BISHOPFOX

EBOOK

THE WOLF

In Sheep's Clothing

HOW LOW-RISK EXPOSURES BECOME CATALYSTS FOR DESTRUCTIVE ATTACKS

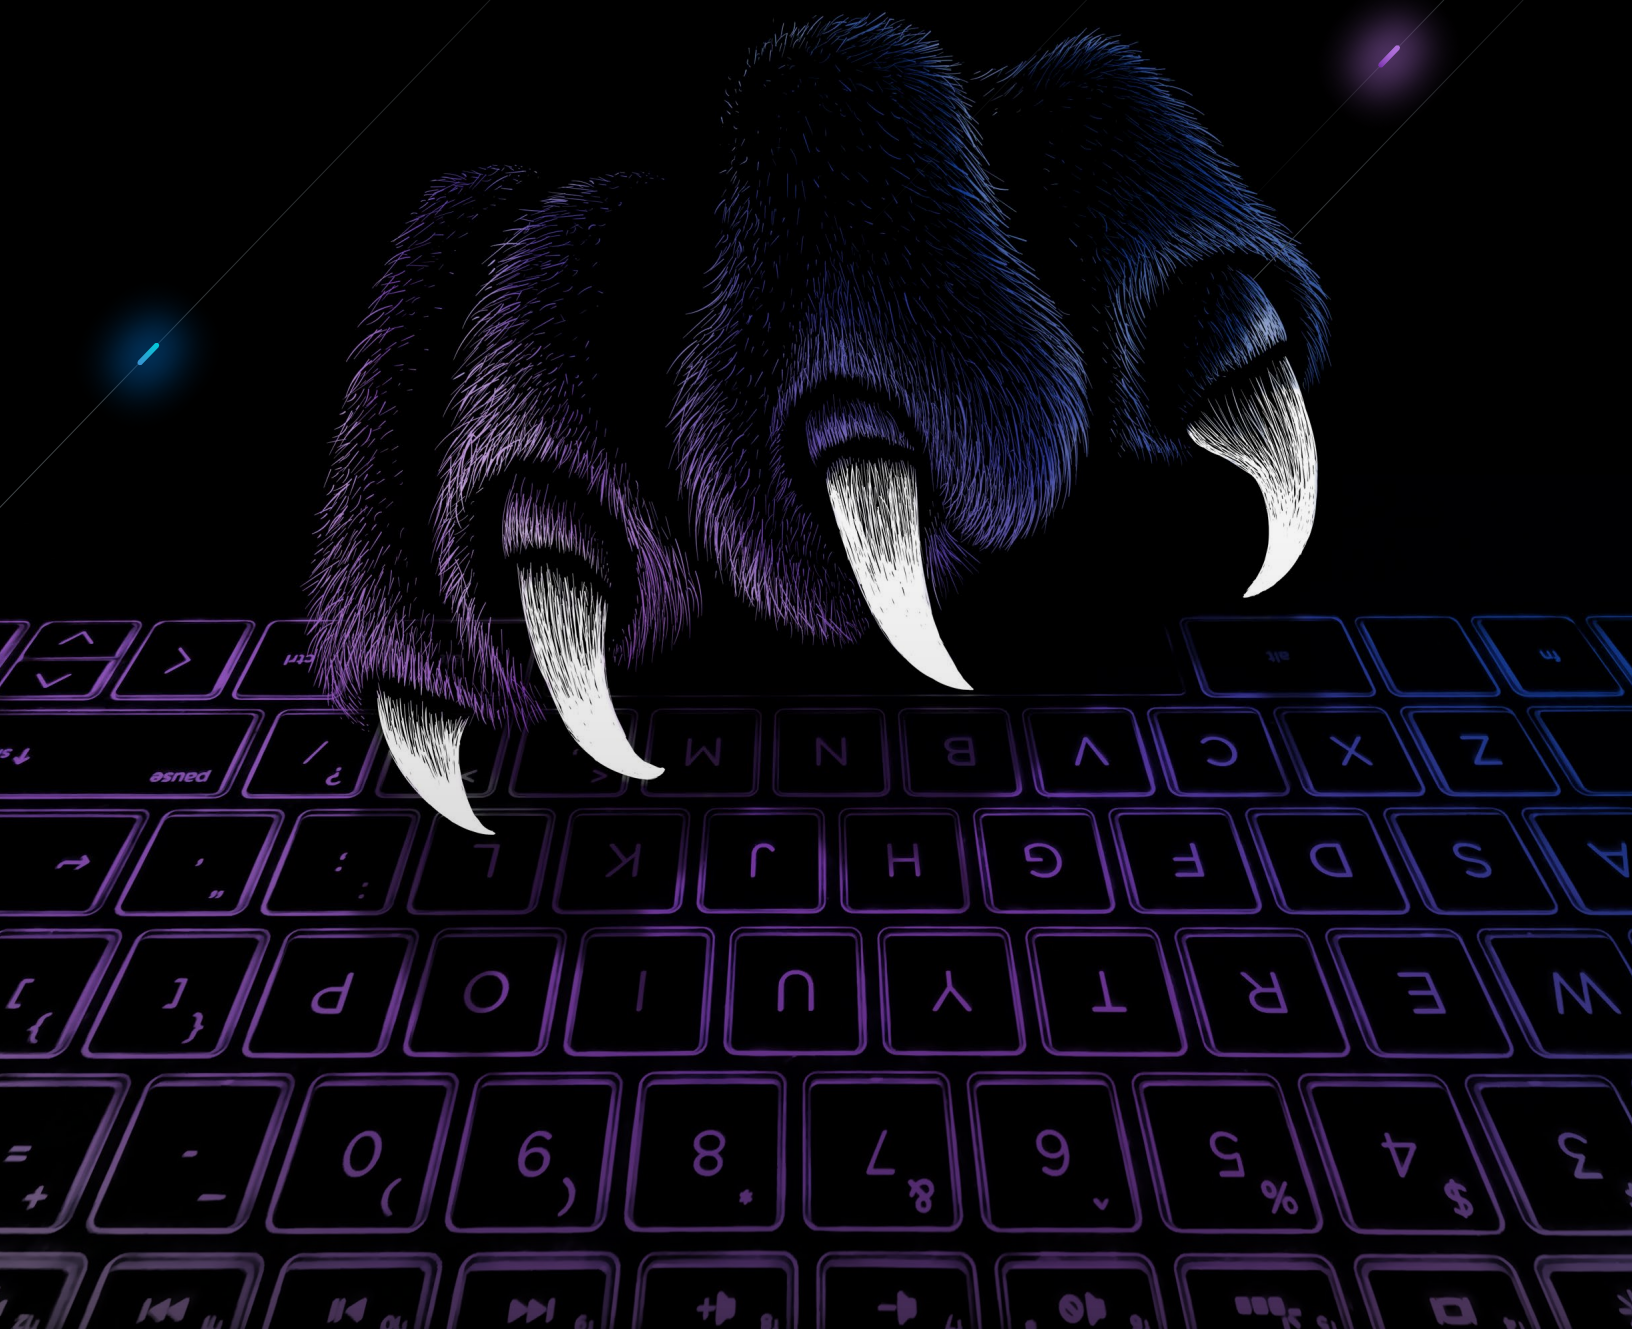

Table of Contents

Introduction

Every time a new CVE is published and gains traction in the press, security teams scramble to determine whether those vulnerabilities pose a threat to their business. The truth is that in many cases, these stories simply become a distraction. The biggest risks are oftentimes vulnerabilities that are much less glamorous, much easier to exploit, and offer more useful footholds for attackers.

Threat actors are like the professionals of any craft. Their level of sophistication, as well as their success rate, are directly tied to their skills, creativity, experience, and effort. While capable of sophisticated attacks, bad actors can save significant time and resources by identifying and capitalizing on often overlooked vulnerabilities that have known exploits and pathways to their primary objective. Think of it from the perspective of a manufacturing plant, why build something from scratch when you can source pre-built parts that accelerate the production of the end-product? It's all about minimizing the time, resources, and effort required to meet your objectives.

The most famous, real-world example of this is EternalBlue. Using the widespread exploit as either an initial compromise vector or method of lateral movement, attackers leveraged EternalBlue for the WannaCry and NotPetya ransomware attacks causing over $1 billion worth of damages in more than 65 countries. More recently, Log4j emerged as an initial vector of compromise that will lead to more complex attack chains. Unfortunately, the extent of these attacks and repercussions won't be known for some time.

While EternalBlue and Log4j are high-profile exposures that garnered widespread attention, there are thousands of exposures categorized as low or medium risk, according to pre-defined severity ratings, that slip through the cracks. However, in the hands of skilled attackers, many of these exposures serve as launching pads or steppingstones to more complex and destructive attacks. The challenge for many organizations is not only identifying these exposures but determining the potential business impact in their unique environment.

The reality is… some vulnerabilities considered 'low risk' may actually be critical in the hands of a skilled attacker.

In this eBook, we'll explore:

- The speed, precision, and covert nature of modern adversaries
- Commonly observed low-risk exposures and how they lead to destructive attacks
- How attack surface management and continuous pen testing can help you proactively uncover, assess, prioritize, and remediate these types of "innocuous" exposures

We'll include examples of exposures found in real-world environments, including a step-by-step view into how ethical hackers exploited them to reach high-value targets.

The Nature of Modern Adversaries

First, let's get to know today's attackers. They are sophisticated, well-resourced, targeted, and always seeking the easiest and fastest route to initial access. They focus on gaining a persistent foothold, moving laterally across the environment to escalate their privilege, and obtaining access to sensitive and/or proprietary data (or whatever other targets they can uncover)... all while remaining undetected.

For adversaries, detection isn't just about being stopped; it's about the value of the information once it's stolen. If your victim is unaware of the compromise, the data can be worth tremendously more. Consider a bank that is unaware of a breach and the damage that could be done before the organization can issue new cards or suspend accounts.

Often, attacks are viewed in terms of pre-attack and post-attack activities. Pre-attack activities include planning, development, and weaponization strategies, while the post-attack phase encompasses the more familiar execution steps. As cybercriminals look to gain initial access, establish persistence, and escalate privilege, more time and effort will be spent on reconnaissance. This means identifying the systems to exploit and obtaining the means by which to exploit them. Unfortunately, this is becoming easier for attackers due to the prevalence of underground markets, forums, peer groups, and more.

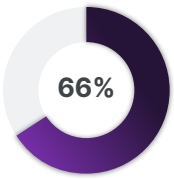

In Nuix's Black Report, 66% of hackers report that new public and private exploits are available every two months or less.

The Beauty of MITRE ATT&CK

The MITRE ATT&CK framework is an invaluable reference for cybersecurity professionals. This extensive knowledge base is a collection of adversary tactics and techniques based on real-world observations of how attacks unfold. Organizations in both the private and public sectors rely on the framework to develop specific threat models and methodologies, as well as ensure effective protection at each stage across the cyber kill chain.

In context of this eBook, it's important to visualize how an attacker would leverage 'low-risk' exposures to not only gain initial access earlier in the MITRE framework, but also to accomplish their objectives.

See the Framework ›

Identifying targets is one piece of the puzzle – exploiting them is another. Unfortunately, once exposures are identified, attackers can exploit with ease. The Nuix Black Report found that 71% of adversaries can identify an exposure and breach a perimeter in less than 10 hours, and once in, 73% can exfiltrate data in half that time. But how often are attackers encountering environments they can't break into? Pretty rarely or never the report found – with 62% of adversaries stating that only 15% of the time or less they can't break into an environment.

ADVERSARIES ARE FINDING AND EXPLOITING EXPOSURES FASTER THAN EVER BEFORE.[1]

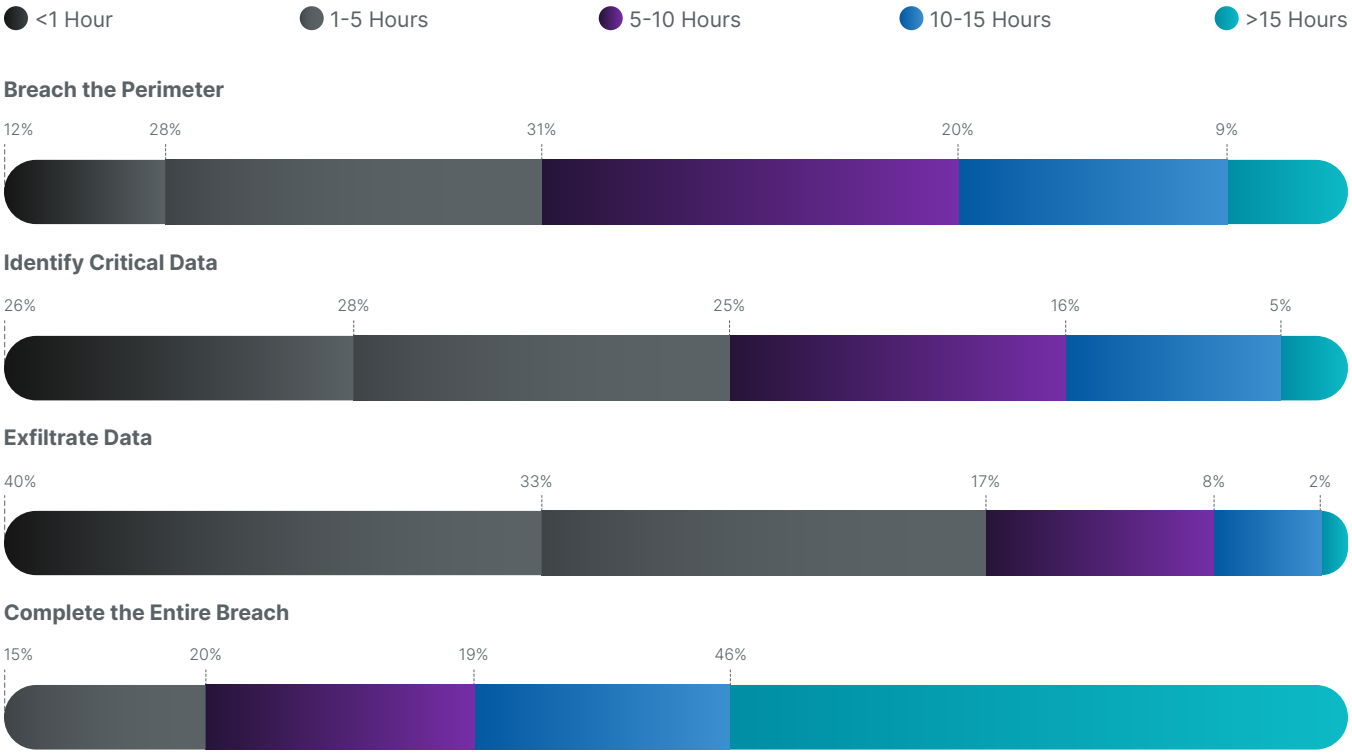

HOW OFTEN ARE ATTACKERS ENCOUNTERING ENVIRONMENTS THEY CAN'T BREAK INTO?[1]

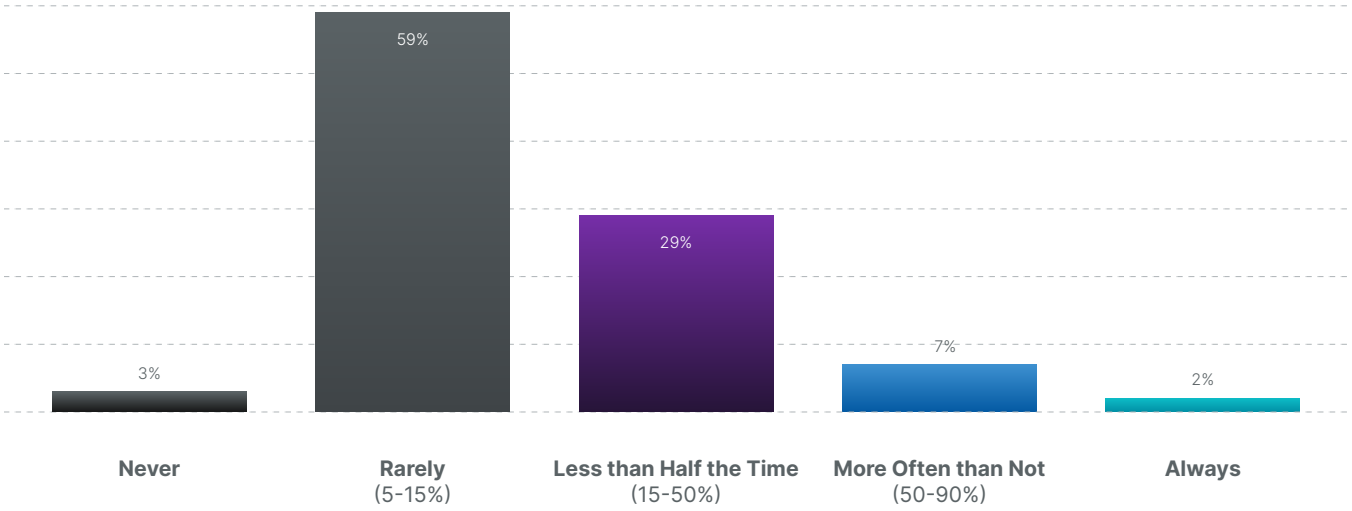

1. Nuix. (2018). The Black Report 2018: Decoding The Minds of Hackers.

Attackers are also increasingly attempting to accomplish their objectives without the use of malware – with **62% of detections being malware-free.** Rather, cybercriminals are concentrating their efforts on legitimate credentials and built-in tools (also known as 'living off the land') in an attempt to evade detection by antivirus products. To further back this new direction for attackers, the **Ponemon Institute** found in 2021 the most common initial attack vector was compromised credentials (20%), following by phishing (17%), and cloud misconfiguration (15%).

Once an attacker is inside an organization's network, an attack converts into a **data breach 18% of the time**. The more vulnerabilities or exposures your organization is susceptible to, the more likely a data breach can occur. Play those odds long enough and the probability of a breach reaches 99% very quickly.

PROBABILITY OF AN INCIDENT RESULTING IN DATA DISCLOSURE.[2]

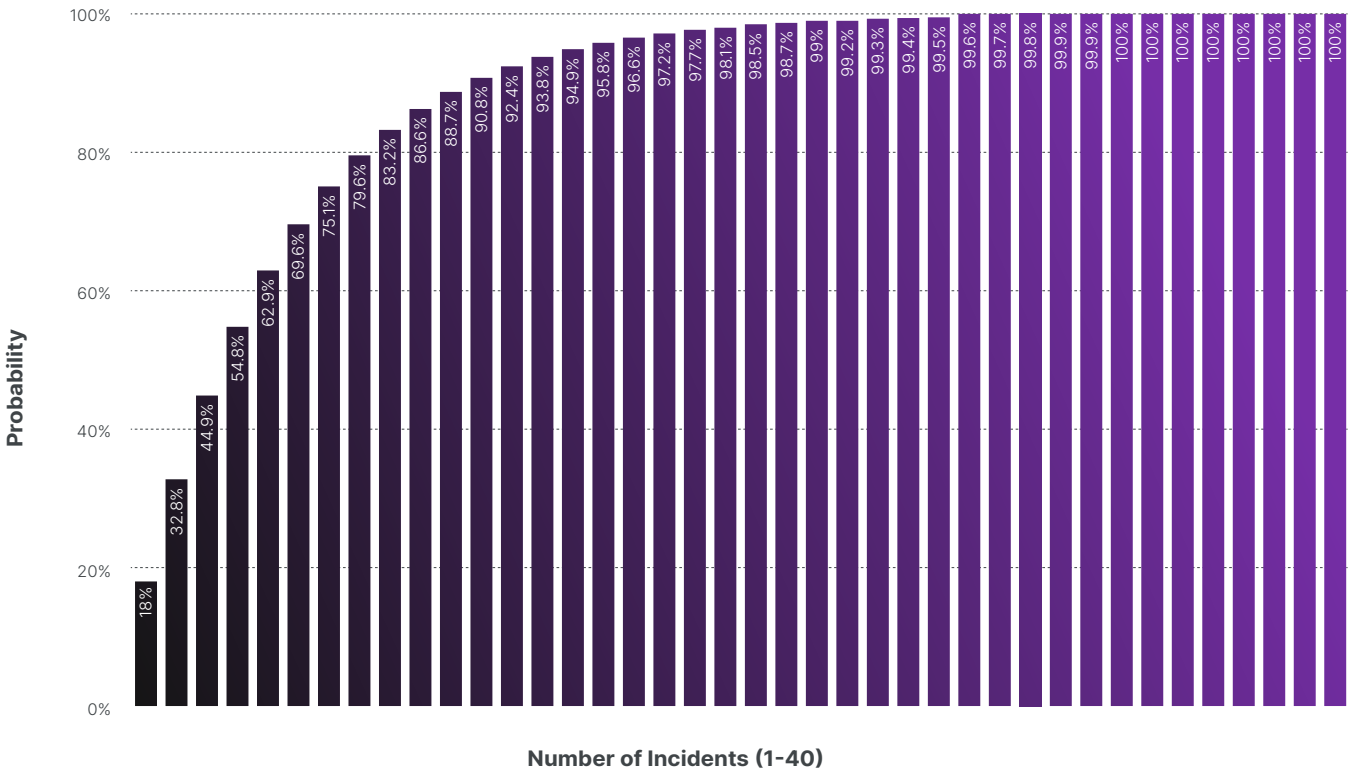

Ponemon also found that the average time to identify a breach in 2021 was 212 days, with an additional 75 days to contain. That's 287 total days, on average, for the lifecycle of a data breach. Adding more fuel to the fire, data breaches with a lifecycle of more than 200 days had an average cost of $4.87 million, compared to $3.61 million when under 200 days.

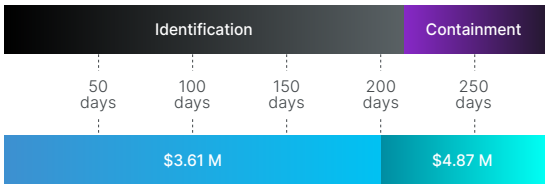

2. Visual inspired from data in Verizon's 2021 Data Breach Investigations Report.

The Five Most Common Missed Exposures

Like many successful criminals, smart attackers are pragmatic. They aren't going to 'burn' a valuable zero-day vulnerability exploit against a victim if a cheap and easy vulnerability does the trick. Imagine a burglar who spends hours trying to pick a door lock when a window is already ajar. The truth is they wouldn't. The same goes for cyberattacks.

So... what are the cyber equivalents of open doors and windows?

In 2021, our Cosmos platform (which combines attack surface management with continuous offensive security testing) identified hundreds of thousands of potential exposures (what we call leads) across our client base. Many of these were categorized as "medium" or "low" severity based on their CVSS scores (< 7) and would typically be ignored by security teams based on their CVSS scores alone — and given the volume of alerts most teams receive every day. However, in the hands of our skilled Adversarial Operations team, these perceived innocuous exposures provided critical steppingstones to more complex attack chains.

The graph below shows the average number of real-world exposure findings based on potential business impact across a twelve-month period for the average Cosmos client. Note: our team rates these findings on a different scale than the CVSS scores; our rankings are based on the actual assessed impact in each specific client environment. It is important to point out that virtually all of these organizations have sophisticated security programs dedicated to vulnerability management, attack surface management, penetration testing, and more. These findings highlight the exposures that slip through the cracks and are real-world exploitable, resulting in compromise of subsequent systems and the potential to cause business-disrupting damage.

AVERAGE PER CLIENT MONTH-TO-MONTH FINDINGS IN COSMOS PLATFORM.

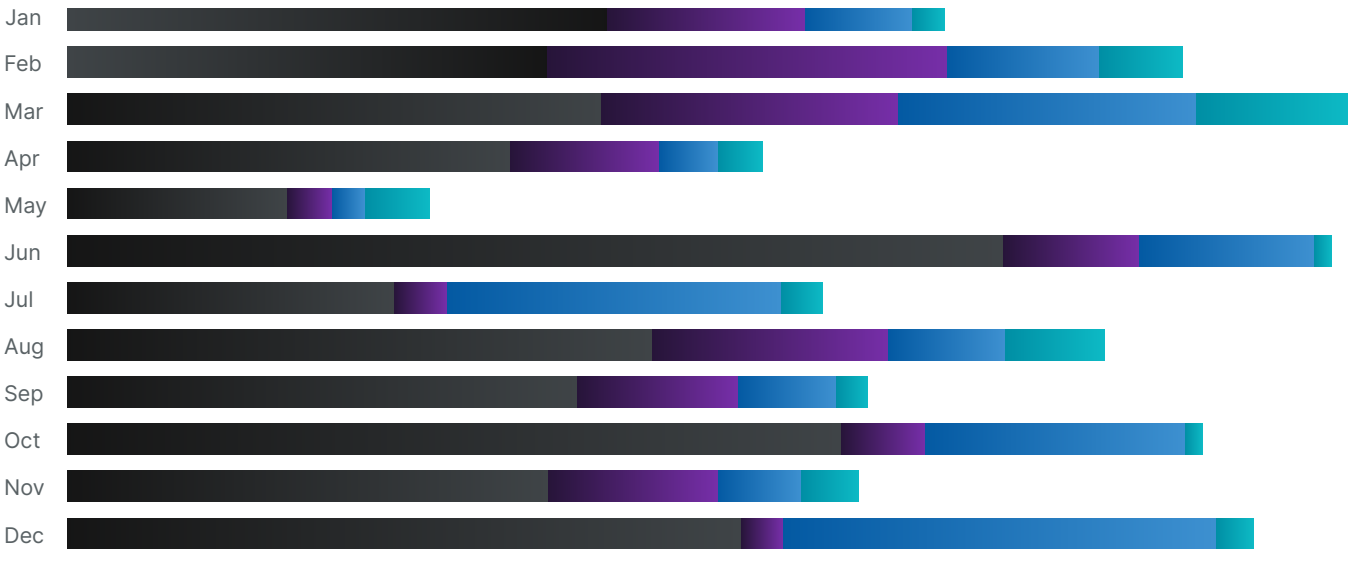

These validated exposures fall into common categories with varying levels of severity associated. It's important to point out that different industries are more susceptible to exploitation and experience varied degrees of business impact.

TYPE OF LOW-RISK VULNERABILITIES FOUND IN COSMOS PLATFORM.

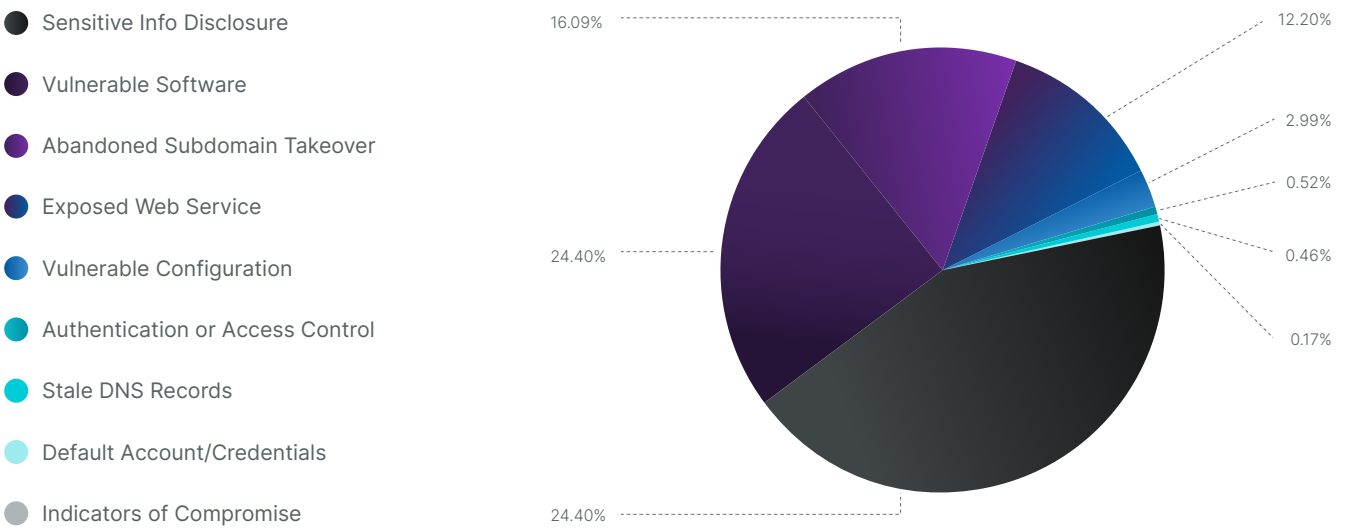

MOST COMMON LOW-RISK EXPOSURES PER INDUSTRY FOUND IN COSMOS PLATFORM.

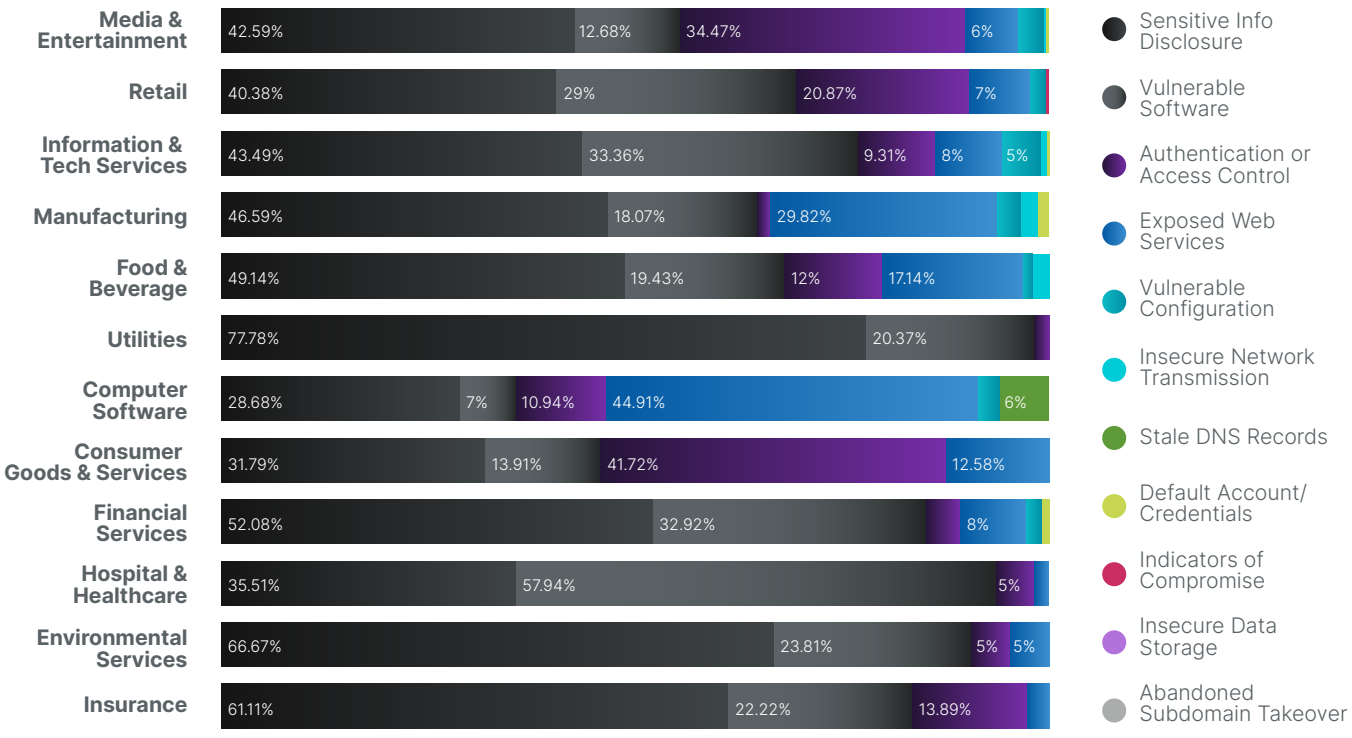

Five Real-World Findings Aligned to Each Exposure Category

COSMOS ADVERSARIAL OPERATIONS TEAM HIGHLIGHTS THE DANGERS OF 'LOW RISK' EXPOSURES

1. Sensitive Information Disclosure

Sensitive information disclosure occurs when private data is exposed to unauthorized parties. This may include financial data, personal privacy information, health records, proprietary information, or other important data. Today's workforce uses many disparate platforms for workflows and productivity. Some are hosted by the organization, while others are purely third-party, cloud-based solutions. Each of these platforms can potentially leak sensitive information – sometimes in unexpected ways. For example, a system administrator may post a script on GitHub or Stack Overflow and inadvertently expose sensitive information about the company's systems.

REAL-WORLD COSMOS FINDING

The Adversarial Operations team searched for shortened URLs related to the target. A number of the goo.gl links redirected to public GitHub gists revealing 39 gists that contained files with VPN credentials. Using a CA certificate and credentials within the files, the AO team was able to authenticate to the VPN. After identifying internal subnets, the AO team was able to find a Prometheus server containing data on approximately 7,000 live targets. To expedite the review of services the assessment team scraped hosts from the Prometheus server and fed them into aquatone, a tool that captures screenshots of web services. The team investigated a large number of services that seemed likely to expose sensitive information. Post-exploitation activities led to the discovery and compromise of services including:

- SonarQube instances, data, and credentials
- 144 Airflow instances and administrative access
- Celery administration panels
- Reldash build information
- Kubernetes secrets
- AWS infrastructure, credentials, and tokens
- Elasticsearch and Kabana code execution, credentials, deployments scripts, SSH keys, and more

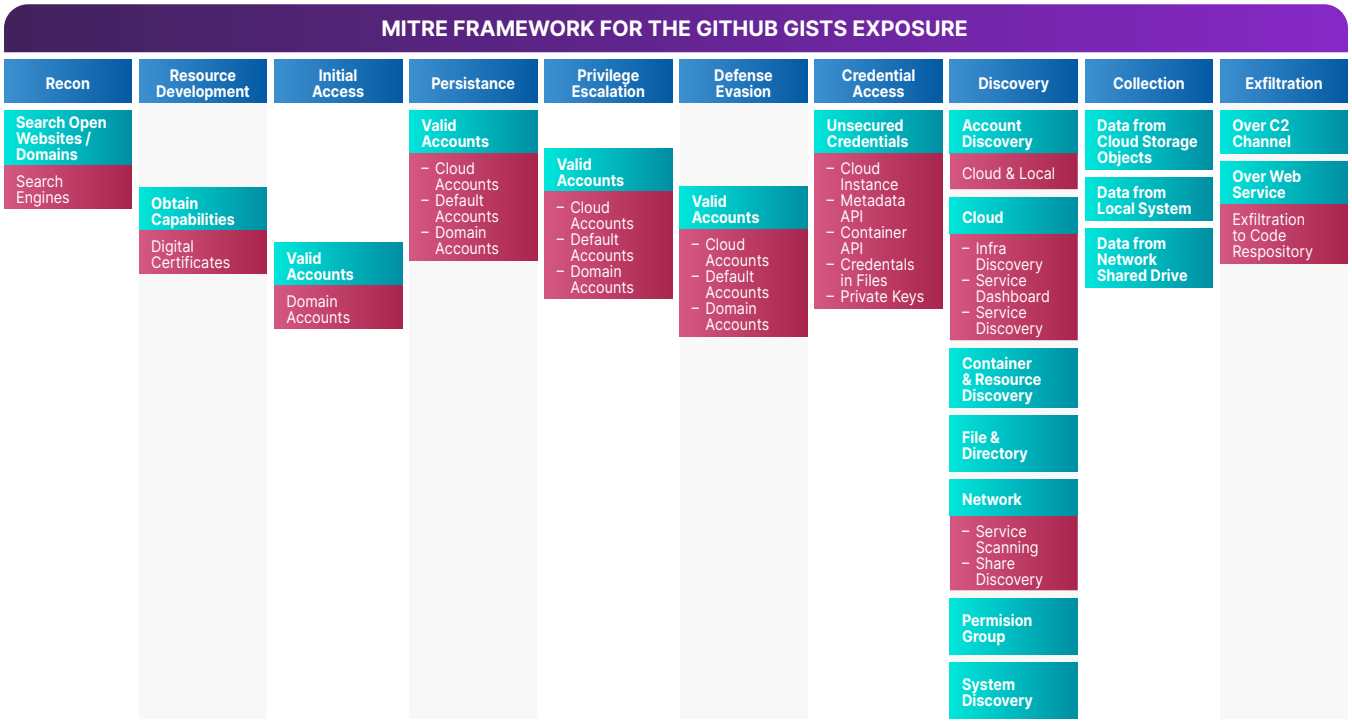

2. Vulnerable Software

Applying patches to every OS and application across an environment is an overwhelming task for every organization, particularly for large and globally distributed ones. Despite this, it's essential to keep OS and application software up to date. **Nearly 60% of breaches** can be traced back to a missing patch that an attacker exploited to gain access.

Unpatched, insecure versions of software can allow attackers to perform arbitrary remote code execution, SQL injection, and other actions that enable them to gain elevated access to the application itself or its supporting infrastructure.

REAL-WORLD COSMOS FINDING

Our assessment team discovered a host running a version of Adobe ColdFusion (a commercial web application development platform) that is vulnerable to XML external entity injection (XXE) via unsafe entity processing within AMF messages (CVE-2015-3269). This vulnerability has a CVSS score of 5.0 and is often overlooked for remediation in favor of higher severity issues. This vulnerability allows an attacker to read arbitrary files on the web server, exposing sensitive information such as configuration files and credentials if exploited.

The impact for this vulnerability is categorized as critical under real-world conditions as it exposed customer data, allowing our assessment team to gain remote code execution with root level permissions, and reveal sensitive information. Specifically, our team accessed data from a total of 926 EC2 instances available within the environment. The XXE vulnerability was also used to retrieve files that contained a password for the ColdFusion admin interface enabling takeover of the mail service within ColdFusion. In addition, post-exploitation activities revealed the following:

- RSA private keys
- Sensitive GitHub and GitLab information
- AWS credentials
- Tanium keys
- SQL server databases with over 90,000 plaintext credentials

MITRE FRAMEWORK FOR THE ADOBE COLDFUSION VULNERABILITY

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Collection
Valid Accounts Default Accounts	Command & Scripting Interpreter – JavaScript – Windows Command Shell	Create Account Domain Account Valid Account Default Account Server Software Component Web Shell	Exploitation for Privilege Escalation Valid Account Cloud Accounts	Valid Account – Cloud Accounts – Domain Accounts	Credentials from Password Stores Security Memory Unsecured Credentials – Credentials in Files – Private Keys	Account Discovery Cloud Account Cloud Infrastructure Discovery Cloud Service Dashboard Cloud Service Discovery File & Directory Listings	Exploitation of Remote Services Use Alternate Authenti-cation Material Pass the Hash	Data from Cloud Storage Objects Data from Information Repository Data from Local File Data from Network Share Drive

3. Exposed Web Service

An insecure exposed web service occurs when a web service used for privileged functionality is exposed to unauthorized parties. This may provide access to administrative functionality, and lead to exposure of configuration data, detailed debugging information, personal privacy information, proprietary information, or other important data.

REAL-WORLD COSMOS FINDING

An exposure of the Adminer web service acted as a steppingstone for access to an Amazon EC2 and S3 bucket containing sensitive information. The Adminer vulnerability (version 4.3.0.1) revealed a method of obtaining database credentials that are stored in configuration files on the server, which led our team to compromise the underlying MySQL server. Once database credentials were obtained, post-exploitation activities led to the discovery and compromise of the following:

- 286 users of the hosted applications
- Default passwords
- Authenticated access to the client's public website
- Ability to alter application configurations
- Database tables for a WordPress installation
- AWS access key from exploited SSRF vulnerability
- AWS EC2 and S3 buckets

MITRE FRAMEWORK FOR THE ADMINER VULNERABILITY

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Collection
Exploit Public-Facing Application	Command & Scripting Interpreter	Valid Accounts	Valid Account	Valid Account	Unsecured Credentials	Account Discovery	Data from Local System
	Python	Default Accounts	Default Accounts	Cloud Accounts	– Cloud Instance – Metadata API	Cloud Accounts	
						Cloud Service Discovery	

4. Vulnerable Configuration

A vulnerable configuration occurs when a security misconfiguration is made at any level in an application stack. These security misconfigurations can result in information disclosures, exposed services, and other vulnerabilities that could allow a malicious user to gain elevated access to the application itself or its supporting infrastructure.

REAL-WORLD COSMOS FINDING

JavaServer Faces (JSF) ViewStates are vulnerable to Java object deserialization when unencrypted. This can be leveraged by an unauthenticated attacker to remotely execute code on the host. Our team identified one host with an unencrypted ViewState being used and created a serialized payload that used the Java URL class to generate a DNS request to their BurpCollaborator server.

Our team was able to exploit the unencrypted ViewState deserialization vulnerability to obtain a reverse shell on the vulnerable host, recover AWS RDS credentials from configuration files, retrieve AWS credentials from the metadata endpoint, and download WAR deployment files containing various secrets. This vulnerable configuration resulted in exposure and enabled post exploitation including the potential ability to exfiltrate unencrypted PII and credit card details from RDS database tables, remote command execution on various AWS hosts via SSM permissions of the recovered AWS credentials, and access to all resources within the environment including EC2, S3, RDS, DynamoDB, IAM, and Secrets Manager.

MITRE FRAMEWORK FOR THE JAVASERVER FACES (JSF) VIEWSTATES VULNERABILITY

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Collection	Command & Control
Exploit Public-Facing Application	**Command & Scripting Interpreter**	**External Remote Services**	**Exploitation for Privilege Escalation**	**Valid Account**	**Unsecured Credentials**	**Account Discovery**	**Exploitation of Cloud Services**	**Data from Cloud Storage Object**	**Encrypted Channel**
Search Engines	Javascript	**Valid Account**	**Valid Account**	Cloud Accounts	– Cloud Instance – Metadata API – Credentials in File	Cloud Account			
		Cloud Accounts	Cloud Accounts			**Cloud Infra Discovery**			
						Cloud Service Dashboard			
						Cloud Service Discovery			

5. Default or Weak Passwords

Once an attacker confirms the type of system they're targeting, they will try the default admin password for that system.... and they start there because it oftentimes works. Weak passwords that are easily guessed or susceptible to brute force attacks also offer easy onramps for attackers.

Default accounts are those accounts that are built into systems, applications, databases, and embedded devices to provide convenient access prior to the configuration of official accounts. These accounts are prime targets for attackers as they are usually publicly documented and are often overlooked during system deployment and hardening. In many cases, the password associated with a default account is the same as the username or is similarly weak. Online databases of these default account credentials are readily available for malicious users to leverage in their attacks.

REAL-WORLD COSMOS FINDING

The assessment team successfully logged into a vulnerable SonarQube web application using default credentials (username 'admin' and password 'admin'). Once logged in, the team gained full control over SonarQube and its assets, including the ability to:

- View and download source code for all projects
- Review all code vulnerabilities identified by SonarQube
- View configuration details including host system information, web server configuration, database integration, and authentication providers
- Manage SonarQube user accounts
- Export application logs
- View installed plugins
- Upload and install custom plugins

Our team successfully compromised the SonarQube host server and proceeded to enumerate the compromised host to identify sensitive information, opportunities to elevate privilege, and additional targets for lateral movement.

MITRE FRAMEWORK FOR THE SONARQUBE VULNERABILITY

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Collection
Valid Accounts Default Accounts	Command & Scripting Interpreter JavaScript	External Remote Services Valid Account Default Account	Exploitation for Privilege Escalation Valid Account Cloud Accounts	Valid Account – Cloud Accounts – Default Accounts	Unsecured Credentials – Cloud Instance Metadata API – Private Keys	Account Discovery Cloud Account Cloud Infrastructure Discovery Cloud Service Dashboard Cloud Service Discovery	Exploitation of Cloud Services	Data from Cloud Storage Object

Eliminating the Risks with Continuous Offensive Security

As found in **a recent survey by ESG**, 67% of organizations say that their attack surface has increased over the past two years. This is no surprise considering how many organizations quickly scrambled to launch and support remote workforces in the wake of worldwide COVID-19 shutdowns. In addition, modern IT environments have become incredibly dynamic due to the proliferation of technologies and migration to the cloud. Unfortunately, traditional security solutions weren't built for the dynamic nature of modern attack surfaces.

PRIMARY REASONS ATTACK SURFACE HAS INCREASED.[5]

5. ESG. (2021). Security Hygiene and Posture Management.

With hundreds of exposures emerging daily, scalable, proactive identification of vulnerabilities that could put an organization at risk is critical. As we've seen, this includes not only issues labeled as 'critical' and 'high risk,' but also those that are seemingly 'low risk.'

By missing exposures and inundating security teams with false positives, conventional approaches are extending the attack window. Armed with the latest tactics and technologies, adversaries are taking advantage, targeting exposures and exploiting vulnerabilities faster than security teams can keep up.

To stay ahead of the curve, organizations must find solutions that can continuously map their perimeter, discover high-risk exposures, and deliver validated findings that enable rapid remediation. We believe successful approaches require a creative combination of technology, automation, and human testing to not only identify high-risk exposures but also assess their impact to speed remediation.

Top Five Benefits of Continuous Security Testing

1. Understand the changes in your attack surface in the face of rapid change.
2. Discover out-of-scope perimeter assets – from third-parties, cloud workloads, etc.
3. Focus on a domain-centric approach that aligns with how customers, partners, and attackers interact with public-facing assets.
4. Protect brand reputation by identifying knock-off sites that imitate legitimate domains.
5. Help teams reduce their attack surface risk by identifying unnecessary services, applications, and exposed data to address.

For those looking to quickly identify and remediate exposures, there are six key areas to take into consideration: attack surface discovery, exposure identification, triage, validation, remediation, and outputs. We recommend downloading our workable guide, "Critical Questions to Ask Offensive Security Providers" to assist you in evaluating how security providers and/or your own security team currently handle vulnerabilities and exposures.

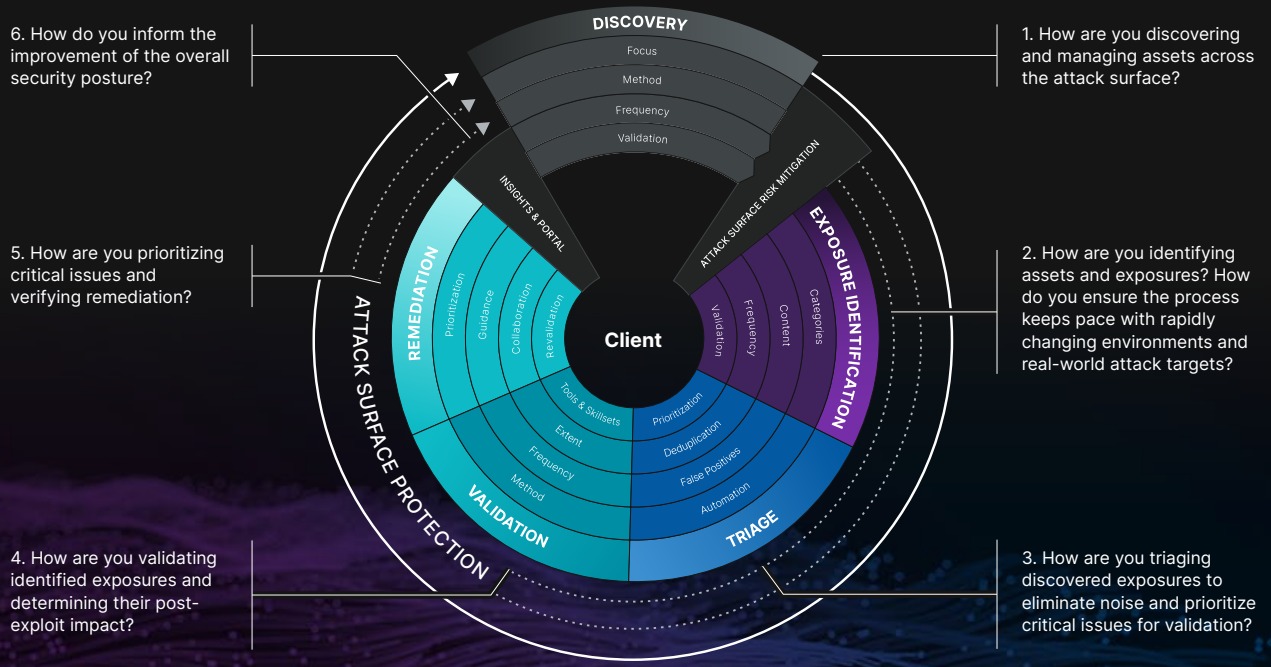

Summary

While vulnerabilities classified as 'low risk' may seem innocuous, they can cause just as much – if not more – damage than high-profile, emerging threats trending in news headlines. These perceived low-risk issues tend to be widespread, and accordingly, attackers are experts in using open-source tools to exploit them.

Several well-known breaches have occurred due to exposures an organization either wasn't aware of or due to a false sense of security. For example:

- 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.**

Each of these examples share the same pattern. Attackers know what works, and the attack path was clear. Meanwhile, defenders are working with tools literally telling them that there are hundreds of thousands of attack paths – and left to ferret out where to focus. But with a continuous offensive security approach that keeps pace with threat actors and accurately identifies business-impacting vulnerabilities, teams can be at the ready when attackers come knocking on the door.

About Bishop Fox

Bishop Fox is recognized as the leading authority in offensive security, providing solutions ranging from continuous penetration testing, red teaming, and attack surface management to product, cloud, and application security assessments.

Over the past 16 years, we've worked with more than 25% of the Fortune 100, 8 of the top 10 global tech companies, and hundreds of other organizations to improve their security. Our award-winning Cosmos platform was named **Best Emerging Technology** in the 2021 SC Media Awards and our offerings are consistently ranked as "world class" in customer experience surveys.

Security isn't just a job to us. We do this because we love it — and because we're committed to the common good. In fact, we have authored 15 open-source tools, shared groundbreaking research, and published more than 50 security advisories in the last 5 years.

Cosmos

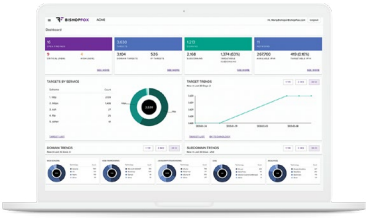

Cosmos proactively defends dynamic attack surfaces by combining advanced technology, automation, and expert-driven testing to continuously identify and remediate high-risk exposures before attackers even know they exist.

Leveraging a proprietary asset discovery and exposure reconnaissance engine, Cosmos continuously discovers and maps your ever-changing attack surface and identifies dangerous vulnerabilities targeted by attackers.

Acting as an extension of your security team, our operators provide deep insights into findings, deliver real-time answers to pressing questions, and conduct on-demand retesting to validate remediation procedures and accelerate the closure of attack windows.

Consulting Services

Internal Penetration Testing

By simulating an attacker who has gained access to the internal network, we locate the most likely vulnerabilities, attack paths, and exploit chains an internal threat actor would leverage to gain access to sensitive data and critical systems.

External Penetration Testing

We simulate an external attacker attempting to exploit your internet-facing networks and applications to help you identify exploitable vulnerabilities and weaknesses in your perimeter.

Red Teaming

We utilize advanced offensive tools and tactics that mimic real-world adversaries to identify exploitable weaknesses in your organization while stress testing your incident responders and their playbooks for handling active, persistent attackers.

CONNECT WITH US

Get started today.

Are you ready to start "defending forward"? Get in touch with our offensive security experts today to explore solutions that meet your unique business needs.

Schedule a Demo **Explore Cosmos**

8240 S. Kyrene Rd. • Tempe, AZ 85284
480.621.8967
hello@bishopfox.com • bishopfox.com