Exposure Management Metrics Primer

The emergence of exposure management is already shifting how organizations approach cyber risk, and an essential element of that shift is how cyber risk is measured. This primer presents an assortment of metrics that organizations can use to analyze and guide every element of their exposure management programs.

Using these metrics, security teams can improve the efficiency and effectiveness of their operations by enhancing the way they prioritize risk. They'll also be better able to quantify and communicate risk to C-level executives, boards of directors and other business teams involved with exposure management.

Operational metrics

These metrics are used by security teams to guide day-to-day exposure management activities.

Attack surface visibility	
Description	Cyber asset visibility is the foundation of exposure management. Organizations have three types of cyber assets: ▪ Known known: Cyber assets that you know are part of your attack surface and are formally accounted for in some fashion. ▪ Known unknown: Cyber assets that you know are part of your attack surface but don't have visibility of. ▪ Unknown unknown: Cyber assets that may or may not be part of your attack surface — you don't know. An attack surface visibility metric tells you how much of your attack surface you can see. Measuring that proportion helps security teams determine where they have visibility gaps they must close so they can better protect their organizations from cyber attacks targeting these blind spots. The goal — though it's likely unachievable — is 100% visibility.
How to measure	Divide the approximate number of cyber assets you have visibility of through the products or platform you're using for exposure assessment by the approximate number of total "known known" and "known unknown" assets in your environment, then multiply by 100 to convert the figure to a percentage.

Note that "unknown unknown" assets are not accounted for in attack surface visibility, as they make the metric impossible to calculate, but you can't just ignore them. Turn those "unknown unknowns" into "known knowns" and move one step closer to 100% attack surface visibility with Ivanti's Attack Surface Checklist.

Assets under management	
Description	One step beyond calculating attack surface visibility is calculating assets under management. This metric has traditionally been used to determine what percentage of devices are being actively managed through a unified endpoint management (UEM) tool, which lets IT administrators control the devices' settings and security policies. In exposure management, "assets under management" are the assets that are generating exposure findings for your Security team. Remember: asset visibility is just the first step — exposure findings are needed to calculate risk.
How to measure	Divide the number of assets from which you're collecting exposure findings by the number of visible assets (use "attack surface visibility" figure from above). Multiply by 100 for a percentage.

Use Ivanti's Exposure Management Readiness Checklist to determine which technologies you need to add to your cybersecurity stack to bring 100% of your assets under management.

Exposure risk ratings	
Description	Once you have visibility of your cyber assets and exposures, you can start assigning risk ratings to those exposures. With exposure risk ratings, security teams can identify the threats most likely to be exploited against the organization and prioritize them for proactive remediation. You can also gauge which threats warrant the attention of the C-suite, board and other exposure management stakeholders.
How to measure	Exposure assessment products and platforms automate the process of calculating risk ratings for vulnerabilities, misconfigurations, security control gaps and other exposures using a range of factors: ▪ Severity (e.g., CVSS and/or vendor score) ▪ Exploitability (e.g., vulnerability intelligence and/or EPSS) ▪ Potential impact (e.g., asset context and/or business impact) Comprehensive risk ratings are part of what makes exposure management different from traditional vulnerability management (which only considers severity) or risk-based vulnerability management (which considers threat context, but not business impact).

Decision-making metrics

These metrics empower C-level executives, boards of directors and other executive stakeholders to make informed decisions around exposure management.

Cyber risk level	
Description	Cyber risk level is simply a numerical score for an organization's cybersecurity posture. Paired with the organization's risk appetite — the level of cyber risk the organization is prepared to accept in pursuit of its business objectives — an organization's cyber risk level provides important context for any executive-level exposure management decision.
How to measure	Exposure assessment products and platforms quantify organizational risk levels by starting at the asset level. They assign risk scores to each asset in an organization's environment based on factors such as: ▪ Risk ratings of exposures associated with the asset (as described above). ▪ Accessibility of the asset (how easily an asset can be accessed via an external or internal pathway). ▪ Asset criticality (the importance of an asset to the organization in terms of potential impact from loss or compromise). Asset risk scores can then be rolled up for decision-making at the relevant level, whether that's the entire organization or a given business unit or sales region. Risk score ranges and meanings differ from vendor to vendor, but many use a 0-to-100 scale like the one used for tests and quizzes — lower numbers equate to higher risk and a weaker security posture, while higher numbers indicate lower risk and a stronger security posture.

Before presenting exposure management stakeholders with any official data, educate them on cyber risk level and risk appetite with this introduction to exposure management deck.

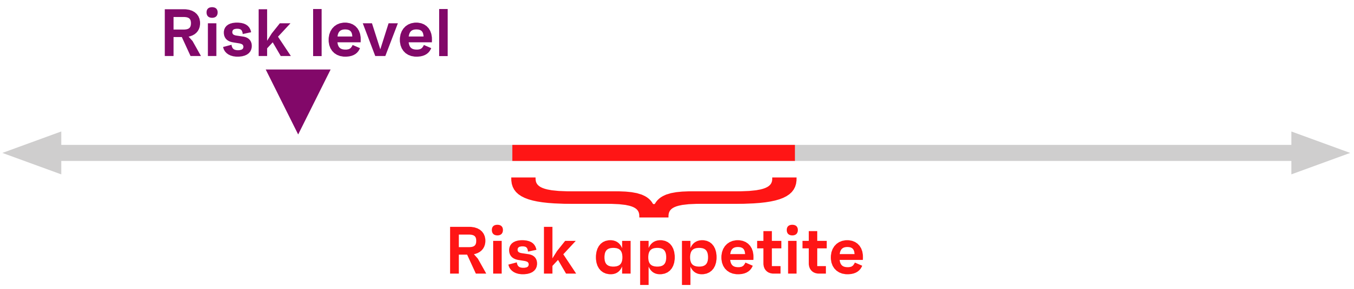

Projected cyber risk level	
Description	An organization’s current risk level is only one part of the decision-making process. It’s equally important to project how a given decision might affect that risk level. ▪ If your cyber risk level is below your risk appetite, and remediation of a given exposure is projected to bring your risk level within your risk appetite (or close to it), then that activity warrants further consideration. ▪ Remediation also warrants consideration when an exposure stands to bring an organization’s risk level below its risk appetite if left unaddressed. ▪ When an exposure isn’t projected to substantially impact risk, either positively or negatively, it can usually be accepted and reevaluated later. While metrics around cyber risk level are helpful for at-a-glance analysis — they allow organizations to quickly understand their current situation and signal when action may be needed — major decisions take more advanced analysis. Metrics like annual loss expectancy (covered below) can help with that.
How to measure	Projecting risk level changes based on recommended remediations is an advanced capability found in some exposure assessment products and platforms. Projections are calculated by estimating how remediating an exposure would impact the factors that go into calculating an organization’s current cyber risk level (as described above).

Annual loss expectancy (ALE)	
Description	Business executives ultimately want things explained to them in the language of business. Annual loss expectancy does just that by quantifying the risk of an exposure in monetary terms.
How to measure	ALE is calculated by multiplying the following variables: ▪ Single loss expectancy (SLE) — a measure of business impact — assigns a dollar value to a single cyber attack by multiplying the value of the asset at risk by the loss percentage that a threat could have on that asset. ▪ Annualized rate of occurrence (ARO) — a measure of likelihood — estimates the frequency at which the threat will strike over a one-year period. By comparing current ALE to residual ALE — the estimated ALE if a risk were to be mitigated or remediated — executives can conduct a cost/benefit analysis to decide whether the organization should address a risk or accept it. Of course, as with most exposure management decisions, the organization's risk appetite should be considered.

Check out Ivanti's Evaluating Cyber Risk Objectively: A Guide to Data-Driven Risk Assessments e-book for step-by-step instructions on how to use ALE and other metrics to guide exposure management decisions at your organization. You can use these risk reporting slides to present your assessments to stakeholders.

Performance metrics

These metrics show the results of exposure management activities. Proving positive impact can increase organizational support for and investment in exposure management.

Retired threat debt	
Description	Security teams rarely get credit for the proactive work they do to protect the organization. All too often, the bulk of the recognition they get is negative: when a security incident strikes and the team responds. Exposure management stands to change that by involving the business side of the house in cybersecurity decisions, though they're still only likely to be involved in more significant decisions, like whether a "celebrity" vulnerability requires remediation or an investment should be made to implement a new security control. By reporting on retired threat debt, security teams can show the full scope of exposure management work they've completed. One important thing to keep in mind when it comes to retired threat debt metrics: they're activity-based, and exposure management is all about outcomes. While they help convey how much work a team did, it's more important to show the results of that work.
How to measure	Exposure assessment products and platforms should provide data you can use to construct retired threat debt metrics. Filtering can usually be applied to that data to zero in on specifics. Some examples: ▪ Retired 9,423 instances of malware in Q2. ▪ Retired 399 vulnerabilities tied to ransomware in Q1. ▪ Reduced the number of critical vulnerabilities by 40% in the last quarter.

Cyber risk level changes

Description	While current and projected cyber risk levels can be used to guide near-term decisions about whether to remediate an exposure, changes to cyber risk levels over time show the long-term impact of those decisions. One way security teams can prove the results of their exposure management work is by demonstrating an ability to keep the organization's cyber risk level within its risk appetite, or to improve it when needed. Reviewing current and projected cyber risk levels with executive stakeholders on an ongoing basis ensures there won't be any surprises when you report the changes, or lack thereof, over time. Rather, reports on cyber risk level changes can be opportunities to reinforce that exposure management is working as designed.
How to measure	Cyber risk level changes can be measured and reported in various ways: ■ Absolute change (the exact numerical difference between cyber risk levels at two different points in time — for example, if the cyber risk level was 72 on January 1 and 83 on December 31, that's a change of +11 over the course of a calendar year). ■ Percentage points (the percent difference between cyber risk levels at two different points in time — sticking with the example above, the change would be 15.3% assuming risk levels are being measured on a scale of 0-100). ■ Line chart (shows cyber risk levels at different points in time and shows any trends — for example, whether levels are generally going up or down). Some exposure assessment products and platforms will automatically present this data on dashboards.

Return on security investment (ROSI)	
Description	Again, business executives tend to want things explained in business terms. There is perhaps no more prominent business metric than return on investment (ROI). Return on security investment (ROSI) is the Security team's version, and it may be the most powerful tool in their exposure management reporting toolbox. Ascribing dollar values to exposure management efforts allows security teams to report on the value of individual investments as well as the cumulative value of all exposure management efforts over a given time period.
How to measure	Unlike traditional ROI, which measures an investment's financial gains relative to its costs, ROSI measures projected financial losses averted thanks to an investment relative to its costs. Those projected losses can be estimated with reasonable accuracy using a mixture of an organization's historical data and industry research. Traditional ROI is calculated by subtracting the cost of an investment from its current value, dividing by the cost and then multiplying by 100 to come up with a percentage. By making one modification to that equation — changing the current value of an investment to the value of projected losses averted because of an investment(s) — organizations can calculate return on security investment (ROSI) to measure the financial impact of exposure management investments.

About Ivanti

Ivanti breaks down barriers between IT and security so that Everywhere Work can thrive. Ivanti has created the first purpose-built technology platform for CIOs and CISOs — giving IT and security teams comprehensive software solutions that scale with their organizations' needs to enable, secure and elevate employees' experiences. The Ivanti platform is powered by Ivanti Neurons - a cloud-scale, intelligent hyper automation layer that enables proactive healing, user-friendly security across the organization, and provides an employee experience that delights users. Over 40,000 customers, including 85 of the Fortune 100, have chosen Ivanti to meet challenges head-on with its end-to-end solutions. At Ivanti, we strive to create an environment where all perspectives are heard, respected and valued and are committed to a more sustainable future for our customers, partners, employees and the planet. For more information, visit ivanti.com and follow @GoIvanti.

ivanti

For more information,
or to contact Ivanti,
please visit ivanti.com.