AssetsCompliance

AssetsCompliance v1.1 — Current Demo Release Documentation

CURRENT DEMO Current demo release

The current AssetsCompliance Practitioners Lite demo: an empty-by-design, practitioner-mounted workspace built around three independent Control Positions, explainable Exposure and Intelligence, and a redesigned review journey.

Version 1.1 Documentation generated 2026-08-27

OFFICIAL DEMO VIDEO

AssetsCompliance v1.1 Demo | Control Infrastructure for Regulated Organisations
https://www.youtube.com/watch?v=_-Otvhzyty4

Contents

1. Version overview

AssetsCompliance v1.1 (AssetsCompliance Practitioners Lite) is the current demo release of the practitioner application (`ac-en`). Its code boundary begins immediately after v1.0's explicit `chore: checkpoint v1.0 end` / `chore: checkpoint v1.1 start` commit pair in the `ac-en` repository history.

2. Release purpose

v1.1 exists to make the demo's connected control model defensible under independent professional review: it rebuilds the review experience around three independently reasoned Control Positions, an empty-by-design mounted workspace, and explainable evidence, exposure, and reporting.

3. What changed from v1.0

v1.1 kept v1.0's connected data model — obligations, controls, ownership, evidence, actions, reporting, and integrations — but changed how a practitioner arrives at and reasons about that data:

- Every workspace now starts **empty by design**; the practitioner mounts (and can unmount) simulated compliance data sources before any record appears, and workspace identity is derived from what is actually mounted.
- v1.0's single generic control view was replaced with three independently reasoned **Control Positions** — Organisation, Practitioner, and Regulator — that can agree, partially agree, or diverge.
- v1.0's four separately named `review-v1` journeys were replaced with one redesigned **seven-step validation journey** (`practitioners-lite-v1-1-validation` , review schema `2.1`) that opens with "Mount your compliance data" ahead of six professional-decision steps.
- Evidence gained an explicit quality basis (sufficiency, relevance, reliability, confidence, conflict); control assessment became frequency-aware; Exposure gained structured causes and full traceability; Intelligence became explainable without an opaque score; and remediation gained an explicit effectiveness judgement with append-only reassessment history.

4. Demo narrative

The v1.1 demo starts from an empty workspace, walks through mounting a simulated data source, and then follows the practitioner through establishing the three Control Positions, investigating a material Exposure, evaluating remediation and reassessment, testing Intelligence as decision support, and defending the professional report.

5. Verified capabilities

Confirmed by the `ac-en` v1.1-era commit history and its `CLAUDE.md` project record:

- Empty-by-design workspace with simulated data mounting and unmounting
- Independent Organisation, Practitioner, and Regulator Control Positions, with agreement, partial agreement, and divergence preserved as meaningful information
- Data Readiness and Workspace Health assessment
- Frequency-aware control execution assessment
- Evidence-quality basis: sufficiency, relevance, reliability, confidence, and conflict, visible on the evidence record itself
- Explainable Exposure with structured causes and full traceability
- Explainable practitioner Intelligence without an opaque score
- Remediation-effectiveness judgement, distinct from completion, with append-only reassessment history
- A professional three-position report with remediation effectiveness, reassessment, change history, traceability, and print presentation
- A redesigned seven-step practitioner validation journey opening with data ingestion

6. Empty-by-design behaviour

A v1.1 workspace shows no obligations, controls, ownership, evidence, actions, or reports until the practitioner deliberately mounts a data source. This is a review-journey design choice: it lets a practitioner see exactly how much of the control picture depends on what has actually been connected, rather than reviewing a pre-populated demo that could obscure that dependency.

7. Data mounting and unmounting

Mounting a source is a simulated ingestion journey: selecting a source type, walking through a deterministic mock dataset, reviewing a data-readiness and mapping assessment, and mounting the result into the workspace's live records. Unmounting reverses this. No real file uploads, real OAuth connections, or external network calls occur during this step — it stays entirely within the demo boundary.

8. Workspace and source scoping

Workspace identity, Open Work, and every downstream view are scoped to whichever sources are currently mounted for that workspace. Nothing is fabricated or assumed to be present because a route exists; presence is always derived from mounted state.

9. Open Work, Data Readiness, and Workspace Health

Open Work is the practitioner's workspace-scoped worklist, derived from currently mounted data. **Data Readiness** assesses, capability by capability (obligation mapping, control coverage, evidence confidence, execution assessment, each Control Position, exposure analysis, and reporting), how ready the mounted data is — rated Ready, Limited, Partial, or Not established, each with a plain-language reason. **Workspace Health** gives a workspace-level view of that same readiness.

10. Regulator, Organisation, and Practitioner Control Positions

Control Position is v1.1's central concept, and it is deliberately three independent positions rather than one:

1. **Practitioner Control Position** — independent professional judgement: rationale, confidence, evidence reliance, assumptions, limitations, challenge, reassessment.
2. **Organisation Control Position** — the organisation's actual operational control state: obligations, policies, risks, controls, execution, evidence, ownership, actions.
3. **Regulator Control Position** — position assessed against applicable regulatory obligations and expectations: what must be demonstrable, uncertainty, exposure, potential breach.

These positions may agree or diverge, and divergence is treated as meaningful product intelligence, never silently resolved away.

11. Control assessment

Control assessment is frequency-aware: it evaluates whether a control executed at the frequency its design requires, not only whether it exists on paper.

12. Evidence assessment and quality basis

Each evidence record carries an explicit quality basis: sufficiency, relevance, reliability, a qualitative confidence rating, and any recorded conflict with other evidence — visible directly on the evidence record, not only inferred from downstream Data Readiness or Intelligence flags.

13. Exposure, intelligence, and traceability

Each material **Exposure** carries a structured cause, a priority basis, and a traceability chain back through the obligation, control, execution, and evidence that produced it, using live, working links. **Practitioner Intelligence** prioritises these exposures and states its trigger, significance, and required action in plain language, without an opaque displayed score.

14. Reporting and review workflow

The professional report presents all three Control Positions together with their divergence, connected evidence, exposure, remediation effectiveness, reassessment history, and change over time. Reporting distinguishes what AssetsCompliance can generate right now from the current mounted workspace state from any separate mounted report-record catalogue, so availability is always derived from real data.

15. End-to-end practitioner journey

The v1.1 validation journey (`practitioners-lite-v1-1-validation` , product `1.1` , review schema `2.1`) has seven steps:

1. Mount your compliance data
2. Establish the three Control Positions
3. Investigate a material Exposure
4. Evaluate remediation and reassessment
5. Test Intelligence decision support
6. Defend the professional report
7. Record professional reliance

Historical v1.0 (`review-v1`) journey definitions and storage remain separately labelled and are never migrated into v1.1 state.

16. Review-stage boundaries and limitations

- v1.1 is a demo/review-stage release open for independent practitioner review, not a certified production application, not legal advice, and not a compliance certification.
- Three formal practitioner reviews and a follow-up review meeting informed this release, but none of them has approved v1.1 — it is release-ready for review, not reviewer-signed-off.
- The "Mount your compliance data" journey is entirely simulated: no real file uploads, real OAuth, or external network calls occur.
- Deeper Pro-only capabilities (for example lines-of-defence permissions, outsourced-control assurance, bulk remediation workbenches, or continuous monitoring) are deliberately excluded from this release as requiring enterprise configuration, not because they are unfinished Lite work.

17. Official v1.1 demo video

AssetsCompliance v1.1 Demo | Control Infrastructure for Regulated Organisations
https://www.youtube.com/watch?v=_-Otvhzyty4

18. Approach to documenting future releases

Later demo releases, and their documentation, will be added to this same `docs/releases/` structure and to the release history on `/demo-access` and `/resources` as they are shaped through practitioner review. No version number, release date, or feature set for a future release is committed to in this document.