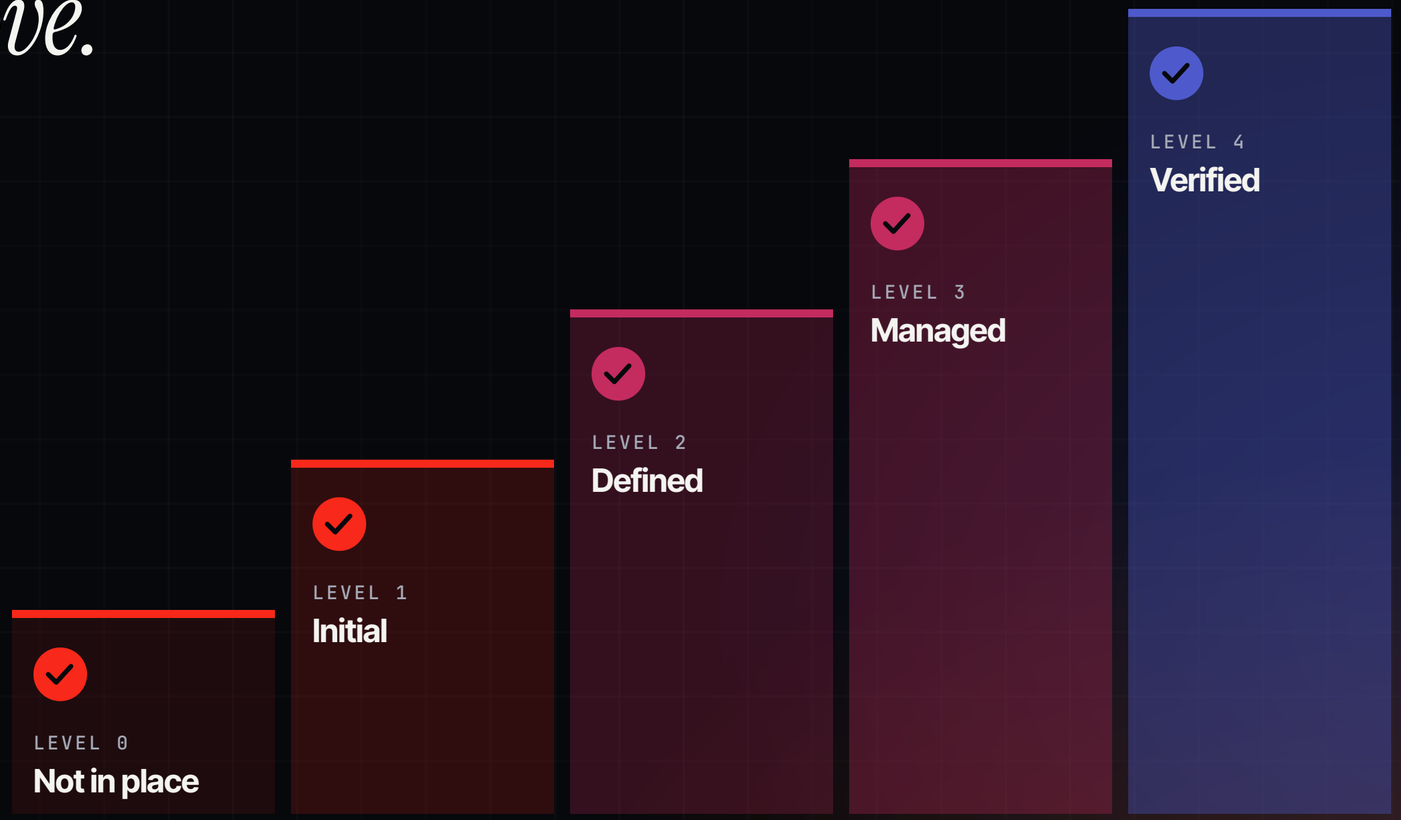
COMPLIANCE AND SELF-ASSESSMENT GUIDE

Know where you stand. Show what you can *prove*.

A control model, seven frameworks and a 42-question self-assessment.



QODIQA-CG-1.0

VERSION 1.0 · 2026-10-29 · NOT LEGAL ADVICE · NOT A CERTIFICATION

7	42	0–4
FRAMEWORKS	QUESTIONS	MATURITY LEVELS

CONTENTS

What is *inside.*

01 **Control model and crosswalk** 04

1.1 Seven capability areas 05

1.2 From request to evidence 06

1.3 Coverage matrix 07

1.4 Crosswalk: EU AI Act, GDPR, HIPAA, 42001 08

1.5 Crosswalk: ISO/IEC 27001, NIST AI RMF, OECD 09

02 **Self-assessment and audit readiness** 10

2.1 Questionnaire: identity and consent 11

2.2 Questionnaire: policy and semantic inspection 12

2.3 Questionnaire: execution control and evidence 13

2.4 Questionnaire: verification, and the score sheet 14

2.5 Reading your result: levels and a worked example 15

2.6 Gap-to-action table 16

2.7 Audit-readiness checklist 17

2.8 What evidence to export 18

2.9 Legal boundary 19

ABOUT THIS DOCUMENT

A guide to *evidence*, not a promise of compliance.

It shows how a QODIQA deployment produces evidence, how that evidence maps to seven frameworks, and how to assess your own readiness.

DOCUMENT ID	QODIQA-CG-1.0
VERSION	1.0
DATE	2026-10-29
COMPANION FILES	QODIQA-Self-Assessment.csv and QODIQA-Self-Assessment.xlsx: the same 42 questions with scoring formulas.
HOW TO USE	Part 01 explains the control model and maps it to seven frameworks; Part 02 is a printable, fillable self-assessment. The PDF runs no scripts: total the score on the score sheet or in the spreadsheet.
SOURCES	The QODIQA specification corpus (QODIQA-* documents); the documentation of qodiqa-gateway v1.0.0; qodiqa-verify and qodiqa-anchors v1.0.0; the earlier QODIQA vs Regulation volume (mapping wording reused).
LICENCES	All rights reserved for this guide (licence pending the owner's decision). Source material, as stated in the repositories: the specification corpus (qodiqa) is CC BY 4.0; the gateway, the verifiers and the anchors are Apache License 2.0.
CONTACT	contact@qodiqa.com · www.qodiqa.com

FOUR DISTINCTIONS KEPT THROUGHOUT

1 **Implemented**

In the gateway code and its documentation (qodiqa-gateway v1.0.0). The limits in its scope page still apply.

2 **Specified**

Required or described by the specification corpus, whether or not a given deployment implements it.

3 **Planned**

Roadmap only: automatic per-category retention applied by a policy rule (today DATA_RETENTION only declares retention settings), restart-safe revocation de-duplication, persistent anomaly baselines. None is supported today.

4 **Out of scope**

Outside a runtime enforcement layer: model training, fairness, organisational governance, legal determinations.

Names of regulations, standards and bodies identify the frameworks only. No relationship, endorsement or partnership with any named body is stated or implied. See the legal boundary on page 19.

PART 01

01

Control model and crosswalk

Seven capability areas, how runtime enforcement turns every decision into evidence, and what that evidence supports across seven frameworks.

1.1 Seven capability areas 05

1.2 From request to evidence 06

1.3 Coverage matrix 07

1.4 Crosswalk: EU AI Act, GDPR, HIPAA, 42001 08

1.5 Crosswalk: ISO/IEC 27001, NIST AI RMF, OECD 09

1.1 SEVEN CAPABILITY AREAS

Seven areas, one control *model.*

AREA AND OBJECTIVE	IMPLEMENTED IN THE GATEWAY	EVIDENCE IT LEAVES	SPECIFIED IN
Identity Know who or what is asking, and who operates the system.	Callers use per-identity API keys; optional OIDC, SAML 2.0 or LDAP single sign-on (tested against test identity providers only). Operators use RBAC keys and TOTP MFA step-up.	Caller context in the signed artifact; enabled and revoked state per signing identity.	QODIQA-RISK-2026-001 QODIQA-TM-2026-001
Consent Act only within what the data subject allowed, and stop when it is withdrawn.	Consent, with purpose and lawful basis, is checked on every request; after a revocation the next request is blocked. Consents can be suspended, resumed and revoked.	Append-only consent audit log and history; a BLOCK decision with its reason.	QODIQA-CORE-2026-001 QODIQA-GEI-2026-001
Policy Decide by written, versioned, approved rules.	Tenant policy is evaluated on every request: eleven rule types, ten of which decide a request; DATA_RETENTION declares retention settings. Changes need review and approval; the policy hash is in the audit chain.	Versioned policy history; the policy hash on each decision.	QODIQA-CORE-2026-001 QODIQA-GEI-2026-001
Semantic inspection Screen what a request asks for, knowing the limits of screening.	Optional and best effort: deterministic keyword classifiers that call no external model, an embedding gate whose model files are not distributed, and policy-driven data loss prevention.	Screening evidence attached to the audit trail; DLP findings as rule ids and text hashes, never the values.	Outside QODIQA-CORE-2026-001 (§1.6); described in the gateway documentation
Execution control Let nothing reach the provider without verified authorization.	Signed Ed25519 artifact (60 seconds at most) bound to method, path and body; single-use nonce in Redis; fail-closed steps. Kill switch, human approval and narrowing delegation for agents.	An ALLOW or BLOCK decision with a reason; an INTENT entry before the provider is called.	QODIQA-CORE-2026-001 QODIQA-GEI-2026-001 QODIQA-ARCH-2026-001
Evidence Record every decision so it can be checked later.	Hash-chained audit log and an RFC 6962 Merkle transparency log; signed checkpoints and receipts; sealed archive export. Enforcement-chain audit data is never deleted automatically.	Audit entries, transparency leaves and roots, receipts, checkpoints, sealed archives.	QODIQA-AUDIT-2026-001 QODIQA-ARP-2026-001
Verification Let a third party check the evidence without trusting the operator.	Offline verifiers in JavaScript and Python (qodiqa-verify) pinned by published test vectors; checkpoints can be anchored outside the gateway (qodiqa-anchors) and co-signed by witnesses.	Verifier verdicts with exit codes; bundles and proofs that verify with keys obtained out of band.	QODIQA-CVS-2026-001 QODIQA-CTSS-2026-001 QODIQA-CL-2026-001

1.2 FROM REQUEST TO EVIDENCE

Every decision leaves a *checkable* trail.

01
Sign
A client asks the gateway to sign an authorization for one request. The artifact expires within 60 seconds.

ARTIFACT

02
Verify
Signature, binding, tenant and expiry are checked. The single-use nonce is recorded atomically, so a replay is refused.

NONCE

03
Consent and policy
Consent and tenant policy are checked on every request, then content screening if enabled. A refusal becomes a BLOCK with a reason.

BLOCK

04
Intent
Before the provider is called, an INTENT entry is written. A refused request has no INTENT, only its BLOCK.

INTENT

05
Record
The decision is appended to the hash-chained audit log and to the Merkle transparency log.

CHAIN · LEAF

06
Receipt
When both entries were stored, the response carries a signed receipt. Streaming responses carry none yet.

RECEIPT

WHAT THESE RECORDS CAN SHOW

- **That the gateway recorded this decision.** A receipt proves recording.
- **That entries were not changed after chaining.** Each hash includes the previous one.
- **That a checkpoint or leaf was signed by a key you can check,** if you hold the public keys from outside the gateway.
- **The reason, the consent and the policy version** recorded with the verdict.

WHAT THEY DO NOT SHOW

- **That the decision was right.** A receipt proves recording, not correctness.
- **That the operator behaved well,** unless checkpoints are anchored with an independent witness.
- **That a writer-capable party did not regenerate the chain,** without an independent tail, anchor or signed transparency leaves.
- **That every rejection is chained.** Early rejections (input guards, artifact decoding, signature check) are counted and sampled, not chained; COMMIT entries need audit schema v8, off by default.

RECORDS YOU CAN EXPORT (SEE WHAT TO EXPORT)

AUDIT ENTRIES | TRANSPARENCY PROOF | SIGNED CHECKPOINT | RECEIPT | EVIDENCE PACKAGE | EVIDENCE BUNDLE | REQUEST PROOF | SEALED ARCHIVE

1.3 COVERAGE

The full *map.*

	CONSENT	POLICY	SIGNED	AUDIT	ISOLATION	SEMANTIC
EU AI Act	○	●	●	●	○	○
GDPR	◐	●	●	●	◐	○
HIPAA	◐	●	○	●	◐	○
ISO/IEC 42001	○	●	●	●	○	◐
ISO/IEC 27001	○	●	●	●	◐	—
NIST AI RMF	○	●	●	●	○	◐
OECD AI Principles	◐	○	●	●	○	○

● **Direct**
The control enforces the requirement at runtime, deterministically.

○ **Supporting**
The control produces verifiable evidence for the requirement, or covers part of it. Complementary controls are needed for the rest.

— **Not claimed**
The requirement lies outside the runtime boundary. No coverage is claimed.

The columns are the six primitives: Consent Gate, Policy Engine, Signed Decisions, Audit Chain, Tenant Isolation & Keys, Semantic Guard. The RAM (QODIQA-RAM-2026-001 §2.3) uses four values: Direct, Structural, Partial, Out of Scope. Where the two differ, read the more cautious one.

ISOLATION

Row-level security is not a boundary on its own: the tenant context is a session setting that the application role can change (gateway limitations).

KEYS

KMS providers are tested only against mocks and local emulators, never against a real cloud KMS or HSM (gateway cryptography guide §3.1).

SEMANTIC

Screening is optional, English-oriented and best effort; its model files are not distributed, and no detection rate is claimed (gateway limitations).

● DIRECT ○ SUPPORTING — NOT CLAIMED

TECHNICAL CONTROLS SUPPORTING COMPLIANCE. NOT A CERTIFICATION OR LEGAL ADVICE.

1.4 CROSSWALK 1 OF 2

Four frameworks, one set of *controls.*

EU AI Act

REGULATION (EU) 2024/1689 · 6 MAPPED

REF.	REQUIREMENT	QODIQA ENFORCES
Art. 9	**Risk management**	Risk rules enforced on every request
Art. 10	**Data governance**	Art. 10 training data: not claimed
Art. 12	**Record-keeping**	Automatic, hash-chained event logs
Art. 13	**Transparency**	Signed receipt for every decision
Art. 14	**Human oversight**	Override and revoke, effective next request
Art. 15	**Robustness & security**	Deterministic, signed, isolated

RAM ADDS

Direct: Art. 12 logging, Art. 26 use within purpose. Partial: 9, 14. Structural: 13, 17. Art. 10 training data: not claimed.

YOUR ORGANISATION

Risk classification, conformity assessment, post-market monitoring, registration, training-data governance, oversight roles.

GDPR

REGULATION (EU) 2016/679 · 6 MAPPED

REF.	REQUIREMENT	QODIQA ENFORCES
Art. 5(1)	**Purpose & minimisation**	Purpose checked on every AI call
Art. 7(1)	**Demonstrable consent**	Consent version bound to each receipt
Art. 7(3)	**Withdrawal**	Enforced on the very next request
Art. 25	**Privacy by design**	Protection inside the execution path
Art. 30	**Records of processing**	Exportable record, per tenant
Art. 32	**Security of processing**	Isolation, KMS keys, signatures

RAM ADDS

Direct: Art. 5(1)(b), 6(1)(a), 7(1), 7(3), 25. Structural: 5(1)(f), 5(2), 30, 32. Notices (Art. 13–14): not claimed.

YOUR ORGANISATION

Legal basis and consent quality, data subject rights, DPO duties, transfers, privacy notices, DPIA, breach notification.

HIPAA

45 CFR PART 164 · 6 MAPPED

REF.	REQUIREMENT	QODIQA ENFORCES
§164.508	**Authorizations**	PHI gated on a valid authorization
§164.502(b)	**Minimum necessary**	Data classes limited by policy
§164.312(a)	**Access control**	Scoped to tenant and subject
§164.312(b)	**Audit controls**	Hash-chained activity record
§164.312(c)	**Integrity**	Alteration is detectable
§164.312(d)	**Authentication**	Identity bound to every decision

CORPUS POSITION

The RAM does not map HIPAA. The corpus states that QODIQA neither guarantees nor certifies HIPAA compliance (QODIQA-PSL-2026-001 §2.2).

YOUR ORGANISATION

The lawfulness of every use of health data. The deploying organisation remains solely responsible (QODIQA-PSL-2026-001 §2.2).

ISO/IEC 42001

AI MANAGEMENT SYSTEM · 2023 · 6 MAPPED

REF.	REQUIREMENT	QODIQA ENFORCES
Cl. 6.1	**Risk actions**	Encoded as enforceable policy
Cl. 8.2-8.3	**Risk treatment**	Applied in-line, per request
Cl. 9.1	**Monitoring**	Continuous decision telemetry
A.7	**Data for AI systems**	Purpose and data-class checks
A.9	**Use of AI systems**	Intended-use boundaries enforced
A.10	**Third parties**	One gateway to every AI provider

RAM ADDS

Direct: Cl. 8.1, 8.6, 9.1. Partial: 6.1.1, 8.3, 8.4. Structural: 9.2, 9.3. Not claimed: 6.1.2 risk assessment, 8.5.

YOUR ORGANISATION

The management system itself: context, leadership, planning, support, internal audit, management review, AI risk assessment.

1.5 CROSSWALK 2 OF 2

Three frameworks, and what stays *outside*.

ISO/IEC 27001

INFORMATION SECURITY · 2022 · 6 MAPPED

REF.	REQUIREMENT	QODIQA ENFORCES
A.5.15	**Access control**	Policy-based access to every model
A.5.23	**Cloud services**	All AI egress through one point
A.5.34	**Privacy & PII**	Consent applied before PII leaves
A.8.15	**Logging**	Append-only, hash-chained
A.8.16	**Monitoring**	Anomalies surface in real time
A.8.24	**Cryptography**	Ed25519 signatures, KMS-managed keys

RAM ADDS

Direct: A.5.33, A.8.15, A.8.24. Partial: A.5.18, A.8.2. Structural: A.5.20, 5.36, 8.16, 8.25, 8.33. Not claimed: A.8.28.

YOUR ORGANISATION

The ISMS: risk assessment, statement of applicability, asset, physical, supplier, continuity and HR security, independent audit.

NIST AI RMF

AI RISK MANAGEMENT FRAMEWORK 1.0 · 4 MAPPED

FUNCTION	FOCUS	QODIQA ENFORCES
Govern	**Accountability**	Policies as code, owned and versioned
Map	**Context**	Tenant, subject, purpose, data class
Measure	**Telemetry**	Verdict rates, reasons, screening
Manage	**Response**	EXECUTE, BLOCK or revoke, in-line

RAM ADDS

Direct: control effectiveness, risk response, documented residual risk. Structural: Govern accountability, Map context, testing.

YOUR ORGANISATION

Governance programme, risk culture and stakeholder engagement; fairness and bias (not claimed). Alignment is not NIST certification.

OECD AI Principles

OECD RECOMMENDATION ON AI · 2024 · 5 MAPPED

PRINCIPLE	VALUE	QODIQA ENFORCES
1.1	**Inclusive growth**	Not claimed — outside runtime scope
1.2	**Human rights & privacy**	Consent enforced before any model runs
1.3	**Transparency**	Reason, consent and policy in every verdict
1.4	**Robustness & safety**	Override, revoke, stop — instantly
1.5	**Accountability**	Traceable decisions in a verifiable chain

RAM ADDS

Direct: 1.4 robustness, 1.5 accountability. Partial: 1.2. Structural: 1.3. Not claimed: 1.1.

YOUR ORGANISATION

Governance structures, organisational commitment and societal engagement. No claim of adherence to the Principles as a whole.

Stays with your organisation

RAM §9.1 · TWELVE AREAS

Model training governance	ML governance; EU AI Act Art. 10
Algorithmic fairness and bias	Fairness evaluation; impact assessments
Content moderation of output	Output moderation; trust and safety
Legal validity of consent	Legal counsel; DPO assessment
Data subject rights administration	DSR platforms; legal operations; DPO
Privacy notice and UX delivery	Product design; privacy engineering
International data transfers	Legal counsel; transfer impact assessments
Business process governance	AI governance structures; board oversight
ISMS completeness	A certified ISMS programme; ISO 27001 audit
Conformity assessment	Notified bodies; EU AI Act process
Incident notification	Incident response plans; authorities
AI system impact assessment	DPIA process; fundamental rights teams

02 SCORING METHOD

How the *score* works.

HOW TO ANSWER

YES	Implemented, and you can show the evidence today.
PARTIAL	Implemented in part, or in place without evidence you could show an auditor.
NO	Not implemented, or not evidenced. A plan is a No.
N/A	Does not apply to your deployment. Left out of the score.

THE FORMULA

Area score = (Yes + 0.5 × Partial) ÷ (Yes + Partial + No) × 100

Example: 4 Yes, 1 Partial, 0 No and 1 N/A give (4 + 0.5) ÷ 5 × 100 = 90.0.

N/A	Removed from the numerator and the denominator. An area where every answer is N/A is not scored.
BLANK	Not scored and not counted. Answer every question before you read a result.
OVERALL	The same formula over all 42 answers. It is not the mean of the area scores.
CAP	The overall level is capped at one level above the level of the weakest area.
USE	A readiness indicator for your own use. Conformance to a CORE profile is binary (QODIQA-CL-2026-001 §6) and is established by the Conformance Test Suite, not by this score.

MATURITY LEVELS 0 TO 4

0 — **Not in place** — 0–19%
No control in place, or none that you could evidence.

1 — **Initial** — 20–39%
Controls are ad hoc and the evidence is incomplete.

2 — **Defined** — 40–59%
Documented and mostly operating; verification is not routine.

3 — **Managed** — 60–79%
Operating and monitored; evidence is exported and sampled offline.

4 — **Verified** — 80–100%
Verified with independent keys and anchors, tamper detection tested, independently assessed.

2.1 QUESTIONNAIRE 1 OF 4

Identity and *consent.*

IDN **Identity** Know who or what is asking, and who operates the system.		BASIS	YES	PART.	NO	N/A	NOTES OR EVIDENCE REFERENCE
IDN-1	Does each application or service that requests signed authorizations use its own identity and credential, with no credential shared between callers?	RISK A-01 · TM §12.4 IDENTITY.md	○	○	○	○	
IDN-2	Is every signing identity limited to a tenant, an environment and a provider scope, and can you disable or revoke one identity without affecting the others?	RISK A-01 IDENTITY.md	○	○	○	○	
IDN-3	Are tenants separated from each other, and have you tested that separation in your own environment instead of relying on the design alone?	CL-50 · CL-52 multi-tenancy.md	○	○	○	○	
IDN-4	Do operators work under role-based access with least privilege, and is MFA step-up enforced for administrative changes?	CL-33 RBAC_MODEL.md	○	○	○	○	
IDN-5	If you use single sign-on (OIDC, SAML 2.0 or LDAP), have you validated it against your own production identity provider? Answer N/A if you do not federate.	RISK A-01 IDENTITY.md · limits	○	○	○	○	
IDN-6	Are signing keys kept outside the application directory under a key-encryption key, with a rotation and emergency-revocation procedure that you have exercised?	SEC §3 · SEC §5 key-management.md	○	○	○	○	

CNS **Consent** Act only within what the data subject has allowed, and stop when it is withdrawn.		BASIS	YES	PART.	NO	N/A	NOTES OR EVIDENCE REFERENCE
CNS-1	Is every AI request tied to a recorded consent with a purpose and a lawful basis before it can reach a provider?	CL-06 · CL-07 · GI-01 overview.md	○	○	○	○	
CNS-2	Does every consent you grant carry an expiry date or duration, so that none stays valid indefinitely by default?	CL-08 · GI-03 REFERENCE.md	○	○	○	○	
CNS-3	Is the very next request blocked after a consent is withdrawn, and have you proved it with a real withdrawal in your environment?	CL-09 · GI-03 overview.md	○	○	○	○	
CNS-4	If you run more than one gateway instance, do you know how long a withdrawal takes to reach all of them? Answer N/A for a single instance.	CL-09 limitations.md	○	○	○	○	
CNS-5	Do you keep the full consent history (grant, suspension, withdrawal, expiry) for as long as the decisions that relied on it?	ARP §3.2 · §5.2 dsr.md	○	○	○	○	
CNS-6	Is there a documented process for data-subject requests (rectification, erasure, restriction, portability) that keeps the audit chain intact?	RAM §9.1 dsr.md	○	○	○	○	

2.2 QUESTIONNAIRE 2 OF 4

Policy and semantic *inspection.*

POL **Policy** Decide by written, versioned, approved rules.		BASIS	YES	PART.	NO	N/A	NOTES OR EVIDENCE REFERENCE
POL-1	Are your enforcement rules held as versioned policy, and is each published version immutable? Changes must produce a new version.	CL-12 · GI-16 policy-governance.md	○	○	○	○	
POL-2	Does a policy change need a second person? Draft, review, approval and publication should not all be done by the author.	CL-45 · CL-46 policy-governance.md	○	○	○	○	
POL-3	Do you test policy changes with the simulation engine on representative scenarios before publishing, knowing it uses synthetic inputs and does not replay history?	RAM §5.3 policy-simulation.md	○	○	○	○	
POL-4	Can you show which policy version applied to any past decision, and reconstruct the policy that was in force on a given date?	ARP §3.3 · GI-16 DATA_GOVERNANCE.md	○	○	○	○	
POL-5	Where your obligations limit data location or data volume, have you written them as policy (residency, redaction profiles) and declared the provider regions yourself?	RAM §3.1 DATA_GOVERNANCE.md	○	○	○	○	
POL-6	Is the emergency policy path access-controlled, with retroactive approval and review, and do you know who may use it?	CL-45 · CL-46 policy-governance.md	○	○	○	○	

SEM **Semantic inspection** Screen what a request asks for, knowing the limits of screening.		BASIS	YES	PART.	NO	N/A	NOTES OR EVIDENCE REFERENCE
SEM-1	Have you decided and written down the role of content screening in your control set? The core standard leaves it outside the enforcement decision.	CORE §1.6 · ARP §9.3 content-screening.md	○	○	○	○	
SEM-2	Do you know which screening layers are active in your deployment? The embedding layer needs model files that are not distributed with the gateway.	CORE §1.6 limitations.md	○	○	○	○	
SEM-3	Has your risk owner accepted in writing that screening is best effort: English-oriented, misses paraphrases and flags some legitimate requests?	PSL §2.4 limitations.md	○	○	○	○	
SEM-4	Is data loss prevention configured for personal data and secrets in requests, answers and tool results, with rules that match the data you really handle?	RAM §3.1 DLP.md	○	○	○	○	
SEM-5	Have you measured screening errors on your own representative traffic instead of relying on the developers' figures, which come from their own data?	CL §6 limitations.md	○	○	○	○	
SEM-6	Do you review screening outcomes (blocks, held requests, misses) on a schedule and feed the findings back into policy?	RAM §5.3 (MEASURE) content-screening.md	○	○	○	○	

2.3 QUESTIONNAIRE 3 OF 4

Execution control and *evidence.*

EXE **Execution control** Let nothing reach the provider without verified authorization.		BASIS	YES	PART.	NO	N/A	NOTES OR EVIDENCE REFERENCE
EXE-1	Is the gateway the only network path from your applications to AI providers, with direct access to the providers closed at the network level?	GI-06 · CL-24 limitations.md	○	○	○	○	
EXE-2	Does every request carry a short-lived artifact bound to its method, path and body? V1 does not bind query or headers; V2 does and is off by default.	SEC §4.3 PROTOCOL.md §4	○	○	○	○	
EXE-3	Is replay protection active, and have you shown that a second use of the same artifact is refused?	SEC §4.5 enforcement-flow.md	○	○	○	○	
EXE-4	When a dependency fails (consent store, policy, audit write, Redis), does the system refuse the request, and have you tested each failure?	CL-18 · GI-12 enforcement-flow.md	○	○	○	○	
EXE-5	Have you rehearsed the global kill switch, and is a named owner authorized to use it?	RAM ENF-HALT overview.md	○	○	○	○	
EXE-6	For agent actions and tool calls, are high-consequence actions held for human approval, and are delegations narrowing and revocable? N/A if you run no agents.	CL-55 · CL-58 DELEGATION.md	○	○	○	○	

EVD **Evidence** Record every decision so it can be checked later.		BASIS	YES	PART.	NO	N/A	NOTES OR EVIDENCE REFERENCE
EVD-1	Is every decision, allowed or blocked, written to the audit chain, with an INTENT entry written before the provider is called?	CL-15 · ARCH §4.3 overview.md	○	○	○	○	
EVD-2	Are signed receipts returned, and do you collect them inside the retrieval window (default one hour) when you need them later?	AUDIT §6.3 claims C-032	○	○	○	○	
EVD-3	Are retention periods for audit and consent evidence defined per category, documented, and set to the longest obligation that applies to you?	CL-17 · ARP §5.2 DATA_GOVERNANCE.md	○	○	○	○	
EVD-4	Do you export sealed archives of audit history on a schedule, keep them off the gateway host, and verify them offline?	ARP §5.2 · CERT §7.3 audit-archival.md	○	○	○	○	
EVD-5	Is the audit chain monitored for integrity continuously, with an alert path and a rule that findings are investigated and never silently repaired?	ARP §5.3 · CERT §7.2 chain-integrity.md	○	○	○	○	
EVD-6	Can you assemble, for any one request, the decision, audit entries, inclusion proof, checkpoint and receipt on demand for an auditor?	CL-32 · ARP §7.1 REFERENCE.md	○	○	○	○	

2.4 QUESTIONNAIRE 4 OF 4

Verification, and your score *sheet.*

VER **Verification** Let a third party check the evidence without trusting the operator.		BASIS	YES	PART.	NO	N/A	NOTES OR EVIDENCE REFERENCE
VER-1	Have you verified a real evidence bundle or proof offline with the independent verifiers, using public keys obtained from outside the gateway?	CL-36 · AUDIT §8.3 evidence-bundle.md	○	○	○	○	
VER-2	Do you run the published protocol test vectors with both verifiers (JavaScript and Python) when you accept a gateway upgrade?	CL-27 · CERT §5.1 verify README.md	○	○	○	○	
VER-3	Are signed checkpoints anchored outside the gateway host, with witnesses whose independence you have recorded? Production needs anchoring to start.	AUDIT §6.4 audit-anchoring.md	○	○	○	○	
VER-4	Do you check the chain against an independently held tail, an anchor or signed transparency leaves, not against the chain alone?	AUDIT §6.2 verify verifiers.md	○	○	○	○	
VER-5	Do you periodically confirm that tampering is detected, for example with the tampered example bundles shipped with the verifiers, and keep the result?	CERT §5.2 verify README.md	○	○	○	○	
VER-6	Is an independent assessment of the enforcement layer scheduled, internal or third party, with findings tracked to resolution?	CL-66 · ARP §8.1 limitations.md	○	○	○	○	

SCORE SHEET		YES	PART.	NO	N/A	SCORE %	LEVEL
IDN	Identity						
CNS	Consent						
POL	Policy						
SEM	Semantic inspection						
EXE	Execution control						
EVD	Evidence						
VER	Verification						
ALL	**Overall**						

HOW TO TOTAL

Count the circles you marked in each area and enter the counts. Leave N/A out of the score.

Score % = (Yes + 0.5 × Partial) ÷ (Yes + Partial + No) × 100.

Level: 0 below 20%, 1 from 20%, 2 from 40%, 3 from 60%, 4 from 80%.

Overall level: no higher than one level above the weakest area.

The companion spreadsheet totals this for you.

2.5 INTERPRETATION AND WORKED EXAMPLE

Reading your *result*.

WHAT EACH LEVEL LOOKS LIKE, AND THE NEXT STEP

0 Not in place
Controls are absent or undocumented, and no evidence can be shown.
NEXT Start with execution control and evidence: close direct provider access and keep the audit chain.

1 Initial
Some controls run but ad hoc; evidence exists in parts.
NEXT Name an owner for each area and list every No in a gap register with a date (QODIQA-CL-2026-001 §6).

2 Defined
Controls are documented and mostly operating; evidence is not routinely verified.
NEXT Export a bundle or proof and verify it offline with keys obtained from outside the gateway.

3 Managed
Controls operate and are monitored; evidence is exported and sampled offline.
NEXT Anchor checkpoints with independent witnesses, test tamper detection and schedule an independent assessment.

4 Verified
Evidence is verified with independent keys and anchors, tested, and independently assessed.
NEXT Hold the level: repeat the questionnaire after each upgrade, policy change and incident.

QODIQA-ARP-2026-001 §2.3 names three audit-readiness levels (Partial, Structured, Continuous). The five levels here are this guide's own; the corpus does not define them.

WORKED EXAMPLE

Illustration only. A fictional mid-sized organisation runs one production gateway for an internal assistant and a document-summarisation service. It has no single sign-on, one instance and no agents, so IDN-5, CNS-4 and EXE-6 are N/A.

AREA	YES	PART.	NO	N/A	SCORE	LEVEL
Identity	4	1	0	1	90.0%	4
Consent	3	1	1	1	70.0%	3
Policy	3	2	1	0	66.7%	3
Semantic inspection	1	3	2	0	41.7%	2
Execution control	3	1	1	1	70.0%	3
Evidence	3	2	1	0	66.7%	3
Verification	1	2	3	0	33.3%	1
Overall	**18**	**12**	**9**	**3**	**61.5%**	**3 → 2**

Level 3 by score, Level 2 after the cap

61.5% reads as Level 3, but Verification (33.3%, Level 1) caps the result at Level 2. The first actions are the verification and semantic-inspection gaps: verify a real bundle offline, anchor checkpoints and measure screening on the organisation's own traffic.

2.6 GAP-TO-ACTION

From a gap to an *action.*

AREA	IF YOU ANSWERED NO OR PARTIAL	SPECIFIED IN (CORPUS)	GATEWAY DOCUMENTS (UNDER DOCS/)
Identity	Give every caller its own scoped signing identity. Put operators behind RBAC with MFA step-up. Keep signing keys outside the application directory and exercise rotation and emergency revocation.	QODIQA-SEC-2026-001 QODIQA-RISK-2026-001 QODIQA-TM-2026-001 CL-33, CL-50 to CL-53	guides/IDENTITY.md · security/RBAC_MODEL.md · security/key-management.md · architecture/multi-tenancy.md · operations/key-compromise-runbook.md
Consent	Record purpose, lawful basis and an expiry for every grant. Test that a withdrawal blocks the next request, measure revocation across instances, and document the data-subject request process.	QODIQA-CORE-2026-001 QODIQA-GEI-2026-001 QODIQA-ARP-2026-001 CL-06 to CL-10; GI-01 to GI-03	guides/dsr.md · api/REFERENCE.md · limitations.md · overview.md
Policy	Move rules into versioned policy. Require a second person to approve publication, simulate on representative scenarios, record the policy hash, and declare residency and redaction rules.	QODIQA-CORE-2026-001 QODIQA-GEI-2026-001 QODIQA-ARP-2026-001 CL-11 to CL-14, CL-45, CL-46; GI-16, GI-17	guides/policy-governance.md · guides/policy-simulation.md · guides/POLICY_LANGUAGE.md · guides/DATA_GOVERNANCE.md
Semantic inspection	Write down the role of screening in your control set. Enable and tune what you rely on, have the risk owner accept its limits, measure errors on your own traffic, and configure DLP for your data.	QODIQA-CORE-2026-001 §1.6 QODIQA-ARP-2026-001 §9.3 QODIQA-PSL-2026-001 §2.4	architecture/content-screening.md · guides/DLP.md · limitations.md
Execution control	Close direct provider paths. Keep binding and replay protection on, test every dependency failure, rehearse the kill switch, and route agent actions through approval and narrowing delegation.	QODIQA-CORE-2026-001 QODIQA-GEI-2026-001 QODIQA-ARCH-2026-001 CL-18 to CL-24, CL-55 to CL-59; GI-06, GI-12	architecture/enforcement-flow.md · protocol/QODIQA-PROTOCOL.md · protocol/DELEGATION.md · security/TOOL_GATEWAY.md · operations/redis-failure.md
Evidence	Define retention per category. Export sealed archives off the gateway host, monitor chain integrity without silent repair, and be able to assemble a proof or bundle for any request.	QODIQA-AUDIT-2026-001 QODIQA-ARP-2026-001 QODIQA-CERT-2026-001 CL-15 to CL-17, CL-32; GI-08, GI-09, GI-15	operations/audit-archival-and-retention.md · operations/audit-chain-integrity-verification.md · security/evidence-bundle-verification.md
Verification	Verify a real bundle offline with keys from outside the gateway. Run the published vectors on each upgrade, anchor checkpoints with independent witnesses, and schedule an independent assessment.	QODIQA-AUDIT-2026-001 §8.3 QODIQA-CVS-2026-001 QODIQA-CTSS-2026-001 CL-27, CL-36, CL-66	security/audit-anchoring.md · qodiqa-verify/README.md · qodiqa-verify/docs/verifiers.md · qodiqa-anchors/README.md

2.7 AUDIT-READINESS CHECKLIST

Before the *auditor* arrives.

EVIDENCE COMPLETENESS — ARP §6.1

- [ ] Decision records for the audit period are exported, blocked requests as well as allowed ones.
- [ ] Consent records and their history are exported for every consent the decisions relied on.
- [ ] Every policy version applied in the period can be retrieved with its hash.

INTEGRITY — ARP §5.3 · AUDIT §8.3

- [ ] The audit chain is verified end to end for the period and no integrity alert is open.
- [ ] Signed checkpoints and witness cosignatures cover the period; the anchor export is off-host.
- [ ] A sample of bundles or proofs is verified offline with keys obtained from outside the gateway.

RETENTION — ARP §5.2

- [ ] Retention periods are documented per category and match the longest obligation that applies.
- [ ] Nothing inside its retention window is deleted or changed; older history is in verified archives.
- [ ] Receipts you may need later were collected inside the retrieval window.

GOVERNANCE AND LIMITS — CL §6 · ARP §9

- [ ] Policy changes in the period show a second person's review and approval.
- [ ] Open No and Partial answers sit in a gap register with an owner and a date.
- [ ] A statement of what the evidence proves and does not prove is ready for the auditor.

GAP REGISTER

Per QODIQA-CL-2026-001 §6: the control point, the nature of the gap, the responsible team and a target date.

QUESTION OR CONTROL POINT	NATURE OF THE GAP	RESPONSIBLE	TARGET DATE

2.8 EVIDENCE EXPORT

What to *export.*

Receipts and bundles can be checked without the gateway. Fetch the public key sets from a source other than the instance under review. Where /v1 and plain paths both exist, they are equivalent.

ARTIFACT	WHAT IT SHOWS	HOW TO OBTAIN IT	VERIFY OFFLINE WITH
Evidence package	One decision: the signature, the canonical audit-chain hash input and a package seal.	`GET /evidence/package/{artifactId}/json` (also `/pdf`, `/audit`)	`bin/verify.evidence.external.js --keys`
Evidence bundle	The package plus a transparency inclusion proof, under one bundle signature.	`GET /evidence/package/{artifactId}/bundle`	`bin/verify.bundle.external.js --keys --transparency-keys`
Request proof	Audit entries, chain path to a checkpoint, inclusion proof, receipt, anchors and keys.	`GET /v1/proofs/{requestId}?tenantId=` (needs the `READ_AUDIT_PROOFS` permission)	`bin/qodiqa-proof-verify.js --proof`
Receipt	A signed statement that the gateway recorded the decision.	`GET /v1/transparency/receipt/{requestId}`; default retrieval window one hour	`bin/qodiqa-receipt-verify.js --keys`
Audit chain	The hash-linked log of every decision.	`audit.log`, or the `audit_chain` rows exported as JSON	`bin/qodiqa-audit-verify.js --audit`
Transparency log	Entries, signed roots, inclusion and consistency proofs.	`GET /v1/transparency/entries`, `/roots`, `/proof/{id}`, `/consistency/{from}/{to}`	`bin/qodiqa-audit-verify.js --transparency --roots --proofs --keys`
Checkpoints and anchors	The signed head of the chain, witness cosignatures and the customer-held anchor export.	`/.well-known/qodiqa/checkpoint-keys`; the file at `QODIQA_ANCHOR_EXPORT_PATH`	`bin/qodiqa-audit-verify.js --anchors`, or `bin/qodiqa-proof-verify.js --anchors-repo`
Sealed archive	A chain segment up to a signed checkpoint, with manifest and seal.	`scripts/audit/archive-export.js`	`bin/qodiqa-archive-verify.js --checkpoint-keys`
Public key sets	The trust anchors for every check above.	`/.well-known/qodiqa/keys`, `/transparency-keys`, `/checkpoint-keys`	Fetch over a channel you trust, never from the bundle you check
History and reports	Policy and consent history; a generated framework report (technical, not a legal assessment).	`/provenance/policies/{id}/history`, `/consent/{id}/history`, `POST /compliance/report/generate`	`POST /compliance/report/verify` checks the report's integrity

2.9 LEGAL BOUNDARY

What this guide is *not*.

Not legal advice

Nothing in this guide is legal advice or regulatory guidance. Interpretation, gap analysis and compliance determinations belong to qualified counsel and competent authorities.

Not a certification

This guide and its scores certify nothing. No QODIQA certification body or programme exists yet, and a self-assessment is not conformance (QODIQA-CL-2026-001 §6).

No endorsement or partnership

Frameworks, standards and bodies are named for identification only. No relationship, endorsement or partnership with any named body is stated or implied.

HOW THIS GUIDE SPEAKS

- QODIQA **supports** a framework's requirements and **provides evidence for** them.
- It never **ensures** compliance, and no sentence here should be read that way.
- Mappings are structural correspondences, not legal opinions (QODIQA-ARP-2026-001 §4).

WHAT STAYS WITH YOU

- Deployers remain solely and fully responsible for the lawfulness of their AI operations (QODIQA-PSL-2026-001 §2.2).
- The gateway enforces only the traffic that passes through it; direct provider access must be closed.
- Independent validation, certification and additional assurance where your domain needs them.

WHAT THE EVIDENCE PROVES

- A receipt proves the gateway recorded a decision, not that the decision was correct.
- A valid verdict means the evidence is consistent and signed by keys you trusted, not that the operator behaved well.
- Evidence is operator-attested unless checkpoints are anchored with an independent witness.

If you report a result

A level from this guide is self-reported. Do not represent it to external parties without having performed the verification procedures first (QODIQA-ARP-2026-001 §2.3, §6).

Claims to avoid

Any statement that deploying QODIQA, or a score in this guide, makes you compliant, certified or endorsed is not authorized (QODIQA-PSL-2026-001 §2).

Questions about this guide: contact@qodiqa.com · www.qodiqa.com.

QODIQA

Evidence, not assertions.

WWW.QODIQA.COM