



SecurityScorecard

Scoring Methodology

AUGUST 2017

SecurityScorecard.com

info@securityscorecard.com

©2016 SecurityScorecard Inc.

214 West 29th St, 5th Floor

New York, NY 10001

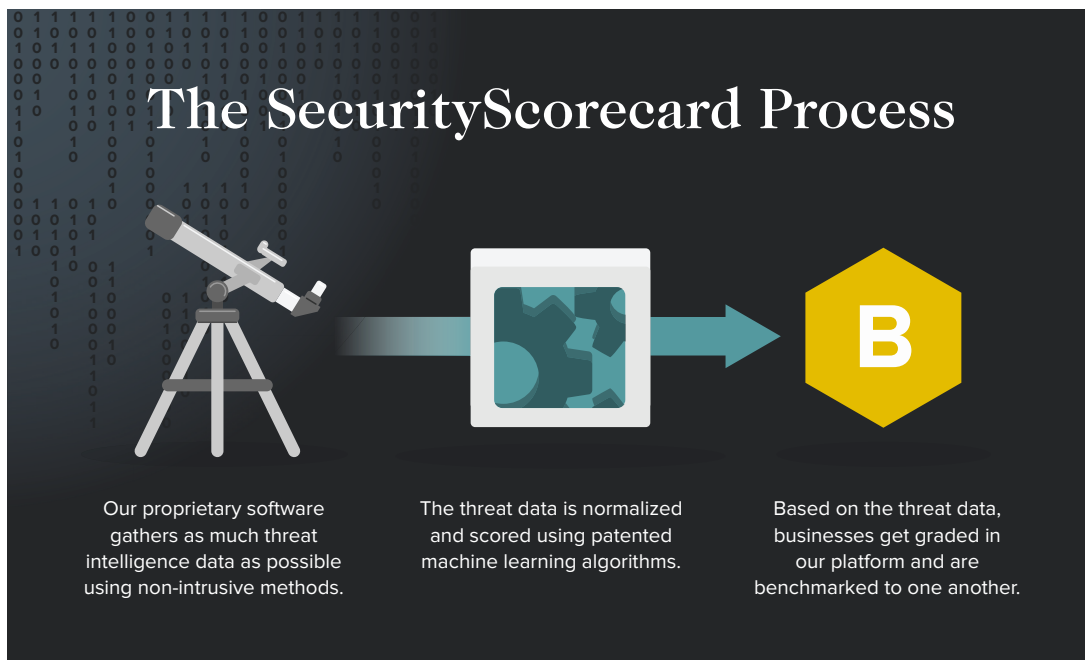
1.800.682.1707

Overview

SecurityScorecard grades the cybersecurity health of organizations based on the information collected by [ThreatMarket](#), our proprietary data engine, as well as our own internal collection activities.

Threatmarket collects information from several sources like data feeds, sensors, honeypots, and sinkholes. Both methods collect data that is externally accessible and public, meaning no intrusive techniques are used to gather the information.

This comprehensive swath of data is then analyzed and appropriately weighted by considering factors such as the severity of the issues, the risk level as defined by industry standards, the overall performance of similar companies, and so on. In particular, comparing the health of a company to that of its peers provides further insights and helps filter out the noise.



This means each company can look at a carefully measured, holistic, and statistically relevant view of the cybersecurity risk associated with its IP footprint and that of its vendors. Ultimately the SecurityScorecard platform reports on whether a company's

behaviors contribute to or mitigate cybersecurity risk over time and provides the user with clear identification of vulnerabilities or gaps in a company's systems. It paints a picture of cybersecurity about a company and its vendors with the appropriate temporal and industry backdrop.

The reality is boards, investors, CISOs and other business leaders need access to all these important details, but they also need a simple and clear language to be able to easily communicate how a company and its third parties manage their security. The SecurityScorecard platform provides this by distilling a company's overall health to a letter grade: something easy to understand whether the consumer is technical or not.

The purpose of this whitepaper is to answer the question: Where does that "A" or "F" letter grade come from?

Understanding Cohorts

As you can imagine, it's important for grades to be contextualized based on company size and industry.

For example, looking at size, companies with a larger digital footprint will generally have a greater 'attack surface' than those with a smaller footprint. As a result, these larger companies have a harder time being able to achieve as low of an absolute issue count as a smaller company.

To correct for problems like this one and to facilitate a meaningful evaluation of cybersecurity risk, SecurityScorecard categorizes companies by IP size and also by industry, putting each company into the relevant bucket, or "cohort."

The IP size buckets are typically structured such that the IP size of the largest company in a given cohort will be no more than 10 times larger than the smallest company in that same cohort.

SecurityScorecard defines which companies are in which industry by using publicly accepted, externally-managed data repositories, such as zoominfo and clearbit.

These cohorts are designed so that there is always a statistically meaningful number of entities in each cohort. (In other words, a large enough sample size is present for each bucket.)

Assigning each company to a given cohort based on its industry and the number of IPs it controls ensures that companies are compared to similar companies—apples to apples.

Confidence in the Data— What are the Inputs?

Before we jump into the more granular aspects like weighting and score changes, let's take a moment to understand what the data actually is. SecurityScorecard collects various types of data on cybersecurity risk.

The majority —about 80 percent-- of the data used in scoring is collected by Threatmarket, our proprietary data collection engine. ThreatMarket collects data in the following ways:

- Scans the entire IPv4 space regularly,
- Operates a battery of sinkholes to track malware infections on client systems on a daily basis,
- Performs a variety of additional collection activities on a non-intrusive basis to identify weaknesses in an entity's cybersecurity posture, such as open ports exposing services that should not be exposed, weak ciphers, out-of-date software with critical vulnerabilities, etc.

In addition, SecurityScorecard supplements its ThreatMarket data with data from public sources and from some third-party commercial sources.

Scoring

Issues: Identification and Weighting

The raw collection data are processed by a set of data analytics to identify potential problems, specifically 77 different cyber-threat related issues.

Overall Grade

SecurityScorecard Security Rating



Issues Graded by Risk Factor



Issue Detection

77 Issue Types

Open ports, Expired or revoked SSL Certificates, Weak Cyphers, Exposed Services, Open DNS Resolvers, Mssing SPF Records, Low/Medium/High CVE Vulnerabilities, Detected Browsers, Malware Events, P2P Activity, CMS Vulnerabilities, Cross Site Scripting, Defacement, Typosquats, Hacker Chatter Mentions...

However, this number can change. SecurityScorecard does, and must, continuously add new issue types and retire old issue types to match the evolving nature of cybersecurity threats. When a new issue type is added, it is initially presented for informational purposes only, with no impact on the score. This policy provides a fair time period to assess a new issue's significance and to inform users in advance before it impacts score.

Additionally, all issues are not equal. Issue types are weighted to account for differences in severity. When possible, the severity of an issue is defined by an industry-accepted standard, such as the NIST Common Vulnerability Scoring System v2. When an issue type does not have a severity ranking available, SecurityScorecard uses

recognized authorities and internal resources to determine severity, considering the opinions of multiple experts to correct for any bias.

It's important to note that once these weights are established for each issue type, they do not change and are the same for all companies. This allows for consistency and reliability in scoring all the way down to the issue level.

Factors: Definitions and Weighting

Now that we understand how things work at the issue level, let's move up a level.

All issue types are classified into 10 different categories or factors. Take a look at the below table to better understand how factors are set up:

Factor scores are calculated in a two-step process:

First, raw factor scores are calculated based on a weighted sum of the underlying issues in the factor. As explained above, the weights are based on issue severity.

Then, for each cohort - a statistically meaningful group of companies with comparable IP size—the raw factor scores are sorted in ranked order and mapped from 50 to 100. (We'll come back to why the scale is 50-100 in the Overall Score section.) These numeric scores are translated to letter grades "F" to "A" and are presented to the user on the platform, along with a list of the issues, organized by severity.

But just like issues, not all factors are equal. While the issue-level weights, discussed above, are severity-based and determined by external (NIST) and internal authorities, the factor-level weights are determined using machine learning.

While all factors have been found to be predictive of breach, SecurityScorecard uses cyber breach data and machine learning algorithms to quantify and rank which factors are more predictive of a cyber breach event. In this process, breach-likelihood ratios are

Network Security	Examples of network security hacks include exploiting vulnerabilities such as open access points, insecure or misconfigured SSL certificates, or database vulnerabilities and security holes that can stem from the lack of proper security measures.
DNS Health	The SecurityScorecard platform measures multiple DNS configuration settings, such as OpenResolver configurations as well as the presence of recommended configurations such as DNSSEC, SPF, DKIM, and DMARC.
Patching Cadence	How diligently a company is patching its operating systems, services, applications, software, and hardware in a timely manner.
Endpoint Security	Endpoint security refers to the protection involved regarding an organization's laptops, desktops, mobile devices, and all employee devices that access that company's network.
IP Reputation	The SecurityScorecard sinkhole system ingests millions of malware signals from commandeered Command and Control (C2) infrastructures from all over the world. The incoming infected IP addresses are then processed and attributed to corporate enterprises through our IP attribution algorithm. The quantity and duration of malware infections are used as the determining factor for these calculations, providing a data point for the overall assessment of an organization's IP Reputation, along with other assessment techniques.
Application Security	Examples of vulnerabilities detected include Cross-site Scripting (XSS) or an SQL injection attacks.
Cubit Score	The Cubit Score factor is SecurityScorecard's proprietary threat indicator that measures a collection of critical security and configuration issues related to exposed administrative portals.
Hacker Chatter	The SecurityScorecard Hacker Chatter factor continuously collects communications from multiple streams of underground chatter, including hard-to-access or private hacker forums. Organizations and IPs that are discussed or targeted are identified.
Information Leak	SecurityScorecard identifies all sensitive information that is exposed as part of a data breach or leak, keylogger dumps, pastebin dumps, database dumps, and via other information repositories. SecurityScorecard maps the information back to the companies who own the data or associated email accounts that are connected to the leaked information, assessing the likelihood that an organization will succumb to a security incident due to the leaked information.
Social Engineering	SecurityScorecard identifies a variety of factors related to social engineering, such as employees using their corporate account information for services, for example, social networks, service accounts, personal finance accounts, and marketing lists that can be exploited. In addition, employee dissatisfaction is monitored through publicly available data.

determined for each factor, by calculating the ratio of the conditional probability of a breach given a poor factor score (C, D, or F) to the conditional probability of a breach given a good factor score (A or B). The greater the likelihood ratio, the more predictive is that factor of a cyber breach.

Factors which are more predictive of breach are correspondingly assigned a higher weight. Just like with issues (and for the same

reason), once these weights are defined, they do not change and are the same for all companies. Factor weight may change when SecurityScorecard periodically re-evaluates factor weights based on updated cyber breach data and changes in the underlying issue types within a factor.

To ensure statistical significance, the machine learning process described above is performed at the aggregate level, across all industries and company sizes.

Overall Score

All the weighted factor scores described above are rolled into the total score which falls on a scale of 50 to 100. As you may have gathered, we pride ourselves on making reliable conclusions and not overstating insights on cybersecurity risk. Even companies with great cybersecurity programs can still get hacked, so no one gets a 100 on this numerical scale.

It should however be noted that scoring, as described throughout this whitepaper, is a data-driven process that ensures that lower scores are always more predictive of breach than higher scores. Put simply, an “F” company has a higher likelihood of getting breached than an “A” company- which is why action should always be taken if the overall score of a company is low.

Additionally, because overall scores are the results of continuous monitoring, they are representative of a company’s health over time. When there is a change in issue count (i.e. when a potential risk shows up) this issue is displayed in the platform on a daily basis but this issue does not make its full impact on the user’s score until two weeks later. This allows the user to get timely notification of a potential problem and encourages them to remediate prior to making a dramatic impact on the score.

SecurityScorecard focuses on providing information about trends in a company’s behavior. Good security health is the result of developing and maintaining the right behaviors over time and is not an absolute measure.

Lastly, as we mentioned earlier we translate these numerical values into a grade school rubric. Details without an effective communication mechanism result in inaction- one of the greatest enemies of proactive cybersecurity program. Using letter grades facilitate easy communication of scores.

Onward to Risk Mitigation

With your new understanding of how cohorts are defined, confidence in the quality of our data, and a better comprehension of our scoring methodology, we hope you now have a better understanding of where grades come from. More importantly, we hope that you'll be able to now use this information to take a proactive approach to reducing risk.