Insight

"Quantum Readiness" in National Infrastructure: From Vendor Declarations to Service-Level Proof

Ofir Joseph and Nadav Cohen[1] | No. 2219 | November 29, 2026

State adversaries can already intercept encrypted communications and store them until a quantum computer becomes available to decrypt them. Consequently, preparing for the post-quantum era is a matter of national security in the present. The first post-quantum cryptography standards were published in 2024, and Israel's allies have set targets and established milestones for transitioning to these standards between 2030 and 2035. However, purchasing a "quantum-resistant" product does not guarantee that the critical service into which it is integrated will undergo the transition safely. The state's central challenge is not merely replacing algorithms, but proving that critical services—including communications, control systems, and software updates—will continue to operate safely after the change, while maintaining security and operational continuity. This article proposes measuring readiness at the service level using evidence from integration testing, and defining in advance how failures' recoveries will be handled and who has authority to accept risks during the transition. The recommendations below are intended for policymakers and regulators, who have the capacity to set uniform requirements and prioritize the most critical services.

The Threat Is Already Here: "Harvest Now, Decrypt Later"

Common mechanisms for key exchange and digital signatures currently rely on mathematical problems that a sufficiently powerful quantum computer is expected to solve efficiently. Although the exact arrival date of such a computer is unknown, the risk begins earlier: an adversary can collect encrypted traffic today and decrypt it in the future. This approach is known as "Harvest Now, Decrypt Later," and it is particularly relevant to intelligence, diplomatic communications, trade secrets, and personal data that must remain confidential for years. Israel is a persistent target of cyber espionage by state actors, and traffic collected in the past and present could become an intelligence asset in the hands of its adversaries in the future. Therefore, the lifespan of information necessitates early, phased preparation. This means that priorities should take into account how long information must remain confidential, even after it is no longer in routine use.

---

[1] Ofir Joseph heads the Cybersecurity Program at The Academic College of Tel Aviv-Yaffo and is a PhD candidate in quantum communication at the University of Haifa. Dr. Nadav Cohen is a visiting senior research fellow at the Institute for National Security Studies (INSS).

A second future risk involves digital signatures, which are used to verify system identities and authenticate software and firmware updates prior to installation. An adversary capable of forging signatures could inject a malicious update into a control component within critical infrastructure or impersonate a legitimate system; consequently, protecting the update chains of control systems and platforms expected to remain in service for many years is of particular importance.

These two risks differ in their timelines. Data exposure is a present-day risk, since what is harvested today cannot be protected retroactively. Forging signatures will become possible only when an adversary has access to a quantum computer, but equipment that verifies updates using a classical key that cannot be replaced will remain vulnerable throughout its operational lifespan. Thus, the urgency of protecting information depends on how long it must remain confidential, whereas the urgency of protecting signatures depends on the lifespan of the equipment being purchased today.

In August 2024, the U.S. National Institute of Standards and Technology (NIST) published its first three standards for post-quantum cryptography (PQC), covering key-establishment mechanisms and digital signatures, and urged organizations to begin transitioning immediately (note that this article focuses on the transition to PQC separately from Quantum Key Distribution—QKD—which requires dedicated physical infrastructure). The United States, the United Kingdom, and the European Union have published targets, recommendations, and guidelines that also include phased milestones through 2035. Some of the published requirements come earlier and are binding: under U.S. National Security Agency (NSA) guidelines, new procurements of national security systems will be required to comply with quantum-resistant algorithms starting in January 2027, while the European roadmap recommends the end of 2030 as the deadline for completing the transition in high-risk use cases. These timelines may affect defense cooperation, interoperability, and requirements imposed on Israeli exporters. Accordingly, when setting targets in Israel, it is advisable to distinguish between binding requirements imposed by partners and milestones intended to guide planning and readiness.

Initial steps have already been taken in Israel, including a January 2025 letter from the Banking Supervision Department requiring banks to prepare and map their encrypted information assets, and a 2025 directive from the Israel National Digital Agency that calls for government ministries to map cryptographic assets. Such mapping provides a snapshot of components, algorithms, and vendors involved and helps determine the scope of the required transition, though it does not indicate whether the service relying on those components can make the transition. The next phase of national policy should focus on proving the readiness of the services that depend on those components. A gradual shift from mapping to integration testing will allow mapping outputs to be used for defining areas of responsibility and creating a comparative baseline across sectors.

The Challenge of System Integration

A vendor may implement the correct software algorithm within its product, but the evidence it provides typically stops at the product's boundaries. Testing is required to demonstrate how the product functions as part of the critical service as a whole. Control systems, legacy communications equipment, and external services may respond differently to larger keys, signatures, and communication messages. For this reason, evidence of a specific product's performance should be incorporated into end-to-end service testing, under configurations and loads representative of its actual operation. Attention should

be paid to interfaces between equipment from different manufacturers, authentication processes, and potential impacts on response times and availability.

In 2025, Cloudflare described cases in which larger messages exchanged during secure connection establishment exposed assumptions built into equipment and software along the communications path. Some connections were dropped even though the cryptographic algorithm itself worked as intended, illustrating the importance of compatibility testing between components. For national infrastructures, testing should also cover industrial control systems, embedded equipment, links with limited bandwidth, and complex supply chains. In these systems, even a minor effect on setup times or message sizes can alter how the entire service operates.

In Israel, reliance on international communications and computing vendors also necessitates coordination between procurement policy and the availability timelines of upgraded solutions. A shared testing framework would enable the state to identify cross-sector dependencies (interconnections and mutual impacts among systems, components, and vendors) and establish priorities based on uniform data. A shared knowledge infrastructure could also help operators understand which technological combinations have already been tested and what limitations were observed during testing.

Readiness Is Measured at the Service Level

Consider, for example, a remote access and control service of a critical infrastructure operator, such as an electric utility operating substations. A new communications gateway may protect one connection using post-quantum cryptography, while other connections or signature mechanisms still rely on legacy solutions. Successful testing of the gateway is useful evidence, but it does not justify defining the entire service as "quantum-ready." Therefore, proving service readiness requires examining the entire set of connections, authentication measures, and recovery mechanisms upon which service delivery relies.

Another example is the update chain of platforms and weapon systems, in which digital signatures authenticate software and firmware prior to installation on critical components. If the transition includes only encrypted channels and not the signature mechanism, the platform will remain vulnerable to forged updates throughout its service life—which is why U.S. guidelines prioritize the transition to firmware signing earlier. Service-level testing examines the entire update path, from the software source to its authentication and installation on the end component. This approach enables policymakers to define readiness in relation to the service and its lifecycle, rather than in relation to an individual component.

Both examples illustrate what constitutes evidence of service-level readiness. Such evidence begins with defining the service's boundaries and dependencies, and detailing the component versions and configurations actually tested. It includes measuring compatibility, availability, latency, and continuity, alongside the level of cryptographic protection, under loads representative of the service's nature. The evidence has a limited period of validity: a substantial update or change of vendor requires re-testing. A readiness picture based on such evidence distinguishes between readiness that has already been proven and a vendor's commitment to provide a future upgrade, making it possible to track tested services, unaddressed dependencies, and upcoming milestones. On this basis, each service can be assigned a readiness level according to the scope of testing completed and the dependencies still requiring attention. Counting upgraded products, by contrast, can obscure an unaddressed critical dependency.

Recovery Is a Security Decision

A phased transition to post-quantum cryptography also requires planning recovery conditions in advance in case of a failure during deployment. Temporarily reverting to a previous configuration may restore service availability, but it removes post-quantum protection: traffic intercepted during this window will remain exposed to future decryption even after protection is restored, meaning a temporary rollback could result in permanent exposure. Furthermore, an attacker may attempt to induce a deliberate failure that triggers a switch to a weaker security configuration—a scenario known as a downgrade attack. **Rollback mechanisms thus become part of the system's attack surface.**

Recovery from a failure therefore involves a decision that weighs information confidentiality, service availability, and infrastructure safety, with the relative importance of each varying by sector. Temporarily reverting to a previous configuration is a security exception, and its parameters can be specified: who authorized it, for how long, under what monitoring, and how it was documented. An exception lacking any one of these safeguards effectively becomes a permanent state. The most serious risk lies precisely in measures intended to restore connectivity: silently disabling authentication mechanisms exposes the service to forgery and impersonation, not just data exposure. The choice between continued operation with reduced protection and suspending the service is also a choice between two different security risks. A vendor does not have the full picture needed to make this decision; in national security systems, it falls to the command or executive level. When the decision rules are set in advance, a decision made during a failure is guided by clear accountability and a known path back to the planned level of protection. When those rules are determined in real time, the decision is driven by the pressure of the moment.

Implications for Companies and Infrastructure Operators

Implementing such a policy directly affects companies' investment planning, continuity of service delivery, and contracts with vendors. Defining phased, predetermined requirements will enable infrastructure operators to coordinate equipment replacement timelines, prepare for upgrade costs, and propose alternatives for addressing key dependencies. A regulatory framework based on shared testing and uniform evidentiary requirements can reduce duplication and create certainty for long-term planning. For companies operating across multiple sectors, consistency in requirements and in how results are presented can make readiness and investment allocation more efficient.

Implications and Policy Recommendations for Israel

The demand for evidence is not an argument for waiting. Infrastructure mapping and integration testing are lengthy processes, so starting early enables the state to plan investments and gradually reduce dependencies. The recommendations below are divided into two spheres. The first sphere lies within the state's direct control: defining readiness, setting priorities, and establishing rollback rules. The second sphere depends on international vendors, and therein lies the unique implication for Israel: the pace of transition for Israeli systems will be determined to a large extent by the roadmaps of foreign vendors, who will likely prioritize the requirements of major markets—foremost among them U.S. government procurement requirements and the European roadmap. An individual Israeli organization may find it difficult to influence these priorities, but the state can increase its influence through coordinated procurement, requirements aligned with those of allied countries, and joint testing.

- **Benchmarks for Post-Quantum Readiness:** The Israel National Cyber Directorate and sectoral regulators should formulate benchmarks for assessing post-quantum readiness at the service level, each within its authority. These benchmarks should define the required evidence, compatibility testing, the validity period of an approval, and mandatory re-evaluation following a substantial change to the service. The depth of testing and reporting cycles should reflect the service's importance and its operational characteristics. Within this framework, the coverage provided by existing directives across all essential sectors, including critical infrastructure and the defense establishment, should be reviewed to ensure they encompass service-level readiness assessment rather than mapping alone. Alongside this, they should establish a measurable reporting format that includes the percentage of critical services tested, material dependencies, and approved recovery plans. Consolidating data within a shared measurement framework will allow oversight bodies to identify recurring barriers across organizations and advance cross-cutting solutions.
- **National Prioritization:** A national prioritization framework should treat the two types of risk differently: in protecting confidentiality, based on the required confidentiality period of the information; and in protecting authentication and updates, based on the expected lifespan of the equipment and the ability to replace its trust anchors. In both cases, the criticality of the service should be taken into account. Such prioritization will make it possible to direct resources to services where the combination of sensitive information, legacy equipment, and vendor dependency justifies an early start. The European roadmap, for example, classifies as high-risk any use case in which exposure of information, even ten or more years from now, would still cause significant harm.
- **Rollback Mechanisms:** The relevant regulators should establish rules for temporary operation using a previous configuration. These should specify who can authorize it, how long the exception may last, and what monitoring is required, while requiring authentication mechanisms to remain in place in every scenario. For national security systems, the authorized command or executive level should be designated in advance to decide between continued operation with reduced protection and suspending the service.
- **Coordinated Procurement and Upgrade Requirements:** The Ministry of Defense and government procurement bodies should coordinate with regulators of critical infrastructure to develop uniform acceptance requirements aligned with those of allied countries, so that vendors face a single, consistent set of Israeli requirements. Acceptance requirements should include proven evidence of interoperability, performance, and a recovery plan. Contractual documents should distinguish between capabilities that are available and have already been tested and those planned for the future, and set a binding deadline for delivery of a post-quantum upgrade in line with allied milestones. They could also establish in advance the vendor's responsibility to provide test results and coordinate necessary updates with other manufacturers in the service environment.
- **Shared Testing Environment:** A shared testing environment should be developed for essential sectors to assess compatibility among new solutions, legacy equipment, and control systems. This environment would also provide a common channel for communicating findings and

remediation requirements to vendors, and help identify early which systems are unlikely to receive timely upgrades from their vendors. This would allow decision-makers to determine whether to replace or isolate those systems, or develop a local alternative. Standardization bodies, procurement agencies, and infrastructure operators should be integrated into this framework, and findings that can be disclosed should be shared across sectors.

Conclusion

Quantum risk already requires decisions today regarding services and data that will remain critical and sensitive for years to come. The test for Israeli policy is whether it can provide evidence that an entire service is ready for the transition, including its dependencies, performance, and recovery mechanisms. This leads to the central question that should guide every discussion on transition progress: Which critical services have already proven readiness, what gaps remain, and who is responsible for closing them and managing the residual risk in the interim? An evidence-backed answer to this question carries far greater weight than any list of products bearing the label "quantum-ready."

Editors of the series: Anat Kurz, Rinat Harash, Eldad Shavit and Keri Rosenbluh