



NxSAM Cloud Security (CSPM) — Quick Guide

Quick reference guide for the Cloud Security (CSPM) module

Product release	**NxSAM 2026.09**
Verified against code on	**2026-09-10**
Document revision	**b441b08-dirty**
Audience	**End users**
Document owner	**eSec Forte Technologies — Product Documentation**

Screenshots are taken from a demonstration data set. Figures are illustrative.

eSec Forte Technologies

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Table of Contents

Cloud Security (CSPM)

- Overview
- Cloud Security (CSPM): every screen and column
- Cloud Security (CSPM): find the right screen

Reference

- Glossary

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Cloud Security (CSPM)

Cloud Security (CSPM)

Cloud Security Posture Management covers how your cloud accounts are configured — storage that is publicly readable, over-permissive roles, unencrypted databases, security groups open to the world.

Why cloud posture is its own problem

Cloud misconfigurations differ from other findings in three ways that change how you work with them:

They are created by anyone with access. A single console click can expose a storage bucket. There is no build or release to catch it.

They are immediately exploitable. A publicly readable bucket is not a weakness someone might exploit — it is already accessible to anyone who finds it.

They drift constantly. Cloud estates change daily, so posture is a moving picture rather than a periodic assessment.

What the module holds

Cloud posture reaches NxSAM from more than one platform, and the module has a set of screens for each, because the platforms describe posture differently.

Prisma CSPM — four screens, fed by Prisma Cloud:

Screen	Shows
Asset Inventory	The cloud resources discovered in your accounts
Alerts	Resources currently violating a policy
Compliance	How your cloud estate measures against frameworks

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Screen	Shows
Policies	The rules being evaluated

CSPM — five screens, fed by Cortex CSPM and SentinelOne CSPM:

Screen	Shows
Assets	The cloud resources discovered, with their open issues and cases
Issues	Individual findings — one rule failing against one or more assets
Cases	Issues grouped into a single piece of work, with a severity of their own
Rules	The individual checks that raise issues
Policies	Named groups of rules, and whether each group is switched on

The two sets do not merge. A resource that both platforms see appears once under each, and the counts on one set say nothing about the other.

Rules, issues and cases

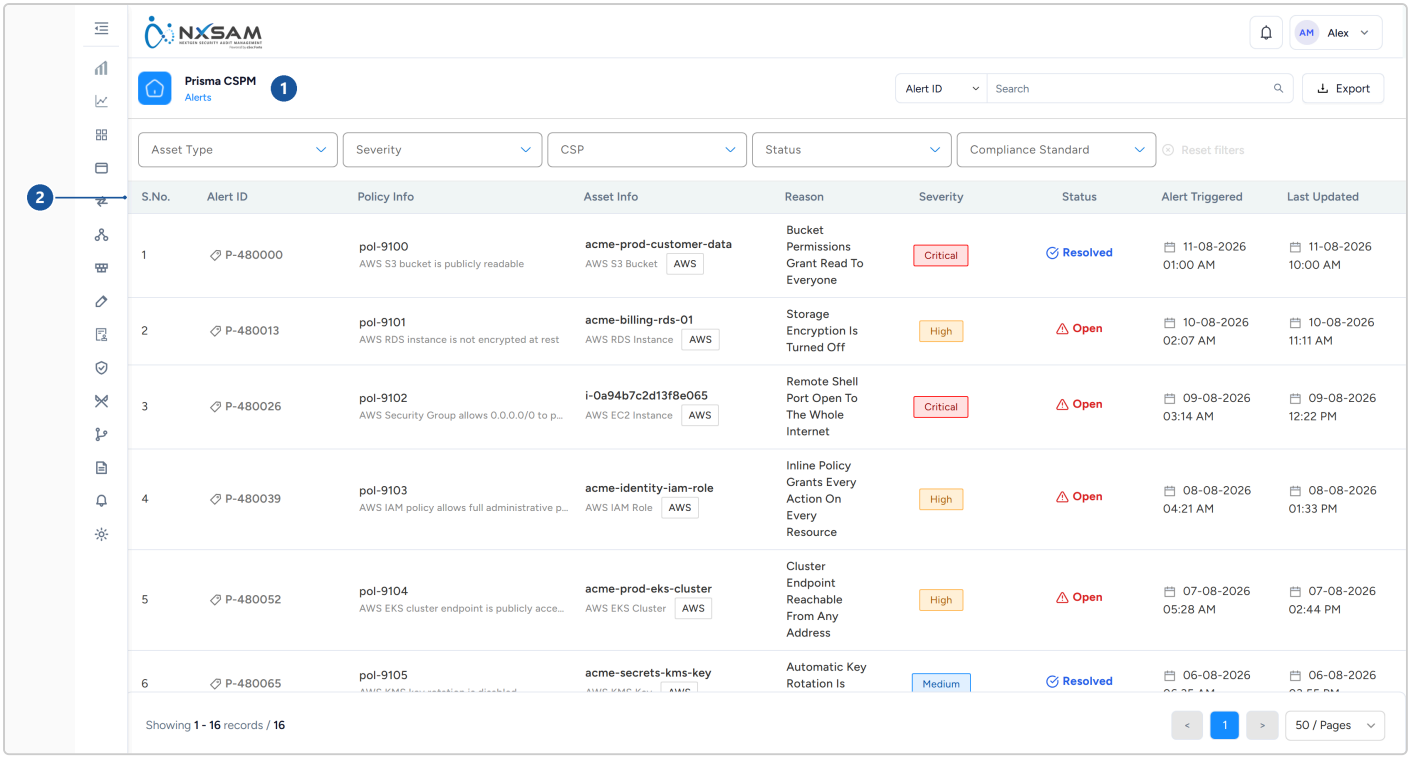
The CSPM screens use three words that are easy to run together, and the module is much easier to read once they are kept apart:

- A **rule** is a single check — "storage must not be publicly readable".
- An **issue** is that rule failing against a resource. One rule can raise many issues; the issue row says how many assets it touches.
- A **case** collects related issues into one thing to work on, and carries its own severity and status. Fixing a case is fixing every issue in it.

Every row on those three screens links sideways — an asset to its issues, an issue to its case and its rule, a rule to its issues, a policy to its rules — so you can start from whichever you were asked about and reach the others without searching.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Alerts, not vulnerabilities



Cloud posture alerts, each tying a policy to an asset and account.

1 where you are · **2** policy, asset and account

A CSPM alert says a resource is currently in a state a policy forbids. That is a statement about *now* — fix the configuration and the alert clears on the next evaluation.

This differs from a vulnerability, which describes a flaw that persists until patched. It is why cloud alert counts move far more than vulnerability counts, and why a spike is often one change affecting many resources rather than a genuine deterioration.

Reading alerts well

- **Public exposure first.** Anything reachable from the internet outranks internal misconfigurations regardless of the severity label.
- **Identity and access next.** Over-permissive roles are the route from a small foothold to a large incident.
- **Look for patterns.** The same misconfiguration across many resources usually means a template, module or provisioning pipeline is producing it. Fix the source and the resources follow.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Fix the pattern, not the resource

Correcting fifty individually misconfigured buckets achieves nothing durable if the template that created them is unchanged — the next fifty arrive the same way.

Compliance view

The compliance screens map cloud configuration against frameworks, which is what makes cloud posture reportable to an auditor rather than only to engineers.

As with policy compliance, read the failures rather than the percentage: a high score can coexist with a small number of severe exposures.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Cloud Security (CSPM): every screen and column

Two families of screens, and no row actions on either — CSPM is a reporting surface rather than a working queue, because the remediation happens in the cloud provider rather than here.

- **Prisma CSPM** — Asset Inventory, Alerts, Compliance Standards and Policies, fed by Prisma Cloud. Described first, below.
- **CSPM** — Assets, Issues, Cases, Rules and Policies, fed by Cortex CSPM and SentinelOne CSPM. Described from CSPM: Assets onwards.

The two families do not share data. A resource both platforms see appears once in each, and the two Policies screens are different lists with the same name.

Asset Inventory

Everything the posture tool discovered across your cloud accounts.

Column	What it holds
Asset details	The resource name and its provider identifier
Asset Type	S3 bucket, EC2 instance, SQL database, storage account, and so on
CSP Details	Which cloud, which account, which region
Status	Whether the resource is currently present
Alerts Distribution	Open alerts against it, split by severity
Total Policies	How many checks apply to this resource type

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Column	What it holds
Last Scanned Date	When posture was last evaluated

Total Policies is the column that explains an empty alert count. A resource with two applicable policies and no alerts is genuinely clean; one with forty and no alerts has probably not been evaluated yet — check Last Scanned Date before concluding it is compliant.

Alerts

The findings. One row is one policy failing against one resource.

Column	What it holds
Alert ID	The provider's identifier, for cross-referencing in their console
Policy Info	Which check failed
Asset Info	Which resource, in which account
Reason	Why it failed — the actual configuration found
Severity	As the policy defines it
Status	Open or resolved
Alert Triggered	When it first fired
Last Updated	When it last changed

Reason is the column people skip and should not. Two resources failing the same policy often fail it for different reasons — one bucket is public to everyone, another is public to a single partner account — and the remediation differs accordingly.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Alert Triggered against Last Updated tells you whether an alert is stable or flapping. An alert triggered months ago and updated yesterday is being re-evaluated and still failing; one triggered and updated at the same moment has fired once.

Compliance Standards

Which benchmarks posture is measured against.

Column	What it holds
Standard Name	CIS AWS Foundations, PCI DSS, ISO 27001, and so on
Standard Description	What the standard covers
Mapped Policies	How many individual checks make it up
Enabled	Whether it is switched on

Enabling a standard does not run anything

Standards select which policies count towards a compliance figure. The policies themselves evaluate regardless. Turning a standard on therefore changes the reported percentage immediately, with no scan involved — which surprises people who expect a delay.

Policies

The individual checks behind the standards.

Column	What it holds
Policy Details	The check and what it looks for
CSP	Which cloud it applies to

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Column	What it holds
Severity	What an alert from it is rated
Status	Whether the policy is active
Alerts	How many are currently open against it
Auto-Fix	Whether the provider can remediate it automatically
Last Evaluated	When it last ran
Created	When the policy was defined

Auto-Fix is the most actionable column in the module. A policy marked auto-fixable can be remediated by the cloud provider without a change request — so a list of high-alert auto-fixable policies is usually the fastest available posture improvement.

It is also worth checking before escalating: an alert on an auto-fixable policy that keeps returning usually means auto-fix is available but not enabled, rather than that the fix is not working.

CSPM: Assets

The cloud resources the posture platform has discovered, with what is currently wrong with each. Rows are the resource; the two count columns are doors into the other screens.

Column	What it holds
Asset Details	The resource name, and the identifier the platform knows it by
Asset Type	The kind of resource — a bucket, an instance, a database, a role
Provider	Which cloud it lives in
Issues →	Open issues against it, by severity. The arrow opens the Issues screen filtered to this asset

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Column	What it holds
Cases →	Cases that include it, by severity. The arrow opens the Cases screen filtered to this asset
Last Observed	When the platform last saw the resource

The severity chips above the table — **Critical**, **High**, **Medium**, **Low** — filter the list to assets carrying an issue of that severity. **Asset Type** and **Provider** are the quick filters; **More Filters** opens the full set, and **Reset Filters** clears everything at once. **Export** downloads the current list as shown.

Clicking the asset name opens **Asset Details**: the provider and its own identifier, the realm and region, the type as a category, class and name, when the resource was first and last observed, and the issues and cases against it broken down by severity.

Last Observed is what tells you whether a clean asset is clean. A resource with no issues and a Last Observed date from last night has been checked and passed. One with no issues and a Last Observed date from three weeks ago has stopped being seen — it may have been deleted, or the account it lives in may have stopped syncing — and its empty issue count is not evidence of anything.

CSPM: Issues

The findings. One row is one rule failing against one or more assets.

Column	What it holds
Issue name	What was found, and the platform's identifier for it
Severity	As the rule rates it
Domain / Type	The area of posture the issue belongs to, and the kind of issue it is
Detection Method	How the platform found it
Status	Where the issue is in its lifecycle

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Column	What it holds
Asset Count	How many resources this one issue touches
Case ID	The case the issue has been grouped into. Opens the Cases screen on that case
Rule ID	The rule that raised it. Opens the Rules screen on that rule

Quick filters: **Severity** and **Status**. **More Filters** adds domain, category and detection method. **Export** and **Reset Filters** behave as on Assets.

Clicking the issue name opens **Issue Details**: the issue's identifier and name, its domain, detection method, asset count, the case and rule it belongs to, and when it was last updated.

Asset Count is the column to sort by when deciding what to fix first. An issue on one asset is one change. An issue on two hundred assets is almost always a template, a module or a provisioning pipeline producing the same misconfiguration repeatedly, and fixing the source clears all two hundred at once — while fixing the resources one at a time clears nothing durably.

CSPM: Cases

Issues grouped into one piece of work. A case is what gets assigned, tracked and closed; the issues inside it are the evidence.

Column	What it holds
Case name	What the case is about, and its identifier
Severity	The case's own rating, which is not necessarily the highest issue in it
Status	Where the case is
Domain	The area of posture it covers

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Column	What it holds
Issues →	The open issues in it, by severity. The arrow opens the Issues screen filtered to this case
Modified	When the case last changed

The severity chips filter cases by their own severity. Quick filters: **Domain**, **Severity**, **Status**. **Export** and **Reset Filters** as elsewhere.

Clicking the case name opens **Case Details**: identifier and name, domain, when it was created and last modified, and its issues broken down by severity.

A case's severity is set by the platform, not summed from its issues

A case rated High can contain a Critical issue, and a case rated Critical can contain only Medium ones — the platform rates the case on what the combination means, not on the worst member. When the two disagree, the case rating is the one the platform intends you to act on.

CSPM: Rules

The individual checks. Every issue on the Issues screen was raised by one of these.

Column	What it holds
Rule name	What the rule checks, and its identifier
Severity	What an issue from it is rated
Type	The kind of rule
Providers	Which clouds it applies to
Status	Whether the rule is enabled

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Column	What it holds
Issues Count	Open issues it has raised. Opens the Issues screen filtered to this rule
Updated	When the rule definition last changed

Quick filters: **Provider**, **Rule Class**, **Severity**, **Type**; **More Filters** for the rest. Clicking the rule name opens **Rule Details**: identifier, name, class, type, providers, issue count, and when the rule was created and last modified.

A disabled rule with a non-zero issue count is a rule that used to run. Its issues are historical — nothing is re-evaluating them — and they will neither clear nor grow. Either re-enable the rule in the platform or treat those issues as closed; leaving them counted as open inflates every total above them.

CSPM: Policies

Named groups of rules, as configured in the posture platform. A policy is how the platform decides which rules run where, so this screen explains why a rule that exists is or is not producing issues.

Column	What it holds
Policy name	The group, and its identifier
Mode	How the platform applies the policy, as set in the platform
Rule Matching	How the policy selects its rules, as set in the platform
Labels	The platform's tags on the policy
Status	Whether the policy is enabled
Rules Count	How many rules it contains. Opens the Rules screen filtered to this policy

Quick filters: **Mode**, **Rule Matching**, **Status**. Clicking the policy name opens **Policy Details**: identifier, name, description, mode, rule matching, status, rule count, and when it was created and last updated.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Mode and **Rule Matching** are the platform's own settings, shown here so you can read them without opening the platform. NxSAM does not change them.

Why there are no row actions

Everything here is evidence about a cloud account that NxSAM reads and does not control. Remediation happens in the provider's console or in the infrastructure code that produced the resource. The module's job is to say what is wrong, where, and how it compares to the standard you are being measured against — not to change it.

The CSPM screens add the sideways links — asset to issues, issue to case and rule, rule to issues, policy to rules — which is a way of *reading* the evidence rather than acting on it. Nothing on any of the nine screens writes back to a cloud account.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

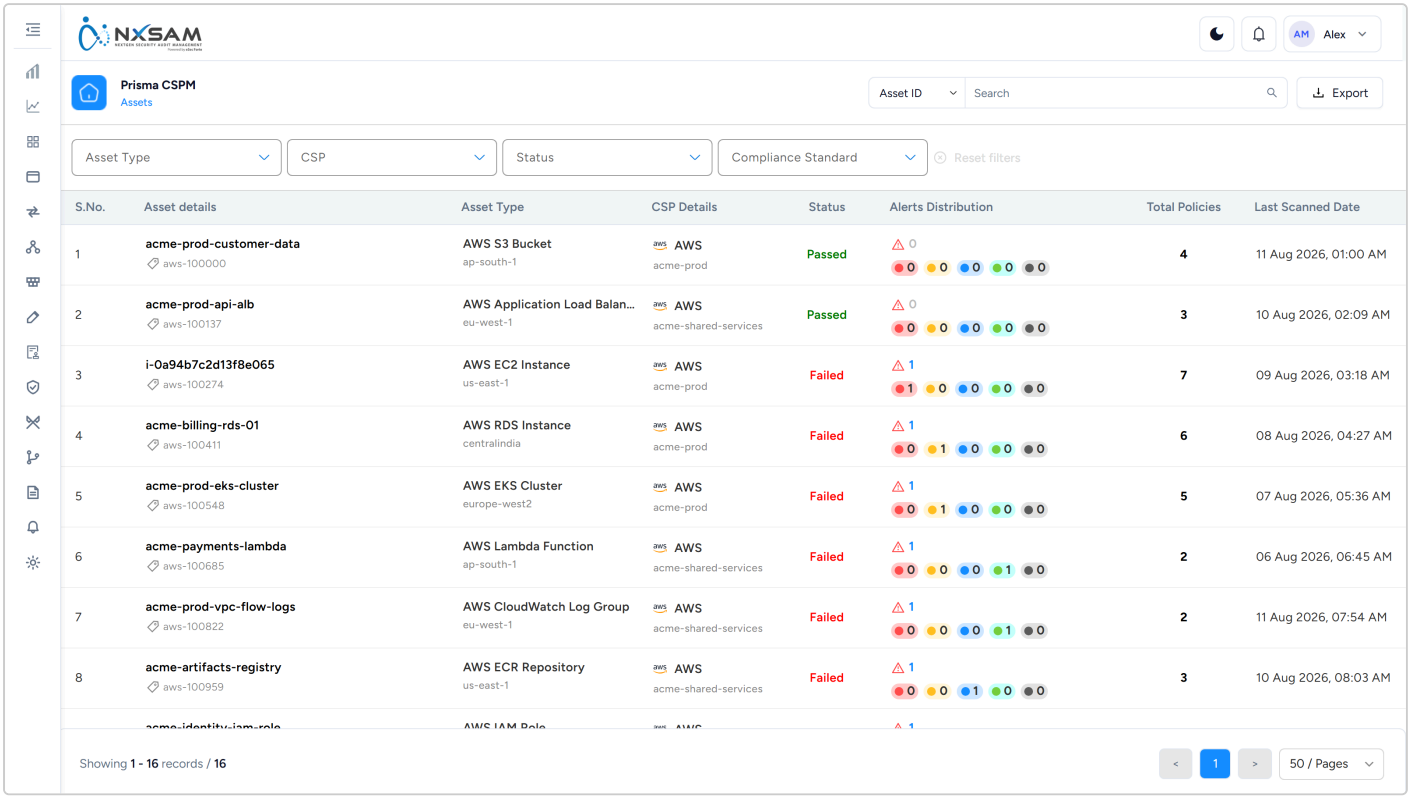
Cloud Security (CSPM): find the right screen

Every screen in this module, what each one is for, and where it lives. Use it to find the right screen quickly, and to confirm you are on the one a colleague means — several screens in this product have similar names and different jobs.

Screenshots come from a demonstration data set, so the figures are illustrative. Click any screenshot to open it at full resolution.

CSPM Asset Inventory

Cloud resources discovered across accounts and regions, with how many policies apply to each.

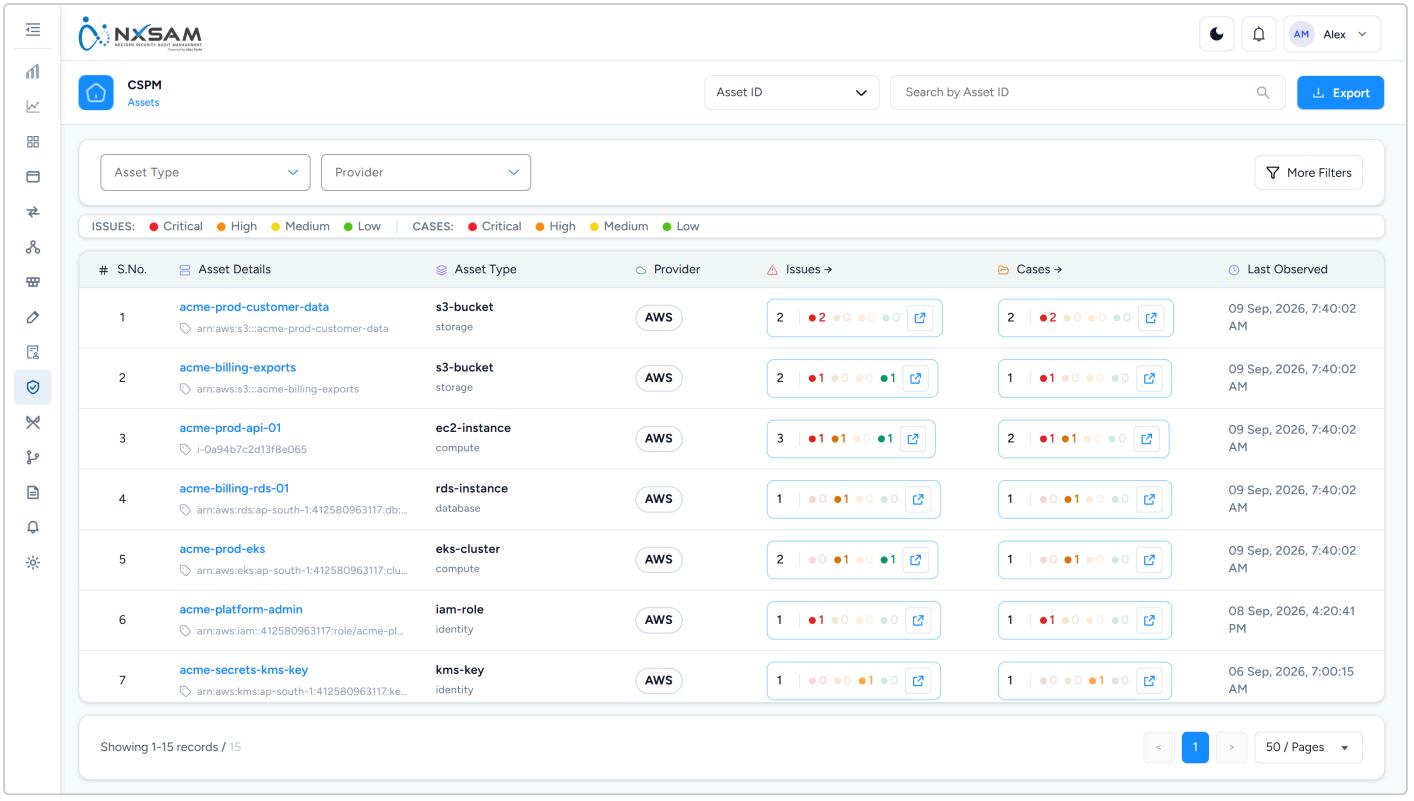


Prisma CSPM → Assets

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

CSPM Assets

The cloud resources Cortex or SentinelOne has discovered, with their open issues and cases by severity. The arrows in the count columns open the other screens filtered to that asset.

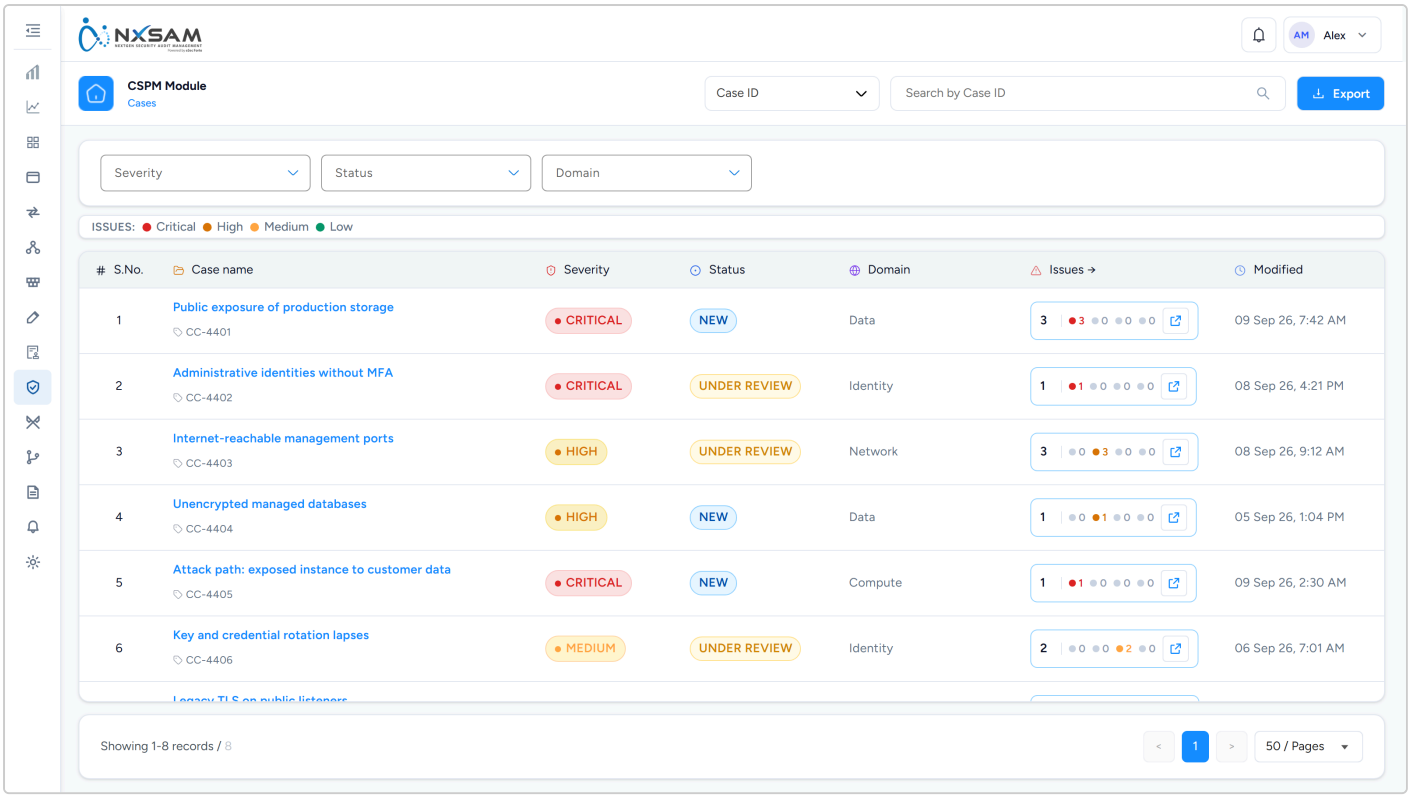


CSPM → Assets

CSPM Cases

Issues grouped into one piece of work, with a severity of the case's own. Fixing the case fixes every issue in it.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

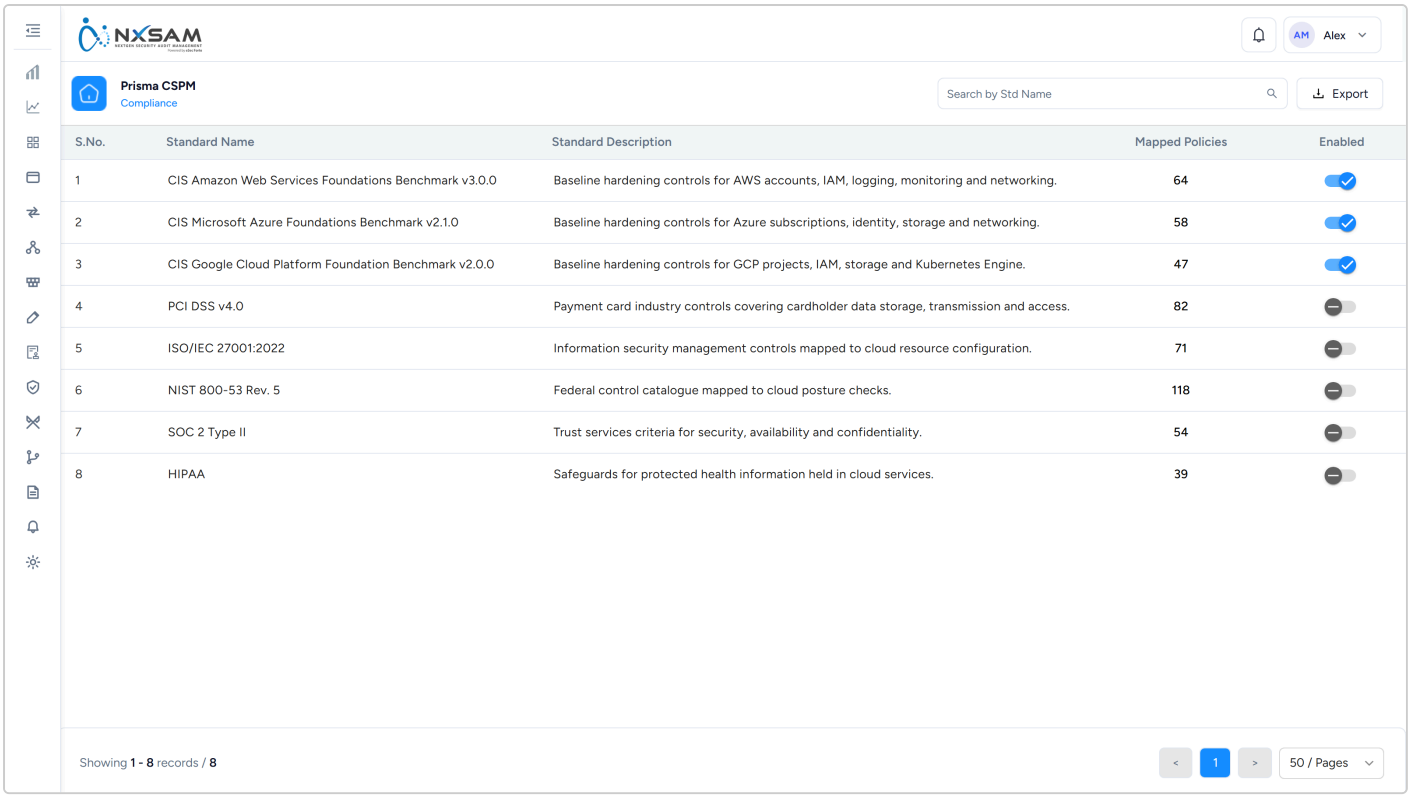


CSPM → Cases

CSPM Compliance Standards

The benchmarks posture is measured against, and which are switched on.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

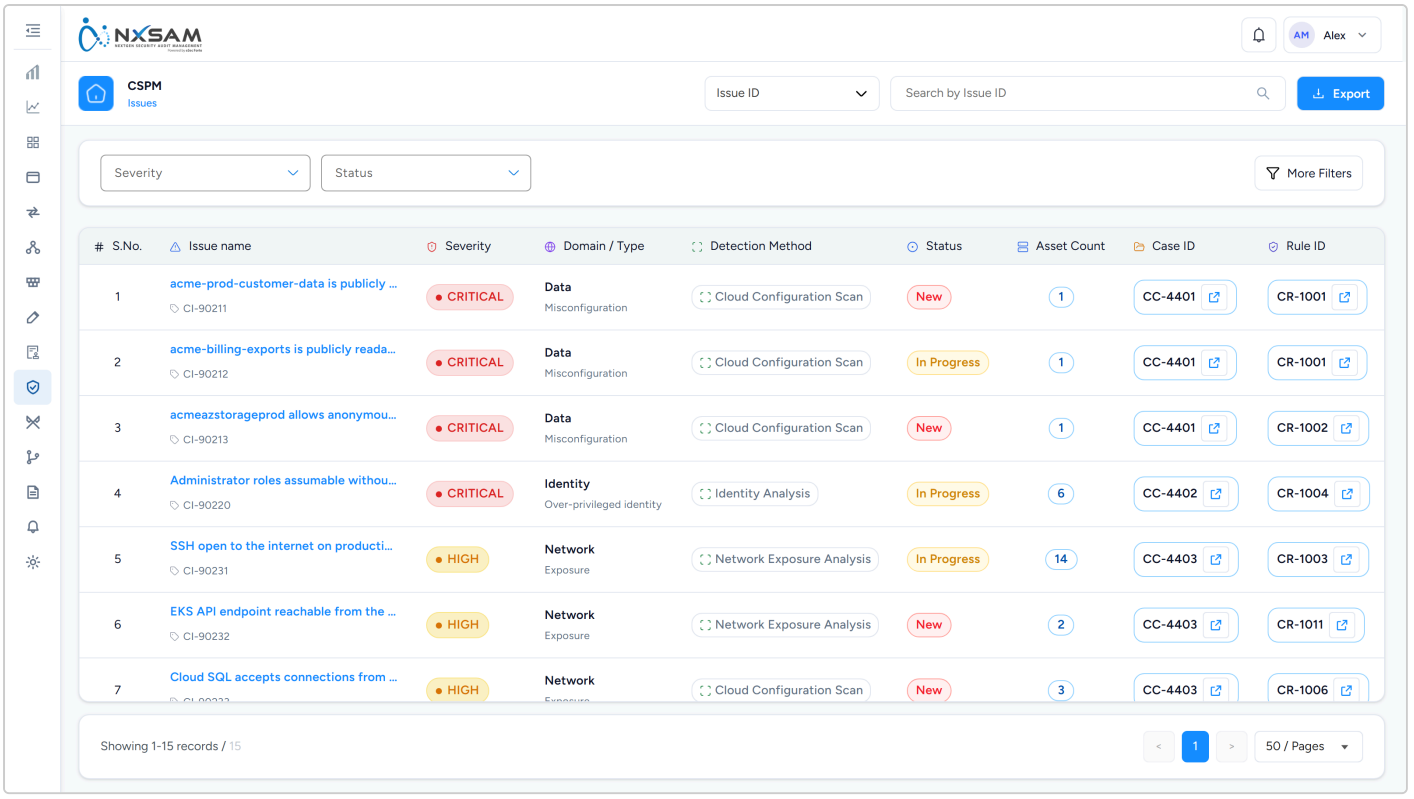


Prisma CSPM → Compliance

CSPM Issues

One row per rule failing against one or more resources. Asset Count is the column to sort by — an issue on two hundred assets is one template to fix, not two hundred resources.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

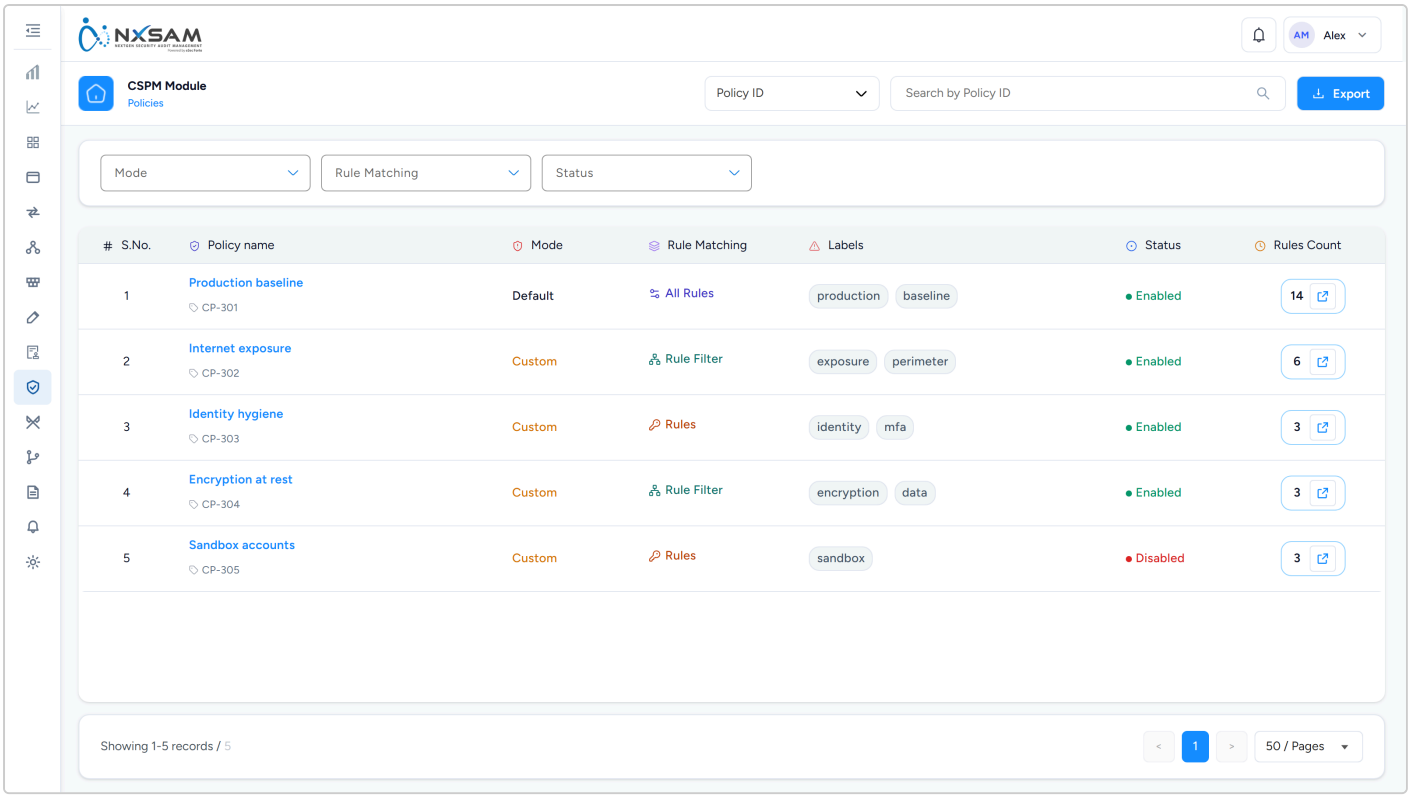


CSPM → Issues

CSPM Policies

Named groups of rules as configured in the posture platform, and whether each group is switched on. Explains why a rule that exists is or is not producing issues.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

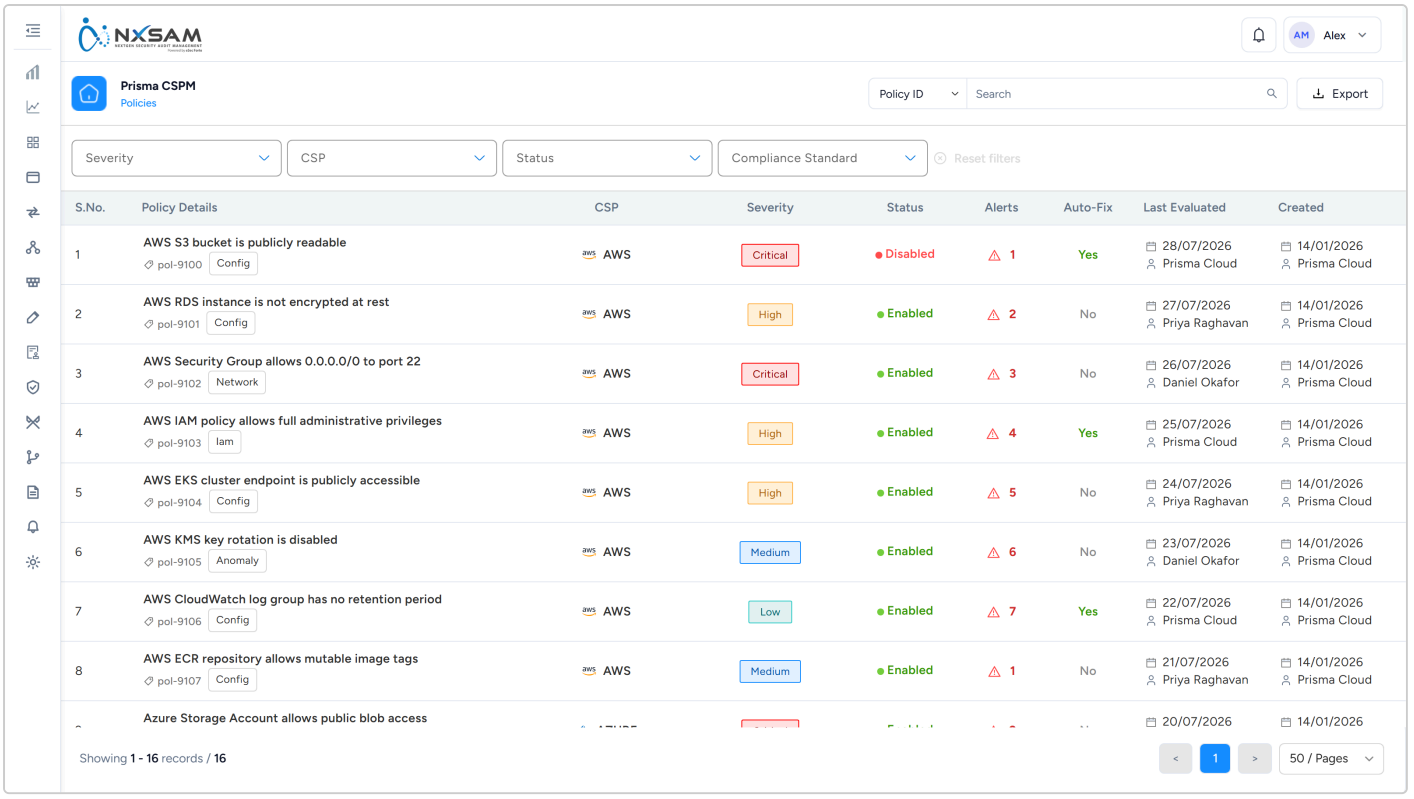


CSPM → Policies

CSPM Policies

The individual checks behind those standards, with the severity each raises.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

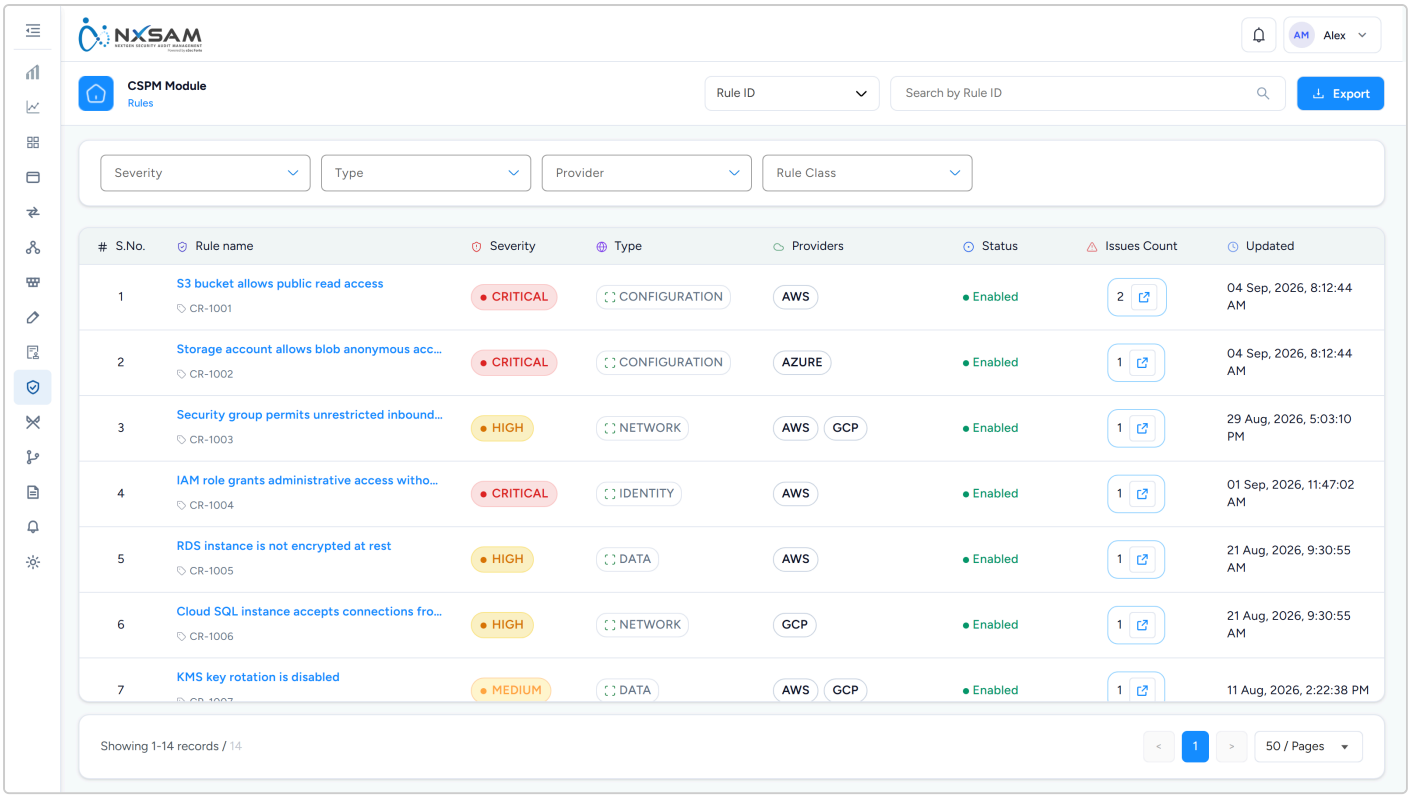


Prisma CSPM → Policies

CSPM Rules

The individual checks behind the issues, with whether each is enabled and how many issues it has raised.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.



CSPM → Rules

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Reference

Glossary

Terms as NxSAM uses them. Several are ordinary security words with a narrower meaning here, and those are the ones worth reading even if the term is familiar.

Assets and scope

Asset — anything that can hold a vulnerability: a web or mobile application, an API, a network host, a thick client, or a container image. Each type is a separate inventory with its own columns and its own licensed limit.

Asset property — an attribute recorded against an asset, such as criticality or whether it is internet-facing. Properties are more than description: they narrow SLA durations and feed the contextual risk score.

Bifurcation — splitting one application into separately assessed parts, so a large system can be tested module by module rather than as one unit.

Collection — a grouping of container images that comes from the scanning platform rather than from NxSAM, and is not editable here.

Scope — what an assessment covers. For applications this is the application; for compliance it is a host list, which is why compliance scope disputes are settled with an exported list of addresses.

Assessments

Audit — one assessment of one asset. It has a type, a category, a cycle, a scope, and a status of its own, separate from any finding's status.

Audit cycle — which round of a recurring programme an audit belongs to.

Confirmatory audit — a re-test created against a parent audit to verify that fixes hold. It is evidence about the original rather than an independent assessment, so the two should not be counted as two rounds of testing.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Considered — the marker deciding whether an audit contributes to aggregate figures. Normally the latest audit per asset that reached a final state. An audit that is *not* considered still shows its findings on the module screen while contributing to no dashboard or report.

Initial audit — a first assessment, as opposed to a confirmatory one.

Findings

Vulnerability (also *finding*, *issue*) — one weakness on one asset, found by one assessment.

Vulnerability master — the catalogue entry (title, description, remediation) inherited by every occurrence of that vulnerability everywhere. Editing the master changes all of them.

Posture — Policy Compliance's unit: one control checked against one host. Every check produces a posture whether it passed or failed, which is why this module's volumes dwarf the others.

Artifact — the evidence attached to a finding: the steps, description and images that let someone else reproduce it.

Instance — which occurrence, where the same control or finding applies more than once to the same asset.

Status

Validation status — whether a finding has been confirmed. Nothing is counted anywhere until this reaches *Approved*, which is also when **Discovered On** is populated.

Client status (also *current status*) — what is being done about the finding. Of its values, only two close it.

Closed type — the flag deciding whether a status stops something counting as open. It is configuration rather than a fixed rule, so a deployment can change what "closed" means.

Posture status — what a compliance check returned: pass, fail or **error**. Error means the check could not run — a scanner or credentials problem rather than a compliance result.

Sub status — the qualifier beneath a main status: which kind of acceptance, or which stage of an approval.

Time and SLA

SLA duration — how long a finding of a given severity on a given asset type may remain open. Configurable, and narrowable by asset property.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

SLA expiry — the resulting deadline. Changing an SLA duration re-derives expiry dates, which can make a batch of findings overdue with no data change at all.

Discovered On — when a finding was first observed. Blank until the finding is approved.

Last observed — the most recent sighting, which is how a finding is known still to exist.

Overdue — past SLA expiry and not closed.

Decisions

Deviation (also *risk acceptance*, *RA*) — a recorded decision not to fix something, with an approval chain and an expiry. It does **not** close the finding: the exposure is still live and still counts as open.

Exception — a status marking an agreed exception rather than a resolution. Reporting that conflates the two overstates remediation.

Auto reject end date — when a temporary acceptance lapses on its own.

Integrations and processing

Integration — a connection to a scanning tool, either pulling from its API or parsing an uploaded report.

Vulnerability key — the columns deciding whether an incoming finding is one already seen. Changing it redefines "the same finding" and splits history.

Queue — asynchronous work: imports, syncs, report generation. A screen that seems not to have updated usually has a queue entry that explains it.

Data process — how far a tool-fed audit has got through the import pipeline; the difference between "the upload finished" and "the data is ready".

Beat — a scheduled job maintaining derived data, such as recalculating SLA states as dates pass. A stopped beat produces figures wrong by a fixed interval and raises no error.

Governance

RLCM — Red Line Compliance Management: assessment of controls against frameworks, rather than of systems for weaknesses.

TPRM — Third-Party Risk Management: the same control assessment applied to suppliers. It shares screens with RLCM and differs by the audit's type.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

UCF — Unified Control Framework: the mapping that lets one control satisfy several frameworks at once, so answering it once covers each of them.

Control — a requirement being assessed, as opposed to a weakness being found.

Standard — the framework a control belongs to, such as ISO/IEC 27001 or PCI DSS.

Reporting

Considered audit filter — the switch on Report Search that reproduces either the dashboard's numbers or the module screens'. Comparing the two settles most disagreements about counts.

Design template — the layout a report renders through, separate from the template that decides its content.

Share view against share copy — sharing a template as a view keeps one definition that your edits reach; sharing a copy creates an independent template that will drift.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.