

QUANTUM CORE INSTITUTE

# AI Systems and Agent Channels

*Applying the AI-cryptographic provisions across the Annex G domains*

---

| Field       | Value                                                                                                                                                                 |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Document ID | QCI-QS1-C4                                                                                                                                                            |
| Version     | V1.2 – September 23, 2026. Clause references are QCI requirement identifiers from QCI-QS1 v2.3 (September 23, 2026).                                                  |
| Applies to  | QCI-QS1 v2.2 QCI-4.4-03, 5.2-03, 5.3-01, 6.4-01; Clause 7 and Annex C section G; Annex G. Companions C1 and C2.                                                       |
| Status      | Informative client handout. Nothing in this handout creates, modifies, or interprets a requirement. Where this handout and the standard differ, the standard governs. |
| Endorsement | None. Providers named are unranked. Quantum Core Institute does not endorse, resell, or certify any provider and has not independently tested any product.            |
| License     | Free. Attribution required. May not be modified and redistributed.                                                                                                    |
| Publisher   | Quantum Core Institute Inc.                                                                                                                                           |

## 1 Read this first

QCI-QS1 v2.2 brought AI systems inside the standard without creating a parallel framework: three incident scenarios under D7 (QCI-4.4-03), linked AI records with separate confidentiality and verification horizons (QCI-5.2-03), AI exposure flags under the same rules as any other (QCI-5.3-01), a Pillar 5 provision (QCI-6.4-01) and Annex C section G. This handout shows where each of those lands in the thirteen Annex G domains, so that an AI system is inventoried and migrated through the same rows as everything else, with the AI fields riding alongside.

The scope guard applies throughout. This is cryptographic risk arising from AI systems: model weights, training data, signed provenance and dense machine-to-machine traffic between agents. It is not model behaviour, output safety or AI risk generally. An organization with no AI systems in scope records “none” and moves on.

### WHY THE REQUIREMENT EXISTS

AI infrastructure concentrates high-value cryptographic exposure in three places at once: artifacts that must stay confidential for years, signatures that a registry trusts without a human in the loop, and agent-to-agent channels that negotiate cryptography with no one watching for a downgrade. Each is an ordinary cryptographic dependency with an unusually short path from compromise to consequence.

## 2 Where the five AI fields land

| AI record element (QCI-5.2-03)             | Annex G domain                       | What to record                                                                                                                            | Remediation path                                                                                                                                                                                                                |
|--------------------------------------------|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>AI artifact class</b>                   | All; a classification, not a control | Model weights, training dataset, model registry, inference endpoint, inter-agent channel, or none                                         | Determines which of the rows below apply                                                                                                                                                                                        |
| <b>Artifact sensitivity horizon</b>        | 7 data at rest                       | Separate confidentiality horizon and provenance/verification horizon, as dates or durations with a start date; drives the HNDL arithmetic | Move artifact-key protection to a KEM-based key-protection scheme: ML-KEM with a key-derivation function and a symmetric key wrap (QCI-4.5-04) in domain 6; approved key-establishment profile on distribution paths (domain 3) |
| <b>Model signing and provenance scheme</b> | 5 PKI and code signing               | The signature scheme on model artifacts, registry entries and attestations; the signing framework, for example OpenSSF Model Signing      | Re-sign critical artifacts with ML-DSA or SLH-DSA; registry verifies the new scheme                                                                                                                                             |

| AI record element (QCI-5.2-03)        | Annex G domain                                  | What to record                                                                                                             | Remediation path                                                                                                                                                                                                                                                |
|---------------------------------------|-------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Inter-agent channel protection</b> | 3 transit, 4 remote access, 9 workload identity | Classical TLS, hybrid TLS with ML-KEM, or none; whether a downgrade is detectable; which identity authenticates each agent | Hybrid key exchange on the channel; workload identity per agent; downgrade detection under D7-AI-3                                                                                                                                                              |
| <b>AI artifact key-wrapping</b>       | 6 key management, 7 data at rest                | RSA-wrapped, ECC-wrapped, ML-KEM-wrapped, hybrid                                                                           | a KEM-based key-protection scheme: ML-KEM with a key-derivation function and a symmetric key wrap (QCI-4.5-04) in the KMS or HSM; sound AES-256 data does not need replacement, but a current rewrap does not undo a key envelope already captured (QCI-4.6-03) |

### 3 The three incident scenarios, by domain

| Scenario (QCI-4.4-03)                      | Domains | What the runbook must cover                                                                                                                                                                                                                                                                                                                      | Evidence for D7                                                                  |
|--------------------------------------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| <b>D7-AI-1 Provenance forgery</b>          | 5, 9    | Detection of a forged model signature or attestation; quarantine of the registry entry; re-signing critical artifacts with a quantum-safe scheme; rotation of registry trust roots                                                                                                                                                               | Runbook; tabletop record; list of critical artifacts and their re-signing status |
| <b>D7-AI-2 Harvested AI asset exposure</b> | 3, 6, 7 | Identification of weights, parameters or training data that transited or rested under quantum-vulnerable key establishment or key envelopes; assessment against the confidentiality horizon; a rewrap, rekey, re-encrypt, retain or retire decision that does not report a current rewrap as remediation of copies already captured (QCI-4.6-03) | Runbook; inventory of artifacts by wrapping mechanism; sensitivity horizons      |
| <b>D7-AI-3 Inter-agent downgrade</b>       | 3, 4, 9 | Detection of downgrade or silent fallback on agent channels; alerting; forced re-negotiation or channel shutdown; identification of which agent identity initiated the fallback                                                                                                                                                                  | Runbook; telemetry showing negotiated groups per channel; tabletop record        |

#### COMMON MISTAKE

Running the AI scenarios as a separate tabletop with the AI team. The scenarios sit under D7 with the organization's other cryptographic incidents, and the same responders handle them. A model registry trusting a forged signature is a PKI incident that happens to involve a model.

## 4 What is available to migrate with

No provider category is specific to AI. The relevant products are the domain 3, 5, 6 and 9 providers in handout C2, applied to AI infrastructure. The entries below are the ones with an AI-specific angle.

| Area                                    | What exists                                                                                                                                          | Notes                                                                                                                                                                                |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Model signing frameworks</b>         | <a href="#">OpenSSF Model Signing</a> specification, PKI- and algorithm-agnostic; Sigstore-based tooling                                             | Referenced in the standard as an example mechanism only; the scheme underneath must be ML-DSA or SLH-DSA to satisfy the flag                                                         |
| <b>Registry signing and attestation</b> | Registry products and artifact stores that verify signatures on pull; SPIFFE/SPIRE-based workload identity for the registry and its clients          | Entrust and others are building PQC-capable machine and AI identity on SPIRE (C2 domain 9)                                                                                           |
| <b>Agent-to-agent channels</b>          | Hybrid TLS (X25519MLKEM768) wherever the agent runtime's TLS library supports it: OpenSSL 3.5+, AWS-LC, BoringSSL, Go, rustls                        | The channel is only as good as the library under the agent framework; inventory the library, not the framework                                                                       |
| <b>Artifact key-wrapping</b>            | AWS KMS (ML-DSA signing; ML-KEM hybrid TLS to the service), Google Cloud KMS (ML-DSA and SLH-DSA signing GA, ML-KEM in preview), HSMs in C2 domain 6 | Protect the data-encryption keys with a KEM-based key-protection scheme: ML-KEM with a key-derivation function and a symmetric key wrap (QCI-4.5-04); sound AES-256 under them stays |
| <b>Hardware for physical AI</b>         | <a href="#">Infineon OPTIGA TPM</a> with post-quantum-secured firmware update, integrated with NVIDIA Jetson Thor (June 2026)                        | Edge and robotics deployments; domain 11                                                                                                                                             |

## 5 Section G in practice

Section G of the Vendor Roadmap Request Pack goes to suppliers of AI platforms, ML infrastructure, model registries, and key management used with AI artifacts. The five questions map to the fields above; the table says what a scored answer looks like.

| Question                                                         | A checkable answer names                                                                           | A non-answer sounds like           |
|------------------------------------------------------------------|----------------------------------------------------------------------------------------------------|------------------------------------|
| <b>G1 Quantum-vulnerable signatures for model signing today?</b> | The scheme (RSA-2048, ECDSA P-256) and where it is used: artifacts, attestations, registry entries | "We use industry-standard signing" |
| <b>G2 Quantum-safe model signing with ML-DSA or SLH-DSA?</b>     | The FIPS 204 or 205 parameter set, the framework, the product version and a date                   | "We are monitoring NIST"           |

| Question                                                                                     | A checkable answer names                                                                                                                                                                            | A non-answer sounds like                                                     |
|----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|
| <b>G3 Hybrid key exchange with ML-KEM for service-to-service and agent-to-agent traffic?</b> | The TLS group (X25519MLKEM768), the library and version, whether it is default or opt-in                                                                                                            | "All traffic is encrypted"                                                   |
| <b>G4 CBOM covering AI artifacts and dependencies, updated on change?</b>                    | A CBOM that meets QCI-5.4-03 and passes QCI-T1                                                                                                                                                      | A component list with no scope statement                                     |
| <b>G5 Protection of symmetric keys used for AI artifacts?</b>                                | The whole scheme: KEM, KDF, key-wrap or authenticated-encryption mechanism, recipient trust, bulk cipher and mode, lifecycle controls; historical-envelope, rewrap, backup and recovery limitations | "Data is encrypted at rest with AES-256", which answers a different question |

#### WHAT THE ASSESSOR WILL ASK

Show me the QASI row for the model registry, its five AI fields, and the D7-AI-1 tabletop record. If the registry's signing scheme is ECDSA and the sensitivity horizon passes the migration date, show me the QCI-5.3-01 AI artifact exposure flag and the decision record. That is the Handbook's fourth worked example, and it is what a well-completed row looks like.

## 6 What to do now

1. Classify every AI system in scope and complete the five conditional fields on its QASI row. A system with an artifact class of "none" is done.
2. Sequence by horizon, consequence and lead time, as QCI-4.2-02 requires. Re-sign critical model artifacts only after establishing their integrity and provenance (QCI-4.6-02), move weight distribution and training data ingestion to an approved key-establishment profile, and move artifact-key protection to a KEM-based scheme, per Clause 9.3.
3. Add the three D7-AI scenarios to the incident runbooks and run one tabletop this cycle.
4. Send Section G to every AI platform, registry and key-management supplier; score the responses; register the exceptions.
5. Give every agent a workload identity of its own (domain 9) so that an inter-agent downgrade can be attributed to an identity, not to "the agents".

## Sources

- QCI-4.4-03, 5.2-03, 5.3-01, 6.4-01, Clause 7 and Annex C section G; Annex E sources S19 (NSA/CISA/FBI AI Data Security CSI) and S21 (OpenSSF Model Signing)
- [AWS post-quantum cryptography](#) · [Google Cloud KMS quantum-safe signatures](#) · [Infineon OPTIGA TPM for Jetson Thor](#) · [Entrust CSP announcement](#)
- QCI-QS1 v2.3 (September 23, 2026), Annex G; handouts C1 and C2