WHITE PAPER

High Fidelity Threat Intelligence

Risk Scores

Dave Dubois, Global Security Product Management
Version: 1.0, Jan 2019

LUMEN℠

Executive Summary

Adaptive Threat Intelligence from Lumen is a market-leading threat intelligence service giving its customers real-time information on interactions with potentially malicious devices. Security analysts – the main consumer of ATI – require High Fidelity Threat Intelligence to prioritize work for themselves, their teams and stakeholder organizations. The Risk Score is a powerful metric to use for such prioritization. This article discusses the various parameters that make up the ATI Risk Score and the dynamic adjustments that are made to it during the threat event life cycle.

LUMEN℠

Terms Used in This Article

Target

An asset or organization that is the subject of a cyber threat.

Asset

A technical computing or communication device — server, laptop, tablet, smartphone or other device that is automated and has the ability to communicate.

Vulnerability

A characteristic of an asset that makes it susceptible to attacks by adversaries. This includes software that has not had security updates applied, computers with easy to guess passwords, connections to unauthorized wireless networks, etc.

Adversary

A person or organization that perpetrates cyber threats for their own gain. The penalty for their actions should be a prison sentence or a hefty fine.

ATI or Adaptive Threat Intelligence

A market-leading product from Lumen that identifies threats directed at target organizations in real-time.

Kill Chain

The progression of activities an adversary perpetrates to infect target assets for their gain.

Category/Threat Category

Types of threats tracked by ATI (examples: Proxy, Bot, Malware, etc.).

Command and Control (C2) Server

A server used by an adversary to control botnets and other infected target assets.

ATI Category Confidence/Category Confidence

A measure of how confident we are that an entity is associated with a specific category or threat type.

ATI Threat Score/Threat Score

An indication of how severe of a threat an entity is. This is based on the confidence in the fidelity of the threat information provided for that entity.

ATI Positive Score/Positive Score

An indication of how positive a reputation an entity may have. For example, an IP address with a high positive score may be associated with a popular domain name as rated by Alexa and other sources.

ATI Risk Score/Risk Score

The overall Risk Score is an accumulation of the Threat Score and Positive Score. It is an indication of the severity of a threat and the confidence in the fidelity of the information provided for that threat.

ATI Threat Set/Threat Set

The intersection of the customer IP space and the known threat events that ATI has in its data base.

Reputation Data

Any information associated with an entity (IP, domain) on the public internet. This data can be threat-based, positive, or neutral and is used to compute the overall Risk Score of an entity.

Indication of Compromise (IoC)

Reputation data for a specific public IP address, domain or sub-domain that indicates the entity is involved in potentially malicious activity or is a confirmed threat.

Source

A feed of reputation data, either external or internal, that provides some reputation information about an entity on the internet. External threat sources typically emanate from a cyber defense organization that researches and tracks cyber threats globally. Internal sources are typically algorithms developed by the ATI Threat Research Team to track and confirm new cyber threats.

ATI Threat Research Team

A team of engineers and data scientists that develop the systems and technology that are used to determine High Fidelity Threat Intelligence information and the technological means to communicate it to customers.

Why Do I Need a Risk Score?

ATI Risk Scores vary from 1–100 and appear in the collection of threat reports available in ATI. Customers find it useful to sort the Threat Events table by whichever column serves their current purpose. Sorting on Risk Score can be an effective way to prioritize threat hunting and incident response activities.

Risk Scores show up in ATI reports in several different areas. On the ATI Dashboard, there is a time-series graph that tracks the maximum risk score trend. Ideally, this should be decreasing over time as security analysts and incident responders prioritize removing the highest priority threats first. The ATI Dashboard also shows the max Risk Score for each ATI Threat Type. These are depicted in the screenshot below.

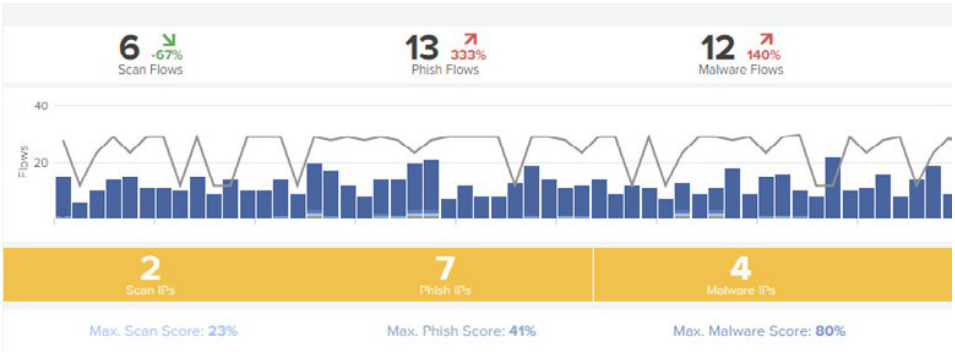

On the ATI Threat Event page, Risk Scores are listed in a column in the Threat Events table as shown in the screenshot below.

Time	Source Tag	Dest. Tag	Source IP	Source Port	Source Side	Source Threat	Source Score	Source AS	Source Org.	Source Location	Dest IP	Dest Port	Dest Side	Dest Threat	Dest Score	Dest AS	Dest Org.	Dest Location
02/06/19 01:27:20 PM			10.105.255.72	80	EXTERN	malware attack	80	3356	Level 3 Communications, Inc.	Mentone, United States	10.68.92.72	59409	INTERN		0	3356	Level 3 Communications, Inc.	Miami Beach, United States
02/06/19 01:26:59 PM			10.105.245.72	80	EXTERN	malware attack	76	3356	Level 3 Communications, Inc.	United States	10.215.47.180	25394	INTERN		0	3356	Level 3 Communications, Inc.	United States
02/06/19 01:33:53 PM			10.242.23.42	30201	EXTERN	attack	72	209	Qwest Communications Company, LLC	United States	10.216.84.192	4500	INTERN		0	3356	Level 3 Communications, Inc.	United States
02/06/19 01:33:38 PM			10.41.78.139	443	EXTERN	attack	72	3356	Level 3 Communications, Inc.	China	10.72.92.135	18909	INTERN		0	3356	Level 3 Communications, Inc.	Seoul, Republic of Korea
02/06/19 01:33:17 PM			10.41.78.139	52149	EXTERN	attack	72	3356	Level 3 Communications, Inc.	China	10.113.3.117	4500	INTERN		0	3356	Level 3 Communications, Inc.	United States
02/06/19 01:32:56 PM			10.41.78.139	443	EXTERN	attack	72	3356	Level 3 Communications, Inc.	China	10.72.92.135	18909	INTERN		0	3356	Level 3 Communications, Inc.	Seoul, Republic of Korea
02/06/19 01:32:48 PM			10.242.23.42	30201	EXTERN	attack	72	209	Qwest Communications Company, LLC	United States	10.216.84.192	4500	INTERN		0	3356	Level 3 Communications, Inc.	United States

Clicking on the Threat Type of any one event will display a view of the snapshot of the most current notification of the threat alongside a snapshot of what the threat looks like now. Positioned below these charts is a time series of the threat that depicts both the severity and the confidence factors of the threat, depicted in the following screenshot.

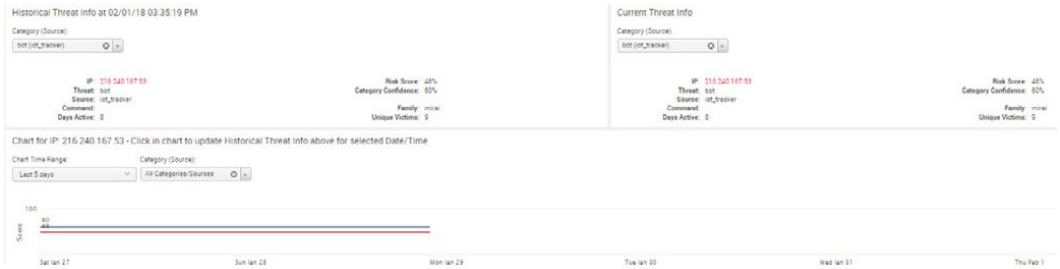

Risk Scores are also incorporated into threat management platforms and applications that consume ATI data. Lumen℠ Secure Log Management (SLM) correlates ATI data with local device logs to indicate which customer-premises devices are under attack.

How is a Risk Score Determined?

If customers are going to use Risk Scores, they typically want to know how they are derived. Broadly speaking, the ATI Risk Score is a combination of three major factors: severity, confidence and time.

Severity

Several subfactors combine to produce the severity factor of the Risk Score, including:

- **Threat Category**
  Primarily determined by how far the adversary has progressed down the Kill Chain. For instance, a "Scan" threat category is not (yet) very dangerous as the adversary is early in the attempt to penetrate the target's peripheral defenses. On the other end of the scale, a high volume (by byte count) conversation between a target asset and an adversarial command and control (C2) server may indicate data exfiltration – a very serious situation that has progressed far down the Kill Chain. The progression of threat categories in ATI include: Proxy, Scan, Phish, Malware, Bot, Attack and Command and Control (C2).

- **Source Provided Risk Score**
  Some Reputation/IoC feeds have a Risk Score associated with the event report. These reports may be updated periodically throughout the lifecycle of the threat event causing changes in the overall Risk Score.

- **Positive rating**
  You must take into consideration the network architecture that the address is in support of when reporting reputation at an IP address level. If the IP address is associated with a single server, then there is a high correlation between the threat event and the IP address. If the IP address is associated with a hosting or CDN service, there may be large quantities of servers hosted behind that IP address, lessening the likelihood that the target asset is interacting with the malicious device that is the subject of the threat event.

Confidence

Several subfactors combine to develop the confidence factor of the Risk Score.

- **Source Confidence**
  The ATI Threat Research Team assess each source before including it in the reputation data and the Threat Set. Each source is assigned a confidence factor based on its reputation and past abilities to correctly predict and determine malicious activities. Sources that are validated and determined to be highly accurate are given a high confidence factor. The higher the confidence we have in the source, the higher the contribution that source provides to the overall Risk Score.

- **Validation**
  When a new malicious entity is reported, the ATI Threat Research Team attempts to validate the new entity and discover more about it. Typically, a surrogate asset is spun off in a sandbox which is used to reverse engineer the dialog to the suspected malicious entity. If the entity responds, then the confidence factor of the Risk Score is elevated. (see “High Fidelity Threat Intelligence: Demystifying Threat Flow Development,” a separate publication).

Risk Score Decay Over Time

It is imperative that threats be removed from the Threat Set in a timely manner. Some sources may remove IoCs as they are remediated while others may have no indication of whether something is still a threat or not. If there is no way to detect the remediation of a threat, then we must observe the last notification we had on it and decay that source’s contribution to the Risk Score over time.

Our research shows that the duration in which a threat is relevant varies between threat categories, indicating that the best practice is to commence the decaying process when the difference between the last notification of the threat (from its source) and the current time exceeds a pre-defined interval that is unique to the category.

The threat continues to decay over an S-Curve. Each ATI Threat Category has its own S-Curve characteristics as depicted in the graph below.

S-Curves
Decaying Logistic graphs
f(t) = 1 - Sigmoid(t)

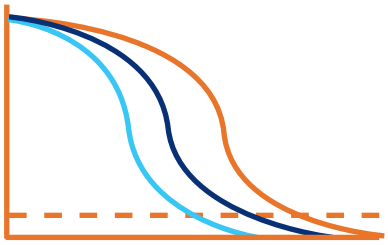

Each threat type gets its own
S-Curve characteristics
t = time

S-Curve characteristics include the steepness of the curve and the total height of the curve. Decaying S-Curves are known as “inverse logistics” curves in mathematics, which are driven by “inverse Sigmoid” functions (a little “geek candy” for readers who are so inclined). Once the Risk Score decays below a predefined threshold, it is removed from the Threat Set.

Conclusion

While there is a significant amount of leading edge research and development that goes into the determination of accurate and relevant Risk Scores, the result for ATI customers is the availability of a simple, actionable metric to use in the prioritization of their work. As always, ATI enables customers to see potential threats before they become breaches, because we are continuously sourcing information from one of the largest IP backbones in the world. The validation and original threat discovery done by the ATI Threat Research Team drives the fidelity of this information to an industry leading level.