



Security Dashboard

Security Dashboard presents actionable security scores by bringing together multiple signals available in Secure Workload. It helps in understanding the current security position and improving it.

Security Dashboard acts as a springboard to many richer drill-downs within Secure Workload such as flow search, inventory search, automatic policy discovery, and forensics.

- [View the Security Dashboard, on page 2](#)
- [Security Score, on page 2](#)
- [Security Score Categories, on page 2](#)
- [High-Level View, on page 2](#)
- [Scope Level Score Details, on page 3](#)
- [Score Details, on page 5](#)

View the Security Dashboard

To view the Security Dashboard, in the left navigation pane, click **Overview**.

Security Score

Security Score is a number from 0 through 100, indicating the security position in a category. A score of 100 is the best score and a score of 0 is the worst. Scores closer to 100 are better.

The Security Score computation considers vulnerabilities in installed software packages, consistency of process hashes, open ports on different interfaces, forensic and network anomaly events, and compliance or noncompliance to policies.

Security Score Categories

There are six different score categories. Most security aspects of a workload are taken into account to come up with these categories.

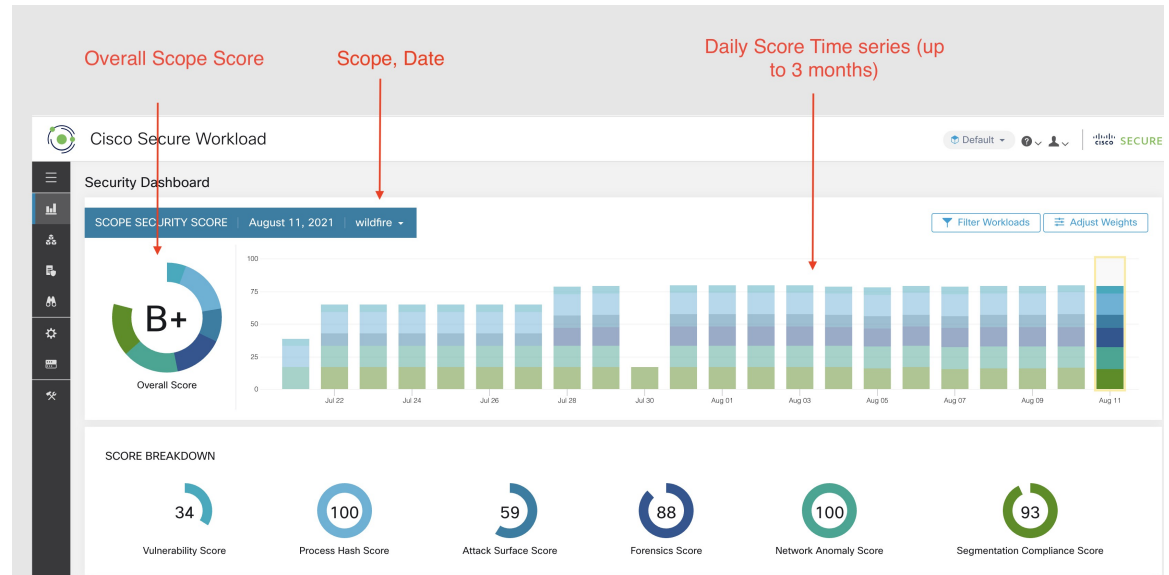
- **Vulnerability Score:** Vulnerabilities in the installed packages on a workload are used for scoring.
- **Process Hash Score:** Process hash consistency (and anomaly) along with Benign and Flagged process hashes is used for scoring.
- **Attack Surface Score:** Process may have one or more ports open on multiple interfaces to make services available. Unused open ports are used for scoring.
- **Forensics Score:** Severity of forensic events on a workload is used for scoring.
- **Network Anomaly Score:** Severity of network anomaly events on a workload is used for scoring.
- **Segmentation Compliance Score:** Compliance (permitted) and violations (escaped) to automatically discovered policies is used for scoring.

High-Level View

Security dashboard has scope level scores for the selected scope. There is an overall score with time series and score breakdown. Score details for the six score categories for the selected scope are displayed.

Scope Level Score Details

Scope Level Score details are on top of the dashboard.




The following details are displayed:

- **Overall Scope Score:** Overall score for the selected scope.
- **Daily Score Time series:** Stacked time series that can go up to 3 months.
- **Score Breakdown:** Breakdown of category scores for the selected day on time series.

Overall Score

The overall score is represented as a letter from **A+**, **A**, ..., **F**, with **A+** considered as the best score and **F** to be the worst. It is displayed as a donut chart with each slice (color-coded) representing a score category.

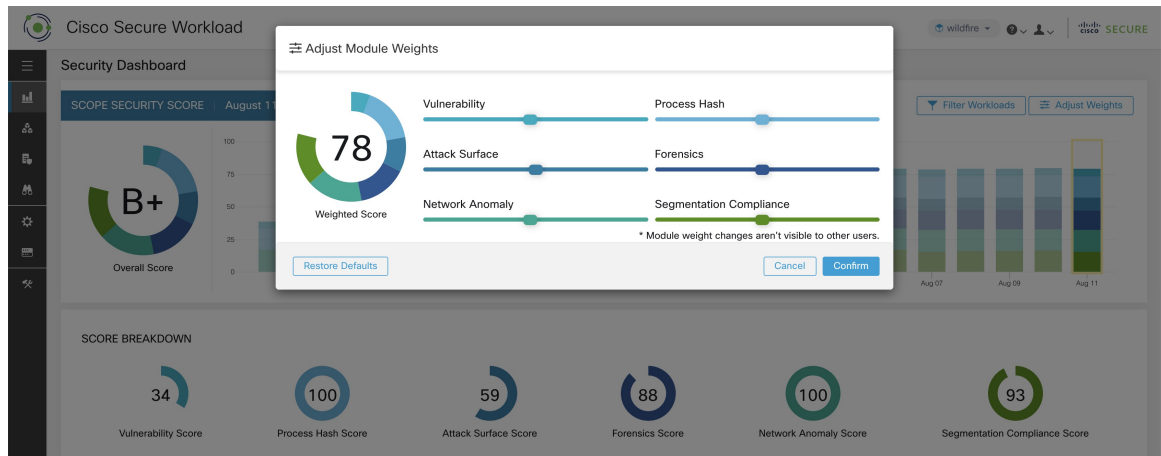


Overall Score

Overall score is the weighted average of the six score categories. By default, all weights are equal. If a score is N/A, it is considered as 0 in the overall score calculation.

$$\text{Overall score} = \frac{\sum W_{\text{category}} \times \text{Score}_{\text{category}}}{\sum W_{\text{category}}}$$

Weights can be adjusted using slides in the **Adjust Weights** module. Each user can set their own weight adjustments, which help in aligning scores with your priorities.



Important: If a score is N/A, it is considered as 0 in the overall score calculation.


Daily Time Series

Stacked time series that can go up to three months. It helps in tracking security position over a long period. Each stack represents an overall score for a day. Each segment in the stack is a category that is represented by different color. You can click on day to get the score breakdown for the day.



Score Breakdown

The Score Breakdown displays the score for all the six categories for the day that is selected on the time series. Score **N/A** indicates that score is not available. It will be counted as 0 for the overall score calculation.

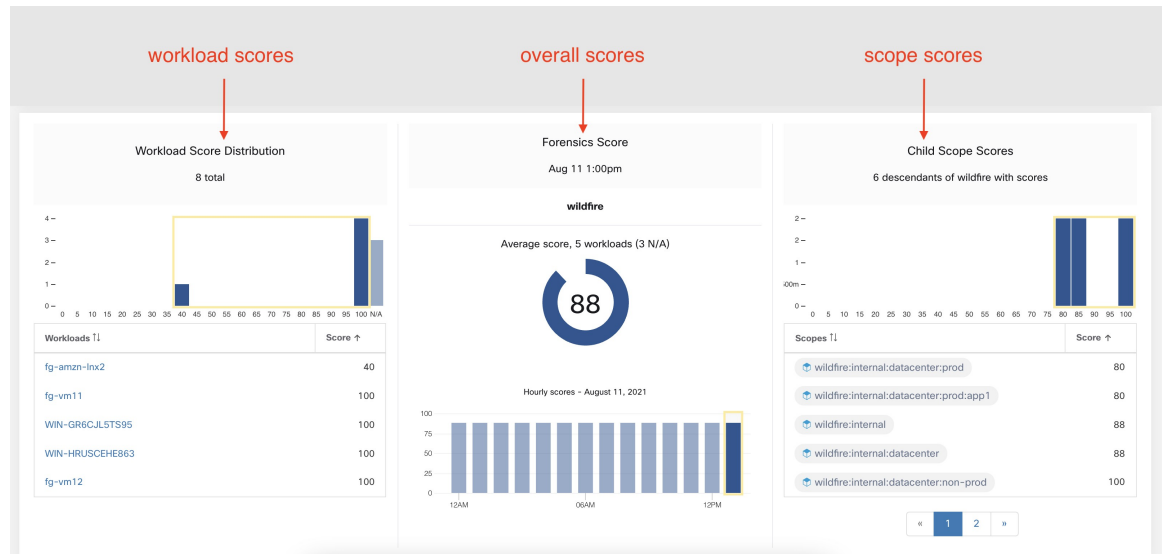


Important

If a score is **N/A**, it is considered as **0** for overall score calculation.

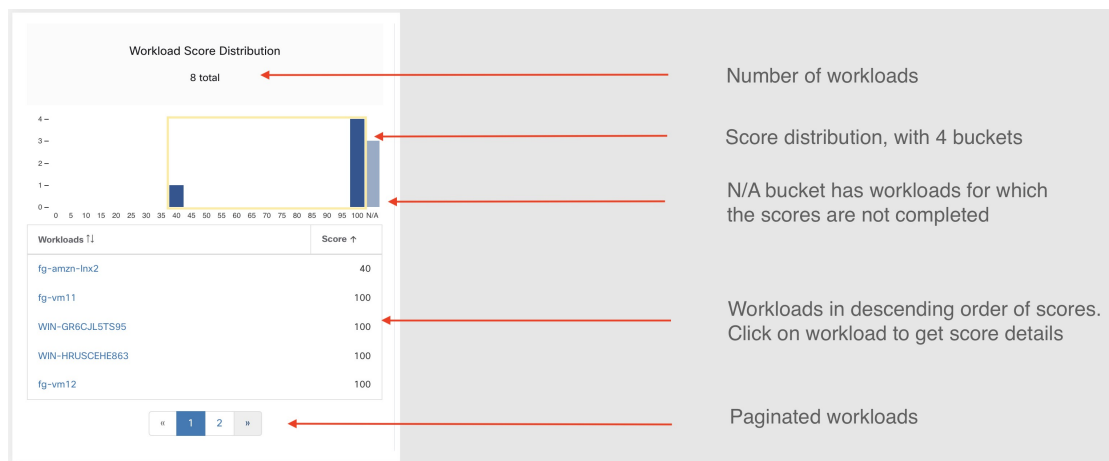
Score Details

Each of the six categories follows the following template. It has workload score distribution, hourly time series, and child scope score distribution.



Workload score distribution provides insight into score contribution from workloads under the selected scope. It helps to bubble up lowest-scoring workloads to expedite corrective actions.

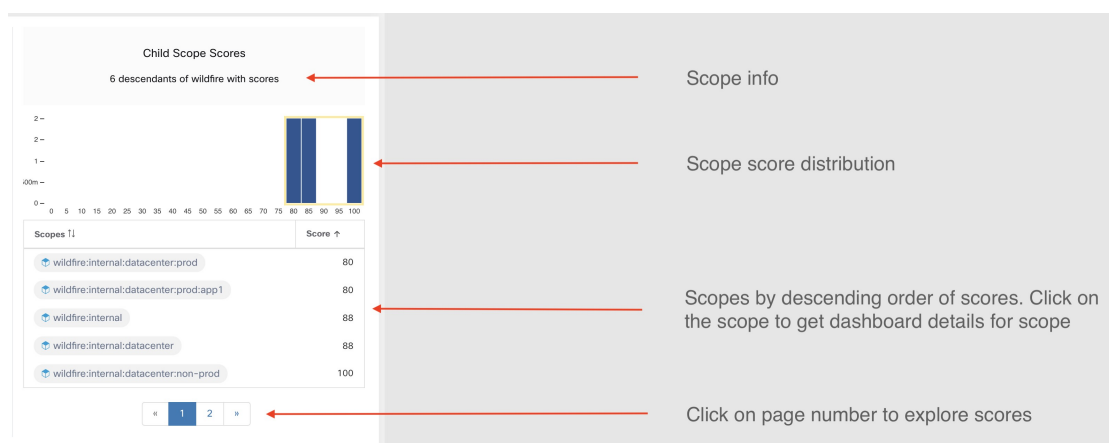
Score Details



Hourly time series helps in getting the hourly score over the course of a selected day. Selecting an hour in the hourly time series updates the workload score distribution and descendent scope distribution to show the selected hour.



Descendent scope distribution provides insight into the score contribution of child scopes of the selected scope.

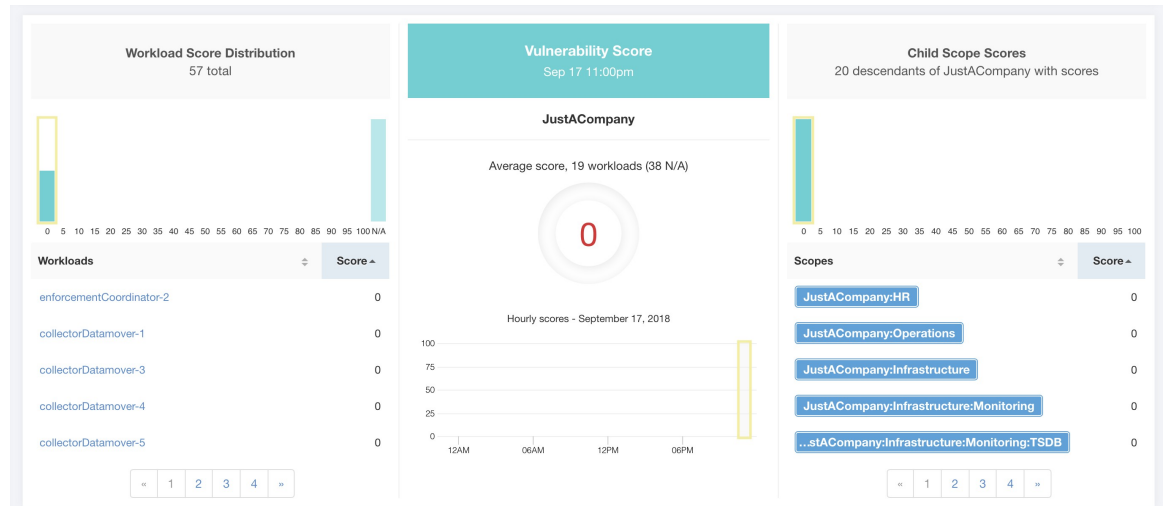


Details of each score category are explained in this section.

Vulnerability Security Score

Vulnerabilities in software packages installed on workloads are used for computing Vulnerability Security Score.

Figure 1: Vulnerability Security Score Details



Lower score indicates:

- One or more installed software packages have serious vulnerabilities.
- Apply patch or upgrade to reduce the chances of exposures or exploits


Software packages on workloads could potentially be associated with known vulnerabilities (CVE). CVSS (Common Vulnerability Scoring System) is used for assessing the impact of a CVE. CVSS score range is 0–10, with 10 being the most severe.

CVE can have CVSS v2 and CVSS v3 score. To compute Vulnerability score, CVSS v3 is considered if available, else CVSS v2 is considered.

Vulnerability score for a workload is derived from scores of vulnerable software that is detected on that workload. The Workload Vulnerability Score is calculated based on the CVSS scores, the vendor data, and may be adjusted by our security research team when data is missing or inaccurate (common for new vulnerabilities). This data is updated every 24 hours when the threat feed is configured. Higher the severity of the most severe vulnerability, lower is the score.

Scope score is average of workload scores in the scope. Improve the score by identifying workload or scopes with vulnerable software packages, and patch or upgrade with safer packages.

Figure 2: Help for Vulnerability Security Score


Vulnerability Score Help

Supported Agent Types
19 supported workloads

✗ Universal Visibility (38)

✓ **Deep Visibility (19)**

✓ **Enforcement (0)**

✗ AnyConnect (0)

✗ Hardware Switch (0)

What is a Vulnerability Score?

A Vulnerability Score is an indicator of security posture in your deployment as it relates to software package vulnerabilities. We use standard [Common Vulnerability Scoring System](#) (CVSS score) to assess the impact of a vulnerability. The Vulnerability Score is calculated based on CVSS scores of vulnerabilities detected on a workload. Like all other Security Scores, a higher score is better, with 0 meaning there is a workload that requires immediate action, and 100 meaning there are no vulnerable packages observed within this Scope.

How is the Vulnerability Score calculated?

A Workload's Vulnerability Score is derived from the scores of vulnerable software detected on that workload. We use the vulnerable package's CVSS score to assess the impact of a vulnerability. Vulnerability score of a workload depends on the most severe vulnerability present in the system; higher the severity of most severe vulnerability, lower is the workload's score. The Vulnerability Score for a Scope is the average Vulnerability score of all workloads within that Scope.

How do I improve my score?


Updating software packages on the most vulnerable workloads to versions without (or with less severe) vulnerabilities is the best way to improve the score.

How do I increase the number of workloads with scores?

Vulnerability Scores can only be calculated when Deep Visibility Sensors are present. Install Deep Visibility Sensors on more workloads to improve your score coverage.

Process Hash Score

Process hash score is an assessment of process binary hash (file hash) consistency across workloads. For example: A web server farm running Apache that is cloned from the same setup config is expected to have same hash for [httpd](#) binaries on all servers. A mismatch is an anomaly.


Figure 3: Process Hash Score Details

Lower score indicates, at least one or both of:

- One or more process hashes are flagged.
- One or more process hashes are anomalous.

Refer to [Process hash anomaly detection](#) for more details.

Figure 4: Help for Process Hash Score


Process Hash Score Help

Supported Agent Types
19 supported workloads

✗ Universal Visibility (38)

✓ AnyConnect (0)

✓ Deep Visibility (19)

✗ Hardware Switch (0)

✓ Enforcement (0)

What is a Process Hash Score?

A Process Hash Score gives an assessment of the consistency of a process binary hash across the system. For example, if you have a farm of web servers running Apache that are cloned from the same configured setup, you would expect that the hashes of [httpd](#) binaries on all servers are the same. If there is a mismatch, it is an anomaly and worth a further investigation. To reduce false alarms, we use the [NIST RDS hash dataset](#) as a whitelist. A whitelisted hash is considered "safe." You can also upload your own hash whitelist and blacklist. A blacklisted hash, if detected, will require immediate action.

Like all Security Scores, a higher score is better, with 0 meaning there is a blacklisted process hash in the system, and 100 meaning there is no hash anomaly observed in the system.

How is the Process Hash Score calculated?

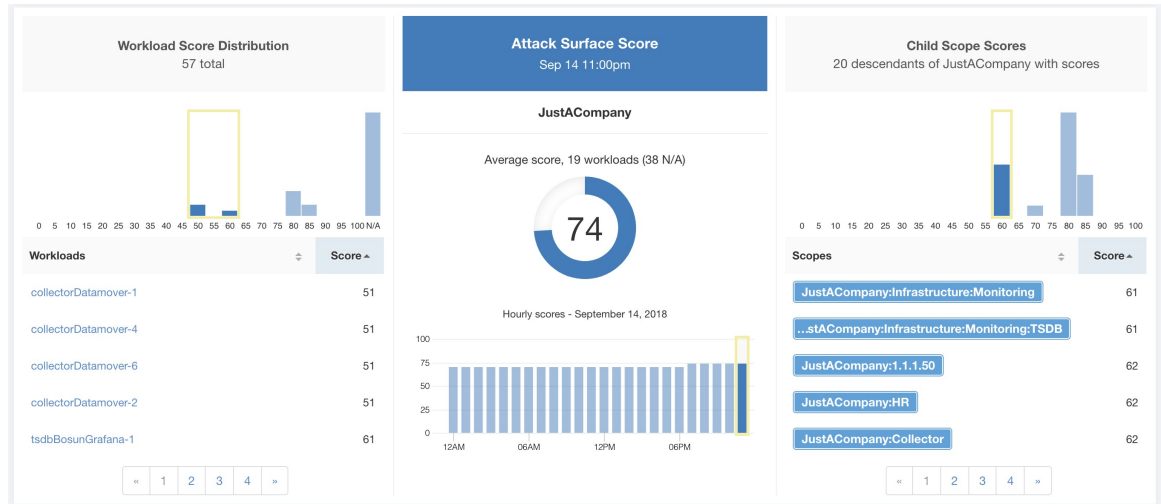
For each process hash we compute a score as follows:

1. If hash is blacklisted: score = 0
2. Else, if hash is whitelisted: score = 100
3. Else, if hash is an anomaly: score is in the range of [1, 99], the higher the better
4. Else: score = 100

Attack Surface Score

Attack Surface Score highlights potential attack surface in a workload. Open unused ports (open ports without traffic) contribute to lowering this score.

Figure 5: Attack Surface Score Details



A lower score indicates:

- Many open ports without any traffic in the last 2 weeks
- Well-known attack ports may be open and unused in last 2 weeks.
- One or more open ports are attached with packages that have serious vulnerabilities.

The attack surface score is a function of unused open ports relative to total ports, with a smoothing factor. Open ports without any traffic over the past 2 weeks are considered “unused open ports”. An extra penalty is applied to unused open ports which are well-known ports that are used in attacks (for example, 21, 22, 8080 and so on).

Figure 6: Attack Surface Score Formula

$$\begin{aligned}
 &\text{Attack surface score} \\
 &= \frac{\alpha + \sum \text{used open ports}}{\alpha + \sum \text{open ports} + (\rho * \sum \text{unused common attack ports}) + f_v(\text{vulnerability pkgs})} \\
 &\quad f_v = \max \left(\left\{ cve_{score} = \begin{cases} CVSS_{v3}, & v3 \text{ exist} \\ CVSS_{v2}, & v3 \text{ not exist} \end{cases} \right\} \right)
 \end{aligned}$$

Laplace smoothing is used with a penalty factor based on heuristic data. Score is computed daily with the past 2 weeks of data.

Tenant score is average of workload scores in the scope. Improve the score by identifying workload or scopes with unused open ports, and closing the unused ports.

When a workload link is clicked an attack surface modal is opened with details on all available ports and interfaces within the context of that workload.

33

Attack Surface Details -

Jun 19 12:00pm to Jun 19 1:00pm

22 Total Ports (12 unused ports on this workload)

Unused Ports Only ☐

These are open ports and interfaces that haven't had traffic in the last 15 days (see help for specifics). Consider closing them to reduce your attack surface (and increase your Attack Surface Score) if they aren't needed.

Port	Package Name	Total Permitted	CVE Max Score	Process Hash	Interfaces	Package Publisher	Package Version
22 (SSH)	openssh-server	16226	None	...cec50428	2	CentOS BuildSystem	5.3p1
25 (SMTP)	None	16254	None	...6ed2d10f	2	N/A	None
53 (DNS)	dnsmasq	36540	9.8	...5d28e929	2	CentOS BuildSystem	2.48
68	dhclient	N/A	None	...69235c25	1	CentOS BuildSystem	4.1.1
123 (NTP)	ntp	100425	7.5	...7c8791b1	6	CentOS BuildSystem	4.2.6p5
631	cups	N/A	7.5	...d417c9ea	1	CentOS BuildSystem	1.4.2
3128	squid	N/A	8.6	...7dc4807b	1	CentOS BuildSystem	3.1.23
5111	collector	15998	None	...a506dd9f	1	(none)	3.4.2.4f
5222	None	7999	None	...524a83d7	1	N/A	None
5640 (Tetration)	collector	N/A	None	...a506dd9f	1	(none)	3.4.2.4f

«

1

2

3

»


Features:

- **Unused Ports Only:** checkbox that when toggled filters out the ports that are used and only shows you the unused ports that are associated with the workload.
- **Columns:** Approved, port, package name, total permitted, CVE Max Score, Process Hash, Interfaces, Package Publisher, Package Version, Total Escaped, Total Rejected, Commonly Hacked Port, Links.
- **Interfaces:** If you click on any one of the line items in the Attack Surface table you can view the interfaces that are associated with each port inside a modal.
- **Approved:** checkbox that when toggled, allows you to intentionally set an “unused port” as “approved” on any one of the scopes on the scope chain that that workload has access to. Note: if a port is approved on a scope and that port is not explicitly approved on any of the children (if that scope has children), then the scope checkboxes are disabled as it is implied that any child scope that the parent scope has access to already is approved in that chain.

Approval Modal:

Edit Approval of port 22

Make sure to be as specific as you can while approving higher up the scope chain as you will be approving this port in all of its children.

- ☐ Tetration : Collector
- ☒ Tetration 
- ☐ Default

[Confirm](#) [Cancel](#)

Interfaces Modal:


Interfaces for port: 4242

Interface	Permitted *	CVE Score	PID	Escaped	Rejected	Links
0.0.0.0	8518443	None	25642	N/A	N/A	None
0.0.0.0	8518443	None	21680	N/A	N/A	None

* Based on Host Firewall

[Close](#)

Figure 7: Help for Attack Surface Score



Attack Surface Score Help

Supported Agent Types

19 supported workloads

✗ Universal Visibility (38)	✓ Deep Visibility (19)	✓ Enforcement (0)
✗ AnyConnect (0)	✗ Hardware Switch (0)	

What is an Attack Surface Score?

An Attack Surface Score is an indicator of security posture in your deployment as it relates to unused open ports on the workloads. Intuitively, the more open ports available to an attacker, the larger the attack surface. Unused ports are ones that can be easily remedied by blocking those ports if they aren't needed.

Ports are considered unused if no traffic is observed on them over the previous 2 weeks. When this feature is initially enabled - either in a new deployment (or upgrade to 3.1) or a new Deep Visibility sensor is installed on a workload - the score will gradually improve over the course of those two weeks as the system stabilizes and learns what ports are in fact unused. Scores are computed daily; newly added sensors will not have scores immediately.

Like all Security Scores, a higher score is better, with 0 meaning there is an open port on a host that needs to be immediately closed, and 100 meaning there are no unused open ports observed in the system.

How is the Attack Surface Score calculated?

The Attack Surface Score is based on the ratio of unused ports to total opened ports, with an additive smoothing to adjust the score so smaller numbers of unused ports will give better scores. E.g. 1 unused port and 2 total ports should give a better score than 100 unused ports and 200 total ports even though the ratio in both cases is 1/2.

The most well-known ports that are commonly hacked are penalized with a much greater weight since they often expose many more vectors of attack. Examples of those ports are 21-FTP, 22-SSH, 23-Telnet, and 8080, 8088, 8888, etc (which are often used for web servers).

How do I improve my score?

Currently, the only way to improve your Attack Surface Score is by closing unused interfaces and/or ports. We will be incorporating more sophisticated approaches in the future, including combining open ports with known vulnerabilities, and allowing unused ports to be present if there are policies that apply to that port.

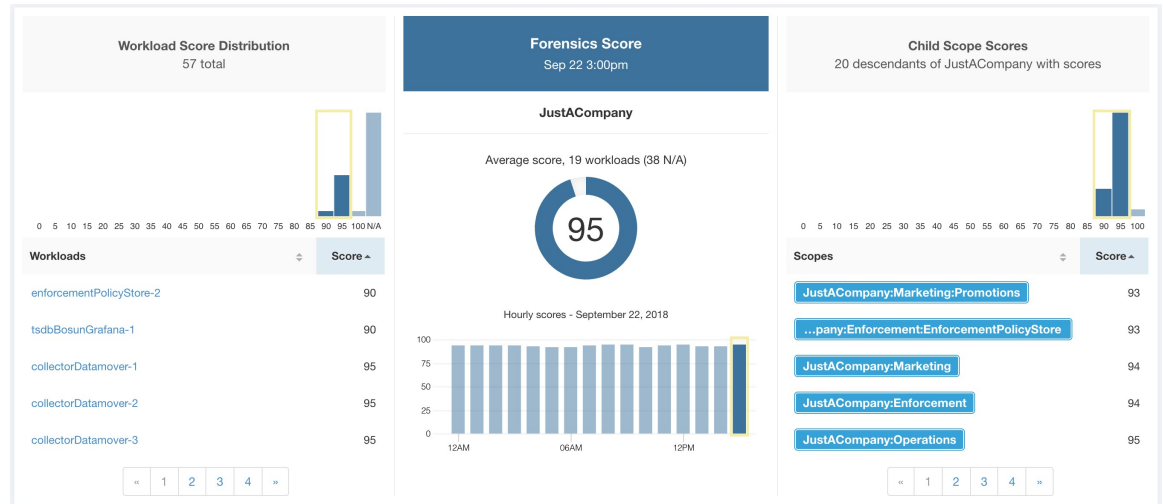
How do I increase the number of workloads with scores?

Attack Surface Scores can only be calculated when Deep Visibility, Enforcement, or AnyConnect Sensors are present. Install more of these sensors to increase your Attack Surface Score coverage.

Forensics Score

Severity of forensics events on workloads is used for computing the scores.

Figure 8: Forensics Score Details



Lower score indicates:

- One or more forensics events were observed on the workload.
- Or one/more forensics rules are noisy and/or incorrect.

To improve the score:

- Fix the issue if any to reduce the chances of exposures/exploits.
- Tweak forensics rules to reduce noise and false alarms.

Forensics score for a workload is inverse function of total impact score of forensics events. Higher is the total impact score of forensics events, lower is the forensics score.


Severity	Impact Score
IMMEDIATE_ACTION	100
CRITICAL	10
HIGH	5
CRITICAL	3

Figure 9: Forensics Score Formula

$$\text{forensics score} = \max(0, (100 - \sum \text{forensics event impact score}))$$

Refer to [Forensics](#) for more details.

Figure 10: Help for Forensics Score



Forensics Score Help

Supported Agent Types

19 supported workloads

✗ Universal Visibility (38)	✓ Deep Visibility (19)	✓ Enforcement (0)
✗ AnyConnect (0)	✗ Hardware Switch (0)	

What is a Forensics Score?

A Forensics Score is one of the Security Scores that when combined will give a simple assessment of your overall security posture. Like all other Security Scores, a higher score is better, with 0 meaning there is a workload that requires immediate action, and 100 meaning there are no Forensic Events observed within this Scope.

How is the Forensics Score calculated?

For each Workload we compute a Forensics Score. A Workload's Forensics Score is derived from the Forensic Events observed on that Workload based on the [profiles enabled for this scope](#). A score of 100 means no Forensic Events were observed, and a score of 0 means there is a Forensic Event detected that requires immediate action. The Forensic Score for a Scope is the average Workload score within that Scope.

- A Forensic Event with the severity **CRITICAL** reduces a workload's score with the weight of **10**.
- A Forensic Event with the severity **HIGH** reduces a workload's score with the weight of **5**.
- A Forensic Event with the severity **MEDIUM** reduces a workload's score with the weight of **3**.
- A Forensic Event with the severity **LOW** doesn't contribute to the Forensics Score. This is recommended for new rules where the quality of the signal is still being tuned and is likely to be noisy.
- A Forensic Event with the severity **REQUIRES IMMEDIATE ACTION** will reduce the Score for the entire Scope to zero.

How do I improve my score?

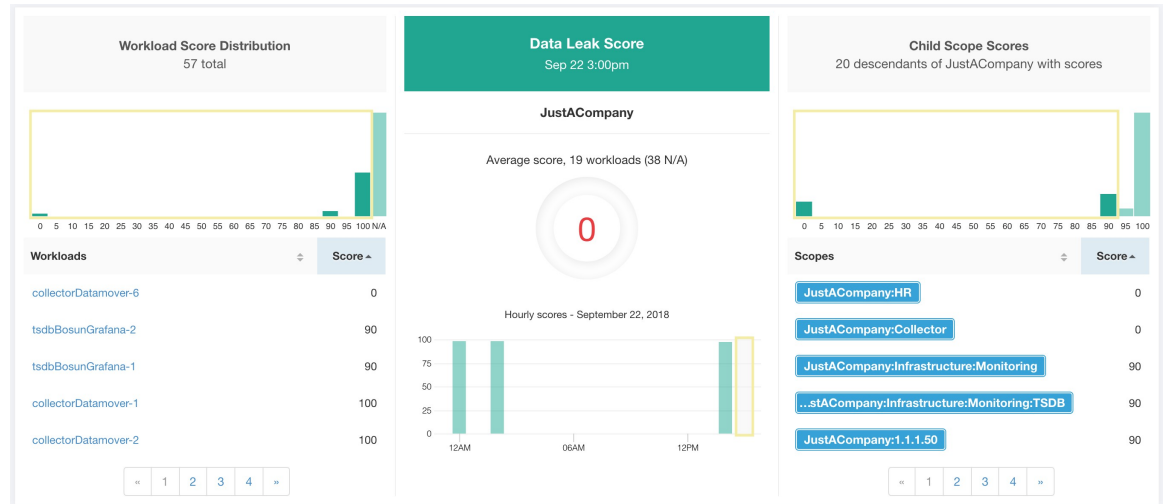
Tuning your Forensics Score can be done by adjusting the Forensic Rules [enabled for this Scope](#). Creating rules that are less noisy will give you a more accurate score. Acting upon and preventing legitimate Forensic Events (events that are evidence of an intrusion or other bad activity) is another good way to improve your Forensic Score.

How do I increase the number of workloads with scores?

See the compatibility chart above for which sensor types are compatible. Installing the supported sensor types on more Workloads will increase your Forensic coverage.

Network Anomaly Score

Severity of Network Anomaly events on workloads is used for computing the scores.

Figure 11: Data Leak Score Details

Lower score indicates:

- Unusually high amount of data is being transferred out of workloads.
- Or Network Anomaly forensic rule is incorrect or noisy.

To improve the score:

- Fix the issue if any to reduce the chances of data exfiltration.
- Adjust Network Anomaly rules to reduce noise and false alarms.

Network Anomaly score for a workload is inverse function of total severity score of Network Anomaly events. Higher is the total severity score, lower is the Network Anomaly score.


Severity	Score
IMMEDIATE_ACTION	100
CRITICAL	10
HIGH	5
CRITICAL	3

Figure 12: Data Leak Score Formula

$$\text{data leak score} = \max(0, (100 - \sum \text{data leak event severity score}))$$

Refer to [PCR-based Network Anomaly detection](#) for more details.

Figure 13: Help for Data Leak Score



Data Leak Score Help

Supported Agent Types

19 supported workloads

✗ Universal Visibility (38)	✓ Deep Visibility (19)	✓ Enforcement (0)
✓ AnyConnect (0)	✗ Hardware Switch (0)	

What is a Data Leak Score?

A Data Leak Score gives you an assessment of whether there are any symptoms of unusually significant amounts of data being transmitted out of your workloads. Like all Security Scores, a higher score is better, with 0 meaning there is a workload that requires immediate action, and 100 meaning there are no Data Leak Events observed within this Scope.

How is the Data Leak Score calculated?

The Data Leak Score is also computed similarly to the Forensics Score. For each Workload we compute a Data Leak Score. A Workload's Data Leak Score is derived from the Data Leak Events observed on that Workload based on the profiles enabled for this scope. A score of 100 means no Data Leak Events were observed, and a score of 0 means there is a Data Leak Event detected that requires immediate action. The Data Leak Score for a Scope is the average Workload score within that Scope.

- A Data Leak Event with the severity CRITICAL reduces a workload's score with the weight of 10.
- A Data Leak Event with the severity HIGH reduces a workload's score with the weight of 5.
- A Data Leak Event with the severity MEDIUM reduces a workload's score with the weight of 3.
- A Data Leak Event with the severity LOW doesn't contribute to the Data Leak Score. This is recommended for new rules where the quality of the signal is still being tuned and is likely to be noisy.
- A Data Leak Event with the severity REQUIRES IMMEDIATE ACTION will reduce the Score for the entire Scope to zero.

How do I improve my score?

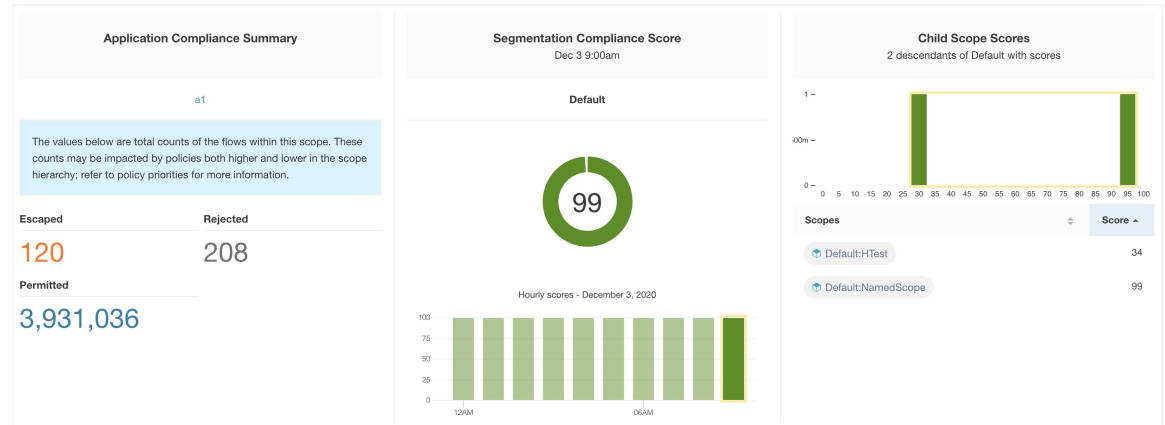
Tuning your Data Leak Score can be done by adjusting the Forensic Rules for Data Leak Events enabled for this Scope. Creating rules that are less noisy will give you a more accurate score. Acting upon and preventing legitimate Data Leak Events (events that are evidence of anomalous exfiltration activities) is another good way to improve your Data Leak Score.

How do I increase the number of workloads with scores?

Data Leak Scores can only be calculated when Deep Visibility Sensors are present. Install Deep Visibility Sensors on more workloads to improve your score coverage.

Segmentation Compliance Score

Segmentation Compliance Score presents a top-level view of policy violations and emphasizes which scopes and workspaces have the most violations.

Figure 14: Segmentation Compliance Score Details

Note Escaped/Rejected/Permitted count that is displayed on security dashboard for root scope does not add up to all the counts respectively displayed for all child scopes. Escaped/Rejected/Permitted count is an evaluation on the policy and not just on source or destination.

Lower score indicates:

- Significant number of escaped flows (policy violations) relative to permitted
- Score is 0 when more escaped flows than permitted.

Segmentation Compliance Score is computed for scopes with an enforced primary workspace. For scopes without enforced workspaces, the score will be computed as the average of descendant scope scores with enforced policies.

Score is computed by using the ratio between escaped and permitted.

Figure 15: Segmentation Compliance Score Formula

$$\text{compliance score} = \left[100 - \frac{100 \times \text{escaped}}{\text{permitted}} \right]$$

Improve score by reducing number of policy violations

- Verify policies correctly cover desired behavior.
- Verify that policies are correctly being enforced.

Figure 16: Help for Segmentation Compliance Score Details

? Segmentation Compliance Score Help

Supported Agent Types 5,059 supported workloads

✓ Universal Visibility (8)

✓ Deep Visibility (23)

✓ Enforcement (25)

✓ AnyConnect (5,002)

✓ Hardware Switch (1)

What is a Segmentation Compliance Score?

A Segmentation Compliance Score is an indication of how effectively enforced Applications are based on observed Rejected and Escaped flows. Rejected and Escaped flows are a sign that enforcement isn't reliable and should be investigated. This score is only applicable if you have Applications with policies that are enforced.

How is the Segmentation Compliance Score calculated?

Segmentation Compliance differs from the other modules in that the score applies only to Scopes and not to specific workloads. If the Scope has an enforced Application, the score is derived from the number of Rejected and Escaped flows relative to the total number of flows observed. The counts are displayed in the left pane, clicking them will take you to the enforced application view. For Scopes that don't have an enforced application, the score is the average of the child scope scores.

How do I improve my score?

Investigating and reducing the number of Rejected and Escaped flows will improve and increase your Segmentation Compliance Score.

How do I increase the number of Scopes with scores?

Create more Enforced Applications will increase your Segmentation Compliance coverage.