

Security Weekly

LABS

A CyberRisk Alliance Resource

SW Labs Product Reviews

Attack Surface Management Product Review:
Bishop Fox Cosmos

Produced by SW Labs, a Security Weekly Resource


CyberRisk
ALLIANCE

Security
Weekly | LABS
A CyberRisk Alliance Resource

 SC MEDIA

About

Contents

Category Overview	3
Bishop Fox Cosmos Product Review	15
Testing Methodology	22

About Security Weekly Labs

Developed by and for security practitioners and professionals, SW Labs aims to guide organizations through the cybersecurity product landscape and help them find solutions that address active problems, narrow selection and confidently make choices. An independent resource — operated by the cybersecurity professionals at Security Weekly and built on the foundation of SC Media’s SC Labs — SW Labs is a clearinghouse for useful and relevant product and service information that enables vendor and buyer to meet on common ground.

At the heart of SW Labs are expertly defined product categorization and product validation methodologies. This framework supports a rich and purposefully organized directory of products and services, a robust calendar of detailed category and product assessments, and a provocative feed of cyber-tech commentary. Laser-focused on the needs of the cybersecurity community, SW Labs is committed to being the essential resource covering the cybersecurity product and service landscape.

This report was developed and written by Adrian Sanabria, who leads the SW Labs initiative for CRA.

Category Overview | Attack Surface Management

Though the term Attack Surface Monitoring (ASM) doesn't specifically refer to external threats, that's what this market currently focuses on. In short, products in this category aim to catalogue and help manage an organization's exposed assets. From this simple definition, the players in this space diverge into various subcategories. The core group we're focused on for the purposes of this group test are products that largely replace the function of an OSINT assessment, an external network vulnerability assessment and some portions of a penetration test.

Introduction

Attack Surface Management is a relatively new category. After discussing it with dozens of practitioners, analysts and founders, it seems clear that this space was born out of a need to fill a gap between vulnerability management tools and penetration testing. At its core, Attack Surface Management is asset discovery and management for exposed assets.

Vulnerability management tools are the most closely related products to ASM and require precise input to give comprehensive output. If we forget to include a website, network segment, API or mobile application - they won't get scanned. If we're not aware of Shadow IT or abandoned cloud projects, they won't be included. Penetration tests will discover some of these gaps, but also have a few shortcomings. First, that penetration tests are periodic in nature: most organizations only have one or two pen tests performed per year. Second, that they are scope and time-limited. Performed on a 'best effort' basis, penetration tests will also potentially miss vulnerable assets.

Asset management is a very different challenge when it occurs on networks and in environments we don't have complete control over. A 'ping scan' isn't going to find an open S3 bucket. An ARP scan isn't going to discover a legacy domain pointing to forgotten cloud infrastructure. A vulnerability scanner isn't going to discover enterprise credentials embedded in a personal GitHub project.

New tools and techniques are required to discover, monitor and manage these assets. They began to emerge years ago as open source projects and reconnaissance tools used by penetration testers. Now, ASM products are using these same techniques (and in many cases, using the same open source tools behind the scenes) to discover this additional attack surface most organizations are currently unaware of.

Currently, the asset discovery and management market is largely focused on identifying internal assets, where an asset is defined as anything with an IP address. In the ASM market, every product is 100% focused on external assets (for now). Perhaps an asset management product could be used to catalogue external assets with IP addresses, but IP-based assets are often a tiny fraction of the external exposed surface. It is a bit like comparing a household vacuum cleaner to a Zamboni. Both products are dedicated to cleaning surfaces, but in very different places using very different methods.

High-profile breach examples: could ASM have helped?

A number of high-profile breaches have occurred due to exposures an organization wasn't aware of. The Buffer breach occurred because a [MongoHQ engineer reused a password](#) that was exposed in an Adobe breach. The Columbia Casualty Company sued Cottage Health System to recover a payout for a cyber breach insurance claim. Why? The insurance company found out that the breach occurred because [patient data was stored on an FTP server with anonymous access enabled](#).

The infamous Equifax breach occurred, not because Equifax was unaware of the danger, but (in part) because they [failed to find struts in their own environment before attackers did](#). There was no software asset inventory that listed Struts as a component of a legacy system exposed to the public Internet. Additionally, dynamic and static code analysis tools failed to identify Struts, despite employees' best efforts.

Each of these examples share the same pattern. There's a risky public exposure that organizations weren't aware of or failed to discover in time. Attackers know what works, so to them, the attack path is clear. Defenders are working with tools literally telling them that there are *hundreds of thousands of paths* the attack could come from! Many of the attack surface management tools we'll be reviewing would have identified and drawn attention to each of these problems and aim to make the most likely attack paths as clear as possible.

While testing these ASM products, two questions kept popping up: how does this differ from existing asset management or vulnerability management offerings, and why aren't those existing vendors active in this space yet? We explore both these questions a bit more below.

How Is ASM Different From External Network Vulnerability Scans?

In the past two decades, vulnerability scanners competed over how much data they could create. More recently, these same vendors are beginning to compete over how much data they can safely ignore. ASM products have emerged in an alert-fatigue-aware era. This awareness is evident in most products we tested.

One of the classic issues with vulnerability management products is their historical close ties to the CVE vulnerability database and CVSS scoring system. Penetration testers know well that some of the most vulnerable findings don't even receive scores — they're often simply listed as "informational", never to be noticed by most organizations. This is especially the case in regulated industries required to remediate all vulnerabilities with a certain score or higher.

The ASM products that aim to prioritize findings (not all do) don't rely on vulnerability databases or CVSS scores. Rather, they look at how relevant a finding is from an attacker's perspective (e.g. Is this something I can use to break in?). They also look at customer-provided and environmental context. For example, has it been marked as a critical asset or is it adjacent to a critical asset. It's worth noting that a market for vulnerability prioritization already exists. Some of these prioritization products already integrate with ASM products.

Here are a few other attributes we found to be unique to ASM products:

- ASM products can start with minimal input — as little as a company name and nothing more. From that starting point, or seed, ASM products will discover and explore other, related properties, subsidiaries, and assets. For example, if you start with a parent company, it's possible an ASM product will discover a one-off project abandoned and forgotten by a subsidiary company three years ago. This concept of 'seed discovery' happens through a variety of methods: website scraping, subdomain guessing, business record lookups, domains with common WHOIS information, information in certificate metadata and many more.
- Several ASM vendors will score findings based on how 'attractive' they are to attackers. While the source for this attractiveness score is part of their secret sauce, it is presumably the product of penetration testing experience and breach analysis (e.g. what gets attacked during actual breaches?).

- Many ASM products gather additional data that an analyst would typically have to enrich through manual processes. For example, an analyst might not recognize the IP address attached to a finding. Is it ours? Is it something we have hosted somewhere? Does it belong to a third party? They'll open another tab to check the ownership records for the IP. Many ASM products do this work for you, automatically tagging assets as Linode or AWS if they are owned by these public cloud providers.
- Most ASM products continuously search for new findings and assets. For example, acquire a new subsidiary or register a new domain and the ASM product will likely begin collecting assets from them on some point, with zero input from the operator (at least, in theory — see the individual product reviews for more information). Keep in mind, this continuous search is doing more than checking existing seeds for new assets, it's looking for new seeds entirely. In theory, if your company acquired another company, some of these ASM products will automatically pick up on this and catalogue the new acquisition's assets as well.

Why Aren't Vulnerability Management Vendors Active In This Space?

Our best guess is that, right now, they don't have to be. The "big three" (Qualys, Rapid7 and Tenable — often referred to collectively as 'QRT') are all large, public companies these days, with the resources to acquire innovation. They could build ASM in-house, or they could decide to wait and see what the market comes up with. Either approach is a valid business strategy. This is an evolutionary, not revolutionary market and we wouldn't be surprised to see vulnerability management vendors make some acquisitions in this space.

One possible reason we haven't seen established vendors step into this space (with the exception of Palo Alto Networks' \$800m acquisition of Expanse) is that it isn't well defined yet. There are so many fringe use cases and techniques to discover and explore assets that hardly any of the vendors currently in this space are even close to feature parity. It is worth mentioning that integrations for ASM vendors already exist on asset management platforms (e.g. Axonius, JupiterOne), in vulnerability prioritization products (e.g. Kenna Security) and in the SOAR space (e.g. Palo Alto Cortex XSOAR).

Scanning The Entire Internet Versus On-Demand Scans

Some competitors perform regular full scans of the entire Internet, so it's worth exploring any potential drawbacks of an ASM product that doesn't. We don't expect this to be a deal killer for most customers unless: 1) the customer needs results within hours (perhaps they're in the midst of an incident) or 2) the customer needs historical data, which are only guaranteed to exist in data sets belonging to ASM vendors that scan the entire Internet on a regular basis and store it indefinitely. It's also worth noting that these vendors tend to scan for different types of attack surface data, so none will be direct apples-to-apples comparisons. Most are also missing large chunks of the Internet, as many organizations don't like being scanned and will automatically send cease-and-desist notices.

There's a debate within this market as to whether ASM vendors will be able to continue scanning the entire Internet as regulatory situations and legal precedents change. Currently, ASM vendors appear to respect requests to stop scanning certain IP ranges, which seems to have kept potential lawsuits at bay. If this changes, we'll likely see these ASM vendors move to an on-demand model, which would break the following use cases:

- Statistical research on the frequency of technology use and exposed vulnerabilities
- Historical research on the same
- The third-party risk monitoring (aka Cyber "Scorecard") business model

Defining The Market

To further define this market, what it is and what it isn't, we'll need to dive further into features and use cases. Finally, we'll define some subcategories and touch on adjacent categories that are related, but not part of the core market.

Terminology

For the sake of simplicity, we'll refer to the components that make up an 'attack surface' as assets. Servers, subdomains, API endpoints, certificates, code repositories, accounts, S3 buckets and much more will all be referred to as assets. The best reason for doing this is simply that nearly every vendor and open source project we explore throughout this group test uses the term asset in the same consistent way. JQuery 2.3.4 is an asset. A subsidiary's Github account is an asset. The IP address of a web server, the web server software running on it and the application hosted on it are all separate assets nested within one another and directly associated with one another.

Attack Surface Management is the primary term we'll use for this space, though we've also seen mapping and monitoring as variations for the Management piece of ASM. Both terms work, but we'll stick with Management as it's most commonly used and best describes how the core tools in this market are intended to be used — for managing assets exposed to the public Internet, which can also be described as 'attack surface'.

Common Market Challenges

False positives

- Assets related to, but not owned by the customer (asset attribution)
- Lookalikes — similar domains and company names, but different organizations

Completeness

- Breadth — finding all the attack surface
- Depth — collecting all the details and metadata related to each entity or asset
- Types — continually adding new types of assets that can be collected (e.g. checking for GitHub accounts associated with a company, mobile apps, etc)

Prioritization

- The more complete these scans are, the bigger the organizational problem becomes. Prioritization is already a key challenge with the products that aim to surface issues in the asset data they collect
- Assigning risk scores — can be done without customer input, but can be much more accurate once asset importance and sensitivity is known

Validation

- Less effective validation methods leading to high false positives (Banner grabbing, keyword searches)
- A few ASM products separate “confirmed” issues from “potential” ones, even providing the proof of confirmed findings. This makes for considerably less work for the analyst tasked with validating these findings.

Categories

The technical approaches and features across vendors in this market vary enough that we felt compelled to break it into a few subcategories. Simply, they can be expressed in terms of how deep they go in terms of discovering assets, prioritizing the results and providing the ability to manage the findings. It is tempting to assign greater value to vendors that go deeper, but the value of greater breadth shouldn't be discounted. Depending on individual needs, use cases and pricing, we wouldn't be surprised to find our readers choosing favorites from more than one category. For example, the historical research use case may not be supported by vendors more focused on prioritization, as most of these vendors don't perform full Internet scans.

Internet Asset Research: The simplest category can be described as a scan of the Internet with an interface allowing the database of assets to be queried. In its simplest form, routable IPv4 address ranges and a limited number of interesting ports are scanned. Services are enumerated and metadata collected. Shodan, SecurityTrails, SpySe, RiskIQ and Censys are examples of these, which tend to have freemium offerings. These tools are widely used by researchers and journalists to explore Internet-wide trends. The results could give an idea of the size and breadth of a zero-day vulnerability, for example.

For those more interested in specific assets (perhaps just the assets they own), these tools are less useful, as results are often missing, incomplete or outdated.

Use Cases:

- Internet Patterns Research
- Historical Research
- Asset Discovery
- 3rd Party or M&A Due Diligence
- Attack surface reduction

Common Features:

- Detailed asset information
- Tagging
- Metadata search
- Complex queries
- API

Pros:

- Quick and easy to perform Internet research or a quick targeted assessment
- Historical data in some cases
- Freemium or low-cost options

Cons:

- Relatively few use cases
- Gaps in coverage due to requests not to scan some networks or dropped probes

External Asset Monitoring: At the next level, vendors have stepped up to also collecting additional, related assets like certificates, DNS records and ‘technologies’ (e.g. software libraries, software frameworks, network software).

They'll also monitor for new assets or changes to existing assets. These tools are tailor built for the long-term monitoring of specific groups of assets. Importantly, they aren't restricted to these groups — it is still possible to use these tools for broad Internet discovery and research outside the customer's organization (something that largely goes away in other categories). BitDiscovery, Shodan Monitor, RiskIQ Digital Footprint, SecurityTrails SurfaceBrowser, and BinaryEdge are examples at this level. Most of these vendors can also be classified as Internet Asset Research as well.

Use Cases:

- Internet patterns research
- Historical research
- Asset discovery and monitoring
- Third party asset discovery and monitoring
- M&A due diligence
- Certificate monitoring
- Competitive intelligence gathering
- Attack surface reduction

Common Features:

- Detailed asset information
- Tagging
- Metadata search
- Alerts on new findings
- Alert on expiring certificates
- Detailed software composition analysis (SCA)
- API

Pros:

- Supports both Internet research use case and asset management use case
- Generally return the most complete dataset on IP-based assets

Cons:

- Large amounts of data to validate with no prioritization
- Missing some asset types (especially non-IP-based)

External Asset Management platforms: The major differentiator at this level is a focus on prioritization and management. Prioritization requires performing some level of risk analysis to separate out risky asset features from the benign. Management functionality adds features like active monitoring, team collaboration, ticketing and commenting.

The concept of seed discovery is significant and worth watching in this market. Simply proving a company name could lead to a news article about an acquisition, which leads the ASM engine to begin collecting assets from both the acquired and the acquirer.

Products in this category more commonly scan for assets on demand and do not retain an Internet-wide asset database (with a few exceptions — see the feature matrix for a detailed list of product features). Randori's Recon product, CyCognito, AlphaWave, Immuniweb, RiskIQ Illuminate, SecurityTrails ASR and Intrigue are the products at this level.

Use Cases:

- Asset discovery and monitoring
- Third party vendor discovery
- External asset management
- Risk prioritization
- Risk validation (via either automated or manual penetration testing)

Common Features:

- Detailed asset information
- Seed discovery
- Tagging (manual and automated)
- Metadata search
- Alerts on new findings
- Detailed software composition analysis (SCA)
- Built for teams with support for commenting
- Issue management with ability to set status, asset importance
- Broad integration support
- API

Pros:

- Identifies issues with assets and prioritizes them
- Discovers risks related to third party vendors
- Issue tracking and management interfaces
- Broad integration support

Cons:

- Generally don't support Internet or Historical research use cases
- False positives are a natural consequence of dynamic asset crawling

Managed External Asset Management Platforms: The primary difference between this category and External Asset Management Platforms is that humans are on staff to validate findings, remove false positives and otherwise ensure the signal to noise ratio is as favorable as possible. While this saves time and effort for the customer, it comes at a price. Bishop Fox's Cosmos (formerly CAST) and Randori's Attack product are the only examples here (though they differ greatly in goals, pricing, and execution; as such, aren't likely to see each other in many bakeoffs — see individual reviews for more details).

Use Cases:

- Everything in the previous category

Features:

- Everything in the previous category, plus
- Outsourced staff to perform validation on any findings

Pros:

- All signal, no noise (in theory — note we did not directly test either of these products)

Cons:

- Higher cost

Notable Adjacent Categories

Third party risk monitoring vendors use similar techniques to gather open intelligence on an organization. However, they use this data to generate risk scores, intended to indicate how safe or risky a business is to work with. The use case is different enough that we've decided to evaluate these vendors (BitSight, RiskRecon, Security Scorecard and a few others) in a separate group test.

The **Data Loss Detection** category scours the Internet for any evidence that a company's private data (credentials, documents, etc) might be exposed to the public Internet. Examples include Digital Shadows, Terbium Labs and Intelliagg.

Vendors in the **Digital Risk Protection** category aim to spot any attempts to impersonate an organization, brand or individual. They often also assist with attempts to take down or disrupt these impersonation attempts. Examples include ZeroFOX, PhishLabs, Constella Intelligence and Digital Shadows.

Asset Reputation is a category that catalogues and reports on the behavior of various assets exposed to the public Internet.

GreyNoise, one example in this category, uses a global sensor network to observe the behavior of assets aggressively scanning the Internet. The most common use case for this data is to separate non-malicious noise from potentially malicious actors. Another use case is related to ASM, however. It is possible for customers to use GreyNoise's database to monitor the behavior of their own assets, or those of subsidiaries or key third-party vendors. If the customer receives an alert that an asset is suddenly behaving maliciously, they can take action.

Also in this category would be services that monitor email and IP blacklists.

Conclusion

There is an immediate need for these products. Nearly every product we tested discovered assets and issues we weren't previously aware of. Additionally, these were assets and issues that traditional vulnerability management products did not alert us to. The dilemma here is that nearly every product we tested surprised us in different ways with different results. One product discovered a branded mobile app. The others don't look for mobile apps. Two spotted an old version of JQuery. None of the others did. One found a few domain names the others didn't find.

Many products had their own niche abilities to discover attack surface that set them apart from the competition. While we can't recommend buying half a dozen ASM products, this is fairly common in new markets and we do believe the market will more or less achieve feature parity over the next year or two.

With respect to adjacent asset and vulnerability management categories, we don't expect to see the market to remain fragmented for long. In the next three years, we'll either see traditional vulnerability management products acquire ASM vendors, or we'll see ASM vendors begin to challenge and even replace external vulnerability scans. We've seen the same trend play out in the endpoint market over the last six years. A simpler, more effective approach, even if incomplete in terms of features, can challenge the incumbents and steal away market share.

After all, we know that complexity is the enemy of security — shouldn't this principle apply to security products as well?

SW Labs Product Review | Bishop Fox Cosmos

Originally founded in 2005 as Stach & Liu and rebranded in 2013, Bishop Fox is a widely recognized security services firm. It employs over 250 and is headquartered in Phoenix, Arizona. Its employees have produced numerous books, research, talks and open source tools over the years. While it's common to see individual creations come out of cybersecurity consulting firms, they occasionally step into the product space as well.

In early 2019, Bishop Fox raised a \$25m series A from ForgePoint Capital to do just that. The initial result of diversifying into the product space is Cosmos, which provides continuous external attack surface testing.

Understanding Cosmos

First and foremost, Cosmos is unlike all the other ASM products we've tested. Other ASM vendors started with the data and have been slowly working towards risk analysis and validating findings, using both humans and automated means.

Bishop Fox has come at the problem from the opposite direction. As a fifteen-year-old consulting firm, they already had the human side of the equation — the talent. They then saw an opportunity to automate all the manual work of gathering, organizing, and analyzing the data - turning it into a product, backed by experts. The final piece was the business model. By leveraging automation, it was possible to deliver a service that constantly monitors companies' exposed assets year-round without labor costs making the service financially prohibitive.

One issue with traditional penetration tests is that they are point-in-time, typically performed only once or twice a year. The gaps left between these assessments was an opportunity for an organization's security posture to decline, or at the very least, become unclear to those managing it.

Another issue with traditional penetration tests is that they are time-bound. Typically, the assessment is scoped for a finite number of hours. The penetration tester gives it their 'best effort' for that given time. It might be enough to assess 30% of the company's external assets, or perhaps only 5%. It's unlikely that a traditional penetration test will come anywhere near touching 100% of an organization's exposed assets, however.

Cosmos addresses both these shortcomings. By running continuously, it is reasonable to expect to find all the assets there are to find (within the asset types Cosmos is designed to discover, of course). As soon as a finding emerges, Bishop Fox's team of offensive security experts can perform a penetration test on that one small element (they call it a micro-pen test).

Another added advantage is that the typical penetration tester mindset doesn't enjoy repetition. Part of the workflow for the offensive team is to continually improve the Cosmos engine to automate and improve testing and decision-making. As a result, the engine can make decisions more and more like a penetration tester does. As it gets more efficient, the Cosmos service can scale more like a traditional software product and less like a human-driven consulting business.

Competitively, in a few years, we may see ASM vendors CyCognito and Randori, who already have offensive experts on staff, meet in the middle and begin to compete with Bishop Fox for larger customers. Currently, however, the customers and priorities Cosmos is designed for differ from the rest of the ASM market. Cosmos also pitches itself as an alternative to private bug bounty programs, where it could make a case that it casts a wider net, finding issues that the bug bounty crowds aren't as skilled in discovering. It's also easy to imagine Cosmos as an alternative to hiring additional internal staff to do the same job, less efficiently.

To be clear, Bishop Fox isn't going to stop selling penetration tests because they're now selling Cosmos. There will likely always be a need for point-in-time security assessments. However, Cosmos makes a *lot* more sense for enough of Bishop Fox's existing customers that it was a no-brainer to raise some funding and develop it.

Review Summary

While the real selling point for Cosmos revolves around the skilled staff that both validates findings and trains the 'brain' of the automated system, the customer will be primarily interacting with the web-based portal application. It should be noted that interaction with the experts behind Cosmos's mini penetration tests is welcomed whenever customers have questions or wish to discuss findings.

While we give an overview of the Cosmos platform in the Usage section of this review, it's likely that most Cosmos customers will log in to research or verify something only occasionally. Bishop Fox's ideal customer will likely integrate Cosmos into an existing vulnerability remediation workflow (e.g. a ticketing system like ServiceNow). This is the primary value of the product — if anything seriously critical shows up, the product will let you know without a security analyst needing to log in every day to tap on dials and gauges.

Target market: Fortune 500 companies with large, complex attack surfaces

Time-to-value: As with other ASM products, the vendor asks for initial seed values (company names, domains) and initial results are available in a few hours.

Maintaining value: As a product with a managed component, maintenance is only necessary if additional companies, subsidiaries, or seed values need to be added to Cosmos for monitoring and analysis.

Total cost: Bishop Fox did not provide pricing, but we'd expect it to be attractive when compared against the costs of building a small internal team to provide the same function as Cosmos.

Strengths: Cosmos effectively reduces the validation work in ASM to zero, leaving some lighter analysis, attribution, and remediation work for internal staff to handle.

Weaknesses: Market limiting price point.

Conclusion: A view into a more efficient, hybrid future for services and product.. When intelligent automation is combined with human services, buyers ultimately win.

Deployment And Configuration

The initial setup process is notable, as Cosmos joins a small club of ASM vendors that not only offer multifactor authentication by default, and they require it to be configured on first login. Be the change you want to see, right?

As with other SaaS ASM products, the initial seed values and provisioning are handled by Bishop Fox for the proof-of-concept, which will become the customer's production instance if they choose to purchase the service.

Usage

The interface is reminiscent of most Attack Surface Management (ASM) products we've reviewed at SW Labs. An intuitive dashboard summarizes findings and assets. Charts break down assets by category and display time-based trends.

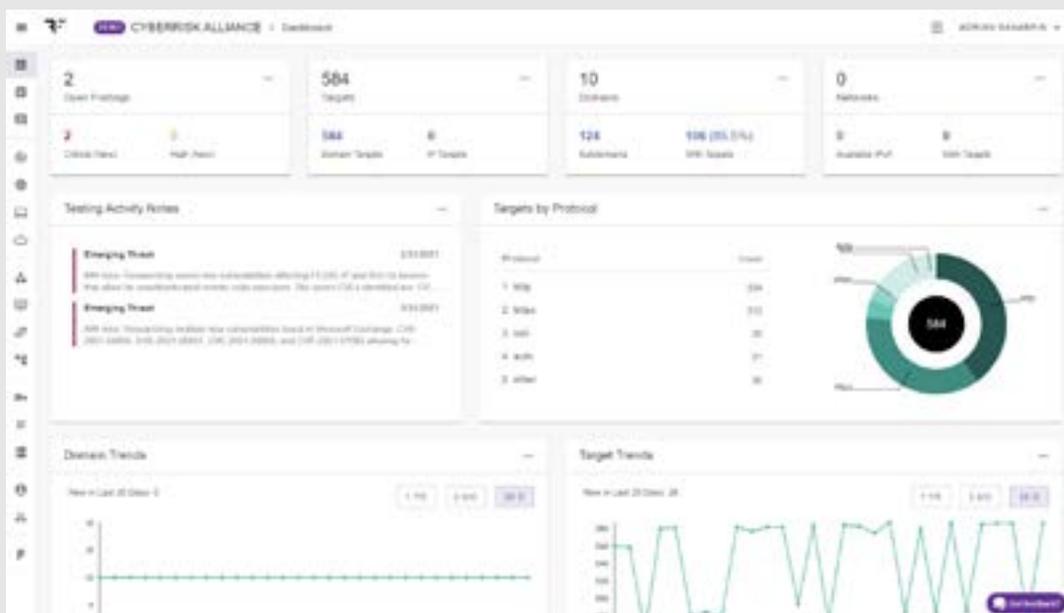


Figure 1 — The Cosmos Dashboard

Findings can be filtered by severity, status, or age. A search interface is also available, though with a service that aims for zero false positives, I wonder if any but the largest customers will have enough findings to justify the use of it (which is generally a good thing, in case that's not clear).

As we drill down into the details of a finding, we begin to see where Cosmos separates from ASM products that are largely automated. Findings read much like they would in a penetration test report. Finding details, screenshots and evidence of exploit or validation are present. For findings that require some narrative, took multiple steps, or succeeded in pivoting further, all these details are included. As we'd expect, credentials and other sensitive details are redacted, but the screenshots make it clear the staff made the 'phone call' from inside the house.

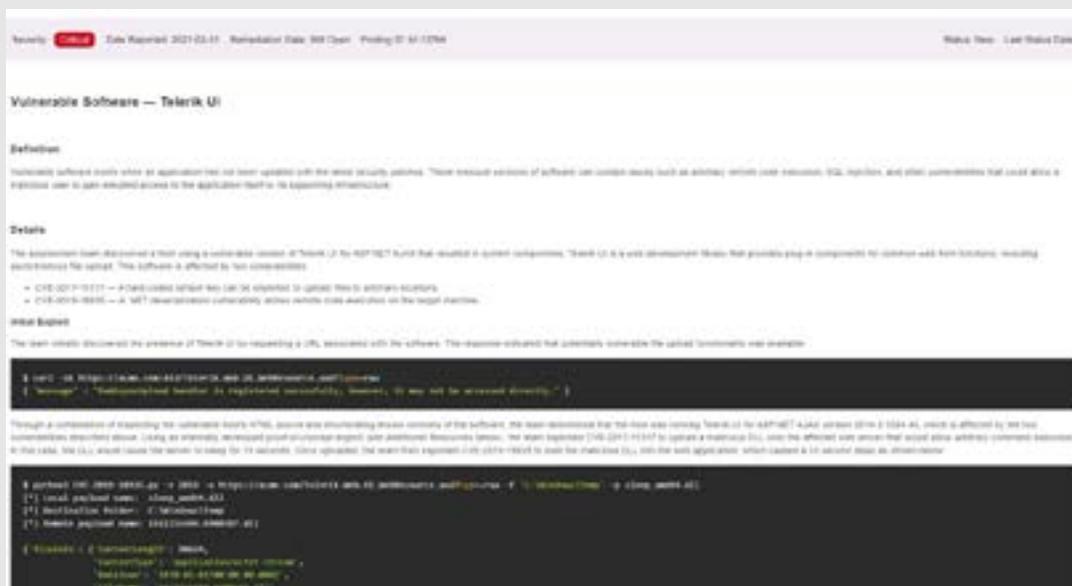


Figure 2 — Detailed findings read like a penetration test report

At this point in the review, it's clear that Bishop Fox isn't selling an ASM product. They're selling a continuous offensive security service — the ASM piece merely enables and feeds that service. Think of Cosmos as an external penetration test that never ends. As soon as new attack surface emerges, it is analyzed by the Cosmos system. The system sends anything interesting to the the Cosmos team and they go to work.

Returning to the interface, assets (called "Targets" in Cosmos) are organized by protocol, URL, protocol, and port. Filters can be applied using any of these elements and it's possible to search by IP or host name.

Like others, Cosmos records screenshots of any websites it finds. Unlike others, however, Cosmos helpfully groups similar screenshots. In large environments, there could be tens of thousands of screenshots, many of which will be identical or have only slight variations (e.g. dev, QA and prod versions). Drop into "quick mode" and the arrow keys can be used to quickly flip through groups of screenshots — a sure pleaser for analysts that click and scan through data all day long.

Cosmos also provides a breakdown of technologies discovered (the classic comparison is [BuiltWith](#) — a service that tells you what technologies are used by a given asset, company, or app). Here, however, the data isn't nearly as detailed as it is in ASM vendors that focus more on comprehensive data collection than risk analysis and prioritization. It lists technologies but isn't forthcoming on versions in use and doesn't link back to the assets (targets) where these technologies were discovered.

One interesting view is a breakdown of targets by hosting provider. At a glance, you can see how many IP-based assets the organization has connected to CloudFlare, hosted by AWS, running on Linode or hiding on DigitalOcean (that's where all the Shadow IT stuff lives, right?).

Lest you think Cosmos is limited to ports and protocols, the next section of the UI is dedicated to domains, subdomains, and networks. Part of the onboarding process with new customers aims to discover all domains used by the organization. We gave Bishop Fox a list of six domains, intentionally leaving some out. Cosmos discovered an additional four domains we never provided to Bishop Fox or the Cosmos software.

The subdomains view includes a list of subdomains, likely discovered and brute-forced using a variety of techniques. Interestingly, domains determined to be abandoned are labeled as such. While the labels seem accurate, there is no reason listed as to why each domain has been tagged as abandoned.

Security Program Fit

Bishop Fox, like other products and services focused on discovering vulnerabilities and misconfigurations, fits solidly within the Identify column of the [Cyber Defense Matrix](#).



Conclusion

For the Fortune 500 (or even the Global 2000), we can see this offering appealing to organizations struggling to attract and retain talent for this kind of work or prefer to prioritize that talent elsewhere. The quality and depth of findings can't be matched by automated approaches today, and this level of experience is priced accordingly.

We're intrigued by this development and can't help but wonder if this is a window into the future of some cybersecurity services. The answer likely lies in the ability for this approach to scale down market, to smaller enterprises.

Methodology

Our aim is to engage with vendors as closely to an actual customer as possible. If a free trial is offered, we take it. If it is necessary to first engage with sales and request an account, we use the contact options provided on the website and wait for a reply, even if we already have contacts at the vendor.

However, while we engaged like a prospective customer would, we made no attempt to hide our identity or intentions at any point. We use real names and identify ourselves as employees of CyberRisk Alliance. We're clear from the very beginning that we intend to perform product reviews and publicly publish the results. No compensation is requested or accepted for any of our reviews.

CyberRisk Alliance monetizes product reviews by licensing product reviews for redistribution after they have been published (commonly known as "reprint rights"). It does occur to us that positive reviews are more likely to sell reprints. We believe that enough vendors are interested in an honest, independent, and unbiased review that we don't have to worry about making everyone happy. With that said, our reviews will be as polite and fair as we can make them.

We try to establish testing methodologies and share them with vendors before testing begins. However, it isn't always possible to make testing methodologies available to vendors with new categories. It's necessary to spend some time with the full range of products to understand the bounds of the categories and how to measure their performance. On the topic of performance, our reviews intentionally highlight product features and the customer experience over technical performance. We believe that technical performance, while important, shouldn't be the focus at the expense of other product attributes.

Finally, vendors are given an opportunity to review drafts before publication. The purpose of this is to ensure the content of our reviews is factually correct, fair and doesn't include any information protected under NDA. We are clear to vendors that this is not an opportunity to insert marketing copy or rewrite our reviews. Any attempt to do so is ignored.

Attack Surface Management Testing

Most attack surface management products require very little input to start the process. We provided each ASM vendor or product with seven domain names, asking that they create the account and kick off discovery as they do for every other customer. If the process is something they kick off, we had them do it. If they had a POC kickoff briefing, we attended that briefing. In cases where hands-off free trials were available, we handled as much as possible by ourselves, unless support was needed to address an issue.

As a basis of comparison, we used the community edition of [Maltego](#) and a few other common OSINT tools to create a baseline for these seven domains, much like an offensive security consultant might do during an OSINT assessment. Approximately two hours was spent manually gathering OSINT data with these tools.

Due to this being a new category we knew little about, we focused our time on understanding the market, how to categorize it and exploring each product's set of features. While we see opportunities for some performance testing, the results would be difficult to compare in a meaningful way. This is due to the lack of feature parity across vendors, which is unsurprising, given the relatively young age of the market.

Instead, our reviews will contain less testing and more explanation of how these products work and compare with one another. We do have some testing metrics that are generally universal, and we try to apply to all products we review. You can read more about those below.

Defining Value

For all product tests, it is necessary to define a tangible "value" to derive some of the metrics we use to evaluate products. Ideally (for us), value would be defined the same for each product within a particular category. However, many products have unique features and key differentiators that may result in a different definition of 'value' from their competitors.

The value of ASM products is derived from a variety of sources, due to the variety of use cases:

- Provide a comprehensive inventory of publicly accessible assets
- Evaluate the risk represented by these assets, noting issues that should be addressed
- Prioritize any issues discovered
- Continuously monitor these assets, reporting any changes or new assets discovered
- Perform 1-4 with as little input from operators as possible (put another way, value can be measured as analyst time saved)

Metrics

Time-to-value is a metric that describes the amount of time it generally takes to get a product from zero to fully deployed and producing value. The clock for this metric begins when the vendor provides access to the product (e.g. an account to a SaaS product or license key + software download).

Labor-to-value is a metric that expresses the effort necessary to *keep* the product at a level of performance where it is providing value consistently.

True Cost is a metric that expresses the total cost of a product, including capital expenditures, operational expenditures, and labor costs. It is effectively product cost + initial deployment cost + maintenance costs, where the following labor cost assumptions are used. We've listed salaries along with the actual cost of the employee to the employer, based on the [US Small Business Administration's most conservative estimate](#) (1.4x of salary). We calculate hourly rates by dividing the actual cost of the employee by 2080 hours (52 weeks multiplied by 40 hour work weeks).

- Junior Security Analyst Salary: \$50k USD (\$70k) - \$33.65/hr
- Security Analyst Salary: \$75k USD (\$105k) - \$50.48/hr
- Senior Security Analyst Salary: \$100k USD (\$140k) – 67.31/hr

To use this in an example, a 1-hour meeting with two senior security analysts and two junior security analysts costs their employer \$201.92.

Other Metrics Considered

1. Account setup process
2. UI/UX navigation
3. Time to discover asset information (some products require a day or two, while others return results in real-time from an existing database)
4. Accuracy of results
5. Usefulness and quality of reporting and dashboards
6. Integration options
7. API functionality