

# Lumen<sup>®</sup> Adaptive Threat Intelligence

User guide | June 2021

LUMEN<sup>®</sup>



This document provides an overview of the Lumen Adaptive Threat Intelligence (ATI) portal and is used to convey to customers and prospects what to expect of ATI. The sections in this document also include internal implementation procedures, support processes, operational guidelines, etc. This document will be updated as necessary to reflect any changes to how the ATI is offered and delivered.

## Service description

The sophistication of cyber threats, and the complexity of maintaining traditional point network security solutions, is driving the adoption of managed security service solutions utilizing threat intelligence. As your organization upgrades your security posture to be more proactive, you want threat intelligence that is actionable and can be integrated with protection actions. Lumen Adaptive Threat Intelligence (ATI) service helps address these business challenges by providing monitoring and alerting of internet-based threats to help protect your users, website or critical applications on the internet.

ATI is an always-on, network-based, near real-time monitoring, threat correlation and alerting service that provides alerts about the traffic to and from your IP addresses monitored by Lumen, and other IPs on the internet. ATI monitors data samples flowing across Lumen's global network infrastructure obtaining information about traffic flows between your network and the other end of the IP communication. The sampled information is subsequently correlated by the ATI against the Lumen database of known malicious IPs. If the sampled information matches a malicious IP, a record is created (an "event") that is forwarded in near real-time to the ATI portal. Information about events is also aggregated and sent to you via email periodically. The service is available in two cloud-based options called Enhanced and Premium Adaptive Threat Intelligence service. If you subscribe to the Premium ATI, events may also be forwarded in near real-time to your security information and event management (SIEM) platform.

Lumen has made a major investment in developing a threat research and engineering group called Black Lotus Labs. The Black Lotus Labs team has developed threat sensing capabilities using one of the world's largest IP backbones. Malicious behaviors are detected off the backbone and classified using sophisticated machine learning algorithms and automated validation infrastructure. Additionally, Black Lotus Labs validates Indications of Compromise (IOCs) that are conveyed using third-party resources. The extra effort pays off in the cultivation of a very high-fidelity threat set. Customers benefit in several ways including:

- Real-time visualization of their interactions with malicious entities
- Botnet research and take-down efforts keeps the backbone safer
- Automated deployment of countermeasures when new threats are discovered via Black Lotus Labs
- Leading botnet research and publication

## Product levels

The ATI Enhanced option includes:

- Monitoring of customer's traffic as it passes through the Lumen infrastructure based on sampled network analysis
- Correlation of traffic against malicious IPs and Domains utilizing Lumen proprietary analysis and threat data
- Near real-time alert notification for events

- Access to near real-time, portal-based reporting and analytics capabilities utilizing event data
- SOC support with 24-hour response time to obtain additional Event information if available
- 99.99% SLA for ATI portal availability, and 2 minutes SLA for event notification, from the time Lumen became aware of the event

The ATI Premium option comes with:

- All features included in the Enhanced option of the Service
- Near real-time event feed to customer's SIEM platform
- Integration with Rapid Threat Defense services that enables automated responses to detected threats based on a customer selected Security Posture
- SOC support with 2-hour response time for high priority events

## Domain analytics

ATI tracks malicious indicators at both the IP address and domain level. Detection of customer interactions with malicious domains requires one of the 3 following conditions:

1. Lumen secure DNS for their primary DNS service. With this option, customers get their own instance of service and Lumen can facilitate selective profiles for each customer, enabling automatic blocking of malicious entities that meet or exceed customer-selected risk score (see Rapid Threat Defense). Interactions with malicious domains will be detected and displayed in portal reports and be conveyed in syslog data transmitted to customers.
2. Lumen public DNS service. This option allows for the detection of interactions malicious domains. These interactions will be displayed in the threat intelligence reports and conveyed in the syslog data transmitted to customers.
3. Customer supplied logs from their DNS service are being ingested into ATI. This option allows for the detection of interactions malicious domains. These interactions will be displayed in the Threat Intelligence reports and conveyed in syslog data transmitted to customers.

## SOC assistance

Adaptive Threat Intelligence (ATI) includes phone-based consulting hours with personnel located in our Security Operations Center (SOC). The quantity of hours per month scales with the product tier and level that you have purchased. (See the below chart.)

Product Tier	Product Level	Monthly Consulting Hours
Small	Enhanced	4
Medium	Enhanced	8
Large	Enhanced	12
Small	Premium	8
Medium	Premium	12

Large	Premium	16
-------	---------	----

In most cases, the SOC can determine more information about the specific threat event and general strategies for mitigating the effect of the event. Please note: The SOC does not have information specific to the architecture of your LAN network environment and cannot recommend specific policy changes to your environment. To initiate a consultation, please submit a portal trouble ticket requesting an ATI consultation.

## Rapid Threat Defense

Rapid Threat Defense is a capability that is being implemented across the Lumen Security product portfolio. Customers specify a “security posture” which has an associated risk score. When malicious entities are discovered that have a Risk Score that meets or exceeds the risk score indicated in the security posture, countermeasures will be automatically deployed to block access to that malicious entity.

The following products have Rapid Threat Defense capabilities:

Product	Level	Capability
Adaptive Network Security	Premium	Block malicious sites by IP address
Adaptive Network Security	Standard	View interactions with malicious sites (no blocking)
Adaptive Threat Intelligence	Premium	Use Secure DNS to block malicious sites by domain
Adaptive Threat Intelligence	Standard	View interactions with malicious sites (no blocking)

Rapid Threat Defense enables a customer to select a security posture that will block malicious sites that have a risk score that meets or exceeds that risk score indicated by the security posture. The selections are as follows:

- No Blocking (default): No indicators will be automatically blocked
- Confirmed Threats: Block contact with indicators having a risk score = 100
- Very High Risk and Confirmed Threats: Block contact with indicators having a risk score > 80

- High Risk Very High Risk and Confirmed Threats: Block contact with indicators having a risk score > 60

Security Posture	
Posture	Description
<input type="radio"/> No Blocking	No indicators will be automatically blocked.
<input type="radio"/> Confirmed Threats	Block contact to indicators with Risk Score = 100
<input checked="" type="radio"/> Very High Risk and Confirmed Threats	Block contact to indicators with Risk Score > 80
<input type="radio"/> High Risk, Very High Risk, and Confirmed Threats	Block contact to indicators with Risk Score > 60

Selecting a security posture sets up automated deployment of countermeasures whenever new malicious entities are discovered by [Black Lotus Labs](#) (The Lumen Cyber Threat Intelligence team). The Black Lotus Labs team has automated the discovery, classification and validation of new malicious entities. Typically, countermeasures are deployed in under 30 minutes from discovery of the new malicious entity.

### Use case 1: Rapid Threat Defense with Adaptive Network Security

A customer has selected a security posture of “Very High Risk and Confirmed Threats” to block access to malicious entities that have a risk score of 80 and above. The customer has Adaptive Network Security – Premium. Black Lotus Labs discovers a new instance of a “trickbot command and control” server at IP address 101.2.3.4 and assigns a risk score of 87 to it. Within 30 minutes, a “Deny” entry will be added to the customer’s instance of Adaptive Network Security, preventing any interaction with this newly discovered malicious site.

### Use case 2: Rapid Threat Defense with Adaptive Threat Intelligence

A customer has selected a security posture of “Very High Risk and Confirmed Threats” to block access to malicious entities that have a risk score of 80 and above. The customer has Adaptive Threat Intelligence – Premium. Black Lotus Labs discovers a new instance of “a gafgyt command and control” server at domain “ExtremeDiscounts.com” and assigns a risk score of 91 to it. Within 30 minutes, a “Deny” entry will be added to the customer’s instance of secure DNS, redirecting any users that attempt to contact the domain to a “redirect page” for an explanation.

In addition to selecting a security posture, customers can also select specific domains or IP addresses to allow or block. The effect is to override any countermeasures deployed by the security posture selection.

- Block IP Address: Always block access to the IP address in this entry
- Allow IP Address: Always allow access to this IP address, even if it is identified as malicious and has a risk score that meets or exceeds the risk score associated with the selected security posture

- Block Domain: Always block access to the domain in this entry
- Allow Domain: Always allow access to the domain, even if it is identified as malicious and has a risk score that meets or exceeds the risk score associated with the selected security posture

Allow and block tables for IP addresses and domains:

### Domain Block List

Item #	Domain	Description		
0	yahoo.com,	No one in my company should go to Yahoo ... ever.		
1	youshouldnotbehere.com	Not here!!		
2	wickedbaddomain.com	This site is wicked bad		
3	donotgothere.com	Malicious website		
4	badsite.com	Really bad site		

Page: 1 Rows per page: 5 1 - 5 of 5

### Domain Allow List

Item #	Domain	Description		
0	wine.com	Never block wine.com. Needed for sanity maintenance.		

Page: 1 Rows per page: 5 1 - 1 of 1

### IPv4 CIDR Block List

▲ Maximum 50 Entries Allowed ▲

Item #	IPv4 CIDR	Description		
0	192.166.166.254/31	Blocked C2 callback		
1	8.8.8.8/32	Evil		
2	87.0.0.0/10	My least favorite subnet		
3	192.167.12.163/32	Malware node		

Page: 1 Rows per page: 5 1 - 4 of 4

### IPv4 CIDR Allow List

▲ Maximum 50 Entries Allowed ▲

Item #	IPv4 CIDR	Description		
0	123.45.67.89/32	My favorite subnet		
1	80.0.0.0/8	Benign subnet		

Page: 1 Rows per page: 5 1 - 2 of 2

---

## Lumen Security Solutions portal

The ATI portal provides you with access to online reports showing near real-time event information for current and past threat events by geography, traffic type, IP address, duration, and much more, as outlined below. These interactive reports allow users to zoom in on location, manipulate queries, and quickly search for critical information based on key information such as an IP address.

**Note:** Only users who have been set up with the managed security services permission and two-factor authentication can access the Lumen Security Solutions section of Control Center.

### Support contacts

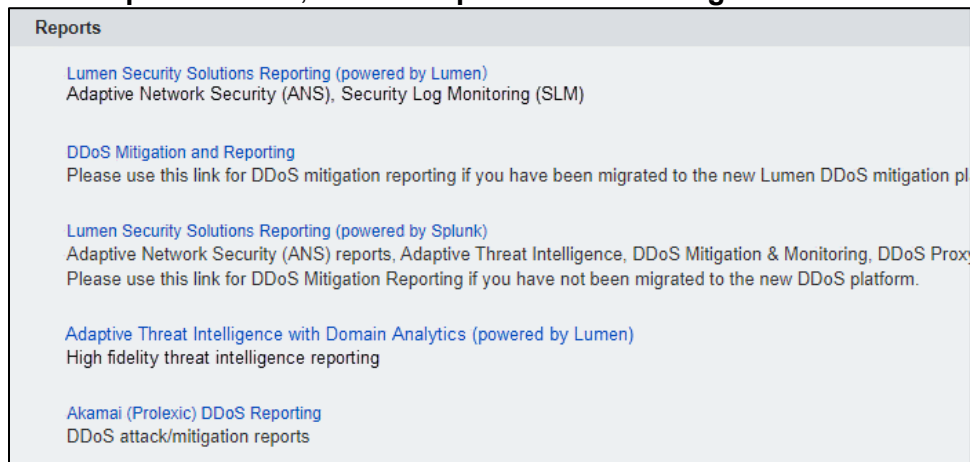
[Access security support contacts](#)

### Accessing the Security Solutions portal

**Note:** Supported internet browsers are Chrome, Safari and Firefox. Use of unsupported browsers will likely result in reduced functionality.

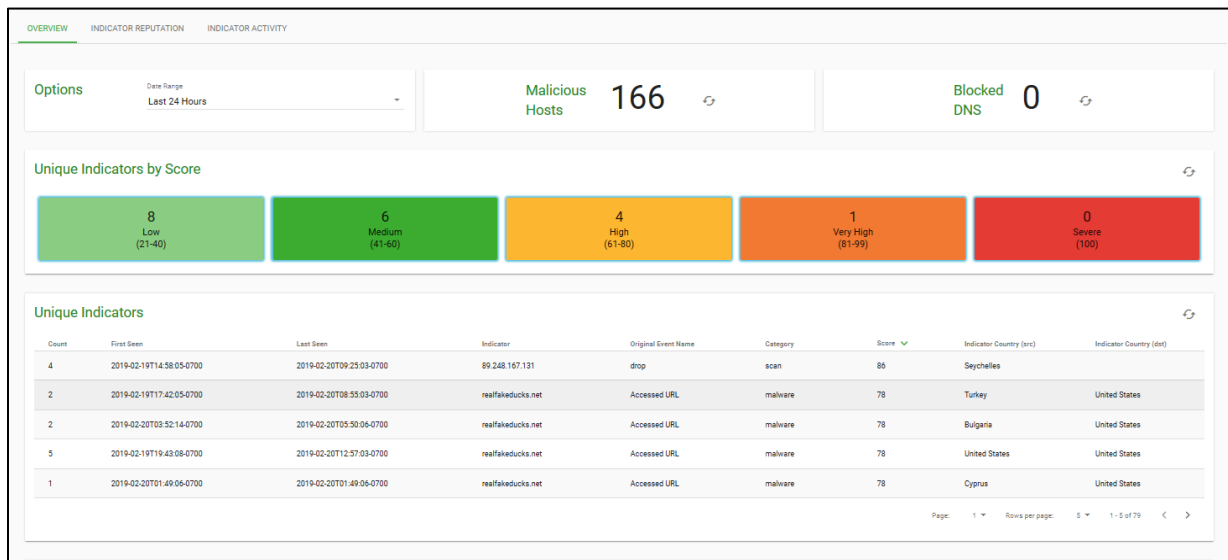
[Learn how to sign in to the Security Solutions portal](#)

1. In the **Reports** section, select **Adaptive Threat Intelligence with Domain Analytics**.



## ATI portal service dashboard

The ATI portal dashboard is the landing page containing the reports and available via the first-tier threat indicators (Figure 5). The dashboard provides at-a-glance status of the threat indicators targeted at the enterprise for the selected time range. The color-coded bar of “Unique Indicators by Score” shows indicators by risk score and provides filtering capabilities. To remove indicators of a specific range, simply click on the color block.



The table of unique indicators column headings are described in the following table.

Column	Description
<b>Count</b>	Quantity of Interactions with this indicator
<b>First Seen</b>	The date and time this indicator was first reported. Note, for clarity, the time zone offset from GMT is included.
<b>Last Seen</b>	The date and time this indicator was last reported. Again, the time zone offset from GMT is included
<b>Indicator</b>	The IP address or domain that is hosting the malicious indicator
<b>Original Event Name</b>	Event as identified from the reporting source for this indicator. Some event names include: <ul style="list-style-type: none"> <li>Threat flow: The set of malicious indicators identified by IP address that has been collected and curated by Lumen Black Lotus Labs</li> <li>Block DNS: Events that have been generated by users attempting to reach Domains that are malicious and meet the customer’s criteria for blocking at DNS or that have been explicitly blocked by the customer administrator.</li> <li>Allow DNS: Events that have been generated by users connecting with domains that have been explicitly allowed by the customer administrator.</li> </ul>
<b>Category</b>	Threat category as explained above



<b>Max Score</b>	The highest risk score presented by this indicator. Note that any indicator may be participating in multiple threat campaigns and hence may have multiple risk core associated with it. The score in this column is the maximum value associated with this indicator
<b>Indicator Country (src)</b>	The geographic location from where this indicator is sourced
<b>Indicator Country (dst)</b>	The geographic location targeted by this indicator

ATI associates a threat category for each indicator that is listed. An indicator may be associated with multiple threat categories. The threat categories are described in the table below:

Category	Description
<b>C2</b>	C2 is shorthand for “command and control”. Each botnet has a C2 entities that manage the activities of the botnet.
<b>Attack</b>	These entities attempt to penetrate the peripheral defenses of an enterprise typically using “dictionary” attacks to crack passwords on publicly addressable assets.
<b>Bot</b>	Entities that have been compromised to participate in the activities a botnet
<b>Malware</b>	Entities that distribute malware for the purpose of compromising assets to gain access to intellectual property
<b>Phish</b>	Entities that proliferate communications for the purpose of collecting credentials to valuable assets. Phishing can use email, phone calls, text, IM and other vectors for this purpose
<b>Scan</b>	Entities that probe the peripheral defenses of an enterprise for the purpose of discovering accessibility, typically pinholes in firewalls.
<b>Spam</b>	Entities that distribute communications for the purpose of attracting attention to services that are generally considered irrelevant to the business of the enterprise targeted.
<b>Anonymous Proxy</b>	Also known as “Proxy” or “TOR (The Onion Router).” Adversaries typically attempt to obfuscate their presence on The internet by positioning behind an anonymous proxy service. Enterprises rarely have legitimate business associated with these entities, so communications with them is typically of interest.

The following indicators may also be present alongside the primary indicators.

Indicator	Description
<b>Popular</b>	Associated with IP addresses that are identified as malicious, but also have many services behind them, lowering the probability that the enterprise is communicating with the specific malicious entity. For instance, the IP address may belong to a hosting provider that has potentially thousands of domains behind it. The risk score associated with this indicator will be diminished to reflect the lower probability of direct interaction.
<b>CDN</b>	Associated with IP addresses that are identified as malicious but are also part of a CDN (content delivery network). This substantially lowers the probability that the enterprise is in direct contact with the malicious entity. The risk score associated with this indicator will be diminished to reflect the lower probability of direct interaction.

Each column can be sorted by ascending or descending with successive clicks. At the bottom of each table are selectors for:

- Tables that exceed the number of entries per page, display a “page” selector is displayed. The default is 1, but clicking the down-arrow icon will list the pages that can be selected
- Number of table entries per page. The default is 5. Select the down-arrow icon to specify how many entries to see in a single display
- The specific entries of the total set are being displayed
- Previous and Next selectors

Note that all tables throughout the ATI reports have these capabilities.

Selecting any unique indicator will bring up the indicator reputation page for that selected indicator.

Scrolling down the dashboard page will display an interactive map that shows the geographic depiction of indicators (Figure 6). The color of the circle represents the highest risk score for all the indicators sourced in that location. The number inside the circle conveys the quantity of unique indicators in that location.



Selecting the +/- controls will cause the map to zoom in or out. Zooming can also be initiated from a mouse wheel. Selecting an indicator group (circle) displays the individual indicators with a presence indicator. Hovering over the presence indicator will display further information on the selected indicator.

## ATI Indicator Reputation page

The ATI Indicator Reputation page can be reached through either the report selector at the top of the page (below), or by selecting a unique indicator from the dashboard.

The top of the Indicator Reputation page displays summary reputation data for that indicator:

The screenshot shows the 'INDICATOR REPUTATION' tab selected. The main heading is 'Results For Realfakeducks.net For Last 24 Hours'. Below this is a search section with a 'Date Range' dropdown set to 'Last 24 Hours' and a 'Type Indicator Here \*' field containing 'realfakeducks.net'. A search icon is on the right. Below the search is a 'Threat Summary' section with a refresh icon. It contains a table with the following data:


Families	Categories	Sources
fake_antivirus_scan_location	malware	symantec

At the bottom right of the table, there are pagination controls: 'Page: 1', 'Rows per page: 5', and '1 - 1 of 1' with navigation arrows.

A different indicator can be selected by typing it into the Indicator field:

This screenshot is similar to the first one but with 'WickedBadSite.com' entered in the search field. A callout box on the right side of the page contains the text 'Enter IP Address or Domain Here' with a line pointing to the search input field.

Scrolling further down the page will display several quantities that have been measured against this indicator. The score is the risk score for this indicator over time. Risk scores can change over time depending on new information that has been released about the threat. Risk scores reflect a combination of severity—an assessment of the potential damage exposure to this threat can cause—and confidence. The confidence component of the risk score is displayed separately, and again, is a graph over time. Several factors can affect the confidence score, including the direct validation of the threat by the Lumen threat research and operations arm, Black Lotus Labs. Here’s a sample of graphs showing threat measurements over time:



Scrolling further will bring up the first of three tables, the passive DNS entries for the indicator:

### Passive DNS

Timestamp	Sources	Entity	Device Host	Request Domain
2019-02-19T16:12:04-0700	42.213.86.3	A	173.0.0.1	realfakeducks.net
2019-02-19T17:07:03-0700	136.8.73.251	A	173.0.0.1	realfakeducks.net
2019-02-20T01:35:12-0700	224.128.195.220	A	173.0.0.1	realfakeducks.net


Page: 1 Rows per page: 5 1 - 3 of 3







The second table displays the Unique (aggregated) Reputation Records:

Count	Entity	Category	Source	Family	Score	Category Confidence	First Seen
12	realfakeducks.net	malware	domain	fake_antivirus_scan_location	78.0	65	

Page: 1 Rows per page: 5 1 - 1 of 1

Clicking on the icon to the left of the unique reputation record, will display all the individual reputation records that make up the unique record:

Individual Reputation Records 

Entity	Category	Source	Family	Score	Category Confidence	Creation Time	First Seen
 realfakeducks.net	malware	domain	fake_antivirus_scan_location	78	65	2019-02-20T01:23:04-0700	2019-01-28T15:26:18-0700
 realfakeducks.net	malware	domain	fake_antivirus_scan_location	78	65	2019-02-19T23:29:03-0700	2019-01-28T15:26:18-0700
 realfakeducks.net	malware	domain	fake_antivirus_scan_location	78	65	2019-02-19T22:06:03-0700	2019-01-28T15:26:18-0700
 realfakeducks.net	malware	domain	fake_antivirus_scan_location	78	65	2019-02-19T19:43:03-0700	2019-01-28T15:26:18-0700
 realfakeducks.net	malware	domain	fake_antivirus_scan_location	78	65	2019-02-19T15:36:05-0700	2019-01-28T15:26:18-0700
 realfakeducks.net	malware	domain	fake_antivirus_scan_location	78	65	2019-02-20T08:08:04-0700	2019-01-28T15:26:18-0700

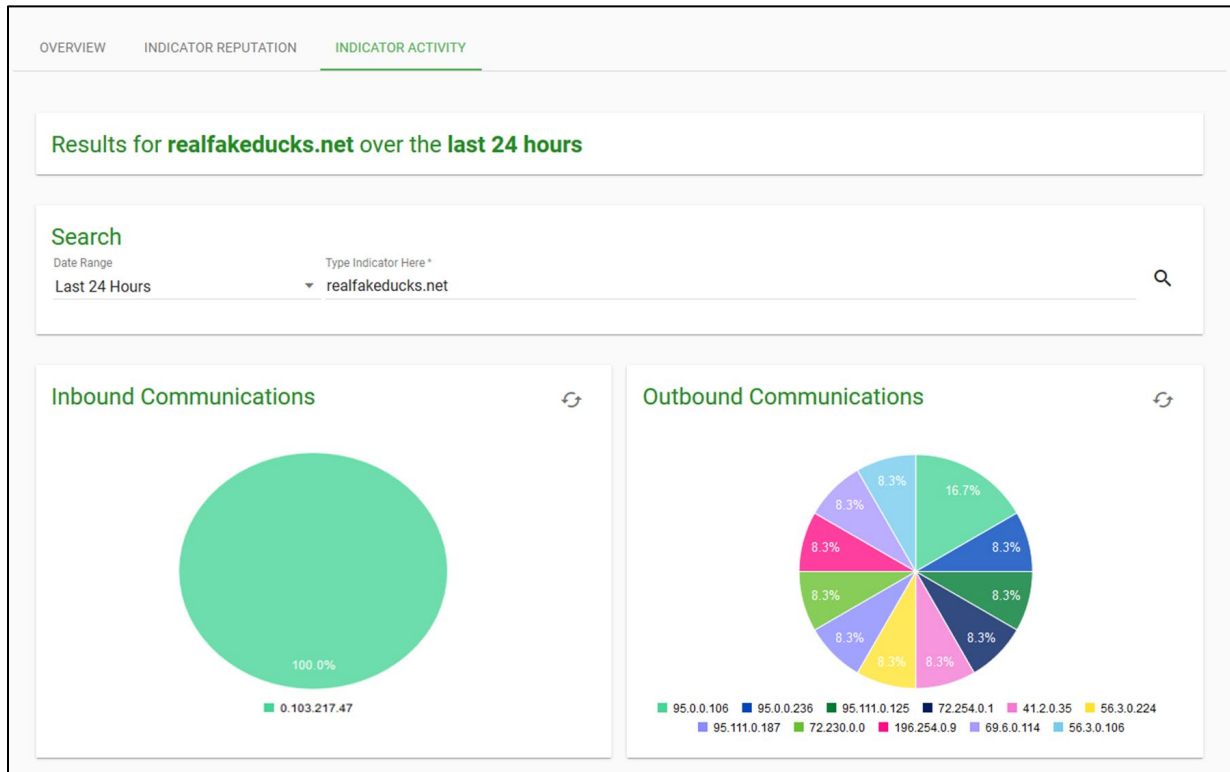
To get all the raw data that makes up the individual reputation record, click the icon to the left of the record selected. A window opens with the data fields. Note that it may scroll for several pages, so only a fraction of available data is shown in below:

View Event

Field	Value
Categories	malware
Category Action	Detect
Category Class	App
Category Confidence	65
Category Context	Domain
Category Device Type	Threat Intelligence
Category Event Type	Threat Intelligence

## ATI Indicator Activity page

Clicking the Indicator Activity page shows more information on the selected indicator. First is a display of inbound vs outbound IP addresses used by the indicator:



Scrolling down, the display depicts aggregated and individual DNS records:

**Aggregated DNS Records**

Source Address	Destination Host	Answer	Type	Destination Score	Count	First Seen	Last Seen
37.30.61.138	realfakeducks.net	127.0.0.1	A	64.0	3	2019-02-14T12:13:03-0700	2019-02-17T11:27:04-0700

Page: 1 Rows per page: 10 1 - 1 of 1

**Individual DNS Records**

Timestamp	Source Address	Destination Host	Answer	Type	Destination Score	First Seen	Last Seen
2019-02-17T11:27:04-0700	37.30.61.138	realfakeducks.net	127.0.0.1	A	64	2019-02-17T11:27:04-0700	2019-02-17T11:27:04-0700
2019-02-14T12:13:03-0700	37.30.61.138	realfakeducks.net	127.0.0.1	A	64	2019-02-14T12:13:03-0700	2019-02-14T12:13:03-0700
2019-02-16T11:10:30-0700	37.30.61.138	realfakeducks.net	127.0.0.1	A	64	2019-02-16T11:10:30-0700	2019-02-16T11:10:30-0700

Page: 1 Rows per page: 10 1 - 3 of 3

The last two tables on this page show the IP data for both the aggregated and individual connection records:

## Aggregated Connection Records



Count	Source Address	Destination Address	Source Score	Destination Score	First Seen	Last Seen	Volume (MB)	Hours
5	241.35.55.164	167.83.229.73	53.0	0.0	2019-02-18T10:51:04-0700	2019-02-20T09:24:04-0700	1365.0	46

Page: 1 ▾ Rows per page: 10 ▾ 1 - 1 of 1 < >

## Individual Connection Records



Timestamp	Source Address	Source Port	Source Score	Destination Address	Destination Port	Destination Score	Octets	Packets	TCP Flags	Protoc
2019-02-19T18:10:08-0700	241.35.55.164	123	53	167.83.229.73	443	0	273	1	24	tcp
2019-02-18T14:04:06-0700	241.35.55.164	123	53	167.83.229.73	443	0	273	1	24	tcp
2019-02-18T17:53:04-0700	241.35.55.164	123	53	167.83.229.73	443	0	273	1	24	tcp
2019-02-18T10:51:04-0700	241.35.55.164	123	53	167.83.229.73	443	0	273	1	24	tcp
2019-02-20T09:24:04-0700	241.35.55.164	123	53	167.83.229.73	443	0	273	1	24	tcp

<  >

Page: 1 ▾ Rows per page: 10 ▾ 1 - 5 of 5 < >

## Appendix A: Threat flow syslog data format

Threat flow events will be sent in pipe delimited (PSV) format for SIEM, each event corresponding to a single row.

Two formats are currently supported: version 1.0 and 1.1. Version 1.0 is only available for backwards compatibility. Version 1.1 eliminates Lumen internal fields (see red highlights in table below), adds threat scores with a range from 0 to 1 (green highlights), and changes the format of an IP address from a long to a dot-decimal notation.

The following shows sample events in PSV format for both versions, with a listing of the fields in sequence below:

### Format for version 1.0

```

1      2      3      4      5      6      7      8      9      10 11 12 13 14 15 16 17 18 19 20 21
1.0|1459952929|71678725|6|123456789|123|49554|16|987654321|321|443|24|802|646|52|1|16|bot| |s|abcde
  
```

### Format for version 1.1

```

1      2      3      4      5      6      7      8      9      10 11 12 13 14 15 16 17 18 19 20
1.1|1459952929|6|10.0.0.1|123|49554|16|10.0.0.2|321|443|24|52|1|16|bot| |0.68|None|s|abcde
  
```

SIEM Pipe Delimited Syslog Format

1.0 Position	1.1 Position	Field Listing
1	1	Version Number
2	2	Epoch Time Stamp
3	N/A	Agent (internal use only)
4	3	Protocol (Enumeration)
5	4	Source IP (1.0 : Long format, 1.1: Dot-Decimal)
6	5	Source ASN
7	6	Source Port
8	7	Source Mask
9	8	Destination IP (1.0 : Long format, 1.1: Dot-Decimal)
10	9	Destination ASN



11	10	Destination Port
12	11	Destination Mask
13	N/A	Input Interface (internal use only)
14	N/A	Output interface (internal use only)
15	12	Octets
16	13	Packets
17	14	TCP Flags
18	15	Source Threat (optional)
19	16	Destination Threat (optional)
N/A	17	Source Threat Score (0-1 or None)
N/A	18	Destination Threat Score (0-1 or None)
20	19	IP Identification (Enumeration: see below)
21	20	Billing Account Number (BAN)

IP identification is either s (source), d (destination), or b (both) and identifies which IP of the flow belongs to the billing account number. Protocol enumeration follows the IANA (Internet Assigned Numbers Authority) convention, with typical values including:

6 -> TCP

17 -> UDP

1 -> ICMP

47 -> GRE

---

## Appendix B: Flow data

Traffic flow information sampled from the Lumen network contains the fields below.

ATI correlation is performed against the traffic flow information below. Events are created and enriched for the traffic flows that contain known bad IPs.

### Traffic flow data fields

Attribute	Description
Time	Time stamp of the sampled flow.
Source/Destination IP	IP addresses of the source and destination
Source/Destination Port	Port number of the source and destination
Source/Destination Location	City (when available) and country of source/destination IP.
Source/Destination AS	Autonomous system number of the source and destination
Source/Destination Threat	Threat category attributed to source and/or destination
Bytes	Bytes sent by this flow
Packets	Packets sent by this flow
Protocol	The protocol field identifies a protocol that sits above the IP layer that is used in the communication. An example of the protocol field value is TCP or UDP.
Service	Service derived from protocol and port

---

## Appendix C: Receiving syslog events using syslog-ng

Lumen uses an open source syslog-ng for the delivery of threat intelligence events. For customers that want to utilize the same to receive the syslog events, the rest of this section provides the steps to setup the syslog-ng server in your own environment to accept communication from Lumen using syslog over TCP with TLS.

Download syslog-ng: <https://syslog-ng.org/>

For a new syslog-ng server setup the following needs to be done to create a TCP TLS source listener.

Your syslog-ng requires a server certificate. You have two options available to obtain the server certificate for your syslog-ng instance:

- Follow your company's process for acquiring a server certificate.
- Follow the steps below to create a *self-signed* certificate.

### Creating a self-signed certificate

1. On your server find your OPENSSL CA installation usually located at /etc/pki/CA. The certs, crl, newcerts and private directories were already present. If the CA directory is not present, you can create it at /etc/pki

```
mkdir CA
```

```
cd /etc/pki/CA
```

2. If the 4 directories above are not present, create them as follows:

```
mkdir certs crl newcerts private
```

3. Execute the following commands in the CA directory:

```
echo "01" > serial  
cp /dev/null index.txt
```

4. Copy *openssl.cnf* into the CA directory. It was found at /etc/pki/tls/openssl.cnf

```
cp /etc/pki/tls/openssl.cnf openssl.cnf
```

5. You need to make a change to openssl.cnf file. You can edit it using *vi* (or another editor) and make the following changes:

```
Find dir = ./demoCA and change it to dir = .
```

6. Save the changes, exit *vi* and execute this command to create the CA.

```
openssl req -new -x509 -keyout private/cakey.pem -out cacert.pem -days 365 -config openssl.cnf
```

7. You will be asked questions for which you will need to provide answers:

- country name
- state or province
- locality
- organization name
- organizational unit name
- common name (name of the server you are on)
- email address (just enter a dot)

8. Next, you will create the client certificate. Execute the following command in the CA directory:

```
openssl req -nodes -new -x509 -keyout serverkey.pem -out serverreq.pem -days 365 -config openssl.cnf
```

9. Then execute:

```
openssl x509 -x509toreq -in serverreq.pem -signkey serverkey.pem -out tmp.pem
```

10. Then execute:

```
openssl ca -config openssl.cnf -policy policy_anything -out servercert.pem -infiles tmp.pem
```

11. You will be asked for a passphrase: use your own. Make sure the key looks like it should. Respond 'y' to the next 2 questions.

12. Execute:

```
rm tmp.pem
```

The above completes the steps to create a self-signed server certificate.

## Updating syslog-ng with the new certificate

1. Make sure the directories `ca.d`, `cert.d`, and `key.d` exist in `/etc/syslog-ng`
2. If the directories do not exist, create them as follows:

- a. `cd /etc/syslog-ng`

- b. `mkdir ca.d cert.d key.d`

3. Copy your server's CA certificate to */etc/syslog-ng/ca.d/cacert.pem*
4. If you created a self-signed cert, use the following command to copy the CA certificate.  

```
cp /etc/pki/CA/cacert.pem /etc/syslog-ng/ca.d/cacert.pem
```
5. Execute the following command to create a hash value for your CA certificate.
  - a. `cd /etc/syslog-ng/ca.d`
  - b. `openssl x509 -noout -hash -in cacert.pem`
6. View the contents of the directory, there will be a file with a hash for a name. Use it to build this command substituting hash with the name of the file found. Leave the .0 on the end.  

```
In -s cacert.pem hash.0
```
7. Additional copies. The cert file has all the meta data and then the actual BEGIN CERTIFICATE. The key.pem starts immediately with BEGIN PRIVATE KEY.
  - a. Copy your server's private key to */etc/syslog-ng/key.d/server.key*
  - b. Copy your server's public key to */etc/syslog-ng/cert.d/server.cert*
8. If you created a self-signed cert, use the following commands to copy the server key and certificate.
  - a. `cp /etc/pki/CA/servercert.pem /etc/syslog-ng/cert.d/server.cert`
  - b. `cp /etc/pki/CA/serverkey.pem /etc/syslog-ng/key.d/server.key`

This concludes the technical set-up tasks.

9. Add a source for the syslog events to the syslog-ng.conf file

```
source fromLevel3 {
  tcp(ip(0.0.0.0) port(you chose)
  tls( key-file("/etc/syslog-ng/key.d/server.key")
  cert-file("/etc/syslog-ng/cert.d/server.cert")
  ca-dir("/etc/syslog-ng/ca.d")) ); };

log {
  source(fromLevel3); destination(d_tcpTls);
};
destination d_tcpTls { file("/{Your Destination}/Level3Syslog.syslog_tcpTls.log");
};
```

Note the IP will stay 0.0.0.0, don't change it. The port value you chose is entered here and must be relayed to Lumen along with the IP or name of your syslog server. Lumen also needs your CA certificate along with any intermediate certificates.

10. Ensure that your firewall rules are set appropriately to allow for passing syslog traffic to the syslog-ng server.
11. You may also need to place the Lumen server's CA certificate in the ca.d directory in order for the syslog-ng server to trust the incoming connection. The following should be placed in a file named *2e4eed3c.0*:

-----BEGIN CERTIFICATE-----

```
MIIEIDCCAawigAwIBAgIQNE7VVyDV7exJ9C/ON9srbTANBgkqhkiG9w0BAQUFADCBqTElMAK
GA1UEBhMCVVMxFTATBgNVBAoTDHRoYXd0ZSwgSW5jLjEoMCIYGA1UECXMfQ2VydGlma
WNhdGlvbiBTZXJ2aWNlcyBEaXZpc2lvdjE4MDYGA1UECXMvKGMpIDlwMDYgdGhhd3RILCBJ
bmMuIIC0gRm9yIGF1dGhvcml6ZWQgdXNIIG9ubHkxHzAdBgNVBAMTFnRoYXd0ZSBQcmItYX
J5IFJvb3QgQ0EwHhcNMjE3MDAwMDAwWhcNMzYwNzE2MjM1OTU5WjCBqTElMAK
GA1UEBhMCVVMxFTATBgNVBAoTDHRoYXd0ZSwgSW5jLjEoMCIYGA1UECXMfQ2VydGlma
WNhdGlvbiBTZXJ2aWNlcyBEaXZpc2lvdjE4MDYGA1UECXMvKGMpIDlwMDYgdGhhd3RILCBJ
bmMuIIC0gRm9yIGF1dGhvcml6ZWQgdXNIIG9ubHkxHzAdBgNVBAMTFnRoYXd0ZSBQcmItYX
J5IFJvb3QgQ0EwggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQCsoPD7gFnUnM
ekz52hWXMJEEUMDSxuaPFsW0hoSVk3/AszGcJ3f8wQLZU0HObrTQmnHNK4yZc2AreJ1CRf
BsDMRJSUjQJib+ta3RGNKJpchJAQeg29dGYvajig4tVUROsdB58Hum/u6f1OCyn1PoSgAfGcq/
gcfomk6KHYcWUNo1F77rzSlmANuVud37r8UVsLr5iy6S7pBOh94ryNdOwUxkHt3Ph1i6Sk/Ka
AcdHJ1KxtUvkcx8cXlxcBn6zL9yZJclNqFwJu/U30rCfSmNZefl2pSy94JNqR32HuHUETVPm4p
afs5SSYeCaWAe0At6+gnhcn+Yf1+5nyXHdWdAgMBAAGjQjBAMA8GA1UdEwEB/wQFMAMBA
f8wDgYDVR0PAQH/BAQDAgEGMB0GA1UdDgQWBRR7W0XPr87Lev0xkhpqtvNG61dIUDAN
BgkqhkiG9w0BAQUFAAOCAQEAeRHAS7ORtvzw6WfUDW5FvIXok9LOAz/t2iWwHVfLHjp2oEz
sUHboZHIMpKnxulvW1oeEuzLIQRHAD9mzYJ3rG9XRbkREqaYB7FViHXe4XI5ISXycO1cRrK1z
N44veFyQaEfZYGDm/Ac9liAXxPcW6cTYcvnlc3zfFi8VqT79aie2oetaupgf1eNNZAqdE8hhuvU5
Hle6uL17In/2/qxAeeWsEG89jxt5dovEN7MhGITINGDrYyCZuen+MwS7QcjBAVIEYyCegc5C09Y/
LHbTY5xZ3Y+m4Q6gLkH3LpVHz7z9M/P2C2F+fpErgUfCJzDupxBdN49cOSvkBPB7jVaMaA==
-----END CERTIFICATE-----
```