2021

ULTIMATE GUIDE TO

Attack Surface Management

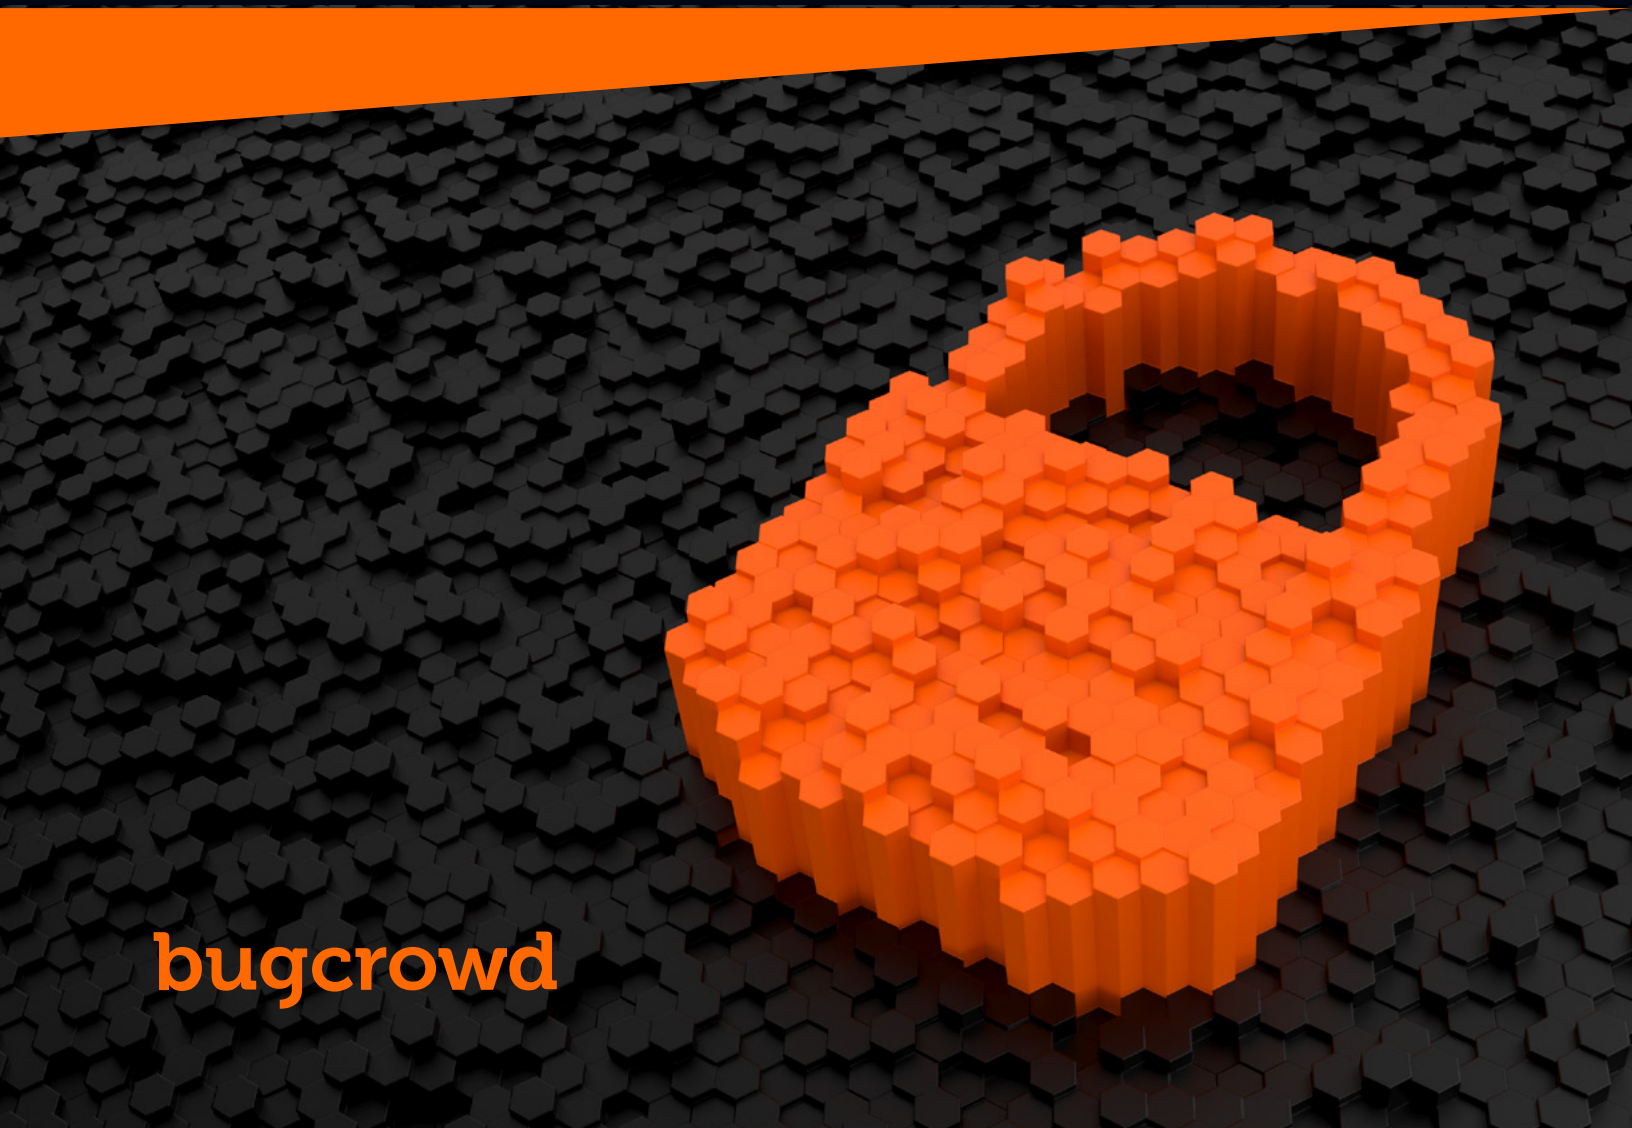

bugcrowd

TABLE OF CONTENTS

INTRODUCTION

Attack surface is evolving faster than ever before. If you think of a business like a human being– living, **breathing, thriving, changing**—then it might be easier to conceptualize the rate of change for the modern attack surface. Consider the software and systems required to manage your business. Now add a layer of complexity for the level of customization you have applied to each. Now think quick—who has access? Has that ever changed? If so, what does that change management process look like? Accountabilities are easy to track initially, until turnover, growth, mergers, digital transformation, and other perfectly normal business operations muddle it all up. These are some of the reasons why **two-thirds of organizations** say attack surface management is **more difficult today** than it was two years ago.

While attack surface drift typically occurs over many years of growth and business change, it can also happen suddenly, and unexpectedly. COVID-19-induced "shelter in place" orders have forced a **quick shift to fully remote work** (where possible). But accelerated timelines for introducing new online services have caused many organizations to shortcut standard security testing protocols. And while the arrangement may not last, **the impact of mismanaged IT** will.

This legacy plus the increase in attack surface highlights a big issue. There was **50x more online data** in 2020 than in 2016. On its own, this increase is not necessarily bad. As organizations mature, they naturally undertake normal growth activities like business transformation, attrition, hyper-growth, and mergers and acquisitions (M&A). These initiatives expand their web of internet-facing assets, but with limited resources and dispersed accountability, the ability to maintain oversight wanes. And in the shadows, **malicious attackers lurk.**

As security teams strive to stay ahead of these attackers, visibility into the attack surface is crucial. This brings us to a key point we'll keep coming back to throughout this guide–how can you secure what you don't know exists?

EXECUTIVE SUMMARY

500+

Average number of unprotected applications per organization

40%

Average amount of unknown attack surface per organization

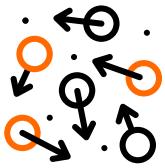

⅔

Percentage of organizations that say attack surface management is harder than it was 2 years ago

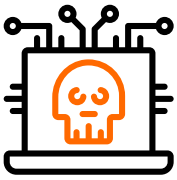

⅓

Percentage of successful attacks against unknown or unprioritized assets

EQUIFAX: WHAT HAPPENED?

When discussing attack surface management, it's hard not to bring up the **2017 Equifax breach**. In this iconic incident, a data breach involving the Apache Struts vulnerability exposed the personal details of **147 million Americans.** However, the story isn't as simple as the media portrayed it.

What you may not have known is that Equifax actually did know about the Apache Struts vulnerability before the now-infamous breach. Equifax relied heavily on automated vulnerability scanners, but failed to maintain a registry of the public-facing technology they owned. This means Equifax failed to find and patch the vulnerability in an unseen asset before the malicious exploit. The problem wasn't just awareness of *external risk*, it was **awareness of at-risk assets.**

This unfortunate story illustrates the importance of visibility into your environment. Unknown attack surface is like leaving your back door unlocked– it's only a matter of time before a bad actor steps through it.

THIS REPORT EXAMINES...

- Why organizations are shifting away from exclusively using attack surface scanners
- Common attack surface management mistakes and how to avoid them
- Top tips for building a business case for attack surface management and actioning on results
- Industry benchmarks comparing asset inventory
- How crowdsourced security combines software-based discovery solutions with the creativity of the world's best reconnaissance hackers to uncover more unknown attack surface

THE OUTSIDE-IN APPROACH
A LOOK AT ATTACK SURFACE MANAGEMENT

WHAT IS ATTACK SURFACE MANAGEMENT?

The term “attack surface management” isn’t new, but it’s often mixed up with the terms “asset management” or “asset discovery.” Unlike these terms, attack surface management isn’t just about monitoring known asset inventory, nor about finding shadow and legacy IT, though those are both important and benefit from being well understood by organizations . Attack surface management covers aspects of both, but most importantly, it’s deployed **within the context of *real* risk.**

The crux of attack surface management is found in this simple sentence:

Define it, prioritize it, and action it

...faster than your attackers.

OUTSIDE-IN VS. AN INSIDE-OUT APPROACH

Organizations typically defend their digital ecosystem using an inside-out approach, but that’s actually **the opposite of how hackers operate**. In other words, while organizations work to secure priority assets, attackers are more focused on whatever fell off the radar.

The outside-in approach is more important as the attack surface continues to expand across the board due to increased data collection, the proliferation of SaaS tools, and a move toward remote work. In 2020, **the pandemic accelerated these trends** considerably at a time when IT and security resources were already stretched.

As hackers began scouring the internet in order to uncover relevant, forgotten assets, this **opened Pandora’s box**, showing an aspect of security that had previously received little commercial attention beyond some scanning tools, but represented a **rich vein of vulnerabilities for attackers.**

> Unknown or unprioritized assets become **ticking time-bombs** when they fail to receive routine maintenance and vulnerability patching, creating opportunity for ill-intentioned hackers to strike. And these aren’t passive risks–Gartner predicts that around **30% of successful attacks this year will be against shadow IT.**

MERGERS AND ACQUISITIONS

Mergers and acquisitions (M&As) will come up a lot in this guide because M&As and attack surface management go hand-in-hand. Simply put, M&A events can wreak havoc on an organization's visibility into and understanding of their assets. M&As are **incredibly complex events**, making asset tracking all the more difficult. Finding the unknown or unprioritized assets belonging to each organization is one matter, but tracking ownership across several layers of change management, hundreds of manually-populated spreadsheets, and several public and private databases is an entirely different ballgame–one that attackers are playing in parallel.

There are also considerations in the due diligence process of a potential M&A. Security teams are often given just a few weeks to perform risk assessments, attempting to **determine risk under immense pressure.** Vulnerabilities in unknown assets can end up being problematic, **expensive "surprises"** that may come up post-sale. The formal announcement of an acquisition is like a gun shot at the starting line for a global network of attackers competing to identify forgotten assets to exploit.

The average amount of unknown attack surface in an organization is 40%. Think of it this way–this means that organizations are "missing" four out of every ten assets. M&A events certainly contribute to the high percentage of unknown assets as two companies work to uncover and combine their respective digital assets. However, they can also serve as a reminder, giving organizations an excellent reason for increased attention to attack surface management.

WHY AUTOMATED DISCOVERY TOOLS AND SCANNERS AREN'T ENOUGH

Let's go back to the Equifax example for a moment.

The investigation found two major gaps in Equifax's security posture:

1. **Equifax failed to maintain a registry of the public-facing technology they owned that matched those technologies known to contain the Apache Struts vulnerability.**

2. **Equifax failed to configure their scanners to search potentially vulnerable public-facing assets.**

Basically, Equifax was doing much in the way of vulnerability scanning within their own footprint, but who was looking outside of that? This brings us back to the question posed at the beginning of this guide–how could Equifax (or any organization) secure what they don't know exists?

ATTACK SURFACE SCANNER LIMITATIONS

To combat the risk in successful attacks against shadow IT, there has been an uptick in scanning solutions designed to automatically identify vulnerable, or otherwise forgotten assets. But is this enough for organizations to **combat dynamic, motivated attackers** looking to do the same?

Scanners do help cover more area than individual humans can do manually. But what does "cover" mean? Some are designed to scan for vulnerabilities strictly within the asset-set you define, or manage and monitor those assets as they evolve; Equifax was using such technology. Others are designed to perform reconnaissance for external-facing connected IT you may have forgotten. Whether rules-based, machine learning (ML), or artificial intelligence (AI) methodologies, all promise reduction in time and resource-drain otherwise required to perform these functions by hand.

However, there are some **inherent limitations of automated solutions** to determine whether the resulting tradeoff between impact and effort will help meet security objectives. **Not all scanners are created equal.**

FIVE MAJOR LIMITATIONS OF ATTACK SURFACE SCANNERS

1. LAG-TIME

Let's start with something that might sound a bit counterintuitive. The fundamental value proposition for most scanners is the ability to provide continual insight into your attack surface, saving time and resources in the process. It's certain you will save hourly effort, but less certain that you'll achieve rapid time-to-value. Unless the scanner in question utilizes continually updated, pre-indexed data, you may be forced to **wait up to a month for an initial scan to complete.** This renders most scanners useless for many critical use cases, including M&A.

2. KNOWN-KNOWNS

Scanners are designed to apply encoded logic or learning frameworks at scale–to cover more ground, with less overhead. For most solutions (other than some relying on Machine Learning), any activity identified as malicious or questionable is derived from what we call "known-knowns," or patterns that have already been identified as warranting further analysis. And while some of the better tools in-market were originally developed by very skilled members of the hacking community, **the ability to stay abreast of the most *recent* attack methods is still a challenge.** In fact, it often takes years before the latest techniques are validated, tested, and incorporated. "Getting a jump on attackers" is almost assuredly not in the cards for organizations relying on such solutions.

3. LACK OF BUSINESS CONTEXT AND INABILITY TO MAKE LOGICAL PIVOTS

The technology landscape for large organizations is often structured in exceedingly complex ways, and no two organizations look the same. Automated scanners are **highly susceptible to getting lost in a maze of interconnectivity,** unable to make sense of logical business structure and priorities. While training is an important requirement for any such technology, it's also time-consuming, and never a one-off. In addition, organizations are always evolving, abandoning behavioral trends as quickly as they are established. This often limits scope and attention to particular areas, or conversely, makes it difficult for scanners to develop a trusted baseline.

4. INABILITY TO SAFELY VERIFY AND PRIORITIZE

While many scanners are tuned to identify assets that may be vulnerable, it is next to impossible for them to verify the accuracy of those initial assessments without serious risk. Scanners often have no concept of scope, nor the implications of various tests across multiple scenarios, where proof of exploitation could cause significant business or security risk to production environments. As a result, **scanners are also highly limited in their ability to truly prioritize discovered assets.** A rollup of 3,000 newly discovered assets, pulled from the shadows, is only as useful as your ability to action them. Scanners can provide preliminary estimates of risk, but the **false positive rate is typically high.**

5. MOST ATTACKERS HAVE BUILT THE SAME, OR BETTER

The good news about automated scanners is that they're designed to rapidly uncover potentially connected assets faster than humans alone can achieve. The bad news–**you're not the only one using them.** Attackers use (and frequently develop themselves), tools that rival and often surpass the power of any commercially-available scanner in-market today. In fact, while you might have the resources to deploy one or two, hackers mapping your attack surface (by the thousands), often use 5-10 or more different scanning technologies, creating a serious disadvantage for defenders everywhere.

As you can see, scanning solutions were broadly intended to replace human resources, but it seems their greatest strength may actually be in supporting them.

HOW HUMAN INGENUITY UNCOVERS MORE ATTACK SURFACE

It's important to recognize the need for automation in attack surface management, in fact, it is crucial for tackling this problem at scale. However, **automation without human intuition is an incomplete solution.**

In the race against ever-evolving attack methodologies, you want to be firing on all cylinders. There are several examples that illustrate how the human element wins over pure software-based solutions.

1. SCOPE

Telling a software-based solution what to test is easy. Telling it what not to test is easy. Having it identify things that might belong to you, and then making a judgement call about whether to perform an exploit? That's where things start to get complicated. By using scanners to identify low-hanging fruit, **humans are left to focus on the edge cases**–identified assets that can be positively attributed to the organization through manual investigation.

2. TAILORING TACTICS BY ORGANIZATION

Many of the activities that cause drift in your internet-facing assets are things that happen seasonally. Think conferences, sales kickoffs, marketing campaigns, etc. For temporal activities like these, new assets must be spun up quickly, and decommissioned just as fast. You know when these things need to happen, but so do your attackers.

Your digital movements have a heartbeat—and it's not that tough to track, and eventually predict the ebb and flow of activity through openly available indicators such as LinkedIn updates, advertising spend, and email campaigns. Attackers are just waiting for you to miss something.

A common example of this is **subdomain takeovers.** These vulnerabilities happen when a subdomain is pointing to a deleted service, which enables a hacker to create a page on the formerly deleted service and point their page to that subdomain. This could lead to potential public relations fallout and cause reputational damage equal to any breach. Scanners aren't built to look for inconsistencies in the brand like this, where this would be easy for a human security researcher to catch.

THE HACKER PERSPECTIVE
SUBDOMAIN TAKEOVER VULNERABILITIES

MICHAEL "CODINGO" SKELTON
GLOBAL HEAD OF SECURITY OPERATIONS

"In my own hunting experience, I've come over a top-tier, publicly traded company which had a subdomain takeover that had been performed by an attacked and turned into a cannabis company (likely for Blackhat SEO efforts).

The site was set up for a team offsite that was then decommissioned, but the DNS record had remained active and then used by the attacker to reclaim that subdomain.

Why didn't scanners pick it up? Ultimately, because once a compromise has already occurred, most scanners will no longer view them as a vulnerability. Scanners aren't built to suss out brand inconsistencies, afterall."

3. CONTEXTUALIZING FOR LESS RISK

Sometimes it might feel like your scanner wasn't made for your environment—the noise and false-positives can be overwhelming. This is because **your scanner *wasn't* made for your environment.** It was made for, and by, everyone else's. It's shaped by the status quo, driven by known knowns, and errs on the side of caution when it comes to making assumptions or logical pivots.

Business logic errors are often great examples of items that fall outside the bounds of those parameters. Given their varied nature, they typically can't be automated. At least not yet.

4. AGILITY IN PIVOTING TOOLS AND TECHNIQUES

Human researchers can pivot their tools and approaches to take action sooner. When a Common Vulnerability and Exposure (CVE) is posted, it can take software weeks or months to test for it, but researchers begin reporting issues in hours because they can quickly refocus tools at the problems. Software companies approach the problem via a route through development, QA, and testing in production and release to ensure their software does no harm. In many cases, very widely known issues never make it into a scanner at all.

Notably, **a scanner can only track what you feed it, and it's going to miss anything you don't**–especially acquisitions and domains not already tightly coupled to the scope you've provided. A competent researcher understands this, and will work to tailor the scope as an engagement grows, fueling future discovery with the seeds of previously identified assets to ensure a greater depth of coverage.

5. DIGESTING AN EXPANSIVE VIEW OF INFORMATION

Scanners just aren't designed to excel in digesting news, social media, and other contextual information wrapped around current events pertaining to an organization. **Those bits of information are breadcrumbs to your business leading attackers toward wherever you're not looking.**

For example, in an M&A event, attackers connect the dots across three or four levels of activity. It's complex, but the payout can be well worth the effort as thousands of potentially vulnerable assets hang in limbo between the owners.

Another example is joining datasets like LinkedIn, GitHub, and Stackoverflow. The challenge here lies in attributing all of these personas to each other. This can easily be done with a unique name, but if they have a common name, you may have to resort to using site favicons, GitHub handles, or even cross-referencing the organization's Twitter. In this case, automation might be possible, but not very practical without inflating noise as attribution changes on a case-by-case basis.

CROWDSOURCED ATTACK SURFACE MANAGEMENT

Leveraging human ingenuity and creativity in your attack surface management strategy is all well and good, but most organizations don't have the resources to hire entire teams to cover this aspect of security. This is where crowdsourced attack surface management comes in.

First, let's cover a little bit of background. **Crowdsourced security** was born out of a need to connect the global security community to the global market, delivering the ability to discover vulnerabilities before an adversary does. You can learn more about crowdsourced security in the report, Priority One: A New Decade in Crowdsourced Cybersecurity.

It is powered by **the Crowd**, a collection of on-demand ethical hackers (aka security researchers) distributed across the world. The Crowd powering this type of security testing consists of researchers with diverse backgrounds and skill sets, allowing organizations to tap into expert testing at scale. If you're interested in learning more about the Crowd, check out Inside the Mind of a Hacker, the most extensive study of global hackers and the economics of security research.

Bugcrowd revolutionized attack surface management by leveraging the power and scalability of crowdsourced security for asset discovery, prioritization, and management. By approaching attack surface management with a crowdsourced solution, organizations **match the effort and scale of attackers with the ingenuity and impact of trusted attack-minded defenders** for the most organic assessment of real risk possible. It gives organizations the **defender's advantage.** We'll talk more about this and Bugcrowd's specific solutions later in the guide.

COMMON ATTACK SURFACE MANAGEMENT MISTAKES

We've already discussed a few mistakes organizations make when fitting attack surface management into their security strategy, for example, exclusively taking an "inside-out" approach. In a recent Research Insights Paper published by ESG, Attack Surface and Vulnerability Management Assessment, ESG evaluated survey data that indicated several winning and losing strategies when it comes to attack surface management. When evaluating vulnerability management priorities and practices, they organized data by security maturity, breaking it up into three groups: Leaders, Fast-followers, and Emerging Organizations. Here are a few takeaways from the report regarding common mistakes to avoid.

1. RELY ON ONLY ONE METHOD

We now know that solely relying on attack surface scanners is not enough. Scanners often find something and then do not refresh or clean up after themselves. This causes a lot of problems with autoscaling groups in a "living" environment.

Organizations must **diversify methods** to find hidden attack surfaces. ESG found that leading organizations are more likely to diversify their methods compared to less mature organizations. The top tools and services used by these leading organizations include:

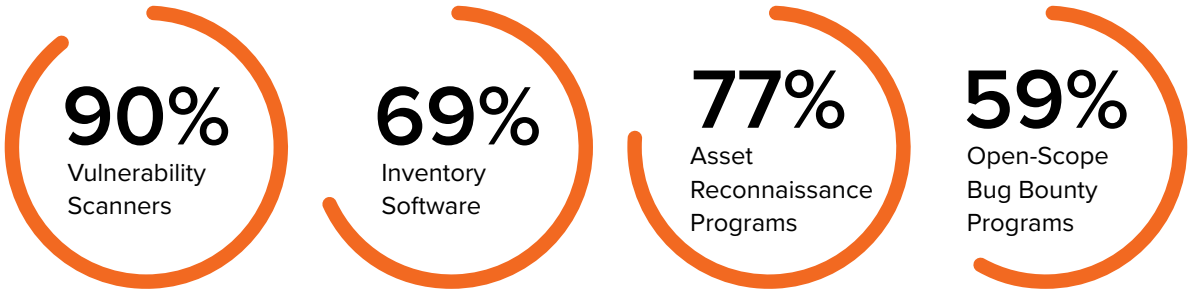

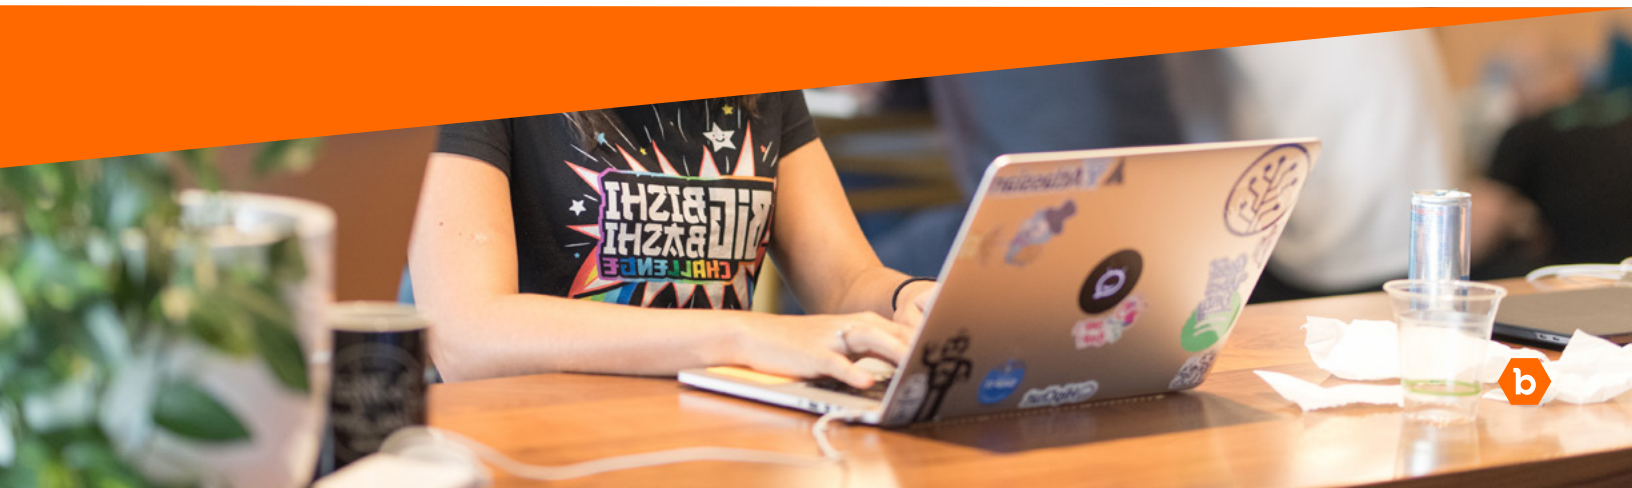

2. WORK IN SILOS

49% of organizations say that attack surface discovery and management is a shared responsibility between security and IT operations. One interesting note is this focus on sharing these responsibilities remains relatively equal across organizations of all maturity levels.

A commitment to a strong cybersecurity culture and focus must be driven from the top down. By building common goals, processes, and strong communication between the security and IT operations team, both groups can work together to create a winning attack surface management strategy.

3. INFREQUENT MONITORING

The average number of unprotected applications per organization is 500+, and all it takes is one exploit to cause big problems. With the high volume of unprotected assets, continuous monitoring is crucial to find these assets and help reduce risk

The report found that the frequency of attack surface discovery is a function of maturity. 72% of leaders conduct attack surface discovery on a continual basis, compared to 52% of fast-followers and 3% of emerging organizations.

4. LACK OF PRIORITIZATION IN FUTURE PLANNING

Attack surface management is no longer a “nice-to-have”–it is a necessity to building a good security posture. This is displayed in the commitment organizations across the board have to increase their attack surface management budget. **70% of organizations plan to increase spending on attack surface management in the next 24 months.** Some ways that leading organizations are prioritizing this include:

STATE OF THE ATTACK SURFACE

Every organization's digital footprint is unique, but it's interesting to see how different industries stack up against each other when it comes to asset inventory. Bugcrowd recently partnered with BitDiscovery to analyze the number of internet-accessible assets of the top organizations in several key industries, to help security stakeholders in the same cohort make better decisions about their own attack surface exposure and risk.

First, let's define an asset. We looked at domain names, subdomains, and IP addresses, including but not limited to web servers, name servers, mail servers, DNS servers, VPN getaways, blogs, IoT devices, and network printers.

We have amassed a database of nearly **4.5 billion internet-connected assets**, which combines meta-data obtained from hundreds of third-parties to create the largest and most complete dataset of its kind. This section of the guide provides the median number of total assets, expired TLS certificates, domain names, and cloud assets over five major industries–insurance, retail, medical devices, personal finance, and investment banking. The median number is taken from the top 10-15 organizations in each industry.

This data provides a benchmark for comparison as you consider your organization's own attack surface.

TOTAL ASSETS

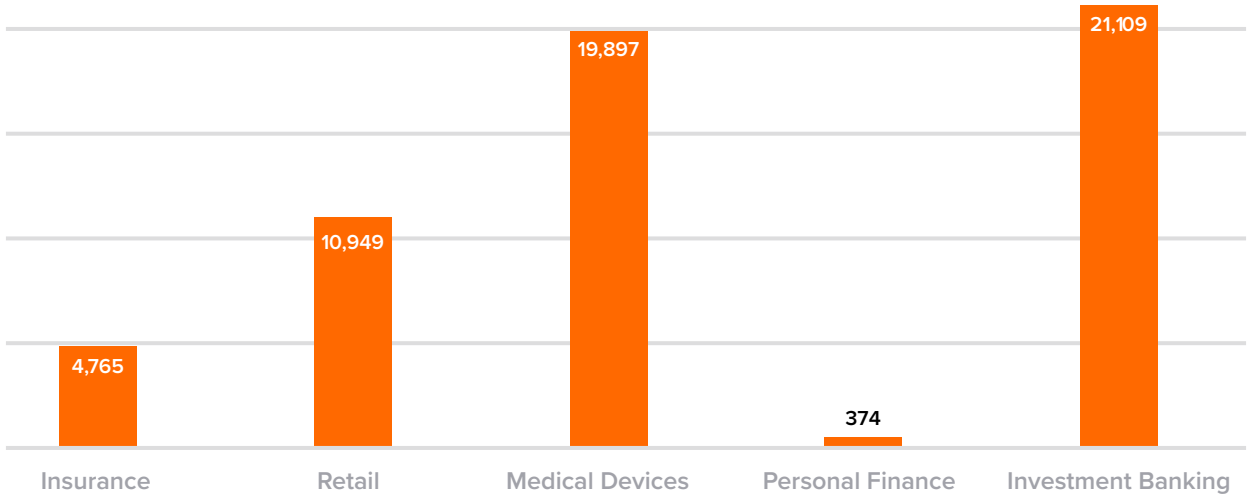

From the data, you can see that investment banking and medical devices have an especially high median number of assets, where personal finance is the lowest. Readers should note that the total number of assets for each organization is affected by many factors, including company brands, subsidiaries, products, partnerships, international presence, IT infrastructure strategy, marketing programs, sales strategy, cyber-attacks, and more.

While these numbers don't tell us anything about the relative size of the company's internet-facing business compared to non-internet facing parts, it does clearly indicate the real size of the company's external attack surface. **The more internet-connected assets, the more opportunities for a motivated adversary to gain entry.** Regardless of how many internet assets your organization has, losing track of even just one could cause serious disruption.

EXPIRED TLS CERTIFICATES

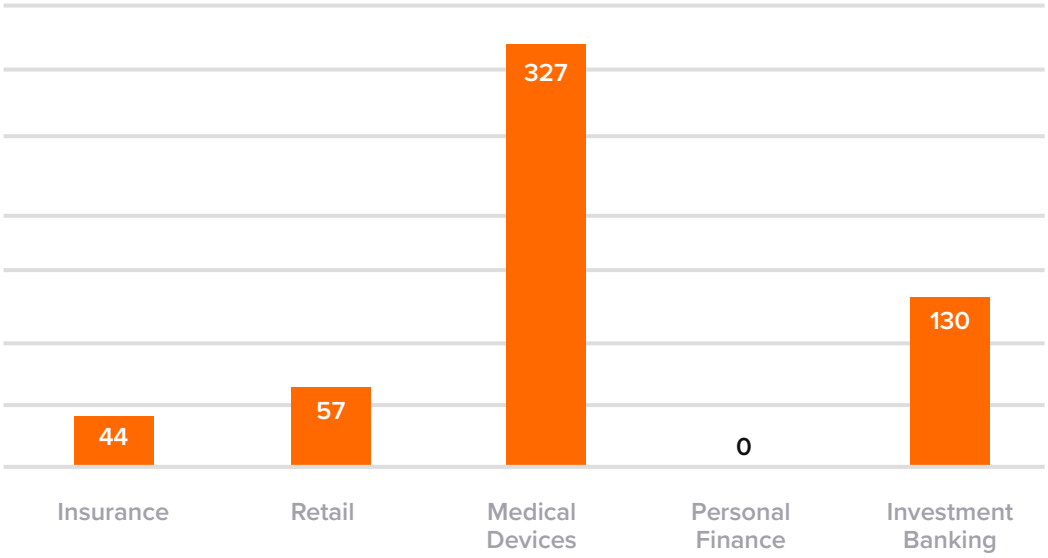

TLS certificates are a type of digital certificate that certifies that a user has verified that it belongs to the owners of the domain name which is the subject of the certificate. As anyone would expect, the more Internet-connected assets that support TLS, the more likely a higher number of expired TLS certificates. So while the number of expired TLS certificates informs on overall IT management hygiene, it's important to view these numbers in context of the oversize of the inventory and the number of assets supporting TLS.

In our industry comparison, medical devices had by far the highest median amount of expired TLS certificates. With the increased availability of free TLS certificates and browser vendors giving preference to TLS support, a larger number of assets initially deployed with TLS has similarly gone up, but many of these assets are often abandoned as they may be testing, development, or staging systems with a limited useful lifespan.

DOMAIN NAMES

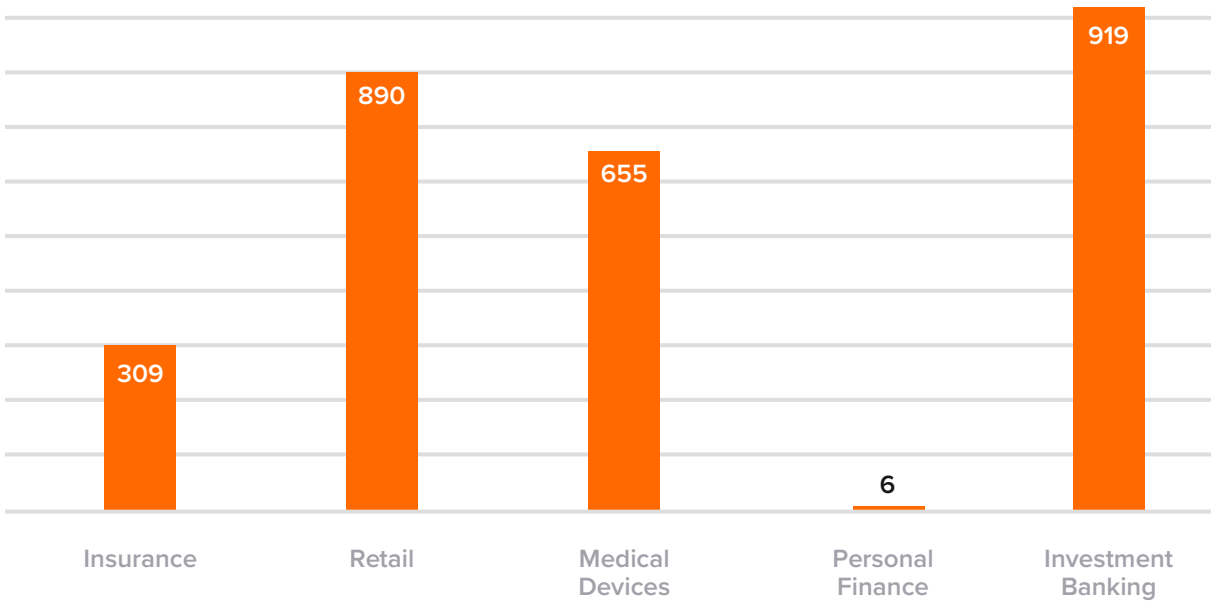

Companies register domain names, often hundreds and even thousands of them, for a variety of reasons. Domain names support brands, subsidiaries, partnerships, company announcements, marketing programs, protect against phishing attacks and typo-squatting, international operations, trademark protection, and even for investment in intellectual property.

It's common for domain name registration information to be hidden behind proxy services for privacy protection and managed by brand protection services. What's also common is for employees to use their work email addresses to register personal domain names, or use even their work or personal email addresses to register important company domain names outside of an approved IT process. **These practices can put the organization at risk**, when access to or memory of these assets are lost when events end or employee's move on. Additionally, domain name registrations can provide leading indications of product releases, impending M&A activity, and legal settlements, all things that motivated adversaries search for when seeking to undermine, expose, or endanger an organization's IP.

To reduce risk of forgotten domains and associated assets, it's important to **develop a strong culture of cross-functional IT accountability and communication.**

CLOUD ASSETS

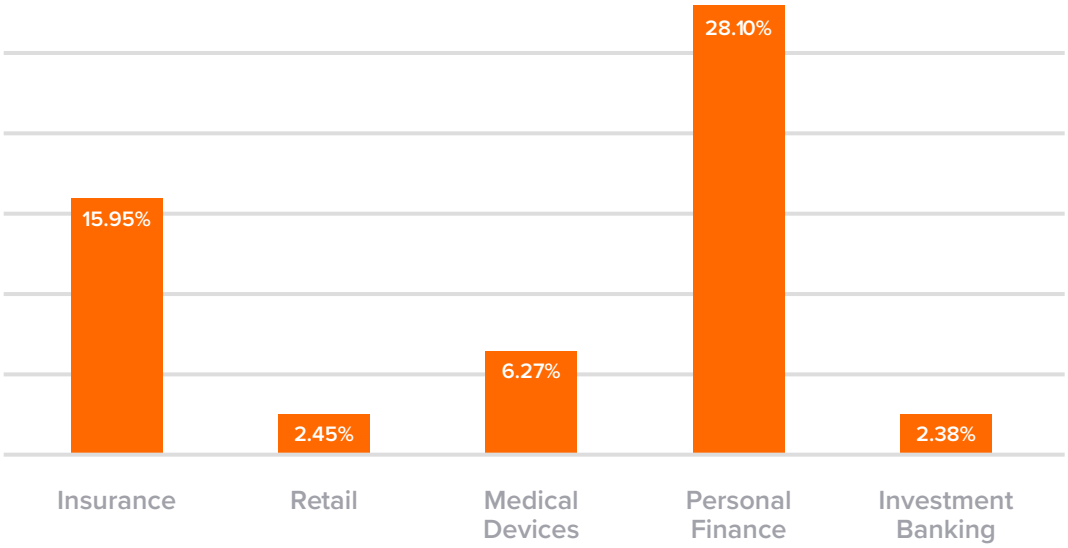

"Cloud Assets" refers to the percentage of internet-accessible and cloud-hosted assets. Cloud providers include Amazon Web Services, Microsoft Azure, Google App Engine, and others. While it's reasonable to assume that a large number of companies have "adopted" or "moved to" the cloud, this phrasing can be confusing or misleading. When a company leverages a third-party cloud service it may mean they've only done so in a tiny area of the overall business and not necessarily broadly across the entire enterprise or IT infrastructure. As we see above, the percentage of Internet-accessible assets hosted within well-known cloud environments remains low in most industries, with some exceptions, despite all the potential benefits of doing so. In terms of overall adoption, cloud vendors have a long way to go if they intend to host a significant number of the overall Internet-accessible assets. When hosting in the cloud, it's important to realize that the organization may have also outsourced security of the assets as well. It's important to understand and monitor the security posture of the cloud providers companies are using.

ACTIONING RESULTS

We've talked a lot about the importance of attack surface management, but what happens once you've already taken action to reduce unknown assets in your attack surface? Once you've built your attack surface map and created a prioritized list of assets, **it's time to take action.** Here are four ways to translate your asset risk outcomes to business value.

1. ELIMINATE ASSETS

The next step for most organizations is to **eliminate assets deemed irrelevant to current operations**, including those you assumed were already gone. Many teams are surprised to see assets they believed were decommissioned years ago show up, but as asset inventory is often tracked manually via spreadsheet, it's not uncommon for accounting errors or ownership shuffles to disrupt offboarding flows leaving now "invisible" assets connected to the network.

2. REMEDIATE HIGH INDICATORS OF RISK

Once you've eliminated the assets that aren't important, take a look at those that contain **clear indicators of risk.** Common indicators of risk include CVEs, configuration, unsecure auth, possible user enumeration, reflected XSS, subdomain configuration, invalid certs, SSL score, and login over HTTP. Some of these are easily remedied, though can cause extreme damage if left unattended.

Let's use CVEs and configuration as an example. Some common configuration problems include many open ports, ports like 22 or 23 that are not encrypted, leaked passwords, or even standard passwords on systems. All of these issues are easily resolved, but if chained together by a malicious attacker that finds them first, can result in something more serious.

3. DIG DEEPER AND KEEP IT SECURE

Depending on the asset's business criticality and associated risk level, you may want to consider adding it to an active testing program like Bugcrowd's Bug Bounty or Next Gen Pen Test solution. These programs incentivize highly skilled security researchers to "dig deep" within a given target, incentivized by vulnerability volume and severity. While offered on-demand or continuous, the latter can **reduce risk without added operational overhead**, making these programs ideal for things like web apps and APIs which might undergo frequent code changes.

4. BUILD A REPEATABLE FRAMEWORK

It's important to perform asset risk assessments after significant business changes, or roughly 9 months-1 year on average. If your organization is experiencing significant changes, initiating deeper assessments like these, every 3-6 months may be advised. Regardless of frequency, it's always important to **work with your internal security and development teams to remediate issues quickly**, and create a plan for securing and tracking newly discovered assets.

BUILDING THE BUSINESS CASE

While most of the security community understands the need for attack surface management by now, sometimes "why" doesn't always neatly translate to "how" and "when" for the Board, your executive team, or your budget. There are obvious reasons for attack surface management, like **protecting your organization's reputation**; it's common knowledge that a breach will likely create bad press and impact customer retention. There are some additional factors beyond reputational damage to include when building the business case for attack surface management.

RISK MITIGATION

The SANS Institute created a comprehensive equation for assessing security investments that personalizes the math to reflect your unique environment, as well as the average impact of the solution in question. As such, this has become a popular method for demonstrating the risk reduction potential for your target investment.

The Return On Security Investment, or ROSI formula requires a business to estimate their annualized loss expectancy (ALE), or the monetary loss from a single incident, multiplied by the number of times such an incident might occur, multiplied by the mitigation ratio, or the expected impact of risk-reduction activities, minus and then divided by cost of solution.

$$\text{ROSI} = \frac{\text{ALE} \times \text{Mitigation Ratio} - \text{Cost of Solution}}{\text{Cost of Solution}}$$

Applying this to attack surface management, let's use the Gartner estimate from earlier, assuming that one-third of successful attacks would be against unknown or unprioritized assets. If that has been, or could be true in your organization, then your ALE might be a little higher. It might also be higher if your attack surface has suddenly expanded due to digital transformation, M&A, or a host of other events that often lead to an explosion of unknown and potentially vulnerable assets.

By using the ROSI formula, you can quickly show security ROI in a way that will speak to your CFO.

RESOURCE SAVINGS

Some large organizations have full-time resources dedicated to hunting for unknown assets in the attack surface, but skill, bandwidth, team size, availability, scope, and more can limit scale and quality of results. In addition, lack of integration, automation, or reporting can slow time to respond.

By partnering with a crowdsourced attack surface management platform like Bugcrowd, you can close this gap by connecting to **a global network of reconnaissance experts** with the skills and experience best suited to your unique environment, allowing you to quickly find and prioritize assets.

But what about organizations that don't have a dedicated function for reducing unknown attack surface? They often have their own set of costly challenges when attempting to spread the work across multiple teams. Asset management alone is often ripe with inefficiencies due to **manual processes and human error**. Automating discovery is great, but automating cross-company alerting and management around at-risk assets *as they evolve*, can save weeks in manual tracking and review for IT teams everywhere.

Crowdsourced solutions like Bugcrowd provide rapid inventory population and categorization, enabling organizations to create customized change-management alerts for things like open ports, vulnerable software versions, or soon-to-expire certificates. Open APIs and publicly available services ensure that information can be disseminated to any number of business functions outside security including IT, marketing, sales, finance, and more.

BUGCROWD PORTFOLIO INTRODUCTION

Bugcrowd's Attack Surface Management portfolio contains a powerful combination of asset discovery, management, and prioritization solutions which when deployed together, can help organizations regain, and maintain control over a dynamic attack surface.

ASSET RISK™

Human-powered discovery infused with platform intelligence for rapid attribution and prioritization

ASSET INVENTORY™

Near-real time asset discovery, alerting and cross-org management for internet-facting assets

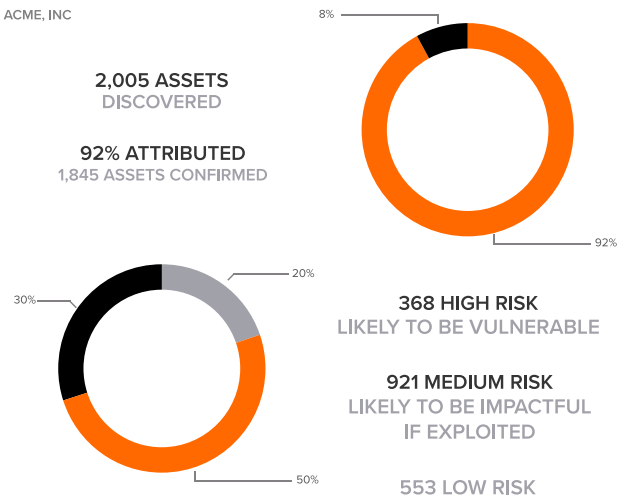

ATTACK SURFACE MANAGEMENT
ASSET RISK™

Asset Risk is best described as ingenuity-driven asset discovery and prioritization. This on-demand offering leverages the power of our global Crowd of vetted security experts to find and prioritize a previously unknown internet-facing attack surface. With access to the latest reconnaissance strategies and tooling from those actively developing them, Asset Risk helps organizations **out-hack digital adversaries before they strike.**

HOW ASSET RISK™ WORKS

1. RESEARCHER ACTIVATION

Bugcrowd utilizes our unique **CrowdMatch™ technology** to locate and activate security researchers best suited for any given reconnaissance program. These individuals are selected by skill, experience, behavior, and trust, and are incentivized in part by the volume and veracity of their findings. This may even include things a scanner is not programmed to surface, like leaked keys, and misconfigured email servers.

2. "SEEDING" THE SEARCH

Organizations supply researchers with information, or seeds, about the known attack surface. Including internal systems and hosting locations can further encourage discovery of things that are inadvertently exposed on the web.

Common seeds include:

- Fully Qualified Domain Names (FQDN)
- Classless inter-domain routing Blocks (CIDR)
- IP Addresses and ranges

3. ASSET DISCOVERY

Commonly called "asset recon," this process of "hunting" for connected and potentially at-risk assets can include use of bespoke discovery technologies in combination with targeted human intuition. Bugcrowd's recon experts don't just find more unique assets,they find more assets full stop, **surfacing 50% more on average**, than programmatic scanners alone.

4. ATTRIBUTION

Domains may be connected to an organization, without necessarily being owned by the organization. This is something that automated scanners struggle to distinguish, creating excessive noise in final results. Bugcrowd's human and technology-driven attribution process helps **verify ownership prior to reporting**, in order to reduce noise and save hours of triage.

5. PROFILING/FINGERPRINTING

Researchers profile attributed assets, examining underlying technologies such as content management systems, CRM, e-commerce platforms, advertising networks, marketing tools and analytics, and more. It's important to note that asset fingerprinting is performed passively, so as not to disturb live systems.

6. PRIORITIZATION

Unlike Bug Bounty programs, Asset Risk™ engagements are focused on going "wide," rather than "deep," to find unknown, at-risk assets. As these can number in the thousands, prioritization is paramount. Bugcrowd utilizes vulnerability and asset data from **over 1,200 managed security programs**, as well as live feedback from the Recon team to assign a priority rating to each asset, even without active testing.

7. REPORTING AND RECOMMENDATIONS

The final report outlines what are believed to be the most vulnerable assets, "colored" by risk rating. While some may be destined for decommissioning (or perhaps were already considered offline!), many may require further analysis. Those assets can be rolled into Bug Bounty or Pen Test programs to fully examine indicators of risk including any high priority vulnerabilities.

ATTACK SURFACE MANAGEMENT
ASSET INVENTORY™

If Asset Risk™ can be summarized by, “human-powered, software assisted,” Asset Inventory™ can be thought of as the reverse. Bugcrowd Asset Inventory™, which is powered by Bit Discovery, is a software-based continuous scanning solution fueled by an ever-growing pre-indexation of (almost) the entire internet. Organizations can configure alerts, filter inventory, and collaborate with other business units to more effectively manage their internet-facing assets. Additionally, extensive APIs help programmatically ensure compliance and security for the business at large.

HOW ASSET INVENTORY™ WORKS

1. ADD DOMAINS

Organizations add domains known to belong to their business, in order to train the search engine on locating additional connected assets. Thanks to pre-indexing, asset inventory can be compiled for virtually any organization within a matter of seconds, not weeks. Subdomains are added instantly, while additional “suggested” domains appear with an option to add to existing inventory.

2. TECHNOLOGY FINGERPRINTING

Extensive technology profiling is applied to each asset to better contextualize information and alert on high-impact security risks like open ports and expired SSL/TLS certificates. Hundreds of known features for technologies ranging from marketing and media, to programming and security are used to automatically characterize each individual asset.

3. AUTOMATE

Continuous discovery keeps organizations apprised of any new or previously unassociated assets, while persistent alerting can inform users when existing inventory becomes at-risk. Public services enable information to be pushed to relevant non-IT business units to bridge compliance and security gaps across the business.

WORKING TOGETHER

While the two solutions can be deployed separately, combining Asset Risk™ and Asset Inventory™ enables insights from one to **fuel and sharpen the activities of the other.** This can improve inventory accuracy, better inform priority rankings, and more rapidly reduce risk across the business.

Let's look back one more time at the question posed early in this guide—**how can you secure what you don't know exists?** The answer is, simply put, you can't. The Bugcrowd Attack Surface Management portfolio solves this problem, combining the scale and persistence of software-based discovery and management solutions with the creativity and impact of the world's best recon hackers. Bugcrowd customers have **uncovered an average of 93% more unknown attack surface**, with clear prioritization and risk-ranking as determined by our "always on" global community of trusted and experienced reconnaissance experts.

	ASSET RISK™	ASSET INVENTORY™	COMBINED
ASSET DISCOVERY		✓	✓
ASSET MANAGEMENT		✓	✓
ATTRIBUTION	✓		
PRIORITIZATION	✓		
FREQUENCY	On-Demand	Continuous	Continuous
DELIVERY	Crowd-Enabled	Software Based	Crowd + Software

Want to learn more about how your organization can uncover and manage an average of 93% more unknown attack surface? Start building your program today at **bugcrowd.com/try-bugcrowd**