SW LABS

Category Overview | Attack Surface Management

Product Review | CyCognito

Produced by SW Labs, a Security Weekly Resource

CyberRisk ALLIANCE SecurityWeekly SC MEDIA

About

Contents

SW 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

Introduction

Though the term Attack Surface Management (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.

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.

Other attributes of 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

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.

Approaches and features

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.

Internet Asset Research at-a-glance

Use Cases

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

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.

External Asset Monitoring at-a-glance

Use Cases

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

Features

- Detailed asset information
- Tagging
- Metadata search
- Detailed software composition analysis (SCA)
- Alert on expiring certificates
- Alerts on new findings
- 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 a some non-IP-based assets

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.

External Asset Management at-a-glance

Use Cases

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

Features

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

Pros

- Issue tracking and management interfaces
- Identifies issues with assets and prioritizes them
- Discovers risks related to third party vendors
- 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 *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).

Managed External Asset Management at-a-glance

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 | CyCognito

This review is part of the April 2021 assessment of the Attack Surface Management (ASM) product category. If you haven't read the category overview, you might want to check it out – it explains the category's basics, use cases and the general value proposition. Our testing methodology explains both how we interact with vendors and how we tested these products. In short, ASM products aim to discover and manage an organization's external digital assets. This approach extends far beyond assets with an IP address, however, including everything from certificates to S3 buckets to DNS misconfigurations.

Company background

CyCognito was founded in January 2017 by Rob Gurzeev, CEO, and Dima Potekhin, CTO. The company takes the tried and true approach of a US business headquarters (Palo Alto) with the tech team located in Israel (Tel Aviv). The company employs around 100 people globally. The CyCognito Platform was launched Nov. 19, 2019 and currently serves over 50 paying customers in the Global 2000, according to the company.

CyCognito has raised three rounds of funding, totaling $53 million. The latest round, led by Accel, was a $30 million Series B announced in mid-2020. Lightspeed Venture Partners, Sorenson Ventures and UpWest participated in the Series B and previous rounds.

Product summary

CyCognito has one of the most functional, mature and stylized user interfaces (UI) of all the ASM tools we tested. Maybe it just comes down to personal preference, but it seems like more security tools should embrace the dark mode trend. The UI isn't all flash and no substance, however, as we explore in more detail in the Usage section of this report.

CyCognito is firmly in the category of ASM products that focus on risk analysis and prioritization. The platform gives an A-F letter grade and it is important that this grade isn't an average of all assets. The issue with averaging risk across findings is that a large number of non-risky assets could potentially raise the overall grade, giving a false sense of security. Even one risky asset could result in a breach, so it makes sense for the overall grade to resemble the biggest individual risk.

A key challenge in this space when attempting to perform risk analysis and identify issues are false positives. Making assumptions based on banner output or keyword searches are common techniques that have a higher tendency to produce false positives. CyCognito is one of the few vendors that perform some active security testing in an automated manner to increase confidence in findings (each finding scores a confidence level on a 100 point scale, in fact). This testing can validate whether an FTP service requires TLS, discover SQL injection, and identify exposed internal network devices, for example. This can cut down significantly on manual validation work, especially at larger scales.

Target market: CyCognito's asset-based subscription pricing is flexible, but starts above what an SMB could afford. The company is currently targeting large enterprises, likely to meet growth goals and maximize the value from a long-ish sales cycle, if we were to guess. There's no functional reason the product couldn't move down market in the future – it's more a question of the company reducing acquisition cost to the point where it starts making sense to go after smaller deals.

Time-to-value: Allow 30 minutes for initial account setup and team configuration. Initial results will be available after a few hours, with full results after a few days. Integrations could take anywhere from a few minutes to several hours, depending on whether any custom work needs to be done. A large enterprise should expect to spend 40-60 hours digging through initial results, which will vary depending on size of the organization's external footprint.

Maintaining value: In a large enterprise, expect to spend 6-12 hours per week reviewing new findings, organizing them and validating issues.

Total cost: Pricing is an annual subscription based on number of Internet-exposed assets with tiered discounts as the number increases. CyCognito declined to share specific pricing information, but our guess is that the annual cost will be comparable to what a large enterprise typically spends on an annual penetration test. As for labor costs, in a large enterprise (10,000+ employees), we've estimated 24 hours of junior analyst time, 12 hours of analyst time and six hours of senior analyst time to do the initial analysis of CyCognito results, some tuning and integration work. Based on our average salary estimates, this labor totals $1,817.22. For ongoing analysis work, we've estimated six hours of junior analyst time per week, which totals $10,498.80 per year.

Strengths: Strong core ASM features. One of the strongest and most usable UI/UX in the ASM market.

Weaknesses: Stops short of validating many types of issues (software exploits, for example), though it does validate some issues with provided evidence.

Conclusion: One of the most comprehensive offerings in the ASM space and one of the most advanced in terms of seed discovery.

Deployment and configuration

As with many other ASM products, the initial seed values are entered by the vendor as the customer's SaaS environment goes through initial provisioning. We found it initially surprising that CyCognito doesn't appear to support native MFA. Then we realized that nearly 100% of their target market will likely be using SSO integrations, which will handle authentication (including MFA) outside the CyCognito Platform.

Initial tuning of the product is straightforward. To share an example, when first exploring the product, we noticed the default dashboard features a world map, showing the geo-location of IP addresses. To our surprise, the map was suggesting we have assets in Brazil and France? That didn't seem right.

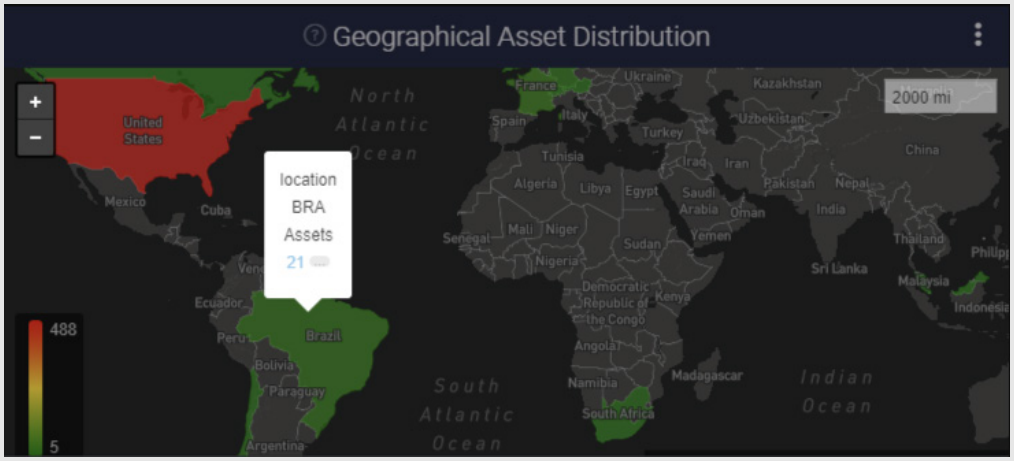

Selecting Brazil, it was easy to open 21 related assets in a new window. It quickly became clear that these were all Office365 resources – they nearly all have the 'autodiscover' CNAME and are helpfully tagged as "Microsoft Office 365" under Platforms.

While these aren't technically false positives, they're noise to me at the moment, so I want to filter them out. A quick filter (Platforms = Office 365) gives me 180 Office 365 assets, which I can then select and exclude from all future dashboards and queries. If I change my mind, it's not a problem – I can either check the "show removed assets" option, or use the same process to bring them back into normal query results, reports and dashboards. The "Exclude asset" dialogue box even prompts to include a comment explaining why these assets are being excluded.

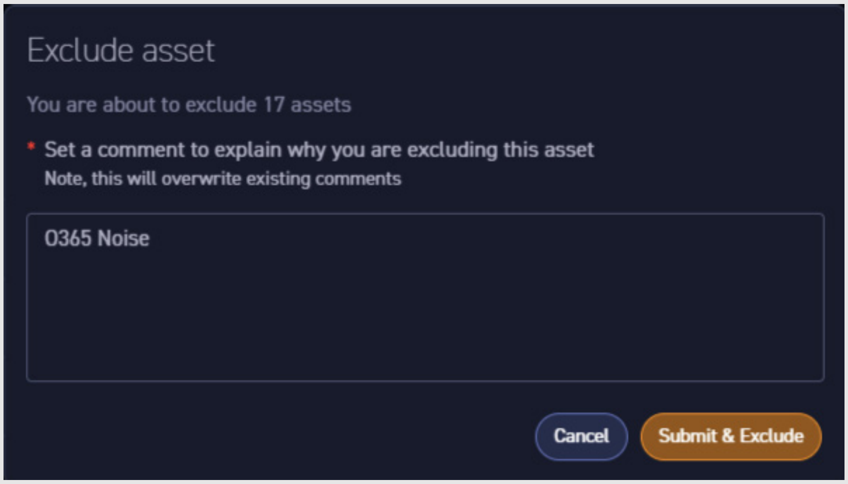
Exclude asset

You are about to exclude 17 assets

* Set a comment to explain why you are excluding this asset
Note, this will overwrite existing comments

0365 Noise

Figure 1 – Excluding third party assets

This is exactly how filtering should work in security tools. There's almost always going to be noise and false positives in security tools, but when it does, it should take seconds to identify it in bulk and remove it from the analysis view. While it makes sense to try to improve scanning and detection engines to reduce that noise, the ability to easily search, filter and exclude results is table stakes.

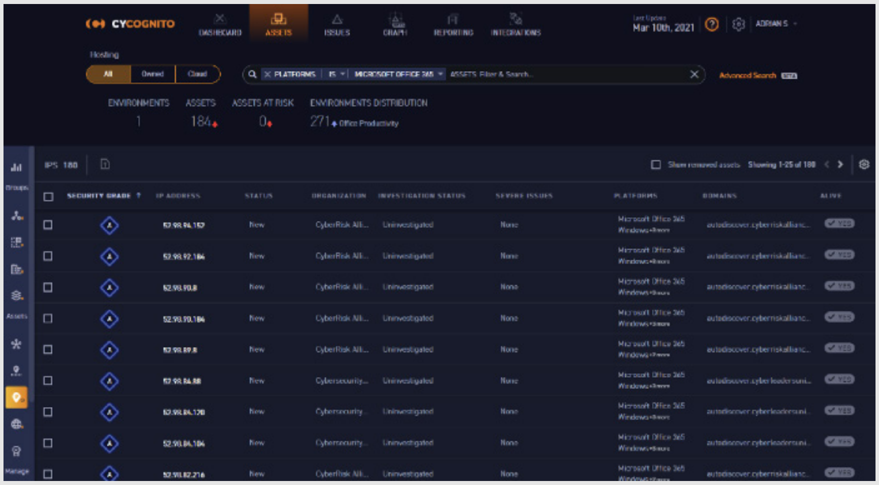

ENVIRONMENTS	ASSETS	ASSETS AT RISK	ENVIRONMENTS DISTRIBUTION
1	184	0	271 Office Productivity

IPS 180

□	SECURITY GRADE	IP ADDRESS	STATUS	ORGANIZATION	INVESTIGATION STATUS	SEVERE ISSUES	PLATFORMS	DOMAINS	ALIVE
□	A	52.98.94.152	New	CyberRisk Alli...	Uninvestigated	None	Microsoft Office 365 Windows [illegible]	autodiscover.cyberriskallianc...	✓ YES
□	A	52.98.92.184	New	CyberRisk Alli...	Uninvestigated	None	Microsoft Office 365 Windows [illegible]	autodiscover.cyberriskallianc...	✓ YES
□	A	52.98.90.8	New	CyberRisk Alli...	Uninvestigated	None	Microsoft Office 365 Windows [illegible]	autodiscover.cyberriskallianc...	✓ YES
□	A	52.98.90.184	New	CyberRisk Alli...	Uninvestigated	None	Microsoft Office 365 Windows [illegible]	autodiscover.cyberriskallianc...	✓ YES
□	A	52.98.89.8	New	CyberRisk Alli...	Uninvestigated	None	Microsoft Office 365 Windows [illegible]	autodiscover.cyberriskallianc...	✓ YES
□	A	52.98.86.88	New	Cybersecurity...	Uninvestigated	None	Microsoft Office 365 Windows [illegible]	autodiscover.cyberleadersuni...	✓ YES
□	A	52.98.86.120	New	Cybersecurity...	Uninvestigated	None	Microsoft Office 365 Windows [illegible]	autodiscover.cyberleadersuni...	✓ YES
□	A	52.98.86.184	New	Cybersecurity...	Uninvestigated	None	Microsoft Office 365 Windows [illegible]	autodiscover.cyberleadersuni...	✓ YES
□	A	52.98.82.216	New	CyberRisk Alli...	Uninvestigated	None	Microsoft Office 365 Windows [illegible]	autodiscover.cyberriskallianc...	✓ YES

Usage

CyCognito Platform defaults to the dashboard, which has some useful trending and summary widgets. However, if the default dashboard doesn't do it for you, eight more dashboard options can be set as the default landing page.

All of CyCognito's dashboards are dynamic and can be filtered by tag, organization or location. It's also possible to pivot from specific elements in the dashboard to the detailed findings. Pivoting deeper into the data helpfully opens in a new tab, so your place on the dashboard isn't lost. In general, CyCognito has one of the most functional interfaces. Large amounts of functionality and detail are baked into the SaaS product, while still supporting a natural, intuitive workflow.

The platform's key app sections include assets, issues, 'graph', reporting and integrations. The assets and issues sections are where most of the work gets done. Both sport their own dashboards, which are also configurable.

Again, it's easy to pivot into the assets given a certain grade, right from the main dashboard (not to be confused with the asset dashboard, which helpfully breaks down asset grades by environments, protocols, technologies and a number of other variables).

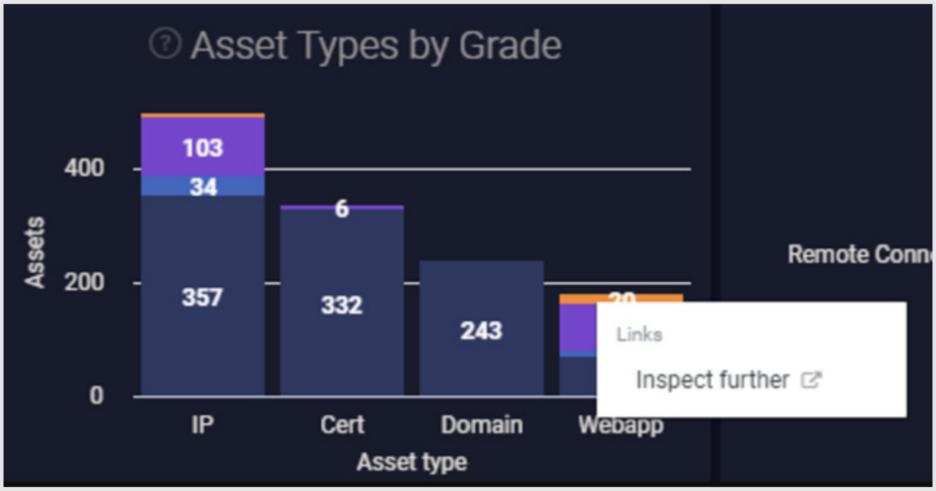

After pivoting, it became clear that most of the "D" grades here are due to another third party we don't have control over, so I'll exclude those as well. It's worth pointing out that third parties aren't necessarily false positives – most organizations entrust data and even direct access to third parties and want to understand their security posture as well. When doing initial analysis, however, most analysts will likely want to focus on first party issues before they turn their attention to third party concerns.

Another excellent workflow design feature is making it easy to create an alert from any search query. It's the ideal way to build out notifications in a new product – naturally creating alerts as discoveries are made during the analysis process. There's no learning curve with a query language – the search bar will automatically create a dropdown that will show you all available options within this stage of the query. Also, no need to build out scripts or learn APIs to enable notifications – just a visual wizard that prompts to complete a few fields and the task is done.

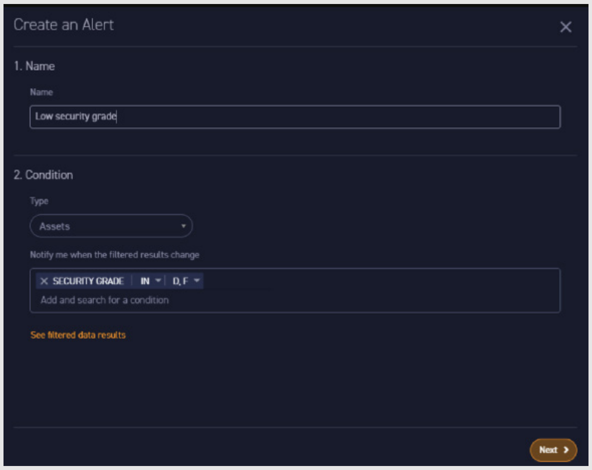

Create an Alert

1. Name

Name: Low security grade

2. Condition

Type: Assets

Notify me when the filtered results change

SECURITY GRADE | IN | D, F

Add and search for a condition

This alert can be delivered to an email address, or to an integrated ticketing system. Current options for self-service integrations include Service Now or Jira. Additional options can be requested right from the integration screen and the company says they also have out-of-the-box support for Slack, GSuite, Zendesk, Splunk, Microsoft Teams and Office365.

Getting into the details

Diving into a specific asset (which happens to be a web application), there's a wealth of information to consume. Basic asset details, a screenshot (for web assets) and a security overview with a grade, discoverability score, attractiveness score and total number of issues associated with that asset. Comments and tags can be entered and associated with the asset.

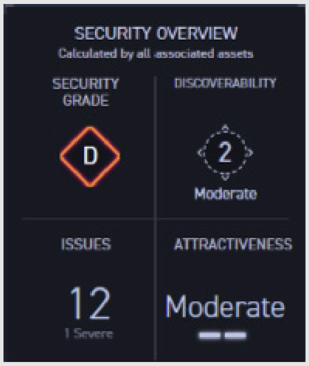

SECURITY OVERVIEW

Calculated by all associated assets

SECURITY GRADE	DISCOVERABILITY
D	2 Moderate
ISSUES	**ATTRACTIVENESS**
12 1 Severe	Moderate

Further down the asset page, the discovery path details how this particular asset was discovered. These discovery paths can be fascinating – the path for this particular case goes 5 levels deep:

- The seed began with a parent company and a subsidiary was discovered by looking at acquisitions on Crunchbase.
- From this new subsidiary company name, it looked for domains owned by the subsidiary.
- It checked the Whois record for one of the domains owned by the subsidiary. In this Whois record, it found that the contact email address had a different domain.
- From the email address's domain, it found a website.
- By crawling this website, it found a Wiki site under a subdomain. This subdomain is the asset we're currently looking at.

The asset page goes on to list all 12 of the issues associated with this asset. Finally, the asset page lists the root webpage for this asset, gives another security overview for the application layer components, details about the page and a list of URLs that lead to this particular page.

The detail view for issues is similarly comprehensive and structured. A particularly useful feature on the issue detail page are all the asset details necessary to determine context and attribution. Everything from open ports to ASN (IP ownership) and, of course, the discovery path, are available here without having to open another tab or leave this particular page. Something as simple as reducing the need to open additional tabs when analyzing findings can yield huge improvements to workflow speed and cognitive fatigue for the analyst.

Back to the big picture

There are number of different asset types and views available. They can be grouped by organization, environment, business units or platforms. Raw lists of assets can be viewed by IP address, web application, IP ranges, domain names or certificates. Finally, under assets, it is possible to manage seeds, inclusion lists, exclusion lists and view the rescan log (every asset and issue detail page includes a rescan option – helpful to check whether the fix you just put in place actually fixed the problem).

Viewing issues includes a dashboard view and views that filter by issue status (open, resolved or archived). Issues can additionally be filtered by 'risk type'. Examples include vulnerable software, abandoned asset, security hygiene and certificate validity, to name a few.

The Graph view is still in beta, but offers some more visual ways of viewing asset and issue data. It starts off with a tree map of all organizations. Selecting a specific organization leads to a hierarchical chart that will be familiar to anyone that has ever seen an org chart or used Maltego. As with other analytical views, this one includes the same filtering and search capabilities, which applies to a graphical view instead of a table view in this case.

Reporting in CyCognito refers to a multitude of different ways to get data out of the platform. It is one of the few ASM products we tested capable of creating an executive summary PDF, suitable for passing up to executive leadership. The executive summary looks clean and strikes a nice middle ground between detail and verbosity. While alerts and exports can be created from nearly any view in the product, summaries of all active alerts and historical exports can be managed under the reporting section of the product.

Performance

For the first ASM group test, we're not doing any formal performance testing, as there is little feature parity in this fast moving market at this point. Any performance results from this test would likely be invalid in just a few months.

With that disclaimer, we can share some anecdotal insights. In our testing, CyCognito found a significant number of domains and software concerns that other ASM products didn't catch. While it's true that nearly every product in this market made unique discoveries over competitors, CyCognito came up with more unique and significant findings than any other.

The techniques used to discover new seeds, which lead to new attack surface, are very important in this market. One clear lesson from past data breaches is that it's often the assets organizations aren't aware of that hurt them. Also, the discovery of unknown assets is arguably the flagship use case for the ASM market. The performance of an ASM vendor's 'discovery engine' should be a key purchasing decision for any buyer in this market.

Specifically, CyCognito excelled at discovering outdated infrastructure and application software – from outdated javascript libraries and content management systems to outdated database services and administrative consoles.

Notable integrations

- ServiceNow (ticketing)
- Jira (ticketing)
- SSO via Auth0 (authentication)

Roadmap

- Discovering SaaS vanity URLs (May 2021)
- Automated vulnerability exploit validation (beta in Q2, GA in Q3)

Support

CyCognito offers email and ticket-based support and has a dedicated knowledge base with answers to common questions, user guides, and tutorials.

Claims

CyCognito advertises it gives customers the ability to "Get the Advantage Over Attackers". We'd buy that – it's really the primary use case for the entire ASM market – improving visibility to the point where attackers don't have an advantage due to 'asset blindness'.

The company's website also has a number of claims about eliminating the risk from shadow IT and remote workforce risks, which is partially true. While Cycognito can certainly discover some shadow IT risks, it isn't able to discover shadow IT use of commercial or consumer SaaS (e.g. Evernote, Dropbox), which gets more into a CASB use case, nor does it search for some common shadow IT asset types that other ASM vendors will discover, like personal GitLab or GitHub repositories that contain company-owned code.

That really highlights the state of this market currently – every vendor gets you part of the way there, but no one gets you all the way there and some get you further than others. It's early days – we anticipate nearly all ASM vendors will add more asset types in the coming months and years.

Security program fit

CyCognito Platform, like other products focused on discovering vulnerabilities and misconfigurations, fits solidly within the Identify column of the Cyber Defense Matrix.

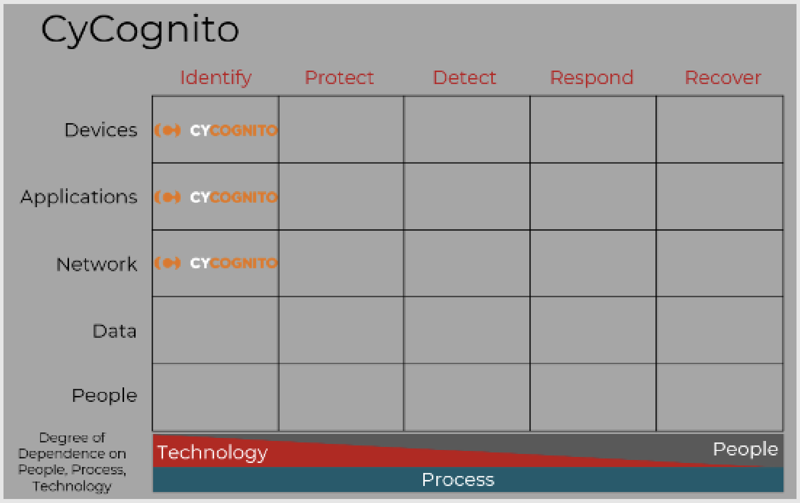

CyCognito

	Identify	Protect	Detect	Respond	Recover
Devices	CYCOGNITO				
Applications	CYCOGNITO				
Network	CYCOGNITO				
Data					
People					

Conclusion

CyCognito is one of the most exciting ASM products currently in this market. The combination of solid technology and user experience provides a solid base to grow and compete from. We can't say enough about what a pleasure the product is to use – the effort put into design and usability is clear. We'd love to see the company grow its validation capabilities in the near future, perhaps adding attack simulations, human testers or both, depending on the path they choose.

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" in order 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:

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

Reviews

Below is a list of product reviews conducted under this category. We recommend reading through the Attack Surface Management Overview in its entirety before digging into individual reviews, but knowing which offerings were evaluated may offer helpful perspective as you do so.

Who We Are

About CyberRisk Alliance

CyberRisk Alliance (CRA) is a business intelligence company serving the high growth, rapidly evolving cybersecurity community with a diversified portfolio of services that inform, educate, build community and inspire an efficient marketplace. Our trusted information leverages a unique network of journalists, analysts and influencers, policymakers and practitioners. CRA's brands include SC Media, Security Weekly, InfoSec World, Cybersecurity Collaboration Forum, our research unit CRA Business Intelligence, and the peer-to-peer CISO membership network, Cybersecurity Collaborative. More information is available at CyberRiskAlliance.com.

About the author

Adrian Sanabria joined CyberRisk Alliance — the parent company of SC Media and Security Weekly — in 2020. He oversees SW Labs, the company's cybersecurity product review and database initiative. Adrian also provides industry commentary for both SC Media and Security Weekly. He brings two decades of industry experience, working as a practitioner, penetration tester, and industry analyst. He spent the last few years as an entrepreneur, challenging norms in sales and marketing for a variety of vendors. Adrian loves to cook, eat, hike, play music and regale his teenagers with stories of what the early days of the Internet were like.