WHITE PAPER

Ivanti RS³®

ivanti

Table of Contents

Executive Summary

Ivanti Neurons ingests information from internal security intelligence, vulnerability scanners, and asset criticality data. This data is then augmented and analyzed with curated threat feeds, account accessibility, and vulnerability risk context in order to calculate the Ivanti RS³.

The Ivanti RS³ is a contextualized, risk-based view of your organization's cybersecurity posture, allowing your security and IT teams to move from detection to remediation in minutes, not months. This document describes the security attributes and terminology used by Ivanti Neurons in quantifying risk and vulnerabilities.

Our proprietary scoring methodology is constantly evolving to provide the most accurate threat context with the widest angle. Score assessments continue to be enhanced with fine-tuned analysis of threat information for broad coverage, as well as input on exploit trends and newly validated exploits.

Vulnerability Risk Rating

Vulnerability Risk Rating (VRR) considers industry-standard CVSS and CWE data, open-source threat intelligence, subject matter expertise, trending information, and more. The VRR indicates how much of a threat a finding poses to an organization or business.

Asset Criticality

An asset's criticality is the importance of an asset to the organization or business in terms of potential impact from loss or compromise. It is derived from user-specified values on a scale from 1 to 5, with 5 being the most critical.

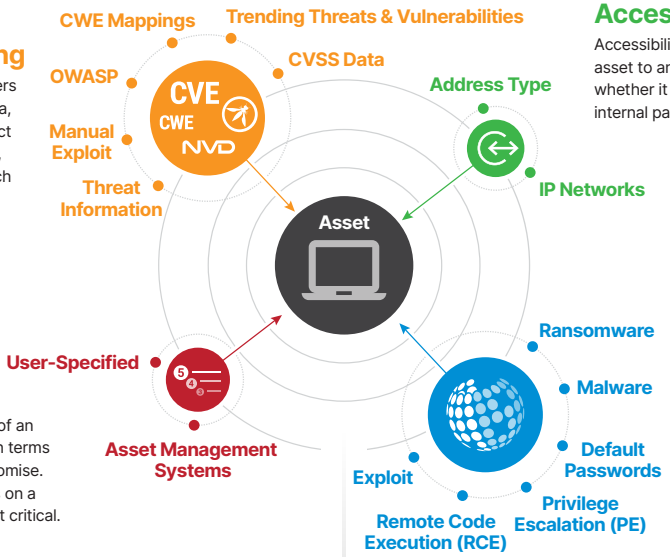

Accessibility

Accessibility reflects the potential exposure of an asset to an organization or business based on whether it can be accessed via an external or internal pathway.

Ivanti Neurons for Vulnerability Knowledge Base

The Vulnerability Knowledge Base is an aggregate of temporal threat intelligence from over 100 sources, trending exploit information, and Ivanti-identified threats. The comprehensive threat surface is then correlated to vulnerability information that provides a rich context for intelligent scoring and prioritization.

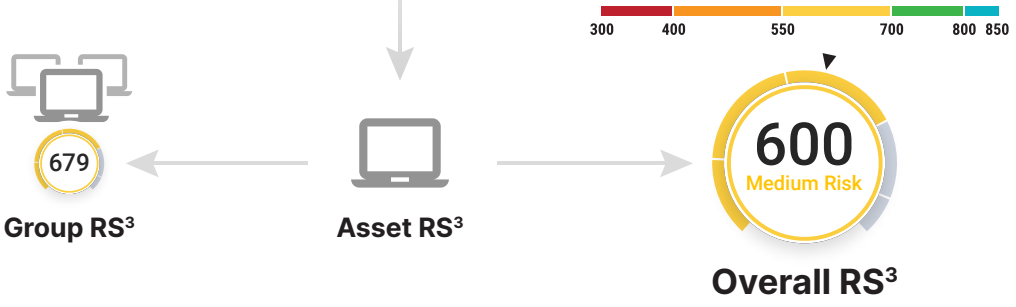

Introduction

Ivanti Neurons powered by RiskSense ingests and correlates data from vulnerability scanners and configuration management systems via our RESTful API. It normalizes the data and minimizes false positives by conducting differential analysis across sources, unifying the data to avoid duplicates, and enriching the findings. At the same time, Ivanti's correlation engine aggregates data across multiple characteristics and maps assessment information to compliance requirements. Together, these processes deliver to security operations and IT teams personalized prioritization of vulnerability remediation in their environment.

Risk is incorporated by contextualizing the internal vulnerability scanner data with external threat data. The platform accomplishes this by unifying this data with threat intelligence feeds from over 100 sources and relevant penetration test data to provide the most accurate threat assessment available in the industry. By doing so, it not only provides valuable threat associations but also directs focus to those security gaps being actively exploited by cyber adversaries.

Ivanti Neurons powered by RiskSense uses an advanced analysis engine powered by machine learning algorithms and human subject matter expertise to continuously measure, monitor, and track an organization's overall exposure to risk, reflected in the Ivanti RS³ and a visual representation of cyber risk posture at the overall, group, and asset levels. The score itself and accompanying in-platform data readouts enable security and IT teams to quickly answer questions from auditors, boards of directors, and the C-suite.

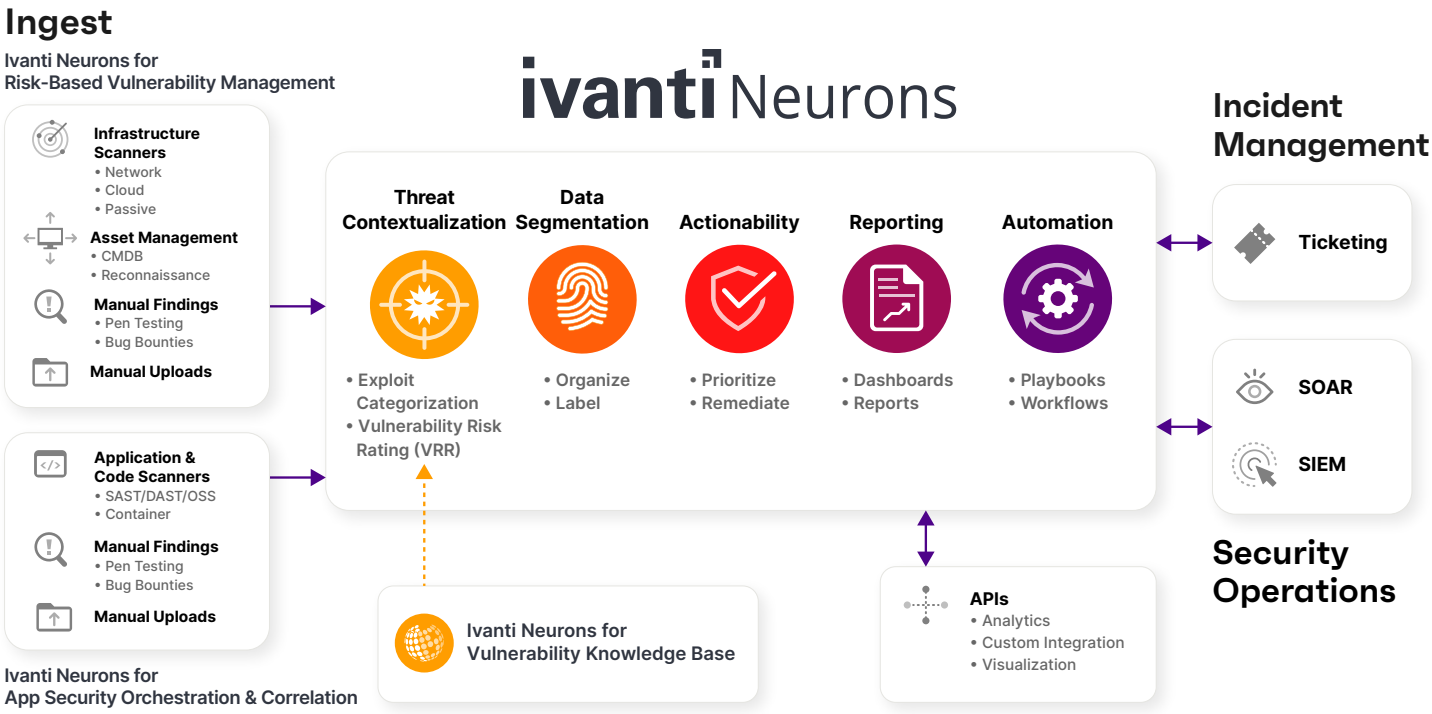

Ivanti Neurons Terminology

The **Ivanti RS³** provides a measure of an organization's overall security and protection against cyber risks and exploits, similar to the way that credit scores provide a measure of creditworthiness. The value of an RS³ falls in the range 300 to 850, with higher scores indicating better overall security against cyber threats and exploits. This range of scores is broken up into five score bands, shown below.

Ivanti Neurons uses the following factors to calculate an organization's Overall RS³:

- **Address Type (Accessibility):** An asset's ability to be accessed either via external or internal pathways.
- **Assets:** The components of an organization's infrastructure, including network hosts (laptops, printers, servers, etc.) as well as web applications. Vulnerabilities that have been identified in an organization are associated with individual assets.
- **Base Score (Severity):** Conveys the base amount of risk posed by vulnerabilities, presented qualitatively as Critical, High, Medium, Low, or Informational. Typically sourced from the NVD severity score or reported by the vulnerability scanner.
- **Criticality:** The importance of an asset to the organization or business in terms of potential impact from loss or compromise, as defined by the user on a scale of 1 to 5, with 5 being the most critical.
- **Vulnerability Risk Rating (VRR):** Represents the risk posed by an individual vulnerability to an asset, presented as a numerical value between 0 and 10, with 10 representing the maximum amount of risk. An organization with a large number of vulnerabilities having high VRR will thus have a lower RS³, while an organization with few low-VRR vulnerabilities will have a higher RS³. Details on VRR are explained further in the following sections.
- **Threat Factors:** (Also known as Exploitability) Parameters that classify the type of attack that can be executed by a given vulnerability, and how difficult it is to execute such an attack. For example, malware is easy to execute, but writing a new kind of attack is more difficult.
- **Vulnerabilities:** Individual findings representing specific instances of weaknesses or indicators of compromise. In Ivanti Neurons, vulnerabilities can be identified by the scanner plugins that found them.

In Ivanti Neurons, every asset (including network hosts and web applications) is given an individual RS³. These asset-level scores are then aggregated together using a weighted averaging scheme to obtain the RS³ for a **Group** (a user-defined collection of assets) as well as the RS³ for the entire **Organization** (all assets).

Vulnerability Risk Rating

VRR represents the risk posed by a given vulnerability, provided as a numerical score between 0 and 10. The higher the risk, the higher the VRR. The score quantifies adversarial risk by leveraging a combination of standardized metrics and knowledge gathered by application scanners through black box testing. Industry-standard sources such as the National Vulnerability Database (NVD), Common Vulnerabilities and Exposures (CVE)[1], Common Weaknesses Enumeration (CWE)[2], and OWASP Top 10[3] are combined with carefully curated high-fidelity threat intelligence sources, as well as subject matter expertise from penetration testers and vulnerability analysts. Together, these components are used to build data-driven models that power and inform the VRR scoring algorithm.

The VRR methodology utilizes this aggregation of information to widen the lens on vulnerabilities, building a larger body of context for each one. Rather than categorically increasing base level scores (resulting in risk inflation), the VRR algorithm intelligently separates and elevates the riskiest weaknesses, allowing consistent prioritization using key indicators of potential compromise across host and application findings.

To assign a VRR to an individual vulnerability, we first begin with the **base score** of the vulnerability. The Common Vulnerability Scoring System (CVSS) provides a baseline evaluation of a vulnerability's risk using a set of standardized metrics reflecting inherent properties of the vulnerability itself, captured in a CVE. Ivanti utilizes the CVSS v3 where available, and substitutes the CVSS v2 when v3 is not provided.[4] If a vulnerability lacks a CVE association, the reported severity from the scanner plugin that identified the vulnerability is used instead. Scanners use a range of different qualitative and quantitative measurements to report severity; these values are normalized to be consistent with our 0 to 10 numerical scale.

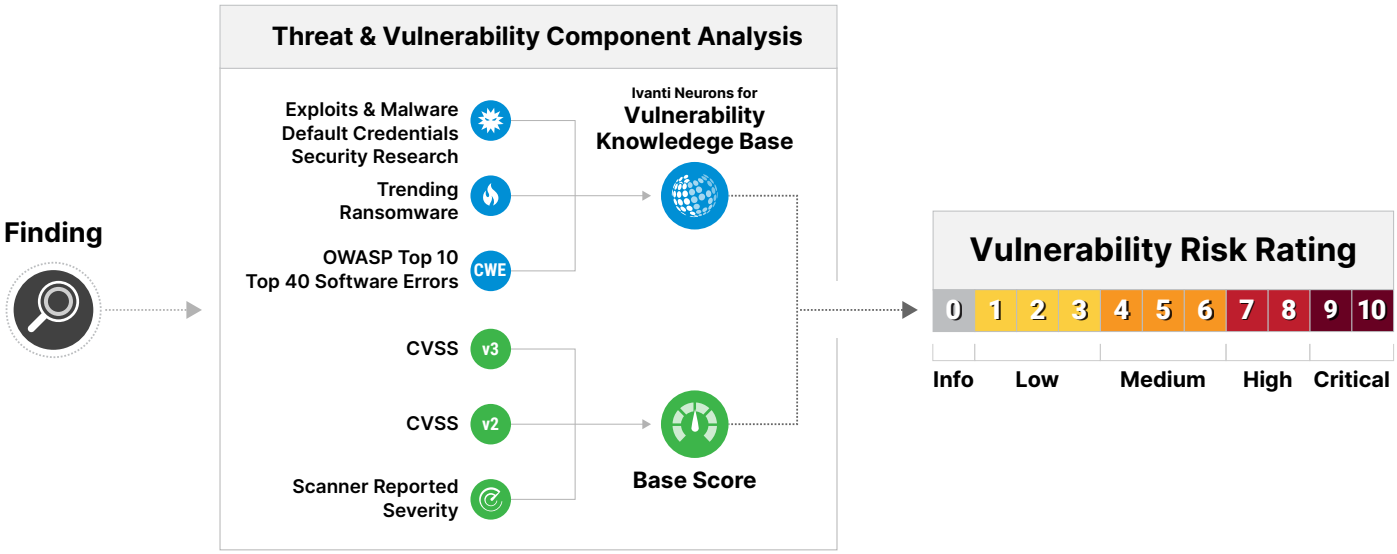

[1] The Common Vulnerabilities and Exposures (CVE) is a catalog of known security threats. The catalog is sponsored by the United States Department of Homeland Security (DHS), and threats are divided into two categories: vulnerabilities and exposures. More information can be found here

[2] The Common Weaknesses Enumeration (CWE) is a formal list or dictionary of common software weaknesses that can occur in a software's architecture, design, code or implementation that can lead to exploitable security vulnerabilities. More information can be found here

[3] The OWASP Top 10 is an industry-standard awareness document that represents the most critical security risks to web applications. More information can be found here

[4] The Common Vulnerability Scoring System (CVSS) is a published standard used by organizations worldwide and provides a way to capture the principal characteristics of a vulnerability and produce a numerical score reflecting its severity. Mentions of CVSS v3 include CVSS v3.0 and CVSS v3.1. More information can be found here

The next step in computing VRR is the identification of the vulnerability's **threat factor**. The CVSS alone cannot provide a complete picture of the severity posed by a vulnerability. Therefore, Ivanti Neurons augments each vulnerability with contextual threat information consisting of data on known exploits, malware and Ransomware associations, and other data seen in the wild and/or utilized by real-world attackers. We gather information from a wide range of carefully chosen threat intelligence sources and knowledge bases, correlating and normalizing it into a unified proprietary database known as the Ivanti Neurons for Vulnerability Knowledge Base, which in turn feeds Ivanti Neurons. In this way, not only is each vulnerability associated with potential and confirmed threats, it is also placed in the context of its environment for a more robust risk profile.

Not all threats are created equal, and thus the VRR methodology is designed to intelligently handle weaknesses and vulnerabilities from both applications and network hosts. When analyzing a given finding, VRR considers the most severe threat among all available vulnerability and weakness information. Ivanti Neurons leverages industry-accepted lists, such as the CWE Top 40 Software Errors and OWASP (Open Web Application Security Project) Top 10, to provide additional context when no threat is available. Organizations require a well-rounded view of vulnerabilities if they are to prioritize remediation. By infusing the approach with subject matter expertise and threat intelligence, VRR becomes essential for optimizing the risk management of your organization's infrastructure and applications.

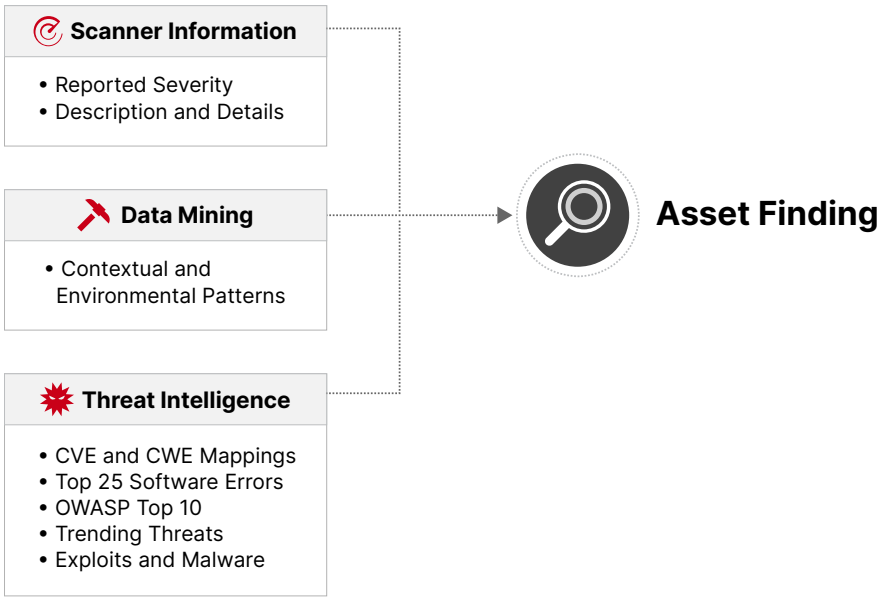

Next in the VRR calculation sequence, because not all threats are exploited with equal frequency, additional parameters are identified: **Trending Threats**, **Ransomware**, **Manual Exploits**, and **Security Research**. Research tells us that certain kinds of newly discovered or existing threats are utilized in the wild with high intensity over a short period of time — we designate these as **Trending Threats** or **Vulnerabilities**. Those findings which are associated with a Trending Threat receive the maximum VRR score and should therefore be addressed as a top priority in an organization. Certain threat categories are uniformly considered to be unusually popular at times as well. In the current security landscape, **Ransomware** vulnerabilities represent critical threats to organizations consistently. For this reason, vulnerabilities associated with any kind of Ransomware are treated similarly as individual Trending Threats and assigned the maximum VRR value.

The second additional parameter, **Manual Exploits**, addresses a known limitation of any automated scoring algorithm: there will always be exceptions to any crafted set of rules. **Security Research** completes the picture, either validating a designation with an intuitive rationale, or by providing qualitative context that can only be identified using subject matter expertise. Penetration tester-validated Manual Exploits single out potential catastrophic compromise to an organization, and are identified to specific vulnerabilities. In addition, organizations such as the Department of Homeland Security, Federal Bureau of Investigation, National Security Agency, and others frequently publish critical-risk CVEs that pose significant threats to organizations. These vulnerabilities are validated by our Penetration testing teams and published under the curated **RiskSense Attack Surface list**.

These extreme-risk vulnerabilities will always have the maximum value VRR, and supersede all other factors in terms of prioritization.

The final step in the VRR calculation methodology is to pass all of the above parameters into a sequence of data-driven algorithms which map each parameter value to a corresponding numerical value, and then compute the explicit VRR using those mapped values. Each vulnerability's risk rating lies within one of a set of **severity** categories according to its numerical value: Critical, High, Medium, Low, or Informational. The equations powering these algorithms have been carefully calibrated using a combination of machine learning techniques and subject matter expertise to guarantee the resultant score is an accurate and actionable representation of the risk posed by a given vulnerability.

Ivanti RS³ Methodology

Asset RS³

In Ivanti Neurons, each **asset** receives an Ivanti RS³ to quantify its robustness against cyber attacks. Both network hosts and web applications are subject to the same analysis and receive an RS³ using the same methodology.

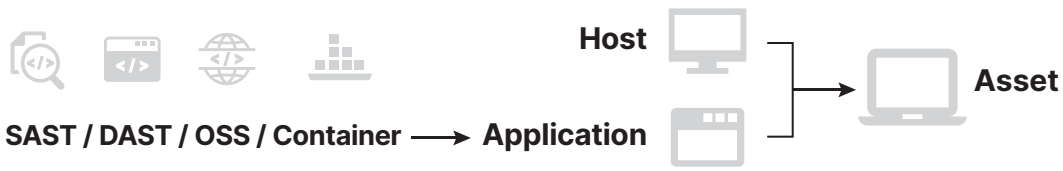

The Ivanti RS³ is provided as a whole number on a scale of 300 to 850; the higher the RS³, the more secure the asset is. To assign this score, we first begin at the highest score of 850, and then apply numerical deductions corresponding to the amount of risk posed to the asset from its various properties and weaknesses.

We begin with the inherent weaknesses an asset is subject to, quantified by the VRRs of its constituent vulnerabilities. This set of VRRs is aggregated together into a single numerical representative by considering the greatest VRR of vulnerabilities that lie in each major severity category: Critical, High, Medium, and Low[5], e.g., we first examine the set of Critical vulnerabilities on an asset, and select the greatest VRR among them to represent that category. This is then repeated for the High, Medium, and Low categories. These four values are combined into a single **representative VRR value** using a "smart" weighted average scheme, in which

[5] Informational vulnerabilities are excluded because they definitionally have a VRR of zero, and thus do not contribute to an organization's risk.

the weights adapt to different configurations of the greatest VRR values. For example, the weights used for a large set of vulnerabilities may be different from those used in a much smaller set. In this way, Critical vulnerabilities greatly impact an asset's security posture, while Low vulnerabilities make a substantially smaller contribution.

In addition to the numerical value of the VRRs of the vulnerabilities, we also consider the **finding concentration** (that is, the number of vulnerabilities) on that asset relative to the entire organization. Although this value has less of an impact on the overall security posture of an asset than its constituent VRRs, this factor ensures that assets which are otherwise equivalent and have similar VRR score profiles can be easily differentiated in priority. Suppose two assets have the same representative VRR value but a different total number of vulnerabilities; the asset with more vulnerabilities is guaranteed to have a lesser or equal security posture score than the other.

The representative VRR value and finding concentration for a given asset are passed into an equation which calculates that asset's **Security Risk**, the primary deductive factor in calculating its RS³. Security Risk is a penalty value in the range of 0 to 440, with a higher numerical value corresponding to a greater amount of risk. If an asset contains zero open non-Informational vulnerabilities, its Security Risk will be 0, i.e., no penalty is applied to the initial 850 value.

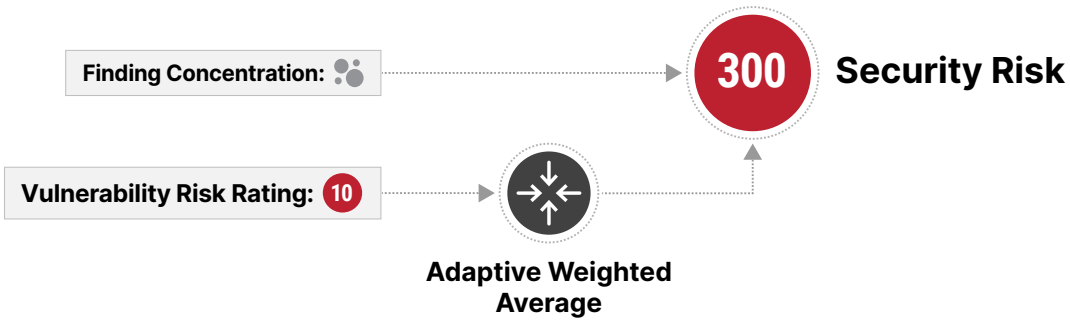

We then move to a contextual view of the asset, taking into account its inherent properties as designated by its own infrastructure as well as the overseeing user. The **Address Type** (also referred to as Accessibility, Externality, or Access Type) of an asset may be either Internal or External, and conveys the additional risk posed to an internet-facing asset. For network hosts, this value is determined by the standard RFC 1918 protocol but may be manually modified by a user under special circumstances. For web applications, this must be specified by the organization. The importance of a particular asset to organizational operations is captured in its asset **Criticality** as a number from 1 (less importance) to 5 (highest importance). This value represents the risk associated with the loss or compromise of an asset, reflecting how important its functionality is to the operations of the overall organization.

The asset's Externality and Criticality are passed into an equation which calculates that asset's **Business Risk**, the secondary deductive factor in computing its RS³. Business Risk is a penalty value in the range of 0 to 110, again with a higher numerical value corresponding to a greater amount of risk. The Security Risk value of the asset is also passed into the Business Risk equation as an auxiliary parameter, to help determine the relative impact of the asset's constituent weaknesses on its business-related properties. A Security Risk value of exactly 0 will result in a Business Risk of 0 and thus lead to a perfect score for that asset (by definition, since it has no non-Informational vulnerabilities).

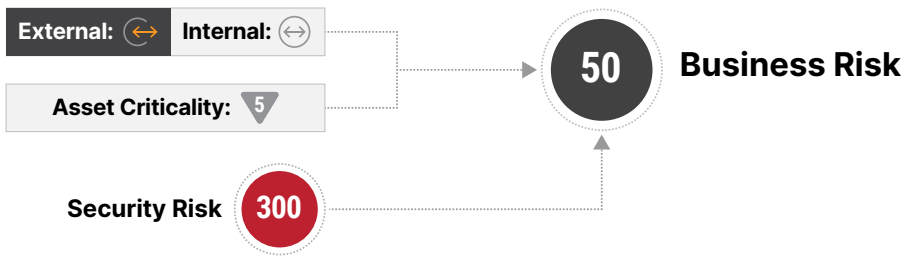

From the initial **Asset RS³** starting value of 850, the Security Risk and Business Risk penalties are applied; the remainder provides the RS³ of the asset. At minimum, Security and Business Risks may be 0, and the asset will have a perfect score of 850. At maximum, an asset may have a Security Risk of 440 and a Business Risk of 110, resulting in the lowest possible score of 300.

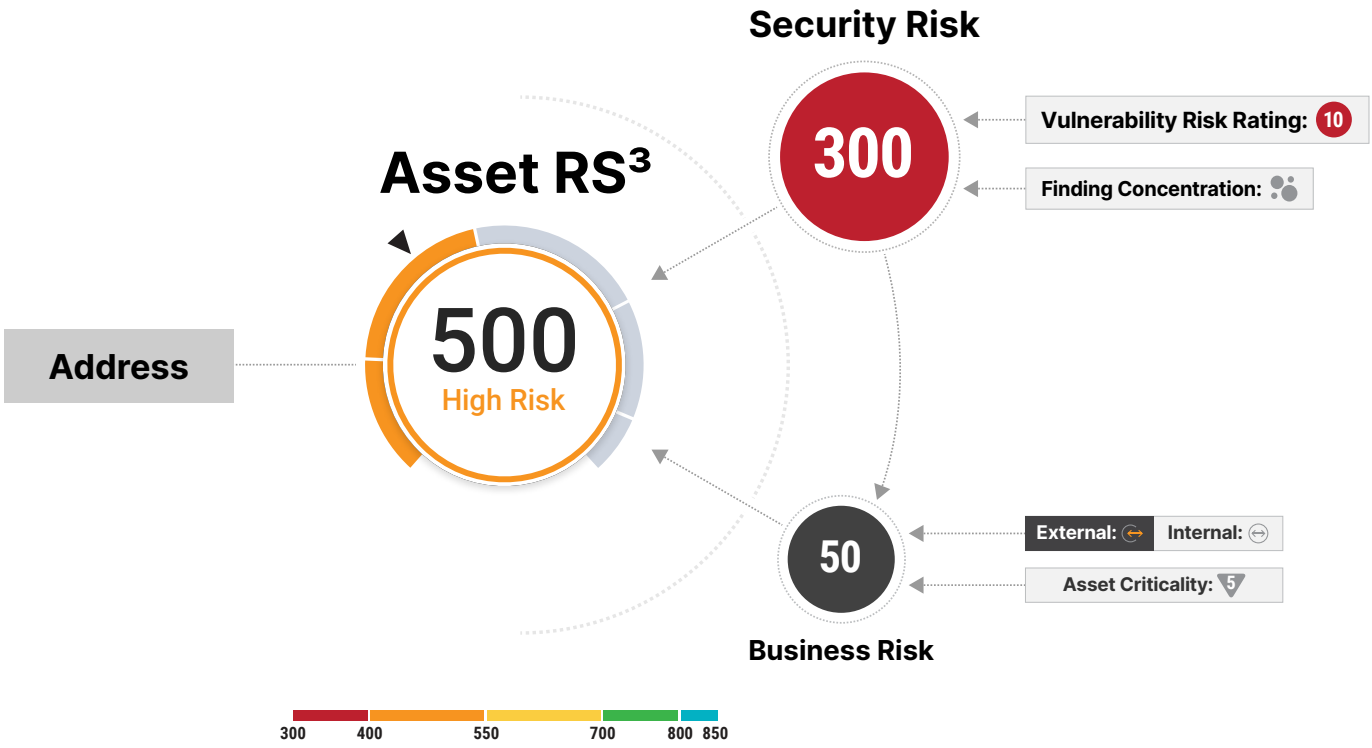

Overall RS³

An organization's Overall RS³ is a single, unified summary statistic of their overall security posture. It is expressed as a weighted average of the RS³ values of all assets in the organization, where the weights are determined by the Criticalities of the assets.[6] Higher Asset Criticality results in greater contribution to the overall score. In other words, the score of a Criticality 5 Asset will have more of an impact (counts five times in the weighted average) on the overall score compared to a Criticality 1 Asset (counts one time in the weighted average).

[6] The overall organizational score seen by users of varying access levels may differ. This is because the Overall RS³ takes into account the scores of all assets to which the user has access. If a user only has access to a subset of assets, they will see the average only of those assets' scores.

Group RS³

Groups are user-defined collections of assets in an organization. Similar to the overall score, the RS³ of a group is given by a weighted average of the RS³s of all assets within that group, where the weights are given by the Criticalities of the constituent assets.

Multi-client RS³

For organizations consisting of multiple clients within Ivanti Neurons, the Multi-client Dashboard presents an aggregated Multi-client RS³, calculated as a standard (unweighted) average of the Overall RS³s of the constituent clients.

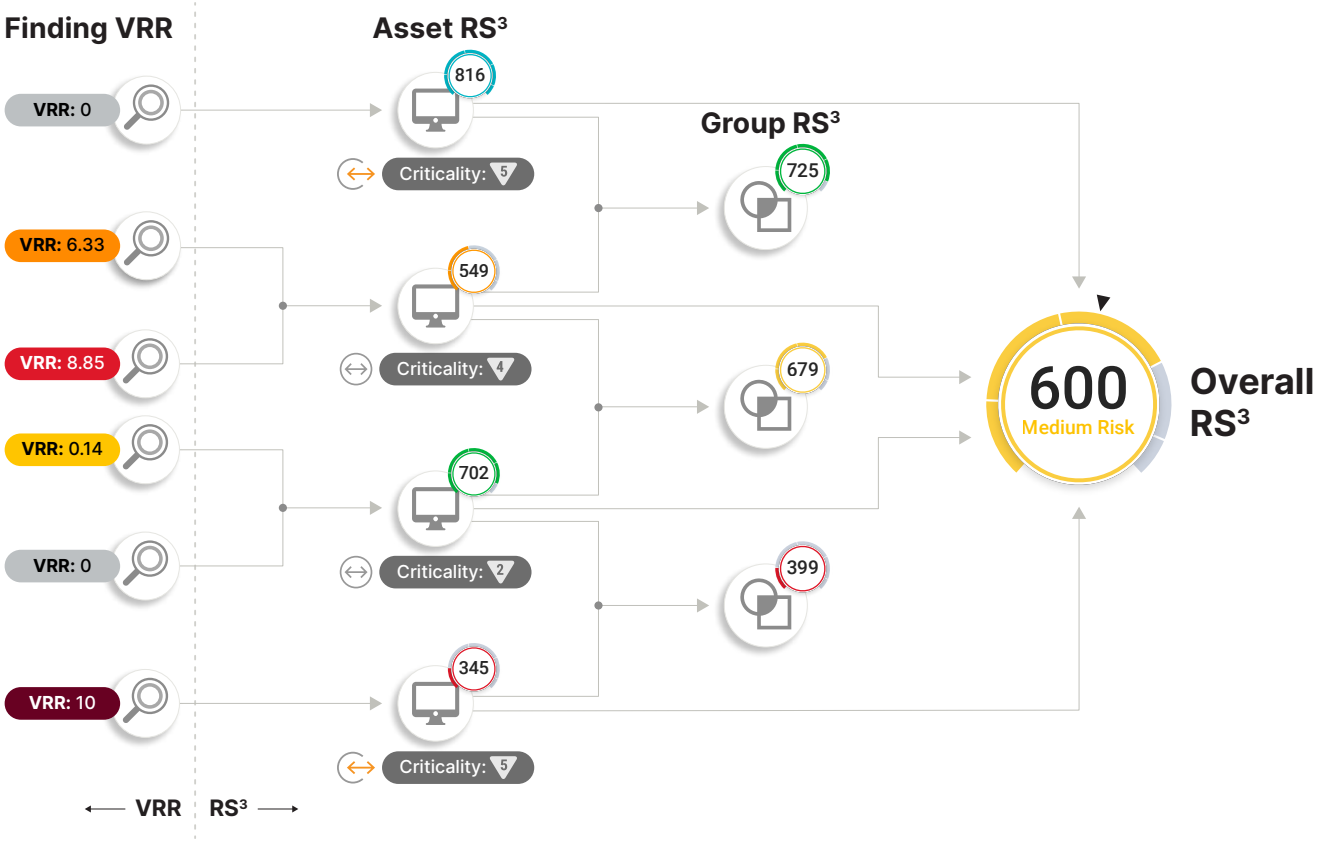

Summary

An organization's Overall RS³ represents its robustness against cyber attacks and information compromise. To derive this score, Ivanti Neurons first examines all of the vulnerabilities detected across the organization's assets. Each vulnerability is assigned a Vulnerability Risk Rating (VRR) on a scale of 0 to 10, with 10 representing the greatest amount of risk posed. This VRR is computed by identifying a base score from CVE or scanner information, and further contextualizing it with threat intelligence, trending data, and human validation. The VRR, in effect, levels the playing field and enables Ivanti Neurons to effectively prioritize asset vulnerabilities across infrastructure and applications. The constituent VRRs of vulnerabilities on an asset are aggregated together according to an adaptive weighting scheme, and then combined with the asset's Criticality and Address Type to assign an RS³ to that specific asset. Lastly, the Overall RS³ is calculated by averaging the scores of all assets (or a subset of assets, in the case of a Group RS³). In doing so, the Ivanti RS³ provides an accurate and actionable metric that continuously tracks security risk posture at an asset, group, or whole-organization level.

Appendix

Cyber Vulnerability Severity Ratings

The Common Vulnerability Scoring System (CVSS) is the industry open standard designed to convey the common attributes of vulnerabilities and assess and score security vulnerabilities in computer hardware and software systems. CVSS is under the custodianship of The National Institute of Standards and Technology (NIST).

Its quantitative model ensures repeatable accurate measurement while enabling users to see the underlying vulnerability characteristics that were used to generate the scores. Two common uses of CVSS are prioritization of vulnerability remediation activities and in calculating the severity of vulnerabilities discovered on one's systems.

The National Vulnerability Database (NVD) provides CVSS scores for almost all known vulnerabilities. The NVD supports both Common Vulnerability Scoring System (CVSS) v2 and v3 standards.

Today, customers use CVSS v3. The numerical score (0.0 to 10.0) can then be translated into a qualitative representation (Informational, Low, Medium, High, and Critical). Vulnerabilities with a base score in the range 9.0–10 are Critical, 7.0–8.9 are High, those in the range 4.0–6.9 as Medium, and 0.1–3.9 as Low. Informational vulnerabilities have a score of exactly 0.0 and do not have any threats or exploits associated with them.

The severity changes based on the scope of the vulnerability and is analyzed with the privileges an attacker needs to exploit it. CVSS Base Score Metrics Attack Vector (AV), Attack Complexity (AC), Privileges Required (PR), User Interaction (UI), Scope (S), Confidentiality Impact (C), Integrity Impact (I), and Availability Impact (A) are provided as an illustration for each vulnerability type.

Vulnerabilities/Findings		Context	
Severity	**Base Score**	**Description**	**Example**
Informational	0.0	Unconfirmed but potential existence of an attack vector which requires investigation: • One or more conditions of a vulnerability and/or security flaw were detected • Conditions require investigation to confirm or deny validity of existence	Accessible Anonymous FTP Server CVE-1999-0497: CVSS Base Score: 0.0 LOW Vector: (AV: N/AC: L/Au: N/C: N/I: N/A: N) https://nvd.nist.gov/vuln/detail/CVE-1999-0497
Low	0.1 - 3.9	Attackers may be able to collect sensitive information from the host, such as the precise version of software installed: • Identify open ports and services • Utilize information to identify known higher severity vulnerabilities • Utilize information to identify potential attack paths	Windows Information Disclosure Vulnerability CVE-2018-0747: CVSS Base Score: 1.9 LOW Vector: (AV: L/AC: H/PR: L/UI: N/S: U/C: H/I: N/A: N) https://nvd.nist.gov/vuln/detail/CVE-2018-0747
Medium	4.0 - 6.9	Attackers may be able to gain access to specific information stored on the host, including security settings: • Partial disclosure of file contents and access to certain files on the host • Disclosure of filtering rules and security mechanisms • Unauthorized use of services	Spectre/Meltdown CVE-2017-5753, CVE-2017-5715, CVE-2017-5754 Spectre/Meltdown CVE-2017-5753, CVE-2017-5715, CVE-2017-5754 Vector: (AV: L/AC: H/PR: L/UI: N/S: C/C: H/I: N/A: N) https://nvd.nist.gov/vuln/detail/CVE-2017-5753
High	7.0 – 8.9	Attackers can possibly gain control of the host, with a high potential leakage of highly sensitive information: • Unauthorized disclosure of information with full read access to files • Potential ability to create backdoors • Information of all the users on the host	WannaCry: CVE-2017-0144, CVE-2017-0145, CVE-2017-0146, CVE-2017-0147, CVE-2017-0148, CVE-2017-0148 CVSS Base Score: 8.1 HIGH Vector: (/AV: N/AC: H/PR: N/UI: N/S: U/C: H/I: H/A: H) https://nvd.nist.gov/vuln/detail/CVE-2017-0148
Critical	9.0 - 10.0	Attackers can easily gain control of the host, which can lead to the compromise of your entire network security: • Complete loss of system protection and full read and write access to files • Remote execution of commands and ability to deploy code • Ability to create backdoors	VMware Guest to Host Escape Vulnerability: CVE-2012-1516 CVSS Base Score: 9.9 HIGH Vector: (AV: N/AC: L/PR: L/UI: N/S: C/C: H/I: H/A: H) https://nvd.nist.gov/vuln/detail/CVE-2012-1516

Comparison of CVSS v2 vs. CVSS v3: The NVD provides qualitative severity rankings of "Low", "Medium", and "High" for CVSS v2 base score ranges and "Low", "Medium", "High", "Critical" for CVSS v3.[7] The severity changes based on the scope of the analyzed vulnerability and the privileges an attacker needs to exploit it.

Metric	CVSS v2 Value	CVSS v3 Value	Description
Attack Vector	Network	Network	Remotely exploitable from one or more network hops away
	Adjacent	Adjacent	Bound to network stack but limited to same shared physical or logical network
	Local	Local	Attack path is via read/write/execute capabilities, from local login or malicious user interaction
	N/A	Physical	Physical contact such as USB or peripheral access is required
Attack Complexity	Low	Low	No special access conditions exist; attacker may expect repeatable success
	Medium	N/A	Some special conditions exist such as constrained system access, or some amount of social engineering
	High	High	Depends on conditions beyond the attacker's control, such as target environment or configurations
Privileges Required	N/A	None	No authorization required to access settings or files
	N/A	Low	Basic user capabilities are required, or the affect is only against non-sensitive resources
	N/A	High	Administrative or equivalent control is required to affect any settings or files
User Interaction	N/A	None	No interaction from any user is required
	N/A	Required	Some action must be taken by a user, such as an application installation
Authentication	None	N/A	No authorization required
	Single	N/A	Attacker must be logged into the system in some way
	Multiple	N/A	Two or more authentications are required by the attacker, even with the same credentials
Scope	N/A	Changed	Resources beyond the authorization privileges intended can be affected
	N/A	Unchanged	Only resources at the same authentication level can be affected
Confidentiality Impact	Complete	High	All resources are divulged to the attacker, or disclosed information results in direct, serious impact
	Partial	Low	Access to some restricted information is obtained, but the amount or kind is constrained
	None	None	No loss of information occurs
Integrity Impact	Complete	High	Complete loss of protection of information, or file modification presents direct consequences
	Partial	Low	Modification of data is possible, but the amount or kind is constrained
	None	None	No compromise of reliability occurs
Availability Impact	Complete	High	Access to resources can be completely denied to a user in a sustained or persistent way
	Partial	Low	Resource performance or availability is reduced or temporarily disrupted
	None	None	No loss of resource availability occurs

[7] The CVSS is owned and managed by FIRST.Org Inc. (FIRST), a US-based non-profit organization, and licensed to the public freely for use. Descriptions of the CVSS metrics in the above table are paraphrased from the original CVSS documentation found here (v3) and here (v2)

About Ivanti

Ivanti makes the Everywhere Workplace possible. In the Everywhere Workplace, employees use myriad devices to access IT networks, applications and data to stay productive as they work from anywhere. The Ivanti automation platform connects the company's industry-leading unified endpoint management, zero trust security and enterprise service management solutions, providing a single pane of glass for enterprises to self-heal and self-secure devices, and self-service end users. More than 40,000 customers, including 78 of the Fortune 100, have chosen Ivanti to discover, manage, secure and service their IT assets from cloud to edge, and deliver excellent end user experiences for employees, wherever and however they work. For more information, visit ivanti.com

Ivanti RS³®

ivanti

Contact us today to learn more about Ivanti

ivanti.com | 1 800.982.2130 | sales@ivanti.com

IVI-2684