AttestHub

INDEPENDENT TECHNOLOGY ASSURANCE · AUSTRALIA

Know what's ready. See what comes next.

Independent security reviews for cloud environments, mobile applications, software products, and cybersecurity programs. Practical findings, clear guidance, and assurance tied to a defined scope.

ASSESSMENT FILE · DECISION SUPPORT

Scope
The systems, period, methods, and evidence covered by the review.

Review
Qualified review before any deliverable or attestation outcome is released.

Outcome
Practical findings with limitations and owner responsibilities stated clearly.

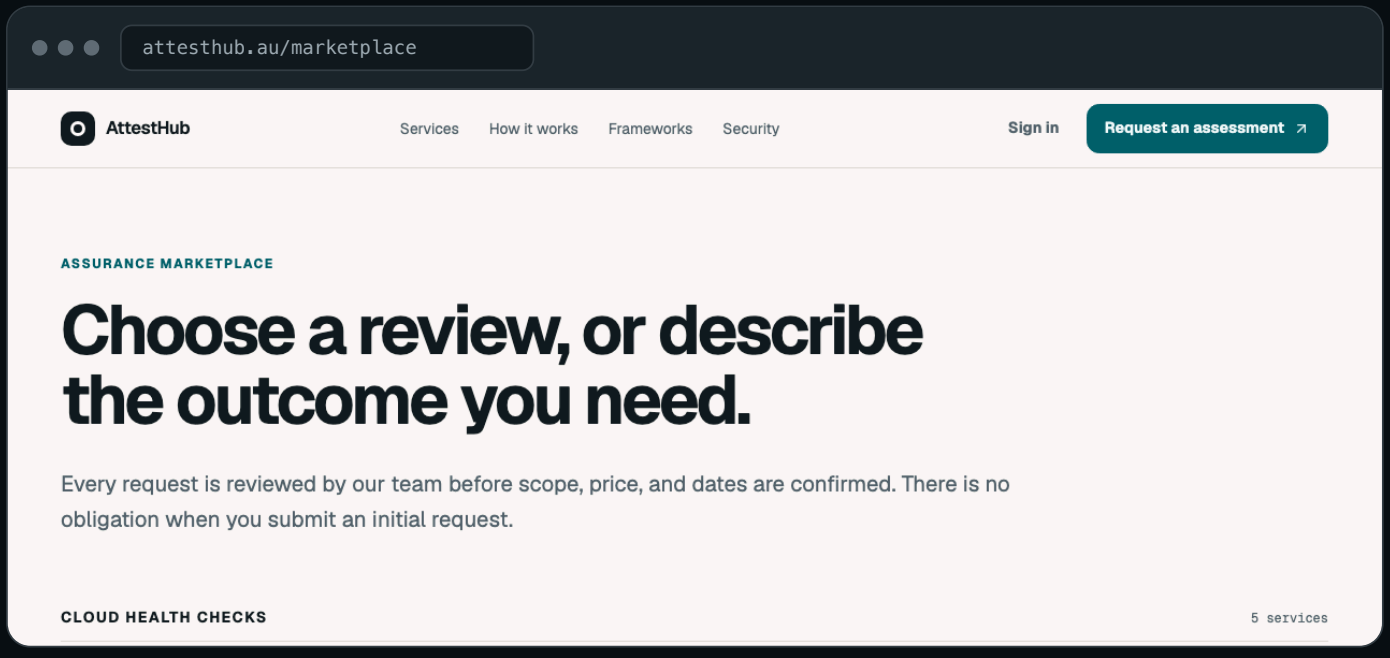

Scope-bound · Evidence-led · Independent quality review · Australian data residency

WHY ATTESTHUB

One assessment file, from request to outcome.

Most organisations know an independent review is required long before they know what to ask for. Start in plain English. AttestHub turns the request into an agreed engagement, keeps the supporting evidence connected, and records what the final outcome does and does not cover. A second qualified reviewer stands between the assessor and the customer.

15	6	9	AU
ASSURANCE SERVICES	SERVICE CATEGORIES	STEP GUIDED REQUEST	DATA RESIDENCY

01	**Describe what needs review**	Choose a service if you know it, or describe the system and the decision you need to make.
02	**Agree the engagement**	Review the proposed scope, price, access, and timeframe before any assessment begins.
03	**Work through the evidence**	Track evidence, findings, and actions, then receive an independently reviewed outcome with its limitations stated.

FOR THE CUSTOMER

- ✓ **Ask in plain English.** You do not need to choose a framework or know the technical name of the review.
- ✓ **No surprise scope.** Scope and quote are versioned and need your explicit approval before work begins.
- ✓ **One next action.** The portal always states what is waiting on you, and what is waiting on us.

FOR THE ASSESSMENT TEAM

- ✓ **A disciplined workspace.** Triage, scoping, quoting, evidence, findings, quality review, and attestation on one record.
- ✓ **Governance by default.** Role-based access, MFA for privileged roles, and hash-linked append-only audit events.
- ✓ **Independence, enforced.** An assessor cannot approve their own deliverable or issue their own attestation.

“Confidence without false certainty.” An assessment applies only to the agreed systems, evidence, methods, and period. The outcome says so, in the same size type as the result.

ASSURANCE MARKETPLACE

Choose a review, or describe the outcome you need.

Fifteen assurance services across six categories, each written so a non-technical buyer can tell what it covers. AttestHub reviews every request before proposing scope, price, and dates. There is no obligation when you submit an initial request.

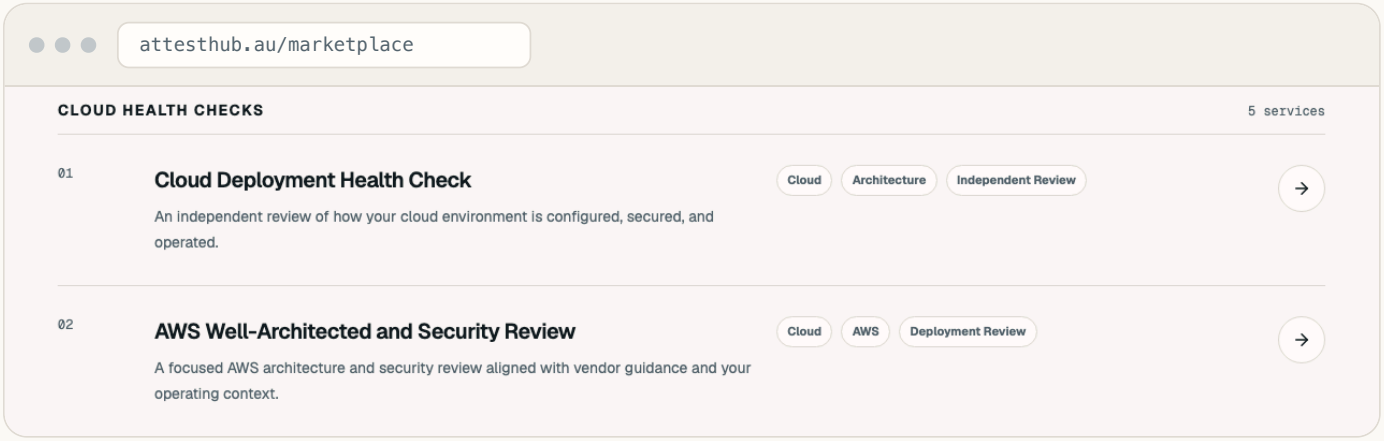

The catalogue is an index, not a checkout. Each service carries a plain-English summary, category tags, an indicative duration and pricing basis, and its own stated limitation. Selecting one starts a request. *This capture predates the AI assurance category listed below.*

CATEGORY	COVERS	SERVICES
Cloud health checks	Cloud deployment health check; AWS Well-Architected and security review; Azure, Google Cloud, and Kubernetes security and production readiness reviews	5
Mobile application reviews	Mobile application security health check; iOS and Android application readiness reviews	3
Product cybersecurity	SaaS product cybersecurity attestation; open-source product security review; AI and MCP product security review	3
Organisation assurance	General cybersecurity attestation; Essential Eight readiness review	2
AI assurance	AI readiness review against the ten Voluntary AI Safety Standard guardrails. Fixed price, no discovery call required	1
Guided assessment	Custom security assessment. Tell us what you built and why assurance is needed, and we recommend a sensible path	1

Stated on every service, before anyone buys

"The final scope, price, methods, and dates are confirmed after triage. The review is point-in-time and does not guarantee that a system cannot be compromised." Durations are indicative (typically two to four weeks for a cloud health check, six to twelve for an organisation-wide attestation), and pricing follows a defined scope rather than a list price.

THE REQUEST

Tell us what needs review.

A nine-step request that adapts its questions to the services selected. It autosaves as a local draft, asks for sensitive detail deliberately, and gives the triage team enough context to scope the work without a discovery call first.

01	**You and your organisation**	This lets us route the request and contact the right person.
02	**What would you like reviewed?**	Choose one or more services. If unsure, choose the guided assessment.
03	**A few service-specific questions**	These questions are tailored to the services you selected and speed up triage.
04	**Tell us about the system**	Plain English is welcome. We will clarify technical detail during triage.
05	**What outcome do you need?**	Understanding why the review matters helps us recommend the right depth.
06	**Scope and access**	Choose the activities you expect. These are preferences, not a final scope.
07	**Timing and commercial context**	Dates and constraints help us identify a viable engagement model.
08	**Supporting evidence**	Optional files can reduce follow-up. They are scanned before anyone can access them.
09	**Review and consent**	Check the summary and confirm the conditions of this initial request.

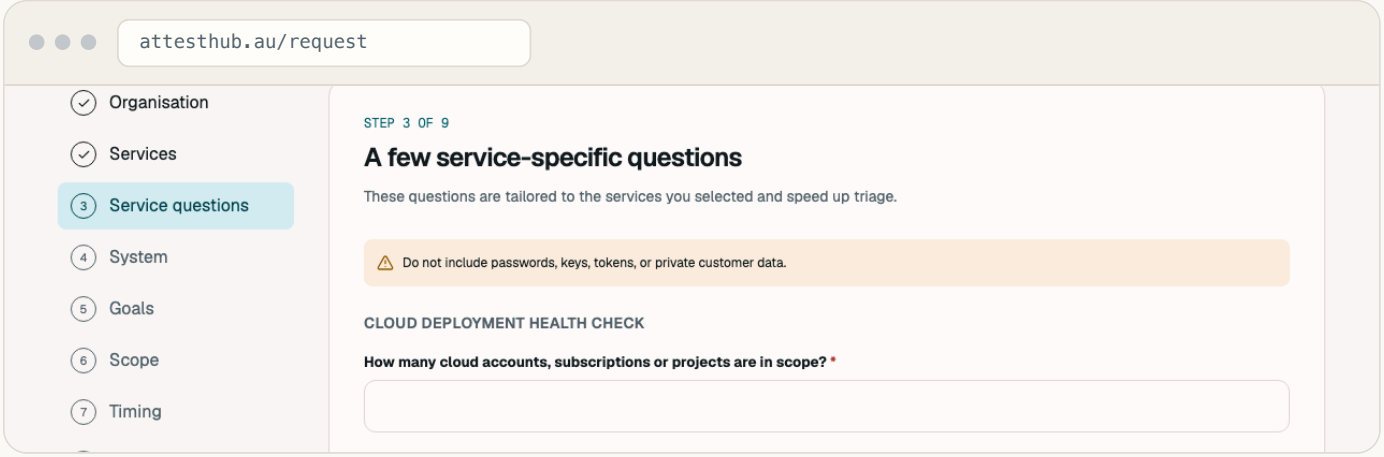

Questions change with the service. A cloud health check asks about accounts, infrastructure-as-code, and console access; a mobile review does not.

Submission is a consent step, not a purchase

Five confirmations are required before the request can be submitted: that the requester is authorised; that the information is accurate; that no credentials or private keys are included; that scope, price, dates, and any attestation remain subject to review; and that the privacy notice and platform terms are accepted.

THE CUSTOMER PORTAL

A secure, trackable engagement, not an inbox thread.

Once a request is submitted it becomes an engagement record the customer can watch. The workspace answers one question on arrival: what needs attention. Everything else (timeline, secure messages, evidence, quote) sits under the engagement it belongs to.

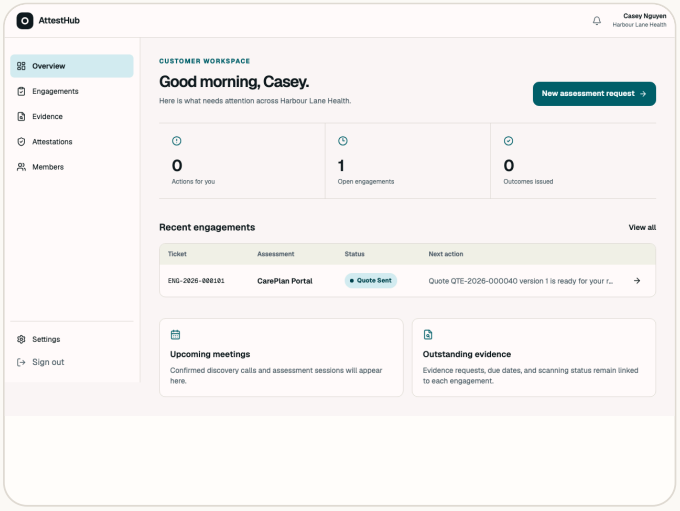

Workspace overview. Actions for you, open engagements, and outcomes issued, with the next action written out in full on every row.

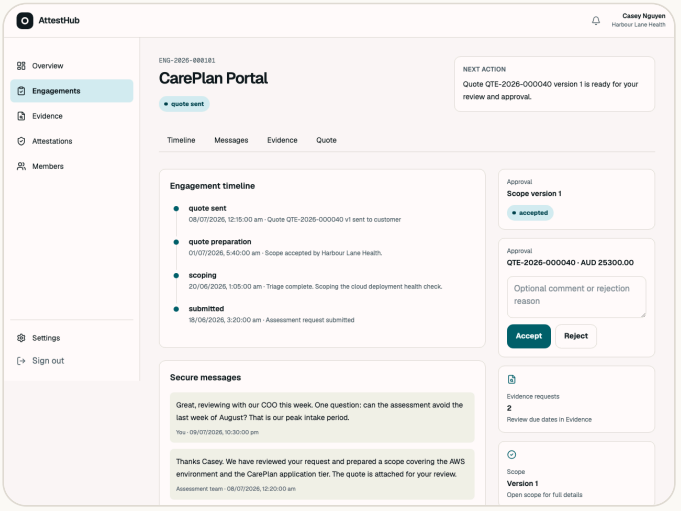

Engagement detail. Timeline, secure messages, evidence requests, and the quote awaiting approval, all against one ticket identifier.

An honest timeline

Submitted, scoping, quote preparation, quote sent. Each entry is dated and attributed, so status is never a guess.

Approval, not assumption

Scope versions and quotes are accepted or rejected by the customer, with an optional comment or rejection reason.

A visible boundary

Secure messages are the customer-visible channel. Internal notes are never shown here, and the portal says so.

Access is role-based and audited, and the participants panel names who can see the file. A customer never has to ask “who else is reading this?”

THE INTERNAL WORKSPACE

Triage, scope, and quote on one record.

The staff view of the same engagement. The current state stays visible while creation and editing sit behind verb-labelled disclosures, so a reviewer reads the file before changing it. Scope and quote are versioned: editing a scope creates a new version rather than quietly rewriting history.

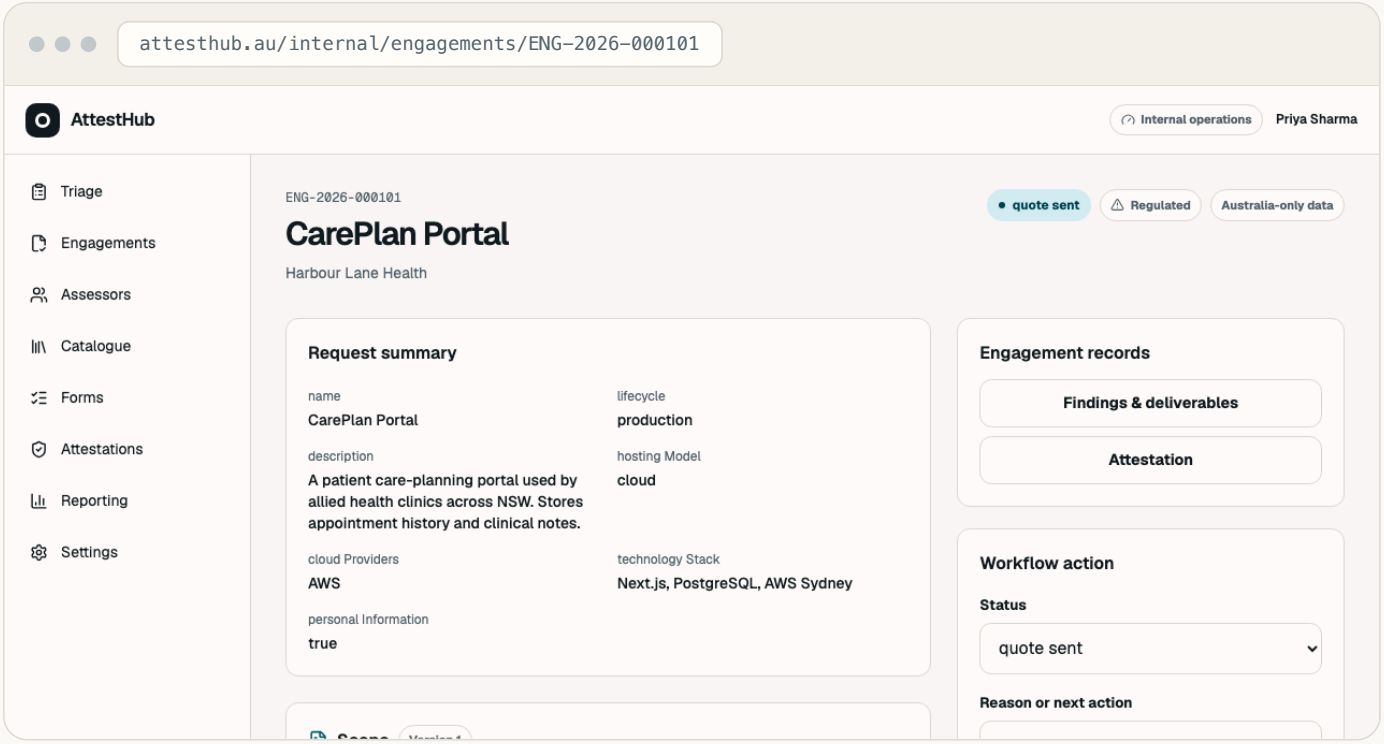

One file, whole history. Request summary, scope version 1, the quote builder with line items, discount and GST, information requests, and the evidence register. The workflow action and internal notes sit to the side.

VERSIONED BY DESIGN

- **Scope versions are immutable.** In-scope, out-of-scope, activities, and deliverables are stated separately, and editing creates a new version.
- **Quote versions.** Priced line items with an expiry date; sending notifies the customer and asks for approval in the portal.
- **Twenty-three workflow states.** Tracked separately for the customer and the internal team, with every transition recorded in status history.

COLLECTED, NOT GUESSED

- **Read-only collectors.** Posture collection from a public repository, with full provenance. No credentials are ever stored.
- **Document imports.** SBOMs as CycloneDX or SPDX JSON, scanner results as SARIF JSON.
- **Conflict of interest.** Triage records carry a conflict review state, so independence is checked before the work is assigned.

EVIDENCE HANDLING

Ask for sensitive things deliberately.

Assurance work needs real artefacts: architecture diagrams, IAM exports, configuration. AttestHub treats every one of them as sensitive by default. Evidence is requested against a named item with a due date and a classification, uploaded to private object storage, and unavailable to anyone until malware scanning succeeds.

Scanned before it opens

Every new object is scanned and tagged with a verdict. A download URL is only issued once the object is tagged clean, and the portal tells the customer so at the upload control.

Short-lived links

Signed upload links expire after five minutes; download links after about one minute. There is no long-lived public URL to leak.

No credentials, ever

Intake and messaging warn against passwords, keys, and tokens, and submission requires that confirmation. Automated collection is read-only and stores no credentials.

Scoped at the key

Object storage is private, and keys begin with the organisation and ticket identifiers. Type, extension, name, and size are checked before upload.

Every request has an owner, a date, and a state

“Current architecture diagram, due 2026-07-22, confidential.” Evidence is not a shared folder. Each item is a tracked request that moves through requested, submitted, under review, and accepted, or else partially accepted, rejected, superseded, expired, or not applicable. Nothing submitted disappears into an inbox, and nothing outstanding is invisible.

Where responsibility sits

Redaction remains the customer’s decision, and AttestHub asks for the least sensitive artefact that answers the question: “redact anything you consider sensitive; we will request specifics later if needed.” File access and classification changes are audited.

EVIDENCE COLLECTORS

Nobody has to hand over the keys to their cloud.

This is the part of an assessment everyone dreads. To review your Microsoft 365 or AWS setup, an assessor normally needs access to it: a new admin account, a read-only role, a key. That single step drags in procurement, legal, and your own security team. AttestHub removes the step entirely.

HOW IT WORKS, IN THREE MOVES

01	**Download the prompt**	Pick your provider and download a ready-made instruction sheet, plus a blank JSON template showing exactly what a good answer looks like.
02	**Your own agent reads your own system**	Give the prompt to an agent that **already has authorised, read-only access** to your environment. It looks, it never touches. AttestHub is never in the room.
03	**Paste the answer back in**	Import the JSON. AttestHub checks it is genuine, works out a preliminary risk posture, and hands it to a human to review.

Why this is the interesting part

Every other way of doing this asks you to trust someone with access. A read-only key is still a key: it can be leaked, it outlives the engagement, and it has to be reviewed, approved, provisioned, and revoked by people who are already busy. The evidence collector inverts it. The only thing that leaves your organisation is a JSON file of non-secret facts you can read before you send it.

WHAT NEVER HAPPENS

- ✓ No account is created for us
- ✓ No key or token is shared
- ✓ No change is made to your systems
- ✓ No access to revoke afterwards

THE SAME EVIDENCE, TWO WAYS

STEP	THE USUAL WAY	WITH AN EVIDENCE COLLECTOR
Getting access	Raise a ticket, provision a role, security-review the assessor's account, approve, and remember to revoke it later.	Nothing to raise. No account exists.
Gathering the evidence	An administrator clicks through admin consoles taking screenshots of settings.	An agent you already trust reads the settings and writes them down.
Getting it right	The assessor asks for the four things that were missed, and waits.	An incomplete file is rejected on import, before anyone waits on it.
Doing it again next year	Start over, and hope it was done the same way last time.	Run the same prompt. The results line up by design.

EVIDENCE COLLECTORS · WHY IT HOLDS UP

A collector that cannot flatter you.

The obvious objection: if the customer runs the collection themselves, what stops the answer being tidied up on the way out? The schema does. AttestHub decides in advance which checks exist and how serious each one is. An import that invents a check, renames one, quietly downgrades a critical to a low, or omits an awkward one is rejected. Not flagged. Rejected.

Severity is ours, not yours

Each check has a fixed severity set by AttestHub. “Severity must be critical.” The importing agent cannot lower it.

Every check, exactly once

A missing check fails the import. You cannot answer the easy five and leave the sixth out.

“I don’t know” is a real answer

An unverifiable check must be marked not available with a reason. Guessing is explicitly forbidden, and it counts against the posture.

TWENTY-FOUR APPROVED CHECKS ACROSS FOUR PROVIDERS

PROVIDER	WHAT GETS LOOKED AT	CHECKS
Microsoft 365	Identity, privileged access, authentication, sharing, email protection and audit configuration	6
Microsoft Azure	Tenant identity, subscriptions, public exposure, encryption, logging and security posture configuration	6
Google Workspace	Administrator access, strong authentication, OAuth apps, external sharing, email controls and audit logs	6
Amazon Web Services	Root and IAM protections, public resources, encryption, logging and security findings across accounts	6

Written into the prompt itself, not left to good intentions

“Work read-only. Before using any command, API, browser action or local tool, explain what it reads and ask the administrator to approve it. Never make configuration changes.” The prompt forbids collecting passwords, cookies, tokens, private keys, or recovery codes; asks for counts and policy names rather than raw exports; and tells the agent to treat everything it finds inside your tenant as untrusted data, so a booby-trapped document cannot redirect it.

Deliberately not overclaimed

One failed critical check makes the whole posture critical; a single unverifiable check keeps it at medium rather than low. The rating is a preliminary signal calculated for human review, and the product says so where you see it: “Evidence is a point-in-time observation and does not itself constitute an attestation.” It speeds up the assessment. It is not the assessment.

VENDORS & SUPPLIERS

The third parties you rely on, and the evidence behind each decision.

Most organisations cannot answer “who has our data, and how do we know they are safe?” without a spreadsheet that went stale a year ago. The vendor register tracks the third parties your organisation relies on, the systems they support, and the assurance evidence behind each risk decision, in the same file as everything else.

START FROM A CURATED CATALOGUE

Sixteen vendors most organisations already use, across nine categories: AWS, Azure, Microsoft 365, Google Workspace, GitHub, Slack, Okta, Stripe, Salesforce, Datadog, OpenAI and more. Each arrives pre-filled with its legal name, category, typical data residency, and recorded certifications.

The product is careful about what that means: “verify the prefilled details against your contract and deployment.” Every field stays editable, and the confirmation asks you to “confirm what your organisation actually shares before saving.”

THEN ANSWER ELEVEN QUESTIONS

A reusable due-diligence questionnaire asks the things that actually decide vendor risk: whether your inputs train their models, where data is processed, whether sub-processors are disclosed, what independent assurance they hold, what happens on a breach, and whether model versions can be pinned before they change underneath you.

Answers are weighted and scored server-side into a rating of low, medium, high, or extreme. The score is always recomputed from the answers. A client cannot declare its own rating.

Supply chain, rolled up honestly

Vendors link to the systems they support, each with a role (model provider, hosting platform, data source, or integrator), and the system view surfaces the highest linked vendor risk. The product is precise about its own authority: “an advisory signal that feeds risk-tier consideration; it does not change the system’s assessor-set tier automatically.” A machine can raise a hand. Only a qualified assessor sets the tier.

TRACKED PER VENDOR

- ✓ Type: model provider, platform, integrator, or data supplier
- ✓ Status: active, under review, suspended, or retired
- ✓ Certifications, data residency, and sub-processors
- ✓ The systems each vendor sits behind

SCORING YOU CAN AUDIT

Low	Medium	High	Extreme
0–25%	26–50%	51–75%	76–100%

Share of the 25-point scale triggered by worst-case answers.

What the register is, and is not

The catalogue is a starting point, copied into your register once and yours to edit from then on. It is not a live feed. AttestHub does not monitor trust centres, re-check certifications, or discover sub-processors for you.

INFORMATION REQUESTS

Reusable questionnaires that never change under a respondent.

Structured questions are how an assessment gets consistent input. AttestHub ships a template library and a builder, with starters for common engagements. Recipients get a secure one-off link (no account needed), and a sent form keeps a snapshot, so editing a template never changes forms already in flight.

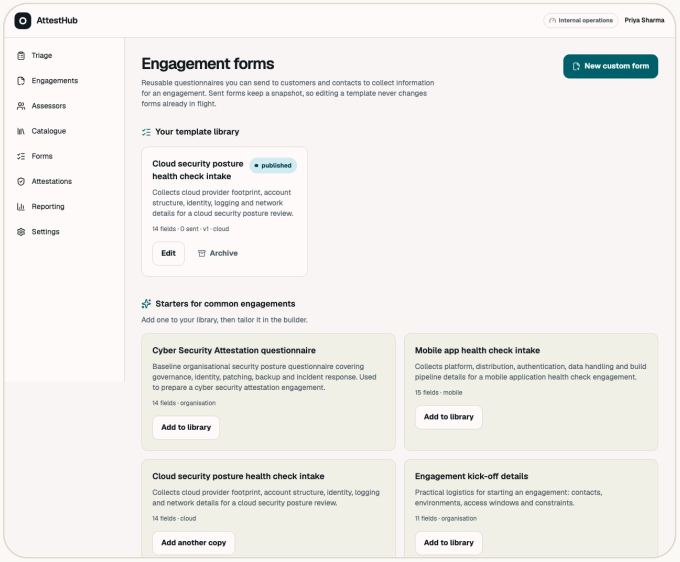

Template library. Published templates with field counts and category, plus starters to add and tailor: cyber security attestation questionnaire, mobile app intake, cloud posture intake, engagement kick-off.

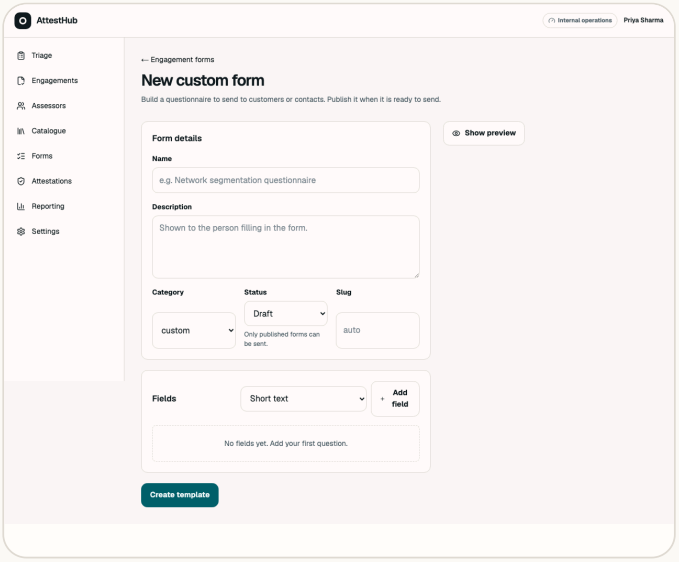

The builder. Compose a custom form for an engagement that does not fit a starter, then publish it into the library for reuse.

Secure one-off links

Send a form to a customer or a third-party contact without provisioning them an account.

Snapshotted on send

In-flight forms are frozen at the version sent, so answers always match the questions that were asked.

Tracked to a due date

Optional message and due date, with the response landing against the engagement it belongs to.

FINDINGS & DELIVERABLES

Nobody releases their own work.

Findings carry a severity from critical through to positive practice, and a lifecycle of their own: open, accepted risk, remediation planned, ready for validation, closed. Deliverables move through a fixed path, and the independence invariant in the platform prevents an assessor from approving their own deliverable or issuing an attestation.

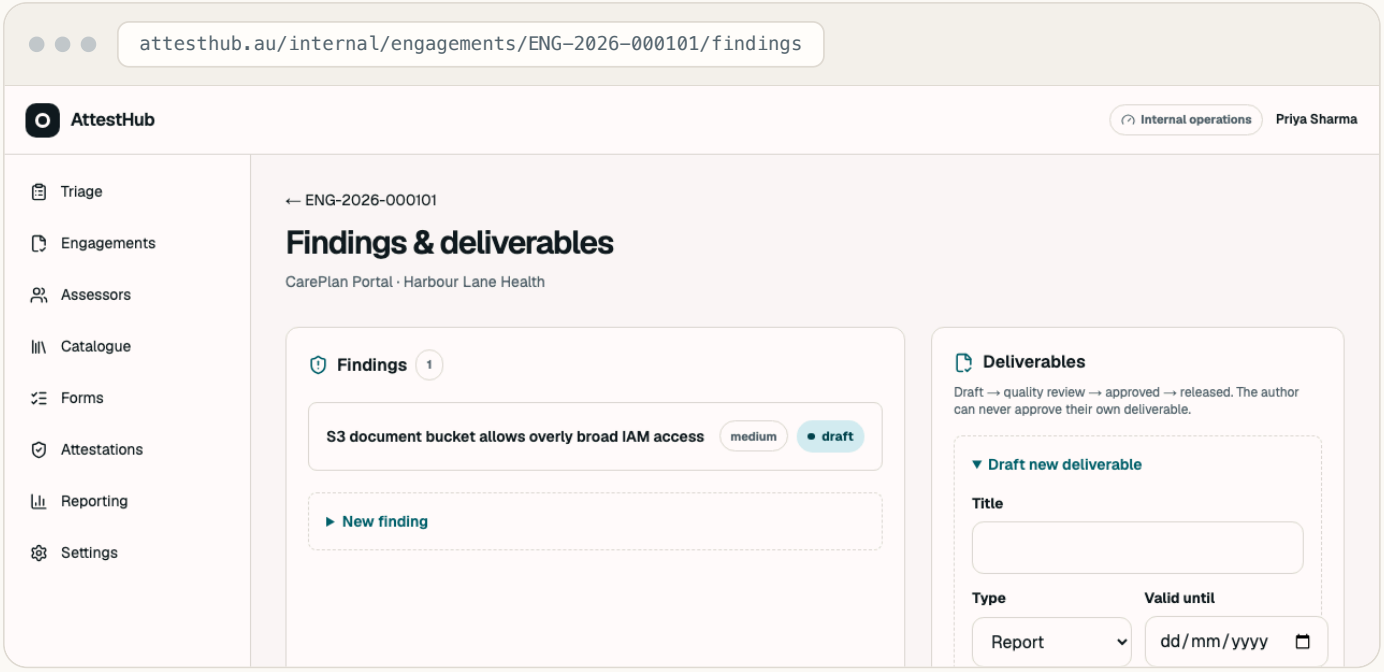

The gate is in the product, not the process document. The deliverable panel states the path and the separation-of-duties rule at the point of authoring.

`draft` `quality_review` `approved` `released` `withdrawn`

Draft	The assessor writes the deliverable against a saved scope version, with a type, an assessment period, and a validity date. The deliverable is pinned to that scope version, so a report can always name what it covered.
Quality review	A second qualified reviewer examines the work. This is a state in the record, not an informal favour.
Approved	Approval is recorded against a person who is not the author. Issuance additionally verifies that the quality reviewer differs from the assessor.
Released	Only then does the outcome reach the customer, carrying its scope, evidence period, and limitations.

ATTESTATIONS & VERIFICATION

A claim anyone can check, and no claim we cannot support.

An attestation is scope-bound and point-in-time. It carries a validity window, a visibility setting that defaults to private, and a revocation state. A third party can confirm an identifier is genuine, current, and unrevoked without seeing anything the organisation did not choose to publish.

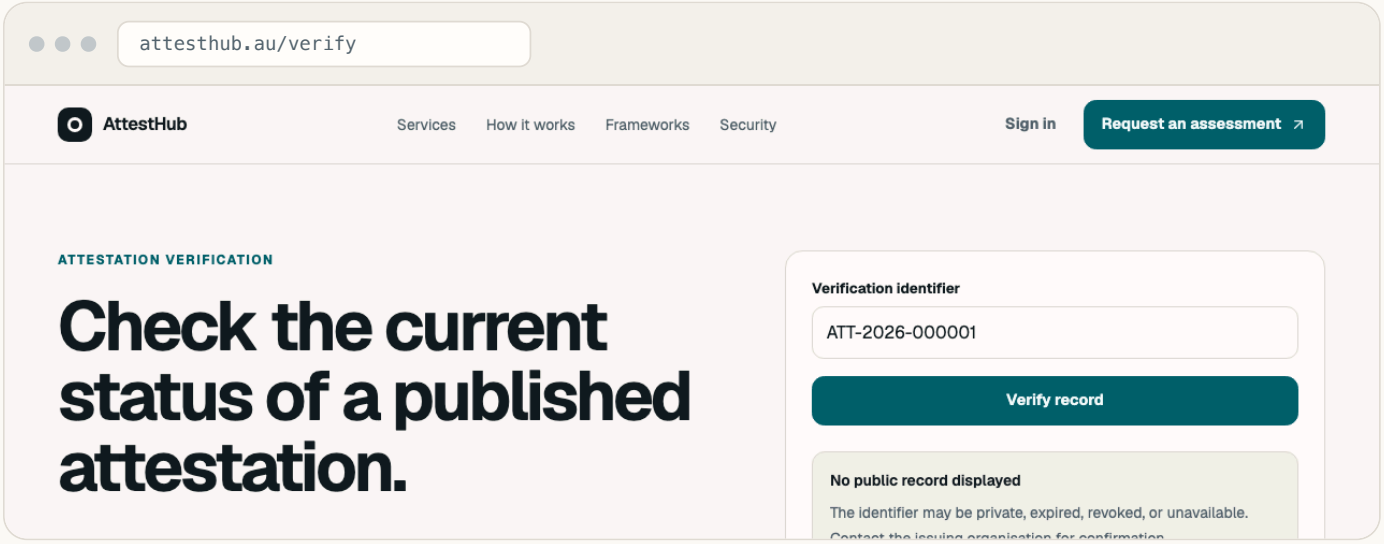

Verification is deliberately narrow. Enter an identifier such as `ATT–2026–000009`. Where a record is private, expired, revoked, or unavailable, the page says so plainly and refers the enquirer to the issuing organisation rather than implying a failure.

A signed certificate

Each record can be opened or downloaded as a signed certificate, with a SHA-256 integrity hash shown alongside it.

Graded outcomes

Attested, attested with qualifications, partially attested, or not attested. The register does not round up.

Revocable

Revocation is a first-class state carrying a reason and a date, and verification reflects it immediately.

What an attestation is not

Verification confirms the status of a record. It does not certify that a product is secure in all circumstances, and it is not a substitute for a certification scheme. We do not claim formal certification unless an authorised certification body performs that work. The assessment applies only to the agreed systems, evidence, methods, and period; ongoing security remains with the system owner.

Checks are logged for abuse detection against a hash of the enquirer's IP address, never the address itself.

AI ASSURANCE

AI assurance against the ten guardrails, checkable by anyone.

AttestHub provides independent, point-in-time AI assurance: readiness reviews against the Voluntary AI Safety Standard, security reviews of AI products and Model Context Protocol tool servers, and publicly verifiable, scope-bound AI assurance attestations.

AI Readiness Review

A fixed-price review of your AI use against the ten Voluntary AI Safety Standard guardrails, with clear gaps, a risk tier, and a prioritised roadmap. Typically one to two weeks, and the only service that needs no discovery call.

AI and MCP Product Security Review

A focused security review for AI products and MCP servers, including tool boundaries, data handling, and deployment controls. Typically three to six weeks, scoped by product and tool surface.

FRAMEWORKS ASSESSED

FRAMEWORK	PUBLISHER	VERSION
Voluntary AI Safety Standard	Australian Government (National AI Centre / DISR)	2024
NIST AI Risk Management Framework	US National Institute of Standards and Technology	1.0
ISO/IEC 42001 (AI management system)	ISO/IEC	2023

What an AI attestation covers

An AI attestation records the models assessed by provider, name, and version; the guardrails covered; the intended and prohibited use of the system; the human oversight and accountability arrangements; and the assessment scope, evidence period, and validity window. A public transparency statement can accompany it, rendering a public-safe subset that omits the owner, business unit, data categories, and private risk detail.

CAPABILITIES

- ✓ Readiness against the Voluntary AI Safety Standard guardrails
- ✓ Security review of AI products, agents, and MCP servers
- ✓ AI system register, risk assessment, and monitoring
- ✓ Publicly verifiable, scope-bound attestations

Stated on the service itself

“The review assesses readiness against the Voluntary AI Safety Standard guardrails; it is not a certification and does not itself make an AI system safe or compliant.”

PLATFORM SECURITY

The platform that assesses security has to demonstrate its own.

AttestHub is multi-tenant. Every customer record carries an organisation identifier, and authorisation is checked in server actions, route handlers, the data-access layer, background jobs, object keys, and exports. Proxy checks are treated as an optimistic redirect only, never as the security boundary.

Authentication that resists reuse

Passkeys, TOTP MFA with recovery codes, magic links, and a fourteen-character password floor. Passwords found in a known breach are rejected, and an account with a passkey cannot fall back to a password.

Tenancy and roles

Organisation tenancy with role-based access. MFA is required for assessors, engagement managers, platform administrators, and organisation administrators with sensitive access. Sessions expire after eight hours.

Hash-linked audit

Material actions append an audit event with actor, organisation, resource, time, reason, and a hash of the preceding event. Append-only, hash-linked per organisation, and verified nightly so tampering or deletion is detectable.

Hardened by default

A nonce-based content security policy, clickjacking and MIME-sniffing protection, cross-origin isolation, and a restrictive permissions policy. Inputs are schema-validated and queries parameterised.

Running it in production

Structured logging redacts passwords, tokens, keys, and document content. Rate limiting is persisted rather than held in memory, and outbound integration secrets are encrypted at rest with AES-256-GCM and decrypted only by the delivery worker. Backups use point-in-time recovery with object versions retained independently of the database, media encrypted with keys managed outside the workload, and quarterly restore tests that verify object checksums and engagement-to-file references, preserving the selected data region and any legal holds.

Stated plainly

Data is stored in the Australian region by default, and an assessment request can require that data remain in Australia as a condition of the engagement. Production deployments add network-level rate limiting and account-lockout alerting at the edge.

INTEGRATION

Metadata where you need it, evidence where it belongs.

AttestHub exposes a versioned REST API with scoped API keys, and a metadata-only MCP server so an assistant can answer questions about engagements without ever touching the underlying files.

Versioned REST API	Organisation-owned API keys, scoped to only what the integration requires.
MCP server	Tools return service, ticket, scope, finding, deliverable, and attestation metadata. They never return evidence or report file contents.
Imports	SBOMs as CycloneDX or SPDX JSON; scanner results as SARIF JSON; read-only posture collection from a public repository.
Verification endpoints	Public verification and signed certificates, rate-limited and served under a locked-down content security policy.
Observability	Structured logging and OpenTelemetry registration, with OTLP export configured for production.
Deployment	Containerised, with reviewed forward-only database migrations. Development secrets are never carried into production.

The boundary is the point

A scoped key and a metadata-only tool surface are not limitations to work around. They are the reason a customer can connect AttestHub to their own systems without widening the blast radius of an assessment. Credentials, private keys, tokens, and sensitive document content never belong in logs or intake fields, and the platform is built so they do not end up there by accident.

NEVER RETURNED BY THE API

- ✓ Evidence file contents
- ✓ Report file contents
- ✓ Internal notes
- ✓ Credentials of any kind

THE COMPLETE PICTURE

What is built, and what is honestly still ahead.

A deck that hides its roadmap is the first thing an assurance buyer stops trusting. This is the current release scope of the platform.

IN THE CURRENT RELEASE

- ✓ Marketplace discovery and adaptive request submission
- ✓ Account, organisation setup, and role-based access
- ✓ Customer portal and internal triage workspace
- ✓ Versioned scopes and quotes with approval records
- ✓ Secure uploads, evidence metadata, forms, and status history
- ✓ Findings, deliverables, and attestation entities
- ✓ Audit logging, REST and MCP extension points
- ✓ Containerised deployment and operational reporting surfaces

ROADMAP, NOT YET BUILT

Payments	Quotes are approved in the portal with purchase order details; payment processing can be added later.
Automated cloud collection	Read-only collection is limited to public repositories and document imports.
Advanced report generation	Deliverables are authored and versioned; generation is manual.
Public badges	Verification is by identifier against the register.
Calendar providers & SAML	Meetings and enterprise single sign-on are planned.
Native mobile apps	The product is responsive on mobile web.

The platform is built so the governance is real before the automation is: versioning, approvals, audit, and separation of duties come first, because those are the parts an assurance outcome actually rests on.

AttestHub

READY WHEN YOU ARE

Assurance should clarify the decision in front of you.

Independent reviews for the systems your organisation depends on. Scope-bound, evidence-led, and written in plain English. AttestHub reviews every request before proposing scope, price, and dates. There is no obligation when you submit an initial request.

Not sure what you need?
Choose the guided assessment. Tell us what you built and why assurance is needed, and we will recommend a sensible path.

Checking someone else's claim?
Verify an attestation identifier against the public register, or browse the attestations organisations have chosen to publish.

An assessment applies only to the agreed systems, evidence, methods, and period. It is point-in-time and does not guarantee that a system cannot be compromised. Ongoing security remains with the system owner.

attesthub.au
hello@attesthub.au

Scope-bound · Evidence-led · Australian data residency