



LOOKING TOWARD THE COSMOS

Making the Case
for Continuous
Offensive Security

Table of Contents

Introduction	3
Defining Continuous Offensive Security	4
Calculating Expected Risk	7
Industry Average Expected Risk and Cost Savings	15
Cost to Replicate a Continuous Offensive Solution	19
Calculating ROI: Three Methods	23
Summary	25
About Bishop Fox	26

Introduction

It's no secret, attack surfaces are increasing at an uncontrollable rate. Fueled by expanding applications, cloud adoption, IoT, and the interconnected nature of modern businesses – security teams face an uphill battle outpacing adversaries to environmental exposures. While automated approaches have rapidly evolved to help security teams achieve scale, they require limited personnel to address an overwhelming number of exposures, many of which often lack real-world exploitability. On the flip side, point-in-time testing uncovers exposures that are real-world exploitable but lacks the scale of continuous discovery. Unfortunately, this imbalance results in a lapse of coverage that is ripe with attacker opportunity.

To give perspective on the scope of this problem, studies indicate that 79% of adversaries can identify and exploit a perimeter exposure of an organization in under ten hours. Furthermore, 54% percent of those adversaries can locate and exfiltrate targeted data in under 15 hours. It's an unfortunate but growing reality that an attacker will find a way in – once they do, they are rarely caught. As a result, organizations have shifted investment to detection and response capabilities. However, this approach is proving too little too late. The most recent **Verizon DBIR report** indicates that 18% percent of all incidents that involve a bypass of security controls result in public data disclosure. The lesson of this story – proactive prevention is the path forward.

While organizations look to close the gap with continuous offensive testing, it is often prohibitive to operationalize with in-house resources. Fortunately, the rise of continuous offensive security solutions has enabled organizations to achieve objectives at a fraction of the cost. However, competing priorities and overlap in existing programs often make it difficult to justify additional investment without quantifiable risk and return that is contextual to an organization's business.



In this eBook:


- We'll cover the factors, inputs, and calculations that are critical to making the business case for continuous offensive testing.
- We introduce our customized return on investment (ROI) calculator that is purposely designed to produce two data points that are critical to justifying spend: cost savings and mitigation of risk associated with a public breach that results in data disclosure.
- Output of the model is intended to draw a direct line from investment to risk mitigation that can be communicated to both technical and non-technical decision makers.



Defining Continuous Offensive Testing

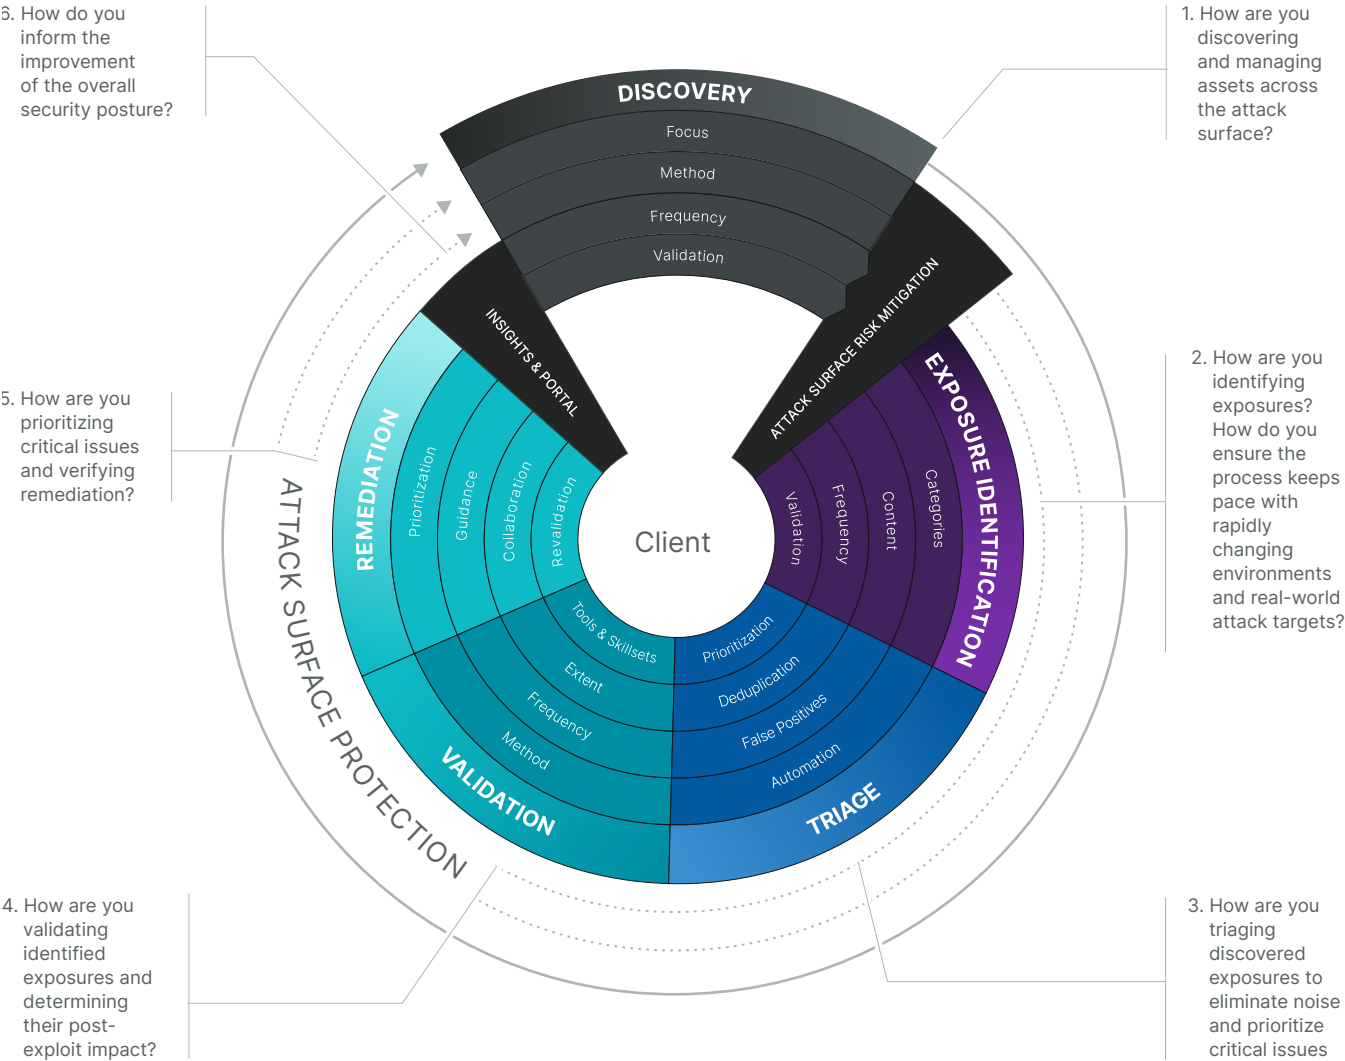
Before we can calculate risk mitigation and ROI for continuous offensive testing, we must first define what continuous offensive testing is. In recent years there's been a renaissance of offensive technologies that look to revolutionize the prevention market. This new wave is a natural rebalancing of the prevention vs. detection and response equation. Prior to this emergence, managed detection and response (MDR) solutions and underlying technologies (EDR, XDR, NDR, etc.) saw an influx of investment as an answer to a growing contingent of advanced adversaries. While MDR proved to be successful to a degree, organizations have begun to complement MDR with a renewed focus on proactive discovery and remediation. Evidence of this shift can be seen in [Gartner's 2021 Hype Cycle for Security Operations](#).

HYPE CYCLE FOR SECURITY OPERATIONS¹



Notice MDR is on the trough of disillusionment, while offensive security represents most solutions in the innovation trigger and peak of inflated expectations categories. While there are a multitude of new and existing offensive solutions that extend beyond the Gartner Hype Cycle, they all share a common goal in helping security teams proactively discover the attack surface and test for weaknesses. The ongoing nature of this is commonly referred to as continuous offensive testing.

6 CRITICAL CRITERIA TO ACHIEVING CONTINUOUS OFFENSIVE TESTING



If we look at some of the emerging and legacy offensive solutions in the marketplace today, you will notice significant gaps in achieving the outcomes of a continuous offensive security solution. While automated solutions help organizations achieve scale, they often produce false positives with missed exposures that place tremendous pressure on already overburdened security teams. On the flip side, point-in-time solutions can pinpoint exploitable exposures; however, they are difficult and costly to operationalize on a continuous basis.

OFFENSIVE SOLUTIONS IN THE MARKETPLACE

An ideal continuous offensive security solution can achieve the desired outcomes by combining the right mix of technology, automation, and human testing to ultimately enable remediation of business impacting exposures at the scale of modern business demands. For the purposes of calculating risk mitigation in the following sections, we will assume the solution that is under consideration can successfully achieve the outcomes outlined in the table below.

	CONTINUOUS OFFENSIVE SECURITY	ASM	BUG BOUNTY	VULN. SCANNERS	MANAGED SCANNING	DAST	BAS
COMPLETE ATTACK SURFACE COVERAGE	✓	✓	Incomplete (Customer Defined)	Incomplete (Customer Defined)	Incomplete (Customer Defined)	Applications Only	Incomplete
CONTINUOUS ATTACK SURFACE DISCOVERY	✓	✓	✗	✗	✗	✗	✗
COVERS EMERGING THREATS	✓	✗	✓	Known CVEs Only	Known CVEs Only	Limited to Applications	✗
CONTINUOUSLY DISCOVERS EXPOSURES	✓	✗	✓	Disruptive & Intermittent	Disruptive & Intermittent	Slow & Intermittent	✗
DEDICATES HUMAN TESTING	✓	✗	Mixed, Unreliable Team, Inconsistent	✗	✗	✗	✗
CONDUCTS POST-EXPLOITATION	✓	✗	✗	✗	✗	✗	✗
DETERMINES IMPACT OF EXPOSURES	✓	✗	✗	✗	✗	✗	✗
ELIMINATES FALSE POSITIVES	✓	✗	✓	High False Positive	Low False Positive	High False Positive	Low False Positive
ENABLES DIRECT INTERACTION WITH TESTERS	✓	✗	✓	✗	✗	✗	✗
PRIORITIZES EXPOSURES AND REMEDIATION	✓	✗	✗	Generic / No Context	Generic / No Context	Generic / No Context	Generic / No Context
VALIDATES REMEDiation	✓	✗	✗	Rescan Only	Rescan Only	✗	✗

CALCULATING EXPECTED RISK

Every potential security investment has two key data points to justify spend: cost savings and expected risk mitigation. It's important to note that expected risk is not the same as potential risk.

For example, when an insurance company calculates the risk of a young driver with a poor driving history, they know the driver is likely to have an at-fault accident within a certain time frame. Using a historical propensity model, the insurance company will look at drivers with the same attributes and formulate a time frame and expected payout for an accident. Let's say that the average time frame is three years at a payout of \$36,000. An insurance company must account for this expected risk on a monthly or yearly basis until the event happens. In this case, they would set aside \$1,000 dollars per month (\$36,000 over 3 years). Of course, the premiums for the driver would be very high in this situation. That's also why insurance companies encourage defensive driving courses or encourage the use of measurement devices that identify safe or dangerous habits to predict expected risk more accurately. In cybersecurity terms, calculating expected risk and the potential to mitigate is similar if you have the relevant inputs.

Historical Quantity of Business Impacting Exposures Discovered Yearly

Unfortunately, there are no vendor-agnostic studies that cover the number of exposures discovered by continuous offensive testing solutions. However, we have extensive data from our **Cosmos solution** that has discovered thousands of public-facing exploitable exposures and confirmed their post-exploitation impact with human testing. The below tables show the average findings based on severity across a twelve-month period across each industry. 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.

FIVE INPUTS NEEDED TO CALCULATE EXPECTED RISK FROM PUBLIC-FACING EXPOSURES:


- Historical quantity of potential business impacting exposures discovered yearly
- Historical percentage of breaches that start externally
- Historical conversion percentage from incident to data disclosure
- Historical cost per record potentially disclosed
- Potential number of records lost in a data breach

Keep in mind the more contextual these values are to your organization's size, industry, and other attributes, the greater the accuracy of calculating expected risk. In the following sections, we will build a table for you to use with the following formula:

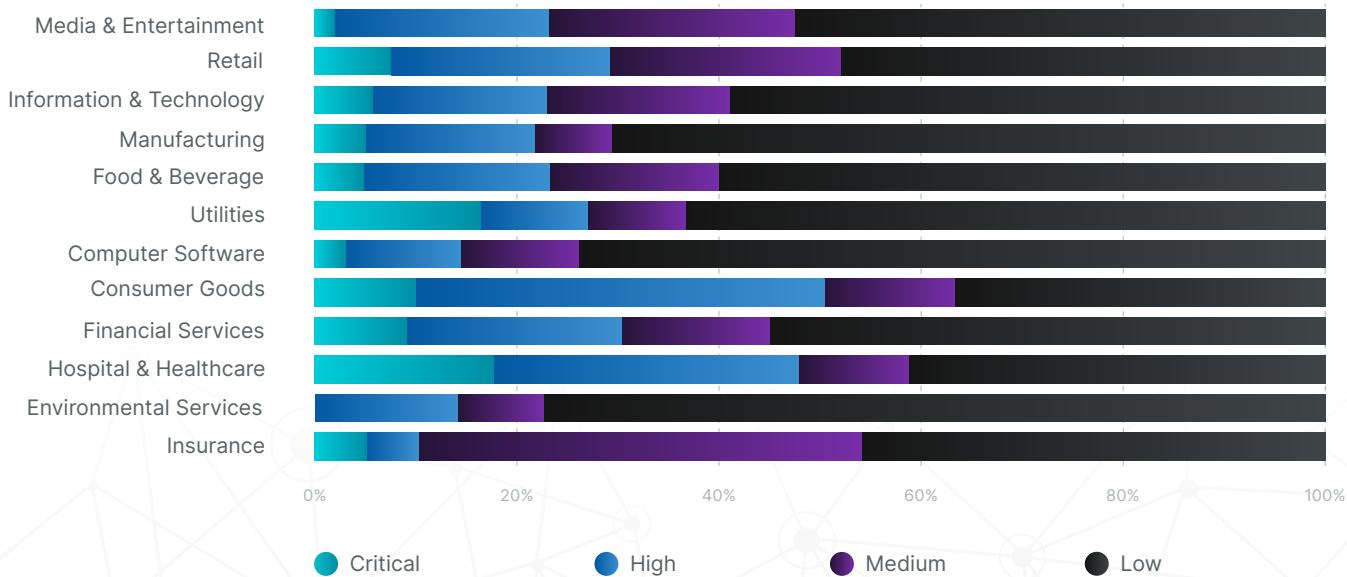
{1-[1-(Probability of an incident to data disclosure * Probability incident originated externally)]^# of Incidents} * Cost per record disclosed * Projected records disclosed

It is important to point out that virtually all the organizations under Cosmos’ protection 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.**

QUANTITY & IMPACT OF INDUSTRY FINDINGS IN COSMOS PLATFORM




SEVERITY DISTRIBUTION AS A PERCENTAGE OF INDUSTRY FINDINGS IN COSMOS PLATFORM

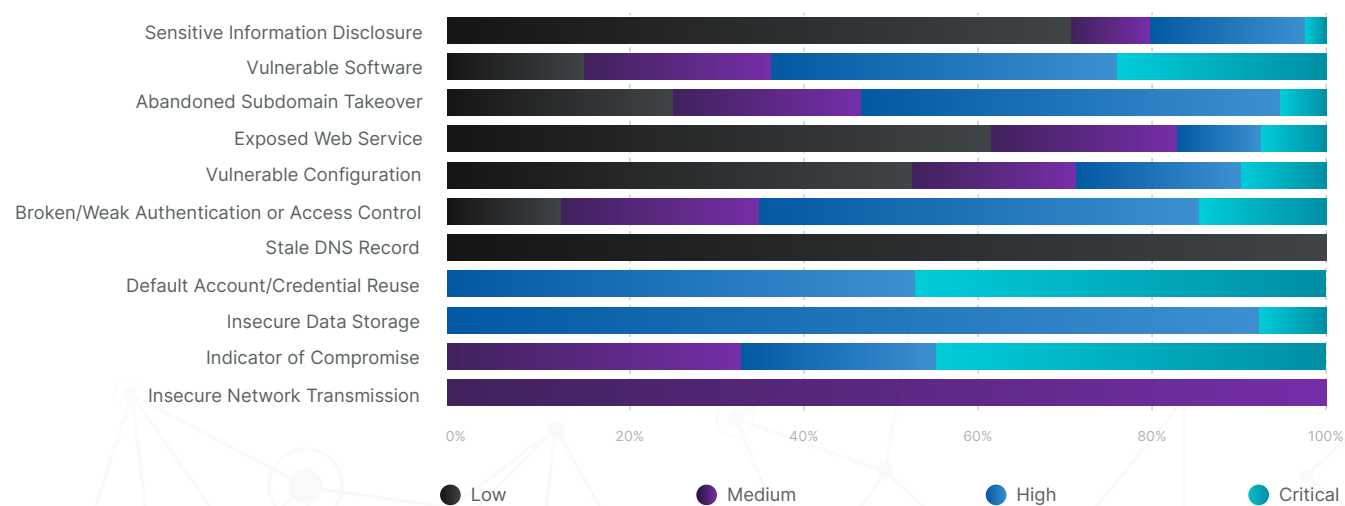


The following tables provide a view across exposure types and severity distribution. We provide this level of detail to help you make an educated decision on the quantity of exposures you feel is applicable to your unique environment. For the purposes of our risk calculation examples, we will only use critical findings due to the business impact they present in post-exploitation activities.

INDUSTRY FINDINGS BASED ON EXPOSURE CATEGORIZATION IN COSMOS PLATFORM



EXPOSURE FINDINGS BASED ON SEVERITY CATEGORIZATION IN COSMOS PLATFORM



In context of our expected risk calculation, **industry averages for critical findings** are as follows (we've included high findings for information purposes):

	CRITICAL	HIGH
MEDIA AND ENTERTAINMENT	3	30
RETAIL	8	20
INFORMATION AND TECHNOLOGY	5	14
MANUFACTURING	3	10
FOOD AND BEVERAGE	3	12
UTILITIES	9	6
COMPUTER SOFTWARE	2	6
CONSUMER GOODS AND SERVICES	4	15
FINANCIAL SERVICES	2	5
HOSPITAL AND HEALTHCARE	3	6
INDUSTRY AVERAGE	4	12

In relation to our expected risk formula, we now have the value for number of incidents.

$\{1-[1-(\text{Probability of an incident to data disclosure} * \text{Probability incident originated externally})]^{\# \text{ of Incidents}}\} * \text{Cost per record disclosed} * \text{Projected records disclosed}$

Percentage of Breaches that Start Externally

Moving forward in our formula, we need to determine what the probability of a potential data breach is from an externally originating threat actor. The **Verizon DBIR report** tracks this information on a yearly basis. While most reports lack a sample set that is statistically significant, the DBIR report tracks over 21,000 incidents across 20+ industries giving us a high degree of confidence at an industry level. For expected risk calculation purposes, the table to the right provides the historical probability that an incident originates from an external threat actor.

Interestingly, the source of security incidents and data breaches are quite different from industry to industry. However, in relation to the type and value of information, distributions make logical sense. For example, financial organizations are often plagued by insider threats whether it be a disgruntled employee looking for financial gain or inadvertent mistakes that lead to data disclosure. On the other hand, utility companies are primarily attacked by outsiders looking to disrupt critical operations via ransomware or other financially motivated means. In relation to our expected risk formula, we now have the value for probability of incidents originating externally.

MEDIA AND ENTERTAINMENT	70%
RETAIL	84%
INFORMATION AND TECHNOLOGY	66%
MANUFACTURING	82%
FOOD AND BEVERAGE	90%
UTILITIES	98%
COMPUTER SOFTWARE	66%
CONSUMER GOODS AND SERVICES	74%
FINANCIAL SERVICES	56%
HOSPITAL AND HEALTHCARE	61%
INDUSTRY AVERAGE	78%

$\{1-[1-(\text{Probability of an incident to data disclosure} * \text{Probability incident originated externally})]^{\# \text{ of Incidents}}\} * \text{Cost per record disclosed} * \text{Projected records disclosed}$

Conversion Percentage from Incident to Data Disclosure

Let's get a common misconception out of the way. Just because a threat actor gains access to an environment **DOES NOT** mean it converts into a data disclosure. Despite the FUD (fear, uncertainty, and doubt) that exists in the marketplace, most incidents do not result in compromise and exfiltration of sensitive data. In fact, historical measures indicate it occurs in 18% of total incidents. For our expected risk calculation, we need the conversion probability of an incident to data disclosure. Again, we will reference the **Verizon DBIR report** by simply dividing the number of incidents by the number of data breaches applicable to each industry. The table to the right represents the values we will use for our formula:

Interestingly, the conversion percentages vary greatly from industry to industry. While media and entertainment had the highest number of exposures from our Cosmos findings, the conversion percentage is the lowest among industries. In contrast, healthcare had the lowest number of discovered exposures but the highest conversion to data disclosure rate. All these values will have significant implications on expected risk in our final calculations. In relation to our overall expected risk formula, we now have the value for probability of incidents converting to data disclosure.

In relation to our expected risk formula, we now have the value for probability of incidents originating externally.

$\{1-[1-(\text{Probability of an incident to data disclosure} * \text{Probability incident originated externally})]^{\# \text{ of Incidents}}\} * \text{Cost per record disclosed} * \text{Projected records disclosed}$

MEDIA AND ENTERTAINMENT	1.5%
RETAIL	23%
INFORMATION AND TECHNOLOGY	13%
MANUFACTURING	46%
FOOD AND BEVERAGE	58%
UTILITIES	42%
COMPUTER SOFTWARE	13%
CONSUMER GOODS AND SERVICES	33%
FINANCIAL SERVICES	65%
HOSPITAL AND HEALTHCARE	72%
INDUSTRY AVERAGE	18%

Cost Per Record Disclosed

When a data breach occurs, it can have far-reaching financial implications. Forensic investigation, downtime, lost business, regulatory fines, class action lawsuits, call centers, and victim identity protection are only a sample of the costs organizations will face once a data breach is publicized. While it's relatively easy to place a value on some of these costs, the long-term implications from reputational damage and abnormal client churn are difficult to quantify. Fortunately, one of the most well-known and widely referenced studies, **Ponemon's Cost of a Data Breach Study**, tracks the cost of a data breach across four short- and long-term categories: detection and escalation, lost business, notification, and ex-post response. In this study, Ponemon analyzes thousands of breaches across a variety of industries calculating the cost of a breach in relation to per record disclosed. For our risk calculation purposes, we will use the following values in the report.

In relation to our overall expected risk formula, we now have the value for cost per record disclosed.

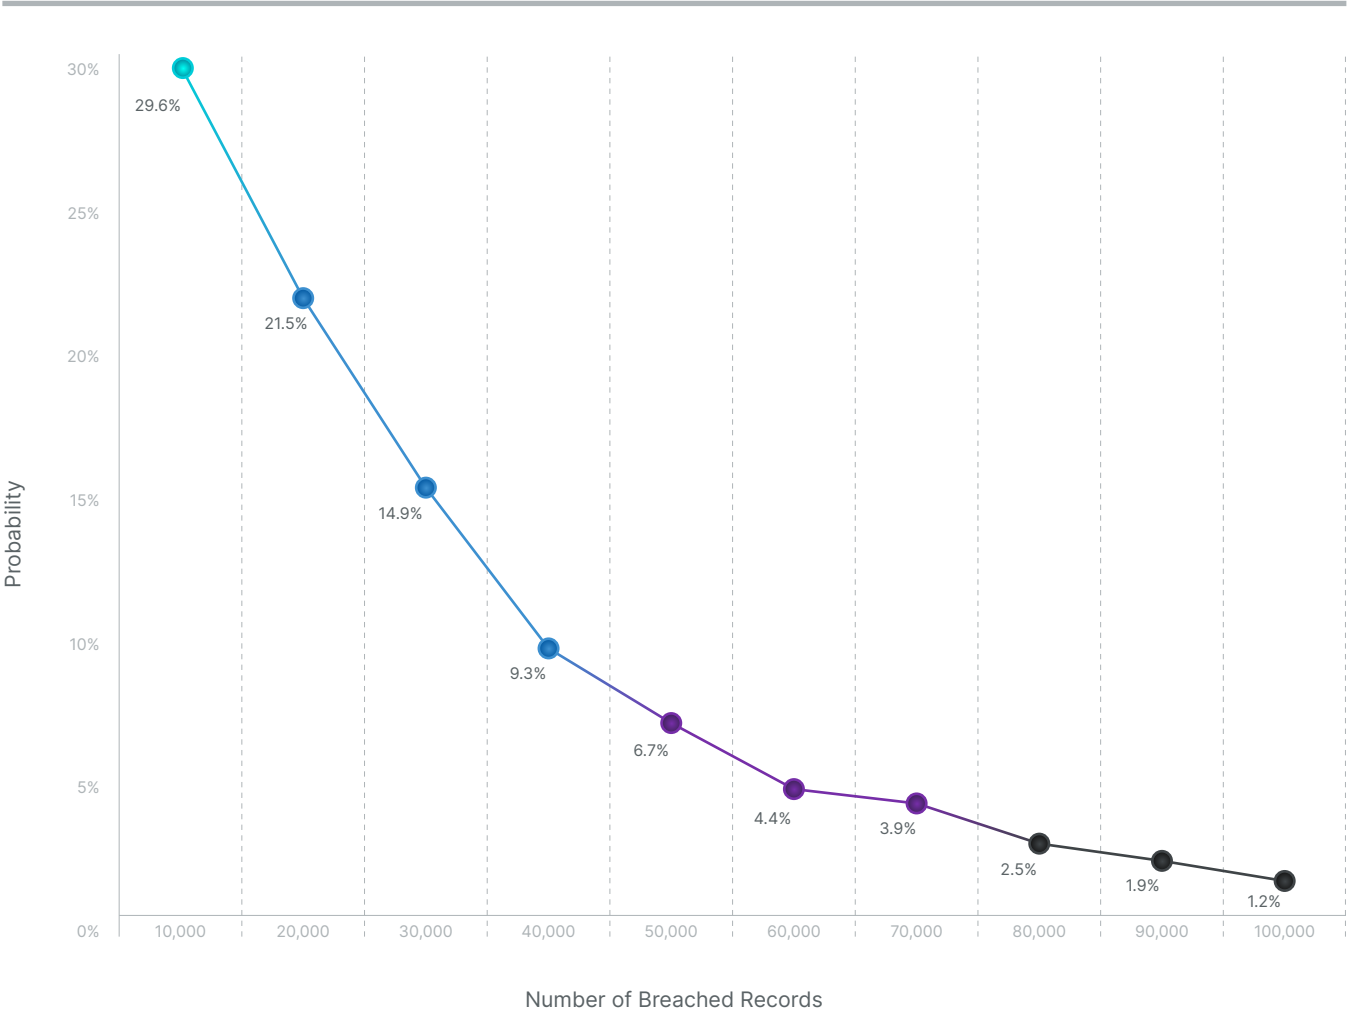
$\{1-[1-(\text{Probability of an incident to data disclosure} * \text{Probability incident originated externally})]^{\# \text{ of Incidents}}\} * \text{Cost per record disclosed} * \text{Projected records disclosed}$

MEDIA AND ENTERTAINMENT	\$123
RETAIL	\$119
INFORMATION AND TECHNOLOGY	\$183
MANUFACTURING	\$160
FOOD AND BEVERAGE	\$123
UTILITIES	\$165
COMPUTER SOFTWARE	\$183
CONSUMER GOODS AND SERVICES	\$131
FINANCIAL SERVICES	\$210
HOSPITAL AND HEALTHCARE	\$429

Projected Records Disclosed

In our expected risk formula, you will notice we need a value for projected records disclosed. Given that no organization's environment, architecture, and data are structured the same, the number of records at risk in a data disclosure event is completely contextual to the organization. For the purposes of overall guidance from an industry comparison perspective, we can estimate value based on propensities from the [Ponemon Cost of a Data Breach Study](#). The graph below represents the probability of a data breach by number of records lost.

PROBABILITY OF DATA BREACH BY NUMBER OF LOST RECORDS



According to the study, the average number of records disclosed is 25,575. In our expected risk calculation, we leverage this number as a guiding point for our expected risk spectrum.

An Important Caveat: Discovery and Exploitation

Before we calculate expected risk, there is an important assumption that we must acknowledge. Our expected-risk formula assumes that critical severity exposures will be identified and exploited by an attacker within a 12-month period. This assumption is important to note as not all exposures will be identified and exploited; some exposures may sit for months, years, or never be located by an adversary. Unfortunately, there is minimal information in the market on the probability of an attacker finding and exploiting public-facing exposures. There have been a few studies by MDR providers; however, the data is applicable to only small and medium businesses and associated with exploitation of physical locations vs. hosted environments. Unfortunately, leveraging this data would likely prove inaccurate for enterprise or highly distributed organizations.

Bringing It All Together

We now have all inputs for our expected risk calculation formula. The graph below is meant to serve as a reference for calculating expected risk in the following examples.

INPUTS FOR EXPECTED RISK CALCULATIONS

	NUMBER OF CRITICAL INCIDENTS*	STARTS EXTERNALLY	CONVERTS TO DISCLOSURE	COST PER RECORD
MEDIA AND ENTERTAINMENT	3	70%	1.5%	\$123
RETAIL	8	84%	23%	\$119
INFORMATION AND TECHNOLOGY	5	66%	13%	\$183
MANUFACTURING	3	82%	46%	\$160
FOOD AND BEVERAGE	3	90%	58%	\$123
UTILITIES	9	98%	42%	\$165
COMPUTER SOFTWARE	2	66%	13%	\$183
CONSUMER GOODS AND SERVICES	4	74%	33%	\$131
FINANCIAL SERVICES	2	56%	65%	\$210
HOSPITAL AND HEALTHCARE	3	61%	72%	\$429
INDUSTRY AVERAGE	4	78%	18%	\$161

*Only critical findings are used in our calculations for number of incidents. If utilizing both critical and high findings makes more sense to gauge risk for your organization, simply add the values found on Page 10.

INDUSTRY AVERAGE: EXPECTED RISK

Let's put the calculator into action. For the following industry-agnostic example, we use industry average data to demonstrate how expected risk is calculated.


THE FOLLOWING VALUES ARE APPLIED TO OUR EXPECTED RISK FORMULA

- 4 critical incidents – Cosmos 2021 Findings Data
- \$161 per record disclosed (breached) – [Ponemon Cost of a Data Breach Report](#)
- 18% incident-to-data disclosure conversion percentage – [2021 Verizon DBIR Report](#)
- 78% of all breaches start externally – [2021 Verizon DBIR Report](#)
- Potential records lost in a data breach. We do not assign a specific value for records disclosed; this value is represented on our X axis as a spectrum for organizations to plot against their estimated number of records that could be at risk from the relevant business impacting exposure(s).

$\{1-[1-(\text{Probability of an incident to data disclosure} * \text{Probability incident originated externally})]^{\# \text{ of Incidents}}\} * \text{Cost per record disclosed} * \text{Projected records disclosed}$

$\{1-[1-(0.18 * 0.78)]^4\} * \$161 * (\text{Projected records disclosed} - \text{Client estimate based on spectrum below})$

INDUSTRY AVERAGE: EXPECTED YEARLY RISK




This leads to the industry average of expected yearly risk in a scenario where continuous offensive security is not operationalized, and exposures that are critical in nature are initially identified and exploited. Based on the above formula, the expected financial risk spectrum ranges from \$1.9M for 25,575 lost records to \$3.3M at the industry average intersection of 45,000 records (intersection of cost and historical propensity). The gray line in the graph represents the average distribution and propensity to lose relevant records based on industry statistics which we covered in the average records disclosed section previously.

While our expected risk is calculated to be \$1.9M - \$3.3M, potential risk if a breach occurs is tremendously higher with a value of \$4.1M for 25,575 records disclosed and \$7.25M for 45,000 records disclosed. Keep in mind that expected risk must be accounted for on a yearly basis if an organization maintains the status quo.

Think of this in the context of rolling a pair of dice. While you may avoid rolling a seven in a game of craps for a limited time frame – on a long enough timeline, the probability approaches a near certainty. Expected risk tells a similar story. The \$1.9M - \$3.3M is what must be accounted for on a yearly basis based on the assumption attackers have found four critical exposures and will exploit them. To paint a scarier picture, if you include the 16 critical and high findings from the Cosmos industry average data, expected risk increases to \$3.7M - \$6.6M. In the context of our dice analogy, the more business impacting exposures, the greater the probability of rolling a seven. The graph below represents this visual, which can be applied to any industry.

LONG-TERM PROBABILITY OF EXPECTED RISK



A Deeper Look at Expected Risk for Individual Industries

To drill further into industry-specific expected risk, it is critical to look at the percentage of breaches started externally and incident conversion rates to data disclosure. In the chart below, manufacturing and food and beverage have high percentages of breaches that start externally (82% and 90% respectively), in addition to high rates of incident conversion to data disclosure (46% and 58% respectively). From an expected risk perspective, this is a dangerous combination that puts them at greater expected risk.

Another perilous combination is that of financial services and healthcare. While these industries experience the lowest percentages for externally triggered incidents, they have the highest data incident conversion to disclosure rates with the highest average cost for data breaches across all industries in 2021 – **with healthcare taking the No.1 title for the eleventh year in a row.**

These examples are key to understanding the expected risk context behind the variances of an individual industry for a potential data breach. These values will be integral in our ROI calculation for business justification.


EXPECTED RISK HEAT MAP

	NUMBER OF INCIDENTS*	STARTS EXTERNALLY	CONVERTS TO DISCLOSURE	COST PER RECORD
MEDIA AND ENTERTAINMENT	3	70%	1.5%	\$123
RETAIL	8	84%	23%	\$119
INFORMATION AND TECHNOLOGY	5	66%	13%	\$183
MANUFACTURING	3	82%	46%	\$160
FOOD AND BEVERAGE	3	90%	58%	\$123
UTILITIES	9	98%	42%	\$165
COMPUTER SOFTWARE	2	66%	13%	\$183
CONSUMER GOODS AND SERVICES	4	74%	33%	\$131
FINANCIAL SERVICES	2	56%	65%	\$210
HOSPITAL AND HEALTHCARE	3	61%	72%	\$429
INDUSTRY AVERAGE	4	78%	18%	\$161

*Only critical findings are used in our calculations for number of incidents. If utilizing both critical and high findings makes more sense to gauge risk for your organization, simply add the values found on Page 10.

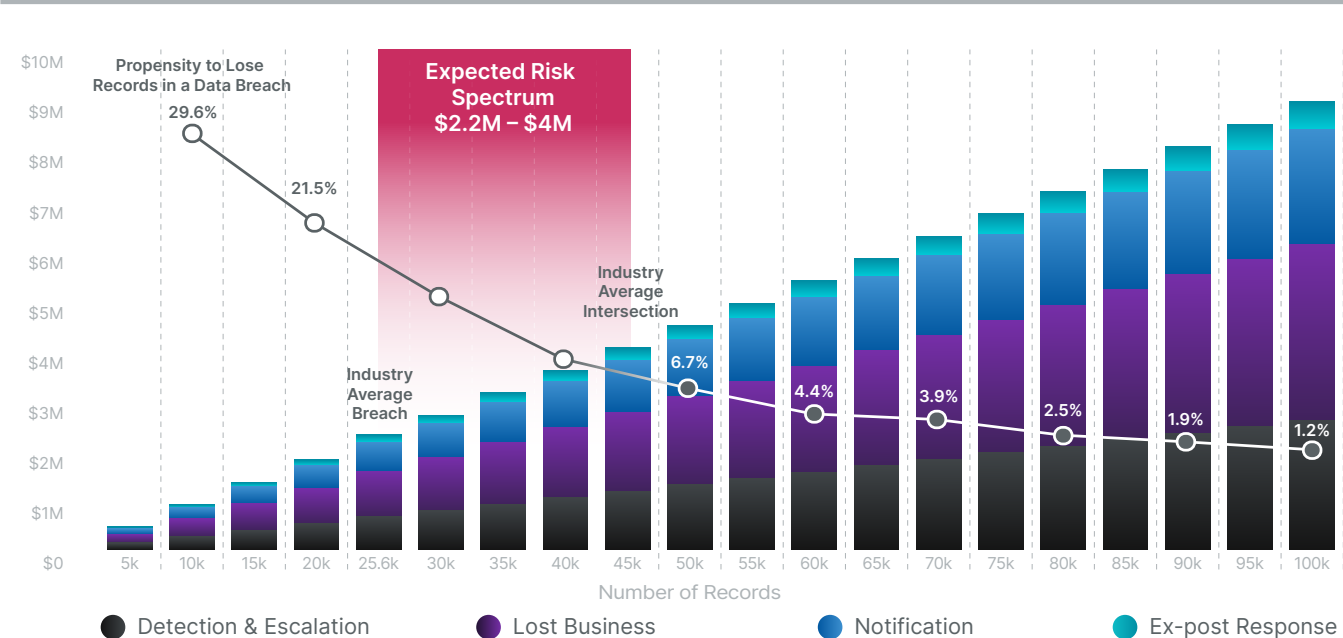
Next is a more detailed look at the expected risk for hospital and healthcare organizations with three yearly critical incidents. The increased financial risk spectrum is a direct result of the higher cost per record and high conversion percentage. The expected risk spectrum jumps up to \$9M to \$10.6M.

HEALTHCARE INDUSTRY: EXPECTED YEARLY RISK



Conversely, consumer goods organizations can potentially set aside less money with a lower expected risk projection, according to the calculated formula. There is a lower cost per record and conversion rate (33%), which is lower than several other industries. For four critical findings in a 12-month period, the expected risk spectrum spans \$2.2M to \$4M, far less than the hospital and healthcare example.

CONSUMER GOODS INDUSTRY: EXPECTED YEARLY RISK



COST TO REPLICATE A CONTINUOUS OFFENSIVE SOLUTION

Now that we have established the expected risk organizations face, we must calculate the cost to replicate a continuous offensive solution that is able to discover the full scope of the attack surface, identify high-risk exposures, and facilitate remediation that outpaces adversaries. For the purposes of our ROI calculation, we have three components we must estimate associated costs for:



Technology and Data Feeds

Identification of the attack surface, relevant exposures, and use of automation/workflow tools comes with a relatively predictable price tag. While there are a multitude of solutions on the market, organizations will need to invest in the appropriate asset discovery tools (commonly referred to as attack surface management) and relevant data feeds that enable them to discover the full scope of their attack surface (not just known IP ranges). Unfortunately, most automated tools often leave portions of the external attack surface undiscovered if a security team is unaware of its presence. Examples would be Github Gists, third parties, subsidies, and shadow IT. Security teams must ensure they have the appropriate technologies and relevant inputs (domains, IP ranges, networks, cloud environments, etc.) that cover the full scope of the attack surface. Unfortunately, as organizations are finding out all too often, a single-missed asset can lead to a business impacting compromise. According to the [2021 ESG Security Hygiene and Posture Management study](#), organizations on average have ten tools dedicated to asset discovery and inventory. For the purposes of our cost to replicate, we estimate the total technology and data feed cost for an average enterprise organization (1,000-5,000 employees) at \$300,000.

All discovered assets should be validated to ensure they are within the scope of the organization; otherwise, security teams risk scanning and discovery of exposures that they do not own. While some discovery and automation tools can validate assets to a degree, there is often an element of human validation needed for accuracy. This is a critical and often overlooked component that must be considered in the replication of a continuous offensive security solution.

Moving ahead, we must look at how exposures are discovered across the attack surface. Oftentimes, attack surface management tools have discovery capabilities that cover a wide range of traditional vulnerabilities and unconventional exposures including exposures within containers, sensitive information leaks, and more. However, these tools are generally supplemented with vulnerability scanning technologies that have long been a cornerstone of security programs. Much like our asset discovery and inventory tools, organizations tend to have more than one solution in place for redundancy and thoroughness. For the purpose of our cost to replicate calculation, we estimate an average enterprise organization would incur a cost of \$50,000 for the relevant tools. However, this cost is widely variable depending on the scope of the attack surface in relation to IPs, assets, domains, etc.

Our final technology component relates to automation workflow tools. It's no secret the number of assets and relevant exposures has expanded beyond the realm of human control. As a result, technology that enables security teams to triage, prioritize, and remediate exposures at scale has become the norm rather than the exception.

In the ESG study, organizations cited automating tasks and processes with security asset management as the highest action that is likeliest to improve security asset management. As a result, the same respondents ranked automating processes associated with security hygiene and posture management as their second highest action, only behind performing continuous testing and validation for gaps within their security controls. While technologies associated with automation can span a wide variety of tools, we have estimated an investment of \$100,000 per year for the average enterprise organization.

It's important to note that many organizations underestimate the scope of their external attack surface. As noted in the ESG study, with 69% of organizations experiencing attacks tied to internet-facing assets, it's no surprise that nearly 70% of those organizations admit that the attack originated through the exploit of an unknown, unmanaged, or poorly managed internet-facing asset. In the context of replication, take careful consideration of the technologies required to discover the complete scope of the attack surface and relevant exposures. The result is often more technologies, more complexity, and the need for more automation to not only cover the wide range of assets but enable the scalable discovery and validation of exposures that present business-impacting risk to the organization.

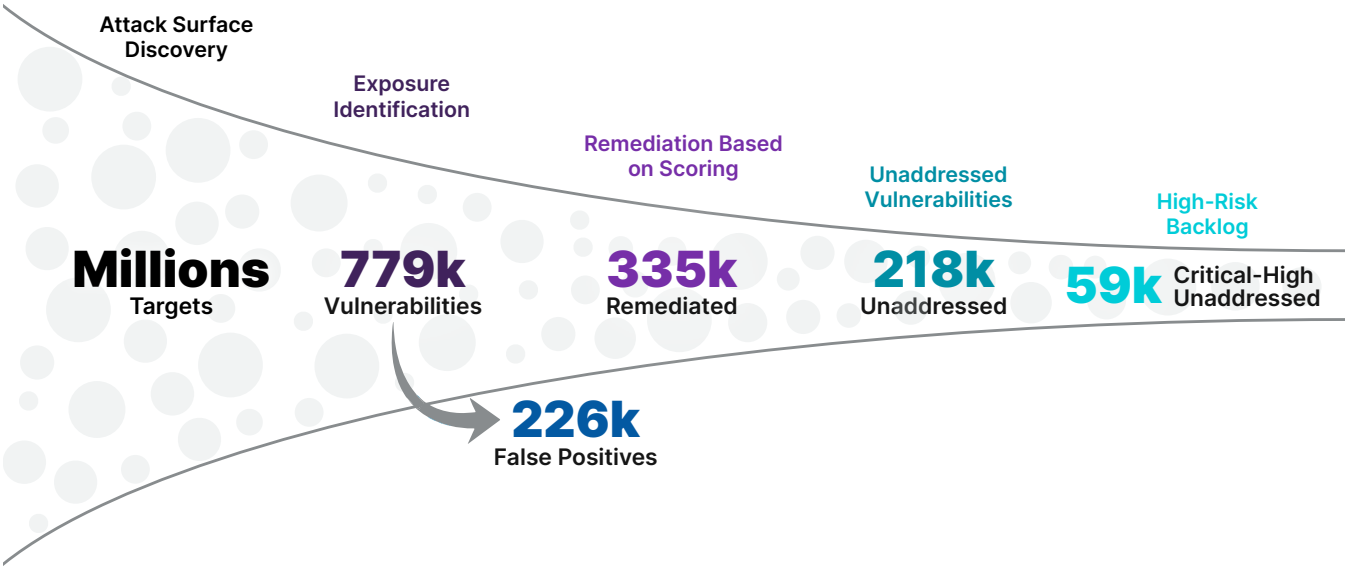
TECHNOLOGY & DATA FEEDS	RECURRING COSTS
ASSET DISCOVERY DATA FEEDS	\$300,000
VULN. SCANNING TOOLS	\$50,000
AUTOMATION WORKFLOW TOOLS	\$100,000

Security Personnel

Acquisition of technologies and relevant data feeds is a relatively straightforward process. The challenge is hiring and retaining the appropriate personnel with the necessary skill sets. Unfortunately, the current supply cannot meet demand. Enterprises are in a constant recruiting and retention battle, resulting in rapid turnover and increasing salary requirements. In recent studies, the average tenure of cybersecurity personnel is a mere 18-24 months with burnout being the primary reason for departure. Unfortunately, replication of a continuous offensive security solution will require overcoming this challenge to effectively operationalize such a large-scale program.

Acknowledging the Time and Effort to Triage

Before we get into the required personnel and associated investment, we must acknowledge the scope of exposures security teams currently face. According to Ponemon's State of Vulnerability Management study, organizations with more than 1,000 employees will average 779,935 individual vulnerabilities identified per scan. Of these vulnerabilities, 29% (226,181) are confirmed to be false positives leaving 71% (553,753) that security teams must spend additional time and effort to triage. Unfortunately, due to resource constraints across prioritization methods, use of automation, and personnel limitations, an astounding 39% (218,381) of vulnerabilities go unremediated. Of these unaddressed vulnerabilities, 27% (58,962) are categorized as high risk on CVSS ratings scales. To paint an even worse picture, these same organizations report a running total of 50,000+ identified and confirmed vulnerabilities in their backlogs. The following funnel illustrates the scope of this ongoing challenge.



In the context of a continuous offensive security solution, ALL exposures must be triaged before testing can occur. While CVSS scores provide a baseline for potential severity, it's not uncommon for seemingly low-scoring vulnerabilities to be overlooked, meanwhile they are leveraged by attackers as steppingstones to more complex attack chains (see our [Wolf in Sheep's Clothing eBook](#)). Focusing testers on the right exposures is critical to identifying risk before attackers have a chance to capitalize.

The associated time and effort to triage these exposures is highly dependent on the right mix of technology, automation, and personnel. To convey the scope of this challenge, let's assume each high-risk and critical vulnerability (which represents 11% of true positive vulnerabilities) takes two minutes to review (assuming limited use of automation), this would amount to almost 2,030 hours of time to review. Taking it a step further, that's 254 workdays (assuming 8 hours a day). Given a 40-hour work week, that would require at least one personnel purely dedicated to triaging scan results year-round – can you imagine the turnover and burnout of doing this every day? While we know this is an extreme example as most organizations have resources to scale this process, it conveys the scope of the challenge required to enable continuous testing at scale.

The Critical Roles.

With our personnel requirements, the following roles are critical to operationalizing a continuous offensive testing solution (assuming acquisition and execution of appropriate technologies and processes):

ATTACK SURFACE ANALYST



Responsibility: Operationalizing asset discovery tools including the categorization of assets, severity, tagging, and grouping of assets. This role is critical in ensuring the full scope of the attack surface is continuously discovered enabling thorough exposure identification.

Cost: \$125,000 USD

VULNERABILITY RESEARCHER



Responsibility: Researching, aggregating, and operationalizing intelligence on the latest exposures and vulnerabilities in the exposure discovery process. This can include the creation of new analyzers and methods that are not covered by licensed technologies. Research and operationalizing of intelligence is highly focused on contextual risks relevant to the organization's unique risk profile.

Cost: \$185,000 USD

VULNERABILITY ANALYST



Responsibility: Triageing scan results to sift out false positives, de-duplicate findings, eliminate noise, and prioritize potential exposures for testing and validation. Note: The number of analysts required is highly dependent on maturity of automation processes that can decrease dependency on human intervention.

Cost: \$125,000 USD

SECURITY ENGINEER



Responsibility: Continuous development of backend platforms and capabilities that enable interoperation of technologies, automation, and personnel across the continuous offensive security solution.

Cost: \$185,000 USD

PENETRATION TESTER



Responsibility: Validating triaged exposures are exploitable using the same tactics, techniques, and procedures observed in real-world attack scenarios. Also, responsible for determination of potential business impact through execution of post-exploitation activities that identify pathways, systems, and data at risk. Supports vulnerability managers in remediation procedures and validation that exposures are no longer susceptible to compromise.

Cost: \$187,500 USD



VULNERABILITY MANAGER

Responsibility: Facilitates remediation of exposures including relevant internal pathways, systems, and data at post-exploitation risk.

Cost: \$187,500 USD



OPERATIONAL MANAGER




Responsibility: Oversees the continuous offensive testing program and relevant personnel, ensuring optimal operation and achievement of targeted outcomes that justifies investment and return on security investment.

Cost: \$218,700 USD

* Personnel costs reflect base salaries, bonuses, benefits, and other means of compensation

The following table outlines our total costs including the estimated number of personnel required to operationalize a continuous offensive testing program for an average size organization with 1,000-5,000 employees. Again, this is for illustrative purposes, accurate estimates are highly contingent upon each organization's unique attack surface and risk profile.

Once the associated costs for replication can be determined, we can combine those values with expected risk to determine our potential return on investment.

	COST PER PERSON	QUANTITY	RECURRING COSTS	ONE-TIME COSTS
<div> TECHNOLOGY & DATA FEEDS</div>				
ASSET DISCOVERY DATA FEEDS			\$300,000	
VULNERABILITY SCANNING TOOLS			\$50,000	
AUTOMATION WORKFLOW TOOLS			\$100,000	
<div> PERSONNEL</div>				
ATTACK SURFACE ANALYST	\$125,000	1	\$125,000	
VULNERABILITY RESEARCHER	\$185,000	2	\$370,000	
VULNERABILITY ANALYST	\$125,000	3	\$375,000	
SECURITY ENGINEER	\$185,000	2	\$370,000	
PENETRATION TESTER	\$187,500	3	\$562,500	
VULNERABILITY MANAGER (REMEDIATION)	\$187,500	1	\$187,500	
OPERATIONAL LEADER / MANAGER	\$218,750	1	\$218,750	
<div> HUMAN RESOURCES</div>				
RECRUITING & HEADHUNTER FEES				\$10,000 pp
TOTAL	\$2,788,750			

CALCULATING ROI: THREE METHODS


We use three different approaches to calculate ROI. This is where it is essential to know the expected risk spectrum, as well as cost to replicate a continuous offensive security solution. For simplicity purposes, we will use \$100,000 as the cost for an outsourced solution, obviously this can vary widely in relation to the scope of service and relevant attack surface.

The first ROI method simply uses the average of the upper and lower end of the expected risk spectrum and divides this value by the solution cost. In the case of our industry agnostic example, this would produce an ROI of 26:1 $\{[(\$1.9M + \$3.3M)/2]/\$100k\}$.

Another method that can be used is the cost to replicate the solution in-house divided by the outsourced solution cost which produces an ROI of approximately 28:1 $(\$2.788M/100k)$.

Lastly, the third method, and the one most often used by security leaders, produces an ROI value using a combination of cost to replicate an outsourced solution and a residual of expected risk. This residual risk is based on the projected time frame it takes to fully operationalize a continuous offensive security solution. During that time period, exposures can still go unaddressed, giving threat actors an opportunity to capitalize on susceptible environments. For example, if we assume a twelve-month window to fully operationalize a continuous offensive security solution, we can take a portion of our expected risk relative to the time frame the solution is not fully operational. Let's say this time frame is six months for simplicity purposes. This approach produces an ROI of approximately 40:1 using the following values from our industry-agnostic example: $\{[(\text{Average risk spectrum}/2) + \text{Cost to replicate}]/\text{Solution cost}\}$.

CALCULATING ROI: THREE METHODS



SUMMARY

Armed with knowledge to identify the unique inputs of their respective organizations, security teams can generate outputs from the customized ROI calculator to demonstrate cost savings and risk mitigation associated with a public breach and data disclosure.

While each method produces a compelling ROI, your target audience should be considered. The method you choose is the one you believe communicates the greatest value to your target audience.



RISK-MITIGATION METHOD

The risk-mitigation method is intended for a non-technical or risk-adverse audience such as security leadership, Chief Risk Officer, Board of Directors, or other executives vested in protecting shareholder value.



COST-REPLICATION METHOD

On the flip side, the risk-mitigation method vs. cost replication is intended for a budget-minded audience, such as Finance or Procurement teams.



BLENDED METHOD

The blended method is intended for a broader audience that could include risk- and cost-minded decision makers.

Using this customized calculation method helps security teams break down costs of technologies, processes, and people associated with in-house teams, and show financial decision makers that investment in a continuous offensive security solution can be more cost and time-effective while ensuring that attack surfaces remain impenetrable to attackers.

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.

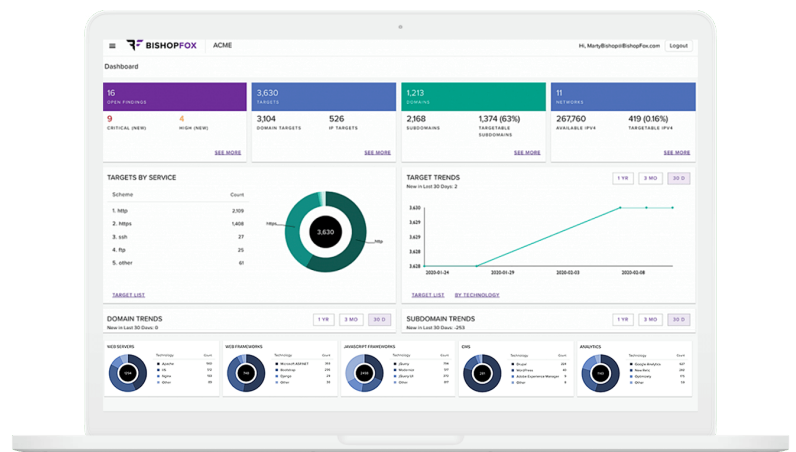
Security isn't just a job to us. We do this because we love it — and because we're committed to the common good. In fact, we have authored 15 open-source tools, shared groundbreaking research, and published more than 50 security advisories in the last 5 years.

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.



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](#)