JANUARY 27, 2026

Digital Integration & Platform Governance Manual

Agenda for Social Equity 2074

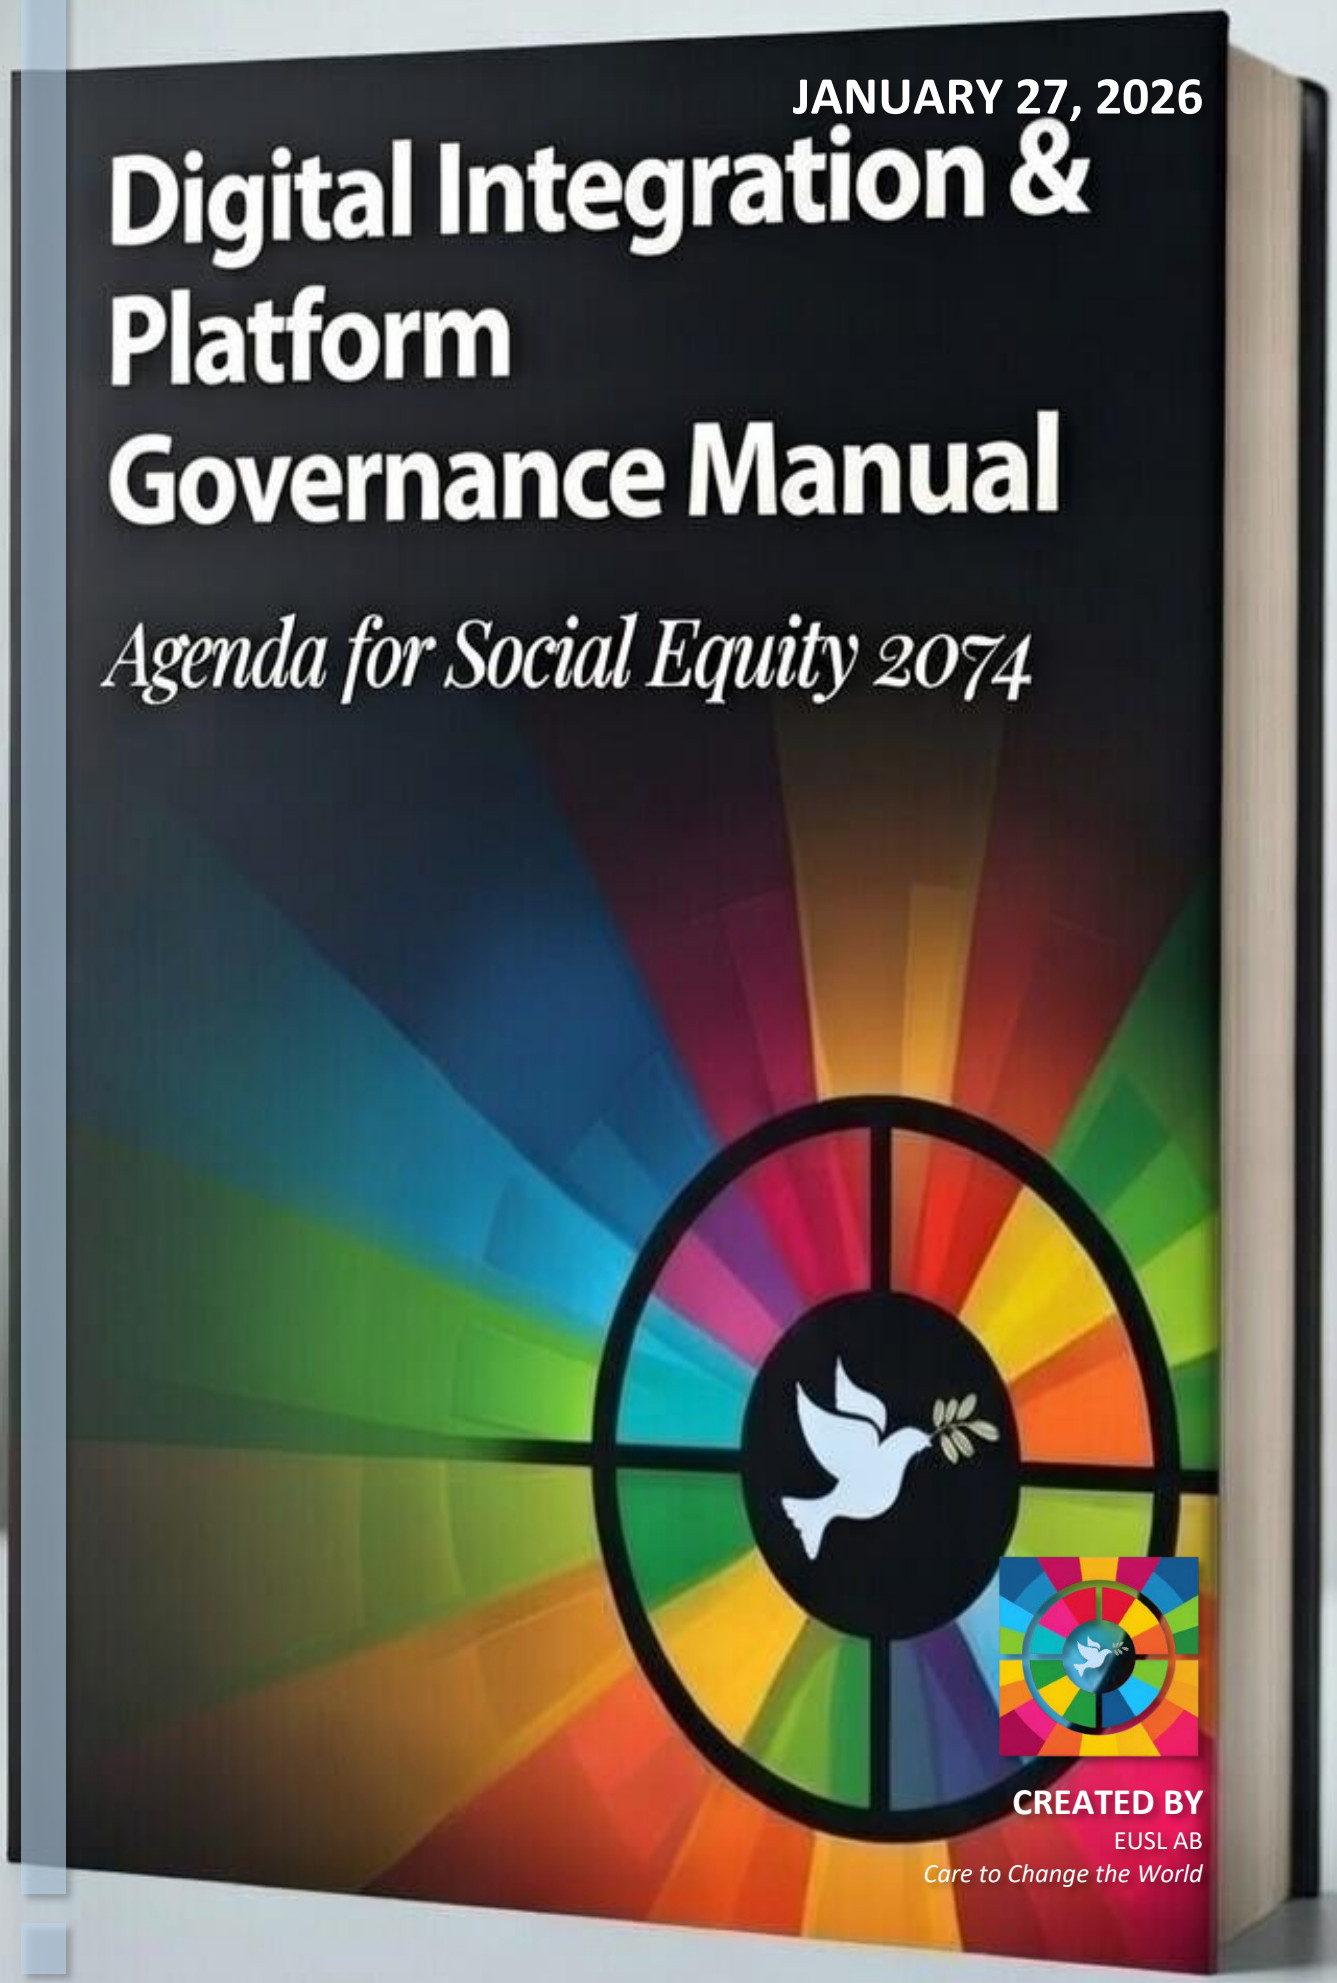

CREATED BY

EUSL AB

Care to Change the World

Table of Contents

Validation System White Paper

Introduction

This Manual governs the technical infrastructure, security posture, and platform governance arrangements through which the A2074-SRS ecosystem is operated. It has binding effect on all Validation Partners, contracted vendors, hosting providers, and other processing agents that participate in the collection, processing, transmission, storage, and disclosure of data related to A2074-SRS validations, including meta-data and derived analytics. The Manual is adopted under the authority of Agenda 2074 as standard-setter, is subject to the oversight and adjudicative jurisdiction of GSIA's ethics and compliance chambers, and is operationalized by accredited Validation Partners in accordance with the Licensing & Accreditation Framework and the Operating Manual (Open Standard).

The institutional architecture is preserved in its entirety. Agenda 2074 defines the universal canon through the seventeen Social Global Goals (SGGs), which constitute the normative pillars for validation. Validation Partners design and operate one or more approved validation models—such as hospitality-style star systems, points, maturity pathways, sector modules, and single-goal deep dives—provided that each model conforms to this Manual and the Multi-Model Validation Framework. GSIA remains structurally independent as the ethics and compliance custodian with powers to supervise, investigate, adjudicate, and, where necessary, to impose corrective measures or suspend access to systems. Patient-level confidentiality governs all results. Records are private by default; disclosure occurs only upon explicit, informed, and revocable consent by the rights-holder, subject to lawful exceptions determined by GSIA. Proportionality applies throughout, ensuring non-comparative evaluation between entities of different sizes and capacities, and leading to fairness in presentation, analytics, and any permitted disclosures.

The Manual establishes a layered and interoperable architecture with privacy-by-design, security-by-design, and minimum-disclosure-by-default as first-order requirements. It defines master registries, a global identity and access regime, cryptographic controls, evidence handling rules, consent ledgering and revocation protocols, and interfaces for secure integration with EUSL, WOSL, DESA, GSIA, and external systems. AI-assisted decision support may be employed for triage and analytics under the guardrails set herein and in the Ethics & Integrity Code, with human oversight, auditability, and bias controls. All provisions herein shall be construed harmoniously with the Foundational Charter, the Governance & Oversight Manual, the Communication & Public Disclosure Protocol, the Legal Compliance & International Law Note, and the web artefacts of A2074-SRS. Change control, backward compatibility, and versioning are governed by Chapter 13, and no derogation from the confidentiality and consent requirements in Chapters 6, 15, and 16 is permitted without GSIA authorization.

EUSL operates as the flagship Validation Partner in Europe, applying a hospitality-style star model aligned to the seventeen pillars. Nothing in this Manual authorizes any claim of ISO certification; any use of ISO 26000 shall be limited to optional self-declaration by applicants as contextual information, without representing or implying certification.

The Manual proceeds from the foundational elements of system architecture and identity to the operational specifics of cybersecurity, privacy, evidence stewardship, auditability, analytics, incident response, vendor risk, change management, and the rights of the "patient-analogue" data subject. Each Chapter is enforceable as a normative control domain. Cross-references are provided to maintain coherence and to anchor decisions in the governing canon and GSIA's oversight.

Chapter 1 — System Architecture and Master Registries

The platform shall be organized as a layered, multi-tenant, and jurisdiction-aware system composed of a presentation layer, a service and API layer, a data layer with logically segregated tenant partitions, and a control plane dedicated to identity, keys, logging, monitoring, consent ledgering, and configuration management. The design shall implement minimum disclosure by default and shall instantiate access only upon a demonstrated lawful basis or valid, informed, and revocable consent. Time synchronization, monotonic counters, and cryptographic hashing of critical records shall be deployed to sustain non-repudiation and traceability as further elaborated in Chapter 9.

Master registries constitute the authoritative sources of truth for entities, events, artifacts, and decisions in the ecosystem. Each registry is singular in its authoritative function, is assigned a steward accountable for completeness and integrity, and exposes versioned read interfaces to approved consumers through secure APIs under Chapter 4. Write interfaces are restricted to designated processes with segregation of duties enforced under Chapter 2.

Global identifiers shall be unique, stable, and non-meaningful. They shall not encode personal data or sensitive semantics. Identifiers for public display, where permitted by consent, shall be distinct from internal identifiers and shall be mapped through protected look-up tables accessible only within the control plane. Registries shall support event sourcing with append-only logs for critical state transitions, retaining prior states for audit and rollback where permitted by Chapter 13 and Chapter 9. Jurisdictional residency constraints shall be respected by physical or logical partitioning, with data location recorded at creation and updated upon any lawful migration.

The following registries are hereby established and shall be maintained as authoritative within their scope. Access baselines are indicative and are to be further constrained by role-based and attribute-based policies defined in Chapter 2 and the consent regime in Chapters 6, 15, and 16.

Registry	Purpose	Authoritative Steward	Primary Identifier	Key Contents	Access Baseline	Retention & Disposition	Cross-References
Partner Registry	Enrolment and status of Validation Partners and sub-entities	A2074 Secretariat with GSIA oversight	PR-GID	Legal entity data, accreditation scope, jurisdictional footprint, contacts, keys	Read: GSIA, Partners (self); Write: Secretariat	Active + 10 years post-termination; archival hashes retained	Licensing & Accreditation Framework; Chapter 12; Chapter 13
Model Registry	Approved validation models and versions	A2074 Secretariat	MR-GID	Model definitions, controls, scoring taxonomie	Read: GSIA, Partners; Write: Secretariat	Indefinite for historical audit	Multi-Model Validation Framework; Chapter 13

				s, change logs			
Engagement Registry	Per-applicant validation engagements	Validation Partner	ER-GID	Applicant profile, scope, timeline, assigned assessors, status	Read: Partner; Conditional read: GSIA; Write: Partner	Active + 7 years; evidence retention per Chapter 7	Operating Manual; Chapter 7; Chapter 9
Validation Result Registry	Canonical record of outcomes	Validation Partner with GSIA read authority	VR-GID	Model used, dated outcome, lineage to evidence hashes, disclosure flags	Read: Partner; Read: GSIA; Public read only if consent	Active + 10 years	Chapters 6, 7, 9, 10, 15, 16
Disclosure Registry	Records of any public or third-party disclosures	Validation Partner; GSIA monitors	DR-GID	Consent reference, scope, audience, channels, timestamps, revocations	Read: Partner, GSIA; Public read where applicable	Active + 10 years; revocation cascades recorded	Chapter 8; Chapter 10; Chapter 15; Chapter 16
Consent Ledger	Immutable ledger of consent states and revocations	A2074 Secretariat (control plane)	CL-GID	Subject identity assertions, consent scope, timestamps, cryptographic attestations, revocation DAG	Read: Partner (need-to-know), GSIA; Write: Controlled workflows	Indefinite for meta-data; personal linkages per Chapter 15	Chapters 6, 15, 16

Evidence Repository Index	Index of evidence objects and chain-of-custody	Validation Partner; GSIA has oversight	EI-GID	Object hashes, storage location, integrity proofs, access controls, versions	Read: Partner; Read: GSIA under mandate; Write: Controlled ingestion	Evidence per legal limits; index indefinite	Chapter 7; Chapter 9
Identity & Role Directory	Authoritative roles and entitlements	A2074 Secretariat	IR-GID	Identities, roles, attributes, delegations, SoD constraints	Read: GSIA, Partner admins (scope-limited); Write: Identity controllers	Lifecycle-bound; audit trails indefinite	Chapter 2; Chapter 9
Key & Certificate Registry	Cryptographic materials inventory	A2074 Secretariat	KC-GID	Key IDs, purposes, rotation history, algorithm policies	Read: Security officers; Write: HSM-mediated	Per crypto policy; logs indefinite	Chapter 5; Chapter 9
Audit & Event Ledger	Immutable audit events and decisions	A2074 Secretariat with GSIA read	AE-GID	Time-stamped events, actor, action, object, result, hash chain	Read: GSIA; Scoped read: Partners; Write: System-only	Indefinite or as mandated by law	Chapter 9; Chapter 11
Vendor & SLA Registry	Third-party risk and obligations	A2074 Secretariat	VRM-GID	Contracts, SLAs, attestations, assessments, issues	Read: GSIA; Scoped read: Partners; Write: Secretariat	Active + 7 years	Chapter 12
Jurisdictional Rules Registry	Data transfer, residency, and lawful bases	A2074 Secretariat (Legal)	JR-GID	Applicable laws, restrictions, adequacy mappings,	Read: All controllers; Write: Legal function	Continuous with versioning	Chapter 6

				derogations			
Public Aggregates Registry	Sanitised aggregates for dashboards	A2074 Secretariat	PA-GID	Statistical outputs, re-ID risk scores, release notes	Public read where published; Write: Secretariat	Versioned indefinite	Chapter 10; Chapter 17

Inter-registry relationships shall be explicit and validated. A Validation Result must reference an Engagement and a Model version; any Disclosure must reference a specific Consent state and scope; Evidence objects must be traceable to their hash in the Evidence Repository Index and linked to the corresponding Validation Result. No publication pipeline shall bypass the Disclosure Registry or the Consent Ledger. Where residency or transfer constraints apply, the Jurisdictional Rules Registry must be consulted programmatically before any cross-border movement, with decisions recorded in the Audit & Event Ledger.

The control plane shall enforce configuration immutability for security-critical settings, with changes executed through signed change requests, peer review, and time-boxed execution windows under Chapter 13. Break-glass mechanisms may be defined only for availability incidents and shall themselves be recorded, temporally constrained, and subject to GSIA post-incident review under Chapter 11.

Chapter 2 — Identity, Access, and Authentication

Identity, access, and authentication are governed by the principles of least privilege, segregation of duties, verifiable provenance of actions, and patient-level confidentiality. Access decisions shall require both a role-based entitlement and, where applicable, attribute conditions that reflect the context, including jurisdiction, tenancy, data residency, consent scope, and the sensitive nature of the requested operation. All access to personal data and evidence materials is denied by default and must be positively authorized through policies binding on the Identity & Role Directory and enforced in the service layer.

Roles are defined to reflect institutional responsibilities. The A2074 Secretariat acts as standard-setter and control plane custodian, with authority over registries, configuration, and accreditation metadata. GSIA exercises independent oversight, with read and investigative access sufficient to discharge ethics and compliance functions, alongside temporary elevated entitlements granted under a documented case mandate. Validation Partners administer their tenant domains, assign assessors, manage engagements, and submit results to the canonical registries. Applicant organizations administer their own users and may access only their own engagement records, evidence they have submitted, and disclosure settings within the limits of consent. Evidence custodians, whether internal or contracted, operate within sealed workflows that bind identity to chain-of-custody events. Public visitors, where a disclosure exists, may access only the specific public materials that have been consented to and released.

Authentication shall employ strong, phishing-resistant multi-factor methods as the baseline for all administrative and assessor roles. Hardware-backed authenticators are required for privileged operations in the control plane and for any action that changes disclosure states, keys, or policies. Step-up authentication is mandatory for sensitive operations including the creation, amendment, or

revocation of consent records; the publication or withdrawal of disclosures; the approval of validation results; and any access to raw evidence materials. Sessions shall be bound to device posture where feasible, and idle and absolute timeouts shall be enforced commensurate with the actor’s risk profile and the operation being performed. Secrets and service credentials shall be managed through a centralized secrets vault, rotated automatically, and never embedded in code or configuration artifacts.

Federated identity may be employed for Validation Partners and applicant organizations through industry-standard protocols with signed metadata, explicit trust anchors, and constrained scopes. Federation must be governed by written trust agreements that specify assurance requirements for identity proofing and authenticator strength, incident notification duties, and deprovisioning timelines. Where federation is used, attribute release shall be minimized, and pseudonymous, non-reversible identifiers shall be issued for any linkage to patient-analogue records to avoid correlability across tenants. The platform shall support just-in-time provisioning subject to pre-approved role mappings and shall enforce joiner-mover-leaver processes with prompt entitlement reviews and revocation upon separation. Periodic access attestation shall be conducted by Partner administrators and audited by GSIA, with attestations and exceptions recorded in the Audit & Event Ledger.

Separation of duties shall prohibit the same individual from initiating and approving actions that materially affect consent, disclosure, cryptographic keys, validation outcomes, or registry schemas. Emergency break-glass accounts shall be technically segregated, stored in sealed escrow, tested on a defined cadence for availability purposes, and monitored with real-time alerts and immediate GSIA notification upon use. Service accounts shall be scoped to the smallest feasible set of operations, restricted by network and time constraints, and instrumented with non-interactive credentials bound to specific workloads. No personal account shall be used for batch processing or integration.

The following matrix expresses the baseline relationship between roles and registry access. It is indicative and shall be tightened by attribute-based policies, consent scope, and case-specific mandates.

Role	Partner Registry	Model Registry	Engagement Registry	Validation Result Registry	Disclosure Registry	Consent Ledger	Evidence Repository Index	Identity & Role Directory	Audit & Event Ledger
A2074 Secretariat (Control Plane)	Read/Write	Read/Write	Read (meta)	Read (meta)	Read (meta)	Read/Write	Read (meta)	Read/Write	Read/Write
GSIA Ethics & Compliance	Read	Read	Read (case-bound)	Read (case-bound)	Read	Read (case-bound)	Read (case-bound)	Read (scope-bound)	Read
Validation Partner Admin	Read (self)	Read	Read/Write (tenant)	Read/Write (tenant)	Read/Write (tenant)	Read (scope-bound)	Read/Write (tenant)	Read/Write (tenant)	Read (tenant)

Validation Assessor	Read (self scope)	Read	Read/Write (assigned)	Write (proposed)	Read (proposed)	Read (scope-bound)	Read/Write (assigned)	Read (assigned)	Read (assigned)
Applicant Org Admin	Read (public)	Read (public)	Read (own)	Read (own)	Read/Write (own)	Read (own scope)	Read/Write (own submissions)	Read/Write (own)	Read (own)
Evidence Custodian	None	None	Read (assigned meta)	None	None	None	Read/Write (assigned)	None	Read (own events)
Public Visitor	Read (public extracts)	Read (public notes)	None	Read (public extracts)	Read (public entries)	None	None	None	None

Assurance requirements shall scale with risk. Control plane operators, GSIA investigators with live-case access, and Partner administrators shall meet the highest assurance standards defined by the platform's authentication policy, including hardware-backed, phishing-resistant multi-factor authentication, constrained network locations where feasible, and periodic re-verification of identity. Assessor and applicant roles shall use strong multi-factor authentication with device binding where feasible. Public visitors shall have no authenticated access unless required for controlled access to semi-private disclosures, in which case contextual and time-limited tokens may be issued.

Identity assertions used to bind consent to a data subject in the patient-analogue sense shall meet heightened verification standards proportionate to the sensitivity of the data and the jurisdictional requirements referenced in Chapter 6. Where legally permissible and proportionate, pseudonymous identifiers shall be used for operational processing, with the re-identification keys held separately under the control plane and released solely under explicit consent or GSIA-authorized lawful basis. All identity lifecycle events, including federation trust changes, authenticator additions or removals, step-up prompts, and denial decisions, shall be recorded in the Audit & Event Ledger with immutable timestamps and actor provenance.

Authentication failures, anomalous access requests, and privilege escalation attempts shall be monitored in real time, with automated containment where feasible and mandatory notification to Partner security officers and the A2074 Secretariat under the thresholds and timelines set forth in Chapter 11. No identity or access modification shall be effective without successful policy evaluation in the control plane and durable recording of the outcome.

Chapter 3 — Data Model and Schemas

The data model standardises core entities, relationships, and exchange formats to ensure unambiguous interoperability across all Validation Partners and platform services. It is constrained by the principles of patient-level confidentiality, minimum disclosure by default, and non-comparative evaluation. All schemas are versioned, forward-compatible, and governed by the change procedures in Chapter 13.

No schema may encode personal data into identifiers. All personal data shall be compartmentalised, with explicit consent states referenced by immutable consent tokens recorded in the Consent Ledger, as described in Chapters 6 and 15.

The canonical entities are defined to mirror the master registries in Chapter 1 and to provide a single source of semantic truth. Each entity bears a globally unique, non-meaningful identifier and lifecycle metadata, including creation time, last modification time, and the hash of the prior version to support chain integrity under Chapter 9. Relationships are explicit, directional, and validated at write time. Mandatory referential constraints are enforced such that no Validation Result can exist without an Engagement, no Disclosure can exist without a matching consent scope in the Consent Ledger, and no Evidence Index entry can exist without an associated Engagement and an integrity proof.

The following table enumerates the canonical entities, their purposes, identifiers, and principal linkages. Field lists are indicative of the minimum required set and may be extended by partner-specific namespaces that are isolated and non-interfering with the canonical namespace.

Entity	Purpose	Primary Key	Required Relationships	Selected Canonical Fields (non-exhaustive)	Versioning
Partner	Represents an accredited Validation Partner and its tenant boundary	partner_id	—	legal_name; jurisdiction; accreditation_scope; contact_points[]; cryptographic_keys[]; status	schema_version; record_version; prev_hash
Model	Defines an approved validation model and its versioned parameters	model_id	—	name; version_tag; control_catalog[]; scoring_taxonomy; applicability_scope; release_notes	As above
Applicant	Represents the applicant organisation within a specific engagement context	applicant_id	—	legal_name; sector_codes[]; jurisdiction; size_attributes; contact_points[]	As above
Engagement	Captures a specific validation	engagement_id	partner_id; applicant_id; model_id	scope; start_date; assessors[]; status;	As above

	instance between a Partner and an Applicant			lawful_basis; residency_zone	
EvidenceIndex	Registers evidence artefacts, hashes, and locations	evidence_id	engagement_id	object_type; content_hash; storage_locator; created_by; access_policy_ref; chain_of_custody[]	As above
ValidationResult	Canonical outcome for an engagement and model version	result_id	engagement_id; model_id	issued_at; outcome_payload; evidence_hashes[]; disclosure_flags; lineage[]	As above
ConsentState	Immutable states of consent as recorded in the Consent Ledger	consent_id	—	subject_ref; scope; issued_at; expires_at; state_enum; attestations[]	As above
Disclosure	Records a disclosure event and its scope, bound to a consent state	disclosure_id	result_id; consent_id	audience; channels[]; published_at; revocation_ref; non_display_controls	As above
AuditEvent	Immutable log of actions, decisions, and changes	audit_id	—	actor_ref; action; target_ref; timestamp; outcome; event_hash; signature	Append-only
JurisdictionRule	Governs transfers, residency, and lawful bases	rule_id	—	jurisdiction; subject_scope; transfer_constraints; adequacy_mappings; effective_period	As above

VendorProfile	Third-party service profile and risk posture	vendor_id	—	services; attestations; SLA_refs; risk_rating; monitoring_refs	As above
PublicAggregate	Anonymised aggregate published for public dashboards	aggregate_id	—	methodology_ref; release_notes; reid_risk_score; measures[]	As above

Enumerations shall be controlled centrally through the Model entity and the JurisdictionRule entity to avoid drift across partners. The minimum controlled lists include sector codes, jurisdiction codes, residency zones, lawful bases, consent states, evidence object types, outcome classifications, and disclosure channels. Controlled lists are published via read-only API endpoints described in Chapter 4 and are cached locally subject to time-to-live directives; stale enumerations shall block writes that depend upon updated values.

The exchange formats shall be JSON as the canonical wire format and XML as an optional alternative for legacy integrations. Both formats share the same semantic contract and versioning regime. All payloads include headers for schema_version, record_version, record_hash, and prev_hash where applicable. Signatures are carried out-of-band at the transport layer and in-band via detached JWS artifacts referenced in the header, providing layered non-repudiation under Chapter 9. Binary evidence content is never transported inline in integration payloads; only integrity proofs and locators appear in the EvidenceIndex, with access mediated by controlled retrieval workflows.

To preserve privacy, personal data fields are minimised and isolated. Where contact points are unavoidable, they are attributed to organisational roles rather than named individuals, unless strictly necessary and consented. Pseudonymous identifiers shall be used within operational processing, with reversible keys held exclusively in the control plane and accessible only under the conditions set forth in Chapters 6 and 15.

Data quality controls are mandatory at ingress. Validation includes schema conformity, referential integrity, enumeration validity, jurisdictional compatibility checks against the JurisdictionRule entity, and consent resolution checks against the ConsentState entity. Failure paths are explicit and recorded as AuditEvents with associated remediation guidance. Partial writes are prohibited for canonical entities; transactions shall be atomic with compensating events recorded for any rollback.

Extensions are allowed through namespaced custom objects. A partner may add eusl:* fields within the Engagement or ValidationResult payloads to support EUSL’s hospitality-style star model, provided these do not conflict with canonical fields and are hidden from public exposure by default. Namespaced fields are not propagated to public aggregates unless explicitly mapped and approved through the change procedures in Chapter 13, with re-identification risk reviewed under Chapter 17.

Chapter 4 — APIs and Integration Standards

Interfaces are designed to expose controlled, auditable, and versioned access to the canonical entities, while enforcing consent, residency, and security controls. All APIs are RESTful over HTTPS with mutual

TLS for partner-to-platform interactions and OAuth 2.1-compliant authorisation with JWT bearer tokens. Event-driven integration is supported through signed webhook callbacks and a secure message bus for high-assurance partners under curated topics. Idempotency is required for all mutating operations and is enforced through idempotency keys tied to the requesting identity and payload hash.

The baseline interface taxonomy aligns to the master registries and enumerations. Read endpoints are broadly available within the actor's scope; write endpoints are tightly restricted and require role, attribute, and consent checks in the control plane. Public endpoints expose only explicitly published, anonymised aggregates and selected disclosures. No endpoint shall permit retrieval of raw evidence; evidence retrieval occurs through time-bound, scope-limited pre-signed URLs or mediated viewers with watermarking and access logging, pursuant to Chapter 7.

The following tables summarise the principal endpoint families and their control characteristics. Paths are illustrative and may be prefixed by region or tenant selectors to satisfy residency requirements.

Endpoint Family	**Selected Operations**	**Auth Requirements**	**Consent & Residency Enforcement**	**Notes**
/partners	GET /partners/{id}; PATCH /partners/{id}	mTLS + OAuth; Partner Admin or Secretariat	Residency not applicable; consent not applicable	Write limited to Secretariat; Partner self-service scope is narrow
/models	GET /models; GET /models/{id}	OAuth; all internal roles	Not applicable	Authoritative catalogue for model versions
/engagements	POST /engagements; GET /engagements/{id}; PATCH /engagements/{id}	OAuth; Partner Admin/Assessor; Applicant Admin (own)	Jurisdictional check at create; consent pre-check optional	Create requires lawful basis declaration
/evidence	POST /evidence/index (metadata); POST /evidence/upload-ticket; GET /evidence/{id} (metadata only)	OAuth; Partner Assessor/Custodian	Residency enforced at upload; access mediated	No raw bytes through API; only locators and proofs
/results	POST /results; GET /results/{id}; POST /results/{id}:approve	OAuth; Partner Admin; GSIA read	Consent gating for any publication flag	Approval step-up authentication required

/consent	POST /consent; GET /consent/{id}; POST /consent/{id}:revoke	OAuth; Secretariat (write); Partner (initiate workflows)	Consent ledger rules; jurisdictional limits	Writes execute only via control plane workflow
/disclosures	POST /disclosures; GET /disclosures/{id}; POST /disclosures/{id}:withdraw	OAuth; Partner Admin; GSIA read	Consent scope validation; revocation cascades	Immediate withdrawal orchestration per Chapter 15
/jurisdictions	GET /jurisdictions/rules; GET /jurisdictions/{code}	OAuth; all internal roles	Not applicable	Used for pre-flight checks
/audit	GET /audit/events; GET /audit/events/{id}	OAuth; GSIA; scoped Partner read	Not applicable	Immutable, append-only reads
/public	GET /public/disclosures/{id}; GET /public/aggregates/{id}	Public or tokenised	Consent and non-display rules enforced upstream	Read-only, sanitised outputs

All write operations use a canonical request envelope that includes a detached JWS signature reference, idempotency key, and actor context. The platform enforces replay protection and evaluates tenancy, role, attributes, and consent in the control plane before committing state. Representative request and response envelopes are shown below.

Integration with EUSL, WOSL, DESA, and GSIA is accomplished through partner-specific adapters that map namespaced fields and enforce each counterpart’s assurance expectations without derogating from the canonical standards. EUSL’s hospitality-style star system integrates by reading the Model catalogue, submitting Engagements and proposed Results with eusl:* annotations, and invoking publication only when a valid consent state exists. WOSL public-facing portals read from /public/aggregates and from specific public Disclosures, never from internal registries. DESA’s programmatic systems, where acting as Applicant or as an ecosystem contributor, use federation with attribute minimisation and submit evidence through the standard EvidenceIndex workflows. GSIA interfaces are read-dominant, with case-bound elevation to retrieve specific AuditEvents, ConsentStates, and EvidenceIndex metadata necessary for ethics review; GSIA does not write to partner-tenant registries and exercises its adjudicative powers by issuing binding directives recorded as AuditEvents and carried out by the Secretariat or the relevant Partner.

Eventing is available through signed webhooks and, for high-assurance partners, through a message bus. Topics include engagement.statusChanged, result.approved, disclosure.published, disclosure.withdrawn, consent.revoked, and jurisdiction.ruleUpdated. Subscribers must register public keys in the Key & Certificate Registry. Delivery is retried with exponential backoff and dead-lettered upon persistent failure, with GSIA visibility into dead-letter queues for matters pertaining to ethics and compliance. Webhook payloads exclude personal data and carry only identifiers and context sufficient for the subscriber to fetch scoped detail via pull APIs.

Rate limiting and fairness are enforced to uphold proportionality and to protect availability. Throttles scale with partner tier and use token buckets per endpoint family, with separate budgets for public endpoints to avoid interference between public interest access and partner operations. Abuse detection and anomaly scoring are applied in real time. Denials are precise and logged as AuditEvents with reason codes, enabling structured appeals and GSIA review where systemic impacts are alleged.

Backward compatibility is maintained through additive changes, deprecation windows, and explicit versioning at both the schema and endpoint levels. Breaking changes require a major version increment, public release notes, sandbox availability, and a migration window administered under Chapter 13. No partner may be forced onto a breaking change without a reasonable migration period, and no change may compromise confidentiality, consent enforcement, or GSIA oversight.

Security controls at the interface level are mandatory and layered. TLS uses strong ciphers with certificate pinning for server-to-server integrations. JWT access tokens are short-lived, audience-scoped, and bound to mTLS identities. All mutating endpoints require step-up authentication when they affect consent, disclosure, keys, or registry schemas. Payloads are scanned for structural anomalies, enumeration violations, and jurisdictional conflicts before any state change is accepted. All accepted and rejected calls are recorded in the Audit & Event Ledger with immutable timestamps and actor provenance to satisfy the non-repudiation requirements of Chapter 9 and to feed incident response under Chapter 11.

Cross-border data movement is orchestrated by the control plane using the JurisdictionRule registry. Pre-flight checks are mandatory for any operation that would cause residency changes or third-country disclosures. Where such operations are attempted, the platform enforces lawful bases and technical safeguards, including encryption key residency constraints, tokenisation, or the prohibition of the operation, with detailed rationale recorded. Partners are provided with deterministic error codes and human-readable messages to facilitate remediation without resorting to workarounds that could compromise compliance.

The public interface is limited to non-identifying aggregates and consented disclosures. Aggregates are released only after statistical disclosure controls and re-identification risk assessments described in Chapter 17 are satisfied. Public disclosures are revocable, and withdrawal must cascade within defined service-level objectives across caches, mirrors, and syndication endpoints, as governed by Chapter 15 and Chapter 16. No scraping facilitation or bulk export of public materials shall be enabled beyond what is necessary to fulfil transparency and adoption objectives, and rate limits for public endpoints shall be tuned to protect the rights of data subjects and applicants.

Chapter 5 — Cybersecurity Baselines and Controls

Cybersecurity within the A2074-SRS ecosystem is governed by defence-in-depth, least privilege, resilience by design, and verifiable assurance. Security controls shall be proportionate to risk, yet uniformly enforced across all platform tiers, such that no Validation Partner, vendor, or sub-processor operates below the baseline established herein. Cryptographic protections, secure configuration, continuous monitoring, and independent testing are mandatory. Nothing in this Chapter shall be construed to reduce the primacy of patient-level confidentiality or to derogate from the consent, disclosure, and non-comparative principles articulated across this Manual.

Encryption shall be enforced at rest and in transit using contemporary, well-vetted algorithms and secure operational practices. At rest, keys shall be generated and stored in hardware security modules, with tenant-scoped keys where jurisdictional residency or contractual obligations require cryptographic

separation. Keys shall be rotated on a fixed cadence and upon any suspected compromise, with usage, rotation, destruction, and escrow events recorded immutably in the Key & Certificate Registry and the Audit & Event Ledger. In transit, mutual TLS shall be required for service-to-service integrations, and public endpoints shall use TLS with strong cipher suites and certificate pinning for platform-controlled clients. Detached, signed artifacts shall be used for critical payloads and decisions to support end-to-end non-repudiation pursuant to Chapter 9.

Secure configuration is mandatory for infrastructure, platform services, and application components. Baseline hardening shall disable insecure protocols, enforce least-privilege service accounts, and mandate tamper-evident configuration stores managed exclusively through the control plane and Chapter 13 change procedures. Secrets management shall be centralized, with automatic rotation, strict access policies, and prohibition on embedding secrets in code, container images, or configuration files. Network segmentation shall isolate control plane, data plane, and public presentation tiers, with deny-by-default policies and explicit allow-lists for east-west and north-south traffic. Content security policies, origin restrictions, and anti-tamper protections shall be applied to public-facing assets to mitigate injection and supply chain risks.

Logging, monitoring, and detection shall be comprehensive, immutable, and privacy-preserving. Security-relevant events, including authentication outcomes, access denials, privilege escalations, configuration changes, failed integrity checks, policy evaluation results, and data egress attempts, shall be recorded to the Audit & Event Ledger. Logs must be time-synchronised with monotonic counters, chained by cryptographic hashes, and retained for investigation and adjudication under GSIA oversight. Monitoring shall include anomaly detection, behavioural baselines for privileged accounts, data access pattern analysis, and alerting thresholds commensurate with the sensitivity of the impacted assets. Telemetry collection shall avoid unnecessary personal data and shall be subject to the privacy requirements of Chapter 6.

Vulnerability and patch management shall operate on a continuous cycle. Attack surface management shall include software composition analysis, dependency health monitoring, container image scanning, and infrastructure as code validation. Critical vulnerabilities shall be triaged immediately with compensating controls where patching is not yet feasible; high-severity issues shall be remediated within defined service-level targets. All remediation actions shall be executed through the change management controls in Chapter 13 and recorded for audit. Partners and vendors shall maintain equivalent processes and provide attestations upon request; failure to do so shall constitute grounds for suspension under the Governance & Oversight Manual.

Testing and assurance shall include periodic penetration testing, red team exercises, purple team collaboration for control validation, and static and dynamic application security testing throughout the development lifecycle. Findings shall be risk-rated, tracked to closure, and subject to GSIA review in aggregate. Where national security, state-secrets, or comparable sensitivities are implicated by an Applicant’s evidence environment, bespoke test plans may be authorized by GSIA, provided that such plans do not reduce the effective security level below the controls set herein.

Data loss prevention shall be enforced through technical controls that prohibit direct export of raw evidence, constrain screenshotting within mediated viewers, watermark sensitive displays, and monitor anomalous volumes or destinations for data egress. Cross-tenant and cross-jurisdiction transfers shall be blocked unless explicitly permitted by the Jurisdictional Rules Registry and a valid lawful basis exists per Chapter 6. Tokenisation and format-preserving encryption may be used to enable selective processing without exposing cleartext data to components that do not require it.

Resilience shall be achieved through documented recovery point and recovery time objectives, multi-zone deployment, and tested backup and restore procedures aligned to Chapter 11. Cryptographic key recovery shall be tested without exposing material outside of HSM boundaries. The platform shall degrade gracefully under load, denying non-essential functions while preserving confidentiality and integrity, and shall prioritize withdrawal of public disclosures upon consent revocation even during partial outages.

Supply chain security is mandatory across source code, build pipelines, artifacts, and runtime environments. Build systems shall be isolated, reproducible, and verifiable, with provenance attestations attached to deployable artifacts. Third-party libraries shall be vetted and pinned, with continuous monitoring for malicious insertions or critical vulnerabilities. Container registries shall enforce signature verification and policy compliance before deployment. Vendors shall provide software bills of materials and support rapid incident response coordination under Chapter 12.

Human factors shall be treated as a core control domain. Privileged users shall undergo enhanced security training, periodic re-verification, and simulated phishing resilience exercises. Social engineering playbooks shall be maintained, and high-risk workflows—such as consent publication, key rotation, and schema changes—shall require dual-control and step-up authentication. Violations shall be adjudicated under GSIA ethics procedures, with corrective actions proportionate to intent and harm.

No derogation from the controls in this Chapter is permitted without a documented risk acceptance approved by the A2074 Secretariat and subjected to GSIA review. Temporary exceptions shall be time-boxed, monitored, and remediated with urgency. Any exception that implicates confidentiality, consent enforcement, or the integrity of master registries is presumptively unacceptable.

Chapter 6 — Privacy, Consent, and Jurisdictional Compliance

Privacy, consent, and jurisdictional compliance form the normative foundation of the A2074-SRS digital regime. Records are private by default, disclosure is exceptional and contingent upon explicit, informed, and revocable consent, and processing must be lawful, fair, and proportionate. Nothing herein authorizes comparative publication that could disadvantage micro-enterprises relative to large corporates; proportionality governs both processing and any permitted disclosure.

Lawful bases and purposes shall be specified at the inception of each Engagement and shall be recorded as part of the canonical Engagement entity. Processing for validation, adjudication, audit, and security operations is restricted to the minimum necessary to achieve those defined purposes. Secondary use requires either a compatible purpose assessment documented by the controller or fresh consent recorded in the Consent Ledger. Marketing or public relations usage of results is prohibited without specific, granular consent that can be revoked without detriment to the underlying validation process.

The Consent Ledger is the exclusive mechanism for recording, verifying, and executing consent and revocation events. Consent must be informed, specific, freely given, and documented with scope, duration, and the identities or classes of recipients. Revocation shall be at least as easy as granting and shall trigger prompt withdrawal of disclosures and suppression of downstream displays, syndicated feeds, caches, and partner mirrors in accordance with Chapter 15 and Chapter 16. Consent states are immutable records; transitions are expressed as new states with cryptographic lineage linking them for audit under Chapter 9. No disclosure action may proceed without a real-time consent resolution check by the control plane.

Data minimisation and purpose limitation shall govern schema design, collection practices, and retention schedules. Personal data shall be segregated and pseudonymised wherever feasible. Named individuals shall not be required unless necessary for verification or adjudication, in which case processing shall be confined to the smallest practicable scope and duration. Public disclosures must not contain personal data except where an individual has provided explicit consent for named attribution, which shall be treated as exceptional and revocable. Aggregated analytics released under Chapter 10 must pass de-identification and re-identification risk assessments in accordance with Chapter 17.

Jurisdictional compliance is administered through the Jurisdictional Rules Registry. Controllers and processors shall consult the Registry programmatically before initiating cross-border transfers, creating new processing locations, or engaging sub-processors that alter the data residency footprint. Where transfers to jurisdictions without adequate protections are proposed, the platform shall enforce the applicable safeguards, which may include standard contractual clauses, supplementary technical measures such as encryption with key residency controls, or, where safeguards are not feasible, prohibition of the transfer. Residency constraints shall be enforced at design time through tenant partitioning and at run time through policy evaluation in the control plane.

Rights of the “patient-analogue” data subject shall be implemented consistently across the platform and partner systems. Access, rectification, portability, restriction, and erasure rights shall be available through authenticated, auditable workflows, with identity assurance commensurate with the sensitivity of the records. Where the exercise of rights would impair the integrity of the Consent Ledger or the Audit & Event Ledger, the platform shall implement logical deletion or tombstoning that preserves forensic integrity while suppressing operational use and public display. Requests and outcomes shall be recorded immutably and be reviewable by GSIA upon complaint or appeal. The detailed operationalisation of these rights is set forth in Chapter 14.

Special category data and sensitive evidence require heightened protections. Evidence repositories shall apply sealed workflows, access watermarking, and strict purpose binding. Where evidence contains confidential third-party information or state-level sensitivities, Partners shall employ layered redaction and differential access policies, with independent verification that redactions are sufficient to prevent re-identification or unauthorized disclosure. Any necessity to process children’s data or data of vulnerable populations must be flagged at Engagement creation, subjected to enhanced scrutiny, and approved through GSIA-defined safeguards or prohibited if adequate protection cannot be assured.

Transparency obligations shall be fulfilled without compromising privacy. Applicants must be provided with clear notices describing processing purposes, recipients, retention periods, rights, and complaint channels, including GSIA’s adjudication pathway. Public transparency portals may describe the standard, validation models, and aggregate outcomes without enabling entity-level comparisons unless specifically consented to by the respective entities. Any permitted public registry entries must be promptly withdrawn upon revocation and must avoid cross-linking that would frustrate the right to be forgotten in public contexts.

Accountability requires that controllers and processors maintain records of processing activities, data protection impact assessments where high-risk processing is envisaged, and incident registers with remedial actions taken. GSIA shall have the authority to require submission of such records for ethics review, to order corrective measures, and to suspend participation in the ecosystem for material or

persistent violations. Partners shall designate privacy officers with authority to enforce this Chapter, to coordinate with the A2074 Secretariat, and to liaise with GSIA during investigations.

Where legal demands seek access to platform records, Partners and the Secretariat shall follow a documented process that includes verification of jurisdiction, scope limitation, challenge where lawful and appropriate, and notification to affected parties unless legally prohibited. Production shall be the minimum necessary and shall preserve chain-of-custody guarantees. Any compelled disclosure shall be recorded with sufficient detail to enable GSIA review, without exposing privileged or confidential legal strategy in the record itself.

No practice that coerces consent, conditions essential services on public disclosure, or penalizes revocation is compatible with this Manual. Incentive structures may reward participation in validation, but they may not condition essential access to core benefits or services on publication of results. Any alleged coercion or retaliation shall be investigated under GSIA oversight with remedies proportionate to the harm and may include suppression of disclosures, corrective communications, suspension, or expulsion from the ecosystem.

Chapter 7 — Evidence Repositories and Chain of Custody

Evidence management within the A2074-SRS ecosystem is governed by confidentiality, integrity, provenance, and proportionality. Evidence repositories shall operate as sealed subsystems with narrowly defined interfaces, strict purpose binding, and denial by default for all access not expressly authorized. Evidence in this Chapter is understood to include any artefact, data file, document, transmission, image, recording, log excerpt, attestation, or derived analytic used to support validation activities, adjudication, oversight, or appeals. Operational practices shall not derogate from patient-level confidentiality or the consent and disclosure regime articulated in Chapters 6, 15, and 16.

Evidence ingestion must follow authenticated, role-bounded workflows initiated from the Engagement context. Each artefact shall be registered in the Evidence Repository Index with a globally unique identifier, cryptographic content hash, canonical media type, declared sensitivity, originator identity, creation timestamp, residency zone, and a pointer to the storage locator. Storage shall be content-addressable or otherwise integrity-verifiable. Under no circumstance may raw evidence be routed through public interfaces or included inline in integration payloads, as clarified in Chapter 4. Where a Partner operates a local repository to meet residency requirements, a synchronized index shall be maintained in the canonical Evidence Repository Index, with cryptographic proofs sufficient for GSIA verification without cross-border material transfer.

Chain of custody is the definitive record of provenance for each evidence artefact. It shall be expressed as an append-only, cryptographically linked event series capturing, at minimum, submission, validation, classification, access grants and denials, transformations, redactions, exports (where permitted), and final disposition. Each event shall include the actor identity, role, purpose code, timestamp, location, event hash, and, where applicable, a detached signature. Transformations—including format conversion, hashing algorithm upgrades, redaction, and derivation of analytic summaries—shall preserve the original artefact unaltered and shall create a new evidence version linked to the prior state with an explicit description of the transformation method. No transformation may occur outside a controlled workflow subject to step-up authentication and double-control where the transformation affects disclosure readiness.

To preserve confidentiality, evidence access shall be mediated through viewer services that enforce watermarks, session binding, click-through confidentiality notices, and copy-out restrictions

proportionate to sensitivity. Downloads of raw artefacts shall be disabled by default and permitted only where strictly necessary for adjudication or legal compliance, with time-bound pre-signed retrieval tokens and fine-grained scope recorded in the chain of custody. Screenshots shall be deterred through application-level controls and monitored through heuristics; while technical measures are not absolute, any suspected exfiltration shall raise immediate alerts under Chapter 11 and be subject to GSIA oversight.

Retention and disposition policies shall be proportionate to the validation purpose, legal obligations, and applicant expectations. Minimum necessary retention shall be defined per evidence class, with defaults linked to Engagement and Validation Result lifecycles and maximums stipulated in partner accreditation conditions. Where legal or contractual requirements mandate longer retention, the necessity shall be recorded with a verifiable reference in the chain of custody. Disposition shall be executed through cryptographically verifiable deletion processes appropriate to the storage medium and sensitivity, with tombstones retained in the Index to preserve auditability without operational exposure of the underlying content.

Evidence repositories shall implement locality and residency controls. Artefacts shall not be moved across residency zones absent a lawful basis and a pre-flight decision recorded by the control plane in reliance on the Jurisdictional Rules Registry. Where cross-border verification is required by GSIA, Partners may be compelled to expose proofs, redacted derivatives, or in-situ viewer sessions, provided such exposure meets the minimum necessary standard and does not export raw content in contravention of applicable law or consent.

The following table establishes the canonical event types to be recorded in the chain of custody, with indicative mandatory fields. The vocabulary is exhaustive for canonical processing; Partners may extend with namespaced events provided extensibility does not obscure or weaken canonical events.

Event Type	Mandatory Fields	Purpose Binding	Notes on Controls
submitted	actor_id; role; timestamp; content_hash; media_type; origin	Validation purpose	Origin check, malware scanning at ingress; quarantine until validated
validated	validator_id; timestamp; hash_verified; classification	Integrity and classification	Confirms hash and assigns sensitivity level and residency zone
access_granted	grantor_id; grantee_id; role; scope; expiry; timestamp	Least privilege	Step-up required for sensitive classes; consent verification where applicable
access_denied	reviewer_id; subject_id; reason_code; timestamp	Denial by default	Captures failed attempts and rationale; triggers monitoring
viewed	viewer_id; session_id; timestamp; purpose_code	Transparency and deterrence	Watermarking enforced; rate and volume limits monitored

transformed	operator_id; method_ref; input_hash; output_hash; timestamp	Controlled derivations	Redaction, format conversion, hashing upgrade; creates version link
exported	approver_id; recipient_ref; scope; lawful_basis; timestamp	Exceptional cases	Requires legal basis; pre-signed tokens; minimal scope
transferred	approver_id; new_residency; safeguards_ref; timestamp	Residency governance	Jurisdictional pre-flight recorded; safeguards applied
disposed	operator_id; method_ref; timestamp; proof_ref	Lifecycle closure	Cryptographic deletion or secure wipe; tombstone retained

Evidence repositories shall be continuously monitored for integrity and availability. Tamper-evident storage, write-once or append-only controls for indexes, and periodic re-hashing for bit-rot detection are mandatory. Where deduplication is used, it must not introduce cross-applicant correlability or inference risks; deduplication domains shall be strictly tenant-scoped unless GSIA authorizes broader scopes for specific security purposes.

Any compromise, suspected tampering, or unexplained inconsistency in the chain of custody or content integrity shall trigger incident response under Chapter 11 with immediate containment, notification to the A2074 Secretariat and GSIA, and suspension of affected disclosure pipelines until integrity is re-established. Findings shall be recorded for adjudication and learning without exposing sensitive content beyond those with a verified need-to-know.

Chapter 8 — Decision Support and AI Guardrails

Decision support tools, including artificial intelligence and statistical models, may be employed to enhance triage, consistency, and efficiency in the A2074-SRS ecosystem, provided that human oversight remains decisive, explainability is assured to an appropriate standard, and bias and fairness controls are demonstrably effective. No automated system may issue binding validation decisions, adjudications, or sanctions. Algorithmic outputs are advisory and shall be treated as inputs to human deliberation documented in the Audit & Event Ledger.

Model governance shall follow a full lifecycle regime encompassing design, training, validation, deployment, monitoring, and retirement. Use cases shall be catalogued with explicit problem statements, data sources, lawful bases, consent dependencies, and potential risks to privacy, fairness, and proportionality. Training and evaluation datasets must be documented with provenance, curation methods, representativeness assessments, and known limitations. Sensitive categories and protected characteristics shall be excluded from model features unless their inclusion is necessary to mitigate bias and is explicitly authorized with appropriate safeguards. Where synthetic data is used for development or testing, its generation method, limitations, and separation from production shall be recorded explicitly.

Explainability requirements shall be commensurate with impact. For triage and routing, post-hoc explanation techniques or transparent rule-based systems are acceptable. For risk scoring and recommendation engines that may influence outcomes, inherently interpretable models or robust, validated explanation methods are required, together with documentation intelligible to assessors,

applicants, and GSIA. Explanations shall identify salient factors, their directional influence, and uncertainty bounds, and shall avoid disclosure of sensitive personal data or proprietary content beyond what is necessary for accountability.

Bias management shall be continuous and evidence-based. Pre-deployment fairness evaluations shall assess disparate impact, calibration, and error rates across relevant groups, including size bands to preserve proportionality between micro-enterprises and large corporates. Thresholds for acceptable disparities shall be defined in policy, with mitigation strategies such as reweighting, constraint optimization, post-processing, or human review triggers. Monitoring in production shall track drift, performance degradation, and emergent biases, with automated alerts and rollback plans. Any sustained disparity exceeding thresholds shall trigger a corrective action plan and, where warranted, suspension of the affected model pending GSIA review.

Data protection and confidentiality are paramount. Feature engineering shall respect data minimisation; personal data shall be pseudonymised, and linkage to identifiable records shall occur only within the control plane and only where strictly necessary and consented. Training and inference pipelines shall operate within secure enclaves or equivalent protections, with strict access controls, audit logging, and prohibition on exporting training data to environments outside residency constraints. Models and their artefacts (weights, parameters, metadata) shall be treated as sensitive assets, versioned, signed, and stored under the Key & Certificate Registry with integrity verification before deployment.

Human oversight shall be structured and documented. Assessors shall receive model outputs with confidence measures, explanation artefacts, and clear guidance on when to defer, escalate, or override. Overrides shall be recorded with rationale, and systematic patterns of overrides shall feed back into model improvement or policy adjustment. No assessor shall be penalized for good-faith overrides based on contextual knowledge or ethical concerns; such practice is integral to preserving non-comparative fairness and proportionality.

The following control catalogue applies to all AI and decision-support deployments and is enforceable by the A2074 Secretariat and GSIA through accreditation and oversight.

Control Domain	**Mandatory Controls**	**Evidence of Compliance**
Use-Case Registration	Catalogue entry with purpose, lawful basis, consent dependencies, impact level	Registered record; GSIA-reviewable summary
Data Governance	Provenance logs; minimisation; residency compliance; pseudonymisation	Data sheets; lineage in Audit & Event Ledger
Model Development	Versioned code and artefacts; reproducible training; bias testing	Signed artefact manifests; fairness reports
Explainability	Model cards; explanation artefacts; intelligible documentation	Human-readable briefs; sample explanations
Security & Access	Enclave or equivalent; key management; access logging	Security architecture; access logs

Monitoring & Drift	Performance dashboards; drift detectors; alerting	Monitoring snapshots; incident tickets
Human Oversight	Override workflows; escalation paths; training	Training records; override analytics
Change Control	Controlled promotion; rollback plans; deprecation	Change tickets; rollback tests
Auditability	Immutable logs; decision traceability; dataset references	Audit trail extracts; GSIA inspection readiness

Vendors and third-party model providers shall be held to equivalent standards. Black-box models may be used only if explanation sufficiency, bias controls, and security assurances can be independently verified by the Partner and, upon request, by GSIA. Where such verification is not feasible, the model shall not be used for any function that materially influences validation outcomes or public disclosures.

No deployment may present or imply “scores” or rankings for public consumption without explicit, informed, and revocable consent, and even then only within the bounds of proportionality and non-comparative principles. Public-facing analytics under Chapter 10 shall rely on anonymised aggregates, with statistical disclosure controls under Chapter 17 and clear separation from any internal risk or triage scores.

Incident response for AI systems shall be integrated into Chapter 11. Model-specific incidents include harmful recommendations, unfair disparate impact exceeding thresholds, data leakage through training artefacts, and adversarial manipulation. Detection shall trigger immediate containment, notification to the A2074 Secretariat and GSIA, suspension of affected functionalities, and a post-incident review to adjust datasets, features, thresholds, or deployment policies. Recurrence without adequate remediation constitutes grounds for suspension of Partner accreditation in the relevant scope.

Nothing in this Chapter authorizes automated adverse action against any applicant or publication of any AI-generated classification or score. All outputs remain advisory and subject to human judgment, and all operations remain subordinate to the confidentiality and consent architecture established elsewhere in this Manual.

Chapter 9 — Audit Trails and Non-Repudiation

Auditability and non-repudiation are foundational to the integrity of the A2074-SRS ecosystem. All material actions, decisions, configurations, and data state changes shall be durably recorded in an immutable audit substrate and rendered traceable to authenticated actors, time references, and cryptographic attestations. The objective is to ensure that any contested state can be reconstructed, verified, and adjudicated under GSIA oversight without reliance on unverifiable testimony or volatile system memory.

Audit events shall be complete, consistent, and privacy-preserving. Completeness requires that every material action—successful or denied—be recorded with sufficient context to support reconstruction of intent, scope, and effect. Consistency requires that event schemas be stable and version-controlled, that time is monotonic and synchronized, and that event sequences are causally ordered. Privacy-preserving design requires that audit payloads avoid unnecessary personal data and that

sensitive elements be tokenised or hashed in a manner that remains useful for correlation while preventing re-identification outside of authorized workflows.

Time shall be authoritative and verifiable. Systems shall use authenticated time sources and maintain monotonic counters at the control plane. Each audit event shall carry both a wall-clock timestamp and a monotonic sequence number scoped to its log stream, with cross-stream correlation supported by cryptographic hash-linking. Any detected time skew beyond defined tolerances shall be treated as a security incident under Chapter 11.

Non-repudiation is achieved through layered cryptographic controls. Each event record shall be hashed and linked to its predecessor (or to a block header) to form an append-only hash chain resistant to tampering. Critical events—such as consent creation or revocation, disclosure publication or withdrawal, model promotion, schema change, key lifecycle operations, and break-glass usage—shall additionally carry detached digital signatures bound to the initiating actor and to a platform key. Periodic anchor hashes of audit segments shall be sealed to the Key & Certificate Registry and independently attested, enabling external verification without exposing event content.

The canonical audit taxonomy shall encompass, at a minimum, the following event classes, each with mandatory attributes. The taxonomy is exhaustive for canonical purposes; Partners may extend with namespaced classes that do not dilute or obscure the canonical set.

Event Class	Illustrative Sub-Types	Mandatory Attributes	Notes
Identity & Access	login_success, login_failure, mfa_stepup, role_assignment, role_revocation, federation_change	actor_id; actor_role; outcome; target_ref (if applicable); timestamp; monotonic_seq; source_ip/device; event_hash; prev_hash	Step-up prompts and grants shall be explicitly marked; failures retained for detection and adjudication
Consent & Disclosure	consent_created, consent_revoked, disclosure_published, disclosure_withdrawn	actor_id; subject_token_ref; consent_id; disclosure_id (if applicable); scope; timestamp; signatures[]	Detached signatures required; revocation cascades must reference orchestration outcomes
Data & Evidence	engagement_created, result_proposed, result_approved, evidence_ingested, evidence_viewed, evidence_exported	actor_id; target_ref; purpose_code; residency_zone; timestamp; integrity_hash; event_hash	Evidence export is exceptional and must include lawful basis reference
Configuration & Schema	schema_updated, feature_toggle_changed, policy_updated, enumeration_refreshed	actor_id; change_ticket; old_value_ref; new_value_ref; timestamp; approvals[]	Subject to Chapter 13 dual-control and change windows

Cryptographic Lifecycle	key_generated, key_rotated, key_retired, certificate_issued	operator_id; key_id; algorithm; purpose; timestamp; HSM_ref; event_hash	Inventory synchronized with the Key & Certificate Registry
Security & Operations	vuln_triaged, patch_applied, incident_opened, incident_contained, breakglass_used	operator_id; severity; asset_ref; timestamp; case_id; outcome	Break-glass requires immediate GSIA notification; see Chapter 11
Jurisdiction & Transfer	residency_set, crossborder_request, transfer_permitted, transfer_blocked	actor_id; jurisdiction_ref; lawful_basis; safeguards_ref; timestamp	Pre-flight checks tied to the Jurisdictional Rules Registry

Storage and retention of audit data shall be designed for durability, confidentiality, and verifiability. Audit streams shall be written to append-only stores with integrity verification, replicated across failure domains, and backed up in a manner that preserves chain properties. Access to raw audit data shall be tightly restricted; GSIA shall have read access sufficient to discharge its oversight mandate, with case-bounded elevation where necessary. Redacted views for Partner administrators and Applicants may be provided where appropriate, provided that redaction does not impair GSIA's ability to reconstruct events for adjudication.

Correlability and traceability shall enable end-to-end decision reconstruction. A binding between audit events and canonical entities shall be maintained through stable references, enabling a reviewer to traverse from a public disclosure back to the originating consent, validation result, engagement, evidence items, and identity and configuration events that influenced the outcome. Decision trails for approvals shall include the identity of approvers, their roles, the policies evaluated, the explanation artifacts (where AI was consulted), and any dissenting notes or overrides.

Tamper detection and response are mandatory. Any alteration attempt, chain break, time regression, or unexpected hash mismatch shall raise an immediate alert, trigger containment procedures, and result in a GSIA-notified case. Post-incident recovery shall include verification of unaffected segments through independent anchor checks and attestations. Where an audit segment is deemed unreliable, it shall be quarantined and preserved for forensic analysis, and dependent processes—especially publication and consent revocation flows—shall be paused until integrity is re-established.

Export and disclosure of audit data shall be tightly controlled. External disclosures—for example, to regulators or courts—shall be the minimum necessary and shall preserve cryptographic verifiability. Any compelled disclosure shall be recorded with scope and legal basis, without revealing privileged content in the audit record. Public transparency does not extend to audit content; only high-level, anonymised operational statistics may be surfaced, consistent with Chapter 10 and Chapter 17.

No operation that affects consent states, disclosures, cryptographic keys, registry schemas, or validation outcomes shall be accepted into the platform unless its corresponding audit write succeeds. In the event of downstream audit storage latency, the platform shall queue and confirm only upon durable write acknowledgement. This coupling preserves non-repudiation and ensures that state transitions are inseparable from their audit provenance.

Chapter 10 — Analytics and Public Dashboards

Analytics and public dashboards exist to advance the adoption and credibility of A2074-SRS while preserving confidentiality, proportionality, and the primacy of consent. Analytics shall be designed to deliver aggregated, anonymised insights that support learning, policy development, and ecosystem transparency without enabling re-identification or comparative harms. Public dashboards are instruments of responsible transparency; they shall never reveal entity-level results absent explicit, informed, and revocable consent, and they shall avoid league tables or constructs that undermine non-comparative evaluation principles.

Data sources for analytics shall be limited to canonical registries and sanctioned aggregates. Raw evidence and personal data are excluded from analytical datasets, except where de-identified, risk-assessed, and authorized for internal quality improvement under strict controls. Analytical pipelines shall operate within the jurisdictional constraints recorded in the Jurisdictional Rules Registry and shall respect the minimum disclosure by default principle. All analytical processes shall be reproducible, version-controlled, and recorded as AuditEvents with dataset lineage and methodological references.

Aggregation and de-identification standards shall be proportionate to risk and audience. Minimum cell counts, k-anonymity thresholds, l-diversity where appropriate, and noise injection or differential privacy methods may be employed to reduce re-identification risk. Suppression rules shall be deterministic and documented in public methodology notes to avoid inference through release variance. Where the risk of re-identification cannot be sufficiently mitigated, the measure shall not be published.

Publication governance shall follow a structured release workflow with GSIA oversight. The following table establishes the canonical steps of the analytics and dashboard release process and the control objectives at each step.

Release Step	**Control Objective**	**Mandatory Outputs**	**Gate**
Scoping & Design	Define purpose, audience, measures, and risk posture	Scoping brief; list of candidate metrics; jurisdictional review	A2074 Secretariat approval
Data Preparation	Build de-identified dataset with lineage	Dataset manifest; de-identification report; residency confirmation	Partner data steward + Secretariat co-sign
Risk Assessment	Evaluate re-identification risk and proportionality	Risk assessment report; suppression/noise plan	GSIA ethics desk review
Draft Visualisation	Create preliminary dashboards and narratives	Draft visuals; methodology notes; accessibility check	Secretariat analytics lead
Quality & Bias Review	Assess for misleading or comparative bias	Bias review memo; alternative presentations if needed	GSIA review and conditions

Consent Cross-Check	Verify that any entity-level disclosures are consented	Consent resolution report; display scopes	Control plane automated check
Publication	Release to public endpoints with rate limits	Release package; version tag; release notes	Secretariat publish authority
Post-Release Monitoring	Monitor traffic, anomalies, and complaints	Monitoring snapshots; feedback log	Secretariat + GSIA oversight
Withdrawal/Update	Execute revocation or update as needed	Withdrawal log; updated release notes	Control plane orchestration

Dashboards shall present measures that are constructive, context-aware, and resistant to misuse. Suitable measures include ecosystem-level adoption counts by sector and jurisdiction, distribution of validation models in use, average time to validation, frequency of consent revocations and publication withdrawals (expressed without identifying entities), and aggregated performance against SGG pillars at regional or sectoral levels, provided that the re-identification risks are addressed. Visuals shall avoid rank-ordering unless all entities displayed have explicitly consented to comparative publication. Confidence intervals or data quality indicators shall be included where appropriate to prevent over-interpretation.

Technical delivery shall prioritise accessibility, security, and revocability. Public endpoints shall be read-only, rate-limited, and instrumented for anomaly detection. Caches, mirrors, content delivery networks, and syndication channels shall be enlisted with contractual and technical means to honour withdrawal instructions promptly pursuant to Chapter 15 and Chapter 16. Where dashboards embed third-party scripts or analytics, only privacy-respecting, first-party controlled tools may be used; no tracking or cookies shall be set beyond what is strictly necessary for security and performance.

Methodological transparency shall be maintained through publicly accessible documentation that explains data sources, transformation steps, de-identification methods, limitations, and release versioning. Each dashboard release shall be tagged and archived with release notes that are retrievable by the public to ensure interpretability over time. Historical series shall be preserved only if they do not frustrate revocation; where a revocation affects a prior series, the series shall be revised or suppressed, with a visible note indicating the reason and date of change.

Commercial and communications usage of analytics shall conform to proportionality and fairness. Materials that reference public dashboards must avoid implying endorsements or certifications beyond what is factually supported. Partners and Applicants may cite aggregates to illustrate sectoral trends, but they may not infer entity-level standing absent consented disclosure. Where misuse or misrepresentation is detected, the A2074 Secretariat shall issue corrective communications and may suspend access to syndication feeds.

Internal analytics for quality improvement are permitted within strict confines. They must use de-identified datasets, be accessible only to authorized personnel, and be segregated from publication pipelines. Findings that could materially affect fairness, bias profiles, or system integrity shall be

escalated to GSIA for review and, where appropriate, publication as ecosystem learning without exposing sensitive or identifying details.

No analytic or dashboard artefact may be released if it would undermine patient-level confidentiality, contravene consent, or enable comparative harm to micro-enterprises. In all cases of doubt, the presumption shall favour non-publication and referral to GSIA for determination. The public value of transparency shall be advanced without compromising the rights and interests protected by this Manual.

Chapter 11 — Incident Response and Business Continuity

Incident response and business continuity within the A2074-SRS ecosystem are governed by preparedness, proportionality, rapid containment, and verifiable restoration, under the supervisory jurisdiction of GSIA. The objective is to safeguard patient-level confidentiality, preserve the integrity of master registries and audit trails, and maintain essential services for Validation Partners and Applicants, while ensuring transparent, accountable handling of adverse events.

Incidents are any events that compromise, or reasonably threaten to compromise, confidentiality, integrity, availability, or lawful processing. Categories include security incidents (such as credential compromise, attempted or actual data exfiltration, privilege escalation, malware activity, or supply-chain corruption), privacy incidents (such as unauthorized access or disclosure of personal data, consent misapplication, or unlawful cross-border transfer), operational incidents (such as outages, data corruption, or queue backlogs affecting publication or revocation), and integrity incidents (such as audit chain breaks or evidence tampering). All incidents shall be handled through a unified case management process with unique case identifiers, severity classification, and traceable actions recorded in the Audit & Event Ledger.

Preparation is mandatory. The A2074 Secretariat shall maintain a current Incident Response Plan, a Business Continuity Plan, and a Disaster Recovery Plan aligned to defined Recovery Time Objectives and Recovery Point Objectives for each critical function, with explicit prioritization of consent revocation orchestration, disclosure withdrawal, control plane policy evaluation, and GSIA access. Validation Partners and vendors shall maintain congruent plans compatible with platform requirements. Contacts, on-call rosters, authority matrices, and communication templates shall be maintained and tested. Table-top exercises and technical simulations shall be conducted on a defined cadence, with lessons learned documented and integrated into control improvements through Change Management under Chapter 13.

Detection and triage shall be prompt and evidence-based. Alerts from security monitoring, anomaly detection, audit chain verifiers, consent orchestration failures, and user reports shall feed a central intake. Each event shall be assessed for scope, affected assets, potential impact on confidentiality and consent, jurisdictional implications, and likelihood of propagation. Severity shall be assigned using a documented rubric that favors caution where consent, audit integrity, or evidence repositories are implicated. False positives shall be closed with rationale; near misses shall be recorded for trend analysis.

Containment and eradication actions shall be proportionate, rapid, and auditable. Short-term containment may include account lockdown, token revocation, key rotation, network segmentation, service throttling, suspension of publication or revocation pipelines, and disabling of non-essential integrations. Eradication shall remove malicious artifacts, patch exploited vulnerabilities, and restore secure configurations through controlled changes. All actions shall be executed under dual-control and

step-up authentication when they affect cryptographic keys, registry schemas, consent states, or disclosures. Break-glass mechanisms may be invoked solely to preserve availability or to execute urgent revocation or withdrawal; invocation shall be time-boxed, tightly scoped, immediately alerted to GSIA, and subjected to post-use review.

Forensic integrity is a non-derogable requirement. Volatile evidence shall be captured where appropriate; logs, audit trails, and chain-of-custody records shall be preserved, hashed, and sealed. Access to forensic data shall be restricted to the minimum necessary, with handling recorded in the Audit & Event Ledger. Where law enforcement or regulatory engagement is necessary, disclosures shall be the minimum necessary and shall not compromise privileged materials or the rights of data subjects except as required by law. Any compelled production shall be recorded with legal basis, scope, and date.

Communication and notification obligations shall be fulfilled in a timely, accurate, and non-speculative manner. Internal notifications shall inform affected teams and leadership on a need-to-know basis, enabling coordinated response without unnecessary dissemination of sensitive details. External notifications shall be made to affected Partners, Applicants, vendors, and, where applicable, regulators and data subjects, taking into account jurisdictional timing and content requirements. GSIA shall be notified without undue delay for incidents that materially affect confidentiality, consent enforcement, audit integrity, or public trust; GSIA may require interim updates and corrective directives during the response.

Business continuity shall prioritize functions essential to uphold rights and prevent harm. The following restoration sequence shall be used absent case-specific deviation: control plane policy evaluation and authentication; consent ledger write and revocation orchestration; disclosure withdrawal pathways (public endpoints, caches, mirrors, syndication feeds); GSIA investigative and read access; Engagement creation and evidence index registration; Validation Result submission and approval; analytics pipelines. Degraded modes shall explicitly prefer non-publication and revocation capacity over publication or non-essential reads. Where partial service is possible, the platform shall communicate service limitations clearly to Partners and Applicants.

Disaster recovery shall achieve restoration to pre-incident integrity within defined objectives. Backups shall be encrypted, integrity-checked, and geographically redundant within residency constraints. Restoration shall include verification that master registries, consent states, audit chains, and evidence indices are intact and consistent. Any inconsistency shall be investigated before resuming publication or cross-border transfers. Post-restoration validation shall be recorded with anchor hashes and attestation artifacts.

Learning and improvement are mandatory. Each major incident shall undergo a post-incident review with root-cause analysis, contributing factor analysis (technical, human, process), and a corrective and preventive action plan with accountable owners and deadlines. Findings that bear ecosystem-wide relevance shall be summarized for GSIA and, where appropriate, shared with Partners as advisories without exposing sensitive details. Recurring incidents or failure to implement corrective actions within agreed timelines constitute grounds for sanction under GSIA oversight.

No practice that obscures, delays, or minimizes material incidents is compatible with this Manual. Good-faith reporting, rapid containment, and transparent remediation are mitigating factors in any adjudication. Willful concealment, destruction of evidence, or retaliation against whistleblowers shall trigger escalated remedies, including suspension or expulsion from the ecosystem.

Chapter 12 — Vendor Management and Third-Party Risk

Vendors and third parties are integral to the operation of the A2074-SRS ecosystem and shall be governed through rigorous due diligence, contractual controls, continuous assurance, and proportional oversight. No vendor engagement may result in a reduction of the security, privacy, or governance posture below the baselines established herein. The A2074 Secretariat is responsible for maintaining the Vendor & SLA Registry as the authoritative catalogue of vendors, services, attestations, risk ratings, and monitoring status, with GSIA oversight over ethics and compliance matters.

Vendor classification shall be risk-based and service-specific. Critical vendors are those that operate or materially influence control plane services, master registries, identity and key management, consent ledgering, evidence repositories, audit substrates, or public disclosure pathways. High-impact vendors process or store personal data, confidential evidence, or cryptographic materials, or affect availability of essential services. Standard vendors provide ancillary services without access to sensitive data or critical paths. Classification shall dictate the depth of due diligence, contractual safeguards, and monitoring cadence.

Due diligence shall assess legal, financial, technical, and ethical fitness to operate within the ecosystem. Minimum requirements include corporate identity verification, jurisdictional footprint mapping, ownership and control checks sufficient to identify sanctions and conflict-of-interest risks, security and privacy posture evaluation, software supply chain controls, data residency capabilities, incident history, and references. For critical and high-impact vendors, independent attestations (such as security certifications or SOC-style reports), secure development lifecycle evidence, penetration testing summaries, and software bills of materials shall be required. Where a vendor relies on sub-processors, the vendor shall disclose and maintain an accurate sub-processor list, and shall flow down equivalent controls and audit rights.

Contracts shall embody the operational controls of this Manual. Mandatory clauses include scope and purpose limitation; confidentiality; data protection with residency and transfer restrictions aligned to the Jurisdictional Rules Registry; consent and non-publication obligations; security baselines; cryptographic key handling; segregation of duties; vulnerability and patch management expectations; logging and auditability; incident notification timelines and cooperation duties; rights to audit and to obtain independent assurance; subcontracting restrictions; termination assistance; data return and deletion with verifiable proofs; and remedies including suspension and termination for cause. For critical vendors, step-in rights or contingency arrangements shall be defined to preserve essential services during vendor failure.

Onboarding shall include technical and procedural integration steps. Vendors shall register keys and certificates, integration endpoints, and contact points in the Vendor & SLA Registry; complete security and privacy training aligned to their role; and pass connectivity, authentication, and policy evaluation tests in a sandbox environment. Data flows shall be documented, minimized, and instrumented with monitoring and alerting. Any data exchange shall use the APIs and standards in Chapter 4 and shall comply with the data model and schemas defined in Chapter 3; no vendor shall be permitted to bypass canonical interfaces or to maintain shadow registries for canonical entities.

Continuous monitoring shall sustain assurance over time. The A2074 Secretariat shall track SLA performance, incident reports, vulnerability disclosures, patch compliance, and material changes in ownership, location, or sub-processors. Critical and high-impact vendors shall provide periodic attestations and make available redacted results of relevant security tests. The platform shall

instrument vendor interfaces for anomaly detection, rate limiting, and abuse prevention. Risk ratings shall be recalibrated upon new information, with corrective action plans issued where deviations from baseline are identified. Persistent non-compliance shall trigger escalation to GSIA and may result in suspension or replacement.

Third-party software and components integrated into platform services shall be treated as vendor risk. Software provenance shall be verified; components shall be pinned, monitored, and patched; and build pipelines shall enforce signature verification and policy checks. Where open-source components are used, licensing, maintenance status, and community health shall be considered. Known vulnerabilities shall be triaged and remediated according to the severity timelines defined in Chapter 5, with temporary compensating controls documented and approved through Change Management under Chapter 13.

Data protection and confidentiality obligations extend fully to vendors. Vendors shall process personal data and evidence solely for documented purposes, employ pseudonymisation and minimisation, and comply with consent and revocation orchestration. Vendors with access to evidence or personal data shall implement sealed workflows, viewer-based access, watermarking where applicable, and chain-of-custody logging compatible with Chapter 7. Cross-border transfers by vendors shall be pre-authorized through the control plane with residency enforcement and recorded decisions, and shall be prohibited where lawful safeguards are not attainable.

Incident management duties for vendors shall mirror the platform's standards. Vendors shall notify the A2074 Secretariat without undue delay of any incident that materially affects confidentiality, integrity, availability, consent, or auditability, and shall cooperate in containment, forensics, and remediation. Vendors shall not communicate externally about incidents that implicate the A2074-SRS ecosystem without coordination with the Secretariat, except where legally mandated; parallel regulatory notifications shall be harmonised to avoid conflicting disclosures. Failure to notify or cooperate constitutes a material breach.

Exit and termination procedures shall preserve continuity and rights. Upon termination or at contract end, vendors shall provide data export in canonical formats, support transition to successor providers or internal services, and execute verified deletion of any residual data, backups, and derived artifacts not required by law to be retained. Proofs of deletion shall be provided and recorded in the Vendor & SLA Registry. Where immediate termination is required to protect confidentiality or consent, contingency arrangements shall be invoked to maintain essential services while severing the vendor connection.

Ethical alignment is required. Vendors shall agree to the Ethics & Integrity Code and to GSIA oversight within their engagement scope. Activities that undermine the non-comparative, proportional, and privacy-first principles of this ecosystem are prohibited. GSIA may require remedial measures or disallow vendor participation where conflicts with these principles are identified. The Secretariat shall publish guidance on acceptable vendor practices and maintain a watchlist for ecosystem participants.

No vendor arrangement may dilute the obligations or accountability of controllers and processors under this Manual. Controllers remain responsible for vendor conduct within the scope of delegated processing. Oversight is continuous, evidence-based, and enforceable through contractual and accreditation mechanisms, with GSIA holding ultimate adjudicative authority over disputes and material non-compliance.

Chapter 13 — Change Management and Version Control

Change management is the formal mechanism through which the platform evolves without compromising confidentiality, consent enforcement, auditability, or partner interoperability. Changes are permitted only where they preserve or enhance security and privacy baselines, maintain backward compatibility within defined windows, and are fully observable, auditable, and reversible. All changes are subject to the authority of the A2074 Secretariat, with GSIA empowered to review, condition, or veto any change whose risk profile threatens patient-level confidentiality, proportionality, or ethical compliance. No derogation from Chapters 5, 6, 9, 15, or 16 is permitted by virtue of any change.

Change scope includes application code, APIs, schemas and enumerations, model definitions, security configurations, infrastructure as code, cryptographic materials, partner adapters, analytics pipelines, and public dashboard artefacts. Configuration is treated as code; all material configuration resides in tamper-evident, version-controlled stores accessed exclusively through the control plane. Build and release pipelines shall be reproducible and attestable, with provenance and signatures attached to all deployable artefacts, and with software bills of materials available for inspection under Chapter 12.

Changes are proposed through signed tickets containing purpose, risk assessment, dependency mapping, residency implications, consent and disclosure impact analysis, rollback strategy, and validation plan. Tickets are assessed by a Change Authority that includes security, privacy, platform engineering, and, where impact is material, GSIA ethics review. Dual-control and step-up authentication are mandatory for approvals affecting cryptographic keys, consent or disclosure logic, master registry schemas, or audit substrates. Where changes interact with incident remediation, the incident case identifier shall be referenced, and post-change verification shall be part of the case closure under Chapter 11.

Versioning follows a semantic discipline at the schema, API, and model layers. Additive and non-breaking changes are introduced as minor versions and shall be announced with release notes and deprecation timers where relevant. Breaking changes require a major version increment and the concurrent operation of old and new versions for a reasonable migration window. Partners shall receive sandbox access, migration guides, and test datasets. Forced migrations are prohibited without an objectively reasonable window that accounts for partner tier and jurisdictional constraints. Enumerations shall be extended additively; redefinition or removal of values requires a major version bump and an automated translation layer during the migration window.

Deployment proceeds through a controlled pipeline: development, integration, staging, canary, and production. Canary exposure is mandatory for material changes, with automated reversion if error budgets, anomaly thresholds, or consent/disclosure checks regress. Blue-green or rolling strategies are used to avoid downtime. Under no circumstances may a deployment proceed to production if audit writes would be impaired or if consent revocation orchestration would be degraded; priority of revocation over publication must be preserved during all change windows.

Data model and schema migrations are designed for zero- or near-zero-downtime. Techniques include additive schema evolution, dual-write and dual-read strategies, backfills that preserve referential integrity, and feature flags that gate new behavior without exposing sensitive data. Flags are never permitted to bypass consent, privacy, or GSIA oversight logic. Flags shall be scoped, time-boxed, and observable; any flag affecting disclosures or evidence handling requires Change Authority approval and explicit rollback parameters.

The following table sets out canonical change categories, authorization requirements, indicative lead times, and rollback obligations. The table is illustrative and does not dilute the general obligations in this Chapter.

Change Category	Examples	Minimum Authorization	Indicative Lead Time	Rollback Requirement	Notes
Emergency Security Fix	Critical vulnerability patch; key rotation for suspected compromise	Security lead + Platform lead; notification to GSIA; post-hoc CAB review	As needed (hours)	Immediate, automated fallback if instability detected	Incident-linked; audit coupling mandatory
Privacy/Consent Logic	Consent evaluation changes; revocation orchestration updates	CAB approval; GSIA review; dual-control	10 business days	Tested rollback with data consistency checks	No reduction in rights or speed of withdrawal
API Non-Breaking	Additive fields; new read endpoints	Platform lead; privacy review	5 business days	Revert to prior minor version	Deprecation notice if future break implied
API Breaking	Field removal; semantics change; auth scope change	CAB approval; GSIA review; partner sign-off plan	30–90 days migration window	Parallel stacks; roll-backable routing	Sandbox and migration aids required
Schema/Enumeration	New entities/values; value retirement	Platform + Data governance; privacy review	10–20 business days	Backfill and mapping; reversible transforms	Enumerations retirements need major version
Model Catalogue	New validation model or major revision	Secretariat approval; GSIA review	15–30 business days	Model freeze and revert path	Multi-Model Framework alignment
Analytics/Dashboards	New measures; method changes	Analytics lead; GSIA ethics desk	10 business days	Withdraw/replace release; public notes	Chapter 10 controls apply

Infrastructure & Network	Segment rules; control plane scaling	Platform SRE; security review	5–10 business days	Slot-by-slot rollback	No impact to audit or consent pathways
Vendor Adapter	Third-party integration change	Vendor manager; security + privacy review	10–20 business days	Disable adapter; quarantine queues	Chapter 12 alignment required

Communication is integral to change safety. Partners shall receive timely notifications tailored to their integrations and jurisdictions, including effective dates, impact summaries, deprecation schedules, and required actions. Release notes shall be clear, versioned, and permanently accessible. Public changes that affect dashboards or disclosures shall include methodology notes and revision histories to preserve interpretability, with suppression or revision where consent withdrawals necessitate change, in line with Chapters 10, 15, and 16.

Change freezes may be declared by the Secretariat when ecosystem risk is elevated or during high-stakes windows, such as widespread publication cycles or mass revocation events. During a freeze, only emergency security fixes and consent-related corrections may proceed. Any attempted bypass of a freeze constitutes a governance breach subject to GSIA adjudication.

All change activities, approvals, deployments, and rollbacks shall be recorded in the Audit & Event Ledger with immutable timestamps, actor identities, artefact hashes, and outcome status. No state transition implicated by a change is valid unless its corresponding audit write is durably committed, preserving the non-repudiation guarantees of Chapter 9.

Chapter 14 — Patient-Analogue Data Rights

The ecosystem guarantees enforceable rights to the “patient-analogue” data subject and, where applicable, to applicant organizations in respect of their records, in a manner that is lawful, fair, proportionate, and compatible with the confidentiality and consent architecture. These rights include access, rectification, portability, restriction, and erasure, together with objection to incompatible secondary use and the right to withdraw consent. No automated adverse action is taken within A2074-SRS; algorithmic outputs are advisory, and the right not to be subject to solely automated decisions with legal or similarly significant effects is structurally preserved by Chapter 8.

Rights requests shall be authenticated to assurance levels commensurate with sensitivity and jurisdiction. Identity assurance is heightened where personal data or sensitive evidence may be exposed. Requests must specify scope, including the Engagements, periods, and types of records sought. Controllers shall record requests, verification steps, decisions, and fulfilment artifacts in the Audit & Event Ledger. Acknowledgement shall be prompt, and responses shall be provided within timelines set by the Jurisdictional Rules Registry, with the platform’s internal target being acknowledgement within seven calendar days and substantive response within thirty calendar days unless applicable law requires a shorter period.

Scope and modality are bounded by proportionality and confidentiality. Access to raw evidence is exceptional; where feasible, the platform shall provide structured summaries, metadata, and

viewer-based access with watermarking in lieu of bulk transfers. Where records contain third-party personal data or confidential content, redaction or layered access shall be applied. Where access would compromise the integrity of the Audit & Event Ledger or the Consent Ledger, the platform shall provide an intelligible account of the relevant records without altering immutable substrates, using tombstoning or logical suppression to reconcile rights with forensic guarantees.

The following table maps principal rights to their operational mechanisms, verification levels, and typical timelines. It establishes minimum guarantees without limiting stricter jurisdictional requirements.

Right	**Operational Mechanism**	**Verification Level**	**Typical Timeline**	**Notes**
Access	Authenticated portal workflow; structured extracts; viewer session for sensitive evidence	High for personal data; Medium for applicant-level records	Acknowledge ≤ 7 days; respond ≤ 30 days	Redaction for third-party data; audit log extract by reference
Rectification	Append corrective record; supersede displays; maintain immutable prior state	Medium–High	≤ 30 days	No destructive edits to audit or consent ledgers; public disclosures updated
Portability	Export of applicant-level records in canonical JSON; consent tokens for cross-partner portability	Medium–High	≤ 30 days	Evidence bytes excluded unless legally required; detached signatures provided
Restriction	Flag records as restricted; suppress processing beyond core purposes	Medium–High	≤ 14 days	Overrides require GSIA-authorized lawful basis; reflected in control-plane policy
Erasure	Logical deletion/tombstone; suppression from operational use and displays	High	≤ 30 days	Not absolute: legal holds, GSIA cases, or audit substrate preserved; public entries withdrawn
Objection	Suppress incompatible secondary use; cease marketing-oriented processing	Medium	≤ 14 days	No marketing or comparative publishing is permitted without explicit consent

Consent Withdrawal	Revoke via Consent Ledger; orchestrate withdrawal across displays, caches, and APIs	Medium	Near-real time; SLA governed by Chapter 15	Confirmation receipts issued; residual meta-data preserved for audit

Requests by natural persons (data subjects) shall be distinguished from requests by applicant organizations. Applicant organizations may exercise rights over their Engagements, Validation Results, and Disclosures, subject to the confidentiality of personal data contained within. Natural person requests shall be routed to the appropriate controller—whether a Partner or the Secretariat—based on processing context; joint controller scenarios shall be transparently disclosed with a single-front-door process to avoid burden shifting.

Special cases require heightened safeguards. Requests involving children's data or vulnerable populations shall be prioritized, with additional verification of legal authority for guardians or representatives, and with content minimization in fulfilment. Requests that intersect with state-level sensitivities or classified information shall be handled through bespoke workflows that provide meaningful access without breaching lawful secrecy obligations; GSIA shall supervise such determinations where fundamental rights are implicated.

Rectification preserves historical integrity. Corrections are implemented by appending a corrective record and superseding operational displays and analytics, while retaining the immutable prior state in the Audit & Event Ledger for forensic reconstruction. Public disclosures impacted by rectification shall be updated or withdrawn promptly, with public release notes where appropriate under Chapter 10, avoiding disclosure of personal data in the explanation.

Erasure is implemented to the maximum extent compatible with law, safety, and auditability. Where erasure cannot fully remove an immutable record without compromising forensic guarantees, the platform shall apply logical deletion, cryptographic tombstoning, and suppression from all operational use, search, analytics, and public displays. Any legal holds, GSIA investigations, or regulatory retention mandates shall be recorded with scope and duration, with periodic review and automatic lift upon expiry.

Portability is effected through signed, machine-readable bundles conforming to the canonical schemas of Chapter 3, excluding raw evidence unless a controlling law requires inclusion and transfer is technically safe. Portability bundles may include consent tokens enabling reconstruction of permitted disclosures by another accredited Validation Partner, with explicit scoping and expiry. Receiving parties must be bound by equivalent safeguards and shall not infer certification status or rankings beyond what the bundle expressly contains.

Restriction and objection are enforced by control-plane policy evaluation. Restricted records are shielded from processing beyond validation, adjudication, security, and legal compliance. Objections to incompatible secondary use shall result in suppression of the offending processing path. Any override requires a GSIA-authorized lawful basis and is exceptional, time-bounded, and auditable.

Dispute resolution and complaints follow a structured path. Where a request is denied in whole or in part, the controller shall provide a reasoned explanation and information on escalation to GSIA. GSIA may order corrective measures, mandate disclosures to the requester, or uphold the denial with conditions. Retaliation or service degradation in response to the exercise of rights is prohibited and, if

alleged, shall be investigated with remedies proportionate to harm, up to and including suspension from the ecosystem.

All rights workflows shall be measured for timeliness and quality. Metrics, excluding personal data, shall be aggregated for internal improvement and, where appropriate, reflected as anonymised operational statistics in public dashboards consistent with Chapter 10 and Chapter 17. In all cases of doubt concerning the reconciliation of a rights request with confidentiality or auditability, the presumption shall favor protection of the data subject's interests and referral to GSIA for determination.

Chapter 15 — Consent Ledger and Revocation Protocols

The Consent Ledger constitutes the singular, authoritative mechanism for recording, validating, resolving, and executing consent and revocation across the A2074-SRS ecosystem. It binds disclosures to explicit, informed, and revocable permissions granted by the rights-holder and enforces minimum-disclosure-by-default across all registries, displays, and APIs. The Ledger operates within the control plane with immutability guarantees, cryptographic attestations, and programmatic orchestration to ensure that any consent-dependent display or data flow is continuously conditioned on a current, valid consent state.

Consent shall be specific, granular, intelligible, freely given, and separable from core validation services. It shall define scope, audience, channels, duration, jurisdictional implications, and linkage to the affected canonical entities, including Engagements, Validation Results, and Disclosures. Each consent record shall include a subject reference (pseudonymous where feasible), a scope descriptor with precise bindings to data categories and outputs, time parameters (issue, expiry, review), the lawful basis, and attestations capturing the identity assurance method, the information presented at the time of grant, and the mechanism used to evidence voluntariness. No consent may be bundled with conditions that penalize refusal or withdrawal. Any purported consent obtained under duress, coercion, or unequal bargaining power is voidable and shall be treated as invalid upon GSIA determination.

The Ledger shall implement state immutability through append-only transitions. A consent "state" is never altered retroactively; new states supersede prior states and establish a cryptographically verifiable lineage for audit and adjudication as set forth in Chapter 9. State transitions include creation, amendment (narrowing or broadening scope), suspension (temporary hold), renewal, and revocation. Each transition event shall be recorded with actor identity, authentication strength, timestamp, prior state reference, and detached signatures from both the initiating party and the platform. Where an Applicant organization acts through an authorized representative, the representative's authority shall be verified and logged, with explicit indication that the consent pertains to organizational disclosures rather than personal data.

Consent resolution is mandatory at the point of every potential disclosure. Before any publication, syndication, public API response, or third-party data share, the control plane shall evaluate whether a valid consent state exists for the exact scope requested. Resolution must consider jurisdictional restrictions, residency, and audience, and must map the proposed output to the consent's enumerated channels and visibility settings. Cached decisions shall honor strict time-to-live limits; critical operations such as publication and withdrawal shall never rely solely on cache and must re-query the Ledger.

Revocation is a first-class right and shall be at least as simple, accessible, and rapid as consent grant. Upon revocation, the platform shall orchestrate immediate withdrawal across all internal displays and public endpoints within defined service-level objectives, followed by coordinated takedown requests

to caches, mirrors, syndication partners, and archival services. The orchestration plan shall prioritize public endpoints, search indexation suppression through non-persistent directives where possible, and downstream partner notifications that include binding obligations to delete or suppress previously received materials. Orchestration outcomes and residual risks shall be recorded against the revocation event, with ongoing monitoring until completion or the best attainable suppression consistent with law.

The following table defines canonical consent states and their operational effects. The vocabulary is exhaustive for canonical purposes; additional partner-namespaced modifiers may be layered provided they cannot dilute the effect of a canonical state.

Consent State	Description	Permitted Actions	Prohibited Actions	Notes
granted	Valid, in-force consent for defined scope, channels, and audience	Publish or share strictly within scope; render public extracts as specified	Any out-of-scope or comparative display; onward use beyond channels	Time-bound; requires periodic review if duration is long
suspended	Temporary hold by subject or controller pending review	Maintain existing display if still within legally required window; hold new publications	Any new publication or syndication; expansion of scope	Auto-expires or escalates to revoked after defined period
amended	Narrowed or broadened scope replacing prior state	Actions aligning with amended scope	Actions aligned only with prior state	Prior state persists only for audit lineage
revoked	Withdrawal of consent for the defined scope	Execute withdrawal orchestration; maintain tombstoned records	Any further display or sharing	Immediate effect; downstream takedown triggered
expired	Time-limited consent lapsed naturally	Cease publication; treat as revocation for orchestration	Any continued display predicated on the lapsed state	Expiry schedules must be monitored automatically

Execution of revocation and expiry shall occur within strict service-level objectives calibrated to risk. The internal objective for removal from A2074-controlled endpoints is near-real time, with completion targets measured in minutes for primary sites and hours for globally distributed caches, subject to technical constraints. For third-party recipients, contractual SLAs shall impose prompt deletion and confirmation, and non-performance shall constitute a material breach handled under Chapter 12. Where complete technical erasure is infeasible, suppression, de-indexing, and visible withdrawal notices shall be employed to mitigate residual exposure, with justifications recorded for GSIA oversight.

Consent for special categories of actions—such as comparative displays, league tables, or entity-level rankings—shall be collected as a separate, heightened consent tier and shall default to “not granted.” Even where granted, proportionality and non-comparative principles apply; GSIA may restrict or

condition such displays to avoid ecosystem harms. Consent to identify named individuals is exceptional and shall require explicit, separate agreement with clear revocation effects and narrow scope.

Auditability of consent operations is mandatory. Creation, amendment, suspension, and revocation events shall be written to the Audit & Event Ledger with chain hashes, detached signatures, and references to any displays or publications affected. The control plane shall refuse publication if the corresponding audit write is not durably committed. Any failure in revocation orchestration—whether technical, contractual, or jurisdictional—shall open an incident under Chapter 11 with accelerated remediation and GSIA notification where residual risk persists.

User experience must be clear, honest, and reversible. Consent interfaces shall present non-misleading summaries, layered disclosures, and easy access to full terms, avoiding pre-ticked boxes, bundled selections, or dark patterns. Confirmation receipts shall be issued for both grant and revocation, including a machine-readable token referencing the Ledger state and human-readable instructions for follow-up. Where consent interacts with language or accessibility needs, localized and accessible variants shall be provided.

Nothing in this Chapter authorizes publication without consent or continuation of display after revocation or expiry. Where a lawful exception requires temporary retention or limited disclosure notwithstanding revocation—such as legal obligations or GSIA-ordered preservation—such exception shall be narrowly construed, documented with scope and duration, and insulated from public display or secondary use.

Chapter 16 — Minimum Disclosure by Default

Minimum disclosure by default is a cardinal principle and operational policy that governs every component of the A2074-SRS platform. The baseline state for all records, results, evidence, and analytics is non-display to the public and non-sharing with third parties, unless and until a valid, specific, and current consent state authorizes an exception, or a narrowly defined lawful obligation compels limited disclosure. This Chapter translates that principle into enforceable technical and procedural controls across registries, APIs, caches, and user interfaces.

All canonical entities are created as private. Engagements, Validation Results, Evidence Index entries, and AuditEvents remain invisible to any public interface and inaccessible to third parties. Disclosure is an explicit, affirmative act that requires a valid granted consent state (Chapter 15), a Disclosure record that binds the consent scope to the intended audience and channels (Chapter 3), and a control-plane decision that emits a time-bound authorization for the specific display or transmission. Absent this tripartite alignment, any attempt to publish must fail deterministically with a recorded denial event.

Publication pathways shall be gated by policy evaluation that incorporates consent, jurisdiction, and role attributes. The rendering layer shall fetch only those artifacts that have active display tokens issued by the control plane for the requesting context. Tokens shall be short-lived, audience-scoped, and revocation-sensitive. Static exports are prohibited except where absolutely necessary and explicitly consented; where produced, they shall carry embedded expiry and revocation hooks and must be registered in the Disclosure Registry for downstream orchestration.

The platform shall embed preventive controls to avoid accidental disclosure. Public endpoints shall query sanitized, publication-scoped views rather than canonical registries. Development, staging, and test environments shall employ synthetic or redacted datasets by default, with strict separation from production. Administrative consoles shall present public-facing previews that are watermarked and

clearly labeled “non-public until consent validated.” Any function that changes a display state shall require step-up authentication, dual-control for high-risk contexts, and a visible summary of the consent parameters being invoked.

De-publication is prioritized over publication. Systems shall always prefer and preserve the capacity to withdraw disclosures upon revocation, even during degraded operating states or maintenance windows. Change Management (Chapter 13) shall not permit deployments that reduce withdrawal performance or reliability. Incident Response (Chapter 11) shall prioritize containment of any unintended displays and immediate suppression of public materials. Public caches, mirrors, and syndication endpoints shall be under technical and contractual obligations to honor de-publication signals rapidly.

The following table sets out the canonical disclosure decision matrix. It is binding on all components that may present or transmit data externally.

Requested Action	**Consent Requirement**	**Jurisdiction & Residency Check**	**Additional Guards**	**Default Outcome**
Public display of entity-level result	granted with explicit channel “public web”	Mandatory pre-flight; deny if constraints unmet	Non-comparative presentation; redactions; rate limiting	Deny unless all checks pass
Third-party syndication feed	granted with channel “syndication:`”	Mandatory; ensure safeguards and contracts	Signed payloads; takedown SLA registered	Deny unless all checks pass
Public aggregate publication	Not required if de-identified under Chapter 17	Residency applies to processing location	Re-identification risk assessment; GSIA review	Deny unless risk ≤ policy threshold
Comparative or rank-ordered display	Separate, heightened consent per entity	Mandatory	GSIA ethics pre-approval; proportionality test	Deny unless explicitly approved
Disclosure involving personal data	Explicit personal consent; exceptional	Mandatory	Minimise, pseudonymise; revocation priority	Deny unless strictly justified

User interfaces shall default to privacy-preserving choices. Public-facing pages shall not reveal or infer entity identities unless a live display token exists. Search engines shall not be facilitated to index beyond consented materials; robots directives and meta-tags shall be generated dynamically based on consent state. Public APIs shall never return internal identifiers; only public surrogates associated with the Disclosure record may be exposed, and these shall be invalidated on withdrawal.

Partners and vendors shall be contractually bound to the minimum disclosure policy. No partner may create shadow registries or public mirrors that bypass the control plane’s consent and policy evaluation. Any marketing, communications, or press materials that reference validations must be checked against current consent states immediately prior to release and must include mechanisms to update or

withdraw content upon revocation. Violations constitute material breaches subject to GSIA adjudication and remedial measures, including suspension of public interface privileges.

Operational telemetry and logging related to public display shall avoid personal data and shall be retained only for security, performance, and accountability purposes. Analytics derived from public traffic shall employ aggregation and de-identification, with no tracking cookies or cross-site identifiers beyond what is strictly necessary to defend against abuse. Rate limits shall protect against scraping and bulk extraction that could reconstruct sensitive patterns, and abuse response shall include dynamic blocking and legal notices where warranted.

Exceptions to minimum disclosure are narrow, explicit, and supervised. Where a legal obligation compels limited public posting—such as a mandated notice—only the minimum necessary content shall be displayed, for the minimum duration, and with revocation hooks to ensure automatic removal upon expiry or supersession. Such exceptions shall be logged with legal basis, scope, and timeline, and shall be reviewable by GSIA. No exception may be interpreted to allow comparative displays or the publication of personal data absent specific statutory authority and proportional safeguards.

The ethos and enforcement of minimum disclosure by default extend to the culture of the ecosystem. Training, onboarding, and periodic attestations for administrators, assessors, and communications teams shall emphasize non-publication as the baseline, consent as the exception mechanism, and revocation as a priority obligation. Any ambiguity or doubt concerning publication shall be resolved in favor of non-disclosure and escalation to the A2074 Secretariat and GSIA for determination.

Chapter 17 — Pseudonymisation and De-Identification Standards

Pseudonymisation and de-identification within the A2074-SRS ecosystem serve to reduce re-identification risk while preserving the analytical and operational utility required for validation, oversight, and learning. These processes are risk-based, context-aware, auditable, and reversible only under tightly controlled conditions authorized by the control plane in accordance with consent and lawful basis. Pseudonymisation is distinguished from de-identification in that the former preserves a protected linkage to identity through separately held keys, whereas the latter aims to remove or transform identifiers and quasi-identifiers to the extent that re-identification risk is demonstrably below defined thresholds. No de-identification shall be represented as absolute anonymisation unless a structured, documented risk assessment demonstrates that re-identification is not reasonably likely by any party using reasonable means within the relevant context.

Pseudonymisation shall be the default for operational processing that touches personal data or patient-analogue attributes. Pseudonymous identifiers must be non-meaningful, non-derivable, and generated through cryptographic or high-entropy mechanisms under the control plane. Re-identification keys shall be stored separately in the Key & Certificate Registry and accessible only under explicit consent or a GSIA-authorized lawful basis, subject to step-up authentication, dual-control, and immutable audit. Pseudonymisation alone does not remove the data from the scope of privacy protections; all controls in Chapters 5, 6, 7, 9, 15, and 16 remain fully applicable.

De-identification for analytics and public dashboards shall follow a documented methodology commensurate with the sensitivity of the data, the richness of auxiliary information reasonably available to potential adversaries, and the intended audience. Techniques may include, individually or in combination, removal or masking of direct identifiers, generalisation of quasi-identifiers, aggregation to minimum cell sizes, noise addition calibrated to utility and risk (including differential privacy where appropriate), suppression of small cells, and derivation of synthetic data for non-sensitive use cases.

The selection and calibration of techniques shall be justified in a risk assessment recorded in the Audit & Event Ledger, with GSIA review for public releases under Chapter 10.

Risk assessment is mandatory and continuous. Re-identification risk shall be evaluated against adversary models that consider realistic capabilities, including access to public records, commercial data, or previously released aggregates. Risk shall be quantified using accepted measures and expressed alongside uncertainty bounds. Where risk exceeds policy thresholds, publication shall be denied or the dataset further transformed. Risk assessments shall be versioned and re-evaluated upon changes in data scope, auxiliary data availability, or methodology.

The following table sets out the canonical de-identification toolkit and its governance guardrails. The toolkit is authoritative; Partners may propose extensions subject to Secretariat approval and GSIA ethics review.

Technique	**Typical Use**	**Controls and Limitations**	**Governance Requirements**
Direct Identifier Removal (names, emails, IDs)	Baseline step for all releases	Does not protect against linkage via quasi-identifiers	Mandatory; recorded in methodology notes
Generalisation (banding ages, region to NUTS-level, time to month/quarter)	Reduce linkage risk from quasi-identifiers	Excessive generalisation harms utility; calibrate to audience	Documented thresholds; GSIA review for public
Suppression (cell counts below *k*)	Prevent sparse disclosures	Choose *k* to balance risk and utility; deterministic rules to avoid inference across releases	Published *k* policy; enforce across time
Noise Addition (Laplace/Gaussian)	Aggregate metrics in dashboards	Accumulates across queries; track privacy budget if differential privacy is used	Privacy budget ledger for repeated queries
Top/Bottom Coding	Prevent outlier identification	May bias statistics; disclose methodology	Required note in release documentation
Pseudonymisation (tokenisation)	Operational analytics with potential re-linkage under control	Not anonymisation; keys must be isolated	Dual-control for re-identification; audit required
Synthetic Data (model-generated)	Prototyping, training, demos	Risk of memorisation or leakage; validate utility vs. risk	Leakage testing; ban for public inference on entities

Differential Privacy (ε-bounded)	High-assurance public statistics	Choose epsilon conservatively; cumulative accounting	Publish ε, composition method, and budget usage

Longitudinal releases and linkage risks require special safeguards. Repeated publication of overlapping aggregates can enable reconstruction and re-identification through differencing. The platform shall maintain a release ledger for public aggregates, recording schemas, cohort definitions, suppression thresholds, noise parameters, and privacy budgets to detect and prevent risky compositions. Where longitudinal comparability is necessary, stable aggregations shall be designed that preserve privacy guarantees across time without enabling reverse inference.

Contextual integrity governs what may be considered “reasonably likely” re-identification. The potential for linkage through publicly known partnerships, sector-specific registers, or media coverage must be considered. Where contextual risks are elevated—such as for small jurisdictions, niche sectors, or singular high-profile Applicants—additional generalisation, aggregation, or non-publication shall be imposed. GSIA may set sectoral or regional guardrails to prevent inadvertent exposure despite technical compliance.

Testing and attestation are mandatory. Prior to public release, de-identified datasets and dashboards shall undergo adversarial testing by an internal or independent team separate from the creators, attempting linkage and re-identification using plausible auxiliary data. Results, including attempted vectors and outcomes, shall be recorded and reviewed by the Secretariat and GSIA. Material weaknesses must be remediated prior to release. For internal analytics, periodic spot-checks shall be performed to ensure that de-identification practices remain effective as data distributions and auxiliary landscapes evolve.

No de-identified dataset may be combined with other datasets by the platform in ways that would raise re-identification risk above approved thresholds. Third parties receiving de-identified outputs under consented syndication must be contractually prohibited from attempting re-identification, must implement equivalent safeguards, and must honor revocation-triggered suppression where feasible. Violations constitute material breaches under Chapter 12 and may trigger GSIA-directed remedies.

Re-identification, when authorized, is exceptional and controlled. It may occur only under explicit, current consent by the data subject, to fulfill a data rights request, or under a GSIA-authorized lawful basis such as a substantiated safety concern or legal obligation. Re-identification requires step-up authentication, dual-control, granular scope definition, and immediate audit logging with justification. Any unauthorized re-identification attempt shall be treated as a security and privacy incident under Chapter 11 with mandatory notification and remedial action.

Methodological transparency is required for public trust. Each public release shall include documentation of the techniques employed, parameters chosen, known limitations, and a contact channel for concerns. Documentation shall avoid revealing sensitive thresholds that could aid adversaries while providing sufficient clarity for accountability and interpretability. Historical releases shall be reviewed periodically; if changes in auxiliary data render prior releases riskier than intended, the platform shall suppress or revise those releases and issue public notes consistent with Chapter 10.

In all cases of doubt, the presumption shall favor stronger protection, reduced granularity, or non-publication, with referral to GSIA for ethics review. The integrity of patient-level confidentiality and

the non-comparative ethos of A2074-SRS shall not be compromised for analytical or communications convenience.

Final Word

This Manual establishes the technical and governance architecture through which the A2074-SRS ecosystem operates lawfully, securely, and ethically. It binds Validation Partners, vendors, and platform operators to a standard of conduct that centers patient-level confidentiality, explicit and revocable consent, proportionality in evaluation, and non-comparative presentation. The provisions are interdependent: master registries and identity controls uphold integrity; data models, APIs, and evidence stewardship ensure coherent and secure operations; cybersecurity and auditability protect against malfeasance and error; analytics and de-identification enable responsible transparency; incident response and vendor governance safeguard resilience; change management ensures orderly evolution without degradation of rights; and data-subject rights, consent ledgering, and minimum disclosure by default anchor the ecosystem in dignity and trust.

Agenda 2074, as the standard-setter, defines the universal canon; GSIA, as independent custodian, supervises compliance and adjudicates disputes; Validation Partners innovate within the boundaries of this Manual and the Multi-Model Validation Framework to serve Applicants of all sizes with fairness; and Applicants retain control over the visibility of their engagements and results. EUSL, as flagship Validation Partner in Europe, exemplifies these obligations through its hospitality-style star model aligned to the seventeen pillars, implemented strictly within the privacy-by-design and consent-first architecture defined herein.

Nothing in this Manual authorizes coercive consent, comparative harm, or disclosure without lawful basis. All actors are obligated to resolve ambiguity in favor of confidentiality and to escalate uncertainties to the A2074 Secretariat and GSIA. As technology, jurisdictions, and practices evolve, the Manual's change controls and oversight mechanisms shall ensure that progress does not erode rights or integrity. Compliance is not a point in time but a continuous discipline, evidenced by auditable practice and accountable to an independent ethics jurisdiction.

By adhering to these provisions, the ecosystem advances credible social responsibility validation while preserving human dignity, institutional trust, and the public interest that the Agenda for Social Equity 2074 is constituted to protect.