WHITEPAPER

Hybrid Cryptography: Bridging Classical and Post-Quantum Security

Why most organisations should deploy classical and post-quantum algorithms together, not switch overnight — and what that actually looks like in production.

The Case Against a Hard Cutover

Replacing a classical algorithm with a post-quantum one, all at once, in a production system is a higher-risk change than most organisations should be willing to make in a single step. The post-quantum algorithms are newer and have had less real-world deployment time than RSA or ECC, systems on the other end of a connection may not support the new algorithm yet, and a single misconfiguration can break connectivity or, worse, silently weaken security in ways that aren't immediately visible.

Hybrid cryptography — running a classical algorithm and a post-quantum algorithm together, so security depends on both being broken rather than either alone — is the pattern most security-mature organisations are converging on for the transition period.

How Hybrid Actually Works

In a hybrid key-establishment scheme, a connection derives its shared secret from both a classical algorithm (such as ECDH) and a post-quantum algorithm (such as ML-KEM) simultaneously, combining both outputs. An attacker would need to break both the classical and the post-quantum component to compromise the connection — meaning even if an unexpected weakness were later found in the still-young post-quantum algorithm, the classical component continues providing the security level organisations already trust today.

This is why "hybrid-first" is a genuinely different strategy from "post-quantum-only," not simply a slower version of the same plan. It removes the single point of failure that a pure algorithm swap would introduce during exactly the period when the new algorithms have the least real-world track record.

Where Hybrid Deployments Actually Show Up

Three deployment patterns account for most real hybrid implementations organisations are running today:

- Edge TLS — client-to-server connections through a CDN, load balancer, or API gateway, the most common and lowest-risk starting point.
- Service-to-service mTLS — internal traffic between services, often via a service mesh, with more controlled and predictable traffic patterns.
- Dual-signing supply chains — software and firmware signed with both a classical and a post-quantum signature through the CI/CD pipeline.

What to Actually Measure During a Hybrid Pilot

Hybrid cryptography does add measurable overhead — larger handshake payloads, additional computation, and in some cases longer signatures or certificates. None of this is usually prohibitive, but it should be measured under real conditions rather than assumed. A well-run pilot tracks handshake latency, CPU and memory overhead, and any downstream effects on systems sensitive to packet or payload size, before wider rollout — with a classical-only configuration kept ready as a rollback throughout.

Hybrid Is a Bridge, Not a Destination

Hybrid deployment is the right strategy for the transition period, not necessarily the permanent end state. As post-quantum algorithms accumulate more real-world deployment history and broader interoperability support, many organisations will eventually move to post-quantum-only configurations for at least some use cases. Treating hybrid as a deliberate, time-bound bridge — with a clear plan for what comes after — keeps the transition moving rather than settling permanently on the added overhead of running two algorithm families indefinitely.