting to other tools
- deep dive sends you straight to this tab with this IP address selected

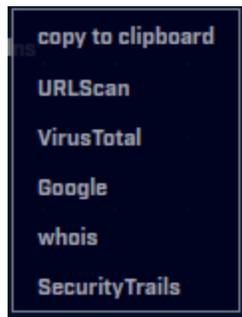
Internal IP Investigation Sources



The default sources are described below:

- Direct HTTP will visit the IP in question directly on port 80. WARNING: This option is not OPSEC safe since it will cause you to send traffic to the IP address being investigated.
- Direct HTTPS will visit the IP in question directly on port 443. WARNING: This option is not OPSEC safe since it will cause you to send traffic to the IP address being investigated.
- copy to clipboard puts the IP address on the clipboard for ease of pasting to other tools
- deep dive sends you straight to this tab with this IP address selected

FQDN (domain) Investigation Sources



The default sources are described below:

- whois will visit the whois lookup tool from MXToolbox to provide ownership and status information of the provided domain
- SecurityTrails' passive DNS service displays historical DNS records for the fqdn

- URLScan displays recent hits and scans for the domain address
- VirusTotal's passive DNS service displays historical DNS records for the fqdn
- Google searches the web for the fqdn

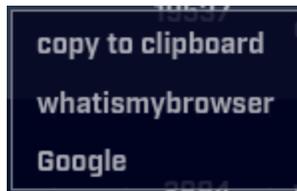
ASN (autonomous system number) Investigation Sources



The default sources are described below:

- IPInfo displays various pieces of information about ASNs and IP addresses, including registered owner data and geographical locations.

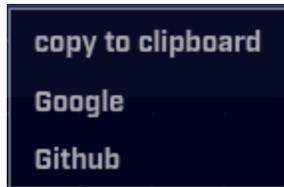
User Agent String Investigation Sources



The default sources are described below:

- Google searches the web for the user agent string
- whatismybrowser copies the selected user agent string to the clipboard for you to paste into the user agent string parser to view a human-readable user agent for said user agent string

JA3 Hash Investigation Sources



The default sources are described below:

- Google searches the web for the ja3 hash
- GitHub does a global (across all repositories) search on GitHub for the ja3 hash

Canary Token Setup for Cyber Deception

The setup process for file and user access canary tokens uses a .zip archive that is downloadable from the Cyber Deception module within the AC-Hunter user interface. Clicking the “CREATE TOKEN” button will download the archive to your system.



The downloaded zip archive needs to be transferred to the system(s) on which the monitoring will take place. These systems will need to be able to communicate with AC-Hunter via TCP ports 39839-39860.

For both user-access tokens and file tokens, you will need to transfer the archive to all Domain Controllers in the domain that have the Active-Directory role enabled.

For file tokens, the archive will also need to be transferred to any system that is hosting file tokens if the tokens are being hosted on systems other than the Domains Controllers with the Active-Directory role enabled. If the system is not a Windows Server distribution (e.g., Windows 10), then you will need to install the Active Directory RSAT tools by performing the following steps first:

- Open a PowerShell prompt as an Administrator
- Run: `Add-WindowsCapability -Name "Rsat.ActiveDirectory.DS-LDS.Tools~~~~0.0.1.0`

After transferring the archives, perform the following steps on one of the systems as a Domain Admin user:

1. Ensure that outbound connections are allowed on TCP ports 39839-39860.
2. Unzip the archive.
3. Open a Powershell prompt as an Administrator.
4. Navigate to the unzipped archive folder.
5. Run: `powershell -exec bypass`
6. Run: `Import-Module .\ACEventService.psd1`
7. Run: `Install-ACEventService`
8. Enter the IP address of the AC-Hunter server.
9. Enter the name of the existing dataset that will hold the deception data (e.g., localhost-rolling). The dataset name **must** match the name of an existing dataset in AC-Hunter
10. To create and register a new user to be monitored, run: `New-ACESUser`
11. To create a new file to be monitored, run: `New-ACESFile`

Perform **Steps 1-9** above on any system to which you transferred the archive.

Step 10, creating a user token, can be performed from any of the systems and only needs to be done once for any user token.

Step 11, creating a file token, must be performed on the system that is hosting the file token.

After **Step 7** is performed, all of the files in the archive folder will be moved to **C:\Program Files (x86)\Active Countermeasures\AC Event Service**. The ACEntService.psd1 import does not persist between PowerShell sessions. To add deception tokens or make configuration changes after exiting the initial PowerShell session, you will need to:

1. Open a PowerShell prompt as an Administrator
2. Navigate to `C:\Program Files (x86)\Active Countermeasures\AC Event Service\`
3. Run: `powershell -exec bypass`
4. Run: `Import-Module .\ACEntService.psd1`

There is an intentional 5-minute delay between account/file creation and alerting. This is to prevent false-positives from scanners and other sources. Additionally, it may take up to 1 hour for the configuration changes to propagate back to agents to inform them to monitor newly-created tokens. Alerts will not be generated for newly-created deception tokens until the agents have received the latest configuration. You can force the agents to grab the latest configuration by manually restarting the **ACES** service on each system.

BeaKer Installation

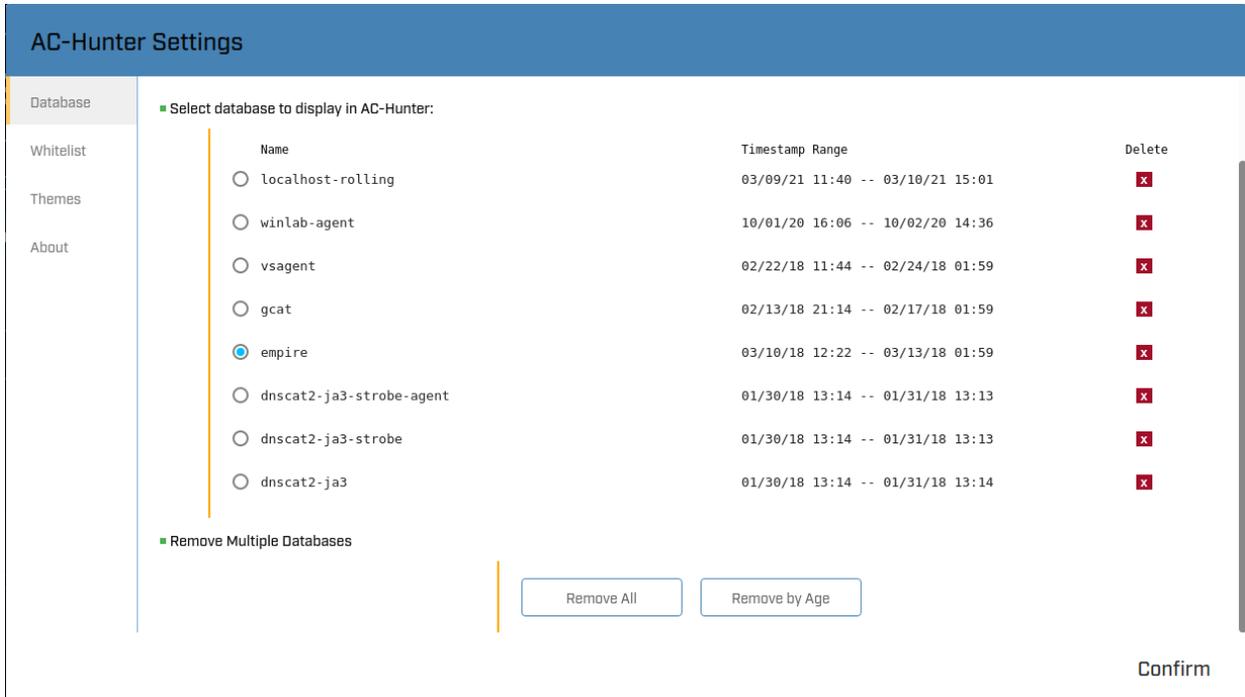
The easiest way to get set up with BeaKer is to use the official installer, which will connect BeaKer to your AC-Hunter install automatically. If you need to modify the location of your BeaKer system after install you can do so by finding and modifying the line containing "BeakerHost" in AC-Hunter's config file (`/etc/AC-Hunter/config.yaml`).

To apply your changes be sure to save the config file and then restart AC-Hunter with:

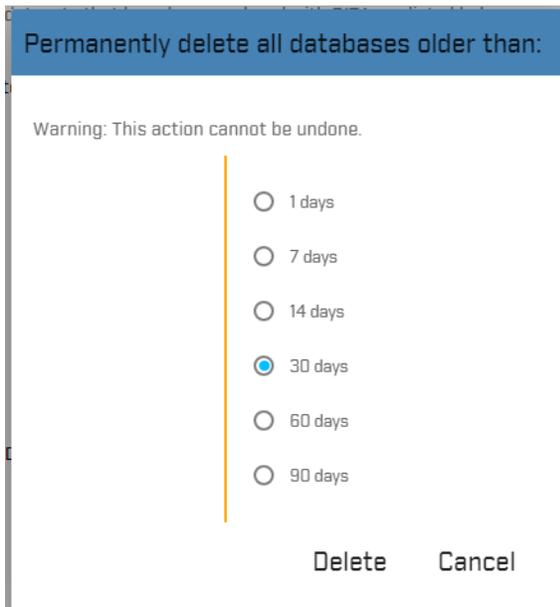
```
hunt up -d --force-recreate
```

Managing Databases

Earlier in this document we walked you through clicking the gear icon followed by "Database" to load different databases for analysis. This screen can also be used to manage your databases.



Note that to the right of each database is an "X". Clicking this icon will prompt you to confirm deleting that specific database. At the bottom of the screen are two buttons. One provides a quick way to delete all databases, while the other allows you to delete databases based on when they were created. Clicking "Remove by Age" will produce the following pop up:



This provides a quick way to implement retention policies and delete data that is no longer needed. For example, if I always want to make sure I have 30 days worth of data available for

review, I can select the "30 days" option which will delete all databases older than this time period.

Note that the "rolling" database should never be deleted as it is used to store current data as it is collected. Neither the "Remove All" nor the "Remove by Age" button will remove the rolling database.

Note that the Community Edition is limited to saving 10 databases, so we encourage you to delete any unneeded databases when you get close to the limit.

Importing Packets from a PCAP File

If you have a packet capture file and want to analyze the content, we include a utility that will create a database from it. To use it, run:

```
/usr/local/bin/import_pcaps.sh -p /full/path/to/pcap_filename.pcap -d database_name
```

Managing Safelists

We have changed the way that safelists function in AC-Hunter. In the past, safelists were associated with a specific day's worth of data. Safelists are now persistent in later databases. So whatever safelists you created on Monday will be carried forward into Tuesday's data and then on into Wednesday, etc. Further, we expect most folks will simply work with the rolling database as that will always include the most current threat hunt info. Obviously any safelists that get created here will remain persistent until they are manually removed.

You can manage your safelists by clicking the Dashboard icon and then the gear icon in the top right of the screen. From the AC-Hunter Settings screen, click the Safelist tab. You should see a screen similar to the following:



There is also a button to remove all safelist entries. This is handy in case you want to backup your safelists but then remove them all to ensure you can review all available data.

By clicking on the View/Edit button you can see all of the safelist entries that are being applied to the current dataset. By clicking the drop-down-triangle to the right of the safelist entry, you can see the details associated with that safelist entry.

VIEW / EDIT GLOBAL SAFELIST

Global Safelist Entries

Search: Ex. 10.10.10.10 type: --- scope: ---

name ↑	type	scope	comment	actions
8075	asn	dst		▼ ✕
Alice Late: 10.55.200.10	ip	src		▼ ✕
Carol: 10.55.182.100--Public: 208.185.50.90	pair			▼ ✕
MICROSOFT-CORP-MSN-AS-BLOCK	asn_org	dst		▼ ✕
Public: 64.4.54.253	ip	dst		▼ ✕
Public: 66.235.0.0/16	cidr	dst		▼ ✕
baddns.r-ix.com	domain_literal			▼ ✕

|< < 1 / 2 > >|

Close

You can edit the safelist from this screen by clicking on the pencil (edit the comments). When you've finished editing the entry, make sure you press the "save" button before pressing "Close".

You can also delete just the single selected safelist entry by clicking the "X" icon. Be careful as AC-Hunter will delete this safelist entry without performing a secondary prompt.

Since the Community Edition is limited to 50 safelist entries, you can delete old safelist entries here to make space for new ones.

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.

Modifying the Sensor Name

On a Zeek, Active-Flow, or Espy sensor create the file `/etc/rita/agent.yaml` and put "Name: YourCustomNameHere" in it. Make sure the file is readable by at least the same user that is set to run the `zeek_log_transport.sh` script. Replace "YourCustomNameHere" in the commands below with your desired name.

```
sudo mkdir -p /etc/rita
```

```
echo "Name: YourCustomNameHere" | sudo tee -a /etc/rita/agent.yaml
sudo chmod +r /etc/rita/agent.yaml
```

After these changes are made, the next time the `zeek_log_transport.sh` script runs it will use this new name. Within a few hours you should see the new name show up in the AC-Hunter interface.

Configuring User Accounts for the Web Interface

Managing Internal AC-Hunter User Accounts

The "`manage_web_user.sh`" script is used to create, delete, and reset internal user accounts. Run `manage_web_user.sh -h` to see the available options.

System Maintenance

Log Maintenance

We suggest you check the amount of free space on both the Zeek and AC-Hunter systems weekly. The logs and databases used can quickly fill up the available drive space.

Deleting Zeek Logs

Zeek logs that accumulate on the Zeek system can be configured to expire and be automatically deleted after a certain amount of time. The setting can be found in the `zeekctl` configuration file (`/opt/zeek/etc/zeekctl.cfg`).

```
# Expiration interval for archived log files in LogDir.
# Files older than this will be deleted by "broctl cron".
# The interval is an integer followed by one of these time units:
# day, hr, min. A value of 0 means that logs never expire.
LogExpireInterval = 0
```

For instance, if you wanted logs to automatically be deleted after 30 days you would modify the setting to be:

```
LogExpireInterval = 30 day
```

This will automatically remove log files located in the `/opt/zeek/logs/` directory on your Zeek system.

Deleting RITA Logs/Databases

Once the Zeek logs are copied to the RITA/AC-Hunter system, they are processed by RITA and then passed to AC-Hunter. In the "Managing Databases" section, we identified how to delete databases from within AC-Hunter. This method is if you cannot access the AC-Hunter interface or prefer to use a script.

First, connect to the RITA/AC-Hunter system with ssh. Run the command:

```
df -h
```

to see how much space is available. If this amount is low, run

```
rita list
```

The command will then display all of the databases that are currently stored in RITA. Once you've decided which ones to delete, you can specify either a specific database to delete on the command line or a pattern that matches a number of them. For example, to delete all databases from July 2018, run:

```
rita delete -m RITA-2018-07
```

You can specify more than one pattern like this:

```
rita delete -r ' (^RITA-2017|^RITA-2018-0[123]) '
```

which would delete all databases from 2017 and January through March, 2018.

To see if you have enough free space, run the following command again:

```
df -h
```

You can repeat the above steps as required.

More information on log and database management

For far more detail on managing log files and databases, please see the FAQ entry at https://portal.activecountermeasures.com/support/faq/?Display_FAQ=4970 .

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.



References

This project uses data from the GeoLite 2 Database created by [Maxmind](#) and distributed under the [Creative Commons Attribution-ShareAlike 4.0 International License](#).



Online FAQ: <https://portal.activecountermeasures.com/support>

Email Support: support@activecountermeasures.com