JANUARY 24, 2026

Agenda for Social Equity 2074

Validation Partner Development Guide

CREATED BY
EUSL AB
Care to Change the World

Table of Contents

Validation Partner Development Guide

Introduction

This Validation Partner Development Guide establishes the procedural, methodological, and governance requirements that any organisation must satisfy to operate as a Validation Partner within the A2074-SRS ecosystem. Validation Partners are entrusted with designing and operating models—stars, points, badges, deep dives, sector modules, and maturity ladders—that must remain anchored to the Rules for Interpretation of the 17 Social Global Goals, uphold patient-level confidentiality, respect consent and revocation rights, and maintain strict proportionality across enterprise sizes and sectors.

This document must be read in conjunction with the Foundational Charter, the Licensing & Accreditation Framework, the Rules for Interpretation, the Governance & Oversight Manual, the Operating Manual (Open Standard), the Digital Integration & Platform Governance Manual, the Ethics & Integrity Code, the Communication & Public Disclosure Protocol, and the Legal Compliance & International Law Note. Together, these instruments establish non-derogable requirements: confidentiality by default; prohibition on comparative ranking; independence of GSIA ethics oversight; and an approval framework ensuring that all models are trustworthy, fair, and methodologically sound.

This Guide proceeds through nine chapters. The first two chapters address the design of validation models and the translation of the SGG canon into measurable, verifiable criteria. The subsequent chapters govern proportionality algorithms, branding permissions, approval dossiers, confidentiality patterns, ethics impact assessment (EIA), piloting and calibration, and long-term maintenance and change control. A Final Word closes the document by affirming the development philosophy governing the A2074-SRS ecosystem.

Chapter 1 — Designing a Validation Model

This Chapter describes the complete methodology design process that a prospective Validation Partner must undertake prior to submitting a model for GSIA review under the Licensing & Accreditation Framework. A compliant model must be anchored in the 17 Social Global Goals, operate proportionately across enterprise sizes and sectors, preserve confidentiality, avoid any comparative posture, and incorporate least-intrusive evidence practices. The design process is iterative and must be documented in a model dossier forming part of the approval package under Chapter 5.

The starting point for any model design is doctrinal anchoring. Each component of the model—criteria, indicators, scoring logic, recognition levels, maturity thresholds, badges, or deep-dive modules—must map clearly to one or more SGG pillars, using the Rules for Interpretation as the authoritative source. The mapping must demonstrate interpretive fidelity: the model may add specificity or granularity, but it may not alter meaning, infer obligations beyond the SGG canon, or dilute protections. Any ambiguity must be resolved in favour of the Rules for Interpretation.

Once anchoring is secured, the model must define its methodological architecture. This includes its recognition format (such as stars, points, badges, or levels), its internal weighting logic, its evidentiary expectations, and its progression structure. The architecture must be transparent, documented, version-controlled, and fully auditable. No implicit ranking logic, percentile thresholds, or comparative algorithms may be incorporated. Recognition is descriptive of alignment to the SGG pillars, not comparative across peers.

Evidence design follows the least-intrusive doctrine. The model must specify when declarations suffice, when desk review is required, when sampling is proportionate, and under what narrow conditions field checks may be invoked. Sensitive artefacts must be handled within secure enclaves using “view-only, no-extract” principles. Identity-linked evidence requires explicit approvals from internal ethics controls and must be minimised in scope and duration. AI-enabled evidence analysis must include model cards, guardrail documentation, bias monitoring, and human-in-the-loop safeguards.

Proportionality must be embedded at the design stage. The model must ensure that microenterprises, small entities, and large organisations can attain recognition fairly, with burden calibrated to capacity. Evidence demands, complexity, and documentation volume must scale appropriately, and any use of quantitative indicators must incorporate normalisation or context-adjustment rules to avoid structural bias.

Confidentiality-by-design must be incorporated explicitly. All outcomes must be private by default, disclosure must be voluntary and revocable, and all external communication must comply with the Communication & Public Disclosure Protocol. Public-facing artefacts (for example, star icons) must only appear with ledger-verified, revocable consent.

The completed model design must be captured in a master design document that includes doctrinal mapping, evidence taxonomies, weighting logic, proportionality rules, data-flow diagrams, UX flows demonstrating confidentiality protections, redaction protocols, calibration plans, and change-control assumptions. This document becomes the core of the approval dossier under Chapter 5.

Table 1: Core Design Obligations for All Validation Models

Obligations	Required Design Feature	Approval Link
Doctrinal Anchoring	Mapping to 17 SGG pillars; interpretive fidelity	Agenda 2074 interpretive compliance
Confidentiality-by-Design	Private-by-default outcomes; revocable consent	GSIA privacy and ethics review
Least-Intrusive Evidence	Redaction defaults; enclave protocols; minimisation	GSIA verification pathways
Proportionality	Scale-adjusted burden; fairness rules	Proportionality tests in GSIA review
Non-Comparative Operation	No benchmarking or ranking	Communication protocol compliance
Transparency & Auditability	Documented logic, versioning, data-flows	Approval dossier and monitoring

Chapter 2 — Translating SGGs into Criteria and Indicators

This Chapter provides the structured methodology for converting the substance of the 17 Social Global Goals and their detailed Rules for Interpretation into measurable, verifiable, and proportionate criteria that can be used across diverse validation models. The translation process must preserve the normative integrity of the SGG pillars, avoid reductionism, and ensure that indicators do not distort or obscure the underlying purpose of each pillar.

The translation process begins with a close reading of the relevant SGG pillar, including its interpretive narrative, examples of expected conduct, safeguard requirements, sector variants, and cross-pillar interdependencies. These components form the normative foundation from which criteria and indicators are derived. A Validation Partner must then identify the specific behaviours, safeguards, processes, and outcomes that collectively constitute alignment with the pillar's intent. Each of these elements becomes a candidate criterion.

Criteria must be substantive, not procedural: they measure whether an organisation fulfils the intent of an SGG pillar, not whether it produces documentation for its own sake. They should focus on governance, behaviours, controls, impacts, and protections directly linked to the pillar's normative core. Criteria may be qualitative, quantitative, or mixed, but must always remain verifiable without requiring intrusive or identity-linked data unless absolutely necessary.

Indicators convert criteria into measurable units. Indicators must be constructed to permit verification through declarations, desk review, sampling, or field checks, depending on proportionality and risk. Indicators must be framed to avoid implying competitive comparison. Indicators must have clear definitions, boundary conditions, permissible evidence types, and verification modes. The use of numeric thresholds is permitted only where context and proportionality are preserved; when numeric thresholds risk bias, models should rely on qualitative judgements grounded in evidence.

To prevent fragmentation, the translation must follow a crosswalk process that shows the relationship between the SGG pillar, its interpretive rules, each criterion, and each associated indicator. The crosswalk must also document where sector-specific adaptations have been introduced, demonstrating interpretive fidelity and ensuring that sector modules do not contradict or dilute the pillar.

Criteria and indicators must support proportionality. For micro- and small-enterprises, indicators may emphasise the presence of foundational practices and safeguards, rather than complex systems or formal structures that are impractical at small scale. For large entities, criteria may require more comprehensive controls, documentation, and monitoring systems, provided that the burden remains reasonable and does not impose disproportionate demands on evidence.

Confidentiality must govern indicator design. Indicators must not require identity-linked data as a default. Where such data is necessary for a given pillar—such as individual-level safety protections or consent governance—the indicator must explicitly incorporate minimisation and secure-enclave review rules.

Table 2: Translation Requirements from SGG Pillars to Criteria and Indicators

Translation Step	**Required Outcome**	**Oversight Reference**
Interpretive Reading	Full extraction of normative content	Rules for Interpretation
Criteria Formation	Substantive, verifiable, proportionate criteria	Agenda 2074 doctrinal checks

Indicator Definition	Clear, measurable, non-comparative indicators	GSIA evidence and proportionality review
Evidence Mapping	Declarations, desk review, sampling, field verification	Operating Manual alignment
Sector Adaptation	Contextual but non-dilutive adjustments	Agenda 2074 sector consistency check
Confidentiality Safeguards	Minimisation; enclave review; no coercive disclosure	Digital Governance Manual; GSIA privacy review

The translation process must be version-controlled and must produce a traceable lineage from SGG text to model criteria. Any subsequent model updates must retain this lineage, ensuring that all refinements remain consistent with the SGG canon and the broader governance instruments of A2074-SRS.

Chapter 3 — Proportionality Algorithms and Qualitative Adjustments

This Chapter establishes the methodological requirements for constructing proportionality algorithms and qualitative adjustment mechanisms that enable Validation Partners to generate fair, non-comparative, and scale-appropriate validation outcomes. Proportionality is a cardinal doctrine of the A2074-SRS ecosystem. It ensures that microenterprises, small organisations, medium-sized entities, and large institutions can each demonstrate meaningful alignment with the 17 SGG pillars without being penalised for scale, resource intensity, structural complexity, or sector characteristics. It also ensures that indicators and scoring logic cannot inadvertently produce competitive rankings or comparability, which are expressly prohibited under the Standard.

Proportionality algorithms must be explicitly documented, readily auditable, and free of any weighting logic that correlates capacity with merit. They must reflect the principle that "everyone can do something," and that alignment should be evaluated relative to what is reasonable, feasible, and proportionate for an entity's size, risk profile, operational complexity, and stage of organisational development. These algorithms therefore rely on multidimensional scaling factors rather than linear quantitative adjustments.

A proportionality algorithm must include at least three foundational elements: a size-scaling dimension, a sector-materiality dimension, and a maturity-progression dimension. The size-scaling dimension must differentiate evidentiary expectations without diluting the integrity of confidentiality and consent safeguards. Microenterprises and small entities must be permitted to meet criteria through foundational practices and simplified documentation, whereas larger entities must demonstrate greater procedural robustness, internal controls, and monitoring systems. These differences must not imply superiority or inferiority; they must reflect proportional feasibility alone.

The sector-materiality dimension must integrate the risk and impact realities of the organisation's sector. High-impact sectors—such as extractives, healthcare, or agriculture—may require deeper evidence of safeguards, whereas lower-impact sectors may meet equivalent alignment through simplified controls. Sector materiality must never weaken the underlying obligations associated with the SGG pillars; it may only modify the nature and depth of evidence required to demonstrate compliance.

The maturity-progression dimension must recognise improvement over time. Organisations at early developmental stages must be able to demonstrate meaningful alignment even if advanced systems are not yet in place. Maturity may be acknowledged through staged criteria, graduated levels, or narrative milestones, provided these mechanisms are non-comparative and internally coherent.

Qualitative adjustments are used where algorithmic scaling is insufficient to capture contextual nuances. These adjustments must follow a documented judgement protocol that includes criteria for invoking qualitative review, the evidentiary basis for adjustment, internal ethics review before application, and auditable records. Adjustments must always be justified by demonstrable contextual factors—such as geographic constraints, supply-chain structure, or legal environment—and must not open avenues for subjective bias or arbitrary decision-making. All adjustments must comply with confidentiality-by-design, using redacted or pseudonymised evidence wherever possible.

AI-assisted proportionality tools are permissible only under strict safeguard conditions. Model cards, intended-use statements, bias testing, human review, and fallback procedures must be documented before deployment. No AI model may generate binding determinations or qualitative adjustments without human oversight, and no model may incorporate training data that could compromise confidentiality or enable entity-level comparison.

The final proportionality algorithm must be assembled into a structured methodology note and incorporated into the model dossier for approval under Chapter 5. It must include explicit references to scaling logic, adjustment protocols, sector-materiality matrices, and ethical guardrails, and must be supported by calibration plans demonstrating the algorithm’s fairness, reliability, and resistance to comparative inference.

Table 3: Elements of a Compliant Proportionality Algorithm

Element	**Purpose**	**Required Safeguard**	**GSIA Oversight Link**
Size-Scaling Dimension	Ensures equitable burden across entity sizes	Calibrated evidence tiers; burden caps	Thematic audits on fairness
Sector-Materiality Dimension	Reflects inherent risk profiles	Sector-specific adaptation notes	Agenda 2074 interpretive review
Maturity-Progression Dimension	Recognises improvement over time	Time-bound milestones; renewal checks	Monitoring cycle verification
Qualitative Adjustment Protocol	Addresses contextual nuances	Documented judgement rules	GSIA ethics review
AI-Assisted Scaling (Optional)	Enhances consistency	Guardrails; human oversight	Digital governance audit
Confidentiality Constraints	Prevents data over-collection	Minimisation; enclave access	GSIA privacy checks

All proportionality constructs must ultimately be aligned to the Rules for Interpretation and operate without enabling ranking, benchmarking, or competitive positioning. Where ambiguity arises, proportionality must default to the most protective interpretation of confidentiality and fairness.

Chapter 4 — Branding, Co-Branding, and Canonical References

This Chapter governs the permissible use of naming, branding, iconography, and co-branding elements associated with A2074-SRS, Agenda 2074, and Validation Partner identity. Its purpose is to ensure that public communications maintain doctrinal integrity, prevent consumer or stakeholder confusion, uphold confidentiality, and avoid any impression that validation outcomes constitute certification, ranking, or endorsement beyond what the Standard permits.

Validation Partners may reference the Agenda 2074 Social Responsibility Standard only as a standards framework under which they operate pursuant to licence and GSIA oversight. They may not suggest ownership of the Standard, modify canonical naming, or imply the existence of "Agenda 2074-approved" or "GSIA-endorsed" products or services beyond the formal licensing and accreditation status granted. All public references must be accurate, non-comparative, proportional, and compliant with the Communication & Public Disclosure Protocol.

Branding must follow strict canonical rules. The terms "Agenda 2074 Social Responsibility Standard," "A2074-SRS," "17 Social Global Goals," and "GSIA Ethics Oversight" may only be used in their complete, unaltered form. Abbreviations may be used only where they do not create ambiguity or diminish the recognised identity of the Standard. Iconography, including any star system, badge emblem, or maturity insignia, must conform to approved templates, colour codes, minimal size ratios, and usage restrictions established by Agenda 2074. No Validation Partner may create derivative marks that resemble or compete with canonical marks.

Co-branding arrangements—such as placing an Agenda 2074-aligned emblem next to a Validation Partner's logo—may be used only when the emblem represents a privately held and consented validation outcome. Co-branding must never be used to imply partnership, joint venture, sponsorship, or endorsement by Agenda 2074 or GSIA. The entity whose outcome is recognised must have provided explicit, informed, revocable consent for public display, verified through the consent ledger. Upon revocation, all co-branded materials must be removed immediately.

References to the SGG pillars must maintain doctrinal accuracy. Validation materials may describe which pillars are addressed by the model, but must not infer that a model covers all pillars unless that coverage has been demonstrated and approved. When referencing deep-dive models, Validation Partners must use precise names and avoid implying that deep-dive alignment with one pillar equates to alignment with the Standard as a whole.

Public communications must avoid comparative or competitive phrasing. Phrases implying superiority, sector leadership, or ranking—such as "top-tier," "best-in-class," "leading performer," or "highest-rated"—are prohibited. Validation outcomes must be communicated strictly as levels of alignment to SGG-anchored criteria.

All marketing and public-facing artefacts must undergo internal ethics review before dissemination, and must be retained in an auditable repository aligned with the Digital Integration & Platform Governance Manual. GSIA conducts periodic spot-checks to ensure that public materials comply with canonical naming, confidentiality protections, and non-comparative rules.

Table 4: Branding and Co-Branding Rules

Category	Permitted	Prohibited	Oversight Mechanism
Canonical Naming	Exact use of A2074-SRS, SGGs, GSIA	Altered names, abbreviations implying ownership	Agenda 2074 brand governance
Iconography	Approved star/badge templates	Derivative or altered marks	GSIA review; partner audits
Co-Branding	Consent-verified outcome display	Implying endorsement, certification, or partnership	Communication Protocol enforcement
Pillar References	Accurate mapping to covered pillars	Suggesting holistic coverage without approval	Agenda 2074 interpretive checks
Public Claims	Descriptive, non-comparative language	Ranking, superiority, competitive framing	GSIA sanctions for breach

Validation Partners are required to protect the integrity of the Standard's identity as carefully as they protect confidential data. Any misuse of naming, iconography, or public claims is a material breach and may lead to corrective action, sanctions, or suspension under the Licensing & Accreditation Framework.

Chapter 5 — Submission Dossier and Approval Workflow

This Chapter prescribes the mandatory content of the submission dossier and the corresponding approval workflow by which Agenda 2074 licenses a Validation Partner's model, acting on GSIA's binding ethics and privacy determinations. The objective is to provide a complete, auditable record that demonstrates doctrinal anchoring to the 17 SGG pillars, proportionality, confidentiality-by-design, non-comparative operation, and operational readiness for pilot and scale. The dossier is a living set of artefacts; upon approval, its versioned contents become the baseline against which monitoring, thematic audits, and subsequent changes are assessed.

A submission dossier must be complete, internally coherent, and consistent with the Rules for Interpretation, the Operating Manual (Open Standard), the Governance & Oversight Manual, the Digital Integration & Platform Governance Manual, the Ethics & Integrity Code, the Communication & Public Disclosure Protocol, and the Legal Compliance & International Law Note. Any omission or ambiguity on a matter material to confidentiality, proportionality, or interpretive fidelity is grounds for deferral pending remediation. Where AI components are proposed, the dossier must include explicit model governance documentation with guardrails, bias testing, intended-use statements, human-in-the-loop controls, and fallback procedures.

Table 5: Submission dossier — required artefacts

Dossier Section	Minimum Required Contents	Oversight Reference
Model Overview & Purpose	Model family (stars, points, badges, deep dive); scope; recognition logic; intended users	Operating Manual; Communication Protocol

Doctrinal Mapping	Crosswalk from relevant SGG pillars and interpretive rules to criteria and indicators	Rules for Interpretation; Agenda 2074 standards review
Evidence Taxonomy & Methods	Least-intrusive plan (declarations, desk review, sampling, field checks); sampling rationale; enclave use	Operating Manual; GSIA privacy checks
Proportionality Method	Size-scaling, sector-materiality, maturity-progression; qualitative adjustment protocol	Governance & Oversight Manual; GSIA fairness review
AI Governance (if applicable)	Model cards; training data lineage; bias testing; human oversight; guardrails	Digital Governance Manual; GSIA digital audit
Confidentiality-by-Design	Data minimisation; redaction standards; secure enclave architecture; consent ledger schema	Governance & Oversight Manual; Legal Note
Consent & Disclosure	Layered notices; consent scopes; revocation workflow; UX mock-ups; takedown procedures	Communication Protocol; Digital Governance Manual
Branding & Claims Controls	Canonical references; approved iconography; disclaimers; anti-comparison language	Branding rules (this Guide, Ch. 4)
Ethics Controls	Internal Ethics Control Function mandate; COI management; escalation triggers to GSIA	Governance & Oversight Manual; Ethics & Integrity Code
Pilot & Calibration Plan	Pilot design; inclusion/exclusion criteria; calibration exercises; readiness gates	Operating Manual; GSIA thematic review
Monitoring & KRIs	KRIs for consent integrity, redaction error rate, model drift, training coverage	Governance & Oversight Manual (Ch. 6)
Change Control & Versioning	Version identifiers; change log template; backward-compatibility policy	Governance & Oversight Manual; Operating Manual
Legal & Jurisdiction	Cross-border transfer assessment; lawful bases; regulatory carve-outs	Legal Compliance & International Law Note
Third-Party Dependencies	Processors; assurance providers; data-flow diagrams; exit plans	Governance & Oversight Manual; Ch. 11 interoperability

Training & Competence	Assessor curricula; annual refresh plan; role-based authorisations	Operating Manual; Ethics & Integrity Code

The approval workflow is disciplined, time-bound, and transparent. Agenda 2074 is the licensing authority; GSIA's ethics and privacy determinations are binding in respect of approval conditions. Where relevant, advisory input from Pillar or Sector Panels may be solicited to refine doctrinal fidelity without compromising due process or confidentiality.

Table 6: Approval workflow — stages, responsibilities, and indicative timelines

Stage	**Lead & Role**	**Actions & Assessments**	**Indicative Timeline**	**Outcome**
Intake & Completeness Check	Agenda 2074 Secretariat	Verify dossier completeness; register version; COI pre-screen	10 business days	Proceed to review or request for information (RFI)
Doctrinal & Design Review	Agenda 2074 Standards Unit	Assess SGG crosswalk; interpretive fidelity; proportionality design	20 business days	Doctrinal memo; redlines for alignment
Ethics & Privacy Review	GSIA Ethics Chambers (Advisory/Audit)	Evaluate confidentiality-by-design; consent regime; AI guardrails; escalation pathways	20 business days (parallel)	Binding ethics memo; required mitigations
Technical & Digital Review	GSIA Digital Oversight	Assess secure enclave, logging, key management, KRI instrumentation	15 business days (parallel)	Digital adequacy note; remediation list
Branding & Communications Check	Agenda 2074 Communications	Verify canonical naming; iconography; disclaimers; non-comparative framing	10 business days	Branding clearance or corrections
Conditional Letter & CAP (if needed)	GSIA with Agenda 2074	Issue conditions; define CAP milestones and evidence	5 business days post-review	Conditions of approval set
Decision & Licence Issuance	Agenda 2074 Licensing	Approve, conditional approve, defer, or reject; assign Model ID & version	10 business days	Licence letter; obligations schedule

Pilot Authorisation	GSIA/Agenda 2074 (joint)	Approve pilot parameters per submitted plan; set readiness gates	Within 10 business days of licence	Pilot go-ahead or revision request

Decision outcomes are reasoned and documented. Conditional approvals specify the milestones, evidence requirements, and timelines that must be satisfied prior to pilot or scale. Deferrals identify deficiencies without prejudice, enabling resubmission upon remedy. Rejections are reserved for material doctrinal conflict, irreparable confidentiality risk, or persistent non-cooperation.

Table 7: Decision outcomes and conditions

Outcome	Description	Typical Conditions	Re-entry Path
Approval	All criteria satisfied; licence issued with monitoring obligations	Standard monitoring cadence; calibration reporting	Immediate pilot per plan
Conditional Approval	Substantive alignment with manageable gaps	CAP items (e.g., consent UX refinement; enclave hardening; branding edits)	Pilot permitted upon CAP completion
Deferral	Incomplete or ambiguous submission	Additional artefacts; clarified crosswalk; revised proportionality	Resubmission within set window
Rejection	Material conflict with SGG canon or confidentiality doctrine	Not applicable	New application after redesign

Upon licensing, the Model ID, version, effective date, and obligations schedule are entered into the official register. Validation Partners are notified of reporting cadence, KRI submission expectations, and the calibration timetable. Any material change post-approval must follow the change-control and re-approval provisions defined in Chapter 9 of this Guide and the Operating Manual.

Chapter 6 — Confidentiality by Design Patterns

This Chapter codifies the mandatory confidentiality-by-design patterns that Validation Partners must implement at process, data, and user-experience levels. These patterns translate the non-derogable doctrine—private by default; explicit, informed, and revocable consent; least-intrusive evidence; secure enclave handling; immutable audit trails—into operational designs that are testable, monitorable, and enforceable under GSIA oversight.

Partners must institute privacy-preserving defaults across the validation lifecycle. Evidence collection is minimised to what is strictly necessary; identity-linked artefacts are avoided unless indispensable; and any inspection of sensitive materials occurs in controlled secure enclaves under "view-only, no-extract" rules with dual-control access and immutable logs. Consent is layered, granular, time-bounded, audience-specific, and revocable at will; revocation triggers immediate suppression and takedown across all channels under the Partner's control. All flows that touch personal or identity-linked organisational data are logged end-to-end with time-stamps, purpose, actor, and duration to support GSIA monitoring without exposing raw data.

The user experience must make privacy choices clear and free of coercion. Dark-pattern designs are prohibited. Consent prompts must be unbundled from service access, employ neutral language, and present declines or revocations as first-class options. Consent scope must be visible and modifiable at any time. Public recognition artefacts (such as star emblems) must never be displayed by default; they are contingent upon valid, current consent and must be withdrawn immediately upon revocation.

Table 8: Confidentiality pattern catalogue

Pattern	Objective	Minimum Controls	Verification
Data Minimisation Template	Collect only what is necessary	Data inventory; purpose specification; necessity test; lawful basis	GSIA sampling; records inspection
Redaction & Pseudonymisation	Remove identity from evidence	Redaction standards; tooling; quality checks; no re-identification	Thematic audits; error-rate KRIs
Secure Enclave Access	View-only inspection of sensitive artefacts	Time-boxed sessions; dual-control; immutable logs; no extraction	Access log review; enclave attestation
Consent Ledger	Verifiable consent governance	Ledger schema; cryptographic integrity proofs; scope & duration fields	Ledger integrity tests; revocation drills
Revocation Workflow	Immediate suppression and takedown	One-click revoke; suppression within defined SLA; takedown confirmation	Incident simulation; SLA metrics
Role-Based Access Control	Least-privilege operational access	Role catalogues; segregation of duties; periodic recertification	Access recertification reports
Cross-Border Transfer Guard	Equivalent protection across borders	Transfer assessment; SCCs or equivalents; risk memo	Legal Note alignment checks
Incident Response	Rapid containment and notification	Playbooks; containment scripts; GSIA notice triggers	Post-incident review; closure evidence

The consent ledger constitutes the auditable backbone of disclosure control. Its schema must capture at minimum the entity identity or authorised signatory, purpose and scope of consent, audiences, channels, duration, effective and expiry dates, notices presented, and revocation mechanics. Integrity must be provable through cryptographic means or equivalent tamper-evident logging. Any public-facing disclosure must reference a valid ledger entry; absence or expiration of consent prohibits disclosure.

Table 9: Consent and revocation — UX and process requirements

Requirement	UX/Process Implementation	SLA/Standard
Layered Notice	Summary banner with link to full terms; plain-language options	Notice visible at decision point
Unbundled Consent	Separate toggles per use (public emblem, press, web listing)	No pre-ticked boxes; equal prominence
Easy Revocation	Persistent “Revoke” control in dashboard; confirmatory notice	Immediate effect; takedown proof ≤ 72 hours
Disclosure Auditability	Display consent ID and scope on internal record	Traceable to ledger; immutable
Non-Retaliation	No functionality loss for withholding consent	Product parity preserved

Secure enclaves must be engineered and operated to a verifiable standard. Enclaves host any viewing of sensitive evidence; they record access metadata; they prohibit export; they enforce short session lifetimes; and they employ hardware-backed key management with separation of duties. Partners must maintain enclave attestation artefacts (for example, configuration fingerprints) and submit them during GSIA digital reviews.

Table 10: Access, logging, and key-management controls

Control Domain	Minimum Standard	Evidence of Sufficiency
Authentication	Multi-factor; phishing-resistant where feasible	Auth logs; configuration policies
Authorisation	Role-based, least-privilege; dual-control for enclaves	Role catalogues; approval records
Logging	Immutable, time-synchronised, tamper-evident	Log integrity checks; hash chains
Key Management	HSM-backed storage; rotation schedule; split knowledge	KMS logs; rotation attestations
Retention & Deletion	Policy-driven; cryptographic erasure for sensitive stores	Deletion certificates; key-revocation proof

All confidentiality patterns must be accompanied by measurable KRIs, including consent-ledger anomaly rate, revocation SLA compliance, redaction error rate, enclave access exceptions, and unauthorised disclosure incidents (with zero-reporting). Breaches of thresholds trigger internal escalation to the Ethics Control Function and external notification to GSIA under the Governance & Oversight Manual.

Partners must document these patterns in their operating procedures, include them in staff training and assessor handbooks, and evidence their operation during monitoring cycles. Any departure from these patterns requires prior written justification, compensating controls, and, where material, GSIA approval before use in production validations.

Chapter 7 — Ethics Impact Assessment (EIA)

This Chapter establishes the mandatory requirement for conducting a comprehensive Ethics Impact Assessment (EIA) for every validation model, module, or significant methodological update prior to deployment, piloting, or scaling. The EIA functions as the principal pre-deployment safeguard ensuring that each model operates within the ethical boundaries of A2074-SRS, preserves patient-level confidentiality, respects consent, maintains proportionality, avoids structural bias, and does not inadvertently introduce comparative dynamics or discriminatory outcomes. It is a non-negotiable prerequisite for approval under the Licensing & Accreditation Framework and remains subject to GSIA's binding ethical review.

The EIA is a structured evaluation of risk, impact, control sufficiency, and mitigation effectiveness across seven domains: confidentiality and data protection; consent architecture; proportionality and fairness; sector-specific harm risks; AI and automation risk (if applicable); conflict-of-interest vulnerabilities; and public-facing communications. It must demonstrate rigor equivalent to a regulatory-grade assessment and include a narrative explanation of mitigations, residual risks, and monitoring plans. Where any domain presents material, unmitigated risk, the model may not proceed to pilot until corrective actions are applied and validated.

The confidentiality domain examines whether the design imposes unnecessary evidence demands, whether secure enclaves are available for sensitive artefacts, whether redaction protocols are adequate, and whether minimisation has been applied to every indicator. It must demonstrate that the model can be verified without exposure of identity-linked information except where strictly necessary, and that all flows comply with the Digital Integration & Platform Governance Manual.

The consent domain evaluates whether disclosure, revocation, and suppression mechanisms are clear, accessible, unbundled, and fully auditable; whether UX patterns are free from coercive design; and whether consent governance (ledger integrity, audience specificity, time-bound scope) adheres to the confidentiality doctrine. Any ambiguity in revocation effects must be remedied prior to approval.

The proportionality and fairness domain evaluates the size-scaling logic, sector-materiality adjustments, maturity-progression rules, and qualitative adjustment protocols described in Chapter 3. It must demonstrate that the model does not privilege larger entities, does not impose excessive burdens on micro-enterprises, and does not embed quantitative thresholds that generate implicit comparison or competitive inference.

Sector-specific risk evaluation identifies potential harms associated with the entity's operating context, including labour, safety, procurement, supply chain, community, environmental, or political exposure. It must consider risks of re-identification, retaliation, regulatory conflict, or misinterpretation of public-facing outcomes, and must propose mitigations proportionate to sector realities.

If AI or automated tools are involved, the EIA must include a full AI risk assessment, including model description, intended use, training data lineage, bias testing results, fairness evaluations, human-in-the-loop controls, fallback mechanisms, and monitoring plans. High-risk AI components may be denied approval or restricted to supervised pilot phases until safeguards mature.

The conflict-of-interest domain must detail how the model avoids assessor bias, how assessors are assigned, how COI declarations are maintained, and how firewalling from commercial incentives is enforced. Any conflict must be actively mitigated; passive declarations are insufficient.

The communications domain must ensure that public-facing artefacts (such as icons or badges) cannot be misinterpreted as certification, ranking, comparative excellence, or endorsement. It must verify compliance with branding rules, disclaimers, and the Communication & Public Disclosure Protocol.

The EIA concludes with a consolidated mitigation plan, including timelines, evidence of implementation, and KRI monitoring indicators. This plan becomes part of the approval dossier and is binding upon model operators.

Table 11: Ethics Impact Assessment — mandatory components

Assessment Domain	Required Outputs	GSIA Review Lens
Confidentiality & Data Minimisation	Evidence mapping; minimisation test; enclave design	Privacy & secure-handling integrity
Consent & Revocation	UX patterns; ledger schema; revocation drill results	Voluntariness; revocability; anti-coercion
Proportionality & Fairness	Scaling logic; bias analysis; burden assessment	Non-comparative fairness
Sector-Specific Risk	Harm scenarios; mitigation controls	Sectoral safeguards; non-discrimination
AI/Automation (if applicable)	Model cards; bias tests; guardrails	AI ethics; human oversight
Conflict-of-Interest	COI matrix; assignment rules	Independence & impartiality
Communications & Branding	Claims controls; iconography	Non-comparative public interpretation

An EIA must be conducted for every major version of a model. Material changes, including new indicators, revised evidence demands, new AI components, or modified recognition structures, trigger a new EIA cycle prior to release.

Chapter 8 — Pilot, Calibration, and Scale Up

This Chapter defines the requirements for piloting, calibration, and scaling a validation model once it has received conditional or full approval under the Licensing & Accreditation Framework. Piloting is a mandatory stage that tests the real-world functionality, confidentiality integrity, proportionality performance, and interpretive fidelity of the model before it becomes available for general deployment. No model may be marketed, offered, or applied at scale until it passes the pilot and meets readiness criteria verified jointly by Agenda 2074 and GSIA.

The pilot phase must implement the exact methodology, evidence rules, proportionality algorithms, and confidentiality patterns documented in the approved dossier. Deviations are prohibited unless authorised through the change-control provisions in Chapter 9. The pilot cohort must be selected to represent diverse organisational sizes, sector contexts, and geographical conditions relevant to the model. Participants must be informed that they are part of a pilot and must receive all consent, revocation, and privacy protections applicable to full-scale operations.

The pilot must generate evidence of methodological validity, fairness, confidentiality performance, operational feasibility, and user comprehension. It must test the effectiveness of secure enclave access, redaction error rates, evidence workloads for different size tiers, the operational realism of sampling and field-check protocols, AI performance where deployed, and the behaviour of consent and revocation flows.

Calibration occurs concurrently with and immediately following the pilot. Calibration tests must include blind parallel assessments, scenario comparison exercises, inter-assessor variance checks, threshold stability tests, and proportionality back-testing. Where quantitative indicators exist, calibration must ensure that they do not concentrate outcomes at extremes or disproportionately penalise particular categories of entities. Where qualitative judgement is applied, calibration must demonstrate consistency, repeatability, and resistance to bias.

The pilot must also test public-facing elements under controlled conditions where consent is granted. This includes iconography placement, communication disclaimers, and the public interpretation of outcomes. It must ensure that no public-facing artefact can be misinterpreted as certification, superiority, ranking, or authoritative endorsement beyond what is permitted under the Standard.

The scale-up decision requires evidence of readiness across four domains: model integrity, organisational readiness, technical readiness, and ethical-legal readiness. Model integrity reflects fidelity to the dossier and the ability to operate without unintended comparative or discriminatory effects. Organisational readiness evaluates staffing, assessor competence, internal ethics controls, and monitoring systems. Technical readiness verifies enclave infrastructure, consent ledger functionality, data-flow mapping, and security integrity. Ethical-legal readiness ensures that the model has demonstrated compliance with confidentiality principles, revocation mechanics, lawful bases for data handling, and sector-specific risk mitigations.

Agenda 2074 and GSIA jointly determine whether a model has met readiness gates. Approval for scale is issued only when the model demonstrates stable calibration, maintained proportionality, operational feasibility, and intact confidentiality safeguards.

Table 12: Pilot and scale-up Gate Criteria

Gate Domain	Pilot Requirements	Scale-Up Requirements	Oversight
Model Integrity	Accurate implementation of approved methodology; indicative calibration	Stable calibration; no doctrinal deviations	Agenda 2074 standards review

Confidentiality & Consent	No breaches; redaction error ≤ defined threshold; revocation tests passed	Sustained enclave performance; ledger integrity; repeatable suppression	GSIA privacy & digital review
Proportionality & Fairness	No adverse size/sector bias; burden within design limits	Documented long-form fairness verification	GSIA fairness verification
Operational Readiness	Assessor competence; workflow stability	Full staffing; ethics controls fully functional	GSIA ethics oversight
Technical Infrastructure	Enclave uptime; log integrity; UX stability	Scalable infrastructure; KRIs implemented	GSIA digital oversight
Legal & Jurisdictional	Lawful bases confirmed; cross-border measures tested	Compliance verified across intended markets	Legal Note compliance
Communications	Correct disclaimers; no misinterpretation observed	Stable, compliant public communications	Agenda 2074 communications unit

Only after all gate criteria are satisfied may the model be authorised for general deployment. Any deficiency triggers a corrective cycle, additional calibration, or a revised pilot phase. Approval for scale does not diminish monitoring obligations; it increases them. Calibration becomes periodic, KRIs become mandatory, and readiness for future versions becomes dependent on documented stability during scale.

All pilot and calibration records must be preserved as part of the permanent model archive. These records inform monitoring, future revisions, thematic audits, and potential sunset decisions in accordance with the Operating Manual and the Governance & Oversight Manual.

Chapter 9 — Maintenance, Change Control, and Versioning

This Chapter establishes the obligations that all Validation Partners must follow to maintain the integrity, continuity, and transparency of their validation models once they have been approved and deployed. Maintenance encompasses all operational, methodological, and ethical responsibilities required to ensure that a model remains accurate, proportionate, confidential, and aligned with the evolving interpretive canon of the 17 SGG pillars. Change control governs how modifications are designed, assessed, communicated, and approved. Versioning ensures that every update is traceable, auditable, and historically comparable, preserving stability for validated entities and clarity for external oversight.

Model maintenance is a continuous duty. Validation Partners must monitor real-world performance, KRIs, calibration results, thematic audit findings, consent ledger anomalies, enclave logs, redaction error rates, assessor feedback, and user experience signals. When patterns emerge indicating potential drift, burden imbalance, confidentiality risk, or interpretive misalignment, the Partner must initiate a structured review. Maintenance encompasses updating documentation, refining criteria, adjusting proportionality algorithms, strengthening privacy controls, enhancing assessor training, and remedying any operational weaknesses identified under GSIA oversight.

Change control is required for any modification that may affect criteria, indicators, scoring logic, recognition levels, evidence demands, proportionality rules, AI components, confidentiality patterns, branding, public-facing artefacts, or client obligations. All changes—minor, moderate, or major—must be documented, justified, tested in a controlled environment, and submitted to Agenda 2074 and GSIA for review when material. Each change must undergo an updated Ethics Impact Assessment if it introduces new risks, new evidence requirements, or new technological components. Changes affecting confidentiality, consent, secure enclaves, or public-facing outputs require mandatory GSIA ethics review prior to implementation.

Versioning creates permanent, auditable lineages of every model. Each approved version must bear a unique Model ID supplemented by a semantic version number reflecting the nature of the change: a major version for structural or doctrinal updates; a minor version for methodological adjustments not altering recognition; and a patch version for errata or clarifications. Partners must maintain a complete version archive, including design documents, EIA outcomes, calibration notes, pilot evidence where relevant, crosswalks demonstrating doctrinal fidelity, and decision records issued by Agenda 2074 or GSIA.

Backward compatibility must be preserved unless there is an overriding confidentiality or ethical justification for retirement. Where a change materially affects validation outcomes, Partners must provide transition plans including clear timelines, stakeholder notices, updated training, and proportional grace periods. Entities must not be disadvantaged by changes imposed without adequate notice or transition support.

All changes must be communicated in accordance with the Communication & Public Disclosure Protocol. Communications must avoid comparative language, must not misrepresent the nature of the update, and must safeguard confidentiality by ensuring that no entity can be indirectly identified through examples, transition matrices, or data illustrations.

Table 13: Change Control Categories and Required Process

Change Category	Description	Required Process	Oversight
Patch (Non-Substantive)	Editorial corrections; clarity improvements; documentation updates	Document change; update version; notify Agenda 2074	Agenda 2074 optional acknowledgement
Minor Method Change	Adjustments not altering recognition or evidence burden	Internal testing; updated EIA note; advance notice to GSIA	GSIA review for proportionality and privacy
Major Method Change	Structural updates; new indicators; altered recognition logic; new AI	Full EIA; pilot or sandbox; formal re-approval	Agenda 2074 licensing; GSIA binding ethics decision

Confidentiality-Critical Change	Any change affecting consent, enclaves, secure flows	Mandatory ethics review; updated KRI plan	GSIA privacy oversight

Partners must treat versioning as a compliance obligation rather than an administrative task. Failure to maintain accurate version lineage or to follow approved change-control procedures constitutes a material breach subject to GSIA's escalation system, including suspension of the affected model.

Maintenance and change-control records must be bound into a single "Model Maintenance Ledger," which serves as the authoritative source during monitoring cycles, thematic audits, and adjudicative proceedings. The ledger must be immutable, timestamped, and accessible to GSIA upon request. It must contain all version histories, EIA updates, calibration results, risk assessments, change rationales, and supporting artefacts demonstrating that updates preserve doctrinal alignment and confidentiality protections.

Model maintenance is perpetual. No model may enter a dormant state without prior written notice to Agenda 2074 and GSIA, nor may an unmaintained model continue operating once significant drift or risk is detected. Where a model reaches obsolescence due to technological, sectoral, normative, or risk-based evolution, a sunset pathway must be executed consistent with the Operating Manual and Governance & Oversight Manual. Partners must notify validated entities of retirement timelines and ensure continuity of private recognition until the sunset period concludes or a successor model is available.

In all cases, change control must favour stability, privacy, and fairness. Progress is welcome; disruption is not. Every model must evolve deliberately and transparently, ensuring that the ecosystem remains robust, trustworthy, and compliant with the Standard's ethical and doctrinal foundations.

Final Word

This Guide concludes by affirming that the development and operation of validation models under A2074-SRS is both a technical discipline and a fiduciary responsibility. Validation Partners are stewards of a public-interest standard anchored in the 17 Social Global Goals, safeguarded by GSIA's independent ethical authority, and grounded in a confidentiality doctrine that protects participants from harm, coercion, exposure, or comparative misuse.

The development process is not merely procedural. It is an ethical commitment to design models that are doctrinally faithful, proportionate, least-intrusive, and transparent. It is a commitment to ensure that scaling does not compromise privacy, that evidence demands do not exceed necessity, that public-facing artefacts do not mislead or misrepresent, and that entities of all sizes can demonstrate meaningful alignment without unfair burden or comparison.

Throughout this Guide, the principles of interpretive fidelity, confidentiality-by-design, proportionality, non-comparative evaluation, and ethical independence form the backbone of every obligation. Model development begins with doctrinal anchoring, proceeds through rigorous evidence design and Ethics Impact Assessment, undergoes structured piloting and calibration, and is maintained through disciplined change control and versioning. At every stage, GSIA and Agenda 2074 act as guardians of fairness, integrity, and trust.

A2074-SRS is not a certification scheme, a reputational ranking tool, or a competitive marketplace for superiority claims. It is a governance standard for responsible conduct, supported by a validation ecosystem designed to help organisations understand, measure, improve, and privately demonstrate alignment with social responsibilities that matter to communities, stakeholders, and society. Everything in this Guide reinforces that mission.

Validation Partners who adopt these principles and follow these processes contribute to a global architecture of equity, accountability, and confidentiality. They help build a system where organisations of all sizes can advance responsibly, where validation adds understanding rather than pressure, and where learning prevails over comparison. In such a system, trust is generated not by publicity or competition, but by quiet, rigorous, safeguarded verification carried out with independence, care, and respect for those who entrust their information to the Standard.

This Guide enters into force upon issuance and applies to all Validation Partners operating or seeking to operate within the A2074-SRS ecosystem. Its fidelity will be measured by the consistency with which Partners honour its obligations and by the trust it sustains across the many actors who rely on this Standard.