



NxSAM System & Health — Quick Guide

Quick reference guide for the System & Health 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

System & Health

- Overview
- Tune the risk scoring rules
- System & Health: every screen, column and action
- System & Health: find the right screen

Reference

- Glossary

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

System & Health

System & Health

These screens are for keeping NxSAM itself healthy. They are administrative rather than security-facing, and most organisations look at them only when investigating a problem.

Health Check

Health Check compares what NxSAM holds against what it should hold — reconciling imported data against its source, and flagging discrepancies.

It answers the question that matters most about any consolidation platform: *are we confident nothing was lost on the way in?* If a scan reported 1,200 findings and NxSAM shows 1,150, health check is where that gap surfaces rather than going unnoticed.

Worth running after any large import, and after changing how findings are matched or classified.

Server Status and Process Manager

Server Status shows whether the platform's components are running. Process Manager shows background work — imports, report generation, synchronisation.

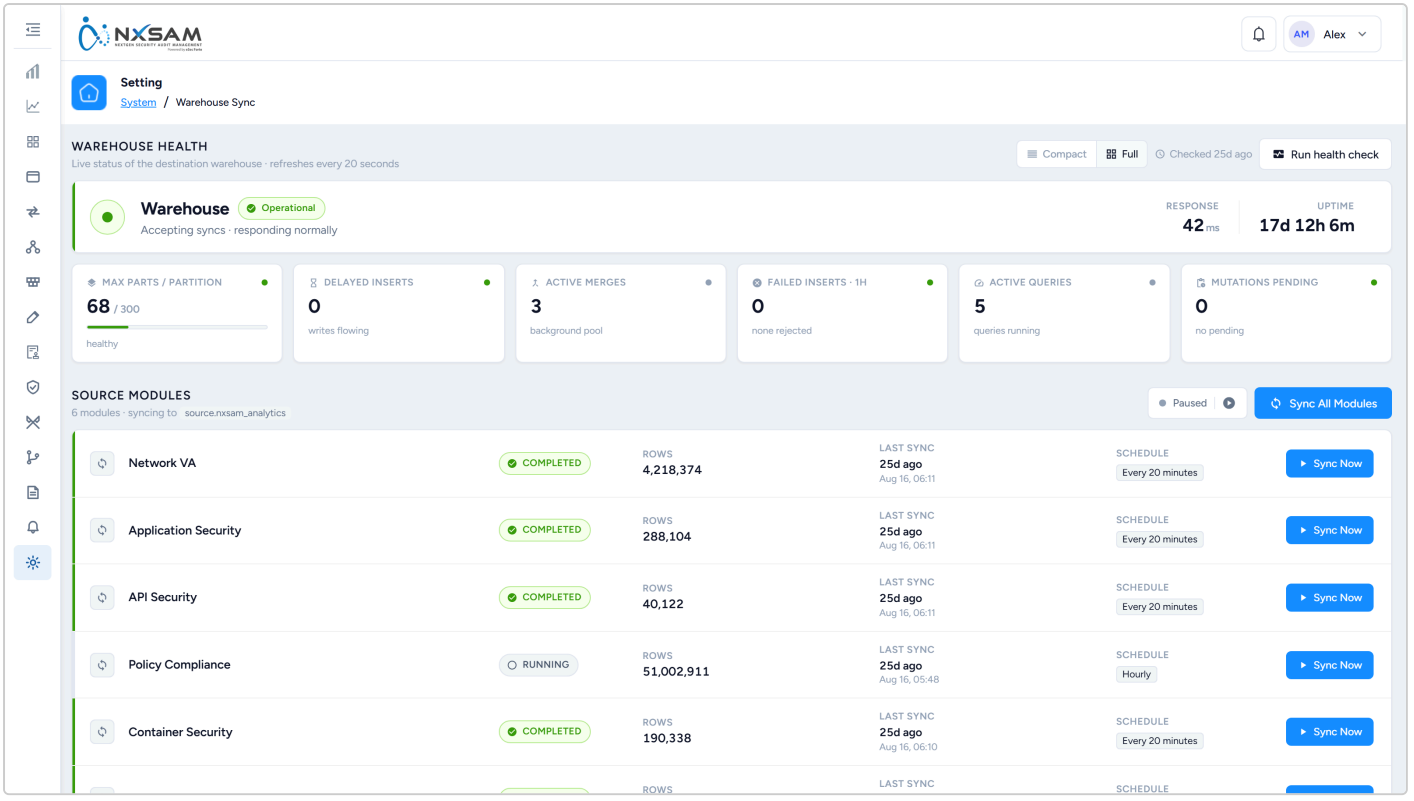
Most long-running work in NxSAM happens in the background: a report is queued rather than produced instantly, and a large import is processed over time. When something "hasn't happened yet", these screens show whether it is queued, running, or failed.

Queued is not failed

A report or import that has not appeared may simply be waiting behind other work. Check whether it is queued before treating it as an error.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Warehouse Sync



Warehouse sync status between the live database and the analytics store.

Where analytics are served from a separate reporting store, this screen shows the synchronisation between the live database and that store.

This is the usual explanation for analytics figures lagging the module screens: the module reads live data, analytics reads the synchronised copy. A gap between them is a sync timing question, not a data integrity one.

Licence

Shows your licence and its limits. Worth checking before a large onboarding — asset limits are enforced, and discovering the ceiling midway through a bulk import is avoidable.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

AI configuration

Where an administrator connects the platform's AI service and decides what it may be used for. Today that is one thing: reading the evidence uploaded against a compliance or third-party control and suggesting a finding to the auditor, with a confidence score.

The screen holds the provider and credential, the model, how often analysis runs, a spending cap per audit, and which modules it is switched on for. A second screen, **AI Queue Details**, lists the analysis runs and lets one be re-run.

Two points worth knowing before switching it on:

- **Nothing is written by it.** The auditor still answers every control. The evaluation appears beside the evidence as a suggestion and is not counted anywhere.
- **The cap is a hard stop.** Analysis pauses for an audit once its spend reaches the cap, and the remaining controls stay *pending* until the cap is raised. A cap set low to be safe therefore looks, from the assessment, like analysis that never finishes.

When to look here

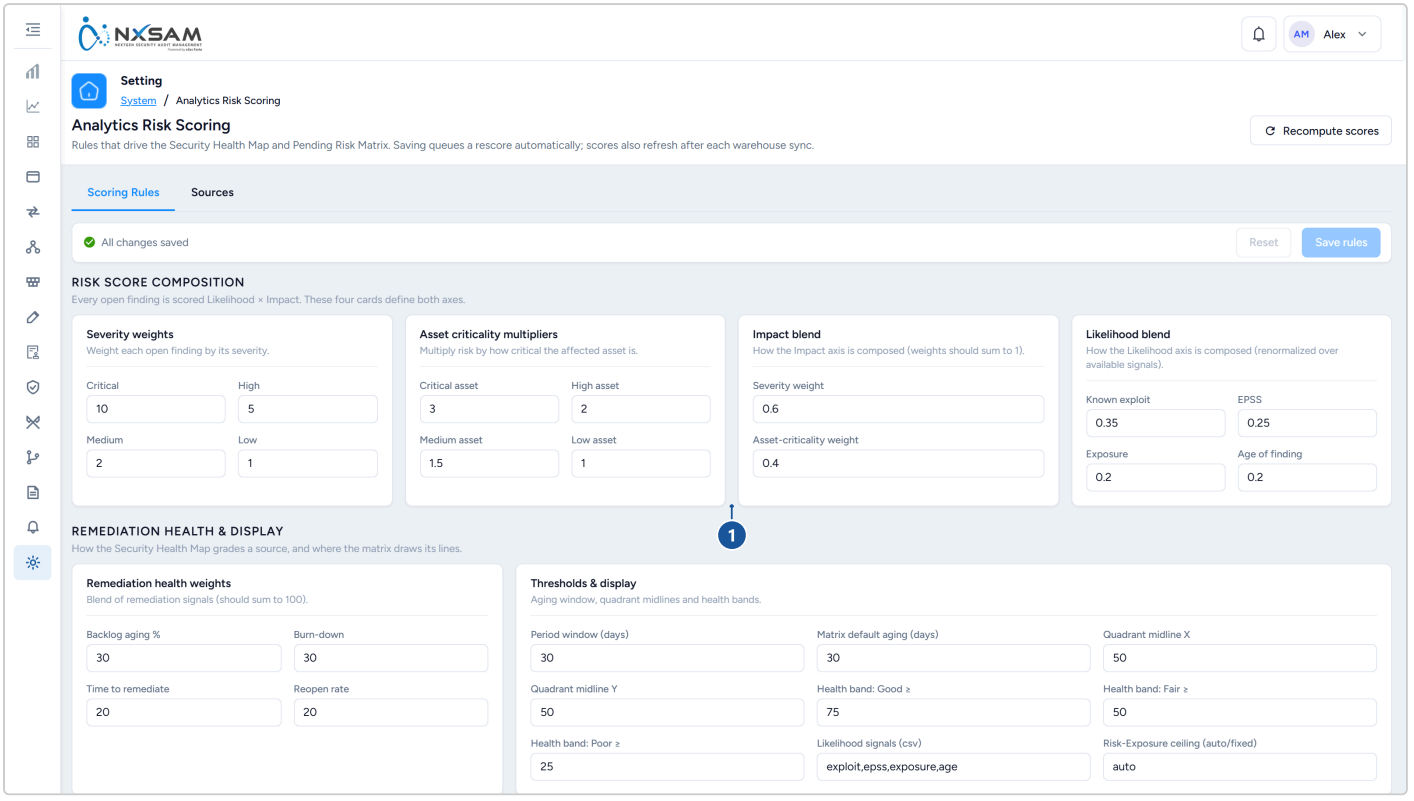
Symptom	Screen
An import seems incomplete	Health Check
A report never arrived	Process Manager
Analytics disagree with a module	Warehouse Sync
Something is broadly not working	Server Status
Cannot add more assets	Licence

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Tune the risk scoring rules

Setting → System → Analytics Risk Scoring holds the rules behind the Security Health Map and the Pending Risk Matrix. Every open finding is scored **Likelihood × Impact**, and this screen defines both axes, how remediation is graded, and where the charts draw their lines.

Nothing here is hard-coded. If the ranking does not match how your organisation actually prioritises work, this screen is where that is fixed — not a support request.



Setting → System → Analytics Risk Scoring → Scoring Rules

1 Likelihood × Impact — the two axes of the risk score

How the score is built

Four cards define the two axes.

Impact

Impact blends two things, and you choose the balance:

Card	What it sets	Default
Severity weights	What each severity contributes	Critical 10 · High 5 · Medium 2 · Low 1
Asset criticality multipliers	How much the affected asset's criticality multiplies the risk	Critical 3 · High 2 · Medium 1.5 · Low 1
Impact blend	The split between the two above	Severity 0.6 · Asset criticality 0.4

The blend weights should sum to 1. Shift it toward asset criticality and a medium finding on a crown-jewel system outranks a high finding on a sandbox — which is usually what people mean when they say the ranking "does not reflect reality".

> ⚠ **Asset criticality has to be set for this to mean anything**
>
> If criticality is unset across the estate, every multiplier is the same and the Impact axis collapses to severity with extra steps. The scoring rules are not the fix for that; populating asset criticality is.

Likelihood

Likelihood blend distributes weight across four signals — defaults in brackets:

- **Known exploit** (0.35) — a public exploit exists for this vulnerability
- **EPSS** (0.25) — the published probability of exploitation in the wild
- **Exposure** (0.20) — how reachable the affected asset is
- **Age of finding** (0.20) — how long it has been open

The weights are **renormalised over whichever signals are available**, which matters more than it sounds. A finding with no EPSS score does not silently score lower; its remaining signals are rescaled to fill the gap. That is why a mixed estate, where only some sources carry exploit data, still ranks sensibly.

Likelihood signals on the Thresholds card names which signals participate at all, comma separated. Removing one there is different from setting its weight to zero: a removed signal is excluded from the renormalisation entirely.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Remediation health

The Security Health Map grades each source on how well findings are actually being closed, not just how many exist. **Remediation health weights** blends four signals and should sum to 100 — defaults are 30 / 30 / 20 / 20:

Signal	What it measures
Backlog aging %	How much of the open backlog is old
Burn-down	Whether the backlog is shrinking
Time to remediate	How long closure takes
Reopen rate	How often a closed finding comes back

Reopen rate is the one worth raising if closures are being recorded optimistically: a team that closes quickly and reopens constantly scores well on the first three and badly on the fourth.

Thresholds and display

These do not change any score. They change where the charts draw their lines, which changes what people conclude from the same numbers.

Setting	Effect	Default
Period window (days)	The window burn-down and time-to-remediate are measured over	30
Matrix default aging (days)	The ageing cut the Pending Risk Matrix opens on	30
Quadrant midline X / Y	Where the matrix splits into quadrants	50 / 50
Health band: Good ≥	The score at or above which a source is Good	75

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Setting	Effect	Default
Health band: Fair ≥	Fair	50
Health band: Poor ≥	Poor; below this is the worst band	25
Risk-Exposure ceiling	`auto` scales the exposure axis to your data; `fixed` pins it	auto

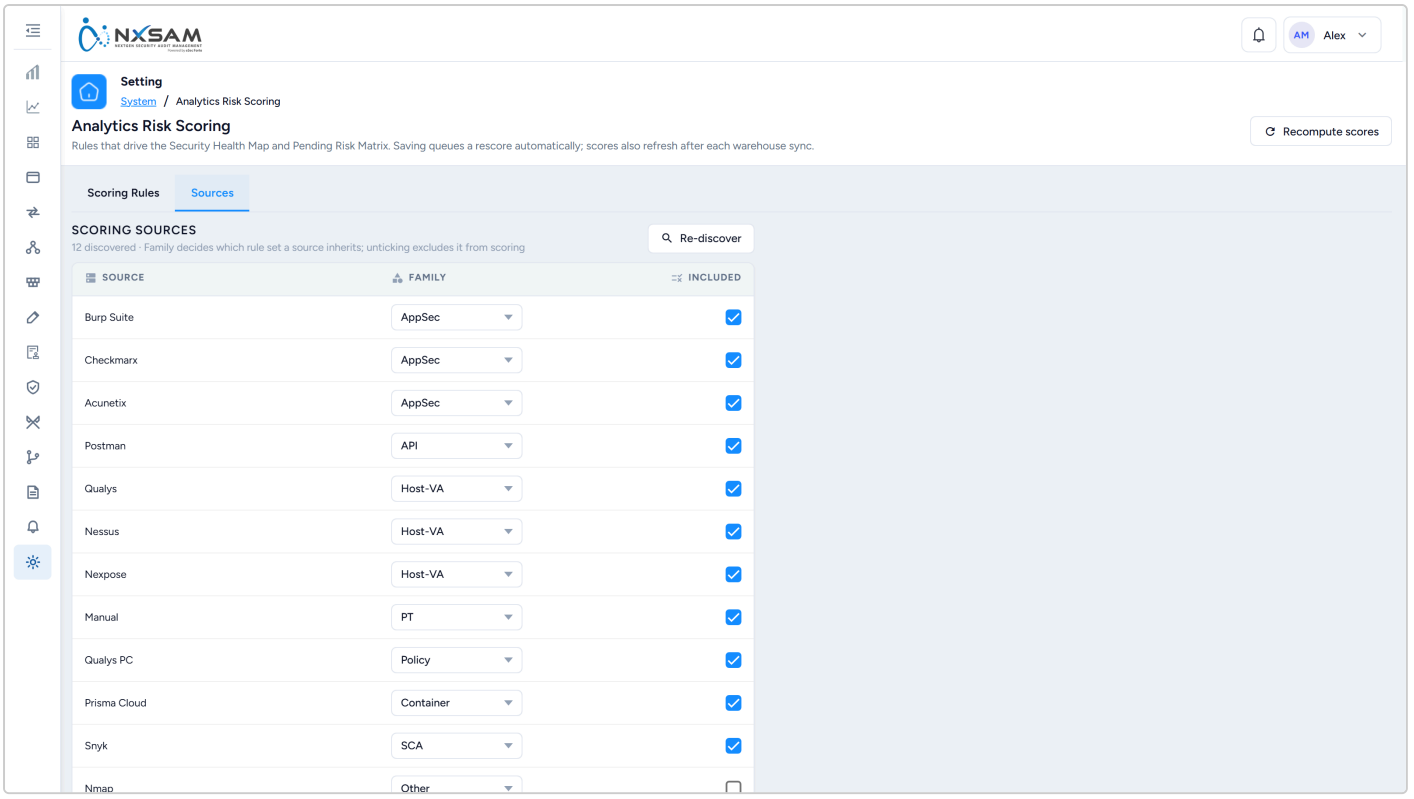
Moving a midline moves sources between quadrants

Nothing was rescored, but a source that was in the acceptable quadrant yesterday can be in the attention quadrant today. If the health map changes character overnight and no data changed, this screen is the first place to look.

Sources: what gets scored, and by which rules

The **Sources** tab lists every scanner and tool discovered in your data.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.



Setting → System → Analytics Risk Scoring → Sources

Each row has two controls:

- **Family** decides which rule set the source inherits: *AppSec*, *Host-VA*, *PT*, *Policy*, *SCA*, *API*, *Container* or *Other*. Families are assigned automatically from the shape of the data, and the assignment is a starting point rather than a verdict — change it where you disagree.
- **Included** takes the source out of scoring entirely when unticked. Use it for a tool whose findings are informational, or one being trialled, so it stops distorting the health map without anyone having to delete its data.

Re-discover rescans for sources that have appeared since. A tool integrated last week will not be scored until it has been discovered.

Saving, and when the numbers change

The bar above the rules shows either *All changes saved* or a count of unsaved changes, and each changed field is marked while you edit. **Reset** discards them; **Save rules** commits.

Saving **queues a rescore automatically** — you do not have to trigger one. Scores also refresh after each warehouse sync. **Recompute scores** rebuilds every score without changing a rule, which is what you want after correcting asset criticality or importing a backlog of findings.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

A rescore is not instant on a large estate. If one is already running, a second request is told so rather than queueing behind it.

If the screen says no rules were found

The default rule set has not been seeded in that environment. The screen is doing the right thing by saying so rather than showing every input at zero, which would read as a configuration where nothing counts. Raise it with whoever provisions the environment.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

System & Health: every screen, column and action

These screens are how you tell whether the platform itself is working, and they are the first place to look when a symptom spans several modules.

Process Manager

Actions: **Logs**, **Restart**, **Stop**.

Background work — imports, report generation, synchronisation, outbound email — runs as processes here. **Logs** is where a failure explains itself; **Restart** brings back a worker that has died; **Stop** halts one that is misbehaving.

> ⚠ **Restart is not free**
>
> A worker restarted mid-message leaves that job in `IN PROGRESS` with nothing consuming it — the job neither completes nor retries. Prefer letting a job finish where you can, and check what is running before restarting.

Import

`Import Type`, `Import Status`, `Log`, `Created At`, `Summary`, with a child-process view of the same columns.

The child view matters for large imports: one import spawns several processes, and the parent completing does not mean every child did. An import reporting success with fewer rows than expected usually has a child that failed, visible only here.

Health Check

The widest diagnostic in the product, and the one that answers "did we lose anything on the way in":

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Column	What it holds
Baseline	What the source reported
Current	What NxSAM holds now
New	Present now, absent from the baseline
Fixed	In the baseline, resolved since
Repeated	Seen in both
Re-open	Closed previously, present again
Total	The reconciliation total
Type	What kind of assessment produced the row
Match Key	The field used to decide two findings are the same
Classification	How the comparison classified the row
Severity, **Status**, **Sub Status**, **SLA Expiry**	The finding's own values
Issues	The findings behind the number

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Match Key is the field to check when a comparison looks wrong

Everything on this screen depends on how two findings are decided to be the same one. If **New** and **Fixed** are both unexpectedly high while **Repeated** is near zero, the match key is not matching — findings are being seen as fixed and new rather than as the same finding twice.

That is a matching-configuration problem, not a data-loss one, and the numbers look alarming in exactly the way real data loss would.

Re-open distinguishes a finding that came back from one that never left. A high re-open count against a specific asset usually means a fix that does not survive a rebuild or a redeploy.

Beats — scheduled jobs

`Module`, `Sub Module`, `Last Sync Status`, `Last Sync Type`, `Last Sync Initiated At`, `Last Sync Processed At`, and per-run `Sync Sub Module`, `Sync Type`, `Sync Status`, `Sync Initiated At`.

These are the periodic jobs that keep derived data current — recalculating SLA states as dates pass, refreshing pre-computed statistics.

Initiated At against Processed At is the useful pair. A job initiated and never processed is stuck; one where the gap is growing run on run is falling behind, and the figures it maintains are increasingly stale.

Figures wrong by a *fixed interval* rather than randomly are almost always a beat that stopped. Nothing else in the product reports an error for it.

Active Directory

Sign-in can be delegated to Active Directory, and **more than one directory can be configured**. Each is held as a named *account* and the screen lists them as cards rather than as a single form.

Pick an account from **Select Account**, then **Configure**. An account with nothing configured shows an empty state rather than a blank form — which is the difference between "not set up" and "set up and broken".

Field	What it holds
AD Type	Which directory product and protocol variant

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Field	What it holds
Domain	The directory host
Port	389 for LDAP, 636 for secure LDAP — the form states both defaults
User Name Prefix	The prefix prepended to a username when binding
Admin User Name	The service account used to search the directory
Users Base DN	Where in the tree user accounts are looked up
Certificate	The uploaded certificate — **shown only on Secure LDAP directories**

Choosing **Secure LDAP** is what reveals the certificate upload on the form and the Certificate row on the card. A directory that is supposed to be encrypted and shows no Certificate row is configured as plain LDAP, and its bind traffic is in the clear.

Each card carries **Edit**; a directory that is not the last remaining one can also be deleted.

User field mapping

The lower half of the form maps directory attributes onto NxSAM's **First Name**, **Last Name**, **Email**, **Mobile** and **Username**.

⚠ The Username mapping must be the AD login name

The form says so, and it is the one mapping that cannot be corrected later without orphaning accounts. Map it to the attribute users actually sign in with, not to a display name that looks right in a directory browser.

Everything else is cosmetic and can be remapped. This one decides whether a returning user is recognised as themselves.

Test Connection

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Test Connection reports *Connection Tested Successfully* or *Connection Failed* before you save. A directory that saves but does not bind produces sign-in failures for real users, and it surfaces at their next sign-in rather than at yours.

AD logs

The panel above the accounts shows the **last 100 lines** of directory activity, with **Refresh** and **Load Logs**. It is where a bind failure explains itself — wrong base DN, expired service-account password, an untrusted certificate — and the first place to look when authentication works for some users and not others.

SLA configuration

`Severity`, `Asset Type`, `SLA Duration`, `Asset Properties`.

SLA duration is set per severity **and** asset type, and can be narrowed further by asset property — so a Critical on an internet-facing production system can carry a shorter deadline than the same Critical internally.

Changing a duration re-derives expiry dates, which can make a batch of findings overdue with no data change at all.

Asset limits

`Client`, `Web Limit`, `Mobile Limit`, `Network Limit`, `Image Limit`, `Think Client Limit`.

Licensed ceilings per asset type per client. Worth checking before a large onboarding — discovering the ceiling midway through a bulk import means a partial inventory and a reconciliation afterwards.

Risk score properties

`Asset Type`, `Properties` — which asset attributes feed the contextual risk score, per asset type. An estate with these unset produces a risk score that is severity with extra steps, and a "top risky assets" list that ranks nothing meaningfully.

Tags

Four tables, one per asset domain: **Application**, **API**, **Network** and **Audit** tags. Each carries `Tags Id`, `Tags Name`, `Created By`; audit tags additionally carry `Clients`.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Export Tags produces the list. Tags are shared within their domain, so renaming one renames it everywhere it is applied — there is no per-use copy.

System configuration

Screen	Columns
Modules	`ID`, `Module`, `Created By`
Controllers	`ID`, `Class`, `Method`, `Created By`
Permissions	`ID`, `Module`, `Permission`, `Code`, `Created By`
Queries	`ID`, `Query`, `Task`, `Status`
Task info	`Info Id`, `Info Name`

These define what the platform can do and who may do it. **Permissions** map a code to a module — roles grant those codes, so a permission missing here cannot be granted however the role is configured.

Queries are maintenance operations bound to tasks. They are administrative tooling rather than a working screen, and running one without knowing what it does is not recoverable through the interface.

System settings

The switches that change how the whole platform behaves. Most are described with the module they affect; one is easy to miss because it lives here and bites elsewhere:

Only 1 On-going VA Audit per Asset — when on, a network VA audit cannot be requested for an asset that already has one in progress. The request fails at the point of saving with a message naming the rule. It exists because two concurrent scans of one host produce two sets of findings for one state, and reconciling them afterwards is worse than waiting. Turn it off only where genuinely parallel assessments of one asset are intended.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

AI configuration

Two screens. The first is a form; the second is the queue it feeds.

AI Configuration

Section	What it holds
Provider & Credentials	The **AI Provider** and its API key. The key is **stored encrypted** and never shown again once saved
Evidence Analysis Model	Which model reads the evidence. Larger models read complex evidence better and cost more per run
Analysis Schedule Interval	How often queued evidence is analysed — every 15, 30 or 60 minutes
Per-Audit Cost Cap ($)	Spending limit per audit. *Analysis pauses when spend hits this cap.* Negative values are not allowed
Evidence Analysis for RLCM	The master switch: *let the platform review uploaded audit evidence and suggest compliance findings*
Apply to Modules	Which assessments it runs for — **Compliance**, **TPRM**, or both

If the AI service is not reachable from this deployment the screen shows **AI service not configured** and a **Retry**; nothing below it can be saved until that clears, and the message says to contact the administrator of the deployment rather than of NxSAM.

Set the cost cap from the evidence, not from the budget. A cap is spent by reading files, so an audit with a large evidence pack exhausts a small cap early and its remaining controls stay *pending* indefinitely. An auditor reporting that "the AI stopped halfway through" has almost always met the cap.

AI Queue Details

Reached from the **AI Queue Details** button on the configuration screen, which opens it in a new browser tab so the configuration stays where it was.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

Column	What it holds
Queue Details	The queue's key and what kind of job it is
Status	Where the run is
Created At	When it was queued
Actions	**Re-analyse this queue** — runs it again

Re-analysing is safe to repeat: it reads the same evidence again and replaces the previous evaluation. It also spends against the same cap, so a queue re-analysed to get past a cap will stop in the same place.

Analytics custom charts

Custom chart cards carry **Refresh** always, and **Edit** and **Delete** only on charts you own.

Charts on a shared profile are read-only to everyone but their owner

A dashboard profile shared with you shows its charts with **Refresh** alone. That is deliberate: a shared profile is one definition seen by many people, so an edit would change what every viewer sees, and a delete would remove it from all of them.

To change a shared chart, ask its owner — or take your own copy of the profile, where the charts become yours to edit.

Refresh re-runs that chart's query alone. Useful when one widget looks stale and the rest of the dashboard is current — which points at that chart's query rather than at the warehouse sync.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

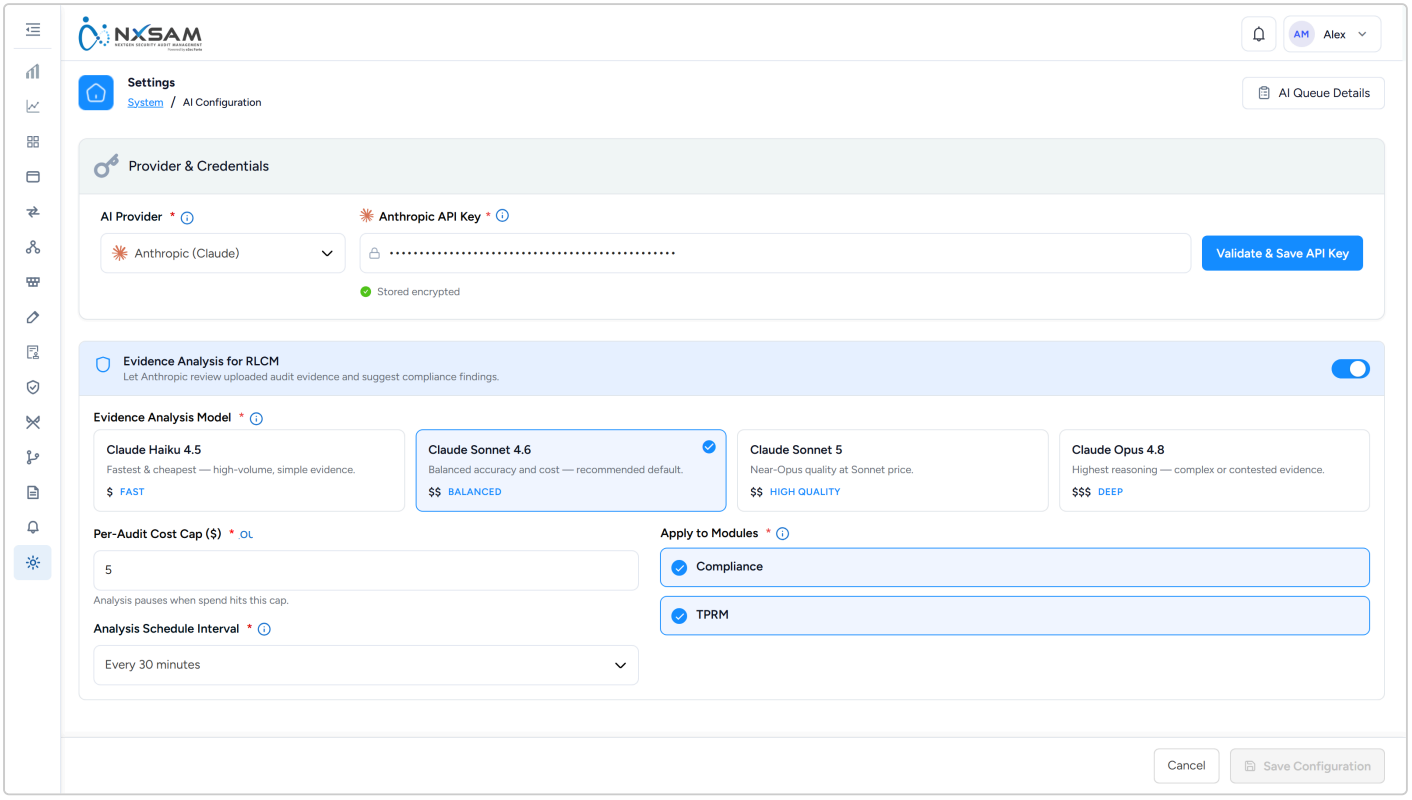
System & Health: find the right screen

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.

AI Configuration

Where the platform's AI service is connected and told what it may do — today, reading uploaded compliance evidence and suggesting a finding to the auditor. Provider, model, schedule, a per-audit cost cap, and which modules it is on for.

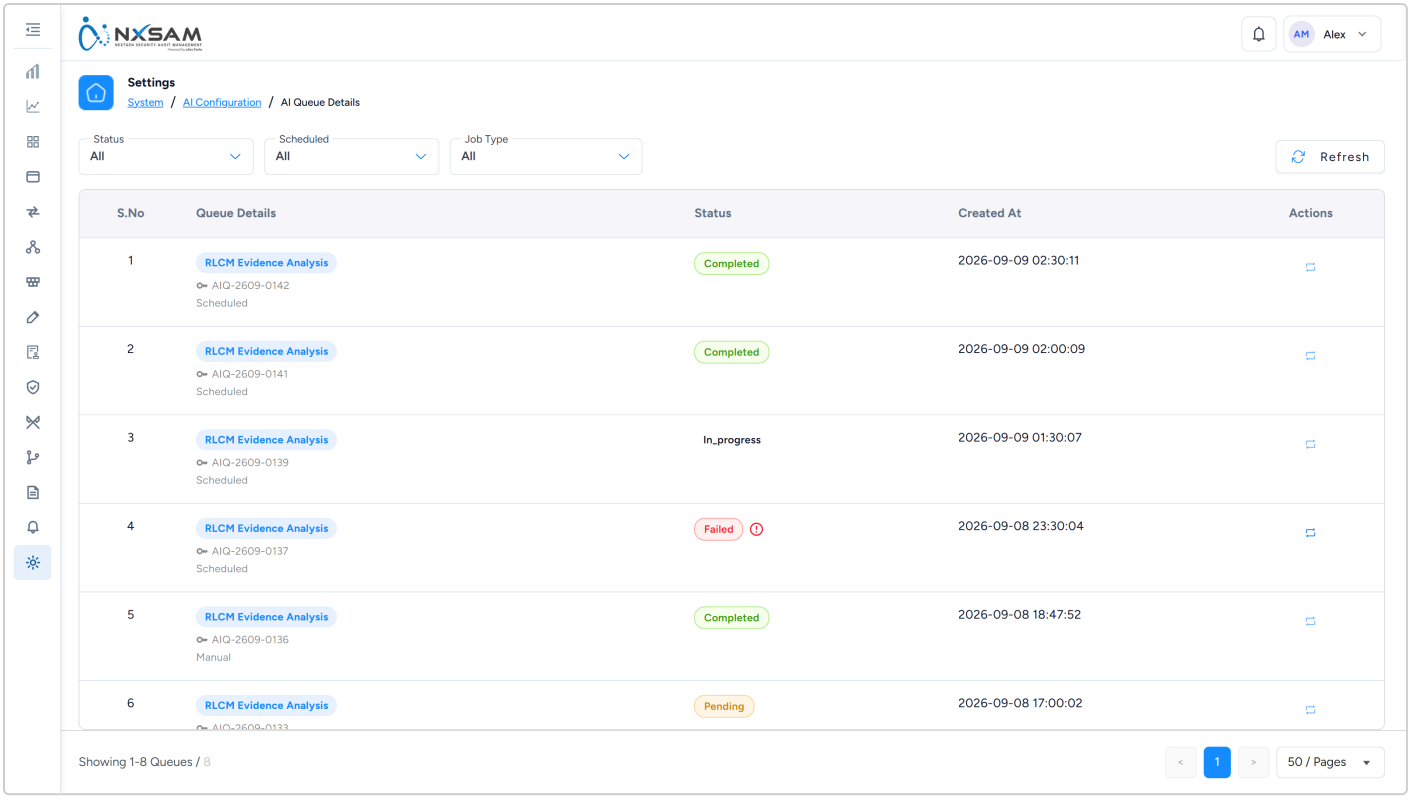


Setting → System → AI Configuration

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

AI Queue Details

Every evidence-analysis run with its status, and a re-analyse action. A run that stays pending has usually met the audit's cost cap.

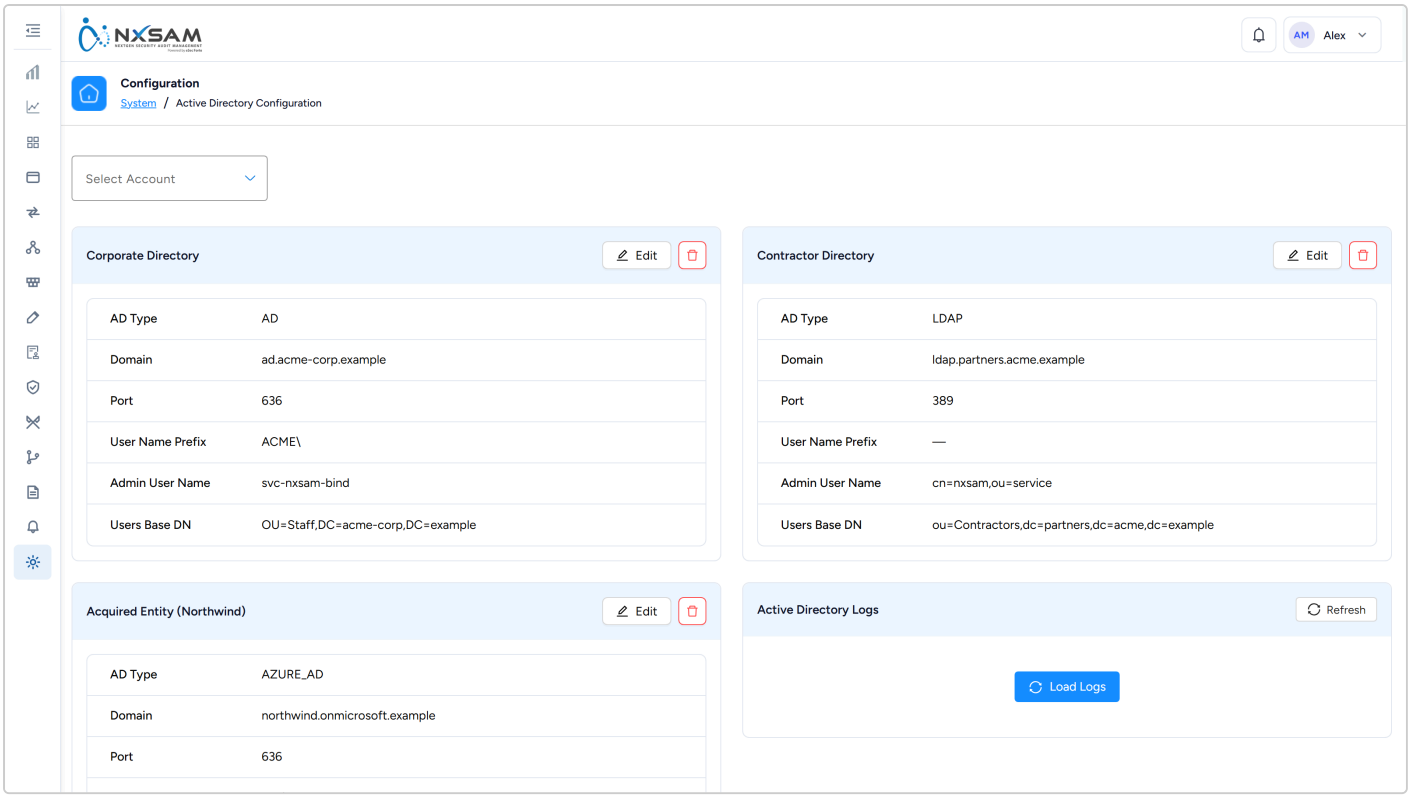


Setting → System → AI Queue Details

Active Directory

Directory integration for sign-in and user provisioning.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

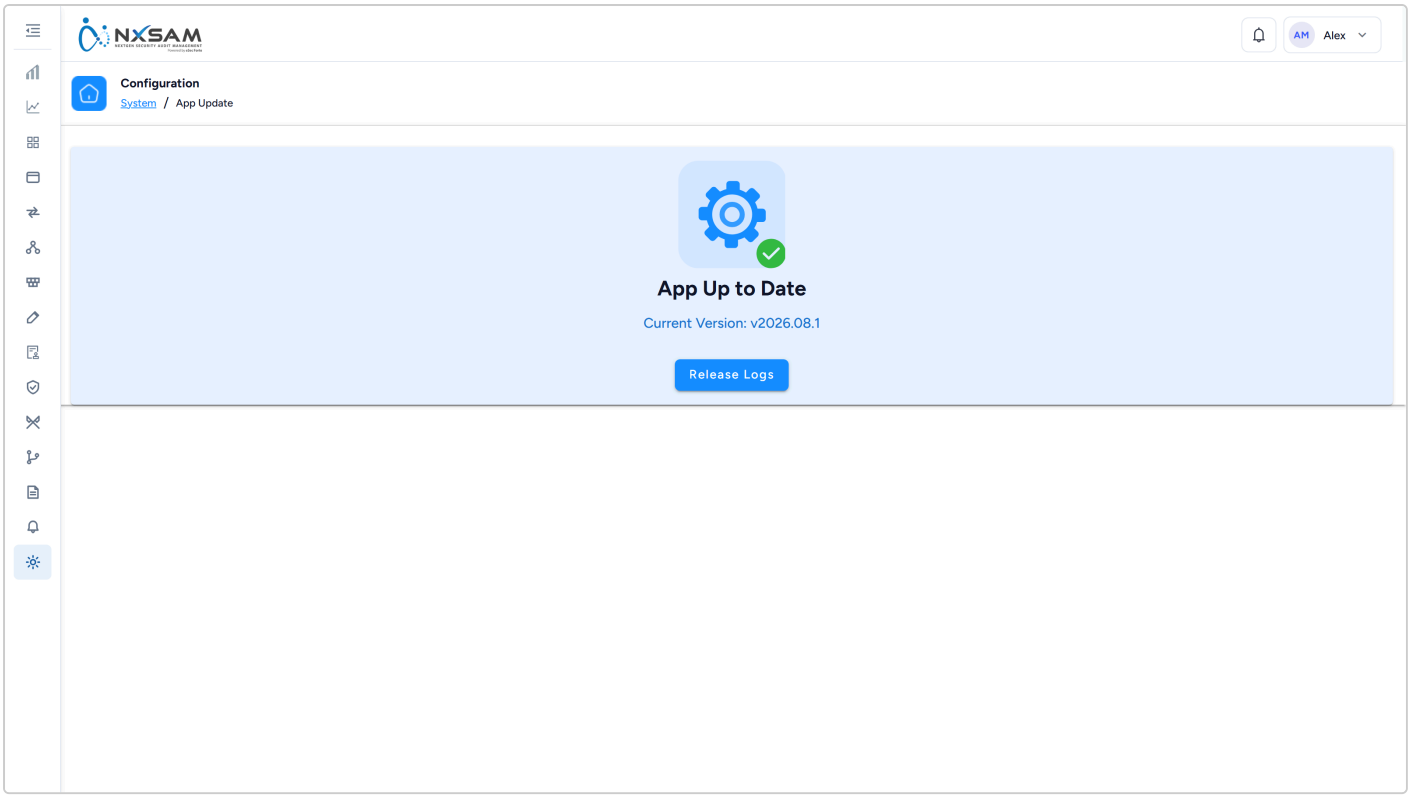


Setting → System → Active Directory

App Update

The running version, and whether a newer one is available.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

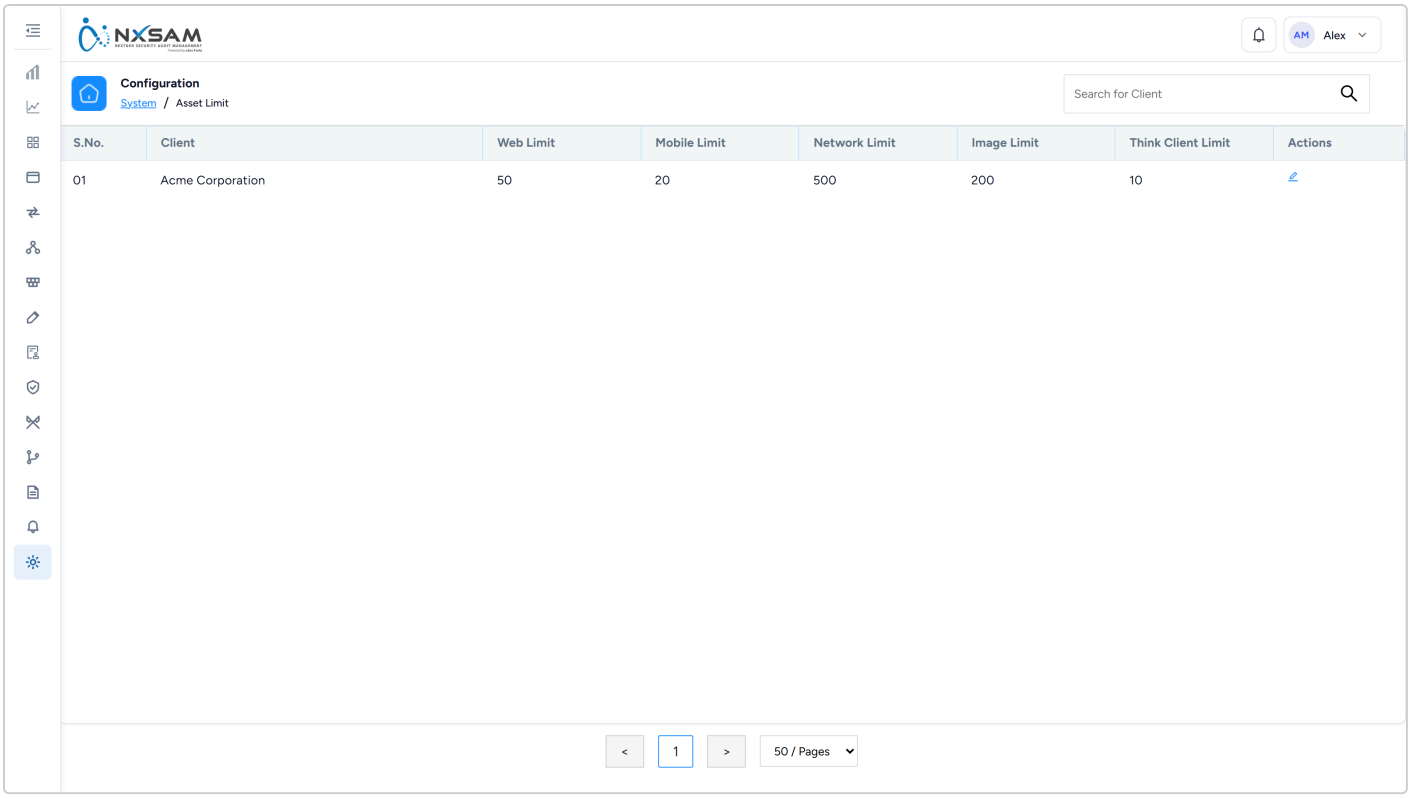


Setting → System → App Update

Asset Limit

Licensed asset ceiling and current usage. Worth checking before a large onboarding.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

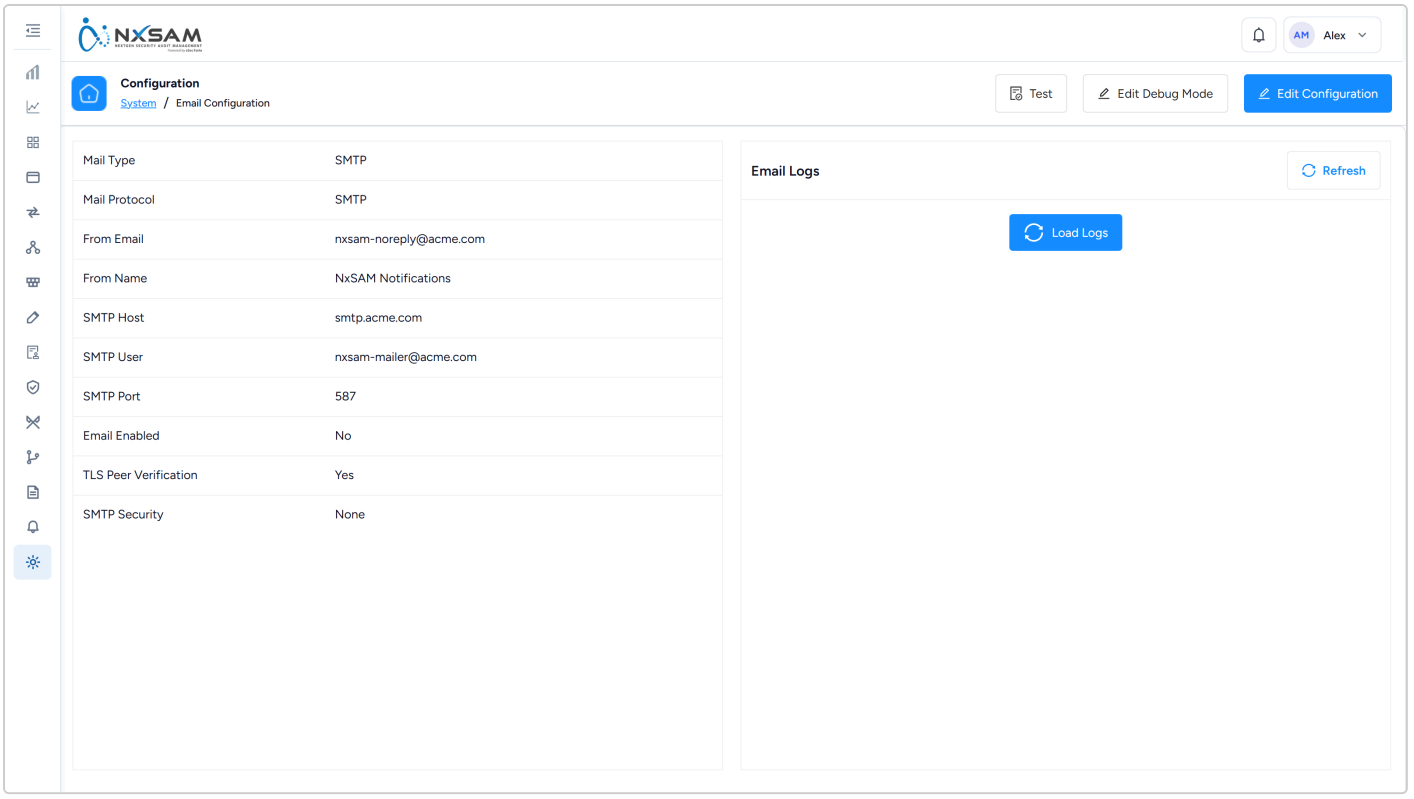


Setting → System → Asset Limit

Email Configuration

The mail server used for notifications and scheduled reports.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

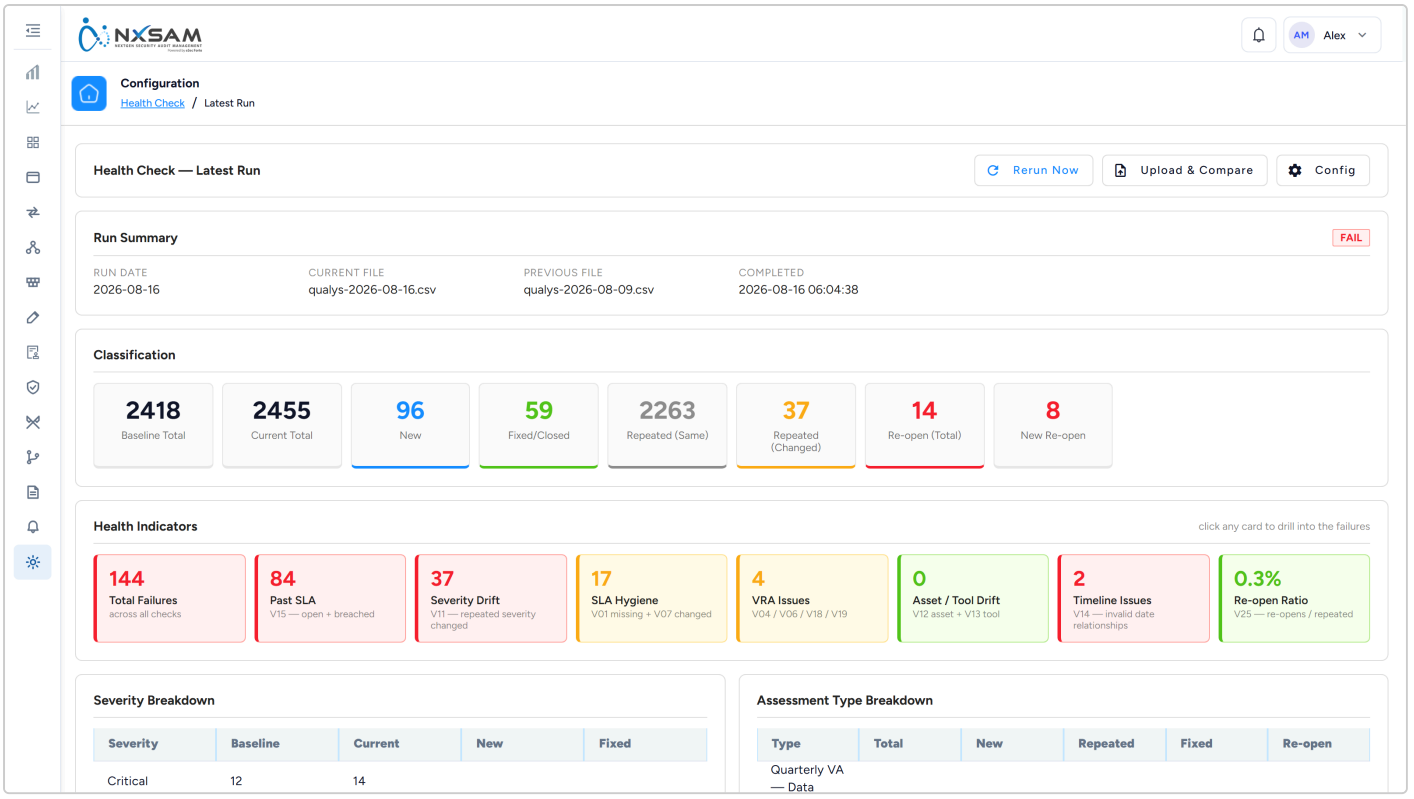


Setting → System → Email Configuration

Health Check

Reconciles what NxSAM holds against what the source reported, so an incomplete import surfaces rather than passing unnoticed.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

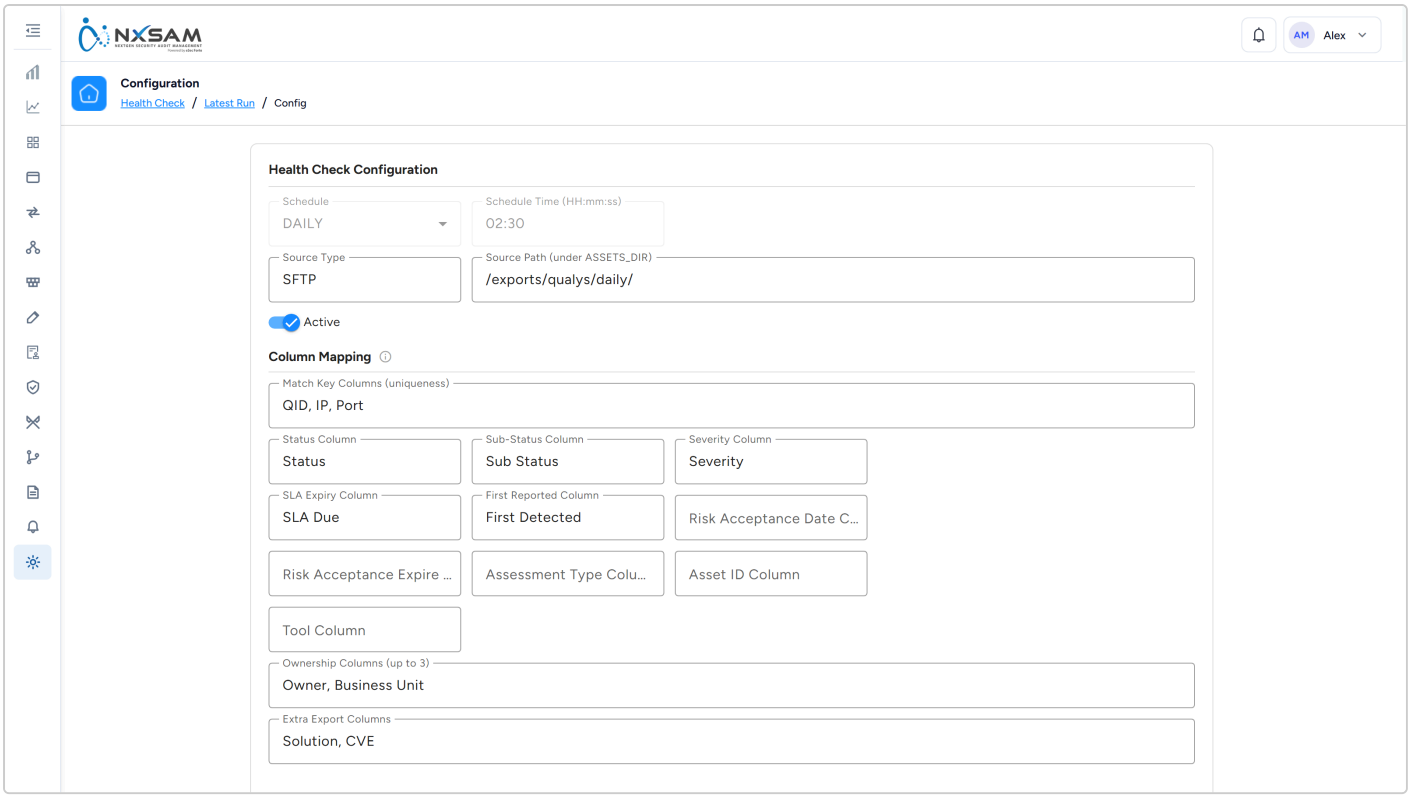


Setting → Health Check → Health Check

Health Check Configuration

The comparison rules that reconcile a tool's own export against what NxSAM holds — which columns identify a finding, and what counts as a mismatch.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

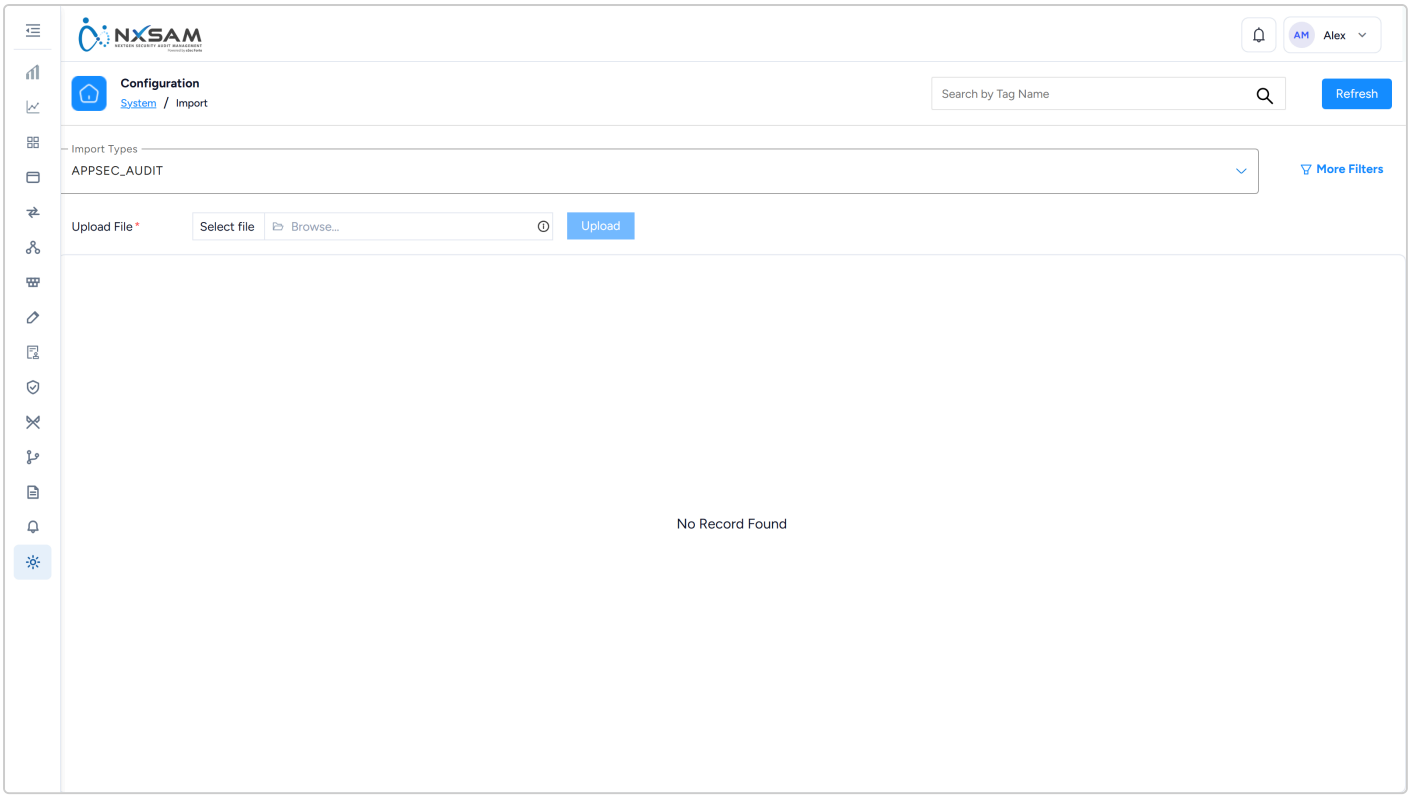


Setting → Health Check → Health Check Configuration

Import Configuration

How bulk imports behave, including what to do with rows that do not match.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

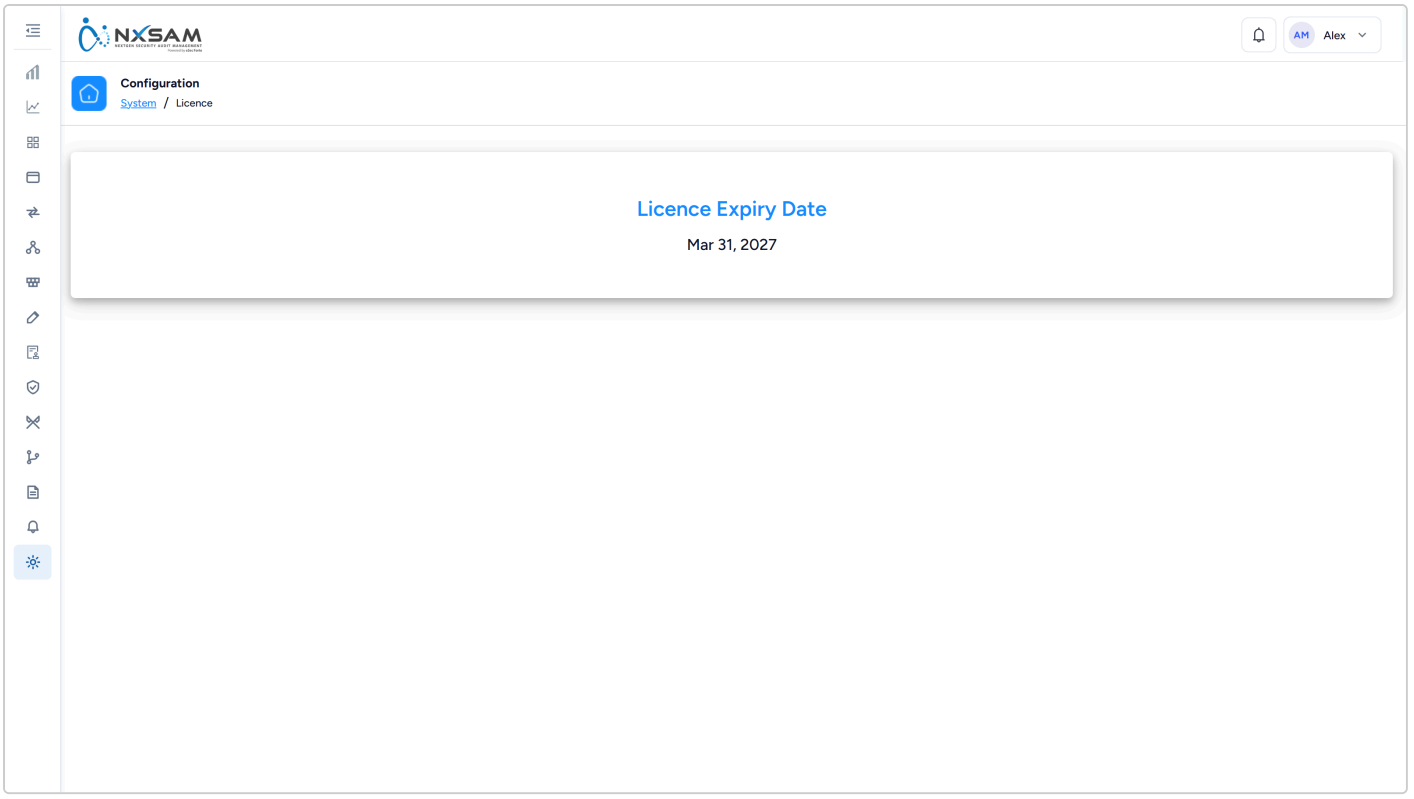


Setting → System → Import Configuration

Licence

The same licence detail, reached from Configurations.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

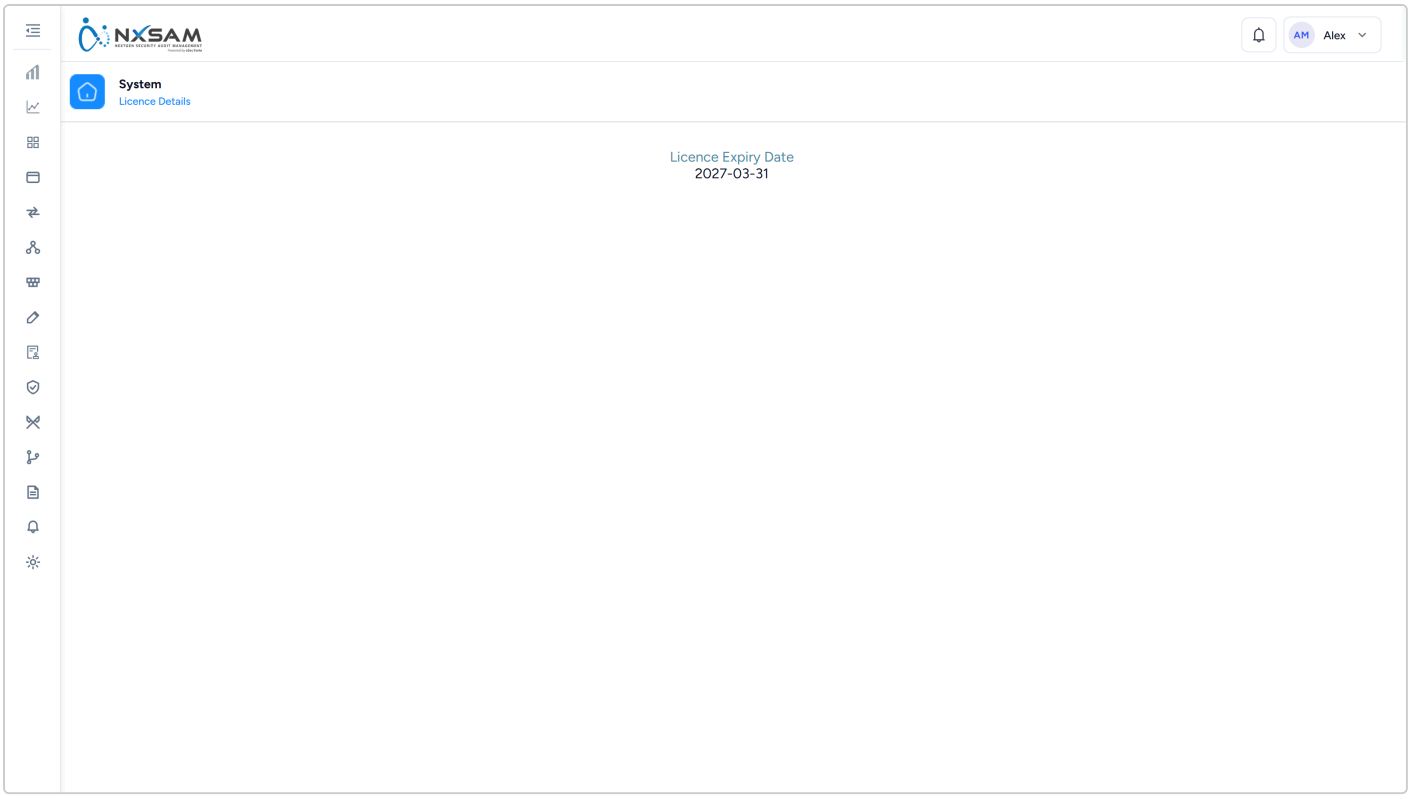


Setting → System → Licence

Licence Details

When the licence expires.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

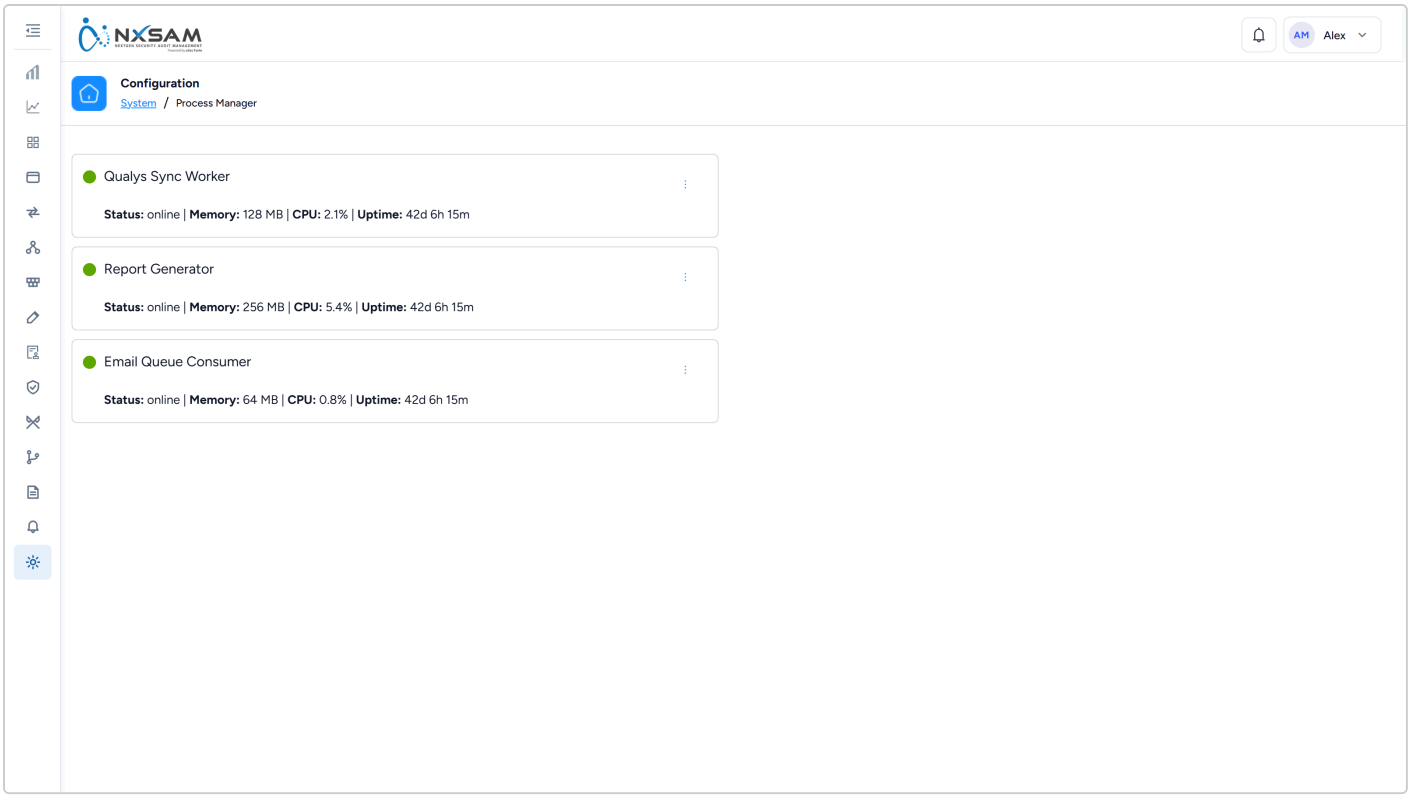


System → Licence Details

Process Manager

Background work — imports, report generation, synchronisation — and whether each job is queued, running or failed.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

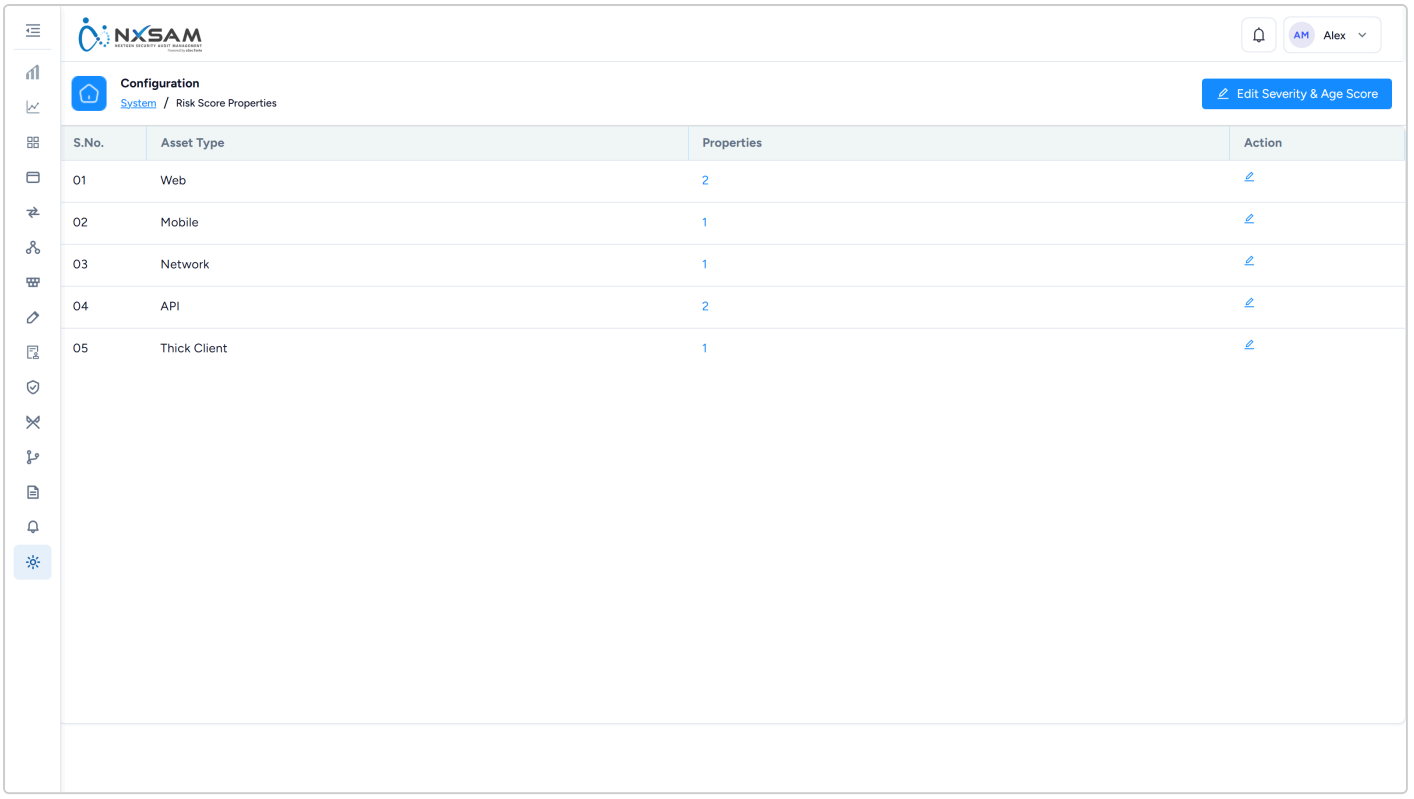


Setting → System → Process Manager

Risk Score Properties

Which asset attributes feed the contextual risk score.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

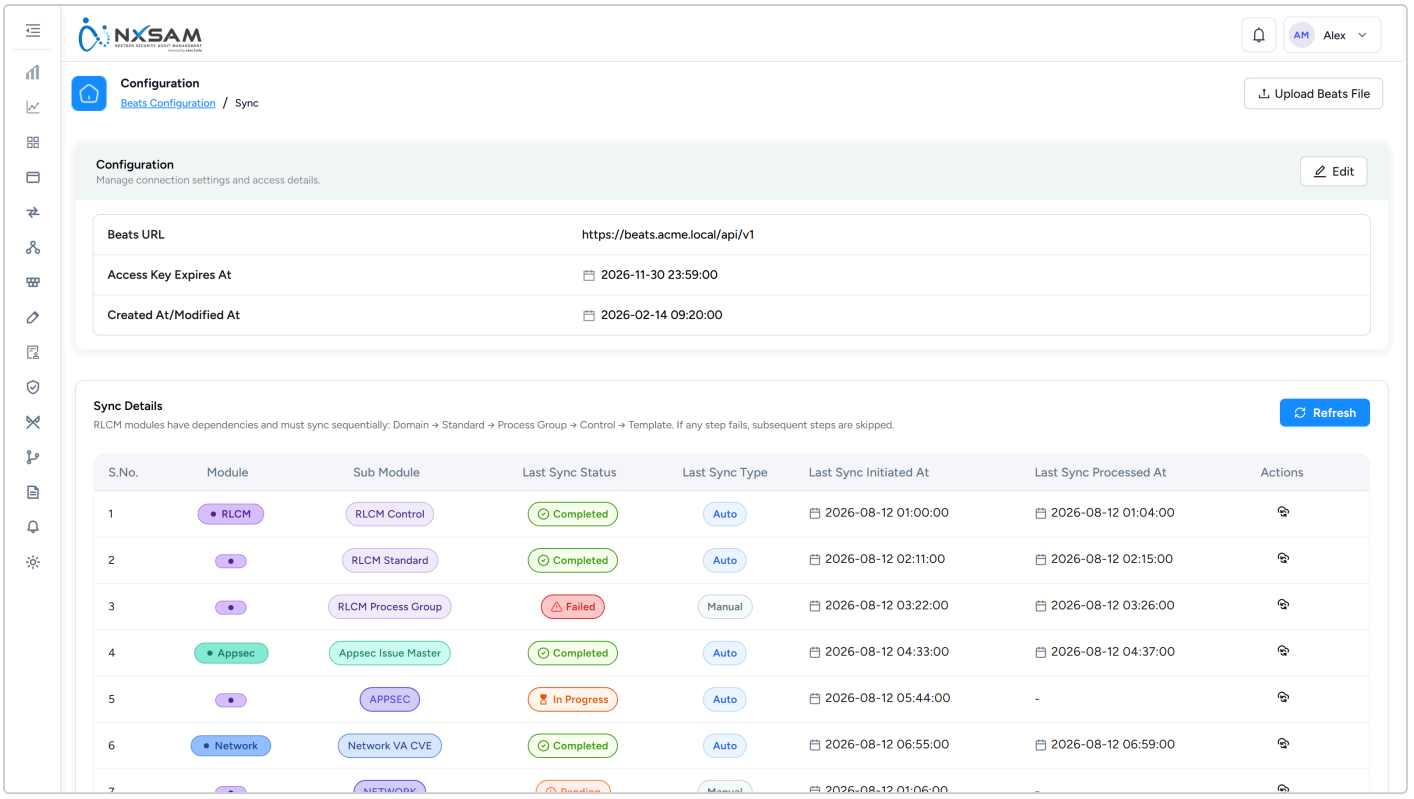


Setting → System → Risk Score Properties

Scheduled Jobs

The periodic background work that keeps derived data current. Figures stale by a fixed interval usually mean a job stopped.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

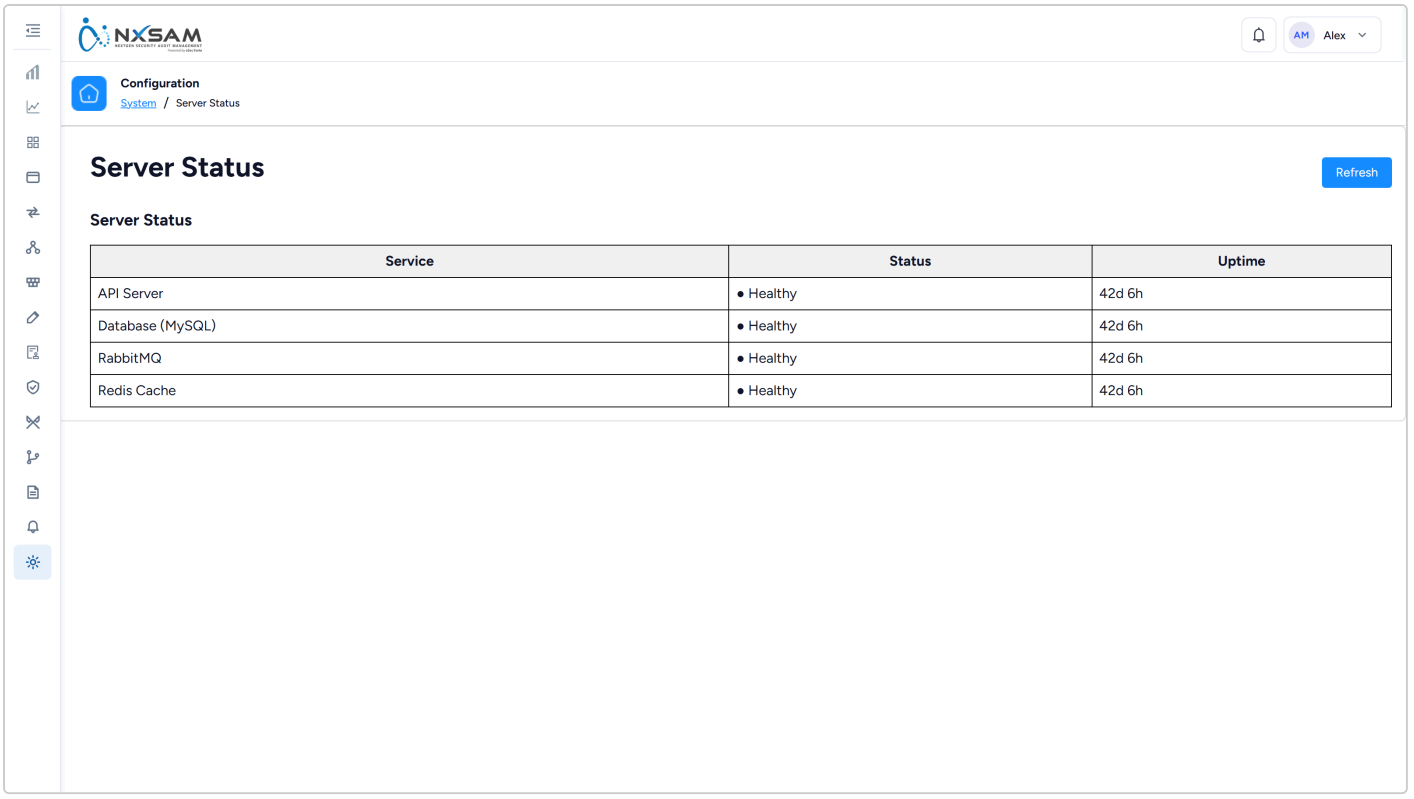


Beats → Scheduled Jobs

Server Status

Whether the platform's components are running.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

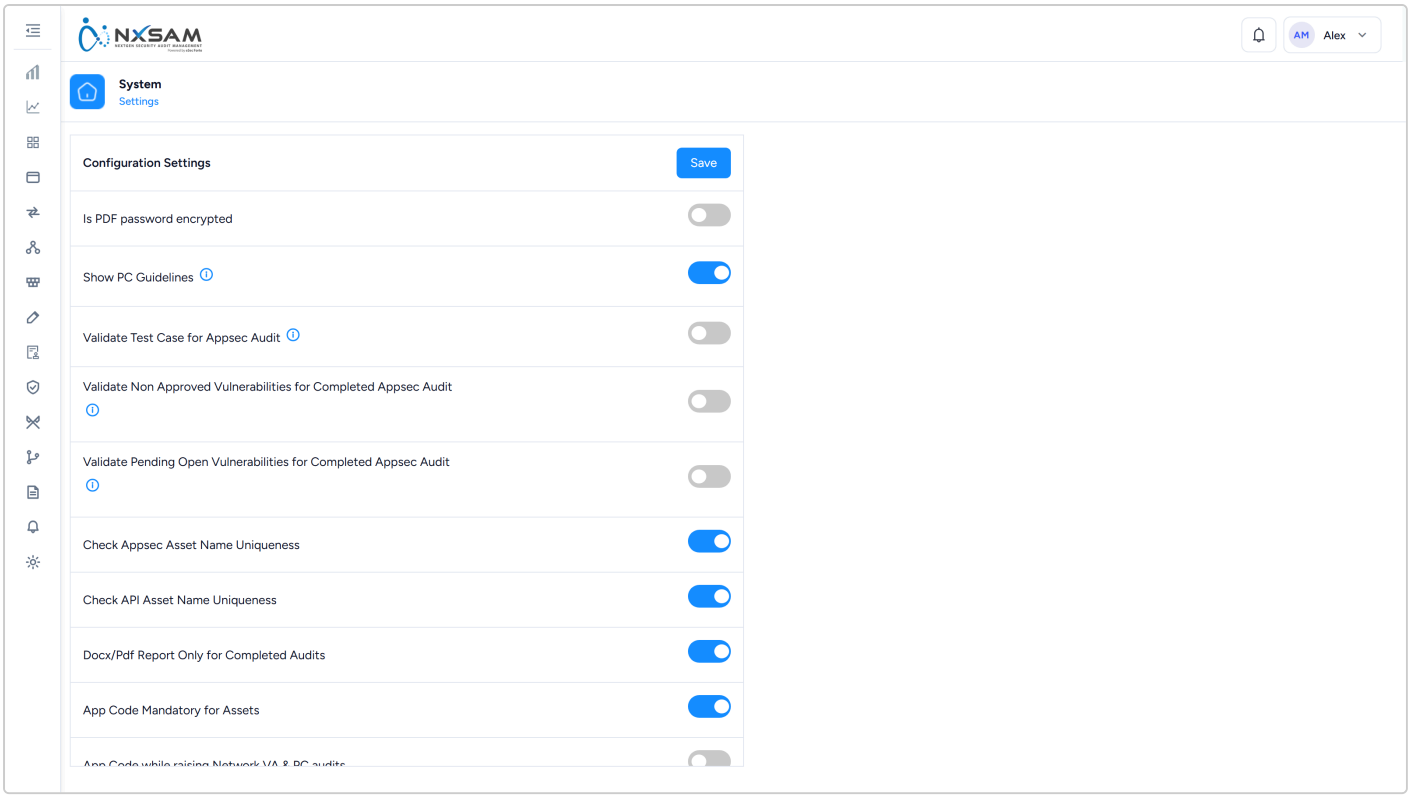


Setting → System → Server Status

System Settings

The platform-wide switches: naming, mandatory fields, report contents, session behaviour. Changes here alter how every module behaves.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

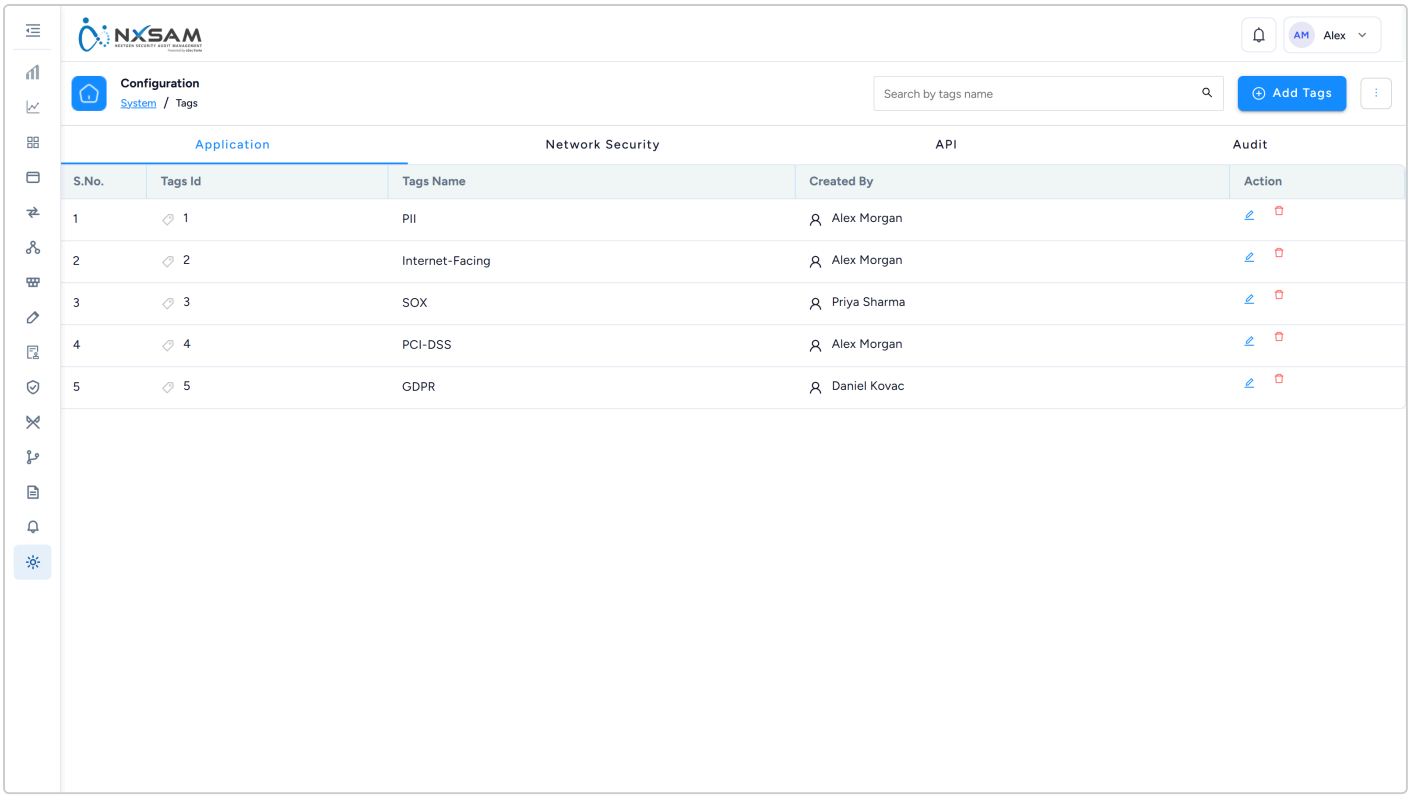


System → System Settings

Tags

Free-form labels applied across assets and audits for grouping.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.

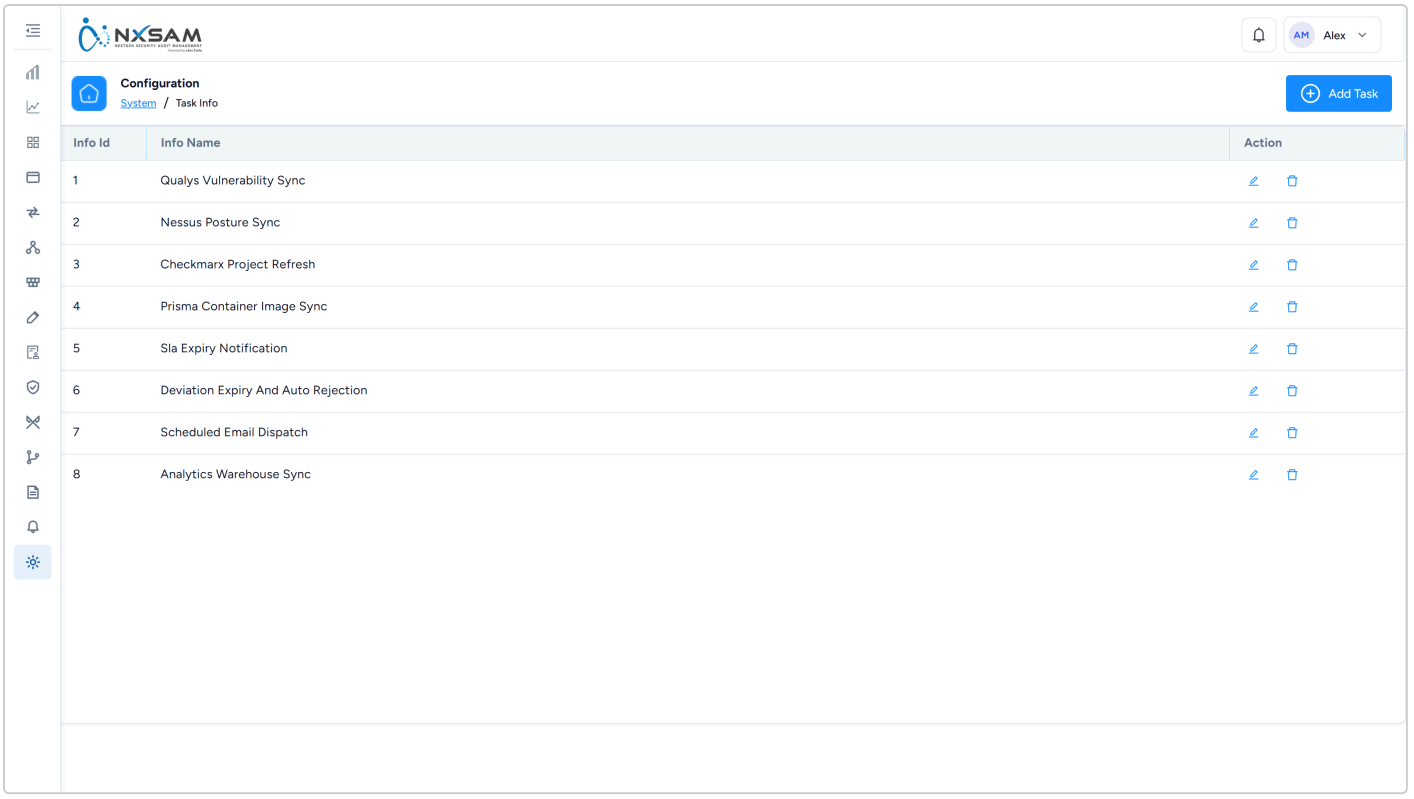


Setting → System → Tags

Task Detail

The detail behind one background job, including the error where one failed.

Copyright 2010-2026 eSec Forte® Technologies Private Ltd.



Setting → System → Task Detail

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.