

POLICY REVIEW BOARD

A small orange triangle is positioned to the left of the title text.

How Bitsight Calculates Security Ratings

Bitsight was founded with the goal of increasing transparency about cybersecurity, enabling dynamic, informed interactions between global market participants and incentivizing a more trustworthy and secure global ecosystem.

We are committed to creating trustworthy, data-driven, and dynamic measurements of organizational cybersecurity performance and being transparent about our processes and methodology.

In this article, we'll describe how Bitsight Security Ratings are calculated, and why.

Security Ratings: Objectives

As the framework for creating the methodology behind our ratings, Bitsight uses the US Chamber of Commerce's [Principles for Fair and Accurate Security Ratings](#), which we helped develop.

US Chamber of Commerce Principles:

- **Transparency:** Rating companies shall provide sufficient transparency into the methodologies and types of data used to determine their ratings, including information on data origination as requested and when feasible, for customers and rated organizations to understand how ratings are derived. Any rated organization shall be allowed access to their individual rating and the data that impacts a change in their rating.
- **Dispute, Correction and Appeal:** Rated organizations shall have the right to challenge their rating and provide corrected or clarifying data. Rating companies should have an appeal and dispute resolution process. Disputed ratings should be notated as such until resolved.
- **Accuracy and Validation:** Ratings should be empirical, data-driven, or notated as expert opinion. Rating companies should provide validation of their rating methodologies and historical performance of their models. Ratings shall promptly reflect the inclusion of corrected information upon validation.
- **Model Governance:** Prior to making changes to their methodologies and/or data sets, rating companies shall provide reasonable notice to their customers and clearly communicate how announced changes may impact existing ratings.
- **Independence:** Commercial agreements, or the lack thereof, with rating companies shall not have direct impact on an organization's rating; any rated organization will be able to see and challenge their rating irrespective of whether they are a customer of the rating company.
- **Confidentiality:** Information disclosed by a rated organization during the course of a challenged rating or dispute shall be appropriately protected. Rating companies should not publicize an individual organization's rating. Rating companies shall not provide third parties with sensitive or confidential information on rated organizations that could lead directly to system compromise.

Furthermore, Bitsight uses these additional guidelines when considering how to build our ratings model and governance practices:

- **Comparability.** Ratings must allow meaningful comparisons of security performance between organizations — even if they are in different industries or locations, or if they differ greatly in size. As we shall see, this has important consequences for how ratings are calculated and normalized. The ratings should also be comparable over time. That is, a rating of 500 last year should mean roughly the same thing as a rating of 500 today. This makes it possible to observe trends and to track performance over time.
- **Ubiquity.** Ratings should be readily available for large numbers of organizations, in all industries, and across the world. This enables comparison against industry and global benchmarks.
- **Empiricism.** Ratings should be based on objective, verifiable data, rather than opinion or subjective judgements. They should be correlated with real-world outcomes.
- **Stability.** Significant shifts in security posture take time, and so Security Ratings should be relatively stable (free from spurious fluctuations).

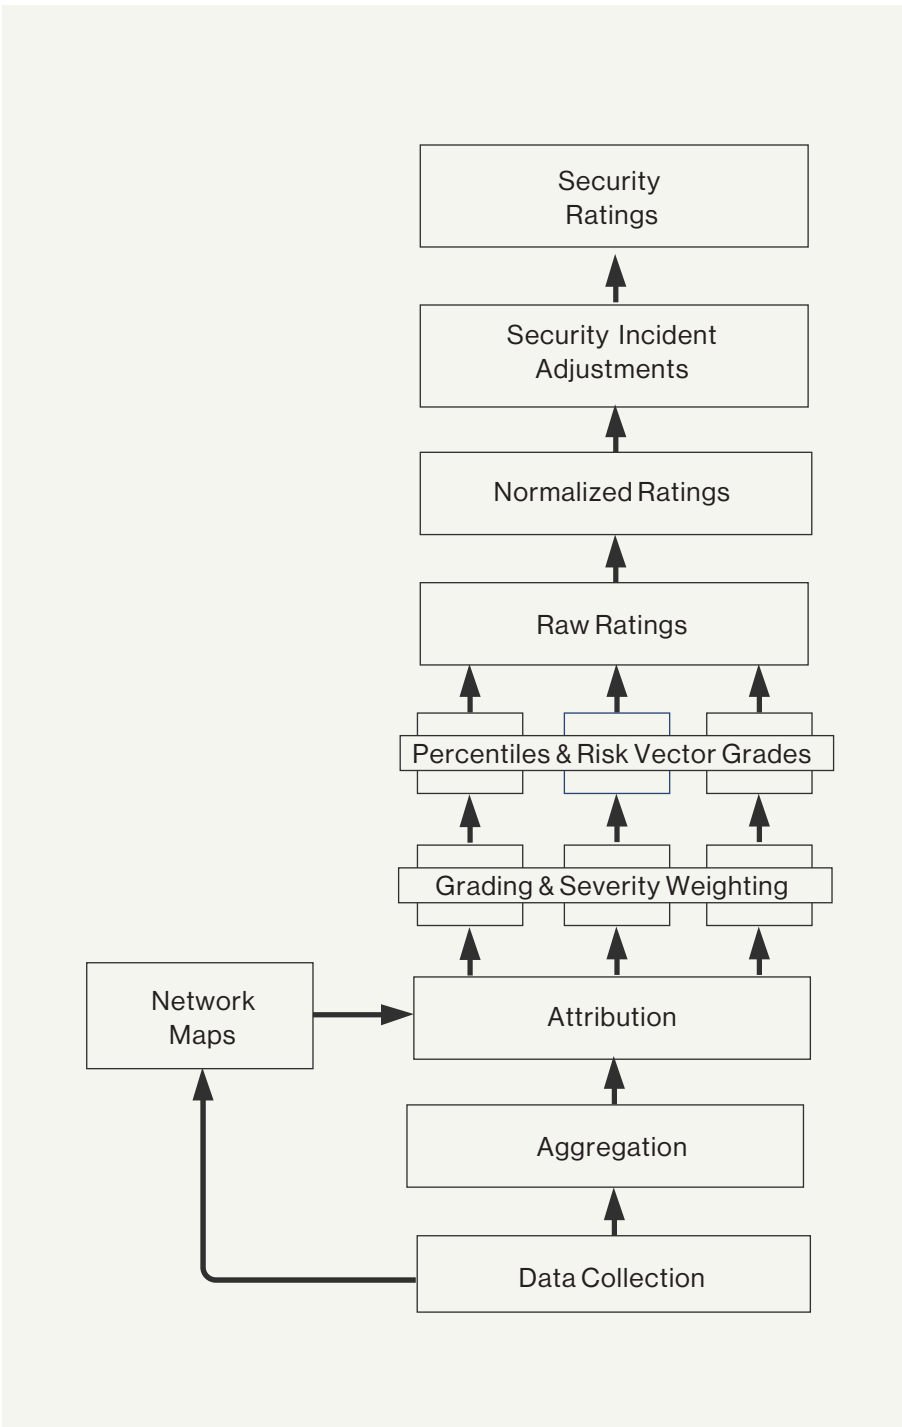
Balancing these principles is challenging, but over the past decade we have continuously refined our methodology to adhere as closely as possible to them. In the following sections, we'll walk through the ratings algorithm and explain it in terms of the above principles.

Comparison With Other Rating Systems

Bitsight defined the cybersecurity ratings industry, but our approach drew inspiration from many successful ratings systems in other domains:

Domain	Negative outcome	Observable risk factors
Consumer credit	Loan default	<ul style="list-style-type: none">• Missing a payment• High credit utilization• Previous default
Restaurant food safety grade	Foodborne illness	<ul style="list-style-type: none">• Poor sanitation• Not following best practices for food handling
Auto insurance	Accident	<ul style="list-style-type: none">• Speeding ticket• Previous accidents
Property insurance	Property damage	<ul style="list-style-type: none">• Missing smoke detectors• Claim history
Cybersecurity	Security incident / breach	<ul style="list-style-type: none">• Poor security hygiene (diligence)• Compromised systems• Risky user behavior• Previous security incidents

The Ratings Process



Data Collection

Security ratings are built on data from over 100 different sources. We collect much of the data ourselves, and we also work with numerous best-in-class data partners (many exclusive) who specialize in various types of telemetry. To date, we have collected petabytes of security relevant data and are adding billions of new observations every day.

The quality of the data is paramount, and so we have invested heavily in curating and refining all of our raw data. Real-world data at Internet scale is noisy and often challenging to interpret. Over the past decade, we have developed techniques and processes to separate signals from noise. We use a combination of human and machine intelligence (including a sophisticated rules engine) to screen out false positives and to ensure that the data we process is accurate.

While all of our data is collected externally, from the Internet (vs. internal networks), that's not to say that our data sources are all public. Much of what we observe relies on sophisticated and proprietary techniques and infrastructure, and these differentiate us from others in this space.

However, Bitsight does not conduct penetration testing or any other intrusive activity. This external perspective enables us to rate hundreds of thousands of organizations worldwide, and also allows us to maintain independence and objectivity.

What we learn by listening

We have an extensive network of sensors deployed at key locations across the Internet. With these, we can see:

- Communications from compromised systems
- DNS queries and responses
- Malicious traffic; e.g. DDOS attacks
- Attempts at brute force attacks
- File sharing
- Endpoint device identifiers
- Traffic from IOT devices
- BGP announcements

What we learn by actively looking

We use non-intrusive probes and queries to observe:

- Open ports
- Server software, configuration and versions
- Known vulnerabilities (CVEs)
- DNS records, including SPF and DKIM
- Web applications

Network Mapping

The heart of the security ratings platform is mapping out the assets that belong to each organization's network. Primarily, these comprise IP addresses (both IPv4 and IPv6) and domain names that the organization owns exclusively. We use both public data (e.g. Regional Internet Registry records and Domain Name system entries) and proprietary techniques to identify these assets. Here, too, we use a combination of human and machine intelligence to make the best possible decisions. Network maps are dynamic, and constantly change as assets are bought, sold, or moved (especially as cloud computing becomes more widespread), so our processes also constantly monitor and update the network maps.

With the network maps in hand, we attribute each day's new observations to the relevant organizations, based on the IP address or hostname where the observation was made.


Risk Vectors, Grading, and Weighting

Each observation has potential implications for an organization's security posture. To assess this, observations are first mapped onto a set of risk vectors, each of which measures a particular area of security performance. (A single observation may result in findings in multiple risk vectors.) Within the risk vector, the finding is then assigned a grade (in the case of diligence) or a severity weight.

In deciding how to evaluate a finding, we rely on:

- Empirical studies of the correlation of outcomes with the issue in question. E.g. Do organizations that use outdated SSL protocols experience breaches at a higher rate?
- Recommendations from authorities and standards bodies, e.g. NIST
- Databases of known security vulnerabilities, e.g. the National Vulnerability Database
- Severity and risk level of security issues associated with the finding
- Industry (best) practices and recommendations from security practitioners.

Currently, 23 risk vectors are included in the security rating. Grouped into categories of security controls, they are as follows:




Compromised Systems (27%)

Compromised Systems are devices or machines in an organization's network that show symptoms of malicious or unwanted software. This often reflects a serious gap in security controls, and so the Compromised Systems category is weighted heavily in the overall security rating.

- **Botnet infections:** devices on a company's network were observed participating in botnets as either bots or Command and Control servers. Botnets can be used to exfiltrate corporate secrets and sensitive customer information, repurpose company resources for illegal activities, and serve as conduits for other infections. Botnet detections are detected by capturing traffic from malicious software, using techniques such as sinkholing.

- **Potentially exploited systems:** Devices observed to be running potentially malicious or unwanted software; e.g. greyware or adware. These events are often indicative of other infections, and, like botnet infections, reflect insufficient device controls.
- **Unsolicited communications:** Systems observed to be scanning other hosts in patterns that are typical of malware seeking new hosts to infect.
- **Spam propagation:** Systems that have been used to propagate spam email (which is a common cybercriminal use for compromised machines). Legitimate email senders are excluded, however.
- **Malware servers:** Servers that are hosting malicious software.



User Behavior (2.5%)

This category measures how often employees at an organization are observed engaging in potentially risky behaviors.

- **File Sharing:** Exchange of media over peer-to-peer networks (e.g. BitTorrent). Since these files come from untrusted sources, they pose a high risk of malware infections.



Diligence (70.5%)

This category measures how effective an organization is in following security best practices and proactively defending against threats. We grade based on whether controls are implemented, and how effectively they are implemented.

Network Services - Protections against network attacks, impersonation, or eavesdropping.

- **TLS/SSL Certificates:** TLS/SSL certificates are used to encrypt traffic over the Internet. Bitsight analyzes certificates and provides information about their effectiveness; e.g. whether they are signed using a secure algorithm.
- **TLS/SSL Configurations:** Whether a company's servers have correctly configured security protocol libraries, and support strong encryption standards when making encrypted connections to other machines.
- **Open Ports:** Which port numbers and services are exposed to the Internet. Certain ports must be open to support normal business functions; however, unnecessary open ports provide ways for attackers to access a company's network.
- **DNSSEC:** It's a protocol that uses public key encryption to authenticate DNS servers. BitSight verifies whether a company is using DNSSEC and if it is configured effectively.

Software Assets - How well the organization follows best practices in managing its software, keeping it updated, and patching against known vulnerabilities.

- **Server Software:** The types and versions of server software that the organization exposes to the internet. Unsupported or outdated software often suffers from known, exploitable vulnerabilities.
- **Desktop Software:** Whether browser and operating system versions are kept up to date for laptops, servers, and other non-tablet, non-phone computers in a company's network which access the internet.
- **Mobile Software:** Similar to the above, except for mobile devices.
- **Patching Cadence:** How many systems within an organization's network are affected by critical vulnerabilities, and quickly the organization patches them (vulnerabilities are publicly disclosed holes or bugs in software that can be used by attackers to gain unauthorized access to systems and data).
- **Insecure Systems:** Devices within the organization's network observed to be unintentionally communicating with a third party (e.g. IoT devices reaching out to expired domains).

Application Security - The company's best practice implementation and risk mitigation as it relates to securing company applications.

- **Web Application Headers:** This risk vector analyzes security-related fields in the header section of HTTP request and response messages. If configured correctly, these fields can help provide protection against malicious behavior, such as man-in-the-middle and crosssite scripting attacks.
- **Mobile Application Security:** If an organization publishes mobile applications on the Apple App Store or Google Play, we evaluate the security of those applications. This risk vector will be included in the security rating in 2021.

Email Security - Controls to protect against email forgery. Email-based attacks such as phishing are often one of the most effective ways for attackers to gain access to an organization's assets.

- **SPF records:** Properly configured SPF records help ensure that only authorized hosts can send email on behalf of a company by providing receiving mail servers the information they need to reject mail sent by unauthorized hosts. Bitsight verifies that a company has SPF records on all domains that are sending or have attempted to send email, and that they are configured in a way that helps prevent email spoofing.
- **DKIM records:** Properly configured DKIM records can help ensure that unauthorized parties can't send email that appears to originate from the organization's domains. Bitsight verifies that a company uses DKIM and has configured it in a way that prevents email spoofing.

Information Exposure

We collect information on data breaches and other security incidents from a large number of verifiable sources; e.g. reputable news organizations and regulatory reporting (obtained via Freedom of Information Act requests or local analogs). Sufficiently severe incidents are factored into the overall security rating, as described in the “Overall security rating” section.

- **Breaches:** Publicly disclosed events of unauthorized access, often involving data loss or theft. These events are graded based on several factors, including the number of lost or exposed data records.
- **General Security Incidents:** a diverse range of events related to the undesirable access of a company’s data (which are considered more severe than Other Disclosures, below). Some categories of General Security Incidents are Ratings-impacting, while others are informational only and do not impact the rating.
- **Other Disclosures:** Other Disclosures are considered the least severe group of events within Public Disclosures and are generally minimal in their impact to business continuity were they to occur. All categories of Other Disclosures are informational only and do not impact the rating.
- **Exposed Credentials:** indicates if employees of a company had their information disclosed as a result of a publicly disclosed data breach. Exposed Credentials is an informational risk vector and does not affect a company’s Security Rating (many websites do not validate email addresses, which makes it difficult to establish that certain exposed records are in fact associated with a company’s employees).

The Time Component: Performance Vs. Instantaneous Risk Exposure

Security ratings are computed one day at a time; there is a new rating for each organization, for each day. However, findings typically affect the rating for longer than a single day.

Why is that? Consider a couple examples from other rating domains. An accident affects auto insurance premiums for several years. A loan default remains on a consumer credit report for seven years. The reason is that, statistically speaking, past negative events can be predictive of current risk. An at-fault accident is evidence that a driver engages in risky behavior, and that behavior is unlikely to change overnight. The older the event, however, the less predictive it is; an accident 30 years ago is not as worrisome as one last week.

The same is true of cybersecurity risk. Our data indicate that a negative event, such as a botnet infection, is indicative of potential deficiencies in an organization's security performance, even several months after it occurred. This is likely because it takes time to make significant improvements to an organization's security program (though the timescale is certainly shorter than years, as in the credit rating example).

For these reasons, compromised systems (malware) events and security incidents (breaches) have an impact on the rating which is greatest on the date they occur, and then gradually decays away as the events age.

In contrast, diligence records (e.g. open ports or SPF records) are measurements of the current state of an organization's systems. In most cases, if we can reliably confirm that the state has changed (e.g. the open port was closed), the rating reflects that immediately. (An analogy from consumer credit ratings: these typically incorporate the current ratio of credit utilization.) Otherwise, the record continues to affect the rating for 60 days. This duration was chosen (again, based on analysis of our data) to balance ratings stability against responsiveness, and aligns with typical update cadences.

Performance Vs. Exposure

Bitsight Security Ratings measure security performance: an organization's effectiveness in preventing cybersecurity incidents. This differs from the notion of exposure, which might be defined as the organization's current level of risk. Exposure may change rapidly as assets are created or taken down, or as configurations change; in contrast, performance tends to change relatively slowly, and reflects organizational practices and programs. Over time, however, good performance tends to reduce exposure.

Size Adjustment

All else being equal, large organizations have more opportunities for things to go wrong. It wouldn't be fair (or accurate) to give a company with two employees and two botnet infections the same rating as a company with 100,000 employees and two botnet infections. The latter likely has better security hygiene. Thus, size adjustment is necessary for the ratings to be comparable between large and small organizations.

As an analogy, consider the gross domestic product (GDP) of countries. Comparing based on GDP alone favors countries with more people. On the other hand, comparing based on GDP *per capita* places large and small countries on an equal footing.

To further motivate the size correction, consider what happens when two companies A and B, with roughly equivalent security postures, merge to form company C. Presumably, the security posture of C remains approximately the same as A or B, at least at first. However, C has events from both of the A and B, and so if we simply used raw counts, it would have a worse security rating, which isn't correct.

In engineering our ratings algorithm, we examine the distribution of event frequency vs. organization size to ensure that ratings aren't unduly skewed by size. The details of size correction vary from risk vector to risk vector, but the methodology in all cases is data-driven.

Ranking and Percentiles

For each risk vector, we compute a raw score. In some cases, this is simply a weighted count of findings (e.g. botnet infections), including time decay. In others, it is a combination of features used to evaluate that risk vector (e.g. expired certificates for the SSL certificate risk vector).

After size adjustment, we have a raw score for each risk vector. To determine the risk vector's grade (A-F), we first convert the score to a percentile, by ranking all the organizations we rate (minus a few outliers), across all industries and locations.

Why use percentiles? First, it's difficult to quantify security performance in an absolute sense. Since we rate such a large number of organizations, however, we can say with confidence how a given organization is performing, compared to the rest of the population. An organization in the top 10% is likely a strong performer, and receives an A.

Percentiles also help ensure stability. There are many natural variations in the data we collect. For example, when a new malware family appears, we may see a spike in infections across large numbers of organizations. If we used raw counts instead of percentiles, security ratings would drop, on average. But in fact, security performance likely stays about the same; it's the external circumstances that are changing. Percentiles maintain a steady distribution of ratings despite variations in events or our visibility into them.

Overall Security Rating


The risk vector ratings are multiplied by risk vector weights and summed to compute a raw overall rating. To compute the security rating that we show in the product, we first normalize the raw rating. This produces the desired distribution of security ratings, on a scale of 250-900, with 250 being the lowest measure of security performance and 900 being the highest. A portion of the upper and lower edge of this range is currently reserved for future use. As of the publication of this document, the effective range is 300-820.

What is the desired distribution? First, it should reflect the fact that the security performance of most organizations is fairly good. Data breaches, fortunately, are still relatively uncommon. Therefore, most organizations' ratings should fall towards the top end of the scale.

Second, the ratings should be spread out across the scale as much as possible, to provide greater contrast between stronger and weaker performers.

Finally, the numerical rating shouldn't convey a false sense of precision. Small changes are unlikely to be statistically significant. For that reason, we round the rating down to the nearest ten-point boundary.

Our normalization process is updated daily to reflect shifts in the underlying distributions. The figure [below] shows the distribution of the transformed ratings. While the raw ratings are clustered towards the top end of the scale, the final ratings have the desired distribution, and are well spread across the scale.



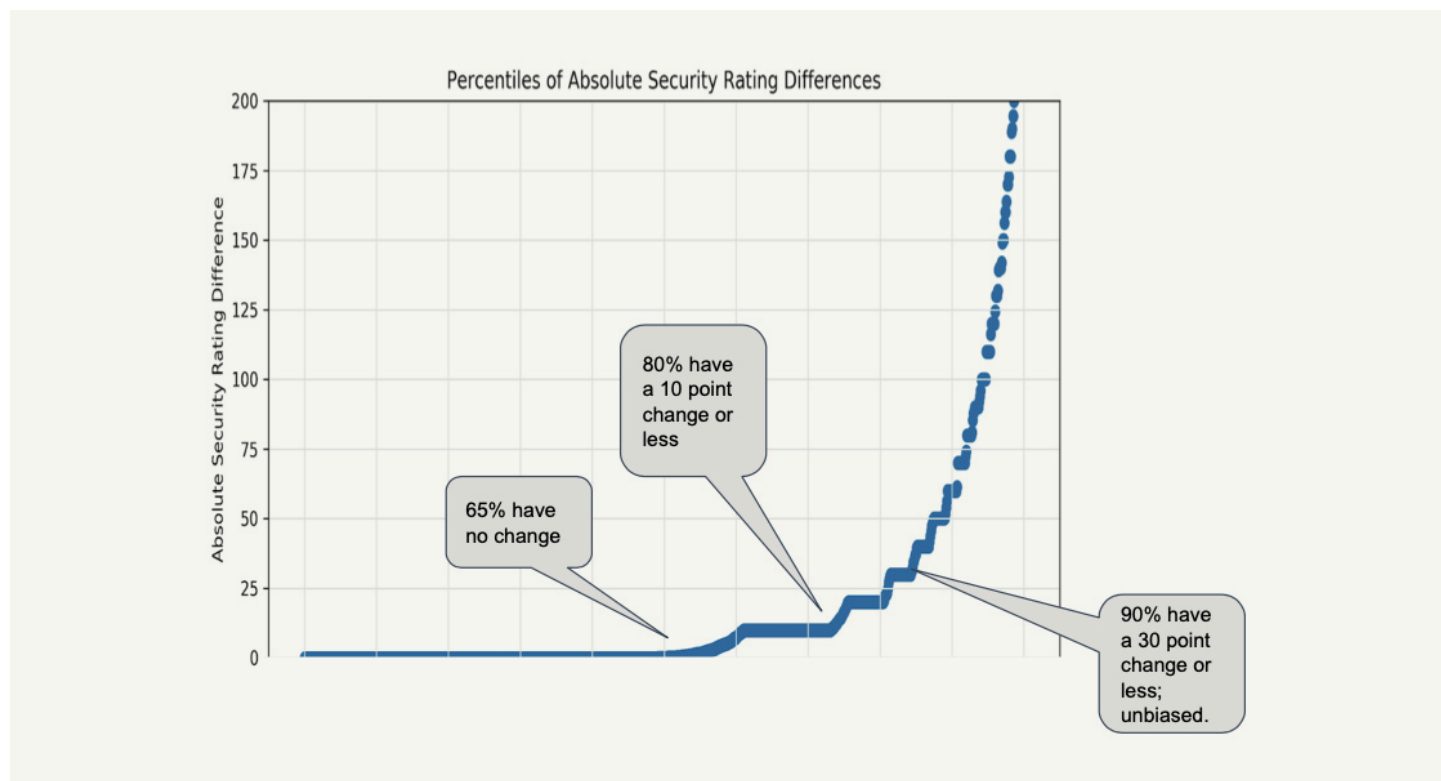
The final step is to adjust the rating for past security incidents (e.g. breaches), if there were any. Security incidents often provide strong evidence of gaps in an organization's security performance, and our research shows that the occurrence of one such incident is correlated with further incidents in the future. The impact of security incidents is corrected for company size, and also depends on the severity of the incident, and diminishes over time.


Dispute Resolution Process, Corrections, And Governance

Bitsight is committed to creating the highest quality and most accurate security ratings in the industry. We are also committed to allowing all rated organizations—not just customers—the opportunity to challenge the assets, findings, and interpretation of those findings used to determine a Bitsight Security Rating, and to provide corrected or clarifying data. As a signatory and contributing author, we are firmly committed to upholding the [Principles for Fair and Accurate Security Ratings](#).

Bitsight has a formal [dispute resolution process](#) that allows rated organizations to dispute findings. Bitsight seeks accurate and prompt remediation for any dispute. The dispute resolution process is governed by the Bitsight Policy Review Board (PRB), a committee created to govern the ratings algorithm and associated policies, and to ensure that they are aligned with our principles. As the highest level of ratings governance, the PRB also adjudicates appeals related to data accuracy and evaluation methodology. It is charged with providing a consistent, transparent, and systematic dispute resolution process that is available to all rated entities. For more information, please visit the [Policy Review Board](#) description.

In addition, to ensure that the ratings are accurate even in the presence of small errors in network maps, we have studied the effects of corrections to network maps (either additions or removals of assets). In most cases, they don't affect the rating at all, and when they do, the errors are unbiased (equally likely to increase or decrease the rating).



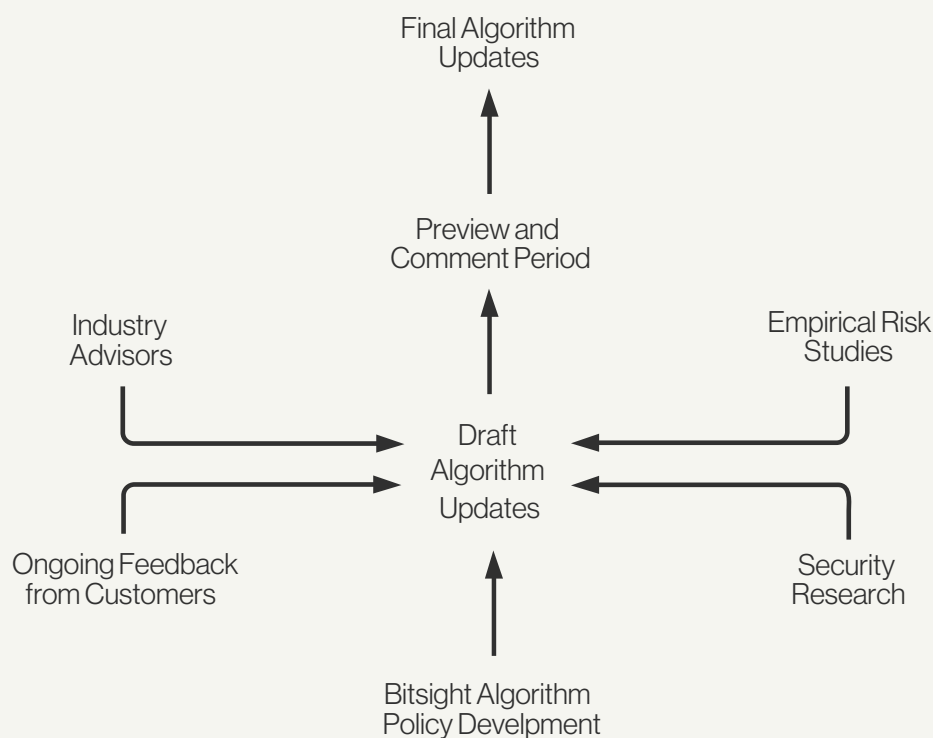


Confidentiality

We take the confidentiality of rated organizations' data very seriously. We follow responsible disclosure procedures for all security findings. Only the rated organization (or others with legal permission) has access to the full finding details (e.g. the IP address where it occurred). Additionally, we do not publicize organizations' ratings, and our terms of service also prohibit our customers from doing so.

Algorithm Updates

We periodically make improvements to the ratings algorithm. These updates often include new observation capabilities, enhancements to reflect the rapidly changing threat landscape, and adjustments to further increase accuracy and correlation with outcomes. These changes are all rigorously governed by our Policy Review Board to ensure that they adhere to our principles and policies. Additionally, we provide a preview of the changes to our customers (and what the likely impact on their rating will be), well before they affect the live ratings, and we invite comments and feedback on them.




Limitations - What Security Ratings Are Not

While Bitsight provides unparalleled visibility across hundreds of thousands of organizations, our data is limited to what we can observe externally. For example, we generally can't see how a company's network is configured internally, or what compensating controls may be in place, so those are not part of our security ratings. Furthermore, Bitsight complements -- but does not replace -- traditional network monitoring, vulnerability scanners, or intrusion detection systems (IDS). Finally, to avoid conflicts of interest, we do not provide incident response or vulnerability remediation services.

Validation

We validate our ratings algorithm by examining how our ratings and risk vectors relate to realworld security outcomes. We have compiled a database of more than 16,000 data breaches and other security incidents. Using this, and our historical ratings data from 2015 onwards, we have consistently found that organizations with low security ratings are more than five times as likely to experience data breaches than those with high ratings. These correlations have been verified by independent third parties, including the Marsh McLennan Cyber Risk Analytics Center, AIR Worldwide, and IHS Markit.



Bitsight is a cyber risk management leader transforming how companies manage exposure, performance, and risk for themselves and their third parties. Companies rely on Bitsight to prioritize their cybersecurity investments, build greater trust within their ecosystem, and reduce their chances of financial loss. Built on over a decade of technological innovation, its integrated solutions deliver value across enterprise security performance, digital supply chains, cyber insurance, and data analysis.

BOSTON (HQ)

RALEIGH

NEW YORK

LISBON

SINGAPORE

BUENOS AIRES



BITSIGHT