

LAB L6

STEALTH

CRYPTOGRAPHY &amp; TRUST

# Post-Quantum Trust Lab

ML-KEM - ML-DSA - Hybrid - C2PA - Confidential Compute

Crypto-agile architectures, ML-KEM / ML-DSA migration paths and verifiable AI provenance.

Thesis: Crypto-agility is now a design property, not an upgrade project. Build it in or pay later.

PQ-READY NEW SERVICES

**100%**

HYBRID TLS OVERHEAD

**< 2 ms p95**

PROVENANCE COVERAGE

**100% AI outputs**

## MANIFESTO

### Why this lab exists

Harvest-now-decrypt-later is the threat regulators won't say out loud. The Lab ships crypto-agile services, hybrid PQ TLS, ML-KEM / ML-DSA migration, and verifiable AI provenance via C2PA - so when the cliff hits, you're already over it.

## KPIS

### Outcomes we measure

- PQ-ready new services: 100%
- Hybrid TLS overhead: < 2 ms p95
- Provenance coverage: 100% AI outputs

## ACTIVE EXPERIMENTS

### What the lab is testing now

#### > Hybrid TLS at scale

X25519 + ML-KEM-768 in production load; latency, CPU and rotation impact measured.

#### > ML-DSA signing pipelines

Code-signing, container-signing, and SBOM signing migrated to ML-DSA-65.

#### > C2PA on AI outputs

Every AI artefact carries a verifiable provenance manifest - model, prompt class, lineage, signer.

#### > Confidential agents

Agent runtime in TEEs (SEV-SNP / TDX / Nitro) for sovereign tenants and regulated data.

## SHIPPABLE ARTEFACTS

# Everything that ships

---

### > Crypto-agility framework

Algorithm-agnostic SDK so primitives can be swapped without touching app code.

### > Hybrid TLS rollout

Fleet-wide hybrid PQ TLS with rollback, telemetry and FIPS-mode toggles.

### > Signing migration kit

ML-DSA pipelines for code, containers, SBOMs, model weights and prompts.

### > Provenance plane

C2PA manifests on every AI output; verifier service for downstream consumers.

### > Confidential compute patterns

TEE-backed agent runtime for sovereign and regulated workloads.

## LAB TEAM

# Who you'll work with

---

- Cryptography Principal
- Confidential Compute Engineer
- Provenance / C2PA Lead
- Migration Architect

## ENGAGEMENT TIMELINE

# Weeks 1-12 - first hybrid TLS in prod by week 6

---

### 1 Weeks 1-4 - Crypto inventory + agility

Inventory primitives, ship algorithm-agnostic SDK, agree migration windows.

### 2 Weeks 4-8 - Hybrid TLS + signing

Hybrid PQ TLS in production, ML-DSA signing pipelines live, telemetry on overhead.

### 3 Weeks 8-12 - Provenance + confidential

C2PA manifests on every AI output, confidential agent runtime in TEEs for sovereign tenants.

## FLAGSHIP PODS

# Squads that productionise this lab

---

- PQ TLS Migration Pod
- Code & Model Signing Pod
- AI Provenance Pod
- Confidential Agent Pod

## PARTNERS

# Who we collaborate with

---

NIST PQC - Cloudflare Research - Microsoft Confidential - AMD SEV - Intel TDX - AWS Nitro - C2PA

PUBLICATIONS

## Receipts

---

### Hybrid PQ TLS at production scale: latency under 2ms

Real World Crypto - 2026

### C2PA for AI outputs: a deployment study across 4 sectors

AXP Internal Whitepaper - 2026

FAQS

## What partners actually ask

---

### Q. Is this real, or theatre?

A. Real. NIST has standardised ML-KEM (FIPS 203) and ML-DSA (FIPS 204). The hybrid TLS rollout is a concrete engineering programme, not a slide.

### Q. What about latency?

A. Hybrid TLS adds < 2ms p95 in our load tests at production CPU profiles. We publish the numbers per fleet.

### Q. Why C2PA on AI outputs?

A. Downstream consumers - regulators, customers, auditors - increasingly demand verifiable provenance. C2PA is the open standard winning that race.

### Q. Confidential compute, really?

A. For sovereign tenants in finance, defence and health it's table stakes. We ship reference patterns on AMD SEV-SNP, Intel TDX and AWS Nitro.

## Partner with Post-Quantum Trust Lab

Outcome-priced. Sovereign by default. Refund-backed if the contracted KPI isn't hit.

Apply: [alfaxprienz.com/labs/post-quantum-trust-lab#partner](https://alfaxprienz.com/labs/post-quantum-trust-lab#partner)