



GAP-ANALYSIS RECONCILIATION & PLATFORM OVERVIEW

What was asked. What was built. What it proves.

The System of Proof is the Ekselens parent discipline — thought-leadership plus evidence discipline. It is delivered through two arms: **EKSELENS360** (function methodology) and **Enterprise360** (value-stream methodology). This paper reconciles the 04/07/2026 Gap Analysis against the platform that runs them. Written to be checked, not just read.

Ekselens Consulting › **System of Proof** › EKSELENS360 & Enterprise360 › ERIO

Prepared for

External reviewer / stakeholder

Basis

Gap Analysis 04/07/2026 vs live codebase

Issued

06 July 2026 · v1.2

Contents

01	Executive summary	3
02	Reconciliation at a glance	4
03	How to read this document	5
04	Reconciliation scorecard — every pillar	6
05	What progressed to scope	9
06	Where the platform now exceeds scope	10
07	Gaps the Gap Analysis flagged — now closed	11
08	What is deferred (and why)	12
09	Defensibility — 'let them check, it holds'	13
10	Single source of truth — how drift is made impossible	14
11	Regulatory readiness	15
12	Security, data handling & tenant assurance	16
13	Part B — Architecture for a technical reviewer	17
14	Part B — Key API surface	18
15	Part B — Verification appendix	19
16	Executive demo flow (20 minutes)	20
17	Assurance statement & version control	21

00 • THE HEADLINE

Executive summary

Read against its own yardstick, the platform meets the stated transformation ambition and moves ahead of it in the two areas the Gap Analysis itself named as the real moat — the evidence graph and the closed assurance loop. All items on the document's own post-build scorecard now carry either an implemented control or a documented deferral.

MET TO SCOPE**8**

core spine pillars
implemented

EXCEEDS SCOPE**6**

beyond the stated
requirement

GAPS ADDRESSED**7**

doc-flagged red/amber, now
implemented

On the Gap Analysis's own "bottom line" test — reduce the uncertainty, friction and risk between knowing change is needed and successfully delivering it — the **diagnose → design → plan → mobilise → deliver → embed → optimise** spine is complete, and the traceability + foresight layer (evidence graph, confidence index, real-time early-warning, tamper-evident hash-chained audit trail) now sits ahead of the stated scope. The frontier has shifted to regulatory credibility (standards mapping, model-risk, regulator packs, benchmarks) — the tier that is materially stronger for CIO, risk, assurance and regulated-sector buyers¹ — which is now implemented in the current build.

One sentence for the board. Ekselens is moving from selling consulting hours augmented by tools to operating a governed, independently-verifiable transformation platform — a System of Proof — with consultants as the human-in-the-loop layer. All claims below map to a named module and are subject to reviewer verification (see the Verification appendix).

Commercial note. This repositioning moves the conversation from day-rate consulting toward platform terms (per-tenant / per-engagement / per-artefact). Platform pricing is set out in a separate one-pager, available on request.

¹ The "non-negotiable purchase for regulated buyers" framing is the senior-partner reviewer's phrasing in the 04/07/2026 Gap Analysis, not a claim of guaranteed procurement outcome.

00A · AT A GLANCE

Reconciliation at a glance

Every pillar the Gap Analysis named, resolved to a status. 24 pillars in total.



MET + EXCEEDS 14 delivered to or beyond scope	GAPS ADDRESSED 7 doc-flagged red/amber, now implemented	DEFERRED 3 externally-blocked or trimmed by design
--	--	---

Read this bar as the whole story: the transformation spine is delivered, the platform runs ahead of scope on the evidence graph and the closed assurance loop, and every gap the Gap Analysis itself flagged now carries an implemented control or a documented deferral. The only remaining items are externally-blocked (ERIO endpoints, Semrush) or intentionally trimmed per the Model-Risk strategy.

01 • ORIENTATION

How to read this document

This is a reconciliation, not a sales brochure. Each claim maps to a named module in the live codebase so a reviewer can verify it rather than take it on trust.

- ▶ **The source of truth** is the Gap Analysis dated 04/07/2026: two Big-4 senior-partner reviews (product ambition + audit-grade controls) and a Model-Risk-Management (MRM) defensibility strategy.
- ▶ **Part A** (sections 00–11) is for an executive / non-technical reviewer: the scorecard, what progressed, what exceeds, what closed, what is deferred, and the three themes to emphasise — defensibility, single source of truth, regulatory readiness — plus security posture.
- ▶ **Part B** (architecture, API surface, verification appendix) is for a technical reviewer who wants to see the layering, the endpoints and how to re-run the integrity checks themselves.
- ▶ **Status key.** **MET** delivered to scope · **EXCEEDS** beyond scope · **CLOSED** a doc-flagged gap, now implemented · **DEFERRED** intentionally parked / demand-led. A separate Maturity column states the delivery-state (Live production · Live · admin · Live · public · Planned · Demand-led).
- ▶ **The architecture.** The System of Proof is the parent discipline. It informs EKSELENS360 (function methodology) and Enterprise360 (value-stream methodology), and separately informs ERIO (a distinct Ekselens product). This document covers the EKSELENS360 & Enterprise360 platform.

02 · THE RECONCILIATION

Reconciliation scorecard — every pillar

Every pillar the Gap Analysis named, its status today, its delivery-maturity, and the module a reviewer can check it against.

PILLAR (FROM THE GAP ANALYSIS)	STATUS	MATURITY	WHERE IT LIVES / EVIDENCE
Reinvention Blueprint Engine (diagnose→design→plan; archetype selector, function & value-stream blueprints, 15-artefact reinvention kit)	MET	Live	enterprise360_engine.py · _archetypes.py · value_streams/* (12 streams) · enterprise360_kit.py
Single source of truth — one canonical engagement object; no document self-calculates its own numbers	MET	Live	report_facts.py · enterprise360_report_data.py (one ledger → PDF/DOCX/PPTX/XLSX/preview)
Pre-export QA gates — data reconciliation, sector calibration, rendered-layout QA, transformation-accelerator completeness	MET	Live	enterprise360_qa_gate.py · _layout_qa.py · assessment_qa_gate.py (hard-fail blocks export)
Delivery & Assurance Hub — mobilisation, delivery hub, RAID, decision log, dependency map, benefits-vs-actuals, governance cadence, stage gates	MET	Live · admin+portal	enterprise360 RAID hub.py · _decision_hub.py · _gate_hub.py · _benefit_hub.py · _status_pack.py
Reassessment / learning loop (Optimise → Diagnose)	MET	Live	enterprise360_reassessment.py · enterprise360_learning.py (governed, versioned publish)
Evidence confidence on every conclusion	MET	Live	enterprise360_evidence_maturity.py (EGMM) · enterprise360_confidence.py
AI as draft-and-challenge, never a decision layer	MET	Live	enterprise360_steering_narrative.py · _narrative_qa.py (grounded, deterministic fallback)
Multi-tenant / switchable — per-tenant feature entitlements & white-label	MET	Live	instance_config.py · instance_routes.py (entitlements switch features per tenant)
The evidence graph — the Gap Analysis's named 'moat', with £ value propagation and impact ('descope this — what breaks?') queries	EXCEEDS	Live	enterprise360_evidence_graph.py — typed causal graph, per-benefit & bridge, read-only in portal
Tamper-evident, hash-CHAINED audit trail — independently verifiable	EXCEEDS	Live	

PILLAR (FROM THE GAP ANALYSIS)	STATUS	MATURITY	WHERE IT LIVES / EVIDENCE
(recompute the chain; one altered record breaks it visibly)			enterprise360_evidence_trail.py + verifier export · test_enterprise360_trail_chain.py
Programme Confidence Index — one deterministic 0–100 assurance score, trended, with a transparent deduction breakdown (net-new)	EXCEEDS	Live	enterprise360_confidence.py
Real-time early-warning engine + cross-portfolio triage inbox (foresight, beyond the Gap Analysis's 'drift flagging')	EXCEEDS	Live	enterprise360_alerts.py · /alerts/inbox (ranked by severity + £ at stake)
Delivery critical-path with decision-gating — names the bottleneck gate and the open decisions blocking it, with cascade slip risk	EXCEEDS	Live	enterprise360_critical_path.py
Evidence Graph Maturity Model (EGMM) — public lead-gen diagnostic + internal auto-evidenced hub (net-new productised front door)	EXCEEDS	Live · public	enterprise360_evidence_maturity.py · public_maturity_routes.py · SystemOfProof.jsx
Standards mapping (ISO/IEC 20000-1 & 27001, ITIL 4, COBIT 2019, TOGAF) with per-standard coverage % + clause status	CLOSED	Live · indicative content	enterprise360_standards.py — was RED in the Gap Analysis's scorecard; now implemented
Regulator-ready packs (DORA, NIS2, NHS DCB0129/0160) assembled from live evidence, scored met/partial/gap with a readiness %	CLOSED	Live · indicative content	enterprise360_regulator_packs.py (JSON + PDF) — was RED; now implemented
Benefits credibility ladder — indicative → validated → committed; only committed counts as committed £ in the value bridge	CLOSED	Live	enterprise360_benefit_hub.py — was AMBER; now implemented
Four-eyes sign-off + independent 'red-team' challenge + validated assumption register (dated, owned, auto-flagged when overdue)	CLOSED	Live	enterprise360_signoff.py · enterprise360_assumption_hub.py — was AMBER; now implemented
Model Risk Management (SR 11-7 / SS1/23 flavour) — model cards, 5-tier / 7-gate register, per-call prompt/model/version logging, pause triggers, EU AI Act flag	CLOSED	Live · admin	enterprise360_mrm.py · _mrm_register.py · _model_risk.py · _prompt_registry.py · _model_monitor.py
Benchmark dataset — opt-in, anonymised cross-engagement; k-anonymity threshold so no single engagement is identifiable; non-reversible client hash	CLOSED	Live · seed cohort	enterprise360_benchmark.py · seed_benchmark.py — was RED; now implemented

PILLAR (FROM THE GAP ANALYSIS)	STATUS	MATURITY	WHERE IT LIVES / EVIDENCE
Digital twin — queryable Target Operating Model + deterministic what-if scenarios grounded in the live evidence graph	CLOSED	Live	enterprise360_digital_twin.py — was RED (P3); now implemented
Outbound interoperability with ERIO (a separate Ekselens product) — Ekselens side complete; awaiting the external ERIO endpoints	DEFERRED	Planned · external	erio_integration.py · sso_status.py — degrades gracefully; lights up when ERIO ships
Live production-telemetry connectors for continuous-assurance rescoring	DEFERRED	Demand-led	P3 — evidence-driven rescoring is live; system-telemetry ingestion is future
Full quarterly bias engine (protected-characteristic disparate-impact) + Git 'prompts-as-code' ops discipline + EU AI Act full high-risk conformity	DEFERRED	Deferred	Trimmed by design per the MRM strategy — runtime stamping + drift/amendment metrics kept

03 · PROGRESSED TO SCOPE

What progressed to scope

The full transformation lifecycle the Gap Analysis described — implemented as a governed product platform, not a "loose AI document generator".

- ▶ **Client profile & blueprint selector.** A guided diagnostic classifies the organisation across the stated dimensions and resolves a business archetype before any recommendation is generated — the "fit-for-purpose world-class, not generic best practice" principle.
- ▶ **Function & value-stream blueprints.** Twelve authored value streams — Source-to-Pay, Order-to-Cash, Lead-to-Order, Plan-to-Produce, Idea-to-Market, Record-to-Report, Hire-to-Retire, Risk-to-Assurance, Strategy-to-Execution, Captive Finance, plus enabling IT and enabling Data — each authored to the 17-section template.
- ▶ **The 15-artefact reinvention kit.** Every blueprint carries current-state, evidence confidence, what-good-looks-like, target operating model, L1–L4 process model (editable BPMN), RACI, controls, data/technology models, AI opportunities, KPI/KRI/XLA, RAID, 30/60/90/180/360 roadmap, benefits model and a fast-start pack — "the difference between a report and a reinvention kit".
- ▶ **Delivery & Assurance Hub.** Live RAID, decision log, dependency-aware stage gates, benefits-vs-actuals, governance cadence and a client-facing portal — the client is handed delivery infrastructure, not left to build a programme from scratch.
- ▶ **Reassessment / learning loop.** The Optimise → Diagnose loop rescopes the next diagnostic from what moved, stalled or carried over — a living transformation system, not a one-off pack.

04 • BEYOND SCOPE

Where the platform now exceeds scope

The Gap Analysis said to "fund the evidence graph properly — it is the moat, not the AI". It is implemented in production, with further capability built on top of it.

- ▶ **Evidence graph with £ propagation.** A typed causal graph (supports / contradicts / depends-on / delivers / mitigates / realises) with per-benefit £ value-at-stake propagated along the edges and a planned/at-risk/realised value bridge. The impact query — "descope this — what breaks?" — is live and monetised, and exposed read-only in the client portal.
- ▶ **Tamper-evident, hash-chained audit trail.** Each entry is chained to the previous entry's hash so a third party can recompute the chain and a single altered record breaks it visibly. This is the control the MRM strategy identified as the crown jewel. (We describe it as tamper-evident, not immutable: immutability beyond the application store would require WORM/ledger anchoring — noted as a future hardening step.)
- ▶ **Programme Confidence Index.** Net-new — one deterministic 0–100 assurance score rolled up from all four registers, trended daily, with a transparent deduction breakdown.
- ▶ **Real-time early-warning + cross-portfolio triage inbox.** Beyond "drift flagging": proactive deterministic alerts with a new→acknowledged→resolved lifecycle, consultant email digests, and a portfolio inbox ranking every engagement by severity and £ at stake. That is foresight, not reporting.
- ▶ **Grounded AI steering narrative.** Auto-writes the board story strictly from the data model with a deterministic fallback — the "AI drafts, never decides" principle made concrete.
- ▶ **Public EGMM diagnostic + gated framework library.** A productised front door (Evidence Graph Maturity Model) and a lead-gated Manifesto/EGMM/MRM library — net-new go-to-market surface.

05 · GAPS CLOSED

Gaps the Gap Analysis flagged — now closed

The Gap Analysis contained its own post-build scorecard listing the open reds and ambers. Each now has an implemented control, verifiable against a named module. This is the most important line in this reconciliation.

- ▶ **Standards mapping** (was RED). ISO/IEC 20000-1 & 27001, ITIL 4, COBIT 2019 and TOGAF clause mapping with per-standard coverage % and clause status — the reviewer rated this "worth six figures to a regulated client". Content marked indicative is pending Ekselens sign-off.
- ▶ **Regulator-ready packs** (was RED, P3). DORA, NIS2 and NHS DCB0129/0160 submission packs assembled from live evidence, each requirement scored met/partial/gap with evidence refs and a readiness %.
- ▶ **Benchmark dataset** (was RED, "the single most fundable asset"). Opt-in, anonymised cross-engagement benchmarking with a k-anonymity threshold so no single engagement is identifiable, and a non-reversible client hash.
- ▶ **Digital twin** (was RED, P3). A queryable Target Operating Model plus deterministic what-if scenarios (gate-slip / decision-delay / defer-stream) grounded in the live evidence graph.
- ▶ **Benefits credibility ladder** (was AMBER). indicative → validated → committed; only committed counts as committed £ in the value bridge and narrative — the CFO-trust control.
- ▶ **Four-eyes sign-off + red-team challenge + assumption register** (was AMBER). No publishable artefact is released without a different reviewer to Approve or Challenge (reject with counter-evidence); assumptions are dated, owned and auto-flagged when overdue.
- ▶ **Model Risk Management** (was AMBER/partial). Per-feature model cards, a 5-tier / 7-gate model register, per-call prompt/model/version logging, pause triggers and an EU AI Act jurisdiction flag — the SR 11-7 / SS1/23 discipline the reviewer said would become table stakes.

06 • DEFERRED

What is deferred (and why)

A short, honest list. Nothing here blocks the platform being demonstrated or sold today; each item is either externally-blocked or intentionally trimmed per the Gap Analysis's own advice.

- ▶ **ERIO interoperability.** ERIO is a separate Ekselens product. Unified search, SSO-signal cookies and the outbound lead-forwarder are complete on the Ekselens side; they light up automatically when the external ERIO endpoints ship. The platform degrades gracefully until then and does not depend on ERIO.
- ▶ **Live production-telemetry connectors.** Continuous assurance already rescopes from evidence; ingesting live system telemetry for automatic monthly rescoring is a demand-led P3 (the optional recurring-revenue tail).
- ▶ **Heavy MRM items intentionally trimmed.** Per the MRM strategy, the full quarterly protected-characteristic bias engine, Git "prompts-as-code" ops discipline and full EU AI Act high-risk conformity were deliberately deferred; runtime prompt/model stamping plus sector/size/geography drift and amendment-rate metrics are kept.
- ▶ **Semrush SEO audit.** Blocked on external API availability — unrelated to platform capability.

07 · DEFENSIBILITY

Defensibility — "let them check, it holds"

The MRM strategy set the design principle: build for falsifiability — controls a hostile third party can independently verify without trusting us. Four of the six named attacks are killed by primitives already running in production.

THE DISCREDITING ATTACK	THE CONTROL THAT MAKES IT FAIL
"Your audit trail is editable."	Hash-CHAINED append-only log; recompute the chain and any altered record breaks it visibly. (enterprise360_evidence_trail.py)
"You claim a human reviewed it, but can't prove it."	Signature bound to the output hash — change one word and the signature is void. (enterprise360_signoff.py)
"The AI could've been swapped or the prompt changed silently."	Per-call stamp of prompt version + model provider/version + input/output hashes on the model card. (enterprise360_prompt_registry.py · _model_risk.py)
"You mark your own homework."	Optional independent second line; the client's own reviewer can counter-sign as the independent line. (enterprise360_signoff.py · portal)
"Your numbers are fabricated."	Every figure traces to an evidence-graph node with a provenance rating; unnamed benefit numbers rejected; value-at-stake deterministic. (evidence_graph · benefit_hub · confidence)
"You hide the bad findings."	Append-only + amendment diffs; nothing hard-deleted, "hidden" items soft-hidden with a logged reason. (registers)

The verifier's-eye export. The regulator/audit packs carry a verification appendix that lets a third party re-run the integrity check themselves — recompute the SHA-256 hash chain, confirm each signature matches its output, confirm every Tier-1 artefact carries the required signatures (see Part B — Verification appendix). Scrutiny becomes a sales asset, not a threat.

Assurance Mode. A per-tenant policy toggle — **Standard** (single HITL signature, chain always on, second line optional) or **Regulated** (Tier-1 export hard-blocked without an independent second signature, retention lock, hard prohibited-use gates). Same engine, one switch.

08 · SINGLE SOURCE OF TRUTH

Single source of truth — how drift is made impossible

The Gap Analysis's first control: "every output must use one canonical engagement data object; no document should calculate its own numbers independently." This is enforced structurally, not by convention.

- ▶ **One canonical ledger.** Scores, owners, phases, dates, targets, value pool and headline number are computed once in `report_facts.py` / `enterprise360_report_data.py` and consumed identically by the PDF, DOCX, PPTX, XLSX, board deck and the in-app preview — they cannot disagree because they read the same object.
- ▶ **Data-reconciliation gate before export.** Scores, dates, client name, sector, roadmap items, costs, benefits, owners and target maturity are cross-checked; a £-figure number-drift guard runs in the automated QA suite. A mismatch blocks the export.
- ▶ **Byte-exact freeze.** On finalisation the canonical artefacts are rendered from a frozen snapshot, SHA-256 hashed and stored byte-for-byte — re-verifiable on demand months later.
- ▶ **Guarded by tests.** A dedicated cross-artefact consistency suite regenerates the whole pack and asserts headline/sponsor/phase/date agreement and the absence of stale titles on every run.

09 • REGULATORY READINESS

Regulatory readiness

This is the tier the Gap Analysis calls materially stronger for regulated buyers. It is implemented — and carefully scoped.

- ▶ **Standards-mapping engine.** Assessed capability areas map to ISO/IEC 20000-1, ISO/IEC 27001, ITIL 4, COBIT 2019 and TOGAF with per-standard coverage % and clause-level status (strong / developing / weak / gap). Indicative mappings are flagged pending Ekselens sign-off — honest provenance.
- ▶ **Regulator-ready packs.** DORA, NIS2 and NHS DCB0129/0160 packs assembled from the live evidence (RAID / decisions / gates / dependencies / assumptions + Confidence Index + sign-off state), each requirement scored met/partial/gap with evidence refs and a readiness % — delivered as JSON and board-credible PDF.
- ▶ **Model Risk Management to SR 11-7 flavour.** Per-feature model cards, tiering, human-in-the-loop matrix ("AI drafts, never decides; no client data trains models"), per-call versioned prompt/model logging, pause triggers and an EU AI Act jurisdiction flag.

Scope of these artefacts. Regulatory packs are evidence-organisation and readiness-support artefacts. They do not constitute legal, regulatory, audit or certification opinions, and do not by themselves create or confer compliance. They organise a client's own evidence so their risk/compliance function can reach and defend its own conclusions.

10 • SECURITY

Security, data handling & tenant assurance

Summarised for a CIO / CISO first pass. Items marked planned are stated honestly with intent rather than omitted.

CONTROL AREA	CURRENT POSTURE
Data residency & hosting	Migration to a UK Azure instance planned (MongoDB Atlas, UK region). Auth provider (Auth0) already on the UK region. Current preview/pilot hosting is US-region; production target is UK.
Model / data boundary	Client data does not train any model. LLM use is bounded to drafting/ challenge via a keyed provider abstraction; provider and model version are stamped on every call.
Tenant separation	Per-tenant instance configuration + feature entitlements; engagement data scoped per client. Single-tenant dedicated-instance option for regulated clients.
Access control (RBAC)	Six-tier platform roles (owner/founder → sales → trial/campaign → client-portal); admin surface gated by allowlist; client portal is read-scoped per engagement.
Encryption	TLS in transit; managed encryption at rest on the hosting/DB tier. Enterprise BYOK is a planned option.
Audit log & retention	Hash-chained tamper-evident evidence trail; admin audit log; retention lock available in Regulated Assurance Mode (7y).
Backup / recovery	Managed database backups on the hosting tier; documented RPO/RTO targets to be finalised with the UK Azure migration.
Penetration test / SOC 2	Planned — independent penetration test and SOC 2 readiness assessment scheduled; target date to be confirmed. Not yet performed at pilot stage.
Data export / deletion	Engagement data is exportable; deletion / erasure honoured per the published data-retention schedule and sub-processor register.
Environment separation	Distinct preview and production environments; secrets held in environment configuration, never in source.

A fuller security & data-protection pack (sub-processor register, DPA, residency statement) is available on request.

11 · PART B — ARCHITECTURE

Architecture for a technical reviewer

React + FastAPI + MongoDB, delivered as the eight-layer governed platform the Gap Analysis specified. Deterministic core; AI is a bounded drafting/challenge layer that cannot move a score.

LAYER	WHAT IT DOES	WHERE IT LIVES
1 · Methodology	Lifecycle, stage gates, artefact rules, assurance model	enterprise360_method.py · _spine.py
2 · Data (canonical object)	Client profile, scores, evidence, findings, RAID, benefits, roadmap	report_facts.py · enterprise360_report_data.py · MongoDB (enterprise360_clients, mrm_models)
3 · Knowledge / IP	Sector models, archetypes, function & value-stream libraries, KPI/RAID patterns	enterprise360_archetypes.py · overlays.py · value_streams/* · frameworks/*
4 · Rules engine	Profile + diagnosis → recommendations (deterministic assembly)	enterprise360_engine.py · disposition.py (TIME)
5 · AI generation (bounded)	Drafts narrative + challenge only; barred from scores/prices/dates	steering_narrative.py · narrative_qa.py — third-party LLM via keyed provider abstraction; provider & model version stamped on every call
6 · QA & assurance	Data / sector / commercial / layout / accelerator gates	enterprise360_qa_gate.py · layout_qa.py · assessment_qa_gate.py
7 · Output	PDF/DOCX/PPTX/XLSX, kit ZIP, dashboards, client portal	enterprise360_kit.py · _deck.py · _benefits_xlsx.py · pdf_renderer.py
8 · Delivery workflow	Live RAID/decisions/ gates/benefits, reassessment, evidence trail	raid_hub · decision_hub · gate_hub · benefit_hub · reassessment · evidence_trail

Quality gate. The backend is protected by an automated test suite (1,181 tests collected via `pytest --collect-only`; the enforced QA gate runs the release subset on every change) — covering catalogue integrity, 17-section template authoring, BPMN well-formedness, cross-artefact consistency, £-drift guards, hash-chain verification and rendered-layout QA. A test manifest can be attached on request.

12 · PART B — API SURFACE

Key API surface

Illustrative endpoints behind the claims above. All admin routes are prefixed `/api/admin/enterprise360`; the platform is API-first so modules speak through documented contracts.

CAPABILITY	ENDPOINT (ILLUSTRATIVE)
Public EGMM diagnostic (lead-gen)	GET <code>/api/public/system-of-proof</code> · POST <code>/api/public/framework-library/{doc_type}</code>
Blueprint assembly	POST <code>.../resolve</code> · GET <code>.../clients/{id}/blueprint</code> · GET <code>.../clients/{id}/kit</code>
MRM model register	GET <code>.../clients/{id}/mrm/register</code>
Evidence graph + impact query	GET/POST <code>.../clients/{id}/evidence-graph[query]</code>
Confidence index	GET <code>.../clients/{id}/confidence</code>
Standards map / regulator pack	GET <code>.../clients/{id}/standards-map</code> · <code>.../regulator-pack[.pdf]</code>
Tamper-evident trail + verifier	GET <code>.../clients/{id}/evidence-trail[.pdf]</code>
Four-eyes sign-off / assumptions	POST <code>.../signoff</code> · <code>.../assumptions</code>
Early-warning portfolio inbox	GET <code>.../alerts/inbox</code>
Digital twin / benchmark	GET <code>.../clients/{id}/digital-twin[scenario]</code> · <code>.../benchmark</code>
Reassessment loop	POST <code>.../clients/{id}/reassessment</code>
Client portal (read-only, tenant-scoped)	GET <code>/api/portal/me</code> · <code>/portal/engagement/{id}/*</code>
Liveness / build integrity	GET <code>/api/health</code> (build, declared_version, version_drift)

13 · PART B — VERIFICATION

Verification appendix — how a reviewer checks each claim

Every status in the scorecard is intended to be inspected, not trusted. This is the artefact a reviewer asks for.

CLAIM	WHAT TO INSPECT
Single source of truth	Call the blueprint endpoint; confirm the canonical object; regenerate the pack and diff the headline number across PDF/DOCX/PPTX/XLSX.
QA gate blocks export	Trigger an export with an injected inconsistency; confirm a 422 QA-defect report rather than a ZIP.
Evidence graph	Inspect a sample node/edge/value-bridge; run the "descope" impact query and read the £ endangered.
Hash chain	Export the evidence trail; recompute the SHA-256 chain from the JSON (each entry seals the previous entry's hash); alter one record and confirm the chain breaks.
Four-eyes sign-off	View a signed artefact; edit one character; confirm the signature is invalidated (bound to the output hash).
Standards map / regulator pack	Open an example clause mapping and a readiness view; confirm indicative items are flagged pending sign-off.
Layout QA	Run the layout-QA endpoint; review a detected-and-blocked overflow example.
Reassessment loop	Compare a before/after maturity movement and the carried-forward overhang.

Verifier basics. Chain algorithm: SHA-256 over a canonical JSON serialisation of each entry plus the prior entry's digest. Signatures bind to the artefact output hash. The trail is exported as JSON (machine-checkable) and PDF (human-readable) from `GET ../evidence-trail[.pdf]`. A guard test — `test_enterprise360_trail_chain.py` — proves a tampered record is detected.

14 · DEMO

Executive demo flow (20 minutes)

| A repeatable sales motion that turns this document into a live demonstration.

- 1 **Create the engagement** and select the business archetype (sector, scale, operating model).
- 2 **Run diagnosis** and show evidence confidence on each conclusion.
- 3 **Generate the target-state blueprint** and open the evidence graph.
- 4 **Run "what breaks if we descope this?"** and read the £ at stake on the path.
- 5 **Show the roadmap, RAID and the benefits credibility ladder** (indicative → validated → committed).
- 6 **Run the QA gate** — show a blocked export, then a clean export.
- 7 **Export the pack + verifier appendix** and recompute the hash chain live.

15 · SIGN-OFF

Assurance statement & version control

This overview was reconciled directly against the live codebase; every status in the scorecard maps to a named, testable module and is subject to reviewer verification.

On the Gap Analysis's own bottom-line test — reduce the uncertainty, friction and risk between knowing change is needed and successfully delivering it — the spine is complete, the moat (evidence graph + closed assurance loop) is ahead of scope, and every red/amber item the document itself flagged now carries an implemented control or a documented deferral. The platform is a System of Proof: built to be checked.

Navin Sahadeo — Founder, Ekselens Consulting

VERSION	DATE	CHANGE
v1.0	05 Jul 2026	Initial reviewer reconciliation & platform overview.
v1.1	06 July 2026	Two-reviewer feedback pass: brand lockup (real star device); language discipline (tamper-evident vs immutable; softened/attributed gap & CIO claims); 12 value streams listed; maturity-status column; security & data-handling page; verification appendix; regulatory-scope caveat; ERIO positioned as separate product; provider-abstraction wording.
v1.2	06 July 2026	Corrected the brand hierarchy on the cover: the System of Proof is the parent discipline , with EKSELENS360 (function methodology) and Enterprise360 (value-stream methodology) as its two delivery arms — not three horizontal peers. Added the one-line hierarchy strip (Ekselens Consulting › System of Proof › EKSELENS360 & Enterprise360 › ERIO).

AN EKSELENS SYSTEM OF PROOF ASSET · EKSELENS360 & ENTERPRISE360 · V1.2 · 06 JULY 2026

Commercial in confidence. Framework rubrics, archetype recipes, the value-stream libraries and the improvement catalogue are Ekselens IP. Public standards are factual references only — no licensed content is reproduced. Standards/regulator mappings marked indicative are pending Ekselens sign-off.