


GoPES 2.0

Integrated Governance, Privacy, Ethics & Security Enterprise Control Architecture and UML Design Specification

Consolidating source diagrams 001–019, with diagram 006 as the governing structural model
UML 2.5.1 aligned • data/AI product lifecycle • generative and agentic AI • runtime decision control

Version 1.0 — 11 September 2026



Source anchor: diagram 006. The four-part GoPES interface (Planning, Documentation, Execution, Monitoring) is retained as the organising structure.

1. Executive summary

This specification consolidates the supplied GoPES diagrams into one coherent enterprise control architecture. Diagram 006 is used as the backbone: GoPES operates as an interface between organisational capabilities, policies and standards, business processes and system capabilities. The remaining diagrams expand that backbone with governance roles, data-quality management, privacy and security controls, ethics criteria, AI/ML lifecycle controls, commercialisation, budgeting, organisational readiness, KPIs, product catalogues, decision engines, API/event integration and immutable decision evidence.

The resulting design treats GoPES not merely as a governance framework but as an executable control plane. Policies are administered in a control plane; decision requests are evaluated against Governance, Privacy, Ethics and Security (G/P/E/S) engines; contextual evidence is obtained from authoritative enterprise sources; high-impact or ambiguous decisions are escalated to named human decision makers; and the final permit/deny/constrain/escalate outcome is issued as a signed decision record that can be enforced by APIs, workflows and AI-agent middleware.

The design is aligned to the current formal UML specification (OMG UML 2.5.1), and the control model is informed by ISO/IEC 42001:2023, ISO/IEC 27001:2022, ISO/IEC 27701:2025, ISO/IEC 23894:2023, ISO/IEC 38507:2022, NIST AI RMF 1.0 and its Generative AI Profile, NIST CSF 2.0, the UK GDPR/ICO guidance as modified by the Data (Use and Access) Act 2025, and current OWASP GenAI/agentive-security guidance including the 2026 LLM Top 10 and Agent Control Standard [R1–R16].




Figure 1 — Consolidated GoPES system context.

1.1 Design outcomes

- One accountable enterprise governance model linking the Board/Data & AI Council, CDO/DGO, business and data owners, technical teams, specialist SMEs, delivery teams and independent assurance.
- One control-plane pattern for data, analytics, AI/ML, generative AI, agentive AI, APIs and data products, rather than separate governance silos.
- Runtime enforcement through policy decision and enforcement points, with human approvals used where law, risk appetite or policy requires meaningful intervention.
- Persistent provenance and signed decision records binding policy versions, evidence, reasons, obligations, approvals, validity windows and execution traces.
- Lifecycle governance from opportunity identification and registration through design, build, validation, release, monitoring, material change, exception management and retirement.
- Traceability from every supplied source diagram into the integrated architecture and requirements catalog.

2. Standards baseline and modelling conventions

The source diagrams use UML-like class, component and actor notation. This specification normalises them to UML 2.5.1 concepts and uses explicit stereotypes where useful (for example «control engine», «policy enforcement point», «data product», «decision record»). UML 2.5.1 remains the current formal OMG UML specification [R1].

Reference	How it is used in GoPES 2.0
OMG UML 2.5.1	Defines the modelling language and semantics used for actors, use cases, classes, components, activities, states, sequences, deployment nodes and relationships.
ISO/IEC 42001:2023	Management-system backbone for governing AI: context, leadership, planning, support, operation, performance evaluation and improvement [R2].
ISO/IEC 38507:2022	Board/governing-body guidance for effective, efficient and acceptable organisational use of AI [R3].
ISO/IEC 23894:2023	AI-specific risk-management process and integration of AI risk into enterprise risk activities [R4].
NIST AI RMF 1.0 + NIST AI 600-1	Operational AI risk functions Govern, Map, Measure and Manage; generative-AI-specific risk considerations [R5–R6].
ISO/IEC 27001:2022 + NIST CSF 2.0	Information-security management, governance, identity, protection, detection, response and recovery [R7–R8].
ISO/IEC 27701:2025 + ICO/UK GDPR	Privacy management, lawfulness, fairness, transparency, minimisation, rights, security and accountability [R9–R11].
Data (Use and Access) Act 2025	Updated UK automated-decision safeguards, including information, challenge/representation and human intervention for significant solely automated decisions [R12].
OWASP GenAI / Agentic AI 2026	AI application and agent threats, runtime control, inspectability, traceability, middleware hooks and control enforcement [R13–R16].

3. GoPES governing structure — diagram 006 expanded

Diagram 006 defines four architectural surfaces around the GoPES interface. The integrated model retains these as the governing structure and makes each surface implementable.

006 domain	Purpose	Required components / artefacts
OC1 Organisational Capabilities	Who is accountable and capable of operating the control system.	Executive sponsorship; Data & AI Council; CDO/DGO; owners/stewards/custodians; specialist SMEs; delivery roles; 2nd/3rd line assurance; competence/training; communications; capacity planning; review cadence.
PS1 Policies & Standards	What rules, principles and success criteria apply.	Enterprise principles; data/AI policies; privacy notices; information classification; access policy; secure development; model/agent standards; third-party controls; documentation standards; exception policy; policy versioning and ownership.
BP1 Business Process	How governance is embedded in operational work.	Intake; assessment; data onboarding; product lifecycle; privacy/DPIA; threat modelling; ethics/fairness review; approval; KYC/AML where relevant; change control; incident/exception handling; retirement; reporting.
SC1 System Capabilities	How the rules are automated and evidenced.	Catalog/repository; workflow; API/event bus; PDP/PEP; identity and secrets; G/P/E/S engines; decision records; evidence store; observability; model/agent gateway; lineage; data quality; dashboards; integrations; immutable logs.



Figure 2 — Integrated capability map derived from diagram 006 and the wider source set.

4. Roles, decision rights and key responsibilities




Figure 3 — Organisational operating model and responsibility structure.

Role	Key responsibilities	Decision boundary
Board / Data & AI Council	Sets risk appetite and strategic principles; approves material policy; receives assurance; resolves enterprise-level conflicts; approves major/time-bounded exceptions; ensures AI use remains effective, efficient and acceptable.	Cannot delegate ultimate oversight; should not perform day-to-day administration.
CDO / Data Governance Office / GoPES Service Owner	Owens the framework, taxonomy, operating model, controls catalog and GoPES platform service; coordinates business, technical and SME communities; maintains metrics and reporting; ensures policy-to-control traceability; manages governance forums.	Accountable for service effectiveness and cross-domain coherence.
Business / Process Owner	Defines intended outcome, customer or operational need, value hypothesis, acceptable service levels and business risk; funds the product; accepts residual business risk within delegated authority.	Owens why the capability exists and whether it remains justified.
Data Owner	Authorises purposes, access, sharing, quality thresholds, retention and critical definitions for owned data domains; resolves stewardship issues; accepts data-quality risk within authority.	Accountable for data use, not merely technical custody.
Data Steward	Maintains metadata, definitions, data-quality rules, lineage, issue logs and evidence; coordinates remediation; supports access and policy decisions.	Operational accountability for data fitness and metadata completeness.
Data Custodian / Platform Owner	Implements technical controls for storage, backup, identity, encryption, retention, recovery, access and platform resilience; produces control evidence.	Technical implementation, not business purpose ownership.
DPO / Privacy Lead	Interprets privacy law and policy; advises on lawful basis, transparency, minimisation, DPIA, rights, ADM safeguards, transfers and privacy by design; provides independent advice/challenge.	Must retain appropriate independence under applicable law/policy.
CISO / Security Lead	Defines security requirements; reviews threat models and agent/tool permissions; oversees identity, secrets, secure SDLC, red teaming, incident response, vulnerability and supply-chain controls.	Owens security control framework and assurance.
Ethics / Responsible AI Lead	Maintains ethical principles and assessment method; evaluates fairness, autonomy, dignity, harm, accessibility, societal impact, sustainability and stakeholder participation; drives mitigation and escalation.	Provides challenge where legal compliance alone is insufficient.
Legal / Regulatory SME	Maps applicable legislation, sector rules, contracts and regulatory commitments; ensures regulatory obligations and records are integrated into policy decisions.	Regulatory rules act as a cross-cutting input, not a separate governance silo.
Enterprise / Solution Architect	Defines approved patterns, target architecture, integration, data flows, control points, non-functional requirements and technical debt constraints.	Ensures GoPES enforcement is designed into systems, not bolted on.
Product Manager / AI Product Owner	Maintains product backlog, value measures, customer outcomes, documentation,	Accountable for complete lifecycle evidence.

	dependencies and release criteria; coordinates assessments and evidence.	
Data Scientist / ML Engineer	Develops/evaluates models; documents datasets and model behaviour; performs reproducibility, bias, robustness, drift and performance testing; supports monitoring and retraining.	Must work within approved data, model and evaluation boundaries.
Agent / Application Engineer	Defines agent identity, tools, scopes, prompts/instructions, memory, action budgets, approval points and failure modes; implements runtime enforcement hooks.	No unrestricted tool access; least-agency and least-privilege defaults.
MLOps / Platform Engineering	Builds CI/CD and model/agent deployment controls; versioning; registries; environment separation; observability; rollback; policy gates; artefact signing.	Ensures repeatable, controlled deployment and evidence capture.
Risk / Compliance (2nd line)	Challenges risk assessments and exceptions; monitors policy adherence; tests control design and effectiveness; provides independent management reporting.	Independent of first-line delivery decisions.
Internal Audit (3rd line)	Performs independent assurance over governance, controls, evidence integrity and management response.	Does not own controls it later audits.
Supplier / Third Party	Provides security, privacy, AI, data, model and operational evidence; complies with contractual controls; reports material changes and incidents.	Subject to third-party due diligence and ongoing monitoring.
End User / Human Approver	Uses the capability within policy; provides meaningful review where required; records rationale for approvals or overrides; raises incidents and suspected shadow AI.	Human-in-the-loop must be competent, informed and able to change the outcome.

4.1 Responsibility matrix (RACI-style)

Activity	Council	CDO/DGO	Owner	Steward	DPO	CISO	Ethics	Tech/AI	Risk/Audit
Policy & risk appetite	A	R	C	C	C	C	C	C	I
Register product/use case	I	C	A/R	C	C	C	C	R	I
Data classification & quality	I	C	A	R	C	C	I	R	I
Privacy/DPI A/ADM assessment	I	C	A	C	R	C	C	R	C
Security/threat assessment	I	C	A	C	C	R	C	R	C
Ethics/fairness assessment	I	C	A	C	C	C	R	R	C
Release approval	I	A/C	A	C	C	C	C	R	C
Runtime decision policy	I	A/R	C	C	C	C	C	R	C
Exception approval	A*	R	C	C	C	C	C	C	C
Monitoring & incident response	I	A	A	R	C	R	C	R	C
Assurance/reporting	I	R	C	C	C	C	C	C	A/R

* Council approval applies to exceptions above delegated materiality/risk thresholds; lower-risk exceptions may be delegated but must remain time-bounded and visible.

5. Logical architecture




Figure 4 — GoPES logical component architecture: policy administration, decision, enforcement, evidence and monitoring.

Component	Required capability
GoPES Orchestrator / Decision Service	Receives decision requests; normalises subject/action/resource/purpose/context; calls authoritative information sources and domain engines; resolves conflicts; determines whether human review is needed; constructs final decision and obligations; maintains trace IDs.
Policy Administration Point (PAP)	Authoring, review, approval, versioning and deployment of policies/rules. Supports effective dates, owners, precedence, scope, test cases, mappings to standards and rollback.
Policy Decision Point (PDP)	Executes policy-as-code using deterministic rules where possible. Returns permit, deny, constrain or escalate plus reasons, required controls and validity window.
Policy Enforcement Point (PEP)	Interposes on APIs, agent tool calls, workflows, data access and deployments. Verifies signed decisions, enforces obligations such as masking, rate limits, tool restrictions, approval requirements and budget ceilings.
Policy Information Points (PIP)	Provide authoritative attributes and evidence: IAM, asset catalog, data classification, lineage, consent/lawful-basis records, risk register, CMDB,

	vendor register, model registry, incident status, data quality and geography.
Governance Engine	Checks ownership, purpose, policy applicability, lifecycle state, delegated authority, documentation completeness, data/product catalog registration and change-control requirements.
Privacy Engine	Checks personal/special-category data, lawful basis, purpose limitation, minimisation, transparency, data subject rights, retention, transfer, DPIA status, automated-decision safeguards and privacy obligations.
Ethics Engine	Evaluates fairness, explainability, autonomy, dignity, beneficence/non-maleficence, accessibility, stakeholder impact, sustainability, contestability and human oversight against risk tier.
Security Engine	Evaluates identity, least privilege, segregation, encryption, secrets, data exfiltration, prompt/tool injection, supply chain, sandboxing, content trust, adversarial threats, incident status and resilience.
Regulatory Rules Adapter	Maps jurisdictions, sectors and contractual rules to executable obligations. Avoids duplicating the four GoPES domains by feeding applicable legal/regulatory constraints into them.
Data Quality Engine	Implements validity, accuracy, completeness, consistency, uniqueness, timeliness and provenance rules; supports profiling, thresholds, issue workflow and quality SLAs.
AI Risk / Model Assurance	Maintains model/agent inventory, risk tier, intended use, evaluation plan, model/system cards, robustness, red-team, drift and performance evidence; aligns to AI RMF/ISO 42001/23894.
Human Approval & Escalation	Routes cases to role-qualified approvers with the complete decision packet; captures rationale, conditions, conflicts of interest, expiry and delegation.
Decision Record Service	Creates immutable, tamper-evident records that bind decision request, policy versions, evidence hashes, domain findings, human approvals, obligations, validity and cryptographic signature.
Evidence & Catalog Store	Stores data/AI product records, assessments, policies, controls, model cards, DPIAs, test results, supplier evidence, exceptions, lineage and audit references with retention controls.
Observability & Continuous Monitoring	Collects security events, data quality, model/agent drift, performance, cost, use patterns, policy breaches and outcome metrics; triggers re-evaluation, suspension or incident workflows.
API/Event Integration Layer	Supports synchronous REST/gRPC authorization decisions and asynchronous event-driven monitoring and evidence ingestion. Kafka-style messaging from the source diagrams becomes an implementation option, not a mandatory dependency.
Identity, Secrets & Workload Trust	Human and machine identities, workload credentials, PAM, key management, signed artefacts and agent identity. Enables attributable decisions and least privilege.
Model/Agent Gateway & Tool Sandbox	Central control for model endpoints, prompts/system instructions, tools/connectors, egress, memory, token/cost budgets, content filtering, sandboxing and per-agent action scopes.
Metrics Layer	Produces business KPIs, operational KPIs, customer measures, employee/adoption measures and control KRIs, preserving definitions and provenance so reporting is reproducible.

6. Data governance and data-quality subsystem

Diagrams 001, 002, 003, 004 and 005 are consolidated into a single data-governance subsystem. Data quality is treated as a governed product characteristic with named ownership, dimensions, thresholds, issue management, evidence and continual improvement, not as a one-time cleansing exercise.




Figure 5 — Consolidated data-quality model.

Dimension	Definition / control expectation	Typical measures
Validity	Conformance to required type, format, syntax, range and domain rules.	schema pass rate; invalid value count; referential integrity failures
Accuracy	Degree to which data correctly represents the real-world object/event for the intended purpose.	verified error rate; reconciliation variance; reference-source match
Completeness	Presence and coverage of required data and records.	null/missing rate; population coverage; mandatory-field completeness
Consistency	Agreement across records, systems, time and taxonomies.	cross-system mismatch rate; conflicting status values; rule contradictions
Uniqueness	Absence of inappropriate duplicates and correct entity resolution.	duplicate rate; merge precision/recall; duplicate key violations
Timeliness	Availability at the required freshness/latency for the process.	age/freshness; pipeline latency; SLA attainment
Lineage / provenance	Traceable source, transformation, ownership and evidence chain.	lineage coverage; untraceable fields; provenance integrity failures

7. Governance, Privacy, Ethics and Security control domains

7.1 Governance

- Purpose and outcome are documented and approved before data/model/tool access is granted.
- Every governed object has a named accountable owner, lifecycle state, risk tier and materiality.
- Policies, standards and controls are versioned, testable and mapped to evidence.
- Decision rights and delegation limits are explicit; exceptions are time-bounded and require compensating controls.
- Material changes trigger re-assessment rather than silently inheriting historic approvals.
- Board/Council receives outcome, risk, exception and control-effectiveness reporting, not only activity metrics.

7.2 Privacy

The privacy model incorporates the UK GDPR principles—lawfulness/fairness/transparency, purpose limitation, minimisation, accuracy, storage limitation, integrity/confidentiality and accountability—and is designed to adapt as ICO guidance is updated following the Data (Use and Access) Act 2025 [R10–R12].

- Classify whether personal, special-category or criminal-offence data is involved and identify controller/processor roles.
- Record lawful basis, purpose, compatibility of reuse, transparency information and data-subject rights handling.
- Determine whether a DPIA is required; bind approved mitigations and residual risk to the product record.
- For significant solely automated decisions, apply the post-DUAA safeguards: provide information, enable representations/challenge, and enable human intervention where required.

- Apply data minimisation, retention, masking/tokenisation, access segregation, transfer controls and privacy-enhancing technologies where proportionate.

7.3 Ethics / responsible AI

The supplied ethics diagrams contain 11 recurring ethical stubs. In GoPES 2.0 these become a structured assessment library: transparency; justice/fairness; non-maleficence; responsibility; privacy; beneficence; freedom/autonomy; trust; sustainability; dignity; solidarity. They are mapped to measurable controls rather than stored as abstract values.

Ethical area	Representative questions / controls
Transparency	Can affected users understand when AI is used, what its role is, its key limitations and how to seek review?
Justice & fairness	Are performance and error rates assessed across relevant groups? Are accessibility and distributional impacts considered?
Non-maleficence	What harms are plausible? What safety constraints, fallback, rollback and incident paths exist?
Responsibility	Who is accountable for outcomes, approvals, monitoring, remediation and user redress?
Privacy	Does processing respect data-protection principles and individual rights?
Beneficence	Is there a credible positive outcome and evidence that the intervention is proportionate?
Freedom & autonomy	Does the system preserve meaningful choice, consent where relevant, contestability and avoidance of manipulative or coercive patterns?
Trust	Are reliability, provenance, limitations, incidents and changes communicated honestly?
Sustainability	Are environmental/resource effects measured when material and balanced against value?
Dignity	Does the design avoid demeaning treatment and preserve human agency and respectful interaction?
Solidarity	Are wider social impacts, vulnerable groups and collective consequences considered?

7.4 Security

Security controls combine the source data-security model with ISO/IEC 27001, NIST CSF 2.0 and current OWASP GenAI/agent guidance. OWASP's 2026 materials emphasise AI-specific threats and runtime agent control; the Agent Control Standard calls for inspectable, traceable and instrumentable agents with portable middleware enforcement hooks [R7–R8, R13–R16].

- Machine and human identity, workload attestation and least-privilege tool/data scopes.
- Data classification, encryption at rest/in transit, secrets management, tokenisation/masking and secure deletion.
- Prompt/instruction injection resistance, untrusted-content separation, output validation and constrained tool invocation.
- Model/data/software supply-chain integrity, signed artefacts, SBOM/model provenance and third-party evidence.
- Agent action budgets, allow-lists/deny-lists, environment sandboxing, network egress constraints and separation of duties.
- Security testing, red teaming, abuse-case testing, vulnerability management, incident detection/response and recovery.
- Resilience and safe fallback: degraded modes, kill switch, revocation, rollback and continuity requirements.

8. Data products, AI/ML and agent lifecycle

Diagrams 015–019 introduce core AI elements, AI product categories, model lifecycle activities and engine/API integration. GoPES 2.0 treats these as governed product types with a common lifecycle but type-specific controls.




Figure 6 — Governed data/AI product lifecycle activity model.

Product / capability family	Examples	Additional control emphasis
-----------------------------	----------	-----------------------------

Natural Language / GenAI	language understanding/generation, speech, translation, chatbots, LLM applications	prompt/data leakage; hallucination/grounding; content provenance; instruction hierarchy; tool access; human disclosure
Computer Vision	image recognition/generation, object detection/tracking, facial recognition, video analysis	biometric/privacy impact; dataset representativeness; adversarial examples; false-match consequences
Robotics / autonomous systems	navigation, planning, sensor fusion, robot arms, HRI	physical safety; real-time constraints; fail-safe states; human override; environment assumptions
Knowledge & reasoning	knowledge graphs, expert systems, planning, logic/inference	source provenance; reasoning trace; rule conflicts; stale knowledge; authority boundaries
Classical ML	supervised/unsupervised/RL, anomaly detection, clustering, SVM	data leakage; evaluation design; drift; explainability; reproducibility; model risk
Data product	curated datasets, semantic products, metrics, APIs, features	contract/SLA; quality; lineage; ownership; privacy/security classification; discoverability
Agentic AI	goal-driven agents using tools, APIs, memory and workflows	identity; least agency; tool permissions; runtime authorization; planning risk; memory isolation; action confirmation; traceability

9. UML use-case model




Figure 7 — GoPES UML use-case diagram.

The use cases below follow UML use-case semantics: each describes an externally observable goal achieved through interaction between an actor and the system. «include» is used for mandatory reused behaviour; «extend» is used for conditional behaviour.

UC01 — Register Data/AI Product

Field	Specification
Primary actor	Business/Data Owner
Goal	Create a governed record before controlled use begins.
Preconditions	Owner and scope known.
Main success flow	Create ID; assign owner/steward; classify product type; capture purpose, users, data, models, agents, suppliers and jurisdictions; set lifecycle state; generate baseline documentation checklist.
Postconditions	Governed object exists in catalog with traceable ownership.
Exceptions / alternate flow	Duplicate/unowned item → reject or merge.

UC02 — Assess & Classify Risk

Field	Specification
Primary actor	Owner / Product Lead
Goal	Determine risk tier and required controls.
Preconditions	UC01 complete.
Main success flow	Assess value, impact, criticality, personal data, automated decisions, security exposure, autonomy, safety, regulatory scope, third parties and change magnitude; calculate risk tier; identify required specialist assessments.
Postconditions	Risk tier and mandatory control set recorded.
Exceptions / alternate flow	Insufficient evidence → status remains Draft/Assessed with blockers.

UC03 — Approve Data Use / Access

Field	Specification
Primary actor	Data Owner
Goal	Grant purpose-bound data access.
Preconditions	Data asset classified; purpose known.
Main success flow	Check purpose, minimisation, quality, privacy/security conditions, user/workload identity, retention, geography and data contract; approve/deny/condition.
Postconditions	Time/scoped authorization with obligations.
Exceptions / alternate flow	Policy conflict → escalate.

UC04 — Perform DPIA / ADM Assessment

Field	Specification
Primary actor	DPO / Privacy Lead
Goal	Assess privacy risks and safeguards.
Preconditions	Personal data or trigger criteria identified.
Main success flow	Map processing; necessity/proportionality; lawful basis; rights; transparency; retention; transfers; automated decision safeguards; mitigations; residual risk; consultation/escalation if required.
Postconditions	Approved privacy assessment and obligations.
Exceptions / alternate flow	High residual risk → reject/escalate.

UC05 — Security Threat & Agent Review

Field	Specification
Primary actor	CISO / Security
Goal	Establish security controls for system/model/agent.
Preconditions	Architecture and data flows available.
Main success flow	Threat model; identities; secrets; supply chain; tool permissions; network/data egress; prompt/content threats; abuse cases; testing; resilience; incident plan.
Postconditions	Security control profile and release conditions.
Exceptions / alternate flow	Critical finding → block release.

UC06 — Ethics / Fairness Assessment

Field	Specification
Primary actor	Ethics/Risk Lead
Goal	Evaluate non-legal human/societal impacts.
Preconditions	Use case, population and outcome defined.
Main success flow	Assess fairness, autonomy, dignity, harm, accessibility, contestability, explainability, stakeholder impact and sustainability; define metrics/mitigations.
Postconditions	Ethics disposition and measurable obligations.
Exceptions / alternate flow	Unmitigable harm → reject/escalate.

UC07 — Build / Validate / Release

Field	Specification
Primary actor	Product/AI Team
Goal	Produce a releasable product under approved controls.
Preconditions	Design approved.
Main success flow	Build/configure; data/model versioning; automated tests; quality checks; performance/fairness/robustness/security evaluations; documentation; release evidence; approval gate.
Postconditions	Signed release artefacts and approved version.
Exceptions / alternate flow	Failed evidence → return to build.

UC08 — Authorize Runtime Action

Field	Specification
Primary actor	AI Agent / Application

Goal	Obtain permission for a material data/tool/action request.
Preconditions	Agent/app identity established; PEP integrated.
Main success flow	Intercept request; gather context; evaluate G/P/E/S policies; aggregate outcomes; route human approval if required; return permit/deny/constrain/escalate.
Postconditions	Enforceable runtime decision.
Exceptions / alternate flow	Timeout/failure → fail closed for sensitive actions; safe fallback for defined low-risk cases.

UC09 — Issue Signed Decision Record

Field	Specification
Primary actor	GoPES Decision Service
Goal	Create attributable tamper-evident evidence.
Preconditions	Decision resolved.
Main success flow	Bind request, policy versions, findings, evidence hashes, approvals, obligations, validity, trace ID and signature; persist; return token/reference.
Postconditions	Verifiable decision record.
Exceptions / alternate flow	Signing/storage failure → do not authorise sensitive action.

UC10 — Monitor / Reassess / Revoke

Field	Specification
Primary actor	DGO / Product / Platform
Goal	Ensure approval remains valid during operation.
Preconditions	Active product.
Main success flow	Collect KPI/KRI/quality/drift/security/cost/user outcome events; compare thresholds; detect material change; re-run policy; suspend/revoke if thresholds exceeded.
Postconditions	Continuous assurance and current status.
Exceptions / alternate flow	Breach → UC11.

UC11 — Manage Exception / Incident

Field	Specification
Primary actor	Owner / DGO / SMEs
Goal	Control deviations and adverse events.
Preconditions	Issue or policy conflict exists.
Main success flow	Classify severity; contain; assess impact; notify required parties; create exception only if justified; define compensating controls/expiry; investigate root cause; remediate; close and learn.
Postconditions	Resolved issue/approved bounded exception with evidence.
Exceptions / alternate flow	Repeated/expired exception → escalation/suspension.

UC12 — Audit / Report / Evidence

Field	Specification
Primary actor	Auditor / Council / Regulator
Goal	Demonstrate governance and control effectiveness.
Preconditions	Evidence retained.
Main success flow	Query catalog/records; validate signatures/provenance; sample decisions; test controls; produce KPI/KRI/exceptions/incidents/compliance reports.
Postconditions	Assurance report and findings.
Exceptions / alternate flow	Evidence gap → control deficiency.

UC13 — Retire / Delete / Archive

Field	Specification
Primary actor	Owner / Steward / Platform
Goal	Safely end a product/data/model/agent lifecycle.
Preconditions	Retirement approved.

Main success flow	Revoke identities/access; stop endpoints; archive required evidence; execute data/model retention/deletion; update dependencies; user/supplier communication; close contracts.
Postconditions	Retired state with verifiable disposal/retention record.
Exceptions / alternate flow	Legal hold → preserve specified evidence.

UC14 — Discover & Govern Shadow AI

Field	Specification
Primary actor	DGO / Security / Business
Goal	Bring unsanctioned AI use into governance or stop it.
Preconditions	Discovery signal exists.
Main success flow	Discover via network/SaaS/browser/expense/user reporting; identify owner and data exposure; triage risk; block/contain if necessary; register legitimate use; migrate to approved platform or retire.
Postconditions	Shadow AI is approved and governed, or removed.
Exceptions / alternate flow	Unknown owner/high-risk data → immediate containment/escalation.

UC15 — Productise & Commercialise

Field	Specification
Primary actor	Product/Commercial Owner
Goal	Move approved capability to a repeatable commercial offering.
Preconditions	Technical/product approval and value case exist.
Main success flow	Define brand/positioning; FR/NFR; pricing; market analysis; financial model; contracts; SLAs; support model; GTM; KYC/AML where relevant; revenue/cost KPIs.
Postconditions	Commercial product with governance obligations embedded.
Exceptions / alternate flow	Unacceptable legal/security/ROI outcome → stop or redesign.

10. UML class and information model




Figure 8 — Core GoPES class/information model.

The class model deliberately separates governance objects from policy, control, assessment and decision evidence. This prevents approvals from being stored as unstructured documents and allows a single product or data asset to be re-evaluated when policies, evidence or operating conditions change.

Class	Key semantics
GovernanceObject	Abstract governed item such as a data asset, data product, AI product, model, agent, API, business process or supplier capability.
DataAsset	Classification, retention, lineage, quality and privacy/security attributes for data.
DataAIProduct	Customer/business outcome, interfaces, SLOs, dependent data/models/agents and lifecycle state.
Agent	Machine identity, tool/action scopes, memory policy, autonomy level, budget and runtime constraints.
Policy	Versioned executable or advisory rule scoped by domain, jurisdiction, product type, risk tier and effective period.
Control	Implementation and testable control objective, mapped to policies/standards.
Assessment	Structured review with score/findings/evidence and validity period.
DecisionRequest / Decision	Runtime or lifecycle authorization request and resolved outcome.
DecisionRecord	Immutable/tamper-evident evidence binding policies, evidence, reasons, approvals and signature.
Exception	Time-bounded deviation with justification, compensating controls, owner and expiry.
Metric / AuditEvent	Outcome and control measurements plus machine-readable audit/telemetry trail.

11. Runtime orchestration and signed decision records




Figure 9 — Runtime authorization sequence.




Figure 10 — Logical schema for a signed decision record.

11.1 Decision semantics

Outcome	Meaning	Typical obligations
PERMIT	Action may proceed as requested.	logging, retention, rate limit, provenance, user disclosure
PERMIT_WITH_CONSTRAINTS	Action may proceed only if attached obligations are enforced.	mask fields, use approved model, tool allow-list, max spend/tokens, human confirmation before commit
DENY	Action is prohibited under current policy/context.	user-safe explanation; incident creation where appropriate
ESCALATE	Automated policy cannot or must not decide.	route to named role; freeze sensitive action until decision
SUSPEND / REVOKE	Existing approval is withdrawn due to change, breach, expiry or threshold.	disable endpoint/agent/tool/data access; notify owner; initiate reassessment

The signed record is the authoritative bridge between the control plane and execution plane. It should be independently verifiable, reference exact policy and evidence versions, have a bounded validity period, and be linked to subsequent execution events. A blockchain is not required; tamper-evidence can be achieved through append-only stores, cryptographic signatures, hash chains, WORM storage or a ledger service. The source diagrams' blockchain token concept is therefore retained as a capability requirement rather than a mandatory technology choice.

12. Deployment architecture and integration




Figure 11 — Deployment architecture separating control, data, execution and observability planes.

12.1 Integration patterns

Pattern	Use	Design note
Synchronous authorization API	Agent tool calls, sensitive data access, release gates, privileged business actions.	Low-latency PDP response; deterministic fail behaviour; signed decision token returned.
Asynchronous event bus	Monitoring, drift, telemetry, catalog changes, incidents, policy-change propagation.	Kafka or cloud event services are suitable; events require schema/version/trace IDs.
Batch assessment	Portfolio reassessment, supplier reviews, policy changes, control recertification.	Evaluate current inventory against changed policy and create remediation backlog.
CI/CD policy gate	Build/release workflows.	Block artefacts that lack evidence, approved dependencies, security/model tests or required signatures.
Agent middleware hook	Agent planning/tool execution.	Enforce at each material action, not only at session start; bind machine identity and context.
Human task/workflow	High-impact or ambiguous decisions.	Approver must see evidence/reasoning and be able to reject/modify outcome meaningfully.

13. Lifecycle states and change control



Figure 12 — State machine for governed products, models and agents.

Approvals are version-specific. A change is material when it can alter risk, affected people, policy scope or performance: new data purpose, new data category, model/provider change, significant prompt/tool/memory change, new jurisdiction, new user population, increased autonomy, new automated-decision effect, security architecture change, major

performance/drift shift or material supplier change. Material change returns the object to assessment before continued operation.

14. Functional and non-functional requirements

ID	Capability	Requirement
FR-01	Inventory	Register and uniquely identify data assets, products, models, agents, APIs, suppliers and governed business processes.
FR-02	Ownership	Require accountable owner, steward/custodian and relevant specialist roles before activation.
FR-03	Policy-as-code	Represent executable rules with version, owner, scope, priority, effective dates, test cases and standard mappings.
FR-04	Context	Resolve authoritative identity, resource, purpose, location, classification, risk, consent/lawful-basis, supplier and incident attributes.
FR-05	Decision	Return permit/deny/constrain/escalate with reasons, obligations and validity.
FR-06	Human approval	Support routed, role-based and recorded approvals with delegation and conflict-of-interest controls.
FR-07	Evidence	Store structured assessments, test results, model/system cards, DPIAs, threat models, exceptions, supplier evidence and policy mappings.
FR-08	Decision record	Cryptographically sign and persist decision records; link to execution events.
FR-09	Monitoring	Ingest quality, security, model/agent, cost and outcome telemetry and evaluate thresholds.
FR-10	Change / revocation	Reassess on material change and suspend/revoke prior authorization where needed.
FR-11	Exception management	Support owner, justification, compensating controls, approver, scope, expiry and review.
FR-12	Reporting	Provide operational, business, control, risk, exception and audit reporting with metric provenance.
FR-13	Shadow AI	Discover/triage unregistered AI use and either onboard or block/retire.
FR-14	Third-party	Maintain supplier/model-provider due diligence, contractual controls, incidents and material-change evidence.
FR-15	Data quality	Measure and enforce validity, accuracy, completeness, consistency, uniqueness, timeliness and lineage thresholds.
FR-16	Privacy	Support DPIA, rights, retention, minimisation, lawful basis and ADM safeguards.
FR-17	AI assurance	Support evaluation plans, benchmark/test sets, robustness/fairness tests, red teaming, drift and model/agent cards.
FR-18	Commercial governance	Link business case, budget, ROI, pricing, productisation and GTM measures to the governed object.

ID	Requirement
NFR-01 Security	Zero-trust identities; least privilege; strong authentication; encryption; secrets isolation; secure defaults; defence in depth.
NFR-02 Availability	Define tiered SLOs. Runtime authorization must support resilient deployment and explicit fail-open/fail-closed policy by action class.
NFR-03 Performance	Target decision latency by use case; cache only context/policies with safe expiry and revocation semantics.
NFR-04 Integrity	Policies, evidence, decision records and deployment artefacts are versioned and tamper-evident.
NFR-05 Auditability	Every material decision is attributable to human/machine identity and exact policy/evidence versions; trace IDs cross system boundaries.
NFR-06 Explainability	Provide policy-level reason codes and user/approver explanations appropriate to audience and legal requirements.

NFR-07 Privacy	Minimise personal data in control-plane telemetry; apply retention and access rules to audit/evidence data itself.
NFR-08 Interoperability	Vendor-neutral APIs/events; adapters for common IAM, catalog, lineage, model registry, SIEM, ticketing and agent platforms.
NFR-09 Portability	Keep domain policy and evidence models independent from specific cloud/model/agent vendors where practical.
NFR-10 Scalability	Horizontal scaling of stateless decision services; asynchronous ingestion for high-volume telemetry.
NFR-11 Usability	Role-specific views and guided evidence collection; avoid governance forms that duplicate system-of-record data.
NFR-12 Accessibility	User and approver interfaces meet organisational accessibility standards; ethics assessment includes accessible design.
NFR-13 Maintainability	Policy test suites, schema versioning, controlled configuration, automated regression and dependency inventory.
NFR-14 Recoverability	Backup/restore, disaster recovery, rollback, key recovery and continuity for control/evidence stores.
NFR-15 Safety	Kill switch, action ceilings, safe fallback, human override and containment for high-impact AI/agent systems.
NFR-16 Cost control	Meter model/tool use and enforce spend/token/action budgets for agents and AI products.

15. Business, financial, organisational and KPI integration

Diagrams 008–011 are retained as first-class governance concerns. GoPES should prove not only compliance but also whether governed capabilities deliver value, remain affordable and can be safely operated by the organisation.

Area	Required records / controls	Core measures
Sponsorship & funding	sponsor; business case; capital/opex; approval authority; benefit owner; growth plan	funding coverage; burn vs plan; benefit owner coverage
ROI tracking	baseline; cost model; benefits; confidence; review period; attribution method	ROI/NPV where appropriate; realised vs forecast benefit; unit economics
Budget allocation & consumption	service/model/tool consumption; chargeback/showback; budget variance	cost per transaction/use case; model/token spend; variance; forecast
Commercialisation	FR/NFR, value proposition, pricing, contracts, SLA, GTM, support, product reviews	revenue; margin; pipeline; adoption; churn; release quality
Organisational readiness	awareness, skills, training, resistance/change plan, communications, roadmap adherence	training coverage; proficiency; adoption; exceptions caused by skills gaps
KPIs/KRIs	metric definitions, owner, source, calculation, threshold, period, lineage	business outcomes; operational efficiency; customer impact; employee/adoption; control effectiveness; risk

16. Core operating procedures

Procedure	Minimum workflow
New use-case intake	Create minimum record; assign accountable owner; identify purpose, users, data, AI/agent capabilities and supplier; initial risk triage within defined SLA.
Policy change	Draft → peer review → SME review → approval → test → versioned deployment → inventory impact analysis → remediation backlog → effective date.
Release gate	Verify required assessments/evidence, test results, unresolved findings, approved data/model versions, security posture, monitoring, rollback and decision-record integration.
Runtime authorization	PEP captures request → context enrichment → G/P/E/S evaluation → conflict resolution → human escalation if required → signed decision → enforcement → telemetry.
Material change	Detect via change request or telemetry; compare against materiality rules; suspend or constrain if necessary; reopen assessments; approve new version.
Incident	Contain → preserve evidence → impact assessment → privacy/security/regulatory notifications as applicable → root cause → corrective action → policy/test updates → closure.
Exception	Document requirement breached, rationale, business need, risk, compensating controls, authority, scope, expiry and monitoring; no permanent waivers.
Periodic recertification	Revalidate ownership, purpose, data classification, supplier/model versions, access, evidence age, metrics, exceptions and retirement candidates.
Retirement	Stop service; revoke identities/keys/tool grants; delete/archive data/artefacts per policy; update catalog and dependencies; preserve required audit evidence.


17. Traceability to the supplied diagrams

Source diagram	Concepts retained	Integrated location
001 Data Quality Overview	Lifecycle loop: define, identify, implement, monitor; data quality strategy and assessment.	Sections 6, 13, 16; Fig. 5/12.
002 Data Quality Framework	Validity, completeness, uniqueness, accuracy, consistency, timeliness.	Section 6; Fig. 5.
003 Data Assessment	Assessment criteria, current/to-be maturity, acceptance, planning, KPIs, reporting.	Sections 3, 14–16.
004 DGO Responsibilities	Council/DGO, business, technical, SMEs and DGO responsibilities.	Section 4; Fig. 3.
005 Data Quality Framework — Data Owner	GoPES interface connected to organisational, policy, business process and system capabilities.	Sections 3–4.
006 GoPES Structure	Primary organising structure: OC1, PS1, BP1, SC1; P-D-E-M interface.	Whole document, especially Section 3.
007 Privacy	Data security/classification, use cases, enforcement, exceptions, references.	Section 7.2/7.4; requirements FR-16.
008 Budget & Financial Tracking	Funding, ROI, budget allocation/consumption.	Section 15.
009 Commercialisation	Productisation, pricing, GTM, business readiness/KYC/AML.	Sections 8, 9 UC15, 15.
010 ORM	Analytics awareness, skills, training, resistance/change.	Sections 4, 15.
011 KPI	Metrics layer and business/operational/customer/employee/marketing KPIs.	Sections 5, 15.
012 Security	Classification, encryption, erasure, masking, resilience.	Section 7.4; NFRs.
012.1 KYC/AML	Business-readiness KYC/AML workflow and evidence.	UC15 / regulated-process extension.
013 Data Privacy	Privacy bus, use cases, enforcement, exceptions, reference policy.	Section 7.2 and Privacy Engine.
014 Ethics	Ethics engine, token/record, 11 ethical stubs.	Section 7.3 and Decision Record.
015 Elements Diagram	AI/data management, training, evaluation, deployment, monitoring, security/ethics, applications/trends.	Sections 5, 8, 14.


016 Engines & Tokens Wider View	G/P/E/S/regulatory engines, data/model pipeline, BCT/token, inference interface.	Sections 5, 11.
017 Product (Data) Requirements	AI product families and product catalogue.	Section 8.
018 Systemic Alignment	MI criteria, ethics stubs, inference, security/privacy/system alignment.	Sections 7–11.
019 Engines	API engine/master bus, rules, Kafka, Hazelcast and build LM integration.	Sections 5 and 12 as implementation options.

Appendix A – Consolidated source diagrams


The following contact sheets reproduce the source diagrams supplied for this consolidation. The integrated model in the main body normalises terminology and removes duplicate or technology-specific assumptions while retaining their intent.




001 Data Quality Overview.ppt




002: Data Quality Framework.rpt




003 Data Assessment only




004 DGO Responsibilities() noc




005: Data Quality Framework: Data Domain(1).png



006: Data Quality Framework: Data Domain(2).png




007: Privacy(1).png




008: Budget and Financial Tracking(1).png


Appendix source sheet 2.




009.CommercialisationClass Diagram01(T).png



010.CRM(1).png




011. KPI(1).png




012. Security(1).png


Appendix source sheet 3.




012.1+1020 KYC and AML Process+T1.png



014. Shioct11.png




013. Data PrivacyT1.png



015. Elements DiagramT1.png

Appendix source sheet 4.



Appendix source sheet 5.

Appendix B — Reference architecture checklist

#	Architecture / governance check
1	Named product/data/model/agent owner and steward/custodian
2	Purpose and affected users documented
3	Risk tier/materiality assigned
4	Data classification and lineage complete
5	Privacy/lawful basis/DPIA/ADM safeguards complete if applicable
6	Security threat model and agent/tool permissions approved
7	Ethics/fairness/accessibility review complete to tier
8	Model/system/agent evaluation evidence complete
9	Supplier/model-provider assessment current
10	Policy and control mappings complete
11	Human approval points explicit
12	Monitoring metrics, thresholds and escalation routes active
13	Signed decision record integration tested
14	Incident, rollback and kill-switch procedures tested
15	Budget/cost limits configured
16	Exception register reviewed and expiries enforced
17	Retirement and retention plan defined

Appendix C — Web sources and standards

Ref	Source	URL
[R1] OMG	Unified Modeling Language (UML) 2.5.1 specification	https://www.omg.org/spec/UML/2.5.1
[R2] ISO	ISO/IEC 42001:2023 — Artificial intelligence management system	https://www.iso.org/standard/42001
[R3] ISO	ISO/IEC 38507:2022 — Governance implications of the use of AI by organizations	https://www.iso.org/standard/56641.html
[R4] ISO	ISO/IEC 23894:2023 — Artificial intelligence — Guidance on risk management	https://www.iso.org/standard/77304.html
[R5] NIST	AI RMF 1.0 Core: Govern, Map, Measure, Manage	https://airc.nist.gov/airmf-resources/airmf/5-sec-core/
[R6] NIST	NIST AI 600-1 — Generative Artificial Intelligence Profile	https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence
[R7] ISO	ISO/IEC 27001:2022 — Information security management systems	https://www.iso.org/standard/27001
[R8] NIST	Cybersecurity Framework (CSF) 2.0	https://www.nist.gov/publications/nist-cybersecurity-framework-csf-20
[R9] ISO	ISO/IEC 27701:2025 — Privacy information management systems	https://www.iso.org/standard/27701
[R10] ICO	Guide to the UK GDPR data protection principles (updated 2026)	https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-protection-principles/a-guide-to-the-data-protection-principles/
[R11] ICO	Artificial intelligence and data protection guidance/resources	https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/
[R12] GOV.UK	Data (Use and Access) Act 2025 — data protection and privacy changes, including automated decision-making safeguards	https://www.gov.uk/guidance/data-use-and-access-act-2025-data-protection-and-privacy-changes
[R13] OWASP	OWASP GenAI LLM Top 10 2026	https://genai.owasp.org/resource/owasp-genai-llm-top-10-2026/
[R14] OWASP	Top 10 for Agentic Applications for 2026	https://genai.owasp.org/resource/owasp-top-10-for-agentic-applications-for-2026/
[R15] OWASP	Agent Control Standard (ACS), September 2026	https://genai.owasp.org/resource/agent-control-standard-ac/
[R16] OWASP	GenAI Security Industry Framework Crosswalk, September 2026	https://genai.owasp.org/resource/genai-security-industry-framework-crosswalk/
[R17] NIST	Privacy Framework 1.1 initial public draft / update status	https://www.nist.gov/privacy-framework/new-projects/privacy-framework-version-11

Reference status note: NIST states that AI RMF 1.0 is being updated, and as of the referenced NIST page the Privacy Framework 1.1 remains an initial public draft/“coming soon”. GoPES should therefore version its framework mappings and not hard-code a single external edition.