

NEW TO YOUR HOME DASHBOARD

Your Security Posture Score

Actionable intelligence, one grade for your organisation, and the specific next steps that improve it.


You will see something new when you sign in to AppCheck and access your Home Dashboard: a single security posture score for your organisation, graded A to F, sitting alongside a short list of the actions that would improve it most. This note explains what the score measures, how it is worked out, and how to get the most out of it.

From findings to actionable intelligence

A vulnerability scanner is very good at telling you what is wrong. It is less good at telling you what to do first. Most teams we work with are not short of findings; they are short of an answer to a simpler question: of everything we could do this month, what would move the needle most?

The security posture score is our answer. It reads the data your account already holds, turns it into one number you can act on, and then does the part that matters: it tells you which few things to change, and what each one is worth. Nothing new is collected, and nothing changes about how your scans run.

What you will see



An illustrative example. Your card will show your own figures.

The card has 4 parts

- **Your grade and score**, out of 100, with a line showing what organisations of a similar size typically achieve.
- **Where the points went**, broken down across five components, so you can see at a glance which behaviour is costing you.
- **Address now**, a panel for open critical and high-severity findings. These do not move the grade, but they are the first thing on the card that asks for a decision.
- **Suggested next steps**, up to three actions ranked by the points each would recover, so the top item is always the cheapest real improvement.

How the score is calculated?

Five components, 100 points. Every measure is expressed per asset, so the grade reflects how well you are covered rather than how large your estate is, a four-asset firm and a three-thousand-asset enterprise are graded on the same scale.

COMPONENT	POINTS	WHAT IT MEASURES
Scope coverage	20	The share of your declared scope that appears in at least one live scan definition.
Scan currency / frequency	25	How recently you last scanned. Full marks within seven days, tapering to nothing beyond a year.
Scan cadence	15	How many separate days you scanned in the last 180. Regular testing scores better than one large burst.
Open exposure	30	Open critical and high findings per asset, with criticals weighted triple. Capped when the most recent scan is stale, because a clean result cannot be claimed from an old scan.
Unreviewed discovery	10	Assets that discovery has surfaced but that have never been accepted into scope or dismissed.

Grades

Grades are absolute, not a ranking, A at 86 and above, B at 70, C at 57, D at 39, and F below that. Nobody has to fall for you to rise, this is not a leaderboard, so any real improvement you make shows up in your grade.

What is not counted against you

- Findings you have triaged as false positives or accepted risks are not treated as open exposure. Triage is the behaviour we want to encourage, not penalise.
- Dismissing a discovered asset clears it from the discovery backlog just as accepting it into scope does. Either decision resolves it.
- If your account is not yet set up, with no scope defined and no scans run, you will see a short onboarding prompt rather than a failing grade.

The part that matters, your next three actions

Every card generates its own actions from whichever components are furthest from full marks, each carrying the points it would recover. In the example above, they read:

Scan within the last 7 days	+4 points
------------------------------------	------------------

Your most recent scan was 15 days ago. Scanning weekly earns the full currency score.

Clear the outstanding critical	+4 points
---------------------------------------	------------------

One critical finding across 235 assets. Resolving it takes open exposure to full marks.

Review the assets awaiting a decision	+2 points
--	------------------

Accept them into scope or dismiss them. Either resolves the backlog.

The list is short on purpose, something you can work to act on between one scan and the next, not another backlog.

Two numbers on your open criticals

Alongside the grade, the card separates your open critical findings into two groups, those a recent scan has confirmed are still present, and those that remain open but have not been re-tested since they were first found.

The distinction matters. It is common to scan diligently every week and still leave the assets carrying your oldest findings untouched for months. Those findings may have been fixed, or they may not: without a re-scan there is no way to tell. Pointing your next scan at them is usually the fastest way to turn a long list into a short one.

Where you stand

Your card shows the typical score for organisations with a similar number of assets, so an A means more than a number in isolation. The comparison is aggregated and anonymous, no other organisation is identified to you, and your score is never shown outside your own account.

What the score is not

A measure of coverage, not a certificate of security

The score reflects how thoroughly you are using AppCheck to cover and re-check your estate. It is not a certification, an audit outcome, or a guarantee that you are secure.

The console can only report on what is in scope, which is why coverage carries weight and why looking properly scores as well as a clean result.

What you need to do

Nothing. The score appears on your dashboard automatically at the end of Q3, and because it is calculated from data your account already holds, you will have a grade and a set of suggested actions from the first day. If you would like to talk yours through, or want help planning the first few actions, your account manager will be glad to walk through it with you.